
EXTROPY #12

1st quarter 1994 (vol.6, no.1)

✓ A PRACTICAL LOOK AT OCEAN COLONIZATION	Bill Eichman	05
✓ THE LAST FREE PLACE ON EARTH	Tom Morrow	14
★ LOGICAL LANGUAGES: A PATH TO POSTHUMAN RATIONALITY?	Simon! D. Levy	15
★ THE OPEN SOCIETY AND ITS MEDIA	Mark Miller, et. al.	18
✓ GOD AND MAN AT YALE: A CONVERSATION WITH DAVID ROSS, PART 1	David Krieger	24
✓ FORUM: NANARCHY (AUTOMATED POLICE & DEFENSE SYSTEMS) Hanson, Finney, Szabo, Drexler, Dinkelacker		32
✓ WORMHOLE WARFARE	Robin Hanson	38

THE TRANSHUMAN TASTE (REVIEWS)

✓ <i>Fuzzy Thinking: The New Science of Fuzzy Logic</i>	Max More	42
✓ <i>The Children's Machine</i>	Harry S. Hawk	44

EXTRO 1 CONFERENCE 1994	37
EXTROPY INSTITUTE UPDATE	40
BACK ISSUES	41
CONTRIBUTORS	47

Cover by Ralph Whelan

EDITORIAL

What do you believe in? — In this: that all things must be determined anew.

Friedrich Nietzsche, *The Gay Science* (1882, 1887)

The Extropian way of living is all about breaking barriers, moving ahead, challenging dogmas and limits of all kinds. Late in the last century Friedrich Nietzsche — one of the greatest (and most frustrating) thinkers of all time — stood out from the common run of philosophers, throwing down a gauntlet before humanity. Nietzsche introduced the concept of the *übermensch* (overman), a concept overlapping the Extropian conception of the transhuman. Those who would become *übermensch*, declared Nietzsche, must first undergo a radical process of investigation and reassessment of all the beliefs, ideals, values, and practices around them — a “revaluation of all values”.

Extropians are the torch-bearers of Nietzsche's radical program of reassessment and self-constitution. The *übermensch* ideal finds expression in the Extropian Principle of Self-Transformation. However, the ability to question everything, to hold up to the bright light of reason every supposition, assumption, and dogma, requires more than just the will to do so. We also require intelligent technology to assist us in this Promethean task.

One of the most promising informational tools, already available in an early form, is *hypertext*. In “**The Open Society and its Media**”, Mark Miller, Dean Tribble, Ravi Pandya, and Marc Stiegler provide an overview of the essential features of a mature hypertext information system. Hypertext, combined with massively-interconnected electronic networks, knowbots (personalized information-retrieval software agents), and the digitizing of practically all (articulate) knowledge, should immensely help us in checking our cognitive models of the world.

Overcoming limits and exploring new paths can be frustrated or constrained by the political, legal, and cultural environment. Boundless expansion and unhampered self-transformation therefore require, some will argue, new environments and cultures located in entirely new locations. *Extropy* has and will continue to explore the potential for such spaces to be established in cyberspace and off-Earth. In this issue, Bill Eichman examines the feasibility of growing **new societies on or under the oceans**. T. O. Morrow contributes his research on the international law of the sea.

In “**Logical Languages: A Path to Posthuman Rationality?**”, regular contributor Simon! D. Levy investigates some fascinating methods for cognitive self-transformation and Nietzsche's task of determining all things anew. Complementing Simon's look at constructed and modified



Photo by Nanc Clark

languages, I review Dr. Bart Kosko's recent book on **fuzzy logic** and its scientific, philosophical, and cultural implications. Harry S. Hawk's review of Papert's *The Children's Machine* deals with similar themes.

David Krieger's interview subject this time is Dave Ross, who astounds and intrigues many of us by claiming to be both Extropian and Christian. Judge for yourself whether this is really possible, and don't miss the conclusion of the Daves' conversation next issue.

This issue's Forum brings together five discussants from the Extropians e-mail list to debate the desirability and feasibility of automated, intelligent police and defence systems designed to be outside human control — the “**nanarchy**” idea introduced in last issue's interview with Mark Miller. Robin Hanson also follows up on last issue — this time developing ideas from Michael Price's **wormholes** article — another crossover from *Extropy* Institute's main electronic forum. Expect to see more connections between this journal and our virtual fora in future.

Finally, I'm delighted to present an invitation to ExI's first conference, **Extro 1** — an encouraging sign of the growth of our movement.

Upward and Outward!
Max More

EXTROPY — a measure of 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) Dynamic Optimism; (4) Intelligent Technology; (5) Spontaneous Order. [See *Extropy* #11 for Extropian Principles v.2.5]

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 life-promoting principles and values, while rejecting dogma and religion. [See *Extropy* #6]

Ocean Colonization

A Practical Approach

by Bill Eichman

copyright ©1993

The Dream

- Colonization of the Oceans — The perennial dream of the red-blooded, raised-on-frontier-stories, post-moonwalk, orbit-denied, SF-reading, technology-loving, adventure-hungry, (ex)nerd.
- The chance to face and develop a truly challenging frontier.
- Trillions of cubic feet of low-cost "real estate".
- "Free Oceania", and the opportunity to start an entirely new nation.
- Freedom (hopefully) from the control of governments. (Not from their influence — this is a planet made small by the jet and rocket engine.)
- The chance to build a new, dramatically more sophisticated society, relatively free of the territorial, economic, personal, and religious tyrannies of our ancestors.
- The chance to live drunk on freedom; the freedom of the frontier. The chance to live as a human being, away from the smell of obedience. The chance to make fortunes.

The Reality

- The most corrosive environment on the planet.
- The most violent weather on the planet.
- Death by drowning an ever-present possibility.
- Rapid vertical pressure change.
- Political & Legal vulnerability. Military vulnerability.
- Extremely expensive infrastructure. High start-up cost.
- Extensive specialized knowledge and skills required.
- Psychologically threatening and physiologically sickening (sea sickness, etc.).

I'm an ocean colonization hobbyist. I "waste" my leisure time tinkering with the complicated social and engineering puzzle of living on and under the planet's oceans. I'm qualified to write for *Extropy* on the topic of ocean colonization primarily by default; for there exist no ocean colonization professionals.

The topic of ocean colonization seems inextricably mixed with the "Start your own Country" movements and projects. This creates some very messy problems, in which law, politics, common sense, engineering, and economic imperatives combine to trip up the unwary social dreamer.

Let me make my prejudices clear. I do not believe that it is practically possible to "Start Your Own Country"; not, at least in the ways that are described in Strauss's book of a similar name, or are suggested in the "Oceania" advertisements that are appearing in the libertarian community. I would argue that the basic idea presented by the "Start Your Own Country" proponents is flawed at the root. "Start Your Own Country" enthusiasts appear to be operating from the axiom that if they can establish a new physical and political territory, that they can then take advantage of that frontier to become wealthy and powerful; whereas I think that the evidence of history is that the opposite is true, that is, that if a people becomes wealthy and powerful, then they can define and defend a new territory, and thus create for themselves a new country.

If it were feasible and possible to easily and openly start a new, independent country, why aren't the multinational corporations and multi-millionaires and billionaires already doing it? Wouldn't the tax benefits alone be irresistible?

For all practical purposes, no attempt to form a new country will be allowed to stand, either legally or physically, by the existing gangs of nations. It is intrinsically against the interests of the existing countries to allow any new competitors to enter the "Government Protection Racket". There are only two things that matter in the game of nations — real wealth, in terms of resources, information, trained workers, and liquid capital; and effec-

tive military force and an effective weapons technology and industry.

If you want to start a country, you must start by amassing a fortune, and building a credible army. Anything less, including a move to "International Waters", is essentially an exercise in eccentricity that will be disassembled immediately if it ever becomes a source of irritation for a genuine wealth-controlling, armed and armored nation. (I'll leave for another time discussions about new forms of wealth and new types of military supremacy — possibly the rapid and sweeping technological changes engulfing our planet will make country-starting more feasible. Certainly I hope so; I'm as interested as the next fellow in escaping the totalitarian thugs who have effectively co-opted our planet's governments, even if I'm entirely skeptical of the current "Start Your Own Country" movement and the concept of a "Free Oceania".)

However, it would definitely be possible to increase the *relative freedom* from government interference in one's life by living at the very fringes of one's "Nation". Many taxes could be avoided, and many regulations circumvented, simply by living in a area and in a way which makes government enforcement difficult, expensive, and inconsequential. This is true if you choose to live in Northern Mexico or the Appalachian backwoods, and it would be doubly true if your home was a mobile ocean platform. As long as you avoid attracting large scale media attention, the odds are extremely high that you would simply be ignored by the government.

The big problem with this *relative freedom* approach to ocean colonization is that it may not make much sense to try to establish a safe floating ocean dwelling, which is likely to cost hundreds of thousands of dollars and still be quite spartan, when that same money spent in Northern Mexico or the North American backwoods could build a fairly substantial estate — an estate with no worries about drowning, or sharks, or storms, or corrosion, which are every bit as real and dangerous an adversary as the taxman.

As far as I can determine, after many years of hobbyist study of ocean colonization, there can be only one really plausible reason to pursue the realization of ocean living at this time; and that is the personal love for the idea, for the oceans, and for the frontier. To do it for the pleasure of it, out of curiosity, and a desire to live a genuinely wild and unusual lifestyle.

Please notice the disclaimer "...at this time." We can easily predict that increasing population pressures and mounting competition for natural resources might make ocean colonization a virtual necessity over the course of the next five hundred years. It's possible that getting in on the ground floor of ocean

colonization now could leave hobbyist ocean colonizers in possession of extremely valuable technical know-how in the future. Still, unless you have a genuine love for the dream of ocean colonization, you would probably be wiser to invest your money in real estate, or mutual funds. There is however one possibility, which I describe in the section "What can Extropians do?", which combines, in a way, investment in real estate with ocean colonization, and this might prove to be a real moneymaker within a few short decades.

Can the Dream be Realized?

If ocean colonization is such a good idea, why aren't the rich industrialists already doing it? They've got the giant ships, the subs, the drilling platforms, floating factory complexes, the whole works — why aren't they colonizing already?

That's a good and important question. Clearly, the answer is that there's not enough profit in it.

Even a very cursory examination of the topic of ocean colonization will reveal its

There can be only one really plausible reason to pursue the realization of ocean living at this time; and that is the personal love for the idea, for the oceans, and for the frontier. To do it for the pleasure of it, out of curiosity, and a desire to live a genuinely wild and unusual lifestyle.

major weakness. Put simply, there are few, if any, *economic* incentives to live full time in the ocean environment. The current needs of industry, science, and politics are fairly well met through the use of our sophisticated fleets of ships and submarines. Only insatiable human curiosity, political desperation, and an appreciation of the needs of future centuries appear to justify ocean colonization efforts in the next few decades; there is unlikely to be much money to be made through ocean colonization that cannot be more easily and efficiently made through the use of our existing ocean technology of ships and drilling platforms.

Note that I am saying that it is *unlikely* that there is enough money to be made to justify ocean colonization. What this means is that there are few existing *industrial* resource-gathering or manufacturing processes that can

be profitably carried out within an ocean colonization program. An ocean colony is unlikely to compete effectively by mining manganese or oil, or harvesting fish, because the existing industrial base, established as it is in landside ports, can carry out the whole process of harvesting, processing, and bringing the product to market at a lower infrastructure and capital cost than that required to build an equivalently productive ocean colony. In almost all cases, ocean living technology is going to be significantly more expensive per cubic foot of dry living-and-work-space than even the most extravagant of land-based buildings, homes, and factories. However, the picture is not entirely grim. As expensive as ocean living is likely to be, there may still exist economic opportunities which could fund and support certain scales of ocean colonies.

Examples of such potential economic opportunities include pharmaceuticals, medical care and technologies which are illegal or hard to obtain, the sale of ocean information and ocean telemetry, biotechnology, and perhaps most importantly, especially in the early stages of ocean colonization, tourism. Other possibilities, of which I am somewhat more doubtful, revolve around data-banking, data-piracy, tax-haven banking, and privacy and security services. (I'm doubtful of the utility of these "piracy" strategies because if they are not mercilessly persecuted, they will tend to become almost ubiquitous. Any ne'er-do-well with a computer and a cellular modem would be able to provide these services at a fraction of the overhead that an ocean colony would require.)

There is one other possibility which should at least be mentioned, and that is the business of organized crime; especially, of course, the multi-billion-dollar industry of manufacturing and smuggling illicit drugs. This is a proven money maker, and the type of manufacturing involved could practically be carried out in an ocean-technology environment. However, it would require enormous cleverness, organization, and guts (or enormous stupidity) to build such a criminal empire, and the distribution network for such a business — which would of course have to operate on land — would be a potentially fatal weak link. The organized crime approach to ocean colonization is either the most hard-headedly realistic, with its billions in profit potential, or the most ludicrously science-fictional of schemes.

None of these economic potentials, with the exception of tourism, can be quickly and easily implemented. All of them will require a considerable investment in equipment, trained specialists, infrastructural support, and marketing. None of them (with the possible

The industrialists are the economic equivalent of the frontier "treasure hunters" Cortez, Pizarro, and Ponce De Leon. They're seeking the treasures of oil, manganese, and fish, and as a side effect, they open the doors for settlers, farmers, and manufacturers to come in and homestead.

It took a hundred years for Europeans to start settling the new world, after its discovery by adventurers and mercenaries. Two hundred more for the industrial revolution to turn the resource rich, frontier-mentality new world into a planetary superpower.

We would be unrealistically optimistic if we didn't expect a similar time curve for ocean colonization (adjusted for the modern engineering innovation curve, of course).

exception of pharmaceuticals extracted from ocean animals and plants, and, again, tourism) gain any special benefit from being carried out in an ocean colony.

This is the harsh reality which faces the prospective ocean colonizer. Unless this basic question of profit drivers can be resolved successfully, ocean colonization is doomed to the scrapheap of impractical and unrealized dreams.

What Can Extropians Do?

With this basic picture of some of the pros and cons (mostly cons) of ocean colonization, the question becomes: What should Extropians and Extropy Institute do with the idea?

Certainly, it would be a mistake to adopt a whole-hearted support for the ocean colonization topic, and especially for the idea of a "Free Oceania". This would place us squarely in the wacko camp.

However, if we present the idea cautiously, skeptically, with the attitude that it makes for an interesting 'hobby' or research project, it might have a level of "memetic appeal" that could attract a number of bright, innovative minds to an extropian philosophy. And then, if it works, and makes money, at least enough to be self supporting, so much the better.

To my mind, there is one specific route which is clearly the most practical, cost-effective, and desirable, and that is the route of "Tourism and Education", and it is this route which I could unreservedly recommend for Extropian consideration and investment. In this plan, Extropian entrepreneurs would build a "model" ocean colony which would serve primarily as a resort for paying customers, secondarily as an education-for-profit school and production studio, and tertiarily as a serious ocean colonization research center.

Here are some suggestions as to what a

realistic and effective Extropian Ocean Colonization Program — ExOCol [?] — might look like:

Short Term Plans

An Extropian Resort, doubling as an Ocean Colonization Research Center. Rented on long term lease, or purchased. Possibly near the US Virgins, to take advantage of the educational and communications facilities on St. John and Puerto Rico.

Project funded by paying vacationers, and Infotainment Productions. Other funding and income sources as they become available/possible.

Investment in boats and boatbuilding facilities, computers, communications, recording equipment, laboratory equipment, and other hard assets that can be used to provide services for the resort, and for infotainment production, and for other projects if the research for ocean colonization proves to be a failure, or resold (choose equipment for depreciation resistance), or used to generate alternative incomes.

Carrying out basic research. Possibly running a leisure-learning vacation 'college', where vacationer-students of all ages could participate in experiments and help in gathering data.

The Turning Point

All practical accomplishments will depend on a minimum of two factors:

- (1) The results of basic research, which will tell us what is feasible.
- (2) The discovery, through human cleverness, of the "profit drivers", the untapped resources and markets, that will fund the relatively expensive floating and undersea structures that would provide the infrastructure for an independent economy.

Without an independent economy and infrastructure, Seacolony efforts will be drastically limited; their best use might be for recreation, creative-logjam-breaking, and as 'hermitages' for artists, software writers, and the like.

Long Term Possibilities (some topics for research)

The construction of larger complexes. The continuing search for profit drivers

Providing medical services to the third world, or to any other profitable market.

Consulting in ocean engineering, agriculture, transportation, ecology, telemetry, etc. (Sale of ocean-adapted engineered goods? Specialized manufacturing?)

Ocean mining, high-return ocean harvest.

Disposal of toxics and nuclear waste in stable clay beds.

Sale of pharmaceuticals harvested from ocean flora and fauna and/or patented(?) synthetics based on oceanic biological mother compounds.

Biotechnological research and manufacturing.

Summary


The dream of ocean colonization is seductive, and, perhaps because of the 'space race', and perhaps because of movies and television and print fictions, it has captured the imagination of thousands. Yet, it remains a dream, for, like the colonization of space, it requires extraordinary efforts and extraordinary expenses.

The greatest obstacle to realizing that dream, or of putting it to rest with the other follies of youth, is specialized knowledge. We simply do not know whether ocean colonization is practical, achievable, or even desirable.

The second greatest obstacle to practical ocean colonization is a lack of existing profitable business and industries which could support the great expense of building an ocean-living infrastructure.

The first thing to be done, then, is to organize and carry out the basic research which is required before even the most rudimentary *realistic* attempts at ocean colonization can be attempted. If we're clever enough, even this 'research & development' expense of basic research can be managed in a way that produces an immediate profit, through the marketing of tourism and the sale of infotainment and education.

This is the test of tests. If we're so smart, can we realize our intelligence, by becoming rich?



*Trimaran Flagship
with Semi-Rigid Sails*

Floating Platform

Based on existing trimaran patterns for safety, capsize-resistance, and maximum mobility. A deck of carefully engineering, light-but strong trusses, using carefully selected materials, some natural, some synthetic.

The development of computer controlled rigid wing sail technology.

Speculative Technology



Important Disclaimer: These technological speculations, and the drawings accompanying this section, are intended only to promote discussion and further research among hobbyists, and are not in any way to be critiqued as definitive. Much of the engineering calculations still not been performed, which might quickly reveal significant flaws in these hobbyist designs.

Ocean Colonization Technology

Tender Flock

Conventional transport boats, houseboats, special purpose boats of all sorts.

*Trimaran Flagship
with Semi-Rigid Sails*




Surface Pod: semi-open, sealable in the event of storms.

Essentially, a type of barge, multipurpose, with redundant flotation systems, a life support 'office' with communications, alarms, & safety systems, and facilities for attaching engine-propellor and tug hauling systems.


Reinforced concrete hull systems, adapted from conventional concrete hull technology, using a gunnite type fiber-reinforced sprayed concrete

*General Purpose Pod
(Housing, Utility, Light Industry)*



Calcium electrodeposition as a construction technique?

Section Through Pod



Submerged Pod

Cylindrical, with domed ends, for pressure resistance. Can operate from floating-to-triple-atmosphere depths: 0 to 30 feet, approximately three stories. 18' diameter?

Constructed of fibred concrete, plasti-rod reinforced, post-tensioned with polymer cables. Multilayer, insulated, with resinated fiber inner and outer coatings.

Steel and stainless steel fittings. Bronze, other alloys, polymers and resins of all types. Ceramics, ceramicized coatings. Biological (bacterial) coatings, should they prove feasible.

Can operate pressurized or unpressurized. Domed ends can be replaced with standardized spherical airlock chambers.

Develop a universal modular connecting system, a combination doorway, airlock, and emergency seal system to standardized dimensions.

A 36' diameter spherical submerged pod system should also be developed to work in tandem with the submerged pod.

Provision for weighting the pods, with a redundant, automated neutral buoyancy system. (Build pod walls 'overthick' and heavy, to reduce need for special weights.)

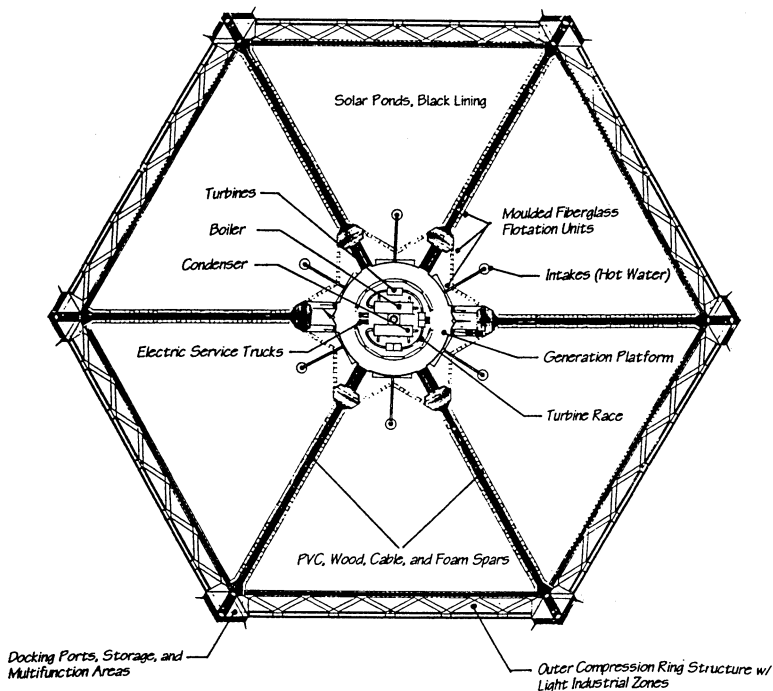
Light buoys, piping sunlight to pods. Light and viewing ports in chosen areas of the pods.

Combined Floating Complex: tethered, and self positioning.

An engineered collision skirt/artificial lagoon interlinked floating barrier, within which floats integrally tethered 'stacks' of floating platforms, surface pods, and submerged pods, with their flocks of tender boats. Ordinarily 'semi-free-floating', either tethered to sea floor, or using computer


S.O.T.E.C.

Solar Pond/Ocean Thermal Energy Generation



S.O.T.E.C.

Solar Pond/Ocean Thermal Energy Generation



Such a complex could position itself near an OTEC/SOTEC/kelp methane plant, or near a gulf stream turbine, should such prove feasible, for 'cheap' power.

Energy Systems, Energy Storage, Energy Use Patterns.

First and foremost, design for energy conservation. Reduce or abandon energy-intensive, low-survival-value systems or habits. Concentrate energy use on productive activities or businesses, and cultivate an ethic of energy efficiency, and reward and praise productive energy use. (Even if an abundant, cheap energy source is discovered, this efficiency ethic will give us a competitive edge, as in the pacific rim countries.)

Solar pond technology, SOTECs, freefloating in the ocean.

Methane digester bladders, freefloating. (biotech)

Wind energy systems, where practical.

