Medical Hypotheses

Medical Hypotheses (1992) 39, 6–16

The Technical Feasibility of Cryonics

R. C. MERKLE

Xerox PARC, 3333 Coyote Hill Road, Palo Alto, CA 94304, USA

Abstract—Cryonic suspension is a method of stabilizing the condition of someone who is terminally ill so that they can be transported to the medical care facilities that will be available in the late 21st or 22nd century. There is little dispute that the condition of a person stored at the temperature of liquid nitrogen is stable, but the process of freezing inflicts a level of damage which cannot be reversed by current medical technology. Could this damage be reversed by future technology? We consider the limits of what medical technology should eventually be able to achieve (based on currently understood chemistry and physics) and whether the repair of frozen tissue is within those limits.

Introduction

There is little doubt that tissue preserved in liquid nitrogen can survive centuries without deterioration. This provides an imperfect time machine that can transport us from the present to the future: we need merely freeze ourselves in liquid nitrogen. If freezing damage can someday be cured, then we can travel to the era when the cure is available. This option is in fact open to anyone who so chooses. First seriously proposed in the 1960s by Ettinger (1) there are now three organizations in the US that provide cryonic suspension services.

Perhaps the most important question in evaluating cryonics is its technical feasibility: will it work?

The reader interested in a general introduction to cryonics is referred to other sources (1, 2, 3). A more detailed version of the present paper is available from the author.

Many isolated tissues (and a few particularly hardy organs) have been successfully cooled to the temperature of liquid nitrogen and rewarmed (4), so freezing is not a process of total destruction. However, the damage done by freezing is beyond the self-repair and recovery capabilities of the issue itself. Will some future technology be able to reverse freezing damage? The laws of physics and chemistry as they apply to biological structures are well understood and well defined, so determining whether the repair of frozen tissue will (or will not) eventually prove feasible is a question which we should, in principle, be able to answer today.

Before we can decide whether future medical technology can repair freezing injury, we must consider what fundamental limits constrain such technologies.

Human tissue and human beings are made of atoms. Whether a person is healthy or ill, alive or dead, depends entirely on the arrangement of those atoms. The fundamental purpose of medicine is to cure the ill and heal the sick. Put another way, the purpose of medicine is to change arrangements of

Date received 13 January 1992
Date accepted 3 April 1992

atoms that are 'unhealthy' to arrangements of atoms that are 'healthy'.

Phrased in this fashion, it is obvious that the limits of future medical technology depend on the limits of our ability to control the structure of matter. The better our tools for controlling the structure of matter, the better our medical technology can be.

The general purpose ability to manipulate structures with atomic precision and low cost is often called *nanotechnology* (also molecular engineering, molecular manufacturing, molecular nanotechnology, etc). There is widespread belief that such a capability will eventually be developed (5–18) though exactly how long it will take is unclear.

We will briefly describe the basic concept of nanotechnology and give references for further reading. We will then examine the implications of this kind of technology for future medical capabilities, and the impact that those future medical capabilities will have on the repair of freezing injury.

Nanotechnology

The central thesis of nanotechnology is that almost any chemically stable structure that can be specified can in fact be built. This possibility was first advanced by Richard Feynman in 1959 (8) when he said: 'The principles of physics, as far as I can see, do not speak against the possibility of maneuvering things atom by atom', (Feynman won the 1965 Nobel prize in physics). Since then, there have been international conferences on the subject (16, 19) and a growing excitement within the research community (14). Rudimentary structures have already been fabricated literally atom-by-atom (15).

J A Armstrong, IBM Chief Scientist and Vice President for Science and Technology, said

> *'I believe that nanoscience and nanotechnology will be central to the next epoch of the information age, and will be as revolutionary as science and technology at the micron scale have been since the early '70s... Indeed, we will have the ability to make electronic and mechanical devices atom-by-atom when that is appropriate to the job at hand' (18).*

The New York Times said (19):

> *'Scientists are beginning to gain the ability to manipulate matter by its most basic components—molecule by molecule and even atom by atom. That ability, while now very crude, might one day allow people to build almost unimaginably small electronic circuits and machines, producing, for example, a supercomputer invisible to the naked eye. Some futurists even imagine building tiny robots that could travel through the body performing surgery on damaged cells.'*

Drexler (5, 11, 12, 13, 17) has proposed the assembler, a small device resembling an industrial robot which would be capable of holding and positioning reactive compounds in order to control the precise location at which chemical reactions take place. This general approach should allow the construction of large atomically precise objects by a sequence of precisely controlled chemical reactions.

The best technical analysis to date of molecular nanotechnology was done by Drexler (17). A more accessible introduction to the subject is his *Engines of Creation* (5). Other suggested readings are (6, 7, 8, 13, 20, 21).

Self replicating systems

The Von Neumann architecture for a self-replicating system (22) consists of a Universal Computer coupled to a Universal Constructor. The computer would direct the activities of the constructor, and the constructor would then build another computer and another constructor. While theoretical in nature, his proposals describe the basic issues involved in such systems.

Further work on self-replicating systems was done by NASA in 1980 in a report that considered the feasibility of implementing a self-replicating lunar manufacturing facility with conventional technology (20). One of their conclusions was that 'The theoretical concept of machine duplication is well developed. There are several alternative strategies by which machine self-replication can be carried out in a practical engineering setting'. They estimated it would require 20 years to develop such a system. While they were considering the design of a macroscopic self-replicating system (the proposed 'seed' was 100 tons) many of the concepts and problems involved in such systems are similar regardless of size.

