Q. J. R. astr. Soc. (1994), **35**, 157–175

Machine Intelligence, the Cost of Interstellar Travel and Fermi's Paradox

Louis K.Scheffer

Department of Electrical Engineering and Computer Science, University of California at Berkeley, California 94720, USA

(Received 1993 November 1; in original form 1993 April 22)

SUMMARY

If machine intelligence is possible, and the computers on which it is based resemble today's computers in some very fundamental aspects, then interstellar travel can be accomplished by data exchange as opposed to the physical movement of matter. Assuming data exchange by radio, such travel is many orders of magnitude cheaper than physical travel. This low cost provides a huge incentive for an emerging society to join an existing galactic civilization as opposed to physically colonizing the galaxy. It is likely, therefore, that there is at most one advanced civilization per galaxy. This civilization may well have unified goals and objectives, thus removing the strongest underpinning of Fermi's paradox.

Also included is a detailed analysis of the cost of interstellar communication by radio, considering both energy costs and fixed asset (antenna) costs. This analysis shows that deliberate communication is quite cheap, and eavesdropping is probably futile.

1 INTRODUCTION

Fermi's paradox is simply stated. If the earth is typical, then civilizations should already have evolved billions of times in the Galaxy, since there are many billions of stars older than the Sun. If any one of these civilizations wanted to colonize the Galaxy, they could have done so by now, even using nothing but technology that is almost now within humanity's grasp. However, there is no convincing evidence that any aliens have even visited Earth, much colonized it: hence Fermi's question, "Where are they?". This argument derives much of its force from the large numbers of civilizations involved. It is easy to imagine that a single civilization rejected colonization, or had strong ecological concerns, or turned to other pursuits, but if a billion civilizations evolved independently then surely some of them would be interested in colonization—and it only takes one to colonize the Galaxy.

Colonization, in this argument, is assumed to be accomplished by physical interstellar travel. This seems possible, although difficult and expensive. Many authors including von Hoerner (1962), Sagan (1963) and Crawford (1990) have studied the feasibility of interstellar travel. These studies have assumed that travel consists of accelerating a spacecraft, including at least one person, to a significant fraction of the speed of light. For such approaches the energy costs alone are very high. Alternative approaches, such as ships that travel much slower than light, are possible, but suffer from extremely long travel times, and are still not cheap.

Another possibility, as yet confined to science fiction, is some way to measure sufficient data about a person, transmit these data to the destination, and then reconstruct the person, including their mental state. This is called 'teleportation' and would presumably seem like travel to the person involved. Prior to 1993, there was a theoretical objection to teleportation—it was believed that measuring the original object would collapse its quantum state and prevent an exact replica from being constructed at the destination. However, Bennett *et al.* (1993) show that by pre-arranging the transfer of correlated but unmeasured particles, the exact quantum state can be reproduced at the destination even though the original's quantum state is collapsed. Nonetheless, this theoretical advance does not address one of the main practical problems of teleportation—reproducing a physical object of any appreciable size requires sending an enormous amount of information, at a correspondingly great expense.

Instead, note that today one can start a computer program on one machine, stop it, send the current state of the program as a stream of bits to a remote location, and resume the execution of the program on an entirely different machine. Consider what this would look like from the point of view of a program that is a conscious entity. It would appear as if the individual had travelled (instantly) to the far location, although neither machine has physically moved. Therefore, if conscious beings can be represented as computer programs, they can travel merely by shipping data. This approach is potentially much cheaper than physical travel for interstellar distances. This paper looks at the engineering problems behind this approach, and some of the implications.

Assuming the main costs are energy and facilities, the difference in travel costs is enormous. Assume that a being of human capabilities takes 3×10^{14} bits to represent (this estimate is justified later). Accelerating a 100 kg mass to 70% of the speed of light takes 10^{19} joules. Representing 3×10^{14} bits against the background of interstellar space (where $kT = 4 \times 10^{-23}$ joules) takes $1{\cdot}2 \times 10^{-8}$ joules. This is a potential reduction of 8×10^{26} in energy cost! Shipping data across 300 light years, with easily foreseeable technology, takes $6{\cdot}70 \times 10^{-4}$ joules bit^{-1}, so the total energy cost becomes 2×10^{11} joules. Electricity now costs about \$0·08 per kilowatt-hour ($3{\cdot}6 \times 10^{6}$ joules), so the energy cost becomes \$4500 for a 300 light year trip. The total trip cost must include the cost of the facilities on each end, which will be dominated by antenna costs at these distances. Using experience with existing large antennas as a guide, we get a total cost of about \$50000 per trip.

On the other hand, the cost of physical interstellar travel is enormous based on energy costs alone. Crawford (1990) discusses a number of possibilities for interstellar propulsion. The lowest energy technology proposed, an electromagnetically launched pellet stream, uses $1{\cdot}8 \times 10^{21}$ joules to accelerate an interstellar probe to $0{\cdot}12\,c$ for a flyby mission. At the same energy cost as used above, the energy for the acceleration phase will cost $\$4 \times 10^{13}$. If the probe has to decelerate to a stop at the end of the journey then the energy cost would be much greater, since a greater mass would need to be accelerated. Thus the cost of a trip by information transfer, including the facilities, appears to be at least a factor of 10^{9} cheaper than the energy cost alone of physical travel.

So the economics of travel by program exchange appear promising. However, a number of questions naturally arise from this definition of travel: (1) Is it possible to represent a conscious being as a computer program? (2) How many bits are needed to represent the state of a conscious being? (3) How much does it cost to send these data over interstellar distances? (4) Are future computers likely to have the features that make this possible? (5) How did the computer get to the destination in the first place? (6) If possible, what are the implications?

Section 2 of this article argues that it is likely that programs can be developed with all the capabilities of human beings (by emulating the neurones of the brain, if by no other method). Section 3 argues that, for various engineering reasons, computers of the future will share the characteristics that allow current computers to stop programs, ship their internal state to another machine, and restart them. In section 4, an estimate is made of the number of bits required to specify the internal state of an intelligent being by using the only known example, human beings. Section 5 shows that using existing and near future radio technology, one can ship this information across interstellar distances quite cheaply. Finally, sections 6 and 7 explore some of the implications of travel by information transfer.

