Matrioshka Brains

by Robert J. Bradbury

This document is a work in progress. It has not been peer reviewed.
Believe it or reference it only if you can convince yourself of its accuracy.

Abstract

Predictable improvements in lithographic methods foretell continued increases in computer processing power. Economic growth and engineering evolution continue to increase the size of objects which can be manufactured and power that can be controlled by humans. Neuroscience is gradually dissecting the components and functions of the structures in the brain. Advances in computer science and programming methodologies are increasingly able to emulate aspects of human intelligence. Continued progress in these areas leads to a convergence which results in megascale superintelligent thought machines. These machines, referred to as Matrioshka Brains[1], consume the entire power output of stars (~10^{26} W), consume all of the useful construction material of a solar system (~10^{26} kg), have thought capacities limited by the physics of the universe and are essentially immortal.

A common practice encountered in literature discussing the search for extraterrestrial life is the perspective of assuming and applying human characteristics and interests to alien species. Authors limit themselves by assuming the technologies available to aliens are substantially similar or only somewhat greater than those we currently possess. These mistakes bias their conclusions, preventing us from recognizing signs of alien intelligence when we see it. They also misdirect our efforts in searching for such intelligence. We should start with the laws on which our particular universe operates and the limits they impose on us. Projections should be made to determine the rate at which intelligent civilizations, such as ours, approach the limits imposed by these laws. Using these time horizons, laws and limits, we may be better able to construct an image of what alien intelligence may be like and how we ourselves may evolve.

Contents

Overview

Matrioshka Brains (MB)[1] are megascale computers constructed out of microelectronic and/or nanoscale components. MB may be constructed as shells around a star (internally powered Matrioshka Brains: IPMB) or may be constructed independently of a star if large amounts of power are harvested from stars in other locations and beamed to the MB (externally powered Matrioshka Brains: EPMB). A third variant, (self-powered Matrioshka Brains: SPMB) which generates power via controlled nuclear fusion or matter/antimatter reactions, is possible.

The two pillars on which the MB arch rests are the extensions of current engineering trends to the largest and smallest scales. At the largest scale, in their initial stages, MB are limited by the mass and energy provided by individual solar systems. At the smallest scale, MB are limited by our ability to assemble materials atom by atom. The terms megascale engineering and molecular nanotechnology are generally used to discuss these different perspectives. The union of construction methods at the small and large scale limits allows the optimal use of locally available energy and matter and is the distinguishing feature of Matrioshka Brains.

Megascale engineering has its roots in science fiction. One of the first scientific examinations of megascale engineering was done by mathematician Freeman Dyson (1960) in which he discussed dismantling Jupiter to construct a shell around the sun to harvest all of its energy and provide a biosphere capable of supporting large numbers of people. Writer Larry Niven addressed some of the problems of gravity in Dyson shells by changing the form of the biosphere from a shell to a rotating Niven Ring. Other examples of megascale engineering exist in fictional literature but these are the most relevant for the discussion of MB.

Nanoscale engineering was first discussed by Richard Feynman in 1959. These ideas were extended by Eric Drexler in his 1981 PNAS paper and *Engines of Creation*. Much of the engineering basis for nanotechnology is documented in *Nanosystems*. Progress in the development of nanotechnology continues and no serious challenges against its ideas have been produced in the last ten years (Merkle, 1998). Estimates of its full scale development and deployment range from 10 to 30 years in the future.

Megascale and nanoscale engineering currently do not exist. Megascale engineering results in the progression of trends in the engineering of large scale structures such as pyramids, oil tankers, suspension bridges, tunnels, sky-scrapers and rockets. Nanoscale engineering results from trend progressions in microelectronic lithographies, micromachining, microvolume and combinatorial chemistry, biotechnology manipulation of genes and proteins, robotics and computer science.

It is paradoxical that many people more easily envision megascale engineering than nanoscale engineering. The most logical explanation for this is that our senses are able to directly interact with megascale structures, while intermediaries such as atomic force microscopes or enzymes are required to sense and manipulate things at the nanoscale level. It is important to remember that atomic scale pumps, motors, engines, power generation apparatus and molecular manipulators (enzymes) exist in every individual reading this document. By mid-1998, the complete genomic DNA sequences (nanoscale programs) for more than 30 different bacteria and yeast (nanoscale assembly and replication machines) were known. Nanoscale technology exists and is rapidly being domesticated by humans.

As has been pointed out by Dyson [1960, 1968], Kardashev [1985,1988,1997], Berry [1974], and Criswell [1985], the progression of existing population and economic growth, power and mass management trends in our society will enable the construction of Matrioshka Brains using existing (non-nanoscale) technologies within at most a few thousand years. *Nanoscale assembly per se is not required*. Current trends in silicon wafer production, if continued, would allow the production of sufficient microprocessors, of current primitive designs, to create a MB by 2250. It would however require most of the silicon in the planet Venus as raw material. A MB built from such processors would have capabilities significantly less than the limits which would be available using nanoscale fabrication. Even so, a computing machine built out of even these primitive components would have a thought capacity in excess of *a million times* the thought capacity of the 6 billion+ people now populating the planet! A small fraction of this thought capacity devoted to extending engineering methods should in a brief period develop nanoengineering and assembly to its ultimate limits.

Background

Computer Trends and Characteristics

To discuss computational characteristics of a MB it is necessary to understand the evolution of computers. This topic is too complex to be discussed in great detail in this paper. In general however, we may assume that current trends in lithography (optical lithography down to 0.08 μm) and computer architecture modifications such as processor-in-memory (PIM), intelligent-RAM (IRAM), content-addressable memory (CAM), etc. should provide approximate human-brain equivalent computational capacity in desktop machines sometime between 2005-2010.

Lithographic methods will continue to improve, transitioning from optical to EUV, to X-ray, to e-beam or nano-imprint (soft lithography), each at smaller levels of resolution. This will culminate in nanoassembly with the ability to manipulate individual atoms. If historic trends were to continue, atomic scale manipulation would be reached by 2050. In the last 5 years however, the introduction of decreased lithographic scales has been accelerating [SIA, 1997]. The formation of a company and continuing work towards the goal of producing a nanoassembler (at Zyvex) and prognostications on nanotechnology development trends confirming earlier projections [Drexler, 1998], provide reasons to believe that nanoassembly may become possible in the 2010-2015 time frame. In fact, Jim Von Ehr, the president of Zyvex LLC, has publicly stated that he believes that Zyvex will be able to do diamondoid nanoassembly by 2010.

Lithographic technologies enable the construction of very powerful computers based on a two-dimensional technology. Systems assembly methods using SIMMs and processor cards (Slot-1) effectively convert 2-D chips into 3-D systems. The addition of optical interconnects (Emcore, Opticomp & others) and high capacity cooling (Beech *et. al.*, 1992; Tuckerman, 1984, SDL Inc.) allow significant increases in communication bandwidth and processing density. Nanotechnology enables the construction 3-D computers which allow the computational, communication, power production and delivery and cooling elements to be tightly integrated into a low cost package using a single uniform assembly process. The development of nanotechnology will be a natural development once the limits of conventional manufacturing and assembly processes are reached. There is no known process that allows efficiencies and capabilities greater than those offered by nanotechnology. It is reasonable to assume that nanotechnology and nanoassembly represents a significant plateau in the development of technological civilizations.

In *Nanosystems*, Drexler outlined the details of a rod-logic computer (essentially a nanoscale abacus). A single rod-logic nanoCPU is a very small computer which consumes very little power with very little capacity. NanoCPUs can be assembled into parallel systems (midi-Nanocomputers) which achieve the processing capacity of current microprocessors at significantly lower power consumption. Further aggregation results in a Mega-Nanocomputer that consumes 100,000 W (~10^4 times greater than a human brain) in a volume of 1 cm^3 (~10^3 times less than a human brain). The high speed, massive parallelism and reduced propagation delays in a Mega-Nanocomputer should result in computational throughput 10^6-10^7 times greater than the human brain.

Merkle and Drexler have also developed helical logic which requires nanoassembly methods to create computers based on the control of the movement of single electrons. The limits on computation are dictated by the size of the computational elements and the heat production associated with the computation. We may assume that the manipulation of single electrons and the use of reversible logic (such as in rod and helical logic) bring us close to the possible limits of computation. These topics are explored in much greater depth in Merkle & Drexler, 1996, Sandberg, 1997 and Frank & Knight, 1998.

Since the details of rod-logic computers (power consumption, size, computational capacity, etc.) are the best defined for Mega-nanocomputers, they will be used for our discussion. Beyond rod-logic, helical logic allows an improvement of 10^{11} in power consumption per operation. Theoretical limits potentially allow improvements of 10^9 in cooling capacities (power density) and 10^4 in operating frequencies. If computers could be produced at these limits, computational capacities from 10^{10} to 10^{20} greater than those presented in this paper may be possible.

Table 1 details the characteristics of some of these computer architectures.

Table 1. Computer Characteristics

Processor Type	Switching speed	Clock Rate	OPS	Power per CPU	Power per logic Op	Mass	Volume	Source
	sec	GHz	sec^{-1}	W	J/Op	kg	m^3	
Circa Y2000 microprocessor (e.g. Merced)	6×10^{-8}	1	1×10^9	50	$\sim 10^{-8}$	0.1	10^{-7}	Intel, Byte
Rod-logic NanoCPU	10^{-10}	10	1×10^9	10^{-7}	10^{-16}	1.6×10^{-17}	10^{-20}	*Nanosystems*
Rod-logic Midi-Nanocomputer	10^{-10}	10	2×10^{17}	10	10^{-16}	2.7×10^{-9}	10^{-11}	*Nanosystems*
Rod-logic Mega-Nanocomputer	10^{-10}	10	2×10^{21}	10^5	10^{-16}	2.7×10^{-5}	10^{-6}	*Nanosystems*, pg. 370
Helical-logic computer	10^{-10}	10			10^{-27}			Merkle & Drexler, 1996, Drexler, 1992
Physical Limits	10^{-14}	10,000		10^{14}				Merkle & Drexler, 1996

Computer and Human Operations Equivalence.

At the simplest level of abstraction, neurons can be considered to be multiplication and adding machines. Neurons multiply the "strength" of a synaptic connection times the "weight" of an incoming signal and sum these values across a number of input synapses. If the result exceeds a certain threshold, the neuron fires and transmits a signal to other neurons connected to its network. Neurons fire very slowly, < 100 times per second. The immense power found in the human brain is due to neuron features other than speed. These include their small size, low power consumption, high interconnection levels (100-10,000 per neuron) and to a large degree sheer numbers. The human neocortex, which is the most highly developed portion of the human brain, and that part which is thought to be responsible for "higher thought", contains ~21 billion (2.1×10^{10}) neurons [Pakkenberg, 1997]. The total number of neurons in the brain is less certain, but since the neocortex contains roughly 1/3 of the brain volume, unless neurons density is much higher in other brain regions, extrapolations from Pakkenberg's data would imply there is a total of 60 billion (6×10^{10}) neurons in the brain. To provide a proper perspective, if current SIA projected trends continue, microprocessors would not have 60 billion transistors circa 2025. Even then, a single transistor does not possess the computational capacity of a neuron. On the other side of the coin, a microprocessor with 60 billion transistors would occupy a volume much smaller than that of the human brain.

