Self-replicating probes for galactic exploration.

Antonia J. Jones

Department of Computing, Imperial College London

 This paper was a talk given at Cranwell in 1991. Note: Images have now been scanned and inserted with due acknowledgements.

Self-replicating probes for galactic exploration.

CONTENTS

List of figures

Foreword.

The invention of agriculture, language and writing were the prelude to the rise of technology. Initially development was slow but with the wisdom of hindsight we can see that the growth was near exponential. Moreover, it was also *irreversible* in the sense that technology once developed as part of the economic infrastructure of the human race is unlikely to be voluntarily abandoned; without it the vast bulk of the human population could no longer survive. Since exponential plundering of the Earth's natural resources will inevitably lead to global degradation of the environment and eventually to catastrophe, three conclusions follow:

- If technology cannot be abandoned it must be utilised to ensure the long term survival of both the Earth's ecosphere and the human population.

- The human population of planet Earth must be stabilised.

- Further population growth should take place off planet.

The time scale for these proposals may be shorter than is generally appreciated, the current global population exceeds 5000×10^6 and despite all efforts to control population it continues to grow at an alarming rate.

It is not currently in vogue to have technological dreams. We live in the most amazing of times but in economic terms the western world of the 1980's has been a rather sordid decade of short term planning mostly characterised by greed. Nevertheless, dreams, analysis and long term commitment are positively required of us if the future of our descendants is not to be "poor, nasty, brutish and short" [Hobbes 1651].

The exploitation of the solar system will be but a further prelude to the diaspora - the eventual goal will be the stars themselves.

Per Ardua Ad Astra

Quiz

Who can tell me the significance of the following two things?

Date: 9 November 1991.

Picture: IBM Logo.

Self-replicating probes for galactic exploration

Introduction.

- The descendants of Mankind will not forever remain in the cradle of the solar system. Eventually, if we survive the multiple hazards of a young technological civilisation, there will inevitably come the great diaspora of our descendants, who (according to [Moravec 1988]) may not even be human by our present standards.

Freeman Dyson [Dyson 1979] gives three reasons why, quite apart from scientific considerations, mankind needs to travel in space.

> "The first reason is garbage disposal; we need to transfer industrial processes into space so that the Earth may remain a green and pleasant place for our grandchildren to live in. The second reason is to escape material impoverishment; the resources of this planet are finite, and we shall not forgo forever the abundance of solar energy and minerals and living space that are spread out all around us. The third reason is our spiritual need for an open frontier." [pp 116 - 117]

The idea of self-replicating star probes is quite simply this - a single probe is constructed and dispatched to a nearby star system. It surveys the system in an intelligent and exhaustive manner, using a high powered laser, which may later serve as a relay, it conveys the results of its survey back to near Earth space, finally using the energy and available raw materials of the system the probe reproduces itself and dispatches its children onwards to repeat its mission in other stellar systems. It may remain or move out of the system depending on what it found.

SRPs (Self Replicating Probes) could be sent to other star systems to reproduce their own kind and spread. Each machine thus created may be immortal (limitlessly self repairing) or mortal. The SRP is so powerful a tool that it could have implications on a cosmological scale. With the SRP humanity could set in motion a chain reaction of organisation sweeping across the Universe at near light speed. This organised part of the Universe could itself be viewed as a higher level organism.

The technology to accomplish this will be third generation. The most probable scenario is extensive exploration and exploitation of the solar system before the first starship is built. The experience gained by developing fusion generators for space flight and self-replicating technologies to begin the terraforming of Mars can then be applied to the construction of a von Neumann probe.

The 1980 NASA report.

In 1980 NASA published *Advanced automation for space missions* an in depth study of the technologies required for

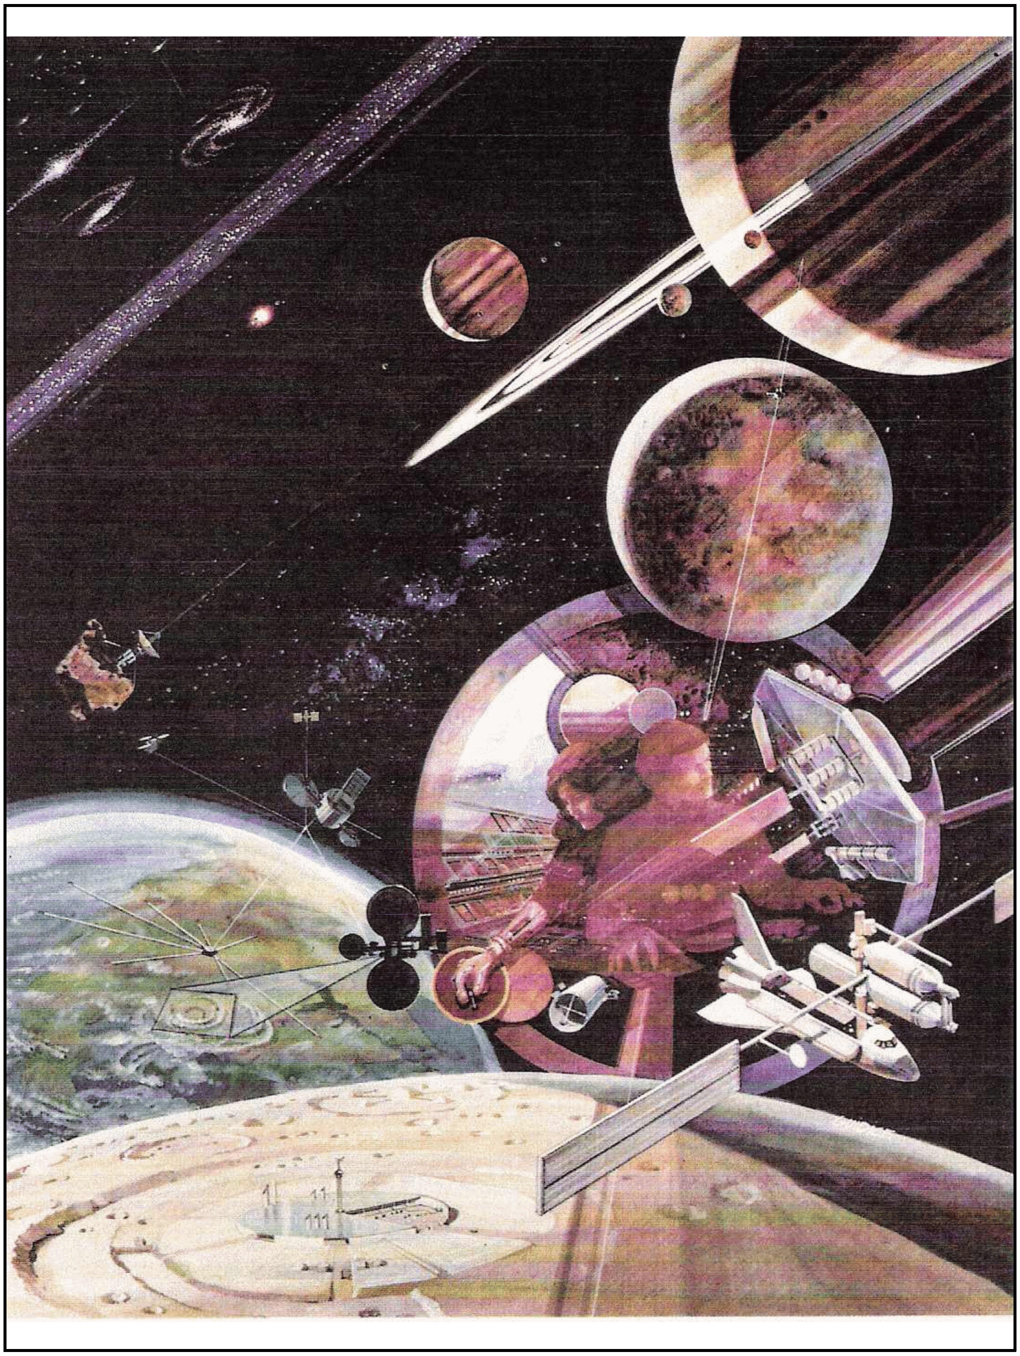

Figure 1 Painting by Rick Guidice. It captures the spirit of the space missions described by the NASA study.

