COVER

Extropy Institute

Editor: Max More

<u>EXTROPY: The Journal of Transhumanist Thought</u> is a journal of ideas, dedicated to discussing and developing themes in the following areas:

- Transhumanism and futurist philosophy
- life extension, immortalism and cryonics
- artificial intelligence (AI) and uploading
- smart drugs (nootropics) and intelligence increase technologies
- nanotechnology applications
- memetics
- space habitation
- spontaneous orders (free markets, neural networks, evolutionary processes, etc)
- science fiction
- extropic psychology
- digital Economy (encryption; digital money)
- scientific eschatology
- artificial life
- futurist morality
- electronic communications

and reviews of media on these topics.

Edited by Max More. Subscriptions and letters can be sent to EXTROPY, P.O. Box 57306, Los Angeles, CA 90057-0306. Phone: (213) 484-6383. E-mail to: more@usc.edu

EXTROPY is published twice a year. Checks payable to "EXTROPY". Subscriptions will start with the current issue in print unless you specify otherwise.

Subscriptions: <u>USA:</u> $9/year. 2-years: $18
<u>Canada and Mexico</u>: $10year.
<u>Overseas</u>: $14/year (air); $12 (surface).
<u>Institutions</u>: USA, Canada, Mexico: $20/year. Overseas: $26/year (air).

BACK ISSUES:
#8: $4.50 #7: $4.50 #5,6: $4 each #1,2,4: $3 each
See insert for details of back issue contents.

EXTROPY INSTITUTE
Membership (Includes subscriptions to *Extropy* and the bi-monthly newsletter *Exponent*.):

USA:	$30
Canada/Mexico:	$33
Overseas:	$40

EXTROPY #9

Summer 1992
(Vol.4, No.1)

EDITORIAL

MORE CHANGES

Extropy sees further evolution with this issue. The most obvious change is the shift to a full size 8.5" x 11" format. The new format, consistent with that of major publications, should help in the expansion of shop distribution and sales. Comments on and suggestions for further improvements in *Extropy*'s format – in addition to its content – are always welcome.

Next issue will see some changes as a result of Extropy Institute's new publication, *Exponent*. *Exponent* will appear bi-monthly in newsletter format. It will feature reviews of fiction and non-fiction books, movies, magazines, and software. *Exponent* will also carry the Intelligence At Work science updates column that has previously appeared here, and a section presenting skeptical analyses of environmental issues.

Extropy will no longer carry many book reviews, though it will still carry longer reviews and review essays. This publication will specialize in detailed explanations and analyses of topics of interest to Extropians: Life extension, cryonics, intelligence increasing technologies, advanced computing, neural-computer integration, memetic analysis, nanotechnology, space habitation, self-transformative psychological techniques, artificial life, electronic communications, computational markets (digital economy), and the limits of physics.

THIS ISSUE

The Extropian Principles is a revised version of the Principles that appeared three years ago in *Extropy* #6. The new version adds the principle of spontaneous order, since it's clear that spontaneously ordering processes are an essential element of the Extropian worldview – as explained at length in "Order Without Orderers" in *Extropy* #7 (Spring 1991). Explanations of the principles have been expanded and some new suggested readings added to the list.

"Extropy Institute Launches" announces the incorporation of the Extropy Institute (ExI) and some of its projects.

David Ross provides a lucid and helpful discussion of uploading – the transfer of a human's consciousness into a computer, responding to those who doubt it to be possible and determining the conditions that are necessary for it to be achieved.

J. Storrs Hall, in "Nanotechnology and Faith" looks at different ways of thinking about nanotechnology, warning us against allowing our thinking to become religious in nature.

"The Making of a Small World" is a delightful piece of short fiction by Mike Perry. You may want to pass this around your non-immortalist friends.

Simon! D. Levy continues his Neurocomputation series with an introduction to genetic algorithms, yet another computational application of spontaneous ordering.

Roboticist Hans Moravec, author of the classic extropian book *Mind Children*, explains the ways in which physics might allow time travel, and shows how this might be applied to solving otherwise intractable computational problems.

The Futique Neologisms series continues this issue, with a further installment of new terms for thinking about the future.

In "Exercise and Longevity", Fran Finney reviews the scientific evidence for the life-extending effects of exercise.

The Transhuman Taste contains four reviews: David Krieger reviews *The Anthropic Cosmological Principle*, a lengthy work examining cosmological questions with a decidedly Extropian perspective. Simon! D. Levy reviews evolutionary biologist Richard Dawkin's *The Blind Watchmaker* – a brilliant exposition of evolutionary theory and how it explains life far better than religious myths. I review three of economist Julian Simon's books in "Economist Against the Apocalyptics." Simon's works provide a needed counterpoint to environmentalist doomsayers. Finally, Harry Shapiro reviews *Bionomics*, which draws parallels between biology and economy.

IN FUTURE ISSUES

The long-awaited articles on the future of electronic communications, and on digital economy (personal communication encryption, digital money), should be appearing soon.

#10 will see an exposition of the Extropian principle of Self-Transformation, examining its psychological, philosophical, and technological aspects.

The next issue will also feature another piece by leading roboticist Hans Moravec on our expansion into the cosmos.

Expect to see a memetic analysis of the spread of Extropian ideas.

Other possibilities include an introduction to molecular-scale computing, analyses of environmental issues, applications of personal identity theory, and discussions of intellectual property and "the Singularity".

[Continued on p.11]
Editorial (cont. from p.4)

EXTROPY - the process of increasing intelligence, information, energy, life, experience, diversity, opportunity and growth. Extropianism is the philosophy that seeks to increase extropy. The Extropian Principles are: (1) Boundless Expansion; (2) Self-Transformation; (3) Intelligent Technology; (4) Spontaneous Order; (5) Dynamic Optimism.

TRANSHUMANISM - Philosophies of life (such as Extropianism) that seek the continuation and acceleration of the evolution of intelligent life beyond its currently human form and human limitations by means of science and technology, guided by progressive principles and values, while rejecting dogma and religion.

The Extropian Principles

V. 2.0

Max More
Executive Director, Extropy Institute

1. Boundless Expansion - Seeking more intelligence, wisdom, and personal power, an unlimited lifespan, and removal of natural, social, biological, and psychological limits to self-actualization and self-realization. Overcoming limits on our personal and social progress and possibilities. Expansion into the universe and infinite existence.

2. Self-Transformation - A commitment to continual moral, intellectual, and physical self-improvement, using reason and critical thinking, personal responsibility, and experimentation. Biological and neurological augmentation.

3. Intelligent Technology - Applying science and technology to transcend "natural" limits imposed by our biological heritage and environment.

4. Spontaneous Order - Promotion of decentralized, voluntaristic social coordination mechanisms. Fostering of tolerance, diversity, long-term planning, individual incentives and personal liberties.

5. Dynamic Optimism - Positive expectations to fuel dynamic action. Promotion of a positive, empowering attitude towards our individual future and that of all intelligent beings. Rejection both of blind faith and stagnant pessimism.

These principles are further explicated below. In depth treatments can be found in various issues of **EXTROPY: The Journal of Transhumanist Thought**. (Spontaneous Order in #7, Dynamic Optimism in #8, and Self-Transformation in the forthcoming #10.)

1. Boundless Expansion

Beginning as mindless matter, parts of nature developed in a slow evolutionary advance which produced progressively more powerful brains. Chemical reactions generated tropistic behavior, which was superseded by instinctual and Skinnerian stimulus-response behavior, and then by conscious learning and experimentation. With the advent of the conceptual consciousness of humankind, the rate of advancement sharply accelerated as intelligence, technology, and the scientific method could be applied to our condition. Extropians seek the continuation and fostering of this process, transcending biological and psychological limits as we proceed into posthumanity.

In aspiring to transhumanity, and beyond to posthumanity, we reject natural and traditional limitations on our possibilities. We champion the rational use of science and technology to void limits on lifespan, intelligence, personal power, freedom, and experience. We are immortalists because we recognize the absurdity of accepting "natural" limits to our lives. For many the future will bring an exodus from Earth – the womb of human and transhuman intelligence – expanding the frontiers of humanity (and posthumanity) to include space habitats, other planets and this solar system, neighboring systems, and beyond. By the end of the 21st Century, more people may be living off-planet than on Earth

Resource limits are not immutable. The market price system encourages conservation, substitution and innovation, preventing any need for a brake on growth and progress. Expansion into space will vastly expand the energy and resources for our civilization. Living extended transhuman lifespans will foster intelligent use of resources and environment.

Extropians affirm a rational, market-mediated environmentalism aimed at maintaining and enhancing our biospheres (whether terrestrial or extra-terrestrial). We oppose apocalyptic environmentalism, which hallucinates catastrophes, issues a stream of doomsday predictions, and attempts to strangle our continued evolution.

No mysteries are sacrosanct, no limits unquestionable; the unknown must yield to the intelligent mind. We seek to understand and to master reality up to and beyond any currently foreseen limits.

2. Self-Transformation

We affirm reason, critical inquiry, intellectual independence, and intellectual honesty. We reject blind faith and passive, comfortable thinking that leads to dogmatism, religion, and conformity. A commitment to positive self-transformation requires us to critically analyze our current beliefs, behaviors, and strategies. Extropians therefore choose to place their self-value in continued development rather than "being right". We prefer analytical thought to fuzzy but comfortable delusion, empiricism to mysticism, and independent evaluation to conformity. Extropians affirm a philosophy of life but distance themselves from religious thinking because of its blind faith, debasement of human dignity, and systematized irrationality.

Perpetual self-improvement – physical, intellectual, psychological, and ethical – requires us to continually re-examine our lives. Extropians seek to better themselves, yet without denying their current worth. The desire to improve should not be confused with the belief that one is lacking in current value. But valuing oneself in the present cannot mean self-satisfaction, since an intelligent and probing mind can can always envisage a superior self in the future. Extropians are committed to expanding wisdom, fine-tuning understanding of rational behavior, and enhancing physical and intellectual capacities.

Extropians are neophiles and experimentalists. We are neophiles because we track the latest research for more efficient means of achieving our goals. We are experimentalists because we are willing to explore and test the novel means of self-transformation that we uncover. In our quest for advancement to the tranhuman stage, we rely on our own judgement, seek our own path, and reject both blind conformity and mindless rebellion. Extropians frequently diverge from the mainstream because they do not allow themselves to be chained by dogmas, whether religious, political, or social. Extropians choose their values and behavior reflectively, standing firm when required but responding flexibly to novel conditions.

Personal responsibility and self-determination goes hand-in-hand with neophilic self-experimentation. Extropians take responsibility for the consequences of our choices, refusing to blame others for the risks involved in our free choices. Experimentation and self-transformation require risks; Extropians wish to be free to evaluate the risks and potential benefits for ourselves, applying our own judgment and wisdom, and assuming responsibility for the outcome. We neither wish others to force standards upon us through legal regulation, nor do we wish to force others to follow our path. Personal-responsibility and self-determination are incompatible with authoritarian centralized control, which stifles the free choices and spontaneous ordering of autonomous persons.

External coercion, whether for the purported "good of the whole" or the paternalistic protection of the individual, is unacceptable to us. Compulsion breeds ignorance and weakens the connection between personal choice and personal outcome, thereby destroying personal responsibility. The proliferation of outrageous liability lawsuits, governmental safety regulations, and the rights-destroying drug war result from ignoring these facts of life. Extropians are rational individualists, living by their own judgment, making critical, informed, and free choices, and accepting responsibility for those choices.

As neophiles, Extropians study advanced, emerging, and future technologies for their self-transformative potential in enhancing our abilities and freedom. We support biomedical research with the goal of understanding and controlling the aging process. We are interested in any plausible means of conquering death, including interim measures like biostasis/cryonics, and long-term possibilities such as migration out of biological bodies into superior vehicles ("uploading").

We practice and plan for biological and neurological augmentation through means such as effective cognitive enhancers or "smart drugs", computers and electronic networks, General Semantics and other guides to effective thinking, meditation and visualization techniques, accelerated learning strategies, and applied cognitive psychology, and soon neural-computer integration. We do not accept the limits imposed on us by our natural heritage, instead we apply the evolutionary gift of our rational, empirical intelligence in order to surpass human limits and enter the transhuman and posthuman stages of the future.

3. Intelligent Technology

Extropians do not denigrate technology, no matter how radically different from historical norms, as "unnatural". The term 'natural' is largely devoid of meaning. We might say that any technological means of altering the environment or the human body is unnatural since it changes the previously existing state of nature. But we can also say that applying our intelligence through technology is natural to humans, and so changing both outside nature and our own biological nature can be regarded as natural.

Extropians affirm the necessity and desirability of science and technology. Practical means should be used to promote our goals of immortality, expanding intelligence, and greater physical abilities, rather than the wishful thinking, ignorant mysticism, and credulity, so common to the New Agers. Science and technology, as disciplined forms of intelligence, should be fostered, and we should seek to employ them in eradicating the limits to our Extropian visions.

We prefer analytical thought to fuzzy but comfortable delusion, empiricism to mysticism, and independent evaluation to conformity.

We do not share common cultural fears of technology, such as those embodied in the story of Frankenstein and the myth of the Tower of Babel. We favor careful and cautious development of powerful technologies, but refuse to attempt to stifle development on the basis of fear of the unknown. Extropians therefore oppose the anti-human "Back to the Pleistocene", anti-civilization rhetoric of the extreme environmentalists. Going backwards means death for billions and stagnation and oppression for the rest. Intelligent use of biotechnology, nanotechnology, space and other technologies, in conjunction with a market system, can remove resource constraints and solve environmental pressures.

We see technological development not as an end in itself, but as a means to the achievement and development of our values, ideals and visions. We seek to employ science and technology to remove limits to growth, and to radically transform both the internal and external conditions of existence.

We see the coming years and decades as being a time of enormous changes, changes which will vastly expand our opportunities, our freedom, and our abilities. Genetic engineering, interventive gerontology (life extension), space migration, smart drugs, more powerful computers and smarter programming, neural-computer interfaces, virtual reality, swift electronic communications, artificial intelligence, neural networks, artificial life, neuroscience, and nanotechnology will contribute to accelerating change.

4. Spontaneous Order

Spontaneous orders are self-generating, organic orders and differ from constructed, centrally directed orders. Both types of order have their place, but spontaneous orders are vital in our social interactions. Spontaneous orders have properties that make them especially conducive to Extropian goals and values and spontaneous ordering processes can be found at work in many fields. The evolution of complex biological forms is one example; others include the adjustment of ecosystems, artificial life demonstrations, memetics (the study of replicating information patterns), computational markets (agoric open systems), brain function and neurocomputation,

The principle of spontaneous order is embodied in the free market system – a system that does not yet exist in a pure form. The free market allows complex institutions to develop, encourages innovation, rewards individual initiative and reinforces personal responsibility, fosters diversity, and safeguards political freedom. Market economies ensure the technological and social progress essential to the Extropian philosophy. We reject the technocratic idea of central control by self-proclaimed experts. No group of experts can understand and control the endless complexity of an economy and society. Expert knowledge is best harnessed and transmitted through the superbly efficient mediation of the free market's price signals – signals that embody more information than any person or group could ever gather.

Sustained progress and intelligent, rational decision-making requires the diverse sources of information and differing perspectives made possible by spontaneous orders. Central direction constrains exploration, diversity, freedom, and dissenting opinion. Respecting spontaneous order means supporting voluntaristic, autonomy-maximizing institutions as opposed to rigidly hierarchical, authoritarian groupings with their bureaucratic structure, suppression of innovation and diversity, and smothering of individual incentives. Understanding spontaneous orders makes us highly suspicious of "authorities" where these are imposed on us, and skeptical of coercive leaders, unquestioning obedience, and unexamined traditions.

Making effective use of a spontaneously ordering social system requires us to be tolerant and peaceful, allowing others to pursue their lives as they see fit, just as we expect to be left to follow our own paths. We can best achieve mutual progress by interacting cooperatively and benevolently toward all who do not threaten our lives, and by supporting diversity of opinion and behavior. Respecting diversity and disagreement requires us to maintain control of our impulses and to uphold high standards of rational personal behavior. Extropians are guided in their actions by studying the fields of strategy, decision theory and game theory. These make clear to us the benefits of cooperation and encourage the long-term thinking appropriate to persons seeking an unlimited lifespan.

5. Dynamic Optimism

We espouse a positive, dynamic, empowering attitude. To successfully pursue our values and live our lives we must reject gloom, defeatism, and the common cultural focus on negatives. Problems – technical, social, psychological, ecological – should be acknowledged but not allowed to dominate our thinking and our direction. We respond to gloom and nay-saying by exploration and promotion of new possibilities. Extropians hold to both short and long-term optimism: In the short term we can cultivate our lives and enhance ourselves; in the long term the positive potentials for intelligent beings are virtually limitless.

We question limits that others take for granted. We look at the acceleration in scientific and technical knowledge, ascending standards of living, and social and moral evolution and project further advances. More researchers today than in all past history strive to understand aging, control disease, upgrade computers, and develop biotechnology and nanotechnology. Technological and social evolution continue to accelerate, leading, some of us expect, to a Singularity – a future time when many of the rules of life will so radically diverge from those familiar to us, and progress will be so rapid, that we cannot now comprehend that time. Extropians will maintain the acceleration of progress and encourage it in beneficial directions.

Adopting dynamic optimism means focusing on possibilities and opportunities, and being alert to solutions and potentialities. And it means refusing to whine about what cannot be avoided, learning from mistakes rather than dwelling on them in a victimizing, punishing manner. Dynamic optimism requires us to take the initiative, to jump up and plough into our difficulties with an attitude that says we can achieve our goals, rather than to sit back and immerse ourselves in defeatist thinking.

Dynamic optimism is not compatible with passive faith. Faith in a better future is confidence that an external force, whether God, State or society, will solve our problems. Faith, or the Polyanna/Dr. Pangloss variety of optimism, breeds passivity by encouraging the belief that progress will be effected by others. Faith requires a determined belief in external forces and so encourages dogmatism and irrational rigidity of belief and behavior. Dynamic optimism fosters activity and intelligence, telling us that we are capable of improving life through our own efforts. Opportunities and possibilities are everywhere, waiting for us to seize them and create new ones. To achieve our goals, we must believe in ourselves, work hard, and be open to revise our strategies.

Where others see difficulties, we see challenges. Where others give up, we move forward. Where others say enough is enough, we say: Forward! Upward! Outward! We espouse personal, social, and technological evolution into ever higher forms. Extropians see too far and change too rapidly to feel future shock. Let us advance the wave of evolutionary progress.

Extropianism is a *transhumanist* philosophy: Like humanism it values reason and sees no ground for believing in supernatural external forces controlling our destiny. But transhumanism goes further in calling us to push beyond the simply human stage of evolution. As physicist Freeman Dyson said: "Humanity looks to me like a magnificent beginning but not the final word." Religion has traditionally provided a sense of meaning and purpose in life, but it also suppressed intelligence and stifled progress. The Extropian philosophy provides an inspiring and uplifting meaning and direction to our individual and social existence, while remaining flexible and firmly founded in science, reason, and the boundless search for improvement.

READINGS

These books are listed because they embody Extropian ideas. However, appearance on this list should not be taken to imply full agreement of the author with the Extropian Principles, or vice versa.

Harry Browne:	*How I Found Freedom in An Unfree World*
Paul M. Churchland:	*Matter and Consciousness*
Paul M. Churchland:	*A Neurocomputational Perspective*
Mike Darwin & Brian Wowk:	*Cryonics: Reaching For Tomorrow*
Richard Dawkins:	*The Selfish Gene*
Ward Dean and John Morgenthaler:	*Smart Drugs and Nutrients*
Freeman Dyson:	*Infinite in all Directions*
Eric Drexler:	*Engines of Creation*
Eric Drexler, C. Peterson with Gayle Pergamit:	*Unbounding the Future: The Nanotechnology Revolution*
Robert Ettinger:	*The Prospect of Immortality*
	Man Into Superman
F.M. Esfandiary:	*Optimism One*
	Up-Wingers
	Telespheres
FM-2030:	*Are You A Transhuman?*
Grant Fjermedal:	*The Tomorrow Makers*
David Friedman:	*The Machinery of Freedom*
David Gauthier:	*Morals By Agreement*
Alan Harrington:	*The Immortalist*
Timothy Leary:	*Info-Psychology*
J.L. Mackie:	*The Miracle of Theism*
Hans Moravec:	*Mind Children: The Future of Human and Robotic Intelligence*
Jan Narveson:	*The Libertarian Idea*
Jerry Pournelle:	*A Step Farther Out*
Ilya Prigogine and Isabelle Stengers:	*Order Out of Chaos*
W. Duncan Reekie:	*Markets, Entrenpreneurs and Liberty*
Ed Regis:	*Great Mambo Chicken and the Transhuman Condition*
Albert Rosenfeld:	*Prolongevity II*
Julian Simon:	*The Ultimate Resource*
Julian Simon and Herman Kahn (eds):	*The Resourceful Earth*
Alvin Toffler:	*Powershift*
Robert Anton Wilson:	*Prometheus Rising*
	The New Inquisition

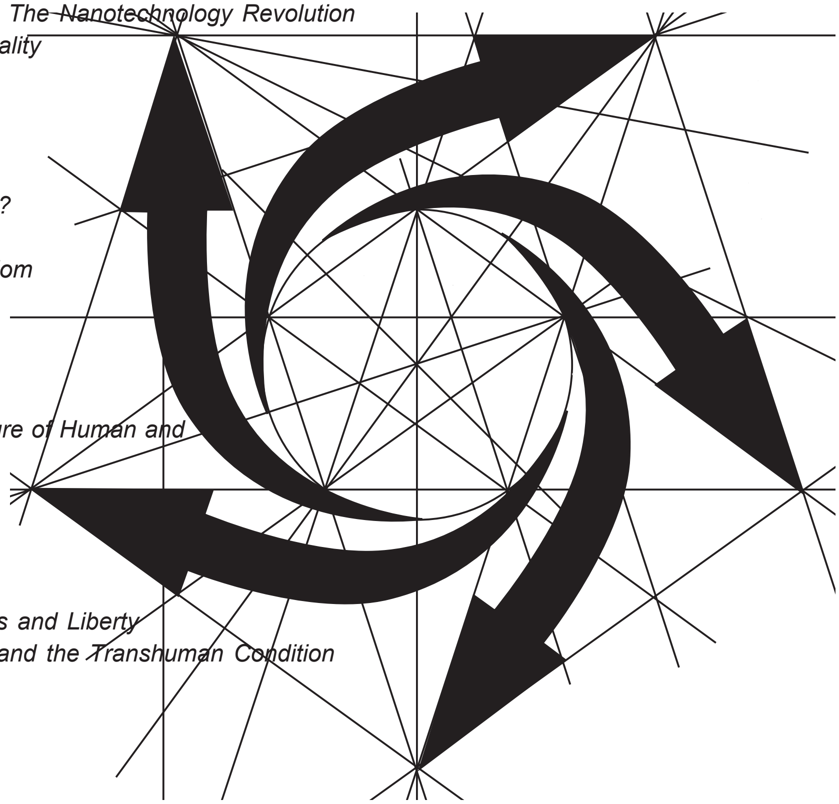

<u>Fiction</u>:

Roger MacBride Allen:	*The Modular Man*
Robert Heinlein:	*Methusaleh's Children*
	Time Enough for Love
James P. Hogan:	*Voyage To Yesteryear*
Charles Platt:	*The Silicon Man*
Ayn Rand:	*Atlas Shrugged*
Robert Shea and Robert Anton Wilson:	*Illuminatus!* (3 vols.)
L. Neil Smith:	*The Probability Breach*
Bruce Sterling:	*Schizmatrix*
Marc Stiegler:	*The Gentle Seduction.*
Vernor Vinge:	*True Names*
	"The Ungoverned" in *True Names...and Other Dangers*

Extropy Institute Launches

Max More
Executive Director

On May 10 1992, Extropy Institute (ExI) became a 501(c)3 California nonprofit corporation. We were able to speed the incorporation process thanks to Fred and Linda Chamberlain, who offered to let us take over the Lake Tahoe Life Extension Festival, change the directors, Articles and Bylaws, and the name. The Festival was incorporated in 1986 and, through it, Fred and Linda (who are also distinguished as the founders of the Alcor Life Extension Foundation in 1972) organized conferences in Lake Tahoe. These conference-festivals were attended by life extensionists and cryonicists from around the world. Extropy Institute intends to restart similar events, though with a broader coverage than life extension and cryonics. ExI also took charge of the records of Lifepact – an outgrowth of the Festival which sought to set up mechanisms ensuring motivations for eventual resuscitation of biostasis patients. We intend to reactivate the Lifepact Project when we can.