If we assume 6×10^{10} neurons × 5×10^{1} firings per second × 10^{3} operations per neuron firing , we end up with a result of 3×10^{15} operations per second (300 Trillion operations per second or 300 TeraOps). This is likely to be at the high end of possible computational capacities since it is assuming that all neurons are being used simultaneously. This is unlikely to be true since the brain clearly has specialized structures for visual, auditory and odor input; speech output; physical sensation and control; memory storage and recall; language analysis and comprehension; and left-right brain communication. It is unlikely that all of these structures will be optimally utilized at any point in time.

A high end estimate of 300 TeraOps for human thought capacity does not significantly differ from those found in the literature as outlined in Table 2.

Table 2. Human Thought Capacity

Brain Capacity	Method	Source
10^{13} calculations per second 10^{14} bits / second	Algorithmic equivalence	Moravec (1987)
10^{14} instructions per second	Extrapolation of retina equivalent computer operations	Moravec (1997)
10^{13}-10^{16} operations per second	Power consumption	Merkle (1989)
10^{17} FLOPS(*)	Arithmetic equivalence	McEachern (1993)

* FLOPS = Floating Point (arithmetic) Operation (addition or multiply) Per Second

The fact that each of these capacity estimates using different methods, computes values within a range of 10,000 demonstrates how poorly understood the brain is at this time. The numbers are however in general agreement. Because of the specialized structures of the brain, it is impossible to focus all of the available capacity on a single problem. Computers, unlike the brain, can devote all of their capacity to a single problem (assuming the problem fits in available memory). This would imply that computers do not require the capacity of the brain to achieve equivalence with specialized areas of the brain. Developing trends in desktop computers are analogous to the multiprocessing occurring in the brain. It is not uncommon for systems may now execute 10-20 processes simultaneously. These might include listening to a network, listening to human speech, recording and compressing information for permanent storage, displaying information for interpretation, and devoting intensive processing power to search, recognition or analytical processes. The available computer power is divided among the tasks at hand in the computer, just as in the brain.

Computer capacity has increased significantly in recent years. Current state-of-the-art computers achieve operating levels as follows:

- Intel Teracomputer: 1.8 Teraflops (1.8×10^{12} FLOPS)
- IBM Teracomputer: 3 Teraflops (3×10^{12} FLOPS)
- IBM ASCI White: 12.3 Teraflops (1.23×10^{13} FLOPS)
- IBM Deep Blue (Chess computer):
  - 200 million (2×10^{8}) positions per second = ~3 million MIPS (3×10^{12} IPS)
- GRAPE (GRAvity PipE) computers for stellar orbit calculations (Taubes, 1997)
  - GRAPE-3 (1991): 600 MFLOPS (6×10^{8} FLOPS)
  - GRAPE-4 (1995): 1.08 Teraflops (1.1×10^{12} FLOPS)
  - GRAPE-5 (1998): 21.6 Gigaflop (2.16×10^{10} FLOPS)
  - GRAPE-6 (2000/1): 100 Teraflops (1×10^{14} FLOPS):
- On the drawing board:
  - IBM Blue Gene (~2003): 1 petaflop (10^{15} FLOPS)

It is clear from these numbers that computers are approaching human brain capacity and will eventually exceed it. As pointed out by Moravec (1997), the Deep Blue computer was able to defeat Gary Kasparov with only 1/30th of the estimated power in the human brain. Either the brain is has less capacity than the estimates above would indicate or humans are unable to devote all of that capacity to a single task.

Computers have always been better than humans in arithmetic. They now seem to be approaching our abilities in tasks which require parallel processing. In recent years, computer systems have demonstrated 'human' abilities such as:

- Besting humans in games such as Blackjack (with random factors) and Checkers & Chess (with non-random principles). The only game where a human can currently compete against a computer is Go.
- Proving theorems previously unproved by humans.
- "Reading" documents (OCR) and "understanding" human speech.
- Driving automobiles.

The realm of activities which are only available to humans is becoming increasingly small so it seems reasonable to assume that computers will match and eventually exceed human capabilities.

Solar Power

Drexler has observed [Drexler, 1992], that with no improvements in device physics, but simply the technology to fabricate small precise structures, it should be possible to construct solar collectors with a mass of $\sim 10^{-3}$ kg/m^2 and a power-to-mass harvesting capability at Earth orbit of $\sim 10^5$ W/kg. The power output of the sun is $\sim 4 \times 10^{26}$ W, implying a mass requirement of $\sim 10^{21}$ kg for solar collectors in Earth orbit. This is approximately the estimated mass of the asteroids in the asteroid belt and significantly less than the mass of the Earth's moon (7×10^{22} kg) or the planet Mercury (3×10^{23} kg). An orbit for solar collectors between Mercury and Venus would reduce the mass requirements still further. Orbits that are very close to the sun could result in a decrease in the lifetime of the solar collectors, that could decrease the quantity of harvested energy due to the requirement for continual reconstruction of the solar cells [Landis, 1998]. Assuming our solar system is typical, harvesting the entire energy output of a star using the material present in its solar system is feasible[2].

Kardashev Civilization Levels

Nikolai Kardashev in his 1964 paper, defined three major evolutionary levels for civilizations. These are outlined in Table 3.

Table 3. Kardashev Civilization Types

Civilization Level	Energy Resource Utilization	Available Power		Available Mass
		erg/sec	W	kg
KT-I	planet	$\sim 10^{19}$	$\sim 10^{12}$	$\sim 10^{24}$
KT-II	star	$\sim 10^{30}$-10^{36}	$\sim 10^{23}$-10^{29}	$\sim 10^{30}$
KT-III	galaxy	$\sim 10^{42}$-10^{43}	$\sim 10^{37}$-10^{38}	$\sim 10^{42}$

For KT-II and KT-III level civilizations, if the material in the stars is excluded then the mass available is perhaps $\sim 10^{-4}$ lower. If the star, and the hydrogen and helium in the solar system are excluded, then the mass available to a civilization is $\sim 10^{-6}$ of the total available mass. There are reasons, discussed below, that suggest that these strength of these exclusions changes with the evolution of the civilization.

Development Trends

In short, we can see that computer power is likely to continue to evolve until it significantly exceeds human intelligence. The ability to harvest stellar power output levels to rearrange the materials of a solar system seems feasible. Unless something occurs to intervene in the course of its evolution, a technological civilization should evolve from a KT-I to a KT-II and perhaps a KT-III level. As the greatest intelligence levels for these civilizations would be accomplished by constructing supercomputers constructed from nanotechnology using star output power levels, it is presumed that at least some civilizations would follow this path. These constructs are examined in more detail below.

Requirements

There are two important requirements for the possible construction of Matrioshka Brains.

1. Designs for the power collectors, computers, communication apparatus and heat radiators required to actually build a MB.
2. A design for self-replicating factories which can build the MB components. These factories do not have to be constructed using nanotechnology. They can be constructed with automata based on existing manufacturing methods utilizing ideas developed by Von Neumann [1966]. In 1982, a team of physicists and engineers organized by NASA produced an extensive study of how to construct automated factories for space manufacturing [Freitas, 1982]. These factories must either be capable of replicating themselves to cover the surface of a planet and be able to construct apparatus to transport their output to appropriate stellar orbits or they must be self-replicating in a space environment. Such self-reproducing factories are not strictly necessary for the production of MBs. They do however serve to decrease the construction time for a MB by many orders of magnitude.

Architectures

Matrioshka Brains have the following components:

Power Collectors

Power collectors or harvesters may have several measures of their effectiveness:

- absolute power harvesting efficiency
- mass or element requirements for a specified harvesting efficiency
- minimal time to construct a relatively efficient collector
- collectors with maximal longevity and minimal maintenance requirements

The simplest power collectors would be high efficiency solar cells or collecting mirrors focusing sunlight on solar cells to convert sunlight into electricity. These devices may typically operate at conversion efficiencies around 30%. Alternate power collectors might focus infrared radiation in ways that would allow a Stirling cycle to be used to generate electricity. Even more elaborate methods would attempt to harvest ultraviolet, X-ray, gamma-ray, microwave and radio energy in ways that would allow their conversion to form which may be used to perform work.

Computing Components

Computing components perform the actual work of the Matrioshka Brain. Computing components may have many architectures which are dependent on the nature of the problem being solved.

The computers would typically be NanoCPUs or Mega-NanoCPUs with a large amount of nanoscale storage and high efficiency, high bandwidth (optical) communications channels to other similar devices.

Storage Components

Storage components are required to store the data being used as inputs to the computing components and to store the results of the computations. We can envision four levels of storage, each requiring less matter and energy:

1. Aggregates of atoms and magnetic or optical sensing (e.g. current high density disk drives).
2. Elements in specific atomic configurations (e.g. DNA).
3. Electron charges (single-electron DRAM) or electrons in specific spin states (atomic NMR).
4. Circulating packets of photons.

Photonic storage seems to be the limit because the smallest amount of matter/energy is required to represent a bit. All of these methods presumably require a complex error-correction-code (ECC) algorithms and redundant backups to protect the data from disruptive events (high energy cosmic rays, etc.).

Heat Disposal Components

Any computing element, no matter how efficient, will require heat disposal. This is a consequence of the second law of thermodynamics. Individual NanoCPUs may be able to radiate away the heat they generate. They would require minimum inter-CPU spacings to avoid being baked by neighboring CPUs. Mega-NanoCPUs require active cooling, typically using a circulating cooling fluid between the CPU and heat radiators. The operating temperature of the CPU is determined by the materials from which it is constructed. The radiator temperature is determined by the allowed temperatures of the CPU, the cooling fluid and the radiator material. Radiators must dispose of the heat produced by the computing elements as well as waste heat remaining from photovoltaic, thermophotovoltaic, or Carnot cycle based power generation activities.

Radiation Protection Components

For most nanoscale devices, radiation damage is a significant hazard. The diagnostics and repair of radiation damage would represent a significant drain on time and energy resources unless significant efforts are made to minimize its effects. Radiation tolerance may be achieved in the following ways:

a. positioning MB in locations where radiation fluxes are minimal;
b. using controlled external power sources (beamed power) instead of uncontrolled internal power sources (stars);
c. active shielding with electric or magnetic fields;
d. passive shielding with bulk mass which is unusable for computer construction or energy production (large quantities of helium, neon or iron are good candidates).

General Architecture

Mega-NanoCPU computers are limited by their heat dissipation requirements and their heat removal capacity determined by the heat removal "fluid". The temperature of the radiators will be determined by the boiling point of the fluid or the melting point of frozen solids circulating in the fluid [Henson, 1988]. In a Matrioshka Brain architecture, the energy collectors, CPUs and radiators are arranged to take advantage of the downhill thermodynamic slope, Each layer of the shell harvests the energy (optical or heat radiation) of the next inner layer, performs what work is possible with that energy and radiates it at an even lower temperature. Appendix A outlines several possible working fluids, radiator materials and operating temperatures.