- Intelligent Earth-sensing information systems.
- Non-Terrestrial utilisation of materials: Automated space manufacturing facilities.
- Replicating systems concepts: A self-replicating Lunar factory demonstrator.

- Exploitation of the solar system: A Titan demonstrator.
- Technology assessment of advanced automation for space missions.

The technologies.

Just how feasible is interstellar exploration based on the physics that we know and the technologies that we either have or can in principle develop? The answer is unequivocally positive - the precise mix of technologies is not yet clear but using a blend of several technologies already existing or proved feasible interstellar travel for robots and possibly humans is possible. The resources required are considerable but possible for a planetary based culture and easily within the capability of a solar system based culture.

There have been a number of suggestions for possible space drives. These include: nuclear pulse rockets [Orion, Daedalus], antimatter rockets, solar sails, the Bussard ramjet, and the vacuum energy drive. Many of these are reviewed in [Forward 1986].

Nuclear pulse rockets.

One of the oldest designs for an interstellar vessel is one that is propelled by nuclear bombs. The first design was called the *Orion* spacecraft and was discussed at Los Alamos National Laboratory in the late fifties. The original goal was to send manned spacecraft to Mars and Venus by 1968 at a fraction of the cost of the Apollo project. Freeman Dyson took these ideas for an interplanetary spacecraft and extrapolated them to an interstellar spacecraft [Dyson 1968]. The ship would necessarily be large, with a total initial mass of 1,600,000 metric tonnes and a payload of some 20,000 metric tonnes delivered to Alpha Centurai after a flight of 130 years. This would be sufficient to support several hundred crew members.

Antimatter rockets.

The antimatter rocket, with an ability to release "200%" of the fuel rest mass, can provide exhaust velocities close to the speed of light but there are several technological hurdles. The first is generation and confinement of sufficient quantities of antimatter. Most people are not aware of it but antiprotons are being produced and stored for days at a time in the particle physics facility at CERN in Switzerland, but only in infinitesimal quantities. The failure of a large quantity, antimatter containment system in an interstellar ship would result in the instant annihilation of the entire probe. But a reliable containment system could probably be designed. At present this is one of the more speculative possibilities for an interstellar drive.

The solar sail.

The most beautiful of the unorthodox methods of space travel is solar sailing. In principle it is possible to sail around the solar system using no engine at all. All you need is a huge gossamer-thin sail made of, say, aluminum-coated plastic film. You can trim and tack wherever you want to go, balancing the pressure of sunlight on the sail against the force of the sun's gravity to steer a course, in the same way as the skipper of an earthly sailboat balances the pressure of the wind in his sails against the pressure of the water on his keel.

This idea was first suggested by the Russian pioneer of space travel, Konstantin Tsiolkovsky. A more recent elegant design for a solar sailboat was the heliogyro invented by Richard MacNeal. Macneal's sail is a twelve-pointed star rotating like the rotor of an autogiro airplane. In 1976 the Jet Propulsion Laboratory in Pasadena, California made a serious attempt with MacNeal to design an unmanned heliogyro ship that could be launched and flown in time to make a rendezvous with Halley's Comet when the comet came by the Earth in March 1986. This was a unique opportunity for the solar sail to prove itself. The space programme managers rejected the Halley's Comet mission as too risky; there was no proof-of-concept test results and the political consequences of a failed mission might have been disastrous to their whole programme.

Unfortunately, the nature and available acceleration of this rather elegant system renders it largely unsuitable for interstellar travel [Drexler 1980].

The Bussard Ramjet.

Consideration of mass ratios for various propulsion systems show that it would be highly desirable to acquire propulsion mass from interstellar space. The only available propellant is the interstellar matter. This is extremely rarefied with typically 1 - 2 atoms/cm^3, or even less. Since these atoms are primarily hydrogen this translates to a density of about 2 X 10^{-24} g/cm^3. In 1960 [Bussard 1960] published his classic paper on the interstellar ramjet. He proposed collecting the matter by ionising it and using an electromagnetic field to concentrate the matter in order to initiate fusion. Because concentrations of interstellar hydrogen are so dilute Bussard found that the scoop intake would have a diameter of around 100 kms.

Controlled thermonuclear fusion was achieved for several seconds on the 9 November 1991 at the Joint European Torus near Oxford, although the temperature reached was insufficient to make the reaction self-sustaining. It will probably take 20-30 years to get the first fusion based electricity generating stations on line but the important point is that it has been proved to be feasible. This, combined with recent advances in high temperature superconducting technology, makes the design and construction of a Bussard Ramjet at least a possibility in principle.

Fishback [Fishback 1969] examined in more detail the collection by a magnetic field and developed equations limiting the speed of the ramjet relative to the plasma, in [Martin 1971] some corrections were made to the numerical results. The limitation is not severe in regions with densities of 1-2 protons/cm^3 but at ten times this density the speed for an aluminium structure is limited to 0.94c and at 100 times to only 0.075c. Obviously the scoop can be considerably smaller in high density regions, but we shall not be able to take advantage from this. Moreover, such regions are not uncommon in the Galaxy [Martin 1972].

Other limitations of the ramjet are studied in detail in [Martin 1973], [Matloff 1977]. In addition the proton-proton reaction is difficult to sustain because of the small cross section [Whitmire 1975]. It may be possible to add a nuclear catalyst, such as carbon-12, which could be recovered, to speed up the reaction [Whitmire 1975] by a factor of 10^{18} which would make the ramjet project more feasible. The success at JET suggest that practical fusion is a possibility but it seems likely that the demonstration of such catalytic reactions will not be accomplished in the near future.

A lot of invention and research is needed, however, before the Bussard ramjet becomes a reality. The fusion reactor must be light-weight and long-lived, it must be able to fuse protons, not the easier to ignite mixture of deuterium and tritium. The reactor must be able to fuse incoming protons *without slowing them down*, or the frictional loss of bringing the fuel to a halt, fusing it, and reaccelerating the reaction products will put an undesirable upper limit on the maximum velocity obtainable.

Cassenti [Cassenti 1982] points out that even without sustained fusion there are variations of this idea e.g. the Ram-Augmented Interstellar Rocket (RAIR) [Bond 1974]. The RAIR uses an onboard nuclear reactor to heat and accelerate the plasma collected by an electromagnetic field. The RAIR cannot out perform the interstellar ramjet but is considerable more efficient than the nuclear pulse rocket.

