The Technical Case for Optical and Infrared SETI

(revised 9-Sept-98)

I. Introduction

In this section we follow the lead of Charlie Townes, who has patiently been trying to explain to the SETI community the good arguments for taking optical SETI seriously. (For full gory technical detail one can do no better than to read his arguments, for example in "At what wavelengths should we search for signals from extraterrestrial intelligence?," *Proc Nat Acad Sci*, **80**, 1147 [1961].)

Historically the Coccini and Morrison suggestion that SETI be carried out at the 21 cm 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 two years – 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 (by Drake and Helou, and more recently by Cordes and Lazio) 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 vs 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 high resolution spectroscopy one must use such a system, mixing the optical frequency 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). 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.

Townes [refs] has made a comparison of received SNR vs wavelength, making reasonable assumptions about antenna apertures and accuracies, detection methods, transmitter power, and so on. The bottom line is 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 the advantages of pulsing and high data rates (which favors optical) – and could easily tip the balance. The conclusion is that the SETI community's historical bias toward microwaves should surely be reconsidered.

Laser technology is in a phase of rapid catchup relative to the mature technology at radio frequencies. Recently lasers with a megawatt of continuous optical output have been built, and picosecond pulses of 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 that would produce repetitively pulsed megajoule nanosecond pulses (see below). As we will show below, optical pulsed beacons formed with that sort of technology permit detection with very simple apparatus – just a telescope with a pair of white-light photomultipliers in coincidence.

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? We consider this question in the next section.

II. Pulses and Carriers

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 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 *milli*hertz 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. In other words, a signal that is a spike in the *frequency* domain is a natural candidate for interstellar signalling at microwave frequencies, for reasons both scientific and technical.

Moreover, interstellar dispersion, and the presence of natural and "cultural" impulsive interference (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 over the same channel.

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. Furthermore, at optical wavelengths the higher carrier frequencies ($\sim 10^{14}$ Hz) result in much larger absolute Doppler shifts; for example, 1 km/s ↔ 5 kHz at 1.4 GHz, whereas 1 km/s ↔ 1 GHz at 1 μm. However, dispersion is negligible at optical wavelengths, even at the nanosecond level. Furthermore, natural and "cultural" sources of nanosecond flashes are probably absent (see §IV). In other words, a signal that is a spike in the *time* domain is a natural candidate for interstellar signalling at optical wavelengths, for reasons both scientific and technical. An added bonus is that, at nanosecond time scales, the stellar background becomes negligible (see §IV).

The above considerations suggest that a pulsed laser beacon is an attractive scenario. Indeed, in §IV below we will sketch plausible parameters for such a beacon, which indeed is entirely reasonable and can accomplish the task of making contact with a civilization out to $\sim$ 1000 light years or more. However, it would be wrong to ignore the possibility of an optical spectral laser line as a means of contact; in fact, Dr. Geoff Marcy (co-discoverer of extrasolar planets) is going to "mine" some of his group's high-resolution stellar spectra for possible laser light.

III. Pulsed Beacon

We seek a workable strategy for a targeted multiplexed optical pulse beacon. That is, we assume that the sending civilization wishes to irradiate, in sequence, the planetary zones of the nearest N sunlike stars (or an intelligent subset, if they know more), going out to a range $R_{\max}$ comparable to the mean distance between advanced civilizations. Order-of-magnitude we might assume N is in the range of 10^3 to 10^6, corresponding to distances $R_{\max}$ of 100–1000 light years (ly). We ask what laser power and pulse width are needed to outshine the sender's star, and to deliver a substantial number of photons to the receiving system. Are such lasers feasible, now or in the foreseeable future? Is there a preferred wavelength regime, or perhaps a "magic" wavelength? What about extinction and dispersive smearing? Are there serious terrestrial or astrophysical backgrounds? And finally, is there a compelling parameter set (analogous, say, to the use of an RF carrier at the hydrogen hyperfine wavelength of 21 cm) that suggests an obvious OSETI search strategy that we could embark upon?