Extropians 1988-1992

Extropy came into existence four years ago, in the Fall of 1988, as a small 24 page publication, titled *Extropy: Vaccine For Future Shock*, edited by myself and Tom Morrow (under our old surnames), while we were graduate students in philosophy. There we set out our view of the future, and our agenda:

> After wandering along at a slow pace for centuries, our world has started to enter a period of change that will far outpace historical standards. The changes occurring in the twentieth century dwarf those of any previous thousand years, but they only hint at what the future holds. We face a turning point in history – a time when computers, artificial intelligence, nanotechnology, self-modification, physical immortality and other factors promise to radically transform virtually every aspect of our existence. We are responsible for preparing ourselves for that future, and for helping others understand that coming age.

Extropy was to provide a forum for the discussion of the future for those who shared certain basic values - the values which are now expressed in the Extropian Principles. The world view that Tom and I shared – and soon discovered that many others affirmed – was neatly summarized in Tom's neologism, 'extropy', which included the ideas of increasing information and intelligence, increasing order, and expanding usable energy. Extropian philosophy was the first to explicitly draw together apparently disparate ideas and interests: Individualistic/voluntaristic political views, enthusiasm for technology, especially life extension, space migration, self-improvement, cognitive enhancement, computers and artificial intelligence, nanotechnology, and so on. The readers of *Extropy* frequently expressed surprise and delight that there were others who shared all these interests and had formed them into a coherent world view.

Over the next few years, *Extropy* grew in size, readership, and sophistication. The writing became longer and more analytical, turning the publication into a cross between a magazine, newsletter and journal. In 1991 I changed the sub-title to reflect the publication's evolution and it became *Extropy: The Journal of Transhumanist Thought*.

Also in 1991, thanks to the efforts of Perry Metzger, the Extropians electronic mail list started up, and is now about to celebrate its first anniversary. The list attracted many more people, who discovered they were Extropians. About the same time, I discussed with Tom and Simon! Levy the idea of forming an organization to make possible projects beyond the journal. The Extropy Institute started up, and three other directors taken on board: David Krieger, Russell Whitaker, and Ralph Whelan.

Extropy Institute brings together several other communities and intellectual groupings: The nanotechnology investigators and enthusiasts, mostly associated with the Foresight Institute, the life extensionists and cryonicists, computer programmers and electronic communication proponents, libertarians and other individualists, the hypertext researchers associated with Xanadu, and now some of those supporting the electronic marketplace of the American Information Exchange (AMIX).

Rather than adopting the traditional corporate structure, ExI will operate as an adhocracy. "Adhocracy" is Alvin Toffler's term for a highly networked, non-hierarchical organization, a type found most often in law firms, consulting companies and research universities. Our networked structure, linking members through electronic mail, computer conferencing and electronic bulletin boards will allow us to adapt efficiently to a changing array of projects, each requiring differing skills and resources.[1]

Extropy Institute Projects:

Elsewhere in this issue you can find the revised statement of the Extropian Principles. Extropy Institute's objectives, and present and future projects are guided by those principles. We feel a pressing need for memetically engineering our culture – intelligently applying ourselves to changing those parts of the intellectual culture hostile or indifferent to our values. We want to increase support for life extension, physical and cognitive augmentation, and combat statism, and paternalism. Especially important in the 1990s is combating the false doom-mongering of the apocalyptic environmentalists. These anti-growth, anti-market, anti-freedom, back-to-the-Pleistocene forces threaten all that we believe in. The crisis-brigade environmentalists have found a way to smuggle through the back door the discredited ideas of socialism, fascism and Malthus. Few people stand up for the side of growth. ExI intends to gather together the many people sharing our viewpoint (as summarized in the principle of Boundless Expansion).

***Exponent*:** September 1 1992 will see the publication of the first issue of *Exponent* – a bi-monthly ExI publication, edited by Simon! D. Levy of the University of Connecticut, and author of *Extropy*'s Neurocomputation series. *Exponent* will publish shorter and more topical items, while *Extropy* will continue to publish longer essays. *Exponent* will include news of ExI activities, updates on advances in the crucial technologies of the future, reviews of books, movies, software, and multimedia, and a section devoted to critically analyzing the doomsday claims of some environmentalists. A subscription to *Exponent* comes with ExI membership.

Introduction to Extropian Philosophy: Producing a booklet containing some essential essays explaining the Extropian philosophy is a high priority. This can be sent to people inquiring about the Institute and serve as one element of our educational outreach.

Environmental Rationality Project: The enemies of growth, progress and the continued evolution of our species are legion. Those supporting our goals are numerous but not organized, leaving almost all the influence over the culture to the entropists. Extropy Institute, which already has numerous contacts in diverse fields, intends to serve as a nexus for the interchange of information and strategy for the pro-growth proponents. We will pursue media attention in order to make our case, hold conferences, distribute essays, books, and videos on the environment, growth, and pollution, and we will give lectures and challenge the anti-growth forces to debates.

Digital Economy Project: Electronic communications and computational markets will be increasingly important as the future progresses. We will educate the public on trends and future technologies and social structures to aid in the transformations already underway, and we will conduct research to stay on the leading edge. We have already established a market on AMIX – the American Information Exchange, an electronic marketplace where documents can be bought and sold, and consulting services hired. We will also support the development of user-friendly and powerful encryption systems, so that electronic communications can be secured against governmental and private invasion. Accessible encryption will accelerate the use of computer networks and computational economies (including electronic money and secure electronic exchanges). The Extropians e-mail list has already seen the exchange of many public keys, thanks to the encouragement of ExI Director Russell Whitaker.

Educational Outreach: ExI will develop educational materials concerning the uses of computer technologies, biotechnology, nanotechnology, and psychological means of enhancing cognitive, physical, and psychological capacities in human beings, and the lengthening of the human lifespan. Our educational outreach will include giving seminars and lectures, and making available audio and video tapes of interest to Extropians. Long-term goals include a television show comprehensively presenting the positive potentials of advanced and emerging technologies. Some of the first topics to be made available as seminars and tapes will be a self-programming course, presenting the most effective means of self-transformation, an effective thinking course, and materials supporting Environmental Rationality project.

ExI will distribute to scientific and educational institutions and to interested professionals, and seek to have published, the results of our surveys and analyses, so that the information derived from this work will have the maximum public accessibility and benefit.

ExPress: At first we will gather from diverse sources any books, tapes, and software that interest us as Extropians. Later we may found a press – the ExPress – to publish works by both new and experienced Extropian writers.

Self-Programming Research Project: In support of the principle of Self-Transformation, ExI will research the most effective techniques for physical, cognitive, and psychological self-improvement. There is an enormous amount of both useful and worthless self-transformation information out there, and many Extropians have been experimenting with it for years. We will bring together the best techniques, offering them in the form of seminars, written courses, and tapes.

ExI Conferences: In addition to the social gatherings that have already been taking place, ExI will organize conferences for the exchange of ideas, discussion of strategy, and debating our differences, as well as providing a way for Extropians to meet in enjoyable surroundings. We would like to hold an annual general conference where a broad range of topics are discussed, in addition to more specialized conferences on intellectual property, computational economies, life extension, etc.

Archive: We have begun, in a small way, to build a library and to archive materials and objects of historical interest relating to the development of life extension and enhancement of cognitive, physical, and psychological capacities. These will be displayed at our events. When resources allow, we will establish a library and a museum, which will be open to the public.

BBS: In order to facilitate communication between ExI and its membership, we will set up an electronic bulletin board (BBS). This will hold various documents, and news updates, and will allow provide a forum for discussions and exchange of information in addition to the two Extropian electronic mailing lists already in operation.

Lifepact Project: Lifepact was started by Fred and Linda Chamberlain to set up mechanisms improving the chances that biostasis patients would eventually be resuscitated. A Lifepact is an agreement between persons so that the first person to be revived from biostasis would take responsibility for ensuring the resuscitation of the other. Eventually a network of these agreements would exist. ExI will restart Lifepact, and provide a repository for the storage of information that might be used to reconstruct an imperfectly preserved personality. ExI members receive a discount on deep underground storage arranged through the Alcor Life Extension Foundation.

Idea Futures Project: Idea futures markets, described by Robin Hanson in *Extropy* #8, allow market mechanisms to help settle controversial empirical forecasts. Idea futures offer powerful advantages over other mechanisms for establishing a consensus on the basis of which to allocate research and investment funds. Extropy Institute intends to investigate the legal and organizational requirements for setting up a functioning idea futures market.

Project Extropia: Extropia is an evolving social ideal, not a place. Project Extropia seeks to develop new virtual and actual communities, grounded in individualistic principles and supportive of Extropian values and goals. The ExI office is now located in what we call the Nexus/L.A. – a house inhabited by Extropians where the inhabitants interact according to individualist principles, providing synergistic cooperation while maintaining autonomy (such as by avoiding common property and contracting for internal services). Other Extropians have expressed interest in setting up Nexi elsewhere. We expect to see the growth of a Nexus Network. In conjunction with the growth of electronic communication, the Nexus Network will be a distributed community of Extropians linked by printed, acoustic and electronic communications. In addition, ExI will encourage and participate in a feasibility study of the Free Oceana/Agora Aqua private Extropian community idea already heavily discussed on the Net.

As you can see, even at this early stage we have an ambitious list of projects. A few of these are underway, others are in the planning stages, and some are just dreams. You can make a major difference to the success and progress of these projects. Your unique skills and knowledge, applied to the project you most want to see move ahead, will have enormous leverage at this time. At this stage we especially need to secure funding in order to develop our programs. If you can support us financially, or by helping in any way with fundraising, we want to hear from you!

In this issue, or inserted into it, is a membership information form. We hope the readers of *Extropy* will want to join us as members, receive the newsletter, *Exponent*, and charge with us into the future.

Max More

Extropy Institute

Extropy Institute
P.O. Box 57306, Los Angeles, CA 90057-0306.
Tel/Fax: (213) 484-6383.
more@usc.edu

Extropy #9 was produced on a 386-20 MHz PC, with an 80 Mb hard drive and 5 Mb of RAM (with some help with conversions from a Mac SE/30), using PageMaker 4.0, Word for Windows 2.0, ATM 2.0, Aldus Freehand 3.1. Fonts used are Arial, Book Antiqua, Eurostile, and EurostileBold. This issue was typeset on a VT-100 at the Electronic Publishing Center at the University of Southern California, and was printed by Customized Computer Publishing at U.S.C.

Formats for Submissions

I can accept most formats, converting with varying degrees of difficulty. Preferred text formats are PC-compatible formats such as Word for Windows 1.1 or 2.0, WordPerfect 4.2, 5.0, or 5.1, Wordstar, ascii text, RTF, or other major word-processor format. EPS, PCX, TIFF and other standard graphics formats are fine, and I can accept PICT and MacDraw files, but prefer to receive PC-compatible files.

Deadline for submissions for #10 (vol.4 no.2):
November 1 1992

Provisional mailing date for #10:
January 15 1993

Advertizers: Contact the Editor for rates.

Print run this issue: 750

Exponent
Keep up with Extropian news and ideas between issues of *Extropy* by reading Extropy Institute's bi-monthly newsletter, *Exponent*. See the insert for information.

Extropians on the Net
For a daily source of Extropian information, discussion, and controversy, join the Extropians e-mail list. To join the list send a request to join to manager Harry Shapiro at:
extropians-request@gnu.ai.mit.edu

[1]For discussion of adhocracies and related issues, see "Computers, Networks and the Corporation," Thomas W. Malone and John F. Rockart. *Scientific American*, September 1991.

Persons, Programs, and Uploading Consciousness

David Ross

Abstract:
The pursuit of indefinitely-extended life falls into two categories: repair of our current bodies and movement of our personalities into other bodies. This article deals with a subset of the latter: the uploading of the human personality into a computer. It is divided into four sections, each of them introduced by a question.

1. What would an intelligent program be like?
2. Is the human mind an intelligent program?
3. What does it mean to upload rather than copy a human?
4. How would you upload a human?

What Would an Intelligent Program Be Like?

The short form of the argument for human uploading goes like this. The human mind is a program running on a computer which is the brain. Silicon-based computers are getting more and more powerful all the time. Before long they will equal human brain power and then it should be possible, at least in principle, to transfer the mind-program from the brain to a silicon computer of equal or superior power. Ignoring for a moment the question of whether the human mind *is* a computer program, this line of reasoning conceals a serious misconception, which much be cleared up at the start.

The argument ignores the important distinction between the processing power of a computer and the IQ of a program running on that computer. This distinction is often lost in discussions of machine intelligence. Several authors, Moravec sometimes among them[1], have pointed to the continual exponential growth in desktop computing power, shown that it passes the human-level ten teraflop range about 2025, and from this argued that human-level intelligent programs will become possible around then.

The truth is that far more progress has been made increasing computer speeds in the past thirty years than in increasing the mind-power of computer programs. The two phenomena (increase in computer power and increase in program intelligence) have almost nothing to do with each other. A program with human-level intelligence does not require a super-megacomputer to run. If any form of intelligent program, human or otherwise, is possible, it can run on the simplest form of computer, provided it has enough memory. It will run equally well, and with equal intelligence, on a computer built of tinker-toys as on one with multiple teraflops of capability. It will just run incomparably slower.

If this is true for intelligent programs, it is also true for human personalities uploaded onto computers. The shortcut argument that all it takes is faster machines to make possible intelligent programs and (by extension) human uploading is false. Establishing the possibility of uploading must go by a more circuitous route. Fortunately, following that route gives quite a bit of insight into how to actually perform the uploading.

What is an Intelligent System?

Intelligence is like pornography: We can't define it, but we know it when we see it. We are probably willing to accept a system as intelligent if it "feels" intelligent in the same way as does another human. That is, if it acts intelligently in most of its dealings with us. The same thing is true for self-awareness. It too must be established from the outside. Until such time as we upload ourselves onto a computer and then observe ourselves, there is no way to prove whether any program is self-aware. For now, a self-aware or an intelligent program must be one that can pass whatever tests would convince us of another human's self-awareness or intelligence. This is one form of the so-called Turing test.

Self-awareness and intelligence are not the same thing. It seems quite possible that an intelligent program might not be self-aware, or that a self-aware program might not have human-level intelligence. Nevertheless, since we want an uploaded human to be both intelligent and self-aware, for the purposes of this article we will lump the two concepts together under the title "intelligent".

Many people are uncomfortable with the subjective form of the Turing test, but all objective tests of intelligence in the form of a finite set of pre-established questions must fail. They must fail because a computer could always be programmed to give the same answers to those questions that a human had already given. In the same way, it is incorrect to say that "If only a computer could do 'X' I would consider it intelligent." Time and again computers have accomplished tasks previously thought to require intelligence — playing chess, reading characters — and done so in clearly non-intelligent ways. Finite task testing, like finite question list testing must fail, and the only real way of establishing the intelligence of another entity is subjective.

One reason why there can never be a finite set of tests that establish the intelligence of a program is because whatever else intelligence may be, it is a function of complexity. A "simple" system is one where its basic components — such as the rules in an expert system — are still individually important. A "complex" system is one where the individual building blocks are submerged in levels of hierarchy and complexity. An intelligent system must be complex. Perhaps unsurprisingly, it shares this property with life and with entropy.

Entropy cannot appear if the number of particles involved is too small or the particles are too simple. The equations of motion of a simple system of colliding particles are time-reversible. It is only when the system is sufficiently complex — when the basic building blocks of the system have all but vanished into the background — that entropy can arise. This phenomenon of a high-level property of a system appearing only when the low-level constituents have effectively disappeared is important. A system that demonstrates entropy contains too many particles to deal with individually; they must be tracked statistically. Note, however, that the arising of entropy is not just a statistical trick. Even a system of trillions of particles is in principle trackable using purely time-reversible Newtonian equations. And yet, despite this, such a system has irreversible, entropy-increasing modes of behavior. Entropy is not the same thing as chaos — the inability to track a system because of round-off error.

Life, too, appears only with a certain level of complexity. Very complex molecules, up to about the level of the simplest viruses, are not fully alive. But the more complex viruses, and the simplest single-cell objects, are alive. The system has passed a level of basic complexity separating life from non-life, and again done so when it is no longer sensible to talk about a group of single constituents, but rather of systems of systems or something even more complex.

So it seems to be with intelligence. Only a 'sufficiently complex' system can be intelligent. Some heuristic arguments for why this is so are possible. An intelligent system is capable of generating new information from old information. This is not just a recombining of the old information, but a recombining of it in ways that demonstrate non-obvious links and consequences — in short, in ways which generate new information. But Shannon has shown[2] that information can be expressed mathematically in the same form as entropy (with a change of sign). In effect, information is the opposite of entropy.

To some degree, the generation of previously unavailable information can be seen as reducing the entropy of the system. For the entropy of a system to be reduced, it must be sufficiently complex to have entropy in the first place. Note that it is not the processing of information that requires either complexity or intelligence. It is the generation of new information that does so.

Probably the single most egregious failure of the Expert System school of Artificial Intelligence, led by Marvin Minsky, has been the failure to understand that an intelligent system must be complex. Thus, a simple rule-following system, no matter how well it may mimic a human in some limited regime, is not intelligent, and adding more discrete rules can never lead to intelligent behavior. This, perhaps, explains the failure of the Expert System school of AI to produce intelligent systems.

This failure is, in a way, unfortunate. If intelligence could be the product of a simple rule-based system, then those rules could, in principle, be programmed into a computer. The resulting system would then presumably be intelligent. Uploading a human (or at least copying a human onto a computer — we'll look at the difference later) would then be a question of determining the complete set of rules that specify the human and programming a suitable computer with those rules. This point of view has been one of the driving concepts behind the expert systems view of artificial intelligence from its beginning. Unfortunately, if intelligence requires complexity great enough that the system building blocks effectively vanish, the only way a rule-based system could be intelligent would require it to have so many rules arranged so complexly that the individual rules are effectively lost in the noise. It is difficult to see how anyone could write such a system. Perhaps one could eventually be grown by an automatic rule-writing program — or by several generations of them.

It should be kept in mind that complexity alone is not sufficient to generate intelligence, any more than it is sufficient to generate entropy or life. The Internet, the interlocking system of computer networks that connects most military and educational institutions and an increasing number of commercial ones, has on it computing power collectively approaching that of a single human brain, and that computing power is arranged in multitudes of very complex systems, but it is shows no sign of either life or intelligence. In conclusion, for a program to be intelligent it must be "complex" and it must be able to convince a human of its intelligence. We do not yet know how to write such programs, but they certainly do not seem impossible in principle. Surprisingly, if the human mind turns out to be such a program, we may be able to copy (or transfer) ourselves without completely understanding how the program works, modify the result, and produce truly intelligent programs that way — a kind of reverse engineering of ourselves.

Is the Human Mind an Intelligent Program?

The Computers Can't Think Fallacy

While most students of spontaneous order, and indeed most computer scientists, would claim that intelligent programs are possible, there are several other schools of philosophy that say intelligent or self-aware algorithms are not possible, and that therefore, by extension, a human cannot be uploaded onto a machine without losing his intelligence or his self-awareness. Because this question is central to the problem of uploading, we'll take a slight detour to attempt to answer some of their objections. More detailed refutations of the "computers can't think" school of philosophy can be found in endnote3.

Simulating intelligence versus real intelligence

One false path the "computers can't think" school goes down is when they discuss the difference between a simulation of something and the thing itself. This argument shows up in several forms, many of them made by John Searle[4]. One form of it is that all computers can do is simulate physical processes. They cannot duplicate them.

Searle has correctly pointed out that no matter how well a computer simulates the chemical processes in a flame, it will not produce the flame's heat. This is because a flame is an example of a process that is "essentially material". The simulation of an "essentially material" process will always be different in kind from the process itself.

There are other kinds of processes, however, which are not "essentially material". For these processes there is little or no difference between the process itself and the simulation of the process. Consider addition. I can add two numbers in my head, and everyone would say I was performing addition. Would anyone say that if I programmed a computer to simulate what my brain is doing that the resulting system would not be doing addition? Addition is a process for which the simulation of the process is the same as the process itself.

A form of simulation that is directly relevant to human uploading is the construction of an emulator. If a programmer has a program which he must run on a different kind of computer than it was written for, but for which he does not have the source code, he will build an emulator. An emulator is a special computer program that translates the system calls of one kind of computer into system calls of the other. Thus, if the pointer to the "ADD" instructions is stored in register 18, say, on the computer for which the program was originally written, and in register 32 on the computer on which it must run, the emulator will intercept all calls to register 18 and redirect them to register 32. Except for some degradation in speed, the executing program will perform the same on the new system as it did on the one for which it was developed.

Although the program-plus-emulator is "simulating" the program running on its original computer, there is no doubt that the program is actually running on the new system. In essence, it is the hardware, not the software that is being simulated. The software is just being run. If we can show that a computer program can be a brain emulator — that is, emulate the hardware at a sufficiently low level — then such a computer-plus-emulator should be able to run a human mind-program.

Processes like addition and computer programs are "essentially immaterial". To simulate them is to run them. Searle's argument that at most computers can only simulate thought requires him to show that the functioning of the human mind is an "essentially material" process. He does not do so and I do not believe he can. In fact, I believe thought can be shown to be "essentially immaterial" under the above definitions.

Another argument made by Searle is that computers cannot think because "they are made of the wrong kind of substance to think". This argument is really a form of the first one — that all computers can do is simulate thought. It is based on the same erroneous assumption that a mind is a physical process on the order of a chemical reaction — that it is "essentially material" by my definition. After making this assumption, Searle then establishes correctly that for essentially material processes simulation and reality are two different things. From this he deduces that computers may be able to simulate intelligence, but they can never actually *be* intelligent. In effect, for Searle, the mind is the brain and the brain is a physical protoplasmic object.

The fallacy in Searle's line of reasoning is that it starts by assuming its conclusion. He assumes that the mind is essentially material, shows that simulation of essentially material processes is different in kind from the processes themselves, and then concludes that a simulation of a mind cannot itself be a mind.

The "The Mind is a Computer Program" Misunderstanding

On the opposite side of the intelligent program question there is also a misunderstanding that must be cleared up. This is the "radical dualist" view of the mind/brain question. This way of reasoning looks on the mind/brain system as directly analogous to a program/brain system. Proponents of this view hold, in essence, that the mind is a program running on the hardware of the brain.

Their syllogism goes as follows. Major Premise: Any computer program can be moved from one system to another system of sufficient power without essentially altering it. Minor Premise 1: The human brain is a computer. Minor Premise 2: The human mind is a program. Conclusion: the human mind can be moved from the computer that is the human brain to another computer of sufficient power.

The brain, if it is a computer at all, is pre-von Neumann. Brain structure at every level determines the functioning of the mind. A given brain, with its structure intact, simply cannot run a different program. The wiring *is* the program.

I believe this line of reasoning also to be in error, though not so seriously as Searle's. Its primary failure is that it ignores how much more important the structure of the brain is to the functioning of the mind than the computer's structure is to the functioning of the program.

In a computer, the physical layout and physical connectivity of the system does not change from one program to another. Many different programs can run on the same computer, with the only differences coming at the micron-scale level on which bits are stored in computer memory.

If a mind, on the other hand, is a program, it is a very peculiar sort of program indeed. Von Neuman invented computer programming because he got tired of changing the physical wiring on the ENIAC for each new program. Instead, he wired the computer a set way and modified the data memory registers whenever he wanted to run a different program. Before von Neuman, the program ENIAC ran was determined by its wiring. In essence the wiring and the program were the same thing. The brain, if it is a computer at all, is pre-von Neuman. Brain structure at every level determines the functioning of the mind. A given brain, with its structure intact, simply can not run a different program. The wiring *is* the program.

The central question in human uploading, then, is can we provide a von Neuman-type modification to the brain? Can we change the essential form of the brain/computer from one where the structure *is* the program to one which will run on a general purpose computer? And can we guarantee that it is the same program when we are done?

The answer to these questions is yes. Such a von Neuman modification is possible in principle. Consider the actual change that he made to ENIAC. Before he learned how to program it, the memory registers were used to store initial conditions and intermediate result data for running the program. How that data was used was determined by which wires were connected to which other wires. After he changed the system, it was no longer necessary to physically modify the wiring by hand to make it run a program. The data in the registers, besides holding the same initial condition and intermediate result data, also held the program information — information that ENIAC read and used to flip electromechanical switches that in turn determined which wires were connected to which others. The programming had replaced the human, but the connectivity of the wiring still changed depending on what was being run on the system. All von Neuman changed was what was doing the rewiring.

It is, after all, strictly incorrect to say that a computer running one program has the same structure as one running another. Throughout the CPU (central processing unit) switches are being set and unset millions of times a second, with their configuration determined by the program steps as it operates. Even though we cannot see it, the physical structure of the computer is changing continuously as a program executes, and in ways that are precisely determined by the program. Even at the beginning, the bits that are the program on disk or in memory, before it starts to execute, are held by transistors in physically different states depending on whether a given bit is a one or a zero.