Solar energy systems, high quality electricity for electronics.

SOTEC, OTEC, Tidal, Wave, and Current generation, where practical. (MHD?)


controlled position systems, could float freely with currents, or under power, if travel is desired.

When not under rapid movement, ordinarily surrounded with floating complexes of solar ponds, food ponds, and 'industrial' systems. These would be left behind, or "reeled in", when high speed travel was deemed necessary.


High speed travel would be with powered tug assistance.

Solar distilled alcohol for liquid fuel backup, possibly as main fuel source for some elements of the colony. 160 proof solar alcohol can be a cost effective fuel in steam, turbine, and modified IC engines. (biotech)

Purchased diesel, gasoline, and bottled gas fuels. Colony-owned ocean floor natural gas wells, if feasible.



SOTEC Isometric View



Collision & Ecology Barrier

Other energy systems, should they become practical and available. (fuel cell, nuclear, hydrogen, etc.)

Food Systems/Biomass Systems.

Algae/seaweed ponds. Water fertilization systems, employing mineral rich bottom water pumped by windmill-buoys, SOTECs, OTECs and other nutrient-enriching systems, including human waste disposal in biomass (not food) ponds. Freshwater ponds on barges?

Kelp farming in union with an OTEC/methane plant? Must be in a deep water location. Ideal end result might be to use catalysts or biotech enzymes to convert methane into a more stable hydrocarbon liquid fuel, to be stored in double-lined floating bladders in tethered cages, for colony and trade use.

Fenced fish ranching. Ocean fishing management w/computers. Nutrient irrigation.

Robot plankton sieves? Human guided plankton sieves.

Greenhouse pods, boats, and platforms. Micro-ranching rabbits/chickens if desired.

Agricultural biotechnology, location of indigenous species suitable for breeding towards domesticity.

Yeasts, algae, and similar "artificial foods" produced with biotech factories.

Trade excesses of sea harvest for land field crops; grains & beans.

A Sample of Ocean Colonization Discussion on the Extropian E-mail list

From: wce@hogbbs.scol.pa.us
Subject: Sea Colonies

>Bill: I read your comments on ocean colonization with a great deal of interest. The idea of floating ARKs or self-sustaining communities is one I've often thought about since I was a little kid. My question is this: Would they HAVE to be mobile? What about used-up ocean oil drilling platforms? They are not as huge as a tanker but stable and designed to take a beating. Is food the limiting factor? Would the area surrounding an ocean platform ever get "Played Out?"

I have learned that there are as many strategies for ocean colonization as there are thinkers who are dealing with the idea.

I am assuming various teams would have to be assembled to study the pluses and minuses of these various strategies. That is, someone will have to study the questions surrounding obtaining and using structures like drilling platforms.

Such platforms could only ever be a tiny part of a real effort, however. They are limited in number and more suited to "retreat" than to colonization.

Stable bases such as these platforms would have to be built as part of

a colonization effort, I would say. These stable bases would be designed as "servers" for a variety of other permutations of colony.

I concentrate more, for now, on shallow water colonization — I haven't studied stilted sea platforms, so, really, I just don't know.

>Have you given thought to the effects of storms/hurricanes on a floating colony?

No one who looks at the idea seriously can fail to give a lot of thought to questions of weather, corrosion, and so forth. The sea is much more dangerous than the land.

I think these problems are solved using multiple strategies, including the following...

Some places are better than others... leeward of islands, and naturally calm regions, will be colonized first.

Breakwater and artificial lagoon technologies, as T. Starr suggests.

High-quality satellite tracking and communications, combined with various degrees of mobility. Sea colonies would be designed to shut down and lock up tight in the face of bad weather.

Structural designing that takes into account the 100-year storms, with combinations of strengthened cores and breakaway peripherals.

The ultimate protection against storms is found, of course, by going *down*. Twenty feet down, a region of comparative stability is found. To my mind, real ocean colonization, as opposed to occupation or hunt-and-gather harvesting, is dependent on achieving an undersea living technology.

Just as with the colonization of the new world, I expect the process to start with the adventurers and the desperate. They will gather enough information to arm the entrepreneurs.

Just as is occurring with 'space colonization' right now... ocean colonization is a 'poor man's alternative' to orbital industry and LaGrange colonies.

What I'm proposing, primarily, is the process of initial R&D. I expect that initial R&D to be salable, even if all we have to sell are warnings, and entertainment for the jaded millions.

I want to emphasize that I am not promoting ocean colonization in and of itself. I feel pretty aware of how tough real colonization is going to be. What I'm suggesting is the very practical and real establishment of an ocean coloni-

zation research center, somewhere in the Caribbean, within the upcoming 5 to 25 year time frame.

I don't think we know enough, or have the right businesses ready, to make ocean colonization at all practical anytime within the next 25 to 40 years. It's another profit desert to be crossed.

But, based on quite a bit of study, I think we could make an ocean colonization *Research Center* profitable immediately, by selling infotainment, while doing the basic research that will lead to patents and future, larger profits.

The ultimate protection against storms is found, of course, by going *down*. Twenty feet down, a region of comparative stability is found. To my mind, real ocean colonization, as opposed to occupation or hunt-and-gather harvesting, is dependent on achieving an undersea living technology.

I see it as a 'doable' experiment in independent science. Most R&D is the province of state controlled or influenced universities or corporate laboratories — which can be fine establishments, but they need competition to keep their bureaucracies and administrations honest. I think there's profit to be made in an independent, capitalist, science-as-business.

An ocean colonization research center can be just the sort of wild collision-of-ideas places that keeps the creative juices flowing. And, it can be a place to explore the huge third world market. (I think it was Charlie Stross who talked about leapfrogging wires and marketing cellular phones in Bangladesh — it's pretty clear to me that this is just the tip of the iceberg of the trillions to be made in using our cleverness to design, manufacture, and sell products that are high profit and focused directly at the third world market.)

So, this is the slant I'm taking with this ocean colonization thread; I don't know if ocean colonization is feasible, I don't know if we could build even a miniature floating/under-sea 'Hong Kong', but I do know that I personally expect to be moving to the Caribbean and trying to set up an ocean colonization center as a part of my retirement.

The thing to do is to rent an inexpensive dock

and a cluster of buildings on some 10 to 50 acre piece of land, presumably somewhere in the Caribbean. We set up solar and other desalination, and start collecting data on their performance as we let cisterns fill with fresh water. And we set up a small marina, a boatbuilding business, and experiment with buildings on the island and on boat platforms, working towards building a first larger 'flagship' platform trimaran that would hold the central computer labs, biotech labs, communications equipment, and the other essentials of the research center.

All through this period we are exploring, doing underwater photography and video for later sale, and studying the profit potentials of the region.

While the onsite people are looking for regional profit, the teams still "back home" in the mainstream study & develop ways that the ocean center's material can be marketed as infotainment for profit. I'm assuming that some combination of education/tourism package would be sold, in which wealthy yuppies pay to visit and play with submarines, and students pay for credits, as one of the cornerstones of the first 25 year's income possibilities. Whatever else, in the form of tapes, magazines, books, stereo VR simulations, etc,

etc, will be decided by assessing the possibilities at the time — but I'm assuming we'd have to market ourselves aggressively from day one, as opposed to counting on grants and patronage.

I think I'm most interested in working on the following:

Water desalination and purification devices, potentially salable throughout the undeveloped countries, based on solar or osmotic filter technology.

A boatbuilding plant (automated?) to produce the functional equivalent of the one family house, in the form of a 'smart' trimaran houseboat, out of resinated organic fibers. The "Model T" of the ocean lifestyle and economy.

Intensive studies of the local biology and ecology, looking especially for chemical compounds with pharmacological and industrial uses, and also for 'agricultural' breeding stock. Gathering baseline data so we can assess ecological impact accurately years later.

An intensive calcium electrodeposition study.

A study of coatings, starting with the best data we can gather from the existing oceanographic

engineering experts.

And so on — pretty humble stuff, actually. It can be done with at first tens, then hundreds of thousands of dollars, which is a capitalization that is within my reach. There is, of course, a tremendous amount of advance work and preparation that has to be done for something like this — we can't afford to waste a drop of energy reinventing the wheel, meaning that we need to be up-to-date on the subject and the literature, and we have to be skilled with the video, computer, and lab equipment, so that we avoid the common business hassle of trying to start a business and learn how to use the equipment at the same time.

As far as something like data piracy or free banking or suchlike, it seems pretty clear that your basic 80 to 240 ft refitted commercial ship, running under a protective flag and docked in a protective country, would be your best bet. (Although, it occurs to me that a factory to produce a "delivery truck" style of submersible could revolutionize smuggling.) Mobility and invisibility is your best protection against the organized killers and torturers (soldiers and police) that the government will turn against you once it becomes aware of you. Assuming you want to get involved in such a thing, which I don't.

Well, this is way long, so I'll just stop here, leaving questions unanswered. As far as I know, there are no definitive texts on ocean colonization, and I've learned about it by reading bits and pieces from a hundred incidental sources.

I'm going to work on assembling some files about it, and I'm working on a comic book exposition of the ocean colonization idea, though that goes slowly so far. I think it's a potent fantasy/meme, getting ripe for harvest.

Later, Bill wce@hogbbs.scol.pa.us

And more discussion of ocean colonization from the Autopia mail list.

From h1r@lems.brown.edu (Henry Robinson)
Subject: Some info I have found...

>Bill Eichman writes:
These are some of the things that need to be done first, imho.

(1) Survey the literature, scan the relevant articles into files, and build a database for education, research, and to attract expertise. ...The place to start is at your local university library...

I've been searching for information and have found a very good reference. It's the 2 volume proceedings of a symposium during 1985 called

Ocean Space Utilization '85. It was held at Nihon University in Tokyo, Japan. Within its 1400 pages the sessions cover the engineering and economics of a seabased village, structural mechanics of ocean structures, dynamic response due to wind and waves, OTEC resources and wave energy design, materials and construction, concrete durability, and corrosion just to name a few relevant topics. J.P. Craven of Univ. of Hawaii wrote the first article titled "A Seabased Village" and mentions Buckminster Fuller, Paolo Soleri, and Kiyonori Kikutake in Japan. I just built a gopher client and have been "travelling around" to University Libraries to see what I could find.

Based on quite a bit of study, I think we could make an ocean colonization *Research Center* profitable immediately, by selling infotainment, while doing the basic research that will lead to patents and future, larger profits.

At the Univ. of Hawaii:

TITLE(s): Floating marine community : research report of the Department of Architecture, University of Hawaii, for Marine Programs, State of Hawaii, Honolulu, Hawaii/ Craven ... [et al. ; Hugh Burgess, editor, Cynthia Ai, assistant editor].

Honolulu : The Department, c1972.49 p. : ill; 22 x 35 cm.

OTHER ENTRIES: Hawaii's Floating City Development Program. Offshore structures Design and construction. Craven, John P. Burgess, Hugh. Ai, Cynthia. University of Hawaii at Manoa. Dept. of Architecture.

TITLE(s): Floating city 2 / [compiled by] Masanobu Kosugi.

1971. 1 v. (various pagings): ill.; 28 cm. P. 14.

Caption title: Floating city 2: data for floating city programs in the state of Hawaii. Photocopies of papers and articles by J.R. Stewart, Kiyonori Kikutake, Eduard Anahory and others. Compiled for Prof. H. Burgess, Architecture 488, Winter 1971.

OTHER ENTRIES: Hawaii's Floating City Development Program. Offshore structures Design and construction. Kosugi, Masanobu. Kikutake, Kiyonori, 1928-

Note: Many references to Japan being very interested in this and funding studies due to the land shortage there. Studies are ongoing in three areas. 1) Ocean Communications City (OCC) which would be built "scores of kilometers off the mouth of Tokyo Bay" with four levels of decks. the overall size would be 5 km. by 5 km. for a total 25 sq.km./deck or 100 sq.km. total living space. It includes a international airport on the roof. It would be supported by 10,000 pillars on the seabed. Est. Population: 1/2 to 1 million. 2) Man-made Islands and land extension. Japan is now building a airport and industrial complex on fill by levelling some near-by mountains gaining space on both sides of the equation. 3)

Floating Villages positioned in the deep ocean areas. This area involves studying ocean dynamics and structures. OTEC power plants form the basis of abundant power and economic viability. The Pacific Ocean off-shore from Japan requires stability in the form of a floating platform type structure such as oil rigs. A large portion of the symposium dealt with corrosion and tests of long term stability of metals and concrete. Also, the dynamics of a long OTEC pipe is studied.

>Bill also says: (in response to the amount of energy to bring up the deep ocean water)... The theory is interesting, but the experiments have to be done before we can be sure. The construction, weight, and stability of the intake tube, and the actual power requirements of pumping, may result in more problems than we might expect. It may, or may not, work as well as we would like. The same is true of solar desalinization, and small scale osmotic desalinization. Everything starts with the experiments — and the experiments start with a detailed survey of existing literature.

60 gallons/minute per horsepower. 4 deg.C. water is 1000 meters deep. I want to correct a error I made earlier about the possibility of fresh water due to the ice caps in deep water thermoclines. ENNNK! Not in the tropics. Oceanographic Surveys in the Sargasso Sea region show that the salinity does not change with depth and temperature.

Some of what I have read to date has changed my opinion about the reinforcement of the concrete. Traditional iron and steel reinforcement will corrode and expand causing the concrete to crack. Other materials should be evaluated. Surrounding the structures with fresh water reservoirs using a plastic barrier would have the added effect of providing fresh water for non-indigenous water fowl and freshwater fish. This requires a benevolent place to have any chance. Does anybody know of any other calm spots in the oceans similar to the Sargasso Sea in the North Atlantic? El Nino?

THE LAST FREE PLACE ON EARTH

by T.O. Morrow

Suppose that you want to start a sovereign territory. You're looking for a little corner of the Earth where you and your friends can build a society based on real consent. You don't plan on being shy about it, either. You want to stand tall, proclaim your independence, and tell statisticians to KEEP OUT! It's a great idea.¹ But where are you going to go?

Forget *terra firma*. Statisticians have claimed all of the Earth's real estate as their own — from entire continents down to rocks that barely rise above the waves. Only one place on Earth remains free of statisticism: the high seas.

Should you therefore set sail under your own flag? That may serve as an important step in the right direction, but statisticians (much less investors and settlers) won't take you seriously until you can claim a fixed site as your own.² That leaves just one option for openly establishing a new sovereign territory on Earth: building an artificial island on the high seas.

To keep things practical you'll want to build where the ocean's floor comes relatively close to its surface, such as on a bank or seamount.³ Fortunately, there are plenty of such shallow spots in Earth's oceans. Unfortunately, many of them fall within areas claimed by statisticians. Figuring out which of the sites suitable for building an artificial island lies outside of statisticians' claims will require that we look into the international law of the sea.

Let's take as our standard the United Nations Convention on the Law of the Sea of December 10, 1982.⁴ Although it is still not officially in force, the '82 Convention codifies many widely recognized customary legal principles and has already begun to shape statisticians' claims to ocean territory.⁵

The '82 Convention grants every coastal State an exclusive economic zone (EEZ) reaching up to 200 nautical miles (370.4 km) offshore.⁶ A State has an exclusive right to construct artificial islands and other installations within its EEZ.⁷ This alone rules out many of the best sites for a new sovereign territory. But the '82 Convention also grants coastal States rights over their continental shelves for at least 200 nmi offshore, and sometimes up to 350 nmi (648.2 km).⁸ Essentially, then, the '82 Convention puts every shallow coastal area out-of-bounds.⁹

Fortunately, seamounts offer many relatively shallow building sites on the high seas. Some 10,000 seamounts appear in the Pacific alone; others appear in the Indian Ocean, the Caribbean, the South China Sea, and the Gulf of Alaska.¹⁰ Of course,

many of these sit too deep for present technology to put to good use. But estimates put over 70 seamounts within 185 m of the ocean's surface — and well outside of any statisticians' territorial claims.¹¹ An appreciable number of these come within 60 m of the surface, and several come within 9 m. Consider, for example, Vema seamount (63 m deep, at 31°38'S, 08°20'E) and Walters Shoal (42 m deep, at 33°13'S, 43°51'E).¹²

Sounds easy, doesn't it? Just find a seamount on the high seas and start building! But it's not that easy, of course. Although free from statisticians' territorial claims, the high seas cannot escape the jurisdiction of the international law of the sea. That law sets up various barriers, some merely nettlesome and others potentially crippling, to founding a sovereign artificial island.

Let's deal with the lower hurdle first. The international law of the sea denies artificial islands all maritime territorial claims except for narrow safety zones (usually of 500 m). This would put a sovereign artificial island at a distinct disadvantage relative to land-based statisticians. But there may be a loophole in this territorial restriction whereby alluvions, such as those deposited in the lee of an ocean current, can form natural islands *even if provoked or guided by human works*.¹³ This suggests that you might grow a natural island next to your artificial one, and thus claim rights to a territorial sea, an EEZ, and a continental shelf.¹⁴

Now let's charge the higher hurdle. International law has traditionally demanded that the high seas remain free for all to use. The authority of even statisticians to build artificial islands on the high seas — much less claim those islands as extensions of their territories — rests on shaky legal foundations.¹⁵ Commentators thus dismiss the possibility that international law would permit a new sovereign territory to claim an artificial island on the high seas.¹⁶

The international law of the sea is hardly writ in stone, however. It remains largely customary, widely untested, and susceptible to diplomatic and military pressures. This holds doubly true with regard to the establishment of sovereign territories on artificial islands. Though they face formidable technical and legal barriers, they still "may be legitimized through general recognition by the existing subjects of international law."¹⁷

Earning such recognition will surely require that newcomers to the international community meet the generally accepted criteria for sovereign status: 1) permanent

population; 2) defined territory; 3) mode of governance; and 4) capacity to enter into relations with other sovereigns.¹⁸ Add to this a fifth criterion, absolutely essential but hard to define: respect. Earn enough of it, and all the other barriers to founding a new sovereign territory on the high seas will wash away.

NOTES:

¹For a more thorough discussion of the benefits of founding an Extropian sovereign territory, see Tom W. Bell, "Extropia," *Extropy* #8, (Winter 1991-92), p.35.

²A defined territory is among the standard criteria for sovereign status. N. Papadakis, *The International Legal Regime of Artificial Islands* pp.114-15 (Sijthoff 1977). For further discussion of these criteria, see below.

³Herein I use "seamount" to refer both to seamounts and guyots, though technically the latter is a flat-topped version of the former.

⁴UN Doc. A/CONF.62/111 with corr.

⁵Renate Platzoder, "Conferences on the Law of the Sea," at 69, 75 in vol. 11 *Encyclopedia of Public International Law* (North-Holland 1989).

⁶Arts. 55-75.

⁷Arts. 60, 80.

⁸Art. 75. What prevents legal continental shelves from always reaching the 350 nmi limit? Art. 76(5) also limits them to within 100 nmi of the 2,500 meter isobath.

⁹There appears to be a notable exception to this rule, however: a roughly 2750 km² (1700 mi²) portion of the Saya de Malha bank, shaped somewhat like a slice of pie pointed east-by-southeast and centered at 10°00'S, 61°75'E. For an excellent chart and descriptions of this area, see R.L. Fisher, G.L. Johnson, and B.C. Heezen, "Mascarene Plateau, Western Indian Ocean," *78 Geologic Society of America Bulletin* 1247 (October 1967).

¹⁰S.K. Keaton and J. Judy, "Note—Seamounts and Guyots: A Unique Resource," *10 San Diego LR* 599, 601 (1973).

¹¹Papadakis, *The International Legal Regime of Artificial Islands* at 15 (cited in note 2).

¹²Raymond R. Wilson, Jr. and Ronald S. Kaufman, "Seamount Biota and Biogeography," at 355, 358-59 in Barbara H. Keating, Patricia Fryer, Rodney Batiza and George W. Boehlert, eds., *Seamounts, Islands, and Atolls* (American Geophysical Union 1987).

¹³Fritz Munch, "Artificial Islands and Installations," at 38, 38 in vol. 11 *Encyclopedia of Public International Law* (North-Holland 1989). Article 121(1) of the '82 Convention defines "island" as "a naturally formed area of land, surrounded by water, which is above water at high tide." For a contrary view based on law prior to the '82 Convention (which readers will recall is still not in force) see Papadakis, *The International Legal Regime of Artificial Islands* at 89-97 (cited in note 2).

¹⁴Article 121(2-3) of the '82 Convention grants islands the same status as any other land form unless they are incapable of sustaining human life or economic activities.

¹⁵Papadakis, *The International Legal Regime of Artificial Islands* at 55-79 (cited in note 2); Craig W. Walker, "Note—Jurisdictional Problems Created by Artificial Islands," *10 San Diego Law Rev* 638, 649-652 (1973); Alfred H.A. Soons, "Artificial Islands and Installations in International Law," 7-12 in *No 22 of Occasional Paper Series* (Law of the Sea Institute at Univ of RI 1974).

¹⁶Papadakis, *The International Legal Regime of Artificial Islands* at 113-14 (cited in note 2); Walker, "Note—Jurisdictional Problems Created by Artificial Islands" at 648-49 (cited in note 16).

¹⁷Papadakis, *The International Legal Regime of Artificial Islands* at 114 (cited in note 2).

¹⁸Id at 114-15.

LOGICAL LANGUAGES: A Path to Posthuman Rationality?

by Simon! D. Levy

I. Introduction

Most of us have heard of artificial languages with names like Esperanto and Interlang¹, that purport to avoid various shortcomings of ordinary natural languages through intentional design. Arguments in favor of these languages have ranged from the supposed necessity for a culturally neutral communication medium, to the desirability of a language whose grammar has no pesky exceptions to memorize.

Despite the collectivist mindset behind many designed languages, two such languages have emerged that seem worthy of Extropian attention. Though these languages, Lojban and E-prime, represent polar extremes in the degree to which they differ from ordinary languages, both seek to increase the rationality of their users by eliminating various types of ambiguity and generalization available in a language like English.

II. Lojban

During the 1930's and 40's a fire prevention engineer turned linguistic anthropologist by the name of Benjamin Lee Whorf conceived a curious and somewhat revolutionary idea about human language and its relation to thought. In Whorf's own words,

... the forms of a person's thoughts are controlled by inexorable laws of pattern of which he is unconscious. These patterns are the unperceived intricate systematizations of his own language. [1956, p. 252]

Named the *Sapir-Whorf Hypothesis* in acknowledgment of the contribution of Whorf's teacher Edward Sapir, this idea has come to exert a profound influence on philosophy and anthropology during the present century.


