Unintentional radio or optical emissions ("leakage") generated by extraterrestrial (ET) technologies are very difficult to detect if they are similar to those we ourselves currently produce. High-power broadcast applications such as TV illuminate most of the sky most of the time, but have equivalent isotropic radiated powers (EIRP) that are too low by two orders of magnitude for even the most sensitive current SETI system (the SETI Institute's Project Phoenix at Arecibo [*Shostak &Tarter 1999*]) to detect at the distance of the nearest star. Artificial optical leakage is even less detectable. The very ambitious Terrestrial Planet Finder will struggle to detect the reflected starlight from a terrestrial type planet and be insensitive to any reasonable degree of ET "street lighting." Much higher EIRPs in microwaves are provided by very directional transmitters, ranging from the widely deployed NEXRAD weather radars ($3\text{x}10^{10}$ W) up to the unique Arecibo planetary radar ($2\text{x}10^{13}$ W). Such transmitters could be detected by current SETI systems, but of course the very narrow beams they produce would need to be pointed at us during our observations. Not only is this circumstance unlikely [*Billingham and Tarter*, *1992*] in the first place, but we could not expect to stay in the beam of such a device (by chance) for long enough to verify that the signal was really of ET origin with our own current search strategies. It is worth noting that *Boyce* [*1990*] found that only three stars had been illuminated serendipitously by the Arecibo planetary radar over two decades of operation. The situation is even worse for the tighter beams of lasers, such as LIDAR systems.

ET signals generated with the express intent of attracting notice over interstellar distances are much easier to detect than leakage. Isotropic beacons, as envisioned in the original *Cyclops Report*, are (as noted therein) very inefficient, since most of the radiated power will never intercept a planetary system. To be more quantitative, adopt a mean planetary system radius of 40 AU (the orbit of Pluto, 1 AU = $1.5\text{x}10^{11}$ m). Under the optimistic assumption that every main sequence star has such a planetary system, the probability of a ray from an omnidirectional beacon passing through a planetary system by chance ranges from $\leq 1\text{x}10^{-6}$ for directions perpendicular to the galactic plane to $\leq 1\text{x}10^{-3}$ for directions toward the galactic center. It is plausible that any ET civilization will value efficiency, so beacons are likely to be targeted at stars, and arguably, a large number of stars. Such a large target list can be supported by illuminating all of the targets at once or by moving the beam of the beacon sequentially through the list. Successful detection strategies must be complementary to the illumination strategy. Most present SETI systems are compatible with only the first illumination strategy. The second, sequential-illumination strategy involves two parameters, the interval that the beacon spends on each target t_{dwell}, and the revisit period $t_{revisit}$. $t_{revisit} \geq M\ t_{dwell}$ if M is the number of targets to be illuminated. If t_{dwell} is sufficiently long, the appropriate search strategy is to cycle through all the search targets in a time shorter than t_{dwell}, and repeat this process for time as long as $t_{revisit}$ (for the next-to-worst case of just one target in common between transmitter and receiver lists). If t_{dwell} is short, the search strategy should be to stay on each search target for a time longer than $t_{revisit}$. Note that if t_{dwell} were very short, say < 1 msec, and the illumination list was not too large, say < $1\text{x}10^{5}$, then $t_{revisit}$ would be < 100 seconds, so the effect is to generate a low duty cycle pulse train at each targeted star all the time, thus this extreme of the sequential illumination strategy becomes equivalent to the simultaneous illumination strategy. Greatly varying estimates of t_{dwell} and $t_{revisit}$ have been made [e.g. *Ross*, 2000; *Howard et al.*, *2000*] but we find none of them particularly convincing. An omnidirectional search system [*Dixon*, *1997*] would obviate the need to make assumptions about these parameters but is likely to be much more expensive, and less sensitive, than a single-target search system, at least with present terrestrial computing technology. Hybrid targeting strategies are possible as well. We may *speculate*, however, that since the purpose of a beacon is to make the SETI acquisition task easier, illumination schemes should produce detectable signals at each target whenever a search is conducted. Such speculation favors beacons generating multiple CW beams, or sparse pulse trains with repeat periods much less than the duration of a typical observation. For terrestrial SETI programs, observations are typically □ 1000 seconds; it is unknown what another civilization might regard as typical.

