

Astronomical Searches for Nanosecond Optical Pulses

A thesis presented by

Andrew William Howard

to

The Department of Physics

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in the subject of

Physics

Harvard University

Cambridge, Massachusetts

May 2006

Astronomical Searches for Nanosecond Optical Pulses

Abstract

With “Earth 2000” technology we could generate a directed laser pulse that outshines the broadband visible light of the Sun by four orders of magnitude. This is a conservative lower bound for the technical capability of a communicating civilization; optical interstellar communication is thus technically plausible.

This thesis considers interstellar communication with nanosecond optical pulses. Its topics are the theory of such signaling, natural sources, two astronomical searches—their search methodologies, experimental implementations, candidate events, and implications—and a custom integrated circuit designed to detect such signals.

The targeted search examined some 6000 Sun-like stars with a sensitivity of ≥ 100 photons/m² in ≤ 5 ns (350–720 nm) using a 1.5 m telescope in Harvard, Massachusetts. It used a pair of hybrid avalanche photodetectors to trigger on coincident pulse pairs, initiating measurement of pulse width and intensity at sub-nanosecond resolution. An identical system on a 0.9 m telescope in Princeton, New Jersey permitted unambiguous identification of even a solitary pulse. Among the 11,600 artifact-free observations at Harvard, the distribution of 274 observed events shows no pattern of repetition, and is consistent with a model with uniform event rate, independent of target. With one possible exception (HIP 107395), no valid event was seen simultaneously at the two observatories.

The all-sky search is a pulsed optical meridian transit survey of the Northern sky ($-20^\circ < \delta < +70^\circ$) with ~ 1 min dwell time and a sensitivity of ≥ 95 photons/m² in ≤ 3 ns (300–650 nm). It uses a 1.8 m spherical telescope to image 1.6×0.2 on two matched focal planes with 512 photomultiplier tube pixels each. Coincident optical pulses trigger custom electronics to record pulse profiles and event timing. No pulses were observed during initial observations of 1% of the sky (which includes $\sim 10^5$ Sun-like stars within range).

Thirty-two PulseNet chips—a full-custom integrated circuit that forms the all-sky instrument’s computing core—digitize 1024 photodetector outputs at ≤ 1 GS/s, filter and store candidate signals, and perform astronomical observations.

For Sarah

Acknowledgments

Graduate school has been a time of great intellectual and personal growth for me. For this I am most grateful to my advisor, Paul Horowitz, whose support, encouragement, enthusiasm, intellect, creativity, and *joie de vivre* have made these past many years so stimulating and enjoyable. That which I have learned from Paul begins with physics, astronomy, and electronics, and extends much farther afield into photography, politics, movies, and myriad other topics. I feel so fortunate to have worked with him.

Thanks also to my committee — Melissa Franklin, Dave Latham, Gu-Yeon Wei, and, before he passed away, Cos Papaliolios — who have supported me through this journey.

The other students in Paul’s lab have been great friends and colleagues. Chip Coldwell showed me the ropes when I first arrived and did excellent work on the targeted search. I would probably be graduating a couple years later were it not for the contributions of Curtis Mead, Jason Gallicchio, Pratheev Sreetharan, Steve Howard, and Chris Laumann in designing and building the all-sky search. It’s been a pleasure working with you guys. Jonathan Wolff (who preceded me by a couple months), as well as Chris Clearfield and Anne Sung, contributed substantially to the SETI instruments. Last but not least of the SETI lab folks is Al Sliski, who has more engineering knowledge than I thought could possibly occupy one human brain; he has been enormously helpful with the all-sky search.

The night observers at Oak Ridge Observatory — Robert Stefanik, Joe Zajac, Joe Caruso, Dave Latham, and Guillermo Torres — stayed up all night, every night so that we could get the results of Chapter 4. Thank you, guys!

A major part of my graduate career was spent designing the custom chip for the all-sky search called PulseNet. I owe a special thanks to Bill Dally and Gu-Yeon Wei for their help designing it. Bill welcomed me into his lab at Stanford and kick-started the design with his can-do attitude, knowledge, and enthusiasm. My time at Stanford was especially enjoyable because the members of Bill’s group — Andrew Chang, Brucek Khailany, Ujval Kapasi, Patrick Chiang, Brian Towles, Mattan Erez, and Ben Serebrin — were great friends, in addition to being great resources for a novice chip designer like me. My thanks also to Dean Liu, who provided the standard cells from which the synthesized blocks of PulseNet were derived. Jon Laven, Amir Frank, and my brother Dan Howard were so generous for letting me sleep in their living room during that year; I enjoyed hanging out with you guys.

Upon returning to Harvard, I have benefited enormously from the patient advice and help of Gu-Yeon Wei. Besides being generous with his time, knowledge, and resources, Gu has been a pleasure to work with.

Throughout my graduate career I have not had to worry about funding for my stipend or research expenses. For this I am enormously grateful to the Bosack-Kruger Charitable Foundation (for supporting the students in Paul’s lab), The Planetary Society (for funding the all-sky survey, and for funding half of the targeted search), the SETI Institute (for funding the other half of the targeted search), and the MOSIS Foundation (for fabricating PulseNet twice, without charge).

My whole life I have felt the love and support of my parents, and of the rest of my family (Dan, Weeze, & Elena). Without this foundation, I very much doubt that I would be completing a Ph.D. in physics.

This dissertation is dedicated to my partner in life, Sarah Muirhead. Her consistent warmth, love, support, and enthusiasm have brightened my days immeasurably. I look forward to a wonderful life with her.

Contents

Abstract	iii
Dedication	iv
Acknowledgments	v
Table of Contents	vi
List of Figures	ix
List of Tables	xi
1 Introduction	1
1.1 Plenitude, ubiquity, antiquity	1
1.2 Overview of thesis	1
1.3 Versions of this document—official and compact	2
2 Search Strategies	3
2.1 With whom might we communicate?	3
2.1.1 Timescales	3
2.1.2 Portrait of the galaxy	4
2.1.3 Number of communicating civilizations	4
2.1.4 Conclusions	5
2.2 Means of signaling	5
2.3 Wavelength choice	6
2.3.1 Trade-offs	6
2.3.2 Pulsed versus continuous	7
2.3.3 The case for optical SETI	8
2.4 Transmission	8
2.4.1 Directivity	9
2.4.2 Aiming precision	10
2.4.3 Extinction	10
2.4.4 Number of photons sent and received	11
2.4.5 Model pulsed transmitters	12
2.5 A brief history of optical SETI	14
3 Backgrounds for Pulsed Optical Communication	16
3.1 Astrophysics on short timescales	16
3.2 Stellar photon pileup	17
3.3 Cosmic-rays and gamma-rays	18
3.3.1 Muons	19
3.3.2 Čerenkov radiation	19
3.4 Instrumental and terrestrial backgrounds	20
3.4.1 Photodetector problems	20
3.4.2 “Cultural” backgrounds	21

4	Targeted Optical SETI	23
4.1	Instrument design	23
4.2	Sensitivity	24
4.3	Harvard observations	26
4.4	Seasonal variation	27
4.5	Data sets	28
4.6	Consistent with Poisson statistics	28
4.7	Other sources of events?	29
4.8	Interesting objects and reobservations	31
4.8.1	Other objects	36
4.9	Conclusions from Harvard observations	36
4.10	Synchronized observations with Princeton	36
4.10.1	HIP 107395	37
4.11	Implications	39
4.11.1	Scenarios	39
4.11.2	Fraction of stars with transmitting civilizations	40
5	All-Sky Optical SETI	42
5.1	Motivation	42
5.2	Overview of experiment	42
5.3	Observatory	42
5.4	Telescope and optics	45
5.4.1	Spherical vs. parabolic	45
5.5	Photomultiplier tubes	48
5.6	Camera and electronics	50
5.7	Gelfand Flashers	50
5.8	Software and user interface	56
5.9	Sensitivity	56
6	PulseNet – Design and Implementation	60
6.1	Overview of Design	60
6.1.1	Notation	62
6.2	Circuits	63
6.2.1	Top-level design	63
6.2.2	Analog samplers	63
6.2.3	Clock distribution	66
6.2.4	Memory and coincidence circuitry	66
6.3	Interface with <code>seti_io</code> module	67
6.4	Interface with <code>astronomy</code> module	68
6.5	Layout	68
7	PulseNet – Testing and Verification	74
7.1	Testing Procedure	74
7.2	PulseNet Rev. 1	74
7.2.1	Problems	74
7.2.2	Measurements	75
7.3	PulseNet Rev. 2	75
7.4	Test waveforms	76
7.5	Input sampler offset voltage measurements	76

8	All-Sky Optical SETI Data Analysis	83
8.1	Calibration of telescope position	83
8.2	Measurements	84
8.2.1	Minimum observing declination	84
8.2.2	Camera Sensitivity	85
8.3	Observing procedures	88
8.4	Initial observations	88
8.5	Implications	91
8.5.1	Methodology	91
8.5.2	Constraints on transmitting extraterrestrial civilizations	93
8.6	Recommendations for future work	93
A	Optical SETI with Terrestrial Planet Finder	95
A.1	Abstract	95
A.2	Introduction	95
A.3	Laser power required	96
A.3.1	Directionality	96
A.3.2	Spectral resolution	97
A.3.3	Fluctuations	97
A.4	Backgrounds	97
A.4.1	Incompletely nulled stellar photons	98
A.4.2	Reflected photons	98
A.4.3	The planet's blackbody spectrum	98
A.4.4	Zodiacal and exo-zodiacal dust	99
A.5	Alternative planetary systems	100
A.6	Conclusions	100
B	Interstellar Gravitational Dispersion of Optical Pulses	101
B.1	Convergent and divergent scattering	101
B.2	Geometric delay	101
B.3	Gravitational delay	103
B.4	Conclusions	104
C	PulseNet – Automated Testing Procedure	105
C.1	AutoTest procedure	105
C.2	SETI functions	105
C.3	Astronomy functions	108
C.4	Results	109
	Bibliography	110
	Glossary	114

List of Figures

2.1	Interstellar extinction as a function of wavelength	11
3.1	Seasonal hit rate for the targeted search	21
4.1	Harvard targeted optical pulse detector block diagram	24
4.2	Harvard targeted search photometer photograph	25
4.3	Seasonal variation in the targeted search event rate	27
4.4	Event probability density vs. zenith angle	31
4.5	Targeted search observations on 2004 November 4	35
4.6	Targeted search observations on 2003 September 17	38
4.7	Limits on the density of transmitting civilizations	41
5.1	Transit survey of the celestial sphere	43
5.2	All-sky observatory building	43
5.3	Telescope and camera	46
5.4	Optical aberrations in spherical and parabolic optics	47
5.5	Multi-anode photomultiplier tube (PMT) dimensional outline	48
5.6	PMT output pulse properties	49
5.7	PMT spectral sensitivity and pixel layout	49
5.8	All-sky search instrument – front	51
5.9	All-sky search instrument – rear	52
5.10	Arrayed PMTs in matched focal planes	53
5.11	Motherboard and daughterboards	54
5.12	PMT pixel-daughterboard correspondence	55
5.13	All-sky software and firmware	57
5.14	Web-based user interface for the all-sky search	58
6.1	Simplified block diagram of PulseNet	61
6.2	PulseNet block diagram	64
6.3	Thermometer code sampling scheme	65
6.4	Flash converters and sense amplifiers	65
6.5	Shift register memory	66
6.6	Two-phase non-overlapping clock generator	67
6.7	PulseNet pin assignments	71
6.8	PulseNet die photograph	72
7.1	PulseNet waveform reconstruction at 1GS/s	77
7.2	PulseNet waveform capture of high information content signal	78
7.3	Uncalibrated sampler offset voltages – sorted by pixel	79
7.4	Calibrated and uncalibrated sampler offset voltages	80
7.5	Sampler offset voltages – Aa/Ab/Ba/Bb	81
7.6	Sampler offset voltages – sorted by threshold	81
7.7	Sampler offset voltage input bias and supply voltage	82

8.1	Calibration of telescope position.	84
8.2	All-sky camera sensitivity	85
8.3	PulseNet pixel dependence on sensitivity	86
8.4	Observing procedures.	89
8.5	Initial sky coverage of the all-sky search.	90
8.6	Constraints on transmitting civilizations	94
B.1	Gravitational scattering in the paraxial plane	102
B.2	Distortion of an annulus of light	102
B.3	Scattering diagram in the paraxial plane	103
B.4	General relativistic time delay	104
C.1	Automated coincidence functionality test – one pixel	106
C.2	Automated coincidence functionality test – all pixels	107
C.3	Automated voltage offset test waveform diagram	108

List of Tables

2.1	Criteria for selecting a means of signaling	6
2.2	Signal characteristics of model pulsed transmitters	13
4.1	Project parameters – targeted optical SETI at Harvard and Princeton	26
4.2	Summary of data from Harvard search	28
4.3	Harvard observations – distribution of events	29
4.4	Distribution of integration times	30
4.5	Distribution of observations	30
4.6	Interesting objects in DS_{clean} – observing statistics	32
4.7	Interesting objects in DS_{clean} – astronomical information	33
4.8	Observations of HD 220077	34
5.1	Project parameters – all-sky optical SETI at Harvard.	44
6.1	PulseNet – summary of capabilities	62
6.2	PulseNet – programming sequence for the <code>seti_io</code> interface	69
6.3	PulseNet – programming sequence for <code>astronomy</code> interface	70
7.1	Incremental PulseNet testing procedure	75
7.2	PulseNet Rev.1 – I_{DD} and I_{DD}' vs. f_{fastclk}	76
8.1	All-sky survey — initial observations	88
8.2	Comparison of all-sky survey and targeted search data	91
A.1	Alternative planetary systems	99
C.1	Serial numbers of PulseNets that failed AutoTest	109

Chapter 1

Introduction

The presence of interstellar signals is entirely consistent with all we know, and . . . if signals are present the means of detecting them is now at hand. Few will deny the profound importance, practical and philosophical, which the detection of interstellar communications would have. We therefore feel that a discriminating search for signals deserves a considerable effort. The probability of success is difficult to estimate; but if we never search, the chance of success is zero.

— Guiseppe Cocconi and Philip Morrison, 1959 [10]

1.1 Plenitude, ubiquity, antiquity

Though not always accepted, the possibility of life on other worlds captivated thinkers through the centuries. In the time of Julius Caesar, in 70 BC, Lucretius of Rome wrote in *De Rerum Natura (On the Nature of the Universe)* that the world around us is open to human understanding, and that “such combinations of other atoms happen elsewhere in the universe to make worlds such as this one . . . so we must realize that there are other worlds in this universe, with different races of men and different animals.” [46]

By the 19th century some suggested sending signals to the presumed inhabitants of the Moon and Mars. The famous mathematician Karl Friedrich Gauss proposed planting large forests in Siberia into which one would sow a giant field of wheat (for contrast) in the shape of a right triangle. An elaborated form would have had squares on each edge of the triangle so as to demonstrate our knowledge of the Pythagorean Theorem. Joseph Von Littrow suggested building canals in the shape of 20-mile wide geometric figures in the Sahara, filling them with kerosene, and setting them afire in night. By century’s end though, enthusiasm abated for signalling through means such as constructing geometric artefact and excitement turned to radio [60].

The second half of the 20th century marked a turning point. It was the first time in the history of the Earth that its inhabitants possessed the technological capability to communicate with other civilizations in our galaxy. The realization that we could build the equipment to send receive a signal over interstellar distances transformed thinking about extraterrestrials from one of speculation to experiment. The search for extraterrestrial intelligence (SETI) had begun.

Along the way, a picture of the galaxy emerged that appears entirely consistent with life elsewhere. This picture, and the arguments in favor of extraterrestrial civilizations—the vast number of stars (and probably planets) in the galaxy, the universality of the physical laws that gave rise to life on Earth, and the aeons for the genesis and evolution of life—can be expressed in many ways, but were perhaps most eloquently summarized by Philip Morrison with just three words: “plenitude, ubiquity, antiquity.”

1.2 Overview of thesis

This thesis describes the scientific basis for pulsed optical communication with extraterrestrial civilizations, as well as two searches for such signals. Chapter 2 develops the rationale for pulsed optical signaling, describes the technical limitations for transmitting and receiving such signals, provides examples of model

transmitters, and speculates on the nature of the signals that we hope to receive. Two topics that would naturally fit in this section are relegated to appendices because their treatment is too long for the main body. Appendix A explores a novel search strategy — using NASA’s planned Terrestrial Planet Finder (TPF) to look for intentional signals in the atmospheric spectra of extra-solar planets. Appendix B shows that temporal broadening of pulsed optical signals due to gravitational scattering is a negligible effect. It was written in response to discussions with Phil Morrison and Paul Horowitz in which we worried that the effect might doom pulsed optical communication over interstellar distances.

Potential sources of pulsed optical signals—astrophysical, atmospheric, terrestrial—are examined in Chapter 3. It is important to understand the sources of these signals as they are the background against which intentional signals from other civilization must compete. It was written in preparation for the Qualifying Oral Examination and was also presented as a talk and paper at the Third International Conference on Optical SETI in 2000.

Chapter 4 describes the design of the Harvard targeted optical search instrument and analyzes five years of data taken with it. The latter half of this search had the benefit of a confirmatory instrument at Princeton, and the results from this joint search are described. Additionally, implications of null observations are considered, and quantitative limits on the density of pulsed optical transmitters are set.

The motivation for and design of the all-sky search for pulsed optical signals—the primary topic of this thesis—is discussed in Chapter 5. Initial results from this instrument are analyzed and implications are considered in Chapter 8.

The photodetectors on the all-sky instrument generate a prodigious quantity of data (3.5 Tb/s, equivalent to the contents of all books in print, every second). Digitizing and processing the photodetector signals required the development of a custom chip called PulseNet. The design and implementation as well as testing and verification of this novel integrated circuit are described in chapters 6 and 7, respectively. Additional material on the automated testing and verification of every PulseNet is in Appendix C.

There is also a glossary to define acronyms, units, and uncommon terms. This was intended to make this document, especially the introductory material, readable by a larger audience (i.e. the author’s family).

Portions of this thesis were previously published:

- Chapter 3 is reprinted with permission from the SPIE proceedings for the Third International Conference on Optical SETI [32].
- Chapter 4 is a slightly reworked version of one section of a paper describing the targeted search that appeared in *The Astrophysical Journal* [35]. Other portions of that paper are sprinkled throughout Chapter 2. These portions are reprinted here under the non-exclusive right of republication granted by the American Astronomical Society. Their original copyright reads: “©2004. The American Astronomical Society. All rights reserved. Printed in U.S.A.”
- Appendix A is reprinted with permission from an article in *Icarus* on the feasibility of using TPF as a SETI instrument [33]. The original copyright notice reads: “Copyright ©2001 by Academic Press.”

1.3 Versions of this document—official and compact

There are two versions of this thesis—“official” and “compact”. The official version is on file in the Harvard library. Unfortunately, because it’s double-spaced, and the margins and fonts are large—all specific requirements for Harvard University PhD theses—it is unnecessarily long and low density.

The compact version solves these problems with single-spacing, narrower margins, and the use of 10-point font for text in the main body. Since page numbers differ between the official and compact versions, citations should reference the official version. Organizational items such as the numbers of chapters, sections, figures, and tables are consistent between the versions.

Electronic copies of the official version (and possibly the compact version) will be posted online¹ for as long as is practicable.

¹<http://www.physics.harvard.edu/academics/phds.html>

Chapter 2

Search Strategies

Portions of this chapter (perhaps 10-20%) were published previously in an article in The Astrophysical Journal (ApJ) [35]. Some of the arguments parallel those in the optical SETI section of SETI 2020 [24], one of whose authors (Paul Horowitz) was also an author of the ApJ paper.

2.1 With whom might we communicate?

2.1.1 Timescales

When imagining the type of civilization with whom we might communicate, it is often assumed that they will be approximately our age, and will have roughly the technological capabilities we possess. The symmetry of this picture is quickly discarded when one considers the great range of timescales involved in the evolution of intelligence. Life emerged on Earth about three billion years ago. Multicellular organisms didn't flower until about 500 million ago. The first anatomically modern *homo sapiens* walked the Earth half a million years ago and developed agriculture 10,000 years ago. Yet it was just in the last century that our civilization acquired the capability to communicate over interstellar distances with radio transmitters and lasers. That is, we are only 50-100 years old as a technological civilization capable of interstellar communication. Surely the exact timing of our technological birth was result of historical accidents and is decoupled from similar milestones in the independent evolution of life on other planets.

If one pictures the history of the Milky Way in fast-forward, one can imagine technological civilizations popping into existence, growing and thriving for some time, and then dying. If SETI is to succeed, civilizations must be born at a sufficient rate and the lifetime of communicating civilizations must exceed the typical interval between the births of such civilizations, and probably also the round trip communication time between two nearby civilizations (at least 10^3 – 10^4 years). If we live in a universe with many civilizations, we are surely not the first.

For the sake of argument, let us take the typical lifetime of the civilizations with whom we hope to communicate to be 10^3 – 10^6 years. On the average, their present age is half of their lifetime. So from their perspective we are a very young technological civilization. Since the Earth only very recently made the transition from a planet full of life to one that also harbors a technological civilization, we can safely assume that any contact will be with a civilization much older, and, consequently, much more advanced.

It is difficult to say how much more advanced they will be, or in what ways. Even on Earth, long term predictions about specific technological developments are usually wrong. To be conservative, we will use the current state of science and technology on the Earth as a minimum bound on the capability of an extraterrestrial civilization. Thus, when considering means of signalling and the engineering of signaling devices, we will limit ourselves to systems that are within the laws of physics as we know them and that we could actually build today (if only we had the money, time, and patience). Their science will also surpass ours, but we will assume no specific discoveries on their part. That is, we should not assume that they know of violations of the laws of physics (as we know them) that allow, for example, faster than light travel. We

assume that the means of communication available to them are also available to us.

We should assume, because of their age, that their science has been *exhaustive*. They have done all of the experiments that we can and would like to do. They have launched space telescopes like the ones that we hope and plan to launch in the near and long term. We should assume that they have exhaustive catalogs of stars and planets, constrained only by the limits of observation. This means that they undoubtedly know of the Sun’s existence and that it is a habitable G dwarf. They may have observed the wobble in the position of the Sun caused by Jupiter’s gravity through instruments like the upcoming Space Interferometry Mission (SIM). They may have even seen the Sun dim slightly as Jupiter passed in front. If they are close (100 pc away? 500 pc?), their scaled-up version of our planned Terrestrial Planet Finder (TPF) may have directly imaged the planets of the Solar System and detected the atmospheric signatures of life on Earth. We can expect to be on their target list.

2.1.2 Portrait of the galaxy

It is important to consider our place in the population of 400 billion stars that make up our galaxy. The Milky Way consists of a pancake-shaped disk of stars (radius ≈ 15 kpc and thickness ≈ 300 pc), a flattened sphere of stars near the center of the disk (radius ≈ 5 kpc), plus a halo of mostly old stars and globular clusters. The Sun is located on the inner edge of a spiral arm, near the mid-plane of the disk about 8 kpc from the center of the galaxy. The closest stars are a few pc away.

The stars near the Sun are all in the disk, which is, more precisely, two disks. The “thin disk” has nearly all of the stars and mass, and its scale height¹ is 325 pc. The “thick disk” is less dense, but extends farther out (a scale height of 1.3 kpc). Both disks have a radial scale height of 3–5 kpc. The thin disk also contains gas and dust which absorb and scatter optical photons. The density of stars (of all types) near the Sun is approximately 0.15 pc^{-3} . Within a radius of 300 pc (≈ 1000 light-years), there are $\sim 10^7$ stars, and $\sim 10^6$ “Sun-like” stars. Out to ~ 1 kpc, there are $\sim 10^8$ stars, of which $\sim 10^7$ are “Sun-like.”

Not all stars are suitable sites for life. O, B, and, A stars are probably too hot and have lifetimes that are too short. The targeted search (Chapter 4) concentrated on F, G, and K dwarfs since they are the most “Sun-like”². These make up about ~ 10 – 20% of the stellar population. The remaining $\sim 75\%$ of the stars are M dwarfs – stars with low mass (0.08 – $0.6 M_{\odot}$) and luminosity, and lifetimes on the order of 10^{11} years. The potential habitability of these stars (for microbes or intelligent life) has long been questioned [36] because of their temporal variability and high UV output during active periods, as well as tidal locking of planets in the Habitable Zone. However, recent calculations suggest that sufficient planetary CO_2 could prevent atmospheric collapse. Segura *et al.* have calculated potential biosignatures for planets around M dwarfs using instruments like Terrestrial Planet Finder [56].

It also appears that planets are common. In the last 15 years, ~ 150 planets have been discovered orbiting nearby stars. This number is expected to grow by orders of magnitude in the near future with astrometric and photometric experiments such as Gaia, Kepler, and SIM. By closely studying our galactic neighborhood, it has been estimated that more than $\sim 20\%$ of stars have planets [44].

2.1.3 Number of communicating civilizations

One usually estimates the number of communicating civilizations in the galaxy, N , by way of the Drake Equation,

$$N = R^* \cdot f_p \cdot n_e \cdot f_l \cdot f_i \cdot f_c \cdot L. \quad (2.1)$$

Drake wrote this equation at the first SETI conference, “The Order of the Dolphin,” in 1961. It is used more for estimation than firm calculation. Nevertheless, it is a useful guide for thought experiments and helps organize our ignorance. The equation is the product of several terms representing everything that has to go right to get a communicating civilization. The uncertainty in the terms grows dramatically from left to right, which are defined as follows:

- R^* is the star formation rate in our galaxy (per year)

¹The distance over which the number density of stars decreases by a factor of e .

²That is, their masses, lifetimes, spectral characteristics, etc. are most like those of the Sun, which is a G2V dwarf.

- f_p is the fraction of stars that have planets
- n_e is the average number of planets which are “Earth-like” enough to support life around those stars with planets
- f_l is the fraction of those that eventually develop life
- f_i is the fraction of those that develop intelligent life
- f_c is the fraction of those that are willing and able to communicate
- L is the expected lifetime of a communicating civilization (in years)

The Drake Equation is the subject of much debate. Optimists have used it to estimate that there are 10^6 communicating civilizations in the galaxy [54], while pessimists have found that number to be 10^0 (=1) [66]. Some find that the product of the first six terms (some of which are completely unknown) is approximately equal to one, so the equation reduces to $N = L$. In this view the number of civilizations with which we might communicate is equal to how long they survive (in years).

2.1.4 Conclusions

An important conclusion from the estimates of N from the Drake Equation is that the source of any signal will be very far away. Even for the optimistic case of $N = 10^6$, the fraction of stars that currently harbor communicating civilizations is very small: $f_{\text{now}} = N/N_{\text{MW}} = 2.5 \times 10^{-6}$ ($N_{\text{MW}} = 4 \times 10^{11}$ is the number of stars in the Milky Way). In this scenario, we must search $\sim 4 \times 10^5$ randomly selected stars just to find one signal (which is, on average, ~ 100 pc away). For an intermediate, though certainly not pessimistic, value of $N = 10^4$, $f_{\text{now}} = 2.5 \times 10^{-8}$ and a successful search will encompass $\sim 4 \times 10^7$ stars to find a signal that is almost a kiloparsec away. If the observed stars are selected by habitability criteria [63, 64] the number that we have to search probably shrinks, but by an unknown amount. SETI is a long-term activity and one should not be disheartened by the lack of immediate results. To speed up the search, we should focus on techniques that can increase the number of stars searched at once (e.g. sky surveys) and ones that increase our confidence that the stars we observe harbor life.

The above arguments—our relative youth, the enormous length and time scales involved, the fraction of stars with communicating civilizations—also suggest that transmission is foolhardy until we have searched the sky for directed communications. Let us go about that easier task and the one with potentially faster returns.

2.2 Means of signaling

In order to design experiments to search for extraterrestrial signals, we must first determine some basic characteristics of the signal carrier. Among the large number of elementary particles, which ones, if any, will another civilization choose to send to us in the form of a message?

The standard SETI answer is contained in the Project Cyclops report [6]. The report (with arguments attributed to Barney Oliver) constructs a list of a criteria for a signal particle (Table 2.1). Criteria *a-b* are requirements for interstellar signalling, and *c-h* are properties of an ideal particle.

Following the logic of this approach, criteria *a*, *c*, and *h* exclude particles of non-zero rest mass (e.g. protons, neutrons, electrons), criteria *d* and *e* exclude all charged particles, and criteria *f* and *g* exclude gravitons and neutrinos³. The only elementary particle that remains is the photon.

Cocconi and Morrison put it succinctly: “Interstellar communication across the galactic plasma without dispersion in direction and flight-time is practical, so far as we know, only with electromagnetic waves.” [10]

³Neutrinos are now known to have mass and would be also be eliminated by criteria *a*, *c*, and *h*.

Table 2.1. Criteria for selecting a means of signaling

Criteria
<i>a</i>) number of particles received must significantly exceed the background
<i>b</i>) signal must exhibit some non-natural property
<i>c</i>) minimal energy per particle, other things being equal
<i>d</i>) negligible absorption by interstellar medium (and atmosphere)
<i>e</i>) negligible deflection by galactic fields
<i>f</i>) be readily collected over a large area
<i>g</i>) permit efficient generation and detection
<i>h</i>) velocity should be as high as possible

Note. — Criteria for selecting a means of signaling. Criteria *a-b* are requirements for interstellar signalling, and *c-h* are properties of an ideal particle.

2.3 Wavelength choice

Having chosen electromagnetic radiation, what wavelength is best for interstellar communication? The Project Cyclops report, along with most of the SETI literature at that time, concluded that the microwave portion of the radio spectrum is ideal. This was the radio paradigm. Cocconi and Morrison spent considerable time discussing the “optimum channel” in their original paper [10]. A year after citing arguments including the transparency of space and of the atmosphere at radio wavelengths and the large source power requirement for optical and other wavelengths, they concluded that it was 1420 MHz (21 cm), the radio emission line of neutral hydrogen. Frank Drake independently reached the same conclusion and began radio observations of τ Ceti and ϵ Eridani at 1420 MHz, the first SETI observations of any object [20]. Indeed, at the time it was not unreasonable to pursue radio SETI at the exclusion of other wavelengths. No one could imagine sending an interstellar signal with optical search lights. Why should they look for one?

The world changed one year later when Charlie Townes invented the laser. After that, in 1961, Schwartz and Townes envisioned interstellar communication with scaled-up versions of these “optical masers” [55]. They pointed out that optical transmission was now conceivable, and that it was an historical accident that lasers were not invented thirty years earlier. They proposed a search for optical signals.

Since the early days of SETI, laser technology has been in a phase of rapid catch-up relative to the mature technology at radio frequencies. Lasers with several megawatts of continuous optical output have been built, and picosecond pulses of more than a petawatt (10^{15} W) have been produced. Progress in solid-state lasers has been impressive, and there are laser designs on the drawing board to produce repetitively pulsed megajoule nanosecond pulses. These beacons permit detection with a very simple apparatus – just a telescope with a pair of white-light photomultipliers in coincidence.

Even with the rapid growth of laser technology, it remains unclear if there is a single preferable wavelength. A prudent approach, in the view of the author, is to search for signals at all wavelengths and modulation schemes for which we can reasonably imagine building a transmitter, and to initially concentrate our efforts in those places that: *a*) are easiest to design and build detection equipment for; *b*) have been searched the least; *c*) optimize some parameters that we believe extraterrestrial civilizations might deem important (such as minimum energy per bit, or transmitter size/weight). These criteria argue for the pursuit of radio SETI (in as broad a portion of the spectrum as is practicable) and optical SETI (including new searches in the near- and mid-infrared). Both are plausible.

2.3.1 Trade-offs

Historically the Cocconi and Morrison [10] suggestion that SETI be carried out at the 21 cm emission wavelength of neutral hydrogen came at a time in our technological development when no other astronomical lines were known in the microwave, and there were no lasers. The rapid development of laser technology since

that time – a Moore’s Law doubling of capability roughly every year – along with the discovery of many microwave lines of astronomical interest, have lessened somewhat the allure of hydrogen-line SETI. Indeed, on Earth the exploitation of photonics has revolutionized communications technology, with high-capacity fibers replacing both the historical copper cables and the long-haul microwave repeater chains. Additionally, the elucidation [13] of the consequences to SETI of interstellar dispersion (first seen in pulsar observations) has broadened thinking about optimum wavelengths. Even operating under the prevailing criterion of minimum energy per bit transmitted, one is driven upward to millimetric wavelengths.

Moreover, there are other considerations that might well encourage the use of shorter wavelengths still. A transmitting civilization might wish to minimize transmitter size or weight, or use a system capable of great bandwidth, or perhaps design a beacon that is very easy to detect.

In comparing the relative merits of radio versus optical, it has sometimes been incorrectly assumed that one would always prefer coherent (heterodyne) detection, and that the noise background is given by an effective temperature $T_n = h\nu/k$. For ultra-high resolution spectroscopy one must use such a system, mixing the optical signal down to microwave frequencies where radio techniques can be used; but if one is interested instead in the detection of short pulses it is far better to use photon-counting detectors (e.g., photomultipliers) [52]. That is because the process of heterodyning and linear detection is *intrinsically* noisy, for fundamental reasons: because heterodyne detection allows a measurement of phase, there must be uncertainty in the amplitude. The added noise is immaterial in the radio region, where there are many photons per mode; but it is serious in the optical, where the photon field is dilute.

Taking these and other factors into account in a comparison of received signal-to-noise ratio (SNR) versus wavelength, and making reasonable assumptions about antenna apertures and accuracies, detection methods, transmitter power, and so on, Townes concluded in 1983 that optical methods are comparable, or perhaps slightly preferred, in the single figure of merit of delivered SNR for a given transmitter power. Other factors are obviously important – for example penetration of an atmosphere (which favors microwave) or high data rates (which favors optical) – and could easily tip the balance. His conclusion was that the SETI community’s historical bias toward microwaves should be reconsidered [62].

2.3.2 Pulsed versus continuous

Are pulses the best beacon? Or should we be looking for laser lines, transmitted continuously at some guessable wavelength, analogous to the microwave searches that have been conducted? What is natural at radio frequencies may not be so at optical. At *radio* frequencies it is easy to do coherent detection, using the ordinary heterodyne techniques of mixing with a local oscillator to a complex (quadrature) baseband. With classical filter techniques, or with contemporary digital processing with discrete Fourier transforms, one can achieve extremely narrow bandwidths, limited only by oscillator stability (a part in 10^9 is routine) and patience (the resolution is the inverse of the coherent integration time). Furthermore, the interstellar medium is kind to carriers—at gigahertz frequencies a carrier is broadened only *millihertz* in its passage through the interstellar medium, if one avoids the most congested region of the galactic center, and even there the broadening is only a few hertz. Scattering and absorption are also small or negligible over galactic distances for such signals [13]. In other words, a signal that is a spike in the *frequency* domain is a natural candidate for interstellar signaling at microwave frequencies, for reasons both scientific and technical.

Moreover, interstellar dispersion, and the presence of natural and “cultural” impulsive interference (switching transients, spark plugs, and so on), make pulses in time less effective. Finally, the relatively low carrier frequency (along with dispersion) prevents high bandwidth communications. By contrast, at *optical* wavelengths the situation is reversed: One cannot realize extremely narrowband systems with optical filters or gratings, but is forced to optical heterodyne techniques, ultimately applying precise radiofrequency spectroscopic methods at the microwave IF. This results in added noise, as mentioned above and well described by Townes [62]. Furthermore, at optical wavelengths the higher carrier frequencies ($\sim 10^{14}$ Hz) result in much larger absolute Doppler shifts; for example, $1 \text{ km s}^{-1} \leftrightarrow 5 \text{ kHz}$ at 1.4 GHz, whereas $1 \text{ km s}^{-1} \leftrightarrow 1 \text{ GHz}$ at $1 \mu\text{m}$. However, dispersion is negligible at optical wavelengths, even at nanosecond timescales [12]. Furthermore, natural and cultural sources of nanosecond flashes of significant intensity appear to be entirely absent (see Chapter 3). In other words, a signal that is a spike in the *time* domain is a natural candidate for interstellar signaling at optical wavelengths, for reasons both scientific and technical. An added bonus is that, at nanosecond time scales, the stellar background becomes negligible.

Nevertheless, one can construct plausible scenarios with continuous optical signals, particularly in the infrared, where the stellar background is reduced. Knowing that this approach works best with large telescopes already outfitted with precise spectrometers, Amy Reines and Geoff Marcy searched archived stellar spectra from the Keck Telescope for laser lines (discussed in greater detail in Sec. 2.5). Appendix A discusses another highly sensitive approach – using TPF to find continuous laser beacons.

2.3.3 The case for optical SETI

Put most compactly, the primary arguments in favor of conducting SETI at optical (rather than radio) wavelengths are:

1. Transmitted beams from optical telescopes are far more slender than their radio counterparts owing to the high gain of optical telescopes.⁴
2. Dispersion, which broadens radio pulses, is completely negligible at optical frequencies.
3. The capability of radio transmitters has reached a stable maturity, while the power of optical lasers has not yet plateaued and has shown an annual Moore’s law doubling extending over the past 30 years.
4. Natural and cultural backgrounds are negligible (though instrumental backgrounds are significant, but manageable in the current optical searches). See Chapter 3.
5. The complexity, 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.

It is also worth noting that scattering and absorption limit the range of transmission in the visible spectrum to a few kpc (see §2.4.3); however, at far-infrared wavelengths (as at microwave wavelengths) transmissions can penetrate nearly the entire galaxy unattenuated. Thus, choice of transmission wavelength may reflect the average separation between civilizations, the number of civilizations in the galaxy, and, more speculatively, the average lifetime of a civilization (by way of the Drake Equation).

2.4 Transmission

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 illuminate, with a beacon distinguishable from astrophysical phenomena and from noise, the planetary zones of the nearest N Sun-like stars within some range R_{\max} (comparable to the average separation between intelligent civilizations), or, more likely, a subset of N that the civilization deems most likely to harbor life. In our region of the galaxy $N \approx 10^6$ for $R_{\max} = 300$ pc.

To send a pulse (or more generally, a packet of information of short duration) to each of $N = 10^6$ 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 ~ 10 ms (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 intelligently selected subset of the targets and/or by using multiple transmitters; it seems altogether reasonable to expect a pulse period of 10^3 seconds or less in this scenario.)

The recipient would be able to observe these pulses only if *a*) the received fluence per pulse corresponds to at least some tens of photons delivered to the receiving telescope aperture, and *b*) the flux of laser photons, during the pulse, exceeds the stellar background. It is a remarkable fact, as we’ll show presently, that *using only “Earth 2000” technology we could generate a beamed laser pulse that outshines the Sun by four orders of magnitude, in white light, independent of range*⁵. One might consider this the “fundamental theorem of optical SETI.”

⁴150 dB for the Keck Telescope at $\lambda=1\mu\text{m}$ versus 70 dB for Arecibo at $\lambda=21$ cm, an 80 dB advantage at optical wavelengths.

⁵The light from the Sun and from the transmitter both fall off as $1/r^2$.

These pulses could be detected with an optical 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 photomultiplier tubes, and all known significant backgrounds disappear at this time scale; see Chapter 3). The electronics can be as simple as 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.

In the sections that follow, we derive several important quantities to demonstrate the feasibility of transmission. Along the way we will calculate these quantities for a model “Earth 2000” system: a Helios laser beamed 100 pc (≈ 300 ly) using a 10 m Keck-like telescope. Helios is a diode-pumped Yb:S-FAP solid-state laser designed at Lawrence Livermore National Laboratory for inertial confinement fusion, potentially capable of generating 3 ns, 3.7 MJ pulses (10^{15} W) at 349 nm (or 4.7 MJ at its native $1.047 \mu\text{m}$ wavelength) at a ~ 10 Hz repetition rate [39].

There appear to be no physical limitations to scaling up Helios to an even greater pulse energy, say 100 MJ. (Its architecture is scalable and highly parallel.) Optical telescopes are also getting larger; for example the Overwhelmingly Large Telescope (OWL) is planned to be 100 m in diameter. In §2.4.5 we calculate the transmission characteristics of several such systems that are even more advanced than “Helios with a Keck,” but still within our technological grasp.

2.4.1 Directivity

An important advantage of optical interstellar communication is the narrow width of beams formed by optical transmitters (telescopes used in reverse) allowing for precise targeting. The angular beam width, θ_b , is a function of the wavelength of the transmitted light, λ , and the transmitter diameter, D_T . For illumination of a circular transmitter aperture by a plane wave, the outgoing intensity pattern is the familiar Airy disk, given by the Fraunhofer diffraction formula,

$$I(\theta_b) = 4 \left[\frac{J_1(\pi D_T / \lambda \cdot \sin \theta_b)}{\pi D_T / \lambda \cdot \sin \theta_b} \right]^2, \quad (2.2)$$

where J_1 is a first order Bessel Function [27]. Eq. 2.2 is normalized so that $I(\theta_b=0) = 1$. The half-power beam width is obtained by finding the solution to $I(\theta_b) = 1/2$, which is $\pi D_T / \lambda \cdot \sin \theta_b = 1.61$. The full width at half maximum is twice this:

$$\frac{\pi D_T}{\lambda} \cdot \sin \theta_b = 3.22 \rightarrow \sin \theta_b = 3.22 \cdot \frac{\lambda}{\pi D_T} \quad (2.3)$$

which gives the standard result for the diffraction-limited beam size,

$$\theta_b \approx \frac{\lambda}{D_T}. \quad (2.4)$$

For a target at range R , the beam is geometrically broadened to a linear diameter,

$$D_b = R\theta_b. \quad (2.5)$$

Transmitted optical beams can be quite narrow. For our example Helios with a Keck system ($\lambda = 1 \mu\text{m}$ and $D_T = 10$ m), the $\theta_b = 2$ milli-arcseconds (mas) beam is only $D_b = 2$ astronomical units (AU) wide at 100 pc. For targets this close, the beam may be artificially broadened to ensure that the beam encompasses the target. The transmitting civilization may even tailor the angular beam width with an adaptive aperture, so as to have a fixed beam size and photon fluence at the target, independent of range [11].

The gain, G , provided by using a telescope, is the ratio of the solid angle for isotropic radiation, $\Omega_{\text{isotropic}} =$

4π , and the solid angle of the beam, $\Omega_b = \lambda^2/A_T$,

$$G = \frac{\Omega_{\text{isotropic}}}{\Omega_b} = \frac{4\pi}{\lambda^2/A_T} = \frac{\pi^2 D_T^2}{\lambda^2}, \quad (2.6)$$

where A_T is the transmitter area. For our example system, $G = 10^{15}$.

2.4.2 Aiming precision

The use of a high-gain antenna puts substantial, although not overwhelming burden on the transmitting civilization. A beam typically of width $\sim 10^{-7}$ must be pointed accurately in order to hit the target. The task is more challenging than just pointing the transmitting aperture at the position where the target appears *now*. Since the target is moving, the aiming must account for the target’s proper motion and range. To accomplish this task, we assume that the transmitting civilization has a catalog of target stars, their current positions (θ_0), proper motions (μ), ranges (R), and radial velocities (v_r), as astronomers do on Earth. How accurately must these quantities be known for successful transmission?

The sky position at which the transmitting civilization must aim (θ) is

$$\theta = \theta_0 + \mu \frac{2R}{c - v_r}. \quad (2.7)$$

Note that $R/(c - v_r) \approx R/c$ is the light travel time. At $R = 100$ pc, 10 AU beaming accuracy corresponds to a proper motion uncertainty of $100 \mu\text{as yr}^{-1}$ and a positional accuracy of 100 mas. To position the beam within 1 AU, the figures are $10 \mu\text{as yr}^{-1}$ of proper motion uncertainty and 10 mas of positional accuracy. The required *range* accuracy depends on the star’s proper motion; for example, to target the planetary zone (say 10 AU) of a star whose proper motion is 10 km s^{-1} , the range uncertainty cannot exceed 5 ly. These requirements are certainly within the grasp of an advanced civilization, given that SIM is expected to achieve astrometric precision of $4 \mu\text{as}$ for single measurements (and down to $1 \mu\text{as}$ for stars with a nearby reference) [65]; and in any case these accuracies are relaxed if the transmitted beam is broadened to illuminate a larger zone, at the expense of received signal strength.

One might also worry that while transmitted pulses may hit their target, they will be significantly dispersed in time due to scattering. Appendix B considers one aspect of this—temporal dispersion due to the gravitational scattering of large bodies—and finds the effect insignificant.

2.4.3 Extinction

The interstellar medium both scatters and absorbs 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 [12]. The prompt pulse (“ballistic” photons) is unscattered (therefore unbroadened in time) and reduced in amplitude. Absorption acts also to reduce the prompt pulse height, so that the total surviving fraction is an exponential function of the total optical depth, $e^{-\tau}$.

Extinction depends strongly on wavelength, distance to the target, and, to a lesser extent, direction. Although extinction maps accounting for the distribution of matter in the galaxy have been made [1, 47], we will use the isotropic mean extinction, $A(\lambda)/A(V)$, which is typically expressed relative to the extinction at V-band (550 nm). The left panel of Fig. 2.1 plots relative extinction as a function of wavelength for two cases of the optical parameter $R_V = A(V)/[A(V) - A(B)]$. ($A(B)$ is the mean extinction in B-band, 440 nm.) The curve labeled R3 applies to the diffuse ISM; R5 is used in dense clouds [7, 15, 25].

Using the approximate value of 1 magnitude⁶ of extinction per kiloparsec at V-band, we can convert the R3 relative extinction curve into curves for the fraction of the signal that remains after traversing various

⁶For non-astronomers: a magnitude is a logarithmic unit of brightness. The brightest stars in the night sky have a visual magnitude of 0 and the dimmest ones are about magnitude five. Five magnitudes is defined a factor of 100 in brightness; one magnitude is a factor of $10^{2/5} \approx 2.5$ (4 dB).

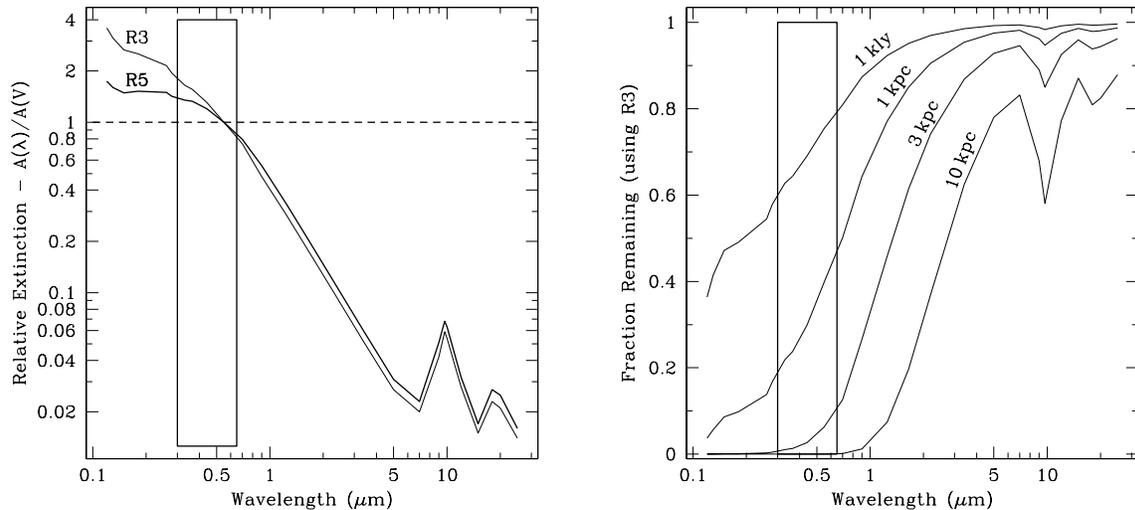


Figure 2.1. Isotropically-averaged interstellar extinction as a function of wavelength. The plot on the left shows the relative extinction as a function of wavelength. The plot on the right shows the fraction of a signal that would remain after traversing 1000 light-years (1 kly), 1000 parsecs (1 kpc = 3260 light-years), 3 kpc, and 10 kpc for R3 extinction. The feature at $10 \mu\text{m}$ in all of the curves is due to a CO_2 absorption line.

distances, as shown in right panel of Fig. 2.1. Note that interstellar extinction exponentially suppresses transmitted signals with distance; the fraction remaining is

$$F_{\text{ext}}(\lambda) = 10^{-2A_m(\lambda)/5}, \quad (2.8)$$

where $A_m(\lambda)$ is $A(\lambda)$ expressed in magnitudes.

Fig. 2.1 implies an effective range of $\sim 1\text{-}3$ kpc for visible-light communication, depending on how much extinction the transmitting civilization is willing to tolerate⁷. It may be the case that in designing a transmitter, a civilization will choose a shorter wavelength, perhaps because of decreased transmitter size and weight or the availability of high-power laser at those wavelengths, even at the expense of photons squandered due to extinction.

Nevertheless, Fig. 2.1 clearly argues that longer wavelengths ($\geq 2 \mu\text{m}$) are optimal if one wishes to minimize extinction and thereby maximize range. Ultimately, the choice of wavelength may depend on the number of communicating civilizations in the Milky Way. If there are few communicating civilizations spaced far apart, their only options may be at longer wavelengths (including radio).

2.4.4 Number of photons sent and received

We also need to calculate the number of photons sent per pulse, the number received, and the number of background photons with which the signal photons compete. The number of transmitted photons per pulse is simply the ratio of the transmitted pulse energy, E_{pulse} , to the energy per photon, E_{photon} :

$$N_{\text{pulse}} = \frac{E_{\text{pulse}}}{E_{\text{photon}}} = \frac{\lambda E_{\text{pulse}}}{hc}. \quad (2.9)$$

The Helios laser generates $\sim 2 \times 10^{25}$ photons per pulse.

The number of photons received is reduced from this number by several factors: the ratio of the receiver area to the beam area at the target, the fraction not lost to extinction, and the fraction not lost to detector

⁷Since Fig. 2.1 is isotropically-averaged, it only represents typical extinction. In some directions, suppression can be *much* more severe.

inefficiencies (q):

$$\begin{aligned} N_{\text{rec}} &= N_{\text{pulse}} \cdot \frac{\pi D_R^2/4}{\Omega_b R^2} \cdot F_{\text{ext}}(\lambda) \cdot q \\ &= q N_{\text{pulse}} \cdot \frac{\pi^2 D_T^2 D_R^2}{16 \lambda^2 R^2} \cdot 10^{-2A_m(\lambda)/5}. \end{aligned} \quad (2.10)$$

If we're interested in the fluence of photons that arrive at the Earth (independent of receiver particulars), we can convert Eq. 2.10 to the number of photons received per square meter (using SI units for the other quantities),

$$N_{\text{rec}} = N_{\text{pulse}} \cdot \frac{\pi D_T^2}{4 \lambda^2 R^2} \cdot 10^{-2A_m(\lambda)/5} \quad \text{per m}^2. \quad (2.11)$$

For our example system, $N_{\text{rec}} = 190$ photons per m^2 .

We must compare N_{rec} with the number from the transmitter's star, $N_{\text{background}}$, that are received during a short interval, τ . We can put an upper bound on $N_{\text{background}}$ by assuming that all energy radiated by the star goes into photons in the wavelength range detectable by the recipient,

$$\begin{aligned} N_{\text{bg}} &= \frac{E_{\text{star}}}{E_{\text{photon}}} \cdot \frac{\pi D_R^2/4}{4\pi R^2} \cdot F_{\text{ext}}(\lambda) \cdot q \\ &= q \cdot \frac{L_{\odot} \tau}{E_{\text{photon}}} \cdot \frac{D_R^2}{16 R^2} \cdot 10^{-2A_m(\lambda)/5}. \end{aligned} \quad (2.12)$$

Note that these photons also suffer from extinction and detector inefficiencies.