2 CAN A PROGRAM BE DEVELOPED WITH HUMAN CAPABILITIES?

Can a program, running on a digital computer, display the complex attributes of a human such as intelligence and consciousness? This is a matter of contentious debate. Searle (1980, 1987) and Penrose (1989) say no; most biologists say yes, in principle; many computer scientists say yes, in practice (Hofstadter & Dennett 1981). Most scientists agree that if the state of every sub-atomic particle of any object could be measured (including a human being), and the evolution of that state under quantum mechanical rules could be simulated, then the behaviour of the object could, in theory, be predicted to an arbitrary degree of accuracy by a digital computer. (Of course, if quantum effects are important, then both the object's behaviour and the prediction will be probabilistic.). Quantum level simulation of humans is impractical because of the large number of particles, the difficulty of measuring the initial state, and possible imperfections in our understanding of quantum mechanics. The human body contains at least 10^{27} atoms, and even one bit per atom would be too much information for any foreseeable computer to handle. Furthermore, we have no idea of how to measure the state of a human on the atomic scale. Although impractical to implement, such a thought experiment shows that in theory a program running on a digital computer could display all human attributes.

On a more practical level, modern biology holds that the nervous system controls behaviour (Ackerman 1992, Hall 1992, Nicholls, Martin & Wallace 1992). The behaviour of the nervous system, in turn, can be understood in terms of the behaviour of the neurones, which are the cells of the nervous system, and synapses, the connections between the neurones. (See Hall 1992 or Nicholls *et al.* 1992 for a description of the neurobiology of neurones and synapses.) This link between the behaviour of the organism and the behaviour of the neurones has been shown directly for small animals whose nervous systems of a few thousand cells are known in great detail. Here researchers can accurately predict behaviour, and have tracked down many of the molecular level changes that occur in the neurones and synapses during learning (Hall 1992, chapter 14; Nicholls *et al.* 1992, chapter 13).

Although the nervous system of higher animals is much larger, the individual components, the neurones and synapses, work in an almost identical fashion (Nicholls *et al.* 1992, p. 2). Therefore, it seems likely that it is possible to simulate the behaviour of a human by simulating the operation of the nervous system on a neurone and synapse basis. (Some other systems found only in larger organisms, such as the immune system and the endocrine system, are probably required as well, but these are simple by comparison.)

Such a simulation is not possible today for various practical reasons, the foremost of which is that the human brain is not understood in sufficient detail. However, it seems possible in principle, and such a simulation should imitate the behaviour of a human, including intelligence and consciousness. The arguments against a computer program such as this achieving all human attributes are very strained. Searle (1980, 1987) argues that such a program would merely be acting conscious, but would not be so. By this solipsist standard, of course, we cannot believe anyone other than ourselves is conscious either. Penrose argues for some currently unknown quantum mechanical effect that is not reflected in the individual synapses. However, this effect has never been seen, and if seen it could presumably be simulated.

The synapse emulation approach takes a huge amount of memory and processing power, but is conceptually simple, and even most detractors of artificial intelligence will concede that this will work. Is it possible to create a conscious program that is significantly simpler than such a simulation? Most of the arguments over whether machine intelligence is possible are really asking this question. This question is not likely to be resolved soon—see Penrose (1989) and Hofstadter & Dennett (1981) for book length discussions. In this article it is assumed that in some cases shorter conscious programs are possible. Fortunately, as will be seen, even the 'brute force' method of simulating each neurone and synapse results in far cheaper interstellar travel than the physical transfer of material.

3 WILL TOMORROW'S COMPUTERS ALLOW A STOP/RESTART CAPABILITY?

Current computers can save the state of a running program and return to the execution of the program at any later time (on the same or a different machine). They are designed to do this because of the large difference in speed between the internal operations of the machine (in 1993, one instruction every few nanoseconds), and the mechanical devices (or humans) with whom the machine must interact. To avoid being idle while waiting for these slow events, existing computers work on several programs at once. When one program is waiting, the computer saves the internal state of that program and resumes work on a different program. When the first program's data arrive, the second program is stopped, and the execution of the first one is resumed using the saved state.

In addition, computers need the ability to save and restore state to respond to interrupts. Here the computer must suspend the current task to handle a time critical event, and return to the original activity once the interrupt has been handled. Another application of the save/restore ability is in a multiprocessor machine, which moves programs from one processor to another to balance the workload. Finally, the ability to save and restore state makes the construction and maintenance of computers much easier, and allows recovery from certain classes of errors. For these reasons the ability to save and restore state is likely to be maintained in computers built in the foreseeable future.

Once it is possible to stop a program and restart it where it left off, the next step is to stop a program, send the information to another computer located elsewhere, and restart the program there. This capability is used routinely in today's computer networks. If this approach is to be used in order to eliminate physical travel, however, there appears to be an obvious limitation—how did the identical machine get to the other end of the communications link in the first place? Fortunately, a fundamental theorem of computer science states that any digital computer can imitate any other digital computer perfectly, although with a potential loss of speed. This is done by writing a small program, written in the instruction set of machine 1, that emulates the instruction set of machine 2. (See Penrose 1989 for a particularly clear explanation of this.) This is not idle computer science theory—there are many commercial programs for doing exactly this. For example, you can buy a program that allows programs designed for an Intel processor to run on a Sun machine, even though these machines use completely different instructions.

This process of emulation can be extended to running on a completely unknown machine, including that built by another civilization. First you send the description of the abstract machine, then the code that runs on this machine. The far end must use a program that reads both the description of the abstract machine and the specified program. By emulating the abstract machine, it can then execute the program. (Penrose 1989 also contains a very clear explanation of this point.) In general, the description of the abstract machine is quite short (perhaps 10^5 bits), and adds a negligible amount to the length of the messages considered here. The speed difference is typically a factor of 5 or 6 for existing computers.

3.1 *Evolution of travel by information transfer*

Travel by information transfer can be regarded as an extension of several existing historical trends that substitute the exchange of information for the exchange of physical objects. Suppose you live in North America, and you have a friend in Japan. You have a story that you both wish to contribute to. Before the advent of writing, you would have had to go to Japan; after writing you could send the manuscript; and now you can send the file. Each of these methods is less costly than its predecessor because only the essential information is being sent. Sending a person's representation rather than their physical body containing that representation is a further step along the same lines.

This approach can also be seen as an extension of the 'virtual reality' work that is presently ongoing. In the near future, you will be able to don a special suit and helmet, and thus be able to direct a robot. You will see what it sees, hear what it hears and so on. It will appear to you as if you had travelled to where the robot is physically present. If there were a supply of such robots at several places, one could 'travel' to any of these places instantly. This is almost possible with current technology; however, it is limited to those distances where the transit time of the information is small compared with human reaction times. If, in the future, a person can be represented as a computer program, robots and computers can be stocked at locations together, thus removing the transit time limitation.