Many so-called "magic" frequencies have been suggested, with the idea that any technology in the galaxy would choose this frequency, and thus remove the need for the receiver to search in frequency. Unfortunately even Earthlings cannot agree on what that special frequency should be, so it seems even less likely that diverse species will converge. Also, if our own experience is any guide, an historically-defined, and hence species-specific, spectrum allocation process may well constrain the possible frequency choices for the beacon.

The maximum useful directivity of the beacon is another controversial issue. The higher the directivity, the less input power needed to create a given flux at the target. The individual velocities and distances to targeted stars must be accurately known in order to take advantage of this directivity to illuminate the position that the target will occupy when the signal arrives. From the terrestrial perspective the requisite parameter is known as the “proper motion” of a star. At present, data from the Hipparcos mission can provide accuracy sufficient to reduce the targeting error to about 6 AU for stars out to about 100 ly. Future space missions, such as SIM and GAIA, should provide data sufficient to reduce the targeting error to about 1 AU at 1000 ly. A much older ET civilization can presumably do much better; however, if they choose to illuminate an area smaller than the planetary system of the target, then they will need to know where the intended receiver’s SETI search systems (if any) are located. One of the first things humans did after developing the ability to place objects in Earth orbit was to start putting our astronomical observing systems into space. For optical systems the dust in the inner solar system limits the sky darkness, while for radio systems interference from radio emitters on the Earth (and on any future inhabited worlds) obscure weak signals. If we continue as we have begun, then we should expect our own major observatories to migrate as far out in our solar system as our technology allows. Thus, there would seem to be three targeting strategies available to the transmitters: A) illuminate the entire planetary system, B) illuminate only the most conspicuous planet (reasoning that recipients will have located their SETI receivers on or near the most detectable object except their star), or C) illuminate the most likely world to have originated life. Strategy C requires the transmitting civilization to be able to find planets, measure their orbits to the required accuracy, and determine whether they support life. We ourselves have plans to obtain at least some of this information out to a range of 50 ly within the next twenty years or so, thus it seems prudent to guess that older civilizations may have enough information to support strategy C) out to the limits imposed by absorption of starlight (of order 1000 ly).

Whatever the size of the area illuminated by the transmitting beacon, the beacon must be powerful enough for the receiving antenna to detect its signal above the noise. In the optical or IR the minimum detectable energy will be a few $h\nu$, ~3×10^{-19} J; in the radio, a few kT, ~1×10^{-21} J. The input energy to the beacon will be larger by the ratio of the area of the illuminated patch to the area of the receiving antenna. If the beacon is intended to be detectable by species in a stage of development similar to ours, we can estimate the receiver apertures based on our own technology to be ~10 m for the optical/IR (Keck telescope) and ~200 m (Arecibo) for the radio. The input energies required for strategy A) are then ~1×10^{5} J for the O/IR and ~1 J for the radio. A conservative ET designer would probably increase these minimal energies by a factor of 10-100 in order to allow for mismatches between detection strategies. On the other hand, in the next two decades the available apertures on Earth for both O/IR and radio should increase by about a factor of 100, suggesting that the above estimates are at least plausible minimum energies. Since the transmitting beacon energies for strategy A) are fairly modest, there would appear to be no pressing need for the beacon designers to adopt the less certain strategies B) or C). To illuminate a 40 AU patch at 1000 ly requires the beacon to have a beam of about a microradian in angular diameter. A diffraction-limited transmitting aperture of ~2 m would be required at 2 microns wavelength and an 100 km array would be required at 10 cm. Smaller apertures could, of course, be used with a corresponding increase in transmitted energy, an option that might be most attractive in the radio where the minimum energy requirements are very low and the (maximum) array size seems awkwardly large from our current perspective.