The ratio of signal to background photons is approximately

$$\begin{aligned} \frac{N_{\text{rec}}}{N_{\text{bg}}} &= N_{\text{pulse}} \cdot \frac{E_{\text{photon}} \pi^2 D_T^2}{L_{\odot} \tau \lambda^2} \\ &= G N_{\text{pulse}} \cdot \frac{E_{\text{photon}}}{L_{\odot} \tau}, \end{aligned} \quad (2.13)$$

where L_{\odot} is the power output of the Sun (which has typical stellar brightness) and E_{photon} is the energy of a typical optical photon. Note that Eq. 2.13 is independent of range, extinction, detector size, and efficiencies.

For our example system, $N_{\text{bg}} = 0.02$ photons (per m^2 in a 3 ns interval), and $N_{\text{rec}}/N_{\text{bg}} = 10^4$.

2.4.5 Model pulsed transmitters

To give a sense of the difficulty (or relative ease) of pulsed optical interstellar communication, we calculate several of the above quantities for a variety of transmission schemes. Table 2.2 lists the characteristics of our model Helios with a Keck system, and three groups of more powerful model transmitting systems (for $\lambda = 500 \text{ nm}$, $1 \mu\text{m}$, and $2 \mu\text{m}$). Within each group, there are systems with ranges from nearby (e.g. 100 pc with a 3 m transmitter) to part or most of the galaxy (e.g. 2 kpc with a 100 m transmitter). While these technically feasible systems would be very expensive by our current standards and set of priorities, they are certainly within reach for an older, more advanced civilization.

Note that the visible-band optical systems (rows *b-e* in Table 2.2) all produce signals that significantly exceed the sensitivities of the targeted and all-sky searches (both are ~ 100 photons m^{-2} in their sensitive regions; the all-sky search is expected to achieve a sensitivity of ~ 17 photons m^{-2} with a planned electronics upgrade). Because of the short wavelength (500 nm), these systems produce very narrow beams, which may need to be broadened if the aiming precision is insufficient.

The bottom two groups of systems in Table 2.2 (rows *f-n*) would operate in the infrared, out of the range of current pulse optical SETI instruments. Their superior range (through reduced extinction) demonstrate a key advantage of longer wavelengths and argue for infrared searches when the detector technology becomes less expensive.

A note of caution when interpreting the equations in this section and the systems in Table 2.2: Eq. 2.8 only

Table 2.2. Signal characteristics of model pulsed transmitters

	Transmitter			R	G	θ_b	D_b	F_{ext}	N_{pulse}	Received
	E_{pulse}	λ	D_T							$N_{\text{rec}} (\text{m}^{-2})$
a)	4.7 MJ	1 μm	10 m	100 pc	10^{15}	21 mas	2 AU	0.96	2×10^{25}	190
b)	100 MJ	500 nm	3 m	100 pc	4×10^{14}	34 mas	3 AU	0.90	3×10^{26}	670
c)	100 MJ	500 nm	10 m	300 pc	4×10^{15}	10 mas	3 AU	0.73	3×10^{26}	670
d)	100 MJ	500 nm	30 m	1000 pc	4×10^{16}	3 mas	3 AU	0.35	3×10^{26}	260
e)	100 MJ	500 nm	100 m	2000 pc	4×10^{17}	1 mas	2 AU	0.12	3×10^{26}	247
f)	100 MJ	1 μm	3 m	100 pc	10^{14}	69 mas	7 AU	0.96	5×10^{26}	360
g)	100 MJ	1 μm	10 m	300 pc	10^{15}	21 mas	6 AU	0.89	5×10^{26}	400
h)	100 MJ	1 μm	30 m	1000 pc	10^{16}	7 mas	7 AU	0.68	5×10^{26}	250
i)	100 MJ	1 μm	100 m	2000 pc	10^{17}	2 mas	4 AU	0.46	5×10^{26}	471
j)	100 MJ	3 μm	3 m	100 pc	10^{13}	206 mas	21 AU	0.99	2×10^{27}	120
k)	100 MJ	3 μm	10 m	300 pc	10^{14}	62 mas	19 AU	0.98	2×10^{27}	150
l)	100 MJ	3 μm	30 m	1000 pc	10^{15}	21 mas	21 AU	0.92	2×10^{27}	110
m)	100 MJ	3 μm	100 m	2000 pc	10^{16}	6 mas	12 AU	0.85	2×10^{27}	290
n)	100 MJ	3 μm	100 m	10000 pc	10^{16}	6 mas	62 AU	0.44	2×10^{27}	6

Note. — Signal characteristics of model pulsed transmitters for the “Helios with a Keck” system and three other groups of transmission systems. The first three columns are transmitter specifications: the energy per pulse (E_{pulse}), the transmitting wavelength (λ), and the diameter of the transmitting antenna (D_T). The remaining columns are the resulting system characteristics, as derived in §2.4: the target range (R), the transmitting antenna’s gain (G) and beam size in angle (θ_b) and extent at the target (D_b), the fraction of the signal not lost to interstellar extinction (F_{ext}), the number of photons in each pulse (N_{pulse}), and the number of photons that hit the target per square meter (N_{rec}).

poorly approximates the spatially-variable extinction of the Milky Way, which may be a serious limitation for optical SETI for short wavelengths and longer ranges. One should not assume that the above systems can deliver the specified number of photons to *all* targets in range. A more detailed model of the stellar and extinction distributions of the galaxy is needed to more accurately assess the capabilities of long range and short wavelength optical signaling.

2.5 A brief history of optical SETI

Science is a collaborative enterprise. The optical SETI experiments described herein build on the work of many physicists, astronomers, and technologists⁸. After the birth of optical SETI with Schwartz and Townes’ groundbreaking paper “Interstellar and Interplanetary Communication by Optical Masers” in 1961 [55], it was many years before serious observations began.

During the 1970s and 1980s a team of Russians led by V. F. Shvartsman and G. M. Beskin searched for non-Poisson statistics in photon arrival times in a project called MANIA. They observed about ~ 100 objects and did not find any significant brightness variations within the temporal range 2×10^{-8} to 10^{-3} s [57]. In the following decade Betz and Townes searched for continuous narrowband CO₂ laser signals from ~ 300 nearby stars using the 1.7 m telescope on Mt. Wilson [4].

During this time there were dozens of radio searches for intelligent civilizations and optical SETI was largely ignored. Stuart Kingsley joined Charlie Townes in advocating optical searches, and also organized three conferences and opened an optical SETI observatory. In 1997-1999, the SETI Institute sponsored workshops to reevaluate SETI search strategies and techniques generally. They recommended and co-funded a set of modest targeted optical searches, including the targeted search described in Chapter 4 and the optical search by Dan Werthimer’s group at UC Berkeley.

Werthimer was the first in optical SETI to use two photomultiplier tubes with outputs wired in a coincidence circuit. (The technique has long been used in high-energy physics.) Background events due to radioactivity, cosmic-rays, photon pile-up, etc. observed in one detector are unlikely to be seen in the other detector, provided the events are *uncorrelated* (see Chapter 3). This technique improves sensitivity since most background events are never recorded. Their two-detector operated for several years as a dedicated search at Leuschner Observatory [67].

The Berkeley group later teamed up with Shelley Wright (UC Santa Cruz), Remington Stone (UC Santa Cruz/Lick Observatory), and Frank Drake (SETI Institute) to develop a three-detector system that measures the statistics of “singles” (pulses in one of three detectors), “doubles” (coincident pulses in one of three detector pairs), and “triples” (coincident pulses in all three detectors) [68]. This technique improves the sensitivity of their two-detector systems and was deployed in targeted search instruments on Lick and Leuschner Observatories [59].

One noteworthy difference between the two- and three-detector systems described above and the Harvard targeted and all-sky experiments is that the former are *statistical*, while the latter are *event-driven*. Statistical experiments count the number of singles, doubles, and triples during a time interval and compare these numbers and their ratios with expected values based on Poisson statistics, target brightness, and event threshold. At the core of these experiments is a set of comparators (to detect signals above threshold) and counters (to count the singles, doubles, and triples), and the data read from these experiments are counts, which are read at regular intervals (typically 10 sec). Event-driven experiments, on the other hand, have a different design philosophy. The detection of signals of sufficient amplitude in each of two detectors constitutes an *event*, which initiates a cascade of actions including measurement of the waveform profiles (with four voltage levels in the targeted search and seven in all-sky survey) and time-tagging using a GPS clock. The data in these experiments are the event details (waveform profile, arrival time, etc.), from which statistics about coincident events (“doubles” in the statistical experiments) can be derived. At the core of these experiments is a larger set of comparators (for several voltage levels) and the circuitry to record the event details.

Returning to the history optical SETI, Geoff Marcy and Amy Reines used the Keck radial-velocity data (which was used to search for planets) to conduct a sensitive search for spurious lines in the spectra of ~ 600

⁸For a more complete history of SETI, see [18] and [61].

F, G, K, and M main-sequence stars. Their search covered 400-500 nm and would have detected laser lines down to the level of a 50 kW transmitter aimed at the Earth with a 10 m aperture at a range of 30 pc [51].

Stuart Kingsley [38] and Ragbir Bhathal [5] have optical SETI programs of their own, and Robert Lodder and his group at the University of Kentucky have looked for 800-3000 nm pulsed signals in the direction of supernovae [45].

Chapter 3

Backgrounds for Pulsed Optical Communication

This chapter was previously published as part of an article in the SPIE Proceedings of the Second International Conference on Optical SETI [32].

3.1 Astrophysics on short timescales

Breakthroughs in astrophysics are often the result of technological advances. As astronomers have broadened the parameter space in which they search (this first happened in wavelength, and then in spatial and temporal resolution), a wealth of new phenomena have presented themselves: pulsars, quasars, active galactic nuclei, just to name a few. Will astrophysics on milli-, micro-, and nanosecond timescales offer similar discoveries?

In preparation for the construction of the Very Large Telescope (the VLT – four 8 m telescopes working in tandem), D. Dravins has reviewed this problem in a paper [21] in *The Messenger*. He notes that using fast detectors, astronomers may learn about the rapid variability of astronomical objects. The scales that short-time techniques hope to probe are remarkably small, and certainly un-imageable – down to perhaps kilometer scales at galactic ranges. Dravins lists the following phenomena as candidates for milli-, or possibly microsecond timescale emission:

1. Plasma instabilities and fine structure in accretion flows onto white dwarfs and neutron stars.
2. Small-scale [magneto-]hydrodynamic instabilities in accretion disks around compact objects.
3. Radial oscillations in white dwarfs ($\approx 100\text{--}1000\text{ ms}$), and non-radial ones in neutron stars ($\leq 100\mu\text{s}$).
4. Optical emission from millisecond pulsars ($\leq 10\text{ ms}$).
5. Fine structure in the emission ('photon showers') from pulsars and other compact objects
6. Photo-hydrodynamic turbulence ('photon bubbles') in extremely luminous stars.
7. Stimulated emission from magnetic objects ('cosmic free-electron laser')
8. Non-equilibrium statistics (non-Bose-Einstein distributions) in sources far from thermodynamic equilibrium.

Note, however, that *none* of these phenomena is expected to produce nanosecond speed flashes of light.

The physical requirements for nanosecond speed optical flashes are quite restrictive. The transmitting region must be centimeters in size (or, if larger, it must be coherent), and yet able to emit an enormous power in the form of optical photons (greater than a solar luminosity in EIRP) in nanoseconds. We cannot think of a region in which such physical conditions exist.

We do however rest easy knowing that the discovery of such a novel phenomenon would be of tremendous astrophysical interest. Until we have evidence of such phenomena, we will have to concern ourselves with more pedestrian astrophysical and terrestrial backgrounds – the topics of the remainder of this paper.

3.2 Stellar photon pileup

One obvious candidate for nanosecond-speed optical pulses is the candidate star itself. This light is spatially unresolved from laser light, which presumably is produced on or around a planet orbiting the target star. On a nanosecond timescale, most stars are observed as a patten of single photons arriving individually; multiple photons rarely arrive during the same nanosecond. For example, a solar luminosity at 1000 ly ($m_V = 12$) delivers only 10^6 photons $\text{m}^{-2}\text{s}^{-1}$, or 1 *milli*-photon per nanosecond into a square meter aperture. Most of these photons are not converted to photoelectrons since photo-counting detectors have peak quantum efficiencies of $\sim 20\%$, and with an average of $\sim 10\%$. Because of this, and further losses in the optical system, it is more useful to speak in terms of the observable quantity: counts of photoelectrons per unit time.

The probability per unit time (“false alarm rate”) of detecting two or more photoelectrons during a time interval τ , with a photoelectron arrival rate r , (assuming that the arrival times are Poisson-distributed) is $r^2\tau = 20$ per second for $\tau = 2$ ns, and $r = 10^5$ Hz. More generally, the false alarm rate for n photoelectrons is

$$R = \frac{r^n \tau^{n-1} e^{-r\tau}}{(n-1)!}. \quad (3.1)$$

In the limit of $r\tau \ll 1$, the false alarm rate for n or more photoelectrons goes to

$$R = \frac{r^n \tau^{n-1}}{(n-1)!}. \quad (3.2)$$

Note that the quantity $r\tau$ is the expected number of photoelectrons in a time τ . The Poisson formula is interpreted as follows: One factor of r gives the arrival rate of single photoelectrons, the factor of $(r\tau)^{n-1}$ comes from the probability of $(n-1)$ additional photoelectrons arriving within τ , the factor of $e^{-r\tau}$ comes from the probability of all of the other photoelectrons *not* arriving in τ , and the factor of $(n-1)!^{-1}$ accounts for the rearrangement of the $(n-1)$ additional photoelectrons. Sometimes the false alarm rate, in the above limit, is quoted as $R = r^n \tau^{n-1}$, without the factor of $1/(n-1)!$. Although this factor is typically less important compared with factors of $r\tau$, it belongs there, and is important for careful calculations, particularly when $(n-1)$ is large.

This means, for example, that for a countrate of 2×10^4 Hz in each of the two photodetectors – which roughly corresponds to observing an $m_V = 0$ star, the brightest object we observe in the targeted search – the rate of detecting two photoelectrons in one photodetector during the same 2 ns is $r_1 = 8 \times 10^{-1}$ Hz. The rate of pileup of these two-photoelectron events in both detectors, by chance alone, is $r_2 = r_1^2 \tau \approx 1 \times 10^{-9}$ Hz, or once every 30 years. To get this false alarm rate up to, say, once per hour the countrate has to be greater than $\sim 10^6$ Hz.

There are at least two different strategies for dealing with stellar photon pileup. Our group sets a fixed threshold of three photoelectrons in the electronics that follow our hybrid avalanche photodiodes (for the targeted search; we use photomultiplier tubes in the all-sky survey). Other groups, such as the optical SETI programs at UC Berkeley and Lick Observatory, have variable thresholds for their multiple photomultiplier tubes. With this strategy, the thresholds are set for each object so as to keep the false alarm rate reasonably low, while maintaining high sensitivity to faint objects.

The above false alarm rate formula immediately demonstrates why two or more detectors, wired in coincidence, are used with most optical SETI experiments. In addition to reducing the rate of stellar pileup, this technique immunizes OSETI experiments to many detector pathologies. As we discuss in §5.1, photon-counting detectors occasionally produce large amplitude pulses due to corona discharge, ion feedback, cosmic-rays, etc., at a rate of, say, 1 per second. With just a single photodetector, the false alarm rate due to these internal detector pathologies is just that, one per second. With two photodetectors wired in coincidence, the false alarm rate is $r^2\tau \sim 10^{-9}$ per second, or about three per century. In practice, we find that the false alarm rate is closer to one per night of observation (~ 5 hours) because of correlations – some of the large

amplitude pulses produced in one detector are seen by the other.

Scattered zodiacal light and airglow are completely negligible when looking for nanosecond speed pulses with narrow field-of-view telescopes. A typical observing site has a nighttime “sky background” of 18–22 magnitudes per square arcsec. Thus, for the targeted search at the 1.5 m telescope in Harvard, MA (sky brightness of 19–20 mag per square arcsec) with a 15 arcsec-diameter field of view, the sky contributes about 13.5–14.5 magnitudes – two to three magnitudes dimmer than the faintest objects we observe. For the all-sky survey, each $1'.5 \times 1'.5$ pixel will see 9–10 magnitudes of sky brightness.

In fact, *daytime* optical SETI is possible. The daytime sky brightness has been measured [9] at 8×10^3 candelas per m^2 , or $\sim 3 \times 10^{-10} \text{ W/m}^2/\text{arcsec}^2$. In astronomical terms, this corresponds to ~ 7 magnitudes per square arcsecond. For a telescope with a rather narrow field of view, the countrates are large, but manageable; for the targeted search, the sky background is ~ 1.5 magnitudes – bright by astronomical standards, but nearly invisible to pulsed OSETI experiments (the false alarm rate is substantially less than once per hour). Our group has not yet observed during the day (the targeted search runs piggyback on existing nighttime observations), but may experiment with it soon. Experiments with larger fields of view could use neutral density filters to attenuate the sky background down to manageable levels, at the expense of sensitivity. Care should be taken to avoid pointing the telescope at the Sun with its 1.4 kW/m^2 (most of which would be focused onto the detectors).

3.3 Cosmic-rays and gamma-rays

Cosmic-rays—the most energetic particles in the known universe—produce optical photons and other particles when they interact with the atmosphere, which form a potential background for optical SETI experiments.¹ Under the broad definition, cosmic-ray primaries are made of individual atomic nuclei (most commonly), electrons, gamma-rays or neutrinos. Their energies range from less than 10^6 eV to greater than 10^{20} eV. The differential flux for these particles is strongly energy dependent:

$$\frac{dN}{dE} \sim E^{-\alpha}, \quad (3.3)$$

where $\alpha \approx 3$ for most of the energy range², meaning that for every factor of ten increase in energy, the flux of particles (which scales as N) goes down by a factor of 100. At an energy of 10^{11} eV, the flux on the Earth’s atmosphere is modest: about one particle per square meter per second. At 10^{16} and $10^{18.5}$ eV, the fluxes are down to one particle per square meter per year, and one particle per square kilometer per year, respectively.

Gamma-rays, although technically part of the cosmic-ray family, are typically lower in energy: gamma-rays in the 3×10^{11} to 10^{14} eV are considered “very high energy.”³ Like cosmic-rays, gamma-rays interact with the Earth’s atmosphere producing an electromagnetic cascade of particles, and a flash of Čerenkov light.

When a cosmic-ray (or gamma-ray) collides with the nucleus of an atom (usually oxygen or nitrogen) in the Earth’s upper atmosphere, the nucleus disintegrates into neutrons, protons, pions, kaons, hyperons, etc., and their antiparticles. These fragments are extremely energetic themselves, given the kinetic energy of the cosmic-ray; they too collide with atoms and produce even more particles. Many of these are unstable and decay (via the weak interaction); pions, for example, decay into muons and neutrinos, if charged, or into a pair of photons, if neutral. Other processes are also at work. Energetic positrons and electrons braking in the electric field of nuclei emit bremsstrahlung radiation (gamma-rays). Pair production generates positron-electron pairs (and positive-negative muon pairs to a lesser extent) out of the energy of neutral particles and gamma-rays. Many of these relativistic particles are also speeding: by exceeding the speed of light in air, they radiate Čerenkov radiation and slow down.

¹For classic and recent reviews of cosmic-rays, see Rossi [53] and Cronin [16], respectively. For gamma-rays, the paper by Catanese and Weekes [8] is relevant and useful.

²For primaries in the range $10^{12} \text{ eV} \leq E_{\text{pri}} \leq 10^{15} \text{ eV}$, the differential flux scales as $\alpha \approx 2.7$; for $E_{\text{pri}} \geq 10^{15} \text{ eV}$ it scales as $\alpha \approx 3.3$. The most energetic cosmic-rays observed to date have E_{pri} of order 10^{20} eV.

³The fact that we observe *charged* cosmic-rays, but not neutral gamma-rays above a certain energy threshold probably implies that the most energetic cosmic-rays are accelerated by very large, extended magnetic or electric fields.

The survivors of these processes (which are observed on the ground) are electrons, positrons, muons, neutrinos and photons. The charged particles and photons are both potential backgrounds for optical SETI experiments. We investigate them in greater detail below.

3.3.1 Muons

Most of the charged particles that survive to sea level are muons,⁴ with a mean energy of 2×10^9 eV. Their total flux (all energies) is given [53] approximately by

$$I(\phi) = I_\nu \cdot \cos^2\phi, \quad (3.4)$$

where ϕ is the zenith angle (muons arriving at angles close to the horizon are attenuated by more atmosphere), and $I_\nu = 8 \times 10^{-3} \text{ cm}^{-2} \text{ s}^{-2} \text{ sr}^{-2}$. These particles are essentially unimpeded by an observatory dome roof, or the 1.25 cm thick aluminum (a few g cm^{-2}) experimental enclosure in the targeted search. Muons pass through individual photodetectors at a rate of once every few seconds. The rate of two muons randomly striking the two detectors in the same nanosecond is therefore of order 10^{-9} per second.

It takes a lucky hit for a single muon to pass through *both* photodetectors. We can roughly calculate the angle-averaged rate for the targeted search as follows: assume that the detectors are 10 cm apart and that each have a 0.25 cm^2 cross-section; the rate of muons traversing both detectors is $\sim 10^{-5}$ per second, or once every ~ 25 hours (also assuming that the average flux is half the maximum). Although it is unlikely that a muon would trigger a false alarm in one night’s observations, this rate is significant for experiments that have observed for many thousands of hours, such as our targeted search. We have not attempted to correlate the zenith angle (which is a function only of the sky coordinates of the object being observed and the time) of the photodetectors for the residual background events (about one false alarm every eight hours of observation) to look for a $\cos^2\phi$ dependence yet, although this is certainly warranted⁵. This background can of course be completely eliminated by placing a scintillator and PMT in anti-coincidence with the two photodetectors.

Another way to have a false alarm is to capture a muon in an atom in the beamsplitter where it will subsequently decay into an electron and a neutrino. The energetic electron will then scintillate—the process of ionization of matter by an energetic charged particle and the subsequent photon emission that occurs as the excited molecules return to their ground states—in the beamsplitter glass and might be detected by both photodetectors. However, such an event would be exceedingly rare since the capture cross section for $\sim 10^9$ eV muons is small. It is also unlikely that a muon would be slowed down to energies where capture becomes more likely; a cosmic-ray muon dissipates $\sim 5 \times 10^6$ eV per g cm^{-2} , and the longest dimension in the beamsplitter (density of order g cm^{-3}) is a few cm. We have further reduced the possibility of this by replacing our cubical beamsplitter with the “thin slide” style beamsplitter in the targeted search.

3.3.2 Čerenkov radiation

As we mentioned above, Čerenkov radiation is formed when a particle exceeds the local speed of light. The radiation is beamed down in a narrow cone with an opening angle

$$\theta_C = \arccos\left(\frac{1}{\beta n}\right), \quad (3.5)$$

where $\beta = v/c$ and n is the index of refraction, and is emitted over a broad range of frequencies in proportion to $1/\lambda^2$ (i.e. blue Čerenkov photons are more common than red ones).

Fortunately for optical SETI the image of a cosmic-ray (or gamma-ray) induced Čerenkov pulse is too diffuse to be detected by the current experiments. A typical 10^{12} eV primary cosmic-ray does produce a short (5 ns duration) optical pulse with about 30 photons/ m^2 falling on the base of the narrow light cone (~ 150 m radius). But, the source appears diffuse – about 2° FWHM. Thus, the narrow field of view of the

⁴The atmosphere is $\sim 10^3 \text{ g cm}^{-2}$ thick, while high-energy photons have a typical “interaction length”—a fraction $1/e$ of particles remain after traversing this distance—of 30 g cm^{-2} .

⁵This statistical correlation has since been done. See §4.7.

targeted search telescope will observe only $\sim 2 \times 10^{-4}$ photons per flash, i.e. rarely one photon, and almost never two or more.

The rate of such events, as seen from an arbitrary point on the ground, is given by $\text{flux} \times A_{\text{footprint}} \times \Omega_{\text{image}} \approx 15$ per second for 10^{12} eV primaries. Scaling the above result (and using the fact that the photon fluence per flash is roughly proportional to E_{pri}), we find that a $10^{17.5}$ eV primary would deliver ~ 100 optical photons to the targeted search telescope; however such events happen about once every thousand years in an arbitrary part of the sky as viewed from an arbitrary point on the ground.⁶

One also has to worry about Čerenkov radiation produced by cosmic-ray muons (or from alpha-particle decays) passing through the beamsplitter glass. The number of Čerenkov photons in one of these pulses is a function of the energy of the relativistic particle, and the distance it traverses in the material:

$$\frac{d^2N}{dEdx} = 370 \sin^2 \theta_C(E) \text{ per eV-cm.} \quad (3.6)$$

For glass ($n = 1.5$, $\theta_C = 0.84$ rad), this means that about 500 “visible” photons (~ 1.5 eV average energy) are produced per muon per centimeter traveled. With the right geometry, the targeted search might be able to see such a flash. The probability of this was reduced though when we installed a lower volume beamsplitter.

Scintillation in the beamsplitter is also a potential source of pulsed light. We have calculated that, as long as the yield is less than $\sim 10^{-3}$ of NaI (a classic scintillator), the flux of optical photons is insufficient to trigger the targeted search.

3.4 Instrumental and terrestrial backgrounds

In our experiments to date, the dominant backgrounds are not astrophysical or atmospheric, but instrumental. We explore these, and others, below.

3.4.1 Photodetector problems

There are a host of potential problems with high-voltage photodetectors (radioactive decay in the PMT glass, ion feedback, scintillation of electron impacts from within). Corona discharge is the largest background in the targeted search though. This process occurs in high voltage environments when sharp points (e.g a dust particle, or a burr on metal) produce an extraordinarily high electric field. This field ionizes the gas between the sharp point and an electrode resulting in corona radiation (a short burst of optical photons) and crackling noise. This is the familiar hum heard around high-voltage power lines. Humidity tends to accentuate this highly non-linear phenomenon. It is also characterized by discharges clustered in time.

The hybrid-avalanche photodiodes (HAPDs) in the targeted search run “negative cathode” (that is, the anode is grounded) at a voltage of -7.5 kV. The common wisdom in the photodetector community is that the negative-cathode arrangement is prone to corona discharge, a tradeoff against its convenient output signal coupling.

As shown in Fig. 3.1, there is a marked systematic seasonal trend in the rate of coincident hits that is consistent with corona discharge. During the cold, dry months of fall, winter, and early spring (October-April), the data exhibits a good hit rate of 0.12 hits per hour of observation and a total hit rate of 0.50 hits per hour of observation.⁷ However, the hit rates are 30-40 times higher during the warmer and more humid summer months (May-September). Furthermore, we see a memory effect: observations following wet weather exhibit hit rates many times higher than the summer average, but drop back after 1-2 nights of clear weather. Opening the camera (which is normally kept tightly closed and flushed with dry nitrogen) for maintenance work similarly raises hit rates, but with a longer decay time constant (~ 15 days). These hits tend to be clustered in time with, say, 10 hits in 3 minutes followed by many quiet tens of minutes.

We believe that humidity promotes corona breakdown in one detector, which affects the other detector via electromagnetic (EMI) and optical coupling. To combat this problem we have added gas lines to the

⁶Detecting the Čerenkov radiation from such energetic cosmic-rays requires effective collecting areas measured in km^2 and a wide angular acceptance.

⁷“Good” hits are a subset of the coincident events that pass basic sanity checks. (After the publication of the paper based on this chapter [32], good hits were later renamed “events” and all hits were renamed “triggers.” See [30] and Chapter 4.)

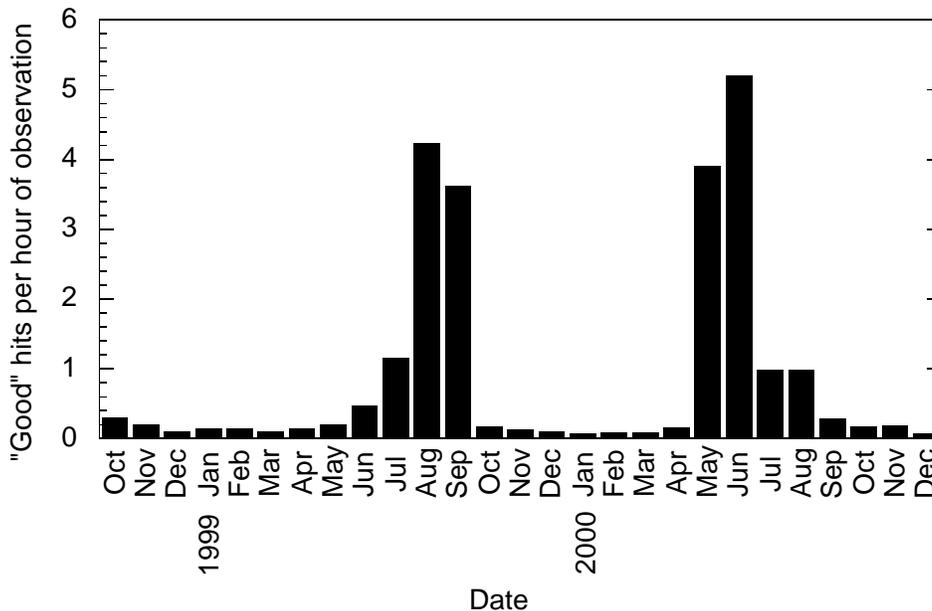


Figure 3.1. Here we show the humidity-induced seasonal trend in the “good” hit rate for the first two years of the targeted search. This is likely the result of corona discharge. (See Fig. 4.3 for the seasonal trend over six years.)

optical and electrical compartments, to keep them under a slight positive pressure of dry nitrogen, and we installed a glass entrance window. We also installed bakeout heaters (250 W total) to the aluminum exterior of the experiment to purge absorbed moisture. Most of these upgrades were completed during the summer of 2000 and the good hit rate appears to have gone down to manageable levels – less than 0.2 good hits per hour of observation.⁸ We believe that we largely mitigated the humidity problem, and that regular bakeouts can reduced it to levels such that no seasonal data needs to be excluded.

To further reduce our background rate, we are collaborated with Dave Wilkinson and colleagues at Princeton University who duplicated the targeted search instrument on their 0.9 m Cassegrain telescope in the Fitz-Randolph Observatory. This telescope will followed the Harvard targeted search telescope through its nightly observing programs for several years. Even with a coincidence rate of 5 good hits per hour, the rate of inter-observatory coincidence is once every 600 years, with a 1 ms time window⁹.

The all-sky survey uses multi-anode photomultiplier tubes that run at 900 V. At this lower voltage, we do not expect, nor have we observed, corona discharge as severe as in the targeted search.

3.4.2 “Cultural” backgrounds

The world is full of pulsed optical lights – sparks, lightning, automobile turn signals, disco lights, etc.; the question though is: are there any cultural phenomena that will deliver of ~ 100 or more optical photons into one of our telescopes during a nanosecond interval? Fortunately, most cultural backgrounds are either insufficiently bright on nanosecond timescales, or they couple poorly into the experiment, i.e. one would never point a telescope at them.

Lightning is of course a source of intense, pulsed light. Measurements [26] have shown that the flashes are $30 \mu\text{s}$ long on average, with structure on the single μs level, and perhaps even faster. However, OSETI researchers do not observe during *local* storms. And it is difficult to imagine lightning reflecting into a telescope from an overhead haze with sufficient intensity, while retaining the short time structure, that would trigger an OSETI experiment.

⁸The origin of this small residual background is unclear.

⁹This was subsequently implemented with a time stamping accuracy of $0.1 \mu\text{s}$ and proved quite effective at vetoing background events. See Chapter 4 for details.

Artificial satellites orbiting the Earth form a background of steady (or transient over a few milliseconds) light. Most of these satellites are small and reflect only modest amounts of sunlight; the Hubble Space Telescope, for example, appears as a magnitude 4.5 object. (Our targeted search program observes stars with $m_V = 0$ –12; the brightest star in the night sky, Sirius, has $m_V = -1.7$.) Satellites with larger surface areas are brighter still: the International Space station and Mir have $m_V = -2.8$ and -3.5 , respectively. The constellation of 66 Iridium low Earth orbit communications satellites are bright enough at times, $m_V = -8$, to be seen during the day. What about planets? They look approximately like the brightest spacecraft. The brightest two, Venus and Jupiter, have maximum brightnesses of -4.4 and -2.7 , respectively, during their peaks.

The question still remains though: will these bodies give a “false alarm” to optical SETI systems? Since they are constant sources of optical photons, we need to worry about pileup. Scaling the result that an $m_V = 0$ star delivers $\sim 2 \times 10^4$ photoelectrons per second in each of the two photodetectors in the targeted search, an Iridium satellite—8 magnitudes or a factor of 1500 brighter—would give countrates of 3×10^7 in each detector. When the latter passes directly through the 15 arcsec field of view of the targeted search (for a few milliseconds), the false alarm rate with 2 and 3 photoelectron thresholds would be 60 and 0.006 per *millisecond*, respectively. Note, however, that satellite and planetary orbits are well characterized and well documented; OSETI observers can simply avoid observing locations where satellites will flare.

NASA is experimenting with pulsed laser communication between Earth-orbiting satellites and the ground, and between deep-space satellites and the Earth [28]. Their conclusions are similar to those of optical SETI researchers: Beamed laser communication offers a low-power, low-mass, and high-bandwidth alternative to RF communication. The tradeoff, for both NASA and optical SETI applications, is that the transmitter has to be aimed very precisely. Consequently, it is unlikely that a narrow beam would be accidentally intercepted by an OSETI experiment. The transmitter and detector would both have to be pointed at each other to within a beam width (each having a probability of order 10^{-9}). On the other hand, laser pulses intentionally beamed from a satellite to an optical SETI experiment is an ideal *test* of the latter.

Could the blinking lights on an airplane cause a false alarm? To calculate this, let us assume that the light is 500 W and radiates isotropically. If the plane is flying at an altitude of 3,000 meters, then it has the same brightness as a solar luminosity 0.3 ly away ($m_V \approx -5$), i.e. somewhat dimmer than an Iridium flare. Although one cannot predict when airplanes will fly overhead (or look them up in a database, as one can for satellites), the probability that they would fly through the beam of the targeted search is quite small; only 15 arcseconds in diameter, the targeted search telescope observes less than $\sim 10^{-9}$ of the sky at any one time.

All of the other potential cultural backgrounds that we have dreamed up so far – local light pollution, electrical sparks, etc – would fail to trigger pulsed OSETI experiments because either (1) they are relatively low power sources of continuous radiation and are therefore insufficiently bright on nanosecond timescales to show multiphoton pileup, or (2) they are short and intense, but do not couple directly into the experiment.

Chapter 4

Targeted Optical SETI

This chapter was published previously as part of an article in The Astrophysical Journal [35].

Based upon the arguments above, and their elaborations (which evolved during a set of workshops sponsored by the SETI Institute in 1997–9 [24]), we designed and built a detector system for pulsed laser beacons. It saw first light on 1998 October 19 at Harvard, and ran continually until 2005 May 10; a second system began synchronized observations at Princeton on 2001 November 17.

4.1 Instrument design

The Harvard system rides piggyback on the CfA Digital Speedometer mounted on the 1.5-m Wyeth Reflector at the Oak Ridge Observatory in the town of Harvard, Massachusetts. The CfA Digital Speedometer supports several dozen research projects, mostly involving radial-velocity measurements of stars [41, 42]. Roughly half the light reflected off the entrance slit of the echelle spectrometer (about one third of the total light) is deflected into our photometer, as shown in Fig. 4.1.¹ Incoming light is re-imaged and passes through a beamsplitter onto two hybrid avalanche photodiodes² (Hamamatsu R7110U-07), whose outputs feed a pair of multi-level discriminators with levels corresponding to roughly 3, 6, 12, and 24 photoelectrons. By time stamping level crossings with a LeCroy MTD-135, we obtain approximate “waveforms” of incoming pulses to a precision of 0.6 ns.³ Coincident pulses seen in the two channels trigger the microcontroller to record the arrival time and waveform profiles. Arrival times are recorded twice—by a GPS clock (0.1 μ s precision and accuracy), and by a computer’s internal clock (1 ms precision, but only \sim 50 ms accuracy, as determined by comparing many GPS and computer time stamps). A “hot event” veto filters out a class of large amplitude, bipolarity signals that appear to be produced by breakdown events in the photodetectors. Pulse counters, threshold adjusting circuitry, and various controls and monitors allow us to test the apparatus to confirm its stability and proper operation. Fiber-coupled LEDs test the detectors and coincidence electronics before every observation. Figure 2 shows the complete photometer.

Each clear night the CfA Digital Speedometer observes typically 20–50 stars with integration times of 2–40 minutes. The observing sequence is determined by the conditions at the telescope and the priorities established in the monthly observing plans. Because of countrate limitations (6000 counts per second for the Digital Speedometer, typically a factor of two higher for the SETI instrument), bright objects are observed only when attenuated (by thin clouds or a neutral density filter) effectively eliminating false events due to photon pileup. Several of the projects involve monitoring variable stars, such as spectroscopic binaries and pulsating stars, so that the targets with the most observations and longest total integration times had

¹The Princeton system has full use of its smaller telescope, hence comparable light-gathering aperture; subsequent instrumentation is identical.

²HAPDs have the advantage of clean pulse height discrimination, at the price of increased corona discharge, as compared with traditional lower voltage multi-stage photomultiplier tubes.

³As configured, the LeCroy chip only timestamps the *last* upward and downward crossing for each level, thus the waveforms of more complicated shapes (e.g. double pulses) cannot be completely reconstructed.

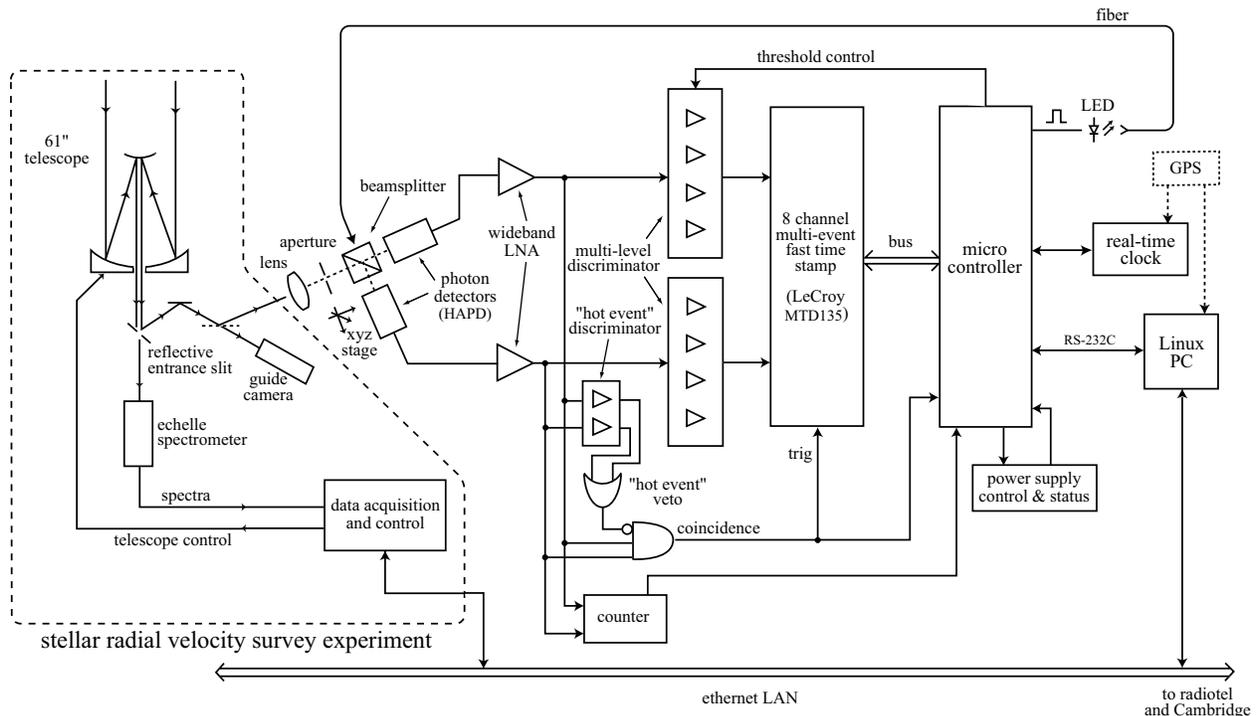


Figure 4.1. Block diagram of the Harvard targeted optical pulse detector. Unused light from the echelle spectrograph is imaged onto a pair of hybrid avalanche photodetectors, whose coincidence triggers fast time-stamping of waveform crossings through four preset levels.

been dominated by variable stars and others unlike the Sun. Soon after the Harvard instrument went into operation we established a new observing project designed to provide a large sample of promising SETI targets. The sample of more than 11,000 stars was drawn from the Hipparcos Catalogue and consists of all the main-sequence dwarfs between spectral types late A and early M with distances less than 100 pc and declinations between -20° and $+60^\circ$. Since November, 2001, the Princeton telescope has simultaneously pointed at the same stars on many observing nights. Coordination is achieved by automatically passing target RA/dec and other parameters to a Princeton computer, where volunteers point the telescope. During each observation, the diagnostic data, along with coincident pulse data, are sent to a PC and recorded in a log file at each observatory. After each night of observations, the log files are incorporated into a web-enabled database to facilitate analysis. We track the data through automated daily emails that 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 (total events, event rate, total observation time, etc.), observational summaries for individual objects, diagnostic data for particular observations, etc. Further details are available in Charles Coldwell's PhD thesis [11].

4.2 Sensitivity

We estimate the sensitivity of the Harvard instrument by following a light pulse through the entire system (Fig. 4.1). The optical path includes four reflections (each $\sim 85\%$ efficient), a lens ($\sim 92\%$), and a beamsplitter ($\sim 92\%$). One of the mirrors is the entrance slit to the echelle spectrometer, which reflects roughly 1/3 of the light into the OSETI instrument on average, depending on seeing. The beamsplitter sends half the light to each HAPD, which has a broad $\sim 20\%$ plateau in quantum efficiency (QE) for $\lambda = 450\text{--}650$ nm (and $\text{QE} > 1\%$ for $\lambda = 350\text{--}720$ nm). The signal must exceed the lowest threshold (three photoelectrons) in each detector



Figure 4.2. The Harvard targeted search photometer, with covers removed. Light enters from the rear of the righthand compartment, focused onto a 30 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 $\sim 50 \mu\text{V}$ negative pulses into 50Ω , 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 $25 \times 25 \times 60$ cm, and weighs 30 kg.

Table 4.1. Project parameters – targeted optical SETI at Harvard and Princeton

Telescopes
1.6 m Wyeth telescope in Harvard, Massachusetts—1/3 of light used for OSETI
0.9 m Fitz-Randolph telescope in Princeton, New Jersey—all light used for OSETI
Photometer
Beamsplit pair of hybrid avalanche photodiodes (350-720 nm response, peaked between 450-650 nm)
Pulse amplitude profiles recorded to 0.6 ns resolution
GPS-derived timestamping of events to 0.1 μ s at each observatory
Objects Observed
15,897 observations of 6176 stellar objects with Harvard instrument
1721 simultaneous observations of 1142 stellar objects with Princeton and Harvard instruments
Objects selected for radial velocity surveys—many FGK dwarfs
Sensitivity
100 photons m^{-2} in the photometers’ waveband and aperture in ≤ 5 ns (80 photons m^{-2} for Princeton)

during the discriminator’s averaging time (5 ns). Accounting for these factors, the Harvard instrument’s threshold sensitivity is 100 optical photons ($\lambda = 450\text{--}650$ nm) per square meter, arriving at the telescope in a group within 5 ns.

The Princeton instrument is identical, except that it receives all of the light from its 0.9 m primary mirror, and the light path includes only three mirrors. Multiplying the appropriate factors, the Princeton instrument is sensitive to signals of 80 optical photons ($\lambda = 450\text{--}650$ nm) per square meter arriving in a group within 5 ns. With both instruments simultaneously observing, Princeton can “veto” a Harvard signal, although the confidence in this veto depends on the signal intensity at Harvard and the observing conditions at the two sites.

4.3 Harvard observations

From October 1998 through November 2003, the targeted search with the Harvard instrument performed 15,897 observations of 6176 stars, for a total of 2378 hours of observation. Our target list is composed of objects being surveyed both for SETI and for other astrophysical interests. Two of the authors [of the original paper – David Latham and Robert Stefanik] are characterizing $\sim 11,000$ F, G, and K dwarfs (2079 observed at least once with the OSETI instrument) for possible observations by next generation targeted microwave SETI, seeking evidence of stellar companions that would prohibit planets in the habitable zone. The radial velocities of a sample of ~ 3000 nearby G dwarfs are being monitored to characterize the population of stellar companions with spectroscopic orbits. Other programs observe a variety of additional targets (very young stars, A dwarfs, and very old stars in the Solar neighborhood, among others).

In its five years of observations, we recorded 4746 “triggers,” i.e., instances when the lowest thresholds are simultaneously exceeded in both channels. Although all triggers are recorded, the reported “waveforms” are passed through a filter that enforces certain validity 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. The subset of triggers that pass this test are labeled *events*; to date, we have registered 1117 events. This filter is unlikely to exclude a genuine pulsed flash—the LED test flashes, which are done before every observation, pass this test with only rare exceptions.

Since the 1117 events are distributed impartially among 6176 objects (§4.6), we have confidence that the majority of the events arise from natural causes. Furthermore, in attempts to identify their source, we

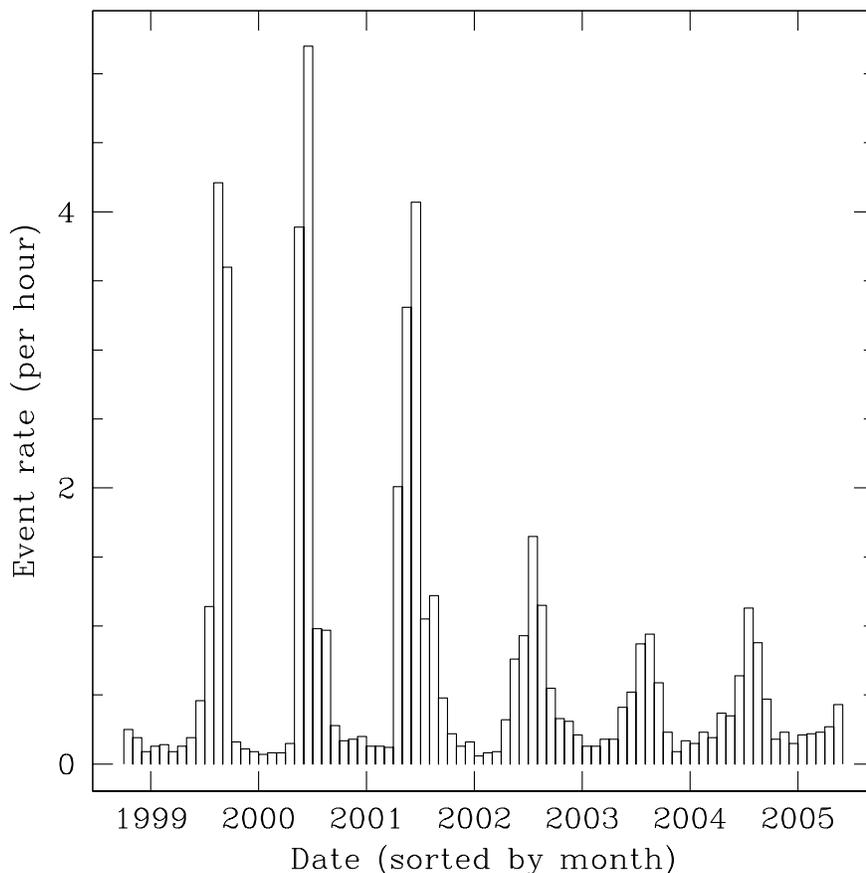


Figure 4.3. Seasonal variation in the event rate at Harvard. Corona breakdown substantially elevates the rate during the humid summer months. This plot includes all observations, regardless of quality.

logged events even during tests with the observatory dome closed. Clearly, instrumental effects contribute background events. In the analysis that follows, we attempt to remove the instrumental backgrounds from the Harvard data to look for residual events, possibly of extraterrestrial origin. We also examined the Harvard observations during which Princeton provided verification through simultaneous observations; with one possible exception, we found no events synchronously occurring at the two observatories (§ 4.10).

4.4 Seasonal variation

A histogram of the event rate by month (Fig. 4.3) reveals the largest source of background events, evidenced by a marked systematic seasonal trend in the event rate, apparently due to ambient humidity. During the dry months of fall, winter, and early spring (October–April), the data exhibit an event rate of about 0.16hr^{-1} and a trigger rate of about 0.5hr^{-1} . However, the event rates are 30–40 times higher during the warmer and more humid summer months (May–September), as shown. Furthermore, we see a memory effect: observations following wet weather exhibit event rates many times higher than the summer average, but recover after 1–2 nights of dry weather. Opening the camera (which is normally kept tightly closed and flushed with dry nitrogen) for maintenance work similarly raises event rates, but with a longer recovery time (~ 15 days). These events tend to cluster in time with, say, 10 events in 3 minutes followed by many tens of minutes of quiet. These symptoms all point to corona discharge, a high voltage breakdown characterized by radiofrequency and optical emission. We mitigated the humidity effect beginning in 2002 by adding dry nitrogen gas lines, bakeout heaters, and an entrance window, but these efforts reduced the summer event

Table 4.2. Summary of data from Harvard search

	DS _{all}	DS _{clean}	DS _{overlap}
Objects	6176	4730	1142
Observations	15897	11600	1721
Observations per object	2.1	2.5	1.5
Integration (hr)	2378	1721	244
Integration per object (hr)	0.39	0.37	0.21
Events	1117	274	130
Event Rate (hr ⁻¹)	0.47	0.16	0.53
Triggers	4746	1066	614
Trigger Rate (hr ⁻¹)	2.00	0.62	2.52

Note. — Summary statistics from the Harvard search for three data sets: DS_{all} consists of all observations made from Harvard. DS_{clean} is a subset of DS_{all} with certain high trigger rate nights removed (see § 4.5). DS_{overlap} is the subset of DS_{all} during which Princeton jointly observed (see § 4.10). Princeton contributed 429 events and 2327 triggers to DS_{overlap}.

rate only by a factor of two or three (see Fig. 4.3).

We examined, and excluded, the possibility that the seasonal variation was due to temperature-dependent gain in the HAPDs: the HAPD bias power supply (Power Technology PD-3) includes temperature compensation matched to the detector characteristics, and furthermore the observed pulse waveforms produced by the HAPDs during hot and cold weather are indistinguishable. In addition, the event rate shows no statistically significant correlation with temperature.

4.5 Data sets

Throughout this paper, we refer to the three data sets shown in Table 4.2: DS_{all} is the data set used above, and consists of all observations made from Harvard. DS_{clean} is a subset of DS_{all} from which nights with anomalous trigger rates have been removed. To be excluded from DS_{all}, a night’s observations must have a trigger rate greater than one per hour, and two or more events spread among two or more objects. Although this cut may seem arbitrary, in practice it cleanly removes nights with corona-polluted data. The data excluded from this set were scrutinized for clear extraterrestrial beacons (e.g., a pulse train of events for one object). DS_{overlap} is the subset of DS_{all} during which Princeton jointly observed (see § 4.10). Roughly 60% of the objects in these data sets are Sun-like stars (late F through early M).

4.6 Consistent with Poisson statistics

DS_{clean} includes 11,600 observations, summarized in Tables 4.2 and 4.3. Of the 4730 objects observed, 95% (4496) had no events at all. Note particularly that among the objects with events, those with more events were observed for longer durations (Tables 4.3 and 4.4) and more often (Table 4.5). The 274 events were distributed among 234 separate objects as shown in the “Observed” column of Table 4.3. The objects at the tail of this distribution are particularly interesting—20 objects with two events, four objects with three events, and four objects with four events. Do any of the objects show an extrastatistical number of events? Is there any evidence of periodicity? Is this distribution consistent with any model?