MULTIPLEXED PULSER

Let us assume that the transmitting civilization has a catalog of target stars, with positions, proper motions, and ranges known with sufficient accuracy to permit aiming to an error no greater than $\sim$10 AU when the beam reaches the target. At a range of 100 ly this corresponds to a positional accuracy of 0.3 arcsec and a proper motion uncertainly of 3 mas/year; at 1000 ly the corresponding figures are 33 mas and 33 μas/year. The required range accuracy depends on the star's proper motion, but not its range: For example, to target the planetary zone (say 10 AU) of a star whose transverse proper motion is 10 km/s the range uncertainty can not exceed 5 ly. These accuracies are relaxed if the transmitted beam is broadened to irradiate a larger zone, at the expense of course of received signal strength.

The positional accuracy is presumably attainable, since if the sending civilization can transmit a beam of the required width they can certainly observe the star's apparent position to the same precision. The required proper motion and range accuracies appear to be completely feasible, certainly for a civilization with good optical technology (at least a century in advance of our own, say) in orbit.

To send a pulse to each of $N = 10^3$ stars (roughly the number of sunlike stars within 100 ly) with a single laser system, the sender would probably use an assembly of fast beamsteering mirrors of relatively small size and weight, in combination with a large objective that is steered slowly. We could build such an arrangement today that would settle in ~0.1 sec to diffraction-limited pointing; so it is surely conservative to assume that they can do as well. This would let them send an optical pulse to 10 stars per second, thus 100 seconds to irradiate 1000 stars. (Of course, to irradiate stars anywhere in the sky they would need either several laser stations on their planet, or an orbiting station; the latter is probably better, anyway.) It is interesting to note that pulsed laser systems on the drawing board today, in particular the LLNL "Helios" laser, has a design repetition rate of 10 pps.

Under this scenario, then, a target star might expect to see an optical pulse, repeated every 100 seconds, coming from a particular one of the 1000 nearest stars. As we shall show below, in a selected (guessable?) wavelength band, and during the nanosecond duration of the pulse, the optical signal outshines the sending star by some six orders of magnitude; in fact, it outshines even the *unfiltered* (broadband) light of the star by a substantial margin. Moreover, with not unreasonable transmitted laser pulse energy, and not unreasonable telescope sizes at both ends, the received optical signal contains ~10^2 photons per pulse. It is not difficult to imagine an OSETI program that spends a few minutes observing each of a thousand stars; in fact, the job could be done in just a few clear nights, perhaps even by amateur astronomers.

For the more difficult scenario of $N = 10^6$ stars the revisit time stretches to 10^5 seconds, roughly a day. The recipient's task is now to observe 10^6 stars, each for a day – no longer a reasonable OSETI program. The searcher's task scales as N^2, with N itself scaling as R_{max}^3. I.e., the effort scales as R_{max}^6. Clearly one needs to exploit some parallelism, perhaps at both ends: The sender can use a network of transmitters; the target can mobilize the amateur astronomers, who have long since become victims of variable-star-ennui and desperately need a new challenge. What amateur wouldn't devote considerable time and effort to a simple observing program that, with something like one in a million odds, could make the greatest discovery in humankind?

PULSE DETECTOR

A receiving system to detect ~100 photons in a nanosecond, coming from an identifiable nearby star, could be extremely simple. Assuming that no great wavelength specificity is needed, one can simply follow the lead of Dan Werthimer at Berkeley, using a reflecting telescope of modest aperture (say one meter), with an optional multilayer filter at the chosen wavelength, followed by a beamsplitter and pair of photomultiplier tubes (PMTs) behind matched apertures in the focal plane. The electronics could be a pair of pulse height discriminators (to reject dark counts and single photoelectron events) driving a coincidence circuit. The observer points the telescope at each star in turn, guiding on the PMT "singles" rates (upstream of pulse height discrimination), and waits for the unique coincidence signature generated by the burst of some hundred photons within the PMT's resolving time of a nanosecond.

Of course, more complicated systems are possible, perhaps desirable: One could time-tag all large multi-photon events, perhaps in a wavelength-dispersive array detector, and look for coincidences post-hoc. In this way you're not committed to any particular threshold, pulse width, or wavelength. The down side is a glut of data – but there is surely a happy middle ground.