Nevertheless, because of its essentially speculative nature, the Sapir-Whorf hypothesis remained largely an attitude toward language, rather than the basis for a scientific research program, until 1960. In that year, Dr. James Cooke Brown wrote an article for *Scientific American* describing Loglan, a language that he had developed to test the hypothesis put forth by Whorf. Brown hoped to

create a "logical language," free of exceptions and irregular forms, based on formal properties of mathematics—properties which he believed gave human beings the ability to reason. If Whorf's idea had any validity, reasoned Brown, then speakers of Loglan would end up with more disciplined and powerful minds than speakers of natural languages.

In attempt to maximize the "target population" of potential Loglan speakers, Brown picked the eight languages with the largest number of speakers (English, Mandarin Chinese, Hindi, Russian, Spanish, Japanese, French, and German) as a basis for the sound structure of Loglan. For each Loglan word, he found the closest corresponding word in each of the eight languages and allowed each language to contribute to some sub-part of its word to the Loglan word. He then computed a "learnability score" for the new word by multiplying the contribution of a given language by its representation in the language set, and summing over the contributions from each language. For example, the Loglan word *blanu*, meaning "blue," contains all of the sounds (/b/, /l/, and /u/) from the English word, and one half the sounds from the Hindi word for blue, *nila*. Therefore, English contributes all of its percentage points, 28%, to the learnability of *blanu*, and Hindi contributes half of its points (1/2 of 11% = 5.5%) to the word. (With the other six languages' contributions added in, it turns out that *blanu* has a learnability score of 76%).

This word-making scheme points out a major design feature of Loglan, viz., its phonemic spelling system[see sidebar], which facilitates its usage in written communication. Unlike most phonemically spelled languages, however, Loglan has complete resolvability between its spoken and written representations: Given a string of Loglan sounds, one can always determine which sounds go together in a word. Contrast this situation with English, in which a given sequence of letters can represent more than one phrase, for example, ANICEHOUSE.

This resolvability rests on two further



design features of the language. First, Loglan builds its words around a regular consonant-vowel "skeletal structure," which differs for each word class. For example, predicate words, which describe observable real-world things like objects, actions, and qualities, always contain a consonant cluster (like the /bl/ in *blanu*) and end in a vowel. Simple operators, such as the pronoun *da* "he," consist of a consonant followed by a vowel.

Second, Loglan grammar lacks syntactic ambiguity, which means that the role of each word in a Loglan sentence has only one possible interpretation. Again, natural languages do not always behave this way, leading to ambiguities like "Flying planes can be dangerous," in which you don't know whether to worry about the flying planes or the act of flying them. More specifically, Loglan's designer modeled its syntax on predicate calculus, a system invented by logicians to represent propositions about the world in a simple and unambiguous way. Each Loglan sentence contains a predicate (roughly, the verb), which describes an action or condition, and a set of arguments, which describe, for example, who

performed the action (the subject) or suffered its consequences (the object). Like predicate calculus, Loglan makes verbs out of things that speakers of English would not normally consider verbs, such as "blue".² Unlike predicate calculus, in which the predicate comes first, Loglan uses the word order subject-verb-object, because of the rarity of verb-first order in natural languages, and because the syntax ends up simpler if the predicate does not come first.

Though it has received scant attention in mainstream academic linguistics, Loglan maintained a core of hardcore followers and now has an incorporated organization – the Logical Language Group, headed by Bob Le Chevalier (a.k.a. lojbab) – to support and promote the language, now called Lojban.³ A variety of on- and off-line documentation, in addition to learning tapes, software, classes, and discussion groups, exist. This brings us to the question of whether Extropians should spend any time learning Lojban.

Despite its possible merits as a means of disciplining our thought, Lojban differs so strongly from English (and any other natural language) that learning it would require a significant investment of time and effort. We might better spend such time and effort spreading Extropian ideas through more traditional channels of communication, like English, which has more speakers than any other language in the world. This problem leads us to wonder whether we can't fix what's wrong with English, without giving the language up entirely.

III. E-prime

Proponents of another artificial language called E-prime have argued that we can still speak English, but eliminate a good deal of what's wrong with that language, if we follow the advice of Dr. Alfred Korzybski, one of the founders of a field known as general semantics. Unlike the traditional academic discipline of semantics, which tends to focus on abstract issues in formal linguistic theory, general semantics concerns itself with the practical implications and consequences of language use and abuse in everyday realms such as teaching, advertising, and news reporting. Dr. Korzybski's student David Bourland contributed to the field by proposing a derived language, "English minus 'be'", or "E-prime", to remedy the fact that "be" in all its forms ("am", "are", "is", "was", "were", "being", "been"⁴) allows for tremendous ambiguity.

Consider, for example, the sentence "John is a liar.". What exactly does someone mean when they say this?⁵ Should we believe that John always lies whenever he says something? Obviously not, or John would have a good deal of trouble getting by in the world. Does the speaker claim that John lies more often than

not – say, 60 percent of the time? Probably not, as the speaker would have a great deal of difficulty persuading us that they had recorded everything that John ever said and divided the number of John's false statements by the total number of statements John made. Instead, when someone says "John is a liar" they probably mean something closer to "I have heard John lie enough times that I don't trust him to tell the truth about matters of importance to me.". This sentence, though much longer than "John is a liar," has the advantage of forcing the speaker to clarify their feelings about John, without making a claim that the listener would find difficult or impossible to falsify.

More generally, eliminating "be" helps speakers avoid the trap of *attributing* a quality or a behavior to a noun (person, place or thing) without *specifying* the conditions under which they observed the quality or behavior. Specificity costs more words but buys the advantage of *falsifiability*, which the philosopher of science Karl Popper considered the criterion for judging scientific hypotheses. For this reason, we have reason to expect that speaking E-prime instead of English will make our dis-

course more scientific, and hence more rational.

Dropping "be" also prevents the employment of the passive voice. The passive voice (as in "He's been killed.") allows speakers to avoid the responsibility of naming the person or group who performed an action. Think of all the times you've heard "I was led to believe that..." as an excuse, and you'll see the benefits of such a constraint.

Some critics of "pure" E-prime have pointed out that "be" can also function more harmlessly as an auxiliary verb to indicate the *aspect* (discrete or continuous) of another verb, as in "I was running around all day yesterday." One such critic, William Dallmann (1992) has argued for another form of the language, which he calls E-Prime_{mod}. According to Dallmann,

The mod version eliminates the is of identification (He is a general semanticist), the is of predication (She is beautiful), but retains is as an auxiliary (She is dancing), the is of existence (To be or not to be), and the denial of identity (The map is not the territory). [1992, p. 134]

Phonemic Spelling

Phonemic spelling, an essential feature of Lojban, can be described to a first approximation as a system of one-to-one correspondence between linguistic sounds and the letters that are used to write them. (We do not wish to call such a system "phonetic spelling," since a phonetic transcription includes details that are not relevant to distinguishing one word from another.) Phonemic spelling, found in languages like Spanish and Korean, contrasts with the confusing muddle that is English orthography. For example, the first sound in the English word "she" can also be spelled with the letters "ti," as in "action," the letters "si" as in "fusion," the letters "sci" as in "conscious," or the letters "su" as in "erasure." This proliferation of spellings, caused mostly by the fact that English has borrowed heavily from several different languages during its history, led George Bernard Shaw to observe that the word "fish" could equally well be spelled "ghoti," with the "gh" derived from words like "tough" and "rough." Of course, Shaw conveniently ignored the fact that "gh" can only sound like "f" at the end of a word, and "ti" like "sh" in the middle, but his point was well taken.

The most serious linguistic objection to a purely phonemic system is that English is to some degree *morphologically spelled*; that is, the same letter or group of letters can be used to spell two different sounds that encode the same meaning. For example, the "s" at the end of "cats" and the "s" at the end of "dogs" are phonemically different (one corresponding to the /s/ in "sue" and the other to the /z/ in "zoo"), but both serve to signify plurality in nouns and singularity in regular verbs. Still, such examples are vastly outnumbered by the situations in which phonemic spelling would represent an improvement.

From a more practical standpoint, it would be tremendously expensive to rebuild keyboards and other such devices to reflect a new spelling system; however, there are straightforward phonemic alphabets (e.g. ARPABET) that use the standard character set of English, so this objection is not all that serious. There is also the problem of what to do with billions of volumes of books and other works spelled in the traditional way. With the increased availability of optical character readers and the proliferation of personal computers, it is not unreasonable to assume that such literature could eventually be converted on-line into ARPABET or a similar system. No matter what language we end up using, it makes sense to push for a phonemic spelling system.

A more troublesome objection involves the possibility of getting around the restrictions of “be”-lessness by the common device of presupposition. So, for example, instead of saying “John is an liar: He told me he’d be here at eight, and it’s now nine-thirty,” I can say “John, that liar, told me he’d be here at eight,” The information conveyed in the former utterance differs little, if at all, from that conveyed in the latter; I have simply chosen to include the claim “John is a liar” as given (pre-supposed) in the latter.

Still, the inability to use any form of “be” can lead to tremendous improvements in our sense of what we actually say (or fail to say) when we speak and write. Furthermore, E-prime has the obvious advantage over languages like Lojban in ease of usage (for English speakers, anyway) and, therefore, in the amount of effort it would take to become a fluent speaker of the language. Though I did not use any form of the verb “be” in the body of this article (except in quotations), I did not find it much harder to compose than other pieces I have written.

IV. What should we do?

On the one hand, both Lojban and E-prime have a strong intuitive appeal on Extropian grounds. The languages represent a deliberate attempt to design a system of communication both logically consistent and—more importantly—potentially capable of improving the rational powers of their users. Since the transhumanist program involves the shedding of old, arbitrary customs and habits in favor of well-thought-out plans of action, combined with active self-enhancement, Lojban and E-prime look like the “languages of choice” for Extropian communication.

On the other hand, artificial languages—especially Lojban—have their origin in philosophical and empirical premises that appear to run counter to extropian values, and perhaps to plain facts. The Sapir-Whorf hypothesis represents an explicit concession to limits on our ability to think, a concession that stands in direct opposition to the principle of Boundless Expansion. Even if such limits exist, must we attribute them to the specific language that we speak? It seems entirely plausible that our ability to reason and our ability to speak come from our genetic endowment, part of which generates the “universal grammar” that forms the basis for all human languages. Under this view of language—promoted by the linguist Noam Chomsky—learning a given language involves the setting of parameters hard-wired into the human species. Of course, this particular form of biological determinism does not rule out the possibility that a given language may influence the way its speakers think, but it does not rule out the “null hypothesis” that your language does not shape your thinking.

In fact, it seems more intuitive to accept the opposite point of view; namely, that your thinking patterns—as determined, for example, by your physical environment—shape your language. For example, most Indo-European languages have two or three genders for nouns, which correlate to some extent with “male”, “female”, and “other”, but some languages spoken in Africa have many more genders, which distinguish among various classes of animate and inanimate things found there. The genders of these languages often embrace categories like “poisonous plant”, suggesting that grammatical systems derive at least in part from the way people think about their environment.

In these sorts of “chicken-or-egg” problems the truth may lie in some uninteresting compromise between the two extremes, or in a new approach. A very exciting possibility for the latter comes from dynamical systems theory, in which circular causality falls out naturally from the underlying assumptions, rather than posing the paradox it has for traditional scientific thought. In other words, our language shapes our thinking, which shapes our language, and so on. This view does not require a Lamarckian version of language evolution; we don’t want to claim that our language or thinking directly affects our genetic endowment. Rather, we acknowledge that given a particular genetic arrangement, we obtain a system in which language and thought can have mutual influence.

This possibility brings out a greater truth about artificial languages. Natural languages serve as an example par excellence of spontaneous orders, so it could very well end up that any artificial language will mutate into a natural language—warts and all—given enough speakers and enough time. Despite the logical unattractiveness of exceptions and ambiguities, they seem to have tremendous appeal to people at some deep level. We might, for example, find future speakers of E-prime committing the sin of the “be of identification” without using that verb or presupposition, as speakers of Indonesian, Hebrew, and many other languages do now: “John a liar.”

This possibility does not bode well for the future of the species, when we consider the thinking of B.L. Whorf himself:

We cannot but suppose that the future developments of thinking are of primary importance to the human species. They may even determine the duration of human existence on the planet earth or in the universe. [1956, p.83]

At this point, many readers—especially those who have followed the Extropian movement for any length of time—will recognize a different possibility. If, as Chomsky argues, the physical makeup of our brains determines the possible form of our language, it seems obvious that we should seek ways to modify our brains so as to lessen the chances for

destructive irrationality. Such efforts, though clearly beyond the scope of this article, have received a good deal of coverage in this journal, and *Extropy* readers can expect to see much writing on this topic in future issues.

Acknowledgments

The author wishes to thank Bob Le Chevalier (lojbab) and the Logical Language Group for making available a large amount of Lojban documentation, and ExI Director Russell Whitaker for contact information and enthusiasm about Lojban. You can reach the Logical Language group at 2904 Beau Lane, Fairfax, VA 22031, or by calling (703) 385-0273. Anyone interested E-prime and general semantics can subscribe to *Etc.*, the journal of the International Society for General Semantics, P.O. Box 2469, San Francisco, CA 94126.

References

- Bourland, D.B. Jr. and Paul D. Johnston, Eds. (1991) *To Be or Not: An E-Prime Anthology*. San Francisco: International Society for General Semantics.
- Brown, J.C. (1960) Loglan. *Scientific American* 202,6. 53-63.
- Chomsky, N. (1965) *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Dallmann, W. (1992) “Is Is Not Is Is Not Is” *Etc.: A Review of General Semantics* 49, 2. San Francisco: International Society for General Semantics.
- Lakoff, G. (1990) *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. Chicago: University of Chicago Press.
- Popper, K. (1968) *Conjectures & Refutations: The Growth of Scientific Knowledge*. New York: Harper & Row.
- Whorf, B.L. (1956). *Language, Thought, and Reality: Selected Writings of Benjamin Lee Whorf*. Cambridge, MA: MIT Press.

Notes

¹and my own personal favorite, Volapük

²This arrangement would not appear so strange to speakers of Japanese, a language in which adjectives can behave like verbs.

³According to lojbab, Lojban arose as dialect of Loglan, because Dr. Brown maintains a copyright on the original language.

⁴American-educated readers may notice that I put the comma outside the quotation marks in this list, and in the rest of the article, breaking with traditional prescribed style in this country. Inspired by my study of Lojban and E-prime, I have decided to take a more pro-active attitude toward my writing habits, and therefore use the more sensible practice favored in Great Britain.

⁵I have also decided, after many years of personal linguistic conservatism, to adopt the common usage of the pronoun “they” to mean “he or she” (and hence “their” for “his or hers” and “them” for “him or her.”) I thank Extropian Rob Michels for pointing out the sensibility of this practice to me. Those who object to it must also have difficulty with French, German, and other languages where an originally plural pronoun became both plural and singular.

The Open Society and its Media

Mark S. Miller, E. Dean Tribble, Ravi Pandya, Marc Stiegler

Electronic media present tremendous opportunities for improving the nature of society. I'll first talk about how discourse affects society, and how changes in media may improve societal discourse. Then I'll describe the Xanadu system and how it was built to achieve these goals.

Improving Society

Improving society is a difficult task. More generally, improving complex systems is a difficult task. Can't figure out which way is up? See if you can figure out which way is down. Engelbart, back in the early 60s, wanted to explain to people why interactive systems would make a significant difference to their lives, and to their ability to express ideas. In Figure 1, the origin on the axis is what people were doing at the time—writing with pencil and paper. When he found himself unable to communicate to people how much better things could be, he contrasted their current experiences with how much *worse* things could be. He tied a pencil to a brick, handed it to people and said, "Okay, now write." People found it very difficult. The unwieldy nature of the tool interfered with their ability to express ideas. With the pencil and brick for contrast, he effectively asked two questions: "What made the difference?" and, "How can we move further in the other direction?" [Engelbart]




Figure 1: Engelbart's Pencil and Brick Experiment

This experiment showed people how important their tools and their media were to their effectiveness, and helped them start to see the next brick to remove.

Karl Marx just performed a similar experiment on society over the course of most of this century. The origin on Figure 2 represents where we are now. Karl Marx tied a very large brick to a very large pencil and the last few years have revealed the result to be far worse than the even his harshest critics imagined [Popper50]. What made the difference between the societies? Two important elements were open markets and open media. How can we move further in the other direction? In this

talk, I'll be addressing the nature of open media, how they differ from closed media, and how social hypertext systems can enhance the advantages of those media. Applying information technologies to the further opening of markets is left as a mission for the reader.

Media Matter

Media matter because it is in media that the knowledge of society evolves. The health of the process by which that knowledge evolves is critical to the way society changes. Karl Popper, the epistemologist, had the insight that knowledge evolves by a process of variation, replication, and selection, much as biology does. Variation of knowledge is what we call conjecture—hypothesis formation, tossing new ideas out there. Replication of knowledge is the spread of ideas through publication and conversation. Selection of knowledge is the discrediting of conjectures through the process of criticism¹. The ability of our knowledge to progress over time depends on an ongoing process of criticism, and criticism of criticism. The ideas that survive the critical process tend, in general, to be better than those that don't.

In closed societies, when arguments can't be spoken, hard truths cannot be figured out.




Figure 2: Marx's Pencil and Brick Experiment

To appear in *Prospects in Nanotechnology: Towards Molecular Manufacturing*, edited by Markus Krummenacker and James B. Lewis. Copyright © 1994 by John Wiley & Sons, Inc. Reprinted by permission of the publisher, John Wiley & Sons, Inc. To order a copy of this work, call 1-800-CALL-WILEY

When people can't openly criticize, can't openly defend against criticism, or can't openly propose ideas that conflict with the official truths, then they are left with mistrust and cynicism as their only defense. This leads to the simple heuristic of assuming the official truth is always wrong. For example, because *science* was promoted by the Soviet propaganda machine, pseudo-science is on the rise in Russia. Because anti-Nazism was promoted by the East German propaganda machine, Neo-Nazism is on the rise in East Germany. The official truth is neither always right nor always wrong; society needs a more sophisticated process for judging claims.

Our society does have open media. Are we in the best of all possible worlds? Is our media good enough? Can they be made significantly better? Among our media TV is so bad it's a joke; only slogan-sized ideas can be expressed. We prize the quality of discourse in our books and journals, but critical discussions in them are only loosely connected; starting from the expression of an idea, it is hard to find articles that criticize that idea. When arguments cannot be found and navigated, the next harder truths *still* cannot be figured out.

Xanadu

I rejoined Xanadu in 1988 largely because of fear about the dangers of nanotechnology, coupled with incredible excitement about the promises of nanotechnology. In looking at the dangers, I saw that none of us individually are clever enough to figure out how to solve those problems. The only hope that I saw four years ago—I no longer believe it is the only hope—is that by creating better media for the process of societal discourse and societal decision-making, we stand a much better chance of surviving the dangers posed by new technologies, so that we may to live to enjoy their benefits.

I'm about to talk through the elements of the hypertext system we built. Xanadu has frequently been called Golden Vaporware, and many people have wondered whether this is a never ending project. One of the things I want to emphasize when I go through all of these features is that I'm only referring to the features that are now running in the software. There are other features that we planned on and anticipate, some of which will be mentioned in the *future plans* slide, but the body of this talk will only cover what is implemented and running.

First I will discuss the four fundamental features—links, transclusion, versioning, and detectors—Marc Stiegler will then present an example using them. Then I will describe the remaining four features—permissions, reputation-based filtering, multimedia, and external transclusion; followed by some concluding remarks.

Links

Hypertext links are directly inspired by literary practice. Literature has many different kinds of links connecting documents into a vast web. Textual examples of these links include bibliographic references, marginal notes, quotation, footnotes, and Post-it notes.

We propose to build engines of citation, so that people can navigate this vast web of literature at the click of a mouse. Most computer text systems are predicated on a misconception: that the meaning of a document is represented purely or primarily by its content. Documents are not islands. Conventional computer text systems put their effort into the appearance of individual documents. My experience in reading documents, especially reading a literature that I'm not familiar with, is that it's difficult to understand documents without their context. A context that helps answer questions such as, "What were the ongoing controversies that the author had in mind?" "What views was he supporting or attacking?" "What attacks was he guarding against?" We need to understand this whole web of connections in order to understand the documents that we're reading. The Xanadu system is built to provide as much support for this contextual information as for content.

With the ability to follow the links in this vast web of documents, isn't it easy to get lost? How does one stay oriented? One answer to these questions is *guides*, a new kind of document that provides an orienting view together with links into the existing literature. I expect guides to come largely from people making their own organizing views of a literature and then cleaning them up for publication, so others may benefit from their work.




Figure 3: Hyperlinks

Hyperlinks

Because "nanotechnology" is now used by many to mean any technology approaching the nanometer scale, we have had to retreat to the term "molecular nanotechnology." Hypertext terminology has gone through a drift similar to nanotech terminology. The Xanadu project is the one that coined the term "hypertext" and originated the notion of the hypertext "link." However, because the term *link* has come to be viewed as something much less capable than what we meant by it, we're now calling it the *hyperlink*. The distinction between the link and the hyperlink is crucial

for supporting active criticism in an open media.

Hyperlinks are fine-grained, bi-directional, and extrinsic. Frequently, an argument is not with a document or chapter as a whole, it's with a particular point that somebody made at a particular place in their text. For example, somebody refers to the fourth law of thermodynamics, and someone else writes a criticism saying there is no fourth law of thermodynamics, linking it to the original. The fine-grained property allows the link to designate the particular piece of text one is taking issue with. Bi-directionality enables readers of the original document to find the criticism, enabling them to exercise fine-grained skepticism, to constantly ask themselves, "What's the best argument against the thing I'm reading *right now*?" and then, "What's the best argument against that, in turn?" Links provided by other hypertext systems have generally been only in the forward direction, enabling a reader to find those documents *referenced* by a given document. However, to find criticism, the reader must find the documents that *refer to* the document they are reading.

Extrinsic linking is the ability to link into a document without editing it. Several other systems support the creation of links that are fine-grained at both ends, but these others do so only by modifying both source and *target* documents². Critics normally will not have the ability to modify the documents they are criticizing. They could spin off their own version into which they attach these links, but then other readers *still* can't find these criticisms from the original documents.

Part of what we mean by an open media is that everyone who is connected to the system can read what they're permitted to read, can write new things, and can make them accessible for others to read. This includes making links to anything that they've read so that anyone else who reads the original can find the material that has been linked to it. All readers of the system are potential authors. We can think of this process as *active reading*. Frequently people make marginal notes to themselves. This is a medium in which readers can share such things with each other. When much writing is commentary about other text, the commented-on text is the best rendezvous point for the authors and readers of commentary to find each other.

Emergent Properties

This kind of accessible criticism can provide a decentralized consumer reports. When people post documents on the system that are either products or descriptions of products, customers of those products can post criticisms of them. What did they think of using them? This commentary can guide the purchasing decisions of others³.