Another variation is the Laser Powered ramjet (LPR) [Whitmire 1977] where the energy to accelerate the plasma is supplied by lasers in orbit about the Sun which in turn draw their energy from the Sun. Moving away from the beam the LPR is not as effective as the interstellar ramjet but moving toward the beam the LPR can be more effective.

The vacuum energy drive.

Physicists such as Wheeler suggest that quantum fluctuations in the energy of the vacuum occur throughout all space over submicroscopic scales of time and distance. Froning [Froning 1986] explores the possibility of exploiting such energies for propelling ramjet-like starships to the further stars. Although a system to do this is surely far beyond our present technologies it is nevertheless a fascinating possibility.

Computing and Artificial Intelligence.

The nature of the SRP concept implies that the probe is unmanned and this in turn requires a very advanced degree of automation, certainly beyond anything we might hope to currently engineer. It is arguable that at present the artificial intelligence requirement is probably the weakest link in the chain. Let us briefly review the current state of AI.

The spectacular success of computing often makes it hard to understand just how far away the attainment of human levels of intelligence in artificial systems may actually be. Broadly speaking we can classify approaches in AI into *symbolic-systems* and *connectionist-systems.*

Symbolic-systems or methods tend to rely on advanced programming in symbolic languages such as LISP or PROLOG and in specifically inserting the knowledge required. Examples are: GENERAL PROBLEM SOLVER (Simon & Newell), EURISKO (Lenat) which had some learning capability, and more generally Expert systems.

This approach is characteristic of classical AI. It has had notable success in areas such as Expert Systems, where expert knowledge in a limited domain is programmed into the system in the form of rules. But these systems are brittle: they work well on problems within their limited domain of expertise but have no realistic 'world model', so fail spectacularly if the slightest piece of 'commonsense' is required for the solution of the problem. They are very time consuming to construct, with problems of maintaining consistency as rules are added, and difficult to update if some of the rules change. Currently Lenat at MCC is engaged in a major project to create a realistic world model by programming over a million rules into a system while developing a host of techniques to maintain consistency. This is an immensely interesting project (indeed it may make or break classical AI) but there are many who think that the entire approach is misguided.

In any event we should consider just how much computing power is likely to be necessary for human-level intelligence.

How can one measure the overall computational power of an information processing system? There are two obvious aspects we should consider. Firstly, information storage capacity - a system cannot be very smart if it has little or no memory. On the other hand, a system may have a vast memory but little or no capacity to manipulate information; so a second essential measure is the number of binary operations per second. On these two scales illustrates the information processing capability of some familiar biological and technological information processing systems.

Considering each axis independently, research in neurophysiology has revealed that the brain and central nervous system consists of about 10^{11} individual parallel processors, called neurons. Each neuron is capable of storing about 10^{4} bits of information. The information capacity of the brain is thus about 10^{15} bits. Much of this information is probably redundant but using this figure as a conservative estimate let us consider when we might expect to have high-speed

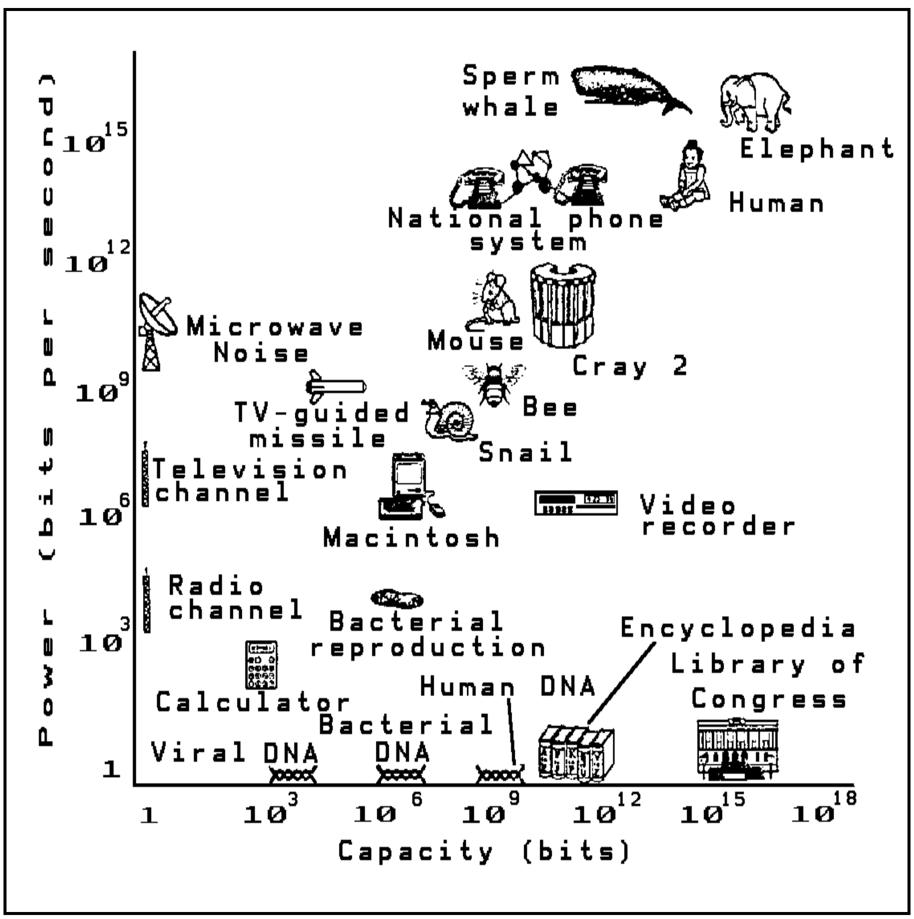

Figure 2 Information processing capability [From: *Mind Children*, Hans Moravec].

memories of 10^{15} bits. The following argument, offered by Bock, shows that the amount of high-speed random access memory that may be conventionally accessed by a large computer has increased by an order of magnitude every six years. If we can trust this simple extrapolation, in generation thirteen, AD 2024-30, the average high speed memory capacity of a large computer will reach 10^{15} bits.

Now consider the evolution of technological processing power. Remarkably, the corresponding graph for processing power follows much the same trend. Of course, the real trick is putting the two together to achieve the desired result, it seems relatively unlikely that we shall be in a position to accomplish this by 2024.

So much for the hardware. Now consider the software. Even adult human brains are not filled to capacity. So we will assume that 10% of the total capacity, i.e. 10^{14} bits, is the extent of the 'software' base of an adult human brain. How long will it take to write the programs to fill 10^{14} bits (production rules, knowledge bases etc.)? The currently accepted rate of production of software, from conception through testing, de-bugging and documentation to installation, is about one line of code per hour. Assuming, generously, that an average line of code contains approximately 60 characters, or 500 bits, we discover that the project will require 100 million person years! We'll never get anywhere by trying to program human intelligence into a machine.