LASER POWER REQUIRED

Contemporary Pulsed Lasers – An "Existence Proof"

At LLNL a pulsed laser with near diffraction-limited beam quality has achieved a kilojoule of pulse energy in a picosecond pulse – that's 10^{15} watts (1 PW, a "petawatt"). This is a flashlamp-pumped Nd-glass laser amplifier ("Nova"), operating in a chirp-pulse amplification mode for extremely high peak power; this particular laser cannot be cycled faster than once per hour. Recently, however, the inertial fusion program has been working on prototypes of highly efficient diode-pumped solid-state lasers ("DPSSL", using Yb-doped strontium fluorapatite). This is an elegant design for a scalable architecture of beamlines that is intended to be combined to produce a high quality beam. The full combination – "Helios" – would produce $\sim$3 ns pulses at $\lambda = 349$ nm (tripled from 1047 nm), with a pulse energy $E_p = 3.7$ MJ and a repetition rate of 10 pps. They're not cheap – of order a billion dollars. However, it is not our mission to develop a *transmitting* system, only to think out what might be plausibly achievable so that we can plan rational *receiving* strategies. Although it's not terribly important for the OSETI mission, it's worth noting that these lasers, intended for fusion power plants, deliver very high efficiency – roughly 10% from wall-plug to photons. When thinking about pulsed optical SETI one should not be unduly influenced by the particular parameter set of Helios; rather, one should assume that a "designer laser" can be tailored (in terms of pulse width, power, repetition rate, and wavelength) to the OSETI mission requirements.

Laser Pulse Energy Needed

Let us calculate what laser power is needed for OSETI. Assume a transmitting aperture of D_t aimed at a receiving aperture D_r at range R. We transmit a pulse at wavelength λ, with pulse energy E_p. The energy per photon is $h\nu$, or hc/λ, thus the number of photons per pulse is $E_p\lambda/hc$. The beam size is approximately $\theta_b = \lambda/D_t$, and the directivity ("gain") of the transmitting aperture is $g_t = 4\pi A_t/\lambda^2 = \pi^2 D_t^2/\lambda^2$. (Note this is approximately $1/\theta_b^2$.) So, the fraction of transmitted photons that are received is $f_r = g_t A_r/4\pi R^2 = \pi^2 D_t^2 D_r^2/16\lambda^2 R^2$, and therefore the number of photons received per pulse transmitted is $N_r = \pi^2 D_t^2 D_r^2 E_p/16\lambda R^2 hc$.

To keep things simple let's scale our result to a set of round-number parameters: $D_t = D_r = 1$ m, $\lambda = 1$ μm, $E_p = 1$ J, $R = 1$ pc. For this set of parameters the transmitting directivity is 10^{13} (130 dB), and the number of photons received per pulse is $N_r^o = 3.6 \times 10^{-3}$. Note that this depends on pulse energy, but not pulse width.

Signal-to-Background Ratio

Obviously we need higher pulse energy. First, though, let's compare with the number of stellar photons collected in, say, 1 ns. (We choose 1 ns because that is roughly the speed of PMT's, and easily achievable performance for pulse electronics.)

We begin by assuming that it is necessary to observe in a narrow wavelength band in order not to be overwhelmed by stellar background; we will reconsider this assumption later. A solar luminosity is $L_\odot \approx 4 \times 10^{33}$ erg/s. Let us assume first that we can guess the wavelength to 1 part in 10^4 (that's $\Delta\lambda = 1$Å at $\lambda = 1$ μm; it might be a Fraunhofer dark line, for example), and filter the starlight accordingly. Then the star puts out about 4×10^{29} erg/s (or 4×10^{22} J/s) in $\Delta\lambda$. Within the θ_b beamwidth of our competing laser transmitter, however, the star puts out a fraction $\theta_b^2/4\pi \approx 10^{-13}$. I.e., the star delivers approximately 4×10^9 J/s within $\Delta\lambda$ and into θ_b; that's 4 Joules per nanosecond.

In other words, for a laser OSETI system using a 1 meter transmitting aperture the number of photons received from the laser transmitter (in $\Delta\lambda$) during the nanosecond pulse exceeds the number of photons received from the parent star if the laser pulse energy is greater than about 4 Joules. This statement is independent of range. It is a bit conservative, as well, if one uses a Fraunhofer dark line where the stellar flux is down by a factor of 5–10 or so.

Looking back at the parameters of the inertial fusion pulsed laser (a 3 ns pulse with 3.7 MJ per pulse at 349 nm, or 4.7 MJ at the native 1.047 μm wavelength), we see that a Helios-class laser with a 1-meter reflector would outshine the parent star by a factor of 3×10^5.