A Matrioshka Brain architecture is highly dependent on the structural materials from which the energy collectors, computational elements and radiators are constructed. At high temperatures the greatest problem is the destruction of the computational elements. Three relatively abundant materials from which high-temperature rod-logic computers could be constructed are diamond (stable to ~1275°K), aluminum oxide (M.P. ~2345°K), and titanium carbide (M.P. 3143°K). Material strength decreases relatively linearly as the operating temperature increases. Thermal expansion must also be taken into account in mechanical designs. So operating temperatures are likely to from 50-80% of the temperatures listed above. There are materials with higher melting points, particularly the elements rhenium and tungsten and the refractory compounds of hafnium carbide and tantalum carbide, but these elements are relatively rare. If the computers are operating at temperatures below ~1200°K, the thermal radiation from the radiators consists primarily of low energy infrared photons (< 0.5 eV) that few materials can harvest in ways that allow direct conversion to electricity. This implies that focusing the thermal energy (via mirrors) and using heat engines with Carnot cycles (e.g. Rakine, Stirling, Ericson, etc.) are likely to be used to generate power in the outer layers of the MB.

The elemental availability must be given consideration. Carbon, oxygen, magnesium, silicon, iron and aluminum are useful structural materials and are much more abundant than nickel, phosphorus, fluorine or tungsten, etc. More abundant elements (C, Al_2O_3, SiO_2, MgO, Fe_2O_3) should be the bulk construction material for MB layers. The most abundant materials being used in the larger outermost (cooler) shells. There is the possibility over long time periods of using breeder reactors to convert significant amounts of a less useful element (e.g. magnesium or iron) into a more useful element (e.g. tungsten or hafnium).

Power Sources

The most probable power source for a young MB is the star itself. The MB would be constructed around a star because it is fast, requires minimum amount of material and wastes the minimal amount of energy in material relocation. Over many thousands to millions of years however, gas giant planets may be disassembled and fed to fusion reactors. Local or remote stars may be disassembled, repackaged in fuel carriers (e.g. high-pressure diamondoid vessels) and transported to the MB location or arranged in locations where the MB orbit will encounter them in the future. The optimal architecture would be one that would eliminate the central star from the architecture and replace it with very high temperature computational elements that can be located in very close proximity (reducing inter-node communications delays) and yet still radiate a large amount of power to outer shells of the MB. However the problem of the controlled delivery of externally generated star output power levels (~10^{26} W) into volumes significantly less than that of the sun (~1,000-10,000 km^3) are not to be taken lightly. For example, a mirrored sphere at the center of an MB would have to be 17% larger than the sun, to receive external power (e.g. laser beams) at a density of 10,000 Suns (10,000 times solar insolation at earth). These mirrors would require phenomenal reflectivity to avoid accumulating heat that would cause them to melt. So to optimize their architecture for the minimum internodal communication delays at their center, an EPMB will likely consume (and radiate) much less than solar power levels.

Figure 1. An example of an integrated "Computronium"-element of a MB.

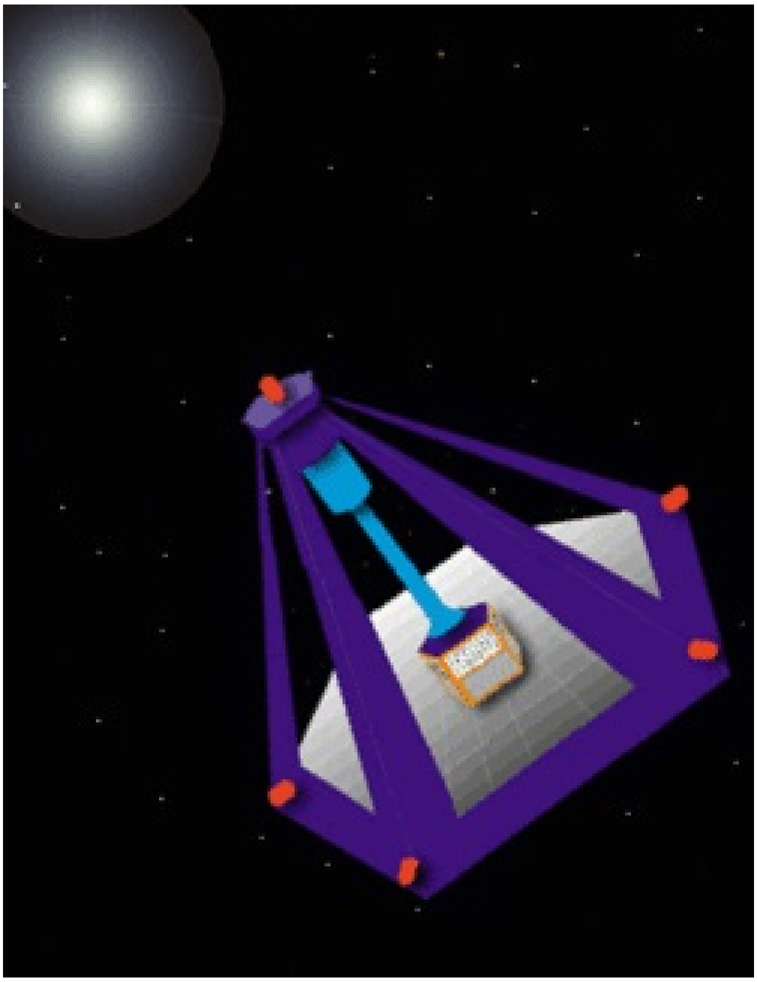

The surface facing the star is a solar array, the surface facing away from the star is a radiator. The hexagonal element in the center of the array is a nanocomputer. The dark blue portion circulates cooling fluid and the light blue portion is a high-pressure turbo-pump. The red bumps are vernier control nozzles for station keeping.

The nanocomputer surfaces are 2-D communication arrays of light transmitters and receivers, composed of VCSELs [Vertical-Cavity-Surface-Emitting-Lasers] and CCDs respectively. These arrays provide high bandwidth communications to adjacent compute-elements (see Figure 2).

Construction Methods and Time Scale

MBs may be constructed very slowly over thousands of years using a small fraction of a civilization's resources. More likely they will be constructed rapidly using leveraged construction techniques to take advantage of the benefits which can be expected from having significantly expanded computational and information storage capacities. In our solar system we can propose the following construction process:

1. Convert one or more asteroids into solar power collectors. A 3 mile asteroid receiving 10^{10} W of solar power can be converted into solar collectors which can harvest 10^{22} W of solar power. Time required: ~several years.
2. Beam the asteroid derived collector power to Mercury where it is used and convert the bulk of the planet into additional power collectors which harvest the entire solar output of the sun. Time required: < 1 month.
3. Use the entire output of the sun to disassemble the remaining asteroids, comets, moons and minor planets to construct the major portions of the MB. Time required: ~20 years.[3]
4. Use some fraction of the material from the moons and minor planets to construct thermonuclear reactors that provide the power required to disassemble the gas giant planets (Uranus, Neptune, Saturn & Jupiter). Time required: 10-1000 years.

An alternate scenario, as discussed in Planetary Disassembly, would skip the first step and convert Mercury and possibly parts of Venus directly into solar collectors and MB computing elements.

The energy requirements for disassembly of asteroids and small planets are dominated by the chemical bond manipulation requirements. In this situation, the best approach is to utilize material in locations which have the highest solar energy flux to construct ever expanding solar collectors. The critical determinant of the time required is the solar collector thickness. Current technologies allow construction of collectors (or mirrors) with masses of 1 kg/m^2. It is envisioned that collectors for solar sails may be as thin as .02 kg/m^2 (Potter, 1996) while Drexler 1992 postulates structures of .001 kg/m^2. The energy requirements for the disassembly of the larger planets particularly Saturn and Jupiter are dominated by requirement of getting the material out of the planet's gravity well. Even if the entire energy output of the sun were used to disassemble Jupiter, it would still take hundreds of years. Faster disassembly requires supplementing solar power with fusion energy derived from manufactured thermonuclear reactors. There are clearly tradeoffs between the amount of solar energy and/or minor planetary matter used for MB computations and the amount devoted to the construction of supplementary thermonuclear reactors and gas giant disassembly. Since the computational benefits derived from the disassembly of gas giants are marginal (relative to the huge benefit derived from dismantling even a single minor planet), civilizations may choose to dismantle the larger bodies at a relatively slow rate.

The mass requirements for the solar collectors and CPUs around the Sun are small compared with the mass available. Only small fractions of the Mercury or the Earth's moon would be required for their construction. Of some concern is whether specific elements required for CPU construction, such as carbon or sulfur would be available in sufficient quantities. If this is not the case, then one can turn to the atmosphere of Venus or the asteroids (esp. carbonaceous chondrites) for further material. The radiator material is of concern since it must have high emissivity. One candidate, likely to be available in high abundance is iron oxide (hematite). It has both a high melting temperature and highly abundant among the inner planets and asteroids.

Construction times are for immature MBs are short. Exponential growth[4] of nanoassemblers would provide sufficient numbers to disassemble and reassemble planets in weeks to months. If non-nano-scale automatons are required the time scale may be years to decades. The construction of a small number of solar collectors near the sun could provide high concentrations of beamed power to any point in the solar system. The strongest limit on construction times are likely to be the time required to move power collectors into the proper positions around the sun or the time required to ship materials from outer solar system locations to inner solar system locations, should essential elements be in short supply. Conversely if non-star centered MBs are desirable (see Location), the limit is on moving sufficient mass from various solar systems to gravitationally balanced or minimally disrupted point between the energy sources.

While many authors have focused on the possibility of moving comets, moons or planets for construction or terraforming purposes, it should be understood that this is not required for MB construction. First, since the elemental requirements of MB should be known, it would be better to disassemble materials on moons or planets and ship only those molecules or atoms which are absolutely necessary. Second, moving a large mass to an alternate orbit requires expending a large amount of energy and mass or waiting a long time or both. Instead the available energy and matter should be used to construct mass-drivers which accelerate material towards positions where optimal energy harvesting and beaming stations may be built. Once operational, these stations return an increased amount of energy to the moons or planets on which mass harvesting operations are taking place. This allows an exponential growth in material breakdown, separation, and transport capacity. Eventually the point is reached where an optimal amount of solar energy is diverted to the transport of materials optimal for MB construction.