Assemblers

Drexler proposed the 'assembler' (5, 11, 17). It adopts the Von Neumann architecture but, unlike the NASA proposal, is both smaller and simpler (having an estimated design complexity of about 10 megabytes). It is composed of a molecular computer and what might be termed a 'molecular constructor'. The molecular constructor has two main elements: a molecular positional device (a small robotic arm) and a well defined set of chemical reactions that take place at the tip of the arm. By positioning reactive compounds, the assembler can synthesize structures by a series

of atomically precise reactions. Conceptually, this is simply an extension of the capabilities demonstrated by the ribosome (23). Instead of bonding a sequence of amino acids into a linear structure under the control of messenger RNA, the assembler bonds a more general set of compounds into a three dimensional structure under the control of a general purpose computer.

Feynman said: 'The problems of chemistry and biology can be greatly helped if our ability to see what we are doing, and to do things on an atomic level, is ultimately developed—a development which I think cannot be avoided.' Drexler's assembler is the smallest embodiment of that dream.

Describing the brain at the molecular and atomic level

In principle we need only repair the frozen brain, for the brain is the most critical and important structure in the body. Other parts of the body can simply be replaced; or, if we wished, the methods employed for the brain could be extended in the obvious way.

The most we could usefully know about the frozen brain would be the coordinates and type of each and every atom in it. (To fully specify the state of each atom would, strictly speaking, require that we specify the states of all its electrons; for the most part these can be readily inferred once the coordinates and type of atom are given). This knowledge, combined with a technology that allowed us to rearrange atomic structure in virtually any fashion consistent with the laws of chemistry and physics, would let us restore the frozen structure to a healthy state.

To do this, we must answer three questions:

1. Where are the atoms?
2. Where should they go?
3. How do we move them from where they are to where they should be?

We shall first consider a simpler problem: how would we go about describing the position of every atom if somehow this information was known to us? The answer to this question will let us better understand the harder questions.

How many bits to describe one atom

Each atom has a location in three-space that we can represent with three coordinates: X, Y, and Z. Atoms are usually a few tenths of a nanometer apart. If we could record the position of each atom to within 0.01 nm, we would know its position accurately enough to know what chemicals it was a part of, what bonds it had formed, and so on. The brain is roughly 0.1 m across, so 0.01 nm is about 1 part in 10^{10}. That is, we would have to know the position of the atom in each coordinate to within one part in 10 billion. A number of this size can be represented with about 33 bits. There are three coordinates, X, Y, and Z, each of which requires 33 bits to represent, so the position of an atom can be represented in 99 bits. An additional few bits are needed to store the type of the atom (whether hydrogen, oxygen, carbon, etc), bringing the total to slightly over 100 bits.

We would fully describe the structure of the brain by storing about 100 bits per atom. If we assume that we require 10 atoms to store a single bit of information, (see 'Memory Requirements' below) then the 100 bits required to describe a single atom can be represented by about 1000 atoms. Put another way, an atom in three-space encodes its own position in the analog value of its three spatial coordinates. If we convert this spatial information from its analog format to a digital format, we inflate the number of atoms needed to record this information by perhaps as much as 1000. If we digitally encoded the location of every atom in the brain, we would need 1000 times as many atoms to hold this encoded data as there are atoms in the brain. This means we would require roughly 1000 times the volume. The brain is somewhat over one cubic decimeter, so it would require somewhat over one cubic meter of material to encode the location of each and every atom in the brain in a digital format suitable for examination and modification by a computer.

While this much memory is remarkable by today's standards, its construction clearly does not violate any laws of physics or chemistry. That is, it should literally be possible to store a digital description of each and every atom in the brain in a memory device that we will eventually be able to build.

Describing the brain more compactly

Because molecules are made of many atoms, we could achieve a more compact description by giving the coordinates and orientation of each molecule, rather than each atom, in the brain. A whole protein molecule might require 150 bits to describe (we would need to give its orientation and type, hence the extra bits), even though it is made of thousands of atoms. We can compress this further by using various other clever strategems (50 bits or less is quite achievable), but the essential point should be clear. We are interested in molecules, and describing a molecule takes fewer bits than describing an atom.

Even more compactly, we could describe an entire cell with only a general description of the function it performs: this nerve cell has synapses of a certain

type with that other cell, it has a certain shape, and so on. We might even describe groups of cells in terms of their function: this group of cells in the retina performs a 'center surround' computation, while that group of cells performs edge enhancement.

What is the most compact description that captures all the essential information? Our memories clearly matter, and estimates of the information content of long term memory are as small as 10^9 bits (24). It's safe to say that we'll need at least this many bits in a satisfactory description of an individual brain. The gap between this lower bound and the molecule-by-molecule upper bound is rather large, and it is not immediately obvious where in this range the true answer falls. We shall not attempt to answer this question, but will instead (conservatively) simply adopt the upper bound.

Criteria of death

> '**death**/'deth/*n*[ME *deeth*, fr. OE *death*; akin to ON *dauthi* death, *deyja* to die—more at DIE] **1**: a permanent cessation of all vital functions: the end of life'
> *Webster's New Collegiate Dictionary*

Determining when 'permanent cessation of all vital functions' has occurred is not always easy. 'Many people in the nineteenth century, alarmed by the prevalence of premature burial, requested, as part of the last offices, that wounds or mutilations be made to assure that they would not awaken... embalming received a considerable impetus from the fear of premature burial' (25).

Current criteria of 'death' are sufficient to ensure that spontaneous recovery in the mortuary or later is rare, but are simply an ad hoc summary of symptoms that have proven resistant to treatment by present methods. Each new medical advance forces a re-examination and possible change. The criteria used 200 years from now will differ dramatically from the criteria used today.