The events in DS_{clean} were modeled with a Monte Carlo simulation of the observations using Poisson statistics. We assumed that the event rate was constant (0.16 hr⁻¹, to generate the observed total of 274 events in 1721 hours)—as one would expect for a random physical process (e.g. radioactivity, cosmic-rays) unrelated to the telescope’s target—and calculated the average number of objects with 0, 1, 2, etc. events

Table 4.3. Harvard observations – distribution of events

Events	Objects		Observations		Integration Time (hr)		Events/hr
	Observed	Simulated Poisson	Total	Per Obj	Total	Per Obj	
0	4496	4495 ± 14	9518	2.1	1327	0.3	0
1	206	207 ± 14	1378	6.7	257	1.3	0.80
2	20	21 ± 4.4	418	21	75	3.8	0.53
3	4	4.9 ± 2.0	81	20	14	3.5	0.86
4	4	1.5 ± 1.1	205	51	48	12	0.33
5+	0	0.7 ± 0.8	0	0	0	0	0
Total:	274	4730	11600	2.5	1721	0.37	0.16

Note. — DS_{clean} sorted by events per object. Note that 95% of the objects (4496/4730) do not have any events. The expected number of objects—derived from a calculation based entirely on the event rate and the distribution of integration times—correlates well with the actual number of events. Also note that both the number of observations and integration time per object is much higher for multiple-event objects, as one would expect for randomly distributed events.

during 10,000 runs (with standard deviations in the averages giving the error bars), as shown in “Simulated Poisson” in Table 4.3. The Monte Carlo event distribution was verified by direct calculation of the Poisson probabilities for each object having 0, 1, 2, etc. events, and summing over objects. Since the total integration time per object varied substantially (Table 4.4), the calculation and Monte Carlo simulation used the actual distribution of observing times.

The observed event distribution in Table 4.3 appears consistent with the model of uniform background rate. The slight discrepancy between the observed and modeled event distributions—more observed objects with four events—can be plausibly explained by postulating that the observations during a few corona-plagued nights (with their highly variable event rates) were included in DS_{clean}. The objects with multiple events are nevertheless of interest, and are discussed further in §4.8.

4.7 Other sources of events?

What, then, caused these 274 Poisson distributed events? Low level corona is a plausible explanation, but several others warrant investigation as well.

Cosmic-ray muons (and other charged particles) are a potential source of events. Could a muon traveling by chance down the axis of the telescope produce enough beamed Čerenkov radiation to be detected by the OSETI instrument? Cosmic-ray muons have an average energy of 2 GeV and a flux of $I(\phi) = I_\nu \cos^2 \phi$, where ϕ is the zenith angle and $I_\nu = 8 \times 10^{-3} \text{ cm}^{-2} \text{ s}^{-2} \text{ sr}^{-2}$ [53]. The number of photons from a single muon is a function of its energy (with a threshold energy of 4.3 GeV), and its pathlength: $d^2N/dE dx = 370 \sin^2 \theta_C(E) \text{ eV}^{-1} \text{ cm}^{-1}$, where $\theta_C = (\beta n)^{-1}$, $\beta = v_\mu/c$, and n is the index of refraction. At sea level, this is ~ 10 optical (2 eV) photons m^{-1} . Although the light cone’s opening angle is small ($\theta_C(E) < 1.4^\circ$ for air), most of the photons either miss the telescope, or don’t couple into the instrument, even for muons traveling down the telescope axis. (See Chapter 3.)

Fig. 4.4 shows the probability density of events and observations as a function of telescope zenith angle for a restricted set of DS_{clean} with a particularly low event rate (9400 observations of 3928 objects with an event rate of 0.15 hr^{-1}). If Čerenkov photons were a significant background, we would expect to see excess probability for events at low zenith angles, which we do not. A histogram derived from triggers (instead of events) from the same dataset (not shown) is qualitatively similar. Thus we conclude, from both calculation and observation, that Čerenkov flashes do not contribute events.

Another potential source of the residual events is energetic muons traveling through both HAPDs. Such muons would have energy sufficient to eject electrons from the photocathodes, which would be amplified and detected. However, it would require a lucky hit for a single muon to pass through *both* photodetectors: Estimating the relevant detector cross-sections to be 0.25 cm^2 , a simple geometric calculation shows that muons traverse both detectors at a rate of $\sim 10^{-5} \text{ s}^{-1}$, or once every ~ 25 hours. We calculated the “detector

Table 4.4. Distribution of integration times

Integration (min)	Objects	Events			
		0	1	2	3+
0-1	36	36	0	0	0
2-3	1207	1203	4	0	0
4-7	1048	1027	21	0	0
8-15	876	845	29	2	0
16-31	821	783	36	1	1
32-63	410	355	53	1	1
64-127	211	172	33	6	0
128-255	83	61	19	3	0
256-511	27	14	6	4	3
512+	11	0	5	3	3

Note. — Distribution of integration times in DS_{clean} . Note that on average the objects with multiple events were observed for longer than those objects with zero or one event.

Table 4.5. Distribution of observations

Observations	Objects	Events
1	3223	65
2-3	767	51
4-7	437	55
8-15	239	46
16-31	42	19
32-63	19	35
64+	3	3

Note. — Summary of the distributions of events and objects from observations in DS_{clean} .

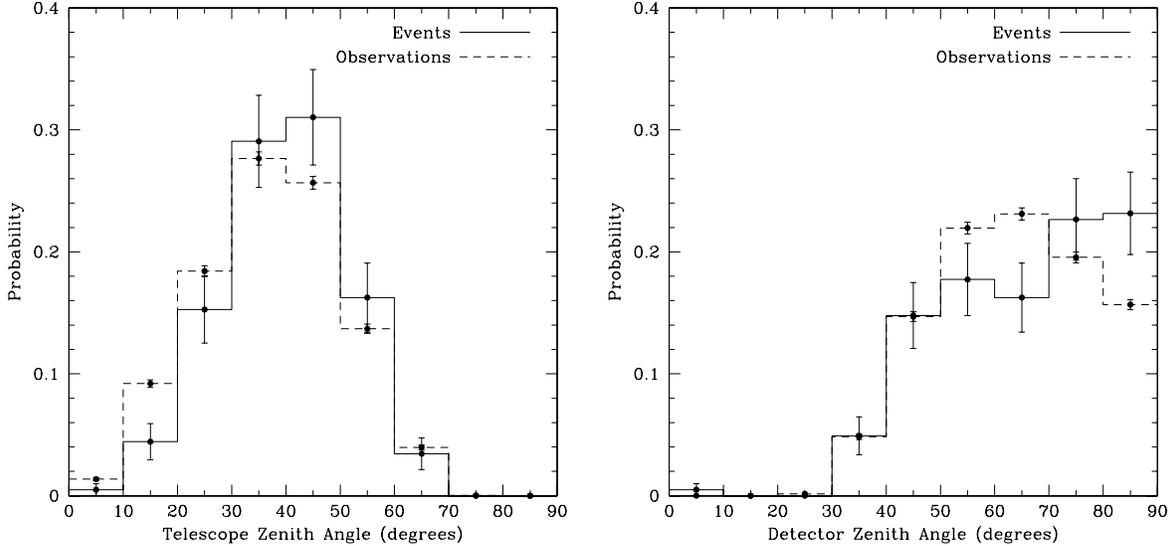


Figure 4.4. Histograms of the probability of events at Harvard for a restricted set of DS_{clean} as a function of telescope zenith angle (left) and “detector zenith angle” (right; the angle between the vector connecting the two detectors and the zenith). If a significant fraction of the events in this data set were due to Čerenkov radiation from muons traveling down the line of sight of the telescope (on the left), or were due to muons passing through both detectors (on the right), then the probability of events would be concentrated at small telescope zenith angles.

zenith angle” (the angle between the vector connecting the two detectors and the zenith) for each observation to test for the excess events that would be expected at small zenith angles, owing to the $\cos^2 \phi$ dependence of muon flux. Fig. 4.4 shows the probabilities of events and observations as a function of the detector zenith angle for the same data set used for Fig. 4.4. We see no evidence of excess events due to this mechanism. Using triggers instead of events in Fig. 4.4 leads to a qualitatively similar histogram.

Čerenkov radiation from muons passing through the 25-mm-cube glass beamsplitter is another potential source of events. For glass ($n = 1.5$, $\theta_C = 0.84$ rad), ~ 500 visible-wavelength photons are produced per muon per cm traveled. It would require a lucky hit to deliver photons from the beamsplitter to both detectors, but it seems plausible that this scenario could produce events. Scintillation in the beamsplitter glass from muons or radioactive decay products is unlikely to trigger an event due to the small size of the beamsplitter and the low scintillation yield of glass. Muon capture in the beamsplitter (where it would energetically decay) is unlikely given the low capture cross-section and small beamsplitter size. To test these beamsplitter scenarios, however, we temporarily replaced the cubical beamsplitter with a thin (1 mm) coated plate beamsplitter. During several nights of tests, we observed no statistically significant change in event rate.

The Čerenkov radiation and scintillation scenarios described above, as well as other potential backgrounds, are discussed in Chapter 3.

4.8 Interesting objects and reobservations

Tables 4.6 and 4.7 show the objects that have two or more events in DS_{clean} , their observing and event/trigger statistics, and astronomical information. We examined the observational histories of these objects for indications that the events from one or more of them was due to intentional extraterrestrial communication—an extrastatistical number of events, a clustering of an object’s events in one night’s observations, or simultaneous event detection at Harvard and Princeton. Objects with any of these characteristics were concurrently reobserved by Harvard and Princeton (“Reobservations” in Table 4.6). These reobservations are not included in the three datasets in Table 4.2.

Note that many of the objects in Tables 4.6 and 4.7 (particularly those with three and four events) were observed often and have long total observation times, typically because they are suspected short period binaries or are variable on short time scales. These objects are therefore deemed less likely hosts for intelligent

Table 4.6. Interesting objects in DS_{clean} – observing statistics

Object	Observations in DS _{clean}							Reobservations			
	Evt	Trig	Obsv	Hrs	Rate	Multiple Evt/Trig	Veto	Hrs	Evt	Trig	Obsv
HD 14535	4	7	52	13.3	0.30	—	0/1	—	—	—	—
ϕ CYG	4	10	59	7.6	0.53	—	—	—	—	—	—
RZ Cnc	4	7	54	15.6	0.26	—	0/0	—	—	—	—
SSLAC	4	8	40	11.7	0.34	—	—	—	—	—	—
α ORI	3	16	40	4.4	0.68	2/3 on 2002 Nov 8	—	1.0	1	4	3
DU Leo	3	8	27	8.3	0.36	3/7 on 1999 May 31	—	2.2	0	0	7
HD 220077	3	3	9	0.3	8.82	3/3 on 2000 Nov 4	—	1.7	0	0	6
LSR2-1471	3	3	5	1.0	3.11	—	—	—	—	—	—
BD+18 2930	2	2	1	0.3	6.00	2/2 on 1999 Feb 15	—	1.7	0	0	5
M67 1221	2	2	8	2.0	1.01	—	—	—	—	—	—
α Equ	2	4	14	1.2	1.70	—	0/2	—	—	—	—
BD+61 1045	2	2	6	1.6	1.23	—	—	—	—	—	—
Capella	2	8	37	5.4	0.37	—	—	—	—	—	—
EU Del	2	9	72	7.1	0.28	—	1/2	—	—	—	—
G65-43	2	2	7	1.9	1.05	—	—	—	—	—	—
HD 32306	2	3	8	2.5	0.80	—	0/0	—	—	—	—
HD 57769	2	4	18	4.4	0.45	—	—	—	—	—	—
HD 72746	2	3	4	1.2	1.62	2/3 on 1998 Dec 27	0/0	3.0	1	2	9
HD 86579	2	4	8	2.6	0.75	—	0/1	—	—	—	—
HD 94292	2	8	26	5.4	0.37	—	1/5	—	—	—	—
ξ UMa	2	7	58	10.6	0.20	—	—	—	—	—	—
HD 18884	2	5	3	0.1	14.49	—	—	—	—	—	—
HD 40084	2	4	13	3.1	0.64	—	0/0	—	—	—	—
HIP 14420	2	2	1	0.1	14.97	2/2 on 2000 Dec 27	—	2.5	1	2	10
RT Lac	2	7	42	12.7	0.16	—	0/1	—	—	—	—
Serge 3151	2	3	3	0.9	2.26	—	1/1	—	—	—	—
TV Psc	2	4	38	2.0	1.01	—	0/0	—	—	—	—
UU Her	2	7	51	10.5	0.19	—	1/2	—	—	—	—

Note. — Objects with two or more events in DS_{clean} listed with their observing statistics (number of events, triggers, observations, hours of cumulative observation, event rate in hr⁻¹), nights with multiple events/triggers from one object (a dash indicates all events occurred on separate nights), and the number of events/triggers that were “vetoes” (not observed) by Princeton during concurrent observations (a dash indicates no concurrent observations in DS_{clean}). Statistics from the Harvard instrument for joint Harvard/Princeton reobservations are listed in the four rightmost columns (which occurred after November 2003 and are not in DS_{all}; a dash indicates no reobservations). No simultaneous Harvard/Princeton triggers were recorded during the reobservations. Table 4.7 lists coordinates and descriptions for these objects.

Table 4.7. Interesting objects in DS_{clean} – astronomical information

Name	Other Name	R.A.	Dec.	V _M	Parallax	Description
HD 14535	HIP 11098	2 22 53	+57 14 43	8	0.37 [0.88]	A2Ia—supergiant, variable
ϕ Cyg	HD 185734	19 39 23	+30 09 12	5	13.00 [0.59]	K0III—binary, 433 day period
RZ Cnc	HD 73343	8 39 09	+31 47 44	9	3.25 [1.56]	cool Algol binary, 21 day period
SS Lac	HIP 108981	22 04 42	+46 25 38	10	1.13 [1.39]	triple, 14.4 and 678 day periods
α Ori	HD 39801	5 55 10	+7 24 25	1	7.73 [1.64]	M2Ib—supergiant, variable
DU Leo	HD 84207	9 44 11	+25 21 11	10	—	G0V—binary, 1.37 day period
HD 220077	HIP 115279	23 20 53	+16 42 39	9	13.07 [1.51]	F7V—visual binary, 0.''23 sep.
LSR2-1471	GSC 03600-00123	21 09 53	+50 49 18	11	—	G0V—binary, 210 day period
BD+18 2930	GSC 03598-00615	14 46 57	+18 18 00	9	—	G8V
M67 1221	GSC 01481-00366	8 51 44	+11 44 26	11	—	K1III—binary, 6394 day period
α Equ	HD 202448	21 15 49	+5 14 52	4	17.51 [0.89]	G0III & A—binary, 99 day per.
BD+61 1045	HD 70050	8 23 16	+61 27 38	8	3.88 [0.95]	G2V—binary, 14.35 day period
Capella	HD 34029	5 16 41	+45 59 53	0	77.29 [0.89]	G5III—binary, 104 day period
EU Del	HD 196610	20 37 55	+18 16 07	6	9.16 [0.99]	M6III—semi-regular variable
G65-43	HIP 69893A	14 18 12	+12 44 29	11	14.24 [2.92]	K3V—binary, 4837 day period
HD 32306	HIP 23422	5 02 01	-5 30 04	7	8.40 [0.84]	F5V—binary, 794 day period
HD 57769	HIP 35919	7 24 17	+36 18 39	7	8.01 [1.18]	F5V—triple, 1.5 day inner per.
HD 72746	HIP 42037	8 34 09	-9 57 10	8	11.22 [1.11]	F2V—visual binary, 0.''21 sep.
HD 86579	HIP 48963	9 59 19	-3 04 30	7	7.57 [1.02]	F5V—binary, 2729 day period
HD 94292	HD 53212	10 53 02	+4 57 43	8	13.02 [0.92]	G5V—double-lined binary
ξ UMa	HD 98230/98231	11 18 11	+31 31 45	4	—	F and G dwarf—quadruple
HD 18884	HIP 14135	3 02 17	+4 05 23	3	14.82 [0.83]	M2III—low level variable
HD 40084	HIP 28343	5 59 22	+49 55 28	6	2.90 [0.79]	G5III—binary, 219 day period
HIP 14420	HD 232747	3 06 11	+51 06 06	10	14.64 [1.79]	K0V
RT Lac	HD 209318	22 01 31	+43 53 26	9	5.19 [1.05]	RS CVn binary, 5 day period
Serge 3151	GSC 03598-00615	21 23 52	+49 07 37	11	—	G0V—spectroscopic binary
TV Psc	HD 2411	0 28 03	+17 53 35	5	6.65 [0.78]	M3III—semi-regular variable
UU Her	HIP 81272	16 35 57	+37 58 02	8	-0.15 [0.91]	F2Ib—semi-regular variable

Note. — Right ascension (R.A. in hours, minutes, seconds; J2000), declination (dec. in degrees, minutes, seconds), visual magnitude (V_M), parallax (mas, with uncertainties in brackets), and spectral types for the objects in Table 4.6. Information was not available where dashes are present.

Table 4.8. Observations of HD 220077

Date	DS _{all}	DS _{clean}	DS _{overlap}	HD 220077				Other Objects				
				Obsvs	Min	Evt	Trig	Objects	Obsvs	Min	Evt	Trig
2000 Nov 4	✓	✓		5	10	3	3	10	11	36	0	2
2000 Nov 18	✓	✓		2	6	0	0	44	49	358	2	3
2001 Jan 26	✓	✓		1	2	0	0	52	57	399	0	0
2002 Jul 7	✓			1	4	0	0	19	23	255	2	14
2002 Sep 13	✓			1	3	0	0	33	34	365	2	10
2002 Oct 21	✓			1	4	0	0	45	51	410	1	11
2003 Jan 15	✓	✓		1	3	0	0	22	29	329	0	0
2003 Jul 16	✓			1	4	1	1	19	23	284	5	17
2003 Aug 21	✓		✓	1	8	0	0	28	29	205	5	27
2003 Sep 30	✓			1	8	0	0	29	37	282	6	14
2004 Feb 12	–	–	–	4	61	0	0	38	45	408	1	1
2004 Feb 14	–	–	–	2	38	0	0	20	30	413	2	7

Note. — Observations of HD 220077. Columns detail the number of observations (Obsvs), the observation duration in minutes, and the number of events and triggers for HD 220077 and other objects observed that night. The datasets that each observation belongs to are checked. Observations on 2004 February 12 and 14 were reobservations with Princeton after the closing date of the three listed datasets. Note that HD 220077 had three events, and no non-event triggers, on 2000 November 4.

civilizations. Although many Sun-like stars were observed ($\sim 60\%$ of the objects and observation time in DS_{clean}), few are represented in Tables 4.6 and 4.7 because they had less observing time per object; that is, Tables 4.6 and 4.7 select for objects with long total observing times, which tend not to be Sun-like. Nevertheless, given the consequences of a confirmed nanosecond pulse, a careful analysis is warranted.

HD 220077

The three triggers (all events) recorded by HD 220077 on 2000 November 4 warrant the greatest attention.⁴ The three events were recorded during 10 min spread over 5 observations. The experiment ran for only 46 min that night, and none of the 10 other objects registered an event (although two triggered), as shown in Fig. 4.5. Although this night was the first time the experiment had run in 25 days, and it did not run during the five nights following, all diagnostic data (count rates, temperature, weather, etc.) appear normal. The event rates during autumn of 2000 were relatively low ($\sim 0.15\text{--}0.20\text{ hr}^{-1}$).

The time differences between successive events ($\tau_1 = 914\text{ s}$ and $\tau_2 = 289\text{ s}$) do not appear to be part of a (perhaps incompletely sampled) regular pattern: $\tau_1/\tau_2 \neq n/m$ for small integers n and m (which we will call the “Rational Period Test”), even when allowing for the $\pm 1\text{ s}$ clock accuracy in 1999 (before the GPS clock was added).

As shown in Table 4.8, HD 220077 was observed 15 times over 50 min in DS_{all}. One additional event was recorded on 2003 July 15 (in DS_{all}, but not DS_{clean}), however the trigger and event rates were both elevated that night. This object was jointly reobserved by Harvard and Princeton for 99 min and no events or triggers were seen at Harvard.

In general, one can calculate the Poisson probability of recording no events in reobservations of a given duration, assuming that an event rate from previous observations should apply. We calculate two such “reobservation probabilities,” p_{r_1} and p_{r_2} , using the event rates for HD 220077 from 2000 November 4

⁴We use Universal Time (UT), and star names from the CfA Digital Speedometer surveys, which in some cases are not in common use. Other names for objects (such as Henry Draper or Hipparcos numbers) and celestial coordinates are given in Table 4.7.

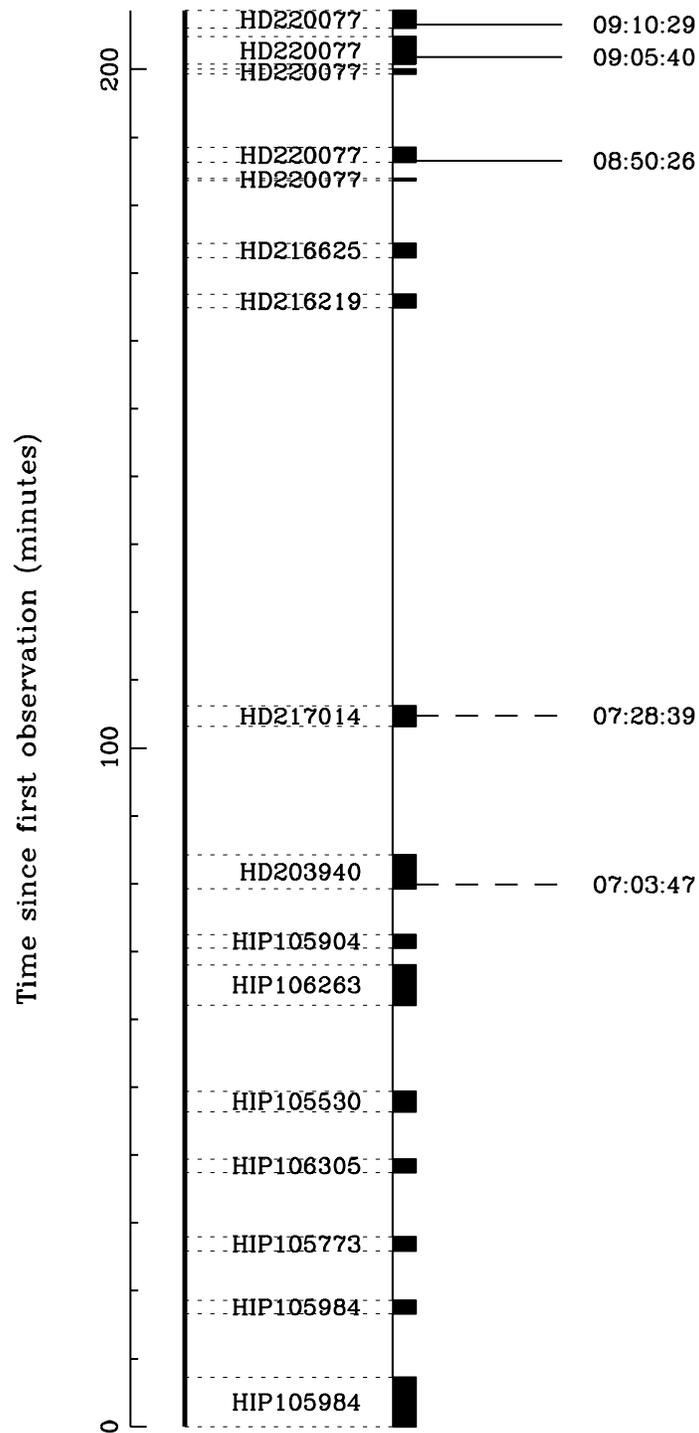


Figure 4.5. Observations on the night of 2000 November 4 beginning at 05:43:53 UT at Harvard. The boxes indicate observation intervals of objects listed below. Events (solid lines) and non-event triggers (dashed lines) are shown above the observations with timestamps labeled. Note that all three events occurred during observations of HD 220077; two triggers that did not meet event criteria were recorded on other objects (HD 203940 and HD 217014). All observations are in DS_{clean} , but not in DS_{overlap} . (These observations are atypically short and sparse; see Fig. 4.6 for a typical night's observations.)

($r_1 = 0.33 \text{ min}^{-1}$), and for all observations of HD 220077 in DS_{clean} ($r_2 = 0.14 \text{ min}^{-1}$), respectively, and the reobservation times listed in Table 4.6. Although instructive, $p_{r_1} \sim 10^{-14}$ and $p_{r_2} \sim 10^{-6}$ are probably unrealistically low since an event rate is poorly defined by so few events — a constant rate may not even be an accurate characterization of a process with such limited statistics. Nevertheless, it remains unlikely that a natural or artificial source would produce three events in ten minutes, and then no events for nearly ten times as long. We conclude that the events in question were probably a statistical fluctuation of background processes. The strength of this conclusion is limited, however, given the modest time spent reobserving. Additional reobservations (for, say, tens of hours) could test this conclusion.

HD 220077 is an F7V dwarf with a visual companion 0.25 mag fainter at a separation of $0''23$. The 38 CfA radial velocities for the composite light of both stars show a hint of a slow drift over the observed span of 4147 days, which supports the idea that the two stars are a physical pair in orbit (for which their separation would be $\geq 17.6 \text{ AU}$).

4.8.1 Other objects

No other object in Tables 4.6 and 4.7 displayed the extrastatistical properties of HD 220077. Several objects had long total observation times and are probably explained as statistical fluctuations in a constant background of Poisson-distributed events.

The six objects with multiple events in one night were deemed the most interesting, and were reobserved for one to three hours. Although some events and triggers were recorded at Harvard during the joint reobservations, none of them were simultaneously seen at Princeton.

Of note in the reobserved group is DU Leo, which recorded 3 events and 7 triggers in 40 minutes over two observations on 1999 May 30. Although this night was included in DS_{clean} , it is worth noting that the 1999 season of corona discharge, with its attendant high event and trigger rates (Fig. 4.3), had begun a few days earlier. We find no set of triggers from DU Leo (or any other object) that passes the Rational Period Test. Taking $r_1 = 0.075 \text{ min}^{-1}$ and $r_2 = 0.006 \text{ min}^{-1}$, the reobservation probabilities for this object are $p_{r_1} \sim 10^{-4}$ and $p_{r_2} \sim 0.5$.

The properties of many of the objects in Tables 4.6 and 4.7 (supergiants, multiple systems) make them less likely sites for Earth-like life. None of them are members of the Habitable Catalog [63], a catalog of $\sim 17,000$ potentially habitable stars from the $\sim 120,000$ star Hipparcos Catalogue [49].

4.9 Conclusions from Harvard observations

Based on the Harvard observations and reobservations, our conclusion is this: Given the low event rates and corresponding small-number statistics, we have found no evidence of clustering or periodicity from any candidate star, and the events are distributed impartially among the targets. There is additionally no correlation of event rate with stellar magnitude, confirming the conclusion that Poisson doubly-coincident “accidentals” do not contribute candidate events at ordinary single-photon count rates. Reobservations of those objects with multiple events in one night did not reveal sources of optical flashes. 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 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 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. It was this motivation that led to the construction of the Princeton experiment in 2001, and to the two years of joint observations, which are discussed below.

4.10 Synchronized observations with Princeton

Given our current background level of roughly one event per night of observation with the Harvard instrument, a single optical pulse from an extraterrestrial civilization would likely be dismissed as a background event. To attract attention, the signal would have to consist of a succession of 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 pulsed beacon.

To address this problem, we duplicated the detector system at the 0.9 m Cassegrain at the Fitz-Randolph Observatory in Princeton, New Jersey. Since November 2001, this telescope has followed the Harvard telescope through its nightly observing programs, synchronized via the internet. Given that the baseline between observatories is $L/c = T \approx 1.6$ ms of light-time, the absolute timing precision of $0.1 \mu\text{s}$ permits us not only to identify approximate coincidences; it further defines an error band in the sky whose width is of order $\Delta\theta \approx \Delta\tau/T = 12$ arcsec. This is comparable to the observed target field, set by the focal plane 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 an event rate $r_e = 1 \text{ hr}^{-1}$ at each observatory, and let us require that each candidate event pair (between the two observatories) be within a broad time window of, say, $\Delta T = 1$ ms to be considered a confirmed detection (recall that our GPS-derived timing accuracy is in fact 4 orders of magnitude better). Then the combined background rate due to “pileup” is $r_{\text{both}} = r_e^2 \Delta T = 3 \times 10^{-7} \text{ events hr}^{-1}$, or 1 event every 3 million observing hours. 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.

While the Princeton observatory provides excellent positive confirmation (simultaneous events would be believed with high confidence), we have less confidence in negative confirmation (“vetos”), particularly of low amplitude events. Sensitivity varies at the two sites due to nonuniform photodetector gain, electronic gain, and environmental factors (e.g., thin clouds at one observatory). These caveats apply most strongly to low amplitude signals; large pulses observed at Harvard should also be seen at Princeton.⁵

Table 4.2 summarizes the Princeton observations that coincided with 1721 Harvard observations of 1142 objects totaling 244 hours of observation ($\text{DS}_{\text{overlap}}$). During these observations, Princeton recorded 2327 triggers and 429 events, while Harvard recorded 614 triggers and 130 events (106 triggers and 17 events in DS_{clean}). The somewhat lower signal thresholds on the Princeton instrument, as well as higher corona rates, may account for the higher trigger and event rates at Princeton.

4.10.1 HIP 107395

During synchronized observations from 2001 November 17 onward, only one pair of triggers was recorded with arrival times that are consistent with an extraterrestrial or astrophysical optical pulse arriving at the geographically separated observatories. On 2003 September 17, during a joint observation of HIP 107395, Harvard recorded a trigger at 06:52:16.944 UT and Princeton recorded one at 06:52:16.943 UT (computer clock times). Unfortunately, the GPS clock at Princeton (with $0.1 \mu\text{s}$ accuracy) was not working for a few months around this date. Thus, unambiguous identification of an astrophysical or extraterrestrial pulse from HIP 107395 is not possible. Since the computer clock times are only accurate to ~ 50 ms, there is roughly a 2% chance that the triggers actually occurred within ± 1 ms of each other; the millisecond alignment may have been the work of chance. Although there are several other reasons to dismiss these as background triggers serendipitously recorded 1 ms apart (described below), it is worth noting that this is the only trigger pair whose arrival times are consistent with a single pulse arriving at both observatories. No other trigger pair arrived with a time separation of less than 0.3 seconds.

As shown in Fig. 4.6, the trigger rates at Harvard and Princeton were both elevated on 2003 September 17. Harvard recorded 23 triggers and 2 events during 32 observations of 25 objects over 6.1 observing hours that night, while Princeton recorded 315 triggers and 64 events during 15 observations of 15 objects over 3.5 observing hours.⁶ During the 20 minute observation of HIP 107395, Harvard recorded 5 triggers and 0 events. Princeton observed HIP 107395 for 26.4 minutes (completely covering the Harvard observation) and recorded 32 triggers and 10 events (8 triggers occurred before Harvard was observing).

The probability of any pair of triggers randomly occurring within ± 1 ms during the observation of HIP 107395 can be estimated as $p_{1\text{ms}} = r_H r_P \tau_{\text{coinc}} T_{\text{obsv}} = 3 \times 10^{-5}$, where $r_H = 15 \text{ hr}^{-1}$ and $r_P = 72 \text{ hr}^{-1}$

⁵The effect of deadtime following corona-induced triggers is insignificant—even under poor conditions (~ 50 corona triggers per hour) the probability of missing a confirming event is $\leq 10^{-3}$

⁶These anomalously high rates are consistent with corona discharge; all observations on 2003 September 17 from Harvard were thus excluded from DS_{clean} by the algorithm described in §4.5.

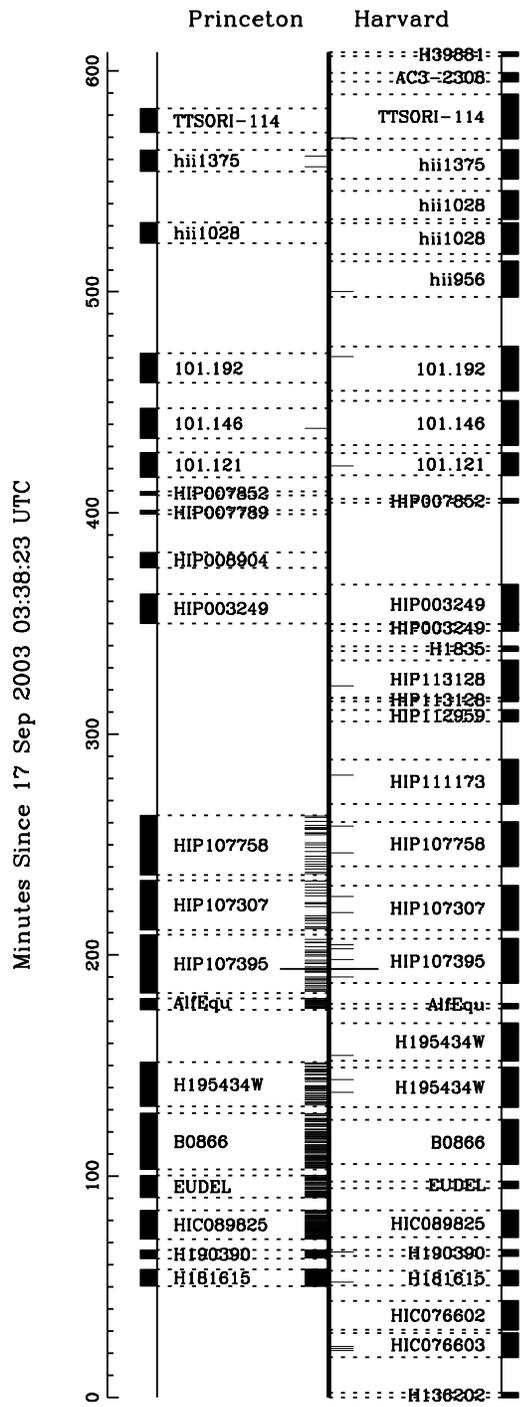


Figure 4.6. Observational diagram for the night of 2003 September 17. Observations at Harvard (right) and Princeton (left), with time increasing upward. Observations are designated by solid boxes and dashed lines to the central vertical axis. Object names are written between the dashed lines. Triggers are shown as short horizontal lines coming out of the central vertical axis (Harvard to the right, Princeton to the left.) These trigger marks are extended for a pair of triggers whose arrival times at Harvard and Princeton were 1 ms apart, as recorded by the computer clocks (see 4.10.1). Note that the trigger rates were elevated at both observatories, particularly during the first half of the night.

are the trigger rates at Harvard and Princeton during the $T_{\text{obsv}} = 20$ minute observation, and $\tau_{\text{coinc}} = 1$ ms is the coincidence window. Of the 1123 observations in $\text{DS}_{\text{overlap}}$, the observation of HIP 107395 on 2003 September 17 has the 5th largest $p_{1\text{ms}}$ (because of high r_H, r_P). The probabilities from all observations can be combined to give the probability of one or more of those observations having a pair of triggers with in ± 1 ms: $p_{\text{all}} = 1 - \prod(1 - p_{1\text{ms}}) = 2 \times 10^{-3}$, where the product is over the observations in $\text{DS}_{\text{overlap}}$.

We considered the possibility that the inaccuracies in the Harvard and Princeton computer clocks were temporally correlated, for example, because the computer clocks are disciplined by identical computer programs and GPS clocks. A comparison of these inaccuracies during normal GPS functioning for Harvard and Princeton events that occurred even within one minute of each other revealed no such correlation.

It is worth noting that not only was this night’s data excluded from DS_{clean} , but furthermore that neither trigger qualified as an event. The Harvard trigger failed because the pulse recorded in one detector was 4 ns long, while the pulse in the other detector was longer than the MTD-135 observation window of 300 ns. The Princeton trigger failed because of an amplitude mismatch in the two detectors (1st and 4th thresholds exceeded). These are both symptoms of corona discharge.

Although the trigger timings are consistent with an astrophysical or extraterrestrial optical pulse, we believe that random coincident background noise is a far more likely explanation. The lack of confirmatory GPS timing, the higher trigger rates on 2003 September 17, and the triggers’ failure to meet event criteria all contribute to this conclusion. Extraordinary claims require extraordinary evidence, which is lacking here.

Nevertheless, even the remote possibility of a world-changing discovery warrants significant reobservations. Over three nights (2003 November 25, 2003 December 2–3), we conducted seven joint reobservations of HIP 107395 (with GPS restored) for a total of 140 additional minutes. No triggers were recorded at Harvard, while Princeton recorded one trigger and no events on 2003 December 2.

The implications of this null confirmation depend on hypothetical source scenarios: If we assume a 50% probability of detecting a pulse during the initial 20-minute observation of HIP 107395 on 2003 September 17 (that is, a signal repetition rate of 1.5 hr^{-1}), then the probability of *not* detecting a pulse in the following seven 20-minute observations is $1/2^7 \approx 10^{-2}$. If we assume (rather unrealistically) that we had a 50% probability of detecting a pulse from *any* object during the 162 hours of joint observations (that is, a rate of $3.1 \times 10^{-3} \text{ hr}^{-1}$), then the probability of *not* detecting a pulse in the following seven 20-minute observations is ~ 0.993 . Thus, if the signal repeats often, the reobservations reinforce our belief that background noise caused the 1 ms trigger pair. If the beacon is broadcast infrequently, the reobservations tell us little.

HIP 107395 (RA = $21^{\text{h}} 45^{\text{m}} 10^{\text{s}}$; dec = $-0^{\circ} 30' 30''$) is an 11th visual magnitude late K dwarf with a parallax of 17.55 ± 2.85 mas (implying a range of ~ 60 pc). It is being surveyed for radial velocity companions for next generation microwave SETI. HIP 107395 is also a member of the Habitable Catalog [63].

4.11 Implications

The foregoing results can be summarized as follows:

1. The Harvard instrument made 15,897 observations of 6176 stars totaling 2378 hours. During these observations (DS_{all}) it detected 4746 triggers with a subset of 1117 events. When we remove the observations characterized by humidity-induced corona with (DS_{clean}), we are left with 11,600 observations of 4730 stars over 1721 hours yielding 274 events and 1066 triggers. These events appear to be Poisson distributed in time and uncorrelated with the target’s brightness and the observing geometry.
2. The Princeton instrument has observed in tandem with the Harvard instrument for 1721 observations ($\text{DS}_{\text{overlap}}$) of 1142 objects totaling 244 hours. The arrival times for one Harvard-Princeton trigger pair are consistent with receiving an optical pulse at the geographically separated observatories (within the accuracy of the computer clocks). For multiple reasons, we believe these triggers resulted from background noise sources.

4.11.1 Scenarios

The implications of our data depend on the model that they are testing. Of the possible intentional optical pulsed signals that an extraterrestrial civilization could generate, let us consider the implications of just two scenarios.

Scenario one: A fraction f of the stars in our region of the galaxy harbor civilizations that transmit optical signals to Earth that our experiment could detect. The signal is composed of multiple pulses in fast succession (less than our minimum observation time of 2 minutes) displaying some hallmark of intelligence (e.g., nonrandom arrival times). This signal is broadcast repeatedly with a period P (greater than our maximum observation time).

Scenario two: A fraction f of the stars in our region of the galaxy harbor civilizations that transmit optical signals to Earth that our experiment could detect. The signal is composed of a single pulse that is broadcast repeatedly with a period P (greater than our maximum observation time).

While these scenarios might be considered simplistic, the implications for our data are nonetheless instructive. The Harvard search, with its sensitivity to multiple pulses, but not solitary pulses, is an excellent probe of scenario one. Scenario two requires a background rate of zero—a good match for the Harvard-Princeton search.

4.11.2 Fraction of stars with transmitting civilizations

An upper bound on the fraction of stars in our region of the galaxy that are signaling Earth with optical flashes can be calculated as a function of P . In the analysis that follows, paralleling Horowitz and Sagan [31], we assume that none of the Harvard or Harvard-Princeton observations detected signals from extraterrestrial civilizations.

The probability of detecting an extraterrestrial signal during an observation of duration t_{obsv} is:

$$p_{\text{obsv}}(P) = \min(1, t_{\text{obsv}}/P), \quad (4.1)$$

where the minimum function bounds $p_{\text{obsv}}(P) \leq 1$. The probability of detecting a signal from any one object is:

$$p_{\text{obj}}(P) = 1 - \prod_{\text{obsv}} (1 - p_{\text{obsv}}(P)). \quad (4.2)$$

The expected number of signal detections during the entire program, S , is the sum of the objects' probabilities times the fraction of objects that are transmitting:

$$S = f \sum_{\text{obj}} p_{\text{obj}}. \quad (4.3)$$

We adjust f so that the Poisson probability of observing zero extraterrestrial signals (e^{-S}) is 0.5; that is, we choose f so that the observing program has a 50% chance of success. Solving for f , we obtain an upper bound on the fraction of transmitting civilizations:

$$f(P) = \min \left(1, \frac{\ln 2}{\sum_{\text{obj}} p_{\text{obj}}(P)} \right), \quad (4.4)$$

where the minimum function limits stars to at most one transmitting civilization.

Fig. 4.7 shows upper bounds on $f(P)$ for scenario one (Harvard; DS_{clean}) and scenario two (Harvard-Princeton; $\text{DS}_{\text{overlap}}$). The two limits of P are of interest. For high repetition rate transmitters ($P \ll T$), $f \rightarrow 1/N_{\text{obj}}$. (The total observing time is $T = \sum_{\text{obsv}} t_{\text{obsv}}$.) There is a cutoff repetition rate, $P \geq T/\ln 2$, above which $f = 1$ and we cannot say anything about the density of transmitters. Also note that the Harvard limit is below and to the right of the Harvard-Princeton limit because the latter derives from fewer objects observed for less time.

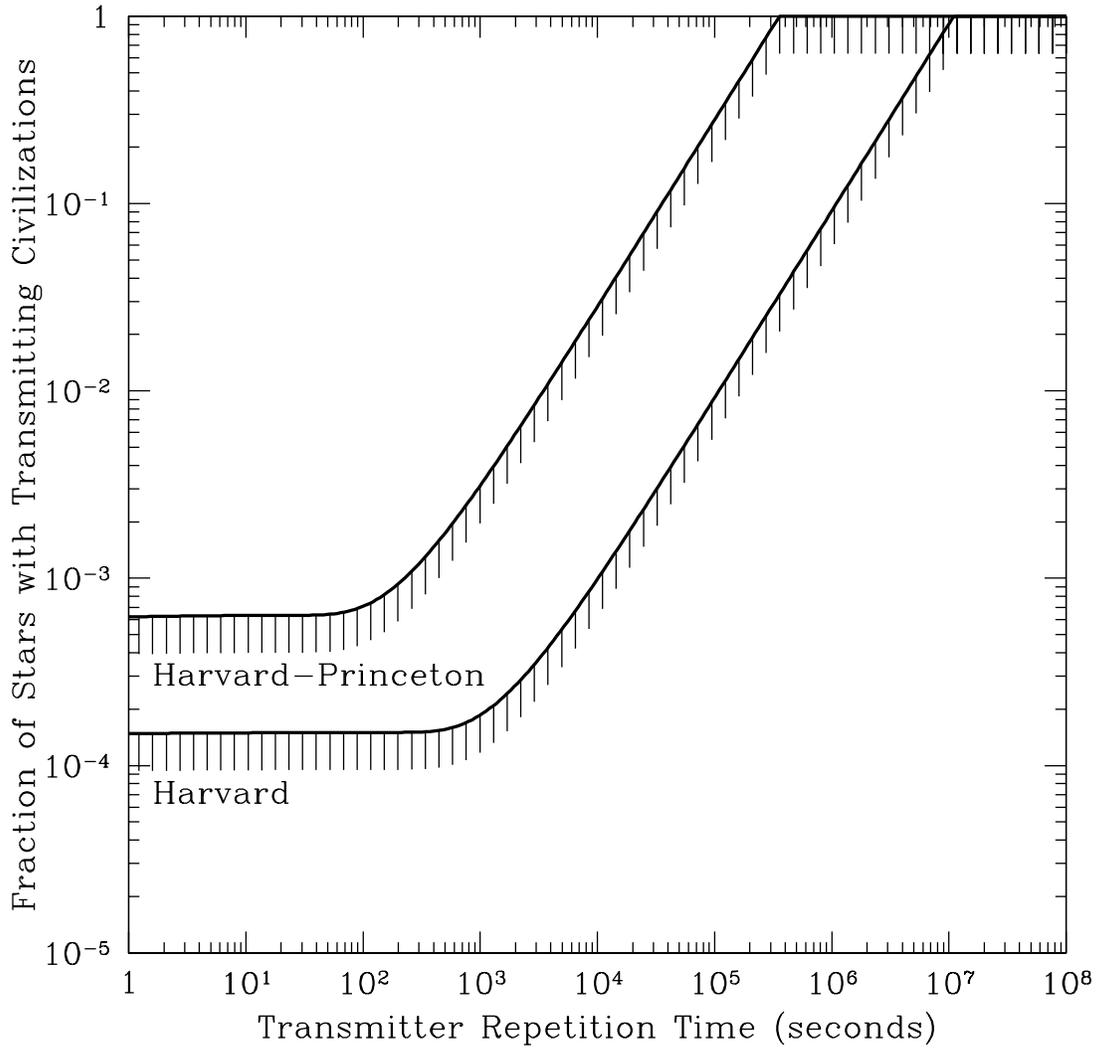


Figure 4.7. Upper bounds on the fraction of stars with transmitting civilizations as a function of transmitter repetition time for the Harvard experiment (scenario one; DS_{clean}) and the Harvard-Princeton experiment (scenario two; DS_{overlap}). Each curve asymptotes to $f = 1/N_{\text{obj}}$ for small P and cuts off at $f = 1$ for $P \geq T/\ln 2$.

Important note: This figure regrettably and erroneously appeared in [35] with the x-axis labeled “Transmitter Repetition Time (hours)” instead of “Transmitter Repetition Time (seconds).”

Chapter 5

All-Sky Optical SETI

5.1 Motivation

The targeted search discussed in Chapter 4 has a significant shortcoming—after several years of data collection, it has covered less than 10^{-5} of the sky area. With $\sim 10^7$ Sun-like stars within 1 kpc, 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.

In contrast to the targeted search, where we are able to choose stars that we believe are likely to harbor life (or are at least good candidates for planetary companions) 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; advanced civilizations may live in, and transmit from, the *voids* between stars [22]. 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 within its sky coverage. 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.

5.2 Overview of experiment

The all-sky search is conducted with a 1.8 m optical telescope that images $1^\circ 6'$ by $0^\circ 2'$ on two arrays of fast, pixelated photodetectors through a beamsplitter. A flash of light in the field of view will illuminate one of 512 matched pairs of photodetector pixels and trigger custom electronics to record the waveform profile and event timing. It is a meridian transit survey; the telescope has settable declination and fixed hour angle. Each point on the sky drifts across the field of view with a minimum dwell time of 48 sec. By observing a different declination each night, we can cover the Northern sky in ~ 200 clear nights, as illustrated in Fig. 5.1.

The all-sky survey is a multi-pixel elaboration of the targeted search. Each of the 512 pair of photomultiplier tube pixels functions like the pair of photodetectors in the targeted search. The parameters of this search are listed in Table 5.1.

This chapter describes the necessary experimental ingredients for building the all-sky search: the observatory building, the telescope, the photomultiplier tubes (PMTs) used to detect the short flashes, the camera (which contains the PMTs, beamsplitter, and electronics), the end-to-end testing devices, and the software to control it all. Finally, we consider the sensitivity of these pieces working together as an instrument.

5.3 Observatory

The all-sky observatory building ($71^{\text{h}} 33^{\text{m}} 26^{\text{s}}$ W longitude, $42^\circ 30' 20''$ N latitude) is located at the Oak Ridge Observatory in Harvard, Massachusetts, home of the 1.6 m optical telescope used for the targeted search (Chapter 4). The telescope sits in a specially-constructed observatory building that measures 9 m (N-S) \times 5 m (E-W), as shown in Fig. 5.2. It is a steel truss structure (capable of supporting itself without a roof)

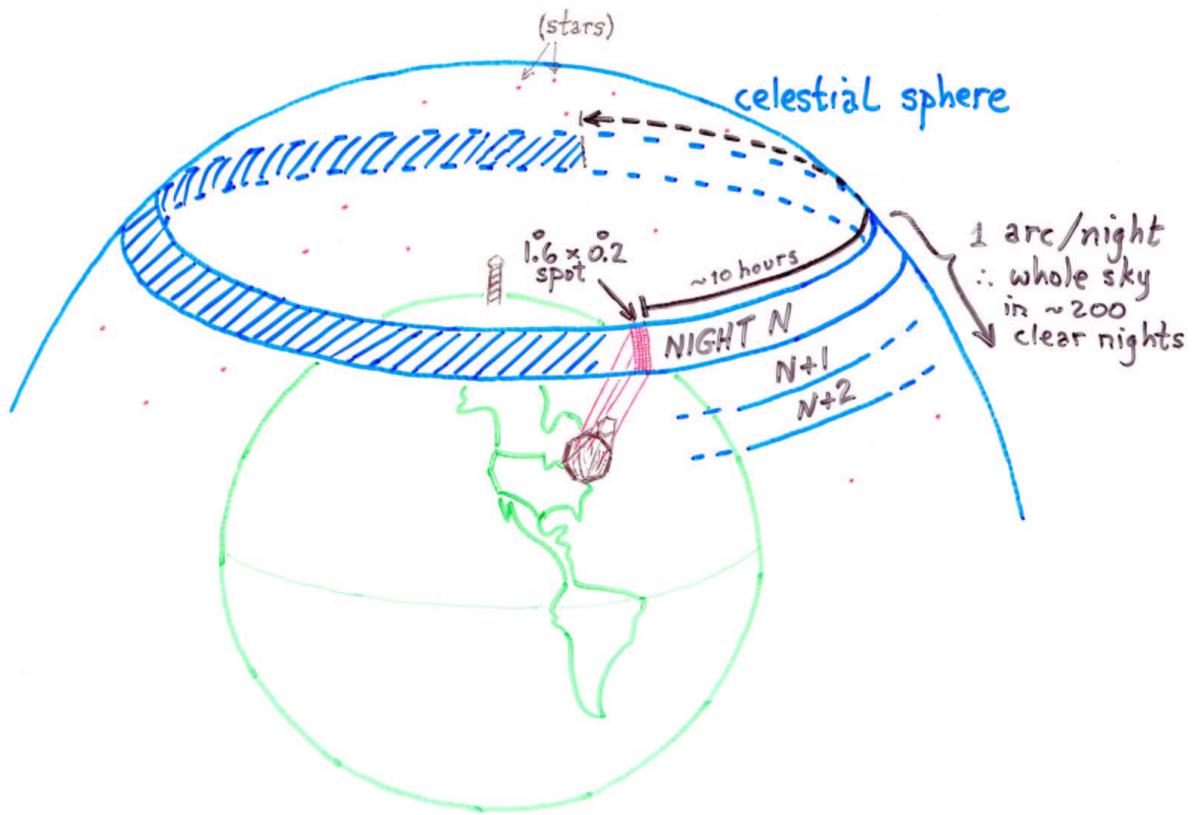


Figure 5.1. Transit survey of the celestial sphere. Drawing by Paul Horowitz used with permission.