Figure 2. An example of the nested shells of an MB

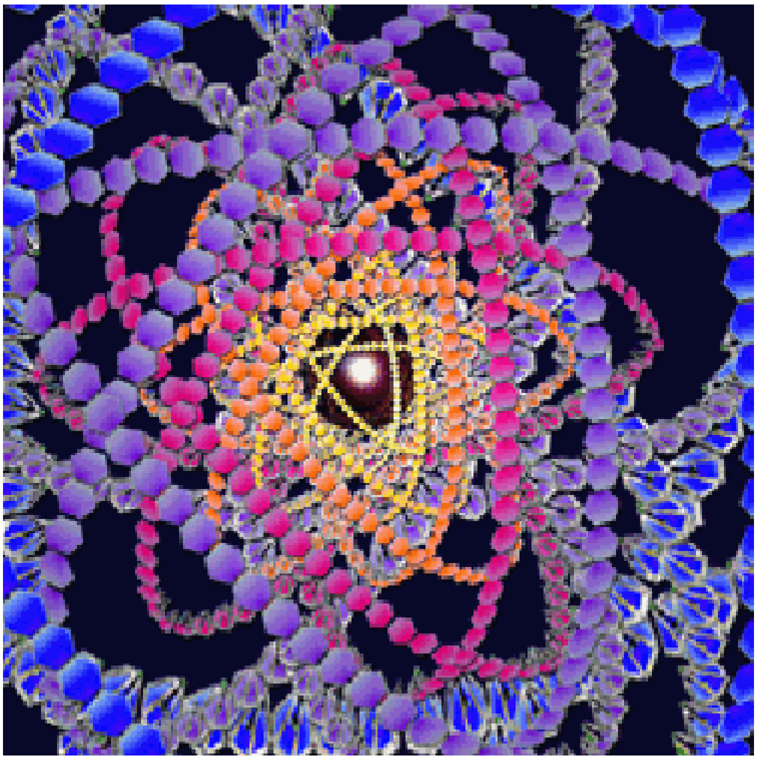

The shells shown are incomplete. A complete set of shells would hide the star. The innermost shells have the highest radiator temperatures and the outermost shells the coolest temperatures.

Limits

Power Limits

Table 4, details the power available from stars of various sizes and the lifetimes for unengineered stars. Stars much less than 0.1 M_{sun} do not become hot enough to burn hydrogen and stars greater than ~100 M_{sun} are both unstable and short-lived.

Table 4. Star power and lifetimes

Star Size	Mass (kg)	Power (W)	Lifetime (years)
0.1 M_{sun}	2×10^{29}	1.2×10^{24}	10 trillion
1.0 M_{sun}	2×10^{30}	3.8×10^{26}	12 billion
10 M_{sun}	2×10^{31}	2.0×10^{30}	20 million
100 M_{sun}	2×10^{32}	2.4×10^{34}	<< 1 million

It can be seen that there is a tradeoff between the amount of power available and the longevity of the power source. If you want to do a lot of thinking in a short time you can construct a MB around a 10-100 M_{sun} star. Unfortunately this massive amount of power increases your cooling requirements significantly and requires such a large diameter for the MB radiators that the amount of construction material available in an individual solar system will likely be insufficient. This then requires importing material from other solar systems or dust clouds creating the requirement that interstellar distance material transit times be incorporated into the construction schedule. As the lifetime of these large stars is short, presumably you would have to plan the construction and begin the material transfers while the star is still forming. This requires transferring the materials against the very strong solar wind of a large mass star during its violent and high radiation output formation stage. Even after the construction of a mega-MB, the large diameter would imply that the transit time for messages between CPUs would be hours or days. Clearly a mega-MB would only be useful if one wanted to solve well-defined problems which required a great deal of thought in a short period of time. Since stars more than 1.5 M_{sun} end their life by becoming supernovae, the MB would have to be disassembled and reassembled elsewhere else unless energy and matter were considered to be so plentiful that the incineration of the megamind is of no concern. These difficulties all argue against the construction of MB around large mass stars.

However, a non-star centered mega-MB can be constructed and be powered by either externally supplied power or internal thermonuclear reactors. This avoids the stellar radiation and lifetime problems leaving only the inter-node travel time problem. If this problem is of no concern, then one might find non-star centered mega-MB in regions where there is a high external energy flux (for power harvesting) and relatively long lived stars. These are the characteristics of globular clusters (GC) which consist of hundreds of thousands to millions of stars in regions of space only a few tens to hundreds of light years in size. The external light flux in GC is many times greater than that available from a single star.

Astronomers believe that GC are at least 8 billion years old with some estimates as high as 12 billion years. These ages are based on two observations:

- Large numbers of low luminosity red stars, implying old small mass stars.
- Low metal abundance in the spectra of the stars in the GC, implying old stars formed before significant numbers of supernova had seeded the stellar dust with metals.

However, the low luminosity of stars in a GC could be due to light harvesting and redirection for power and the low metal abundance due to metal mining for construction projects. There are several paths that lead to the this situation.

The first possibility is that an older external MB civilization would send a robotic nano-probe with the necessary mining plans to a GC "cloud" very early in its formation. The seed then constructs the necessary mining equipment and commences harvesting metals before they have an opportunity to be incorporated into stars. The second possibility is that if star-lifting is possible [Criswell, 1985], then the MB (or their probes) may evolve in the GC or arrive from remote locations and are recycling the stars to harvest all of the available metals. It may even be possible that a MB based civilization (KT-II+) could engineer the formation of a GC by using mega-lasers or focused redirected solar winds (essentially large ion-beams) to direct many large mass interstellar dust clouds towards a common point in space. It is questionable whether the estimated age of the universe would allow sufficient time for such construction efforts however.

A final question remains as to why a MB civilization would not harvest all of the energy available in a GC for the purpose of thinking? If optimal MB architectures are constrained by specific element abundances (e.g. carbon), then the best use of these elements is in the construction of computational machines or long-term memory storage and not power harvesting apparatus. Diverting any of the rare elements for the construction of thermonuclear breeder reactors may be a suboptimal use of resources. It may be much more efficient to allow gravity to serve as the container for stellar thermonuclear reactors and harvest the heavier elements as cheaply as possible. If the GC are breeding grounds for the production of elements required for low-power data-storage devices, then the power loss from the stars in GC radiating into space is of no concern. Over many billions of years, the GC is gradually converted from lighter elements into a massive long-term solid-state memory.

Size Limits

MBs must have minimum sizes determined by the input power and the radiator temperature. As the power input is increased, the MB size must increase as well. A MB has no maximum size limit other than those dictated by the available matter. Since 80+% of the mass of many galaxies appears to be "missing", the question of whether such matter has been incorporated into MB should be considered. It is important to realize that the larger a MB is, the slower the brain-level thought process are due to increased inter-CPU-unit communication distances. A smaller MB, constrained to less power consumption, is similarly capable of less aggregate thought but the thoughts are circulated much more quickly. Mega-MBs have many more thoughts and require days or perhaps months to circulate them internally. Minimal-MBs have fewer thoughts but may circulate them in minutes. It is possible to imagine MB clouds or swarms consisting of various sizes of MB thinking at different speeds or about different levels of a problem. Complex tradeoffs exist in power delivery, heat removal and line-of-sight vs. routed communications in determining optimal MB sizes and configurations. The relatively high power flux in Globular Clusters would enable the construction of a MB community which had relatively small inter-MB communication times on the order of months. In contrast a MB community in the galactic halo (a KT-III civilization) would require inter-MB communication times of tens to hundreds of years. There are advantages to layers of nested communities as shown in Table 5.

Table 5. Brain Communications, Power & Capacities

Brain Level	Intra-entity communication time	Power Consumed (W)	Thought Capacity OPS/sec
Human Brain	milliseconds	10	10^{13}
Mega-nanoCPU	microseconds	10^{5}	10^{21}
Single-layer MB	minutes	10^{26}	10^{42}
Multi-layer MB	hours-days	3×10^{26}	3×10^{42}
Globular Cluster MB community	weeks-months	10^{31}	10^{47}
Galactic Halo MB community	tens-thousands years	10^{36}	10^{52}

If rapidity of integrated thoughts is desirable then smaller, more densely packed CPUs and MBs are preferable. As the volumetric CPU density increases, a greater amount of power and mass must be devoted to cooling functions (e.g. circulating fluid or ballistic atomic transport). In Mega-nanoCPUs, some 10-30% of the mass and power must be devoted to cooling [Drexler, 1998]. The ultimate tradeoff is between the thought capacity lost because power and mass must be devoted to cooling and thought capacity lost due to increased inter-CPU or inter-MB communication times. Without knowing the details of the CPUs, cooling systems, communication systems, thought algorithms and architectures, it is difficult to estimate exactly where the tradeoffs occur. There is a high probability that different types of problems would be solved by different types of system architectures such is now the case with microcontrollers, microprocessors, general purpose computers, specialized instruction sets, chess computers, gravity computation machines, integrated systems to model protein folding, etc.