These ever shifting criteria for 'death' raise an obvious question: is there a definition that *does* have a theoretical basis and is *not* dependent on the technology of the day?

When someone has suffered a loss of memory or mental function, we often say they 'aren't themselves'. As the loss becomes more serious and all higher mental functions are lost, we begin to use terms like 'persistent vegetative state'. While we will often refrain from declaring such an individual 'dead', this hesitation does not arise because we think they are 'alive' but because there is still hope of recovery with memory and personality intact. From a physical point of view we believe there is a chance that their memories and personalities are still present within the physical structure of the brain, even though their behavior does not provide direct evidence for this. If we could reliably determine that the physical structures encoding memory and personality had in fact been destroyed, then we would abandon hope and declare the person dead.

The information theoretic criterion of death

If we knew the coordinates of each and every atom in a person's brain then we would (at least in principle) be in a position to determine with absolute finality whether their memories and personality had been destroyed in the information theoretic sense, or whether they were preserved but could not, for some reason, be expressed.

Considerations like this lead to the *information theoretic criterion of death*. A person is dead according to the information theoretic criterion if their memories, personality, hopes, dreams, etc have been destroyed in the information theoretic sense. If the structures in the brain that encode memory and personality have been so disrupted that it is no longer possible in principle to recover them, then the person is dead. If they are sufficiently intact that inference of the state of memory and personality are feasible in principle, and therefore restoration to an appropriate functional state is likewise feasible in principle, then the person *is not dead.*

A simple example is in order. If a computer is fully functional, then its memory and 'personality' are completely intact. If we took an axe to the CPU, then the computer would no longer be functional. However, its memory and 'personality' would still be present on disk, and once we repaired the CPU we could fully restore the computer.

If the structures encoding memory and personality *have* suffered sufficient damage to obliterate them beyond recognition, then death by the information theoretic criterion has occurred. An effective method of insuring such destruction is to burn the structure and stir the ashes. This is commonly employed to ensure the destruction of classified documents. Under the name of 'cremation' it is also employed on human beings and is sufficient to ensure that death by the information theoretic criterion takes place.

Does the information theoretic criterion matter?

It is normally a matter of small concern whether a physician of 2190 would or would not concur with

the diagnosis of 'death' by a contemporary physician applied to a specific patient in 1990. A physician of today who found himself in 1790 would be able to do little for a patient whose heart had stopped, even though he knew intellectually that an intensive care unit would likely be able to save the patient's life. Intensive care units were simply not available in 1790, no matter what the physician knew was possible. So, too, with the physician of today when informed that a physician 200 years hence could save the life of the patient that he has just pronounced 'dead'. There is nothing he can do, for he can only apply the technologies of today—*except in the case of cryonic suspension.*

In this one instance, we must ask not whether the person is dead by today's criteria, but whether the person is dead by all future criteria. In short, we must ask whether death by the information theoretic criterion has taken place.

Logically, there are two *and only two* ways in which cryonics can fail.

Cryonics will fail if:

1. Information theoretic death occurs.
2. Technologies that are feasible in principle prove impossible to develop in practice, even after centuries of development.

While information theoretic death could in principle occur at any time, the risk will obviously vary with the circumstances. For example, it appears implausible that information theoretic death could occur during storage at the temperature of liquid nitrogen. Peter Mazur, a well known cryobiologist and critic of cryonics, has said: 'Cryobiologists are often asked how long cells can remain viable at -196°C, the temperature of boiling liquid nitrogen (which is the usual cryogenic fluid). The answer is clear—more than 1000 years' (26). Hayflick has kept normal fibroblasts from embryonic human lungs in liquid nitrogen for 28 years (as of June 1990) without noticeable deterioration (27).

Information theoretic death that occurs significantly before suspension is outside the scope of this paper. Information theoretic death that occurs secondary to premature thawing is not an argument against cryonics, but an argument against unreliable refrigeration; this issue is considered elsewhere (2).

This leaves two main types of damage which might result in information theoretic death: freezing damage and ischemia.

Freezing damage

There is an extensive literature on the damage caused by both cooling and freezing to liquid nitrogen temperatures. Some reviews are (28, 29, 30). *Scientific American* had a recent and quite accessible article (31). Of particular relevance in the present context are (32, 33).

Many types of tissue including human embryos, sperm, skin, bone, red and white blood cells, bone marrow, and others (28, 29, 4) have been frozen in liquid nitrogen, thawed, and have recovered. This is not true of whole mammals, although many non-mammalian animals can be frozen to temperatures as low as -50°C and survive (31). The brain seems more resistant than most organs to freezing damage (34, 35). Recovery of overall brain function following freezing to liquid nitrogen temperature has not been demonstrated, although recovery of unit level electrical activity following freezing to -60°C has been demonstrated (35).

The complexity of the tissues that have been successfully frozen and rewarmed is remarkable, and supports the hypothesis that good structural preservation has been achieved. The mechanisms of damage that have been postulated in the literature are sufficiently subtle that information loss is neglibible; that is, death by the information theoretic criterion is unlikely to have occurred.

Ischemic injury

Although modern cryonic suspensions can involve a delay of less than 5 min (when 'death' is declared promptly upon cessation of heartbeat), and the future legalization of cryonic suspensions prior to a declaration of clinical death might eliminate delay entirely, delay is often unavoidable. The most significant damage that such delay causes is ischemic injury.