Thus, the difference between a preand a post-von Neuman computer is one of degree, not one of kind. The structure is still being modified as the program executes. What is different is the level on which the modification takes place. At its most basic level, a program is dependent on the structure of the particular computer on which it is running. If the bits that specify the program are not on and off — physically charged one way or the other — then the program isn't even there. If it doesn't have a physical way of putting intermediate and final results into a memory, it cannot run. In essence two different program cannot be run on the same computer, because the programs modify the computer in essential ways when they run.

The real question, therefore, is can a human mind modify the general purpose switching of a general purpose computer instead of the very special purpose wiring of a human brain. Can there be a level of interface above which is the human mind and below which is an emulation of the human brain hardware?

In the mind plus brain-emulator plus computer system it is not really relevant where you draw the line and say from here up is the mind and from here down is the emulator. All that matters is that the emulation reach low enough down (probably at least to the individual neurons) so that it is emulating systems that are below the essential level of the brain. Remember that a mind is a "complex" system. It arises when the basic constituents are such a small part of the overall system that exactly what they are, provided they work correctly, does not matter. I don't believe anyone would argue that any given neuron or any given synapse is important for thought. Even at the most basic level, it is how that neuron interacts with the tens of thousands of other neurons to which it might be connected that matters.

We can successfully run a mind/brain on the computer if our program emulates the brain at the neuron/synapse level. Below that level, the computer is just transistors and switches — a general-purpose von Neuman computer. What runs above the level where emulation starts can be looked at as a brain emulator with a mind program running above it, but there is really only one program. By running the emulator all the way down to the level of the individual neuron and synapse, we are in effect running both the brain (as an emulator) and the mind (as a program) on the computer, regardless of where you consider the brain to leave off and the mind to begin — or even if you consider the question irrelevant. We have finessed the problem of mind/brain duality by going down to the level where the brain itself leaves off.

What Does it Mean to Upload a Human?

There is some dispute as to whether "uploading" or "downloading" is the proper term for moving a person onto a computer. In general computer parlance, "uploading" refers to moving a file from your disk to your computer, or from your local computer to a remote system. Downloading usually refers to copying a file from a remote system to your local computer. If looked at from the perspective of the person doing the file transfer, if the movement is in some sense away from him, the process is called "uploading", while if it is toward him, it is "downloading". If it is the person himself being moved, the terms become confused. As the first elements of the personality are transferred, from the perspective of most of the person, the process is uploading. As the last elements come across, and most of the personality has already been transferred, it looks like downloading. Neither term is completely accurate. I suppose I could follow Max More's lead and coin a new term, say "transloading", but that seems unnecessary. I will adopt the viewpoint at the beginning of the process and call the procedure "uploading".

The main reason people want to upload themselves onto a computer or other non-biological system is to escape death. The point of uploading is to transfer your personality from your body to the computer. However nice it might be to have a copy of yourself running around, that is not enough. It is not enough even if the copy thinks it is you. The goal is to complete a direct transfer of consciousness, and to do it in such a way that a person can be sure beforehand that *he* will make it across onto the computer. The only way for that certainty to occur, at least until enough people have made the transfer that anecdotal evidence of successful uploading is available and believed, is for the person to be conscious throughout the transference.

In deciding whether uploading is theoretically possible we have to make several careful distinctions, particularly between uploading and copying. There are a wide range of philosophies concerning what the self is. For a quite thorough discussion of the current theories and their applicability to uploading, I strongly recommend Max More's Ph.D. dissertation[5]. In discussing whether a consciousness can be uploaded, we need to be careful about directions in time. Were a person copied identically, both the original and the copy, looking backward in time, would feel that they were the original. They would remember equally well events that happened to the person, and they would both swear that the other was the copy. It is important also to note that neither one of them would feel that they would continue were only the other one to survive.

If it is just looking backwards in time — I feel that I was uploaded is all that matters — then I could let someone copy me and then kill my original and not care. However strong might be philosophical arguments that the original self survives in such a situation, to a person looking ahead in time to the transfer such a situation is not much preferable to just being killed. It certainly does not feel like a transfer.

There has been substantial discussion that one of the desirable consequences of uploading is the possibility of making backup copies of yourself to be activated in event of the destruction of the original. A backup, as normally envisioned, is a copy of the original. While, looking backwards in time, it would likely believe that it was the original, it would not *be* the original, since, at the time the copy was made, the original still existed, and presumably had its own opinion about who was there first. By keeping a backup up to date using techniques similar to those discussed below for uploading, however, it should be possible to "upload" yourself to a new program in all but the very worst of circumstances.

It is easy to picture a copying process where it is not possible from the perspective of either the person being copied or people outside to tell which is the original and which is the copy. We really have no way of knowing when we awoke this morning that we were the same person who went to sleep the evening before. Had someone made a copy while we slept and then destroyed the original, we would never know it.

Nonetheless, because we have lots of experience of continuity of self while conscious, we are willing, looking forward in time, to accept periods of unconsciousness — whether sleep or anesthesia — without worrying about whether it will be "us" that makes it through. But, as in lots of other ways, once copying and uploading become possible, our common sense can easily fail us. I submit that while there may be artifacts of the copying process that enable us to distinguish an original and a copy (someone reliable witnessed the process, say, and states that the copy was made at a location remote from the original without the original being disturbed), if the person is not conscious, then except for those artifacts, there is no way for a person looking forward to such a procedure to be sure *he* would survive.

Being conscious through the process, however, should be sufficient for the person involved (who, after all, is the one who matters) to be assured beforehand that he will feel that he will make it through each part of the process to the end with his personality still intact.

How Do You Upload a Human?

If the arguments presented above are correct, we do not need to argue directly that an entire mind can be transferred to computer hardware. All we have to argue is that a nerve and all its synapses can be. To upload a mind we just 'upload' neurons one at a time until we've uploaded everything. It isn't even necessary to understand the functioning of the brain as a whole to do this. As long as the system can recognize what a neuron is and replace it with a nerve-replacement structure (NRS), uploading can be accomplished.

This NRS is the key to uploading. It has two purposes. First, when it replaces a neuron, it must function just like the neuron it replaced. Externally, this means that it must interact with the remaining neurons with which it is in contact just like the neuron did. This must be done through physical actuators that release and absorb synaptic chemicals and electrical impulses.

Internally, the replacement structure must be a program running on a computer that controls the actuators. This program is, in effect, a neuron emulator that learned, in the process of replacing the neuron, how to successfully emulate it.

The second purpose of the NRS comes into use only when all neighboring neurons have also been replaced. At that point instead of interacting with neighboring NRSs through actuators, it emulates the synapse as well. Now, it is all program.

This procedure continues until the entire brain and all sensory structures have been replaced. Since seeing the transfer first-hand is important, let's look at it from the perspective of a person having it done.

Uploading Jason Macklin

Jason Macklin turns his head slightly but cannot see the tube connected to his neck through his carotid artery. Still, he knows it is there, and he knows what it will do. For years he has resisted the urgings of family and friends to get rid of his natural body and upload his mind onto the Web, to become a creature in Cyberspace like them. Unlike nearly everyone else, he doesn't have even a neural tap to communicate instantly with any person or database on the world-wide Cyberspace Web. He sighs. Maybe he has been old-fashioned. If he'd already acquired artificial senses through direct neural implants, the idea of replacing his entire brain with such structures wouldn't be so frightening. Still, the doctors say he'll feel absolutely nothing while the process is going on.

When human uploading first became possible, there was substantial debate over whether the uploaded person was the original or just a copy. A person uploading himself to avoid dying along with his body needed to know that he would really make to transition to Cyberspace. It might be nice to know that a copy of you would survive your death, but not nearly as nice as knowing *you* would. And the only way to be sure that it was really you who made the transfer was to do it fully conscious.

While Jason ponders his decision, the process is already underway. Nanomachines have crossed the blood-brain barrier and are systematically replacing each neuron with a functionally equivalent artificial structure. This structure is very special. Part of it is a series of nanomachine actuators that interact with neighboring cells just as though the replaced neuron were still in place. The rest of his body cannot tell that anything has changed. The remainder of the structure is a program that emulates the neuron and controls the actuators.

A person is not just brain cells. The nanomachines invading Jason's body will replace all sensory neurons as well, and then replace all the parts of his body that influence the neurons with programs to do the same thing. Thus, muscles that stretch and flex will have their program equivalents, and neither the replaced nor the remaining neurons will be able to tell which is which.

On the main computer there exists an artificial world that, as far as sensory input is concerned, is identical to what Jason sees and hears and feels as he lies on his bed. All of the sensory input — and this includes the food in his stomach, the orientation of his limbs, his sense of breathing, as well as sight and sound — is duplicated in Cyberspace. When every neuron and its surrounding tissues have been replaced by their temporary physical counterparts, the structures on the computer that control each replacement part begin to interact among themselves directly, in direct synchronization with how they perform in his body. Gradually, each synapse in his brain is absorbed into the program structure of the emulation program, its functionality retained but its physical structure gone.

Jason's eyes are open, staring at the ceiling, but even as the cells within his eyes are replaced, he has no sense of disorientation. For a time, the unreplaced cells are getting input from photons in the room while the replaced cells are getting input from the artificial world in Cyberspace. Because the synchronization is complete, he cannot tell which part gets which input. After his eyes have been completely uploaded, everything still looks the same.

After a while, a doctor comes into the room and removes the tap into his neck. She holds out her hand and tells Jason to stand. For a moment, Jason wonders why he feels so normal. The doctor's touch, the pressure of the floor against his feet, all feel completely physical. And yet he knows that all the sees and feels now is part of Cyberspace.

The doctor leads him over to a curtain at the side of the room and draws it back. Through the revealed window, Jason looks into another room, identical to the one he is in. On a bed in the middle of the room lies his body, still connected to its cable. For a moment, he watches it breathe. All of its neurons have been replaced, but his new autonomic nervous system still controls his old body. Just as in a dream, however, his conscious movements no longer make his old body move. The doctor hands him a switch which he knows will turn off his old body. He represses the feeling that he is committing suicide and throws the switch. In the next room the body — he no longer thinks of it as himself — releases its last breath and seems to relax. All nerve and muscle connections are severed at once and the body dies instantly.

He feels less emotion than he thought he would. He knows that if he doesn't like it here in Cyberspace, he can always have another physical body constructed, grown from his original DNA, if he wishes. But right now that is not his concern. The room around him appears ordinary and familar in order to ease his transition into Cyberspace, but outside the door is a vast new world. It will take time to learn all it has to offer. But time is what he now has. As an entity running in Cyberspace, with proper backup programs, he is assured of essentially unlimited life.

References

[1]Hans Moravec, *Mind Children: The Future of Robot and Human Intelligence.* (Harvard University Press, 1988.)

[2]Claude E. Shannon, "A Mathematical Theory of Communication." *Bell System Tech. J.* (July, 1948, October 1948.)

[3]Paul M. Churchland, "Functionalism, Qualia, and Intentionality." *Philosophical Topics* 12, no.1: 121-145. Daniel C. Dennett, *Brainstorms*. (Montgomery,

Nanotechnology and Faith

J. Storrs Hall

The following thoughts appeared on sci.nanotech – the Internet newsgroup dedicated to discussing nanotechnology in all its aspects, from the social to the highly technical. It was written in response to other postings which are not reproduced here.

A recent letter from Thomas Donaldson to the cryonics list contained the following:

> 3. "WHEN NANOTECHNOLOGY COMES..."
>
> This comment is a tangent more or less unrelated to the others. It has consistently disturbed me that many cryonicists, particularly those using the above phrase or its variations, speak about cryonics in very much the way 19th century born-again Christians might speak about the Millenium. I've heard a lot about nanotechnology, but just what is this Nanotechnology which is supposed to be coming? (The Christians, of course, had a different word for it. But that's OK. Just change the words and what do you have?).
>
> Would Eric Klein or some other exponent of this world view (theology?) explain in simple terms just what this Nanotechnology is? And please note the capital letter: as I said, I already know a lot about the uncapitalized form.
>
> For instance, I had a very strange experience not long ago. Someone who (I think) is a Believer claimed that when Nanotechnology came, the tiny critters could be used to cure cancer. When I pointed out that almost the same thing, and to the same effect, was happening now by experimental treatments in which lymphocytes were modified and cultured up in large number to attack a patient's cancer, he seemed not to notice, shrugging it off with the statement that Nanotechnology will do much more.

This is a good question, and one that deserves a lot of thought. A major reason for that is precisely to keep our meme complex trimmed of the pseudo-religious memes that commonly attach themselves to any similar vision.

It is vital to distinguish between a vision, which "Nanotechnology" certainly is, and a Faith, which it is not (and which we must constantly guard it from becoming). A vision, in the sense I'm using it here, is a picture of some wonderful future development for which the visionary has some reasonable grounds of belief. I'm using "Faith" to mean a belief system, like a religion, that is adhered to without any reasonable evidence that it is true or possible.

Now "faith" is often used to denote any belief by people wishing to denigrate that belief. For example, I'm occasionally accused of having a "faith in technological progress". I do believe that technological progress tends to make things better for people in general. Detractors from this point of view exhibit any number of social problems and remind us that technology hasn't solved them. This might be a valid argument *if* I had a Faith of the religious variety; i.e. that technology would solve all problems. Of course it won't—and indeed it does create some new problems: If you cure a disease that was killing half the population (e.g. the Black Plague) you must now find a way to feed all these people you didn't have to worry about before.

Thus one of the most obvious distinguishing characteristics between visions and Faiths is that the object of faith is held to be a panacea. Most religious paradises and many ideological utopias fall into this category: "Once we get to X, there simply won't be anything wrong."

It is all too easy to take a vision and hang this meme onto it, which makes it much less useful for either predicting or designing the future. For example, take the vision which some people had around the turn of the century of universal ownership of motorcars. This is a good vision; we as a society and as individuals are considerably better off than we would be without them. However, it would have been silly to imagine the result as idyllic.

The same is true of nanotechnology. If you think our legal system is a mess now, imagine it after any one, much less all, of the nanotech developments that could greatly affect our way of life. Imagine the scale of industrial accidents or terrorism, much less out-and-out war. It's virtually certain that if we do manage to increase our intelligence, we'll increase the complexity of everyday life more than enough to make up for it. Even if I can buy a newly manufactured body totally free from disease, I dread waking up the morning after the warranty runs out.

The next religious meme that we need to look out for is that of believing that one's vision is unique, or the best, or other characterization that causes you to dismiss alternatives without serious consideration. It is easy to see how this meme is advantageous to a belief system in the fierce competition of a memetic ecology; it is also easy to see how unlikely it is actually to be true. This meme finds its ultimate expression in religious wars.

Note that if one is actually working to develop something, some mechanism like this is necessary to focus the effort; but one should focus the effort because the effort needs to be focused to be effective, rather than from an erroneous belief that all alternatives are bad.

Now one can imagine self-reproducing robots using computer and mechanical technology not greatly different from what currently exists; and molecular manipulation without self-reproducing robots. One can imagine many of the effects we anticipate, being done by biomolecular engineering, others by extensions of conventional chemistry, AI being achieved by ingenious innovations in software instead of simulating brains, etc.

Another view is that nanotechnology is simply a name for any technique or group of techniques that manipulates matter at the molecular level. In this view nanotechnology is unique because it is all-inclusive. The trouble with this definition is that it allows one to call anything nanotechnology, e.g. chemistry. If so, then we have nanotechnology now, and it doesn't do all the things we claim. A reasonable definition of nanotechnology must include the notion of a broad and general ability to design, build, and control molecular mechanisms across a very wide range of possibilities.

A third and final memetic attachment that makes a Faith out of a vision is some hook that relates directly to belief and propagation: "Not only should you believe this simply because it is true, but you'll go to Hell if you don't. Furthermore, it's a sin *not* to try to convince others of this belief." In ideological faiths, this meme expresses itself in the social castigation of the Politically Incorrect.

It is convenient to give names to these three memes (remember that a meme corresponds to a gene, i.e. it is the smallest unit of idea replication that is identifiably separable from the overall meme complex). We'll call them Panacea, Incomparability, and Apostasy, respectively.

Nanotechnology seems mercifully free of Apostasy at the moment: No one seems to be claiming it's any kind of sin to disbelieve the Word of Saint Eric. There is some tendency, unfortunately, for nanotechnology as a set of popular ideas, to accrete a bit of Panacea and/or Incomparability. It is our responsibility to continue to scrub nanotechnology to keep it free of them.

Nevertheless, I do believe that it is reasonable to put a fair amount of faith, in the simple uncapitalized sense of the word, in the ability of advancing technology in general to solve a broad range of well-defined physical problems, among them the cure and/or prevention of certain diseases, or indeed the aging process itself. Such a belief is reasonable not only, or even primarily, because we can posit particular mechanisms for the solutions, but because we have a long history of scientific and technological success for just such problems.

THE MAKING OF A SMALL WORLD

Fiction by Mike Perry

[This story originally appeared in *Venturist Monthly News*, April 1989, and is reprinted with permission.]

"You requested audience, apprentice Gorn?"

"Yes O Great Wizard Snorrl, Lord of Galaxies, ruler of Many Worlds, King of Evolved Immortals..."

"Enough! What can I do for you, young fellow?"

"I'm having trouble playing God."

"Not an uncommon thing, your first billion years (to invoke our ancient and honored time unit)... What is your problem?"

"They don't respect me."

"Your charges? Tell me about it."

"Well, first I made this world, got it peopled it with intelligent life, in a nice setting I had made with forests and meadows, creatures that crawled and flew and leaped and galloped, all the usual things..."

"You got a genome permit?"

"Oh yes, all straight evolved lifeforms, nothing tampered with already...?"

"Very good. Go on."

"So then I went among the inhabitants, did good things, healed the sick, fed the hungry, spoke kind words, and... well, they..."

"Put you to a painful death?"

"Very. Only the backup information saved me, and they would have *eaten* that if they could, the miserable vermin. Why if you could have seen –"

"Tell me about it later. How're they doing now?"

"Oh, fine, just fine, ought to be applying for membership soon, which means I'll be in a jam for overpopulating..."

"I wouldn't worry too much, this time. How about your next world?"

"Yes, I did make another one, and that time, I naturally tried to avoid the public spotlight, went around in secret, showing myself to a few only..."

"And...?"

"Well, mostly they didn't believe I existed. And they'll be applying for membership soon, too, and..."

"Argh! So twice in a row you've lost control after only a few thousand years."

"Uh, about 900 years in the last case."

"Oh, my. Well, as you know, you only have one more try at this thing. Maybe you ought to get out while the getting's good, to avoid discredit. Take up cosmological eschatology or something respecta--"

"No! I want to build a world of primitives and keep them that way as long as possible. I want to lord it over them, century after century, millennium after millennium. I want them to sing my praises. I want it to be a *long time* before they become dissatisfied enough to develop and apply for membership and start playing the games we play..."

"Still haven't grown up, eh? Well, the rules entitle you to one more shot."

"So what you suggest is... ?"

"If you really must know...?"

"Of course, why did I request audience?"

"Yes, I suppose you have to have your way. Well, this'll sound crazy, but about the best strategy is to give 'em a good, severe beating every day of their lives."

"What?"

"They'll fear you, they'll respect you, and they'll love you."

[Long pause.] "Yes, I admit there's a certain logic to that, but I'd have to be many places at once... use robots, of course! Big, metallic buzzing things with wings for hot pursuit and clawed feet for grasping and whiplash antennas for striking *hard*. And I know just the creatures to clone and try it on... picked up some genomes on a nice blue planet that was third out from it's primary, I can even recreate some of their original language and culture -- their year is almost the same as ours, by the way --"

"I see your mindwheels are whirring, so I'll leave you be."

"Yes, I must start building this world at once..."

*   *   *

1,000,000 years later, the Daily Globe*, a leading newspaper on the world created by Gorn, reports:*

Weird scheme to defeat Just Punishment; scientists scoff; ethicists howl; legislators vow to stop it.

A group claiming that Just Punishment is "unjust" say they believe it can be "defeated" through science. Simon Burr, spokesman for the self-styled "Committee for the Overthrow of Physical Abuse" (COPA) claims "the robots that administer our daily beatings could be destroyed through technological means," and cites an example where a robot was held at bay for more than an hour while its intended "victim" escaped. Scientists, however, take a dim view of Burr's proposal. Jeffrey Snag, senior researcher at Applied Mechanical and Aesthetics, a firm specializing in technology for improving the quality of life and justice, says, "The idea of interfering with such superlative machinery is just patently absurd. There's no prospect for defeating the robots in the foreseeable future -- they are simply too swift and powerful. Besides, why try for an empty 'freedom from abuse' anyway? What good would it do? Recently we've developed some tight fitting clothing to better distribute the force of the blows, and that's what I consider progress."

Other voices are being raised in defense of Just Punishment and similar practices among humans. Ezeldadeath Bugler-Boss, spokeswoman for the Committee for Ethical Bruising, declares that "Beatings are beautiful, pure and simple. I just bubble with warm feeling over the worth of welts." She is "looking forward to an expanded role for impact therapy in human life," and argues that "a little hand-to-hand combat from time to time could usefully augment the blessings of Just Punishment." Asked about COPA she indignantly concludes, "Our whole society is predicated on the assumption of daily beatings, which we humbly accept as a foundation of our being and a springboard for spiritual growth. When you consider all the benefits -- the stability, the security, the certainty -- of knowing this meaningful experience will always be with us, I don't see how anyone can raise an objection."

However John Crue, a construction worker, admits he is "not entirely happy with the punishment we get for the crime of being alive" and comments further that "being whipped like a horse by giant flying things may have its advantages, but I like it better when they stop. I don't know how I'd adjust to no beatings, but I do consider it from time to time."

But some authorities are so distressed by what they perceive as an affront to the natural order of things that they are taking legal action. Recently the Department of Proper Behavior filed felony charges against COPA for obstructing due process and attempted sabotage. COPA attorney Anthony Sharp denies that his organization has broken the law, arguing that "laws protect human lives and property but there is no law specifically forbidding the sort of practice that COPA is engaged in. The robots are not human property nor an endangered species. To interfere with or even destroy them is no violation of law but simply an exercise of constitutional rights." But DPB officials are sure COPA can be challenged on legal grounds. As Chief Administrator Wilbur McTwitch put it: "The framers of the Constitution wanted to promote individual rights, but the rights of the individual must ever be subordinate to the machinery of great Gorn. I think there is legal precedent to act against those who would attempt a change on so fundamental a level, and if not it *could* be established. I am looking forward to this case."

❐

Extropy reality check: Risks in perspective.

Cigarette smoking (one pack/day)	1600 days
Being poor vs well-to-do	1400
Working as a miner	1000
Being overweight by 30 lb	900
Motor vehicle accidents	200
Small cars vs large cars	100
Being murdered	90
Falls	40
Drowning	40
Speed limit raised from 55 mph to 65 mph	40
Poison + suffocation + asphyxiation	37
Fire, burns	27
Firearms	11
Nuclear power (UCS)	1.5
Nuclear power (NRC)	0.03

Table from B.L. Cohen and I.S. Lee. "A catalog of risks," *Health Phys*., **36**, 707 (1979). Reprinted in Bernard L. Cohen, "The Risks of Nuclear Power" in *The Resourceful Earth*, p.561.

Neurocomputing 6: Genetic Algorithms

Simon! D. Levy

Imagine the following scenario: You are a traveling salesman who must make a tour of a large number of cities that are interconnected by a network of roads. You need visit each city only once. To save yourself time and gas money, you want to take the most efficient route possible, spending as little time as you can on the road. What kind of decision procedure can you use to come up with the best travel plan?

This scenario, commonly called the "traveling salesman problem," is one of a class of mathematical puzzles known as *optimization problems*. For the traveling salesman problem, as for many other optimization problems, there is no general algorithm that will guarantee you the best solution in a reasonable amount of time. Various methods have been explored, with different degrees of success. For example, you might simply measure every possible route through all the cities, and then pick the one that covered the least distance. Of course, this "brute force" solution would become impracticable for anything more than a handful of cities, because the number of possible routes increases very rapidly as you add more cities. Parallelizing the problem by getting a bunch of computers to try one route each, and then comparing the results, is one way to improve the usefulness of this approach.

A second, more common approach involves *gradient descent* along an "error surface" generated by a particular solution and its neighbors. The idea is to take the error of the current solution and compare it to the error of nearby (i.e., similar) solutions, moving to the neighbor with the lowest error. Gradient descent is a fundamental technique in neural network algorithms, especially back propagation. See my Neurocomputing 3 article in *Extropy 6,* and references therein, for more information.