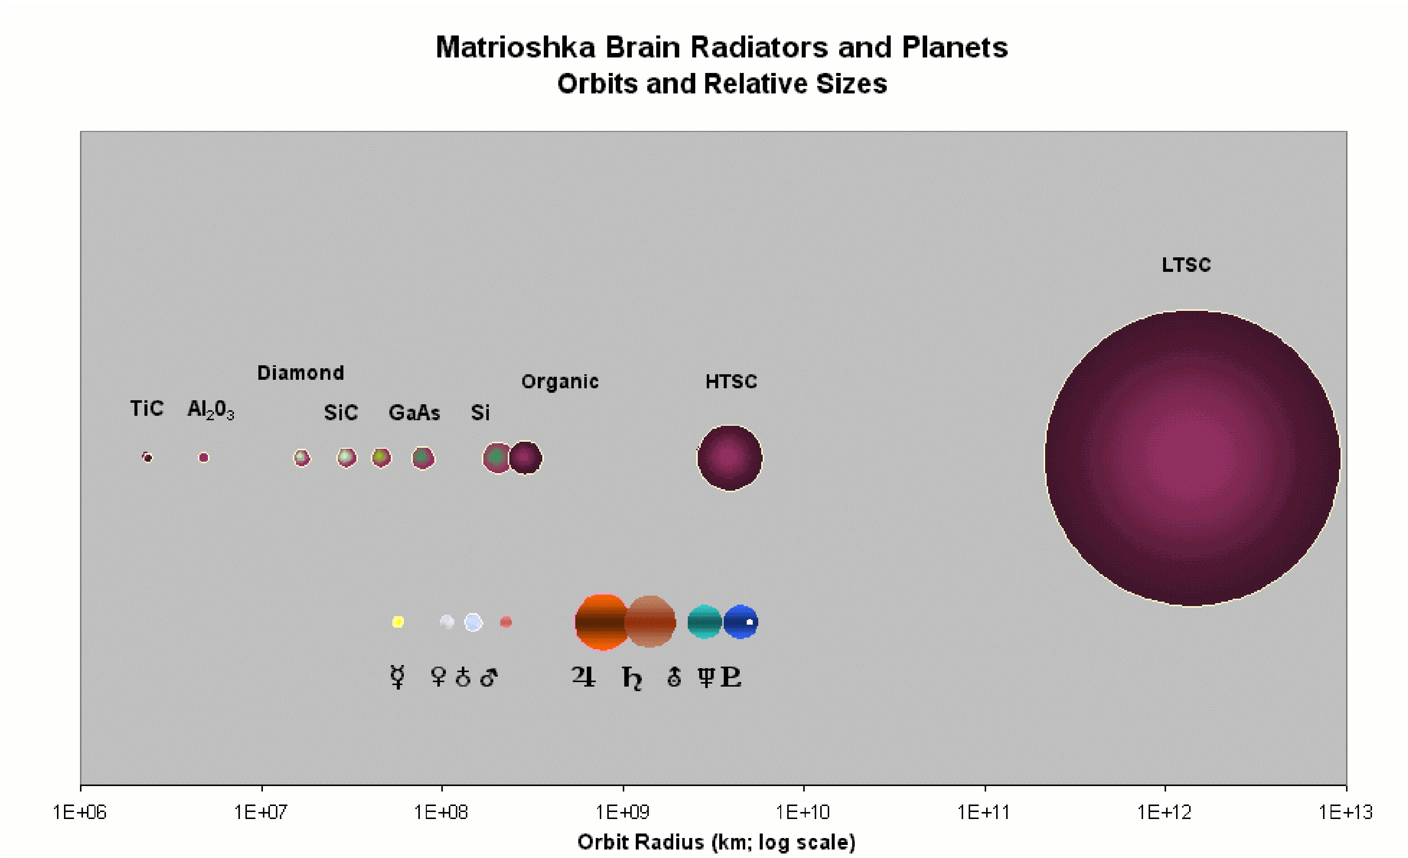

Graph of Matrioshka Brain layers constructed from different "computronium" materials. The computronium orbits are based on the radiator sizes necessary to radiate away the entire power output of the sun at the highest operating temperature that can be tolerated by a specific computronium type. The computronium orbits may be compared with the orbits of the planets to better comprehend the construction scale of Matrioshka Brains. The size of the circles, though not-to-scale, gives a relative impression of the radiator sizes required by the fact that radiator efficiency scales with T^4 (i.e. cold radiators are very poor). This law affects both the inter-node communication times and the radiator mass requirements. It is doubtful that our solar system contains enough mass in the planets to construct a fully populated layer of radiators at LTSC (Low Temperature Super Conductor) temperatures.

Material Limits

Matrioshka Brain construction is limited by the fundamental abundance of elements in their local region of space. Silicon may be excellent for building microprocessors and carbon (as diamond) excellent for building rod logic computers but neither of these elements is highly abundant in the universe as a whole. A major part of engineering MB will be the efficient partitioning of matter into the various components. The table in Appendix B outlines the various abundances of elements and particularly good uses in the construction of MB.

There will be ultimate limits on the MB architecture due to insufficient materials. Possible examples include:

- Insufficient titanium, aluminum or carbon to build the maximum number of nanoscale components, particularly high pressure circulation systems and nanocomputers.
- Insufficient aluminum or magnesium to build solar collecting apparatus capable of harvesting and redirecting the maximum amount of available solar power.
- Insufficient copper, nickel or iron for the construction of highly efficient metal oxide radiator surfaces (though amorphous carbon, e.g. lampblack, may be a substitute).
- Insufficient circulating fluid (Na, NH_3, CH_4, O_2, N_2, Ne, He) for the efficient cooling of computers (rod-logic, semiconductor, helical-logic, superconductor) at specified operating temperatures.
- Insufficient rare elements (Sb, In, Cd, Te, Hg, As, B, etc.) used as semiconductor dopants or as layers in solar cells or semiconductor lasers.
- Insufficient silver or gold to build highly effective telescopes for observing or communicating with other civilizations.

While theoretically, it is possible to fuse lighter atoms into heavier atoms with a net gain of energy, this only solves the problems for elements in short supply which have atomic numbers less than iron. Elements which have atomic numbers higher than iron require a net addition of energy for synthesis. If optimal computers require amounts of these elements greater than their natural abundances in a solar system, optimal MB architectures will be limited by the available energy and element transmutation efficiencies.

For comparison purposes, the following table outlines the elemental composition of three nanostructures designed by Eric Drexler at the Institute for Molecular Manufacturing and a familiar complex of nanomachines.

Table 6. Element Use in Nanoparts and Biological Systems

Element	Nanomachine components						Complex Nanomachine Aggregate	
	Pump		Fine Motion Controller		Differential Gear		Human without water	Range
	# Atoms	%	# Atoms	%	#Atoms	%	%	%
H	1806	29	403	15	864	10	51	10-51
C	1826	29	1433	55	2461	29	25	25-55
N	224	3	536	20	628	7	3	3-20
O	367	5	134	5	367	8	20	5-20

F	0	0	12	<1	0	<1	<<1	0-1
P	77	1	0	0	452	5	0.5	0-5
S	220	3	34	1	356	4	0.5	0.5-4
Si	1645	26	44	1	2792	33	<< 1	<<1-33
Ca							0.3	0-0.3
Totals	6165		2596		8297			

This shows clearly the variability that nanomachine compositions may have and illustrates the difficulty we will have in determining what elemental makeup of MB may be. However, it seems reasonable to say that whatever architectures are chosen, some elements will in excess relative to other elements. While the carbon, silicon, metals, semiconductor dopant atoms and elements with unusual properties (melting point, hardness, density, ferromagnetism, superconductivity, etc.) are likely to be fully utilized, there may be a significant excess of hydrogen, helium, neon, and perhaps even nitrogen and oxygen. Possible uses for these materials could include the construction and maintenance of biological zoos [Ball, 1973] or radiation shields and controlled fusion fuel sources.

Obviously substitutions can and will occur. MB will optimize their structures to make the most efficient use of the readily available elements. Without knowing the specific material requirements for various MB components it is impossible to predict at this time which elements will be the platinum, gold and silver of a MB culture. We can presume that very young MB will however use all available matter within a solar system and commence the study of where additional matter should be mined or whether the local star(s) and local MB architecture(s) should be engineered for long term elemental transmutation activities to create element ratios which are better suited for optimal MB architectures. Element transmutation on massive scales may require large amounts of energy. and has long time scales it is likely that interstellar mining in high density gas clouds will initially be a more rapid and less expensive solution to the accumulation of rare and valuable materials. In the long term, as local resources are exhausted, elemental transmutation will be the only reasonable solution for producing optimal element ratios. As an example, if a highly efficient method were discovered to convert ^{56}Fe into ^{184}W (consuming ~4.5 $\times 10^{13}$ W/mol) and 10 Earth masses of iron (5.9 $\times 10^{26}$ kg) were available, it would take approximately 4000 years using the *entire* power output of the sun to perform the transmutation.

As there exists the possibility, pointed out by Kardashev [1997] and presumably by many others, that intelligent life may have existed for 6 billion years or more. If intelligent life has existed for that long and evolves into MB architectures as postulated here, then interstellar mining activities may have been occurring for billions of years on galactic scales. This has serious consequences for astrophysical theories about the origin and history of the universe as they depend heavily on observed abundances of metals in stars and interstellar space and assume these ratios have not been adjusted by extraterrestrial intelligences optimizing their personal element ratios.

Even very sloppy single-layer MB architectures come relatively close to the most efficient computing structures possible given the physical laws of this universe. They will also be able to utilize most of the energy produced by a local star with only a small fraction of the locally available matter. If computational throughput is the major emphasis of MB (see thought limits), then it may be much more important to construct small-hot MBs and not large-cold MBs (whose radiators would exceed local material requirements). Thus there may be no incentive to go on interstellar mining expeditions and the astrophysicists may still be able to sleep nights.

Longevity Limits

A Matrioshka Brain has a lifetime limitation of tens to hundreds of billions of years around small stars. Criswell [1985] proposed that removing some of a star's mass might be used to extend the lifetime of a star. Such, "star-lifting", if feasible, would allow a MB to have a lifetime of tens of billions of years around most stars. Since MB are constructed of modular subunits, dismantling a MB and reassembling it around another star is a well defined process. Since thought time would be lost during the disassembly, relocation and reassembly process, it is probably preferable to do this as infrequently as possible. The construction of MB around low-mass stars with very long lifetimes would be a natural choice to minimize the frequency of relocation operations. If regions of space exist which have a relatively high matter abundance, such as within interstellar gas and dust clouds, it may be feasible to construct MB fueled from the materials they harvest as they circulate through the clouds. The longevity in this case is a function of the mass consumption rate which in turn is determined by amount of material available for the mass harvesters, the speed through the clouds and the overall rate of thought. The construction of CPUs with completely variable clock rates and therefore power consumption rates is possible now. One expects that a MB constructed of such subunits would do much of its thinking in regions of high mass density and sleep while traveling across low mass density regions. Of course these periods of being awake and asleep may last hundreds to tens of thousands of years.

Dyson [1979] demonstrated that in an open universe it is theoretically possible to live indefinitely, consuming less and less power as one thinks more slowly. Current results in astrophysics lean towards an open universe structure [*REF TBD*]. Though Dyson did not indicate exactly what the physical nature of immortal "beings" would be, it is clear that MB which have tremendously greater thought capacity than we do will have a much longer time than our sun has existed in which to consider and solve this problem.

Thought Limits

It is important to comprehend that MB are constrained by propagation delays. You can easily make a MB which has increased thought capacity by harvesting large amounts of external energy, beaming it to the MB locale and expanding the sphere with the MB uses to radiate energy. As the radius of the radiation sphere is increased, so too does the inter-node communication path. The time to communicate with your nearest neighbor remains the same, but the time to communicate with your most distant member of your layer increases significantly. As a result integrated "thought" *must* slow down.

> *This is clearly seen when imagining the management of three different planetary probes, one on the moon, one on Mars and one orbiting Saturn. The moon probe may be managed from earth in real time. The Mars probe can be given directions between cups of coffee. The Saturn probe can be given directions only several times a day. If you expect the more distant probes to do useful work in a reasonable time you have to build into them increased amounts of intelligence and autonomy.*
>
> *If the thoughts between CPUs in a MB are independent, then the brain can be made very large with little effect. If however the MB is attempting to solve a problem which requires all of its capacity then it must think slower to maintain synchronization between CPUs as their inter-node distance increases. In theory MBs orbiting in the galactic halo, a KT-III civilization would be able to think collectively, but their "thought" time must be on the order of tens of thousands of years or more.*

The two major problems facing MB are how to think more efficiently and how to think smaller.

Thinking more "efficiently", means to solve the problem with less heat generation. If the thought engines generate less heat, they can be placed closer together and can therefore solve a problem more quickly. It might be useful for very complex problems to devote a significant amount of thought and prototyping to the production of thought engines which are optimal for a specific problem.

McKendree [1997] discusses the possibility of nanotechnology based engineering being able to "surge" the production of various components necessary for the minimal solution times for problems which are well-defined from a computational standpoint. Using these methods, all of the CPUs in the MB would then be reconstructed for that specific problem, the problem would be "thought about", and after a solution is produced the process would be repeated for other problems. Current FPGA (field programmable gate array) products from manufacturers such as Xilinx and research in configurable computing (Villasenor & Mangione-Smith, 1997) are the foundations for these MB computational methods. Alternately, CPU groups or complete MBs may have architectures designed for solving specific types of problems, e.g. galactic stellar motion computations as is now done by the GRAPE computers in Japan. Some possible architectures could be:

- rapid prototyping for the construction of novel thought machines,
- monitoring architectures to monitor the development of life on all planets of a galaxy,
- massive communication architectures to interpret the signals from MBs from nearby galaxies,
- radiation-hardened MBs for rapid-thought in high energy flux environments such as those near

the accretion disks of black holes.