4 HOW MANY BITS ARE REQUIRED?

The cost of travel by information transfer depends directly on the number of bits to be sent. To estimate this, an estimate is required of how many bits are necessary to represent a being with at least the capabilities of a human. This estimate is formed by considering the only known example of a mechanism with this capability, the human brain. Two methods are used—the first considers the number of bits of storage required to simulate the physical brain, and the other creates a more abstract estimate by starting with the genetic information and adding the sum total of all experience after conception.

How much information does it take to represent the state of the brain? According to contemporary biology (Ackerman 1992, Hall 1992, Nicholls *et al.* 1992, for example), if the state of every neurone and the state of every synapse is described, there is enough information to predict the behaviour of the brain. It is difficult to get an accurate count of the number of neurones and synapses in the brain, especially since this changes during development, but Ackerman estimates that the adult human brain contains 10^{11} neurones and 10^{14} synapses. (Surprisingly, the brain of an infant contains more neurones and synapses, perhaps by a factor of 2. They are overproduced during the early stages of development, and selectively eliminated during infancy.) To allow for a margin of error, it has been assumed that 10^{12} neurones and 10^{15} synapses must be simulated.

How many bits does it take to represent the state of a neurone and its associated synapses? With apologies to biologists, here is roughly how a neurone works. The firings of other neurones release chemicals (neurotransmitters) at the synapses. This causes the cell membrane directly opposite to admit or refuse entry to ions, which changes the cell's internal potential at that point for a few milliseconds. These potentials all sum (actually a convolution because the cell has considerable resistance and capacitance). If the potential at one certain spot (the initial segment) exceeds a certain threshold, then the nerve 'fires', and generates an action potential that propagates to the output synapses, causing them to affect other neurones in turn. The contribution of each synapse to the total potential may be different, and learning is thought to occur by changing these contributions. This of course is a vast oversimplification—see Hall (1992) or Nicholls *et al.* (1992) for book length discussions of this topic.

How much of this process must be modelled to reproduce the essential features of the network of neurones? The most abstract neural net models, as discussed by Hopfield (1982) and others, use just one state variable for each neurone (often only one bit), and one weight for each synapse. A typical model that attempts to be physically accurate, such as MacGregor (1987, 1993), uses four variables to represent each synapse—one for the synaptic strength, two for the target region on the other neurone (which indirectly determines the contribution to the total potential), and one for the duration of excitation. MacGregor also requires four variables to represent the state of the neurone, plus four more variables to represent the internal state of each 'compartment', a subsection of the neurone.

Since synapses and neurones require similar amounts of state, and there are about 10^3 more synapses than neurones, the synapse data dominate the size requirements. The numbers that model the synapse require relatively few bits to represent. About 10^6 molecules of neurotransmitter are released in each synaptic event, in 200–300 'quanta' of about 5000 molecules each (Hall 1992, p. 76). Specifying the exact number of molecules released would require about 20 bits, because $\log_2(10^6)$ is about 20. More likely, one need only specify the average number of quanta released, so eight bits per synapse weight will probably suffice. In MacGregor's model each synapse also has two compartment identifiers (a few bits each), a target neurone identifier (perhaps 30 bits for a human brain), and a duration, which probably needs to be specified to about the same precision as the weight. Thus 50–60 bits are probably enough for each synapse using this model. Other known features which are not in MacGregor's model include physical location (for longer range neurotransmitters), and exactly which neurotransmitters are emitted. Considering all this information, and using reasonably efficient coding, it seems likely that each synapse can be described by roughly 100 bits. Combining this with the (high end) estimate of 10^{15} synapses implies roughly 10^{17} bits suffice to describe the state of the brain.

A less constructive argument shows that fewer bits are necessary, but unlike the synapse argument it gives no clue as to what those bits should be. For a deterministic machine, the internal state can always be recreated from the initial state and all the input thereafter. The total number of bits in the initial state plus all input may be much smaller than the number of bits in the internal state. This is a special case of Kolmogorov–Chaitlin complexity (Chaitlin 1977), which states that the complexity of an object is not the number of bits it contains, but rather the length of the shortest program that can recreate those bits. If the full object is required one can run the program to create it. The Kolmogorov complexity can be bounded in a number of ways. It is always less than or equal to the number of bits in the object, plus at worst a small constant. For a deterministic machine it is always less than the complexity at any previous time plus the complexity of all the input since that time. This is true since you can start with any previous state and apply all input since then to get the current state.

The same basic idea can be applied to analogue systems. Here the sender sends the best known initial state, and inputs to the system since the initial state. Both the sender and the receiver compute the expected evolution of the state, and as reality diverges from the prediction, the sender sends additional messages describing the differences. The receiver applies these corrections to its internal state to keep in sync with the sender. (Linear predictive coding, as used in speech synthesis, works in exactly this manner.) If the system is well suited to this coding, then the total size of the initial description, plus all inputs since then, plus all corrections, will be smaller than the size of a message that describes the current state without prior information. This type of coding works well for relatively stable systems, but poorly for chaotic analogue systems whose divergence time is short compared with the interval between measurements. In such systems, tiny errors in inputs (which are inevitable in an analogue system) can lead to arbitrarily wrong results, and a large amount of correction data must be sent. If the system is completely chaotic, each correction may be as big as a description starting from scratch, and nothing is gained by predictive coding, since nothing is gained by knowing the past history.

Will predictive coding result in savings when applied to human beings? It would be surprising if a system as complex as a human did not exhibit some chaotic behaviour. However, the amount of chaotic behaviour and the time scale of any divergence are not clear. Identical twins start with identical (genetic) states. Identical twins raised together are more alike than identical twins raised apart, and identical twins raised apart are much more alike than two people picked at random, according to Bouchard *et al.* (1990). Since even identical twins raised together have merely similar, not identical, experiences, this indicates that small differences in experience lead only to small changes in personality, implying that much of human behaviour is not chaotic, or at least has a long time scale for divergence. Thus the approach of sending the initial state, and the inputs since that time, is worth examining.

Assuming that the correction data are small compared with the new experience data, the number of bits in a human can be estimated. People begin their development with about 4×10^9 bits of genetic information. Assume that vision dominates the data content of experience, and assume high definition television (HDTV) has about the same data rate as the eyes. HDTV is currently being compressed to about 10^7 bits sec^{-1} (as of early 1993). An old person has lived about 100 years, or 3×10^9 sec, so their total experience should be less than 3×10^{16} bits. If we include redundancy that is not caught by current HDTV compression, such as scenes that are very similar to those seen the day before, the actual experience is probably much less. How much less is unclear, but an order of magnitude seems reasonable, for an experience of 3×10^{15} bits.