Now consider how Nature handles such problems. You can view the traveling salesman's task as a kind of ecological niche, so that a good solution is an organism that can perform the task without dying of exhaustion or losing out to more successful competitors. Unlike most computer algorithms, though, the solutions developed by Nature are not the result of deliberate planning by a conscious agent. Instead, Nature relies on random mutation and natural selection, as described originally by Charles Darwin in *The Origin of Species*, and more recently by Richard Dawkins in *The Blind Watchmaker* (reviewed in this issue). Good solutions arise by mutation – accidental, minor changes in genes – and are favoredby the environment, allowing the organism bearing the mutant genes to survive, reproduce, and pass these genes on to its offspring. Sexual reproduction, in which an offspring gets half its genes from one organism and half from another, ensures that good solutions will get a chance to combine, producing offspring that may be better than either parent.

Genetic algorithms, first developed by John Holland at the University of Michigan, exploit this cycle of mutation, sex, and natural selection in an attempt to arrive at solutions to optimization problems. The general procedure is delightfully simple and can be described by the following steps:

(1) Start with a set of possible solutions to your problem. If you have no idea of how a good solution would look, just generate these first solutions randomly.

(2) Take each possible solution, apply it to the problem, and examine the results. If you're satisfied with some solution(s), quit here. Otherwise, go on to step (3).

(3) Based on some previously established criterion, reject solutions that fall below a certain level of success at solving the problem.

(4) Create new solutions by splicing together a parts of one successful solution with parts of another.

(5) Every now and then, "mutate" (randomize) some part of some solution.

(6) Go to (2).

As usual, I think it's a good idea to illustrate this procedure with an example. I don't have the paper (or the patience) necessary to fiddle with the traveling salesman problem, so I'll switch to something a little less complicated, namely, a neural network implementation of the Boolean XOR function, an example that I also used for Neurocomputing 3. This function takes two inputs, each of which may be zero or one, and outputs a one if the inputs are different. If the inputs are the same, the output is zero. In other words...

Input 1	Input 2	Output
0	0	0
0	1	1
1	0	1
1	1	0

Now, an interesting problem is how to train a neural network to "become" this function. That is, given the following diagram, where I's are inputs, W's are multiplicative connection weights, sigmas are summations, and thetas are thresholds, we wish to obtain W's and thetas such that the input/output relations of the function are all satisfied.

As readers of *Extropy 6* will recall, one scheme for getting succcessful network parameters is back-propagation, whereby the network's actual output is compared with the desired output, and the difference between actual and desired output is propagated back through the network, to adjust the weights and thetas.

Of course, our present concern is genetic algorithms, so the question arises as to whether we can "breed" a network – call it an XORganism – to do the XOR problem. To examine this question, I wrote a little C program that worked as follows. (As usual, this program is available to *Extropy* readers at no cost, by sending a letter to me care of the magazine.) Each organism was represented, naturally, by five weights and two thetas. These were the seven "genes" of the organism. "Mating" consisted of producing a new organism with half the genes of one parent and half the genes of another. Since organisms had seven genes each, the program flipped a coin to determine which parent the seventh gene came from. The program started with a collection of 100 organisms with random genes. For each generation, the program computed each organism's score on the XOR problem, and killed off any organism that got fewer that three out of four correct on the problem. Then, the program bred the resulting organisms according to the mating scheme just described. In addition, the program mutated (randomized) a randomly picked gene in some randomly picked organisms every few generations. Parameters of the program were (1) the number of matings per generation, (2) the minimum score necessary for survival, (3) the number of generations between mutations, and (4) the number of organisms to mutate in a mutating generation.

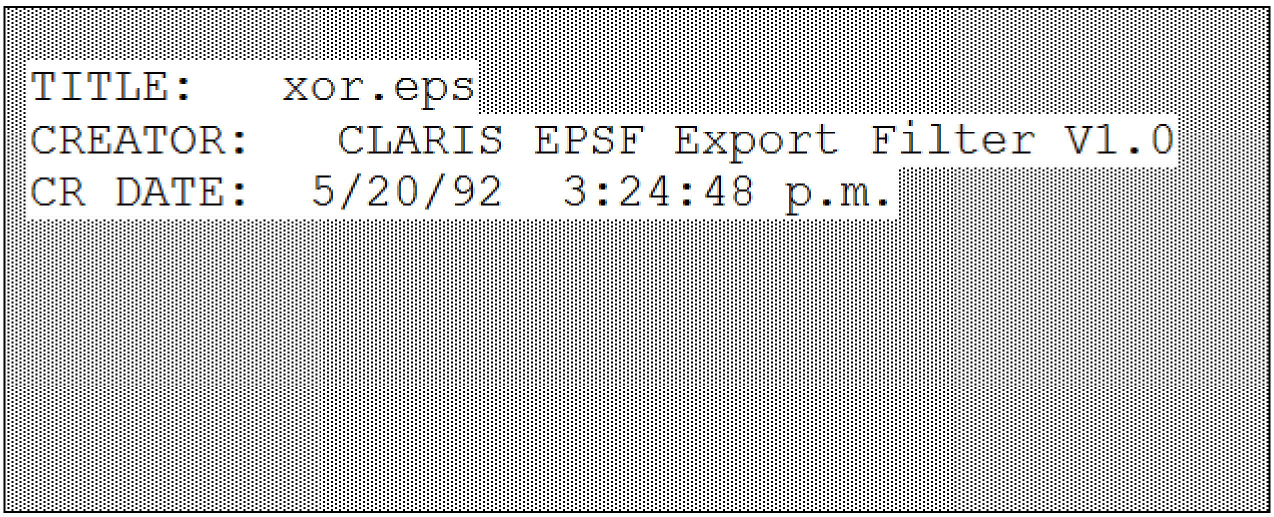

I didn't have time to explore possibilities that would result from fiddling with all the parameters, but some things about the XORganisms seemed pretty clear after several runs of 100 generations or so. First, the number of completely successful organisms (score = 4/4) was always much smaller than the number of organisms that got a score of 3/4. This result is what we expect, because the niche we defined was a score of 3/4, meaning that 4/4 scorers was over-achievers. Second, the emergence of these over-achieving "uber-organisms" seemed to depend a lot on initial conditions. If the initial random creation of 100 organisms produced one with a perfect score, that organism tended to reproduce itself, making lots of other over-achievers, but if no over-achiever existed from the start, chances were fairly good that none would emerge. Finally, adding mutations seemed to have a beneficial effect on the number of over-achievers produced. For example, I compared 10 runs of 100 generations, all with a single organism being mutated every generation, against 10 runs of 100 generations with no mutations at all. The average number of uber-organisms produced with mutations was around 32, whereas the number produced with no mutations was about 9.

Using genetic (or neural net) algorithms to model Boolean functions may be overkill, like earning a Ph.D. in mathematics so that you can balance your checkbook. Nevertheless, the XORganism example is instructive because it shows how a simple neural net problem can be solved by the technique of genetic algorithms, thereby showing one way in which genetic algorithms fit into the rubric of neurocomputing. Genetic algorithms have also caught on for "real-world" applications. One usage, not surprisingly, is the solution of optimization problems similar to the traveling salesman problem described above. These include the design of aircraft and VLSI chips. Other applications include an image-recognizing program that looks at many sub-programs, linking together those that run best, and a program for studying the behavior of ants through simulation.

In the next installation of this column, I'll discuss a topic that seems to interest many members of the Extropians mailing list, namely, the idea of modeling economies computationally, or modeling computers economically. ■

SOURCES

Books

A.K. Dewdney. *The Armchair Universe.* New York: W.H. Freeman. 1988 (See especially the chapter entitled "The Evolution of Flibs")

J. Holland. *Adaptation in Natural and Artifical Systems.* Ann Arbor: University of Michigan Press. 1975.

Articles

M. Antonoff. "Genetic Algorithms: Software by Natural Selection." *Popular Science,* October 1991.

R. Morin. "A Look at Genetic Agorithms." *SunExpert Magazine,* November 1990.

P. Wayner. "Genetic Algorithms." *BYTE,* January 1991.

Time Travel and Computing

Hans Moravec

 Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213
(412) 268-3829

The last few years have been good for time machines. Kip Thorne's renowned general relativity group at Caltech invented a new quantum gravitational approach to building a time gate, and, in an international collaboration, gave a plausible rebuttal of "grandfather paradox" arguments against time travel. Another respected group suggested time machines that exploit quantum mechanical time uncertainty. The technical requirements for these suggestions exceed our present capabilities, but each new approach seems less onerous than the last. There is hope yet that time travel will eventually become possible, even cheap.

Most time-machine fiction deals with the sociological implications of temporal trips or messages—indeed the time travel is often a mere literary device for placing humans in unusual situations. A recent paper from Thorne's group, by contrast, examines time travel at a basic physical level, deriving the quantum mechanical wave functions of systems that contain temporal loops. This article looks at the situation at an intermediate scale—the uses of time travel, packaged as negative time delay elements, on computation. This view is interesting because, on one hand, it predicts colossal improvements in the ability to solve important problems, and, on the other, provides a crisp logical metaphor for macroscopic implications of time travel.

A Brief History of Time Travel

When H.G. Wells wrote *The Time Machine*, his first novel, in 1895, the scientific world was unimpressed. In that Victorian age Science was in the process of crossing the t's and dotting the i's on Newtonian mechanics, and worrying about unemployment in the coming age of fully codified physical knowledge. Time was a rigid universal framework for the clockwork processes of physical law, without the slightest hint of a mechanism for evading its unvarying progress. That time machines were provably impossible was conveyed convincingly to generations of students, and echoes to this day.

The physics revolution at the beginning of the twentieth century shattered objective certainty about the immutability of time, but not most physicists' gut assurances. Special relativity combined space and time into a single continuum, with velocity a kind of rotation that transformed one into the other. A barrier—the speed of light—still separated the "spacelike" from the "timelike", but it was a fragile one. There was nothing in special relativity itself to preclude the existence of particles, now dubbed tachyons, that always moved faster than light. A tachyon message returned by a distant, rapidly receding relay could arrive before it was sent, a consequence construed by the conservative majority of the physics community as an indictment of tachyons. The impossibility of time travel (or closed causal loops, or, in the language of relativity, closed "timelike" loops) had become an axiom of physical law. Tachyons have not, in fact, been detected, even though they should be creatable with arbitrarily little energy (the faster they move, the less it takes), so perhaps the conservative majority is correct in this case (but perhaps they just have a tendency to hide - see below).

General relativity patches together tiny regions of flat special-relativistic spacetime into large gravity-warped structures. Powerful gravity fields imply radically convoluted spacetimes. Kurt Gödel was first, in 1949, to notice that general relativity predicted time travel under certain circumstances. In Gödel's solution to Einstein's equations, the centrifugal tendency of a rotating universe exactly balances its tendency to gravitationally collapse. In such a universe the spacelike and timelike directions are skewed sufficiently that a spaceship accelerating around the universe can arrange to return to the place *and time* of its launch, giving the crew an opportunity to wave *bon voyage* to their departing younger selves. General relativity has been repeatedly confirmed experimentally in the large scale, so those who dislike the prediction take solace in the fact that our universe appears hardly to rotate.

The next major class of solutions, made in the 1960s by Roy Kerr and Ezra Newman and colleagues, are harder to dismiss. The Kerr-Newman solutions of the Einstein equations are for rapidly rotating and/or charged black holes. In the most extreme of these, the rotation of the body counteracts the gravitation enough to expose the twisted viscera of the black hole (normally hidden behind a discreet, one way, event horizon) including regions of *negative* spacetime, from which a spaceship could return to the outside universe before it entered. Attempts by the conservative majority to find independent reasons for a *cosmic censorship* rule to prevent such lewd exposure have been unsuccessful, so far. Their only comfort is that it would take about the mass of a galaxy, with extraordinary spin, to make a practical time machine this way.

In 1974 Frank Tipler published another solution to the general relativity equations, this time for the region around an extremely dense, very long rapidly spinning cylinder—dense as a neutron star, the diameter of a city block, with its surface moving at about one fourth the speed of light, and infinitely long, because that simplified the mathematics. Spacetime wraps itself around such an object like a roll of paper, producing alternate layers of negative and positive spacetime. A carefully aimed spaceship could swing around such an object, staying mostly in a negative region, and come out before it left. A finite length cylinder should also work, and might allow a time machine with only the mass of a star. But, the conservatives say, maybe it's not possible to prevent a finite cylinder from gravitationally collapsing lengthwise.

As yet, no one has devised a satisfactory comprehensive single theory that combines gravity and quantum mechanics—many try, and the implications of such a theory promise to be awesome. In 1988 Kip Thorne and company, patching together partial theories, described a time machine using both quantum mechanics and general relativity. A tiny, spontaneously formed, gravitational spacetime wormhole is pulled out of the hyperactive froth that is the quantum vacuum, and stabilized, by two large conductive plates, resembling an electrical capacitor. Initially these plates are as closely spaced as possible, and each becomes host to one "mouth" of the wormhole. When they are separated in our normal spacetime, they yet remain connected through the wormhole, which is an independent spacetime tunnel. Regardless of their external separation, a message or object entering one mouth appears instantly (by its own reckoning) at the other, as if the mouths were the two sides of a single door. Thorne's group then uses special relativity to differentially age the two mouths. One is taken on a "twin paradox" round trip at near the speed of light, so that less time elapses for it than its stationary counterpart. When it returns, the external separation between the two mouths has a time as well as a space component. A message sent into the itinerant mouth exits from the stationary one after a delay. And a message delivered into the stationary mouth exits from the traveller before it was sent! This kind of machine could perhaps be constructed with a planet's worth of aluminum spread out into plates the area of Earth's orbit, separated by the size of an atom—still beyond our means, but getting closer. The non-linear equations of general relativity are notoriously hard to solve, and only the very simplest cases have been explored. Even more significantly, there isn't any theory of quantum gravity yet at all. It's quite possible that in all this unexplored territory, waiting to be discovered, lie quite feasible ways to build time machines.

Another approach to time travel asks why it isn't observed routinely. There is no intrinsic time direction in Newton's mechanics nor in the differential equations of the new physics. The future determines the past just as fully as the past determines the future. Why, then, can our past selves leave messages for our future, but seemingly absolutely never the other way around? This question has not been definitively answered. Attempted explanations involve "boundary conditions", the initial values of physical quantities at the edges of space and time, which the differential equations of physical law then fill out. The universe must be somehow very different at one end than the other, and this difference orients the arrow of time. Most common is the thermodynamic explanation, which talks about the state of disorder or "entropy" of matter and energy. The universe started in a very rare, highly ordered state, and is running down into increasingly common states of disorder. Though this explains why a ship can't run its engines by separating water into steam and ice, it does not explain why one can't send today's lottery numbers into yesterday by expending a few megawatt hours of energy. One explanation which *does* was offered by John Wheeler and Richard Feynman in 1945. They noted that Maxwell's equations, the first "modern" physical theory, give two solutions for the effect of accelerating an electric charge. One, called the retarded wave, follows the acceleration and describes an electromagnetic disturbance diverging outward at the speed of light—the radio waves which link our civilization. The other, called the advanced wave, is for a similar disturbance that precedes the acceleration and converges on it (or, in another way of looking at it, diverges backwards in time). This latter wave is never seen. Wheeler and Feynman's analysis assumes that the advanced wave is, in fact, produced, and expands outwards into the past. There, eventually, it encounters a condition, perhaps the extreme density of the big bang at the beginning of the universe, that reflects it, producing a retarded wave exactly out of phase, that retraces its spacetime and exactly cancels it out. The retarded wave from an accelerated charge is not cancelled in an analogous manner, because there is no reflector in the future of the universe to reverse it—perhaps because the universe is "open", and expands forever. If this is a correct explanation, it might be possible to send signals backward in time by means of a kind a reflector that acts like the big bang—maybe a black hole. Such a reflector would reverse and return a retarded wave, cancelling and thus apparently preventing it. If the reflector were installed one light year away from a light source aimed at it, it would suppress light from the source one year *before* installation. Similarly the suppression would go away one year before the reflector was removed. Messages could be sent a year into the past simply by moving the reflector in and out of the beam. Recently John Cramer devised an interpretation of Quantum Mechanics, called the transactional model, that uses this approach to explain every interaction. Transmission of a photon from one place to another actually involves two signals, one moving forward in time, the other backward, a "handshake", between the two locations. In Cramer's model a time communicator could be built that works much like the Wheeler-Feynman type, but using a wave absorber to prevent a transaction, rather than a reflector to cancel it.

This kind of machine could perhaps be constructed with a planet's worth of aluminum spread out into plates the size of Earth's orbit, separated by the size of an atom – still beyond our means, but getting closer.

There is a spookier possibility. Suppose it is easy to send messages to the past, but that forward causality also holds (i.e. past events determine the future). In one way of reasoning about it, a message sent to the past will "alter" the entire history following its receipt, including the event that sent it, and thus the message itself. Thus altered, the message will change the past in a different way, and so on, until some "equilibrium" is reached—the simplest being the situation where no message at all is sent. Time travel may thus act to erase itself (an idea Larry Niven fans will recognize as "Niven's Law"). This situation can be modeled quantum mechanically. If the message is a particle sent to the past, the wave function for that particle will subsequently propagate into the future, where it encounters and "interferes" with its original self, cancelling or reinforcing depending on the relative phase (which depends on the round trip length and other things). Now, the resultant wave function indicates the possible places the particle might actually be found in a measurement—large magnitudes are likely locations, near zero values unlikely ones. A round trip that causes the particle wave function to be cancel itself means that a particle is unlikely ever to be positioned to start the trip. It is exactly this kind of quantum probability effect (without reference to time travel) that powerfully confines electrons to discrete shells about the atomic nucleus and causes the "Pauli exclusion principle" that prevents two indistinguishable electrons from being in the same place simultaneously. The Wheeler-Feynman advanced wave theory, Cramer's transactional model of quantum mechanics, and even quantum electrodynamics, the most accurately verified physical theory we have, all involve interactions that reach backwards as well as forwards in time. It may well be that time travel is as common as dirt, and shapes our physical laws, but conspires, by wave function interference, to prevent any operations that would result in logically inconsistent situations, and so makes overt macroscopic time travel difficult. Even with a time machine you will never succeed in preventing your own birth, or changing the antecedents of any present observation—some odd co-incidence or accident (perhaps one disabling your time machine) will always thwart the attempt. But this does not rule out carefully contrived *logically consistent* causal loops (The recent paper by Friedman, Morris, Novikov, Escheverria, Klinkhammer, Thorne and Yurtsever examines the consequences of such constraints on the wormhole time machine mentioned earlier). Wave interference typically appears as banded patterns with a strong central (zero order) peak surrounded by a dark fringe, itself surrounded by the another (first order) bright fringe, and so on. If causal loops behave this way, it will be necessary to make major perturbations in experimental arrangement to skip from the usual zero order (no overt time travel) situation to a first order case of a non trivial consistent causal loop.

Let's suppose that eventually some approach results in a practical, even cheap, way to package time machines to make *negative time delay elements*—which output signals that predict what their inputs will receive some fixed time later. Such devices would have very interesting consequences for computational problems, which increasingly means almost every field of activity.

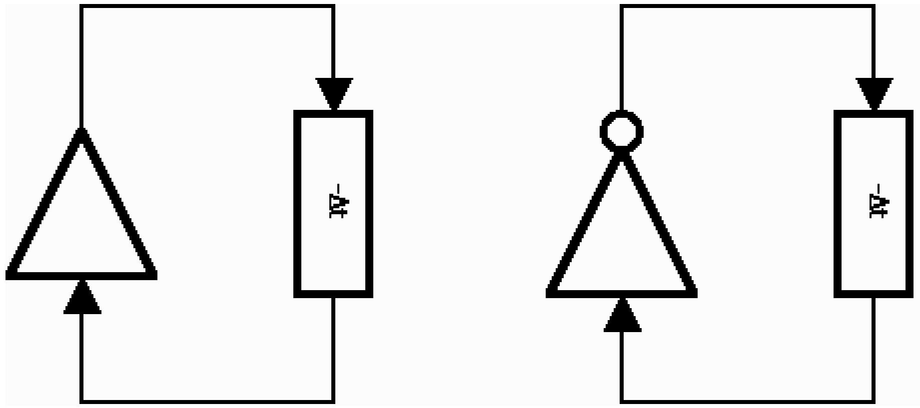

Causal loops involving an amplifier (triangle) and inverter (triangle with circle) in causal loops with negative time delays (-Æt) that cancel their forward delay.

Time Loop Logic

Computer circuitry is composed of devices called gates that combine binary signals to produce other such signals. The simplest gate is an amplifier whose output is identical to its input. Almost as simple is the NOT gate, or inverter, whose output is 0 when its input is 1, and vice versa. All gates take a little time to respond to changes in their input—typically a few billionths of a second. When the input of an amplifier is connected to its output, the circuit will lock up with its output permanently at either 0 or 1. A NOT gate in a similar loop tends to oscillate rapidly between 0 and 1, at a rate that depends on its delay. It is possible to slow down this oscillation by inserting extra delay into the loop. Conversely, a negative delay element would speed up the oscillation. Imagine the two cases of an amplifier and a NOT gate with a certain delays each with its output connected to its input through a negative delay device. (See illustration above.)

The amplifier circuit is in a consistent causal loop—when first switched on, it can permanently assume either 0 or 1 without contradiction. The loop with the inverter, on the other hand, is a simple case of the classical time travel grandfather paradox, a paradoxical causal loop. An input of 1 to the inverter gives an output of 0, which is brought back in time to contradict the input. It takes some quantum mechanics to make sense of the situation, and we will have to say something about how the signals are physically represented. Most digital circuits represent signals as electrical voltages or currents in wires, which is inconvenient because electrons interact with each other and with matter in complex, hard to analyze ways. Some experimental circuits use much simpler space-crossing beams of light . Let's suppose 1 and 0 are encoded as coherent light beams of opposite phase (perfectly out of step with one another—one crests where the other has troughs). In that case a 0 that meets a 1, as in the inverter circuit, will simply cancel. Either alternative would have zero net probability, and the circuit (perhaps containing a charged laser, ready to emit a beam) should simply fail to turn on (ignite) at all, somewhat like a ball balanced on a knife edge that, against all odds, teeters indefinitely instead of falling to one side or the other. This is Niven's law at work in the small. The circuit finds itself perpetually in a dark fringe of an interference pattern.

On the other hand, suppose the two signals are encoded as beams of identical phase but opposite polarization (horizontal or vertical direction of vibration). Then the light in the inverter circuit will be in step with itself, and so should ignite. But what about the polarization? Neither of the possibilities individually results in a consistent situation, but quantum mechanics allows these to be superimposed to produce a *mixed state*. The mixed state that contains equal measures of 0 and 1 is changed by the inverter to an indistinguishable mix of 1 and 0, and so allows the circuit to remain consistent. Light with equal amounts of vertical and horizontal polarization is *unpolarized*. The inverting causal loop insures that the light remains perfectly unpolarized. An attempt to change this, for instance by slipping a polarizing filter somewhere into the circuit, should cause the beam to extinguish, as for phase modulation.

Imagine that the beam in the polarization modulated inverter circuit is sampled, perhaps by a partially silvered mirror, and examined. Quantum mechanics tells us that each extracted photon will be in a mixed state of polarization until the polarization is measured, at which point it "collapses" to one possibility or the other. Sometimes the photon is in a 1 state, equally often it is found to be a 0. But suppose, as in existing computer circuitry, signals are represented not by single particles but by bulk aggregates which behave almost like classical continuous variables. An aggregate of random values will act like a signal intermediate in value between 0 and 1, say around 1/2. In classical terms, it makes sense that a circuit that converts 0 to 1 and 1 to 0 should leave some intermediate value unchanged:

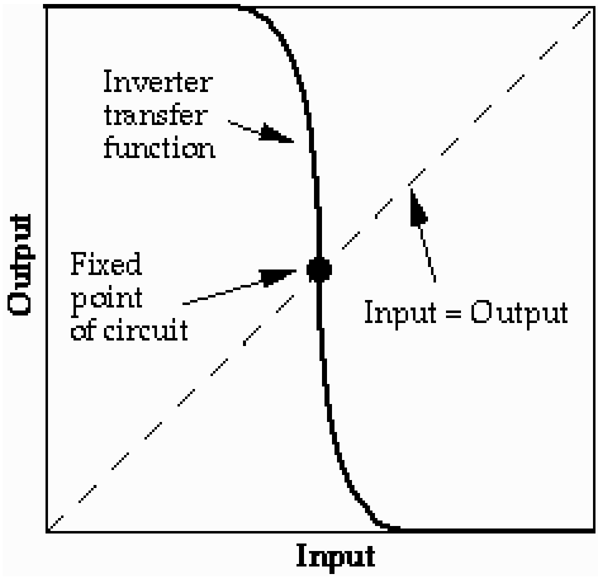

There are other ways to view the situation. An inverting loop without a negative delay will oscillate between 0 and 1 at a frequency that depends on the time delay of the inverter. As we insert increasingly long negative delays in the circuit, the frequency of the oscillation increases. The mixed state is when the frequency reaches infinity, and the circuit is in the 0 and 1 states simultaneously. Note also that the the inverter's 1/2 state is extremely unlikely—a small deviation in the input would cause the output to saturate at one of the two extremes. Yet the circuit indefinitely maintains precisely that unlikely situation. The negative delay connects regions of spacetime in the same way that the orbit of an electron around an atomic nucleus connects different parts of its wave function, and so creates electron shells separated by empty regions. Like a focussing lens that brings to a single spot photons that would otherwise have landed all over a target, it produces interference patterns that make a few states of the world very likely and and the rest unlikely. It's just that, by eliminating the standard possibilities, negative time delays make likely things that are otherwise nearly impossible.

