The Physics of Information Processing Superobjects: Daily Life Among the Jupiter Brains

Anders Sandberg

December 22, 1999

Abstract

Physics limits information processing, and hence the possible forms of intelligent beings and their civilizations. In this review I discuss physical limitations on density, speed, size, energy dissipation and communication, sketching the constraints on very powerful information processing objects.

1   Introduction

The laws of physics impose constraints on the activities of intelligent beings regardless of their motivations, culture or technology. As intelligent life begins to extend its potential, information storage, processing and management will become extremely important. It has been argued that civilizations generally are information-limited [64] and that everything intelligent beings do, not just thinking but also economy, art and emotion, can be viewed as information processing [4, p. 660] . This means that the physics of information processing imposes limits on what can be achieved by any civilization. In the following I will look at the problems of very large computing systems. They represent the extremes of what an individual being or a culture can become.

A growing civilization will need more information processing. If there is internal competition of any kind for a resource, then there will be a strong incentive for expansion to increase the supply, be it a planned move, market competition or just basic tropism. Since all kinds of growth and existence consume resources (be it metal, silicon, energy, memory space or attention) and tend to continue to expand until all of a niche is filled, a positive feedback loop emerges as parts of a civilization compete for limited resources, expand their reach, grow until they fill the available niche and spawn new growth by the new demands created by their expansion (including efforts to find new resources or niches induced by scarcity). This leads to an exponential growth in information processing, until it reaches the limits set by physical law and technology or enters into another regime [73].

More specifically, it appears advantageous in terms of flexibility and efficiency for individual beings to exist as software rather than (biological) hardware [59, 58,

60]: less resources are needed to sustain the being, autoevolution[1] becomes much more realisable and the limits to its existence are determined by the computing system it exists in rather than a constant body; as technology advances the being will be extended too. It is quite probable that computing systems much faster, denser and larger than any genetically evolved being could be built. An information entity (infomorph[2]) can evolve and act on timescales limited only by the computing system it is implemented in, while a physical being has a characteristic timescale limited by its size and composition (which has been determined by available materials for biological evolution and the being's evolutionary past rather than its own goals).

Thus we see that there exists many reasons for the emergence of very powerful information processing systems. But what are the limits to them according to known physics?

This paper will review physical limits of large scale information processing. Section 2 is an overview of the limitations that will be studied. Section 3 deals with bounds on the density of memory storage and to some extent processing elements depending on available construction materials and fundamental physical limits. Section 4 considers the speed of computation, how fast elements of the system can change state, mainly based on available energy densities. Section 5 considers very large systems and ways of sustaining them against the force of gravity. Section 6 deals with the demand for entropy dissipation in computing systems, which place a lower limit on the need for energy to run the computations and heat dissipation necessary. Section 7 discuss limits to communication such as delays, bandwidth and noise. Section 8 considers with exotic physics such as quantum computing, black hole computing, wormholes and physical eschatology, which remain theoretical or speculative but if possible would enable very powerful computing systems. In an appendix three large computing systems are sketched as examples.

2 The Main Factors of Very Large Scale Information Processing

Information processing consists of the storage, modification and transmission of meaningful patterns of information embodied in a physical system.

What constitutes meaningful patterns is of course observer- or user-dependent, so in the following I will consider systems processing general information patterns with no regard for the meaning of these patterns. However, the processing is constrained so that mappings of one meaningful pattern to another is possible, which in terms of general patterns means that general computation is feasible: the system can map one arbitrary stored pattern into another with high reliability.

Important physical factors that limit information processing are:

1. Processing and Memory density

---

[1] Evolution of intelligent beings directed by them instead of natural selection.
[2] The term originally coined by Charles Platt in *The Silicon Man* 1991

The individual computing elements and memory units will have a finite size, which is limited by the laws of physics. This places an upper limit on the processing and memory density of the system.

2. Processing speed

The speed at which information may be processed or retrieved from memory. As we will see, there are many kinds of limits on how fast a computing element can become, mostly based on the natural timescales of physical processes.

3. Communication delays

Since nothing can communicate faster than light there will be delays between different parts of an extended system. The faster the processing speed of a system, the longer the delays will appear from an internal subjective view. To minimize the delays the system has to be as physically small as possible or avoid the need for long-range communication.

4. Energy Supply

Computation (except for reversible computation) requires energy and dissipates heat both in principle and in practice. And even reversible computation is limited by the need for error correction and perception [35].

These four factors influence the shape and structure of possible computing systems. Factor 1 and 2 show that in the long run it is not possible to create ever more powerful individual processors; instead intelligence has to turn towards parallel systems where many elements work together (the human brain neatly demonstrates that even fairly slow and inefficient elements can produce a very powerful computing system)[3].

Factor 3 implies that as we move towards faster and faster systems they have to become smaller to avoid excessive delays, up to the limits set by factors 1 and 2. After that, extension of the system will increase delays which have to be handled in other ways than hardware. Keeping the system concentrated will help, as will modularization. One interesting structure is the "intelligent superobject" (the "Jupiter brain"[4]), a computing system on the planetary scale or beyond.

Factor 4 demonstrate that even very advanced civilizations will have to produce energy somehow and dissipate the resulting heat (which might make it detectable). A concentrated structure is easier to supply with energy, but removing waste heat becomes nontrivial, while a more distributed structure is easier to cool but also induces signal delays and distribution problems. Detecting heat emissions has been suggested as one approach of detecting activities of alien civilizations [28, 29].

It should be noted that there exist "software limits" in addition to the physical limits, such as the fundamental limits of Turing machines, limits due to problem complexity, the ability to create efficient and correct software and limits on information acquisition ability. Also, there are limits to the ability to interact

---

[3] It has been argued in [51] that parallelization is relevant mainly when energy is relatively evenly distributed in the computer, and for compact computers serial processing becomes more efficient since high-energy quantum operations can be used. However, practical limits to energy density, operation speed and energy dissipation limit this approach, as discussed in section 3, 4 and 7.

[4] The term likely originated by Perry Metzger.

with the surrounding physical world. These constraints will not be dealt with here, but are obviously important and worth considering.

3 Density

The density of information processing (the number of processors, gates, registers or bits per unit volume) is limited by several factors, from the range of materials that are available to fundamental properties of physics.

Normally information is encoded in solid matter since it is desirably stable, although the complete system does not have to be solid to process or store information; chemical networks [13] and DNA chains [63] are capable of general computation. Liquid systems are somewhat limited since the signal speed is determined by diffusion unless structures for high-speed communication are included, and in that case the relevant parts are still solid. In the following discussion mainly solid systems will be considered.

3.1 Molecular Matter

Molecular matter has a density on the order of 1 gm/cm^3 and consists of molecules separated by intermolecular forces. Information can be encoded as patterns of molecules, or in the state of molecules, for example by its conformation or by the sequence of subunits as in DNA (in pure DNA the information density is 32 atoms per bit). The bistable molecular switch described in [2] uses just 29 atoms to encode a bit.

According to [9], rhodopsin-based bio-optical memories could theoretically approach 10^{12} bit/cm^3 and in practice maybe $10^{10} - 10^{11}$ bit/cm^3 in the near future. In these memories large numbers of rhodopsin molecules are used to hold a single bit; theoretically each molecule could hold one bit, but access would be cumbersome and thermal noise would degrade the information.

Estimates from nanotechnology [27] suggest that a molecular computer may encode information in fluorinated polyethene molecules, where each bit is marked by the presence of hydrogen or fluorine on a certain carbon atom. Taking into account the need for reading devices and packing, Drexler arrives at an estimate of around 10 atoms per bit, which would correspond to on the order of $5 \cdot 10^{21}$ $bits/cm^3$ if diamondoid densities were to be reached.

3.2 Degenerate and Nuclear Matter

The information density of nanotechnology is limited by the properties of molecular matter. Denser forms of matter are believed to exist, such as the degenerate matter in white dwarf stars (with a density up to 10^6 gm/cm^3) and the nuclear matter in neutron stars with a density on the order of 10^{17} gm/cm^3 [1].

Degenerate matter consists of a lattice of nuclei surrounded by a degenerate electron gas; it is believed to be stable even at high temperatures such as 100,000 K. In naturally developed white dwarf stars the main component would be carbon and/or oxygen nuclei in a bcc lattice [16, p. 596]. Information could be stored

in cold degenerate matter by using different nuclei, such as the stable carbon isotopes C^{12} and C^{13} allowing information densities up to 10^{29} bits/cm^3, around six orders of magnitude more than in molecular matter. Whether information processing can be done is less obvious; one possibility might be to arrange nuclei to act as a quantum dot cellular automaton (similar to [44]) where nuclear spin is used instead of electron polarization.

Nuclear matter consists of hadrons squeezed together, possibly forming a gas, superfluid or a crystalline matrix. The bulk of neutron stars is believed to be superfluid [16], making it unsuitable as an information storage medium. Vortex and flux tubes through the superfluid might be stable enough to sustain information storage and process it through various wave modes and interactions with each other, but the density appears to be fairly low. It has been suggested that quark matter, matter consisting of a soup of quarks, could co-exist with nuclear matter in a mixed phase in the cores of neutron stars . This would result in a complex interior structure, with several different phases consisting of droplets, rods and sheets of quark and nuclear matter [37]. Another possibility is that pure quark matter is stable even in vacuum [77]. It is natural to ask if this kind of very dense matter can be used to sustain information processing.

The crucial question is whether stable structures are possible. This is currently uncertain, as it depends on both the temperature of the system and the detailed properties of quark-hadron mixtures, which are unknown. The rich structure appears promising, but doesn't prove that stable structures can exist (it could be that quark matter is just a homogeneous glop of quarks, although excited states may provide a new layer of complexity). Most studies of nuclear matter assume high temperatures, but low-temperature nuclear matter may have very different properties. If stable nuclear density structures can exist, then the theoretical limit on information storage would be a few quarks or hadrons per bit, resulting in an upper bound on the information density of 10^{40} bits/cm^3.

3.3   Theoretical Limits

Thermodynamics implies that the maximum number of bits of memory that can be stored in a system of average energy E is [51]

$$I = \frac{E}{k_b T \ln 2} + \log_2 \left(\sum_i e^{-E_i/k_B T} \right) \tag{1}$$

where T is the temperature, k_b Boltzmann's constant and E_i the different energy states.