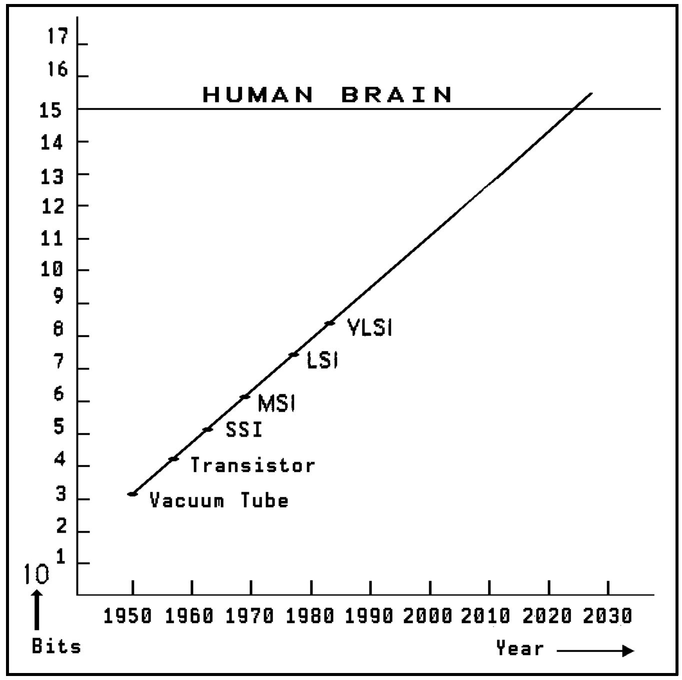

Figure 3 Storage capacity.

What other options are available? One is direct transfer from the human brain to the machine. Considering conventional transfer rates over a high speed bus this would take about 12 days. The only problem is: nobody has the slightest idea how to build such a device.

What's left? There must be another alternative, because intelligence is acquired every day. Every day babies are born, grow up, and in the process acquire a full spectrum of intelligence. How do they do it? The answer, of course, is that they learn.

If we assume that the eyes, our major source of sensory input, receive information at the rate of about 250,000 bits per second, we can fill the 10^{14} bits of our machine's memory capacity in about 20 years. Now storing sensory input is not the same thing as developing intelligence, however this figure is in the right ball park. Maybe what we must do is connect our machine brain to a large number of high-data-rate sensors, endow it with a comparatively simple algorithm for self organization, provide it with a continuous and varied stream of stimuli and evaluations for its responses, and let it learn.

This argument may seem cavalier in some aspects. The human brain is highly parallel and somewhat inhomogeneous in its architecture. It does not clock at high serial speeds and does not access RAM to recall information for processing. The storage capacity may be vastly greater than the estimated 10^{15} bits, since each neuron is connected to as many as 10,000 others and the structure of these interconnections may also store information. Indeed, although we do not know a great deal about the mechanisms of human memory, we do know that is multi-levelled with partial bio-chemical storage. However, none of this invalidates Bock's point.

All this leads to an approach to AI very different from formal rule based systems, the *connectionist approach* based on modelling biological information processing systems. Although still very much in its infancy the connectionist approach has had a number of successes, particularly in pattern recognition and non-linear control. But there is a long way to go before we can expect such systems to begin to display intelligence.

In any event although machines constructed along such lines may exhibit intelligence they will be *truly alien*.

Humans are mammals and mature through a process of socialisation. The child is taken by the hand and experiences the world described, in a moderately consistent fashion, by adults. Eventually the child acquires a similar model of reality and thus gains membership of a cultural linguistic group.

The idea that an intelligent, conscious machine can be constrained to behave according to Asimov's Laws, or some similar prescription, seems to me to be mistaken. The problem with this entire approach is that any machine sufficiently sophisticated to engage in reproduction in largely unstructured environments and having, in general, the capacity for survival must be able to modify it's behaviour - a self-reprogramming capability. Such a SRP in theory may be able to "program around" any normative rules of behaviour gifted it by its creators. Similarly, hard wiring the constraints can be circumvented either by clever software or by simply redesigning and replacing the offending module. It would therefore appear impossible, as with people, to absolutely *guarantee* the 'good' behaviour of any intelligent, conscious SRP.

On the other hand we might argue that also, like people, machines can be taught 'right' and 'wrong', as well as the logical rationales for various codes of behaviour. Possibly they might be expected to remain even more faithful to a given moral standard than people. But what of succeeding generations and how do we accommodate these expectations to the alien nature of such a machine?

Self-replicating automata.

In September 1948 John von Neumann gave a lecture entitled *The General and Logical Theory of Automata*, which is reprinted in Volume 5 of his collected works. Because he spoke in general terms, there is very little in it that is dated. Von Neumann's automata are a conceptual generalization of the electronic computers whose revolutionary implications he was the first to see. He suggested models closely related to the abstract definition of a Finite-State-Machine (FSM). The main theme of the 1948 lecture is an abstract analysis of the structure of an automaton which is of sufficient complexity to be able to reproduce itself (the summary here is abstracted from [Dyson 1979]). He shows that a self-reproducing automaton must have four separate components:

> **Component A**. Is an automaton which collects raw materials and processes them into an output specified by a written instruction which must be supplied from outside. In effect an automatic factory.
>
> **Component B**. Is a duplicator, an automaton which takes a written instruction and copies it.
>
> **Component C**. Is a controller, an automaton hooked up to both A and B. When C is given an instruction it first passes the instruction to B for duplication, then passes it to A for action, and finally supplies the copied instruction to the output of A whilst keeping the original for itself.
>
> **Component D**. Is a written instruction, a program, containing the complete specifications which cause A to manufacture the combined system, A plus B plus C plus D.

Von Neumann's analysis showed that a structure of this kind was logically necessary and sufficient for a self-reproducing automaton, and he conjectured that it must also exist in living cells. Whilst von Neumann did not live long enough to bring his theory of automata into existence he did live long enough to see his insight into the functioning of living cells brilliantly confirmed by the biologists. Five years later Crick and Watson discovered the structure of DNA, and now many children in high school learn the biological identification of von Neumann's four components.

> **Component A**. Is the ribosomes.
> **Component B**. Is the enzymes RNA and DNA polymerase.
> **Component C**. Is the repressor and derepressor molecules and other items whose functioning is still imperfectly understood.
> **Component D**. Is the genetic materials, RNA and DNA.

So far as we know the basic design of every microorganism larger than a virus is precisely as von Neumann said it should be. Viruses are not self-reproducing in von Neumann's sense since they borrow the ribosomes from the cells they invade.

Von Neumann's first main conclusion was that self-reproducing automata with these characteristics can in principle be built. His second main conclusion follows from the work of Turing and, whilst it is less well known, goes deeper into the heart of automata theory. He showed that there exists in theory a *universal automaton*, that is to say a machine of a certain definite size and complication, which, if you give it the correct written instructions, will do anything that any other machine can do! So beyond a certain point, you don't need to make your machine any bigger or more complicated to get more complicated jobs done. All you need is to give it longer and more elaborate instructions. You can also make the universal automaton self- reproducing by including it within the factory unit (Component A). Von Neumann believed that the possibility of a universal automaton was ultimately responsible for the possibility of indefinitely continued biological evolution. In evolving from simpler to more complex organisms you do not have to redesign the basic biochemical machinery as you go along. You have only to modify and extend the genetic instructions. Everything we have learned about evolution since 1948 tends to confirm that von Neumann was right.

Thus there is no theoretical problem with self-replication, just a problem of implementation. If we can construct an AI we can probably equip it with the capability of replicating itself given raw materials, energy and sufficient time.

However, the implications of self-replicating machines go far beyond merely replicating an AI. Such a capability would revolutionise economics and offer a serious possibility of terraforming. The command and control of these technologies would have to be field tested within the solar system - there are serious ethical objections to turning lose self-replicating machines in other beings systems.