Infinite Iteration

Here's a more complicated case. Make a computing box that accepts an input, which represents an approximate solution to some problem, and produces an output that is an improved approximation. Conventionally you would apply such a computation repeatedly a finite number of times, and then settle for the better, but still approximate, result. Given an appropriate negative delay something else is possible:

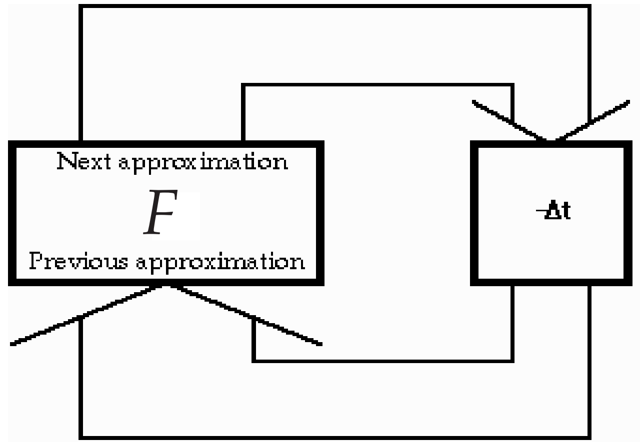

In this arrangement the result of each iteration of the function is brought back in time to serve as the "first" approximation. As soon as the machine is activated, a so-called "fixed-point" of F, an input which produces an identical output, usually signaling a perfect answer, appears (by an extraordinary coincidence!) immediately and steadily, just as either 1 or 0 appears in the simple amplifier loop. If the iteration does not converge, that is, if F has no fixed point, the computer outputs and inputs will shut down or hover in an unlikely intermediate state, like the inverting loop.

The Compleat NP Machine

Generally speaking, whether a problem can be solved in a computer depends on its *computational complexity*, which describes how the difficulty grows as the problem increases in size. Finding the largest number in a list of numbers takes time proportional to the list length. A graph of computing time versus list length yields a straight line—the complexity is *linear*, and this is an easy problem. Sorting the list into numerical order is harder. Simple methods take time proportional to the square of the list length—the graph is an upward curving parabola—and even the best possible sorting methods give slight upward curves. The cost of solving a system of *n* equations in *n* unknowns grows as the cube of *n*, and other problems grow as the fourth, fifth or even higher powers of the problem size. But any problem whose difficulty can be expressed as a fixed power of size is of *polynomial* complexity, and even large instances are soon solvable in a world where computer power doubles every few years. Not so problems of *exponential* complexity. These multiply in cost with each fixed increment of problem size. Whereas a linear problem grows as 1, 2, 3, 4, 5, 6, ..., and a cubic as 1, 8, 27, 64, 125, 216, ..., an exponential problem's cost may grow as 1, 10, 100, 1000, 10000, 100000 and so on to the astronomical. An important class of apparently exponential problems are called NP, short for Nondeterministic Polynomial, meaning exponential problems that could be solved in polynomial time given an exponential supply of computers to examine alternatives. Examples include many design problems such as finding the best arrangement of logic gates, or the fastest program, to compute a given function, and also finding limited length proofs of mathematical theorems. Solving such problems could synergistically increase the power of the machines that solve them. Even poetry or music writing might fit, if approached as the problem of finding the best sequence of constrained words or sounds to express an idea or a state of mind.

The hard core of NP problems are called NP complete, and it 's been shown that a fast solution for any NP complete problem can be translated into a fast solution for any other. One famous and convenient NP task is the *traveling salesman problem*—given *n* cities and the distances between pairs of them, find the shortest tour that passes through each city exactly once.

A version of the time loop iteration box of the last section can quickly solve such a task. The F box for this problem takes as input a particular tour, that is, a permutation of the cities. It also has a knob whose position specifies a limit on the length of the tour. The box calculates the input tour's length, and outputs the same tour if that length is less than or equal to the limit.

If the length exceeds the limit, the box generates the next permutation (by some permutation counting scheme). When the circuit is activated it settles down immediately to a tour shorter than the limit, or

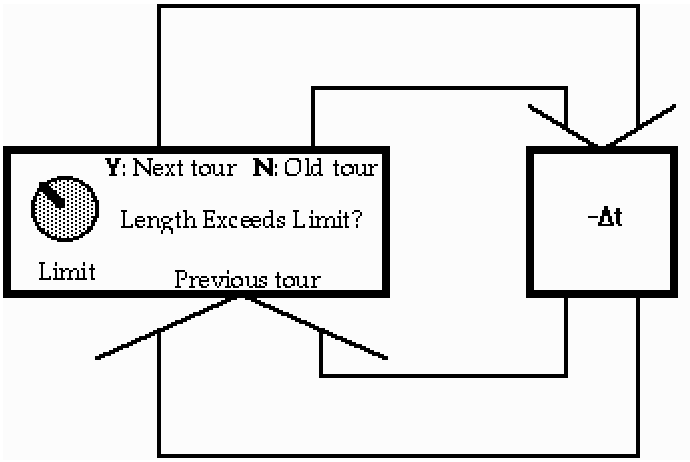

sticks in an undecided state if no such solution exists—i.e. if the limit was set too low.

Turning the knob to find the boundary between these two conditions solves the problem. The "knob turning" can, of course, be done automatically by additional machinery, for instance an ordinary computer. In fact, one can imagine building a very general "chronocomputer" by simply slipping a sufficiently large "negative delay register board" into a peripheral slot of any computer.

Solving Chess

NP problems quickly exceed conventional computer capacity, but of exponential problems they are the easiest to solve. In an NP problem there are an exponential number of candidate solutions to be considered, but the correctness and cost of each candidate (a salesman tour, a circuit for a function, a mathematical proof) can be checked easily, a fact our time-loop NP computer exploits. But there exist harder exponential problems where evaluating a single candidate answer is itself an exponential problem.

Finding the best move in a game like chess is an instance of such a problem. Your best move can be found by considering all possibilities for your move, then for each of those, all possible responses by the opponent, and all your possible responses to all of those, and so on, until a "tree" of all possible games has been mapped out. Then, working backwards from the "leaves" of this tree, i.e. the final moves of the games, all but the best moves can be pruned away. It is easy to evaluate possible last moves—from the point of view of the player whose turn it is, some are wins, some are losses and some result in draws, and no finer distinctions need be made. When all but the best last moves are discarded, the relative merits of second-to-last moves become apparent—each is given by the best last move that results from it. This pruning proceeds to earlier and earlier moves until the beginning of the game, and the best first move, is met. The chain of best moves from first to last is a perfect game of chess.

One can imagine a kind of generalized chess played on a variety of board sizes, with 8 by 8 giving the standard game. Finding a perfect game is easy on small boards, with few possible moves, but becomes enormously difficult as the board grows. The standard game is already so difficult that, despite known mathematical shortcuts, there is no hope that a conventional computer, even one using all the time and matter in the universe, could search the entire game tree. Today's chess computers examine a few levels of the tree, and use a formula to (very imperfectly) guess the value of the rest.

But time loop computers seem to greatly transcend the power of conventional machines. Can the approach used for NP problems be applied here? If, instead of city tours, the input to the computation box is the move sequence for an entire game, the box could be arranged to stabilize only on sufficiently "good" games. But, unlike traveling salesman tours, games have only three values—win, lose or draw (for white, say). There seems no easy way to assign a number to indicate how nearly perfect a game is—and a perfect game itself is recognized only by considering and eliminating all other possible games.

There may be a devious way, however, to use negative delays to fold the massive tree search in time, in a fashion that makes the NP solver look positively pedestrian. This construction is going to resemble a mathematical induction. Suppose we have a circuit (call it the *n* box) that, given a particular chessboard position *n* moves into a game, is able to immediately tell us its value (win, lose or draw), as if it had searched the entire move tree from there. We could then build a box that provides the same services for a position *n*-1 moves into the game by adding a "single move unit" that takes the *n*-1 board, and, one by one, generates the possible next moves (typically about 30 of them), and feeds the resulting boards to the *n* box. It compares the *n* box results to one another, and chooses the best (for the player whose move it is). The time taken to do this is cancelled by a negative time

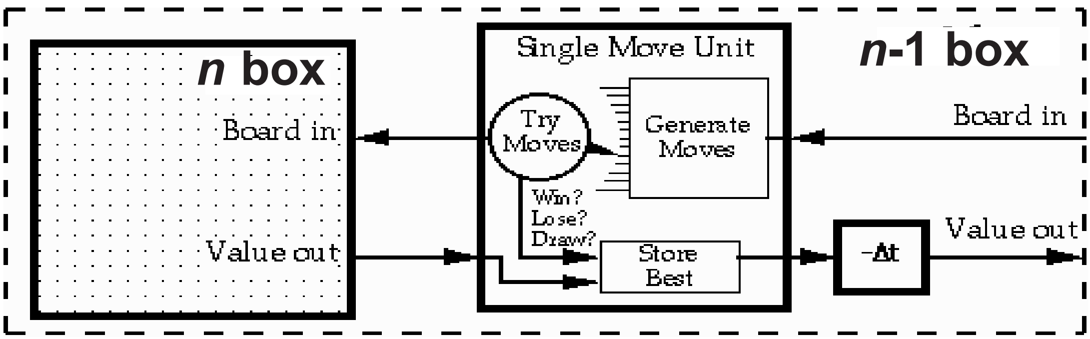

delay element, and so the *n*-1 box, like the *n* box, produces the value of a position as soon as the position is presented. :

The *n* box is itself composed of a single move unit and an *n*+1 box, and so on. The chain can be stopped eventually because each single move unit has a "short circuit" path that it uses when it encounters a position that, because of a king capture or excessive move count, is already a win, loss, or draw and so requires no further search. The whole machine is simply a chain of single move units and negative time delays as long as the longest game (a few hundred moves for standard chess), with the last unit set to permanently return a draw, signifying a stalemate.

Feed it the standard starting chess setup and the machine will immediately indicate whether the game is a forced win for white, black or a draw. To play (perfect) chess, give the machine the result of each of your possible moves, and choose one it finds best. But what is the machine up to in these deliberations? Having accepted time travel logic, the operation of the first single move unit is easy to intuit—it simply evaluates a few dozen possibilities and returns the answer, through the agency of a negative time delay, before its deliberations are complete. The second stage is more mysterious. Its input is changed by the first unit after it has delivered an answer, but while it is still doing the calculations for that answer! It is being asked to do several dozen different computations simultaneously. The third and subsequent units are required to be in even more bewildering states of superposition. The machine would certainly not work if its interior were being observed, as quantum observation collapses states of superposition to particular possibilities. Each single move unit would be seen to receive new boards to evaluate before it had completed the previous ones. Even weak observations (like hearing a whisper through a heavy wall), or indirect or delayed ones, would cause occasional collapses. To have any chance of giving correct answers, the machine must be totally protected from any peeking, perhaps by burying it in many layers of exotic shielding, or by placing it out of range of observation. Even then, its operation depends on aspects of reality hidden from present day experiments, such as the actual independent existence of the alternative worlds that appear in quantum mechanical calculations. Time will tell.

The Non-Computable

Exponential problems are big, but finite. Our universe may not be large or long lasting enough to solve serious instances of them without time travel, but it is possible to conceive of universes that are. But some problems are so hard that no bounded universe would be sufficient—their complexity is infinite. Kurt Gödel, who discovered rotating-universe time travel in general relativity, is much better known for shocking the mathematical community with his incompleteness theorems, showing that in any consistent and sufficiently interesting mathematical theory, there are unverifiable truths.

In the seventeenth century Pierre Fermat wrote that he had discovered a truly remarkable proof that there are no integer solutions to the equation $X^n + Y^n = Z^n$ when X, Y, Z > 1 and n > 2, but unfortunately the margin of the book he was annotating was too small to contain it. Subsequent generations of mathematicians have searched in vain for this (or any) proof to "Fermat's last theorem". Nor has a counterexample been found to prove it wrong. Most probably there was a subtle, unnoticed, flaw in Fermat's "proof". Today the theorem remains a good candidate for a Gödelian true but unprovable statement in arithmetic. If so, we won't be able to prove that it is unprovable, since such a proof would imply that no counterexample (providing a negative proof of the theorem) can be found, thus proving that no such counterexample exists, thus proving that the theorem is true, contradicting its unprovability.

But Gödel's theorems hinge on the finiteness of proofs. We might try to evade them with time loops. Suppose, as in the NP machine example, we build an ***F*** box that takes as input a quartet of numbers {X,Y,Z,n} and tests them to see if they constitute a counterexample to Fermat's conjecture. If so, the same numbers are presented at the output of the box. If not, the "next" quartet is generated and fed via a negative time delay back to the input. When the device is switched on, it immediately shows a counterexample to the conjecture, if such exists, proving the conjecture false. If , on the other hand, the box's signals hover in an intermediate state, the theorem must be true.

There's a problem, of course. A finitely sized machine can examine numbers only as wide as its signal paths. The time loop gives a lot of leverage, since, in an instant, it examines a number of cases exponential in the number of digits in these paths. But the number is still finite, and if the machine fails to find a counterexample, one can't be sure there isn't one that just happens to exceed its capacity. A search for any other kind of proof or disproof founders on this same Gödelian impasse—though a time loop machine can examine all alternatives of a given length, a wonderful proof may yet elude us because the machine we've built is too small to contain it.

In situations like Fermat's last theorem we may willing to forgo an astronomically (or infinitely) large proof or counterexample if we are told by a trustworthy source that the statement is true or false—a compact answer that fits in a very small machine. The time computer techniques we've discussed thus far examine, in a short fixed period, a number of cases exponential in the size of the machine. It may be possible to substitute time for space, to build a small, continuously operating machine that examines an exponentially increasing number of cases as time passes. Take, for instance, a single move unit from the chess machine example, and connect its output to its input via a negative time delay. Observed, the system acts as a simple time loop, as in the NP example. But unobserved, the move sequencing might spawn multiple alternatives, just as in the chess solver, but of the original unit itself rather than units downstream. As time goes by, each copy repeatedly splits itself again, and the number of parallel machines grows exponentially. After a few hundred iterations, this single unit will have covered the same tree of alternative games as the extended chess machine. But by letting it run beyond that point, longer games can be examined. The answer is found by the peculiar procedure of inputting the initial chess board to the single move unit and peeking at its output value at the start of the computation, then tightly closing an observation-proof box around it. The answer you receive right away is trustworthy only if the machine is subsequently left to run, undisturbed and unobserved, for long enough .

For more general problems, let's consider an abstract device conceived in the 1930s by British mathematician Alan Turing, inspired by David Hilbert's concept of "mechanizing" mathematics. A Turing Machine is a finite, usually simple, computer or "finite state machine" connected to a "read/write head" that moves back and forth on an indefinitely long tape. The machine proceeds in regular steps, at each reading a symbol on the square of the tape under the head, then, controlled by an unambiguous finite internal list of rules, writing a symbol in its place, and moving one square forward or back. By convention, the internal rulebook for any given Turing Machine is fixed, but its tape can be initialized with an arbitrary string of symbols. Turing showed that such machines, generally using lots of tape to store and reference initial inputs, intermediate results and final answers, existed for every conceivable computation. (See next page.)

Turing also showed that there exist machines that interpret the initial contents of their tape as a rulebook of another Turing Machine, and proceed, slowly, to simulate this other machine. These are called "Universal" Turing Machines, for their ability to do any computation that another machine can do. A Universal Turing Machine is a good mathematical model for a digital computer, and Turing used the concept to prove "non-computability" theorems that are equivalent to (but more straightforward than) Gödel's "unprovability" results.

By appropriately initializing its tape, one can program a Universal Turing Machine, like a conventional computer, to do just about any computation, for instance to search through {X,Y,Z,n} quartets looking for a counterexample to Fermat's Last Theorem, or through inference chains for a proof or disproof. A literal Turing Machine, though, is very slow—it spends most of its time crawling back and forth

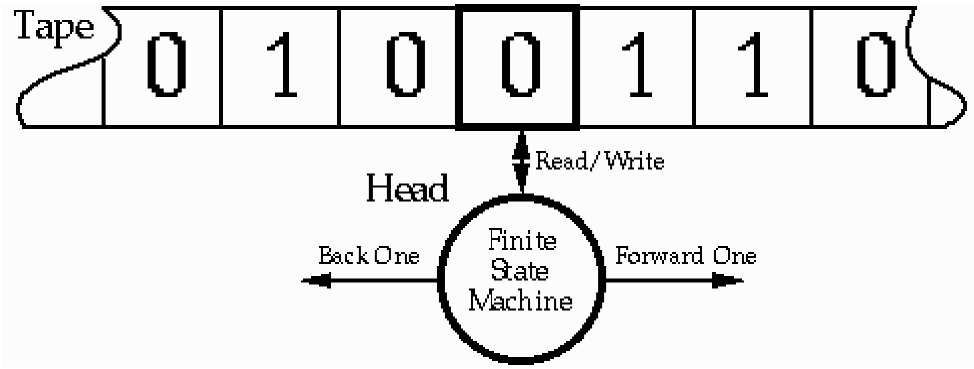

across increasingly long stretches of tape to get at widely separated bits of information. But it is amenable to a time-computer transformation. Imagine that the tape runs not in space, but in time. At any instant it is just a single square (probably a circuit) storing a symbol. Instead of moving left or right, the head sends a

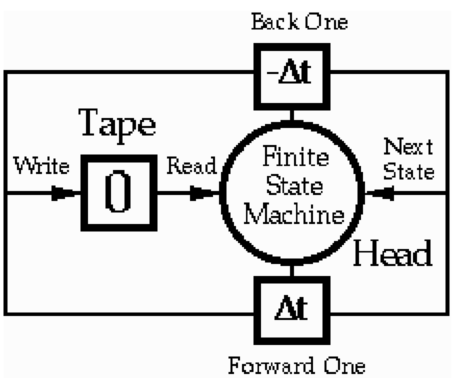

message to its past self through a negative delay, or to its future self through a positive delay:

Many questions could be posed about this design. A few have answers. Because the machine is created at a particular time, and does not exist before then, the "tape" ends abruptly in the backwards direction. This is not a problem because in any computation there are straightforward ways to "fold" tape usage to exclusively the forward portion. Interesting computations will use a lot of tape—i.e. will extend far into the future. But if the program copies its ultimate answer to the cell initially under the head, the answer will be immediately available. As with the the folded chess machine, one must start the "Temporal Turing Machine", quickly peek at its tape cell for the answer, and then seal it up, to let the computation run undisturbed, unobserved. But then how is the program installed on the tape, on cells corresponding to time periods when the machine is sealed? The easiest answer is to encode it in the finite state portion of the machine rather than the tape.

In its sealed box, the head effectively runs forward and backward in time, repeatedly overwriting cells, spawning multiple worlds ever faster. Yet the number of situations the machine can examine (different Fermat's Last Theorem quartets, for instance), can at most grow exponentially in the machine's running time. This is truly prodigious by conventional standards, but still quite finite. So here we are, as far out on a speculative limb as I'm willing to climb, and infinite problems remain hopelessly out of reach. It's time to quit.

Time Travel, Consciousness and Reality

The above constructions are no less strange when applied to "conventional" time-travel scenarios. Imagine that the computation boxes actually contain human beings—a person can quite plausibly evaluate the length of a traveling salesman tour, or enumerate the possible next moves in a chess game. Assuming, as before, that all observations must be logically consistent, what would a human, so embedded in a time loop, experience? When the NP machine produces a solution, its computation box evaluates only one case—a correct one, since quantum interference cancels others. To the person in the computation box, no less than to those outside, a correct answer appears, as if by magic, on the box's input—found correct, transcribed to the output, and relayed back in time, it also creates that input—a circular but logically consistent situation. On the other hand, there is little reason to believe that worlds containing incorrect solutions can be experienced, since such experiences would involve logical contradictions.

Things are not so straightforward for persons inside the chess machine—let's put a friend in each single move unit. The opening position unit feeds first moves one after another to the rest of machine (the "2" box), and examines the reponses—no difficulty. But the second unit is required to be in a superposition of states, simultaneously examining each of the first move's alternatives. When made of dumb machinery, it was reasonable to simply seal it off from observation, to prevent spoiling the computation by collapsing the superposition to a single possibility. But an intelligent friend will probably need to be released sooner or later, and will retain a memory of the computation. In the strange logic of quantum mechanics, this memory, released from the box, constitutes an observation that retroactively collapses the computation. We could, of course, preserve the superposition and the computation by callously keeping the boxes with our friends permanently sealed, or dropping them down a black hole, or accelerating them away beyond observation. But then there will be no way to learn about their experiences. So, apparently, one can either exploit parallel worlds, or experience a single one of them, but not both. "All the world's a stage" wrote William Shakespeare, and apparently we players act in only one story at a time. But when we close our eyes and listen closely, we hear from the wings the echoes of other stories. Whether they really are other plays (with audiences?), or just sound effects by a clever stage manager, remains undecided.

References:

recent time-travel articles:

Michael S. Morris, Kip S. Thorne and Ulvi Yurtsever, **Wormholes, Time Machines, and the Weak Energy Condition**, *Physical Review Letters* **61**#13, pp. 1446-1449, 26 September 1988.

John Friedman, Michael S. Morris, Igor D. Novikov, Fernando Escheverria, Gunnar Klinkhammer, Kip S. Thorne and Ulvi Yurtsever, **Cauchy Problem in Spacetimes with Closed Timelike Curves**, *Physical Review D*, 1990.

Yakir Aharonov, Jeeva Anandan, Sandu Popescu and Lev Vaidman, **Quantum Time-Translation Machine**, *Physical Review Letters* **64**#25, pp. 2965-2968, 18 June 1990.

excellent summaries:

Nick Herbert, ***Faster Than Light: Superluminal Loopholes In Physics***, NAL Books, New York, 1988.

Barry Parker, ***Cosmic Time Travel: A Scientific Odyssey***, Plenum Press, New York, 1991.

Turing machines and computation:

FUTIQUE NEOLOGISMS 3

Compiled by Max More

"Extropians are future shock wave riders." - Jay Prime Positive.

The Futique Neologisms series is intended to provide a compilation of neologisms regarding advanced and future technologies and their applications. Some of them are purely for fun, but others may become entrenched in regular usage. Some of the terms listed are already widely accepted in transhumanist/extropian cultural groups.

ADHOCRACY – A non-bureaucratic networked organization. "This form is already common in organizations such as law firms, consulting companies and research universities. Such organizations and institutions must continually readjust to a changing array of projects, eah requiring a somewhat different combinations of skills and other resources. These organizations depend on many rapidly shifting project teams and much lateral communication among these relatively autonomous, entrepreneurial groups." (*Scientific American*, September 1991, p.133) [Alvin Toffler, *Future Shock*, 1970]

ASIMORT – (a) A dead science fiction writer. (b) A dead secular humanist. (c) Any person who believes it to be their duty to die to make room for later generations. [Mark Plus, April 1992]

BIEVOMECH – (pron. bi-evo-mech) Biological methods and mechanisms evolved through natural selection. Examples: Hearts, lungs, skin, gills, DNA, ribosomes, and analogous non-teleologically evolved mechanisms in xenoevolutures. [Tanya Jones, Jay Prime Positive, Ralph Whelan, December 1991]

BOGOSITY FILTER – A mechanism for detecting bogus ideas and propositions.

BROADCATCHING – "Catching television and other media selectively so that the sum of the collected parts is personalized." (Quote by Nicholas P. Negroponte, *Scientific American*, September 1991, p.112.) [Coined by Stewart Brand, *The Media Lab*, 1987.]