If the computing system is viewed as a number of modes of elementary particles with total energy E in a volume V, [51] derives the maximum information as a function of energy as

$$I = \frac{4}{3 \ln 2} \left(\frac{\pi^2 r V}{30 \hbar^3 c^3} \right)^{1/4} E^{3/4} \tag{2}$$

where $r = \sum_l r_l$ and r_l is the number of particles/antiparticles in particle species l times the number of polarizations, $\hbar = h/2\pi$ the Planck-Dirac constant and c

the speed of light. For a system composed of one kilogram of photons ($9 \cdot 10^{16}$ J of mass-energy) confined to one liter the maximal capacity is $2.13 \cdot 10^{31}$ bits.

Quantum mechanics places a limit on the amount of information that can be stored in an isolated finite region of space with a finite energy content, the so-called Bekenstein bound [5]. Since the region and its energy content is bounded, the phase space of the system in the region is bounded. But due to quantum uncertainty, the phase space cannot be divided into arbitrarily small partitions (if the partitions are too fine, they will be impossible to distinguish and thus cannot encode any information), and an upper bound on the information content can be derived. According to [72, appendix C] we get:

$$I \leq \frac{2\pi E R}{\hbar c \ln 2} \tag{3}$$

where I is the information content, E is the energy of the volume and R is the radius. This can also be written as

$$I \leq kMR \tag{4}$$

where M is the mass in the region and k is $2\pi c/\hbar \ln 2 = 2.57686 \cdot 10^{43}$ bits/(m kg). For the entire solar system we get approximately ($M = 2 \cdot 10^{30}$ kg, $R = 7.375 \cdot 10^{9}$ m) $3.8 \cdot 10^{83}$ bits. On the other extreme, a hydrogen atom might theoretically encode $4 \cdot 10^{6}$ bits, and a proton only 44 bits (!). In a medium with constant density, the upper bound of the amount of information possible to pack in a sphere scales as R^4.

Another limitation on (useful) information densities is black hole formation. If the region is smaller than $2GM/c^2$ (where G is the universal gravitation constant) then a black hole forms, implying an event horizon closing off the region from two-way communication. We get the following bound on the information density $\rho = I/V$ from this and equation 4:

$$\rho \leq \frac{3kc^4}{16\pi G^2} \frac{1}{M} \tag{5}$$

From this we see that while the storage capacity of very dense systems grows as the system is extended, the information density eventually has to decrease since the packing cannot be too great or it will collapse; there is always an upper bound on the size of a constant density system where it will collapse. Of course, the Bekenstein bound may not be the least upper bound of the information content, and the physical properties of the computing system will also limit the information density. In practice equation 2 is a much stronger bound than the Bekenstein bound.

4   Speed

Anyone can write $E = h/t$ on a napkin and persuade you over lunch that it imposes a fundamental limit on the power-delay product. The trouble begins when people publish their napkins. – Rolf Landauer

The speed of computation is limited by the number of transitions the system can perform per unit time without breaking down. This in general depends on the available energy to perform the transitions and how strong the energy barriers of the system are.

Current (1999) semiconductor circuits show switching times down to the nanosecond scale, and MOSFET devices have demonstrated 7.8 picosecond switching [56].

The mechanical rod-logic described by [27] has switching times on the order of nanoseconds, and the gears simulated by [34] can rotate at rates up to 100GHz. However, this is a mechanical system, and limited by the speed of sound in the building material (In Drexler's design diamond, $\approx$17 km/s). Using electromagnetic interactions a computing element can change state even faster, like the quantum cellular automata of [44] which are predicted to stabilize on the order of picoseconds.

Molecular computation is limited by bond energies, above 10^{15} transitions per second the energy involved becomes bigger than the bond energies and the system starts to break up [59]. So molecular computers would have to work in the femtosecond range or slower.

A system using nuclear reactions would have a characteristic timescale on the order of strong interaction reactions, around $4 \cdot 10^{-24}$ seconds ($= h/m_p c^2$).

Here quantum effects become relevant, and to keep errors in switching time down the energy involved in a switch has to be higher due to the Heisenberg uncertainty relation $\Delta E \Delta t \geq \hbar$ (if the energy is smaller than $\hbar/\Delta t$ quantum fluctuations will become larger than the input signal). In the nuclear matter case the required switching energy has to be larger than $3.75 \cdot 10^{-20}$ J per operation. Note that this energy doesn't have to be dissipated, it is just necessary to prevent quantum noise from disrupting the computation.

Quantum mechanics places other limitations on how quickly a system can change state. It can be shown that a system with average energy E above the ground state determines the maximum number of mutually orthogonal states that the system can pass through per unit of time to $2E/h$ [53]. This energy is necessary even in the total absence of noise.

A similar bound on the switching time can be based on the Bekenstein Bound: if a switch transforms n input bits into n output bits, then the switch has to be larger than

$$R \geq \left[\frac{n \hbar c \ln 2}{2\pi} \right] \frac{1}{E} \tag{6}$$

where E is the energy inside the switch, just to contain the information. Since information can only be transmitted at lightspeed, the time it takes for the resulting bits to move over the distance R to the next switch (which is a lower limit on the cycling time ignoring the switching time) is

$$t \geq \left[\frac{n \hbar \ln 2}{2\pi} \right] \frac{1}{E} \tag{7}$$

If we assume n=2 bits in each switching operation, then for energies less than $4.44 \cdot 10^8$ J per switch this bound limits computation, not quantum gravity.

The bound in [53] is a factor $\pi^2/\ln 2 \approx 14.23$ larger than the bound in equation 7 for two-bit systems.

Perhaps the ultimate limit to switching would be the Planck time, $5 \cdot 10^{-44}$ s ($\sqrt{G\hbar/c^5}$), since at this scale spacetime appears to loose its smooth properties and time becomes hard to define. At this switching speed, the uncertainty relations require an operation energy on the order of the Planck energy ($\sqrt{c^5\hbar/G}$), 3 J per operation.
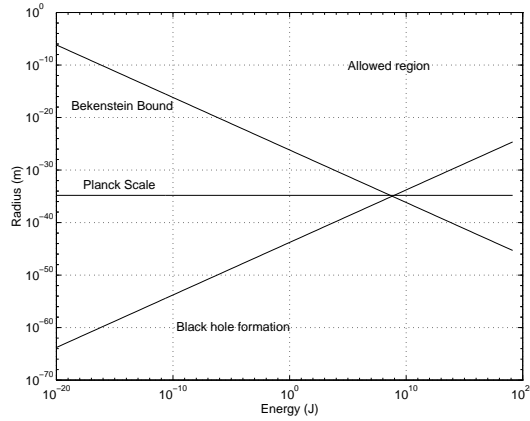

Figure 1: Constraints on the radius of a two-bit system.

5   Size

> I was given to understand that She manipulated whole sciences and thought systems as I might string words into a sentence. But Her 'sentences' were as huge and profound as the utterances of the universe itself. − David Zindell, *Neverness*

So far, we have only looked at the limits of the parts making up the information processing system, not the entire system as a whole. As long as the system is small, its size doesn't matter much, but larger systems obviously require special considerations due to signal delays, mass and thermal dissipation.

If longer signal delays are acceptable the system does not have to be concentrated to a single structure and the design can be made more freely. If delays are costly, such as in very fast systems (where a lot will happen while the signal snails across space), then the system has to be built as densely as possible. For very large systems this will impose constraints due to gravity.

A spherical structure made of diamond or any other material held together by covalent molecular bonds will become unstable if the pressures inside it exceed the binding strength of the bonds ($6 \cdot 10^{-19}$ J/bond). A rough estimate of the maximal pressure: the crystal structure will collapse if the work done when moving together two atoms is equal to the bond energy. Assuming a cubical lattice, we get $1/l^2 \approx N$ atoms/m^2 where l is the bond length of 154 pm.

Breakdown will occur when $l \cdot (P_{crit}/N) \approx E$, which gives $P_{crit} \approx E/l^3$, $P_{crit} \approx 1.64 \cdot 10^{11}$ N/m^2.

The pressure inside a spherical body of constant density is

$$P = \frac{2\pi G \rho^2 R^2}{3} = \frac{3GM^2}{8\pi R^4} \qquad (8)$$

So a compact diamond structure would have a maximum radius on the order of 9760 km, somewhat larger than the Earth. Halving the density doubles the possible radius and quadruples the mass, which suggests a trade-off between internal delays and total computing power.

Although harder materials are possible, molecular bonds are of this order of magnitude so larger structures will not work unless they are stabilized by other forces.

In the same way it is known [16] that electron degenerate matter becomes unstable at the Chandrasekar limit (1.44 solar masses) and nuclear matter at the Landau-Oppenheimer-Volkoff limit (somewhere between 2 or 3 solar masses, depending on the equation of state of nuclear matter [76]). The maximum size of neutron stars and similar objects is small (on the order of 10 kilometers); while the small size and high density is positive for information processing the limits on mass limits the total information capacity to around 10^{60} bits. For both spheres of degenerate matter and neutronium the relation $RM = $ constant holds, making delays shorter as the system becomes denser.

The ultimate limit to size is of course relativistic collapse into a black hole; sufficiently large masses distort spacetime so much that nothing can move outwards and the structure vanishes behind an event horizon. For a non-rotating structure with mass M this corresponds to the Schwartzchild radius:

$$r_s = \frac{2GM}{c^2} \qquad (9)$$

while a structure with angular momentum L and charge Q has a critical radius (In the Newman metric):

$$r_k = m + \sqrt{m^2 - a^2 - e^2} \qquad (10)$$

where $m = GM/c^2$, $a = L/Mc$, $e = GQ/\epsilon_0 c^2$. These formulae suggest two means of increasing the mass of the object while still keeping it stable: rotation and charge.

A rotating structure can indeed become larger than a static structure, since the pressure is partially balanced by centrifugal forces (seen classically); an object with sufficiently large angular momentum or charge can also escape becoming a black hole [57]. While the obvious solution is to rotate the entire structure as a rigid body, a subtler possibility is to use a more complex distribution of angular momentum.