Genetic engineering and the genome project.

In the United States it is possible to patent a gene or new life form, plant or animal. In Europe the EC parliament are currently considering such legislation. The proponents of patenting argue that without this protection the research would not proceed because there would be no effective means to protect the huge investment required. We need to think about these issues carefully. There are potentially huge benefits for mankind but also qualitively new hazards. New genes released into the ecosphere are not analogous to other forms of ecological pollution, they can replicate, proliferate and mutate with unforseen and conceivably catastrophic consequences. No new technology is without risk but in this case the risk/benefit equation is very difficult to quantify.

The relevance of these issues to the SRP debate is clear. Using the short cut of modifying existing life forms genetic engineering is *already* producing self-reproducing 'technology'. The debate about ethics is no longer an abstract issue.

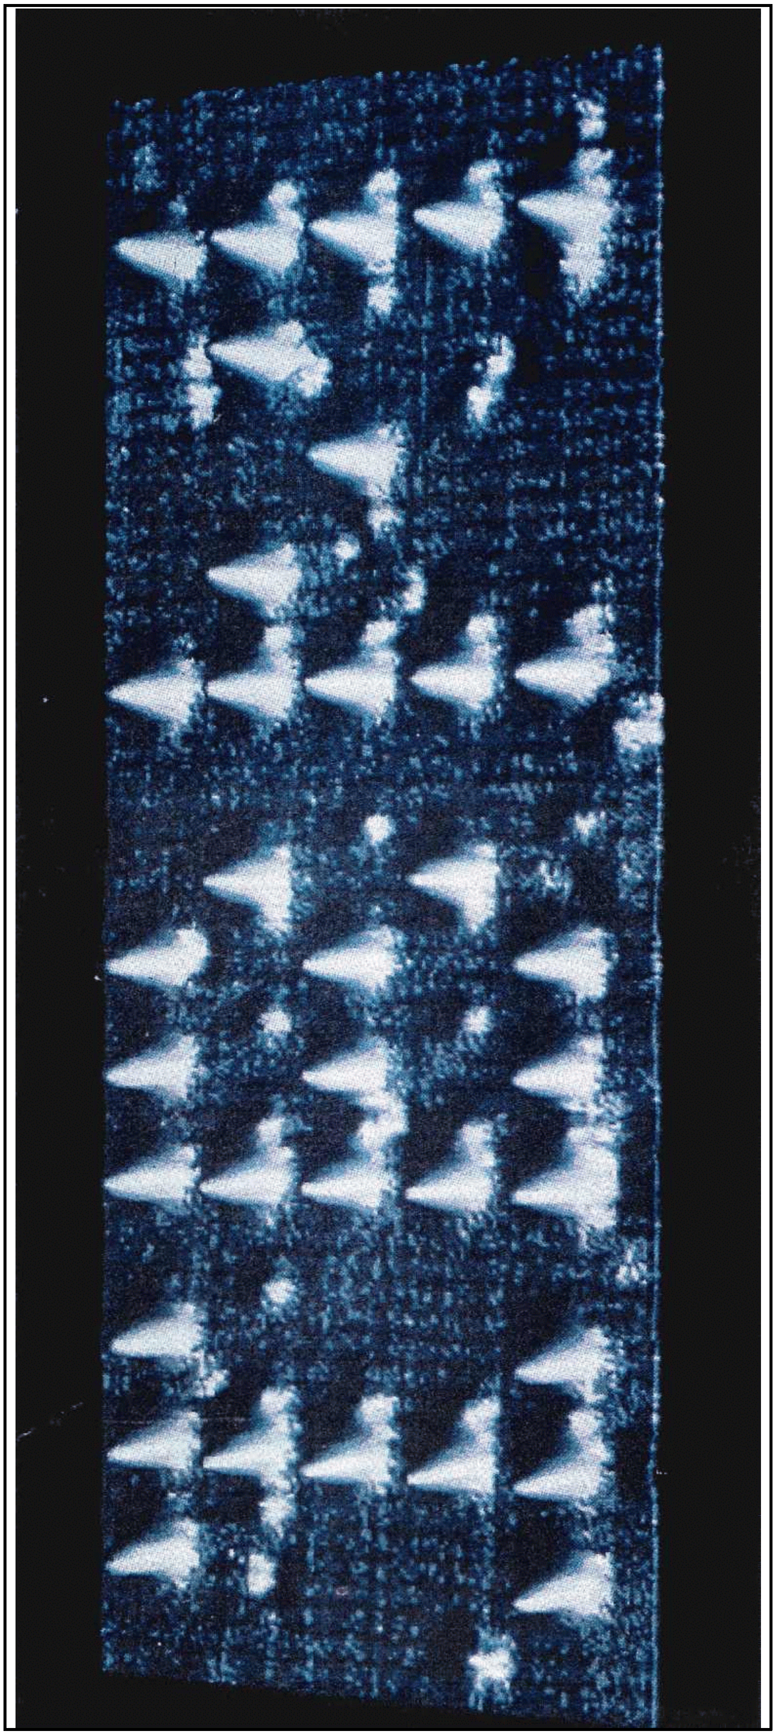

Figure 4 The letters IBM, shown here were the first structures ever assembled one atom at a time (magnification about 8 million).

The *Human Genome Project* is engaged in mapping the entire gene structure of human beings. Many of the pending patent applications relate to the commercial exploitation of *human* genes whose function has emerged from the genome project. Just where do we draw the line? Is it reasonable to have our own genes licensed by a patent holder? What life forms is it ethical to construct and what forms are not ethical?

We stand at a threshold here: for the first time in the biological history of the planet a life form can re-engineer itself. Just as with the earliest industrial revolution we cannot put the clock back and unknow what we have learnt. It is not knowledge itself which is intrinsically good or bad, knowledge only offers more choices. It is the choices which we make that define us as ethical beings.

Nanotechnology.

One intriguing set of possibilities for creating artificial life is based on Richard Feynman's suggestions in 1959 for ultraminiaturization and extension of our industrial manufacturing capabilities all the way down to the molecular level, see [Schneiker 1989].

Just as *Microtechnology* is engineering at the 1 micrometre = 10^{-6} m scale, so *Nanotechnology* is engineering at the 1 nanometre = 10^{-9} m scale. By extension *Picotechnology* is engineering at the 10^{-12} m scale and deals with the manipulation and modulation of individual atomic bonds and orbitals plus the special equipment where atomic or molecular measurements are made to sub-nanometre precision. There are already special cases where mechanical positioning and measurements of spacing between individual atoms has been pushed to a resolution of 1-10 picometres! See .