While not applicable to cryonics as presently practiced, it is interesting to note that the use of chemical fixatives such as aldehydes and in particular glutaraldehyde) would reliably improve structural preservation and would be effective at halting almost all deterioration within minutes of perfusion (36). The utility of chemical preservation has been discussed by Drexler (5) and by Olson (37), among others.

Temporary functional recovery has been demonstrated in optimal situations after as long as 60 min of total ischemia (38, 39, 40). Hossmann, for exam-

ple, reported results on 143 cats subjected to 1h of normothermic global brain ischemia (41):

'Body temperature was maintained at 36° to 37°C with a heating pad. ... Completeness of ischemia was tested by injecting ^{133}Xe into the innominate artery immediately before vascular occlusion and monitoring the absence of decay of radioactivity from the head during ischemia, using external scintillation detectors. ... In 50% of the animals, even major spontaneous EEG activity returned after ischemia..... One cat survived for 1 yr after 1 hour of normothermic cerebrocirculatory arrest with no electrophysiologic deficit and with only minor neurologic and morphologic disturbances.'

Functional recovery is a more stringent criterion of recovery than the more relaxed information theoretic criterion, which merely requires adequate structural preservation to allow inference about the pre-existing structure. The ability to recover function in a significant percentage of animals suggests that structural integrity must be quite good. This view is also supported by the observation that reliable identification of the various cellular structures is possible hours later. Detailed descriptions of ischemia and its time course (42, page 209 et sequitur) also clearly show that cooling substantially slows the rate of deterioration. Even moderate cooling 'postmortem' slows deterioration significantly.

Other evidence that structure does not deteriorate for some time is the observation that lysosomes do not rupture for at least several hours following the onset of ischemia (43, 44); and that messenger RNA and protein structure remain almost completely chemically intact 'postmortem' even when left at room temperature for up to 16h (45).

Present evidence supports but does not prove the hypothesis that information theoretic death does not occur for at least a few hours following the onset of ischemia. It is very possible that many hours of ischemia are required to produce information theoretic death.

Memory

It appears likely that short term memory, which can be disrupted by trauma or a number of other processes, will not be preserved by cryonic suspension. Long term memory, however, is much less volatile. The available evidence supports the idea that memory and personality are stored by alterations in the synapses between nerve cells (46–50).

Greenough and Bailey (51) say:

More recently it has become clear that the arrangement of synaptic connections in the mature nervous system can undergo striking changes even during normal functioning. As the diversity of species and plastic processes subjected to morphological scrutiny has increased, convergence upon a set of structurally detectable phenomena has begun to emerge. Although several aspects of synaptic structure appear to change with experience, the most consistent potential substrate for memory storage during behavioral modification is an alteration in the number and/or pattern of synaptic connections'.

What, exactly, might these changes be? Very strong statements are possible in simple 'model systems'. Bailey and Chen, for example, identified several specific changes in synaptic structure that encoded learned memories from sea slugs (Aplysia californica) by direct examination of the changed synapse with an electron microscope (52).

Electron microscopy is able to recover a wide range of information about the nervous system (53, 54, 55). It is much cruder than the techniques considered here, which literally propose to analyze every molecule in the structure. Further alterations in synaptic chemistry should be detectable when the synapse is examined in more detail at the molecular level.

Human memory, whatever the precise mechanism, must persist over the lifetime of a human being. It must tolerate the natural conditions encountered by humans and the experimental conditions to which primates have been subjected without loss of memory. It almost certainly involves changes in tens of thousands of molecules to store each bit of information. Given that future technology will allow the molecule-by-molecule analysis of the structures that store memory, and given that such structures are large on the molecular scale then it appears very unlikely that such structures will survive the lifetime of the individual only to be obliterated beyond recognition by freezing. Freezing is unlikely to cause information theoretic death.

The repair process

Consequently, the approach that we will consider is simple:

1. Determine the coordinates and orientations of all major molecules, and store this information in a data base. As a side effect, disassemble the tissue into its component molecules.
2. Using the information stored in the data base, run a computer program which determines what changes in the existing structure should be made to restore it to a healthy state.
3. Move the molecules to their desired locations.

We will call this approach 'off-board repair' to distinguish it from other proposals (5, 37). Off-board repair has been previously discussed (56).

We shall presume that step 1, the analysis phase, takes place while the tissue is still frozen. The thawing process itself causes damage and, once thawed, continued deterioration will proceed unchecked by the mechanisms present in healthy tissue. This cannot be tolerated. The temperature at which the other phases takes place is left open.

By describing the intermediate state which must be achieved during the repair process, we reduce the problem from 'Start with frozen tissue and generate healthy tissue' to 'Start with frozen tissue and generate a structural data base. Take the structural data base and the original molecules and generate healthy tissue'. It is characteristic of off-board repair that we disassemble the structure into its component molecules prior to attempting repair.

Off-board repair is the best that can be achieved

Regardless of the initial level of damage, regardless of the functional integrity of lack thereof of any or all of the frozen structure, regardless of whether easier and less exhaustive techniques might or might not work; we can take any frozen structure and convert it into the canonical state described above. Further, this is the best that we can do. Knowing the type, location and orientation of every molecule in the frozen structure under repair and retaining the actual physical molecules (thus avoiding any philosophical objections that replacing the original molecules might somehow diminish or negate the individuality of the person undergoing repair) is the best that we can hope to achieve. We have reached some sort of limit with this approach, a limit that makes repair feasible under circumstances which would astonish most people today.