It has been proposed [52, 38] that streams of magnetic pellets or bars could be used to build dynamic structures for space access. The idea is to accelerate a stream of pellets electromagnetically from the ground, send it towards a receiver

on the underside of a space station where the stream is redirected downwards using electromagnetic fields (and hence, by conservation of momentum, induces a net force on the station keeping it aloft despite being static relative to the ground) and finally received using another station on the ground sending them back to the accelerator.


Figure 2: A space fountain system. Magnetic pellets are shot from a ground station towards a space station where they are deflected downwards to the ground station. The net result is a "pillar" held up by momentum transfer.

This system imparts a net force on the station proportional to the kinetic energy of the pellet stream, and could be used in a modified form for "pillars", "beams" and "arcs" in large structures. For example, two circular streams moving in opposite directions (to keep the net angular momentum zero) around the equator of a spherical structure could function as an "arc" on which other structures could be attached using electromagnetic fields. As the load increased, more momentum could be added to keep up with the stresses. The practical limits to this kind of dynamic structure seem to be energy dissipation and how well the forces from the fields can be distributed in the total volume of the structure. The accelerations of the pellets will produce electromagnetic radiation, leading to energy losses proportional to the square of the acceleration.

Rotation and electromagnetism mesh well in stabilizing ultra-heavy structures. In the neutron star case, a strong magnetic field may [10] increase the maximal mass with up to 13-29% compared to a nonmagnetic system. In this case the fields interact with spacetime in a complex way to stabilize the system, which could presumably be used by intelligent life.

A more basic design would be to give each component of the structure an electric charge large enough to counteract the attraction from gravity; since both fields scale as r^{-2} this can essentially neutralize gravity between similarly charged objects. One of the main problems with this design is that the structure will tend to attract oppositely charged particles from cosmic radiation or the solar wind; some form of shielding or neutralizing screening is necessary. Another problem is that if the fields become strong enough the electrical repulsion between the constituent particles of the units will begin to break them up (not to mention

that the charges tend to accumulate on the surfaces of particles), which forces the use of smaller units. Truly extreme fields, such as those necessary to prevent gravitational collapse into black holes, will also suffer from pair production when $E \geq m_e^2 c^3 / \hbar e$ and likely break down in a way similar to the electric fields around black holes described in [65].

These Rube Goldberg solutions just hint at what can be done by carefully balancing different forces against each other. There don't seem to be any physical limits to the size of well balanced systems, although lack of raw materials, security concerns and energy will by necessity limit their size. And as we will see, thermodynamics places other limits on the size and structure of large computing systems.

6 Energy

One of the most fascinating developments in modern physics is the link between computation and thermodynamics.

In thermodynamics entropy is a measure of disorder, or rather the logarithm of the number of possible states the system could be in that are consistent with the observed macroscopic state, and according to the third law of thermodynamics entropy always increases. If we look at a computation that erases information, such as setting a register to a zero, we decrease the number of possible states (previous values of the register) to one state (zero). This would decrease the entropy of the system, which is not allowed – unless the entropy of some other aspect of the system is also increased at the same time. Thermodynamics leads to the need for energy dissipation to erase bits of information: energy is dissipated into heat to take away the entropy increase.

The cost of erasing one bit of information is given by Brillouin's inequality:

$$\Delta E \geq k_B \ln 2 T \tag{11}$$

where ΔE is the amount of energy expended, k_B is the Boltzmann constant, T is the absolute temperature. At room temperature this cost is $2.9 \cdot 10^{-21}$ J, while at 3 K (the cosmic background temperature) it has decreased to $2.87 \cdot 10^{-23}$ J. Obviously, to be able to process information cheaply the system should be very cold, but that may require extensive cooling. Thus there will be extensive demands for energy in large information processing systems. Very dense and fast systems will dissipate huge amounts of energy; assuming a molecular computing system with 10^{12} bits/cm^3 and a switching speed of 10^{15} Hz would lead to an energy dissipation of $2.8 \cdot 10^6$ W/cm^3, which would obviously vaporise the material.

6.1 Reversible Computing

It may be possible to do computations without having to expend energy at all if no bits are erased, so called reversible computation. In a logically reversible process the input and output can be logically retrieved from each other. A physically reversible process is not just logically reversible, but it can be run

backwards, producing the input from the output and vice versa. According to the second law of thermodynamics it cannot dissipate heat.

Examples of reversible operations are copying a record and its inverse, 'cancelling' a record with an identical record if they are known to be identical. Adding two numbers is reversible as long as a record of one of the terms is kept.

Logically reversible computers could be built from reversible circuits [32] or the reversible Turing machine [6]. Physical reversibility can be achieved using reversible logical circuits [54], mechanical logic [55] or by using quantum computation which by its nature is reversible (see section 8.1).

It has been shown that any irreversible computation can be turned into a reversible computation with a slight increase in memory and time complexity [7]: if the time needed is T and the memory demand S, then the output can be calculated reversibly in time linear in T and space of the order $O(ST^a)$, where a can be made arbitrarily small. It is also possible to communicate reversibly [35] between two reversible minds.

Unfortunately there are limits to the usefulness of reversible computation. Error correction is by necessity irreversible (several erroneous states are mapped to a single correct state), and hence needs irreversible operations. There is a tradeoff between dissipation and decreasing the risk of undetected bit errors; by using error-correcting codes the number of bits in the system is increased (and hence the net number of bit errors) but more can be corrected.

Another way to decrease the problem of bit errors is to make the potential wells of the registers deeper; this makes it less likely that thermal noise or outside interference (such as cosmic rays) will throw the register from one state to the other. If the register is similar to a harmonic oscillator, then the probability that thermal noise kicks it out of the current potential well of height E is proportional to $e^{-E/k_B T}$, and can be made arbitrarily small by increasing E. Fortunately, the depth of the potential well does not matter for reversible computation, so during error-free operation no energy has to be dissipated even when E is large, and correcting the error only requires $k_B T \ln 2$ Joules per bit. If there are n bits in the system, then the total energy that has to be dissipated for error correction is proportional to

$$E_{diss} \propto n \ln 2 k_B T e^{-E/k_B T} \tag{12}$$

which decreases quickly when E grows. We see that in order to get reliable storage and calculation deep, stable potential wells are necessary for any system. Molecular bonds provide one source of potential wells, with a depth on the order of 10^{-19} J, which makes them stable up to temperatures on the order of a few thousand degrees Kelvin. Nuclear bonds are on the order of 10^{-12} J, which makes them stable up to 10^{10} K.

There are also problems for reversible systems when dealing with the rest of the universe, which is highly irreversible and disordered. There is no guarantee that information gained from the environment can be undone to ensure reversibility since the outside world changes unpredictably [35]. A very large computational system, like a planetary sized "solid state civilization" may of course choose to ignore the outside world completely, but if it is necessary to interact with it energy has to be dissipated.

It thus appears likely that large computing systems will have to use energy to do error correction, interact with the environment and to do physical work (such as repairing damage), but the energy requirements are relatively modest.

6.2 Getting Rid of Heat

If the system dissipates energy it will have to radiate away the resulting heat. If the system is a spherical blackbody with dissipation is P watts/m^3, Stefan's law of blackbody radiation gives us a temperature of

$$T = \left(\frac{rP}{3\sigma}\right)^{1/4} \tag{13}$$

Since the volume grows with the cube of the radius while the cooling area only grows as the square of the radius, the temperature increases as the system grows[5]. If the maximum allowable temperature is T_{max}, the maximum size becomes r_{max}:

$$r_{max} = \frac{3\sigma T_{max}^4}{P} \tag{14}$$

For molecular matter, T_{max} would probably be of the order of the melting point (or one or more orders lower if there are fragile subsystems), around 1000 K or so (this fits with the observations of [34], which suggests that nanomechanical gears may begin to fail at 600–1000K). This gives a maximum radius of $170/P$ km (if the 600K limit is chosen, the radius will be just $22/P$ km). The size of P will depend on how much irreversible calculations occur, which is very hard to estimate, but it is obvious that even quite low dissipation densities places a limit in the asteroid size range or forces strict rationing of dissipation.

Another useful measure is the information production, which can be found by deriving equation 11 [4]:

$$\frac{(dI/dt)}{(dE/dt)} \leq \frac{1.05 \cdot 10^{23}}{T} \mathrm{bit/sW} \tag{15}$$

If we plug it into Stefan's law, we get

$$\frac{(dI/dt)}{(dE/dt)} \leq \left(\frac{4\pi\sigma}{k_B \ln 2}\right) T^3 r^2 \tag{16}$$

This implies that the information production of the system grows as T_{max}^3 and r^2.

If we assume the system only needs energy for correcting errors, and all errors are of thermal nature, then we reach the following relation between size and temperature by balancing equation 13 and equation 12 (with bit density ρ):

$$r < \left[\frac{3\sigma}{\rho k_B \ln 2}\right] T^3 e^{E/k_B T} \tag{17}$$

---

[5] This argument was originally pointed out to me by a poster on the extropians mailing list in 1995

This represents the possible sizes of "cold" brains, which use reversibility to the fullest. The maximal radius decreases as T increases up to $T_{crit} = E/3k_B$ (for molecular matter approximately 2400 K, corresponding to an error rate proportional to $e^{-3} \approx 0.05$), and then grows as T^3 beyond that. In the cold region large systems are possible since they have an extremely low rate of errors, avoiding the need for self heating by error correction. In the hot region above T_{crit} errors are so common that self-heating from error correction is significant and the limiting factor of size. Beyond $T_{max} \propto E/k_B$ it is unlikely that the system will be stable, as the thermal noise drowns out all storage. Obviously in this case, the optimal solution is to keep as cold as possible.

If we assume an additional error rate per bit ϵ independent of T, for example cosmic rays or bugs in the system, formula 17 becomes

$$r < \left[\frac{3\sigma}{\rho k_B \ln 2} \right] \frac{T^3}{\epsilon + e^{-E/k_B T}} \tag{18}$$

Unlike the purely thermal system this system has a local maximum at T^* beside the minimum at T_{crit} if $\epsilon < 0.001$ errors per second. At temperatures lower than T^* the system is dominated by the need of removing the heat from the non-thermal errors, and the radius grows like T^3. Above T^* thermal errors become more significant and the radius has to decrease to keep the system cool enough, until T_{crit}. This suggests that for imperfect but nearly reversible systems keeping at a temperature T^* maximizes the possible volume of computing elements.