A lower (and probably more realistic) number is gained by noting that speed reading is an attempt to use the eyes and mind at full speed. An extremely fast speed reader might achieve 100 words sec^{-1}, and that only on material with a familiar style. This represents perhaps 500 bits sec^{-1}. [100 words sec^{-1} times 5 bits $word^{-1}$; 5 bits $word^{-1}$ seems small, but it is correct for a familiar style (Shannon 1951).] Imagine a monomaniacal 100 year old person who has done nothing but read at this rate for 16 h day^{-1} for their entire life, and who remembers all that he or she has read. The internal state of that person can then be described with at most 10^{12} bits.

As other examples, people blind from birth, assuming their experience is dominated by hearing, have a total experience of about 6×10^{14} bits. (Two channels of very high quality audio can be encoded in about 2×10^5 bits sec^{-1}.) People both deaf and blind provide a still lower limit, although how much lower is not clear.

We can bound the number of bits from below as well. However, figures for the *proven* memory of people are disappointingly small. The *Guiness Book of Records* (Matthews 1993) lists as the greatest feat of memorization a person who memorized 16000 pages of Buddhist canonical texts. Assuming there are about 600 words per page and 5 bits per word, this amounts to 5×10^7 bits.

So, representing the mental state of a human being requires at least 5×10^7 bits, and 10^{17} bits are almost surely enough. The energy required to store this state is rather small by conventional standards. Assume that the correct figure for the state of a human is 10^{15} bits. Each bit of information requires at least energy kT to store. At 3 degrees K, this is 4×10^{-23} joules per bit, or 4×10^{-8} joules overall. If you believe that the only real difference between any two human beings is in their behaviour, then you can say that, on a very fundamental level, any two people differ by less than an erg!

5 THE COST OF SENDING INFORMATION ACROSS INTERSTELLAR DISTANCES

Finally, here we are on solid ground. A number of studies such as Oliver (1977), Wolfe (1977) and Drake *et al.* (1973) have shown that the cheapest known way to ship this information, in terms of the energy required, is in the microwave region. Assuming both ends of the link cooperate, then Doppler shift is not a problem, and can be ignored. This differs from most previous analyses, where the receiver was eavesdropping and the Doppler shift was unknown.

These are the constants and variables used in the following cost calculation: (1) k, Boltzmann's constant; (2) c, speed of light; (3) T, temperature in degrees Kelvin; (4) R_t, radius of transmitter antenna; (5) R_r, radius of receiver antenna; (6) λ, wavelength of transmission; (7) D, distance from source to destination; (8) C_e, cost of energy (\$ Joule^{-1}); and (9) C_a, antenna cost (see text for units).

From Shannon (1949), reliably sending one bit requires energy kT (in the case where the bits per second sent equals the bandwidth in Hertz, which we assume here). Assuming we transmit energy E, we can find the received power as follows, assuming a parabolic antenna on both ends. We calculate the energy density at distance D (the first term) and multiply by the gain of the transmitting antenna (the second term). This gives us the energy density at distance D. We multiply this by the area of the receiving antenna (the third term) to get the total energy received. This must be at least kT. (Photon quantization noise also dictates that we need roughly 1 photon per bit as well. For now, this noise and other engineering losses will be ignored, but they will be considered later.)

$$\left(\frac{E}{4\pi D^2}\right)\left(\frac{4R_t^2\pi^2}{\lambda^2}\right)\pi R_r^2 = kT. \qquad (1)$$

Solving for E, the energy per bit, we find

$$E = k\left(\frac{T}{\pi^2}\right)\left(\frac{D^2}{R_t^2}\right)\left(\frac{\lambda^2}{R_r^2}\right). \qquad (2)$$

Next, assume that both ends have antennas of radius R_a, and assume we need to ship Nb bits. Allow a factor of Ke for engineering, and let C_e be the cost of energy. Assume that the fixed asset costs are dominated by the antennas, and the antenna cost scales like the cube of the diameter. These observations are based on experience with the Deep Space Network of JPL, as described in Potter, Merrick & Ludwig (1966). Let the antenna cost be C_a dollars for a 1-m antenna. We transmit at channel bandwidth Bw, which we will also take to be the number of bits sec^{-1} transmitted. (These can actually be chosen independently, but a lower number of bits sec^{-1} per unit bandwidth gains little, and a higher number costs exponentially more in power.) The antenna is assumed to have a lifetime L, so the cost per trip is equal to the total cost of the antenna divided by the total number of trips made during its lifetime (this assumes the antenna is kept busy). The energy cost per trip, and the antenna cost per trip are computed:

$$C_{\text{energy}} = k\left(\frac{T}{\pi^2}\right)\left(\frac{D^2}{R_a^4}\right)\lambda^2 \, Nb \, Ke \, C_e \tag{3}$$

$$C_{\text{antenna}} = C_a R_a^3 \left(\frac{1}{L}\right)\left(\frac{Nb}{Bw}\right). \tag{4}$$

To find the cheapest overall cost, we sum the antenna and energy costs, differentiate with respect to the radius of the antenna, set the result equal to 0, and solve. The best radius is:

$$R_{\text{opt}} = \left(\frac{k\,T\,Ke\,C_e\,L\,Bw}{3C_a}\right)^{\frac{1}{7}}\left(\frac{2D\lambda}{\pi}\right)^{\frac{2}{7}} \tag{5}$$

and the minimum cost is:

$$C_{\text{opt}} = \left(\tfrac{7}{12}\right) Nb(4\,k\,T\,Ke\,C_e)^{\frac{3}{7}}\left(\frac{3C_a}{L\,Bw}\right)^{\frac{4}{7}}\left(\frac{D\lambda}{\pi}\right)^{\frac{6}{7}}. \tag{6}$$

Note that the dependence on all terms is sub-linear, except for the number of bits. Therefore the total cost is not very sensitive to errors in the estimated costs of power, antennas, and so on. The diameter of the antenna is extremely insensitive, varying only a factor of 10 for factor of 10^7 of most assumptions.

What are reasonable values for the remaining parameters? All other things being equal, we want λ and T to be as small as possible. T is limited by the cosmic background noise to about 3 degrees K. The lower limit for λ is determined by two factors not in the basic equations—photon quantization noise and the stability of the physical structures required to send and receive the beam. Of these, photon quantization noise appears to be the stricter limit, and imposes a minimum wavelength of about 5 mm (60 GHz), where each photon has an energy of roughly kT. The physical stability required to operate at this wavelength seems achievable. An antenna tolerance of about 0·3 mm is required (see Ruze 1966 for antenna tolerance theory), but one can achieve roughly this level of accuracy on Earth, in a 1 g field, with wind and temperature gradients, in antennas of 100 m. [The radio telescope at Effelsberg is the best example. See Love (1976) and Hechenburg, Grahl & Wielebinski (1973) for more information.] Costs are reduced by choosing the widest possible bandwidth, and 20 per cent bandwidths are achievable in practice, so a 10 GHz bandwidth seems reasonable. Currently existing technology yields a *Ke* (engineering margin) of about 10. The breakdown is 80 per cent for each aperture efficiency, 70 per cent for converting input energy to microwaves, 6 degree K noise temperature, and 50 per cent coding overhead. These efficiencies are determined by engineering, not physics, and will probably be improved.