Gaining access to every molecule

We can gain access to every molecule in frozen tissue by using 'divide and conquer'. In this method, an object is analyzed by successfully dividing it into smaller objects. Ultimately, the smallest objects are analyzed directly. Dividing tissue which is below the glass transition temperature (at least 130K) will involve making fractures. Tissue at liquid nitrogen temperatures is already prone to fracturing, so it should require only modest effort to deliberately induce a fracture that would divide a piece into two roughly equal parts. Fractures made at low temperatures (when the material is below the glass transition temperature) are extremely clean, and result in little or no loss of structural information. Hayat (57, p398) says 'The fracture plane often follows the contours of membranes and leaves bumps or depressions where it passes around vesicles and other cell organelles. ... The fracturing process provides more accurate insight into the molecular architecture of membranes than any other ultrastructural method'.

The freshly exposed faces can now be analyzed by any of a wide range of atomic resolution surface analysis techniques (58). The ability to rapidly sequence DNA is likely to emerge in the next few years from this kind of technology (59). Scanning probe devices depend on the interaction between a single atom at the tip of the probe and the atoms on the surface of the specimen under analysis. Clearly, future descendants of such devices can be very small.

Near field microscopy allows almost any optical method to be used at a resolution of 12 nm (60). This opens up another wide range of high resolution imaging methods.

Dividing into pieces

The division into halves continues until the pieces are small enough to allow direct analysis. A reasonable size might be small cubes a few tenths of a micron on a side.

One might view these cubes as the pieces of a three-dimensional jig-saw puzzle, the only difference being that we have cheated and carefully recorded the position of each piece. Just as the picture on a jig-saw puzzle is clearly visible despite the fractures between the pieces, so too the three-dimensional 'picture' of the brain is clearly visible despite its division into pieces.

There are a great many possible methods of handling the mechanical problems involved in dividing and moving the pieces. It seems unlikely that mechanical movement of the pieces will prove an insurmountable impediment if we allow for a century or two of development.

Memory requirements

The information storage requirements for a structural data-base that holds the detailed description and location of each major molecule in the brain can be met by projected storage methods. DNA has an information storage density of about 10^{21} bits/cm^3. Conceptually similar but somewhat higher density molecular 'tape' systems that store 10^{22} bits/cubic centimeter (5)

should be quite feasible. We can describe the 2×10^{23} or so significant molecules in the brain in about 10^{25} bits. This is about 1000 cm^3 (1 liter, roughly a quart) of 'tape' storage. If a storage system of such capacity strikes the reader as infeasible, consider that a person has about 10^{14} cells and that each cell stores 10^{10} bits in its DNA. Therefore, every person has (among other things) a raw storage capacity of 10^{24} bits—and people are unlikely to be optimal information storage devices.

Computational requirements

The computational power required to analyze a data base with 10^{25} bits is well within known theoretical limits (61, 62, 63). The cost of computation has been declining steadily and predictably for decades (64, p64). There are design proposals for molecular logic elements that dissipate roughly 10^{-21} joules per gate operation when operating at 50 picoseconds (6, 12), which is sufficient to do the job. Other proposals that dissipate little energy have been made (65, 66, 67) and a wide range of molecular logic elements have been discussed (68). A wide range of computational devices will be developed in the future that are very small in size and dissipate extraordinarily small amounts of energy. Extrapolation of current trends in miniaturization suggest that energy dissipations in the 10^{-21} joule range will be achieved by the year 2015 (69, Fig. 1). There is no presently known reason to expect the trend to stop or even slow down at that time (61, 63).

Energy costs appear to be the limiting factor in these proposals (rather than the number of gates, or the speed of operation of the gates). At today's price of roughly 10 cents per kilowatt hour, 10^{12} joules costs $30 000. This much energy would support 10^{33} gate operations: 10^8 gate operations for each bit in the structural data base, or 5×10^9 gate operations for each of the 2×10^{23} significant molecules present in the brain.

How much is enough?

We can get a rough idea of how much computer power might be required to analyze a data base with 10^{25} bits if we consider image recognition. The human retina performs about 100 'operations' per pixel, and the human brain is perhaps 1000 to 10 000 times larger than the retina. This implies that the human image recognition system can recognize an object after devoting some 10^5 to 10^6 'operations' per pixel. (This number is also in keeping with informal estimates made by individuals expert in computer image analysis.) Allowing for the fact that such 'retinal operations' are probably more complex than a single 'gate operation' by a factor of 100 to 1000, we arrive at 10^7 to 10^9 gate operations per pixel—which is quite in keeping with our estimate of 10^8 operations per bit or 5×10^9 operations per molecule.

More computer power

In fact, even more computational power will be available, so our margins for error are much larger.

Energy loss in computation is related to speed of operation. By slowing down the operating speed from 50 picoseconds to 50 nanoseconds or even 50 microseconds we should achieve corresponding reductions in energy dissipation per gate operation. We can both decrease the energy dissipated per gate operation (by operating at a slower speed) *and* increase the total number of gate operations (by using more gates). Because the gates are very small to start with, increasing their number by a factor of as much as 10^{12} (to approximately 5×10^{26} gates) would still result in a total volume of only 10 m^3. Given that manufacturing costs will eventually reflect primarily material and energy costs, such a volume of slowly operating gates should be economical to manufacture and would deliver substantially more computational power per joule.

Determining the healthy state

In the second phase of the analysis, determination of the healthy state, we determine what the repaired (healthy) tissue should look like at the molecular level. That is, the initial structural data base produced by the analysis phase describes unhealthy (frozen) tissue. We must generate a revised structural data base that describes the corresponding healthy (functional) tissue. The generation of this revised data base requires a computer program that has an intimate understanding of what healthy tissue should look like, and the correspondence between unhealthy (frozen) tissue and the corresponding healthy tissue. As an example, this program would have to understand that healthy tissue does not have fractures in it, and that if any fractures are present in the initial data base (describing the frozen tissue) then the revised data base (describing the resulting healthy tissue) should be altered to remove them. Similarly, if the initial data base describes tissue with swollen or non-functional mitochondria, then the revised data base should be altered so that it describes fully functional mitochondria.