Figure 3: Maximal radius of information processing system as a function of working temperature (as per equation 18, with $\rho = 10^{18}$ bits/m^3, $E = 10^{-19}$ J and $\epsilon = 10^{-6}$).

In the region above the curve in figure 3, the system will tend to heat up (i.e. move to the right), in the region below the curve the system will cool down (i.e move to the left). For a given radius there may exist a single or two stable states, corresponding to a cold system where heat is dissipated faster than it is generated (the branch where $T < T^*$), or a hot system where much energy is used for error correction (the branch where $T > T_{crit}$.

The above calculation has not taken the cosmic background radiation into account; this makes the curve level out for small T, and makes it impossible (given

the assumptions of only passive cooling) to go beneath a 3 Kelvin working temperature.

It should be noted that the above limits all assume a spherical radiating surface, which is pessimal. A more reasonable system would be extended radiators circulating cooling substance into the computational core and radiating away the heat into the cosmic background, or if the delays could be managed, a distributed system where the total surface area would become much larger. A planetary sized structure could be orbited by a circle of large radiator sails in the geosynchronous (cerebrosynchronous?) orbit extending connecting pipes down to the "ground".

By expending enough energy, the system can in principle cool itself down to any finite temperature, including below the cosmic background temperature. One interesting passive cooling device that has been suggested by Wei Dai is to use large black holes; since the Hawking radiation decreases with increasing mass,

$$T = \frac{\hbar c^3}{8\pi G k_B M} \tag{19}$$

they can become extremely cold (a solar mass black hole would have a temperature of 10^{-8} K) if they are prevented from accreting in-falling material. One possibility would be an "inverse Dyson shell" surrounding the black hole.

6.3   Energy Sources

Another major problem is where to get the energy to dissipate in the first place. For small-scale operations energy is plentiful: in near-Earth space there is always around 1000 W/m^2 of solar energy, which is very convenient for satellites and keeps the biosphere running. The trouble begins when the energy needs become larger.

One classic solution to this problem is Freeman Dyson's proposal [28, 29] to englobe the sun with solar collectors. This way a large fraction of the total energy output of $4 \cdot 10^{26}$ W could be used to do work, such as information processing. There is enough material in the solar system to build a complete shell [66], and this could form the basis for a very large distributed processing system where waste heat would be dumped into the cosmic background on the exterior, and signals sent through the interior of the shell.

The total power of energy striking the shell is independent of its radius, but as the shell gets smaller the energy flux becomes larger and the thermodynamic efficiency becomes better. The temperature of the shell (assuming it to be a blackbody and that the cosmic background radiation is at $T_u = 3$ Kelvin) is

$$T_{shell} = \left[\frac{E}{4\pi\sigma r^2} + T_u^4 \right]^{1/4} \tag{20}$$

The efficiency becomes

$$\eta = \frac{T_{shell} - T_u}{T_{shell}} \tag{21}$$

which decreases monotonously as the shell becomes larger.

The shell will become useless at a very large radius since it will be nearly isothermal with the cosmic background and no work can be extracted. A smaller shell will also require less material to build, but must withstand higher temperatures; the logical conclusion would be a tightly fitting shell just above the chromosphere, or even surrounding the stellar core[6] if such structures can be built. But this would require extensive cooling, since the cost of erasing bits scales as $k_B T$. In fact, the optimal radius for a shell cooled solely by blackbody radiation when maximizing information processing is rather large, since the bit-erasure cost has to be balanced against thermodynamic efficiency: the maximal amount of bits that can be erased per second is $I \propto (1/T_{shell})(1 - T_u/T_{shell})$ where the first factor is due to equation 11 and the second the efficiency of equation 21. The maximum occurs when $T_{shell} = 2T_u$, at a radius of

$$r_{Imax} = \sqrt{\frac{E}{60\pi\sigma}} \frac{1}{T_u^2} \tag{22}$$

In the case of the sun this would correspond to $r_{Imax} = 6.8 \cdot 10^{14}$ m, an immense shell 629 light-hours from the sun, with the capacity to erase $1.2 \cdot 10^{50}$ bits per second. This is obviously far too large to be practical due to material requirements and signal delays. This suggests that feasible Dyson shells (be they large and cold or small and hot) either use most of their energy for cooling, or beam the concentrated energy out to external information processors.

Stars are excellent sources of energy, but quite inefficient; most of the bulk is not used for energy production at any time in the stellar lifecycle. Fusion is clearly the most powerful and easily accessible energy source in the present era: hydrogen and helium are the most plentiful elements, and can be harvested from gas giant planets or stars for practical use (for example by "star lifting" as suggested by [23]). Advanced civilizations may dispense with stars altogether, relying only on artificial fusion power and keeping in the cool interstellar.

Another possible (if somewhat speculative) high density energy source would be matter-energy conversion. It has been argued [22] that advanced civilizations could create very small black holes using intersecting beams of extremely powerful gamma radiation just for this purpose. Black holes radiate Hawking radiation with a power of

$$P = \left[\frac{\sigma \hbar^4 c^8}{1024\pi^3 G^2 k_B^4} \right] \frac{1}{M^2} = \frac{8.909 \cdot 10^{31}}{M^2} \mathrm{W} \tag{23}$$

where M is the mass. Once created, a small black hole could be kept stable by a constant supply of matter equal to its radiation loss (taking care to avoid making it too small, which would lead to evaporation, or too large, which would make it less efficient). Gathering the intense radiation in an useful way is left as an exercise for the posthuman engineer; most likely the radiating hole could be surrounded by something similar to a Dyson shell.

The ultimate limits of energy production lie on the cosmological scale, such as extracting shear energy from a collapsing spacetime [4, 72] discussed in section 8.4. Although the total amount of energy that can be extracted in a finite region

---

[6] This idea of submerged "Dyson shells" is due to Nick Szabo.

16

is still limited, it appears quite likely that producing enough energy even for very ambitious projects is possible

7 Communications

> I learned much about the Entity's sense of Herself. Each moon-brain, it seemed, was at once an island of consciousness and a part of the greater whole. And each moon could subdivide and compartmental-ize at need into smaller and smaller units, trillions of units of intel-ligence gathering and shifting like clouds of sand. – David Zindell, *Neverness*

The basic limits that apply to communication are the finite speed of light, noise in the communication channel and finite bandwidths.

7.1 Delays

The finitude of lightspeed inevitably introduces delays in the signals sent across an extended system; even for today's computers it has begun to influence the design.

A measure of how communication delays relate to the size of the system is:

$$S = \frac{(\text{length})}{(\text{speed})(\text{time})} \tag{24}$$

where the time is the time for one "clock cycle" or other characteristic transi-tion, the speed is the speed of signals sent through the system and length the characteristic length scale of the system. The higher this ratio is, the more clock cycles will occur during the wait for information stored elsewhere in the system, and hence a larger "subjective" delay.

In the case of humans, the neural cycling time is on the order of 10^{-3} s, the communication speed around 100 m/s and the length scale 0.1 m, which gives a S value on the order of 1. For a standard microprocessor S is on the order of $10^{-2} - 10^{-3}$ which implies that microprocessors currently do not have to wait long for information from other parts of the chip (as of 1999, but the problem is clearly growing). The mechanical nanocomputer in [27] have $S = 10^{-2}$. The Internet has a size on the order of 10^6 m, communications close to lightspeed and intercontinental delays on the order of hundreds of milliseconds; S becomes around 3, slightly higher than the human mind.

The subjective effects of S depends on the application. For data retrieval and communication, it just creates a subjective delay which may or may not be acceptable (a delay of a minute in delivering an e-mail is usually acceptable; a one-minute delay in delivering a frame of video is not acceptable). Subjec-tive distances increase for very fast minds; for entities exploiting nanosecond timescales at the speed of light distances of centimeters are significant, for fem-tosecond entities micrometers and for nuclear entities femtometers. Structures larger than this will be "large" compared to the processes that go in them.

For infomorphs, delays limit the physical distribution of their component processes: if they are too far apart, the being would have to slow down its rate of subjective time in order to keep synchronized. Even if the processing is infinitely fast lightspeed limits the speed of infomorphs if they wish to interact with the outside environment at a certain rate; since the human mind acts as a whole on a timescale of hundreds of milliseconds, a human-like infomorph running at "normal" speed would at most be able to extend 30,000 kilometers before the delays started to limit its speed.

How can humans react quickly when their brains as a whole are so slow? The answer lies in modularization: low-level systems do most of the basic processing quickly and often manage most of perception and behavior on their own; slower higher level systems consisting of many low-level systems step in to regulate lower levels when needed. The most interesting aspect of this is that the conscious mind consistently seems to misattribute behavior and perception to itself, even when they are done by lower levels and occur several hundred milliseconds before they become conscious [47, 46]. In the same way the conscious mind experiences itself as unitary, despite all internal delays or even removal of central interconnections as the corpus callossum.

This suggests that a mind may exist on a wide range of timescales. One may conjecture that similar hierarchical modularizations with even more levels are possible, which would enable much larger minds with longer internal delays without loosing their high-level unity. By necessity, the highest levels would be much slower than lower levels, but this would not significantly impair their performance since most of it would take place at the quicker lower levels, and the higher levels would experience it as if they were doing things in realtime despite their slowness.

7.2 Bandwidth

> There is an old network saying: Bandwidth problems can be cured
> with money. Latency problems are harder because the speed of light
> is fixed. You can't bribe God. – David Clark

Unfortunately, the amount of information that can be sent over an information channel is limited. According to the Nyquist theorem, the highest signal rate that can be carried over a channel with bandwidth W Hz is $C = 2W$ bits/second. By using multilevel signaling $C = 2W \log_2 M$, where M is the number of discrete signal levels, but the more levels in the signal, the more noise-sensitive it becomes. Furthermore, forcing a physical quantity into one of 2^k possible ranges seems to be 2^k as hard as forcing it into one of two ranges, rather than just k times as hard [8]. In the following we will assume binary signals through the channel.