If MBs are part of a KT-III civilization, then it is likely that some partitioning of activities and architectures may occur to allow individual MB to optimize their activities based on local material resources, energy availability, radiation environment and nearby galaxy proximity.

Thinking "smaller", means to develop new architectures which move through the macro-atomic structural level to the sub-atomic structural level. One can begin to see hints of possible approaches in this arena in single-electron devices, optical computing and quantum computing. Compute engines built using these methods are "faster", in that more computation is done per time interval, but the limits imposed by heat removal and inter-compute engine communication time still impose limits to thought capacity. These solutions may provide several orders of magnitude (10^2-10^4) improvement in macro-atomic scale MB but increases beyond this will either be impossible or will involve magical physics that we do not currently understand well.

Location

Matrioshka Brains, if built from nanoscale components, are susceptible to radiation damage. High radiation areas, such as the galactic center, or the vicinity of black holes, would be unlikely areas to locate Matrioshka Brains. MB may be located in such locations if:

- significant mass resources are used for shielding the MB from radiation; or
- a not insignificant fraction of instruction capacity is devoted to redundant hardware and/or the execution of diagnostic programs *and* energy resources are devoted to the recycling of damaged components.

Development of a long-lived, stable, maximally efficient entity would require environmental characteristics which include:

1. A low radiation flux
2. Large amounts of local, easily harvestable energy
3. A relative abundance of easily utilizable mass
4. A minimum potential for disruption by gravitational or stellar events (e.g. supernova)

It would appear that some of these characteristics (high energy, low radiation, quiescent environment) are satisfied by globular clusters. One can envision non-stellar MBs in globular clusters harvesting the energy from multiple stars and beaming it to a point of low radiation flux. The globular clusters may be surrounded by MBs emitting large amounts of infrared radiation. MBs with an age of several billion years are likely to be optimally constructed MBs. Their outer layers may consist of superconducting logic elements cooled by liquid helium (Likharev, 1993). The radiator temperature of such structures will be only slightly greater than the microwave background radiation and will thus be very difficult for astronomers to detect.

If it is possible for MBs to harvest significant amounts of fusionable mass (H, He, etc.) from either stellar lifting or interstellar gas cloud mining, then the construction of migrating MBs is possible. These MBs may be constructed as solid spheres and may use the harvested elements in large numbers of fusion reactors to generate power. [*Question: Is it possible for a large MB with a solid shell to retain a large mass of H/He as a internal atmosphere as a potential fuel source (or will the H/He collapse into a gas planet)*?] Structures such as this could be found orbiting around the galaxy in the galactic halo.

Evidence for the existence of MBs

A significant majority of the work done in astronomy and astrophysics, unfortunately assumes that none of the observations could be explained by technological civilizations evolved to the limits of known physics. While this is based on the valid scientific principle that the simplest explanation is the preferred explanation there is a problem with this perspective. As we know life can evolve (i.e. we exist), and as this paper discusses, we can envision what civilizations could be like at the limits imposed by physics, for current astronomical interpretations to be valid, they must assume one of two things: (a) that the evolution of intelligent technological civilizations (ITC) has a very very low probability, or (b) that the developmental path of intelligent technological civilizations is uniformly fatal. We will not comment on (a) other than to observe that trends in the discovery of extraterrestrial planets would lead to expectations that there may be many locations in which life can evolve and trends in the discovery of microorganisms in harsh environments lead to the impression that life is quite able to adapt and survive in a variety of conditions and locations. Assumption (b) is even more strict than it seems at first glance because it requires that "fatal events" (e.g. very near supernovae, blazars, star or planetary collisions, huge solar flares, etc.) be either very severe or very frequent. In our case the development of simple microorganisms appears to occur over 1-2 billion years. Subsequent development of complex organisms takes less than a billion years and the evolution of intelligence appears require a few hundred thousand to a few million years. Once a planet has established a foundation of complex life forms, it may have hundreds to thousands of chances to evolve ITC. To prevent the this, the disasters must either be so severe that they eliminate the complex life form foundation or if less severe occur on a ~million year frequency. Another way of looking at this is that the "wisdom" incorporated into genomic blueprints is cumulative and you have to eliminate all the copies of the plans in order to lose that knowledge.

If we assume that even a single civilization makes it past the barriers to a KT-II level, then we may assume that within a few million years, they may take the galaxy to a KT-III level [Newman, 1981 & others]. Whether they choose to colonize the galaxy, or simply allow the galaxy to evolve many independent KT-II civilizations and the gradual emergence of a KT-III civilization occurs remains a matter of much discussion beyond the scope of this document (it requires a significant understanding of the motivations and goals of a KT-II level MB). The path to a KT-III level may be either a single KT-II civilization colonizing the galaxy or multiple KT-II/KT-II+ civilizations developing in local regions over a time-scale of millions to billions of years. Our galaxy is old enough for either of these situations to have occurred.

It is useful to examine some of the unexplained or poorly explained observations in astronomy and astrophysics that could bear witness to many galaxies being at a KT-III level.

1. The Missing Mass. ~80-90% of the mass of nearby galaxies is considered "missing" because the number of observable stars does not correspond to the rotation curves of the galaxies. [See Dark Matter essay by Joe Silk]
2. The Gravitational Microlensing Observations. Estimates are of approximately 4×10^{11} objects with masses of approximately 0.3-0.5 M_{sun} orbiting our galaxy. [Alcock, et. al.; Macho Project; ABC News 970817]
3. Missing stars in galactic halos.
    1. The Stellar Content of NGC 5907's Dark Matter Halo, Michael Liu, Francine Marleau, James Graham, Stephane Charlot, Penny Sackett, Steve Zepf

       Comments include: *"What they did not see - lots of stars - has led them to conclude that the halo is composed of a weird population of stars, mostly dim dwarfs too faint to see from Earth. Most galaxies contain a mix of bright giant stars and dim dwarf stars, with about half of the light coming from each group. If the halo of NGC 5907 contained a mix similar to that in our own galaxy, the team would have seen hundreds of bright giants in the field of view. Instead they saw only a handful of bright stars. The best explanation of the team's observations is that at least 20 times more light comes from dwarfs than giants in the halo of NGC 5907. "Our results force us to turn to more esoteric descriptions of the stellar content of NGC 5907's halo," said Michael Liu, a graduate student at UC Berkeley and lead author. "In particular, our data combined with the measured colors of the halo suggest a very metal-poor stellar population with an enormous excess of faint dwarfs."*
4. Low Surface Brightness and Dwarf Galaxies. There are galaxy classes that are dimmer and smaller than most galaxies. Most dwarf galaxies have rotation curves indicating they are dominated by "dark matter". A low surface brightness and/or small apparent size could be consistent with extensive MB development of the outer regions of the galaxies.
5. The Gamma-Ray Halo or Corona around the Milky Way and the dull "infrared glows" in nearby galaxies. The gamma-rays (20-30 MeV) may be due to cosmic rays (high speed electrons) transferring energy (via the inverse Compton effect) to visible or infrared photons around the galaxy. The infrared glows or halos have been attributed to galactic dust [Smith1996].

    Comments include: "Looking in any other wavelength, there is nothing out there that should be obviously making gamma rays. These gamma rays are providing the first evidence that some high energy process is occurring out there,"
6. An excess of far-infrared light detected in the COBE mission. [STSCI 98/09/01]

Comments include: *"The unexpected preponderance of far infrared light implies that many stars have `fallen between the cracks' in ultra-sensitive visible light probes of the distant corners of the universe, such as the Hubble Deep Field".*

7. The paradox between the age of the universe and the age of the Globular Clusters. Some Globular Clusters have low metal abundances (e. g. M15 & M92) and have ages estimated from 14-16 billion years. These would exceed the best estimates for the age of the universe (~12 billion years). This paradox would not exist if the low metal content of the stars could be attributed to stellar mining activities.
8. Observations that low temperature objects in the IRAS survey data by Kardashev et. al. [Kardashev, 1997] have an anisotropic distribution around the galaxy. These may be due to increased concentrations of MB toward the center of the galaxy.
9. The arrangement of observable galaxies as "walls" requiring the introduction of quantum density variations and inflation in the "Big Bang" theory. The "walls" could be explained by an expanding "front" of MB "locusts" that consume all stellar fuel sources as they expand.
10. The anisotropic variation (1 part in 100,000) in the Cosmic Microwave Background Radiation. This could be explained by varying densities of MB radiating at temperatures very close to the cosmic background.
11. The variation in the brightness of Type Ia supernovae depending on their age and the suggestions that a new force must be responsible for the increasing expansion rate of the universe. [*Science* (30 January, 1998):651, *Science* (31 October, 1997):799]. The calculations for supernovae may not properly correct for the fact that the oldest supernovas occur in "natural" galaxies (where element compositions may be unaffected by the activities of MB) while the nearest, most recent, supernovas may be occurring in "engineered" galaxies.

This list is approximately ordered from those phenomena that are most poorly explained, to those which are best explained using current theories. There is a significant question however, whether astrophysicists may be guilty of making the theories fit the data because they choose to ignore the possibility that "intelligence" may be coloring what they observe. That requires a strong assumption that the laws of the universe are "biased" against intelligence.

Evolution of Civilizations

The evolutionary path from Kardashev Type I civilization to an early Kardashev Type II civilization is rapid. Once nanocomputers, solar collectors and radiators are designed, the construction time of a Matrioshka Brain (KTII) is limited only by the willingness of the KT-I civilization to waste mass in repositioning material in its solar system. If the solar system is mass rich, mass may be "thrown away" in the rapid relocation of material to construct the MB. If it is a mass poor system, mass may be "conserved' using gentle accelerations and gravity assists to create trajectories which ultimately enable MB construction. The construction times depend to some degree on the size and sophistication of the MB being constructed but should be in the range from months to a few thousand years.

Criswell [1985] defined the concept of "stellar husbandry" which consists of the removal of the atmosphere of a star ("star lifting") and gradually returning the stored materials which are capable of undergoing fusion reactions to allow a significant extension of the lifespan of the star at least 1000 times (to 10^{14} years). This activity also provides an extensive source of materials for the construction of larger (and cooler) MBs. If possible, this activity would take tens to hundreds of millions of years. Since star lifting will eliminate many short term material resource constraints as well as provide a greatly extended lifespan for the MB, it would likely be an important goal.

Combining these perspectives provides a reasonable concept of KT-II evolution. Initial MB construction utilizes materials from asteroids or planets with the lowest gravity in closest proximity to the star. Construction is rapid (a few years), may be inefficient in its mass utilization and produces a hot (~500-1500°K) MB relatively near the star. As more material becomes available, larger planets in more distant orbits are dismantled and the MB shell expands or additional layers are constructed which are cooler (~70-300°K). Finally, if star lifting activities are undertaken and large quantities of metals become available, the MB enters its final stages with both a large size (5 AU radius) and cool temperatures (< 30°K).