We have argued that our own level of detection technology may be an appropriate guide to the flux goals of an ET beacon designer interested in attracting the attention of emerging technologies. Terrestrial experience is unlikely to constrain the advanced technology that an ET civilization would use to implement such a beacon. It is, nonetheless, instructive to consider how we might go about implementing a beacon.

At radio wavelengths a phased array implementation now appears to be optimal. Until recently, such arrays have been too expensive for non-military applications. Dramatic reductions in the cost of microwave components, spurred by the explosive development of wireless communications, have now made these arrays more attractive. An array of N elements, each with power P_e and gain G_e, has an EIRP in the array beam of $N^2 G_e P_e$. If the cost of controlling the array phase can be kept small, as seems likely, then the cost of the array should scale as the total cost of the elements, which will in turn be roughly proportional to $N \square P_e$. Thus the cost per unit EIRP should go down as 1/N [*Oliver*, *1995*]. G_e should be low in order to allow the phased array beam to be electronically steered around as much sky as possible; a reasonable value is 3. In mass production, an element with $P_e \sim 1$ W (typical for a pager or cell phone) should cost of order \$10. Today, typical budgets for "big science" projects are of order of 1×10^9 dollars, which would set N to be of order 1×10^8 and the array EIRP ~ 3×10^{16} W. With a conversion efficiency of 50%, the array would consume 200 MW; power that could perhaps be supplied by solar cells located at each element. Electronic steering of the beam could be quite rapid, 1 msec or less to change targets. At present, we have reasonable catalogs of suitable targets with about 1×10^5 stars. This beacon could provide 1 msec pulses with an EIRP more than 1000 times that of the Arecibo planetary radar and revisit all the stars in such a catalog every 100 seconds to produce a low duty cycle pulsed signal. Alternatively the array could form 1×10^5 simultaneous CW beams, each with an EIRP of 3×10^{11} W, which is still ~1×10^5 times brighter than broadcast TV. Such a multi-beam beacon could be detected by our most sensitive current radio SETI systems from thousands of stars to a distance of 200 ly. Presumably an ET technology could do much better, but even the beacon considered here would be sufficient to reach next generation radio SETI systems used on the Square Kilometer Array [*Taylor and Braun*, *1999*] to a range well beyond 1000 ly (see the section on Future SETI).

We do not currently have the technology to build optical phased arrays. The usual suggestion for an optical transmitting beacon [*Townes*, *1997*; *Howard et al.*, *2001*; *Lampton*, *2001*; *Howard and Horowitz*, *2001a*; *Kennan*, *1999*; *Horwath*, *1996*; *and Kingsley*, *1993*] is a ~10 m reflector with a high-power laser, usually pulsed ($\Delta t \sim 1$ nsec) because that is what our own primitive optical reception technology finds convenient. To obtain diffraction-limited performance the mirror would probably need to be above the atmosphere, although conjugate mirror technology might be applicable to a ground based mirror [*Romanov et.al.*, *1999*]. We already have the technology to place such mirrors in orbit at a cost of order 1×10^9 dollars. We also have the technology to make 10-50 kJ laser pulses, and have plans for laser systems that can generate ~2 MJ pulses in the near future [*Campbell and Hogan*, *1999*]. Such laser systems also cost about 1×10^9 dollars, on the ground, but these lasers and their power supplies are far too large for us to place in orbit, at any price. These lasers also do not support high pulse repetition rates, making targeting a large list of stars difficult. Finally, the extremely high fluence of the petawatt pulses pose severe materials problems for today's optics. To provide, say, 1000 one MJ pulses per second, at a typical efficiency of 10%, would require an average input power of order 10 GW. In short, our technology is not up to this task yet, although optical technology is advancing very rapidly. Thus terrestrial experience cannot be used to provide any useful constraints on the parameters of optical ET transmitting beacons.