Using higher and higher frequencies of the electromagnetic spectrum extremely high signal rates can be sent in a directional manner, for example using lasers. Unfortunately there are some problems involved with extremely high frequencies due to pair production: in the presence of another particle or a field, the gamma-ray photon may split up into pairs of electrons and positrons. This occurs at a frequency of

$$\nu = \frac{2m_e c^2}{h} \qquad (25)$$

at a bandwidth of $2\nu \approx 4.9 \cdot 10^{20}$ bit/s. Although this does not necessarily imply a limit on the bandwidth, it implies a growing source of noise. And since the energies of individual quanta become higher, the number of quanta per Watt signal-strength decreases, leading to increasing noise (see below).

Also, there is an upper limit to the rate of information that can be sent using electromagnetic radiation for a given average energy [17, 42]:

$$C = \left(\frac{512\pi^4}{1215h^3c^2} \frac{A_t A_r}{d^2} E^3 \right)^{1/4} \qquad (26)$$

where A_t and A_r are the areas of the the transmitter and receiver respectively, d their distance and E the power of the transmitter. For a transmitter and receiver one square meter each one meter apart and with a 1 J/s energy budget the information rate is $1.61 \cdot 10^{21}$ bits per second. The rate scales as $E^{3/4}$. The optimal spectrum turns out to correspond to blackbody radiation, but if the receiver can only detect the energy and timing of arriving photons the spectrum instead corresponds to the spectrum of black bodies in a one-dimensional world and the information rate becomes

$$C = \left(\frac{4\pi^2}{3h} E \right)^{1/2} \qquad (27)$$

which is independent of transmitter and receiver area and distance. For an 1 J/s energy budget the maximum information rate becomes $2.03 \cdot 10^{17}$.

One obvious way to circumvent this problem is to send information encoded in small pieces of matter at high speed. The energy requirements are much larger when lightspeed is approached, so the energy efficiency

$$\eta = \frac{C}{P} = \frac{km}{(\gamma - 1)mc^2} = \frac{k}{c^2} \frac{\sqrt{1 - v^2/c^2}}{1 - \sqrt{1 - v^2/c^2}} \qquad (28)$$

(where P is the energy used to accelerate the matter, k is the number of bits per kilogram and v is the final speed) decreases towards zero. On the other hand, the efficiency for very low speeds is high, but is balanced by the longer delays.

As always, Bekenstein's bound introduces a constraint on information flow. The message channel can be viewed as a chain of regions of size R containing energy E, in which information flows from one to the next in time R/c (assuming light-speed transmission). This gives a bandwidth limitation of

$$C \le \left(\frac{2\pi c}{\hbar \ln 2} \right) \frac{RE}{(R/c)} = \left(\frac{2\pi}{\hbar \ln 2} \right) E \qquad (29)$$

or around $9 \cdot 10^{34}$ bit/(s J).

Regardless of the amount of energy used in transmitting information, an additional limit is the Planck bandwidth

$$W = 2\sqrt{c^5/hG} = 2 \cdot 10^{43} \text{bit/s} \qquad (30)$$

At this bandwidth, quantum gravity becomes important and the wavelength of individual quanta becomes less than their Schwartzchild-radius.

It should be noted that the above limits apply to single channels; by using several noninteracting channels the information transmission can be increased further.

7.3   Noise

In reality, the channel capacity is somewhat lower due to noise. Shannon demonstrated that the maximal channel capacity (also called the error-free capacity) in the presence of noise is

$$C = W \log_2(1 + \frac{S}{N})$$

(31)

where S is the signal power and N the noise power. Shannon also proved that if the information rate is lower than the error-free capacity, then it is possible to use a suitable coding to completely avoid errors. If energy dissipation is no problem, then noise can be ignored. Otherwise, the bandwidth will at least grow as the logarithm of the power used.

Noise leads to the problem that energy has to be expended in sending the information. In a noiseless channel information can be sent without dissipation [43], but the minimum energy per unit of information required to transmit information over a channel with effective noise temperature T satisfies the inequality

$$\frac{E}{I} \geq kT$$

(32)

as shown by [45]. The dissipation will be $E(T)$ J, where $E(T)$ is the minimum possible energy for the system with a given entropy; not all energy used in the information channel will be lost.

For extremely dense and high-bandwith systems energy dissipation from communication will likely play an important role, a role that cannot easily be circumvented with reversible computing. The exact amount of communications used is however very architecture dependent, ranging from nearly none in passive repositories of information to R^6 in 3D-structures where every node communicates with every other node.

8   Exotica

> Any sufficiently advanced technology is indistinguishable from magic.
> – Arthur C Clarke

So far we have looked mainly at what may be possible according to classical physics. If we turn towards purely quantum phenomena or more speculative areas, new possibilities emerge for information processing systems.

8.1 Quantum computers

We have mainly assumed that information processing is done using classical Turing machines. If quantum computation is taken into account, the potential power grows significantly.

Formally, programs are executed on quantum computers by the unitary evolution of an input that is given by a state of the system. This form of computation uses the counterintuitive properties of quantum mechanics, like placing bits in superpositions of 0 and 1, using quantum uncertainty to generate random numbers and creating states that exhibit purely quantum-mechanical correlations [31]. A famous result by [69] showed how factoring can be achieved in polynomial time on a quantum computer, database searches can be done in $O(\sqrt{n})$ time [33] and many-body quantum mechanical simulations can be run with an exponential increase in speed [11].

It is important to realize that quantum computers are qualitatively more powerful than classical computers, not just quantitatively better. They are at least equivalent to probabilistic Turing machines, and possibly more powerful [71], although they cannot solve general NP-complete problems [8]. It is known that there exist universal quantum computers that can emulate all other quantum computers just as universal Turing machines can emulate all other Turing machines [24].

Quantum computation is reversible except for the irreversible observation step when the state of the computer is measured macroscopically. This is due to the fact that the quantum computer operators are all unitary (and hence logically reversible). At first this suggests that quantum computers will be unusable in reality, since they would lack error correction, but this is surprisingly not true. By splitting the signal across several channels with partial error correction and then merging them, error-correction can be achieved [70]. It is also possible to make the system fault tolerant so that errors during error correction can be avoided [25]. A theorem similar to the Shannon theorem holds for quantum channels [49], although quantum information introduces some new complexities [3].

Quantum computation also relates to the field of quantum cryptography, the use of quantum mechanical effects to transmit information in a way that cannot be eavesdropped, even against an adversary with unlimited computing power. The basic idea is to exploit the existence of pairs of conjugate properties, where one cannot be measured without disturbing the other; the eavesdropper cannot avoid disturbing the communication. Quantum-secured communication has already been demonstrated [18, 14] and will likely be an important part in secure communication in a world with extremely powerful computation. It might be possible that quantum devices and quantum computers could eavesdrop this kind of channel, but it is highly uncertain. While quantum key-distribution is secure, "post-cold-war" applications such as two-party secure computation (where both parts want to know the answer but not reveal their data) have been shown to be breakable [19].

Quantum computers are the natural choice of computing systems on the nanoscale (and even moreso on the femtoscale). A physical implementation requires coherent, controlled evolution of the wavefunction at least until the computa-

tion is completed. Various possibilities for implementing quantum gates have been proposed, such as influencing the excitation states of atoms using external electromagnetic fields or laser pulses, interacting quantum dots [44, 31] or heteropolymers [50]. Prototypes of quantum gates based on nuclear magnetic resonance in bulk liquids have actually been made to work [20], with up to seven qbits [40].

It is hard to tell what importance quantum computation will have in very large computational systems except for the obvious speed, density, security and complexity power advantages. For example, is there a difference in power between minds using quantum information and minds using classical information? The exponential speedups and possibility of simulating physical systems efficiently appears to be a great advantage, but are they generally useful for advanced information processing?

8.2 Black Holes

If black holes do not destroy information (this is currently controversial), then information trapped inside will be released through Hawking radiation, and if they evaporate unitarily they can in principle be used as processing elements as suggested in [51].

A computing system compressed to the Schwartzchild radius would have an information content of

$$I = \frac{4\pi G m^2}{\hbar c \ln 2} \tag{33}$$

Given the quantum-mechanical limitations on state-switching in section 4 the time it takes to flip a bit is

$$t_{flip} = 2\pi \hbar I / E = 4\pi^2 R / c \ln 2 \tag{34}$$

which equals the amount of time it takes to communicate from one side of the hole to the other.

The total lifetime of the hole would be

$$t_{life} = \frac{G^2 M^3}{3 C \hbar c^4} \tag{35}$$

where C is a constant depending on the number of particle species with masses less than $k_B T$ for the hole. For $O(10^1 - 10^2)$ species C is on the order of 10^{-3}–10^{-2}. For a one kilogram black hole the lifetime would be around 10^{-19} s, during which 10^{31} operations can be performed on 10^{16} bits. In one second, this would produce 10^{50} operations, the quantum-mechanical maximum information processing rate of one kilogram of matter.

This would be a highly serial operation, since black holes cannot be packed very closely. However, by charging the black holes so that they become extreme Reissner-Nordström black holes it might be possible to create a dense volume of "black hole processors" where their mutual attraction is balanced by electrical repulsion (but see [65]). What kind of mechanisms would be necessary to create the holes and send/receive information from them remains rather speculative.

8.3 Wormholes

General relativity at least in principle allows wormholes, "tunnels" that connect remote parts of space-time with shortcuts [62]. Morris and Thorne [61] have demonstrated that a static, traversible wormhole is formally possible, as long as the wormhole "throat" is threaded with sufficient negative energy densities (such as Casmir forces between conducting spherical plates). Wormholes that do not violate the energy conditions [68] and self- maintained wormholes might even be possible [39]. On the other hand, the topological censorship theorem [41] appears to forbid wormhole topologies. The theorem is based on the assumptions of a globally hyperbolic asymptotic flat spacetime where the averaged null energy condition holds (the averaged energy over the whole of spacetime is positive); if these conditions do not hold or can be circumvented, then wormholes might be possible.