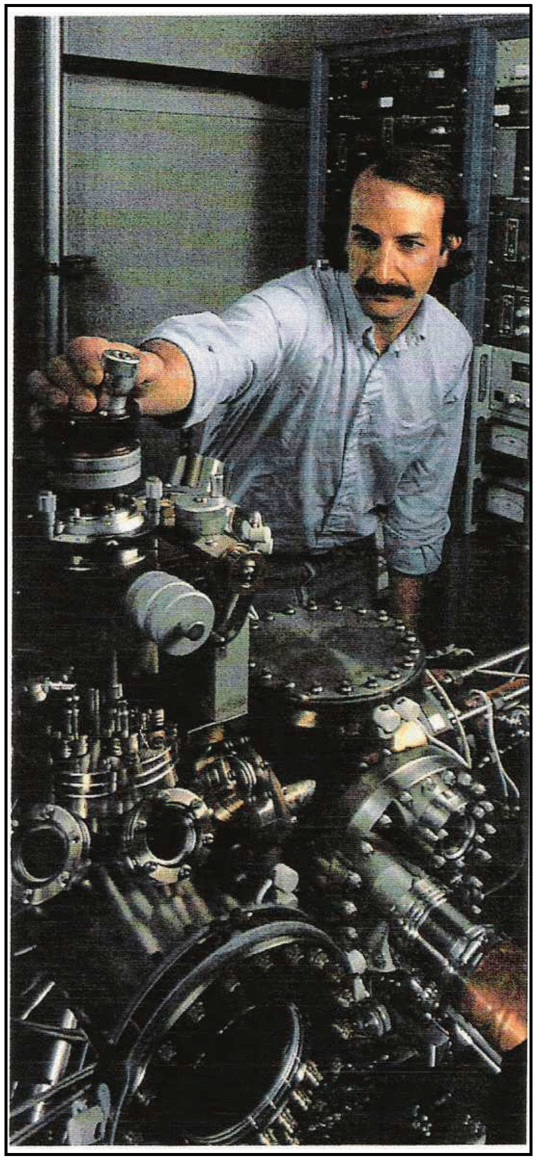

Figure 5 Dr. M. Eigier of IBM used an STM to position the atoms.

The instrument with this phenomenal capability, the *Scanning Tunnelling Microscope* (STM) was invented in 1981 at the IBM Zurich Research Laboratories by Binnig and Rohrer, who received a Nobel Prize in 1986, . The STM scans an ultra sharp conducting tip, using tungsten, for example, over a conducting or semiconducting surface. When the tip is within a few angstroms of the surface, a small voltage applied between the two gives rise to a tunnelling current of electrons, which depends exponentially on the tip-to-substrate separation (about an order of magnitude per angstrom). A servo system uses a feedback control that keeps the tip-to-substrate separation constant by modulating the voltage across a piezoelectric positioning system. As the tip is scanned across the surface, variations in the voltage, when plotted, correspond to surface topography. As an added advantage STMs can function over a wide range of temperature and pressure, and in liquid as well as gas environment. With STM-derived technology, the capabilities for building Feynman nanomachines reduces the originally proposed "building down" sequence to *just one step*.

The following possibilities involve nanotechnology and were all first anticipated over two decades ago. Picking up and moving individual atoms and molecules, building miniature telerobotic surgeons that operate on neurons, unlimited life span and reversal of aging using robot-directed, molecular level repair, boundless wealth, defenses against virtually all diseases and virtually unlimited *in situ* evolution, super compact personal computer databases that could easily store the Library of Congress etc. The possibilities seem endless and even a small fraction of the potential of nanotechnology is realised we can, in principle, be anything we want to be.

In a world where just about anything is possible day to day prediction becomes impossible - humanity will have entered the singularity.

Some ethical questions.

On AIs.

According to former British Agriculture Minister Peter Walker:

> "Uniquely in history, we have the circumstances in which we can create Athens without the slaves."

However, if robots attain intelligence, sensitivity, and the ability to replicate, might they not be considered legal persons, hence slaves? Is mankind creating a new form of life to enthral to its bidding? Is it immoral to subjugate an entity of one's own creation? Modern jurisprudence, in contrast to ancient Roman law, does not permit a parent to enslave his or her child. Questions of "machine rights", "robot liberation", the "artificial life liberation front" are future possibilities. If the intelligence or sensitivity of robots ever exceeds that of humankind, ought we to grant them civil rights equal or superior to our own?

In the long term, there are two possibilities for the future of mankind: Either we are a biological waystation in the evolutionary scheme of things, or else we are an evolutionary dead end. Replicating systems technology or artificial life would offer humanity other options for continued and fruitful evolution. At present, machines reproduce themselves but only with human help. Such mutualism is commonplace in biology, indeed it is arguably the very existence of a rich ecosphere which makes the continuation of any species possible.

This issue has tremendous importance to the question of human survival and long term evolution. Of course, if we succeed in creating artificial life we shall find ourselves co-inhabiting the universe with an alien race of beings. But the ultimate outcome is unknown: we could dominate them, they could dominate us, we could co-exist as separate species, or we could form a symbiotic relationship. This last is the most exciting possibility.

On SRPs.

A number of fundamental but far reaching ethical issues are raised by the possible existence of replicating machines in the Galaxy. For example, is it morally right, or equitable, for a self-reproducing machine to enter a foreign solar system and convert part of that system's mass and energy to its own purposes? Does an intelligent race legally "own" its home sun, planets, asteroidal materials, moons, solar energy, solar wind, and comets? Does it make a difference if the planets are inhabited by intelligent beings, and if so, is there some lower threshold below which a system may ethically be "invaded" or expropriated? If the sentient inhabitants lack advanced technology, or if they have it, should this make any difference in our ethical judgement of the situation?

Why is it so quiet out there?

There is an extensive literature on the likelihood of other life in the Galaxy. Freeman Dyson, however, rejects as worthless all attempts to calculate from theoretical principles the frequency of occurrence of (intelligent) life forms in the universe. He says

> "Our ignorance of the chemical processes by which life arose on Earth makes calculations meaningless. Depending on the details of the chemistry, life may be abundant in the universe, or it may be rare, or it may not exist at all outside our own planet. Nevertheless, there are good scientific reasons to pursue the search for evidence of intelligence with some hope for a successful outcome. The essential point which works in our favour as observers is that we are not required to observe the effects of an average intelligent species. It is enough to observe the effects of the most spendthrift, the most grandiosely expansionist, or the most technology-mad society in the universe." [Dyson 1979] [pp 209-210]

For myself I find the anthropic principle hard to accept and so am forced to give serious consideration to the Zoo hypothesis. However, science is not, in the final analysis, about opinions it is about objective fact and at present there are no verifiable facts concerning the existence of extraterrestrial life.

The Russian astronomer Kardashev suggested that civilisations in the universe should fall into three distinct types.

A type 1 civilisation controls the resources of a planet.

A type 2 civilisation controls the resources of a star.

A type 3 civilisation controls the resources of a galaxy.

We have not yet achieved type 1 status but will probably do so within a hundred years or so. The difference in size and power between type 1 and type 2, or between type 2 and type 3, is a factor of the order of 10^{10}, unimaginably large by present human standards. But the process of exponential growth allows this immense gulf to be bridged remarkably rapidly. To grow by a factor of 10^{10} takes 33 doubling times. A society growing at the modest rate of 1% per year will make the transition from type 1 to type 2 in less than 2,500 years. The transition to type 3 would take longer, since it requires interstellar travel, but the periods of transition are likely to be comparatively brief episodes in the history of any long-lived society. Hence Kardashev concludes that if we ever discover an extra-terrestrial civilisation, it will probably belong clearly to type 1, 2 or 3 rather than to one of the brief transitional phases.

Detectability of alien civilisations

Type 1. Undetectable at interstellar distances except by radio or directed high energy laser. Following the original suggestion of Cocconi and Morrison this is the method of search that our radio astronomers have followed for the last thirty years.