The complexity of the program that determines the healthy state will vary with the quality of the suspension and the level of damage prior to suspension. Cryonic suspension under favourable circumstances preserves the tissue with remarkable fidelity down to the molecular level. If, however, there was significant pre-suspension injury then deducing the correct (healthy) structural description is more complex. However, it should be feasible to deduce the correct structural description even in the face of signficant damage. Only if the structure is obliterated beyond recognition will it be infeasible to deduce the undamaged state of the structure.

Restoration

Given a complete description of the structure to be built down to the level of individual molecules, we now have to build it. It is in principle possible to build a human brain, for this has been done by traditional methods for many thousands of years.

A wide range of methods of restoring the structure should be feasible. A high precision method which allowed accurate restoration might simply stack the molecular components in the desired locations. A problem with this approach is the stability of the structure during restoration. Molecules might drift away from their assigned locations, destroying the structure. There are many possible methods of stabilizing the intermediate structure during the course of restoration. We will consider only one of them here.

One simple method for achieving a high degree of stability of the intermediate structure would be to use low temperatures. If the structure were restored at a sufficiently low temperature, a molecule put in a certain place wouldn't move: it would be held in place by Van der Waals forces.

In this scenario, each new molecule would simply be stacked (at low temperature) in the right location. Because biological systems make extensive use of self-assembly, it would not be necessary to achieve perfect accuracy in the restoration process. If a biological macromolecule is positioned with reasonable accuracy, it would automatically assume the correct position upon warming. Positional accuracy of less than an angstrom has already been demonstrated with current STM technology.

Large polymers, used either for structural or other purposes, pose special problems. The monomeric units are covalently bonded to each other, and so simple 'stacking' is inadequate. If such polymers cannot be added to the structure as entirely pre-formed units, then they could be incrementally restored during assembly from their individual monomers using positional synthesis techniques that use highly reactive intermediates (17, 70).

The chemical reactions required to make a polymer from its monomeric units at reduced temperatures are unlikely to use the same pathways that are used by living systems. In particular, the activation energies of most reactions that take place at 310 K (98.6° Fahrenheit) can not be met at 77 K: most conventional compounds don't react at that temperature. However, it is possible to a) use compounds that are highly reactive and thus do not require any significant activation energy (70) or b) to provide activation energy by using non-thermal energy sources, e.g. mechanical force (17). Addition of monomeric units to the polymer could then be done at the most convenient point during the restoration operation, and could be done at the ambient (low) temperature.

An obvious problem with this approach is the need to re-warm the structure without incurring further damage. There are several equally obvious solutions to this problem, which can be combined into a single robust solution. These solutions are: 1) rapid, highly controlled heating 2) atomically precise introduction of cryoprotectants into the structure at low temperature prior to warming 3) modification of the structure to achieve a thermal expansion coefficient as close to zero as feasible, thus reducing or eliminating thermal strain during warming, and 4) the addition of structural elements (long protein strands, for example) to provide additional strength to prevent fractures.

Conclusion

Cryonic suspension can transport a terminally ill patient to future medical technology. The damage done by current freezing methods is likely to be reversible at some point in the future. In general, for cryonics to fail, one of the following 'failure criteria' must be met:

1. Pre-suspension and suspension injury would have to be sufficient to cause information theoretic death. In the case of the human brain, the damage would have to obliterate the structures encoding human memory and personality beyond recognition.
2. Repair technologies that are clearly feasible in principle based on our current understanding of physics and chemistry would have to prove impossible to develop in practice, even after several centuries of development.

Examination of likely future technical capabilities supports the argument that unprecendented abilities are likely to be developed. Restoration of the brain

down to the molecular level should eventually prove technically feasible. Off-board repair utilizing divide-and-conquer is a particularly simple and powerful method which illustrates some of the principles that can be used by future technologies to restore tissue. A wide range of approaches other than the one considered here are feasible. The present method is not proposed as the 'right' or 'best' method, it is proposed as a conceptually simple and feasible method. A single feasible method of repairing freezing injury establishes the effectiveness of cryonics, regardless of the methods that are eventually implemented.

Off-board repair consists of three major steps: 1) Determine the coordinates and orientation of each major molecule. 2) Determine a set of appropriate coordinates in the repaired structure for each major molecule. 3) Move them from the former location to the latter. The various technical problems involved are likely to be met by future advances in technology. Because storage times in liquid nitrogen literally extend for several centuries, the development time of these technologies is not critical.

The particular proposal discussed here requires only that 1) tissue can be divided by some means (such as fracturing) which does not itself cause significant loss of structural information; 2) the pieces into which the tissue is divided can be moved to appropriate destinations (for further division or for direct analysis); 3) a sufficiently small piece of tissue can be analyzed; 4) a program capable of determining the healthy state of tissue given the unhealthy state is feasible; 5) the sufficient computational resources for execution of this program in a reasonable time frame are available; and 6) that restoration of the original structure given a detailed description of that structure is feasible.

The extant literature supports but does not prove the hypothesis that cryonics is a feasible method of saving the lives of people who would otherwise certainly die. Further study of cryonics by the technical community is needed. As should be evident from this paper multidisciplinary analysis is essential in evaluating its feasibility. Given the life-saving nature of cryonics, it would be tragic if it were to prove feasible but was little used.