Whether such wormholes can exist and be built in practice is uncertain, but if they are possible wormholes would be extremely useful for long-range information processing and communications. A system could be distributed across the universe and still remain a cohesive unit with short internal delays by using communications through a network of wormholes; the limits to size discussed in previous sections would be circumvented.

Wormholes still have some limitations. Since the negative energy densities across the throat have to be larger than a critical value, the amount of mass/energy that can be sent through the wormhole is limited since it partially cancels the negative energy during its passage. It is also likely that an analogue to Bekenstein's bound holds for wormholes, limiting the bandwidth across a wormhole. Assuming a (negative) wormhole mass equal to a similarly sized black hole and a relation of the same order as the Bekenstein bound, the maximal bandwidth across the wormhole would be on the order of

$$C \leq kRE = \frac{kc^4 R^2}{2G} \tag{36}$$

For a one meter wormhole, C is $1.56 \cdot 10^{87}$ bit/s. For a nanometer wormhole C is $1.56 \cdot 10^{69}$ bit/s, so it is likely that other bandwidth limitations become relevant long before the wormhole is saturated.

Another problem with wormholes is the possibility of time travel; by moving one end of the wormhole at a relativistic speed a time differential could be created, so that signals sent into one end would appear in the future or past as seen from the other end. Whether such closed timelike curves (CTCs) can exist is debated; although they might be allowed by physics if the principle of self- consistency holds [15], most physicists think there will be censoring effects (the "Chronology Protection Conjecture" of [36]), such as the quantum-field build-up suggested by [74]. If they hold, either wormholes cannot be created, or they will be forced to form networks that do not allow time travel.

It should be noted that if nonstandard causality is allowed, information processing could turn extremely strange. The halting problem could be solved[7] and all

---

[7] the computer would simply report if it received a signal from the future or not, and then start running the program. If the program halts, the computer will send a signal back in time to itself.

calculations could be done in $O(1)$ time[8]. Inconsistent programs (receive X from the future, send back $X + 1$) would force an outside error (such as transmission noise or circuit breakdown) to keep the universe consistent.

8.4 Physical Eschatology

The ultimate limits of computing overlap with the field of "physical eschatology", the study of the possible long term evolution of matter and life in the universe, originally developed by Freeman Dyson [30] and Frank Tipler [4, 72]. They proposed two different scenarios where life (interpreted as information processing systems) can survive forever by adapting to and modify the universe (the Final Anthropic Principle [4]), as well as perform an infinite number of calculations and store a diverging amount of information

In the Dyson scenario the universe is open and gradually cooling, but life survives by expending less and less energy and working more and more slowly.

In the Tipler scenario the universe is closed but experiences an anisotropic collapse; the resulting usable shear energy grows faster than temperature and allows a divergence of information processing. The subjective time (as measured by the growth of information) of life goes to infinity as the final singularity is approached. While in the Dyson scenario life trades time for energy, in the Tipler scenario life trades energy for time.

Tipler derives an inequality from the Brillouin inequality (equation 11) that determines the amount of information that can be processed between now and the c-boundary of spacetime at t_Ω:

$$I = \int_{t_{now}}^{t_\Omega} \left(\frac{dI}{dt} \right) dt \leq \frac{1}{k_B \ln 2} \int_{t_{now}}^{t_\Omega} T(t)^{-1} \left(\frac{dE}{dt} \right) dt \qquad (37)$$

In open of flat spacetimes the right hand side can diverge even if the total energy used between t_{now} and t_Ω is finite, since $T(t) \to 0$. In closed universes $T(t) \propto 1/R(t) \propto t^{-\frac{1}{3}}$ where $R(t)$ is the universal scale factor. In this case more and more energy has to be used to counteract the rising heat. Tipler calculates that the total available shear energy in late phases of the collapse will be $E \propto t^{-1}$, giving $dE/dt \propto t^{-2}$ and

$$I \leq c \int t^{-2} t^{\frac{1}{3}} dt \propto t^{-\frac{2}{3}} \qquad (38)$$

which diverges as $t \to 0$. This shows that there is theoretically enough energy for an infinite amount of computation in both open, flat and closed shear-dominated anisotropic universes. Neither of the scenarios as published takes reversible computation into account, so there is even a slight affordance in cold universes.

Dyson doesn't assume life will spread beyond a finite region, which limits the available number of possible states; in the long run the system inside the region will necessarily begin to repeat previous states (which Tipler [72] calls "The Eternal Return"). In order to develop indefinitely and without repetition life must expand the region it inhabits in phase space faster than $\log(t)$ (this is

---

[8] by receiving the answer and then doing the calculation at some future time

the number of bits necessary to encode a simple counter counting upwards each tick). In the Tipler scenario the Bekenstein bound is circumvented by having the energy used in information storage increase faster than the radius of the universe shrinks, allowing stored information to diverge as $t \rightarrow t_\Omega$.

In the Dyson scenario, memory appears to be a problematic constraint. How to create more memory as it is needed while minimizing energy usage is a serious problem for indefinite survival. Dyson proposes the use of analog memory, but as we have seen the Bekenstein Bound is still a limit if life is confined inside a finite region; again, life must spread in order to be able to avoid the Eternal Return. At the same time, long range communications in an open universe require energy since the redshift lowers the energy of signals sent between different systems; the cost of sending a signal seems to rise exponentially. Balancing these two constraints against each other may be very hard, although Dyson calculates that by slowing down the exchange of signals enough indefinite communication can be upheld.

In the Tipler scenario information is stored in more and more energetic particles in order to prevent it from being lost in the thermal noise. This means that if the density of particle states is too low, information cannot grow indefinitely, and if it is too high the shear energy will be dampened by the spontaneous creation of high energy particles. This means that dN/dE, where N is the number of particle states, diverges, but is asymptotically bounded by E^2. This is a testable prediction of the scenario.

One of the major problems for the Tipler scenario is how to avoid the creation of horizons that limit communications and to tap the shear energy of spacetime. He points out that general relativity allows chaotic solutions, and by carefully balancing the collapse life can manipulate it to avoid horizon formation and maximize shear energy production, but exactly how this is to be done in practice is left open. That megascale manipulation of matter is possible has been shown by Dyson [30], but that the necessary planning and coordination required can be achieved across the universe remains to be seen.

Another, more severe problem, is the question whether quantum gravity will prevent the indefinite collapse assumed in the above calculations. If the collapse rebounds or changes characteristics at the Planck scale or beyond, then the scenario will have to be revised and may be entirely impossible.

Ćircović and Bostrom [21] point out that observational results suggesting the existence of a large positive cosmological constant may be the biggest threat to this kind of scenarios. If the cosmological constant is large, then the universe will eventually enter an era of exponential expansion which will likely make any form of organized structure impossible in the long run due to the approach of horizons and increase of entropy. The only way of upholding the Final Anthropic Principle in this case is to exploit the 'Linde Scenario' [48]: if chaotic inflation occurs, then there should exist an infinite number of different inflation domains where life could exist, and intelligent life could in principle escape from a dying domain to another using wormholes or possibly by creating "baby universes" [12, 22]. Tipler [72] has criticized this escape scenario on the grounds that the amount of information that can be transferred between each domain or into the baby universe is finite, which means it cannot grow indefinitely as needed to avoid the Eternal Return.

These eschatological scenarios deal with what appears to be the ultimate problems of long-lived, very advanced civilizations. At present they remain mainly in the domain of speculation, mostly because they just describe necessary conditions on cosmological models to sustain life indefinitely; as our knowledge of cosmology advances, it is very likely that these models can be significantly refined and possibly even expressed as proofs of construction rather than necessary conditions.

9   Conclusions

As we have seen, the known laws of physics allow extremely powerful information processing systems. Whether these may be built in practice is of course unknown, but many of them appear within the reach of an advanced technological civilization [26, 30]. Within these very broad limits tremendous diversity is possible: practically all combinations of compact or distributed, hot or cold, tiny or gigantic, fast or slow systems appear possible in principle.

Their exact uses cannot be predicted, since they will be highly dependent on the needs or desires of their builders, who will most likely themselves be posthuman (or postalien) beings. They appear to provide an ideal environment for intelligent life to develop and diversify in.

In this review constraints due to software or the mathematics of information processing has not been considered much, just direct computing power. In a full treatment of the subject these must be included, both fundamental limits such as limits of computability and practical limits to learning, algorithmic complexity and software reliability. They might provide more fundamental or insurmountable barriers to information processing than physics, and make some forms of information processing superobjects uneconomical even if they are physically feasible.

It is my hope that this preliminary review of the subject will stimulate others to more detailed and stringent studies of the technological limits of thought.

9.1   Acknowledgments

10   Appendix: Three Superobjects

10.1   Jupiter brain: "Zeus"

Zeus is a $r = 9000$ km sphere of nearly solid diamondoid, consisting mainly of reversible quantum dot circuits and molecular storage systems. Surrounding the central sphere is a concentric shield protecting it from radiation and holding radiators to dissipate heat into space. Energy is provided by fusion reactors distributed outside the shield.

The total mass is around 10^{25} kg, 1.8 times the Earth's mass. This amount of carbon could be gathered from gas giant cores or through star lifting [23].

The long-range internal connections are optical, using fiber optics/waveguides or directional signals sent through vacuum. Each node is a processing element and/or memory storage system assumed to act as a semi-independent unit; if we insert lightspeed, a switching speed on the order of picoseconds and a desired $S = 1$ into equation 24 we get a diameter of $3 \cdot 10^{-4}$ m for the nodes.

The connectivity between the nodes is assumed to be a "small-world" network structure [75] which allows a sparse connectivity where two arbitrary chosen nodes will be connected by a short series of intermediary links. Each processing node (of which there are N) will have a number of links $k \gg \log(N)$ to other nodes, mostly connected to neighboring nodes but a few ($p \approx 0.01$) to remote nodes. The total number of links is $kN/2$, but only $pkN/2$ will be long-range and take up significant space. Assuming communications links of constant cross section a and length $\propto r$ and a number of nodes $\propto r^3$, the total volume of communications in the system is $V_{comm} \propto pkr^4$. This means that even for very low k the system will be dominated by the volume of communications links.