Type 2. May be a powerful radio source or it may not (depends on the life style of its inhabitants) but there is one source of emission which a type 2 civilisation cannot avoid making. By the second law of thermodynamics, a civilisation which exploits the total energy output of a star must radiate away a large fraction of this energy in the form of waste heat - infrared radiation at around ten microns (about twenty times the wavelength of visible light) [Dyson 1959]. Up to now, the infrared astronomers have not found any objects that arouse suspicions of artificiality.

Type 3. A type 3 civilisation should produce emissions of radio, light and infrared radiation with an apparent brightness comparable with those of a type 2 civilisation in our own Galaxy. In particular a type 3 civilisation should be detectable as an extra-galactic infrared source. However, since we are rather vague about the nature of type 3 civilisations and our knowledge of galactic evolution is weaker than our knowledge of stellar evolution a type 3 civilisation might be harder to spot than a type 2.

It has been pointed out [Oliver 1975] that the number of intelligent races that have existed in the past may be significantly greater than those presently in existence. Specifically, at this time there may exist perhaps only 10% of the alien civilisations that have ever lived in the Galaxy (type 1 or type 2) - the remaining 90% having become extinct. If this is true, then 9 out of every 10 replicating machines we might encounter in the solar system could be emissaries from long-dead cultures.

If we do in fact find such machines and are able to interrogate them successfully, we may become privy to the doings of incredibly old alien societies long since perished. These societies may lead to many others, so we may be treated, not just to a marvellous description of the entire biology and history of a single intelligent race, but also to an encyclopedic travelogue describing thousands or millions of other civilisations known to the creators of the probe we are examining. Probes will likely contain at least an edited version of the sending race's proverbial *Encyclopedia Galactica*, because this information is essential if the probe is to make the most informed and intelligent autonomous decisions during its explorations.