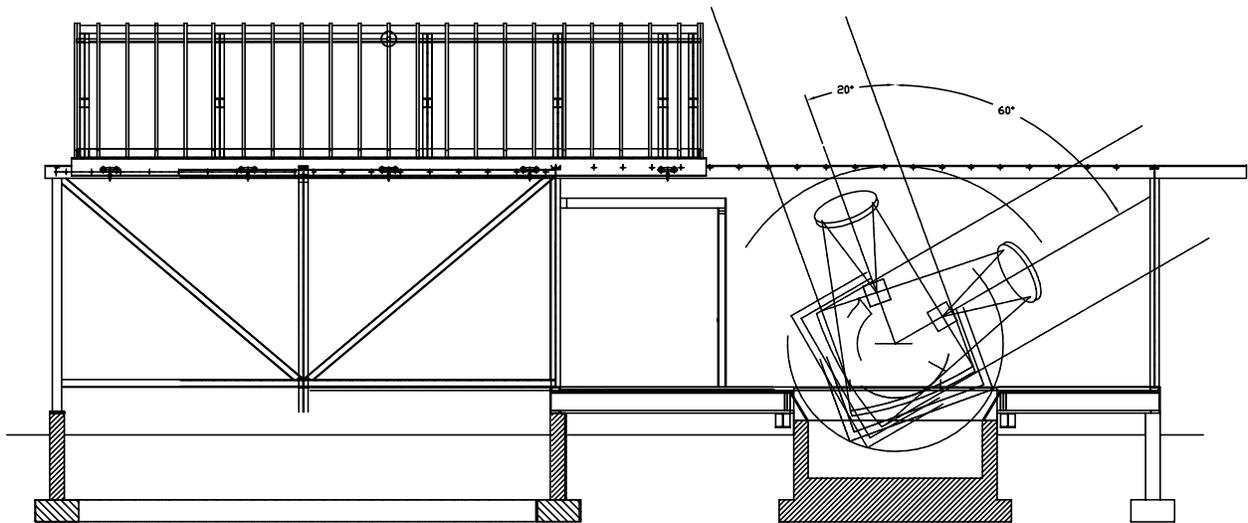


Figure 5.2. All-sky observatory building viewed in profile from the west. The roof rolls on rails that span the southern telescope room (right), control room (middle), and extended northern roof support structure (left). The telescope sits on a separate concrete pier.

Table 5.1. Project parameters – all-sky optical SETI at Harvard.

Telescope	1.8 m f/2.5 spherical "quasi-Newtonian" telescope in Harvard, Massachusetts
Survey Mode	Survey Northern sky ($-20^\circ \leq \text{declination} \leq +70^\circ$) in ~ 150 -200 clear nights Telescope points at fixed nightly declination (transit mode) Sky drifts through $1.6^\circ \times 0.2^\circ$ focal stripe with a minimum dwell time of 48 sec
Photometer	16 64-pixel photomultiplier tubes (PMTs) arrayed in two matched focal planes, using a beamsplitter The PMTs have response in the 300–650 nm band, peaked between 300–450 nm with a quantum efficiency of up to 20% Each $1.5' \times 1.5'$ sky pixel is observed by two PMT pixels – a signal must be observed simultaneously in a pixel pair to trigger action by a PulseNet Pulse amplitude profiles with up to 1 ns resolution GPS-derived timestamping of events to 0.1 μs precision
Electronics	1024 wide-band amplifiers for PMT signals 32 PulseNet chips for analog-to-digital conversion, coincident pulse recognition and storage, and astronomy functions 12 microcontrollers and 12 PALs for PulseNet I/O, telemetry, diagnostics, etc. 41 custom PC boards (of 4 types) PC104 for instrument control and data transfer, via dual-ported SRAM
Sensitivity	≥ 95 photons m^{-2} in the photometers' waveband in ≤ 3 ns; a planned circuit upgrade will improve this to ≥ 17 photons m^{-2}
Initial observations	1% sky coverage ($\sim 10^6$ stars) in 17 hr of observation 0 optical pulses detected above threshold

with an attached wood facade. The telescope sits on an isolated concrete pier in the southern portion of the building, while the northern part comprises an environmentally sheltered control room for electronics, computers, equipment, and operators. A rolling roof is suspended on inverted C-shaped rails that span the length of the building and extend another 7 m to the North on a steel support structure. A custom controller powers a DC motor that translates the roof along the North-South rails by turning a helical U-groove worm gear that is wound by steel cable attached to each end of the roof. During observations, the roof is parked to the North; it can also extend slightly over the south wall to lift heavy equipment into the building. Two “barn doors” open the top portion of the south wall for low-declination observations. The bottom portion of this wall is attached with bolts, and can also be removed.

The control room houses computers and control electronics for several devices. The “Shulsky box” is an 18-port power-controlling device with manual switches and computer control. Cameras allow for remote monitoring of the control room, telescope and roof positions, and building exterior. Infrared lights ensure good illumination for these cameras while maintaining darkness in the PMT-sensitive visible spectrum. A commercial weather station (Davis Vantage Pro2) records and monitors environmental conditions and will allow automatic roof closing during inclement weather.

5.4 Telescope and optics

As shown in Fig. 5.3, the telescope is a “quasi-Newtonian” with a 1.8 m spherical primary mirror and a flat 0.9 m secondary mirror inclined at 22°5 from perpendicular to the optical path. The dashed green lines show the optical path from primary to secondary through the an approximately 60–40 plate-glass beamsplitter and onto two matched focal planes within the camera.

As a transit survey instrument, the telescope is steered only in declination; the hour angle is fixed on the meridian. The declination drive system is a 1.2 m-diameter aluminum arch rigidly attached the telescope frame that is friction-driven by a stepping motor. Computer controls enable precise positioning of the telescope (§5.8) and mercury limit switches prevent the it from rotating past critical angles.

The telescope is covered in a mylar coat with a transparent window to minimize dust accumulation. The primary and secondary mirrors both have heaters to prevent condensation.

5.4.1 Spherical vs. parabolic

The spherical primary mirror does not produce astronomical-quality images. Because the telescope was not intended for high-resolution imaging, this is not a problem. The camera pixel size is relatively large, and was chosen to match the mirror’s ~ 1.5 arcmin point source blur circles.

Fig. 5.4 illustrates the tradeoffs between a spherical and parabolic primary mirror (which we also considered). Uncorrected spherical mirrors are rarely used in astronomy because of spherical aberrations, which are essentially independent of the angle off-axis, θ , because of their symmetry. For these mirrors, the angular size of the blur circle due to spherical aberration is given by

$$B_{SA} = \frac{1}{128(f/d)^3}, \quad (5.1)$$

where f/d is the ratio of the mirror’s focal length to diameter, and B_{SA} is expressed in radians. For the $f/d = 2.5$ all-sky telescope’s primary, $B_{SA} = 1.5$ arcmin, as shown in the top ray-trace in Fig. 5.4.

Parabolic mirrors focus an on-axis point source to a perfect point in the focal plane, and nearly on-axis sources are only slightly distorted. Many optical telescopes use parabolic primaries because of the importance of high-resolution, small-field imaging. With these mirrors, the dominant optical aberration is coma, which causes point sources at infinity to spread into comet-like shapes, as in Fig. 5.4. The angular size of these cones is given by

$$B_{coma} = \frac{3\theta}{16(f/d)^2}, \quad (5.2)$$

where B_{coma} has the same units as θ . For an $f/d = 2.5$ parabolic primary, $B_{coma} = 1.8$ arcmin at 1° off-axis (as shown in the ray-trace in Fig. 5.4).

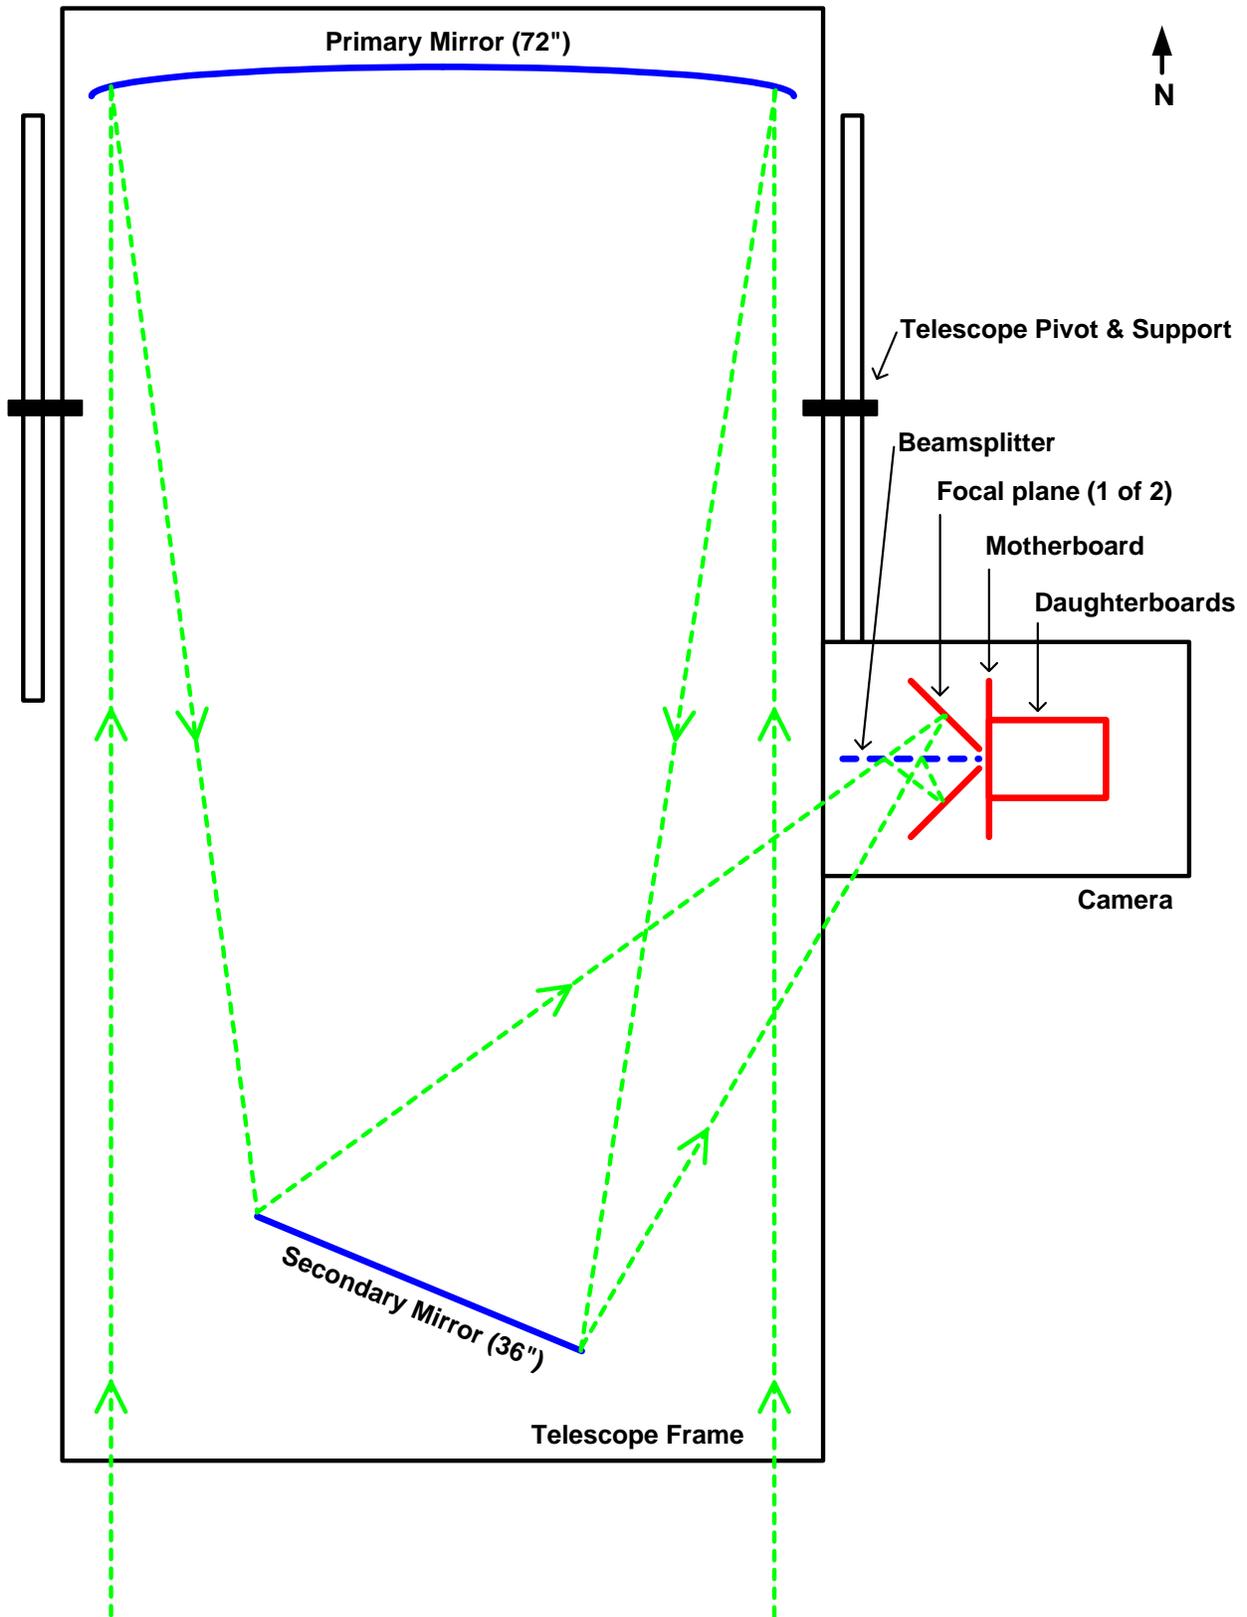


Figure 5.3. Telescope and camera viewed from above. The dashed green lines trace the optical path through the system.

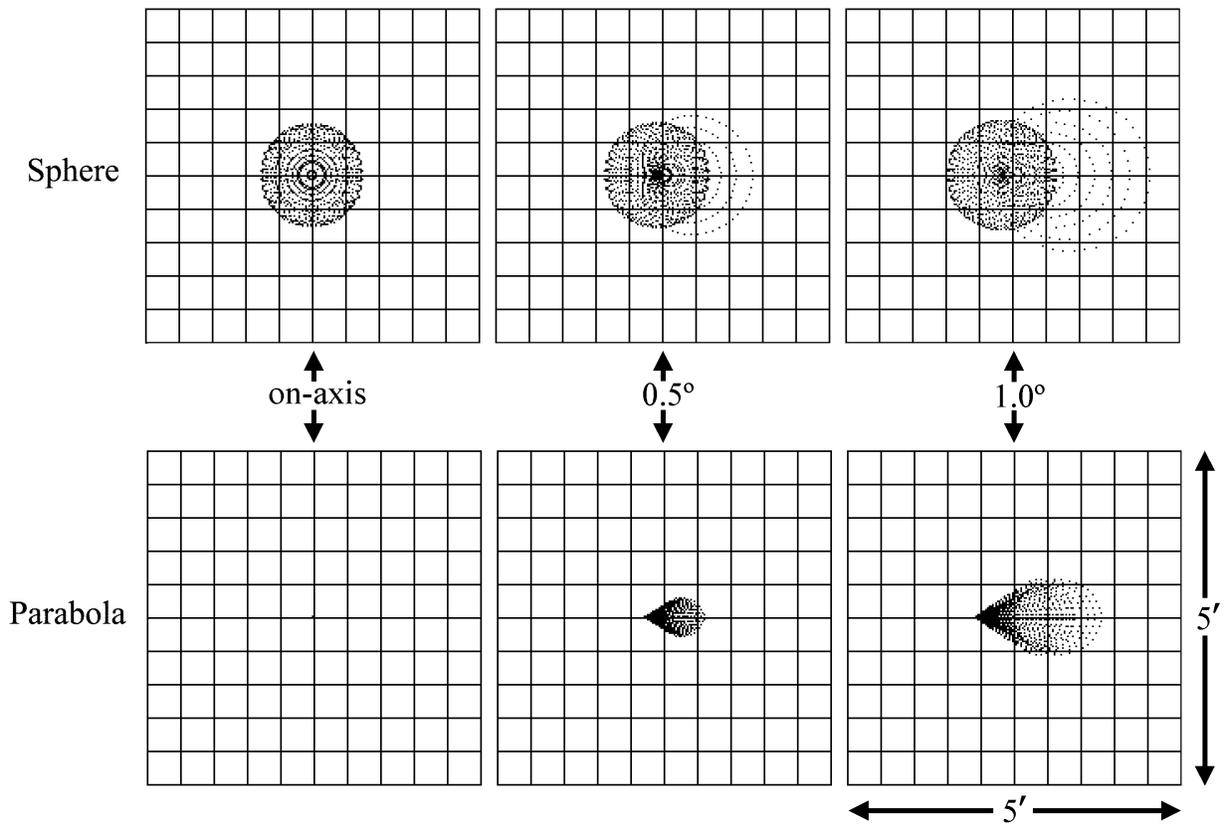


Figure 5.4. Ray-traces showing optical aberrations in the image plane for a point source at infinity (for three off-axis angles). Results are shown for a spherical primary mirror (top three grids) and parabolic primary mirror (bottom three grids). The scale for each grid is 0.5 arcmin per square.

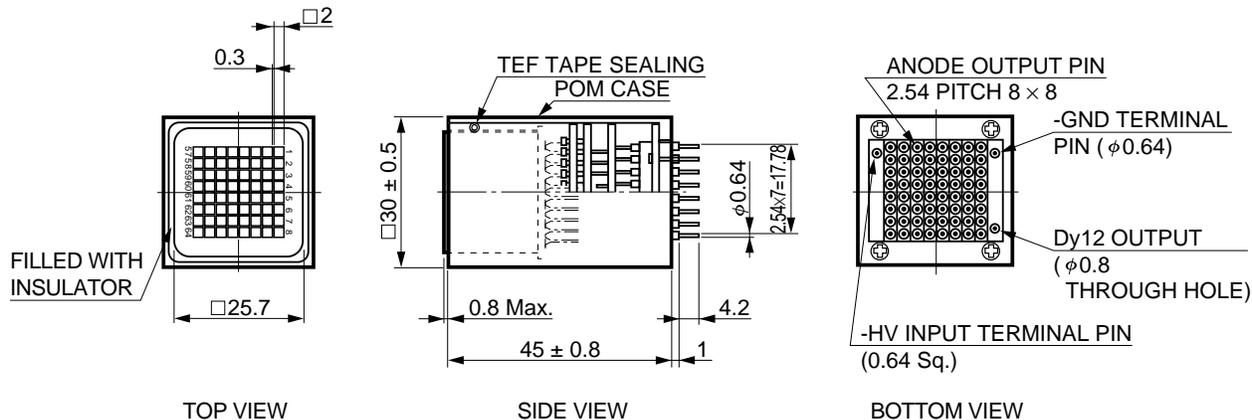


Figure 5.5. Dimensional outline of the Hamamatsu H7546. Sixteen of these photomultiplier tubes form the detector array in the all-sky survey. Note the 64 independent cathodes (top view) and anodes (bottom view). Dimensions in the drawing are in mm.

Since the all-sky search only images coarsely, and a wide field of view is essential for the survey, a spherical primary is a reasonable choice. Furthermore, because of their symmetry, spherical optics are much easier to manufacture and are therefore much less expensive. Given the price difference (at least an order of magnitude for a 1.8 m mirror), a spherical primary was the only rational choice.

5.5 Photomultiplier tubes

The photodetectors for this experiment are Hamamatsu H7546 64-pixel photomultiplier tubes (PMTs). They each have a single photocathode and 64 independent dynode chains and anodes, as shown in Fig. 5.5. They function like 64 independent tubes, although they have roughly the size and weight of traditional, single-pixel PMTs.

The primary reason for choosing a PMT over other photodetector technologies is the fast response shown in Fig. 5.6. Their fast response (~ 3 ns electrical pulse width) and quick recovery time allow for single photon detection. Other performance specifications are similar to single-pixel tubes. The quantum efficiency (QE) peaks at $\sim 20\%$ between 300 and 450 nm, but doesn't fall off completely until ~ 650 nm (Fig. 5.7). We operate the tubes at 900 V, where they have a gain of $\sim 10^6$ (also Fig. 5.7).

All photomultiplier tubes suffer from the relatively soft pulse height spectra (Fig. 5.6) that results from cascading many low-gain stages. This makes it difficult to set a threshold for an integer number of incident photons since their pulse height distributions significantly overlap. These properties contrast with the single-stage gain of Hybrid Avalanche Photodiodes (HAPDs), used in the targeted search, that can cleanly resolve small integer numbers of incident photons. On the other hand, HAPDs run at much higher voltages and suffer from corona breakdown, as described in §3.4.1.

One drawback with these multi-anode tubes is non-uniform pixel response¹, which can vary by up to a factor of 2–3, although typical pixels vary by ~ 30 –50%. We do not currently have a way to compensate for this.

Each PMT has an active area of 18.1 mm on a side, with each of 64 square pixels measuring 2 mm on a side ($1.5'$ in sky angle) with 0.3 mm gaps between pixels. Because packaging overhangs the active area by another 6 mm on a side the tubes are staggered diagonally in the two matched focal planes.

¹The non-uniformity results from the manufacturing challenge of evenly depositing the photocathode material in vacuum when the glass cover is already in place.

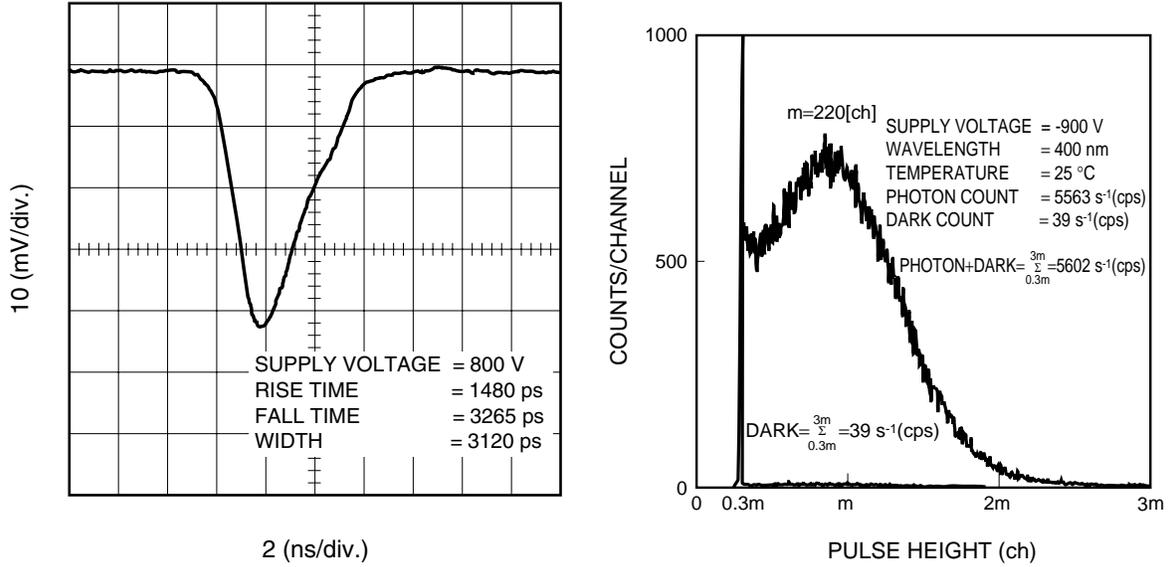


Figure 5.6. Typical time response from a single photoelectron (left), and pulse height spectrum (right) for the 64-pixel Hamamatsu H7546 photomultiplier tube.

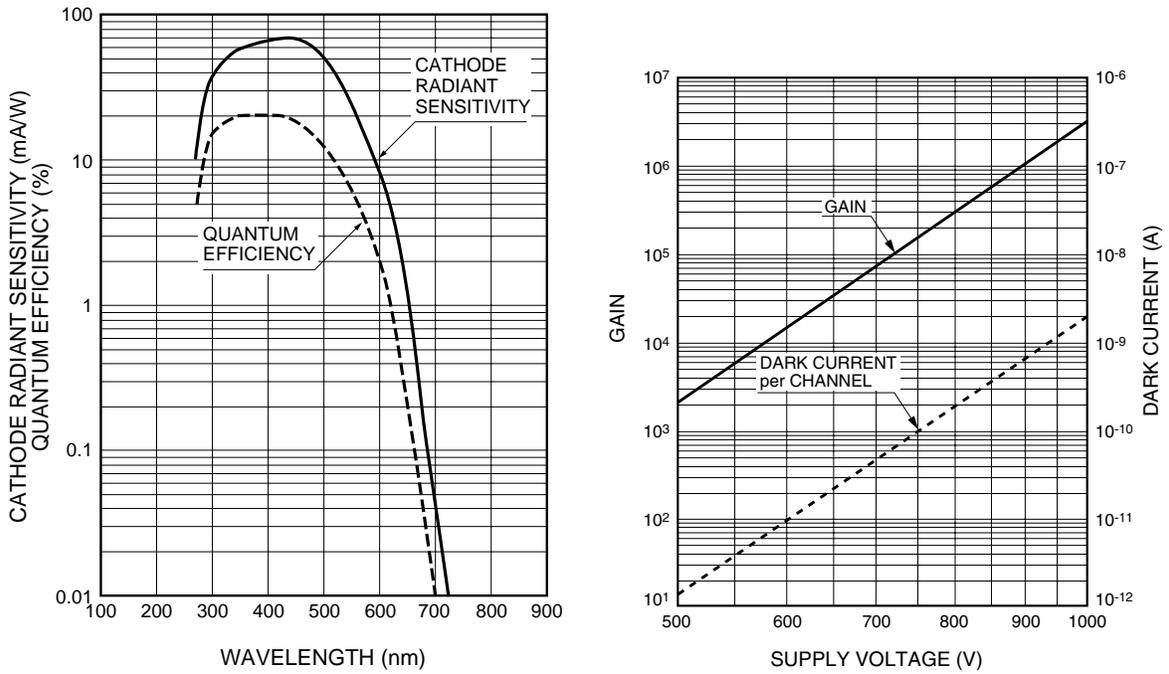


Figure 5.7. Spectral sensitivity (left) and gain as a function of supply voltage (right) for the 64-pixel Hamamatsu H7546 photomultiplier tube.

5.6 Camera and electronics

The camera contains the all-sky search’s photodetectors, beamsplitter, and electronics in a frame of Rexroth extruded aluminum bolted to the East side of the telescope. The camera’s optical compartment is shown in Fig. 5.8, while the electronics compartment and chassis are shown in Fig. 5.9. When the Shulsky Box-controlled shutter is open, the telescope’s converging optical beam passes through an entrance window and impinges on the dielectric-coated plate glass beamsplitter² captured in two grooved Delrin supports. The PMTs are socketed in their printed circuit boards that are diagonally arrayed on vertical Elma rails with PMT pixels in the two matched focal planes numbered according to Fig. 5.10. An ambient light detector below the left PMT array automatically disables the PMT high voltage power supplies during non-dark conditions. The Gelfand Flasher I (§5.7) produces optical flashes on the back side of the beamsplitter that are reflected and transmitted into matched detector pixels for testing.

When photons strike PMT pixels, they generate photoelectrons with some probability. Hugely amplified in number through the PMT dynode chains, they are captured by anodes in the form of current pulses, and terminated in $50\ \Omega$ loads. These signals pass from the PMT boards through multi-coax ribbon cables³ and on to the daughterboards. From there the PMT signals are amplified by arrays of NEC 2771TB discrete amplifiers with 21 dB of non-inverting gain and are routed into full-custom PulseNet chips that look for coincident pulses in matched pixel pairs. As shown in Fig. 5.9, the eight daughterboards are held in place in horizontal rails and plug into sockets in the vertically-oriented motherboard, which lies between the PMTs and daughterboards.

The arrangement of daughterboards on the motherboard is shown in Fig. 5.11. Enhanced 8051 microcontrollers on the motherboard oversee I/O with the 32 PulseNets on eight daughterboards. Each daughterboard is serviced by one SETI microcontroller, which oversees the SETI-related tasks of programming PulseNets and shuttling their coincident event captures to the host computer. Every two daughterboards are serviced by an astronomy microcontroller, which handles the PulseNets’ astronomy functions — measuring countrates for the purpose of coarse but diagnostic photometry.

Fig. 5.12 shows the complicated correspondence between PMT pixels and PulseNet pixels. Note that since four PulseNets service a given PMT (indicated by the PMT pixels of four different colors), adjacent PMT pixels always communicate with different PulseNets insuring that flashes extending over pixel boundaries will not be missed.

A PC104 single-board computer sits atop the electronics chassis and communicates with the microcontrollers and PALs via dual-ported RAMs an ISA bus on the motherboard. Other electronics include chips to generate clocks for the 32 PulseNets, programmable logic (PALs) that timestamps PulseNet coincident events with a GPS clock accurate to $0.1\ \mu\text{s}$. Additionally, the microcontrollers telemeter and report power supply voltages, daughterboard temperatures, and camera humidity.

The camera is enclosed with aluminum panels attached to the Rexroth frame. Fans circulate air and transport heat from the electronics to heatsinks on two of the panels. The only connections between the camera and the rest of the observatory (besides isolated power supplies) are with fiber optic cables that provide network connectivity to the PC104 and GPS-accurate time to the time-stamping PALs. The lack of copper signal connections to the instrument make it robust against the common lightning strikes at Oak Ridge Observatory.

5.7 Gelfand Flashers

The Gelfand Flasher I (GF1) and Gelfand Flasher II (GF2) both emit short optical pulses to conduct complete end-to-end tests. The GF1 is a fixed position flasher that resides inside the camera and sends its pulses through the “fourth port” of the beamsplitter, as shown in Fig. 5.8. The GF1’s blue LED emits a $\sim 50\ \text{ns}$ optical flash, generated by a user-controlled logic-level transition (which is connected to a microcontroller

²The 24 cm by 36 cm beamsplitter (Edmunds Optics 72502) is optimized for 45° use and has a 40/60 reflection-to-transmission ratio that varies by 5–10% over the visible spectrum. Transmission depends weakly on polarization, which can vary by up to 20% from nominal.

³The 40-pin Samtec cables (seen in blue in Figs. 5.8 and 5.9) have the electrical properties of single coaxial cables, but the convenience and density of mass-terminated ribbons.

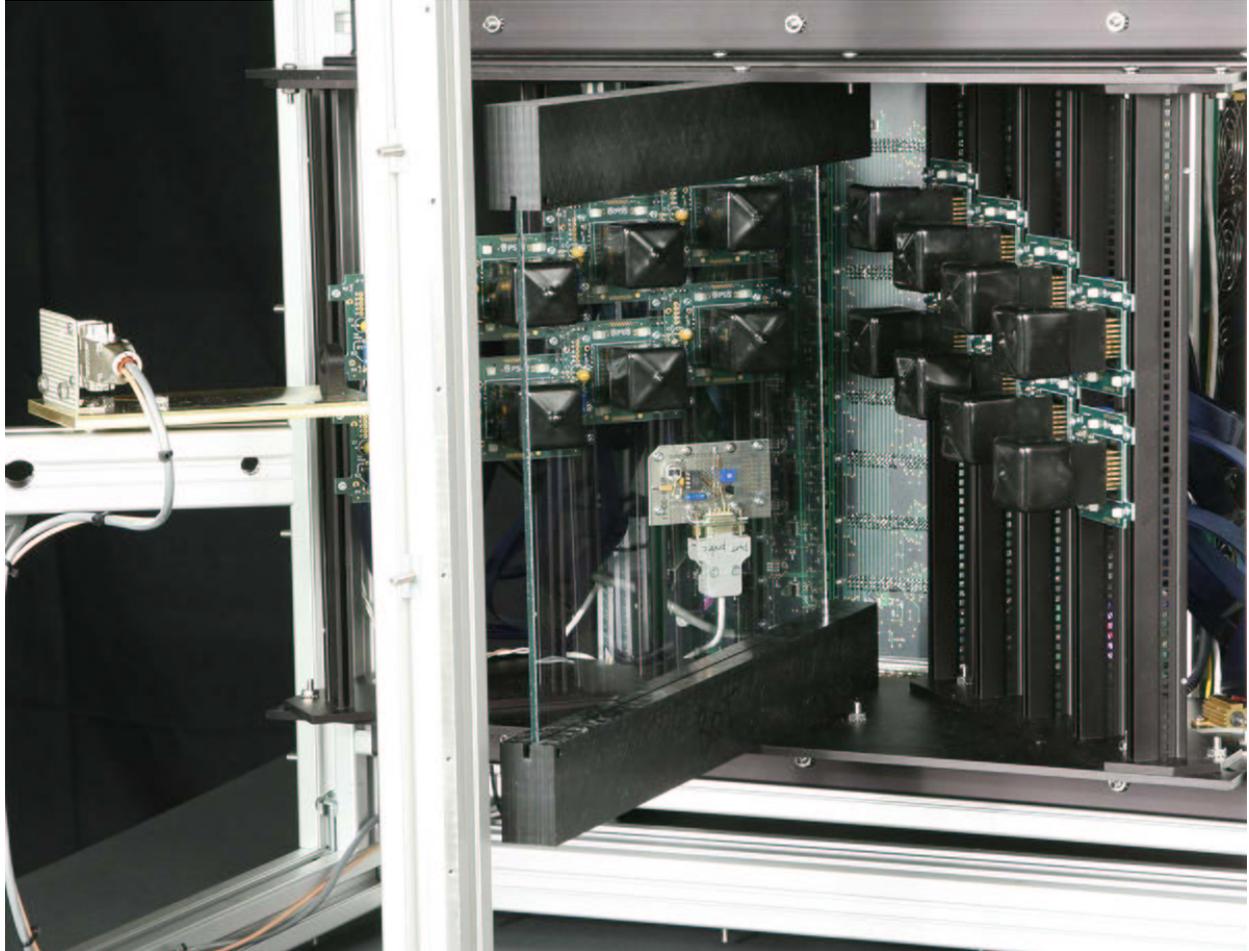


Figure 5.8. All-sky search instrument – front. Two arrays of eight 64-pixel PMTs tessellate matched focal planes formed by a plate beamsplitter (coming out of the page) viewed approximately from the perspective on the converging optical beam. Each PMT is socketed in a custom printed circuit board that holds a power supply, termination resistors, and sockets for two 40-pin ribbon coaxial cables that carry the unamplified PMT signals to the daughterboards (not shown). The motherboard sits vertically behind the PMTs with daughterboards connected on the opposite side. An ambient light detector below the left PMT array automatically disables PMT high voltage during non-dark conditions. The Gelfand Flasher I is seen on the far left with a logic-level cable connection; this device produces an optical flash on the back side of the beamsplitter that is reflected and transmitters to matched detectors.

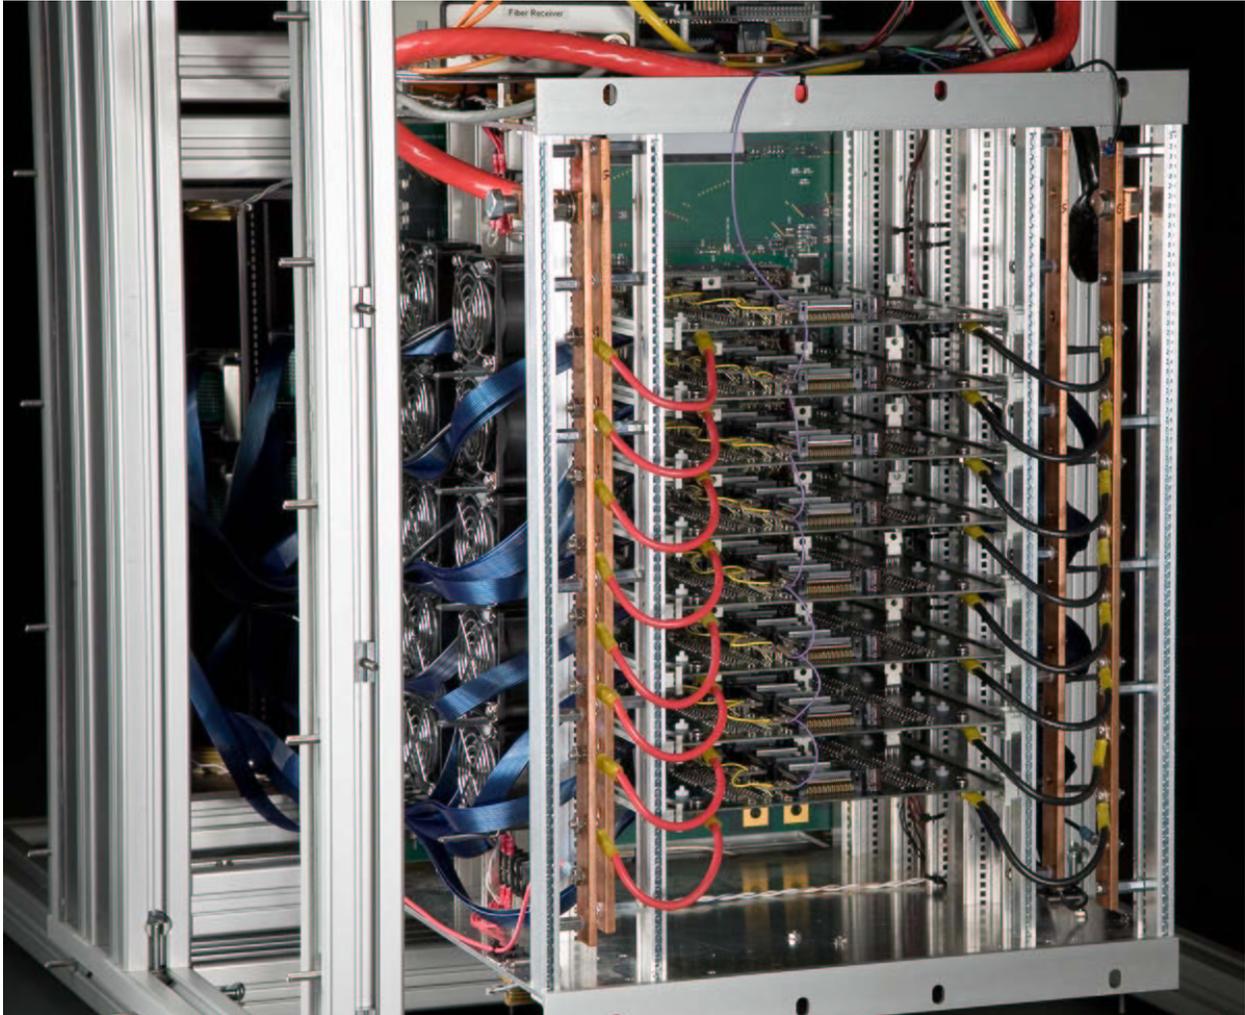


Figure 5.9. All-sky search instrument – rear. Eight daughterboards are seen in the electronics chassis with power connections in front and blue 40-pin ribbon cables connecting to the PMTs (the ribbons are visible on the left side of the chassis; see also Fig. 5.8). Each daughterboard supports four PulseNet chips and processes signals from one pair of matched 64-pixel PMTs. Fans circulate air and, when fully enclosed with panels, transport heat from the electronics to heatsinks on two of the panels. The PC104 single board computer and fiber transceiver are partially visible on top of the chassis.

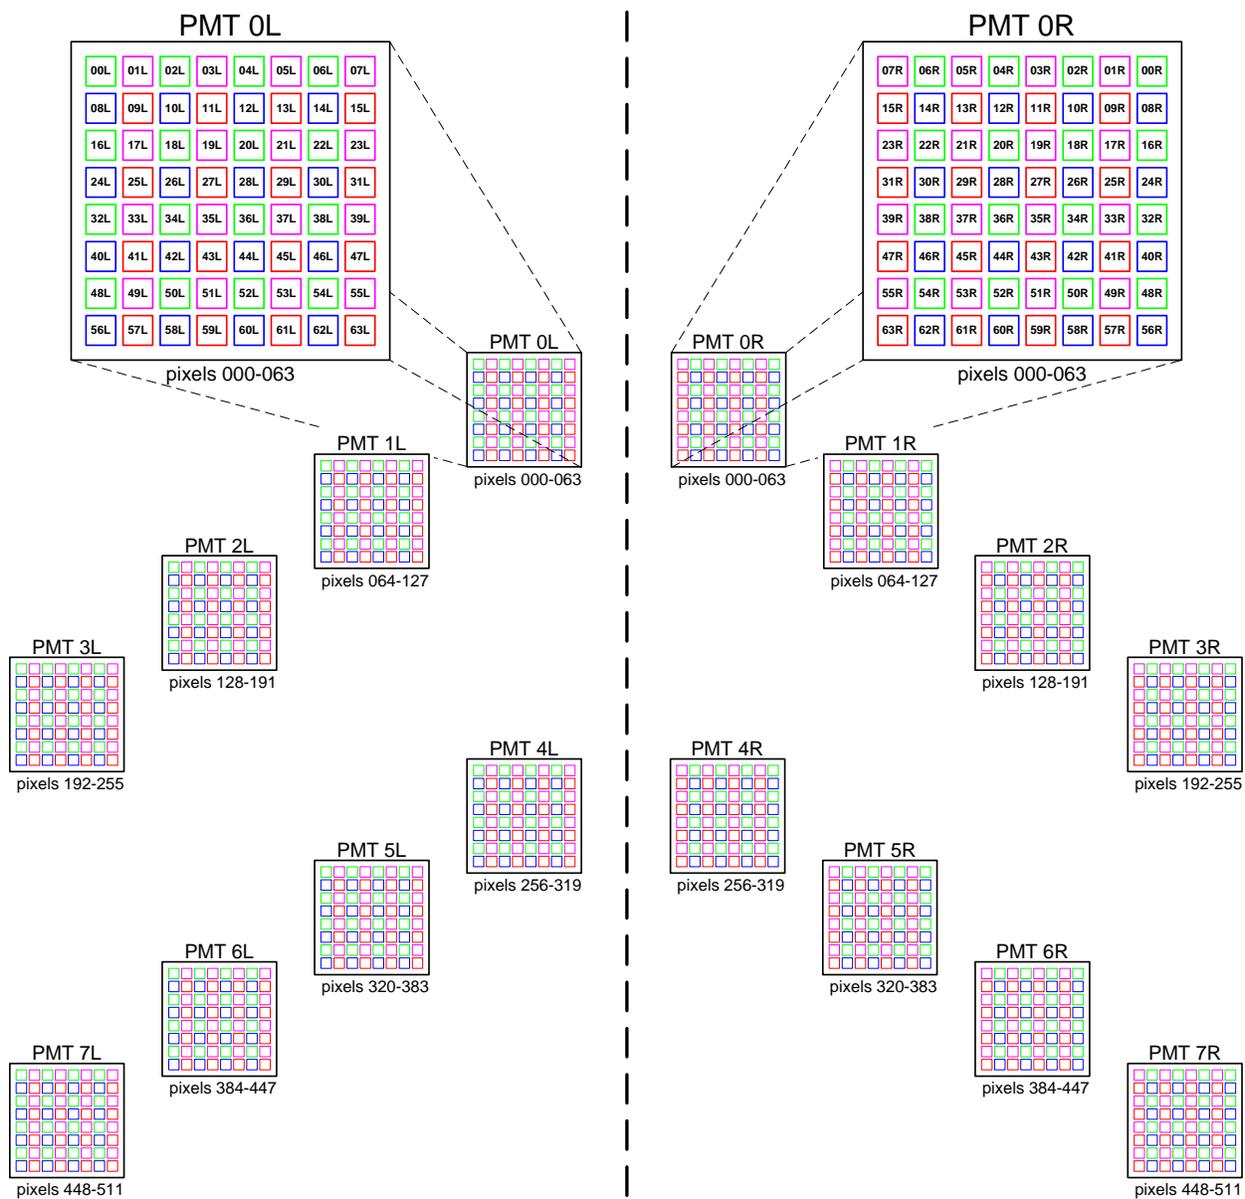


Figure 5.10. Photomultiplier tube arrays in the left (transmitted through the beamsplitter) and right (reflected) image planes.

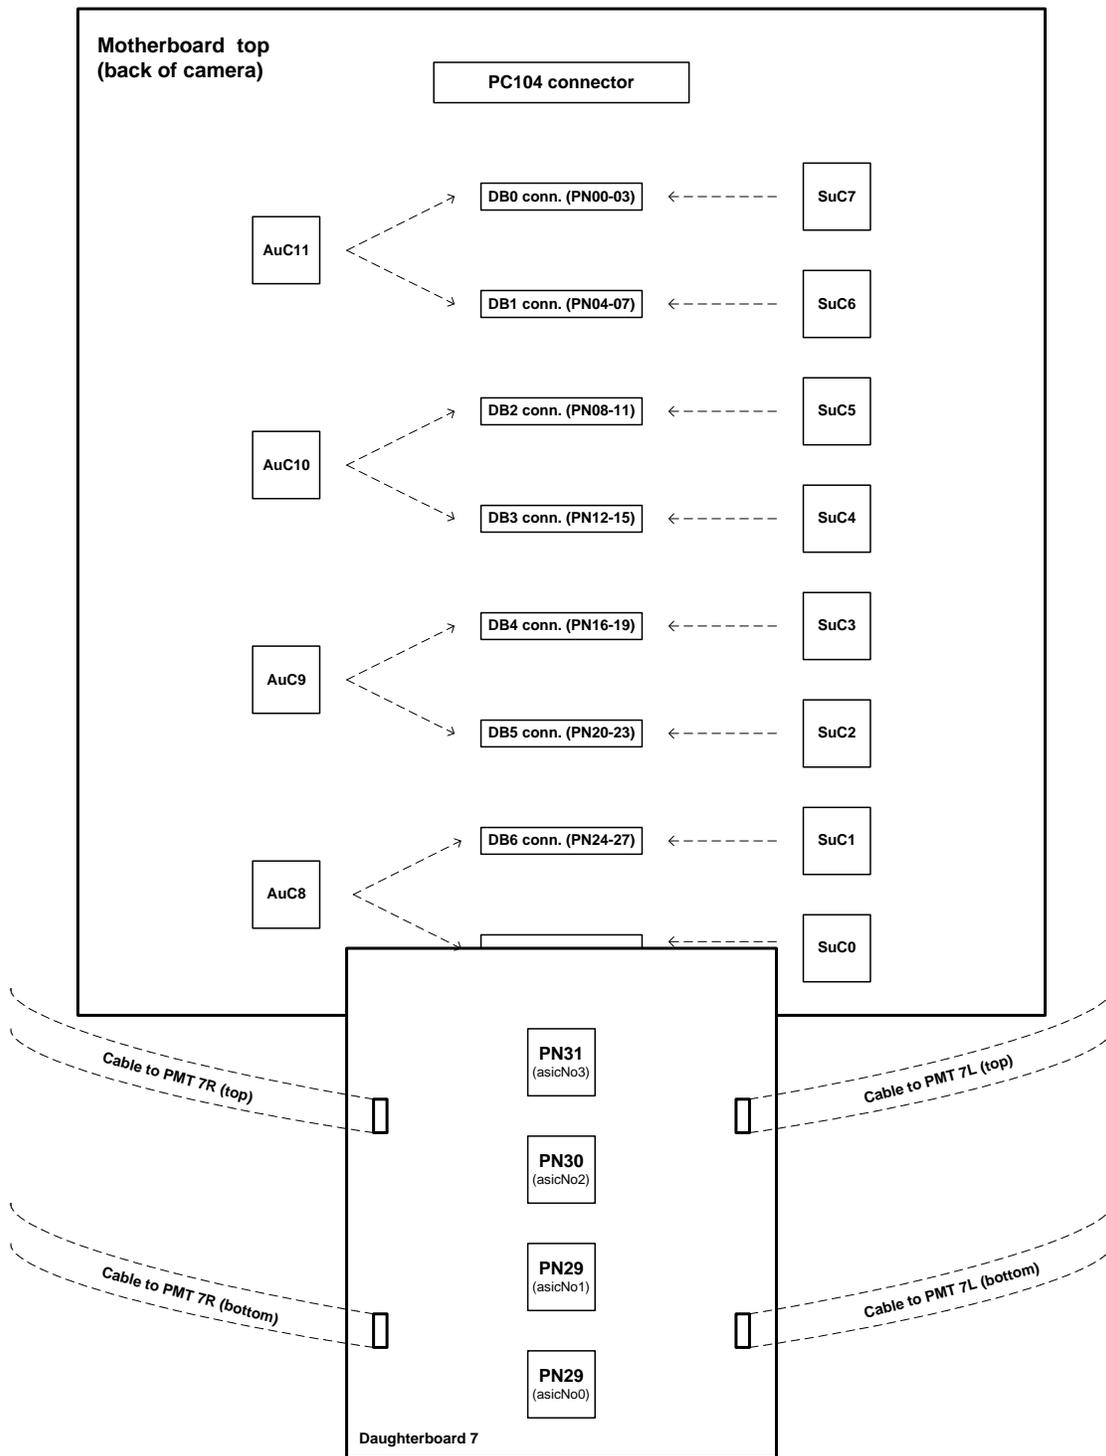


Figure 5.11. Schematic of motherboard and one of eight daughterboards, which fit in the DB0–DB7 sockets. Each daughterboard is serviced by one SETI microcontroller (SuC0–SuC7; note the opposite numbering sense from the daughterboards), and every two daughterboards are serviced by an astronomy microcontroller (AuC8–AuC11). Chips on the motherboard communicate with the PC104 through an ISA bus bridged on the PC104 connector (top). Each daughterboard receives signals from two PMTs which are processed by four PulseNet chips per daughterboard.