Since the processing/memory nodes need to be close to each other due to the many short-range connections, the possible distributions are either a central core surrounded by connections, a cortex with connections through the interior, or distributed clusters of nodes in the interior. Of these the cortex model is most volume-efficient, and it will be assumed that Zeus is organized in the same way. Incidentally, it also provides the simplest solution for cooling if we assume that mainly the nodes produce waste heat.

The average distance between two random points in a sphere is $(36/35)r$ and the average chord between two points on the surface $1.33r$ [67], producing a total volume of long-distance connections of cross-section area $V_{comm} = p(1.33ra)kN/2$. The total volume of nodes is $V_{node} = vN$

Adding some numbers, $a = 10^{-14}$ m^2, $p = 0.01$, $k = 100$, $v = 2.7 \cdot 10^{-11}$ m^3 we get a ratio between connection and node volume as ≈ 2200, i.e. a very thin layer (1350 meters) of nodes on top of an interior filled with connections. There are $5 \cdot 10^{37}$ nodes in Zeus; if they all were storage only, the capacity would be up to 10^{47} bits. The number of "operations" per second if they were all single processors would be up to 10^{49} (since the diameter is much larger than the nanocomputers discussed in [27] they would most likely be clusters of parallel processors, and this estimate is up to 12 orders of magnitude too low).

Total S would be on the order of $4 \cdot 10^{10}$ for the nodes, showing that latencies for remote information are very large compared to their speed. For the whole

system the characteristic timescale is $r/c = 0.03$ s.

If the whole structure holds a temperature of 4 Kelvin, then it can radiate away 10^{10} W into the surrounding 3 Kelvin universe. This energy corresponds to $2.6 \cdot 10^{32}$ bit erasures per second. Thermal errors due to equation 12 at this temperature are close to zero as long as the energy barriers are larger than $10^{-21} J$. This suggests that the major source of errors will be nonthermal errors such as cosmic rays; these can in principle be kept very low through the shielding (except for neutrinos; however they do not appear to contribute a significant amount of dissipation due to their very low interaction).

A major contribution to energy dissipation will likely be communication. Equation 32 shows that a sizeable amount of energy needs to circulate in the system just to enable communication. Dissipation will likely be a few orders of magnitude less, but even changes in bandwidth usage would require buffering of communications energy. The effective noise temperature in the communications links is likely $\ll$ 4K, but since there are $2.5 \cdot 10^{39}$ links the maximal temperature-bandwidth product (with 10^{10} W budget) becomes $2.9 \cdot 10^{-7}$ K bits/second. If the bandwidth per channel is a modest 10^6 bits/s, the noise temperature has to be less than 10^{-13} K. At a given dissipation power, there is a trade-off between bit erasure and communication dissipation.

Taken together, Zeus is an example of a "cold" information processing superobject. The main limiting factors in this design is the availability of carbon, the material strength of diamond and the need of keeping the system cool.

10.2 Dyson brain: "Uranos"

Uranos gradually emerged when the matter of a solar system was converted by intelligent life into a Dyson sphere surrounding its sun-like star at a distance of 1 AU. It consists of numerous independently orbiting structures, ranging from large (hundreds of kilometers) solar collectors to microscale devices moving between the structures for repair and adjustment.

The efficiency of converting solar energy to work is around 30%, giving $3 \cdot 10^{25}$ Watt of available energy. The working temperature for an unshielded object in an 1 AU orbit is 395K. The number of bit-erasures that can be achieved under these conditions is $7.9 \cdot 10^{45}$ bits/second.

The total amount of matter available in the solar system (disregarding hydrogen and helium) beside the sun is $\approx 1.7 \cdot 10^{26}$ kg [66]. If the energy collecting system is assumed to hold a fairly minor fraction (1%) of the total mass, and assuming molecular densities, then Uranos can contain up to 10^{52} bits. Assuming processing nodes of the same type as Zeus, we get 10^{39} nodes and 10^{51} operations per second.

The internal delays between distant nodes are on average 660 seconds. Assuming the same picosecond switching as in Zeus gives $S \approx 6.7 \cdot 10^{14}$, suggesting even less synchronization than Zeus.

Where Uranos really outperforms Zeus is information production/destruction; the high energy throughput makes it possible to dissipate 10^{22} times as many bits as Zeus. It might make sense to keep Zeus-like structures in orbit outside

Uranos to act as information repositories and the Dyson shell itself for processing.

The main limitation of Uranos is the availability of matter, and the amount of energy that can be extracted from the sun.

10.3   Neutronium brain: "Chronos"

Chronos was originally created by the carefully orchestrated collapse of a globular cluster. By manipulating the orbits of the stars and organizing close encounters half of the stars were ejected from the cluster and the other half dropped into the core. During this process star lifting was used [23] to redistribute mass in order to produce a maximal amount of iron. The iron was merged with a central neutron star kept stable by strong mass flows. Iron was used to avoid inducing fusion reactions, and later moved inside where neutron drip and eventual conversion to quark matter occured. Energy for the merging process and cooling was supplied by the matter-energy conversion in a series of micro-black holes surrounded by Dyson spheres. The result is an extremely massive body delicately balanced between gravity and rotation, surrounded by a huge system of support systems.

Note that without stabilizing the system using large amounts of angular momentum, just combining all available mass into a single system does not maximize information processing; in order to avoid gravitational collapse the density has to decrease as more matter is added, and beyond a certain point the desirable nuclear densities are no longer available. On the other hand, if communications delays are acceptable, spreading out the mass into a number of neutronium spheres orbiting each other would enable better energy dissipation and hence a higher bit erasure rate.

Assuming an original mass of 10^{36} kg, half of it ends up in the core. Seeking a density of 10^{20} kg/m^3 produces a 100 kilometer sphere of quark matter. This corresponds to a maximum $5 \cdot 10^{61}$ bits of potential storage capacity, although in practice only part of it is available due to the need of using a significant amount of the mass for support, communications and processing.

Given the timescale of nuclear reactions, Chronos would be able to perform on the order of 10^{85} operations per second. The S value becomes $3 \cdot 10^{20}$, making it even more dispersed than Zeus and Uranos despite its smaller size. Subunits with $S = 1$ would be just $3 \cdot 10^{-16}$ meters across. These subunits roughly correspond to one or a few bits of storage each, so rather than processor clusters as in Zeus and Uranos they would likely be individual processors in Chronos.

Since nuclear bonds are stable up to around 10^9 K, Chronos can operate in the high temperature region. If it dissipates energy through blackbody radiation, it could have a power on the order of 10^{39} W, similar to a quasar. This would correspond to $7 \cdot 10^{53}$ bit erasures per second. However, other sources of dissipation would be communication and maintenance of the momentum flows keeping the system stable; since these flows would be highly relativistic dissipation losses would likely themselves have a noticeable mass-equivalent.

The major limitations of Chronos is the initial amount of mass in the globular cluster and the strength of nuclear bonds. Chronos outperforms Zeus and Ura-

nos, but the performance might not be worth the cost. The energy demands are extreme, corresponding to a swarm of 10^9 1 kg black holes converting matter to energy; the remaining mass of the globular cluster would be exhausted after a million years. A smaller sphere of nuclear matter able to support itself would have a better efficiency in converting power into computation.

References

[1] AUDOUZE, J., AND ISRAEL, G. *The Cambridge Atlas of Astronomy*. Cambridge University Press, 1985.

[2] AVIRAM, A., CLEMENTI, E., DUPUIS, M., AND FARAZDEL, A. *J. Am. Chem. Soc.* (1990).

[3] BARNUM, H., NIELSEN, M., AND SCHUMACHER, B. Information transmission through a noisy quantum channel. http://xxx.lanl.gov/abs/quant-ph/9702049, 1997.

[4] BARROW, J., AND TIPLER, F. *The Anthropic Cosmological Principle*. Oxford University Press, 1988.

[5] BEKENSTEIN, J. *Phys Rev D 23 287* (1981).

[6] BENNETT, C. The thermodynamics of computation – a review. *Int J. Theoret. Phys 21* (1982), 905–940.

[7] BENNETT, C. Time/space trade-offs for reversible computation. *SIAM J. Comput. 18*, 4 (August 1990), 766–776.

[8] BENNETT, C., BERNSTEIN, B., BRASSARD, G., AND VAZIRANI, U. Strengths and weaknesses of quantum computing, 1994. quant-ph/9701001, http://feynman.stanford.edu/qcomp/bennett-strengths/index.html.

[9] BIRGE, R. Protein-based computers. *Scientific American 272* (March 1995). 66–71.

[10] BOCQUET, M., BONAZZOLA, S., GOURGOULHON, E., AND NOVAK, J. Rotating neutron star models with magnetic field. *Astron. astrophys. 301*, 757 (1995).

[11] BOGHOSIAN, B., AND TAYLOR, W. Simulating quantum mechanics on a quantum computer. Expanded version of a talk given by WT at the PhysComp '96 conference, BU, Boston MA, November 1996, quant-ph/9701019, 1996. http://babbage.sissa.it/abs/quant-ph/9701019.

[12] BORDE, A., TRODDEN, M., AND VACHASPATI, T. Creation and structure of baby universes in monopole collisions. *Phys.Rev. D 59* (1999). http://xxx.lanl.gov/abs/gr-qc/9808069.

[13] BRAY, D. Protein molecules as computational elements in living cells. *Nature 376* (27 July 1995), 307–312.

[14] BUTTLER, W., HUGHES, R., KWIAT, P., LAMOREAUX, S., LUTHER, G., MORGAN, G., NORDHOLT, J., PETERSON, C., AND SIMMONS, C. Practical free-space quantum key distribution over 1 km. *Physical Review Letters 81*, 15 (1998), 3283–3286.

[15] CARLINI, A., AND NOVIKOV, I. Time machines and the principle of self-consistency as a consequence of the principle of stationary action (ii): the cauchy problem for a self-interacting relativistic particle. *Int.J.Mod.Phys D5 445* (1996). gr-qc/9607063 http://xxx.lanl.gov/abs/gr-qc/9607063.

[16] CARROLL, B. W., AND OSTLIE, D. A. *An Introduction to Modern Astrophysics*. Addison Wesley, 1996.