The ultimate fate of MBs is unclear. A tradeoff must be made between active thought and information storage. Material returned to the star (or consumed in thermonuclear reactors) to enable active computation cannot be utilized in information storage. A means of utilizing all of the potential energy available and gradually converting most of the mass to iron may be developed. The iron could then be arranged, perhaps utilizing other required elements, in the form of a massive static information store. The last energy available could be utilized in accelerating these information stores in the direction of untapped energy sources where they could regenerate new MBs.

Interaction with and between Matrioshka Brains

Communication between humans and MB is essentially pointless. The computational capacity difference between a MB and a human is on the order of 10^{16} (ten million billion) times ***greater*** than the ***difference*** between a human and a nematode (~10^9)! A single MB can emulate the entire history of human thought in a few microseconds. It is important to consider that intelligence may not be a linear process. There is a rather large difference between the intelligence of a human and a chimpanzee or parrot, yet their computational capacities are not separated by more than a few orders of magnitude. Accumulated knowledge (language, history, teaching methods, scientific theories and data) significantly leverage the intelligence of individual humans. We may therefore expect, that the intelligence gap between a MB compared to a human or even human civilization could be significantly worse than that which might be expected from differences in computational capacity alone.

If one supposes that the transition period between KT-I civilizations and KT-II civilizations is short (thousands of years) relative to the lifetime of a MB (billions of years), then it makes little sense for MBs to concern themselves with creatures which are much much lower than insects. Perhaps they may take an interest once a civilization has progressed to the KT-II (MB) level as one now has the equivalent of a "child" which may be rapidly educated. The mass and energy resources available to MB are so large that they may observe us quite closely for a very long time from a large distance waiting to see if we will make the transition to a MB level .

It seems silly for them to interact with us at the level which we are now at. More likely, possible future outcomes of pre-KT-I civilizations, like our own, have been computed in some detail (it only takes seconds to compute thousands of thousand-year scenarios for us). We should not feel too bad however. A single MB has the same problem relative to a KT-III civilization which we have with them. KT-III civilizations made up of 10^{11} or more MBs would think on a radically different time scale than individual MBs. Since it is likely that the MBs of a KT-III civilization would be separated by light years, the propagation delays between them become are a significant problem. What does one think about when you yourself can compute an answer to most questions you might transmit before the answer can be received?

Are there galactic MB *Oracles* which have utilized their design and simulation capacities and mass transmutation or star lifting activities to construct optimal architectures for solving specific types of problems? The travel time to ask a question and receive an answer from such an *Oracle* may be 10^4 - 10^5 years. Nanotechnology enables surge construction of optimal problem attack architectures [McKendree, 1997]. So questions must involve problems which cannot be solved with an optimal architecture in the time to send a question and receive an answer from an *Oracle*. Presumably, individual MB would perform a return-on-investment analysis to determine whether it is more efficient to ask an *Oracle* or use local resources and reconstruction activities to produce an optimal architecture for thinking about the problem and producing a local solution. Obviously, utilizing ones resources to attack one problem means that those same resources cannot be used to solve another problem. The would be significant cost-benefit tradeoffs involved in asking the *Oracle(s)* or consuming local resources.

A single MB may use a fraction of 1% of its available mass to construct 100 billion telescopes with mirror diameters equal to that of the moon. These telescopes would fill a planar space corresponding to roughly the orbit of Jupiter. Using this number of telescopes they should be able to monitor most of the solar systems in the galaxy. If we assume some reasonable fraction of the galactic dark matter constitutes a KT-III civilization with billions of MB, then we may also assume they can monitor to a significant degree many activities occurring in nearest old galaxies within Kardashev's (1997) "civilization window". Major activities of MBs may be the monitoring of developing local KT-I civilizations and the nearest remote KT-III civilizations and contributing this information to the galactic gossip.

Given the possible existence of a galactic MB/KT-III civilization for 3-6 billion years, there should be a large directory of problems and answers computed and stored by MBs from preceding times. There should be a large amount of information about galactic history (stellar births & deaths, civilization histories, lifeform blueprints, etc.). The galactic knowledge base is potentially huge, but it plagued by the problems of long latency times for information retrieval as well as bandwidth limitations if the volume of information is large. While waiting for the retrieval of answers to questions, MBs may devote their time to devising complex problems which have not been solved and can only be solved in millions of years by a dedicated MB or closely linked MB cluster. It is difficult to imagine what these problems might be since even one MB has sufficient computational capacity to easily solve problems far beyond our current capabilities.

Conclusions

This paper discusses the background of technology trends that will lead to the development of nanotechnology and computers constructed at the limits of physics. Nanotechnology also allows technological civilizations to transition from the pre-KT-I level to KT-II levels, utilizing all of the local power and material resources in hundreds of years or less. Such civilizations would likely devote their attention to engineering their solar systems into the longest lived and most intelligent entities possible, e.g. Matrioshka Brains (MB). The age of the universe and our galaxy allow sufficient time for this to have occurred many times over. Current astronomical observations and unanswered astronomical puzzles may be explained by galactic civilizations at the KT-III stage. These observations also suggest that the Fermi Question, "Where is everyone?" should be answered with, "Almost everywhere we look, we just have a difficult time seeing them!".

The huge difference between the capacities and intelligence of a MB and our feeble human minds provides an explantion for why we have no contact with "them". It also implies that CETI (SETI) of the radiowave search type is likely to fail. The pre-KT-II stages of technological civilizations are likely to be so short (e.g. hundreds of years), that there will be few of them in our galaxy. Optical SETI searches might succeed if the path of the Earth or Solar System were to transit a direct communication path between two communicating MB. However interstellar distances are so great the probability of that seems small. They have little reason to waste time or energy transmitting signals directly to us. Searches that are likely to yield signs of MB include the gravitational microlensing searches, near and far infrared searches and occultation astronomy.

Self-preservation and self-structural optimization are the only goals or activities that we may easily imagine MB pursuing. Due to speed-of-light limits, growth in MB physical size or interstellar colonization yield diminishing thought returns. Instead they are likely to focus on becoming, smaller, faster, and more efficient. Their activities, based on technologies that originated with Richard Feynman's observation "There is plenty of room at the bottom", will transcend his observation because they understand that "There is *more* room at the bottom".

Open Questions

1. What are realistic time estimates for turning Mercury into a Solar Power Array?
2. How rapidly could all of the carbon in the atmosphere of Venus be harvested?
3. How long does it take for a reasonable fraction of solar output (RFSO) to break all chemical bonds in the solar system?
4. How long does it take for a RFSO to transport all non-solar matter in the solar system out of individual planetary gravity wells?
5. What are the energy/mass/time tradeoffs to relocate significant fractions of the mass in the solar system into the inner solar system using a RFSO?
6. What do Matrioshka Brains "think" about?
7. If MB do not exist, then one is faced with a serious problem of what makes it so difficult for intelligent life to evolve and achieve the MB level?
8. Even if MB do not exist, is it a reasonable path for humanity to follow?
9. Which architectures are really better - IPMB or EPMB?
10. Can the radiation damage problem for nanocomputers be solved by electro/magnetic shielding or must it be solved by mass shielding or relocation to regions remote from radiation sources?
11. How close can real architectures come to theoretical limits of computational capacity?

12. Are there problems which require so much computer processing power that they can only be answered by suicide MBs such as those who would harvest the energy output of supernova or Black Hole?
13. Is there any way a MB could tolerate the radiation flux associated with harvesting energy from matter falling into a black hole? Would this be the best use of the matter (its destruction in the pursuit of energy production)?
14. Is there any point to interstellar travel (MB voyaging or MB probe research expeditions) if you can observe all galactic activities with large numbers of telescopes, have high bandwidth communication channels to other MB and can precompute the probable end points of observed developmental paths?
15. If the dark matter is not MB, then what is it?

Appendix A. Layers of Matrioshka Brains

Phase Change Coolant		**Radiator Temperature**	**Radiator Material**
Solid (ice)	Coolant fluid	(^{o}K)	
Iron	Aluminum or Silicon	1808	Nickel Oxide
Silicon	Aluminum or Calcium	1683	Nickel Oxide
Calcium	Aluminum	1112	Nickel/Iron Oxide
Aluminum	Sodium or Potassium	934	Nickel/Iron Oxide
Magnesium	Sodium or Potassium	921	Nickel/Iron Oxide
Lithium	Sodium, Potassium or Phosphorus	454	Nickel/Iron Oxide/Graphite
Sulfur	Sodium, Potassium or Phosphorus	385	Nickel/Iron Oxide/Graphite
Sodium	Potassium	371	Nickel/Iron Oxide/Graphite
Ice	Pentane	273	Iron Oxide/Graphite
Ammonia	Methanol	195	Iron Oxide/Graphite
Methanol	Ethanol	179	Iron Oxide/Graphite
Pentane	Ethane	143	Iron Oxide/Graphite
Methyl Silane	Ethane	117	Iron Oxide/Graphite
Argon	Oxygen	84	Iron Oxide/Graphite
Nitrogen	Oxygen	63	Iron Oxide/Graphite
Oxygen	Fluorine	55	Iron Oxide/Graphite
Neon	Hydrogen	~24	Iron Oxide/Graphite

This table is based on the concept of the high heat capacity of a phase-change coolant in which the heat is absorbed by solid ice particles circulating in a fluid coolant [*Nanosystems*, Section 11.5.2]. The melting temperature of the "ice" should be between the melting point and the boiling point of the coolant fluid. The case of Ne in H_2 does not quite meet this criteria, but may be possible with appropriate coolant pressurization.

Appendix B: Element Abundances and Uses

Element	Relative Abundance	Application
Hydrogen	1,000,000,000,000	Radiation shield (LH_2), coolant (LH_2), power source (fusion)
Helium	97,723,722,096	Radiation shield (LHe), superconducting electronic or astronomical detector coolant
Oxygen	741,310,241	Ceramics (Al_2O_3), Coolants (LO_2, alcohols), high temperature superconductors
Carbon	354,813,389	High strength structural elements and thermal conductors (diamond), nanocomputers, coolants (CH_4, alcohols)
Neon	120,226,443	Coolant (LNe)
Nitrogen	93,325,430	Coolants (LN_2, NH_3), Ceramic nitrides, nanoscale components
Magnesium	38,018,940	Lightweight structural component, high temperature structures (Mg, MgO)
Silicon	35,481,339	Semiconductor materials, ceramics (SiC, Si_3N_4)
Iron	31,622,777	Radiation shield, strong structural components
Sulfur	19,952,623	Nanoscale components
Argon	3,162,278	Coolant (LAr)
Aluminum	2,951,209	Low efficiency solar reflectors and mirrors, lightweight structural components
Sodium	2,137,962	Coolant (LNa), nanoscale components
Nickel*	1,778,279	Radiator emitters (NiO)
Chlorine	316,228	Chemical reactions, coolants (CFCs)
Phosphorus	281,838	Nanoscale components, semiconductor lasers
Fluorine	39,811	Chemical reactions, nanoscale components, fluorocarbons (FCs), coolants (CFCs)
Copper*	16,218	Radiator emitters (CuO), high temperature superconductors
Germanium*	3,616	High speed semiconductor electronics
Selenium*	2,113	Solar cells (CdSe)
Gallium*	1,509	Solar cells (GaAs), high temperature semiconductors (GaN), semiconductor lasers and electronics
Boron	603	Semiconductor electronics
Arsenic*	208	Solar cells (GaAs), semiconductor lasers and electronics
Tellurium*	202	Solar cells (CdTe), astronomical detectors (HgCdTe)
Cadmium*	47	Solar cells (CdSe, CdTe), astronomical detectors (HgCdTe)
Niobium*	44	Low temperature superconductors
Silver*	14	Low electrical resistance, visible reflectors
Mercury*	13	Astronomical detectors (HgCdTe)
Antimony*	10	Astronomical detectors (InSb)
Gadolinium*	9	Magnetic refrigeration

Indium*	6	Semiconductor lasers, astronomical detectors (InSb)
Gold*	6	Infrared reflectors

Green rows (elements with an atomic number lighter than iron) are those elements that may be produced via fusion of lighter elements with the possibility of a net gain of energy. Pink rows(*) (elements with an atomic number greater than iron) can only be synthesized with a net loss of energy. Abundances are for our solar system, other solar systems may vary significantly.