This suggests that it is not necessary to use a narrowband filter at all: if we recalculate the laser-to-starlight ratio above, assuming the photomultiplier received unfiltered (broadband) light, we still outshine the star by a factor of 30 during the pulse.

But advanced civilizations are expected to be more *advanced* than we are! So, for example, a Keck-class beam director (10 m diameter) with a Helios-class laser outshines the unfiltered star by a factor of 3000. This represents our technology in the first decade of the 21st century. With a modest extrapolation of another 2–3 orders of magnitude in delivered flux, which can hardly be considered daring since we will probably attain that sometime in the next century, we conclude that a pulsed optical beacon from a moderately advanced civilization can outshine its unfiltered star background by 6 orders of magnitude.

Optical pulsed beacons thus appear to be practicable, at least in terms of outshining their parent star.

Signal Photons per Pulse

Signal-to-noise ratio (SNR) is apparently no problem in OSETI. But what about "S"? What does it take to generate a received pulse of 100 photons?

From above we can scale N_r as $N_r \propto D_t^2 D_r^2 E_p R^{-2}$. Taking $E_p = 3.7$ MJ (Helios), we receive 1.3×10^4 photons per pulse at $R = 1$ pc with 1 meter apertures at both transmitter and receiver. For the larger ranges needed we must use larger apertures: At 30 pc (100 ly, within which there are roughly 10^3 sunlike stars) we would receive 75 photons per pulse from Helios if we used 1.5 m apertures at both ends. One could use instead an asymmetrical strategy – the same received fluence would result from a 2.3 m aperture transmitting to a 1 m receiver. At 300 pc (1000 ly, within which there are roughly 10^6 sunlike stars) we would need 5 m apertures at both ends, or comparable scaling, to deliver 90 photons per pulse.

R	N_{stars}	D_t	D_r	N_{photons}/pulse
1 pc	-	1 m	1 m	1.3×10^4
100 ly	1000	1.5 m	1.5 m	75
		2.3 m	1 m	75
1000 ly	1,000,000	5 m	5 m	90

The laser pulse widths we have assumed – of order 1 ns – are comparable to photomultiplier speeds, and so the ~100 photons are detected as one large photoevent. In practice this is desirable, because the PMT's linearity means that we can discriminate the laser pulse from stellar background by setting a pulse height threshold high enough to exclude the smaller response from the expected maximum number of stellar photons within the coincidence interval. If the laser outshines the unfiltered star by 6 orders of magnitude, this is not difficult.

At very high count rates we would be using the PMT in "current" mode, in which the individual photoevents are not preserved at the output, but in which the output waveform tracks the incident intensity, with an averaging time of order nanoseconds. At count rates significantly below 10^9/sec the PMT operates in "pulse counting" mode, in which individual photoevents are seen as output pulses; if multiple photons arrive within the ~ns response time of the PMT the result is an output pulse of proportionally larger amplitude. Good PMTs cleanly discriminate single photoelectron events from multiple photoelectron events; even a poor PMT will discriminate a single photoelectron event from a 50-photoelectron event.

Which mode is applicable to optical SETI? A solar luminosity (4×10^{33} erg/s) delivers 10^6 photons per second into a square meter aperture at 1000 light years; at that distance it has visual magnitude $m_V = 12$. So a 1 m receiving telescope with a beamsplitter but no filtration can observe a G2v star in pulse-counting mode at distances down to ∼100 ly ($m_V = 7$), at which magnitude it will have count rates of 5×10^7/sec. There will be no pileup from the Poisson-distributed arrivals, contrasted with the very large pulses produced by an intense laser of the sort described above. Significantly brighter stellar candidates bring us into the current mode. In either mode there should be no difficulty discriminating an intense laser pulse from the much weaker stellar background, providing of course that the number of photons delivered by the laser pulse is adequate (at least 25, say). Below we consider backgrounds that can interfere with this simple picture.

Note that a realistic calculation of detector count rates must take into account atmospheric and instrumental losses, as well as quantum efficiency of the detector.

Beamsize vs Planetary Zone

As we scale up the transmitting aperture it's necessary to check that the resulting diffraction limited beamsize doesn't become smaller than the planetary zone we wish to irradiate. At 30 pc the diffraction limited beam from a 1.5 m aperture transmitting at 1 μm is 4 AU; at 300 pc a 5 m transmitting aperture produces a 12 AU beam. At less than maximum range less transmitting gain is needed, such that the beam size remains the same in absolute units. These figures suggest pretty accurate shooting; perhaps the advanced civilization doing the transmitting would choose to spread the beam a bit, raising the laser power to compensate.