The antenna cost C_a is harder to estimate, since large antennas have been built over a wide time span, all have different specifications, and technology is changing rapidly. In Potter *et al.* (1966), the cost of the NASA deep space network (DSN) antennas scaled like diameter raised to the 2·78 power, with an implicit assumption of constant surface accuracy. (An exponent of 3 was used in the analytic expressions for mathematical simplicity, and because the actual value is not clear, as discussed below.) Using this rule of thumb, and recasting historical cost figures into 1993 dollars by using the USA consumer price index, the cost of a 1000-m antenna suitable for interstellar communication can be estimated.

The largest DSN antennas cost $\$50 \times 10^6$ ($\$12 \times 10^6$ in 1966) for a 64-m antenna. Potter *et al.*'s DSN scaling would predict $\$100 \times 10^9$ for a 1000-m antenna. More recently, in 1992, the replacement 100-m antenna at Greenbank cost $\$55 \times 10^6$ (Tietelbaum 1992), implying a cost of $\$33 \times 10^9$ for a 1000-m antenna. However, two of the major design constraints for this size of ground based antenna are gravity and wind (Hechenburg *et al.* 1973), factors that are absent or much smaller in space. The observatory at Arecibo, where the 305-m reflector is supported by the ground, was upgraded to an accuracy comparable with the DSN antennas in 1974, at a cost of $\$25 \times 10^6$ [$\$8·5 \times 10^6$ in 1974, as reported in LaLonde (1974)]. Using DSN scaling, this would predict a cost of $\$700 \times 10^6$ for a 1000-m antenna. Other alternatives are an array of antennas, rather than a single antenna, or a large space based antenna, or antennas on the Moon. These alternatives were studied by Basler (1977), in the context of a search for extraterrestrial intelligence, and the conclusion was that space and ground based antennas were of similar cost for similar capabilities, if kilometre scale apertures are required. The estimated cost was roughly $\$30 \times 10^9$ ($\$11·4 \times 10^9$ in 1975) for an array with a collecting area of 7 km^2. Using DSN scaling, this would imply a cost of $\$1·5 \times 10^9$ for a 1000-m antenna; however, it is not clear this scaling is appropriate since the cost of such an array should scale as the effective radius squared (or less, once economies of scale are included). Still another alternative, at least for point to point communication, involves positioning the antenna where it can utilize gravitational focusing from the Sun (any distance at least 550 AU from the Sun, on the line that passes from the target star through the Sun, will work). This is a rather inconvenient location; however, the effective aperture of an antenna located there is greatly increased. This obviously trades transportation cost for antenna cost, but a rough calculation (see Appendix A) implies that by using such focusing a factor of roughly 10^6 can be gained in collecting area, and hence a 1-m antenna at the gravitational focus is the equivalent of a conventional 1000-m antenna. The cost of the antenna itself is negligible in this case, but the cost of getting it into position might well be on the order of $\$10^9$–$10^{10}$.

So, given only moderate advances over what is achievable now, perhaps a

TABLE I

Low, medium, and high cost cases

Var:	*D*	*Ce*	*Nb*	*A*	*Ke*	Size	Cost trip^{-1}	E bit^{-1}
Units:	ly	\$ kwh^{-1}	bits	\$ km^{-1}		m	\$	joule
Low	10	\$0·01	10^{12}	10^9	5	700	2	6×10^{-5}
Med	300	\$0·08	3×10^{14}	10^9	5	2100	52500	7×10^{-4}
High	10^4	\$1·00	10^{17}	10^{10}	10	10800	7×10^8	1×10^{-3}

reasonable cost for a space faring civilization is $\$10^9$ for the equivalent of a 1000-m antenna. The useful life *L* of such an antenna is also uncertain. Without being specifically designed for a very long life, our existing spaceborne antennas, such as those in the Pioneer and Voyager probes, have lasted 20 years, and are expected to last at least 25 years more, according to Kerr (1993). Therefore, a 100-year lifetime seems like a safe minimum, with much longer lifetimes perhaps possible.

What are costs with all these assumptions? Because there is a wide spread of possible values for the parameters, we will consider a typical case, a high cost case, and a low cost case as shown in Table 1. For typical assumptions a 300 light year trip is used, $\$10^9$ for a 1000-m antenna, 3×10^{14} bits to be sent, and \$0·08 kwh^{-1} energy cost. Since this is at least a 300-year project, enough research is assumed to reduce the engineering margin *Ke* to 5, and design the antenna for a 300-year life. In this case, it can be concluded that the antenna should be about 2100 m in diameter. The total cost is \$52500 per trip, which is many orders of magnitude cheaper than any cost estimate of physical interstellar travel.

Instead, try some pessimistic assumptions. Assuming a trip of 10000 light years, a high antenna cost of $\$10^{10}$ for a 1000-m antenna, a high estimate of the number of bits to be sent (10^{17}), high energy costs (\$1 per kwh), and no improvements in the state of the microwave art. We assume a depreciation time equal to the one way trip time, however. The optimum antenna is then 11000 m in diameter and a trip costs $\$7 \times 10^8$. This is quite expensive, but at the same energy costs the cost of accelerating a 100 kg person to 70 per cent of the speed of light is $\$1{\cdot}2 \times 10^{12}$ at 100 per cent efficiency. Factoring in the weight of life support, shielding, and engineering inefficiencies will make the physical travel costs much higher.

On the other hand, optimistic assumptions yield stunningly low costs. If we assume 10 light year trips, 10^{12} bits to be sent, $\$10^9$ for a 1000-m antenna, 1000-year life of the antenna, \$0·01 per kwh, and a factor of 5 for engineering, we get an antenna size of 700 m and a cost per trip of less than \$2. If it is assumed that the antennas are using gravitational lensing as well, so that the antenna area for a given cost is multiplied by 10^6, and 10000-year depreciation, then the cost per trip is $\$6 \times 10^{-5}$! Looking at the details of the solution in this case shows the power of radio technology when combined with gravitational lensing. Each 10^{12} bit message is sent in 100 sec using a 100 watt transmitter (10^4 joules total). This signal is transmitted through a 4·5-m antenna, which looks like a 4·5-km diameter antenna thanks to gravitational lensing. The energy advantage of information transfer over physical travel is on the order of 10^{17} in this case.