There's a particular capability we're used

to in conversation that is almost impossible to successfully attain using paper-based literature, which is hearing the absence of a good response to an argument. A reader can not only see what the most compelling arguments are against some statement; they can also see when there aren't any, or when all the seemingly compelling arguments have been successfully refuted. Such absences are quite obvious in conversation. Electronic media can make these absences obvious as well, but in a context where the absence will be much more telling because the missing argument could have come from a much larger audience over a more extended period of time.

Other hypertext systems with their uni-directional links reproduce the asymmetry present in our paper-based media—it's much easier to find something that a document cites than it is to find those documents that cite a given document. One of the effects of this asymmetry in paper media is the pathological division of scholarly fields into disjoint "schools." Instead of healthy intellectual engagement, debate, and cross-fertilization of ideas, we see a process of increasing inability to communicate between schools, and more preaching to the converted within a school. The terrible irony of attempting scholarship with uni-directional links is that *the* very attempt to engage in healthy debate across schools accelerates the pathological division process. How does this occur?

Let us consider two schools within a discipline. Generally students within a school see the documents supporting the positions of that school. The students also see criticisms of documents in the other school. Intellectually eager and honest students, seeking to know both sides, will occasionally follow these criticism links forward. The result is that they will see the parts of the other school's literature that is most soundly criticized by their own school, immunizing them more and more against the foreign ideas. With bi-directional links, they can also find the greatest challenges to their own school. Bi-directional links let them also find the most telling criticisms of the ideas they are inclined to accept.

Transclusion

Before there were modern economies, there were many little villages, each with their own little manufacturers having to go through

a large amount of the production process themselves. These economies were, therefore, much less productive. An individual baker or shoemaker, for example, would reproduce the same kind of work that was being reproduced in many other villages and would have to

span in an arrangement. Therefore, when somebody writes a criticism of content as it appears in one arrangement, that criticism is visible for the same content as it appears in all other arrangements, including arrangements that were made before the criticism was at-

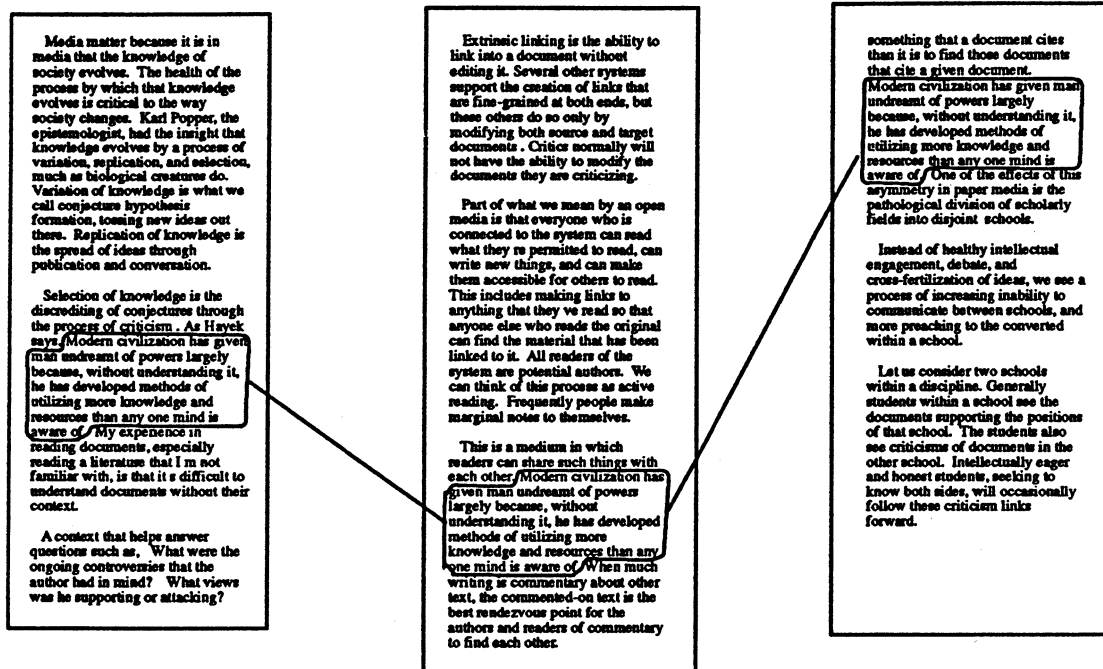


Figure 4: Transclusion - Different Arrangements of Shared Contents

fashion a shoe from, not quite raw materials, but without intermediate goods. In extended economies, people can build on one another's work, and there can be a finer grained division of labor and knowledge, with better specialization.

Now, with respect to literature, authors are frequently faced with the task of re-explaining and restating background material that has been explained well elsewhere. If you could just borrow that material, those existing good explanations, and incorporate them (with automatic credit where due), your efforts could be spent stating what's new. We introduce the concept of transclusion to separate the arrangement of a document from its content. There is an underlying shared pool of contents, and all documents are just arrangements of pieces from that pool. In Figure 4, the three circled appearances of the same text are actually just one piece of text in the underlying shared pool of contents, and it just happens to appear in three different arrangements which constitute three different documents. We refer to the three documents as transcluding that piece of text. The separation of content and arrangement also leads to good support for incremental editing. Different versions of a document are just different arrangements of mostly shared content.

This is more than just a hack to avoid the storage cost of making separate copies; hyperlinks are linked to the content, not to a

tached. The normal incremental editing process of a single document is analogous to evolution by point mutation. The ability to transclude text from other documents allows the analogue of sexual recombination. Were links visible only from the arrangement into which they were made, the variation processes would destroy selection pressures by leaving criticisms behind.

Remembering the Past: Historical Trails

As you are editing, an historical trail gets left behind—bread crumbs in history space. The historical trail is simply a sequential arrangement of the successive arrangements of contents. This is yet another kind of context important for understanding. "Things are the way they are because they got that way." [Weinberg] Understanding how they got that way often aids our understanding of what they are.

Preparing for the Future: Detectors

Besides looking into the past, one also reads a literature knowing it will be changing. How can one keep up? To keep track of what's happening, to keep up with changes, we introduce detectors. One can post a revision detector to find out when things are edited, when new versions of something appear, and then one can use version compare to find out how

they're different. With version compare one can engage in differential reading—reading just the differences between the current version and the version one has most recently read.

Link detectors are a way of finding out when new links are made to existing material. Let's say that you published something and you want to find out when others post comments on it. You'd like to be informed of comments, but you don't want to have to go back and constantly recheck all the things that you've written, so you post a link detector on all the things that you've written as well as on other documents you're interested in seeing further comments on. You want to see what people will say about them. As new comments are posted on those documents, you are continually informed.

E-mail is just the special case where you establish a canonical point in the literature, for each person, a place others link to in order to send that person a message. That person simply has a link detector there saying, "Show me all new things that are attached to here." This generalizes to treating any shared point of interest in the literature as in some sense, a mailbox, or a meeting room for further conversation or conferencing about a topic. Canonical documents become meeting places. Should two disjoint discussions about the same topic spontaneously form in two places, anyone who notices can just make a link between them. The link detectors of each community will then inform them of the existence of the other.

At this point I'll shift over to Mark Stiegler and Dean Tribble who will demonstrate, with the Xanadu software, an example involving exactly the elements so far discussed.

The WidgetPerfect Saga

This is a true story about how a hypertext system was able to save several thousand jobs. There is one special characteristic about this true story. It is a true story from the year 1995. It is a story about one of the events that took place at the company, I'm sure most of you have heard of it, called WidgetPerfect. WidgetPerfect is the second largest manufacturer of widgets in the world, second only to their big competitor, Microwidget. The guys at WidgetPerfect in the year 1995 had identified a really significant opportunity in the upcoming expanding environment of widget components technology.

They were developing the world's first fully modular widget. They had a team working on it. Dan was in charge of the preparation of the marketing materials for the modular widget. Ruth was in charge of the technical work team and John was in charge of the budget and finance and all the costing. At this point, the modular widget was in prototype stage when a very unfortunate thing happened. Microwidget, the big competitor, came out

with a partially modular widget hitting the marketplace first with an inferior product. It was technically inferior, but nonetheless it was in the marketplace first.

Dan was examining this Microwidget partially modular widget and it was overall inferior, but nonetheless it had one really striking improved feature. It had a funculator

was reaching for the telephone to call Ruth when Boeing, the largest purchaser of widgets in the world, called him up about a \$15 million widget order. He got distracted with this purchase and he never quite got around to calling Ruth.

We have good news.

Ruth, knowing that the success of her

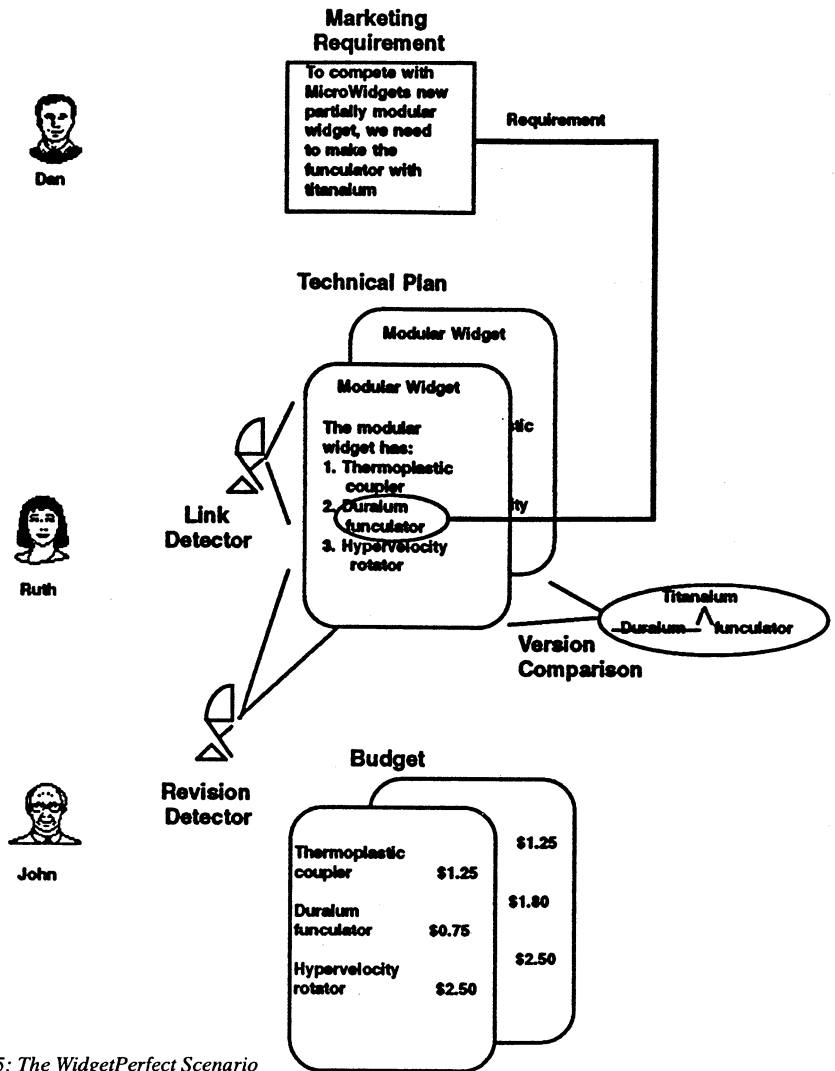


Figure 5: The WidgetPerfect Scenario

made out of titanalum, whereas the fully modular widget that was being developed by Ruth only had a duralum funculator. This was an important improvement for certain key market sectors. Even though the partially modular widget did not have anything comparable to a thermoplastic coupler or a hyper-velocity rotator, they had to make a change.

So Dan created a new document in the marketing requirements describing this titanalum funculator and he attached a link to the part of the technical plan that specifically referred to the duralum funculator that was in the current plan. He made that a new requirement.

Now Dan knew that in order to get anything to happen with improving the widget prototype, he would have to talk to Ruth and he

technical design depended on her being able to respond promptly to new requirements, had attached a link detector to her technical plan. This link detector would be constantly watching for new links of the link-type requirement to be attached. When Dan had attached the new requirement to the duralum funculator, Ruth's link detector went off. Ruth was alerted. She followed the link detector out to the link, followed the link back to the new requirement, saw what the change required was, and modified the technical plan to reflect the use of a titanalum funculator. Well, this is all very fine except for an additional problem which is, as I think everyone here knows, that titanalum is considerably more expensive than duralum, and so this had some significant effect on the manufacturing cost. Well, Ruth

knew that this was going to have an impact on the budget, and she was reaching for the telephone to call John when smoke started billowing from the laboratory where the prototype of the modular widget was being manufactured. She ran off to deal with the emergency and she never quite got around to calling John.

We have good news.

John, knowing the success of his budget was completely dependent on his responding to modifications to the technical plan, had attached a revision detector to the technical plan and this detector was constantly watching for updates. So when the technical plan was indeed updated, John's revision detector went off. He followed the revision detector up to the technical plan, used the hypertextual version compare capabilities based on the transclusion relations, to compare the new version of the plan to the old, found that the change was that duralum had been deleted and titanalum had been put in, and then he went back into the budget and updated the budget documents to reflect the increased costs due to the use of titanalum.

As a consequence of this, the modular widget program was completed on time with a fully adequate specification. It was a completely superior product. It blew Microwidget off the face of the earth. As a consequence, thousands of jobs at WidgetPerfect were saved.

At this point there was a video demonstration of Dean Tribble walking the Xanadu software through the WidgetPerfect scenario.

Permissions

A social system is, to a large extent, a system of rights and responsibilities. Xanadu has an extensive permission system called the club system, intended to deal with some of these issues. In Figure 6 we have a document which Bob can edit. Bob has sent it as a mail message to various people in a blind carbon copy ("bcc") relationship. Alice and Chuck are both members of the bcc club of people who have permission to read this document. Bob, though, is the only member who can read or edit the bcc club. If this were a cc list, Bob would still be the only person who could edit it but it would be self-reading. Everybody who was a member of such a cc club could see who else was a member of that same club.

This demonstrates a principled answer to permissions meta-issues: one can distinguish between who can read a document, who can read the list of people who can read a document, who can read that list in turn, out to any

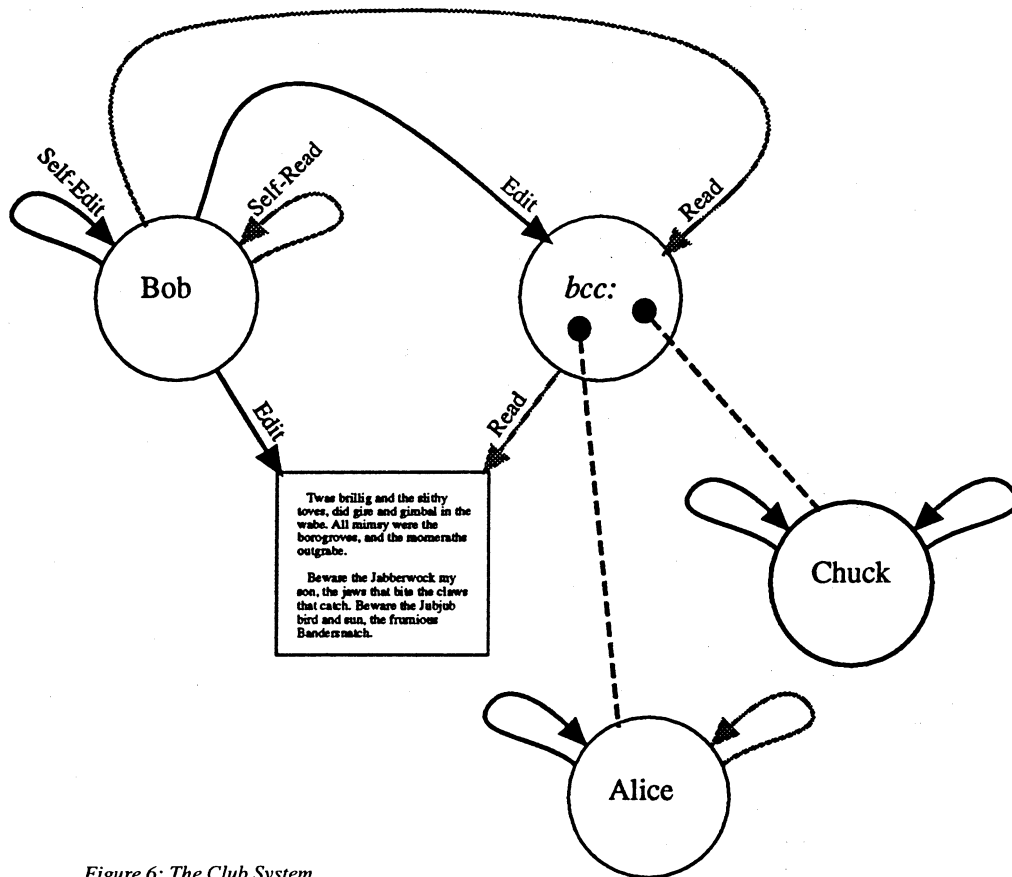


Figure 6: The Club System

desired degree of distinction (and similarly for the editing dimension). However, infinite regress and needless complexity are avoided by using clubs that are self-reading and/or self-editing whenever further distinction is currently not necessary. Should such distinction later become necessary, it can always be introduced by someone with appropriate edit permission to the club in question. Users only grow meta-levels on an as-needed basis.

Our permission system also supports the notion of accountability. All actions in the system are taken by someone. When you look at information in the system, you see some identity attached to them. There are no official truths; there is only who said what, and the structure of the system reflects that.

Reputation-based filtering

One of the potential pitfalls of an open hypertext system is the junk problem. The ability to find good commentary and criticism will be especially important when reading very important documents, but it is precisely on these documents that one expects to be inundated with tons of worthless or irrelevant links. Without a filtering mechanism, it would be on exactly the documents for which one most needs good commentary that the provision of commentary would be most useless. For example, imagine how many links there would be onto the First Amendment to the

Constitution.

Links can be endorsed as worth reading by various readers; however, no one may endorse with the identity of another. Different endorsers will establish varying reputations with different readers, much as with movie reviewers. Readers can filter their view of links into a document both by who endorsed as well as by link-type. When even this mechanism gives too coarse an answer, one can rely on documents such as a hypothetical Guide to the Citations to the Bill of Rights endorsed by a reputable publishing house. This very same link filtering ability is also what allows one to find such guides in the presence of a swamp of links.

Hypertext + Multimedia = Hypermedia

Increasingly ideas are being expressed in media other than text, and increasingly computers are used to handle these other media. We usually refer to *hypertext* because text is the most important case and the clearest example. But nothing I have presented, none of the things you have seen the system do is in any way specific to text or even to media that have linear flow to them. It all applies equally well to a variety of other media, such as sound, engineering drawings, postscript images, and compressed video. In all cases, one can make fine-grained links, edits, transclusions, and

version compares (even if the data is block-compressed or block-encrypted). Although the implementation has some optimizations targeted at text, in no way does the *architecture* make any special cases for text. Documents can, of course, be composite arrangements in which several media are mixed together.

External Transclusion

No software system is an island. We don't imagine that once the product is available everyone will instantly take all information they want access to and transfer it into Xanadu. We have to coexist with lots of other systems for many good reasons.

We handle that with *external transclusion*. Our documents are able to transclude into arrangements that are within the system. These, in turn, are able to represent transclusions of materials that are stored elsewhere. By perceiving other systems through the window of Xanadu, you can see those other systems as if all those documents were within the Xanadu system. Through Xanadu, I could follow a link from a WAIS document into a Lexis document even though neither system has any notion that such a link even exists. It is not just that the Xanadu system is not an island, that we have to coexist with everything else, it is that through Xanadu *those* systems are able to coexist with each other in a way they are unable to now, making *them* into non-islands.

Conclusions

When we started building the system, we were thinking purely in terms of paper based literature—of writing. What we've built is something that has many of the best aspects of both writing and conversation. Many of the aspects of each are complimentary, many of virtues of conversation make up for flaws in writing and vice versa. We found ourselves building a system that supports the dynamic give and take of conversation and the persistence and thoughtfulness of literature. [See Tanoe 1]

Our status is that we currently have a working portable server. It has some bugs in it, including some performance bugs, but we are working on it. However, all the features that I talked about so far work. We are continuing ahead with the effort on both the server and the front end. The front end is in a preliminary stage; we consider it adequate to show that the server is real, and to exercise its features. We plan to do a much better front end. The protocol between the front end and the server is very stable, and has been stable for a long time now. Our plans are to get investors and to finish both the front end and the server. The target for our first product is small-to-medium-size workgroups within com-

Writing

**Persistent
Cumulative
Large expressions
Revision before publication
Large audiences
Freedom of entry
Reputation-based filtering**

Conversation

**Visibility of arguments
Can "hear an absence"
Small expressions
Revision after publication
Small audiences
Fast publication
Feedback from audience**

Table 1: Combining the Best of Writing and Conversation

panies that have a large body of documents that they need to be managing and evolving.

There is one major feature that our first product lacks. We provide hypertext because documents are not islands. We make the system inter-personal because people are not islands. We provide for the transparent windowing into other systems because no product is an island. However, for the moment each server is still an island with respect to the other servers, and so each workgroup is also an island. We have architected the system so that, soon after first product, we will be able to weave all the servers together into a transparent distributed system. When you follow a link from one document to another, if the other document isn't here but in some server in Tokyo, it will be transparently fetched for you and the only thing you'll notice is that following that link took longer.

For any media to radically improve the process of opinion formation in society, we believe it *needs* features equivalent to fine-grained, bi-directional, extrinsic, filtered links. These links must not get lost when the documents they are attached to change. Issues of authority, privacy, and responsibility must be handled in a robust and secure fashion. Open entry of readers and editors is crucial for open discussion. Open entry of server providers is less obvious, but equally important, in order to make centralized control impossible. We will be providing support for people who want to do on-line services based on our software. All of this is necessary to achieve our open electronic publishing dream. In so doing, we hope to improve the quality of public debate, in order to obtain the benefits of the open society yet again.

Acknowledgments

We thank the whole extended Xanadu team for having struggled together for many years on a project that has been at least as much a cause as a business. We thank Eric Drexler for exploring the relationship of hypertext publishing to evolutionary epistemology [Drexler], thereby inspiring this talk. We thank Anita Shreve for extensive help editing this talk into a paper.

Notes

¹Karl Popper originally proposed that selection proceeds by a process of refutation [Popper59]. His student William Bartley generalized this to criticism [Bartley].

²World Wide Web *anchors*; Microsoft Word *bookmarks*; Lotus Notes and Folio Views *Popup text*.

³The use of bi-directional links for a decentralized consumer reports is already happening on the American Information Exchange.

Bibliography

Douglas C. Engelbart, "Augmenting Human Intellect: A Conceptual Framework," *SRI Project No. 3578*, October 1962.