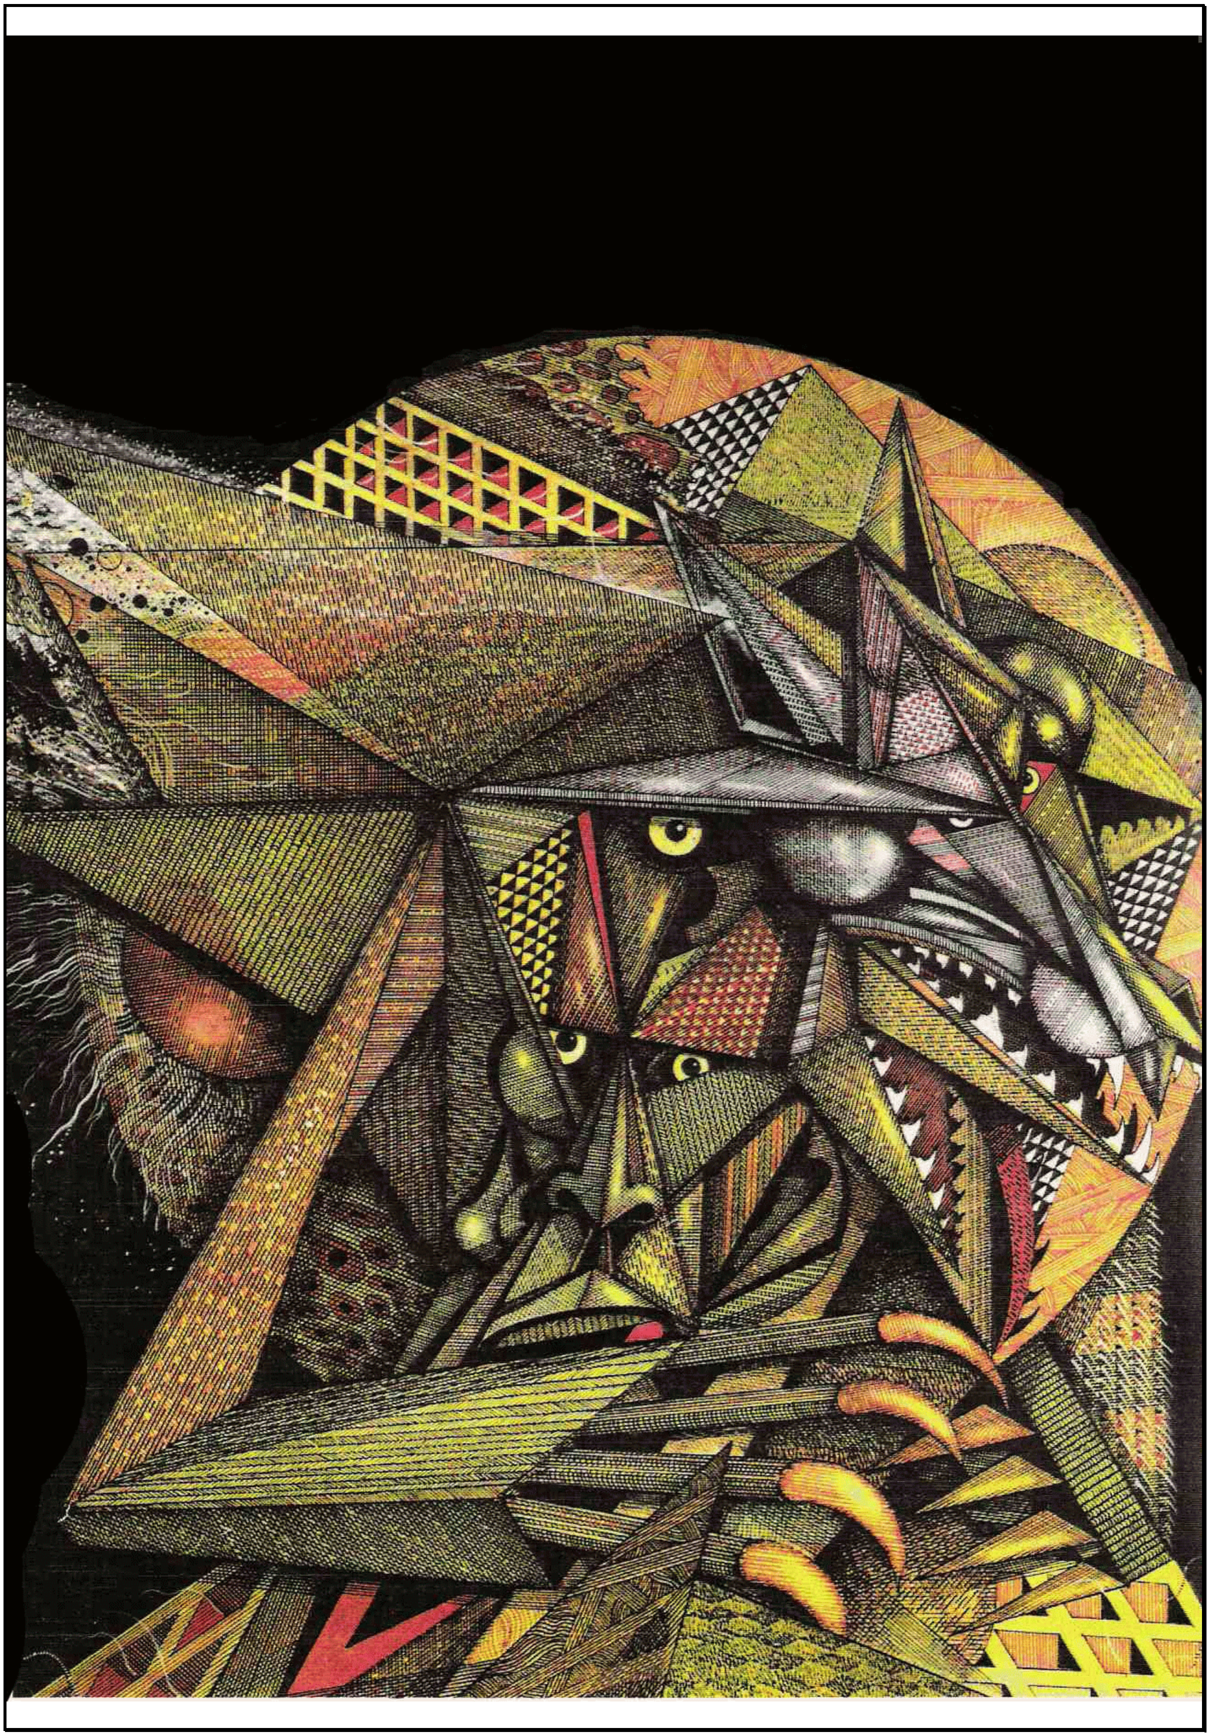

Figure 6 *The Werewolf Principle* by Cliford D. Simak (Pan). Cover by Ian Miller.

If we ever encounter extraterrestrials they are liable to be *very* different, see Figure 6.

References

[Bock 1985] P. Bock. *The Emergence of Artificial Intelligence: Learning to Learn*. AI Magazine, Vol. 6, No. 3, Fall 1985, pp. 180-190.

[Bond 1974] A. Bond. *An analysis of the potential performance of the ram-augmented interstellar rocket*. Journal of the British Interplanetary Society **27**, 674-685, 1974.

[Bussard 1960] R.W. Bussard. *Galactic matter and interstellar flight*. Astronautica Acta **6**, 179-194, 1960.

[Cassenti 1980] B. N. Cassenti. *A comparison of interstellar propulsion systems*. AIAA paper 80-1229, AIAA/SAE/ASME 16th Joint Propulsion Conf., Hartford, Conn., 30 June - 2 July 1980.

[Cassenti 1982] B. N. Cassenti. *A comparison of interstellar propulsion methods*. Journal of the British Interplanetary Society **35**, 116-124, 1982.

[Drexler 1980] K. E. Drexler. *Lightsails - An overview*. In discussion meeting on Gossamer Spacecraft (Ultralightweight Spacecraft): Final report, Ed. R. G. Brereton, JPL Publications 80-26, Jet Propulsion Laboratory, Pasadena CA. pp 8-12, 15 May 1980.

[Dyson 1959] Freeman J. Dyson. *Search for artificial stellar sources of infrared radiation.* Science **131**, pp 1667-1668, 1959.

[Dyson 1968] Freeman J. Dyson. *Interstellar transport*. Physics Today. October, 41-45, 1968.

[Dyson 1979] Freeman J. Dyson. *Disturbing the universe*. Harper & Row, New York, 1981.

[Fishback 1969] J. F. Fishback. *Relativistic interstellar spaceflight*. Astronautica Acta **15**, 25-35, 1969.

[Forward 1986]. R. L. Forward. *Feasibility of interstellar travel: a review*. Journal of the British Interplanetary Society **39**, 379-384, 1986.

[Freitas 1980a] R. A. Freitas. *Interstellar probes: A new approach to SETI*. Journal of the British Interplanetary Society **33**, pp 95-100, March 1980.

[Freitas 1980b] R. A. Freitas. *A self-reproducing interstellar probe*. Journal of the British Interplanetary Society **33**, pp 251-264, July 1980.

[Freitas 1980c] R. A. Freitas. *Xenology: An introduction to the scientific study of life, intelligence, and civilisation in the universe.* Xenology Research Institute, 1980.

[Froning 1986] H. D. Froning. *Use of vacuum energies for interstellar space flight*. Journal of the British Interplanetary Society **39**, 410-415, 1986.

[Hobbes 1651] T. Hobbes. *Leviathan*, Published by Andrew Crooke, London 1651.

[Kahn 1976] H. Kahn, W. Brown and L. Martel. *The next 200 years*. William Morrow & Co., New York 1976.

[Martin 1971] A. R. Martin. *Structural limitations on interstellar spaceflight*. Astronautica Acta **16**, 353-357, 1971.

[Martin 1972] A. R. Martin. *The effects of drag on relativistic spaceflight*. Journal of the British Interplanetary Society **25**, 643-653, 1972.

[Martin 1973] A. R. Martin. *Magnetic intake limitations on interstellar ramjets*. Astronautica Acta **18**, 1-10, 1973.

[Matloff 1977]. G. L. Matloff. *Interstellar applications and limitations of several electrostatic/electromagnetic ion collection techniques*. Journal of the British Interplanetary Society **30**, 213-222, 1977.

[Moravec 1988] H. Moravec. *Mind Children: The future of robot and human intelligence*. Harvard University Press, Cambridge, Mass. 1988.

[NASA 1980] Eds. R. A. Freitas and W. P. Gilbreath. *Advanced automation for space missions*. NASA Conference Publication 2255. Proceedings of the 1980 NASA/ASEE summer study, University of Santa Clara, CA. July 23 - August 29, 1980. For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington DC 20402.

[O'Neill 1977] Gerard K. O'Neill. *The high frontier, human colonies in space*. William Morrow and Co., New York, 1977.

[Oliver 1975] B. N. Oliver. *Proximity of Galactic civilisations*. Icarus **25**, 360-367, June 1975.

[Parkinson 1974] Bob Parkinson. *Starship as third generation technology*. Journal of the British Interplanetary Society **27**(4) pp 295-300, April 1974.

[Schneiker 1989] C. Schneiker. *Nanotechnology with Feynman machines: scanning tunneling engineering and artificial life*. In *Artificial Life*, Ed C. G. Langton, Santa Fe Institute, Addison-Wesley, 443-500, ISBN 0-201-09356-1, 1989.

[Whitmire 1975] D. P. Whitmire. *Relativistic spacecraft and the catalytic nuclear ramjet*. Astronautica Acta **2**, 497-509, 1975.

[Whitmire 1977] D. P. Whitmere and A. Jackson. *A laser powered interstellar ramjet*. Journal of the British Interplanetary Society **30**, 223-226, 1977.

00382339 E.I. Monthly No: EI7407044623
Title: STARSHIP AS THIRD GENERATION TECHNOLOGY.
Author: Parkinson, Bob
Source: Journal of the British Interplanetary Society v 27 n 4 Apr 1974 p 295-300
Publication Year: 1974
CODEN: Journal of the British Interplanetary Society AW ISSN: 0007-084X
Language: ENGLISH
Journal Announcement: 7407

Abstract: By its nature an interstellar vehicle will use 'third generation' technology based on the research and experience gained during the next generation of space flight. Consideration of what may be possible by this time, together with some physical limitations, favors continuous thermonuclear propulsion, a laser-riding electric propulsion system or possibly a Bussard interstellar ramjet as the propulsion system for a starship. A review of progress to date in space flight is given and the potential for further advances through research emphasized.

9 refs.
Descriptors: *SPACE FLIGHT; ROCKETS AND MISSILES
Classification Codes: 656 (Space Flight); 654 (Rockets & Rocket Propulsion) 65 (AEROSPACE ENGINEERING)