Acknowledgements

It is the authors pleasant duty to acknowledge the many people who have commented on or encouraged the work on this paper as it evolved. The reviewers were not selected because of their opinions about cryonics: some support it, some don't, and some reserve final judgement. While the quality of the result could not have been achieved without their help, the author must accept responsibility for any errors in the final version. The author would like to thank: Dave Biegelsen, Arthur C Clarke, Mike Darwin, Thomas Donaldson, Eric Drexler, Greg Fahy, Steve Harris, Leonard Hayflick, Hugh Hixon, Peter Mazur, Mark Miller, David Pegg, Chris Peterson, Ed Regis, Paul Segall, Len Shar, Irwin Sobel, Jim Southard, Jim Stevens and Leonard Zubkoff.

References

1. Ettinger R C W. The Prospect of Immortality. Sidgwick and Jackson, 1965.
2. Wowk B, Darwin M. Cryonics: reaching for tomorrow. Alcor. 12327 Doherty St, Riverside, CA 92503, USA, 1991.
3. Harris S B. Many are cold, but few are frozen: A Humanist looks at Cryonics. Free Inquiry, 9: 2, Spring. 19–24, 1989.
4. Fahy G M. Analysis of solution effects injury: rabbit renal cortex frozen in the presence of dimethyl sulfoxide. Cryobiology 17: 371–388, 1980.
5. Drexler K E. Engines of Creation. Anchor Press, 1986.
6. Dewdney A K. Nanotechnology: wherein molecular computers control tiny circulatory submarines. Scientific American, 100–103, Jan 1988.
7. Foresight Update, a publication of the Foresight Institute, Box 61058, Palo Alto, CA 94306, USA.
8. Feynman R. There's Plenty of Room at the Bottom. A talk at an annual meeting of the American Physical Society, December 1959. Reprinted in Miniaturization (H D Gilbert, ed), Reinhold, New York. 282–296, 1961.
9. Hansma P K, Elings V B, Marti O, Bracker C E. Scanning Tunneling Microscopy and Atomic Force Microscopy: Application to Biology and Technology. Science, 209–216, Oct 14 1988.
10. Foster J S, Frommer J E, Arnett P C. Molecular manipulatoin using a tunneling microscope. Nature, 331, 324–326, Jan 1988.
11. Drexler K E. Molecular Engineering: An Approach to the Development of General Capabilities for Molecular Manipulation. Proceedings of the National Academy of Sciences (USA), 78: 5275–78, 1981.
12. Drexler K E. Rod Logic and Thermal Noise in the Mechanical Nanocomputer. In: Proceedings of the Third International Symposium on Molecular Electonic Devices, (F Carter ed), Elsevier 1988.
13. Drexler K E. Machines of Inner Space. In: 1990 Yearbook of Science and the Future, 160–177, Encyclopedia Britannica, Chicago 1989.
14. Pool R. A Small Revolution Gets Under Way. Science Jan 5 1990.
15. Eigler D M, Schweizer E K. Positioning Single Atoms with a Scanning Tunnelling Microscope. Nature 344, 524–526, April 5, 1990.
16. The Invisible Factory. The Economist, 91, December 9, 1989.
17. Drexler K E. Nanosystems: Molecular Machinery, Manufacturing and Computation. John Wiley 1992.
18. Armstrong J A. New Frontiers in Computing and Telecommunications. Creativity 10, 2 June 1991.
19. Pollack A. Atom by Atom, Scientists build 'Invisible' Machines of the Future. The New York Times, Science section, B7, Tuesday November 26, 1991.
20. Freitas R A, Gilbreath W P. Advanced Automation for Space Missions, Proceedings of the 1980 NASA/ASEE Summer Study. NTIS, US Department of Commerce, National Technical Information Service, Springfield VA 22161.
21. Merkle R C. Self Replicating Systems and Molecular Manufacturing. Journal of the British Interplanetary Society. (submitted).