[17] CAVES, C., AND DRUMMOND, P. Quantum limits on bosonic communication rates. *Rev. Modern. Phys. 66*, 2 (1994), 481–537.

[18] C.H. BENNETT, F. BESSETTE, G. B. Experimental quantum cryptography. *J. Cryptography 5*, 3 (1992).

[19] CHAU, H., AND LO, H. Why quantum bit commitment and ideal quantum coin tossing are impossible. PhysComp96, quant-ph/9605026, 1996. http://babbage.sissa.it/abs/quant-ph/9605026.

[20] CHUANG, I., AND GERSHENFELD, N. Bulk spin-resonance quantum computation. *Science 275* (17 Jan 1997).

[21] ĆIRKOVIĆ, M., AND BOSTROM, N. Cosmological constant and the final anthropic hypothesis. submitted to Am. J. Phys, 1999. http://xxx.lanl.gov/abs/gr-qc/9906042.

[22] CRANE, L. Possible implications of the quantum theory of gravity. hep-th/9402104, 1994. http://xxx.lanl.gov/abs/hep-th/9402104.

[23] CRISWELL, D. Solar system industrialization: Implications for interstellar migration. In *Interstellar Migration and the Human Experience* (1985), B. Finney and E. Jones, Eds., University of California Press, pp. 50–87.

[24] DIVINCENZO, D. Two-bit gates are universal for quantum computation. *Phys. Rev. A 51* (1995), 1015. http://xxx.lanl.gov/abs/cond-mat/9407022.

[25] DIVINCENZO, D., AND SHOR, P. Fault-tolerant error correction with efficient quantum codes. quant-ph/9605031, 1996. http://babbage.sissa.it/abs/quant-ph/9605031.

[26] DREXLER, K. *Engines of Creation*. Anchor Books, 1986. http://www.foresight.org/EOC/index.html.

[27] DREXLER, K. *Nanosystems: Molecular Machinery, Manufacturing, and Computation*. Wiley, 1992.

[28] DYSON, F. Search for artificial stellar sources of infrared radiation. *Science 131* (1959), 1667–1668.

[29] DYSON, F. *Perspectives in Modern Physics (Essays in honor of Hans Bethe)*. John Wiley & Sons, NY., 1966, ch. The Search for Extraterrestrial Technology.

[30] Dyson, F. Time without end: physics and biology in an open universe. *Reviews of Modern Physics 51* (1979), 447–60.

[31] Ekert, A. Quantum computation, 1993. http://feynman.stanford.edu/qcomp/ekert/index.html.

[32] Fredkin, E., and Toffoli, T. Conservative logic. *Int. J. Theoret. Phys. 21* (1982), 219–353.

[33] Grover, L. A fast quantum mechanical algorithm for database search. In *Proceedings of 28th Annual ACM Symposium on the Theory of Computing (STOC)* (May 1996), pp. 212–219. http://babbage.sissa.it/abs/quant-ph/9605043.

[34] Han, J., Globus, A., Jaffe, R., and Deardorff, G. Molecular dynamics simulations of carbon nanotube-based gears. *Nanotechnology 8* (1997), 95–102.

[35] Hanson, R. Reversible agents: Need robots waste bits to see, talk and achieve? In *Proceedings for Workshop on Physics and Computation* (October 1992). http://www.hss.caltech.edu/hanson/reverse.html.

[36] Hawking, S. Chronology protection conjecture. *Physical Review D 46*, 2 (15 July 1992), 603.

[37] Heiselberg, H. Quark matter structure in neutron stars. Talk given at Intl. Symposium on Strangeness and Quark Matter, Chania, Greece sept 1-5 1994 HEP-PH-9501359, 1994. http://xxx.lanl.gov/abs/hep-ph/9501359.

[38] Hyde, R. Earthbreak: Earth to space transportation. *Defense Science 2003+ 4*, 4 (1985), 78–92.

[39] Khatsymovsky, V. Towards possibility of self-maintained vacuum traversible wormhole. Preprint gr-qc/9612013, December 1996. http://xxx.lanl.gov/abs/gr-qc/9612013.

[40] Knill, E., Laflamme, R., Martinez, R., and Tseng, C.-H. A cat-state benchmark on a seven bit quantum computer. http://xxx.lanl.gov/abs/quant-ph/9908051, 15 Aug 1999.

[41] Kristin Schleich, D. M. W. Topological censorship. In *Proceedings of the Lake Louise Winter Institute, Particle Physics and Cosmology* (Feb. 20-26 1994), World Scientific. http://xxx.lanl.gov/abs/gr-qc/9903061.

[42] Lachmann, M., Newman, M. E. J., and Moore, C. The physical limits of communication. Santa Fe Institute working paper 99-07-054, 1999.

[43] Landauer, R. Zig-zag path to understanding. In *Proceedings of the Workshop on Physics and Computation PhysComp'94* (1994), IEEE Computer Soc. Press, pp. 54–59.

[44] Lent, C., Tougaw, P., Porod, W., and Bernstein, G. Quantum cellular automata. *Nanotechnology 4*, 1 (January 1993), 49–57.

[45] LEVITIN, L. Energy cost of information transmission. *Physica D 120* (1998), 162–167.

[46] LIBET, B. Unconscious cerebral initiative and the role of conscious will in voluntary action. *The Behaviorial and Brain Sciences 8* (1985), 529–566.

[47] LIBET, B., JR., E. W., FEINSTEIN, B., AND PEARL, D. Subjective referral of the timing for a conscious sensory experience. *Brain 102* (1979), 193–224.

[48] LINDE, A. *Phys.Lett. B 211*, 29 (1988).

[49] LLOYD, S. The capacity of the noisy quantum channel. http://xxx.lanl.gov/abs/quant-ph/9604015, 1996.

[50] LLOYD, S. A technologically feasible quantum computer. http://feynman.stanford.edu/qcomp/lloyd/index.html, 1996.

[51] LLOYD, S. Ultimate physical limits to computation. Submitted to Nature. http://xxx.lanl.gov/abs/quant-ph/9908043, 1999.

[52] LOFSTROM, K. The launch loop – a low cost earth-to-high orbit launch system. AIAA Paper 85-1368, 1985.

[53] MARGOLUS, N., AND LEVITIN, L. B. The maximum speed of dynamical evolution. *Physica D 120* (1998), 188–195.

[54] MERKLE, R. Reversible electronic logic using switches. *Nanotechnology 4* (1993), 21–40.

[55] MERKLE, R. Two types of mechanical reversible logic. *Nanotechnology 4* (1993), 114–131.

[56] MII, Y., RISHTON, S., TAUR, Y., KERN, D., LII, T., LEE, K., JENKINS, K., QUINLAN, D., JR., T. B., DANNER, D., SEWELL, F., AND POLCARI, M. Experimental high performance sub-0.1 /spl mu/m channel nmosfet's. *IEEE Electron Device Letters 15* (1994), 28.

[57] MISNER, C., THORNE, K., AND WHEELER, J. *Gravitation.* W. Freeman & Co, 1973.

[58] MORAVEC, H. Pigs in cyberspace. http://www.aleph.se/Trans/Global/Posthumanity/pigs.txt (included in [60]).

[59] MORAVEC, H. *Mind Children: The Future of Robot and Human Intelligence.* Harvard, 1988.

[60] MORAVEC, H. *Robot : Mere Machine to Transcendent Mind.* Oxford Univ Pr., 1998.

[61] MORRIS, M., AND THORNE, K. Wormholes in spacetime and their use for interstellar travel. *Amer. J. Phys 56* (1988), 395.

[62] PRICE, M. Traversable wormholes: Some implications. *Extropy 5*, 1 (1993). http://www.aleph.se/Trans/Tech/Space-Time/wormholes.html.

[63] PUDLÉK, P. Complexity theory and genetics. In *Proceedings of the 9th Conference on Structure in Complexity Theory* (1994), pp. 183–195.

[64] ROBERTSON, D. *The New Renaissance: Computers and the Next Level of Civilization.* Oxford Univ. Pr., 1998.

[65] RUFFINI, R. On the dyadosphere of black holes. In *Proceedings of Yamada Conference* (Kyoto, Japan, April 1998). http://xxx.lanl.gov/abs/astro-ph/9811232.

[66] SANDBERG, A. Dyson sphere faq. http://www.student.nada.kth.se/ãsa/dysonFAQ.html, 1996.

[67] SANTALÓ, L. *Integral Geometry and Geometric Probability.* Addison-Wesley, 1976.

[68] SCHEIN, F., AND AICHELBURG, P. Traversable wormholes in geometries of charged shells. *Phys.Rev.Lett. 77*, 4130 (1996). http://xxx.lanl.gov/abs/gr-qc/9606069.

[69] SHOR, P. Algorithms for quantum computation: Discrete log and factoring. FOCS, 1994.

[70] SHOR, P., AND SMOLIN, J. Quantum error-correcting codes need not completely reveal the error syndrome. http://xxx.lanl.gov/abs/quant-ph/9604006, 1996.

[71] SIMON, D. On the power of quantum computation. FOCS 1994, http://feynman.stanford.edu/qcomp/simon/index.html, 1994.

[72] TIPLER, F. *The Physics of Immortality.* Macmillan, London, 1994.

[73] VINGE, V. What is the singularity? presented at the VISION-21 Symposium, March 30-31, 1993., 1993. http://www.student.nada.kth.se/ñv89-nun/offloading/vinge.

[74] VISSER, M. From wormholes to time machines: Remarks on hawking's chronology protection conjecture. *Physical Review D 47*, 2 (15 Jan 1993), 554.

[75] WATTS, D., AND STROGATZ, S. Collective dynamics of 'small-world' networks. *Nature 393*, 6684 (1998), 440–442.

[76] WEBER, F., AND GLENDENNING, N. Dense stellar matter and structure of neutron stars. Presented by F. Weber at the 3rd Mario Schoenberg School on Physics, Joao Pessoa, Brazil, July 22 – August 2, 1996. To be published in the Brazilian Journal of Teaching Physics, 1996. http://xxx.lanl.gov/abs/astro-ph/9609074.

[77] WITTEN, E. *Phys. Rev. D* (1984).