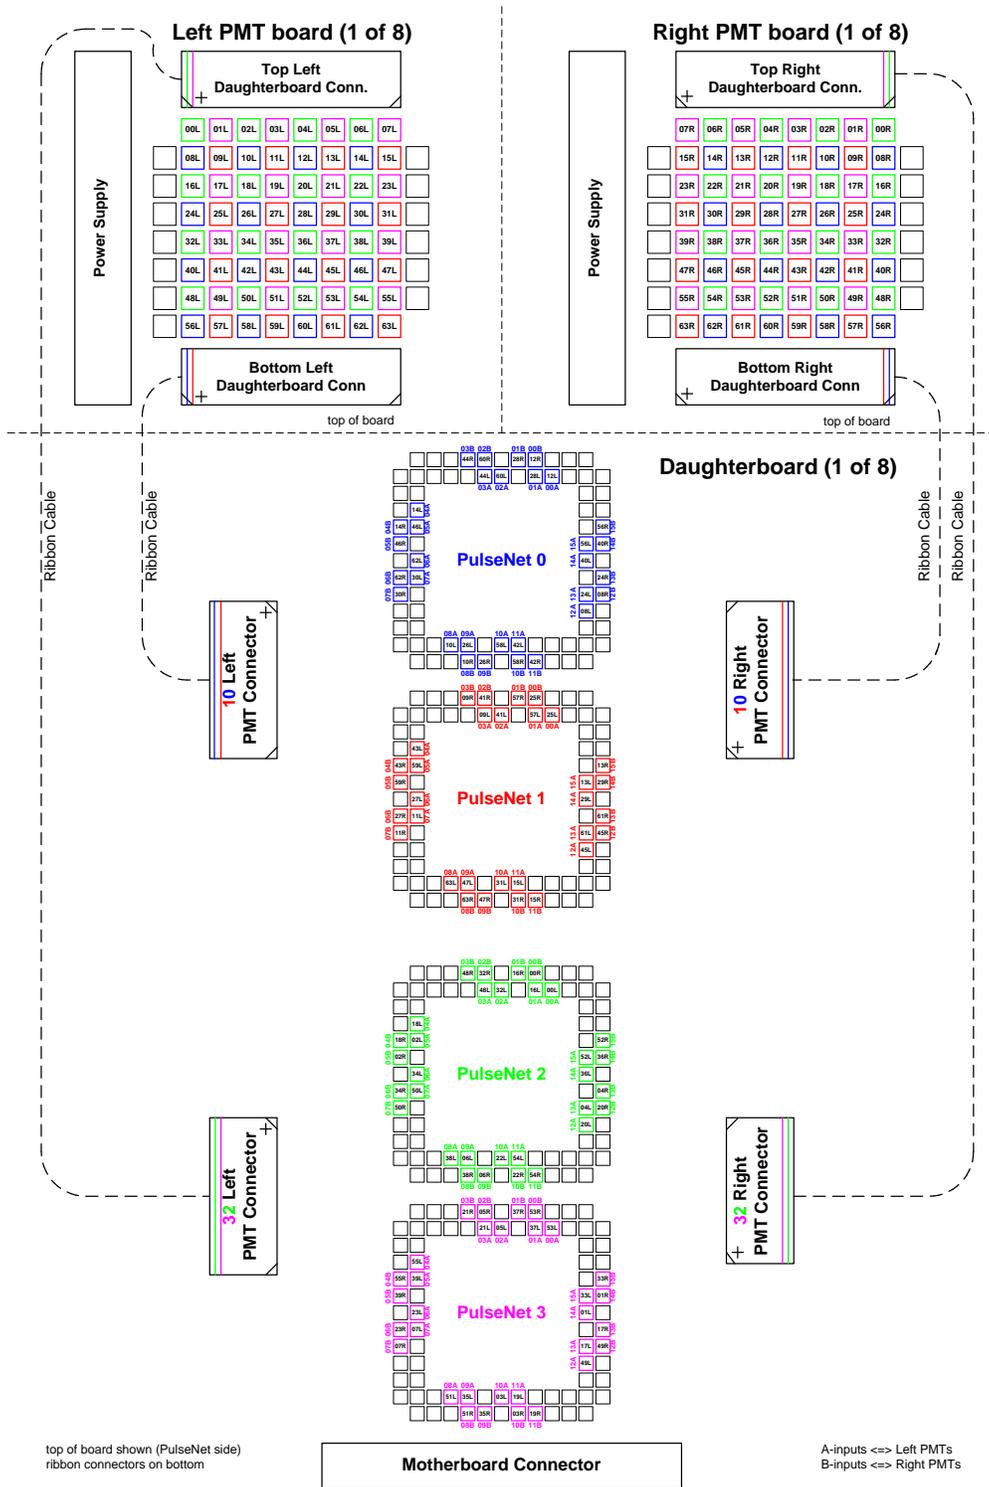


Figure 5.12. Correspondence between PMT pixel numbers and PulseNet pixel numbers. Signals from a pair of PMTs (top) pass through ribbon cables, terminate on a single daughterboard (remainder of plot), and are serviced by four PulseNet chips. Note that adjacent PMT pixels are serviced by separate PulseNets (indicated by the four pixel colors) so that optical pulses extending over the edge of one pixel will also be detected and serviced by neighboring pixels. The PMT pixel numbers (0–63L, 0–63R) and PulseNet pixel numbers (00–15A, 00–15B — color coded) are labeled on PulseNet pins.

output in the camera). It has a mounted lens, but the unfocused flash is bright enough to trigger coincident events in 6–12 of the 32 PulseNets. PMTs below the midplane are obstructed and rarely see GF1 flashes. It was not intended to test every pixel pair in the array; that is the job of the GF2.

The GF2 was mostly built, but not installed, at the time of this writing. Its pulsed blue laser focuses flashes onto one pixel pair, with 2-dimension positional control. Photons from the GF2 travel through a lens and down a ~ 1 m-long tube, cantilevered from the telescope drive plate out into the unfocused primary beam, where they reflect off of a stepper motor-controlled 45° tip-tilt mirror and then follow the traditional optical path off of the secondary mirror, through the beamsplitter, and focus on ~ 1 pixel pair. Computer control should allow for testing of the entire active focal plane, perhaps before each night’s observations.

5.8 Software and user interface

Control over the instrument, the observatory, and their subsystems is achieved through several layers of electronics and software (Fig. 5.13). At the base are subsystems that have their own dedicated controls. The roof drive controller, for example, has physical buttons on a panel interface in the control room, but can also be controlled electronically through its serial connection to control software on a Linux PC (Costas). The telescope control is similar; a paddle interface allows for hand-held control, but computer control allows for position calibration and accurate positioning with respect to celestial coordinates. The telescope and roof controllers are conservative in relinquishing control to computers; each have pairs of limit switches that first indicate a dangerous position and then cut power to the drive motors if the warning is unheeded. A daylight sensor, fabricated from a strand of wavelength-shifting fiber, collects photons on the outer north observatory wall and delivers them to a photodiode in the control room. During daylight hours, the sensor prevents the roof from moving north (moving south is always allowed) via a relay interlock in the roof drive, unless manually overridden by a push-button. The control software indirectly supplies power to all of the electronics and subsystems through the Shulsky Box, an 18-port custom-built power control device that also offers manual control through a switch panel.

User control over the control software is achieved through a web interface that gives nearly complete control over the observatory and camera from any networked computer (Fig. 5.14). The control software is written in Python and uses CherryPy to serve dynamic content from Kid templates. The AJAX-based (Asynchronous Javascript and XML) interface automatically updates state and telemetry displays. SVGs (Scalable Vector Graphics) display real-time content including telescope and roof position (upper right in Fig. 5.14). This is backed by a searchable PostgreSQL database of events and associated instrument/observatory states. Besides the features mentioned above, one can also program and query camera components such as SETI and astronomy microcontrollers, clocks drivers, DACs, and the GF1. Regular telemetry of the camera and subsystems update displays on the user interface. Cameras with infrared illuminators display pictures of the observatory to remove observers. This computing infrastructure is backed by a database that logs all events⁴ and the complete instrument state in which they occur. Manual observations are conducted by a sequence of commands on the web interface.

While observations are currently conducted manually by a sequence of user-initiated commands on the web interface, we plan to make the nightly observations a fully automated task. The control software will follow a decision tree and a set of standard observing procedures somewhat like the ones developed for the initial observations (§8.3).

5.9 Sensitivity

We can estimate the sensitivity by following a light pulse through the entire system, as we did with the targeted search (§4.2). The optical path includes reflections on the primary and secondary mirrors (each $\sim 85\%$ efficient) and a beamsplitter ($\sim 92\%$). The beamsplitter sends half the light to each PMT array, which have broad $\sim 20\%$ plateaus in quantum efficiency for $\lambda = 300\text{--}450$ nm (QE $> 10\%$ for $300\text{--}525$ nm; see

⁴“Events” are defined more broadly here than just coincident events in the camera. They are the full set of actions that the control software initiates and observes. Examples include: roof control commands, programming of camera components, regular telemetry requests, and serial port communications.

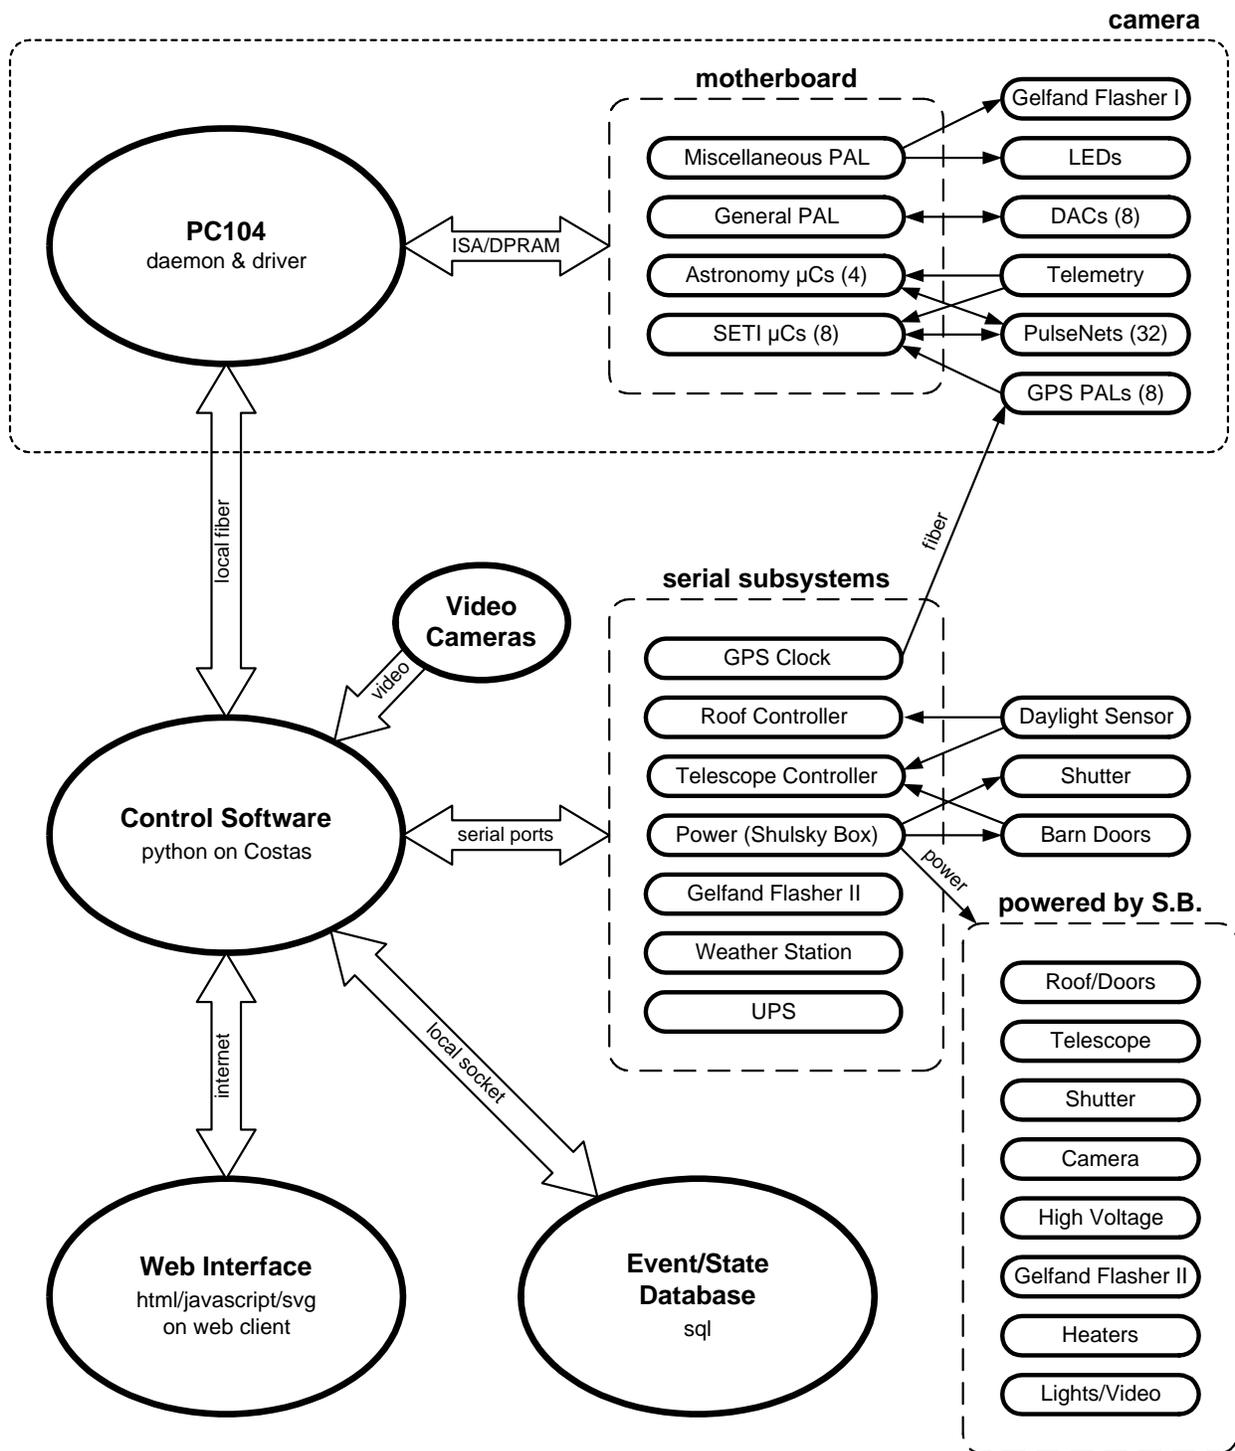


Figure 5.13. All-sky software and firmware. Control software on a Linux PC (Costas) is the communications and signaling nexus for the all-sky search. Using a web interface, with events and state information logged to an SQL database, it powers nearly all electronics and subsystems through the Shulsky Box, orchestrates observatory functions (roof and telescope movement, weather monitoring, heaters, etc.), and communicates with all electronics in the camera through the PC104.

Fig. 5.7). The sensitivity of PulseNet depends on the value of $V_{\text{bias}} - V_{\text{ref}} = V_{\text{trigger}}$ and on the pulse height of a single photoelectron, V_{pe} . The initial observations described in Chapter 8 used $V_{\text{trigger}} = 250 \text{ mV}$ and $V_{\text{pe}} \approx 15 \text{ mV}$ and triggered on $250 \text{ mV} / 15 \text{ mV} = \sim 17$ photoelectrons. Combining these factors, for the most sensitive waveband of $\lambda = 300\text{-}450 \text{ nm}$, the all-sky instrument requires ~ 250 photons on the primary mirror to trigger. These photons must arrive within $\leq 3 \text{ ns}$ (the PMT pulse width) so that their outputs pile up. The primary has an area of $\pi(0.91 \text{ m})^2 = 2.6 \text{ m}^2$ so the overall sensitivity, with the current set of daughterboards and initial observations, is $95 \text{ photons m}^{-2}$ in $\leq 3 \text{ ns}$.

The threshold will likely improve for future observations if/when new daughterboards are installed. Assuming that their higher gain amplifiers produce $V_{\text{pe}} = 50 \text{ mV}$ and we can use a $V_{\text{trigger}} = 150 \text{ mV}$, the improved system will trigger on ~ 3 photoelectrons and will have an overall sensitivity of $17 \text{ photons m}^{-2}$ within $\leq 3 \text{ ns}$ in the sensitive band.

Chapter 6

PulseNet – Design and Implementation

During the early design phase of the all-sky survey (c. 1999-2000), we considered the electronics challenge of digitizing 1024 photodetector outputs at gigahertz speed, looking for a pair of those signals that are simultaneously above a certain threshold, and setting a switch to steer the coincident signals into a memory. The level of parallelism in this design leads to an enormous data rate: 3.5 Tb/s, the equivalent of the contents of all books in print, every second. A primary limitation in meeting this challenge was in high-speed chip-to-chip communication. If the digitizing of matched input pairs was done in different chips, then a thousand or so ~ 1 Gb/s signals would have to be piped around a printed circuit board to other chips that detect coincident pulses, which in turn would need to set fast switches to stream the digitizer outputs from the coincident inputs into memory. Implementing this scheme with the speed necessitated by the intrinsic qualities of photomultiplier tubes was simply not feasible with commercially available chips communicating on printed circuit boards. The solution was to do the difficult parts—parallel digitizing, coincidence detection, fast switching of unusual and infrequent signals into memory—on a full-custom chip where the timescales are intrinsically much shorter and where the high-speed digital communication can occur on traces with much smaller loads. Digital communication with this custom chip could be done at the leisurely pace of microcontrollers. It was in this context that PulseNet was conceived.

6.1 Overview of Design

The primary purpose of PulseNet is to detect and record coincident signals from 16 matched pairs of PMT outputs, as shown in Fig. 6.1. These analog signals are digitized at up to 1 GS/s by comparing each to seven external voltages ($V_{\text{ref}}[6:0]$) on the rising edges of two interleaved clocks (*fastclocka/fastclockb*). The resulting 7-bits/pixel/clock of thermometer code are encoded to 3-bits/pixel/clock and are delayed by an 8-bit deep array of 2-phase shift registers. Meanwhile, a coincidence trigger circuit looks for coincident pulses in a matched input pair: it selects one thermometer code bit from each of the 32 PMT signals and looks for instances when an input pair simultaneously exceed a given threshold—a “coincidence” (e.g., pixels 14A and 14B both exceed $V_{\text{ref}}[3]$ on a rising edge of *clka*). Coincidences cause the rest of the chip to spring into action; specifically, they trigger switches that steer the streaming samples from the coincident input pair into 256-bit long shift registers. The SETI I/O controller then passes the waveforms and coincident pixel pair address off chip.

PulseNet’s parallel and independent function is to measure “countrates”—the number of pulses exceeding a certain V_{ref} in a time interval. With the appropriate V_{ref} , a countrate is proportional to the photon flux on the PMT pixel, i.e. it measures the brightness of the star on that pixel. Countrate measurements are orchestrated by the astronomy I/O controller, which sets switches to select a pixel pair and thermometer code bit for a particular measurement, starts and stops four 32-bit ripple counters, and passes the counts off chip.

PulseNet contains $\sim 250,000$ transistors and was fabricated through the MOSIS Corporation on TSMC’s

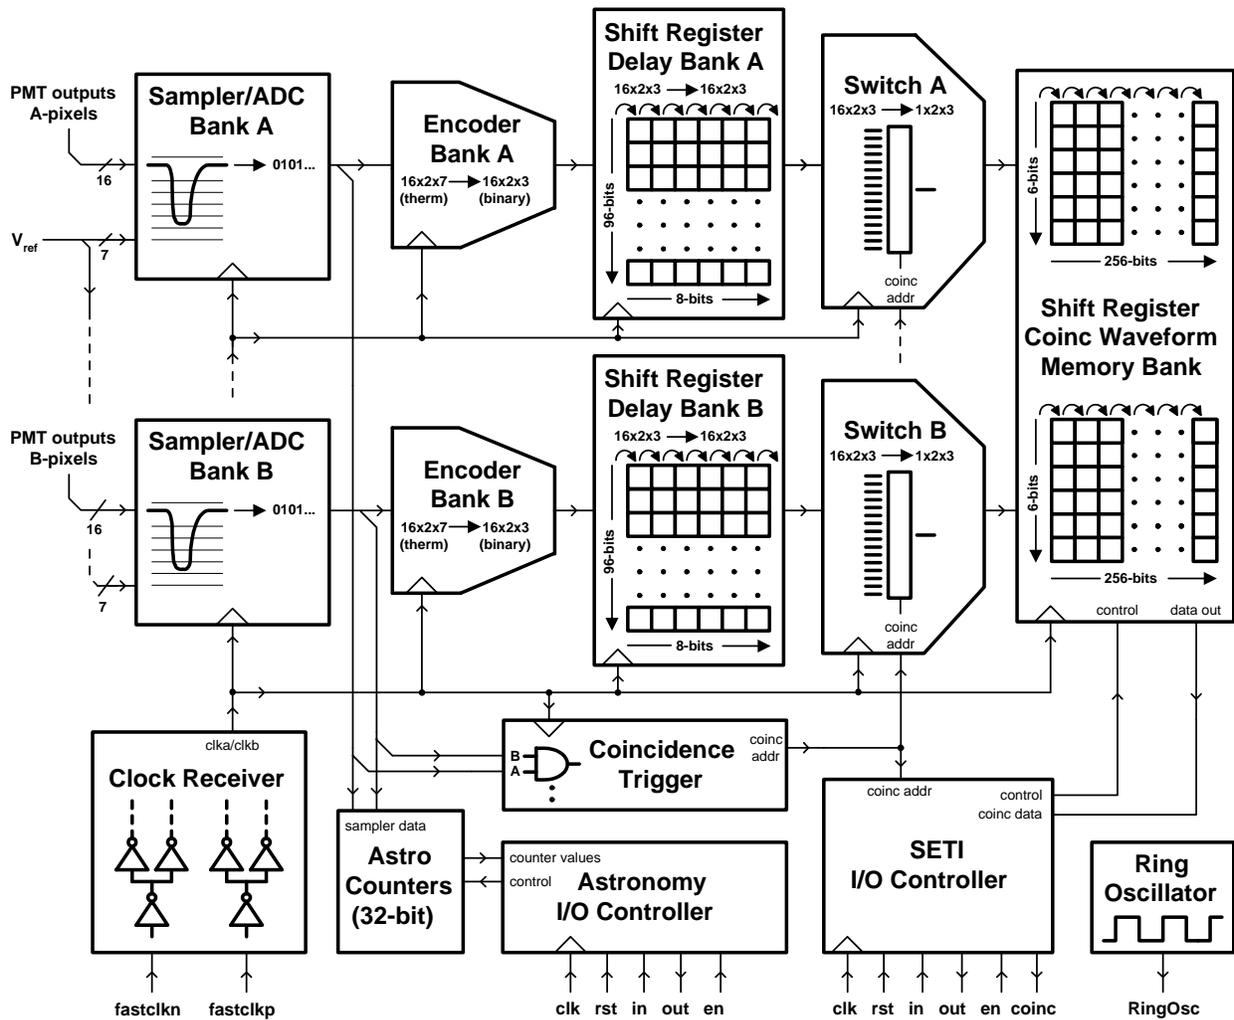


Figure 6.1. Simplified block diagram of PulseNet. Compare with Fig. 6.2, the block diagram showing major blocks and top level signals as they are named in the PulseNet design files.

Table 6.1. PulseNet – summary of capabilities

Purpose
Digitize 32 analog inputs at up to 1 GSps
Detect coincident pulses in 1 of 16 matched analog input pairs
Store 512-bit long sample of coincident signals
Measure number of times an input exceeds a voltage threshold (astronomy countrates)
SETI Capabilities
Detect coincident pulses in matched input pair)
Store 512 samples (including “pre-trigger” samples)
Astronomy Capabilities
Measure countrates on 1 of 7 voltage thresholds ($V_{\text{ref}}[6:0]$) using four 32-bit counters (for both clock samples (a/b) on one pixel pair (A/B))
Countrates proportional to photon flux on PMT pixels
Analog Samplers
32 flash analog to digital converters
Each compares input to 7 voltage references ($V_{\text{ref}}[6:0]$) on the rising edges of two interleaved ≤ 500 MHz clocks for ≤ 1 GSps
Memories
12-bit wide/256-bit deep shift register memory for storing coincident waveforms
192-bit wide/8-bit deep shift register memory for delaying all waveforms prior to coincident pair trigger
Data rate
~ 100 Gb/s per PulseNet – ~ 3.5 Tb/s in all-sky survey
Miscellaneous
$\sim 250,000$ transistors
Fabricated on TSMC $0.25 \mu\text{m}$ process through the MOSIS Foundation
$3.1 \text{ mm} \times 3.1 \text{ mm}$ chip packaged in an 84-pin ceramic leadless chip carrier (CLCC)

$0.25 \mu\text{m}$ CMOS process. The chip measures $3.1 \text{ mm} \times 3.1 \text{ mm}$. All circuits were full-custom designs, with the exception of the three synthesized state machines (`seti_io`, `astro_io` (a module within `astronomy`), and `memcontroller`). At 400 MHz and 2.5 V (standard operation) PulseNet dissipates 1.1 W, but has been shown to work at as high as 500 MHz and 2.87 V. PulseNet is fully functional and thirty-two of the chips are used in the all-sky optical SETI experiment (Chapter 5). A summary of PulseNet’s capabilities is given in Table 6.1.

6.1.1 Notation

In the chapters relating to PulseNet, **sans serif type** will be used to indicate the names of signals and circuits that appear in the PulseNet design. Many of these names are unusual, and might otherwise look like typographic mistakes.

There are several types of signal parallelism in PulseNet. A consistent notation is essential to keep things straight. As is standard practice, a signal that is $(n + 1)$ -bits long will be written “`signal[n:0]`”, and the m th bit of that signal is “`signal[m]`”. The blocks (groups of repeated circuits) in PulseNet also use this notation. For example, the third instance of the block `sampencbank` is written `sampencbank[3]`. When referring to an arrayed signal that comes from an arrayed block, there is potential confusion regarding what the index refers to. In this case, the block numbers are moved outside the square brackets. For example, `cp_add3[3:0]` is the signal `cp_add[3:0]` from `sampencbank[3]`.

Another type of parallelism has to do with analog inputs coming in matched pairs, e.g. `in00A` and `in00B`.¹ In general, capital A and B refer to matched input pairs. Additionally, PulseNet also uses two interleaved

¹A-channel inputs come from PMTs in the left focal plane and B-channel inputs are from PMTs in the right focal plane.

clocks to achieve a high sampling rate. Signals relevant to one of the clocks will include lowercase **a** or **b** in their names, e.g. `clka` and `clkb`. Note that clock references are always lowercase, while pixel pair references are uppercase. To refer to the set of all four combinations of pixel pairs and clocks, superscripts and subscripts are used, e.g. `memoutBBAa`.

These notations may all be used in combination. For example, the `membank` module accepts encoded samples from all four `sampencbank` models. Each sample is 3-bits wide for each combination of clock and pixel pair member. This set of 48 signals is compactly written as `GBBAa0..3[2:0]`.

6.2 Circuits

6.2.1 Top-level design

The top-level block diagram for PulseNet is shown in Fig. 6.2. Unlike Fig. 6.1 (the simplified block diagram), Fig. 6.2 shows the top level blocks and signals organized as they are in the actual implementation of PulseNet.

Thirty-two analog inputs are sampled in four `sampencbank` modules or “blocks”. Within each block, there are fourteen comparators per input, one for each of seven voltage references (`Vref[6:0]`) for both clocks (`clka`, `clkb`). As described below, the comparators produce 7-bit thermometer code, where each bit is the 1-bit comparison between the input and voltage reference. The thermometer code output follows three paths. One path is immediately encoded from 7-bits per sample to 3-bits per sample, and is delayed in an 8-bit long shift register. In the second path, one of the seven bits is picked off in a mux from each input and for each clock. The data that are selected by the mux correspond to the one high bit in `c_coinc[6:0]`. These bits go into a coincidence detection circuit (an array of AND gates) which produces the signal `coinc_addba[3:0]`, the address of the coincidence pixel with that block. The third path for the streaming thermometer code bits is used for the astronomy mode. Another set of muxes (addressed by `c_ast[6:0]` for threshold number, `pix_blk_ast[3:0]` for block number, and `pix_add_ast[3:0]` for pixel number within the block) pick off four streaming comparator outputs and send them to the astronomy module where the number of samples that exceed the the address threshold are counted in four 32-bit counters.

When a coincidence is detected in one of the `sampencbank` modules, the message propagates to `membank` on `coinc_addba[3:0]`, which contains twelve 256-bit deep shift register memories (3-bit samples for `BBAa`) and a state machine controller called `memcontroller` that latches the coincident pixel address, informs the `sampencbank` modules which input samples they should stream on `GBBAa[2:0]`, steers those samples into the memories, and communicates with the `seti.io` module.

PulseNet has two clock inputs, `fastclocka` and `fastclockb`, which are received, fanned-out, and repeated to drive nearly all of the modules shown in Fig. 6.2. A 29-stage ring oscillator operates independently of all other circuits in PulseNet. The back-to-back inverters of `ringosc` oscillate at ~ 1.6 MHz; the exact frequency is an excellent probe of the process parameters associated with manufacturing the chip, and of the on-chip temperature.

6.2.2 Analog samplers

Fig. 6.3 shows the sampling scheme in greater detail. An analog input (V_{in}) is compared with `Vref[6:0]` on the rising edges of the interleaved `clka` and `clkb`, producing two 7-bit thermometer-coded outputs: `Therm_a[6:0]` and `Therm_b[6:0]` (upper left pane). Note that `Vref[6:1] < Vin` since PMTs produce negative pulses; in order to detect V_{in} samples that exceed its bias voltage (an artifact of spurious PMT signals), we set `Vref[0] > Vin`. The right pane of Fig. 6.3 shows V_{in} and `Vref[6:0]` wired to comparators that produce the thermometer coded samples. Note that the comparators for `Vref[0]` are wired in the opposite sense, that they detect $V_{in} > V_{ref}[0]$. As an example, the largest amplitude samples of V_{in} on the left have `Therm_a[6:0] = Therm_b[6:0] = 0011110`.

Fig. 6.4 shows one of the comparators in greater detail; V_{in} is compared to `Vref` on the rising edges of `clka` and `clkb` in modified `strongarm`² sense amplifiers, whose outputs drive a second stage of reduced-size sense amplifiers (improving gain and reducing hysteresis), followed by RS latches. The sense amplifier [17, 43] is

²They are called `strongarm` because the design was originally used as a flip-flop in the StrongArm microprocessor.

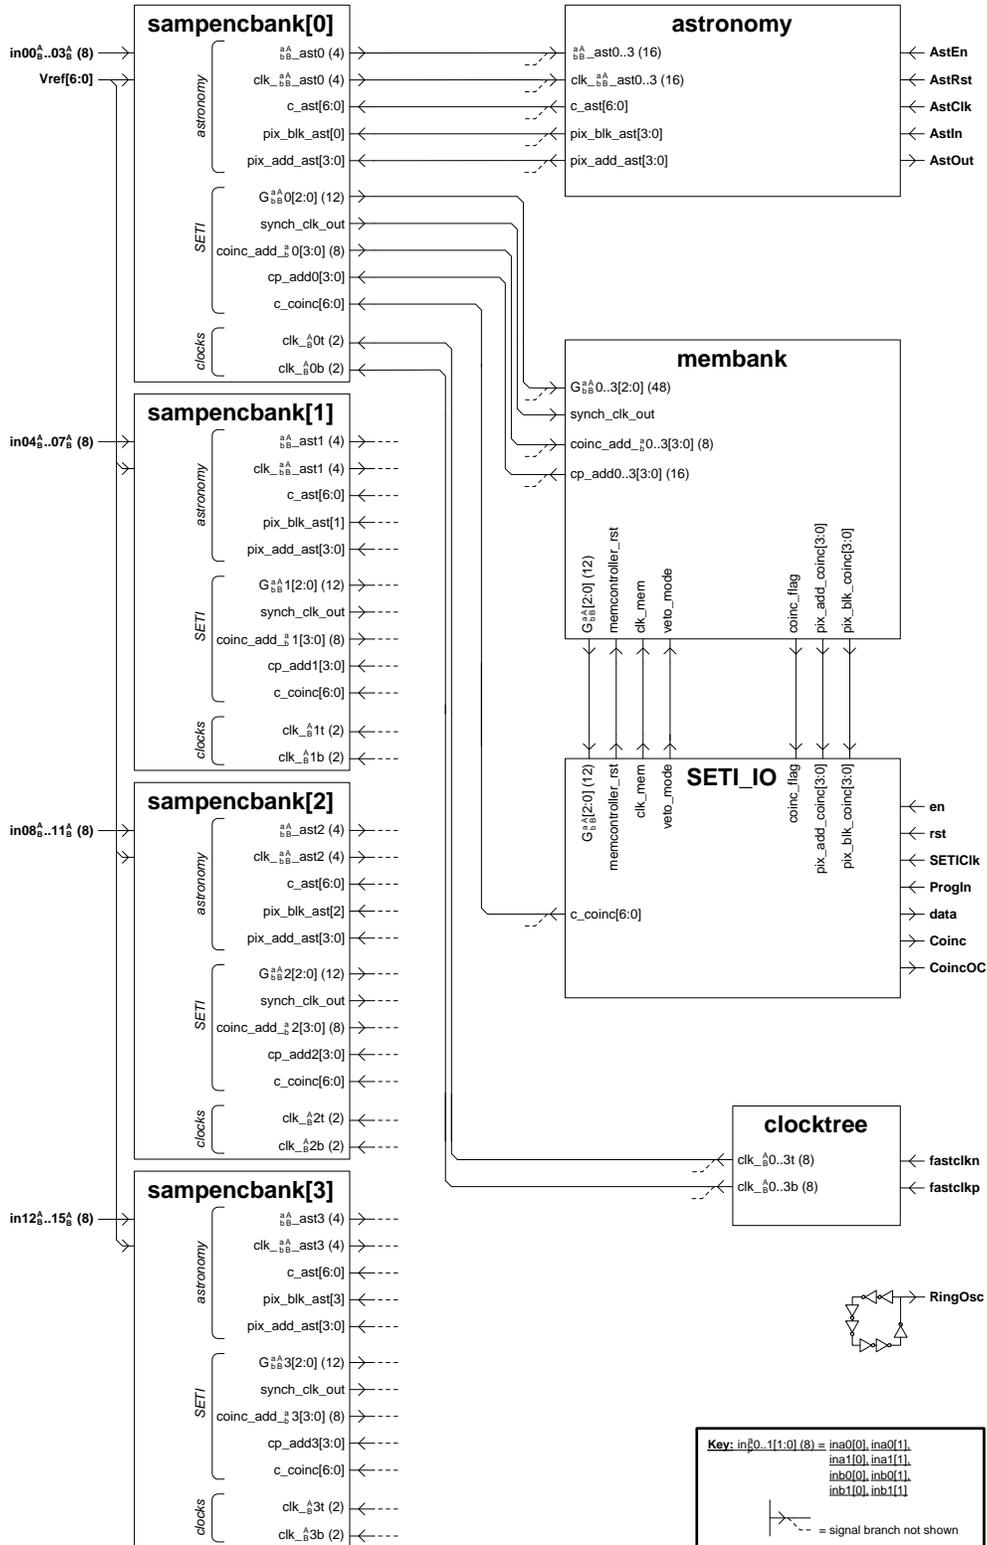


Figure 6.2. Block diagram showing the major blocks and top level signals as they are named in the PulseNet design files. Blocks and signals in the diagram, such as **seti_io** and $\text{G}_{bb}^{aa}[2:0]$, are indicated in the text by sans serif type. Compare with Fig. 6.1, the simplified block diagram.

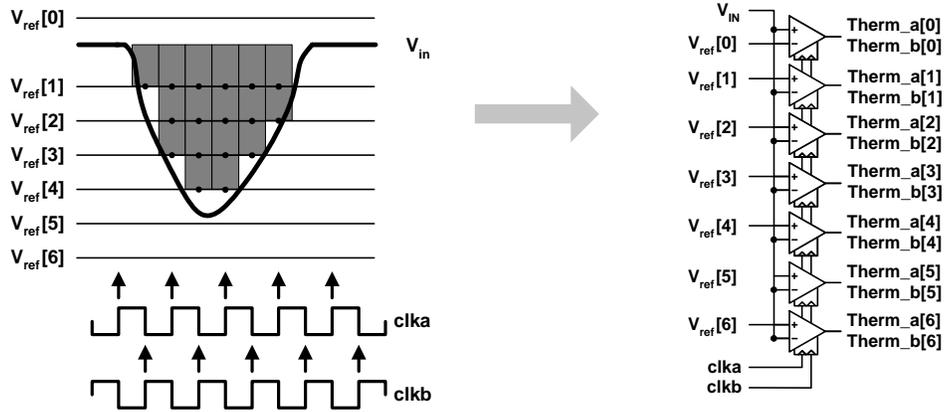


Figure 6.3. Thermometer code sampling scheme.

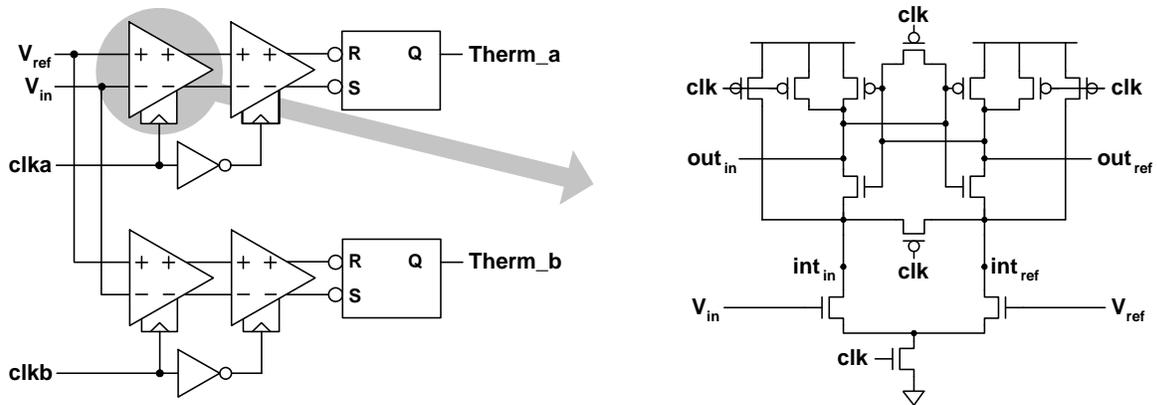


Figure 6.4. Flash converters and sense amplifier details. The comparators in Fig. 6.3 are shown on the left; for each clock phase there are two cascaded StrongArm sense amplifiers followed by a latch (left). Circuit details for the sense amplifier are shown on the right.

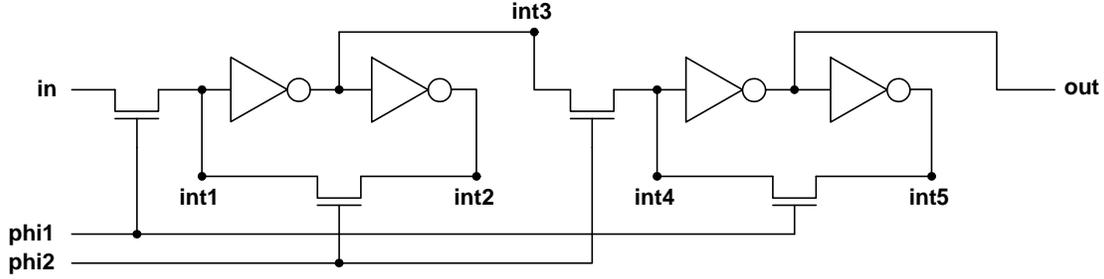


Figure 6.5. Shift register memory unit cell – srcell.

a clocked, regenerative, gate-isolated differential amplifier (right pane of Fig. 6.4). The design provides good input isolation and a small aperture time (tens of ps) that can resolve small voltage differences (tens of mV). The 448 sense amplifiers on PulseNet are grouped in four blocks with staggered clocking to mitigate power supply noise. The output nodes of **strongarm** (out_{in} and out_{ref}) are pre-charged high when clk is low. The input pair (V_{in} and V_{ref}) converts the differential input voltage into a differential current, which is integrated on int_{in} and int_{ref} . When the source node corresponding to the high input, say int_{in} , reaches a threshold drop below V_{DD} , the transistor above this node begins conducting, transferring charge imbalance to out_{in} and out_{ref} . This starts the regenerative action of the cross-coupled inverters at the top and the circuit quickly latches the state.

6.2.3 Clock distribution

PulseNet is clocked by **fastclocka** and **fastclockb**, which are complementary. PulseNet was designed to work at $f_{fastclk} = 500$ MHz and this was later demonstrated under conditions of $V_{DD} > 2.5$ V.

The clocks are dc-coupled on inputs that are weakly biased to the midpoint. The inputs feed an “inverter horn”, a circuit that exponentially amplifies the output drive strength using a chain of successive inverters where the inverter size increases by a factor of 2–4 with each stage. The inverter horns feed repeaters, whose outputs are destined for the **sampencbank** modules. These outputs are delayed, however, by inserting 2, 6, 10, and 14 inverters for the clocks that go to **sampencbank**[0], **sampencbank**[1], **sampencbank**[2], and **sampencbank**[3], respectively. This offsets the clocks in successive samplers by ~ 200 ps so as smooth out the load on I_{DD} . Spikes in I_{DD} —caused, for example by all of the samplers clocking synchronously—will cause ground bounce ($\Delta V = L dI/dt$) because of inductance in the bond wires that connect power and ground to the chip.

For the ~ 5 mm long wires that connect the clock drivers to the four **sampencbank** modules, there are repeaters every ~ 1 mm to boost signal amplitude. The clock lines are equalized in length and also shielded from parallel signal wires by ground planes to avoid coupling to those signals.

6.2.4 Memory and coincidence circuitry

There are two main memories on PulseNet. They are both based on cascading the two-phase shift register cell in Fig. 6.5 and clocking it with the two-phase clock generator in Fig. 6.6. The first memory delays for eight clock cycles the 3-bit encoded samples from every sampler. Since the shift register delays samples taken with both $clka$ and $clkb$, it continuously stores the previous sixteen 3-bit samples, for every sampler on the chip. This gives the coincidence detection circuit time to work and still allows for pre-trigger samples for coincident waveforms.

The second memory stores the samples of coincident waveforms. When triggered, **memcontroller** sets muxes that steer 3-bit samples taken on both clock phases from the coincidence input pair into four 3-bit-wide 256-bit-deep shift registers. The **memcontroller** module prevents the shift register from overflowing by stopping the clock after 256 cycles. It waits to start clocking the shift register again told to by **seti_je**.

In Fig. 6.6, The two-phase non-overlapping clock generator, **2phasegen**, takes the clock ϕ as its input and generates the complementary phase ϕ_{bar} with an inverter. The outputs, $\phi1$ and $\phi2$, are guaranteed

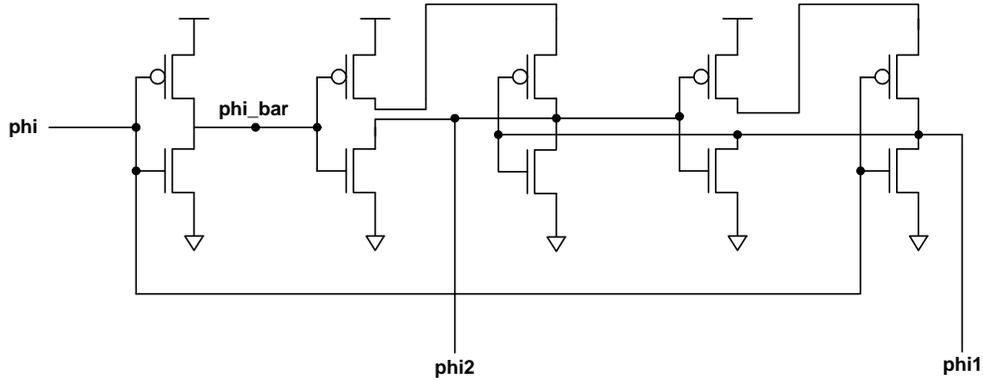


Figure 6.6. Two-phase non-overlapping clock generator – 2phasegen.

to be non-overlapping. This is because `phi1` for example, can only go high when the inverter that drives it is powered by a p-channel fet that turns on only when `phi2` is low.

The non-overlapping phase relationship between `phi1` and `phi2` makes `scell` robust to data racing through. In Fig. 6.5, first note that `scell` has two storage nodes: `int1` and `int3`. In the rising edge of `phi1` the state of `in` is transferred to `int1`, and the complement of `in` appears on `int3`. On the next rising edge of `phi2`, which is guaranteed to be after the falling edge of `phi1`, the state of `int3` is transferred to `int4`, and the complement of `int4` appears on `out`. The net result of the two inversion is that the data passes from `in` to `out` on one full cycle of `phi1`.

6.3 Interface with `seti_io` module

The `seti_io` module is a state machine and the interface for SETI-related tasks. It performs two main functions: it receives programming data and sets registers affecting the whole chip, and it transfers to the outside world the coincident event waveforms and related data recorded by other modules.

The `seti_io` module has seven I/O pins: an enable (`en`), a clock (`SETIclk`), a reset (`rst`), a data input (`ProgIn`), a data output (`data`), and two coincidence pins (`Coinc` and `CoincOC`, which is an open-collector version of `Coinc` that is not used). The module must be enabled (`en` high) for any I/O. It processes information on the positive and negative edges of `SETIclk`: input (on `ProgIn`) must be valid on positive edges, and output from the `seti_io` channel of PulseNet is valid on negative edges of `SETIclk`.

The programming sequence for `seti_io` is listed in Table 6.2. The module must first be initialized with a low-high-low sequence on `rst` in order to reset internal registers that keep track of the programming state (e.g. `mode`, `IO_counter`, and `memcount`). The module is then programmed with the following five bits that completely set the state of PulseNet for programming tasks:

- `thresh_enc[2:0]` – the encoded threshold address; this is decoded into a 1-of-7 address, `c_coinc[6:0]`, which selects a $V_{ref}[6:0]$ as the coincidence threshold.
- `veto_mode` – the veto mode state; when enabled (`veto_mode = 1`), PulseNet will reject coincidence events that have a waveform sample $> V_{ref}[0]$ during the the first 16 samples).
- `clock_half` – the clock half state; when enabled (`clock_half = 1`), PulseNet enters an unusual mode where the samplers for odd-numbered pixels are not clocked. The mode was intended as a way to run on lower power, but is never used in practice.

The `seti_io` module responds by echoing the decoded, 1-of-7 threshold address (`c_coinc[6:0]`), and the veto mode and clock half states. Echoing verifies that PulseNet latched and properly decoded the programming data. This is a nice verification for daily operation, and was extremely useful for initial testing. At this point in the sequence, `seti_io` is fully programmed, has reset the main memory (`membank`), has set the coincidence

threshold level (`c_coinc[6:0]`) and various other internal registers. It switches to `mode=1` (another state machine variable) and is ready for a coincident event.

When such an event is received, data from samplers has been completely processed and stored in `membank` by the time `seti_io` is notified on the internal node `coinc_flag`. The `seti_io` module responds by setting the `Coinc` pin and sequentially passing data regarding the coincidence, as shown in Table 6.2. These data are: the address of the coincident pixel pair in the form of two 1-of-4 addresses (`coinc_blk_add[3:0]`, the `sampencbank` number that registered the coincidence, and `coinc_pix_add[3:0]`, the pixel within that block), the decoded threshold address, veto mode and clock half states, as before, and the 3072 samples in the waveforms from the coincident pixel pair ($\text{memout}_{Bb}^A[2:0] \times 256$).

The `seti_io` module must be reset and reprogrammed before it is sensitive to additional coincidences.

6.4 Interface with astronomy module

The `astronomy` module handles I/O for astronomy-related tasks. It contains four 32-bit ripple counters and the `astro_io` module, which functions in a similar manner to `seti_io`. The `astronomy` module performs one main function (through several steps): it measures counts. Counts are simply the number of times that the samples for a particular pixel pair are equal to or greater than a particular threshold level. The pixel pair is chosen with `pix_add_enc[3:0]` and the threshold level is chosen with `thresh_enc[2:0]`. The four counters are for the clock edge (a/b) and pixel pair (A/B) combinations. Setting `pix_add_enc[3:0]` and `thresh_enc[2:0]` causes muxes in the `sampencbanks` to steer particular thermometer code bits from one pixel pair into lines that directly feed the counters.

The `astronomy` module has five I/O pins: an enable (`AstEn`), a clock (`AstClk`), a reset (`AstRst`), an input (`AstIn`), and an output (`AstOut`). Like `seti_io`, it must be enabled (`AstEn` high) for any I/O, and it processes information on the positive and negative edges of `AstClk`.

The programming sequence for `astronomy` is similar to the sequence for `seti_io`, and is listed in Table 6.3. The module must first be initialized with a low–high–low sequence on `AstRst` in order to reset internal registers that keep track of the programming state (e.g. `IO_counter`, and `count32`)³.

The module is then programmed with the following seven bits:

- `thresh_enc[2:0]` – the encoded threshold address; selects one `Vref[6:0]` as threshold for astronomy measurements. Note that this threshold address is independent of the one used for `seti_io`.
- `pix_add_enc[3:0]` – the address of the pixel pair whose inputs are to be counted.

It responds by echoing the addresses of the threshold (`c_ast[6:0]`, which is 1-of-7) and the input pair (`pix_blk_ast[3:0]` and `pix_add_ast[3:0]`, each of which are 1-of-4 addresses). At this point, `astronomy` is fully programmed and has set `c_ast[6:0]` and `pix_blk_ast[3:0]`/`pix_add_ast[3:0]` so that the appropriate sampler outputs will be counted. On the next rising edge of `AstClk`, `astronomy` enables its four counters, which increment on rising edges of the selected sampler thermometer-coded output, until `AstClk` is toggled again, when it stops the counters. The values on the four counters are passed out on the next 128 clock cycles. The module must be reset and reprogrammed to measure additional counts.

6.5 Layout

The pin assignments and a die photograph of PulseNet are shown in Figs. 6.7 and 6.8, respectively. It is important that the `sampencbank` blocks be close to the chip’s edges so as to minimize on-chip wire length for the analog inputs. Some of the details of `sampencbank` are labeled in the instance along the bottom edge of the chip in 6.8. Just above the pins are eight vertical rectangles, each with fourteen horizontal structures. Each rectangle is a 2-phase 7-level sampler, with the circuit topology shown in the right pane of Fig. 6.3. Signals mostly flow from the edges of chip toward the center. The 7-bit thermometer-coded samples are encoded to down to 3-bits per sample and digitally delayed in the `shiftmem` modules. After identification in

³For unknown reasons, this module sometimes does not reset with just one of these reset sequences, but always does for multiple reset sequences, which are now part the camera’s control software.

Table 6.2. PulseNet – programming sequence for the seti_io interface

pins				state registers			comments
rst	Progn (input)	data (output)	Coinc	mode	IO_counter	memcount	
high	–	low	low	0	0	0	<u>reset module</u>
low	–	low	low	0	0	0	reset internal registers
							<u>program module</u>
low	thresh_enc[0]	low	low	0	0	0	3-bit threshold address
low	thresh_enc[1]	low	low	0	1	0	"
low	thresh_enc[2]	low	low	0	2	0	"
low	veto_mode	low	low	0	3	0	veto mode
low	clock_half	low	low	0	4	0	clock half mode
low	–	c_coinc[0]	low	0	5	0	echo 1-of-7 threshold addr
low	–	c_coinc[1]	low	0	6	0	"
low	–	c_coinc[2]	low	0	7	0	"
low	–	c_coinc[3]	low	0	8	0	"
low	–	c_coinc[4]	low	0	9	0	"
low	–	c_coinc[5]	low	0	10	0	"
low	–	c_coinc[6]	low	0	11	0	"
low	–	veto_mode	low	0	12	0	echo veto mode
low	–	clock_half	low	0	13	0	echo clock half mode
							<u>receive coincidence</u>
–	–	–	high	1	0	0	asynchronously raise Coinc
							<u>output coincidence data</u>
low	–	coinc_blk_add[0]	high	1	0	0	1-of-4 block address
low	–	coinc_blk_add[1]	high	1	1	0	"
low	–	coinc_blk_add[2]	high	1	2	0	"
low	–	coinc_blk_add[3]	high	1	3	0	"
low	–	coinc_pix_add[0]	high	1	4	0	1-of-4 pixel address
low	–	coinc_pix_add[1]	high	1	5	0	"
low	–	coinc_pix_add[2]	high	1	6	0	"
low	–	coinc_pix_add[3]	high	1	7	0	"
low	–	c_coinc[0]	high	1	8	0	1-of-7 threshold address
low	–	c_coinc[1]	high	1	9	0	"
low	–	c_coinc[2]	high	1	10	0	"
low	–	c_coinc[3]	high	1	11	0	"
low	–	c_coinc[4]	high	1	12	0	"
low	–	c_coinc[5]	high	1	13	0	"
low	–	c_coinc[6]	high	1	14	0	"
low	–	veto_mode	high	1	15	0	veto mode
low	–	clock_half	high	1	16	0	clock half mode
low	–	memout_Aa[0]	high	1	17	0	Aa sample ‡
low	–	memout_Aa[1]	high	1	17	1	"
low	–	memout_Aa[2]	high	1	17	2	"
low	–	memout_Ba[0]	high	1	17	3	Ba sample
low	–	memout_Ba[1]	high	1	17	4	"
low	–	memout_Ba[2]	high	1	17	5	"
low	–	memout_Ab[0]	high	1	17	6	aB sample
low	–	memout_Ab[1]	high	1	17	7	"
low	–	memout_Ab[2]	high	1	17	8	"
low	–	memout_Bb[0]	high	1	17	9	Bb sample
low	–	memout_Bb[1]	high	1	17	10	"
low	–	memout_Bb[2]	high	1	17	11	"
low	–	low	high	1	17	12	loop 256 times to ‡

Table 6.3. PulseNet – programming sequence for astronomy interface

AstRst	pins		state registers		comments
	AstIn	AstOut	IO_counter	count32	
high	–	low	0	0	<u>reset module</u>
low	–	low	0	0	reset internal registers
					<u>program module</u>
low	thresh_enc[0]	low	0	0	threshold address
low	thresh_enc[1]	low	1	0	"
low	thresh_enc[2]	low	2	0	"
low	pix_add_enc[0]	low	3	0	pixel address
low	pix_add_enc[1]	low	4	0	"
low	pix_add_enc[2]	low	5	0	"
low	pix_add_enc[3]	low	6	0	"
low	–	low	7	0	no-op; decode addresses
low	–	high	8	0	no-op
low	–	c_ast[0]	9	0	echo 1-of-7 threshold address
low	–	c_ast[1]	10	0	"
low	–	c_ast[2]	11	0	"
low	–	c_ast[3]	12	0	"
low	–	c_ast[4]	13	0	"
low	–	c_ast[5]	14	0	"
low	–	c_ast[6]	15	0	"
low	–	"	16	0	no-op; set internal registers
low	–	pix_add_ast[0]	17	0	echo 1-of-4 pixel address
low	–	pix_add_ast[1]	18	0	"
low	–	pix_add_ast[2]	19	0	"
low	–	pix_add_ast[3]	20	0	"
low	–	pix_blk_ast[0]	21	0	echo 1-of-4 block address
low	–	pix_blk_ast[1]	22	0	"
low	–	pix_blk_ast[2]	23	0	"
low	–	pix_blk_ast[3]	24	0	"
					<u>start counters</u>
low	–	low	25	0	send start signal to counters
					<u>stop counters</u>
low	–	high	26	0	send stop signal to counters
					<u>read counters</u>
low	–	high	27	0	no-op; set internal registers
low	–	aA_data[0]	27	0	start aA counter output
low	–	aA_data[1]	27	1	
low	–	aA_data[2]	27	2	
		:	:	:	
		:	:	:	
low	–	aA_data[31]	27	31	finish aA counter output
low	–	aB_data[0]	27	32	start aB counter output
low	–	aB_data[1]	27	33	
low	–	aB_data[2]	27	34	
		:	:	:	
		:	:	:	
low	–	aB_data[31]	27	63	finish aB counter output
low	–	bA_data[0]	27	64	start bA counter output
low	–	bA_data[1]	27	65	
low	–	bA_data[2]	27	66	
		:	:	:	
		:	:	:	
low	–	bA_data[31]	27	95	finish bA counter output
low	–	bB_data[0]	27	96	start bB counter output
low	–	bB_data[1]	27	97	
low	–	bB_data[2]	27	98	
		:	:	:	
		:	:	:	
low	–	bB_data[31]	27	127	finish bB counter

Note. — PulseNet programming sequence for astronomy interface. See §6.4.