PREFERRED WAVELENGTH

Microwave SETI was energized by Cocconi and Morrison's 1959 suggestion to use the 21-cm hydrogen hyperfine line as a guessable beacon frequency. No analogous wavelength has been suggested for OSETI, although Townes has pointed out that extinction in the visible makes infrared a preferred band. There have been suggestions in the past to exploit the Fraunhofer stellar dark lines to reduce stellar background; moreover, a unique dark line might serve as a guessable wavelength. Possibilities are the Ca H- and K-lines, or the Fe lines, both at the blue end of the spectrum, or the yellow Na D-lines. A dark line at longer wavelength (less problem of extinction and smearing) would be better. Ideal, perhaps, would be a dark line in the near-IR that coincides with the wavelength produced by a particularly favorable lasing material. If it proves feasible to construct wavelength dispersive optical pulse detectors, the idea of a "magic wavelength" may lose some of its luster, just as it has with microwave SETI.

However, as calculated above, stellar background is not a problem for the pulse energies that are needed to produce detectable count rates. There seems to be little reason to consider narrowband filters, or multichannel spectroscopy.

EXTINCTION AND SMEARING

At optical wavelengths the interstellar medium both scatters and absorbs. That is why we cannot see our own galactic center (25 magnitudes of extinction!), or indeed much of the galactic plane at all, and that is why maps of galactic structure must be made from observations at longer wavelengths, primarily from radioastronomy.

Unlike at radio frequencies (where plasma scattering is the dominant effect), optical scattering is caused by interstellar dust grains. The effects of scattering can be severe, and limit optical SETI to distances of order a kiloparsec, at which the visual extinction is some 2 magnitudes. The effect of scattering on a laser pulse is to reduce the "prompt" pulse height, simultaneously producing delayed tails on two time scales – a close-in tail (seconds later) from forward scattering by large grains, and a much longer tail from diffuse scattering. The prompt pulse is unscattered (the term "ballistic photons" has been used in a similar context), therefore unwidened; its amplitude is reduced by a factor of roughly $\exp(-\tau_{\rm scatt})$, where $\tau_{\rm scatt}$ is the scattering optical depth. Absorption also reduces the prompt pulse, so the surviving fraction is roughly $\exp(-\tau)$, where τ is the total optical depth. At visual wavelengths scattering and absorption are comparable, and the visual extinction is about 2 magnitudes per kiloparsec. So for OSETI beacons out to 1000 ly (300 pc) the energy in the unscattered pulse is reduced by at most 40%. If optical signaling is a local activity (less than a kiloparsec), extinction and smearing do not seem to be serious problems. For greater range it is necessary to go to longer wavelengths; at $\lambda = 10\mu$m optical beacons would be effective over distances of 10 kpc or more, according to Jim Cordes at Cornell, who has considered this problem recently.

INSTRUMENTAL, TERRESTRIAL, AND ASTROPHYSICAL BACKGROUNDS

What are the backgrounds – both instrumental and natural – that interfere with the optical beacon search sketched above?

Instrumental

Preliminary experiments attempting to detect short optical pulses from a few target stars, done by Dan Werthimer at Berkeley, have demonstrated the need for a pair of PMTs in coincidence (rather than simply using a single PMT with pulse height threshold set to reject events of few photoelectrons), apparently because there are occasional large pulses in the dark current of a PMT. These may arise from radioactive decay events in the PMT glass, from ion feedback, from scintillation in the glass caused by electron impacts from within, and/or from cosmic-ray muons. A beamsplitter and coincidence circuit effectively eliminates all but muons (a muon can pass through both PMTs); if need be, the latter can be eliminated by an anti-coincidence arrangement with external muon scintillator.

Terrestrial

It would be interesting to know if any atmospheric phenomena produce nanosecond light flashes, at a level that interfere with the detection of multi-photon nanosecond pulses. The gamma ray telescope on Mt. Hopkins in Arizona detects atmospheric Cerenkov flashes in the night sky; however this instrument operates with time scales of many microseconds, not nanoseconds.