CALCUTTA SYNDROME – The condition in which the ratio of available mass to population falls below the minimum level necessary to support a given quality of life ($M/P < m_c$). [David Krieger, November 1991]

CYBRARIAN – Net-oriented information specialist. [Jean Armour Polly, 1992]

EXTROPIATE – Any drug that has extropic effects, including all cognition enhancing and life extending drugs. [David Krieger, December 1991]

EVOLUTURE – An organism produced through evolution; the antonym of *creature*. [Mark Plus, June 1991]

KNOWBOTS – Knowledge robots, first developed by Vinton G. Cerf and Robert E. Kahn for National Research Initiatives. Knowbots are programmed by users to scan networks for various kinds of related information, regardless of the language or form in which it is expressed. "Knowbots support parallel computations at different sites. They communicate with one another, and with various servers in the network and with users." (*Scientific American*, September 1991, p.74.) [Corporation for National Research Initiatives]

MORPHOLOGICAL FREEDOM – The ability to alter bodily form at will through technologies such as surgery, genetic engineering, nanotechnology, uploading. [Max More, April 1992]

NEG – Someone who typically complains, moans, and whines. Someone practicing the contrary of dynamic optimism.

PARTIAL – A computer simulation of part of a person's personality, created in order to carry out a task not requiring the entire person. [Greg Bear, *Eon*, 1985]

PERICOMPUTER – Any small, portable computing device such as a laptop, an electronic pocket calendar, or Apple's forthcoming Newton. [Lawrence G. Tesler]

POWERSHIFT – A transfer of power involving a change in the nature of power, from violence to wealth, or from wealth to knowledge. [Alvin Toffler in *Powershift* 1990]

SMART BAR – A bar at which smart drinks are sold. Smart drinks usually contain choline and/or an amino acid precursor (such as phenylalanine, tyrosine, or glutamine) for stimulatory neurotransmitters.

SMART-FACED – The condition resulting from social use of cognition-enhancing drugs: "Let's get smart-faced." [Russell E. Whitaker, December 1991]

UBIQUITOUS COMPUTING – Also known as "embodied virtuality." Computers that are an integral, invisible part of people's lives. In some ways the opposite of *virtual reality*, in which the user is absorbed into the computational world. With ubiquitous computing, computers take into account the human world rather than requiring humans to enter into the computers methods of working. [See Mark Weiser, "The Computer for the 21st Century" in *Scientific American*, September 1991.

VITOLOGY – The study of any life-like system, including biology and artificial life. [Max More, December 1991]

XENOEVOLUTURE – An evoluture from a planet other than Earth. [Jay Prime Positive, December 1991]

The source of unattributed neologisms is not known to the compiler. Please send any corrections and additions, c/o the editor.

Exercise and Longevity

Fran Finney

Who wants to grow old? If you are a regular *Extropy* reader, it is a fairly safe assumption that aging is not one of your favorite fantasies. Many of us have looked at length into various methods to slow down (ideally to stop) the aging process. A large number of these methods are experimental – some are quite expensive, some have no hard data to confirm that they actually do much good, and some even have mixed data suggesting that they could cause the opposite of what we wish to achieve. There is, however, one simple factor that we can add to our lifestyle that can be arbitrarily cheap, has little, if any risk of shortening our lifespan, and, increasing evidence shows that by making this simple factor a daily part of our lives, we can certainly retard some aspects of the aging process. That factor is regular physical stress to the cardiovascular and musculoskeletal systems, of long enough and intense enough duration to evoke a physiological response without causing irreparable damage, or, in a nutshell, ***exercise*****.** This article is not meant to suggest that exercise alone should be considered as an adequate life extension regimen, but, added to other methods, it is certain to have positive effects.

A large number of physiological changes are associated with getting older. If we focus on physical deterioration we see decreases in athletic performance, caused by losses in:

- Muscle - Mass, strength, and performance speed;
- Aerobic capacity (V02max); and
- Joint flexibility;

And changes in body composition evidenced by:

- Decreased muscle mass;
- Decreased bone density; and
- Increased percent body fat

(see, e.g. [Cisar 91], [Crammer 87], [Pickles 83], [Shepard 90], [Shephard 91])

Although all of the above changes have been unequivocally associated with growing older, the causes are not clear. Are these all normal developments within the aging process itself, or could they be caused by the decreased activity levels that usually accompany advancing age?

And, of course, there is the big question of life expectancy - not only do we want to remain as youthful as possible throughout our lives; we want to live longer as well! What effect, if any does exercise have on life expectancy?

Muscle Mass, Strength, and Performance Speed

Most people believe that the human body achieves its maximal muscular strength between the ages of 25 and 35 years. Although there is a great deal of variation in performance from person to person, typically, an individual's strength declines at a steady rate after this age. In a sedentary adult, over time, muscular losses in mass, strength, and speed appear to be interrelated.[Costill 90A]

Let us first look at some of the physiological changes to muscle tissue that accompany aging. At birth, a person is bom with essentially all the muscle fibers he/she will ever have. These muscle fibers increase in length as the child grows, and increase in cross-sectional area as demands are placed on them. (The cross-sectional area of a muscle is roughly proportional to its strength.) As a typical, sedentary adult ages, we see a decrease in the actual number of fibers in a muscle, a decrease in the size of the individual fibers, and a decrease in total muscle mass. The loss in number of fibers seems to be selective - most of the fibers that are lost are the fast-twitch, or type II fibers - fibers associated with quick bursts of strength and speed.[Costill 90A]

But, these changes are not universal. Barrie Pickles states in *Biological Aspects of Aging:*

> "If skeletal muscles are used frequently, they show remarkably few structural and functional changes with age. The majority of changes noted in the muscles of elderly persons are characteristics of disuse rather than age."[Pickles 83]

Regular physical activity appears to minimize the mass and strength losses seen with aging A person who exercises regularly throughout his/her life may not show a significant decrease in either number or size of muscle fibers.[Costill 90A] Even the selective loss in fast-twitch fibers might not be necessary Costill speculates they are reabsorbed by the body in response to the inactivity observed with aging. Older adults who are physically active have been reported to have faster movement times than younger inactive adults.[Cisar 91]

How about a sedentary individual who becomes more active? Can such a person undo the effects of his/her previously inactive lifestyle? Once a muscle fiber has been lost, it cannot be replaced. However, individual muscle fibers can grow thicker. Furthermore, when type I fibers are subjected to chronic stimulation at a high frequency (i.e. speedwork) they alter their structure to resemble type II fibers. [Pickles 83]. So theoretically, a person should be able to replace bulk, strength, and speed in atrophied muscle. How does research bear this out?

Maria Fiatarone showed that significant gains in muscle strength could be shown in older adults who exhibited "atrophy of disuse". Volunteers aged 86 to 96 participated in an 8-week exercise program consisting of three sessions per week of progressive-resistance exercises. Participant's strength improved by an incredible average of 174%, muscle girth increased by 9% and speed improved by 48%.[Fiatarone 89,90] Other researchers have also reported marked improvements in muscular strength and mass.[Cisar 91] (Other cases of improvements in speed at which an activity is performed have also been reported, but it is not clear whether these improvements were due to true physiological changes, or whether they were due to a "practice effect".)[Pickles 83]

Aerobic Capacity (V02max)

V02max is considered to be the best measure of cardiovascular capacity, and is referred to by many sports medicine experts as the single most important indication of fitness level.[Shangold 88] It is defined as oxygen consumption at the point at which it fails to rise despite increasing exercise intensity. V02max is a function of the maximum rate that oxygen can be carried to and be utilized by body tissues.

An athlete has a much higher V02max than a sedentary individual. Average V02max in men declines with age almost 50%, from 48ml/kg-min at age 25 to 25.5ml/kg-min at age 75.[Costill 90B]

Two longitudinal (long-term) studies demonstrate how regular cardiovascular exercise as an individual ages can help prevent age-associated decrease in V02max. Fred Kasch demonstrated this in a study taking place over a span of 23 years. Thirty athletes were in the study - fifteen who had worked out regularly and vigorously over the entire period - and fifteen "controls" who had stopped training for at least 18 years. The average age at the start of the study was about 47, and at the end about 70. V02max declined 13% in the group that continued to exercise; in the de-trained group, the decline was 41%.[Kasch 90] A different study done by David Costill showed a 0.5% decline in V02max for competing track athletes (aged 50 to 82) over a period of 10 years, and a 14% decline in a similar group that stopped competing over the same 10 year period. [Costill 90B] According to Costill:

> The decline in aerobic endurance seen throughout life appears to be affected more by the intensity and volume of regular exercise than by aging *per se.*

All of the groups in both longitudinal studies, even the active ones, did show at least a slight decrease in V02max. Most researchers attribute this to an age-related drop in maximum heart rate, or the fastest rate that an individual's heart is physiologically capable of pumping at effectively. Can enough intense exercise prevent a drop in maximum heart rate? At this point the jury is out on this! At any rate, the slight drop in V02max noted in the competing track athletes in Costill's study was not functionally significant.

Cross-sectional studies (comparing different groups) also show that exercise does affect cardiovascular capacity. Endurance athletes consistently have higher V02max values than sedentary individuals of the same age.[Kavanagh 90], [Cisar 91], [Shephard 91].

Joint Flexibility

To understand changes in joint flexibility that accompany aging, first we need to look at the connective tissue that holds joints together and provides cushioning and lubrication for the joints. All connective tissue cells contain a variety of substances, including: collagen, which provides structure and tensile strength; elastin, which provides elastic flexibility; glycoproteins that enable tissue to retain fluid and stay well-hydrated; and hyaluronic acid, to provide lubrication. As a healthy, active adult ages, dynamic collagen production continues, in response to stresses placed on the body. However, production of elastin, some glycoproteins, and hyaluronic acid diminishes. Clinically, we see the following:

1. The diameter of collagen fibers in a given location of the body increases. This increases the tensile strength of the connective tissue, but reduces its elasticity.

2. The amount of elastin in the skin, bronchial tree, and large arteries is reduced; the elastic properties of these areas are therefore reduced.

3. Lowered production of fluid-holding glycoproteins causes progressive tissue dehydration.

4. Reduced hyaluronic acid secretion causes increased joint friction.

To make matters worse, elastin is not produced at all in an older adult, so when a person's elastic connective tissue is injured, the tissue is replaced with inelastic, collagenous based connective tissue. Cartilage (which provides cushioning and some stability to the joints) also cannot be replaced. Therefore, as a person ages, cartilage is worn away.[Pickles 83] All of the above changes can decrease flexibility and make movement more difficult. Exercise, while it does not seem to alter these physiological changes, can make degenerative changes in the joint functionally less restrictive. Although damaged cartilage can not be replaced, regular non-stressful exercise has been shown to reduce degeneration in articular canilage. [Pickles 83] True elasticity in the joints does diminish, but functional flexibility as demonstrated by range of motion can be maintained through a program of stretching. [Cramer 87], [Cisar 91]. And increased joint activity, although it does not cause production of elastin, will increase the tensile strength of the ligaments holding the joints together, by increasing the collagen production - which will strengthen the joints and make them more resistant to injury.

Changes in Body Composition

Three measurable structural components of the human body are muscle, bone and fat. [Shangold 88] In the general population, with increasing age, we see a decrease in muscle mass, a decrease in bone density, and an increase in adipose tissue (fat.) Regular exercise can Fevent loss of muscle mass as a person ages, and can also increase muscle mass in older individuals. Let us now look at how exercise can affect bone density and percent body fat as a person ages.

Bone Density

Although people tend to view bone as an inert substance, it is actually a very active tissue. Bone is constantly being remodeled, in response to stresses and forces placed on it. Assuming an adequate diet and no hormonal abnormalities, the stronger the force applied to a section of bone tissue, the stronger that section of bone will become. Likewise, when less force is applied to a section of bone, excess material is reabsorbed by the body,and the bone becomes lighter and more fragile.

Bone mineral content starts to decline at about age 40 to 50 years, with a subsequent loss of about 10 percent per decade. It is possible that (aside from the bone loss associated with decreased estrogen in women at menopause) thesmaller amount of bone in the elderly could be entirely caused by an ageassociated decrease in activity.[Pickles 83] Less activity leads to a reduction in the stresses and strains placed on the bone, which triggers the body's reabsorption mechanisms.

It appears that both weight bearing and muscle action are important in the maintenance of normal bone density. Studies where weight bearing alone (in paralyzed patients) or muscle movement alone (in astronauts) were attempted to reduce bone loss, showed that either by themselves was not very effective. However, combining compression of a bone with the simultaneous activity of its overlying muscles (i.e. weight bearing exercise) does produce increased density and cortical thickening of the bone.[Pickles 83]

Leslie Pruitt worked with post-menopausal women to see how a weight training program would affect their bone mineral content. At the end of nine months the bone mineral content in their spines had increased, while it decreased in a control group of women who did not exercise.[Pruitt 90] Other researchers show similar results in both sexes.[Cisar 91], [Crammer 87], [Shephard 91]

Body Fat

Numerous studies have shown that regular aerobic exercise can effectively maintain or reduce percent body fat.[Cisar 91], [Kasch 90], [Shangold 88], [Shephard 91] In a 23 year longitudinal study discussed earlier in this paper, lead by Frank Kasch, the exercising group lost an average of 7.5 pounds, compared with their individual weights taken at the start of the study. The nonexercising group, in contrast, gained an average of 6.8 pounds. Percent body fat was not taken at the start of the study. But at age 70, the exercisers averaged 15.9% body fat and the nonexercisers averaged 25.7%.[Kasch 90]

Cross-sectional studies demonstrate the same effect. For example, female distance runners in their 30's average about 15% body fat, while a sedentary 30-year-old woman averages 27%. [Shangold 88]

How effective is exercise as a mechanism for losing body fat? Can a previously sedentary, obese individual lose excess body fat through increased activity? The American College of Sports Medicine has concluded that to lose a significant amount of body fat (greater than 5%) through exercise *alone,* an individual must exercise at least 20 minutes per day, 3 days a week, at a sufficient intensity and duration to bum 300kcal per session. Combining increased exercise with a modest reduction in calories can be much more effective in reducing percent body fat. However, reduction of calories accompanied by a *reduction* in physical activity results in a higher percent body fat - most of the weight loss being due to loss of muscle tissue and fluids.[Shangold 88]

Human Growth Hormone and its Relationship to Exercise and Aging

It has been suggested that the changes in body composition associated with aging are due at least in part to a decreased secretion of human growth hormone (GH). Secretion of human growth hormone generally decreases with age after the third decade of life.[Rudman 90], [Lancet ed. 91] Decrease in lean body mass, and increase in adipose tissue have been correlated with decreased levels of this hormone; these changes in body composition can be reversed by replacement doses of the hormone. [Rudman 90] Increased GH levels have also been associated with fewer by-products of tissue breakdown after intense exercise and decreased muscle soreness.[Elam 89] Administration of growth hormone by injection three times a week to GH-deficient men over sixty was shown over a six month period to significantly increase lean body mass, decrease percent body fat, increase skin thickness, and increase lumbar vertebral bone density.[Rudman 90] However, currently such administration, besides costing $14,000 a year [Smith 90] is not available to the general Extropian. In addition, there are potential health complications. Short-term side-effects of existing GH-replacement on GH-deficient adults include edema, arthralgia, hypertension, and carpal tunnel syndrome.[Lancet ed. 91] There is also an increased risk of cancer in acromegalic adults - adults with abnormally high secretion of GH. (Unfortunately, cancer cells seem to respond to growth factors as well as normal cells.) This has not been ruled out as a serious possible complication of long-term high-level GH therapy.

It has been suggested that the anti-aging effects of exercise could be due to increased secretion of GH. Studies have shown a pronounced increase in levels of human growth hormone immediately following exercise - both in men in their 20's, who went from a pre-exercise average level of 2.5 mcg/l to a post-exercise average level of 12.5 mcg/l[Quirion 88] and in men over 50 (pre-exercise .8 mcg/l; post exercise 9.3 mcg/l, and one hour later 3.7 mcg/l). [Metivier 88] The documented effects of exercise alone, as compared to growth hormone administration alone, show similar effects on body composition and muscle strength. Exercise studies have been able to demonstrate greater gains in less time – perhaps since there is not as much need to be concemed with potential complications in moderating "dosage" level! Also, so far studies on growth hormone have not shown any effect on functional flexibility, aerobic capacity, or muscle performance speed.

There is some indication that diet can affect physiological secretion of growth hormone in response to exercise. Men on very high carbohydrate, hyperinsulin-inducing diets (diets that cause high levels of insulin to be present int the blood) did not increase their GH levels after exercising as much as men on low-carbohydrate diets with associated low insulin blood levels, although GH levels still increased significantly.[Quirion 88] There are serious problems associatedwith being in a perpetually hypoinsulemic (low insulin) state, and with the excess of fat and protein that accompany a low carbohydrate diet. However, a much safer and effective altemative is available. Supplementing otherwise unaltered diets with the amino acids L-arginine and L-omithine shortly before exercise and before bed (nommal times for increased GH secretion) showed significant evidence of increased growth hormone secretion – and a more pronounced increase in muscle mass and strength and decreased tissue breakdown by-products than exercising, placebo- supplemented controls.[Elam 89] Combining exercise, a healthy diet, and amino acid supplementation would show all the benefits of an exercise program, with perhaps an increased effect on increased muscle mass and strength, decreased adipose tissue, and decreased muscle soreness as a result of enhanced GH release in response to exercise.

Other Age-Related Traits

Besides purely athletic deterioration, we see many serious health problems associated with aging. Of these age-related problems, several can be reduced and controlled through a regular program of exercise. Exercise alone has been shown to have protective effects against high blood pressure[Kasch 90], [Tanji 90], coronary heart disease[Shephard 91], [Omish 90], stroke[Blair 89], and non-insulin dependent diabetes[Staten 91].

As an individual ages, his/her proprioceptive senses become less accurate, causing losses in balance and coordination. Although the mechanisms for these losses are poorly understood at this time [Pickles 83], research has shown that decreases in balance and coordination are reduced by regular exercise. [Cramer 87], [Cisar 91]

People also associate general decreases in mental functions and reaction time with aging. Robert Dustman showed that sedentary adults ages 55 to 70 who were put on a "vigorous" exercise program showed significant improvement in response time, memory, and mental ¨"flexibility", when compared with a control group that was not put on the exercise program.[Dustman 91]

And what about lifespan? Regular exercise apparently can help us function at a "younger" level, but can it prolong life? Research indicates that it can. In one study, involving 16,000 men, those who walked 9 or more miles a week had a lower mortality rate than those who walked less than 3 miles. [Dustman 91] Another study involved comparing mortality rates of 13,000 men and women at different fitness levels. Fitness was categorized according to V02max - and subjects were grouped into five categories. After adjusting for age differences, smoking, and cholesterol, the subjects in the least-fit category still had death rates 3.4 times higher for men, and 4.6 times higher for women than subjects in the most-fit category.[Blair 90]

It might seem from some of the information presented in this paper that we have found a virtual panacea to keep our bodies from aging physically: muscle strength, speed, and mass can be maintained; bone mass need not decrease; percent body fat can be kept at a youthfully low level; aerobic capacity can for the most part be maintained - and losses in flexibility, coordination and balance can be minimized. However, exercise alone cannot stop some aspects of aging: As a person ages, his immune system deteriorates, making him more susceptible to disease and its associated periods of inactivity. Injuries take longer to heal. Cartilage breaks down and is not replaced. This, combined with the typical age-related decrease in secretion of hyaluronic acid, causes progressive osteoarthritis in the joints, which also makes adherence to an exercise schedule more difficult. Barrie Pickles states in *Biological Aspects of Aging:*

> The key to preventing the normal age-related changes from affecting functional capacity for ease of movement during the entire lifespan is a gradually increasing adherence to a program of physical fitness. As a person ages, physical fitness becomes more and more important. [Pickles 83]

As a person ages, physical fitness also becomes more and more difficult to maintain - requiring greater deterrnination and willpower. But the rewards - enjoying a functionally more youthful, productive, and potentially longer life, are certainly worth the extra effort! ❒

REFERENCES

Blair, S.: Exercise Can Enable You to Live Longer. *Randal Wellness Newsletter* 1990; 2(1):1

Blair, S.: Low Fitness May Mean Higher Risk of Stroke. *The Physician and Sportsmedicine* 1989;17(9):37-4

Cisar C.J., Kravitz, L.: Tuming Back Time: Exercise and Aging. *Idea Today* 1991; 9(1):28-36

Costill D.L.: Muscular Strength and Aging. *Sports Medicine Digest* 1990A; 12(8):4

Costill D.L.: Endurance Performance and Aging. *Sports Medicine Digest* 1990B; 12(10): 7

Crarnmer E., Mummert C.: Reach Out To Seniors. *Dance Exercise Today* 1987; 5(10):5-8

Dustman R.: Successful Aging. *University of California at Berkeley Wellness Letter* 1991; 7(11):1

Elam R.P., Hardin D.H., Sutton R.A.L., Hagen L.: Effects of Arginine and Ornithine on Strength, Lean Body Mass, and Urinary Hydroxyproline in Adult Males. *The Journal of Sports Medicine and Physical Fitness* 1989; 29(1): 52-56

Fiatarone M.A.: Strength Improvements at any Age. *Running & FitNews, American Running and Fitness Association* 1989; 7(11): 1

Fiatarone M.A., Marks E.C.,Ryan N.D.: High-intensity Strength Training in Nonagenarians: Effects on Skeletal Muscle. *Journal of the American Medical Association* 1990; 263(22):3029-3034

Kasch F.W., Boyer J.L., Van Camp, S.P., et al: The Effect of Physical Activity and Inactivity on Aerobic Power in Older Men (A Longitudinal Study). *The Physician and Sportsmedicine* 1990; 18(4):73-83

Kavanagh T., Shephard R.J.: Can Regular Sports Participation Slow the Aging Process? Data on Masters Athletes. *The Physician and Sportsmedicine 1990;18*(6):94-104

Lancet Editorial: Growth Hormone Therapy in Elderly People. *The Lancet* 1991; 3B7: 1131-1132

Metevier G., Gauthier R.: The Effects of Acute Physical Exercise on Blood Serum Cholesterol, Triglycerides, Human Growth Horrnone (HGH) and Free Thyroxine (T4) in Men over Fifty Years of Age. *The Journal of Sports Medicine and Physical Fitness* 1988; 28(1):7-10

Omish D., Brown, S.E., Scherwitz, L.W. et al: Can Lifestyle changes reverse coronary heart disease? The LIfestyle Heart Trial. *Lancet 1990;* 336:129-133

Pickles B.: Biological Aspects of Aging. *Physical Therapy of the Geriatric Patient* 1983; (edited by Osa Jackson):27-63

Pruitt L.: New Findings To Keep Bones Strong. *Running & FitNews, Arnerican Running and Fitness Association 1990;* 8(9):4-5

Quirion, A., Brisson, G., De Carufel, D., et al:Influence of Exercise and Dietary Modifications on Plasma Human Growth Hormone, Insulin and FFA. *The Journal of Sports Medicine and Physical Fitness* 1988; 28(4):352-353

Rudman, D., Feller, A.G., Nagraj, H., et al: Effects of Human Growth Horrnone in Men Over 60 Years Old. *The New Englend Journal of Medicine 1990;* 323(1):1-6

Shangold, M.M., Mirkin, G.: *Women and Exercise; Physiology and Sports Medicine.* 1988 Philadelphia: FA Davis Co.

Shepard, J.G., Pacelli, L.C.: Why Your Patients Shouldn't Take Aging Sitting Down. *The Physician and Sportsmedicine 1990;* 18(11):83-90

Shephard, R.J.: Exercise and Aging. *Mayo Clinic Health Letter,* special supplement, June 1991

The Transhuman Taste

Reviews of Extropian interest

The Anthropic Cosmological Principle

by John D. Barrow and Frank J. Tipler. Oxford: Oxford University Press, 1988. 706 pp.; £9.95 UK, $15.95 U.S (paper).

Reviewed by David Krieger

This book is almost too big to review. This is perhaps the single most relevant book to Extropian thought about the future. Barrow and Tipler present convincing and rigorous answers to questions like: What is the place of humanity in the universe? Are there other civilizations in our galaxy? What is the ultimate fate of the cosmos? Must life and intelligence eventually come to an end?

The anthropic cosmological principle of the title comes in three flavors, described in detail in the book. The Weak Anthropic Principle (WAP) is the simple statement that the values of physical constants are not random, but restricted by the requirement that it must be possible for intelligent observers to arise, for we are here to make the observation. This version has some limited predictive value. Knowing that carbon atoms are a necessary ingredient of our makeup, and being able to measure the rate of expansion of the universe, we can infer its size: the universe must be large enough to be old enough for the concentration of carbon to be great enough for carbon-based life to evolve – at least ten billion light years in extent.

The WAP is almost a tautology; the Strong Anthropic Principle goes much farther:

> Strong Anthropic Principle (SAP): The universe must have those properties which allow life to develop within it at some stage in its history.

This controversial statement provides a *reason* for all we observe around us: the universe is constrained to accomodate our existence. Barrow and Tipler consider three possible cases which could give rise to the SAP. They rightly dispense with the first - that there exists one possible Universe "designed" with the goal of generating and sustaining "observers" - as immune to scientific proof or disproof, noting that "Indeed it is a view either implicit or explicit in most theologies."