5.1 *How accurate is this cost model?*

How do limitations in the analytic cost model affect the above analysis? The main point of the analysis is that interstellar communication is relatively cheap. Because the costs above are for the radio region, then any argument claiming that another portion of the spectrum is cheaper (Townes 1983) strengthens the overall thesis. Errors are only relevant if the above estimates are low, meaning that interstellar communication cannot be achieved (by any means) for the costs quoted above.

How could the computed costs be low? Aside from Nb (already discussed in detail) the main errors appear to be the cube law scaling of antenna cost with size, the lack of dependence of antenna cost on wavelength, the omission of labour costs, and the assumption that the link is busy all the time. The antenna cost versus size is probably the biggest approximation. This is an empirical relationship, and there is a wide range of plausible exponents. For example, DSN reported an exponent of 2·78 (derived from 2 data points!), but an array of antennas might lead to an exponent of 2. However, the total costs are insensitive to the exact scaling; repeating the above analyses with exponents ranging from 1·5 to 3·5 gives costs within a factor of 2 of those shown. This is because the antenna size is insensitive to almost all parameters (equation 5), so under a very wide range of scaling, roughly 1 km antennas are optimum. Since antennas of this size form the basis of the cost estimates, errors resulting from a wrong scaling law exponent will be small. The exact value of the multiplicative constant A is less clear since we have no experience with antennas of this size. Fortunately the total cost only varies as the 4/7 power of this constant, so even an error by a factor of 100 would only lead to a cost difference of a factor of 14. This is not nearly enough to affect the main conclusions of this article.

Considering the wavelength/cost trade off might lead us to pick a different optimum wavelength (Townes 1983), but as discussed above this only strengthens the main argument. Keeping the antenna busy should be possible—the capacity is 10^{10} bits sec^{-1}, which does not seem excessive for a link between two civilizations. Finally, for large antennas and short (less than 100 years) lifetimes, labour costs are a small correction. However, if long lifetime antennas require labour to keep them working, this could become an appreciable cost, and would need to be considered.

6 APPLICATION TO THE FERMI PARADOX

The Fermi paradox is simply stated: if humans are typical, then life begins soon after a planet is formed, and achieves spaceflight soon after evolving intelligence. Since the Sun was not formed until late in the life of the Galaxy, then (if we are typical) many other civilizations would have achieved spaceflight billions of years ago. If even one was interested in colonizing the Galaxy, the Galaxy (including Earth) would have been colonized long ago, since any reasonable spaceflight speed allows the Galaxy to be colonized in only a few million years. However, the Earth is not colonized by aliens now, nor is there any evidence it ever has been. So either we are not typical (an assumption that has almost uniformly proved wrong in the past) or some factor makes galactic colonization extremely unlikely, because out of millions of civilizations and billions of years, no one has yet colonized the Earth. See Brin (1983) for a much more detailed discussion of this problem.

How does travel by information transfer affect this argument? The first observation is that the same technology (machines with human class capabilities) that enables easy travel by information transfer will also offer many other possibilities that may radically change the nature of society. These include, but are not limited to, personal immortality, easy duplication of persons, unlimited hibernation, backups in case of accidents, mechanical telepathy, increased intelligence, expansion of acceptable environments (high vacuum and lower temperatures, for example), and being in multiple places at the same time. On a higher order, there are biologists who argue that the major advances in evolution occur when two previously separated forms are merged, as when eukaryotic cells picked up mitochondria [see, for example, the theories of Margulis, as described in Mann (1991)]. With machine based beings this merging might be possible between two beings with wildly different evolutionary backgrounds, because it is easy for any digital computer to imitate any other. This could result in a drastic speeding up of evolution by combining the most favourable forms of all existing species in the Galaxy. These arguments simply show that the social structures, ethics, and all other aspects of such a culture could be completely different from ours.

However, for the purpose of argument, assume that galactic races have the same basic motivations as humans, and are motivated to explore and travel. Travel by information transfer is better in many ways than physical travel. It happens at the speed of light with no loss of time to the traveller. It is a factor of perhaps 10^8 to 10^{17} cheaper than spaceflight. It is probably safer as well. The only drawback is that there must be a receiver at the far end, and this is only a serious drawback for the first society to evolve in the Galaxy. Even for the first society, the ease of travel makes the job easier. One strategy would be to send out the lightest possible receiver-robot combination to nearby systems. If a probe arrives successfully, scientists and engineers can visit, perform their investigations, and start a colony if conditions are favourable and the home world decides to do so. This is similar to the self-replicating probe strategy of Tipler (1980) except that the probes are not autonomous—instead new planetary systems become outposts of the existing civilization as soon as a probe arrives. Self-replicating probes would not be any faster or cheaper, are more difficult to build since all the required knowledge must be built in, and are harder to control.

A factor of roughly 10^9 in cost is hard to overlook. Here on Earth, we are contemplating a manned mission to Mars costing roughly \$$10^{11}$. If anyone who wanted to could travel there for \$100, would we ever spend the money to develop the ship? There would be no scientific motive, no economic motive, and no political motive.

The first civilization to evolve spaceflight may well have colonized the Galaxy. We may see no evidence of this here on Earth because it may have happened before the Sun was born, or because there is a probe in the solar system but we have not found it yet (a small probe at 550 AU, for example, would be extremely hard to detect), or perhaps because the galactic civilization has chosen not physically to visit every stellar system. (Once a civilization has seen a few million planetary systems the patterns may become clear, and a telescopic investigation is enough. This argument seems to have been overlooked.) The second civilization in the Galaxy, after it discovers the first, has a choice (presuming the first civilization is willing to cooperate). It can develop slow, expensive, physical travel, or, for a tiny fraction of the expense, it can use travel by information transfer to go anywhere in the Galaxy at the speed of light. Ships are not needed for most scientific work; the scientists can already travel to all of the interesting places discovered by the first civilization. There is no engineering challenge; the first civilization has already done the work and the methods are in the library. There is no need to spread to another planet to protect against catastrophe on the home planet, since the inhabitants of the second civilization would spread over the whole Galaxy at just under the speed of light. There is unlikely to be a military motive, because the first civilization is both more advanced and more numerous. The second civilization might develop interstellar spaceflight out of planetary pride, but if human civilization is typical the factor of 10^9 cost differential will suppress this very effectively.

One of the strongest arguments of the original Fermi paradox was the 'large number' argument. This holds that if even one civilization (out of presumably a very large number) was interested in expanding into every possible niche then the Earth would have been colonized by now. Thus, instead of explaining why one culture did not colonize the Earth, any explanation must show why every single culture (out of billions), over billions of years, was not interested in colonizing the Galaxy.