Some data on cosmic-ray-induced Cerenkov flashes in the atmosphere suggest that these are innocuous: A 10^{12} eV primary cosmic ray produces a flash of roughly 5 ns duration, with the light falling on a "footprint" on the ground of radius roughly 150 m. From within that footprint the source "image" looks like a diffuse blob in the sky, of FWHM roughly 2°, producing a fluence over the pulse of roughly 30 visible-light photons per square meter. The primary flux goes roughly as $E_{\rm pri}^{-1.7}$ for energies between 10^{12} and 10^{15} eV, and as $E_{\rm pri}^{-2.3}$ at higher energies; the photon fluence per flash goes roughly as $E_{\rm pri}$.

From these facts one quickly finds that, for 10^{12} eV primaries, the rate of Cerenkov events, seen from an arbitrary point on the ground, is given by $\text{flux} \times A_{\rm footprint} \times \Omega_{image}$, about 15 per second. However, because the Cerenkov image is diffuse on the sky, the fluence for each event, as seen by a focal-plane aperture corresponding to a 10 arcsec field of view, is only about 6×10^{-5} photons per flash; i.e., the PMT rarely detects even a single photon, and never two or more.

Lightning is of course a pulsed terrestrial source. But it is difficult to imagine pulse widths less than microseconds in flashes that scatter from distant storms (one doesn't observe in *local* storms). However, for "cultural" terrestrial events such as electrical sparks it is altogether plausible to imagine the production of sub-microsecond pulses; it requires some creativity to imagine how to direct such pulses into the telescope without broadening them.

A source of *steady* background is airglow and scattered zodiacal light. These are entirely negligible – roughly 100 photons per pixel per second in a telescope of 1 m^2 aperture. In fact, it may even be possible to do *daytime* optical SETI, under conditions of very clear sky and good seeing (which permits small a focal-plane aperture).

Astrophysical

Of course, the stellar background, which we assume is not resolvable from the laser beacon, is of astrophysical origin; we have discussed the rates above. It would be interesting to know if there are astrophysical phenomena that can produce nanosecond pulses. That requires coherence on the distance scale of tens of centimeters. One might take solace in the thought that detection of such phenomena would be a worthy discovery in its own right.

A COMPELLING *A PRIORI* STRATEGY?

As always in SETI, They can't tell us how to search before we have detected Them! In this guessing game the SETI establishment has generally worked in the centimetric microwave portion of the electromagnetic spectrum, based on favorable considerations of energy efficiency, atmospheric transparency, available technology, and the existence of a near-unique wavelength marker at HI. Of course, since SETI hasn't succeeded yet it is hard to argue that this choice of search strategy is optimum – but at least it has motivated searches, of increasing sophistication, sensitivity, and thoroughness.

An analogous argument for some unique search methodology – pulses? if so what width? wavelength? revisit time? – is needed in order to make OSETI attractive, and to get the ball rolling. We've mentioned megajoule nanosecond pulses at stellar dark lines, as an example to see if the whole idea makes sense. It apparently does; but it needs a good look to see what else makes sense, and whether there is a combination of parameters compelling enough to get some really serious searching going. In the meantime, groups at Harvard/Smithsonian and at Berkeley are going ahead with modest targetted searches for nanosecond pulsed laser beacons. These searches will target a few thousand nearby solar-type stars, dwelling on each for at least a few minutes.

Note that the treatment in this paper is not an attempt at a "parametric tradeoff" analysis; rather it is just an attempt at an OSETI reality check.

SUMMARY

The calculations above suggest that a multiplexed OSETI pulsed beacon is altogether reasonable for establishing contact out to perhaps 100 ly, in which case a Helios-class laser and a 5 m transmitting aperture, along with appropriate beam-steering optics, would suffice to contact a reasonably sophisticated amateur astronomer. For ranges out to 1000 ly amateur astronomers would have to defer to the professionals in a serious OSETI receiving effort, if nothing more than a Helios-class laser were used as a transmitter. There is no reason to assume, however, that the laser LLNL plans for the year 2030 is an upper bound on what an advanced extraterrestrial civilization might use to establish contact via targeted optical pulsed transmissions.

The SETI community needs to consider alternative OSETI strategies – choice of wavelength, pulse widths and repetition rates, revisit times, etc. – in an attempt to identify a particularly compelling *a priori* strategy, involving both sender and receiver, that could be the basis for major earth-based OSETI receiving efforts in the near term.