22. Von Neumann J, (Burks A W ed). Theory of Self Reproducing Automata. University of Illinois Press, 1966.
23. Watson J D, Hopkins N H, Roberts J W, Steitz J A, Weiner A M. Molecular Biology of the Gene. 4th ed. Benjamin Cummings, 1987.
24. Landauer T K. How Much Do People Remember? Some Estimates of the Quantity of Learned Information in Long-term Memory. Cognitive Science 10: 477–493, 1986.
25. Walker A E. Cerebral Death. Second Edition. Urban and Schwarzenberg 1981.
26. Mazur P. Stopping Biological Time: The Freezing of Living Cells. Ann NY Acad Sci 541: 514–531, 1988.
27. Hayflick L. Personal communication, 1991.
28. Pegg D. Principles of Tissue Preservation. In: Progress in Transplantation. Vol 2 (P J Morris and N L Tilney eds) 69–105.
29. Mazur P. Freezing of living cells: mechanisms and implications. American Journal of Physiology 247 (Cell Physiology 16): C125–C142, 1984.
30. Pegg D E, Karow A M. The Biophysics of Organ Cryopreservation. Plenum, 1987.
31. Storey K B, Storey J M. Frozen and Alive. Scientific American, 263: 6, 92–97, 1990.
32. Darwin M, Hixon H, Leaf J. Postmortem examination of three cryonic suspension patients. Alcor Life Extension Foundation. 12327 Doherty Street, Riverside, California, 92503.
33. The Cryobiological Case for Cryonics. Cryonics, March 1988. Alcor, 12327 Doherty Street, Riverside, CA 92503.
34. Suda I, Kito K, Adachi C. Viability of Long Term Frozen Cat Brain in Vitro. Nature 212, 268, October 15, 268.
35. Suda I, Kito K, Adachi C. Bioelectric discharges of isolated cat brain after revival from years of frozen storage. Brain Research 70: 527–531, 1974.
36. Hayat M A. Fixation for Electron Microscopy. Academic Press, 1981.
37. Olson C B. A Possible Cure for Death. Medical Hypotheses 26: 77–84, 1988.
38. Nordstrom C H, Siesjo B K. Neurochemical Determinants of Ischemic Cell Damage. In: Protection of the Brain from Ischemia. (P R Weinstein and A A Faden eds). Williams and Wilkens p49–66, 1990.
39. Hossmann K A. Hemodynamics of Postischemic Reperfusion of the Brain. In: Protection of the Brain from Ischemia. (P R Weinstein and A A Faden eds). Williams and Wilkens, p21–36, 1990.
40. Hossman K A. Post-ischemic resuscitation of the brain: selective vulnerability versus global resistance. In: Progress in Brain Research (K Kogure, K A Hossman, B K Siesjo and F A Welsh), 63: 3–17, 1985.
41. Hossmann K A. Resuscitation potentials after prolonged global cerebral ischemia in cats. Critical Care Medicine, 16: 10: 964–971, 1988.
42. Bowen I D, Lockshin R A. Cell Death in Biology and Pathology. Chapman and Hall, 1981.
43. Kalimo H, Garcia J H, Kamijyo Y, Tanaka J, Trump B F. The Ultrastructure of 'Brain Death', II. Electron Microscopy of Feline Cortex after Complete Ischemia. Virchows Archiv B Cell Path. 25, 207–220, 1977.
44. Hawkins H K, Ericsson J L E, Biberfeld P, Trump B F. Lysosome and Phagosome Stability in Lethal Cell Injury. Am J Pathol 68: 255–288, 1972.
45. Morrison M R, Griffin S T. The isolation and in Vitro translation of Undegraded Messenger RNAs from Human Postmortem Brain. Analytical Biochemistry 113, 318–324, 1981.
46. Shepherd G M. Neurobiology. Oxford 1983.
47. Kandel E R, Schwartz J H. Principles of Neural Science. 2nd ed. Elsevier, 1985.
48. Squire L R. Memory and Brain. Oxford University Press, 1987.
49. Lynch G. Synapses, Circuits and the Beginnings of Memory. MIT Press, 1986.
50. Turner A M, Greenough W T. Differential Rearing Effects on Rat Visual Cortex Synapses. I. Synaptic and Neuronal Density and Synapses per Neuron. Brain Research, 329: 195–203, 1985.
51. Greenough W T, Bailey C H. The anatomy of a memory: convergence of results across a diversity of tests. Trends in Neuroscience 11: 4: 142–147, 1988.
52. Bailey C H, Chen M. Morphological basis of long-term habituation and sensitization in Aplysia. Science 220: 91–93, 1 April 1983.
53. Mize R R. The microcomputer in cell and neurobiology research. Elsevier 1985.
54. Sasaki-Sherrington S E, Jacobs J R, Stevens J K. Intracellular control of axial shape in non-uniform neurites: a serial electron microscopic analysis of organelles and microtubules in AI and AII retinal amacrine neurites. Journal of Cell Biology 98: 1279–1290, April 1984.
55. Merkle R C. Large Scale Analysis of Neural Structures. Xerox Technical Report CSL–89–10, November 1989. Xerox Corporation, Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304, USA.
56. Merkle R C. Molecular Repair of the Brain. Cryonics 10: 21–44. Alcor Life Extension Foundation, 12327 Doherty Street, Riverside, California 92503, USA, October 1989.
57. Hayat M A. Principles and Techniques of Electron Microscopy: Biological Applications. 3rd ed. CRC Press, 1989.
58. Pool R. The Children of the STM. Science 634–636, Feb 9, 1990.
59. Williams C C, Wickramasinghe H K. Microscopy of Chemical Potential Variations on an Atomic Scale. Nature 344: 317–319, 22 March 1990.
60. Betzig E, Trautman J K, Harris T D, Weiner J S, Kostelak R L. Breaking the Diffraction Barrier: Optical Microscopy on a Nanometric Scale. Science 251: 1468, 22 March 1991.
61. Bennet C H, Landauer R. The fundamental physical limits of computation. Scientific American 253: 48–56, July 1985.
62. Fredkin E, Toffoli T. Conservative Logic. International Journal of Theoretical Physics. 21: 3/4: 219–253, 1982.
63. Benneu C H. Notes on the History of Reversible Computation. IBM Journal of Research and Development, 32: 1: Jan 1988.
64. Moravec H. Mind Children. Harvard University Press, 1988.
65. Likharev K K. Classical and Quantum Limitations on Energy Consumption in Computation. International Journal of Theoretical Physics 21: 3/4: 1982.
66. Merkle R C. Two Types of Mechanical Reversible Logic. Nanotechnology (submitted).
67. Merkle R C. Reversible Electronic Logic Using Switches. Nanotechnology (submitted).
68. Carter F. Proceedings of the Third International Symposium on Molecular Electronic Devices. Elsevier 1988.
69. Landauer R. Dissipation and noise immunity in computation and communication. Nature 335: 779, 27 October 1988.
70. Musgrave C, Perry J, Merkle R C, Goddard W A. Theoretical analysis of a site-specific hydrogen abstraction tool. Nanotechnology, April 1992.