The other two interpretations of the SAP are founded on two competing interpretations of quantum mechanics: second, that observers are necessary to bring the universe into being; and third, that an ensemble of other universes is necessary for the existence of our universe. The third, springing from the "Many-Worlds" interpretation proposed by Wheeler, Everett, and Graham, is explored in depth by Barrow and Tipler, who provide a mathematical expression of this interpretation and show a number of potentially testable consequences.

Along the way, the authors explore the application of and the evidence for the WAP and SAP in the fields of biochemistry, physics, astrophysics, cosmology, and of course quantum mechanics. I guarantee that once you read and understand this book, you will know a great deal more about each of these sciences than you did before.

Perhaps the most relevant variant of the anthropic principle for Extropians is the Final Anthropic Principle: essentially, that life will never die out. More precisely, the amount of cogitation accomplished and the amount of information processed must increase without bound over time. This variation requires the universe and certain elementary particles to have specific properties, providing an experimental test of the FAP.

Barrow and Tipler demonstrate that the FAP can be satisfied in either a closed or an open universe. If the universe is open, that is, if its mass is not sufficient to cause it to collapse back to an inverse Big Bang, the Second Law of Thermodynamics decrees that it will suffer a Heat Death. The authors show that careful "energy husbandry" will permit intelligences to continue to operate, at slower and slower rates, for an infinite amount of time. Satisfying the FAP in this way will require that intelligent life survive the disappearance of all baryonic matter due to proton decay; the surviving leptons must be capable of making up some form of information-processing machinery. Leptonic minds, if they are possible at all, will be cool, slow, and vast, exchanging photons sparingly among individual beta particles light years apart, over unimaginably long epochs.

If, however, the universe is closed, the time available to us is finite. Barrow and Tipler demonstrate that the intelligences of that time still have the potential to do an infinite amount of thinking in that finite amount of time, thanks to the increasing density of shear energy (not to be confused with "Sheer Energy," a brand of ladies' pantyhose) in the universe during the era of collapse. From thermodynamics, they derive an integral expressing the maximum amount of information that can be processed and show that, over our remaining interval of time and even with a finite amount of available energy, this integral diverges - an infinite amount of thought can occur. Good news for devotees of Boundless Expansion!

The present state of scientific knowledge has neither confirmed nor denied the Strong Anthropic Principle. In their survey of the sciences mentioned above, the authors point out a wide variety of suggestive evidence. For example, if the resonance level of the C^{12} nucleus were not almost exactly its actual level of 7.656 MeV, no carbon could arise, as almost all carbon has, through stellar nucleosynthesis. If it were slightly lower, the nucleosynthetic process would stop with beryllium; if slightly higher, all carbon would swiftly fuse further with helium to form oxygen.

Furthermore, the criteria for a suitable solvent for organic reactions are many and, in many cases, mutually contradictory. Barrow and Tipler show that only water fits into the very narrow window left by these many requirements. In fact:

> Water is actually one of the strangest substances known to science. This may seem a rather odd thing to say about a substance as familiar but it is surely true. Its specific heat, its surface tension, and most of its other physical properties have values anomalously higher or lower than those of any other known material. The fact that its solid phase is less dense than its liquid phase (ice floats) is virtually a unique property.... Indeed it is difficult to conceive of a form of life which can spontaneously evolve from non-self-replicating collections of atoms to the complexity of living cells and yet is not based in an essential way on water.

Meanwhile Barrow and Tipler consider the Fermi question (Where is everybody?) and reach Fermi's perhaps unpopular conclusion, "If they existed, they would be here." - that we are alone in the galaxy, and perhaps in the universe. Barrow and Tipler present a convincing argument that a civilization capable of launching a single self-replicating Von Neumann probe would have such an artifact (and the industrial base to build the next generation of probes) in every solar system of the galaxy within 300 million years. (Barrow and Tipler have reckoned without the development of nanotechnology, which would greatly strengthen, not weaken, their argument - by my calculations, nanotech reduces this time-frame to three million years.) Also, the cost relative to wages of raw materials is always dropping in a society that is advancing technologically - for example, a project within the means of the United States government today will, in a few generations, be within the means of individuals. (Nanotech shortens this period considerably as well.) Thus, it would take only *one nut, anywhere,* to blanket the galaxy with self-replicating space probes. Until we discover evidence of such developments in our own solar system, the authors argue, we must conclude that We Are, in fact, Alone.

Why should you run out and buy this book right now? First and foremost, it will definitely expand your understanding of the nuts-and-bolts mechanisms that make life as we know it, and human consciousness, possible. The prospect of immortality (to coin a phrase) has now been given a solid theoretical basis. The dilemma of the Heat Death or the Big Crunch has always been dangled mockingly in the faces of immortalists. Barrow and Tipler show that while such events might lead to the end to life and consciousness, they need not. With 10^{32} (or 10^{98}, depending on the flavor of proton decay you prefer) years to work on the problem, if we can't design a beta-particles-only platform for consciousness, then we don't deserve to live for the rest of the Long Haul.

The most important aspect of the book, to me, is the light it sheds on what we may call the God Axiom. For our entire history, discussion on this topic has been dominated by the ravings of madmen (the Revelation of St. John the Divine comes to mind) and the prevarications of those in a position to profit (medieval Popes, "televangelists," and similar snake-oil salesmen). Science (as opposed to individual scientists) has remained aloof from the fray, because the God Axiom was immune to disproof, and therefore an unscientific notion.

The SAP, however, is the first statement relevant to the God Theorem that has testable consequences. If the SAP holds up under experiment - say, to the same extent as the Law of Conservation of Energy - then it ties the hands of the Deity: either God requires an audience, or God is required to keep making universes until we (or some intelligent species) show up. So much for the omnipotence of God. If the definition of God includes omnipotence, then so much for God. (Arthur C. Clarke has proposed a similar application of Gödel's Theorem to God's omniscience.)

If the SAP is disproven (i.e., if it is physically possible to have a uni(multi)verse with no consciousness in it), then the God Axiom is dealt a different blow. God may exist, but cannot be the kind of personal God who wants burnt offerings and goes peeking in people's bedrooms. If the SAP is disproven, then humanity was not part of some Divine Plan; we arose by accident, in a universe where we could just as easily have never existed. Disproof will leave only the various sects of Deists standing, while proof will knock the crutch out from all theologies. If the SAP is true, then there is no omnipotent God - *yet*:

Barrow and Tipler point out that the infinity of consciousness is not optional. In order to continue to process information, whatever intelligence exists in the remote future must continue to expand and to take control of greater and greater quantities of matter. The "intelligence community" of that time will expand until, in the authors' words (and with their emphasis),

> At the instant the Omega Point is reached, life will have gained control of *all* matter and forces not only in a single universe, but in all universes whose existence is logically possible; life will have spread into *all* spatial regions in all universes which could logically exist, and will have stored an infinite amount of information, including all bits of knowledge which it is logically possible to know. And this is the end.

The academic and scientific credentials of both authors are impeccable. Barrow is an astronomer; Tipler, a physicist and mathematician. Science fiction fans may be familiar with Tipler's name from his time-travel paper "On Rotating Cylinders and the Possibility of Global Causality Violation" (Larry Niven later borrowed the title for a short story about a time machine such as Tipler proposed).

I found that this book simultaneously expanded my mind and disabused me of long-held reassuring notions. The book is a gold mine for futurists and science fiction authors, containing enough ideas to fuel the average writer's career for decades. Anyone seriously interested in our future, particularly the Long View, must read this book.

The Blind Watchmaker

by Richard Dawkins

New York: W. W. Norton & Company, 332 pages.

Reviewed by Simon! D. Levy

Richard Dawkins is my kind of writer. In an age when academics like to express themselves in terms such as "not inconsistent with," Professor Dawkins starts his book about evolution with the statement that "our existence... is a mystery no longer because it is solved. Darwin and Wallace solved it...." After those refreshingly audacious beginnings, Dawkins goes on to present the clearest, most convincing, and most entertaining explication of a scientific theory that I have read.

The fundamental question that preoccupies Dawkins and shapes his argument is the question of complexity. Choosing the example of the human eye, he echoes the incredulity of the creationists, against whom he convincingly argues, in asking how such a fantastically intricate structure could ever have evolved spontaneously. In a nutshell, the answer is that the difficulty we have in understanding evolution is the result of an incorrect metaphor. According to this metaphor, which Dawkins calls the "Boeing 747 macromutation," evolution proceeds like a hurricane blowing through a junkyard: What are the chances that, among the inconceivably huge number of ways the junk could get blown around, the final arrangement of the junk will be a fully assembled and operating jetliner? Or, to use the argument of eighteenth-century theologian William Paley, if we stumble across a rock in a field, we are perfectly satisfied to believe that the rock has always been there, without wondering about where it came from or who made it. If, however, we were to stumble across a watch in a field, the same answer would hardly suffice. Instead, we would be forced to attribute the existence of the watch to a watchmaker. In much the same way, argued Paley, the existence of fantastically complex structures in Nature must force us to admit the existence of a Divine Watchmaker, who is responsible for the creation of us and of all living things.

The answer, of course, is that such a macromutation from functionless disorder to functional order is effectively impossible. Watches don't spring full-blown from rocky fields, and 747's don't get put together by a hurricane from scraps in a junkyard. Clearly, then, evolution could not have worked this way. A complete human eye couldn't spontaneously evolve from a bare patch of skin. Instead, as Dawkins shows, evolution is understandable in terms of its gradualness. A random mutation, caused for example by cosmic rays, may lead to a small change in the genes that an organism passes on to its offspring. If this change is beneficial to the offspring — if, for example, the change produces a body better able to avoid predators or detect prey — then the offspring will have a greater chance of surviving long enough to pass their mutant genes onto their own offspring. These two forces, mutation and natural selection, have acted (and continue to act) together as a "blind watchmaker," producing structures of fantastic complexity.

In a chapter called "Accumulating Small Change," Dawkins shows how these little mutations could build up to something really big: Imagine a piece of graph paper on which all possible genetic configurations are represented. Similar configurations will be closer together in the graph; for example, the genetic configuration of a monkey is very close to that of a human being, so these two primates will be close together on the graph. Primates have a very different genetic makeup from viruses, so primates and viruses will be far apart on the graph. Large changes in genetic makeup — for example, the change from a bare patch of skin to an eye — represent long-distance movements across the graph.*

A major point of *The Blind Watchmaker* is that such movements necessarily require movement through every intervening stage. You can't make large jumps from one stage to another, any more than a hurricane can assemble an airliner out of a pile of scrap. So the first movement on the path from bare skin to eye probably involved some sort of small mutation, which made a patch of skin more light-sensitive than the skin on the rest of the animal. This increased sensitivity gave the mutant animal a competitive edge, so that it was able to pass the sensitive-patch genes onto its descendants, one of which mutated further in the advantageous "eye" direction, and so forth.

Now, this *gradualist* view of evolution is not without its detractors. First are those like C.E. Raven who argue, somewhat persuasively, that individual small steps on the evolutionary graph are of dubious value to the organism. It is clear that an organism with eyes has a big advantage, *ceteris paribus,* over its blind counterparts. But how could some very small change — for example, a mutation that makes the a proto-eye just a bit more sensitive to light — provide a major benefit for survival?

Dawkins' answer to this question is twofold. First, he points out that even the tiniest mutation may be advantageous: "A simple, rudimentary, half-cocked eye... is better than none." This part of the answer is made more convincing by the second part, which is the amount of time involved. If we think of the trip from bare skin to eye in terms of a few generations, or a few dozen, it is extremely unlikely that the proper mutations would arise quickly enough to make the journey possible. As Dawkins points out, however, the time spans involved in evolution are on the order of several hundred million years. If we acknowledge the enormous number of generations that must have elapsed between successive small mutations, the development of complex structures such as the eye becomes far more comprehensible.

The second attack on gradualism comes from the so-called punctuationist school of evolutionary theory, whose most famous exponent is Harvard paleontologist Steven Jay Gould. Based on their observation of fossils dated by radioactive testing, the punctuationists have argued that evolution does not seem to take place gradually. For example, the size of the brain of *Homo sapiens* represents a tripling of the size of the brain of this species' ancestors, in a period of three million years. If this size increase were gradual, i.e., spread out evenly over the years, the difference in brain sizes between successive generations would be far too small to provide any advantage. Furthermore, the fossil record fails to reveal this kind of gradual change.

In his counter-attack on the punctuationists, Dawkins points out that this view of gradualism is essentially a straw-man. No sensible biologist, least of all Dawkins, would take such a position. In fact, Dawkins seems particularly angry at the media attention that Gould and his ilk have been given because of the erroneous belief that they were challenging a widely held view. According to Dawkins, the punctuationists, like all reasonable evolutionists, do embrace some form of gradualism; the alternative is to believe in the Boeing 747 macromutation. Instead, the significant way in which punctuationists differ from people like Dawkins is that the punctuationists believe in long periods with no evolution, interrupted by brief periods of accelerated evolution. Dawkins, on the other hand, seems either to stick to the steady-rate-of-evolution view, or to consider punctuationism a "minor ripple" in Darwinian theory. Usually, I have no patience for this sort of "we're all saying the same thing" approach, because it is only through controversy and disagreement that science proceeds. Nevertheless, Dawkins makes his point convincingly. Not having read Gould, I am in not in a position to evaluate Dawkins' criticisms further, but it would not surprise me in the least if what he has said is true, given the way that television and the popular press report on science.

All in all, I found *The Blind Watchmaker* a thoroughly enjoyable and worthwhile book. It gave me a sense of why Charles Darwin, clearly a hero in Dawkins' eyes, is so revered by thinking people everywhere, including a past Alcor president who adopted the nineteenth-century genius' surname as his own. Dawkins himself pulls no punches in demolishing the creationists, showing even their most plausible arguments to be full of holes; his book is therefore indispensable for Extropians trying to explain the origin of life (and hence the material nature of consciousness) to bewildered loved ones and students.

To be fair, there were one or two glaring errors that were surprising, given the generally high quality of scholarship and writing of this Oxford professor. For example, *candelabra* is plural, not singular; this kind of mistake makes one wonder about the worth of editors nowadays. But such peccadilloes are more than offset by an engaging style that is rare in science writing. Dawkins has an ear for language; he tells us how animals "make their living" and that "however many ways there may be of being alive, it is certain that there are vastly more ways of being dead." This charming way of writing, combined with Dawkins' obvious passion for the issues in his field, make for truly pleasurable reading. In short, I recommend *The Blind Watchmaker* without reservation.

* Dawkins even wrote a program to illustrate this "evolutionary space" idea with little stick-figure creatures that he calls "biomorphs." We had a copy of the program (Apple Macintosh version) lying around at the lab where I work, and I had a good time running a few of Dawkins' attractive simulations, though I couldn't get the thing to work on System 6.05 or higher. In any event, it's a bargain at $10.95, and you can have fun while supporting a worthy cause.

Economist Against The Apocalyptics:

The wisdom of Julian Simon: Three books. Reviewed by Max More

The Ultimate Resource (Martin Robertson, Oxford, 1981). 415 pages. ISBN: 0-85520-563-6

The Resourceful Earth: A Response to Global 2000 (Basil Blackwell, Oxford 1984). Edited by Julian Simon and Herman Kahn. 585 pages. ISBN: 0-631-13467-0

Population Matters: People, Resources, Environment and Immigration (Transaction Publishers, 1990). 577 pages. ISSN: 0-88738-300-9

Dualistic thinking is always tempting. I'm tempted to portray Julian Simon as the heroic rebel fighting Paul Ehrlich's Evil Empire of environmental crisis-mongers, people-haters, and coercion enthusiasts. Forcing viewpoints into diametrically opposing positions usually distorts the situation and obscures information. Yet, from an Extropian perspective, so bad are the facts and values of Ehrlich and those like him, and so perceptive and agreeable the writings of Simon, that dualistic treatment might be close to the truth.

Julian Simon is a professor at the University of Maryland, College of Business and Management, and a researcher for the Hudson Institute. He is one of the valiant few standing against the tide of irrationalist, apocalyptic environmentalism and is deeply unpopular with the dark forces of anti-growth, pro-statist environmentalists. Recently, in addition to besting his opponents intellectually, he backed his principles with money. Simon challenged Ehrlich to make a bet on the real cost of raw materials ten years in the future. According to Ehrlich's view, these prices should rise greatly, due to their purportedly increasing scarcity. On this view, resources are a fixed stock, and they will be gradually consumed at an accelerating rate as population grows. Simon's economistic viewpoint holds that resources are effectively unlimited, and that substitution and technological innovation (boosted by population growth) will hold down the prices of resources. Recently, the ten years expired. Simon won the bet and Ehrlich paid up. Did this prevent Ehrlich from continuing to make more incredible and unfounded claims? Of course not. But neither would he renew the bet.

Format of the Books

Each of the three books reviewed here differ in format. The earliest, *The Ultimate Resource* (TUR), is a book-length treatment of resource trends and population effects, written for the intelligent layperson, but solidly based in economic theory. Some of the same issues treated in TUR and the other books are analyzed in full professional detail in Simon's *The Economics of Population Growth* (1977). TUR explains the theory of scarce resources, looks at technological and economic forecasts, questions the finiteness of resources, and delves into issues of food supply and famine, availability of land, energy supplies, pollution, and the negative and positive effects of population growth.

The Resourceful Earth (TRE) is not written by Simon (apart from one of the essays), but his editorial light shines through (nicely complemented by co-editor Herman Kahn). TRE is an unparalleled source of information on environmental and resource issues by a collection of experts. D. Gale Johnson examines world food trends and argues that "the prospects for the long run are in the direction of gradual declines in the real prices of the primary sources of calories for poor people." The essay on global forests, followed by the discussion of the data (or lack of it) regarding species loss by Simon and Wildavsky, shows just how far from the evidence popular beliefs about environmental issues stray. (More on this example below.) Other authors sharply probe the available information on supplies of agricultural land, soil erosion, water availability, global climatic trends, trends in non-mineral resources as well as in oil and petroleum, nuclear power, solar energy, coal, environmental quality, air and water quality, nutrition and health trends, and cancer rates. The overall message of the collection is certainly not complacent, but does show that most trends are in the right direction, and solutions to current problems lie in market mechanisms, not centralized coercion.

Population Matters, the most recent book, is a collection of essays written by Julian Simon. These 58 essays are a concentrated source of invaluable information for enlightening yourself and for gathering intellectual ammunition. The collection consists of eight parts: General Overview; Natural Resources; Population Growth; Population Policies, Programs, and Beliefs; Immigration; Failed Prophecies and the Doomsaying Establishment; Progress, World Views, and Modes of Thought; Publication, Funding, and the Population Establishment.

Far more than the other books, PM reveals in horrifying detail the power and unity of the doomsday establishment view. Simon relates his difficulties in getting published, and shows how the various anti-growth environmentalist groups are interlocked. He also explains how funding imperatives drive researchers into exaggeration and distortion. Funds are more accessible to those claiming their work to be vital to the future of all life! This effect demonstrates the ineffectiveness of feedback mechanisms linking researchers' claims and their costs and benefits. (In *Extropy* #8, Robin Hanson proposed a system of Idea Futures to improve incentives to make reasonable claims.)

Scarcity

Having given an idea of the format of each book, rather than go through them chapter by chapter, I will give an idea of Simon's approach throughout the books. First, in any discussion of resource issues, the concept of scarcity must be clarified and its consequences drawn out. Increasing scarcity of a resource will (in a free market) be reflected in a persistently rising price. A related important measure of scarcity is the relationship between price and income. If the price of aluminum remains constant while our income rises, then we will feel that aluminum is becoming less scarce.

This method of measuring scarcity is the economic measure, and differs importantly from technological tests of scarcity. The economist's approach relies on price mechanisms and on long-run cost trends. The technological method begins by estimating the currently known quantity of the resource on or in our current planet. Second, it calculates the future use rate of the resource on the basis of the current rate and, finally, calculates the numbers of years, given the prior calculations, before the resource is exhausted. Such technological estimates generally suggest that resources will become increasingly scarce and will eventually run out entirely. If this were so, we should expect resource prices to climb. Yet, as Simon shows:

> Considerable data showing trends in raw-material prices are available, as seen in the Appendix to this book. The overwhelming impression given by these figures is that costs for extractive materials have fallen over the course of recorded price history. The economist's first-approximation forecast is that these trends toward less scarcity should continue into the foreseeable future unless there is some reason to believe that conditions have changed, that is, unless there is something wrong with the data as a basis for extrapolation. [TUR: 21]

Since technological progress is accelerating, not slowing, projecting future price trends on the basis of the past is likely to understate the decreasing scarcity of resources. We should also note that current prices contain information about *future* scarcity: If speculators have reason to believe a resource will become more costly to acquire in the future, they will buy it now to hoard and resell in the future. This action will raise the current price of the resource.

Two Types of Forecast

Technological forecasts, in contrast with economic forecasts, suffer from several shortcomings. They rely on the assumption that "a certain quantity of a given mineral 'exists' in the earth, and that one can, at least in principle, answer the question: How much (say) copper is there?" Intuitively plausible as it may be to the economically unsophisticated, this assumption is replete with difficulties. Attempting to define the available quantity of a resource, such as copper, is hopeless in principle, let alone in practice. The grades of a resource differ dramatically, varying in difficulty of extraction and amounts at low concentrations (such as metals in sea water) vastly exceed the amounts normally counted as 'proven reserves'. New sources may arise outside the system considered by the forecaster, such as resources from the seas, from other planets, or by processes such as nuclear breeding of fuel or nucleosynthesis of elements. Trying to define total availability of a resource is a vain endeavor.

> Consider the definition of the potential supply of oil that is implicitly or explicitly used by many: the amount that would be recorded if someone conducted an exhaustive survey of all the earth's contents. This supply is apparently fixed. But such a definition is thoroughly non-operational, because such a survey is impossible even in principle. The operational supply is that which is known today, or that which we may forecast as being known in the future, or that which we estimate will be sought and found under varying conditions of demand. These latter two quantities are decidedly not fixed but rather variable, and they are the ones relevant for policy decisions. [TUR: 31-

Table 1: Number of Years of Consumption Potential for Various Elements

	Known reserves/annual consumption	U.S. Geological Survey's estimates of "ultimate recoverable resources" (=0.1% of materials in top kilometer of earth's crust/Annual Consumption	Amount estimated in earth's crust/ annual consumption
Copper	45	340	242,000,000
Iron	117	2,657	1,815,000,000
Phosphorus	481	1,601	870,000,000
Molybdenum	65	630	422,000,000
Lead	10	162	85,000,000
Zinc	21	618	409,000,000
Sulfur	30	6,897	na
Uranium	50	8,455	1,855,000,000
Aluminum	23	68,066	38,500,000,000
Gold	9	102	57,000,000

SOURCE: William D. Nordhaus, Resources as a constrain upon growth. *American Economic Review* 64 (1974), p.23.

Simon identifies five major difficulties in technological forecasting. Table 1 (previous page) illustrates the growth in known reserves between 1950 and 1970, reinforcing the economist's view that use of the known-reserve idea is misleading and worthless.

Second, the supply of minerals tends to be highly price elastic. That is, a small increase in price greatly increases the potential resources that can be profitably extracted. Many technological forecasts are based on current prices and current technology and so inevitably show rapid exhaustion of resources. Third, technological forecasts that attempt to go beyond 'known reserves' must make highly uncertain guesses about future discoveries of new reserves and about technological innovations. The economic approach only needs to assume that the long-run cost trend will continue. Fourth, the mineral resources of the Earth have not been thoroughly inventoried because it has never been worth anyone's effort to do this. Finally, technological forecasts depend on how imaginative a forecaster is in thinking up future extraction methods.

Pollution

The popular view is that pollution is getting worse in just about all respects. The truth is very different. With some exceptions pollution is becoming less of a problem. Some previous pollution blamed on human activity is really the result of natural forces, such as the falling oxygen content of the Baltic Sea.

Despite propaganda about the dead Great Lakes, since the 1970s their quality has been improving. The fish catch in Lake Erie, one of the two most polluted of the Lakes, increased from its low in the '60s, with 10 million pounds of fish caught there in 1977. "Lake Superior's purity seems to have been increasing rather steadily or holding constant, at least in terms of the dissolved solids for which we have data." [Baumol & Oates, TRE: 444] Water quality in other areas is heading in the right direction. The oxygen content of the Thames river in England has risen since the 1960s, and the Hudson river is cleaner now than it has been in decades. "By almost every measure available – amount of money spent, number of sewage treatment plants constructed, number of crabs returning, number and size of fish, visibility of sewage, number of people swimming – the 155-mile long main stem of the Hudson River between New York City and Troy is improving." [TRE: 449]

Air is improving: Levels of carbon monoxide, sulfur dioxide, and suspended particulates have all been falling, and unleaded gas has contributed to reduced lead levels. For example, in New York by 1975, soot levels in Brooklyn fell to one-sixth of their 1945 level, and in Manhattan the level declined by two-thirds. Chicago also showed dramatic improvements, and the air in other cities became cleaner, if not by as much.