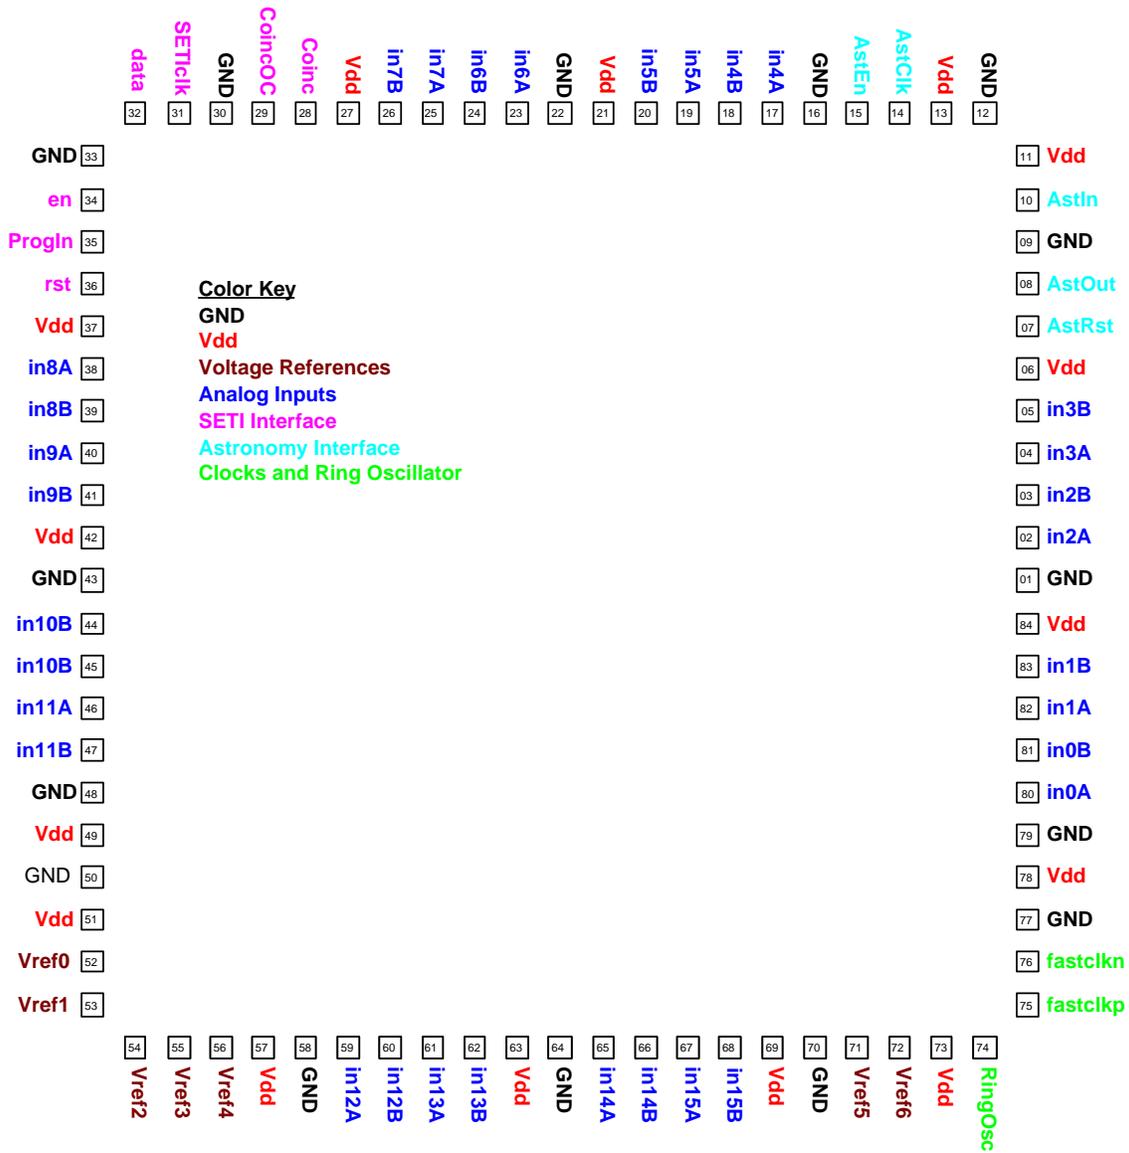


Figure 6.7. PulseNet pin assignments with color key.

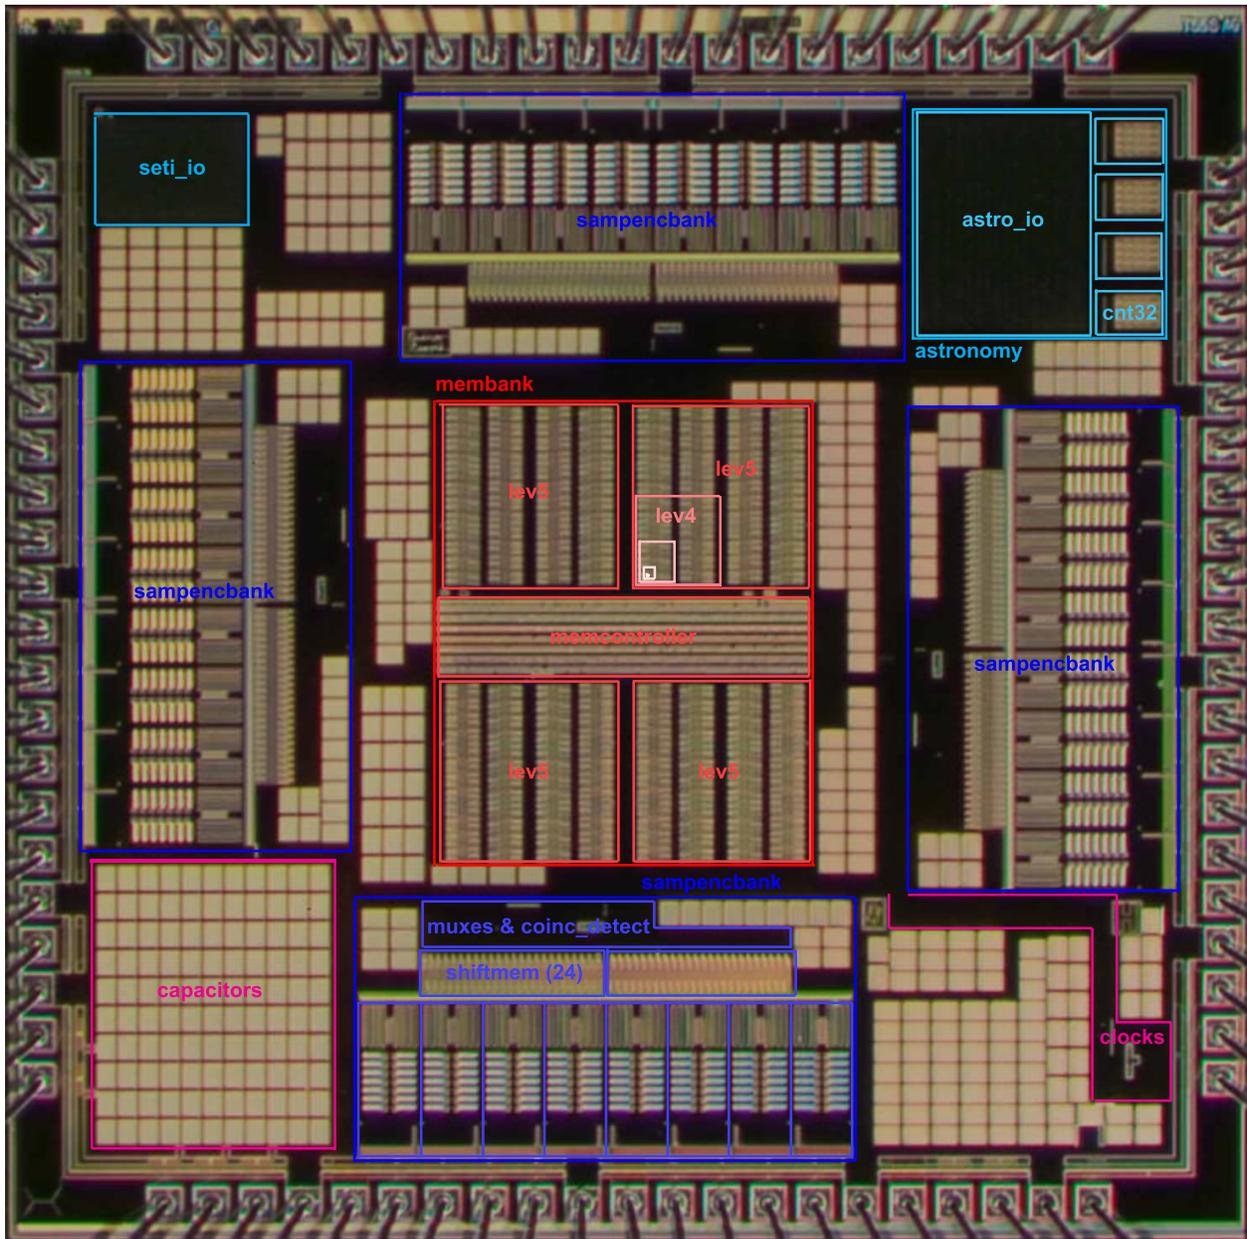


Figure 6.8. PulseNet Rev. 2 die photograph with color-coded block names overlaid.

`coinc_detect` the signals from inputs with coincident pulses are muxed out of `sampencbank` and into `membank`, in the middle of the chip.

The `membank` `membank` module contains four 3-bit wide and 256-bit long shift register memories (`lev5`) and the synthesized state machine (`memcontroller`; its seven long rows of gates are visible in Fig. 6.3) that controls the flow of data into the memories and communicates with `seti_io`. Each 3-bit wide sample is recorded in one `lev5` module. Inside these are the 256-bit-long hierarchical `lev4` memories nested like Matryoshka dolls: four `lev3` per `lev4`, four `lev2` per `lev3`, etc. down to `lev1`, which contains four 1-bit shift-register cells and clock drivers for those cells.

The I/O modules and their pins are relegated to the corners, `seti_io` in the upper left and `astronomy` in the upper right. The four 32-bit counters (`cnt32`) and synthesized state machine (`astro_io`) are visible. (Note that `seti_io` and `astro_io` are less visible in Fig. 6.3 because they are covered in an uninterrupted plane of the top metal layer, in contrast to the other modules which have regular holes in the top metal layer that reflect and diffract light.)

Most of the dark narrow lanes between blocks are full of global signal traces and their repeaters. These lines typically go between the `sampencbanks` and the I/O controllers (upper right and left corners) or the clock driver circuit (lower right corner). Capacitors are ubiquitous in Fig. 6.3; there is a block of them in the lower left and they essentially fill any other free space.

There are five metal layers (`m1`–`m5`) above the transistors and polysilicon in the TSMC 0.25 μm process for routing power, clocks, and signals. “Ground bounce”—the inductive effect from current surges, $\Delta V = L dI/dt$ —was a serious concern for a highly synchronous design like PulseNet so it was important that all circuits have low impedance power and ground connections. (There are fifteen of each.) Because of the need for low resistance, virtually all of `m4` is dedicated to GND and all of `m5` to V_{DD} . The resistance of each stacked via connection ($\sim 20\text{--}30\ \Omega$ for the full stack) is large compared to the resistance of the ground and power sheets ($\sim 0.07\ \Omega$ per square). The other three metal layers were used for signal and clock routing. Global signals and clocks were typically routed on `m3` and local ones on `m2`/`m1`. Modules with many data lines usually had `m1` dedicated to either vertical or horizontal traces and `m2` dedicated to the other.

Because of the worry of signal coupling, the analog reference voltages ($V_{\text{ref}}[6:0]$) are bussed around the chip in a ground-enclosed ring (like a 7-strand coaxial wire) on the outer edge of the chip just inboard of the pads. They are also low-pass filtered at every sampler.

Chapter 7

PulseNet – Testing and Verification

7.1 Testing Procedure

Prior to submitting PulseNet for fabrication, an incremental testing procedure (Table 7.1) was devised to verify functionality of features in order of increasing complexity.

There are two versions of PulseNet. Both versions were generously fabricated on the TSMC 0.25 μm process through grants from the MOSIS Corporation. “Rev. 1” was fabricated in January 2003 and did not fully work. “Rev. 2” was fabricated in May 2005 and thirty-two of these chips are used in the all-sky search.

7.2 PulseNet Rev. 1

7.2.1 Problems

PulseNet Rev. 1 was not a fully functioning chip. In hindsight, it was submitted for fabrication well before design reviews, testing, and simulation inspired high confidence in a working design.

The first problem with Rev. 1 was an apparent V_{DD}/GND short in all chips tested. Although it initially seemed serious, this minor problem was traced to an on-chip design flaw—two of the fifteen V_{DD} and two of the fifteen GND pins (pins 21/22 and 42/43; see Fig. 6.7) were connected to the wrong rail. This was easily fixed by taping Kapton polyimide film over the relevant pin contacts to prevent those pins from electrically connecting in the test socket.

The second problem with Rev. 1 was much more serious and ultimately prevented the chip from being fully tested. While working through the incremental test procedure in Table 7.1, test #6 (the test of SETI coincidence logic) failed because PulseNet failed to complete the programming sequence. Specifically, when programming the SETI module (Table 6.2), it would successfully echo all of the decoded 1-of-7 threshold address bits (`c_coinc[6:0]`), but failed to echo the next bit in the sequence, `veto_mode`, and became unresponsive. The problem was independent of the values of the programming parameters and was reproduced on several chips.

It was later discovered that there was a race condition in `seti_io`. The logic in this state machine was written in Verilog and synthesized from a gate library. In the programming sequence (Table 7.1), when `mode = 0` and `IO_counter = 11`, it should advance to the state (`mode = 0` and `IO_counter = 12`) on the rising edge of `SETIclk`. However, the code to do this was written in such a way that `seti_io` checked the value of `IO_counter`, and, if was 12, it *immediately* incremented `IO_counter`. This set up a race between the gates that check the `IO_counter` registers and those that set them. This is especially problematic because `IO_counter` is a 5-bit register and three of the bits have to change on the transition from 11 to 12 (01011 to 01100 in binary). So when one of the bits in `IO_counter` changes, its value is no longer 11. The other bits of `IO_counter` that should transition may not.

The design flaw was especially unfortunate because simulations using good models for gate delays (Nanosim) verified the malfunctioning behavior after this problem was discovered. (Simple simulations (Verilog and IRSIM) done before fabrication did *not* reveal this problem.)

Table 7.1. Incremental PulseNet testing procedure

Test	Procedure and Comments
1 Smoke test	Power PulseNet with $V_{DD} = 2.5$ V. Look for smoke and elevated I_{DD} .
2 Ring oscillator test	Check for signal. Its frequency should be $f_{\text{RingOsc}} \approx 1.6$ MHz.
3 Clock PulseNet	Slowly increase f_{fastclk} from 50 MHz to several hundred MHz. Watch I_{DD} and make sure that it is linear with f_{fastclk} .
4 Program Astronomy port	Program Astronomy unit and have it echo back the programming data. This will verify that the flip-flops and one I/O port are working.
5 Test samplers	Use variable amplitude square waves and verify functionality. Make sure that the Astronomy module counts these properly.
6 Test SETI coincidence logic	Send large negative pulses to a matched input pair. Verify that coincidence is recorded and that waveforms match input.
7 Test sampler sensitivity	With feedback through the Astronomy port, decrease input amplitude. Note minimum voltage differences that trigger samplers.
8 Check for cross-talk	Verify that large pulses on physically adjacent pins do not trigger.
9 Test with PMT outputs	Verify that PulseNet can trigger on real signals.
10 Verify veto	Verify that the veto mode feature works.

The solution was to rewrite and re-synthesize the `seti_io` and `astronomy` modules so that every variable is stored in two registers, one valid on the negative edge of the relevant clock (denoted with a “_n” suffix), and one valid on the positive edge (“_p” suffix). So, for `seti_io` in PulseNet Rev. 2, there are two versions of `IO_counter`. On a transition from 11 to 12 on the positive edge of `SETIclk` the module checks if `IO_counter_p = 11`. If that is true then it sets `IO_counter_n = 12`. On the next negative edge of `SETIclk`, `IO_counter_p` is set to the value on `IO_counter_n`. Thus checking and setting `IO_counter` still happen in clock cycle, but the two events never race, because they happen at opposite clock edges.

7.2.2 Measurements

The problem with `seti_io` prevented Rev. 1 from being used in the all-sky instrument and from being fully tested. However, a small number of measurements were made.

In test #2 in Table 7.1, f_{RingOsc} was measured at $V_{DD} = 2.4$ V, 2.5 V, and 2.6 V and found to be 1.663 MHz, 1.748 MHz, and 1.828 MHz, respectively. These frequencies are consistent with HSPICE simulations showing that the transistors in that batch of chips are in the “fast-fast” corner. That is, the transition times of n-channel and p-channel transistors are at the low end of the expected distribution.

Measurements of the power consumption (Table 7.2) showed that the current (I_{DD}) is approximately linear with clock frequency (f_{fastclk}), as one would expect from the $P = I_{DD}V_{DD} = f_{\text{fastclk}}CV_{DD}^2$ model of power dissipation (C is the capacitance that has to be charged on the average clock cycle). The quiescent current at $f_{\text{fastclk}} = 0$ MHz is mostly due to other components on the tester board. It is worth noting that the chip was *not* hot to the touch after being clocked $f_{\text{fastclk}} = 500$ MHz for a few minutes.

7.3 PulseNet Rev. 2

The fix to `seti_io` described in §7.2.1 worked, and PulseNet Rev. 2 passed all of the tests in Table 7.2. Measurements of f_{RingOsc} were similar to those from Rev. 1.

Table 7.2. PulseNet Rev. 1 – I_{DD} and I_{DD}' vs. f_{fastclk}

f_{fastclk} (MHz)	I_{DD} (A)	I_{DD}' (A)
0	0.15	0.15
63	0.23	0.20
125	0.30	0.23
250	0.44	0.31
500	0.70	0.46

Note. — PulseNet tester board current (I_{DD} and I_{DD}') as a function of clock frequency (f_{fastclk}). I_{DD}' is the current with half of the samplers disabled (`clock_half = 0`). I_{DD} and I_{DD}' include contributions from chips on the tester board besides PulseNet.

The only aspect of Rev. 2 that has not performed up to specification is the clock interface. With Rev. 1, the chip could be clocked up to $f_{\text{fastclk}} = 500$ MHz at $V_{DD} = 2.5$ V with no problem. For an unknown reason, Rev. 2 has poor input clock coupling at high frequencies. For $f_{\text{fastclk}} > \sim 250$ MHz, the measured fastclock amplitude on the PulseNet pin trended downward for a fixed fastclock power input (using a Hewlett Packard H4677A Signal Generator, which is good up to 1 GHz). It is not clear where this power was going because there was only one substantive change to the clock receiver circuit from Rev. 1 to Rev. 2. Input protection diodes were added to protect against surges, but this would not have added enough capacitance to produce the observed clock coupling.

The result of this problem is that PulseNet cannot be clocked above $f_{\text{fastclk}} = \sim 300\text{-}350$ MHz using the clock driver on the tester board. With signal generator mentioned above, most chips could be clocked up to $f_{\text{fastclk}} = \sim 450\text{-}500$ MHz at $V_{DD} = 2.5$ V, and some chips were clocked up to $f_{\text{fastclk}} \approx 575$ MHz at $V_{DD} = 2.75$ V.

Some additional results from the testing procedure are described below. See also Appendix C, which describes the automated verification procedure to which 52 chips were subjected. (All chips in the all-sky instrument were verified.)

7.4 Test waveforms

Fig. 7.1 shows a PulseNet waveform reconstruction of a fast double pulse with a closely agreeing oscilloscope trace overlaid. For this test, PulseNet was clocked at $f_{\text{fastclk}} = 500$ MHz so that the sampling rate was 1 GS/s.

Fig. 7.2 demonstrates an important capability of PulseNet in the context of optical SETI. It shows the reconstruction of a complicated waveform, the recording of which during an optical SETI observation would be, to put it mildly, momentous. However, the information-containing aspects of this signal would be completely lost on previous optical SETI instruments. The Harvard targeted search, for example, would record the rise and fall times of one of the pulses, and possibly multiple events if the signal was broadcast frequently. The event-driven experiments (e.g. at Leuschner and Lick Observatories) would detect extra “doubles” and “triples,” but if reobservations of the source did not reveal an increased event rate, the original events would probably be attributed to one background or another. Thus, PulseNet provides a unique and important experimental tool — the ability to precisely measure the waveforms that trigger the instrument.

7.5 Input sampler offset voltage measurements

The final set of measurements presented in this chapter are the statistical properties of the sampler offset voltages. V_{offset} measures the extra voltage that must be applied to one sense amplifier input so that the circuit thinks that the inputs are equal. The scatter in this bias translates into scatter in the trigger thresholds for the 1024 PMT pixels, which ultimately increases the rate of background events and limits sensitivity.

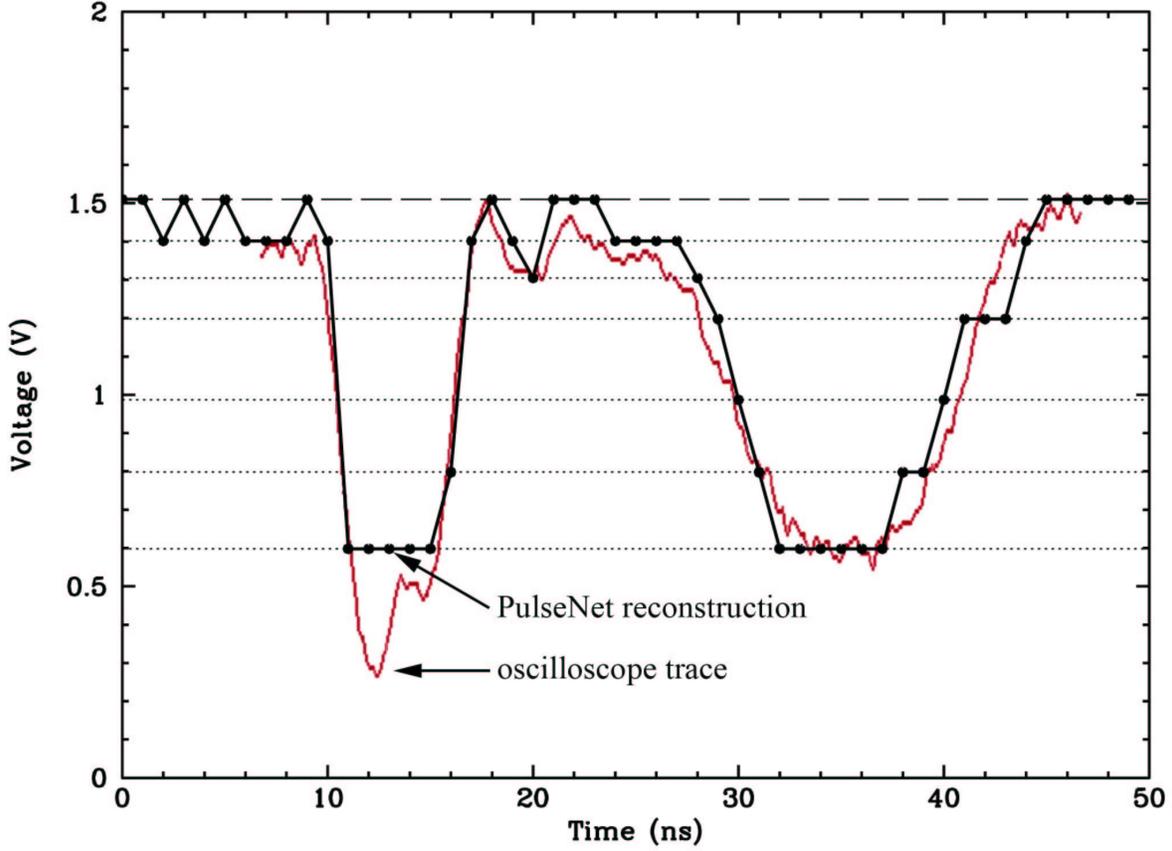


Figure 7.1. PulseNet waveform reconstruction at 1GS/s with oscilloscope trace for comparison.

There are 448 1-bit comparators on PulseNet (16 pixel pairs \times 2 inputs per pixel pair \times 2 clocks \times and 7 thresholds) and 42 chips were tested using the automated procedure described in Sec. C.3 for a total of 18816 measured values of V_{offset} on different sense amplifiers, a statistically-rich data set.

The first step in analyzing this data was to remove the pixel-dependent offsets that resulted from variations in PCB trace lengths and impedances for the samplers' inputs, as well as the consistent offsets due to clock coupling (for inputs whose traces happen to be near the clock traces). This effect is shown in Fig. 7.3 and its removal is illustrated in Fig. 7.4.

Having removed the experimental bias, we can examine the distribution of V_{offset} , shown in Fig. 7.5 with subplots for combinations of pixel pair member (A/B) and clock phase (a/b). Fig. 7.6 is similar, with subplots for the seven threshold voltages. Uncalibrated offsets are shown in red, and calibrated offsets in blue. The $1-\sigma$ variation for the whole dataset is about $\sim 11\text{mV}$, which is consistent with the $\sim 15\text{mV}$ offset for the sense amplifier input pair (Fig. 6.4) predicted by the model in [48]. There appear to be no consistent offsets in the subplots of Fig. 7.5, so we must conclude that the clock phases and input pair members are equivalent, at least in terms of offset voltage.

Fig. 7.6 reveals a consistent small offset for samplers associated with two reference voltages, $V_{\text{ref}}[0]$ and $V_{\text{ref}}[6]$. This makes sense upon considering design choices made when organizing the layout of the samplers arrays in PulseNet. At times it was necessary to run a clock line (which transitions frequently, and has large drivers) next to the voltage references. So as to minimize the impact, traces for the V_{ref} 's were kept in a tight grouping, with $V_{\text{ref}}[0]$ and $V_{\text{ref}}[6]$ on the outside. The sensitivity of the samplers associated with these V_{ref} 's is far less important than those associated with $V_{\text{ref}}[1]$ and $V_{\text{ref}}[2]$, which measure smaller voltage differences. In any case, 3-5 mV is a very small offset.

Fig. 7.7 shows offset voltage variation away from the nominal supply voltage of $V_{DD} = 2.5\text{V}$ and the nominal bias voltage of $V_{\text{bias}} = 1.5\text{V}$ used in Figs. 7.5 and 7.6.

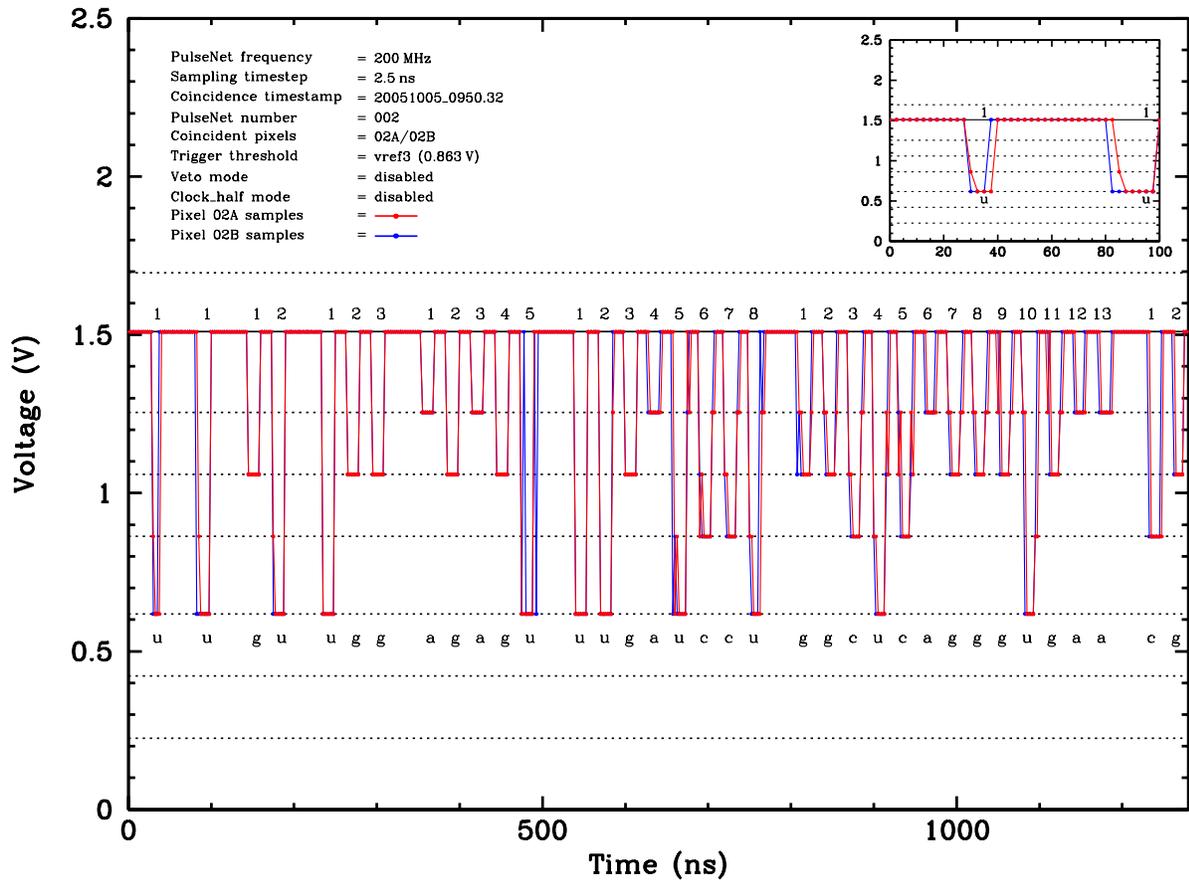


Figure 7.2. PulseNet waveform capture demonstrating the ability to trigger on and resolve the details of a signal with high information content. This test signal shows characteristics that would generate enormous excitement if it were obtained while observing an astrophysical source. Information is encoded to two ways: by the grouping of pulses in the Fibonacci Sequence (1, 1, 2, 3, 5, 8, 13, ...), as indicated by the numbers above the pulses; and in the four pulse amplitudes, a genetic code whose sequence is the beginning of the ribosome genome [69], as indicated by the RNA base abbreviations (a, u, g, c) below the pulses. Such a signal would clearly be of intelligent origin (pulsars don't know integer arithmetic), and would have profound biological implications, possibly demonstrating a common biochemistry in independent evolutions of life.

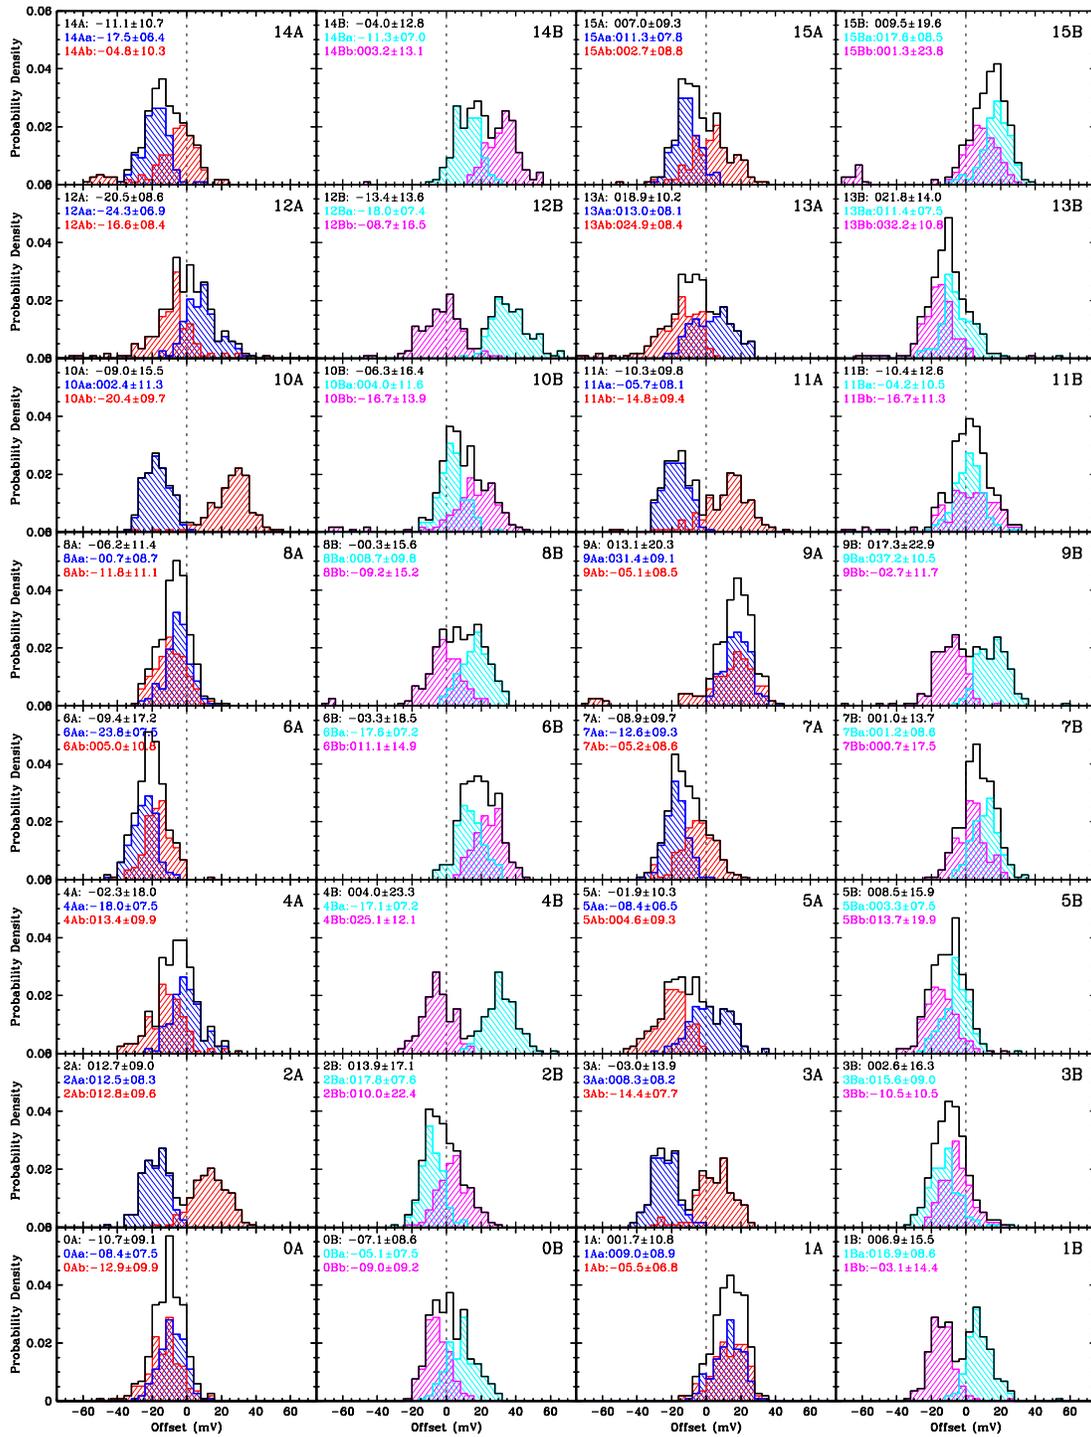


Figure 7.3. Uncalibrated sampler offset voltages for 52 PulseNets shown in 32 panels, one for each of analog input. Each panel shows data from samplers for both clock phases (fastclockb = blue or cyan; fastclocka = red or magenta). The narrow distributions, offset from their common mean, indicate that a substantial portion of the variation in the offset voltage distribution for all pixels (Figs. 7.5 and 7.6) is due to static offsets associated with pixel-dependent signal trace length and impedance variations, as well as timing-dependent clock coupling (e.g. some pixels always have a certain amount of clock coupling at the sampling time).

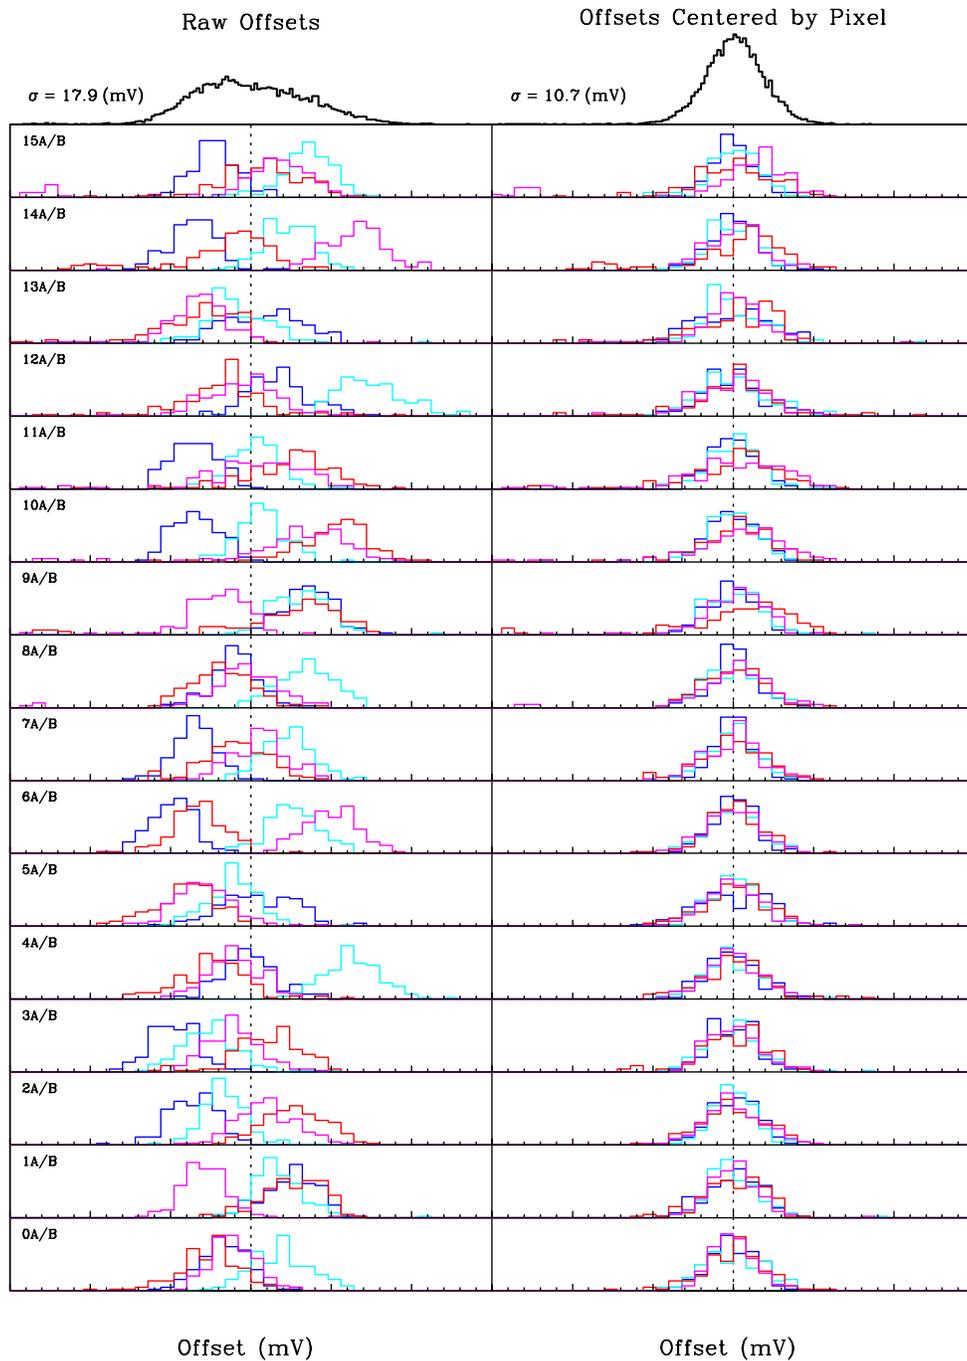


Figure 7.4. Calibrated and uncalibrated sampler offset voltages sorted by pixel. Uncalibrated distributions for all combinations of pixel number (0–15), pixel pair member (A, B) and clock phase (a, b) are shown in the plots on the left, with the sum of those distributions shown in the top left. The plots on the right are same distributions, but with their mean values subtracted (so as to remove offset voltage contributions from the test setup). The resulting calibrated distribution for all pixels is shown in the top right.

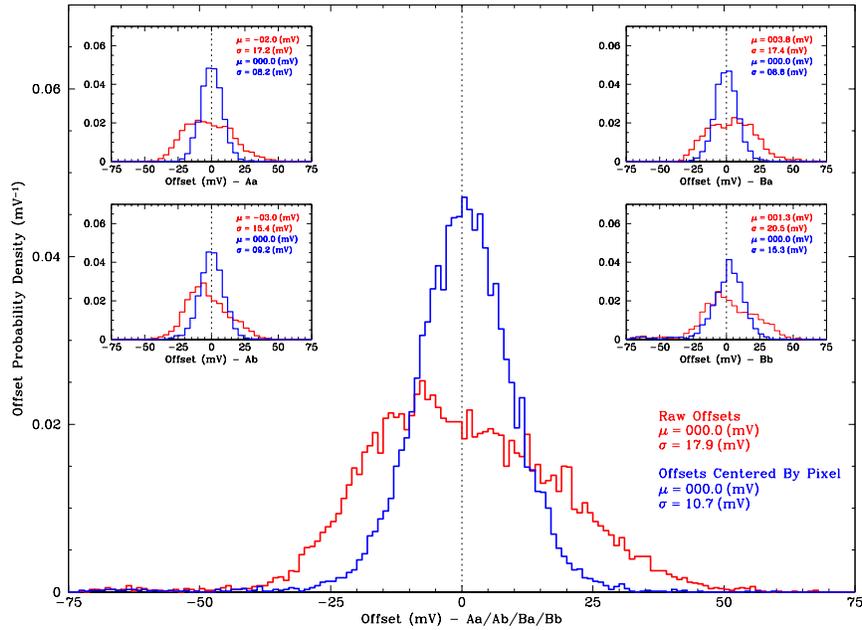


Figure 7.5. Sampler offset voltages sorted by pixel pair member (A/B) and clock phase (a/b) in the subplots and for all samplers in the center. Statistics for calibrated (blue) and uncalibrated (red) offsets are also given.

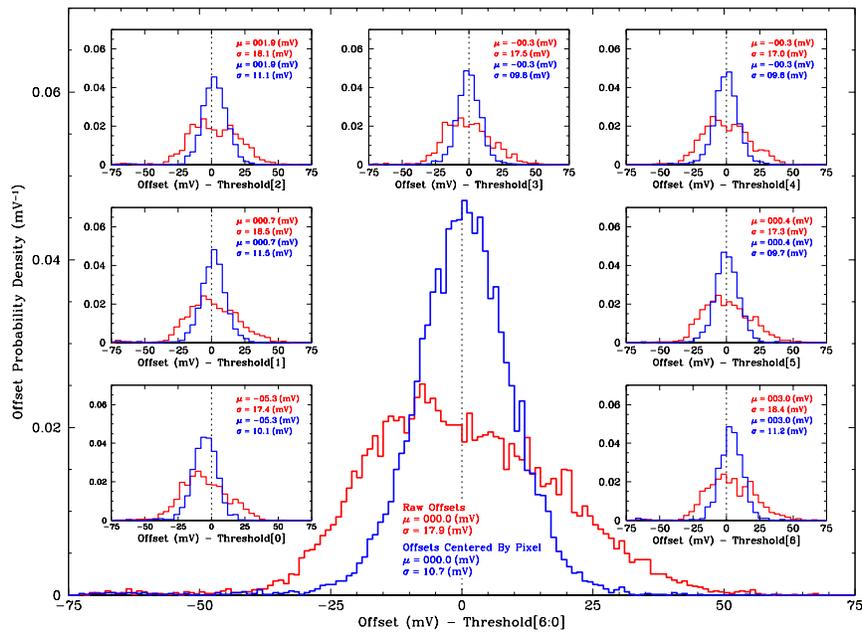


Figure 7.6. Sampler offset voltages sorted by threshold in the subplots and for all thresholds in the center. Statistics for calibrated (blue) and uncalibrated (red) offsets are also given.

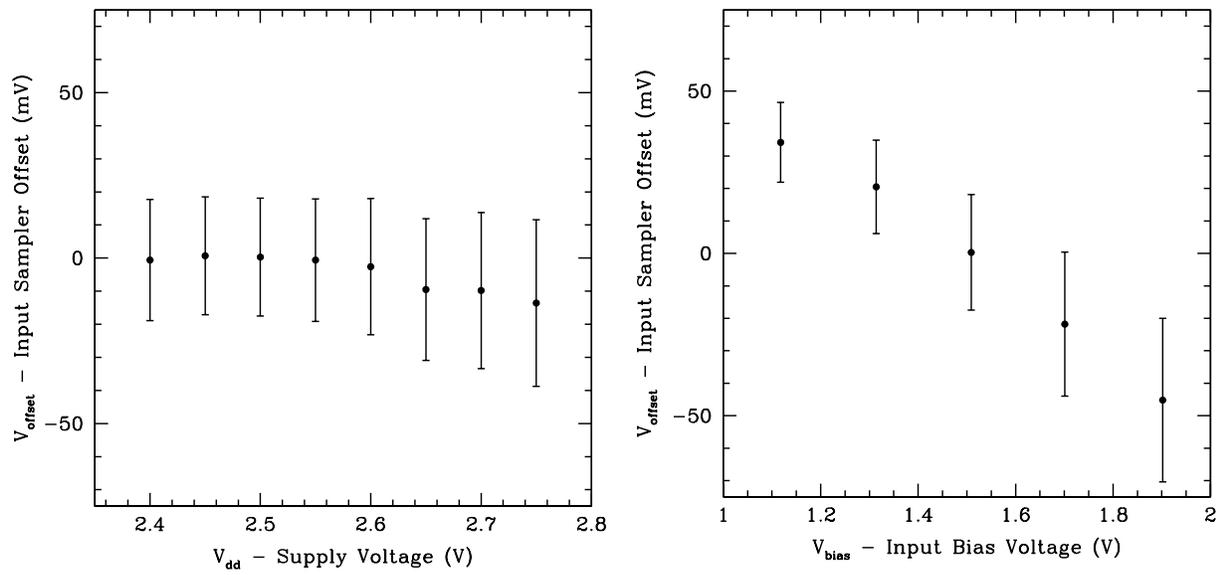


Figure 7.7. Offset voltage (V_{offset}) versus V_{DD} and input bias (V_{bias}). The points on each plot are mean values, and the error bars are $1\text{-}\sigma$ variations, for 448 sense amplifiers on one chip.

Chapter 8

All-Sky Optical SETI Data Analysis

This chapter summarizes sensitivity measurements and calibrations of the all-sky camera, as well as initial observations and their implications. The observations totaled 17 hours over three nights. While these totals are small compared with the targeted search (~ 2400 hours over five 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 targeted search. The initial observations also provided a way to develop general observing procedures and to plan for automated observations.

8.1 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. 8.1, 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. 8.1, the Moon's leading limb grazes the left edge of a PMT at time t_1 . Measurement of t_1 yields the offset between the midline of the PMT array and the local meridian via the Eq. 8.1, which is derived by inspecting Fig. 8.1:

$$\Delta t = t_1 - t_0 + t_M - t_{PMT}, \quad (8.1)$$

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:

$$\Delta d = d_1 - d_0 - d_M - d_{PMT} \quad (8.2)$$

where d_M is the Moon's radius.

These measurements were performed on PMT 4L and PMT 5L for the Lunar transit on 7 May, 2006 yielding $\Delta t = 3^m 13^s (= 0^\circ 48')$ and $\Delta d = -1^\circ 5'$. 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.