Sir Karl R. Popper, *The Open Society and its Enemies*, Princeton University Press, 1950

Sir Karl R. Popper, *The Logic of Scientific Discovery*, Harper & Row, 1959


William W. Bartley, III, *The Retreat to Commitment*, Open Court Publishing, 1962

Gerald M. Weinberg, *The Secrets of Consulting*, Dorset House Publishing, 1985

K. Eric Drexler, "Hypertext Publishing and the Evolution of Knowledge," *Social Intelligence*, Volume 1, Number 2, 1991

**To order a copy of this work,
call 1-800-CALL-WILEY**

PROSPECTS IN NANOTECHNOLOGY



Toward
Molecular
Manufacturing

Editors:
Markus Krummenacker
James B. Lewis

Foreword by K. Eric Drexler

God and Man at Yale

A Conversation With Dave Ross

by David Krieger with Gayle Pergamit

Dave Ross founds companies for a living. He co-founded Palantir Corporation, which became Calera Recognition Systems, the longtime leader in optical character recognition software. He also founded Arkenstone, a non-profit corporation which makes reading machines for the blind, and his current venture, RAF Technologies. Previously, at NASA's Jet Propulsion Laboratory, he formulated the "prime rib" technique of selecting orbital rendezvous trajectories, still the standard method for planning solar system missions. In 1992, Dave presented his talk "Seven Paths to Immortality" to the annual Eris Conference in Aspen, Colorado, in which he talked about longevity technologies, uploading, cryonics, and how he views these technologies from a Christian perspective. Later in 1992, I interviewed Dave at the home of Gayle Pergamit, co-author of Unbounding the Future and co-founder, with her husband, economist Phil Salin, of the American Information Exchange online information marketplace.

I guess we could start with the question "How did you get this way, and what are you planning to do about it, if anything?"

Reading way too much Robert Heinlein as a teenager probably started it... like it did for almost all of us, I think. It gave me a healthy dose of libertarianism. Growing up... my father was a bootlegger, among other things, during Prohibition, so I have, shall we say, a little bit of anarchy in my blood. So that probably is the origin of some of it.

Gayle: Was everybody's father involved in anarchy?

Certainly my children's father...never mind. I got interested, I guess, because of an overdose of science fiction in the "space biz," and sort of followed the traditional lines that all of us, I guess, of Gayle's and my generation did, which was: space in the Seventies and Eighties, computers and nanotechnology in the Nineties. That's how I met most of the people in the "Palo Alto School of Applied Austrian Economics," through the space biz. I had an undergraduate degree in physics—at which point I spent four years in the Navy as a legal officer, which made as much sense as anything else in the Navy did—and then went to graduate school in aerospace engineering. From there I went to Jet Propulsion Lab, and from there I was hired by Gary Hudson to do rockets. I had connected up with Gayle at a space colonization class; I worked for Gerry O'Neill¹ in the summer of '77 doing a NASA-Ames space study. My doctoral work is in orbital mechanics; I was doing the orbital mechanics for asteroid retrieval for the '77 summer study. Went to Jet Propulsion Labs

and did pretty much the same thing, mission design and mission analysis stuff. That's where Gary Hudson showed up while I was inebriated at a Christmas party, proceeded to interview me, and the rest was history—I went to work doing rockets.

Gayle: Actually, the tradition of filling someone full of alcohol and interviewing them for a job—they wake up, employed and on the high seas—

dV/dt: That's a long and venerable tradition.

I believe the term is shanghai'd, yes. He didn't even have to administer the alcohol. I worked for a while for Hudson on the Percheron rocket; when I later left, [Hudson] blew up. I was in charge of the test facility and the guy that was second was Jim Fruchterman. He and I and another guy, Eric Hannah, with whom I had written papers on space settlement, started what is today Calera Recognition Systems (it used to be Palantir Corporation), doing optical character recognition, which seems like a long step from rockets. It's not quite as long a step as it sounds, because my degrees are really in applied math and algorithm design; I did that stuff at JPL, and algorithms for space flight and algorithms for character recognition, though different, are not different in kind, just different in degree.

So we did that, and I was there for a long time. I got more and more interested in computer science—that was starting around 1981, '82—I hit the computer wave at that point and got to thinking about machine intelligence because we were making algorithms to do things that I do not consider artificial intelligence but that traditionally have been

considered "intelligent acts"—recognition. That got me to thinking more and more about "What are the potentials for, and the problems with, traditional concepts of artificial intelligence?" That, and the evolution of this group of people, thinking more about manipulating things on very small scales, as with nanotechnology, and the potential of an unbounded future for humanity off the planet from the rocket phase. Both fields are unbounded, from the point of view of what you can do with limited resources and the possibility that your resources may be almost limitless. This got me thinking about future things more and more, and coming to the conclusion, after thinking about it, that though the problems of making machines think are very, very hard, they're not impossible, and that we are liable, if current trends continue, to see within even our "normal" lifespans, our own creations beginning to rival us and in some danger of surpassing us in intelligence, and this got me increasingly to thinking, "Well, that's very nice, what do we do about it?"

And, of course, I read all the "usual suspect" books: *Engines of Creation*, *Unbounding the Future*, Hans Moravec's *Mind Children*, and then a lot of the other things: *The Emperor's New Mind* by Penrose²; *Gödel, Escher, Bach*³; and a bunch of other stuff. Plus, working in what could be called "artificial intelligence" led me to give serious consideration to which way I thought things were going to go, in terms of machine and human intelligence and machine-human alliances and analogues of parasitism and symbiosis and so on. That sort of got me interested in what has come to be called "extropian" ideas. I actually made contact with the Extropians through two almost unrelated channels: first, Jim Bennett said I ought to be on the Extropians [e-mail] list; then just before I got on the Extropians list I ran into Max More at a talk I gave. He asked me to write an article for *Extropy*, which became the article on human uploading.⁴ So, that's the circuitous path by which I got here.

Was it at the ERIS gathering that you ran into Max?

No, it was at the Nock Forum in Los Angeles.

Before we go on to the farther-out Extropian things, talk a little bit about the "prime rib"

curves.

They derive from stuff that I worked on as part of my doctoral dissertation, which I defended very nicely and never turned in, that were essentially optimal ways of moving around in the solar system. One of the techniques that I developed was an analytic way of determining the optimal orbits: if you're in orbit around one body in the solar system and you want to go to some orbit around another body in the solar system, you want to do that with a two-impulse transfer—you burn at the beginning, you burn at the end. It turns out that there's a very nice analytic formula that just gives you the result; it's very, very nice, and I worked that out, extending some other people's work, as part of my dissertation. That problem, and an orbit around and trajectory fly-by, turns out to be an interesting problem for NASA. I went to work for Jet Propulsion Labs—my thesis advisor informed me I was going to go work for the Advanced Projects Group at JPL—so I said fine, all right—

"Twist my arm."

Twist my arm! I and a man named Dr. Neal Hulkower, who is and remains a good friend of mine, were the co-authors of several papers on the subject of prime rib curves.⁵ What prime rib curves are: NASA had a traditional method of determining possible orbits between two bodies, which were called pork chop curves. The reason they're called pork chop curves is they plot out shapes in space that actually look kind of like pork chops; they have sort of a rounded area and a tail off at one end. What we did was to plot vertically the position around the sun of one of the bodies and horizontally the position around the sun of the other body, in this particular case the Earth (it doesn't have to be the Earth), so you'd have Earth horizontally and an asteroid vertically, then do a plot at each point of the sum of the two delta-vees, for example, that you'd get from my analytic formulas. These give you curves, and the interesting thing about them is—probably as an artifact of the plotter we used—they looked like cuts through a piece of prime rib; there are the ribs and then these contours around them going out and out and out, and that's why we called them prime rib curves. What's nice about them is they will tell you the best possible trajectory between these two bodies, which means "Don't bother if it's not good enough."

Right.

So this is used for questions like, "I want to do a fly-by of an asteroid, I only have a certain amount of propulsion mass that's available in my rocket, is this particular rendezvous possible?" And it's easy to run; these things run in no time, so you make some of these plots—you make hundreds of these plots and look at the results, and they'll tell you what the best possible trajectories are. The nice thing



is that, because you know the Earth's location, you know what day you have to take off on, because it's the day the Earth is in that location. That will then tell you where the thing must be in its orbit to arrive there. It doesn't tell you whether or not the body is in that position after the number of days of flight, but it will give you the number of days of flight and will tell you where the body has to be. What you then do is to look at each year on that day how close the target is to that location, and then go back and use more traditional methods for saying "Ok, what's the real optimum around here", and of course it'll always be worse. The interesting thing is you can usually find one within a twenty- or thirty year launch date that is within a few percent of the optimum. So that was prime rib curves.

Which came first, immortalism or libertarianism, in your case?

Libertarianism clearly did. It hit in probably early high school. I didn't, of course, have a name for it then. I guess I didn't start calling myself a libertarian until college, because, you know, I graduated from college in '71; what was a "libertarian"? A friend of mine once introduced me to somebody else as a libertarian, and suddenly it dawned on me that, yes, that was true, but I had never considered myself that; I had always called myself a conservative. But I realized that I really was a libertarian and not a conservative.

Were you ever a "Randroid"?

No, I was never a Randroid; I was fortunate enough to escape Rand until I was too old to be infected that badly [laughs]. I mean, I've read almost all of Rand, and was old enough and mature enough in my thought to pick and choose.

You mentioned Heinlein, what other influences brought you to libertarianism?

Natural cussedness, family background. My mother had a friend who would come over during school for lunch—I lived right near the school, so I would go home for lunch—and she would often be over there, a good friend of my mom—who was a *rabid* conservative, but of a very decidedly libertarian and anti-government streak, and she sat there for *years* indoctrinating me [laughs]. So that was a good part of it.

Another part of it was simply a desire to myself be left alone by society, and the willingness to say, "I want this for myself; I should want it for everyone else." I didn't come at it from an economic perspective. I was in college from '67 to '71 as an undergraduate; I was at Yale, and there during the May Day riots and the Black Panther trial and all that stuff. The two largest organizations on campus, and they were just about the same size, were the SDS on the one hand and the Party of the Right on the other. The Party of the Right was the most conservative party in the Yale Political Union. We were what you today would call classical liberals, or classical libertarians as well, though we considered ourselves traditionalists.

The interesting thing was, there was a very high Catholic content—I became a Christian under those circumstances—a very strong influence from Catholicism; some of the major leaders of the group were very activist Catholics. We had, at the same time, both a gut and an intellectual reaction against the SDS and their totalitarianism. I remember we were sitting in a work area of the library while there was a demonstration going on right across the street in Beineke Plaza. And I remember to this day, listening to one of the leftists who

said, "Let's give 'em the chant!" and the crowd went into some chant, and he said, "No, the other one!" and like turning a switch they switched to the other one.

At that moment, it was an apotheosis: "This is brain dead! These people's minds are turned off, and these people's minds are turned off in something they consider extremely important," which seemed like exactly the wrong place.

So the Party of the Right was the main opposition on campus to the Gestapo tactics of the Left. And we did lots of things—removing red flags and getting bottles and rocks thrown at us. A group of students occupied one of the buildings on campus and got expelled for doing so, so the Left put up "Reinstate the 45", and we went with paint cans and painted "Repudiate the 45" and they put back "Reinstate the 45," so we put "Liquidate the 45," always believing in performing *escalatio* on them.

[The Party of the Right] was essentially a large group of people who believed in individual freedom and individual responsibility, being the primary opposition to a group of people who believed in neither. It wasn't that we were in favor of the war in Vietnam and the SDS was opposed to it; the war in Vietnam was by this point really peripheral. This was just rebellion and anarchy and so forth on campus.

Anarchy in the chaotic sense.

Anarchy in the sense of chaos, not anarchy in the sense of anarchism. I guess I didn't like the tactics of the left. One example was a man on campus named Alex Spinrad who was head of the Young Socialist Alliance. He was an Austrian, and a Democratic Socialist, which was in those days, as today, a rare thing. He believed very strongly in individual responsibility and individual rights; he happened to be a Socialist, so his economics were screwed up, but his politics weren't so bad, and he really hated the SDS. He decided that he was going to put out a series of parody pamphlets. Every time the SDS put out a pamphlet, he put out a parody pamphlet. Well, since there were about three Democratic Socialists at Yale, he came to us, because he needed help, he needed people to distribute things, so we said, "Heh heh, we'll sign up for that one!" They were brilliant; he had an organization called SUDS, instead of SDS, and it was liberating the washing machines in the basements of local colleges. It was one thing after another. The fun point was having to change about three words in an SDS pamphlet to turn it into screaming hilarity.

We of course handed these things out. And of course, where do you go to hand them out for most effect? You go to the college where the SDS is the most powerful. So we would go to their colleges, and we would go around and put them on the dinner tables

during the dinner hour. Of course the SDS didn't like this at all, so they'd follow us around trying to take them back. There's nothing that will get somebody to read something faster than telling them they can't. So they would try to take them back, and now they're getting into fights, but not with us. It was perfect; it was beautiful; it was great.

We had a lot of support from the campus police. I remember we were out removing red flags one night, and I got a call the next morning. A friend and I spent the previous day in the cemetery practicing with bow and arrow. We went and shot an arrow with a thread attached to it over this flagpole that came out from a fifth-floor window; used the thread to hoist a rope, then tied the rope around and broke the flagpole from the ground. Well, while we were breaking the flagpole, the people whose room it was awoke and came out and started throwing bottles at us. We dove into the car a friend of mine had and sped away.

The next morning I got a call from the head of the campus police, asking me to please come see him. So I went over, and he said, "You people were out last night," and I said, "Nice we're having weather, isn't it?" He said, "No, no. I want you to know it's really easy to tell the signature of your group when you people are out. I want you to know that when we have a call and we figure it's you, we're going to respond to that call—in a month or two." So we had good relations. The university hated us, of course, because Kingman Brewster was busy capitulating to the terrorists and shutting down classes and not letting people go to class, and here we were, a very verbal and vocal opposition. I guess we were trained in street tactics by fighting the Left.

You mentioned that you were a libertarian also on economic grounds—you mentioned that Alex Spinrad was very much on the side of individual responsibility and rights—

But he was not a libertarian.

So at that time were you economically aware?

I was always free-market; I was always fervent for free-market economics—purely at a gut level, rather than from any intellectual basis, through high school anyway—because I didn't know what was right, but I knew the Keynesian crap we were getting made no sense—but I didn't have any intellectual basis for challenging it. I was instinctively distrustful of the idea that the government could tune the way society behaved. I saw where it clearly did not work in the political realm, which I understood something about, and it made no sense to me that it could therefore work in the economic realm, not so very distinct from the political.

My understanding of things from a more Hayekian sense came from meeting Phil Salin years later, when I was in graduate school.

Because as an engineer and a would-be entrepreneur, I sort of had a gut understanding of this, but the intellectual basis for it was much more from the political than the economic, which came much later.

You said before we got started that you are on the way from being a libertarian to an anarchist.

Yeah. Libertarians as constitute, in particular, the Libertarian Party, and also as shown by *Reason* magazine, tend to concentrate on, "How can we change society in directions that increase human options and increase freedom, on a more global scale?" I'm increasingly coming to the conclusion that, at least within a reasonable amount of time, that isn't the way to go about it. We can't reform society. You're not going to get the government to go away, or get smaller, by making the government get smaller; you're going to make the government go away or get smaller by ignoring it, which is essentially the Soviet model. The state withered away when no one paid any further attention to it.

Harry Browne's idea of how to find freedom in an un-free world.

Basically. If we all ignore it, it will go away. Of course, it won't matter if it doesn't! That's the nice thing about it. Increasingly, I've come to the conclusion that, at least for myself, I have to modify my immediate surroundings to match what my goals are and to do what I want, and I'm not going to successfully change how the world is. I can only deal with the few people around me, my family, myself; and that I can be much more effective doing that—I can be much more effective, from my own point of view, at figuring out how to shelter income than I am at trying to convince the IRS not to try to get it.

Removing yourself from the realm of the government rather than trying to remove the government from your realm.

That's right.

I'm planning to interview Tim May on crypto-anarchy next in this series. You also said earlier that you're getting more interested in cryptography.

Yes. I guess, as I've become increasingly an anarchist, I've gotten interested in it, but it's not just that. If we're going to move into a world that's more cyberspace, that's more information-based—I don't mean "Information Age," that trite stuff, but I mean it really is based on information.

Where information is the structure, or in this case the substance.

In this case, the medium is the message, right? The more that you do that, the more interesting it is to me how you preserve that information from being stolen or copied or

corrupted. Not only by government and so forth, but by competitors, or just because you don't want somebody to know something that's personal. So cryptography, secure communications systems—Jim Bennett and I are working on a global data services idea, a secure worldwide network where I can save a file here and it's automatically stored holographically all over the landscape, and can be reconstructed from anywhere else on Earth, but only by people who know how to do it. Those are the ideas that interest me because I see us moving more and more to this sort of interconnected web, as I call it, the network of networks that's the backbone of cyberspace. If the directions that I think we're going to go in are the directions that we do go in, the ability to live in this [information-based] world has got high evolutionary survival value, and that's one of the reasons I'm interested in acquiring it, but the other is that it's just plain interesting.

*Speaking of ignoring the state and it will go away, have you read Snow Crash?*⁶

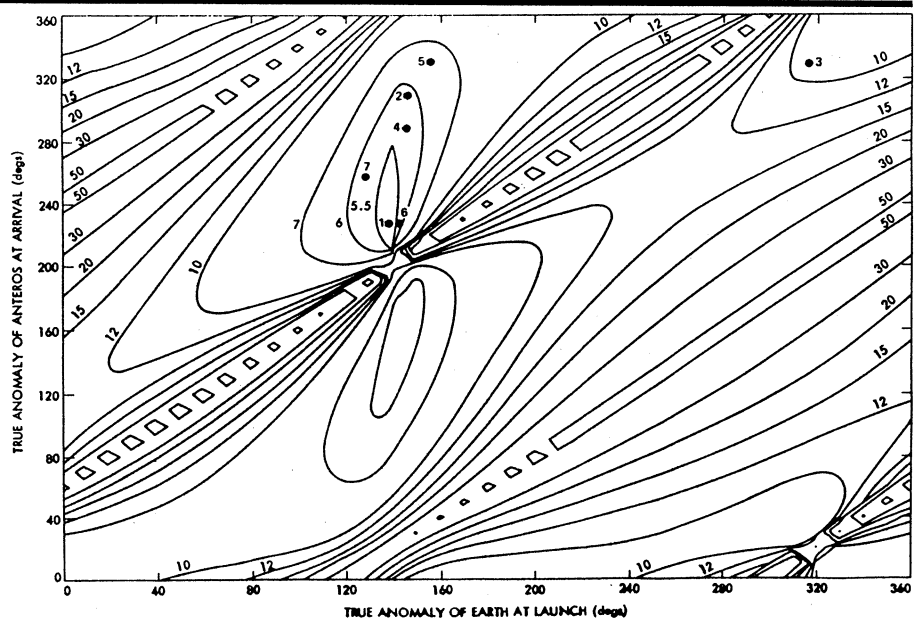
Yes, of course I've read *Snow Crash*. *Snow Crash* was the fastest a book has ever gone through this community. A model that I've always thought better than *Snow Crash*, which of course was tongue firmly planted in cheek in many many places, was Vinge's "The Ungoverned." It's probably the best model of "Ignore the state and it will go away."

So, I am increasingly turning into an anarchist, and of course most people think anarchists are the ones who want to go around throwing bombs at people, and I have no interest in doing that whatsoever. Nor do I have any interest in having bombs thrown at me, because the government tells me to go get shot at, either. I'd just like to come to a mutually agreeable pact in which I will ignore them and they will ignore me. This is unlikely to occur, so I'll just sneak under the rug and they won't find me.

Shading over from the political to the technical, when did you first suspect that some of the advances you read about in Heinlein and others were actually feasible?

Well, it never occurred to me that they weren't. That's part of the result of sticking to hard-science science fiction, which I always did. Back in the old days, when they had "New Wave" science fiction, I rejected it out of hand; I didn't like it. Not because I didn't understand dealing with the emotions, but that it didn't have the hard science. I liked the hard science.

It never struck me that the things were not practical. Even in the earliest times, reading things by, say, John W. Campbell, where the science isn't very good—it's hard science, but wrong—you got the idea of many different things that were possible. In Heinlein, you got the hard science and nothing but—



With actual calculations behind it.

With the calculations behind it. As with almost everybody else, my first contact with the idea of an intelligent machine, of course, was Mike, in *The Moon is a Harsh Mistress*⁷. I was fifteen when that was written. That was my first contact with mass drivers, which I would later see again with Gerry O'Neill; of intelligent machines; of a libertarian form of government; of lots and lots of things that were right there and hit right then. It was never a case of saying, "Oh, wow, all this stuff is possible," because I never thought that it wasn't; it just was, "Yes, of course, this is the way things are going to be." It took me years to realize there were people who *didn't* think that way, because the few friends that I had in high school, and the large number of friends that I had in college, pretty much all felt that way. It was only later that I realized that there were people who had these self-imposed limits on the possible.

I never had to go through a breaking of a boundary or a barrier; the idea that there were limits never occurred to me on a gut level. Certainly not self-imposed limits, where you'd say, "No, I won't go do this." Well, why not? "Well, I don't know... because no one's ever done it before." That sort of attitude was not one that I ever had to break out of, because it had never occurred to me. It certainly never occurred to my parents; I was never indoctrinated that way by them. So it was never a case of, this was something different or new, it was a case of occasionally I found out—well, let me give you an example:

I was in high school. One thing I will never forget—I was just laughing about this with my wife Heidi the other day—the math teacher said to me in a fairly snotty way—I realize now snotty, it didn't seem that way at the time—she said, "Here, you're special. Here, you're smarter than everybody else. You're gonna get into college"—she knew I

was going to Yale—"You're going to get into college, and you're going to discover that everyone else there is just like you, just as bright and just as capable as you are." And my response at the time was [gleefully] "Really? You mean I am not alone?" And it was true. It was great. It was truly astounding to discover.

That was the moment. That gave me no revelation about myself, but that gave me a tremendous revelation that there were people who thought that thinking beyond limits, thinking of unusual things, working on unusual things, was strange, and that the idea that you might like having a lot of other people who were interested in the same things you were in, would somehow be daunting, instead of fascinating. That was [snaps fingers] a tremendous revelation when she said that to me, and I still remember it though it's been 25 years ago.

Among the non-limits that had been part of your intellectual makeup all your life, that would include non-limits on lifespan. What was the first exposure to immortalism that you had, from your readings, and when did you first—not necessarily begin to realize that it could carry over, but when did you first encounter other people who were putting that into practice?