Some choice descriptions of pollution in the past, making our own environment appear pristine and pure by comparison can found in both TUR and TRE. In the streets of London in 1890, red-jacketed boys would run around trying to collect the horse manure than threatened to drown the city. Passing wagons would spray passersby with the filth and it accumulated along the sides of the road where it would be thrown. In addition, the noise of the incessant horse-drawn traffic was deafening. Many examples of terrible pollution in the past are available; the point is that we cannot realistically estimate our pollution situation without comparing it to the past. Another point well made by Simon and undoubtedly obvious to most readers of this journal is that zero pollution is not a viable option. Pollution is a by-product of desirable production, reduction of pollution is costly in terms of other desired goods and services forgone (the 'opportunity cost').

Population and Technology

Population projections have been made for decades and have almost invariably turned out to be wide of the mark. For example: "As of 1969, the U.S. *Department of State Bulletin* forecast 7.5 billion people for the year 2000, echoing the original UN source. By 1974, the figure quoted in the media was 7.2 billion. By 1976, Raphael Salas, the executive director of the UN Fund for Population Activities (UNFPA) was forecasting "nearly 7 billion." Soon Salas was all the way down to "at least 5.8 billion." And as early as 1977, Lester Brown and the Worldwatch Institute (which the UN is supporting) dropped it down again, forecasting 5.4 billion people for the year 2000." [TUR: 169-70] These and other examples should make us wary of making drastic policy decisions on the basis of forecasts.

Simon's position has sometimes incorrectly been characterized by critics as holding that population growth is always and in every way a good thing, but his position is more complicated. The reader should consult Simon's books to get an accurate idea of his position. For present purposes my understanding of Simon's position is that he claims that whether population growth has positive effects depends in part on the values of the people involved, and that while population growth can be expected to have powerfully positive long-run effects, it does bring some short-run costs.

Discussions in *The Ultimate Resource* and *Population Matters* (and in great detail in his technical work, *The Economics of Population Growth*) demonstrate that population growth tends to accelerate technological progress. One reason for this is an increased rate of innovation. The more people there are, the more minds are working on finding to solutions, whether grand or mundane: "...the data show clearly that the bigger the population of a country, the greater the number of scientists and the larger the amount of scientific knowledge produced; more specifically, scientific output is proportional to population size, in countries at the same level of income." [TUR: 202]

In addition, faster population growth speeds up the growth rate of industries, and faster-growing industries have faster rates of growth of productivity and technological practice. Simon even provides graphs refuting the contention that population growth and scientific discoveries were not correlated in ancient Greece and Rome. Rising population not only increases the rate of innovation, it spurs the adoption of existing technology. New methods, in agriculture for example, may be initially more laborious (though more efficient and productive later) and so will only be adopted under pressure of population growth.

Economies of scale provide another reason for the acceleration of technological progress in a growing population. A bigger population implies a bigger market; Greater division of labor occurs and therefore more specialized skills develop; A wider variety of services is offered; A larger population means larger production runs, and this means more learning by doing; Better infrastructure becomes affordable.

Simon demonstrates other surprising and perhaps counter-intuitive results of population growth. I will only briefly mention the effects of population growth on the availability of land for recreation. In the U.S.A., "Land dedicated to wildlife areas, national and state parks and forests, and recreational uses has risen from 8 million acres in 1920 to 61 million acres in 1974. And the President's Commission in 1972 foresaw a further rise of about 37 percent from 1964 to 1980 in 'pure recreation land outside towns.'" [TUR: 235] Long-term population growth has increased income and transportation systems, making recreational land more accessible.

Species Loss

"20% of all species will be extinct by the year 2000." "One hundred species are becoming extinct every day." You have probably seen these figures, or similar, bandied about in the press and stated with conviction by numerous commentators. (Even the Nature Conservancy, of which I was a member, and which protected land mostly by the proper method of buying it, used these figures in a recent fund-raising effort.) Julian Simon and Aaron Wildavsky's essay, "On Species Loss, the Absence of Data, and Risks to Humanity," reinforced by Roger Sedjo and Marion Clawson's "Global Forests" (both in *The Resourceful Earth*) show these figures for the irresponsible fabrications they are.

The falsely alarming statistics come from the *Global 2000* report (to which *The Resourceful Earth* is an effective rejoinder). Lovejoy's figures in Global 2000 are based on Myer's *The Sinking Ark* (1979) "which was written under the auspices of a committee of which Lovejoy was one of three members, and whose prologue is a motto of the World Wildlife Fund, on whose staff Lovejoy serves." [TRE: 173] Myers estimates an extinction rate of one species every 4 years between 1600 and 1900. Myers then estimates an extinction rate of one species per year from 1900 to the present (then 1980), but gives no source for this estimate). To continue, in Simon and Wildavsky's words:

> Some scientists have (in Myers' words) "hazarded a guess" that the extinction rate "could now have reached" 100 species per year. That is, the estimate is simply conjecture and is not even a point estimate but rather an upper bound... Even this guessed upper limit is then increased and used by Myer's, and then by Lovejoy, as the basis for the "projections" quoted above. In *Global 2000* the language has become that economic developments "are likely to lead" to the extinction of between 14 and 20 percent of all species before the year 2000 (*Global 2000*, vol.II, p.328), which calculates to about 40,000 species lost per year. Observe that an upper limit for the present that is pure guesswork, and that is 100 times the observed rate in the recent past, has become the basis of a forecast for the future which is 40,000 times greater than at present, and which has been published in newspapers to be read by tens or hundreds of millions of people and understood as a scientific statement.

The authors go on to further undermine the credibility of the Global 2000 statements by analyzing that report's use of deforestation data. The essay concludes with a discussion placing the risks from species loss into proportion. As the authors note, they do not intend to suggest that we should ignore possible dangers to species.

> Individual species, and perhaps all species taken together, constitute a valuable endowment, and we should guard their survival just as we guard our other physical and social assets. But we should strive for a sensible view of this set of assets in order to make the best possible judgements about how much time and money, and human welfare, should be spent in guarding them in a world in which this valuable activity must compete with other valuable activities, including the guarding of other valuable assets and of human life.

Immigration

The environmentalists not only try to restrain technology and growth, they also oppose immigration. The free flow of persons is increasingly important to the dynamism of economies. If the environmentalists, allied with traditional anti-immigration groups, have their way, yet another chain will have been wrapped around economic growth, not to mention around the lives of those trying to better their condition.

Simon provides helpful data to show the overall beneficial impact of immigrants on the economy. After putting the likely numbers of illegal immigrants into perspective, Simon refutes the contention that immigrants abuse welfare and government services. Small proportions of illegal immigrants use government services, being afraid of getting caught if they apply for welfare. Legal immigrant families make use of welfare about as much as do citizens, and much *less* if services for the elderly are included.

One reason why immigrants are a net boon to the economy is their demographic composition: "On average, it is the young, strong, and single who make the move. Of the illegal Mexicans, more than 80 percent are male, half are single (most of the married men leave their wives and children in Mexico), and most are youthful – less than 20 percent of the workers are over 35, and they average perhaps 27. Among the Vietnam refugees, only 12 percent are 45 or over, compared to 32 percent for the 45-and-over group in the U.S. population as a whole." [TUR: 273]

I would like to have seen a statement from Simon to the effect that immigration laws are unjust and should be abolished entirely. However, despite Simon's strongly free market approach, he writes as an economist and so we cannot expect such directly political pronouncements. He does state one of two ideological positions as "The freedom to move across national boundaries is a human right that ought to be recognized," [PM: 282] but goes on to cite economic data in favor of this view rather than explicitly agreeing with it as a political view.

You know that Simon must be doing a good job of attacking apocalyptic environmentalism when you see how upset they become at the mention of his name. I found it hilarious and witty for Simon to include on the back cover of *Population Matters* some comments from his critics, in addition to the usual favorable comments (including one from F.A. Hayek): "The man's a terrorist" – Mark Plotkin, World Wildlife Fund, Cox Newspapers. "Schizophrenic nonsense and baloney...sabotaging the human race." – Dr. L.E. Marshall, Estherville, Iowa. How nice to see the balanced thought and intellectual care of these critics!

The works reviewed here should be on the shelves and reading lists of all who think of themselves as Extropians. They make fascinating and enlightening reading, and will serve as invaluable sources of intellectual ammunition for years to come. Future works by Julian Simon will be reviewed in *Extropy* or ExI's new newsletter *Exponent*.

Select Bibliography

Issues in the Economics of Advertising (1971)
The Effects of Income on Fertility (1974)
The Economics of Population Growth (1977)
The Ultimate Resource (1981)
Theory of Population and Economic Growth (1986)
Effort, Opportunity, and Growth (1987)
The Economic Consequences of Immigration (1989)
Essays on the Effects of Population Growth in LDC's (forthcoming)

Edited Books
Research in Population Economics: Vol. I

Bionomics: The Inevitability of Capitalism

By Michael Rothschild, New York: Henry Holt and Company, Inc., 1990. 423 pp; $24.95 (hardback). ISBN: 0-8050-1068-8.

Reviewed by Harry Schapiro

This review is in both a short and a long form. If you want to read the short form, stop at the end of this paragraph. *Bionomics*, written by Michael Rothschild is philosophically compatible with many Extropian views. In simple terms, bionomic theory combines Austrian economics with classical and neo-classical Darwinian biology to explain the current state of world economics (circa 1990). "Bionomics argues that what we call capitalism (or free-market economics) is not an ism at all but a naturally occurring phenomenon."[1] Rothschild was a business consultant and the book is crafted to appeal to an educated but non-extropian readership. For its clever synthesis of biological and economic theory and its promotion of spontaneous order, I highly recommend it. Extropians should consider promoting this book as it promotes a weakened set of extropian meme's.

Extropian principles of Boundless Expansion, Dynamic Optimism, Self-Transformation, and Spontaneous Order fit well into bionomic theory. Rothschild brings out these and other views by discussing the following eight topics in the book:

* Evolution and Innovation
* Organism and Organization
* Energy and Value
* Learning and Progress
* Struggle and Competition
* Feedback loops and Free Markets
* Parasitism and Exploitation
* Mutualism and Cooperation

In as simple terms as possible, bionomics holds that replicating information makes economies dynamic, that organizations making better marginal use of this information will thrive, and that because new information is always being created, Malthusian/doomsday economics will never produce predictive or useful theories.

A Malthusian-orientated economist might hold that a healthy business will reach diminishing returns at high levels of productivity. Bionomic theory holds that a healthy business will learn more about production, will innovate, and will gain greater economies of scale and scope. This is similar to the Extropian principles of Boundless Expansion and Dynamic Optimism. The learning process healthy businesses go though can be likened to the Extropian principle of Self-Transformation. In bionomics, it is not life forms that are self-evolving but businesses. As businesses grow and learn and transform themselves, bionomics predicts that those businesses that focus on a niche in the economic system, thus avoiding direct competition, will thrive. This is not unlike species in biology evolving so they can live near each other without competition for the same resources.

My own background is in general studies, Liberal Arts, Management, and Writing. Since I am not trained in economics I will not analyze each page of *Bionomics*. However, I would like to show how Rothschild's bionomics diverges from Extropian philosophy. I will example, briefly, his chapters on: Theories of Change, Brains and Tools, Surplus and Genes, and Profits and Technology.

Chapter 2 — Theories of Change

Rothschild relates how Hegel and the dialectic serves as a basis for Marx's economic theory. This is an orthodox view; Rothschild compares the dialectic to the theories of biological catastrophe" like Cuvier's.

Rothschild develops a position that Marx was the first to propose to the effect that economic systems change throughout history in the same manner that Cuvier (Catastrophe) and Lamarck (inherited variation) first proposed that biological systems also change/evolve.

Rothschild holds that Smith, Ricardo and Malthus viewed economics with a "Mechanical/Newtonian" world view. Although not stated by Rothschild, he thus views Marx as having a more organic world view. While Newtonian-minded Ricardo may have been a mechanist, and while this seems to be true of Malthus, I don't see such a view as being held by Smith. True, Smith did hold that economics reached its maturity in the age of Commerce and he was not fully laissez-faire[2] but rather partly interventionist[3] in view (perhaps revealing a component of "economics as physics" within in his otherwise organic view[4]) Smith did write and demonstrate that economics had changed as a result of environmental influences. Smith states that these changes took place in a self-ordering manner. (An important test for "organic" theories.) The four "economies" that human life moved through were "Hunters, Shepherds, Agricultural and Commerce."[5] In this regard, Smith's views can be seen as an economic version of punctuated equilibrium[6] which Rothschild describes in chapter 5. Why Smith did not believe that economies would further change is unknown. What is clear is that he believed that economies had changed. In this regard, Smith is likened to Cuvier or Lamarck in the limited view or scope of his theory.

Chapter 6 — Brains and Tools

Thoughout his book, Rothschild makes a insightful point in talking about the power of DNA. He views it as the biological "library", relating it by analogy to the written language used by humans. He also presents the orthodox Darwinian view that DNA evolved spontaneously. At the end of chapter 6 he goes further in stating:

> Our lineage is unique because our anatomy allowed our ancestors to supplement their genetic evolution with technological evolution. Through creativity and innovation - behaviors made possible by the vastly enlarged brains of a strain of juvenilized apes - our ancestors were able to satisfy their most fundamental economic needs. The brains of our forebears became a living bridge connecting the ancient process of genetic evolution with the brand-new process of technological evolution. Up to this point in earth's history, the only form of living information was nature's - the mechanism of DNA. But, once the Habilis brain, itself a product of DNA, began to innovate, it launched an entirely new realm of living information." (Emphasis mine.)

It is at this point, if not sooner, that Rothschild could have introduced the concept of the *meme*. He did not feel the need for its inclusion. This raises several questions:

1) Does meme theory (memetics) have any effect upon the theory of bionomics?
2) Does memetics affect bionomics? In what way?

3) Was Rothschild aware of memetics? In the article written by Rothschild in *Liberty*, he talks about his lack of knowledge of the Austrian school of economics and then delineates how bionomics interacts with Austrian economics. It is very likely that Rothschild likewise never heard of meme theory. However, the 8th footnote in his book is: "See Dawkins, Richard, *The Selfish Gene* (Oxford: Oxford University Press, 1976)." It is in this book, Chapter 11: Memes the New Replicators, that the concept of memetics was first proposed. Since Rothschild is so focused upon encoded information, and since memes are nothing if not examples of physically encoded, stored, and transmitted information, surely bionomics and meme theory must integrate.

In personal conversation Rothschild indicated that he is aware of memetics. He finds that cultural aspects of memetics are too soft, or "mushy." Because written information is stored outside of the body, Rothschild indicated that the transfer and evolution of such information can be tracked through historical records. It remains my personal belief that memetics is compatible with bionomic theory. Today we lack the ability to locate and study the storage sites of memes. Thus Rothschild is correct that in a direct and focused manner memetics cannot help bionomics. One day we may well gain the ability to find receptor sites in the brain. On that day, meme theory and bionomics will converge. Today they remain convergent only in theory. The theory being the effect of coded and replicating information on human existence and evolution.

Chapters 12 and 13: Surplus and Genes; Profits and Technology

Every genetic mutation has an effect upon the survival of that gene. The aggregate effects of various genetic mutations effect the survival vehicle in which those genes reside. The product of genes are their phenotypic expression; Richard Dawkins extends phenotypical effects to artifacts which are outside the survival vehicle. The size and general construction of dams are among the extended phenotypes of beavers.

Whether visible like a beaver dam, or "invisible" like microscopic changes in a survival vehicle's "eye," phenotypic expression affects the ability of a survival vehicle to thrive. Apparent advances in phenotypic expression can lead to extinction rather than reproductive success. Poor expressions might result from genetic mutations that lead to "investments" in wings that are larger than needed, or a larger brain that consumes too much of a survival vehicle's food. Exploring how genes compete within a species, Dawkins finds that marginal changes often work best. Marginal changes provide "niche" advantage combined with reasonable energy costs. Natural selection within a dynamic environment favors the fittest survival vehicles. Fitness can be defined as possessing marginal survival advantage. Winning genes have the best evolutionarily stable strategy (ESS).

The previous paragraph is paraphrased from Richard Dawkins books, *The Selfish Gene* and *The Extended Phenotype*. In chapters 12 & 13 of *Bionomics* Rothschild covers similar ground. Drawing upon his business consulting background, Rothschild is able to present the previously mentioned material, in a highly practical form. All the central data is there. What are absent are the related, more radical "memes" associated with Dawkins, "Libertarians," and perhaps even *Extropy*. This is not to imply Rothschild is white-washing the material, nor to imply that Rothschild is an Extropian at heart. In Extropian terms, the meme set presented in these chapters contains all the important "infectious" thoughts in a manner that is *not* likely to cause a meme-based immunological response. Proof of his success is evidenced by an article by John Hillkirk in the September 20, 1991 Money section of *USA Today* in which Donald Peterson, Retired Ford Motor Company CEO lists *Bionomics* as one of three "favorite books."

Rothschild, in a manner likely to appeal to a "Big Three" CEO, draws detailed case studies from both biology and economics. His analysis shows that successful survival vehicles benefit from the same type of marginal efficiencies. He presents his compelling conclusion in both text and in the form of a typical business income statement. These income statements compare bumblebee hives and super-warehouse grocery stores. Rothschild created his income statements from University of Vermont Zoology professor Bernard Heinrich's physiological and ecological study of bumblebees[7] and case studies by Willard Bishop Consulting Economists Ltd. Not only is Rothschild convincing in presenting the benefits of marginal effects of advantage in evolutionarily stable systems (ESS) but he is able successfully to prove the pivotal function of replicators (like genes or memes), without the language of Dawkins. In the case of the bumblebees this is the phenotypic expression of its genes and the resulting ESS employed by successful hives. In the case of Super Valu/Cub Food, he demonstrates how the grocery business knowledge memes) of a post-Civil War Saint Paul area butcher (Grandfather Hooley) mutated into specific blueprints, training manuals and classroom exercises for the Cub Food franchisees.

Rothschild details how information expressed (phenotypically) in business strategy and practice is the basis for Cub Food's success. Grocery stores are a genus. The species include "24-hour convenience," "neighborhood supermarket," and "mom & pop." Hooley's Butcher Shop, the antecedent species of Cub Food, evolved into a tiny grocery store. It continued evolving, faster than competing locals of the same species. Hooley family groceries evolved steadily into new species. The memes of Grandfather Hooley (minimize costs, maximize revenues, innovate) fostered a new species: the nations first super-warehouse grocery.

By the end of chapter 13, the reader, in the language of bionomics understands the importance of survival strategies [read: ESS] and the role of DNA and information [read: genes and memes]. Rothschild promotes this information with vigor and scholarship while tuning it to his readers, educated people who are business minded. His success in this effort will undoubtably lead to the success of *Bionomics*.

[1]From the front flap of the Jacket to *Bionomics*.

[2]In the modern "libertarian" sense (he believed in "regulation of paper money banking, the compulsory registration of mortgages, government participation in education, the granting of temporary monopolies to merchants engaged in enterprises of risk..., government stamps of quality on plate and on linen and woollen cloth, and the establishment of a maximum rate of interest.)" Ronald Hamowy, *The Scottish Enlightenment and the Theory of Spontaneous Order* (Illinois: Southern Illinois University Press, 1987) p. 21.

[3]Calling Adam Smith, often referred to as the father of laissez-faire economics, an interventionist is an unorthodox view at best. Perhaps it is more fair to say Smith was not an anarchist and his beliefs in the "legitimate role of government" result in visible governmental "hands" affecting the marketplace. If laissez-faire economics is defined in terms of a free and unfettered marketplace, Smith can be called a (limited) interventionist.

[4]Ronald Hamowy, *The Scottish Enlightenment and the Theory of Spontaneous Order* (Illinois: Southern Illinois University Press, 1987) p. 13. In as much as the central theme of Smith's work in *The Wealth of Nations* is based upon his theory of spontaneous order — a theory first delineated in his prior work, *Theory of Moral Sentiments*. "Probably the clearest exposition of the idea of spontaneous order as it related to economic phenomena is offered in the work of

Adam Smith. It should be emphasized, however, that the theory that complex social patterns are self-coordinating and need no deliberate ordering applies as much to Smith's moral theory as to his analysis of the market." Hamowy's footnote: See T.D. Campbell, *Adam Smith's Science of Morals* (London: George Allen & Unwin, 1971), 94-106.

[5]Ronald Hamowy, *The Scottish Enlightenment and the Theory of Spontaneous Order* (Illinois: Southern Illinois University Press, 1987) p. 15

[6]This is a theory put forth by Niles Eldredge and Stephen Jay Gould (Eldredge & Gould 1972). Darwin's theory of Natural Selection is not readily detectable in the fossil record — the fossil record does not reveal gradual change but rather reveal sudden change. The theory of punctuated equilibrium attempts to explain this by showing that a small sub-group of a species can evolve gradually, but when through natural catastrophe, the parent species dies out, the new species can move into that niche. See Dawkins, Richard, *The Extended Phenotype*, (Oxford: Oxford University Press, 1982) pp. 101-109.

[7]*Bumblebee Economics* (Cambridge: Harvard University Press, 1979).

Electronic *Extropy*

Extropy #8 and #9 will be available on diskette (3.5") by the end of July, in various formats, including ascii text and PageMaker.

We will be uploading these issues to AMIX – The American Information Exchange – where anyone can buy them and download them. We will gradually stock the private Extropians market on AMIX with all back issues, and with additional essays, probably including our new bimonthly newsletter *Exponent*. More details will be announced later in *Exponent*, and *Extropy* #10 (Winter issue 1992-93).

Can You Help?

We are looking for a talented person willing to help design attractive covers for *Extropy* twice per year. The sharper our image, the more copies will sell on the shelves.

Extropy Institute has only just got underway and is looking for talented people to help out in several capacities. Of course, financial help is wanted; the first thing to do here is to join as a member. Some of the early members of ExI have joined as sustaining members (Thanks Harry Shapiro and J. Storrs Hall). Other types of help would be valuable too, such as graphic design, fund-raising skills, and help with our projects. We could do with various items of office equipment, especially a good fax, and some additional filing cabinets. If you can help out in any way, please contact Executive Director Max More at the *Extropy* address, or phone (213) 484-6383, or send e-mail to more@usc.edu.

CONTRIBUTORS

Fran Finney: Fran received her BS in Biology from CalTech, did some graduate work in Neuroscience at the University of California, San Diego, and is studying for a Masters in Physical Therapy at the University of Southern California.

J. Storrs Hall: Moderator of the USENET sci.nanotech newsgroup and *Extropy*'s Nanotechnology Editor.
josh@cs.rutgers.edu

David Krieger: David is System's Librarian at RAND, *Extropy*'s Science Editor, an Extropy Institute director, and a former Technical Consultant to *Star Trek: The Next Generation*.
dkrieger@netcom.com

Simon! D. Levy: Editor of ExI's newsletter, *Exponent*, Computing and Linguistics Editor for *Extropy*, and a director of Extropy Institute, Simon! is studying for a Ph.D in Linguistics at the University of Connecticut while working at Haskins Laboratories.
levy%gary@venus.ycc.yale.edu

Hans Moravec: Director of the Mobile Robot Laboratory at Carnegie Mellon University, Hans is author of *Mind Children: The Future of Robot and Human Intelligence*.
moravec@cs.cmu.edu

Max More: Editor of *Extropy* and Executive Director of Extropy Institute, Max received his BA in Philosophy, Politics, and Economics from St. Anne's College, Oxford University in 1987, and now is writing his Ph.D dissertation on *The Diachronic Self: Identity, Continuity, and Transformation* at the Philosophy Department, University of Southern California.
more@usc.edu

R. Michael Perry: Co-Editor, *Venturist Monthly News*, historical columnist for *Cryonics*, and Patient Caretaker at Alcor Life Extension Foundation, Mike has a Ph.D in Computer Science, and is writing a book on death and immortality: *Forever for All*.
alcor@cup.portal.com

David J. Ross: Founder of Calera Recognition Technologies, Vice President of Research for Arkenstone, Inc., and Chairman of RAF Technology, Inc.
davros@sri.com

Harry Shapiro: Manager of the Extropians e-mail list, Harry is Manager of Computer Services for Warwick, Baker, Fiori, has consulted at Bell Laboratories, and worked in technical positions at E.F. Hutton and Shearson Lehman Brothers – the brokerage division of American Express.
habs@panix.com

Printed by Customized Computer Publishing, University of Southern California, Los Angeles. (213) 740-6262

P.O. Box 57306
Los Angeles
CA 90057-0306

Address correction requested