Footnotes

1. A *Matrioshka Brain* is a term invented by the author to more accurately describe an optimal computing architecture for megascale structures that use stars as power sources. Earlier papers and discussions have referred to megascale thought machines as "superintelligences" or "Jupiter Brains". The term "Jupiter Brains" creates the impression of a semi-solid brain the size of Jupiter. There are a number of problems with this image. First, such a structure cannot consume stellar power outputs without melting. Second, a computer with this architecture would consume a significant fraction of its energy budget circulating cooling fluids. Finally, if collections of such structures were arrayed around a star, they would have to resolve the problem of not radiating heat onto each other (acerbating the already difficult problem of staying cool). The term "matrioshka" is derived from the nesting wooden dolls (Matryoshkas), found in Russia, and is intended to bring to mind the nested shells of "computronium" from which a Matrioshka Brain is constructed.
2. An expanded discussion of the harvesting of the entire power output of a star may be found in the Planetary Disassembly discussion.
3. Though not discussed here, there appear to be MB designs that would allow sunlight to continue to reach the Earth. The most obvious would be to place all construction outside of Earth's orbit. Other designs could contain orbital "holes" synchronized with the Earth's orbit or collectors that reorient themselves so as to minimize light interception when between they pass between the sun and the Earth.
4. Most people have a difficult time believing the growth rates allowed by self-replicating, self-assembling machines. Bacteria have an optimal replication time of 20 minutes and supplied with sufficient energy and nutrients a single bacteria could grow into an Earth equivalent mass in less than 2 days. Drexler [*Nanosystems*, 1992] estimates the doubling time for general purpose assemblers at 1-3 hours. Hall [1999] and Freitas [1999] have proposed that special purpose assembly line arrangements would have much faster production times.

References

- Ball, John A., "The Zoo Hypothesis", *Icarus*, **19**:347-9 (1973).
- Beech, R., Benett, W. J., Freitas, B. L., Mundigner, D., Comaskey, B. J., Solarz, R. W., Emanuel, M. A., "Modular Microchannel Cooled Heatsinks for High Average Power Laser Diode Arrays", *IEEE Journal of Quantum Electronics* **28**(4):966-976 (1992).
- Berry, Adrian, *The Next Ten Thousand Years: A Vision of Man's Future in the Universe*, Saturday Review Press, E. P. Dutton (1974).
- Criswell, D. R., "Solar System Industrialization: Implications for Interstellar Migrations", pp 50-87 in *Interstellar Migration and the Human Experience*, Ben R. Finney & Eric M. Jones (Eds.), University of California Press, Berkeley, CA (1985).
- Drexler, K. E., "Molecular engineering: An approach to the development of general capabilities for molecular manipulation," *Proc. Natl. Acad. Sci., U.S.A.* **78**:5275-5278 (1981).
- Drexler, K. E., *Engines of Creation*, Anchor, (1987).
- Drexler, K. E., "Molecular Manufacturing for Space Systems: An Overview," *J. British Interplanetary Society*, **45**:401-405 (1992).
- Drexler, K. E., *Nanosystems: Molecular Machinery, Manufacturing and Computation*, John Wiley & Sons (1992).
- Drexler, K. E., Presentation at the Foresight Institute's 1998 Senior Associates Meeting.
- Drexler, K. E., personal communication February 22, 1998 and *Nanosystems*, pgs 330-332.
- Dyson, F. J., "Search for Artificial Stellar Sources of Infrared Radiation," *Science*, 3 June,

1960, **131**(3414):1667-1668.
- Dyson, F. J., Interstellar Transport, *Physics Today*, **21**(10):41-45 (10 October 1968).
- Dyson, F. J., "Time without End: Physics and Biology in an Open Universe," *Reviews of Modern Physics*, Vol. **51**, No. 3, July 1979
- Feynman, R. P., "There's Plenty of Room at the Bottom", (1959).
- Frank, M. P. & Knight, T. F., "Ultimate Theoretical Models of Nanocomputers," *Nanotechnology*, to appear, (1998).
- Freitas. R. A., W. P. Gilbreath (eds.) *Advanced Automation for Space Missions*, Proceedings of the 1980 NASA/ASEE Summer Study held at the University of Santa Clara, Santa Clara, California, June 23-August 29, 1980. NASA Conference Publication CP-2255, November 1982. 393 pp. [STIS Abstract]; [Online copies of this study: @IslandOne; @GeoCities/CapeCanaveral]
- Freitas. R. A., personal communication 1999.
- Froman, D., "The Earth as a Man-Controlled Space Ship," *Physics Today*, July 1962, 19-23.
- Hall, J. S., "Architectural considerations for self-replicating manufacturing systems", *Nanotechnology* **10**(1999):323-330.
- Henson, H. K., and K. E. Drexler, Heterodensity heat transfer apparatus and method. U.S. Patent Office: #4,759,404 (1988).
- Kardashev, N. S., "Cosmology and Civilization," preprint, Moscow, (1997).
- Kardashev, N. S., "Supercivilizations as Possible Products of the Progressive Evolution of Matter," Marx, G. (Ed.) *Bioastronomy - The Next Steps*, Kluwer Academic Publishers (1988), pp 295-302.
- Kardashev, N. S., "On the Inevitability and the Possible Structures of Supercivilizations," Papagiannis, M. D. (Ed.) *The Search for Extraterrestrial Life: Recent Developments*, IAU (1985), pp 467-504.
- Kardashev, N. S., "Transmission of Information by Extraterrestrial Civilizations," *Astronomicheskii Zhurn.*, **41**:282 (1964). [*Soviet Astronomy* **8**(2):217-220 (1964).]
- LaBarbera, M., "Principles of Design of Fluid Transport Systems in Zoology", *Science*, **249**:992-1000 (1990).
- Landis, G., Personal Communication (1998).
- Likharev, K. K., "Rapid Single-Flux-Quantum Logic", in: *The New Superconducting Electronics*, H. Weinstock and R. Ralston (eds.), Kluver, Dordrecht, pp. 423-452 (1993).
- Lytkin, V., Finney, B., Alepko, L., "Tsiolkovsky, "Russian Cosism and Extraterrestrial Intelligence", *Q. J. R. Astr. Soc.***36**:369-376 (1995).
- McEachern, R. H., *Human and Machine Intelligence: An Evolutionary View*, R & E Publishers, 1993.

McKendree, Tom, "Short, Medium, and Long-Term Space Development Using Nanotechnology", Presentation at the Foresight Institute's1998 Senior Associates Meeting.
- McKendree, Tom, "The Logical Core Architecture", Foresight Institute Fifth Foresight Conference on Molecular Nanotechnology, 1997.
- McKendree, Tom, "Implications of Molecular Nanotechnology Technical Performance Parameters on Previously Defined Space System Architectures," Foresight Institute, Fourth Foresight Conference on Molecular Nanotechnology, 1995.
- Merkle, R. C., Presentation at the Foresight Institute's 1998 Senior Associates Meeting.
- Merkle, R. C., "Energy Limits to the Computational Power of the Human Brain," Foresight Update 6 (1989).
- Merkle, R. C. & Drexler, K. E., Helical Logic, *Nanotechnology* **7**:325-339 (1996).
- Minsky, M., comments on pg. 190 in "Astroengineering Activity: The Possibility of ETI in Present Astrophysical Phenomena", from *Communication with Extraterrestrial Intelligence: CETI*, Proceedings of the conference held at the Byurakan Astrophysical Observatory, Yerevan, USSR Sept. 5-11, 1971, C. Sagan (ed.), MIT Press (1973).
- Moravec, Hans, "Mind Children: The Future of Robot and Human Intelligence," Harvard University Press (1987).
- Moravec, Hans, "When will computer hardware match the human brain?," *Journal of Transhumanism*. 1998. Vol. **1** (Received Dec. 1997).
- Newman, W. I., and C. Sagan, "Galactic civilizations: Population dynamics and interstellar diffusion", *Icarus* **46**:293-327 (1981).
- Niven, L, *Ringworld* (1970), *Ringworld Engineers* (1980), *Ringworld Throne* (1996).
- Pakkenberg, B. & Gundersen, H. J. G., "Neocortical Neuron Number in Humans: Effect of Sex and Age", *J. of Comparative Neurology* **384**:312-320 (1997); Neocortical neurons: 19.3 billion in females, 22.8 billion in males.

- Partain, L. D. (ed.), *Solar Cells and Their Applications*, Wiley, 1995.
- Potter, S. D. & Matloff, G. L., "Light Sail Propulsion Using Thin-Film Photovoltaic Technology", *JBIS* **49**:329-334 (1996).
- Sandberg, Anders, "The Physics of Information Processing Superobjects: Daily Life Among the Jupiter Brains", *Journal of Transhumanism* **5** (1999). Preprint was available in 1997.
- Semiconductor Industry Association, 1997 Roadmap
- Smith, B. J., Harvey, P. M., "High Spatial Resolution KAO Far-Infrared Observations of the Central Regions of Infrared-bright Galaxies", *Astrophysical Journal* **468**:139 (Sep 1996);

  See also: http://casa.colorado.edu/~beverly/research/FIR.html
- Taubes, G, "The Star Machine", Discover **18**(June, 1997):76-83.
- Tuckerman, D. B., *Heat-Transfer Microstructures for Integrated Circuits*, Ph. D. Thesis, Stanford University (1984) [@UMI]
- Ultralight Optics: http://snoopy.gsfc.nasa.gov/~lunartel/lun7.html
- Villasenor, J. and Mangione-Smith, W., "Configurable Computing", *Scientific American* (June, 1997)
- Von Neumann, J., *Theory of Self-Reproducing Automata*, Banks, A. W. (ed.), University of Illinois Press, Urbana (1966).

Created: July 21, 1999; Derived from a discussion of Jupiter Brains created ~November 1997
Last Modified: August 16, 2004

Author: Robert Bradbury