Again, it's a case of coming to the realization that people are on the opposite side. All teenagers are immortalists. That was natural. Almost everybody is, when they're a teenager; the concept of a finite lifespan has no meaning for them at all. Doubtless, my first contact with indefinitely extended lifespan was Lazarus Long, of course; Heinlein, in *Methuselah's Children*. That struck me as kind of a neat idea. Again, that's typical—Heinlein does things in a sort of macro way, whereas these days we're looking at doing the same things in sort of a micro way, but still, because he didn't have the technology—he *couldn't* have had the technology increase that we have 25 or 30 years later.

I have been for many years an immortalist because I am a Christian. I became a Christian in college, and came to the idea that everyone is immortal anyway; so the idea of living forever, of indefinitely extended lifespan, never struck me as particularly alien. Nor have I ever encountered anywhere (that I felt people were talking sense) the idea that there's something good about breaking down, getting old, dying. You don't like it in a car, why should you like it in yourself? It has never made any sense to me that those things oughtn't to be resisted and fixed and so on. Nor have I ever seen any, nor do I now, see any conflict between physical and spiritual immortality. I don't believe that immortality is possible in a physical sense, but let's debate that again in another 10,000 years. [laughs] I would prefer that we be around to debate it in another 10,000 years.

Gayle: Do you mean by that that there are some things staring at us like the heat death of the Universe, that there's no escape from?

Yes, or just that statistical likelihood of accident is non-zero, and non-zero chance of termination over an infinite amount of time results in certainty. But, being already an immortalist from a spiritual point of view, that doesn't particularly bother me. One way or the other. I remember saying once, to Mark Miller, that since I already believe that I'm going to live forever, I'm not worried about the other part.

As far as understanding that there were people interested in physical immortality, or at least indefinitely extended physical lifespan, I read Ettinger's book when it came out. So I was familiar with *The Prospects for Immortality* when it was published in 1963. I read it then; it's sitting on my shelf at home, an original edition of it. So I knew about the cryonics movement before there was a cryonics movement, and thought it seemed a quite sensible way of going about it, at least as a short-term solution.

Are you an Alcor member?

No.

Are you contemplating signing up?

[sighs] I have certainly contemplated signing up. It is not a matter of significant concern to me one way or the other. The standard arguments that a person will make regarding whether or not a person should become an Alcor member are utterly useless on me, because I completely agree with all of the reasons that most people give [for signing up]. It is a personal decision that I must come to on the basis of my Christianity, because I'm looking at it from a very different perspective, and it's very difficult to give me insight on that

particular front. So I am not an Alcor member, though I certainly—obviously—have nothing but admiration for Alcor's intents and goals. *An Alcor sympathizer.*

I am clearly an Alcor sympathizer, yes.

You became a Christian as an adult, rather than being raised in it. Your family was—?

Nor have I ever seen any, nor do I now, see any conflict between physical and spiritual immortality. I don't believe that immortality is possible in a physical sense, but let's debate that again in another 10,000 years. I would prefer that we be around to debate it in another 10,000 years.

My family was nominally Christian; my parents are no longer *nominal* Christians, they are full-fledged Christians. They probably always were, at least my mom, but I grew up in the First Church of the Well-Dressed; and was confirmed because that's what everybody in the town did, and didn't particularly believe—at 14, it was compartmentalized, and by 17 it was gone. A good healthy dose of Heinleinian agnosticism—decidedly not atheism; atheism never made any more sense to me than any other form of absolutism, and as Heinlein himself pointed out, it's an absolutism based on negative proof, and I'm a good enough scientist and mathematician to know that negative proof is, although not impossible, extremely difficult.

I was raised in a nominally Christian background—we went to church most of the time, not all of the time; the church was not particularly edifying and certainly not anything that would be called a quote "Bible-believing, spirit-filled," whatever, church.

The usual kinds of complaints that people have about organized, quote-unquote, religion—it's observance of the outward forms and not so much—

Right. "I don't belong to any organized religion; I'm an Episcopalian." Actually, I'm not an Episcopalian. Yeah, it was outward form, and whether or not there was any inward belief, it didn't communicate to me. When I went to Yale, as I said before, a good number of the people who were in the Party of the Right—let me explain for a moment the Yale Political Union. In those days, there were such things as liberals who believed in freedom of speech—this has pretty nearly vanished—but in those days, the Political Union was a debat-

ing organization that also had speakers. It was rapidly becoming a speaker's bureau, but its emphasis was still on debate, and a lot of it. Some, like the Party of the Right, which had elected membership for life at least, was purely a debate organization. There was the Party of the Right, a Conservative Party; at various times, a Progressive Party and a Federalist Party, a Liberal Party, and sometimes a Party of the Left. There would be debates as well. There would be speakers in the whole forum, but the Party of the Right was purely debate-oriented in their own caucuses.

The Party of the Right was usually the largest or second to the Liberals or sometimes occasionally third to the Conservatives, but all pretty large parties in the Union. Plus, the Party of the Right had this whole bunch that the Yale Daily News called "the interlocking directorate of the Right." We controlled, at various times, the Young Americans for Freedom chapter, the Young Republicans, and the Young Democrats. Of course, the Young Socialists were very closely allied with us, which was because they were democratic. Plus, we had the Calliopean Society, and we had the Coalition for a Free Campus, we had ad hoc organizations that the SDS would have been proud to have been able to conjure up for as many different reasons as we did.

A lot of the people who were in that were members of an organization called the Ultramontane Society. The Ultramontane Society was, is a traditionalist Catholic organization. "Ultramontane" is the concept of the Papal authority sent over the mountains, out of Italy, into the political realm in Europe, during the late Medieval and early Renaissance times. Many of the leaders in that group were very strongly based in Thomistic and other Medieval philosophical strata. This was the first time I had seen that beliefs of any kind could be systematized. Probably, from that origin, is why I am a libertarian, why I am a free market person, why I can now express that in philosophical terms. What were previously gut reactions became intellectual, because that was the first time I had seen knowledge systematized. I'd never been exposed to that in high school. Mathematics, yes, but that was very different. Even there, it wasn't presented as a base of systematized knowledge, even though of course it clearly is. So that was the first time I had encountered knowledge systematized, and the interesting thing is that it was knowledge systematized in a realm I had not given much thought to, namely the spiritual realm.

What I saw there, in looking at the philosophy of Thomas Aquinas, a bit of Anselm

and Augustine and so on, was a philosophical basis that held together and made sense. A philosophy must have two driving things, at least: it must make sense internally, and it must connect to the real world in terms of axioms verified. At that point, I was willing to say that Christianity, at least as extended with Thomistic and Platonic philosophies, had the systematized base, whether or not it had the legs that held it up as being pinned to data in the real world. As time went on, I came to believe that that was also true, and became a Christian accordingly when I did that, probably one of the few people in the history of the world to become a Christian from purely intellectual reasons.

I am a Christian because I believe that Christianity is true. I am a Hayekian because I believe that Austrian free-market economics is true, i.e., is that it corresponds to the real world.

It's a system with predictive value.

It's a system with predictive value, exactly. And that influenced my thoughts on a lot of things—obviously on immortalism, but also—if you've got a philosophical problem that you want to embrace, the question may be Extropianism in general, it may be Randian philosophy, it may be "How do you deal with a market? How do people exchange ideas and verify ideas as good or bad? How do you organize a group to accomplish a set of ends?"—All of those things could be approached from the point of view of systematized knowledge. Not that we know everything; not that we've incorporated everything that can ever be incorporated, but that there are tools we use. We know that ad hoc "I feel this way, I want it to be this way" notions are bad ways of approaching those problems, and there are instead evolved systems that have been developed by other people that give you good rules of thumb sometimes, good absolute rules other times, to approach things—that's a much more reasonable and much more rational approach to use for solving a whole wide range of problems. And I first got that systematic approach from the Thomistic and Catholic philosophy. Now, I'm not a Catholic, and I never have been. That isn't what I got out of it. I didn't get the specifics; I got the general.

Do you consider yourself to belong to any specific denomination?

No. The church I go to is Presbyterian; the church I was raised in was Dutch Reformed, which might as well have been Presbyterian in terms of it was presbyterian in the sense of an organization—a presbytery is a group of people who run the church, so it's a presbyterian organization in that. Insofar as the religious beliefs of the Dutch Reformeds are apparent, they are similar to the Presbyterians, but no, I go to the church that happens to be one that I like and that is in my opinion

teaching the Bible and is alive, in that the faith of the believers who are in the church means something—they do things with that faith, it's not a dead church. I try not to go to dead churches. But I don't particularly care about the denomination; I'll go to charismatic Catholic services; I'll go to Baptist churches and Pentecostal churches and so on. I try to avoid stepping on the toes of the denominations that don't agree with each other, so, for example, I won't take communion in the Catholic church because they don't want to offer it to me, not because I think there's anything wrong with it; but that's fine, that's their right.

You became a Christian for intellectual reasons, because you saw evidence of the truth of Christianity in the world. How do you reconcile that with the idea that proof denies faith?

I tend to agree far more with St. Augustine, which is that reason bolsters faith, rather than the opposite. It is a philosophical truism that you cannot believe something that you have proof is true. That's like saying you can't have a door that's both open and closed at the same time; there's no emotional content to that statement, it's simply a true statement. There is a wonderful line in a book by Morris West who wrote *The Shoes of the Fisherman*⁸; the book is called *The Clowns of God*, and in it, a person shows up who is asked, "Are you a believer? Are you a Christian believer?" And he makes the statement that, "That option is impossible for me." They, of course, take it that he's an atheist.

It turns out that the person that they're talking to is Christ, and the option of faith is denied him; you cannot have faith in what you know to be true. So in that sense, of course, proof denies faith; but so what? It does not bother me in the least that there are many things—nothing to do with religion—that I believe to be true that I can't prove. I believe my wife to be faithful—I cannot prove it, but I will nonetheless live as though that were the case, because to the best of my evidence, it is. It makes far more sense based on the evidence that I have than the opposite does, or any other variant of it does, and yet I can't prove it. So, in that sense, were I to prove it, I would no longer require the faith. But faith is sufficient for those things, and there are many, many areas of life, not just religion, where there are things that you or I believe to be true, and act on that belief—faith is more than just belief—in other words, put my trust in that belief, because I cannot have proof.

If you encountered an open-minded person—they're rare these days, but there are a few—to whom you were interested in offering the evidence that made you a Christian, what proof in the world would you present to someone to demonstrate to them the truth of the Christian faith?

Of course, this is a topic for hours and

hours and hours of discussion; I can go into it in several different directions. Understand that I am necessarily summarizing. They are several directions. We have to go back for a moment and consider what a person is. I don't mean what makes a person an individual, but I mean what a human being is, what a human being is like. As a scientist, I am very familiar with the mindset that is very common among scientists, that the side of the human that is purely rational, deductive, axiomatic, linear, is all that there is to the person. We all know that that isn't true, but it is a very common mind-set among scientists. It's also a very common mind-set among technoids, whether or not they themselves are scientists. It's a common mind-set; I believe that that mindset is often followed both by scientists, for more well-understood (by them) reasons; and by technoids, for more intuitive reasons, that the opposite, which is, "No, everything is touchy-feely," is repulsive to them, and in fact is a lousy way to go about the world. If the choice is between cold reason and squishy "feelingism", I will stick with the cold reason, thank you very much. The problem is, a lot of people—myself included, for some number of years—get into making that a dichotomy, and that's not a dichotomy. Humans are far more than the sum of the molecules that make them up; in particular, they're the organization of those molecules, which is not itself a material thing.

So, as far as evidence in the world, I see in the world things that are inexplicable by standard reason/science-based arguments. I see, for example, the idea, very common, and well-shared among societies as a whole, that there exists a standard of behavior that humans ought to follow and do not follow. Both of those things seem to go hand in hand. And, having looked at it from many, many, many different points of view—religions, philosophies, cultures, and times—those core things are extremely similar across all of humanity, which is weird, because I can see that there would be consistency of actual behavior—evolved reasons why actual behavior matches—but why there should be consistency of ought-to-be-followed-but-aren't-followed behavior, is a bit more of a mystery. As a scientist, it is a complete mystery to me.

You don't believe that in a sense that's a spontaneous order, in that these are a set of rules that work best, and they emerge through a process of variation and selection?

No, I don't agree with that, for two reasons. First of all, I don't agree that it's a spontaneous order, and second of all, I don't believe that they're the rules that work best, so the answer is no on both counts. I believe that there are many things that it is possible to account for in a lot of different ways. The fact that you can account for them in different ways doesn't make any of those ways true; it makes

the *observation* correct, or at least it grounds the observation. An order may be spontaneous or an order may be caused; if you can't see the cause, you can't tell them apart.

I'm a very strong believer in evolved systems, but I believe evolved systems are caused. They may not be intentional—please understand the difference—they may not be intentional, but they are caused, in the sense that there are active agents producing things, and maybe unintentional consequences. But the point is that, even if they are intentionally caused, it's very, very often difficult to tell, by looking at it from the outside, which of the two is going on, and it's dangerous to make too much of a statement about what's going on underneath. That's not what I mean by it not being a spontaneous order.

Human beings generally have a *why* behind even the spontaneous orders that have grown up. It may be the one you alluded to, "Because this makes society function better." Well, one, "It may make society function better, but why should I care?" Two, "Why should I ever, not as an ideal, but ever, involve in behavior that results in my death?" and yet throughout history, throughout the ideals of culture after culture after culture, there are plenty of ideals that say, "Under these circumstances, it is worthwhile to die for somebody else." It may be "Greater love than this hath no man, that he lay down his life for his friends," or it may be the fireman who rushes into the building to save a trapped child at grave risk of his life; he's not doing that because it's going to protect his genes, because it won't. In no rational sense will it protect his genes, and the idea that somehow this is going to increase the likelihood that somewhere else, someone's going to rush into a building and save his genes, is so unlikely and so far-fetched that I find that real hard to believe. I prefer to say, no, it's in some sense, like much else that evolves, a manifestation of physical law.

The process of evolution says nothing about the driving forces of evolution. Evolution proceeds the way that it does because it is constrained in the way that it can proceed by the physical laws that are running underneath. It may have many, many, many different ways it can go; but it also has many, many, many ways it cannot go. We cannot evolve a system of antigravity as a natural process of nonthinking creatures, because it violates physical law, or at least it appears to. You can't evolve that, however much you shake things around, there are things you cannot evolve. I'm not so much interested in the mechanism of how things go or which direction they may have gone, as I am, as a scientist, in what the underlying structure is. I see, therefore, these behaviors that are not explainable anywhere else; I have a belief that things

don't come from nowhere. Organization may come from nowhere—I don't even believe that's really true; I think it's a manifestation of underlying physical law even when it's spontaneous order. That's what I meant by "caused," that evolution is caused, even in non-rational systems.

I'm trying to clarify in my own mind what you mean by that. The results of the spontaneous order are—I hate to say pre-determined—but there is an agent that is producing the results,

I have a belief that things don't come from nowhere. Organization may come from nowhere — I don't even believe that's really true; I think it's a manifestation of underlying physical law even when it's spontaneous order. That's what I meant by "caused," that evolution is caused, even in non-rational systems.

even though the mechanism is spontaneous order?

There are two different things occurring. In a case where there are rational agents operating, they're all intending things, they're intending consequences. The fact that they can't see all the results of their intentions doesn't stop the fact that there are things being intended, and in that sense, what happens is caused, but not intended. In the case of non-rational things, like for example, ice melting, clearly there's nobody there to have an intent. What I meant there by "the process is caused," water evaporating from a puddle is caused in the sense that it's consistent with physical law, not that I mean that there's rational behavior on the part of the individual molecules. There's statistical behavior in accordance with physical law. As a scientist, I'm interested in what the physical law is. So I look for explanations.

I look at the radical materialists, and I see no explanation that makes any sense to me. I hear them stating that there are explanations, but I do not hear them stating explanations that make sense to me. I look at the concept that "the world is the way it is because it was constructed that way," and—

That's a tautology.

Well, if you believe it was constructed. My house is the way that it is because it was constructed that way, so at least the model makes some sense; I can make analogies that make some sense. My house is not a sponta-

neous order, it was designed. It has multitudes of examples of spontaneous order in that design, but that makes it no less designed. To me as a scientist, the world looks far more like a construct than it looks like something that just happened. I also see other things—besides what people call moral behavior versus what they actually do, and the fact that they know that there's a conflict—I see the existence of something called love, that is unrelated or at least need not be attached to reproduction, to family, even to presence of the other person.

And I see no materialist reason for this. I can rationalize it away, and say, "It's hormones; it's trained behavior; it's all of those things," but when I look at it, any order is more than the sum of its parts. It is, among other things, the order, itself, which wouldn't be there—you wouldn't talk about a spontaneous order forming from the parts if the parts were the order. So there's more to it than that. And I see no place where that can come from; it looks to me much more "caused"—it's something in the substratum of the Universe, rather than an illusion. Even illusions have substratum causes.

So I see that, and over and over and over again I see causes of that. Those are sort of external. I also know, as a scientist, that creating an explanation for something, that holds together, is extremely difficult. Creating a new scientific theory is very, very, very hard, and creating one that holds together as well as Newtonian gravitation or general relativity is one of those things that occurs a few times in the history of humanity. It's a very, very hard thing to do. The fact that a theory—for example, Newtonian gravitation—holds together, regardless of attachment to externals, is itself a very powerful argument for its truth, because it's very hard to do. I make the distinction between self-consistency and axioms that connect to the real world, but that's really a false dichotomy. The two are very much tangled together and it's not as clean a cut as that.

Because our notion of self-consistency is largely shaped by our experiences in the real world.

Exactly. So, when I see philosophical structural systems that have been pounded on for a thousand years and still hold together, to my judgment, as self-consistent systems—self-consistent and rich systems; it's very easy to make a group of three things and make them self-consistent; those are self-consistent but not rich—but to make self-consistent, rich systems is a very, very hard thing to do if there's nothing plugged into the world holding them into place all over, and of course experiment is what holds scientific theories in place.

Aristotle, trying to reason from whole cloth, or even Kant trying to reason from whole cloth, didn't come up with relativity. You had to have Maxwell's equations and the Michelson-Morley experiment and measurements of Lorenz contractions to have the anchors.

Aristotle didn't even come up with the principle that things fall at the same speed; he thought heavy objects fell faster.

Yes; at some point you've got to go out and plug into the real world. Those are, in grotesque summary, the externals. Plus—if I have a friend who has an ailment, and that friend takes a particular nutrient, and takes the nutrient for a period of time, and gets over the ailment, and ascribes the getting over the ailment to having taken the nutrient; if I have the same ailment I'm liable to take that nutrient, not because of anything external, except the fact that, from a history of knowing that person, I believe the person to be honest—again, I have faith that the person is honest; I can't prove it, but I can give many examples. Giving an infinite number of examples is not the same thing as proving—but nonetheless I have *faith* that the person is honest.

The proposition that the person is honest is consistent with all your available evidence.

That's correct. And again, that's a perfect example of a non-falsifiable, but verifiable, circumstance. You can see it over and over; you're verifying it, but you can never guarantee that you can falsify it based on finite evidence. You *may*—a single false example will prove that it's wrong—but you may never get that false example. And yet—just because you can't prove it false, you're picking up bits and pieces of verification, and that's pretty good. Even if occasionally you do falsify it, that doesn't really change it.

Do you subscribe to Popperian kinds of epistemology?

As a scientist my approach to most theories is, "Theories have to be falsifiable." But I will go beyond that and state that there are things that are not falsifiable that are verifiable. I am willing to accept the idea that there are ways of looking at things that may not be falsifiable, for many reasons. For example, the theory that the universe is going to end in the next hundred thousand years is unlikely to be falsifiable by me. I still don't believe it's true, and I will live accordingly, that I do not believe it's true, but it is not falsifiable, by me, because it's unlikely that I'll be here in a hundred thousand years to tell. It may be falsifiable, but I have no way of knowing that, and yet I have to live consistently with that.

Now I'm going to go into more of the sort of internal [evidence]. I see other people that I have faith in, that I *believe* when they tell me things to be true. I saw—I see—in friends who are Christians, who would tell me of the

changes that took place internally, of the things that happened to them when they became a Christian. They were people whose honesty and low likelihood of self-delusion I had quite a bit of faith in, though obviously I couldn't prove anything.

But you had seen many examples.

I had seen plenty of examples of it, in a large number of people. Between those two, was sufficient cause for me to say "I believe these things are true," with a lot of other things—evidence from eye-witness accounts of the Resurrection, people who then later were killed because of their beliefs and because of what they were saying, and from the observation that though I had never met these people, I had met plenty of people who it seems unlikely would be killed for a known fraud. Lots of things along those lines said, "Okay, I believe." And have seen in myself the same changes, the same phenomena—the same answered prayers—that I had seen and that my friends had told me about before. So, from an internal perspective, I have seen phenomena that are consistent with phenomena that I had been told about, and that are internally verifiable.

Because Christianity is a matter of *faith*, there will never be, this side of eternity, proof. There will never be, this side of eternity, proof of a scientific theory, either. That doesn't stop me—I would rather continue improving scientific theories, even if the theory is wrong in some parts—in some sense, Newtonian gravitational theory is completely wrong. And yet it's certainly practical for sending rockets to the moon, and if what I want to do is send a rocket to the moon, it's quite sufficient. The fact that I may not understand the wrinkles and how things work, on an absolute level, either from a physical or from a spiritual point of view, doesn't mean that I have any excuse for not casting my bets with what I *believe* to be the truth. So I do as a scientist, and so I do as Christian.

You mentioned earlier that when, for example, a moral code that appears to be global and emerges in widely-separated societies, emerges over and over again, in the sense that if it comes out of a process of evolution, you seem to imply that it was inherent in the physical principles that underlie the process of evolution. In that sense, is the Invisible Hand the Hand of God?

In some sense, the Invisible Hand is the Hand of God, because of course I believe that God created the universe and gave it the structure that it has, insofar as things that evolve are constrained by physical law, then, yes, in that sense, it is. What I'm not saying is that any particular social or political or economic system is developing with direct manipulation by God or anything of the sort. I do believe that God acts in everyday life and

through individuals and so on. I also believe in miracles. But I don't believe that God is such a poor carpenter that He needed to set up a system of physical law that He tinkers with continuously. But, because I believe in first causes, and I believe that God is the First Cause of the universe, in that sense, yes, your statement is true, the Invisible Hand is the Hand of God. I do believe in the scriptural verse that says, "Where the spirit of the Lord is, there is liberty." In that sense, I definitely think it's properly called the Hand of God.

The interesting thing is, literally the only place besides the Church that I have found a consistent group of people who live moral lives is libertarians/anarchists. They consistently live moral lives. After all, the only political or economic system that I've ever seen that I consider moral at its core is capitalism. The very idea that I'm going to spend my life attempting to create things because *you* want them—that's a wonderfully altruistic system at its core. "I'm not producing things because I want them, I'm producing them because I think *you* do." That's wonderful. It's not a zero-sum game; we can all get rich. I think that's a very wonderful thing.