In this new scenario, all civilizations after the first have a huge barrier physically to colonizing new worlds. The first civilization, by the time the Sun is born, might well restrain itself. For all civilizations after the first, it is roughly 10^9 times easier to join the existing civilization than to engage in physical travel. Thus easy galaxy-wide travel might well lead to only one, widely spread, culture in the Galaxy. If this culture has decided against colonizing Earth (for whatever reason), then the 'large number' argument loses its force.

7 CONCLUSIONS AND OBSERVABLE CONSEQUENCES

There are two main conclusions—travel by information transfer is probably possible, and much cheaper than physical travel over interstellar distances. From this it follows that if there is intelligent life in the Galaxy, it probably comprises one large civilization.

What might one observe if this hypothesis is true? Even if it is assumed that the communication is by radio (as opposed to some more advanced method which we do not know of yet), we would not expect to see much. Simple engineering considerations dictate that even if the Galaxy is saturated with radio traffic, minimization of costs by the communicating civilizations would result in signals which are extremely difficult to detect from Earth. No conspiracy is required—presumably each civilization's engineers try to minimize costs, and this will result in communications which are hard to detect. Minimum cost radio communications are very well directed (each beam covers roughly 10^{-10} steradians), and even to the intended receiver are just above the background noise when observed by a roughly 1-km diameter antenna. This will be extremely hard to eavesdrop on. Furthermore, as total costs decrease as λ decreases, cooperative links would be expected to run with the highest frequencies possible (at least until photon quantization starts increasing the energy requirements). A 50–60 GHz range seems likely to minimize total costs for a spacefaring civilization, and this frequency range is difficult to explore from the ground.

Other evidence of a galactic civilization as proposed here would be subtle at best. If gravitational focusing is not used, one would expect to find space based structures roughly a kilometre in size in each occupied stellar system. At least one ship would also be expected at every 'interesting' place in the Galaxy which is unlikely to develop life (such as around black holes and relatively recent novae), but these ships may well be small because, as a remote outpost, visitors would accept a higher cost to go there. We would expect at most one ship in each stellar system; and if one is present in the solar system it is probably small. In addition, it might be quite far from the Sun (by our standards) if it is using gravitational focusing to make communication easier. One would not expect to see ships on the move; most of the interesting places have had a receiver long ago.

The worst case from a detection point of view might be low temperature machine civilizations in the Oort clouds, using gravitational focusing to communicate by radio. Currently such a civilization could not be detected in our own solar system, much less anywhere else in the Galaxy. This site might well be attractive to a machine civilization—low temperatures and low gravitational forces should lead to reliable machines, easy access to superconductivity, and easy access to raw materials.

For a final conclusion, Fermi's paradox can be used in reverse. If humans are typical, then intelligent life has evolved many times in the Galaxy. As we have not been colonized, however, there must be some explanation of why each civilization chose not to colonize the Earth. One possible explanation for this is a unified galactic society. This argument explains a strong economic motivation for each civilization to join such a society, and explains why no signs can be detected of the trade and travel that surely goes on if such a society exists.

APPENDIX A

Gravitational focusing by a star for communication purposes

From Einstein, we know that light is deflected by gravity. A ray passing at distance b from a mass M is deflected through an angle α, where

$$\alpha = \frac{4GM}{bc^2} \tag{7}$$

G is the gravitational constant. Using the mass and radius of the Sun, we find that initially parallel grazing rays come to a focus at about 550 AU from the Sun. Consider a 1-m antenna at this focus point. If the Sun was perfectly symmetrical, then this antenna would receive all the radiation that goes through a ring 0·5 m wide extending all the way around the Sun. Since the Sun is about 4×10^9 m in circumference, this is a potential gain of 2×10^9 in collecting area. The same focus is obtained at any further distance as well; rays with larger impact parameters are focussed at greater distances from the Sun.

However, the Sun is not perfectly symmetrical. Assuming the main non-symmetry is oblateness due to rotation, the lowest gain will occur when we are in the plane of the equator, and at the closest possible focus (grazing rays). From Blandford *et al.* (1989), geometrical optics and the Einstein formula can be used with Newtonian gradients in the case where the deflection is small. We start with the Newtonian potential of an oblate sun in spherical coordinates from Misner *et al.* (1973). m_{Sun} is the mass of the sun; r_{Sun} is the radius of the sun, and $j2$ is a dimensionless constant describing the oblateness:

$$U(r,\theta) = \frac{Gm_{\text{Sun}}}{r}\left[1 - j2\,\frac{r_{\text{Sun}}^2}{r^2}\,\frac{(3\cos(\theta)^2 - 1)}{2}\right]. \tag{8}$$

We convert this to cylindrical coordinates (d, ϕ, z), compute the gradient in the radial direction, and find a radial deflection $\mathit{defl}_{\text{rad}}$ of:

$$\mathit{defl}_{\text{rad}} = \left(\frac{4Gm_{\text{Sun}}}{c^2}\right)\left(\frac{d^2 - 2j2\,r_{\text{Sun}}^2\cos(\phi)^2}{d^3}\right) \tag{9}$$

and a deflection $\mathit{defl}_{\text{perp}}$ perpendicular to this of

$$\mathit{defl}_{\text{perp}} = \left(\frac{8Gm_{\text{Sun}}}{c^2}\right)\left(\frac{j2\,r_{\text{Sun}}^2\cos(\phi)\sin(\phi)}{d^3}\right). \tag{10}$$

From this, a straightforward simple geometrical argument shows that the only rays that will strike a centred antenna come from four patches near the poles and the equator, where either the cosine or the sine terms are near zero. What is the area of these patches? Using small angle approximations to the trigonometric functions, and assuming a centred circular antenna of radius r_a we find that the patches are ellipses, bent to follow the Sun's curvature. Each ellipse has a minor radius of

$$r_{\text{minor}} = \frac{r_a}{2} \tag{11}$$

and a major radius of

$$r_{\text{major}} = \frac{r_a}{2j2}. \tag{12}$$

So the combined area A of all four elliptical regions is

$$A = \left(\frac{r_a^2}{j2}\right)\pi. \tag{13}$$

So the collecting area has been increased by a factor *gain*, where

$$\mathit{gain} = \frac{A}{\pi r_a^2} = \frac{1}{j2}. \tag{14}$$

What is the value of $j2$ for a typical star? Once again, we assume the Sun is typical. Because the value of $j2$ for the Sun is important for one of the main tests of general relativity (the perihelion shift of Mercury), it has received considerable attention (see section 7.3 of Will 1993 for example). Despite the attention, however, little agreement has emerged. If the Sun rotated uniformly, then $j2$ would be about 10^{-7}. However, as sunspots at different latitudes rotate at different rates, the Sun is known not to rotate as a solid body. If the interior rotates faster than the surface, the value of $j2$ might be considerably greater. Different measurements (all difficult and indirect) have given very different answers, some as high as 3×10^{-5}. The best current estimate of $j2$ is about $1{\cdot}7 \times 10^{-7}$ (Brown *et al.* 1989). Furthermore, if general relativity is correct, then $j2$ must be less than about 10^{-6} to agree with the experimental evidence. So we can safely say that the gain from solar focusing is at least 3×10^4, probably about 6×10^6, and almost surely less than 10^7 for an antenna in the plane of the Sun's equator. Still higher gains can be achieved at higher solar latitudes, or at focal points further from the Sun.