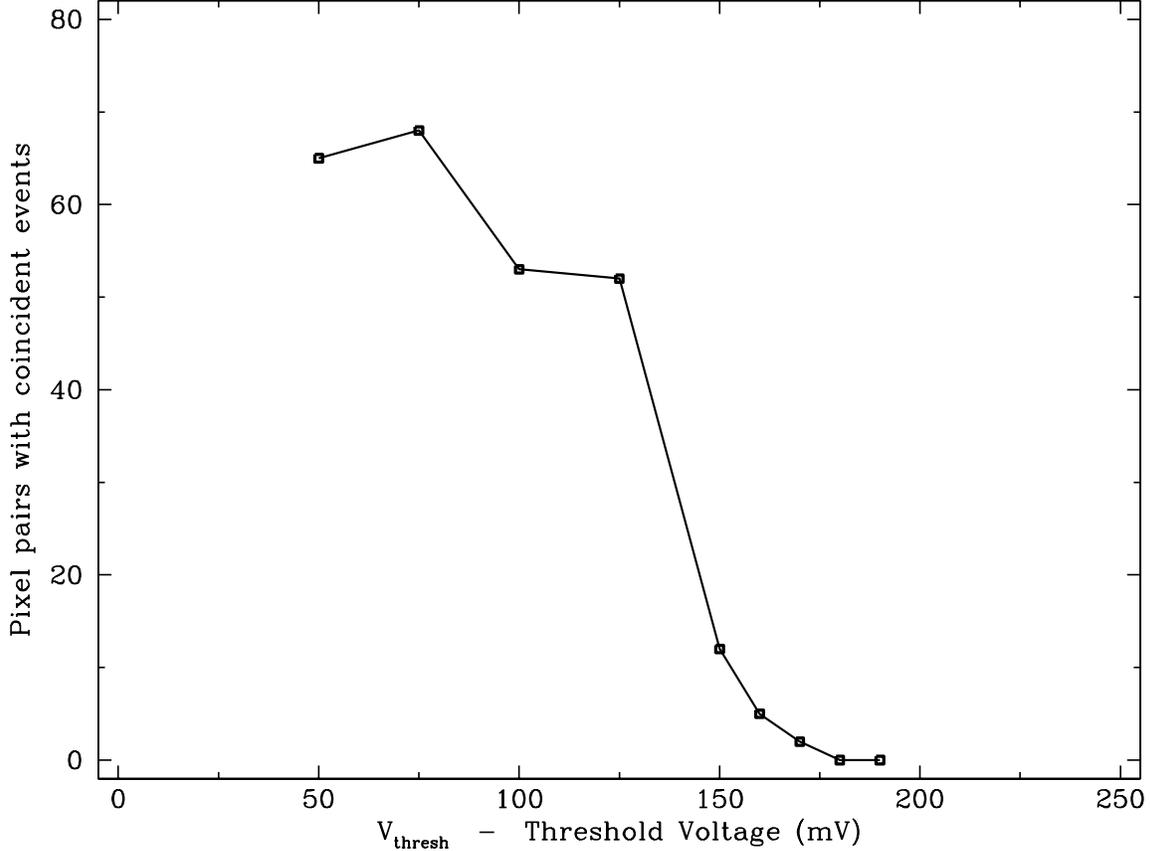


Figure 8.2. Number of pixel pairs (out the 512 total pairs in the PMT pixel array) with coincident events during 1-minute observations as a function of threshold voltage. Compare with Fig. 8.3, which plots these measurements sorted by PulseNet pixel number.

8.2.2 Camera Sensitivity

The sensitivity of the all-sky instrument is determined largely by the threshold voltage used to trigger coincident PMT signals, e.g. $V_{\text{thresh}} = V_{\text{bias}} - V_{\text{ref}}[1]$. 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 on other signals may couple into the PMT signals (before or after amplification). The list of candidates includes:

- 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 11mV (see Fig. 7.5)
- amplifier gain variations — the specification for small signal gain at 900 MHz¹ is 21_{-2}^{+3} dB. This is probably a conservative estimate, but on the high side of the range, 3 dB is a $\sim 40\%$ increase in gain.

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

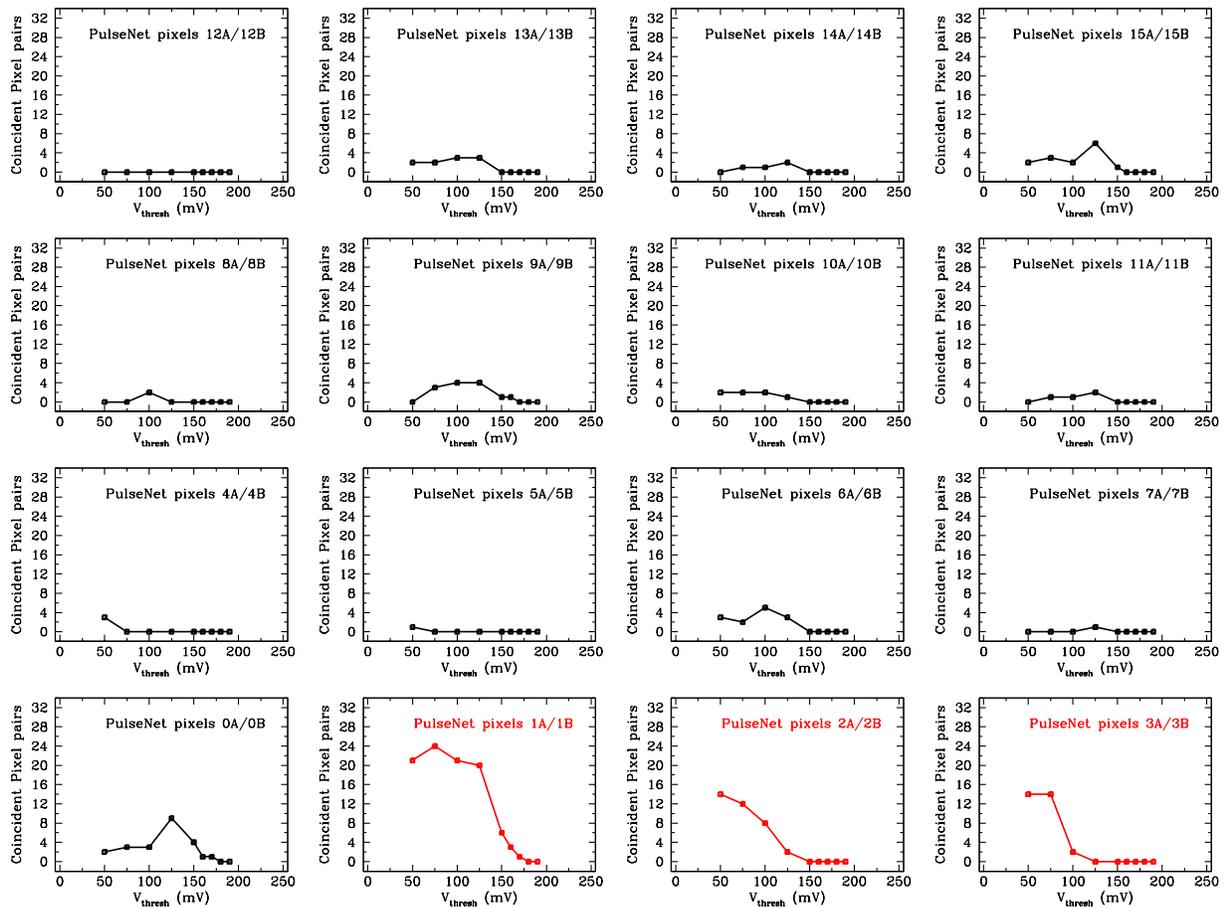


Figure 8.3. Dependence of PulseNet pixel number on sensitivity. For each numbered PulseNet input pixel pair (of which there are 16 per PulseNet numbered 0–15), a plot shows the number of pixel pairs (from all 32 daughterboards) with coincident events during 1-minute observations as a function of threshold voltage. The three red traces highlight an excess of coincident pixel pairs in PulseNet pixels 1, 2, and 3. Compare with Fig. 8.2, which combines these 16 plots.

- variable PMT cathode sensitivity — the variation can be as high as a factor of 2–3 on a tube, but typical pixels vary by $\sim 30\text{--}50\%$; note that this variation will not affect the single photoelectron pulse size, just the rate of such pulses.

To determine the threshold voltage for observations and to check for sources of noise, the camera sensitivity was measured. Each measurement was a 1-minute 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 were counted. Many observations were made at several threshold voltages to obtain Fig. 8.2.

Note that above $V_{\text{thresh}} = 170\text{ mV}$, no coincident events were recorded. Between 125–150 mV Fig. 8.2 shows a sharp increase that levels off to 50–60 pixel pairs for $V_{\text{thresh}} < 125\text{ mV}$. For $V_{\text{thresh}} < 150\text{ mV}$, the total number recorded in each 1-minute 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 ($V_{\text{thresh}} \geq 150\text{ mV}$), the events were concentrated in one to three pixel pairs in the entire array, each of which received 20–50 events. For observations with $V_{\text{thresh}} < 150\text{ mV}$, 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. 8.3 plots the same data sorted into 16 subplots according to PulseNet pixel number². Note the traces for PulseNet pixel numbers 1, 2, and 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 photomultiplier tubes. 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. 8.2 and 8.3. For $V_{\text{thresh}} \leq 150$ 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*. 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 *in01B* (the pixel with the largest number of pixels with coincident events in Fig. 8.3) 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 doesn't 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 *sampencbank0*, 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 `coinc_add_b0..3[3:0]`) in their number order 0–3 (see Fig. 6.2 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 one 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 ($G_{\text{bB}}^{\text{A}}0..3[2:0]$) so that samples from the coincident pixel pair (e.g. $G_{\text{bB}}^{\text{A}}0[2:0]$) are passed out of *sampencbank* and stored in memory. However, if two 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. 8.3. 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. 8.2 and 8.3. For $V_{\text{thresh}} \geq 150$ 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 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

²The 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.

³This is because software in the host computer calculates `int(log2(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 two 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 `log2(coinc_add_a0[3:0])` is integer truncated. The software can and should be rewritten to identify and handle such multiple coincidence events.

Table 8.1. All-sky survey — initial observations

Date	Declination Range	Observations	Total Observing Time	Triggers	Events
26 April, 2006	0° 17′ – 1° 53′	5	2 hr 33 min	2	0
28 April, 2006	1° 53′ – 3° 29′	12	6 hr 9 min	0	0
29 April, 2006	3° 29′ – 5° 05′	17	8 hr 19 min	0	0

Note. — Initial observations using the all-sky instrument and following the procedures in Fig. 8.4.

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 reflections).

8.3 Observing procedures

For the sake of consistency, safety, and meaningful results, it was necessary to develop a set of manual observing procedures (Fig. 8.4) before conducting the initial observations described in §8.4. An important aspect of this is the tests on the camera functionality done each night, and other set of tests before each 30-minute observation. Additionally, telemetry data provide constant feedback on the 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 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 $V_{\text{ref}}[6:0]$ and V_{bias} by setting $V_{\text{ref}}[1] > V_{\text{bias}}$ and setting the PulseNets to trigger on $V_{\text{ref}}[1]$. 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.

The Gelfand Flasher I test sends an optical flash at PMT arrays from a de-focused 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.

For nightly operations to survey the whole Northern sky, these procedures will likely be automated.

8.4 Initial observations

The initial observations listed in Table 8.4 were conducted over three nights in late April, 2006 using the procedures outlined in §8.3. The 34 observations of ~30 min duration each totaled 17 hr 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 $V_{\text{ref}}[2]$ as the trigger with the reference voltages spaced as follows: $V_{\text{ref}}[0] = 1.700 \text{ V}$, $V_{\text{bias}} = 1.500 \text{ V}$, $V_{\text{ref}}[1] = 1.375 \text{ V}$,

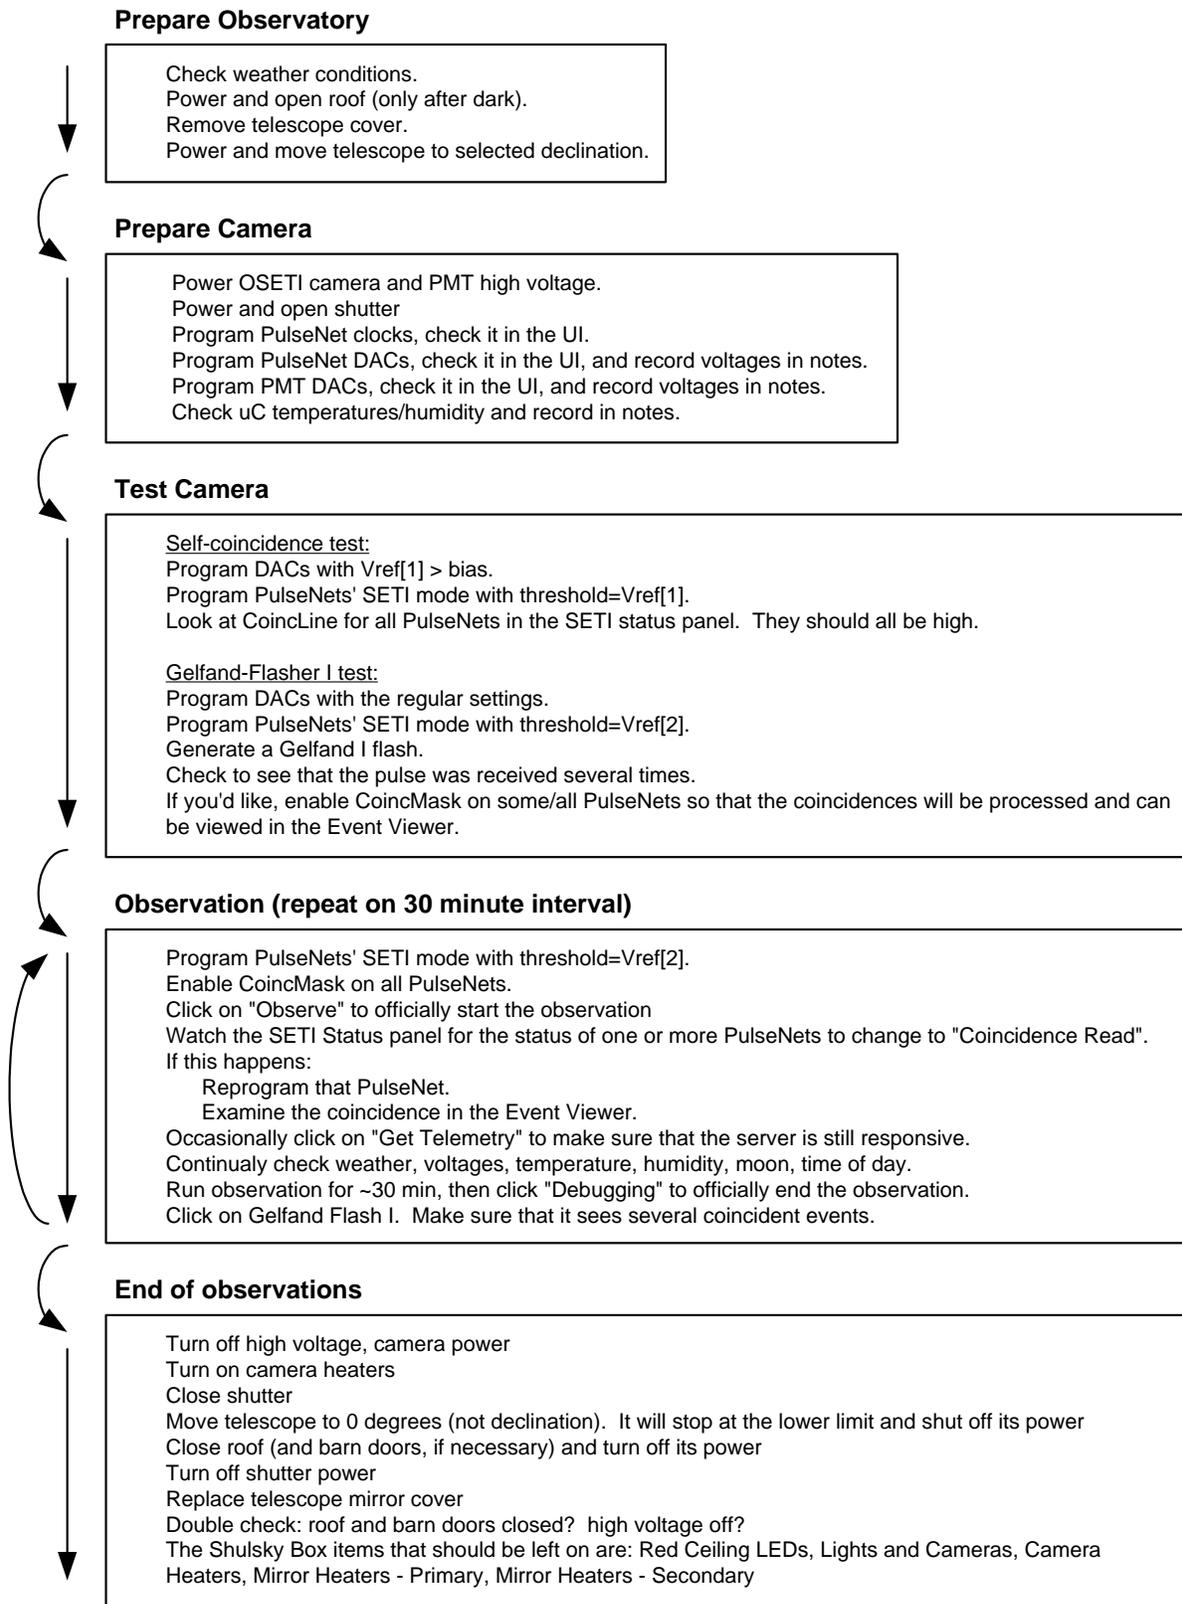


Figure 8.4. All-sky observing procedures for initial manual observations.

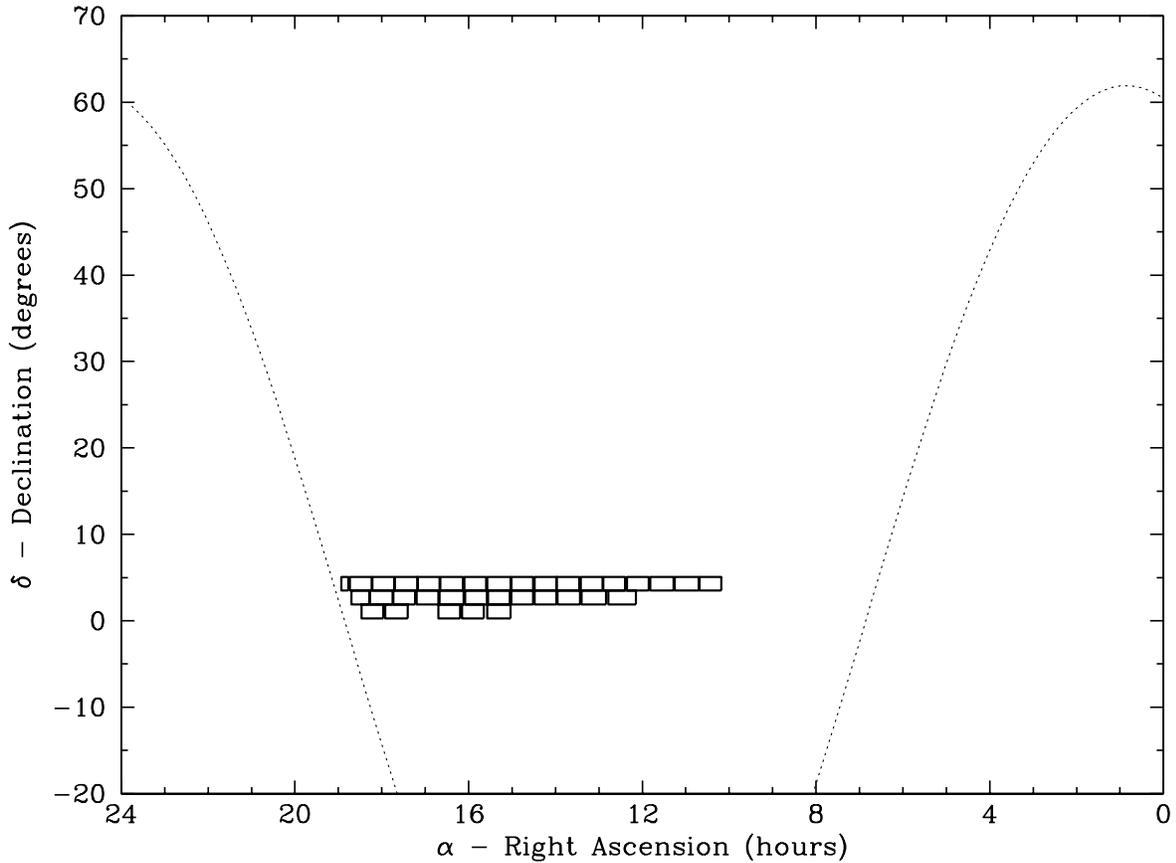


Figure 8.5. Initial sky coverage for the all-sky search in equatorial coordinates. The 34 observations over 17 hr, each represented by a rectangle in right ascension and declination, are summarized in Table 8.1. The dashed line represents the galactic equator.

$$V_{\text{ref}}[2] = 1.250 \text{ V}, V_{\text{ref}}[3] = 1.100 \text{ V}, V_{\text{ref}}[4] = 0.950 \text{ V}, V_{\text{ref}}[5] = 0.75 \text{ V}, V_{\text{ref}}[6] = 0.500 \text{ V}$$

- 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’ observations⁴. 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 anti-signature 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 the result from §3.3.2, the minimum energy cosmic-ray that would deliver ~ 100 photons m^{-2} into a PMT pixel is $\sim 10^{16}$ eV, which has a frequency about once every several years from an arbitrary point on the ground.

⁴Note 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.

Table 8.2. 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 ≤ 5 ns	100 in ≤ 5 ns	95 in ≤ 3 ns	17 in ≤ 5 ns
Sky Coverage	0.0002%	0.00005%	1%	64–75%
Objects	4730	1142	$\sim 10^6$	$\sim 10^8$
Observations	11,600	1721	17	~ 2700
Observations per object	2.5	1.5	1	1
Observing time	1721 hr	244 hr	17 hr	1350–1800 hr
Observing time per object	22 min	13 min	48 sec	48–140 sec
Events	274	130	0	–
Event Rate	0.16 hr ⁻¹	0.53 hr ⁻¹	0.00 hr ⁻¹	–
Triggers	1066	614	2	–
Trigger Rate	0.62 hr ⁻¹	2.52 hr ⁻¹	0.11 hr ⁻¹	–

Note. — Statistics comparing the all-sky survey and the targeted search. Data for the targeted search are from Table 4.2. 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 $-20^\circ < \delta < +70^\circ$ to $-30^\circ < \delta < +90^\circ$. The sensitivities listed are for the wavelength ranges listed in §4.2 and §5.9 and, in the case of the full survey, assume a successful daughterboard upgrade.

These initial observations covered 1% of the sky (400 sq-deg), as illustrated in Fig. 8.5. 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 in the targeted search — were surveyed for 48 sec each. Even with a conservative threshold voltage setting of 250 mV, the initial observations had approximately the same sensitivity as the targeted search (95 photons m⁻² in ≤ 3 ns), and *no* observations produced events that were consistent with genuine interstellar optical flashes.

8.5 Implications

8.5.1 Methodology

By adapting the procedure used in the targeted search (§4.11.2), 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.

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. The observational target of this search is the sky itself. As such, the methodology for constraining the number 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 turns 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-like⁵. When integrated over the whole sky, yields 10^8 stars:

$$\iint \rho_{\text{stars}} \cos \delta \, d\delta \, d\alpha = 10^8 \text{ stars.} \quad (8.4)$$

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 $\rho_{\text{stars}} \sim 2400$ 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:

$$T_i(\alpha, \delta) = \begin{cases} t_{\text{drift}}(\delta) & \text{for all } (\alpha, \delta) \text{ observed in the } i\text{th observation} \\ 0 & \text{otherwise} \end{cases} \quad (8.5)$$

where $t_{\text{drift}}(\delta) = (48 \text{ sec})/\cos \delta$. That is, when plotted, each $T_i(\alpha, \delta)$ looks like one of the rectangles in Fig. 8.5, and has an approximately constant value of the drift time through the array for the patch of the sky observed. There are 17 T_i 's in the initial observations; to fully cover the the Northern sky ($-20^\circ > \delta > +70^\circ$) once with ~ 30 min. observations, there will be ~ 2700 of them.

Recall that in the targeted search analysis (§4.11.2) 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 i th observation, $T_i(\alpha, \delta)$ has an associated probability density function,

$$p_i(\alpha, \delta, P) = \min(1, T_i(\alpha, \delta)/P), \quad (8.6)$$

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

Although no point in the sky was observed more than once during the initial observations, they will be during longer runs. In the targeted search, folding together the probabilities of detection for multiple observations of the same object was accomplished with Eq. 4.2. For this search we have

$$p(\alpha, \delta, P) = 1 - \prod_i (1 - p_i(\alpha, \delta, P)). \quad (8.7)$$

The final step in this analysis is to calculate the expected number of signal detections for all observations, S . Recall that in the the targeted search analysis, this is just $S = f(P) \sum_{\text{obj}} p_{\text{obj}}$ 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:

$$S = f(P) \iint p(\alpha, \delta, P) \rho_{\text{stars}}(\alpha, \delta) \cos \delta \, d\delta \, d\alpha. \quad (8.8)$$

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

$$f(P) = \min \left(1, \frac{\ln 2}{\iint p(\alpha, \delta, P) \rho_{\text{stars}}(\alpha, \delta) \cos \delta \, d\delta \, d\alpha} \right), \quad (8.9)$$

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

⁵For 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.

8.5.2 Constraints on transmitting extraterrestrial civilizations

Applying the methodology of §8.5.1 to the initial observations of 1% of the sky yields the plot of $f(P)$ in Fig. 8.6 (red trace). For comparison, results from the targeted search (Fig. 4.7) are plotted in blue (Harvard search) and cyan (Harvard-Princeton joint observations). The dashed lines from these plots represent the data reported in Chapter 4 and [35], while the solid lines represent the addition of ~ 1.5 yr 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 trade-off 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. 8.6 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. 8.6 in the context of estimates for the number of transmitting civilizations N (see §2.1.3). Estimates for N range from 1– 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 $f(P) \leq 7 \times 10^{-7}$ for $P < 48$ sec, which corresponds to a limit of $N \leq 3 \times 10^5$, *using the above model for extinction and stellar density*. A full sky survey has the potential to constrain $f(P) \leq 10^{-8}$, achieving a sensitivity of $N \leq 4 \times 10^4$.

8.6 Recommendations for future work

The all-sky experiment took 6–7 years to develop and build (including designing the observatory building, telescope, instrument, PulseNet, etc.). Surveying the Northern sky is the obvious next step. This task will be made easier by automating the observatory and instrument for unattended observations; the observing procedures (Fig. 8.4) should serve as a guide when developing the automation.

A full survey of the Northern sky will comprise, by far, the most comprehensive optical search ever. Our ability to calculate the implications from such a search will improve if a detailed, spatially-dependent model of the number density of stars by spectral type is folded into a similar model for interstellar extinction. Such a model would allow for more meaningful application of the methodology in §8.5.1.

The sensitivity of the all-sky search could be significantly improved. In the initial observations, the voltage threshold was set at $250 \text{ mV} \approx 17$ photoelectrons. By identifying and fixing hot pixels, this might be brought down to $\sim 150 \text{ mV} \approx 10$ photoelectrons. However, the residual background noise of *fastclock*-coupling on PulseNet pins in01B, in02B, and in03B will limit the impact of this solution. Second generation daughterboards with flying clocks may cure the *fastclock*-coupling problem. However, before replacing the daughterboards, all sources of background triggers should be more fully understood so that they can be minimized in the replacement design.

The longer term future of optical SETI lies in the infrared and with very large collecting areas. Pulsed infrared SETI has the advantages of greater signaling range (due to decreased scattering and absorption in the IR; see §2.4.3), lower energy cost per photon, and decreased stellar background. These are offset by the challenges of building a sensitive instrument to detect infrared pulses. Terrestrial sources of background photons (including the atmosphere and the experimental apparatus) are significant. Contemporary detector technology (e.g. solid-state photomultiplier) requires cryogenics, and is often limited by small sensitive area and slow response. Major improvements in detector technology could enable pulsed IR searches.

Optical SETI can benefit from collaborations with cosmic-ray and gamma-ray observatories [14, 23]. Such instruments have large collecting areas (tens of square meters), sensitivity to nanosecond-scale optical pulses, and wide fields of view — precisely the demands of a sensitive optical SETI. Presently these experiments routinely discard pulses coming from point sources because such events are inconsistent with the larger source size of diffuse air showers. Optical SETI programs could use the large aperture and fast detectors of these experiments *with no hardware modifications at the site*. The improvement in collecting area compared with the current searches is dramatic. For example, the Solar Tower Atmospheric Cherenkov Effect Experiment

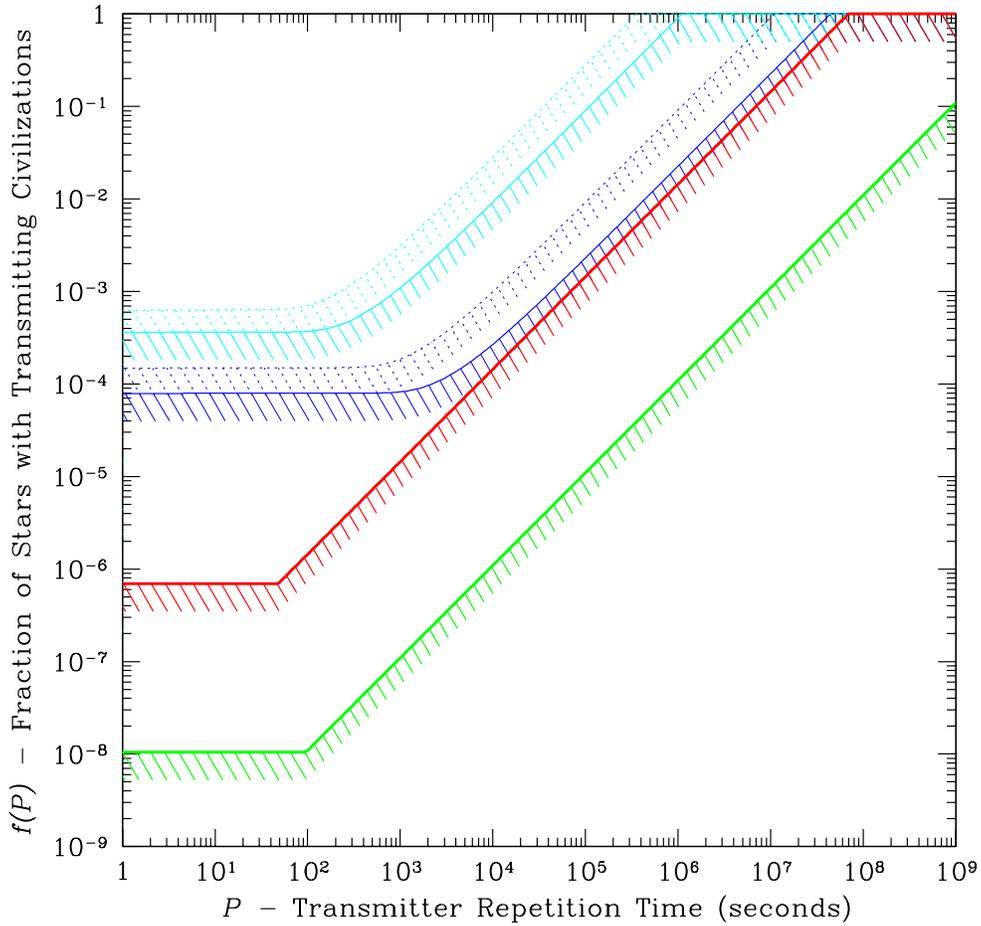


Figure 8.6. Upper bounds on the fraction of stars with transmitting civilizations as a function of transmitter repetition time for three searches — the Harvard targeted search (DS_{clean}) in blue, the Harvard-Princeton targeted search (DS_{overlap}) in cyan, and initial observations from the all-sky survey (1% of the sky) in red — along with projected results from the all-sky instrument after surveying the Northern sky twice in green. Two curves are shown for each targeted search: the dashed lines are those reported in Chapter 4 and [35]; solid lines incorporate data from ~ 1.5 yr of additional observations. Each curve asymptotes to $f = 1/N_{\text{obj}}$ for small P and cuts off at $f = 1$ for $P \geq T/\ln 2$.

(STACEE) uses a research solar power plant with 200 large mirrors (37 m^2 each). STACEE uses 64 of these mirrors to focus light onto PMTs (one per mirror) and triggers on coincident pulses in multiple PMTs. It's sensitivity is ~ 2 photons m^{-2} , a factor of about 50 greater than the targeted searches and the all-sky search in its present configuration. The mining of the data archives from cosmic-ray and gamma-ray experiments has just begun. Jamie Holder from the University of Leeds, and colleagues, recently begun searching the Whipple 10 m gamma-ray telescope data archive for signals with characteristics (compact and circular source images) consistent with pulsed optical communication. This instrument has several hundred PMTs and can achieve a sensitivity of ~ 10 photons m^{-2} . The data archive contains ~ 5000 hours of observing time [29].

Appendix A

Optical SETI with Terrestrial Planet Finder

This appendix was previously published as an article in Icarus [33].

A.1 Abstract

NASA's space-borne nulling interferometer (the Terrestrial Planet Finder – TPF) will look for the traces of early life in the infrared spectra of extrasolar planets, beginning in roughly 2010. We point out that this instrument will also be sensitive to deliberate laser transmissions from a technologically advanced civilization. A kilowatt-class infrared laser with a 10-m beam director would produce a signal visible to TPF at a range of 15 pc that is distinguishable from astrophysical phenomena and noise.

A.2 Introduction

Slated for construction in roughly a decade, TPF will be a set of infrared (3-30 μm) telescopes whose combined light forms a nulling interferometer. Although details have yet to be worked out (separated spacecraft *vs.* single spacecraft, and 1 AU *vs.* 5 AU orbit), its high angular resolution (0.75 mas at 3 μm and a free-flying 1000 m baseline, which is equivalent to 0.01 AU resolution at a range of 15 pc) will allow it to examine extrasolar planetary systems while nulling the light from the parent star. TPF will hunt for planets and will examine their structure, formation, and evolution. It will also search for the chemical signatures of life on these planets, in the form of CO₂, H₂O, CH₄, and O₃ absorption bands. We argue below that the features of this spacecraft also allow it to receive and identify intentional laser transmissions of modest power from extraterrestrial civilizations.

Historically, the Cocconi and Morrison [10] suggestion that SETI be carried out at the 21 cm line of neutral hydrogen came at time when no other astronomical lines were known in the microwave spectrum. In the following two years, the laser was invented, and Townes [55] noted that these relatively low power “optical masers” could be used for interstellar communication. The suggestion has received increasing attention as lasers have continued to show an annual Moore’s Law doubling in power over the past 40 years (during which time radio technology has remained relatively static). Today, there are optical SETI¹ programs at Harvard [34], Berkeley [40], Columbus [38], and elsewhere.

The merits of optical SETI are well documented in other articles [62, 37, 52], but we highlight a few advantages here. Given the high gain of optical telescopes, optical beacons can be narrowly focused on target systems. The bandwidth-limiting dispersive broadening observed in radio pulses is negligible in the optical

¹“Optical” SETI is to be interpreted in the broad electromagnetic sense – including the near IR and UV – and is characterized by photon counting; this is to be contrasted with the heterodyne techniques of microwave SETI.

regime. The computational power and sophistication required for broadband microwave SETI searches is unnecessary in optical SETI. Also, the pace of laser development on this planet has made possible optical interstellar communication *from* Earth; shouldn't we be looking for such communication *to* Earth?

TPF will be one of a handful of instruments in the Solar System sensitive to optical interstellar communication. Clearly, such instruments have to mitigate the background light from stars. One approach in optical SETI is to build instruments capable of detecting optical pulses on a very fast timescale, where the received optical pulse manifests itself as many photons arriving at the detector in an unresolved (\sim ns) time interval, against a background pattern of single, Poisson distributed photon arrivals from the host star. A simple calculation, assuming only "Earth 2000" technology – a transmitter capable of delivering nanosecond speed, megajoule optical pulses at a 10 Hz repetition rate attached to a Keck-class telescope – would outshine our Sun by a factor of more than 1000 in broadband visible light, and could be easily detected by another Keck-class telescope at distances of up to 300 pc [34]². Several optical SETI programs monitor nearby stars for optical pulses of this type in roughly the 300-600 nm band. Our group is now developing a pixelated wide-field camera to search the Northern sky for such high-intensity pulsed signals.

An alternative approach in optical SETI is to reduce the stellar background by high resolution spectroscopy and/or nulling interferometry. In such a strategy, one looks for unexplained lines in the spectra of stars and their planetary disks. Such lines would be either continuous-wave (CW) lasers (possibly modulated) or pulsed lasers. These would be detectable at considerably lower transmitted power levels than is possible without interferometry and spectroscopy, as we will demonstrate below. TPF is the first interferometer with the angular resolution capable of separating starlight from planetlight and has a modest spectral resolution (up to $R = \nu/\Delta\nu = \lambda/\Delta\lambda = 10^5$). Thus it can probe the entire zodiacal disk of a nearby star for planets and life.

The discovery of extraterrestrial communication from a nearby planetary system using TPF might proceed along the following lines. A planet candidate is discovered after several hours of observation at minimal spectral resolution (R of order unity). Followup observations at modest spectral resolution ($R \sim 100$) reveal a rich absorption spectrum with a large spectral peak in a particular wavelength bin. The spectral purity is then probed at still higher resolution ($R \sim 10^5$), revealing an unresolved narrow emission line³. Ground-based telescopes then look for modulation in the signal, down to the scale of nanoseconds; they also examine the full range of wavelengths accessible to terrestrial telescopes.

A.3 Laser power required

Extraterrestrial lasers would have to compete with at least four backgrounds in potential TPF searches: (1) incompletely nulled stellar photons, (2) reflected photons from the extrasolar planet, (3) photons from the extrasolar planet's blackbody spectrum, and (4) light scattered by zodiacal dust in both the target solar system and in ours. We will examine these backgrounds in greater detail below, but first note that they are all isotropic (roughly) and spectrally broad. Therefore, as Geoff Marcy [51] pointed out, the laser power necessary for interstellar communication is reduced by the following factors:

A.3.1 Directionality

Lasers are directional and therefore have gain. A laser of wavelength λ transmitted through an aperture of diameter D emerges as a beam of width $\theta_b \approx \lambda/D$, with a gain of $g_t = \pi^2 D^2/\lambda^2$. At $\lambda=10 \mu\text{m}$, a 10-m telescope has a gain of $g_t \approx 10^{13}$, or 130 dB. (Note that this is a factor of 10^6 greater than the gain of the 305-m Arecibo dish at $\lambda=21 \text{ cm}$.)

²Note that extinction limits optical SETI to a range of ~ 1 kpc in the visible part of the spectrum and ~ 10 kpc in the infrared. These distances correspond to roughly the thickness and radius of the galactic disk and volumes enclosing $\sim 10^7$ and $\sim 10^{10}$ sun-like stars, respectively.

³The time requirements to detect microscopic *or* intelligent life on an extrasolar planet are lengthy. NASA estimates that planet detection will require observations with 2.0 hour integration times (with $R = 3$, $S/N = 5$). Detection of atmospheric gases such as CO_2 and H_2O would require integrations of 2.3 days ($R = 20$, $S/N = 10$), and for detection of life-indicating O_3 or CH_4 the corresponding figure is 14.7 days ($R = 20$, $S/N = 25$).

A.3.2 Spectral resolution

The spectrometer onboard TPF will sort the incoming radiation into frequency “bins” (of width $\Delta\nu$) thereby reducing the background in any given bin by approximately $\Delta\nu/\nu = 1/R$ (assuming a flat spectrum). Because lasers are spectrally narrow, all of their photons will fall into one bin. The spectral resolution R of TPF varies widely, from $R = 3\text{--}20$ for “planet detection and spectroscopy,” to $R = 3\text{--}300$ for “continuum and spectral line imaging,” and extending to $R \sim 10^5$ for “specific lines” [3]. Let us assume that extraterrestrial lasers would be discovered in TPF’s “planet detection mode,” and take $R = 20$ for this calculation.

A.3.3 Fluctuations

We must not forget that the laser signal competes only with the *fluctuations* in the background, and not the background signal itself. Pixel-to-pixel fluctuations in the spatial image will inhibit TPF’s planet detection power. These include physical processes such as emission from the incompletely nulled star and dust, as well as instrumental effects such as telescope jitter, intrinsic detector noise, and emission from the cooled telescope. Diffractive intensity scintillations should be insignificant at TPF’s high frequency (compared to radio), and long integration times [13]. Spectral fluctuations will limit the spacecraft’s ability to identify atmospheric gas absorption bands and extraterrestrial laser emission lines. Noise sources of this type include spectrometer noise, and noise in the reflected and blackbody light from the planet. On large spectral or spatial scales, these fluctuations are characterized by a signal-to-noise ratio; Appendix A of [3] finds $\text{SNR} \approx 7$ during a typical run of TPF.⁴ However, for the purposes of detecting extraterrestrial lasers – which deposit all of their light into one spectral bin – the figure of merit is the bin-to-bin variation in the spectrum, that is, the deviation from a smooth spectrum. Note that fluctuations of this type include not only the above noise sources, but also spectral *features* (such as absorptions lines) whose width is comparable to the bin size. For the purpose of the calculations in the paper, let us assume that the “signal-to-fluctuation ratio” is the same as the signal-to-noise ratio; in practice they may differ by perhaps a factor of two.

Taking account of these factors, the laser power sufficient for detection by TPF is

$$P_L \geq \frac{\lambda^2}{\pi^2 D^2} \cdot \frac{\Delta\nu}{\nu} \cdot \frac{5}{\text{SNR}} \cdot L_T \equiv \alpha \cdot L_T, \quad (\text{A.1})$$

where L_T is the total isotropic radiated power of background sources seen in a spatial TPF pixel, and where we have assumed that a $5\text{-}\sigma$ signal is required for detection.

To get a sense of the order-of-magnitude of α , let us assume that the transmitting device is simply a CO_2 laser ($\lambda=10\ \mu\text{m}$) coupled to a Keck-class telescope ($D=10\ \text{m}$). Combining this with the above results, we find $\alpha=4 \times 10^{-15}$. This means, for example, that a laser of $\approx 10^{11}\text{W}$ average power would be visible over the background of its G2V host star, *without the use of interferometry*. In other words, the high transmitting aperture gain at IR wavelengths, combined with only modest wavelength specificity, already reduces the required transmitting power (relative to solar luminosity) by some 14 orders of magnitude. Moreover, as we shall now demonstrate, TPF’s exquisite interferometric nulling capability greatly reduces this figure, bringing it within the range even of modest contemporary lasers.

A.4 Backgrounds

As discussed above, the background for extraterrestrial lasers can be broken down into several distinct sources. Incompletely nulled stellar photons, as well as inhomogeneities in the zodiacal and exo-zodiacal dust, are important during spatial imaging, while reflected and blackbody photons from the planet are important during spectroscopic integrations. In the discussion that follows, we calculate these backgrounds using Earth/Sun values for the physical parameters of planets and stars in units of broadband visible equivalent isotropic radiated power (EIRP).

⁴This calculation is for either 2 m mirrors at 5 AU or 3.5 m mirrors at 1 AU imaging a planet at 10 pc and integrating for $10^5\ \text{s}$ (roughly a day) at $12\ \mu\text{m}$ with $R = 20$. It includes (i) the effects of zodiacal and exo-zodiacal emission, (ii) the galactic cirrus, (iii) leakage signal and jitter from the target star, (iv) the telescope properties, (v) and detector noise.

A.4.1 Incompletely nulled stellar photons

The most technically challenging aspect of TPF is nulling light from the parent star. Although a nascent field today, astronomical nulling is being developed at various observatories on Earth, and will soon be developed on the Space Interferometry Mission in orbit. After this decade of research, TPF is projected to achieve a null depth of $N=10^{-5}$ – 10^{-6} . Taking the more conservative estimate of $N=10^{-5}$, the nulled stellar luminosity is

$$\begin{aligned} L_N &= N \cdot L_\odot \\ &= 4 \times 10^{21} \text{ W}. \end{aligned} \tag{A.2}$$

Note that L_N is only important for spatial imaging near the parent star (assuming that L_N does not saturate TPF's spectrometer). During spectroscopic integrations on a planet separated from its parent star by many pixels on the TPF image, L_N is unimportant.

A.4.2 Reflected photons

Light reflected off of the extrasolar planet will be a background for discovering a planetary atmosphere and detecting extraterrestrial lasers. Assuming an albedo of unity, the power reflected from a planet of radius R_\oplus at a distance r_\oplus from its sun is approximately

$$\begin{aligned} L_R &= \frac{\pi R_\oplus^2}{4\pi r_\oplus^2} \cdot L_\odot \\ &= 5 \times 10^{-10} \cdot L_\odot \\ &= 5 \times 10^{-5} \cdot L_N. \end{aligned} \tag{A.3}$$

This estimate may be high by perhaps an order of magnitude since the hot star's blackbody peak is in the visible part of the spectrum, rather than at TPF's infrared wavelength. We have also not taken account of the variations in apparent planetary brightness due to orbital phase as seen from Earth.

A.4.3 The planet's blackbody spectrum

Although cooler and smaller than its parent star, the planet emits a blackbody spectrum of its own, which peaks in the infrared. This will be spatially resolved from the parent star, as well as from other planets in the system, and will be seen as a bump on the zodiacal background. In total power, the planet emits far less than its nulled parent star,

$$\begin{aligned} L_B &= \left(\frac{T_\oplus}{T_\odot}\right)^4 \cdot \left(\frac{R_\oplus}{R_\odot}\right)^2 \cdot L_\odot \\ &= 4 \times 10^{-10} \cdot L_\odot \\ &= 4 \times 10^{-5} \cdot L_N. \end{aligned} \tag{A.4}$$

However in the infrared, where TPF is sensitive, the nulled stellar and planetary blackbody power per unit frequency are more nearly comparable. For example, at $\lambda=10 \mu\text{m}$ ($\nu=3 \times 10^{13}$ Hz),

$$\begin{aligned} \frac{L_B}{\Delta\nu} &= N \cdot 4\pi R^2 \cdot \frac{2h\nu^3}{c^2} \cdot \frac{1}{e^{h\nu/kT} - 1} \\ &= 8 \times 10^4 \text{ W Hz}^{-1} \text{ (star)} \\ &= 1 \times 10^3 \text{ W Hz}^{-1} \text{ (planet)}, \end{aligned} \tag{A.5}$$

where we take the null depth to be $N=10^{-5}$ for the star and $N=1$ for the planet. These signals should be within the dynamic range of the spectrometer on board TPF. Such results should not surprise us, given that TPF is being designed to resolve and identify extrasolar planets.

Table A.1. Alternative planetary systems

Planet	Minimum P_L (kW)	L_R/L_\odot	L_B/L_\odot	r/r_\oplus	R/R_\oplus	T/T_\oplus
Mercury	1	4×10^{-10}	4×10^{-10}	0.39	0.38	1.61
Venus	2	8×10^{-10}	7×10^{-10}	0.72	0.95	1.17
Earth	1	5×10^{-10}	4×10^{-10}	1	1	1
Mars	0.2	6×10^{-11}	5×10^{-11}	1.5	0.53	0.81
Jupiter	6	2×10^{-9}	2×10^{-9}	5.2	11	0.44
Saturn	1	4×10^{-10}	4×10^{-10}	9.5	9.4	0.32
Uranus	0.05	2×10^{-11}	1×10^{-11}	19	4.0	0.21
Neptune	0.03	8×10^{-12}	1×10^{-11}	30	3.9	0.20
Pluto	0.01	1×10^{-14}	9×10^{-15}	39	0.18	0.16

Note. — Minimum power of a 10 μm laser with a 10-m diffraction limited beam director on a Solar System model planet at a range of 15 pc for detection by TPF under the assumptions of Sec. A.3. This power is calculated using Eq. A.1 with $L_T = L_R + L_B$ (L_N is unimportant when doing spectroscopy on a planet well separated from its sun). Physical data in the three rightmost columns are from Zeilik and Gregory [70] (T is equilibrium blackbody temperature). L_R/L_\odot and L_B/L_\odot were calculated using Eqs. A.3 and A.4, respectively. Note that Pluto would not be detected by TPF since it emits less power than a background 0.03×0.03 AU patch (TPF’s pixel size at 10 μm) of a 1 Zodi cloud. Extraterrestrial lasers on such planets would have to exceed the radiated power of the local zodiacal dust (this directed laser power is roughly 10 W).

A.4.4 Zodiacal and exo-zodiacal dust

Dust in the Solar System will create a diffuse infrared glow that will cloud, but not block TPF’s view.⁵ Exo-zodiacal dust also presents a significant challenge. A “1 Zodi” cloud of Solar System-like dust is only 0.3 AU in diameter, yet it emits and scatters roughly the same amount of infrared and optical radiation as the Earth. This dust is warm (275 K and decreasing with distance as $r^{-0.4}$), small (~ 40 μm grains), and smoothly distributed over the ecliptic, except for wakes and rings due to gravitational effects from planets, and bands due to recent asteroid or comet collisions [2]. TPF’s high angular resolution is therefore essential to subtract out this largely uniform exo-zodiacal background from the image of extrasolar planets (and their possible inhabitants’ lasers)⁶.

With detail on a scale down to 0.01 AU, the TPF image will be a central bright star surrounded by a diffuse zodiacal disk. Within this shroud of dust we will see the reflection and emission from a planet. An extraterrestrial laser affixed to or orbiting the planet will therefore only have to overcome the background from the planet itself. Any techniques used to spatially resolve the planet from the zodiacal disk and the parent star will also resolve the extraterrestrial laser. Taking the total background to be the sum of the reflected and blackbody backgrounds, $L_T = L_R + L_B = 9 \times 10^{-5} \cdot L_N \approx 4 \times 10^{17}$ W, we estimate that the laser power sufficient for interstellar communication, $P_L = \alpha \cdot L_T \approx 1$ kW. In just 40 years of development on Earth, we have managed already to produce *megawatt* infrared CW lasers [58]. Taking “Earth 2000” technology as a lower bound on extraterrestrial technological sophistication, we conclude that infrared CW lasers are an altogether reasonable way to achieve interstellar contact.

⁵This is the primary motivation to reduce TPF’s mirror size and place it in a 5 AU orbit.

⁶Note that the above quantities were calculated assuming that the ecliptic plane of the imaged planetary system is perpendicular to TPF’s line of sight. As this angle varies to an edge-on view, the background from exo-zodiacal dust will increase by a factor whose maximum value is roughly the ratio of the diameter to the thickness of exo-zodiacal disk. In a beautiful measurement using COBE, Reach et al. [50] found this factor to be ~ 3 in the Solar System.

A.5 Alternative planetary systems

In this rather simple treatment, we have calculated various quantities for detecting laser transmissions from the vicinity of an Earth-Sun system with TPF. Although it will not dramatically change the above results, the bodies of extrasolar systems studied by TPF will almost certainly differ in size, temperature, and relative separation. For example, suppose that the extraterrestrial civilization lives on the moon of a Jupiter-like planet in an Earth-like orbit, where its lasers would have to compete with a larger planetary background. Substituting $R_{\oplus} \rightarrow R_J$, we find that the required laser power, $P_L \approx 100$ kW, is still well within our present capability. Table 1 shows that P_L is at most a few kW for Solar System planets in their natural orbits as viewed by TPF from 15 pc (its specified maximum range for planet detection). Moreover, as our arbitrary choice of “Earth 2000” technology for the transmitting device suggests, this discussion was intended not as a blueprint for Earthly transmission, but rather as a sanity check for optical SETI with TPF.

We also note that the extraterrestrial laser could be on a satellite in a wide (say 5 AU) orbit and therefore spatially uncorrelated with planets in the system. In this case, the signal would be a narrow spectral peak on the exo-zodiacal background. Such signals might be missed if planetary systems are first scanned in low spectral resolution and only followed up with higher resolution spectroscopy if planets are discovered.

A.6 Conclusions

Although not designed for optical SETI, TPF is uniquely sensitive to infrared extraterrestrial signals originating from extrasolar planets. This sensitivity in an unexplored corner of frequency/location space is particularly important in the logic of SETI where we have no *a priori* knowledge of the signal. Given NASA’s congressionally mandated inactivity in SETI, we do not expect stand-alone search programs in the near term. However, TPF scientists should be aware of the possibility that their experiments may serendipitously extend the reach of humanity far beyond the Solar System.