There's another direction in this, and this is something that I've been thinking about very recently, and was of course the origin of my talk at Eris this year: looking more at what a human being is, and tying together the Extropian things, from life extension through cryonics through human uploading and all of that, in the Christian viewpoint.

In Part Two, Dave tells me about the "Seven Paths to Immortality"—how Extropian technologies like life extension, cryonics, virtual worlds, and uploading, are seen from the Christian viewpoint—as well as the anthropic cosmology of Teilhard de Chardin, the meaning of the Singularity, why Extropians change jobs so frequently, and who gets the "essence" when three philosophers share an artichoke. Stay tuned.

¹O'Neill, Gerard K. *High Frontier*. New York: Bantam, 1978.

²Penrose, Roger. *The Emperor's New Mind*. New York: Penguin, 1989.

³Hofstadter, Douglas R. *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Vintage Books, 1980.

⁴Ross, David J. "Persons, Programs, and Uploading Consciousness," in *Extropy* #9 (Summer 1992).

⁵For example: Hulkower, Neal D., and Ross, David Justin. "Missions to the asteroid Anteros, and the space of true anomalies." *Acta Astronautica*, Vol. 10, No. 3, pp. 133-141, 1983.

⁶Stephenson, Neal. *Snow Crash*. New York: Bantam, 1992.

⁷Heinlein, Robert A. *The Moon is a Harsh Mistress*. New York: Berkley, 1968.

⁸West, Morris L. *The Shoes of the Fisherman*. New York: Dell, 1964; *The Clowns of God*. New York: St. Martin's Press, 1981.

FORUM

Automated police & defense ("Nanarchy")

This issue's Forum features just a fraction of the 16,500-word debate from the Extropians e-mail list, sparked off by Mark Miller's comments in his discussion with David Krieger last issue.

Following the Forum is another Extropians e-mail list follow-up, this one on Mike Price's wormholes article in *Extropy* #11.

All writings appear here with the permission of their writers.

enforced rules such that ... that kind of military instability cannot happen"

They create, and then relinquish control over, "a dispersed system of communicating nano-Gorts" which is "monitoring for certain inter-boundary activities that may be coercive, and stops any that fall within the possibility of coercion", the boundaries being those determined by some "initial division of property" in the whole universe. Some set of mutual funds would be created, every person alive on "inheritance day" would get an equal share in each fund, with trade in shares then allowed.

Mark sees the central design problem as how to "engineer ... a mutually constraining development process for designing a secure mechanism ... such that ... we can be confident that the system as a whole does not have any trapdoors in it". And Mark wants to find a minimal kernel, with "minimal set of constraints, on top of which we can bootstrap a system of enforcement mechanisms that are capable of enforcing such a system of property rights." This is "extremely difficult", but "cryptographic techniques as well as the progress on program proving ... give me hope that we could actually carry this thing out".

MY CON ARGUMENTS

First let me respond to specific points, and then I will comment more generally.

Post-enforcement is not "trashed" by AIs and upload copies. One viable alternative is to (except for human biological children), always hold "parents" always retains responsibility for the "children"; if parents can't be found then siblings take their place. Parents might sell this obligation to an insurance company, but if that company goes bust, they are responsible.

The risk that any anarchy we might create now could evolve into something else, and then evolve again, must be weighed against the risk that any "permanent" solution may turn out to be terribly wrong.

enforcement depends on punishment creating an incentive not to commit a crime, and that gets trashed by post-Singularity confusions of identity" revealed by considering "If you create an AI, and it goes out and kills somebody; are you responsible?" and that "the whole process of thinking about agoric systems made clear that you want to assign rights to lots of little things". So instead we need "pre-enforcement" where "when the coercion is attempted, it is prevented".

The third problem Mark wants to address is that:

"In the absence of ... nanarchy ... if you go with the homestead model, and in the presence of the possibility soon of self-replicating, space-faring machines that are able to arrange for their own military defense and able to use the resources that they're acquiring by spreading to engage in that defense, what results is a terrible winner-take-all race ... whoever gets there first takes the entire universe, and the rest of us are left with essentially nothing. Alternatively, ...you might end up with .. a very extreme oligarchy, and it's also not necessarily stable because of the logic of military power in a system where whoever expands fastest or expands in the direction of more available mass-energy, gets to have more mass-energy at his disposal to beat on the other guys. There's a positive feedback in there that probably still ends up with one winner taking all.

To solve these problems, Mark proposes "central planning and central authority". Some unspecified power ensures that "the first wave that explodes out there into the universe be the minimum framework of

Date: Mon, 30 Aug 93 15:01:32 PDT
From: Robin Hanson
<hanson@ptolemy.arc.nasa.gov>
Subject: WAR/NANO/LAW: Against Nanarchy

Derek Zahn suggests:
>I hope we can have a debate about "nanarchy".

OK, I'm ready to debate (my *Extropy* arrived Friday), and offer to take the CON side. Will any worthy opponent take the PRO side? I will now respond to Mark Miller's discussion in the *Extropy* interview, but to go beyond that, we'll need a proponent. (This idea has been published no where else, to my knowledge).

First let me summarize Mark's case.

SUMMARY OF PRO ARGUMENT

Mark offers nanarchy as an alternative to anarchy and minarchy, and as a solution to 3 problems.

The first problem is that anarchy "relies on the ability to use coercion", and while its answer to the question "who will watch the watchers" is self-consistent, "the paradox is that" it requires "the proper activities of those users of force" in order to "turn market forces in on the users of force. (I think he means to say that there are other semi-stable situations besides anarchy, and so an anarchy might evolve to something else, and other arrangements need not evolve to anarchy.)

Second, Mark thinks that "post-

Unless the permanent solution is clearly better than the "average" expected future political regimes, risk averse folks should prefer the mixture.

If we are considering whether to offer our political support to some growing "movement" in favor of nanarchy, we must consider the possibility that this movement will get out of hand, implementing something that looks to most people like nanarchy, but not to us.

If this nanarchy must be designed and implemented, and a global political consensus formed in its favor, all before the first wave of intersellar colonization, there is likely not enough time left.

If technology continues to improve, it's not clear how a nanarchy built on old technology could prevent more advanced attempts at "coercion". If technological improvement must "run out" before nanarchy, then this seems unlikely before the first intersellar colonization.

A nanarchy trying to implement tradable universe shares would have to be able to tell who was the "rightful" owner of some shares. If these rights could be split arbitrarily, the task of the nano-Gorts in detecting and preventing all violations could quickly become impossible. So nanarchy would have to be limited to enforcing some limited and clear concept of "coercion", and I'm not sure there is a natural choice here.

Now let's get to the central point. Mark clearly thinks that colonization of the universe naturally leads to a single military power, and this seems the primary reason to create a single power now, and "do it right". However, this central point is not so much poorly defended as hardly defended at all. Sure, "whoever expands fastest ... gets to have more mass-energy at his disposal to beat on the other guys". But the same could be said of ordinary economic growth, yet few argue this implies a single power.

The question is: why would a power controlling more mass tend to grow at a faster percentage rate, or find it in their interest to wage war? I would actually be most convinced by a board game, plausibly modeling the colonization process; where I could play and see that a single power was the natural result.

If at least one of the largest expanding powers were "open", allowing us all to buy shares in it, or to immigrate into it, then there is little risk that the universe will be shut off from us.

Mark says he developed most of these ideas with Eric Drexler years ago, but

only now does he have "a community to say these things to." But they need not only to talk, but also to listen. I encourage one or both of them to now submit their ideas, in detail, for critical scrutiny by a larger community.

Robin Hanson

hanson@ptolemy.arc.nasa.gov 415-604-3361 MS-269-2, NASA Ames Research Center, Moffett Field, CA 94035 510-651-7483 47164 Male Terrace, Fremont, CA 94539-7921

Date: Tue, 31 Aug 93 22:46:07 PDT
From: hfinney@shell.portal.com
Subject: WAR/NANO/LAW: Against Nanarchy

I will take issue with some of Robin's points, but I cannot support Mark Miller's vision of nanarchy. I agree with Robin's criticism that nanarchy would not be possible, because of the difficulty of anticipating and preventing all forms of coercion which may ever be developed. I think that the universe is too complex to be able to anticipate things as fully as this. Godel's theorem and the related computer-science proofs of the impossibility of solving the halting problem both point to the limitations of any finite decision-making process in the face of the full range of complexity which nature may create.

However, I do think that Miller's criticisms of a non-nanarchy future have more validity than Robin suggests.

> The risk that any anarchy we might create now could evolve into something else, and then evolve again, must be weighed against the risk that any "permanent" solution may turn out to be terribly wrong. Unless the permanent solution is clearly better than the "average" expected future political regimes, risk averse folks should prefer the mixture.

The difference between anarchy and nanarchy in this context is that the first is inherently fluid and capable of shifting to other states. Nanarchy, on the other hand, is designed not to be capable of such shifts. We might even say that nanarchy is DEFINED to be a stable, self-enforcing political system which provides minimal guarantees against coercion.

It is not clear how such a system could be terribly wrong; in fact, it would appear that such a system would be inherently right. In fact, it could be said to be a "libertarian utopia", a universe where libertarian principles are in effect part of the laws of nature, as Miller says.

This is the main attraction of the nanarchy concept, that we would have once and for all a guarantee of a non-coercive universe.

With even a Friedman-type anarchy system, there is always the danger that the anarchy might be unstable, that some group of protection agencies could form a secret coalition, and in a surprise move destroy their competition. It is not at all clear that this would be impossible. Nanarchy is an attempt to make this impossible.

> If we are considering whether to offer our political support to some growing "movement" in favor of nanarchy, we must consider the possibility that this movement will get out of hand, implementing something that looks to most people like nanarchy, but not to us.

This is a valid practical consideration, but IMO it should be considered only after we settle the question of whether nanarchy, if successfully implemented, would be a superior system. I presume from Robin's comments that he does not think that it would be.

> Now let's get to the central point. Mark clearly thinks that colonization of the universe naturally leads to a single military power, and this seems the primary reason to create a single power now, and "do it right". However, this central point is not so much poorly defended as hardly defended at all. Sure, "whoever expands fastest ... gets to have more mass-energy at his disposal to beat on the other guys". But the same could be said of ordinary economic growth, yet few argue this implies a single power.

I do believe that this could easily occur, because of the speed and ease with which nanotech self-reproducing machines could expand from star system to star system. The standard scenario involves small nano-seeds entering a star system and touching down on some asteroids. They start building copies of themselves, including seed launchers for nearby stars. They also start building detection equipment and weaponry to destroy rival seeds which attempt to enter the same system.

Using this approach, the first seeds to enter a system, if they beat their rivals by only a few years or perhaps even a few months, would end up with full possession of the planetary resources of that system. They would then be in a position to expand outward all the faster. You would have an initial wave of exponential growth (which would settle down to third-power growth) and in such a framework an initial head start could become a numerically overwhelming advantage.

(Our intuition of how growth works, based on experience with biology, does not fit a model with such easy growth. A better model would be a super-saturated solution into which you drop a few seed crystals.)

Harder to defend is the notion that the

possessor of the majority of star systems in some region would be able to use his numerical majority effectively against his smaller rivals in order to steal possession of their planetary systems. We don't know enough about interstellar war to judge! But it is certainly plausible, and even without this effect I think the initial exponential growth phase, when all that is going on is claim-staking and colonization, is going to allow only a very small number of competitors.

Perhaps we could agree that, given enough of a head start, one agent could effectively take permanent possession of all non-Solar star systems. Then if we agreed on that, we might discuss how much of a head start is needed, and whether such a head start was likely to occur.

> If at least one of the largest expanding powers were "open", allowing us all to buy shares in it, or to immigrate into it, then there is little risk that the universe will be shut off from us.

It is not clear why such a power would allow us to do this. What would our bargaining position be to encourage it to allow us such access? What can we offer a power which expects to gain possession of a decent fraction of a whole Universe?

I for one don't buy Miller's argument that owning 10^{-10} of the universe is 5% as good as owning the entire universe. I am astonished that most people supposedly feel that this is true (at least according to Miller). This simply shows a failure of imagination, in my book, like the people who think that if they keep buying lottery tickets they're bound to hit the jackpot one of these days.

Assuming these expanding powers think like this (which seems valid since people who think like this would be most likely to try the expansion strategy, just as people who want to make zillions of copies of themselves will be among the first to try uploading), then I don't think any of them are going to cheerfully sell off shares.

In sum, I think it is all too likely that a nanotech race for possession of the universe will occur. I don't think nanarchy will work, but if it did it might be better than what will happen without it.

(It's also worth noting the many structural parallels between this debate and the 'grey goo' issue, with nanarchy playing the part of Drexler's 'active shields' from EoC. Despite the revisionist thinking which suggests that grey goo will not be a problem because of the presumed difficulty of its creation, I think these

arguments only suggest that it might come five years later than was first thought.)

Hal hfinney@shell.portal.com

Date: Wed, 1 Sep 93 0:03:23 PDT
From: szabo@netcom.com (Nick Szabo)
Subject: WAR/NANO/LAW: Against Nanarchy

Robin Hanson:

>> Sure, "whoever expands fastest ... gets to have more mass-energy at his disposal to beat on the other guys".

This might be, well, immaterial. A better strategy could be to spend cycles on increasing computing mem*cycles with

10^{-10} of the universe is 5% as good as owning the entire universe.

The fraction of the universe owned will likely be unimportant compared to how well one can put it to use.

Initial definition of property rights raises many of the same bound-rationality problems as central planning of economies. Only much worse: in this case the planners must forecast both future technology changes and the actions of agents not just as smart as themselves, but quadrillions of times smarter. Imagine, for example, defining property rights for the radio spectrum in the 18th century.

Nick Szabo

szabo@netcom.com

Some strategies will lead to better cryptanalysis, trojan horse capabilities, physical weapons, and/or economic capabilities within the market than competitors which choose a different brain/brawn and weapons mix. The chance that we can find... the best mix for implementing the nano-Gorts, is practically nil. (Szabo)

Date: Wed, 1 Sep 93 14:12:41 PDT
From: Robin Hanson
<hanson@ptolemy.arc.nasa.gov>
Subject: WAR/NANO/LAW: Against Nanarchy

I proposed that "a viable alternative" to nanarchy is to always hold "parents" responsible for their "children" "(except for human biological children)".

Hal Finney responds:

> why in this alternative would one except humans?

Just as a grandfather clause, because we're used to this, and it doesn't fail terribly. I don't think it makes sense to hold human "parents" always responsible for the actions of their AI "children". ... If you created the AI twenty years ago and it has been living on its own all this time

I didn't say it would be optimal, just "viable". (Though given scenarios like those sketched by Nick Szabo, my proposal may well be optimal.) Crime could be deterred at an acceptable cost. Nanarchy is not *required* to deal with future crime.

> We might even say that nanarchy is DEFINED to be a stable, self-enforcing political system which provides minimal guarantees against coercion.

We might define socialism as successful central planning too. But it is fairer to define nanarchy as a centralized *attempt* to stop coercion, an attempt that could also go terribly wrong (system prevents anyone touching anyone else, system taken over by despots, etc.)

>> Mark clearly thinks that colonization of the universe naturally leads to a single military power ... Sure, "whoever expands fastest ... gets to have more mass-energy ...". But the same could be said of ordinary economic growth, yet few argue this implies a single power.

current materials, rather than obtaining new raw materials in distant places. Many variables — speed of light lag, physical limits to computation, the memory-time cost of reversibility vs. the memory-time cost of building more radiator surface, etc. Some strategies will lead to better cryptanalysis, trojan horse capabilities, physical weapons, and/or economic capabilities within the market than competitors which choose a different brain/brawn and weapons mix. The chance that we can find, in this practically infinite search space, the best mix for implementing the nano-Gorts, is practically nil.

> Perhaps we could agree that, given enough of a head start, one agent effectively take permanent possession of all non-Solar star systems.

I find it hard to imagine competitive agents with lag times over a few seconds. As computing speed grows, even lags of nanoseconds might become prohibitive. As computing mem*cycles grow, agents may become more distant from each other in subjective lag time, without becoming more distant physically. We might get the "interstellar peace effect" with billions of such posthuman agents here on Earth.

> If for one don't buy Miller's argument that owning

> I do believe that this could easily occur ... nanotech self-reproducing machines could expand from star system to star system. the first seeds to enter a system ... would end up with full possession of ... that system. They would then be in a position to expand outward all the faster. You would have an initial wave of exponential growth (which would settle down to third-power growth) and in such a framework an initial head start could become a numerically overwhelming advantage.... I think the initial exponential growth phase, when all that is going on is claim-staking and colonization, is going to allow only a very small number of competitors.

This makes no sense. With two exponential growing things, it is the one with the larger exponent (time derivative of log of amount) that grows to be biggest. If the exponents are the same the ratio between the two is constant. So I ask again, why should the first colonizer have the biggest exponent? And ordinary economic growth is (at least) exponential, so why doesn't this argument apply to that?

As Nick points out, technological innovation may not run out, so those who stay closer to the center of innovation, and then expand later may grow faster, by using more advanced technology.

>> If at least one of the largest expanding powers were "open", allowing us all to buy shares in it, or to immigrate into it, then there is little risk that the universe will be shut off from us. It is not clear why such a power would allow us to do this.

For the same reason that the U.S. should allow immigrants and foreign investors; because they help us grow faster.

Robin Hanson

Date: Thu, 2 Sep 93 23:15:10 -0700
From: drexler@netcom.com (K. Eric Drexler)
Subject: WAR/NANO/LAW, "nanarchy"

Some comments regarding Robin Hanson's recent message regarding "nanarchy":

On terminology: "nanarchy" is a cute but misleading term. I didn't originate it and haven't been using it, so please don't blame it on me.

An idea that I think worth exploring is the use of automated systems to provide a stable framework for security (in a military sense and perhaps a police sense). In the absence of some idea of how a future political system could legitimately decide to violate certain basic principles, one might attempt to build those principles into the system and then throw away the key. These ideas are touched on, rather gingerly, in *Engines of Creation*.

Some specific responses to Robin's

comments:

>The risk that any anarchy we might create now could evolve into something else, and then evolve again, must be weighed against the risk that any "permanent" solution may turn out to be terribly wrong. Unless the permanent solution is clearly better than the "average" expected future political regimes, risk averse folks should prefer the mixture.

If we are in a biological (predation-based) rather than a market (trade-based) ecology, major evolutionary steps are likely to kill us. A risk-averse person might prefer enforcement of conditions that stabilize a market ecology, for example, suppressing the transfer of resources by forcible seizure.

>If we are considering whether to offer our political support to some growing "movement" in favor of nanarchy, we must consider the possibility that this movement will get out of hand, implementing something that looks to most people like nanarchy, but not to us.

If technological means emerge for projecting military and police power with highly automated systems, then it is likely that they will be used in some manner. It seems prudent to formulate a picture of how they might be used beneficially, or at least less destructively than when military and police power has been subject to discretionary political control.

>If this nanarchy must be designed and implemented, and a global political consensus formed in its favor, all before the first wave of interstellar colonization, there is likely not enough time left.

My expectation is that some political entity will (for a time) be able to dominate the world, and will be terrified of the consequences of not doing so, because of the risks associated with an arms race arising in a more symmetrical situation. If this happens, then it would be desirable to have a clear understanding of how this power could be relinquished without turning it over to potential enemies.

>If technology continues to improve, it's not clear how a nanarchy built on old technology could prevent more advanced attempts at "coercion". If technological improvement must "run out" before nanarchy, then this seems unlikely before the first interstellar colonization.

If machine intelligence systems can perform a million years of R&D per calendar year (and it seems they can), then it may well be that a good understanding of the limits of military technology can be developed rather quickly.

>Now let's get to the central point. Mark clearly thinks that colonization of the universe naturally leads to a single military power....

We will seemingly face a time when multiple technological capabilities will expand by orders of magnitude, quite

rapidly, and in a world closely coupled by transportation systems. Why did tiny Britain rule so much of the Earth in the late 1800s? Largely because it had ships and was first out the starting gate in the Industrial Revolution — a relatively slow and small transition in technology. To understand what may be ahead, imagine a history in which the Industrial Revolution had faster payoffs: in which Britain had built aircraft carriers and a substantial nuclear arsenal (and so forth) before other nations managed to build a steam locomotive. Dominance by a single power (or coalition) during the next revolution is not certain, but would be unsurprising.

>If at least one of the largest expanding powers were "open", allowing us all to buy shares in it, or to immigrate into it, then there is little risk that the universe will be shut off from us.

Yes indeed: if the dice fall the right way, there is (then) little risk.

>I encourage one or both of them to now submit their ideas, in detail, for critical scrutiny by a larger community.

Actually, the ideas are insufficiently detailed, at present, to be submitted in detail to anyone. I've spent much of the last ten years trying to explain simple molecular machines, to provide a basis for understanding just how large the coming jump in technology will be. I sympathize with the view that large jumps should be discounted based on historical experience and the prevalence of false alarms ("Boys have cried wolf."), but I am persuaded that, this time, we face one. I would encourage Robin to present ideas for addressing issues of short-term and long-term military stability during and after a rather abrupt transition to a world with molecular manufacturing and machine intelligence, without begging the question by assuming that the effects of this transition will necessarily be small or gradual.

To restate a basic motivating problem:

It is plausible that a political entity or coalition will achieve unilateral military dominance based on a technology so different from today's that past military experience provides no basis for predicting the stability of a multilateral competition. With this (absolutely corrupting) power in hand, how can that political entity or coalition relinquish its power safely? Who could it trust?

To state a motivating question:

Assume that one favors constitutions and law over dicta and force. Assume that we will find ourselves in a world containing entities far more stable and predictable than human beings and able to think orders

of magnitude faster. If constitutional and legal systems are, ideally, impersonal systems of rules and enforcement mechanisms, should one insist that they forever be structured so as to depend on the decisions (or whims) of persons? If so, why?

Date: Fri, 3 Sep 93 13:17:29 PDT
From: Robin Hanson
<hanson@ptolemy.arc.nasa.gov>
Subject: WAR/NANO/LAW: "nanarchy" and war

I'm not sure Eric Drexler and Mark Miller are worried about the same problems, or envisioning the same answer. So I will try to let Eric's comments stand on their own.

Where to start? First let me say that I can certainly imagine treaties between suspicious military powers which are enforced in part by automated systems, and that a single military power with internal divisions might use similar techniques. Specifically, I can imagine a course-grain automated monitoring system, broadcasting the situation at many militarily strategic points to many military powers (or distrusting internal organizations). This would require enough monitoring sites to detect large scale military movements, but not enough to see who stepped on your geraniums.