Note: one additional use for this focusing might be data return from interstellar flybys. In this case, after the flyby of another stellar system, the probe manoeuvres onto the focus line from Earth. This greatly increases the gain of the spacecraft antenna, and would enable it to return data to Earth much faster.

REFERENCES

Ackerman, S., 1992. *Discovering the Brain.* National Academy Press, Washington, DC.

Basler, R., 1977. A preliminary parametric analysis of search systems. In: Morrison *et al.* 1977, pp. 181–183.

Bennett, C.H., Brassard, G., Crepeau, C., Jozsa, R., Peres, A. & Wooters, W.K., 1993. Teleporting an unknown quantum state via dual classical and Einstein–Podolsky–Rosen channels. *Phys. Rev. Lett.*, **70**, 1895.

Bouchard, T., Lykken, D., McGue, M., Segal, N. & Tellegen, A., 1990. Sources of human psychological differences: the Minnesota study of twins reared apart. *Science*, **250**, 223.

Blandford, R.D., Kochanek, C.S., Kovner, I. & Narayan, R., 1989. Gravitational lens optics. *Science*, **243**, 824.

Brin, G., 1983. The 'great silence': the controversy concerning extraterrestrial intelligent life. *Q. J. R. astr. Soc.*, **24**, 283.

Brown, T.M., Christensen-Dalsgaard, J., Dziembowski, W.A., Goode, P., Gough, D.O. & Morrow, C.A., 1989. Inferring the Sun's internal angular velocity from observed p-mode frequency splittings. *Astrophys. J.*, **343**, 526.

Chaitlin, G.J., 1977. Algorithmic information theory. *IBM J. Res. Dev.*, July **21**, 350.

Crawford, I., 1990. Interstellar travel: a review for astronomers. *Q. J. R. astr. Soc.*, **31**, 377.

Drake, F., Kardashev, N., Troitsky, V., Gindilis, L., Petrovich, N., Pariisky, Y., Moroz, V., Oliver, B. & Townes, C., 1973. Techniques of contact. In: *Communication with Extraterrestrial Intelligence* (*CETI*), p. 230, ed. Sagan, C. MIT Press, Cambridge, Massachusetts.

Hall, Z., 1992. *An Introduction to Molecular Neurobiology*. Sinauer Associates, Sunderland, Massachusetts.

Hechenburg, O., Grahl, B. & Wielebinski, R., 1973. The 100 meter radio telescope at Effelsberg. *Proc. IEEE*, **61**, 1288.

von Hoerner, S., 1962. The general limits of space travel. *Science*, **137**, 18.

Hofstadter, D. & Dennett, D., 1981. *The Mind's I.* Bantam Books, New York.

Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. *Proc. Nat. Acad. Sci. U.S.A.*, **79**, 2554.

Kerr, R.A., 1993. Getting in touch with the edge of the solar system. *Science*, **260**, 1591.

LaLonde, L.M., 1974. The upgraded Arecibo observatory. *Science*, **186**, 213.

Love, A., 1976. Some highlights in reflector antenna development. *Radio Sci.*, **11**, 671.

MacGregor, R.J., 1987. *Neural and Brain Modeling*. Academic Press, San Diego.

MacGregor, R.J., 1993. *Theoretical Mechanics of Biological Neural Networks.* Academic Press, San Diego.

Mann, C., 1991. Lynn Margulis: science's unruly earth mother. *Science*, **252**, 378.

Matthews, P. (ed.), 1993. *The Guinness Book of Records 1993*. Guinness Publishing, Reading, U.K.

Misner, C.W., Thorne, K.S. & Wheeler, J.A., 1973. *Gravitation.* Freeman, San Francisco.

Morrison, P., Billingham, J. & Wolfe, J. (ed.), 1977. *The Search for Extraterrestrial Intelligence* (NASA SP-419). US Govt. Printing Office, Washington, DC.

Nicholls, J., Martin, A. & Wallace, B., 1992. *From Neuron to Brain* (3rd ed.). Sinauer Associates, Sunderland, Massachusetts.

Oliver, B., 1977. The rationale for a preferred frequency band. In: Morrison *et al.*, 1977, pp. 65–74.

Penrose, R., 1989. *The Emperor's New Mind*, Oxford University Press.

Potter, P., Merrick, W. & Ludwig, A., 1966. Big antenna systems for deep-space communications. *Astronautics and Aeronautics*, Oct. **4**, 84.

Ruze, J., 1966. Antenna tolerance theory—a review. *Proc. IEEE*, **54**, 633.

Sagan, C., 1963. Direct contact among galactic civilizations by relativistic interstellar spaceflight. *Plan. Space Sci.*, **11**, 485.

Searle, J., 1980. Minds, brains, and programs. In: eds Hofstadter & Dennett 1981, pp. 353–373.

Searle, J., 1987. Minds and brains without programs. In: *Mindwaves*, pp. 208–233, eds Blackmore, C. & Greenfield, S. Basil Blackwell, Oxford.

Shannon, C., 1949. Communication in the presence of noise. *Proc. IRE*, **37**, 10.
Shannon, C., 1951. Prediction and entropy of printed English. *Bell Syst. Tech. J.*, **30**, 50.
Tietelbaum, 1992. Radiation systems. *Fortune*, **126**, 99.
Tipler, F., 1980. Extraterrestrial intelligent beings do not exist. *Q. J. R. astr. Soc.*, **21**, 267.
Townes, C., 1983. At what wavelengths should we search for signals from extraterrestrial intelligence?. *Proc. Nat. Acad. Sci. USA*, **80**, 1147.
Will, C. M., 1993, *Theory and Experiment in Gravitational Physics* (2nd ed.). Cambridge University Press.
Wolfe, J., 1977. Alternate methods of communication. In: Morrison *et al.*, 1977, pp. 103–110.