Appendix B

Interstellar Gravitational Dispersion of Optical Pulses

This appendix demonstrates that one potential technical obstacle to interstellar pulsed optical communication — temporal broadening due to gravitational scattering off of large bodies — is a negligible effect.

B.1 Convergent and divergent scattering

There are two classes of gravitational scattering for optical pulses. If the extremal rays of the optical pulse (in the paraxial plane) pass on either side of the scatterer, then each ray is bent toward the scatterer and the beam converges (see Fig. B.1). If, on the other hand, the beam misses the scatterer and passes entirely to one side of it, then the extremal rays are both deflected toward the scatterer, but the outer ray is deflected less — a net divergence. In the skew plane (perpendicular to the paraxial plane and parallel to the line of sight; see Fig./,refdistort for a head-on view) for this latter case, extremal rays *converge*, but this convergence is much smaller than the expansion in the paraxial plane. Since the location of the scatterer will be random with respect to the line of sight, the smaller skew contractions will be more than compensated for by the larger paraxial expansions. This leads to an overall broadening of the beam, or divergent scattering.

Convergent, or multi-path, scattering is sufficiently rare that we need not worry about it for optical SETI. Gravitational lens searches put the probability of multi-path imaging at roughly 10^{-6} . In fluid mechanics terms, the flow is laminar; paths never cross. Although this effect is rare in the vacuum of space, we still need to consider multi-path scattering due to atmospheric inhomogeneities (for ground based optical SETI programs). A ray which scatters at an altitude h by an angle θ_s will travel an extra pathlength $h\theta_s^2/2$. For $h = 25$ km and $\theta_s = 10$ arcsec, the time delay $\Delta t = h\theta_s^2/(2c) = 10^{-16}$ sec. In reality the photons will random walk through the seeing disk by a larger number of smaller scatterings, but the above single scattering calculation sets the scale for this time delay as entirely negligible. One might also worry that the effective aperture of the telescope will be much larger than the physical aperture when atmospheric seeing is considered, but again the effect is negligible. The “extra aperture” due to seeing is $h\theta = 0.05$ m, using the above numbers.

Now let us return to the subject of divergent gravitational scattering. Rays arriving at the atmospheric seeing disk will be delayed relative to one another for two reasons. First, scattered rays will travel different distances leading to a *geometric* delay. These rays also take different paths through the gravitational potentials of stars so there is a *general relativistic* time delay, as well.

B.2 Geometric delay

Gravitational scattering causes a light ray to deflect by an angle $\theta(b) = 2R_S/b$, where $R_S \equiv 2GM/c^2$ is the Schwarzschild radius of a mass M , and b is the impact parameter. Rays which diverge from the host star by an angle ψ in the paraxial plane will be separated by a distance $db = r_1\psi$ at their closest approach to the scatterer. The ray closer to the scatterer (1) will scatter by an angle $\theta(b)$, while the more distant ray (2)

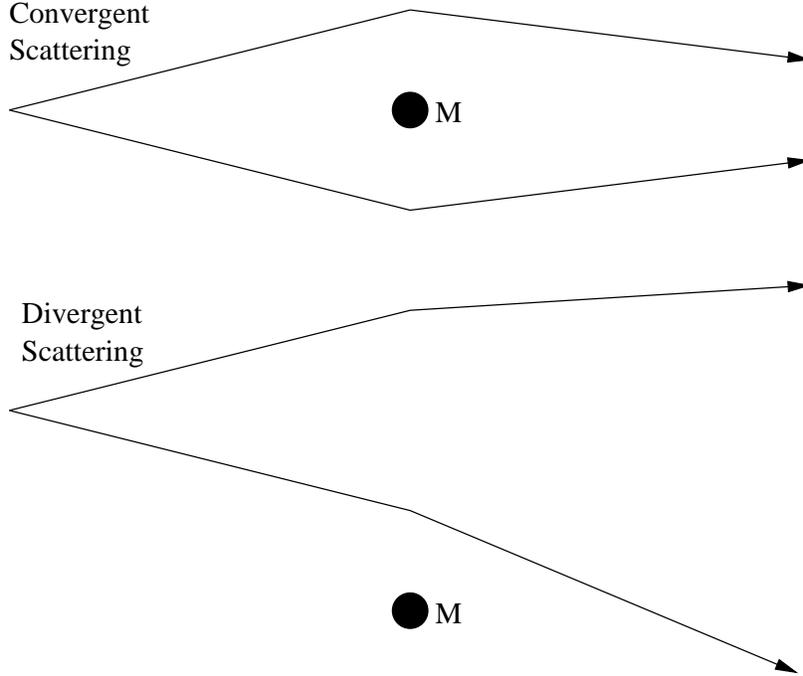


Figure B.1. Convergent and divergent gravitational scattering in the paraxial plane.

will scatter by a smaller angle $\theta(b + db) = 2R_S/(b + r_1\psi) \approx \theta(b) \cdot (1 - db/b)$ (see Fig. B.3). Thus (1) and (2) diverge by an additional angle $\theta(b) \cdot db/b$ when they scatter on the same side of a star.

In the skew plane, the extremal rays converge, but only slightly. The angular width in the paraxial direction goes from ψ to $\psi(1 + db/b)$, while the angular width in the skew direction goes from ψ to $\psi(1 - db/b \cdot R_S/b)$. The additional factor R_S/b is $\ll 1$, given that R_S for a solar-type star is ~ 3 km, whereas the mean solar separation in our region of the galaxy is $\sim 10^{13}$ km (or, roughly speaking, 10 lightmicroseconds vs. 1 lightyear). Since the solid angle of the beam increases by a factor of $\sim 1 + db/b$, there is a net broadening. As mentioned previously, the small contraction in the skew direction is washed out over many scatterings since the locations of scatterers with respect to the beam is arbitrary. We will therefore only consider scattering of paraxial rays.

Consider again the paraxial rays (1) and (2) from Fig. B.3. They travel different distances as they scatter. Ray (1) will travel a distance d_1 and ray (2) will travel a distance d_2 :

$$d_1 = \sqrt{r_1^2 + r_1^2\psi^2} + \sqrt{r_2^2 + r_2^2(\psi/2 + \theta(b))^2}$$

$$d_2 = \sqrt{r_1^2 + r_1^2\psi^2} + \sqrt{r_2^2 + r_2^2(\psi/2 - \theta(b + r_1\psi))^2}$$

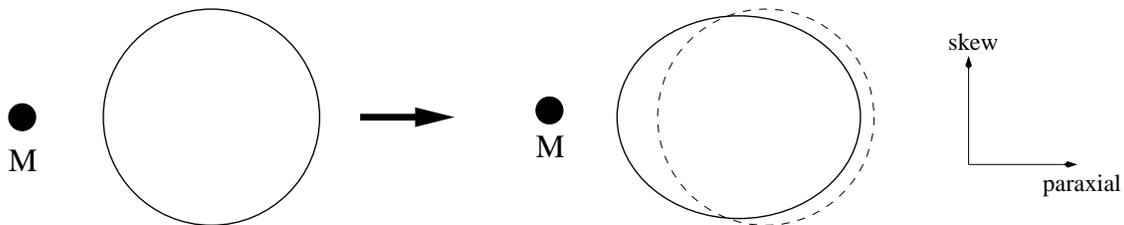


Figure B.2. A circular annulus of light is distorted into an ellipse. Note that the expansion along the paraxial axis is much larger than the contraction along the skew axis. Thus the area enclosed by the beam increases and there is an overall broadening.

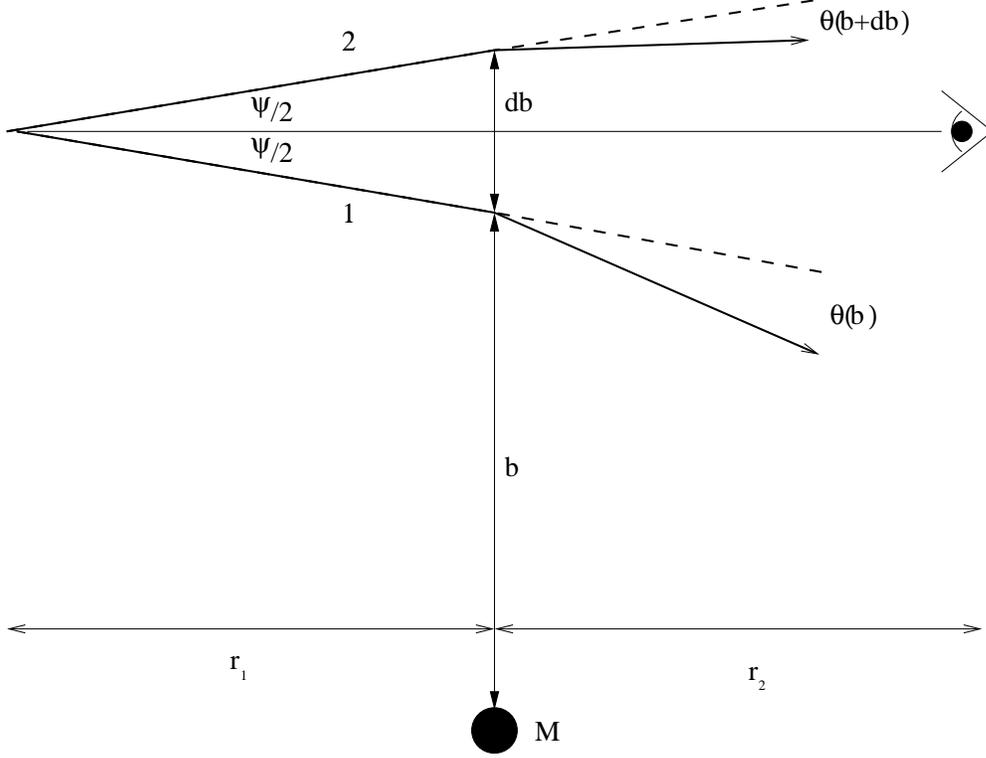


Figure B.3. Scattering Diagram in the paraxial plane.

The path length difference between the two rays is

$$\begin{aligned}
 \Delta \equiv d_2 - d_1 &= r_2 \{ \sqrt{1 + (\psi/2 + \theta(b))^2} - \sqrt{1 + (\psi/2 - \theta(b + r_1\psi))^2} \} \\
 &\approx r_2 \psi \{ \theta(b) - \theta(b + r_1\psi) \} / 2 \\
 &\approx \frac{r_1 r_2 R_S}{2b^2} \cdot \psi^2
 \end{aligned}$$

in the limit of $\theta, \psi \ll 1$. θ and ψ are defined in Fig. B.3.

So the geometric time delay from a single scattering event is

$$\Delta t_{geo} = \frac{r_1 r_2 R_S}{2b^2} \cdot \frac{\psi^2}{c}. \quad (\text{B.1})$$

B.3 Gravitational delay

Light pulses are also slowed down as they dip into the gravitational potential of stars. The travel time between A and B (see Fig. B.4) is given by (see [19]):

$$\begin{aligned}
 ct &= \underbrace{\sqrt{r_1^2 - b^2} + \sqrt{r_2^2 - b^2}}_{\text{geometric path length}} + \\
 &\quad \underbrace{2M \ln \left(\frac{\sqrt{r_1^2 - b^2} + r_1}{b} \frac{\sqrt{r_2^2 - b^2} + r_2}{b} \right)}_{\text{gravitational delay}} - M \left(\frac{\sqrt{r_1^2 - b^2}}{r_1} + \frac{\sqrt{r_2^2 - b^2}}{r_2} \right). \quad (\text{B.2})
 \end{aligned}$$

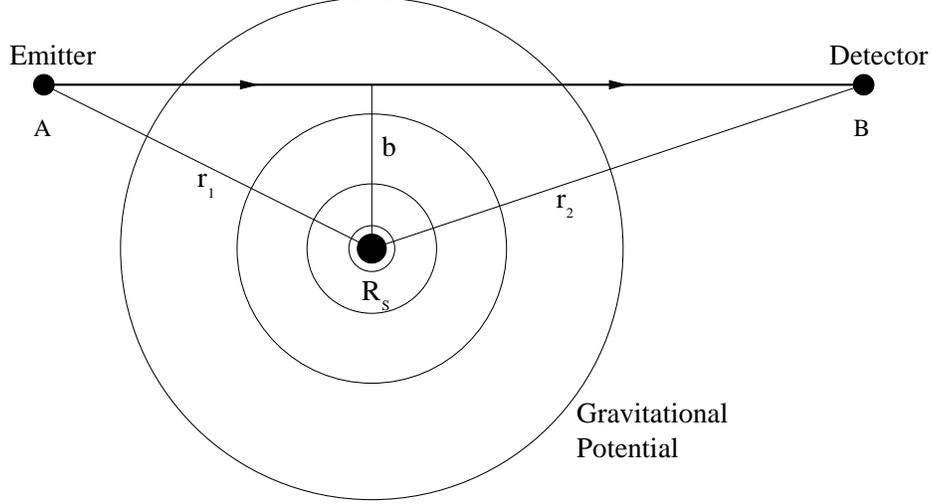


Figure B.4. General relativistic time delay.

We consider the general relativistic part (second bracket) in the limit of $r_1, r_2 \gg b \gg R_S$. The travel time reduces to

$$\begin{aligned} ct_{GR} &\approx R_S \cdot \ln\left(\frac{r_1 r_2}{b^2} [1 + \sqrt{1 - (b/r_1)^2}] [1 + \sqrt{1 - (b/r_2)^2}]\right) \\ &\approx R_S \cdot \left(\ln \frac{4r_1 r_2}{b^2} - 1\right). \end{aligned} \quad (\text{B.3})$$

Now consider two slightly different paths, one with impact parameter b , and the other with $b + db$. The time difference between these two paths is

$$\begin{aligned} \Delta t_{GR} &\equiv t_{GR}(b + db) - t_{GR}(b) \\ &= \frac{2R_S}{c} \cdot \ln \frac{b + db}{b} \\ &\approx \frac{2R_S r_1 \psi}{bc} \end{aligned} \quad (\text{B.4})$$

Note that $\Delta t_{geo} = \Delta t_{GR} \cdot \frac{r_2 \psi}{4b} \approx \Delta t_{GR} \cdot \frac{db}{4b}$, i.e., $\Delta t_{geo} \ll \Delta t_{GR}$.

To get a sense for how large this effect is, let's plug in some numbers. For a close scattering (e.g., $b = 0.01$ ly) off of a large mass (say, $100 M_\odot$),

$$t_{GR} = 3 \times 10^{-6} \text{ ns} \left(\frac{M}{100 M_\odot}\right) \cdot \left(\frac{b}{0.01 \text{ ly}}\right)^{-1} \cdot \left(\frac{r_1}{10 \text{ ly}}\right) \cdot \left(\frac{\psi}{5 \text{ nas}}\right),$$

where we have taken the opening angle between extremal rays that strike our detector to be 5 nanoarcseconds (100 m over a baseline of 10 ly). It is difficult to imagine that millions of such scatterings will conspire in such a way as to give a detectable broadening of nanosecond pulses for ranges of kiloparsecs to megaparsecs.

B.4 Conclusions

Pulse broadening has an insignificant effect on nanosecond optical pulses traversing the galactic medium for both convergent and divergent gravitational scattering. In the convergent case, the probability of aligning a scatterer on the line of sight is extremely small. For divergent scattering, the magnitude of the effect is negligible. The path length dispersion due to atmospheric seeing is likewise entirely negligible.

Appendix C

PulseNet – Automated Testing Procedure

Because PulseNet has a large number of inputs and has two independent modes of operation (SETI and Astronomy), testing all or even more signals paths was impractical when done by a person. An automated test was the only reasonable way to verify the chips.

C.1 AutoTest procedure

Fifty-three PulseNet Rev2 chips were tested using AutoTest, an graphical automated test program written in Python. For consistency, the following sequential protocol was used.

1. **Label chip with serial number** of the form sn001 and glue down lid (unless already sealed by MOSIS).
2. **Measure ring oscillator frequency** at $V_{DD} = 2.4\text{ V}$ (2.41 V as measured on power supply meter), $V_{DD} = 2.5\text{ V}$ (2.51 V on meter), and $V_{DD} = 2.6\text{ V}$ (2.61 V on meter) with clocks disabled. Remember to reset the Astronomy and SETI ports prior to making the measurements.
3. **Measure maximum clock frequency**¹ at $V_{DD} = 2.5\text{ V}$ (2.59 V on meter) and $V_{DD} = 2.7\text{ V}$ (2.78 V on meter). Set the f_{fastclk} output to 17.5 dBm on the signal generator. Start at $f_{\text{fastclk}} = 450\text{ MHz}$ so as to not blow out the clock inputs with a large signal.
4. **Measure minimum clock amplitude** for stable operation at $f_{\text{fastclk}} = 400\text{ MHz}$ and $V_{DD} = 2.50\text{ V}$ (2.59 V on meter). Unstable operation is characterized by a precipitous drop in I_{DD} as the f_{fastclk} is reduced.
5. **Measure input sampler offsets with Astronomy channel** for all combinations of pixels and thresholds (28 minutes). This is explained below in §C.3.
6. **Verify Seti coincidence trigger and pulse recording circuitry** for all pixel pairs (13 minutes). This is explained below in §C.2.

C.2 SETI functions

AutoTest verified the functionality of the SETI components of PulseNet by simultaneously sending a specially tailored waveform to each pixel pair, sequentially. The waveform consists of twenty negative pulses with amplitudes evenly spaced between 1500 mV and 75 mV. AutoTest then retrieved and plotted the coincident waveform data. An example is shown in Fig. C.1. A closeup of the first forty data points appears in the upper left corner of the plot.

¹A Hewlett Packard 8648C Signal Generator was used to generate the fastclock inputs.

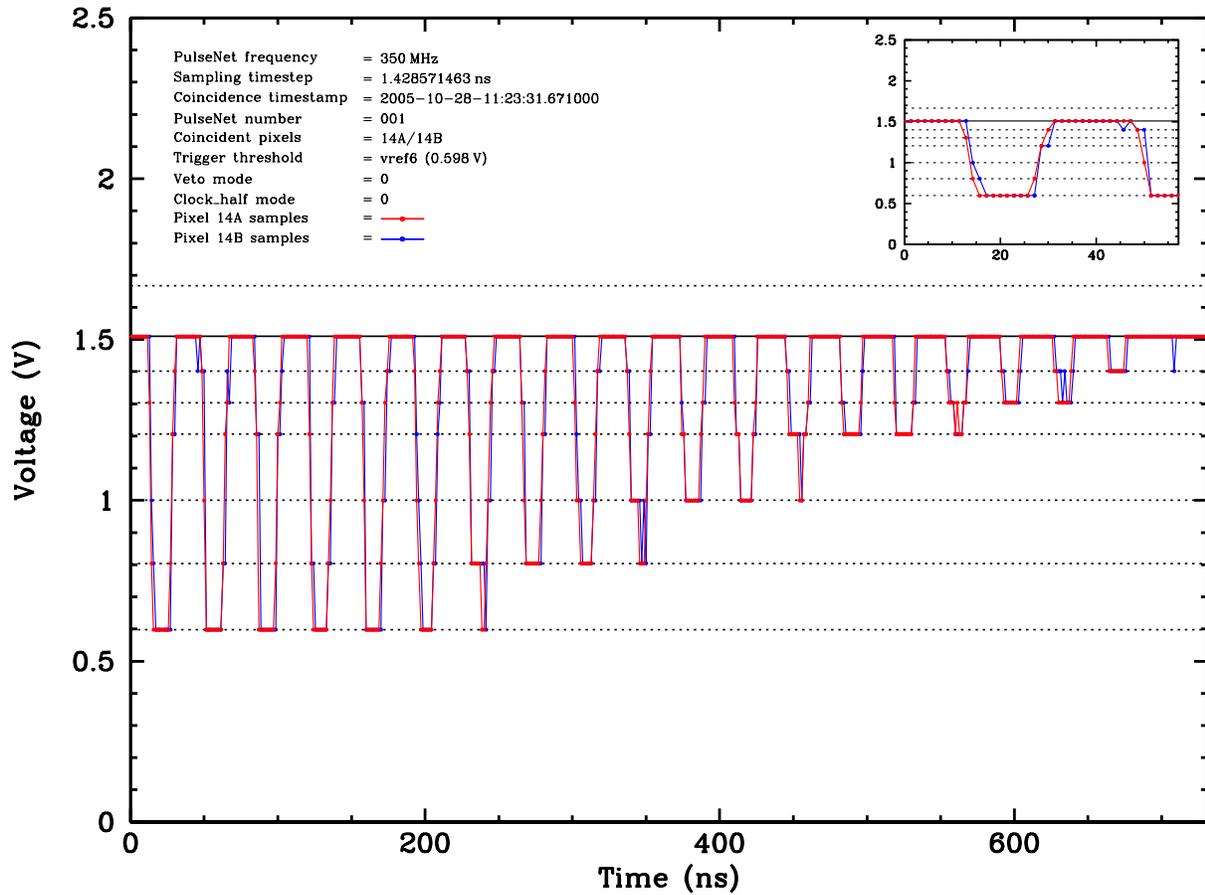


Figure C.1. Example output from automated coincidence functionality test for one pixel pair. The inset plot (upper right) shows the first forty data points.

Note the placement of the voltage thresholds in Fig. C.1 (the dashed horizontal lines). $V_{\text{ref}}[3:1]$ are closely spaced and clustered below V_{bias} to resolve small pulses. $V_{\text{ref}}[6:4]$ are spaced farther apart below $V_{\text{ref}}[3:1]$ to increase dynamic range. $V_{\text{ref}}[0] > V_{\text{bias}}$ for vetoing signals that exceed V_{bias} .

Fig. C.1 also displays useful test parameters, e.g. the PulseNet serial number, the pixel pair and voltage threshold that triggered the coincidence, and the sampling rate. These data are automatically passed from AutoTest to the plotting program.

To test all analog inputs, AutoTest loops through the pixel pairs, each time programming and reading out the coincident waveforms, like that ones shown in Fig. C.1. These sixteen plots are arrayed in another plot, like the one shown in Fig. C.2 which is not designed for detailed checking (the data files are saved if one is really interested), but is intended to verify, at a glance, that the coincidence functions of the PulseNet in test operate properly. One can see immediately by looking at Fig. C.2 that the waveforms recorded by all sixteen pixel pairs are as expected and approximately the same. One can also verify that each pixel pair received a coincidence during the right test by reading the pixel numbers in each sub-plot (e.g. “sn001-0A/0B”) and checking that they are in order on the larger plot.

The SETI coincidence tests were all done at $f_{\text{fastclk}} = 350$ MHz. This speed was a compromise between wanting be above the expected clock speed in the experiment ($f_{\text{fastclk}} = 333$ MHz) and the maximum operating clock speed ($f_{\text{fastclk}} = 500$ MHz when run at $V_{\text{dd}} = 2.87$ V).

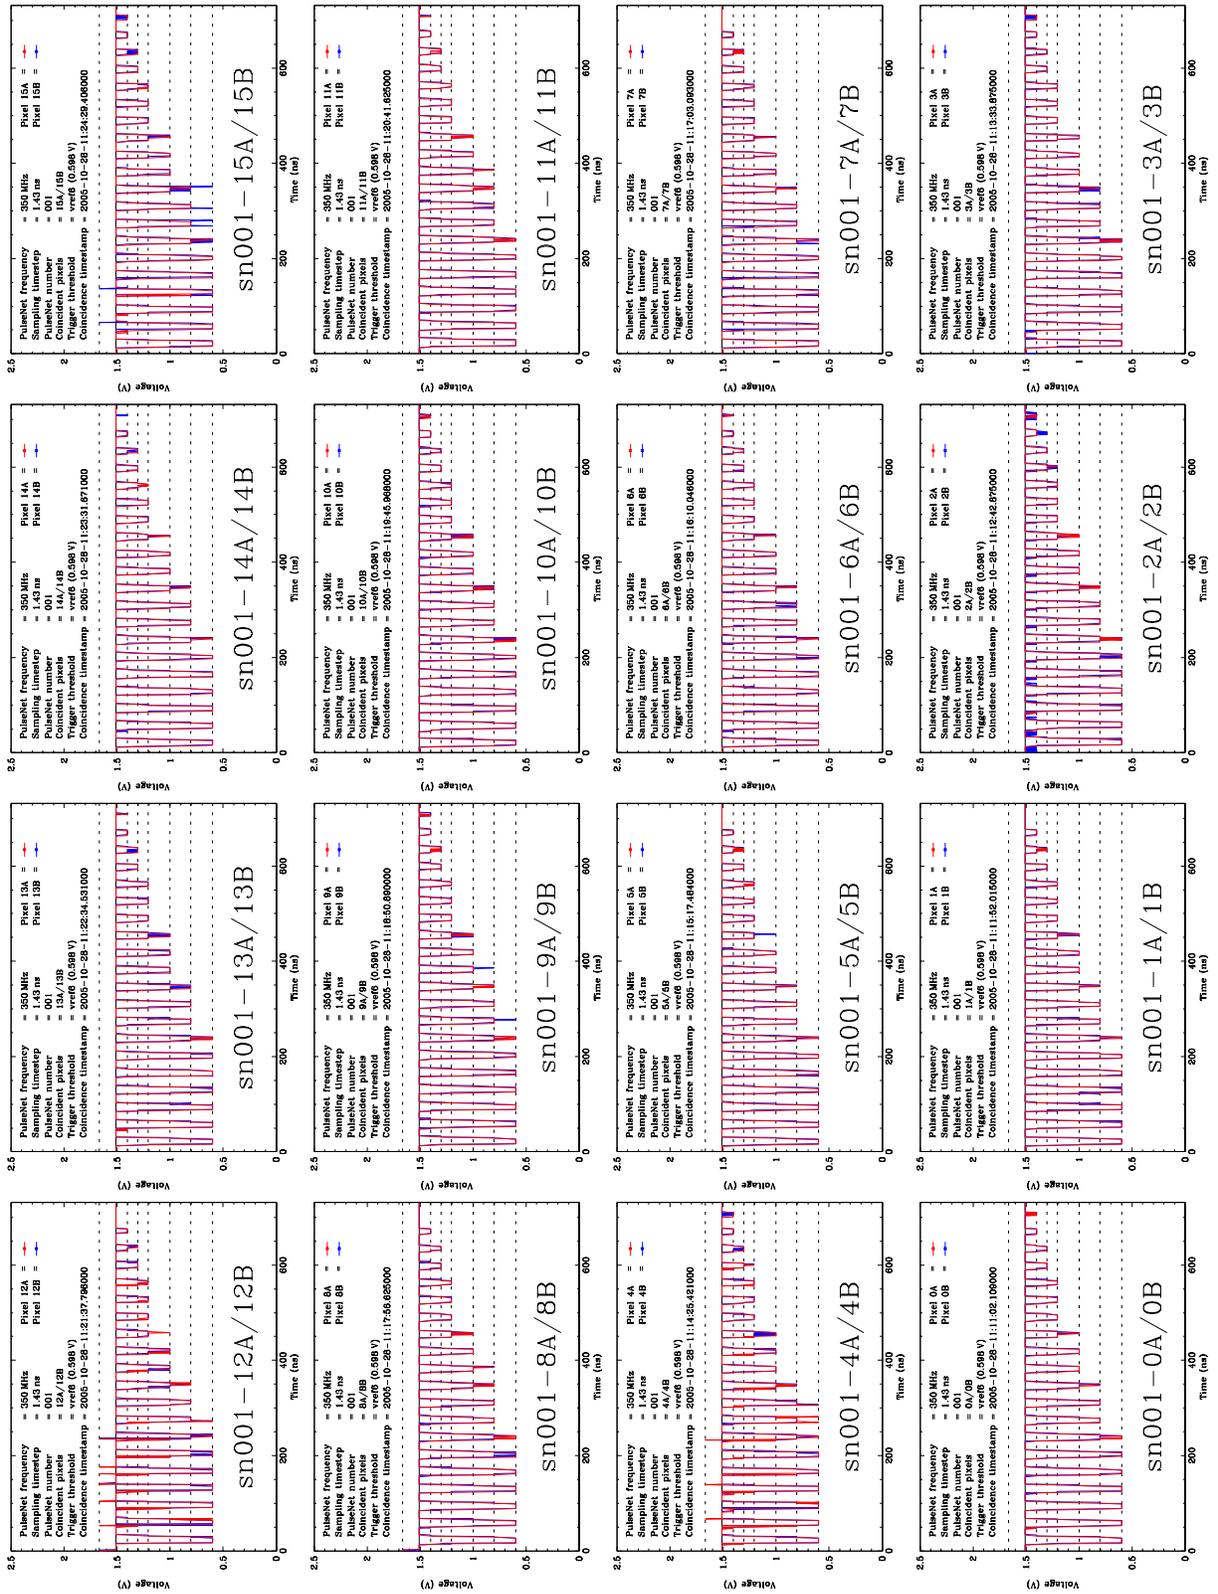


Figure C.2. Example output from automated coincidence functionality test for all pixels.

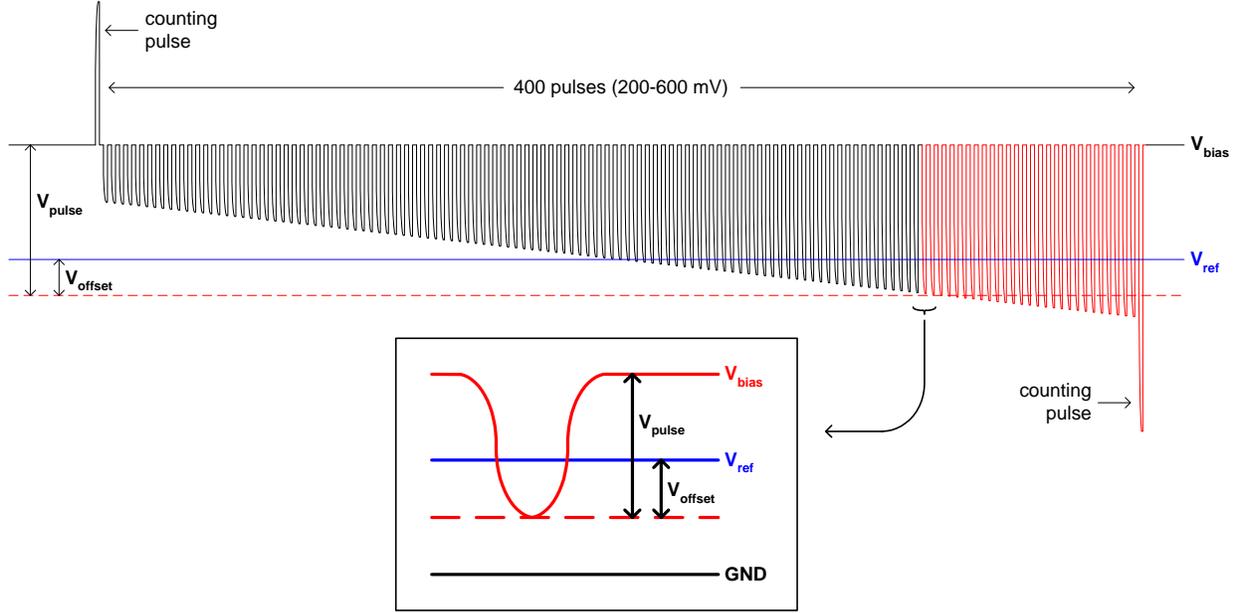


Figure C.3. Diagram of automated voltage offset test waveform (“ramped comb”) consisting of 400 negative pulses with amplitudes between 200-600 mV and two counting pulses. Only those pulses with amplitudes greater than V_{pulse} (shown in red) are counted. The lower inset diagram shows the relationship between various set and measured voltages for the minimum amplitude pulse that triggers PulseNet.

C.3 Astronomy functions

AutoTest checked the astronomy functions on PulseNet by sending a waveform consisting of many pulses of increasing amplitude (a “ramped comb”, illustrated in Fig. C.3) individually to each pixel pair. The amplitudes of the 400 negative pulses were evenly spaced from 200 mV to 600 mV at 1 mV/pulse. Prepend and appended to this sequence were two “counting pulses” (more on these below), positive 500 mV and negative 1000 mV, respectively. This waveform was generated using a Tektronix AFG Arbitrary/Function Generator.

The astronomy counting functions were tested and the sampler offset voltages were measured in one pixel pair and threshold by programming PulseNet to count for a fixed duration (~ 2 seconds) with the reference voltage (V_{ref}) for the threshold in question set between 200 and 600 mV below the input bias voltage, V_{bias} (see Fig. C.3). For each ramped comb, only those pulses with amplitude greater than $V_{\text{bias}} - V_{\text{ref}}$ plus the offset voltage (V_{offset}) trigger the analog sampler and included in the count for one sense amplifier, n_1 (shown in red in Fig. 6.2).

During the ~ 2 second counting interval, the samplers on PulseNet see many (>1000) ramped comb waveforms. The precise number is determined by a second measurement in which the reference voltage is changed to either greater than V_{bias} or less than ($V_{\text{bias}} - 600$ mV), and counting for the same ~ 2 second interval. In this case, only the counting pulses will be counted so that the values on read off of PulseNet’s four counters (n_2) will be identical and equal to the number of ramped combs during each counting interval. Thus we can calibrate out the counter duration: the number of pulses counted in each ramped comb is the quotient of these two counts, $n_{\text{count}} = n_1/n_2$.

Now, the ramped comb waveform is only useful if we can convert n_{count} into V_{offset} . We do this by comparing two expressions for V_{pulse} , the amplitude of smallest pulse in the ramped comb waveform that is counted by PulseNet. The first expression for V_{pulse} comes from examining the upper portion of Fig. C.3,

$$V_{\text{pulse}} = (V_{\text{bias}} - V_{\text{ref}}) + V_{\text{offset}}. \quad (\text{C.1})$$

Table C.1. Serial numbers of PulseNets that failed AutoTest

Serial Number	Reason for Failure
sn011	V_{DD}/GND short — bond wires bent square
sn015	V_{DD}/GND short — foreign material seen on chip through microscope
sn027	no signal on three analog inputs (0A, 1A, 1B); V_{bias} was high (short?)
sn029	V_{DD}/GND short
sn031	flakey performance; occasionally no coincidences on 2A/2B and 3A/3B

Note. — Serial numbers of PulseNets that failed AutoTest for the reasons indicated. All other PulseNets with serial numbers in the range sn000–sn052 passed the tests listed in §C.1.

The second expression comes from examining the inset closeup of the minimum amplitude pulse in Fig. C.3,

$$V_{\text{pulse}} = 200 \text{ mV} + \left(400 \text{ mV} - (n_{\text{count}} - 1) \cdot \frac{1 \text{ mV}}{\text{count}} \right). \quad (\text{C.2})$$

Combining eqs. C.1 and C.2 we get an expression for the offset voltage in terms of measured quantities:

$$V_{\text{offset}} = 600 \text{ mV} - (n_{\text{count}} - 1) \cdot \frac{1 \text{ mV}}{\text{count}} - (V_{\text{bias}} - V_{\text{ref}}). \quad (\text{C.3})$$

Note that although the measurement outlined above was for V_{offset} on one sense amplifier, it actually yields the offset voltages for *four* sense amplifiers — both clocks phases (a/b) on one pixel pair (A/B), for the given threshold — since PulseNet has four counters.

A successful test of the astronomy functions using AutoTest tells the user two things: 1) that all 448 sense amplifiers work properly, and 2) the values of V_{offset} for each sense amplifier. In AutoTest, the user enters the reference voltage levels, the thresholds and pixel pairs to be tested, the number of times each combination should be tested. In twenty-eight minutes, AutoTest loops through the variables and outputs files with the raw counter values, which are converted to offset voltages using Eq. C.3.

C.4 Results

All chips in the serial number range sn001–sn053 passed the tests listed in §C.1, except for the chips listed in Table C.1, which failed for the reasons indicated.

Bibliography

- [1] E. B. Amôres and J. R. D. Lépine. Models for Interstellar Extinction in the Galaxy. *AJ*, 130:659–673, August 2005.
- [2] D. E. Backman, L. J. Caroff, S. A. Sandford, and D. H. Wooden, editors. *Exozodiacal Dust Workshop*, April 1998.
- [3] C. A. Beichman, N. J. Woolf, and C. A. Lindensmith. *The Terrestrial Planet Finder (TPF) : a NASA Origins Program to Search for Habitable Planets*. Pasadena, Calif. (JPL publication 99-3), 1999.
- [4] A. Betz. A Search for Infrared Laser Signals. In S. Shostak, editor, *Third Decennial US-USSR Conference on SETI*, volume 47, pages 373–379, 1993.
- [5] R. Bhathal. Optical SETI in Australia. In S. A. Kingsley and R. Bhathal, editors, *Proc. SPIE Vol. 4273, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*, pages 144–152, August 2001.
- [6] J. M. Billingham and B. M. Oliver. *Project Cyclops: A Design Study of a System for Detecting Extraterrestrial Intelligent Life*. NASA/Ames Pub. CR-114445, Moffet Field, Calif., 1972.
- [7] J. A. Cardelli, G. C. Clayton, and J. S. Mathis. The Relationship Between Infrared, Optical, and Ultraviolet Extinction. *ApJ*, 345:245–256, October 1989.
- [8] M. Catanese and T. C. Weekes. Very High Energy Gamma-ray Astronomy. *PASP*, 111:1193–1222, October 1999.
- [9] Clark R. Champman and Dale P. Cruikshank. *Observing the Moon, Planets and Comets*. Schramm and Groves, Laguna Niguel, CA, 1980.
- [10] G. Cocconi and P. Morrison. Searching for Interstellar Communications. *Nature*, 184:844–846, 1959.
- [11] C. Coldwell. *A Search for Interstellar Communications at Optical Wavelengths*. PhD thesis, Harvard University, 2002.
- [12] J. M. Cordes. *SETI 2020*. SETI Press, Mountain View, Calif., 1st edition, 2002.
- [13] J. M. Cordes and T. J. Lazio. Interstellar Scattering Effects on the Detection of Narrow-band Signals. *ApJ*, 376:123–133, July 1991.
- [14] C. E. Covault. Large Area Solar Power Heliostat Array for OSETI. In S. A. Kingsley and R. Bhathal, editors, *Proc. SPIE Vol. 4273, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*, pages 161–172, August 2001.
- [15] A. N. Cox. *Allen’s Astrophysical Quantities*. New York: AIP Press; Springer, 2000.
- [16] J. W. Cronin. Cosmic Rays: the Most Energetic Particles in the Universe. *Reviews of Modern Physics Supplement*, 71:165–172, March 1999.
- [17] W. J. Dally and J. W. Poulton. *Digital Systems Engineering*. Cambridge University Press, New York, 1998.

- [18] S. J. Dick. *The Biological Universe*. Cambridge University Press, Cambridge, 1999.
- [19] R. D’Inverno. *Introducing Einstein’s Relativity*. Oxford University Press, Oxford, 1992.
- [20] F. D. Drake. How Can We Detect Radio Transmissions from Distant Planetary Systems? *S&T*, 19:140, January 1960.
- [21] D. Dravins. Astrophysics on its shortest timescales. *The Messenger*, 78:9–19, 1994.
- [22] F. Dyson. Looking For Life In Unlikely Places: Reasons Why Planets May Not Be The Best Places To Look For Life. *Int. J. Astrobiology*, 2:103–110, 2003.
- [23] D. Eichler and G. Beskin. Optical Search for Extraterrestrial Intelligence with Air Cerenkov Telescopes. *Astrobiology*, 1:489–493, December 2001.
- [24] R. Ekers, K. Cullers, J. Billingham, and L. Scheffer, editors. *SETI 2020*. SETI Press, Mountain View, Calif., 1st edition, 2002.
- [25] E. L. Fitzpatrick. Correcting for the Effects of Interstellar Extinction. *PASP*, 111:63–75, January 1999.
- [26] R. H. Golde, editor. *Lightning*. Academic Press, London, 1977.
- [27] E. Hecht. *Optics, 2nd edition*. Reading, MA: Addison-Wesley Publishing Company, 1987.
- [28] H. Hemmati. Overview of Laser Communications Research at JPL. In S. A. Kingsley and R. Bhathal, editors, *Proc. SPIE Vol. 4273, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*, pages 190–193, August 2001.
- [29] J. Holder, P. Ashworth, S. LeBohec, H. J. Rose, and T. C. Weekes. Optical SETI with Imaging Cherenkov Telescopes. In *29th International Cosmic Ray Conference, Pune*, pages 101–106, 2005.
- [30] P. Horowitz, C. M. Coldwell, A. W. Howard, D. W. Latham, R. Stefanik, J. Wolff, and J. M. Zjajac. Targeted and All-sky Search for Nanosecond Optical Pulses at Harvard-Smithsonian. In S. A. Kingsley and R. Bhathal, editors, *Proc. SPIE Vol. 4273, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*, pages 119–127, August 2001.
- [31] P. Horowitz and C. Sagan. Five Years of Project META - An All-sky Narrow-band Radio Search for Extraterrestrial Signals. *ApJ*, 415:218–235, September 1993.
- [32] A. W. Howard and P. Horowitz. Is There RFI in Pulsed Optical SETI? In S. A. Kingsley and R. Bhathal, editors, *Proc. SPIE Vol. 4273, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*, pages 153–160, August 2001.
- [33] A. W. Howard and P. Horowitz. Optical SETI with NASA’s Terrestrial Planet Finder. *Icarus*, 150:163–167, March 2001.
- [34] A. W. Howard, P. Horowitz, C. Coldwell, S. Klein, A. Sung, J. Wolff, J. Caruso, D. Latham, C. Paliolios, R. Stefanik, and J. Zjajac. Optical SETI at Harvard-Smithsonian. In G. Lemarchand and K. Meech, editors, *ASP Conf. Ser. 213: Bioastronomy 99*, pages 545–552, 2000.
- [35] 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, and J. M. Zjajac. Search for Nanosecond Optical Pulses from Nearby Solar-Type Stars. *ApJ*, 613:1270–1284, October 2004.
- [36] J. F. Kasting, D. P. Whitmire, and R. T. Reynolds. Habitable Zones around Main Sequence Stars. *Icarus*, 101:108–128, January 1993.
- [37] S. A. Kingsley. The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum: a Review. In S. A. Kingsley, editor, *Proc. SPIE Vol. 1867, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum*, pages 75–113, August 1993.

- [38] S. A. Kingsley. Prototype optical SETI observatory. In S. A. Kingsley and G. A. Lemarchand, editors, *Proc. SPIE Vol. 2704, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum II*, pages 102–116, June 1996.
- [39] W. F. Krupke. Diode-pumped Solid State Lasers (DPSSLs) for Inertial Fusion Energy (IFE). In *Proc. SPIE Vol. 3047, Solid State Lasers for Application to Inertial Confinement Fusion: Second Annual International Conference, Michel L. Andre; Ed.*, pages 73–82, December 1997.
- [40] M. Lampton. Optical SETI: The Next Search Frontier. In G. Lemarchand and K. Meech, editors, *ASP Conf. Ser. 213: Bioastronomy 99*, pages 565–570, 2000.
- [41] D. W. Latham. Surveys of Spectroscopic Binaries at the Center for Astrophysics. In *ASP Conf. Ser. 32: IAU Colloq. 135: Complementary Approaches to Double and Multiple Star Research*, pages 110–118, 1992.
- [42] D. W. Latham. Target Selection for Planet Searches. In G. Lemarchand and K. Meech, editors, *ASP Conf. Ser. 213: Bioastronomy 99*, pages 137–141, 2000.
- [43] M. Lee and P. Dally, W. Chiang. Low-Power Area-Efficient High-Speed I/O Circuit Techniques. *IEEE J. Solid-State Circuits*, 35:1591–1599, Nov. 2000.
- [44] C. H. Lineweaver and D. Grether. What Fraction of Sun-like Stars Have Planets? *ApJ*, 598:1350–1360, December 2003.
- [45] R. Lodder. Near-infrared (IR) Spectrometric Imaging Using a one-meter Fresnel Telescope. In P. Shuch, editor, *Proceedings of SETICon 01*, 2001.
- [46] Lucretius. *De Rerum Natura (On the Nature of Things)*, ca. 70 B.C. In “The Quest for Extraterrestrial Life: A Book of Readings”, ed. D. Goldsmith, pages 3–4, 1980.
- [47] T. Neckel, G. Klare, and M. Sarcander. The Spatial Distribution of the Interstellar Extinction. *A&AS*, 42:251–281, November 1980.
- [48] M. J. M. Pelgrom, A. C. J. Duinmaiger, and A. P. G. Welbers. Matching Properties of MOS Transistors for Precision Analog Design. *IEEE J. Solid-State Circuits*, 24:1433–1439, Oct. 1989.
- [49] M. A. C. Perryman *et al.* The HIPPARCOS Catalogue. *A&A*, 323:L49–L52, July 1997.
- [50] W. T. Reach, B. A. Franz, J. L. Weiland, M. G. Hauser, T. N. Kelsall, E. L. Wright, G. Rawley, S. W. Stemwedel, and W. J. Spiesman. Observational Confirmation of a Circumsolar Dust Ring by the COBE Satellite. *Nature*, 374:521, April 1995.
- [51] A. E. Reines and G. W. Marcy. Optical Search for Extraterrestrial Intelligence: A Spectroscopic Search for Laser Emission from Nearby Stars. *PASP*, 114:416–426, April 2002.
- [52] Monte Ross. Search via Laser Receivers for Interstellar Communications. *Proc. IEEE*, 53:1780, 1965.
- [53] B. Rossi. Interpretation of Cosmic-ray Phenomena. *Reviews of Modern Physics*, 20:537–583, July 1948.
- [54] C. Sagan. On the Detectivity of Advanced Galactic Civilizations. *Icarus*, 19:350–358, July 1973.
- [55] R. Schwartz and C. Townes. Interplanetary and Interstellar Communication by Optical Masers. *Nature*, 190:205–208, 1961.
- [56] A. Segura, J. F. Kasting, V. Meadows, M. Cohen, J. Scalo, D. Crisp, R. A. H. Butler, and G. Tinetti. Biosignatures from Earth-Like Planets Around M Dwarfs. *Astrobiology*, 5:706–725, December 2005.
- [57] V. Shvartsman, G. Beskin, S. Mitronova, S. Neizvestny, V. Plakhotnichenko, and L. Pustil’nik. Results of the MANIA Experiment: an Optical Search for Extraterrestrial Intelligence. In S. Shostak, editor, *ASP Conf. Ser. 47: Third Decennial US-USSR Conference on SETI*, page 381, 1993.

- [58] R. E. Slusher. Laser Technology. *Reviews of Modern Physics Supplement*, 71:471–479, March 1999.
- [59] R. P. S. Stone, S. A. Wright, F. Drake, M. Munoz, R. Treffers, and D. Werthimer. Lick Observatory Optical SETI: Targeted Search and New Directions. *Astrobiology*, 5:604–611, October 2005.
- [60] D. W. Swift. *SETI Pioneers — Scientists Talk About Their Search for Extraterrestrial Intelligence*. Tucson, AZ, University of Arizona Press, 1990.
- [61] J. Tarter. The Search for Extraterrestrial Intelligence (SETI). *Annual Reviews of Astronomy & Astrophysics*, 39:511–548, 2001.
- [62] C. Townes. At What Wavelengths Should We Search for Signals from Extraterrestrial Intelligence? *Proc. Natl. Acad. Sci.*, 80:1147–1151, 1983.
- [63] M. C. Turnbull and J. C. Tarter. Target Selection for SETI. I. A Catalog of Nearby Habitable Stellar Systems. *ApJS*, 145:181–198, March 2003.
- [64] M. C. Turnbull and J. C. Tarter. Target Selection for SETI. II. Tycho-2 Dwarfs, Old Open Clusters, and the Nearest 100 Stars. *ApJS*, 149:423–436, December 2003.
- [65] S. C. Unwin. Precision Astrometry with the Space Interferometry Mission - PlanetQuest. In P. K. Seidelmann and A. K. B. Monet, editors, *ASP Conf. Ser. 338: Astrometry in the Age of the Next Generation of Large Telescopes*, page 37, September 2005.
- [66] P. Ward and D. Brownlee. *Rare Earth: Why Complex Life is Uncommon in the Universe*. New York : Copernicus, 2000.
- [67] D. Werthimer, D. Anderson, C. S. Bowyer, J. Cobb, E. Heien, E. J. Korpela, M. L. Lampton, M. Lebofsky, G. W. Marcy, M. McGarry, and D. Treffers. Berkeley Radio and Optical SETI Programs: SETI@home, SERENDIP, and SEVENDIP. In S. A. Kingsley and R. Bhathal, editors, *Proc. SPIE Vol. 4273, p. 104-109, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*, pages 104–109, August 2001.
- [68] S. A. Wright, F. Drake, R. P. Stone, D. Treffers, and D. Werthimer. Improved Optical SETI Detector. In S. A. Kingsley and R. Bhathal, editors, *Proc. SPIE Vol. 4273, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*, pages 173–177, August 2001.
- [69] M. M. Yusupov, G. Z. Yusupova, A. Baucom, K. Lieberman, T. N. Earnest, J. H. D. Cate, and Harry F. Noller. Crystal Structure of the Ribosome at 5.5 Å Resolution. *Science*, 292(5518):883–896, May 2001.
- [70] M. Zeilik and S. Gregory. *Introductory Astronomy and Astrophysics*. Saunders College Publishing, Fort Worth, 4th edition, 1998.

Glossary

arcminute angle of 1/60th of a degree.

arcsecond angle of 1/3600th of a degree; approximately equal to $5 \mu\text{rad}$.

as arcsecond

AU astronomical unit – The distance from the Earth to the Sun.

CMOS complimentary metal-oxide semiconductor – a major class of integrated circuits that uses complimentary transistors (NMOS and PMOS) to achieve relatively low power and high density designs.

dec declination, sometimes abbreviated δ – north/south coordinate in the equatorial coordinate system. Range = -90° to $+90^\circ$.

DS_{all} data set from the targeted search – all observations made from Harvard. See §4.5.

DS_{clean} data set from the targeted search – a subset of DS_{all} from which nights with anomalous trigger rates have been removed. See §4.5.

DS_{overlap} data set from the targeted search – a subset of DS_{all} during which Princeton jointly observed. See §4.5.

FIFO first in first out – a queue-like memory device.

GHz gigahertz – frequency of 10^9 s^{-1}

GS/s gigasamples per second, 10^9 samples per second.

I/O input/output

IC integrated circuit

ISM interstellar medium

kly 1000 light-years

km kilometers

kpc 1000 parsecs

LST Local Sidereal Time – Greenwich Mean Sidereal Time plus an offset for the longitude of the observer. See also *sidereal time*.

ly light-year – the distance light travels in one year, 9.5×10^{15} m.

m meters

μ as microarcsecond

mas milliarcsecond

MHz megahertz – frequency of 10^6 s^{-1}

ns nanosecond – 10^{-9} s

pc parsec – 3.26 light-years

PCB printed circuit board

PMT photomultiplier tube – a type of photodetector

PulseNet a full-custom integrated circuit designed for the all-sky survey. It digitizes sixteen pair of analog inputs at up to 500 MHz, triggers on and stores coincident pulses, and measures count rates.

RA right ascension – east/west coordinate in the equatorial coordinate system. Range = 0 to 24 hours.

SETI search for extraterrestrial intelligence

sidereal time a measure of the rotation of the Earth with respect to the stars rather than the Sun. One sidereal day is the time it take for the Earth to rotate once with respect to the stars and is approximately 23 hr 56 min 4 sec.

SIM Space Interferometry Mission – A planned space telescope NASA recently changed the name of this mission to SIM PlanetQuest.

SNR signal-to-noise ratio

TPF Terrestrial Planet Finder – currently a set of two proposed satellite telescopes (an infrared nulling interferometer and an optical coronagraph) designed to observe the atmospheric spectra of extra-solar planets in search of life-indicating features.

UI user interface — the web page by which an operator controls the all-sky instrument.

UPS uninterruptible power supply