Even here I find it hard to imagine throwing away the key, though I could see requiring a high degree of unanimity to make changes. I could also see automated systems to manage certain defensive functions, such as to shoot anything that crosses a certain "no mans land". But before I'd even want to think about a permanent very fine grain system to not just monitor, or even punish, but prevent most forms of coercion, I'd want to have seen lots of experience with smaller scale systems, gradually taking on more and more responsibility. You know, like automated fences that know when to let you escape a burning house and that don't try to shoot flood waters.

I too am concerned about the prospect of a single military power; Hitler got too close for comfort. But I don't understand why, in a nanotech era, a single power should be so much more terrified of breaking up into multilateral competition than they would be now.

And I have great problems swallowing Eric's extreme sudden transition scenario. I have tried to keep up with nanotech issues for many years, and greatly admire Eric's efforts to elaborate a plausible and detailed image of how advanced nanotechnology could work. But the publicly visible efforts by Eric and the Foresight Institute have largely ignored key policy issues like estimating the speed or scope of a nanotech transition. I recall no analyses, offered for critical public scrutiny, which suggest such an extreme transition.

If I were to guess, I'd say Eric thinks that soon after replicators one could easily create many cubic meters of nanocomputers, and that within a few years such computers would naturally become advanced AIs, who could then build cubic kilometers of nanocomputers, and then the game is up. I think that AI (and even huge nanocomputers) is much harder than this, and therefore expect uploading well before AIs, slower more incremental growth of nanotech economies and armies, and that there may never be other things that think millions of times faster than uploads.

Eric asks "just what terribly-wrong outcome should we fear?". As I said before, I fear a system trying to prevent too many useful actions it could not tell from potential coercion, it costing too much and looking too ugly, it preventing us from using more familiar punishments to deter types of coercion the system doesn't cover, and most of all the system being taken over by despots.

Eric "would encourage Robin to present ideas for addressing issues of short-term and long-term military stability". I focus on imagining the folks in some region trying to contract for defense services, and looking for good indicators that the folks they contract with won't enslave them or roll over should someone try to invade. My best idea there is for them to look at betting markets on this question, where the market speculators are in distant places and so are not threatened by a bad local outcome. This is not much help, though, if a single military power is the clear global military equilibrium.

Robin Hanson

P.S. I simply don't accept the premise that "constitutional and legal systems are, ideally, impersonal systems of rules and enforcement mechanisms" and therefore ideally would not "depend on the decisions (or whims) of persons".

Date: Mon, 6 Sep 93 23:45:48 -0700
From: jamie@netcom.com (Jamie Dinkelacker)
Subject: WAR/NANO/LAW: Present tense / future imperfect

There's an old adage which says that it isn't true that everyone wants to write the great American novel — everyone wants to have written it. And then collect the rewards after the work is already done. Often, conversations about nanotechnology sound similar. The notion

of "after the Singularity" is much the same as "after the revolution" or "at the Second Coming" or "after the aliens land" ... what have you. The future promised by nanotechnology offers significant hope for the human condition, but will it happen? On what time scale? If so, who will live to see it? During the transition, who are to be targets? What can we do TODAY to foster beneficial outcomes for humanity, and save our skins in the process (at least until uploading)?

The "fast" vs. "slow" onset of nanotechnology is an entertaining discussion point for many enthusiasts. Often, it seems as if those who are most learned and experienced with the technical specifics of nanotechnology anticipate the fast scenario, whereas those who are most learned and experienced with the human and social dimensions harbor expectations of the slower emergence, and sense lurkers in the shadows.

These perspectives may not actually be in opposition to each other, but result from contrasting foci of the discussants, which may well come down to technical vs. social distinctions. This is similar to the difference in some measure between a laboratory demonstration at a conceptual level, and something that actually works in the marketplace. At times, this is seen with pharmaceutical development which may proceed rapidly, but face a tortuous journey of delays through the FDA, negotiating liability aspects, and getting accepted in the market. A broad and deep intellectual chasm separates the technical capacity that will generate the first assembler and its broad-based impact on our economic structure. The time between these two is a point of much speculation and contrasting opinions. This interval may well be a transition of immense risk.

Programming, tool development and design will have to pass far beyond the current state of the Merkle-Drexler bearings before the assembler (whenever it gets here) is actually of use in constructing workable nano-computers, cell-repair mechanisms, diamondoid teeth and the like. These objects are along the development path, but for the time being (and we don't know how long that time is to be) the design and programming is to be done by humans, who have their own concerns, distractions and need to eat & make a living. Who pays the developers, the resources they can amass, and alliances they can build will have determinant effects on the evolution of nanotechnology. Who they alienate can also be a major factor.

continued on p.39.

EXTRO 1

The First Extropy Institute Conference on Transhumanist Thought

Sunnyvale, California, April 30 - May 1 1994

PURPOSE: Extro 1 will be a rich, intellectually invigorating gathering designed to help push outward the boundaries of progress and possibility. It will be both a serious study and a joyful celebration of humanity's limitless potential and how it will be achieved. Besides presentations of accepted papers, the conference will feature lectures by leading thinkers, panel discussions, the first Extropy Awards banquet, and other events.

LOCATION:

The ballroom at the Sunnyvale Sheraton, in Sunnyvale, California.

TIMES:

Saturday April 30, 8am-9am: Registration and welcome.

Saturday, 9am-8pm: Sessions.

Saturday 8.0-10.0pm: Banquet, Extropy Award Ceremony.

Sunday, 8.45am-1.0pm: Sessions.

(Pre- and post-Extro parties can be expected too.)

SESSIONS:

Only those sessions currently (Jan 7) certain are listed here. Many others are under consideration and development. We will not announce our invited speaker until after confirmation.

The Extropians E-mail List: Past, Present, & Future

Is There an Extropian Epistemology? Pan-critical rationality and the Extropian Principles

SIMNET — a neural network simulator for modeling complex dynamical systems

Cryptographic Techniques for Resuscitation of Biostasis Patients

Keynote speaker

Extropy Awards presentation

ACCOMMODATION:

Extro 1 attendees are responsible for making their own accommodation arrangements. If you are on tight budget, you might try asking local Extropians to take you in (but don't expect this for free). Bay Area Extropians can be contacted via Exl's local events e-mail list: exibay@gnu.ai.mit.edu

ATTENDANCE FEES

	Before Mar 1	After Mar 1
Exl Members	\$60	\$70
Non-members	\$75	\$85

These rates include attendance at all talks and panel discussions, one copy of the conference Proceedings volume (which will include considerably more than just the papers delivered at the event), and light refreshments (coffee, juice, fruit), but not the banquet. (Banquet meals will start at \$15.)

REGISTRATION DETAILS:

You may register immediately by mailing your payment (check, cash, or money order in U.S. currency) to:

Extropy Institute — Extro 1
11860 Magnolia Avenue, Suite R
Riverside, CA 92503-4911

Extropy Institute members will see further developments reported in *Exponent* — Exl's members' newsletter. To receive further information on Extro 1 as it becomes available, write to the same address, or phone 909-688-2323.

PROCEEDINGS:

Extro 1 Proceedings will be available at the event and afterwards. If you do not expect to attend but want to reserve a copy, contact us by April 1.

In addition, the Extro 1 team will make available audio and video tapes of the proceedings.

TABLE SPACE:

If you or your organization would like table space at Extro 1, contact the Extropy Institute office as soon as possible.

QUESTIONS & SUGGESTIONS:

Contact either Max More at the Exl office, or at more@extropy.org, or Derek Zahn at derek@cs.wisc.edu

Wormhole Warfare

From: Robin Hanson <hanson>
To: Extropians@gnu.ai.mit.edu
Date: Thu, 17 Jun 93 17:13:06 PDT
Subject: TECH/WAR: Wormhole Wars

Technology changes the face of war. How would wormholes change war?

First, I'd expect defensive redundant booby-trapping of wormholes connecting potential enemy regions. Wormholes are the major transportation and communication channels; folks would invade along them if they could, so if limited in number they would be choke points — fortified against the most advanced invaders one could imagine.

Second, I'd expect military powers to try and control the entry of wormholes into their territory. If war breaks out, and the enemy has lots of wormholes behind your lines, close to targets and to raw materials, they can see what you're doing and hit you fast. Bad news.

So I'd expect mainly bit streams to go through official wormholes; wormhole passage through wormholes would be tightly controlled, if they could manage it. And even bit streams can be dangerous; once aliens had connected up from across the universe, it might be most unwise to run unknown complex software from distant lands, as in Vinge's *A Fire Upon the Deep*.

Regions with too many unknown wormholes in it might be dead zones, the sort of place no one could plausibly defend because attack could literally come from anywhere in great force. Neighboring regions might want to explode a quasar there or something to try and limit the threat of invasion from that direction.

Third, regions which, for the same "empire" or "universal" time, are at an earlier cosmological co-moving time would have strong military advantages. Say war breaks out at some empire time, and existing wormholes are sealed against attack. In this case the "earlier" region can send a cloud of wormholes toward their enemies the old-fashioned way, on rockets, to arrive rather soon in empire time. If any of the wormhole cloud gets through, a beachhead is formed for attack. Similar holes sent the other way would likely be quickly destroyed by threatening to form

causal loops, and even if they didn't they would take a *very* long time in empire years to get there.

If warring regions have empire times at similar cosmological times, as in the meeting aliens example, and wormhole access is denied, and technological/economic growth is at all in force, then defenders have a huge advantage because they can just wait and grow, as Mike Price commented in his paper.

So the major links between and within civilizations might be under tight military control, new additions to the network subject to military veto, with regions at the geographic center of an empire having a strong military advantage. "Empire" doesn't sound so far-fetched in this case.

Robin Hanson

From: "" <dsg@staff.tc.umn.edu>
To: Extropians@gnu.ai.mit.edu
Date: Thu, 17 Jun 93 19:49:02 CST
Subject: TECH/WAR: Wormhole Wars

On Thu, 17 Jun 93 17:13:06 PDT, Robin Hanson wrote:

>Technology changes the face of war. How would wormholes change war? First, I'd expect defensive redundant booby-trapping of wormholes connecting potential enemy regions. Wormholes are the major transportation and communication channels; folks would invade along them if they could, so if limited in number they would be choke points - fortified against the most advanced invaders one could imagine.

You seem to be assuming each polity/society/culture is on one side of a wormhole only. Try this situation: There is a wormhole which is either natural or a relic of a long-dead civilization. On both sides, it's in relatively isolated areas. The first people to find it and make use of it have several hundred years to build around both ends of it, before the first contact with anyone else.

For them, it's not a danger. It's a supply line; and if all else fails, a line of retreat.

Dan Goodman dsg@staff.tc.umn.edu

From: Robin Hanson
<hanson@ptolemy.arc.nasa.gov>
To: Extropians@gnu.ai.mit.edu

Date: Fri, 18 Jun 93 14:37:14 PDT
Subject: TECH/WAR: Wormhole Wars

Dan Goodman writes:

>Try this situation: There is a wormhole which is either natural or a relic of a long-dead civilization. On both sides, it's in relatively isolated areas. The first people to find it and make use of it have several hundred years to build around both ends of it, before the first contact with anyone else. For them, it's not a danger. It's a supply line; and if all else fails, a line of retreat.

Sure, it's just like the one mountain pass between two parts of a nation. Indeed a supply line and retreat route, and if the other side got it you'd be split in two. So it would be top priority for fortification and defense.

Robin Hanson

Date: Sat, 19 Jun 93 18:48:58 GMT
From: price@price.demon.co.uk
(Michael Clive Price)
Subject: TECH/WAR Wormhole Wars

Robin Hanson wrote:

>Regions with too many unknown wormholes in it might be dead zones, the sort of place no one could plausibly defend because attack could literally come from anywhere in great force.

I think this could be the case with most of the universe, as well, after a while. The enemy only has to sneak one nano-scale wormhole past the defenders, expand and then bootstrap more wormholes through to establish a beachhead. All institutions in space seem vulnerable to this.

High security, corporate and military, establishments may move into basement universes, where they are safer. Only a known number of holes to watch and monitor. If they need more space or matter, they inflate their basement universe a bit more.

>Third, regions which, for the same "empire" or "universal" time, are at an earlier cosmological co-moving time would have strong military advantages. Say war breaks out at some empire time, and existing wormholes are sealed against attack. In this case the "earlier" region can send a cloud of wormholes toward their enemies the old-fashioned way, on rockets, to arrive rather soon in empire time. If any of the wormhole cloud gets through, a beachhead is formed for attack.

The neutral zone between two hostile civilisations is open to this wormhole-rocket attack, from *both* civilisations, since the home-worlds both exist at an earlier cosmological, co-moving time than the n-zone. The n-zone lies close to the future light-cone of both home-worlds.

However the home-world does have this advantage *within* its own empire. If we colonise the galaxy and an empire forms,

expect rule from Sol, not Trantor.

> So the major links between and within civilizations might be under tight military control, new additions to the network subject to military veto[...]

Agreed. You can imagine the public-good arguments. :(

"Private wormholes are a threat to national security!"

A bit like the damnable nuclear non-proliferation treaty, powerful states may seek to exert their control by controlling wormholes.

> "Empire" doesn't sound so far-fetched in this case.

Robin Hanson

Mike Price price@price.demon.co.uk

From: Robin Hanson <hanson>

To: Extropians@gnu.ai.mit.edu

Date: Mon, 21 Jun 93 12:22:35 PDT

Subject: TECH/WAR: Wormhole Wars

I wrote:

> Regions with too many unknown wormholes in it might be dead zones, the sort of place no one could plausibly defend because attack could literally come from anywhere in great force.

Mike Price writes: I think this could be the case with most of the universe, as well, after a while.... High security, corporate and military, establishments may move into basement universes, where they are safer. Only a known number of holes to watch and monitor.

Hmm.. Most of the universe littered with tiny hidden wormholes, and largely unoccupied because it can't be defended. When it is occupied for a moment, it is soon destroyed without bystanders knowing who did it. All activity of consequence in "side" universes. Sounds suspiciously like the universe we see :-). Chilling thought.

Of course very spatially concentrated civilizations would have the same advantages as side universes, the ability to police their border. Activity would have to be very concentrated as well of course — probably look something like a quasar ;-).

Robin

Date: Wed, 23 Jun 93 01:27:32 GMT

From: price@price.demon.co.uk
(Michael Clive Price)

Subject: TECH/WAR: Wormhole Wars

> Of course very spatially concentrated civilizations would have the same advantages as side universes, the ability to police their border.

Difficult to stop a well aimed, massive black hole. Hiding in a basement universe has considerable other advantages as well: You're only visible when you choose to be, and if you don't like the location of your exit portals, you move them somewhere else. And you can be connected to all over the universe as well.

> Activity would have to be very concentrated as well of course - probably look something like a quasar;-).

A beacon to the marauding, berserker hordes.... I think hiding in basement universe has a lot going for it. Of course if you had a comms link to a private basement universe (and kept continuous backups) then you'd get all the advantages of roaming freely with thesecurity. Killed, and you wake up 'back at the ranch'.

> Robin

Mike Price price@price.demon.co.uk

Continued from p.36

While it does seem reasonable to expect a convergence between the general application design efforts and assembler emergence, present day social forces may well conspire to maintain dislocations between the times when general assemblers and useful programs/designs for them will emerge. Investment strategies, organizational savvy, public profiles and industry dynamics are just some of the factors to be considered. Given the general political and economic state of things today, I'm skeptical that the assemblers and utility instructions (e.g., for nanocomputers, deep AI, ...) will simultaneously emerge in a clap of Olympian Thunder. Nevertheless, the perception that they may come together, generates its own threat profile. This could be problematic for some people.

It's not uncommon for people in the general public, upon first learning about nanotechnology, to find fear stirring in their emotions. Not everyone across the planet cheers at the thought of having a force or agency that can, in terms of contemporary lyrics, "turn their brown eyes blue." Fear can often lead to dogmatic and drastic reactions. Angry people get even; fearful people get nasty. Significant public education and involvement would be a wise precaution, but has yet to begin in earnest.

Various nanotech enthusiasts strongly advocate an "open" development of nanotech. But let's keep in mind the fact that Admiral David Jeremiah, vice chair of the Joint Chiefs of Staff in the Pentagon, has already shown interest in future implications of nanotech in security contexts. Others, across the planet, would welcome the opportunity to make Orwell's fantasy a reality. So, it's reasonable to consider that "open" means "in addition to secret" but not "instead of secret," at least in real-world pragmatic terms. That kitty is out of the basket. Whether nanotech is developed in an "open" environment or in both "open" and "secret" environments, clear and

present dangers will likely be perceived by the fearful and by competing forces.

To the extent that nanotech is developed "openly" there will be a broad base of information available about the players and their progress. Assassins bullets, commissioned by the fearful, may well find rich targets among these players who will be easy to locate. For the work done in "secret," those entities fearing a nanohegemony may be making plans at this moment to act preemptively when the time is ripe. If they have nukes or CBW agents, running for cover may not offer protection or benefit.

Consequently, discussion that leads people to understand the positive potentials of nanotechnology might somewhat ease the tensions of the transition. Without key benefits for the public mind to rally around, fear and nervousness may well fill the vacuum. Notions of improved health care, a culture of abundance, or a Diamond Age of materials, are wonderful thoughts (and ones I hope to witness) but they are far too vague to allay the general public's fearful tendencies or block off demonic leaders who may rise to rail against nanotech. As some writers and speakers develop a more near term focus wherein the public opinion is led to anticipate nanotech in clear and specific terms (e.g., the end of death), the potential for successfully making it across this transition from now to the nanotechnology enabled society of whim increases.

Spaceship Earth hurtles on. I recall Paul Saffo remarking "never mistake a clear view for a short distance." If we only look far, far away past the Singularity, it's just possible we may miss clues for successful navigation. This transition is going to be a bitch, in human terms. Presently, it seems that the Extropians list is the only ongoing forum where these issues are actually receiving exposure, dialog and analyses. Upwards and Onwards.

Jamie Dinkelacker, Palo Alto CA
Jamie@netcom.com 415.941.4782

Extropy Institute

The December literary supplement to the *Village Voice* described Extropians as "radical humanist technophiles", and referred to "the movement's combustible mix of fringe academics, over-educated computer programmers, and renegade philosophers". A narrow description, but one that nevertheless hints at some of our concerns and enthusiasms. For those of you for whom this issue is your first real contact with Extropian ideas, the short version of The Extropian Principles to the right will help clarify our shared values and goals. (The full text appeared in last issue.)

Extropy Institute (ExI) was incorporated in 1992 as an educational, tax-exempt organization. Like the Extropians e-mail list, ExI was an outgrowth of *Extropy* (founded 1988). We created ExI in order to provide a structure and a network that would facilitate the spread and development of Extropian ideas, values, and culture.

This organizational mission breaks down into two aspects which together explain all our activities: (a) Within our existing Extropian culture refining and developing our ideas, working together to transform ourselves into transhumans and to evolve a radically new culture free of the irrationalities and limitations of the past. (b) To clearly and persuasively communicate our philosophy of life even to those who are not already attuned to the same ideas and attitudes, in order to influence the broader culture in more extropic directions.

In pursuit of these goals Extropy Institute — though yet limited by a relatively small (though rapidly growing) membership, and continual tightness of funds — continually seeks new outlets for its members energies, abilities, and intellects. Our primary publication, *Extropy: The Journal of Transhumanist Thought* is supplemented by our member's newsletter, *Exponent*, edited by frequent *Extropy* contributor Simon! D. Levy. *Exponent* carries shorter articles, membership information such as forthcoming meetings, reports on progress, and reviews of relevant books, software, and other media.


A variety of meetings take place, such as last summer's *Extropy* 5th birthday party, weekly lunch meetings in the N. California Bay Area, and now monthly Idea Forum discussion meetings in the Los Angeles area. In addition, impromptu get-togethers take place all over the country. This year sees an important new development: The initiation of annual "Extro" conferences (see p.37) where ideas can be explored in depth, and bounced off persons of many different specialities and perspectives.

Supplementing printed publications and physical meetings is the Extropian virtual community. The Extropian cybercommunity continues to expand, now encompassing the main Extropians e-mail list (generating 150-300 kilobytes per day), the ExI Essay list, five local e-mail lists for arranging meetings, parties, and other joint activities, and

Continued on page 41

EXTROPY INSTITUTE
1 1860 MAGNOLIA AVENUE, SUITE R
RIVERSIDE, CA 92503
909-688-2323

For details of membership dues, see p.2, lower right.



Extropy Institute

EXTROPIAN PRINCIPLES v.2.5 (short version)

- 1. Boundless Expansion** — Seeking more intelligence, wisdom, and effectiveness, an unlimited lifespan, and the removal of political, cultural, biological, and psychological limits to self-actualization and self-realization. Perpetually overcoming constraints on our progress and possibilities. Expanding into the universe and advancing without end.
- 2. Self-Transformation** — Affirming continual moral, intellectual, and physical self-improvement, through reason and critical thinking, personal responsibility, and experimentation. Seeking biological and neurological augmentation.
- 3. Dynamic Optimism** — Fueling dynamic action with positive expectations. Adopting a rational, action-based optimism, shunning both blind faith and stagnant pessimism.
- 4. Intelligent Technology** — Applying science and technology creatively to transcend "natural" limits imposed by our biological heritage, culture, and environment.
- 5. Spontaneous Order** — Supporting decentralized, voluntaristic social coordination processes. Fostering tolerance, diversity, long-term thinking, personal responsibility, and individual liberty.

ExI Directors

Max More, President, Editor of *Extropy*. more@extropy.org
Tom Morrow, Vice President. t0morrow@aol.com
Simon! D. Levy, Editor of *Exponent*. levy@yalehask.bitnet
Tanya Jones, Treasurer. tanya@alcor.org
Ralph Whelan, Secretary. ralph@alcor.org
David Krieger. dkrieger@netcom.com
Russell E. Whitaker. whitaker@eternity.demon.co.uk

BACK ISSUES

#1, 2, 4, 5, 6, 7, 8, 9, 10, 11: \$5 each.
Available from Extropy Institute (address, p.2)

#11, Vol.5 No.1 (2nd Half '93): Uploading Consciousness, by Ralph Merkle; Extropian Principles v.2.5, by Max More; Traversable Wormholes: Some Implications or Contact! A Post-Singularity Phase Change, by Michael Price; A Conversation with Mark Miller, Part 2: The Day the Universe Stood Still, by David Krieger; "Bunkrap!": The Abstractions that Lead to Scares About Population and Resources, by Julian L. Simon; Reviews of *Theories of Everything*, *In Our Own Image: Building an Artificial Person*, *Mirror Worlds*.

#10, Vol.4 No.2 (Winter/Spring '93): Pigs in Cyberspace, by Hans Moravec; Protecting Privacy with Electronic Cash, by Hal Finney; Technological Self-Transformation, by Max More; Mark Miller interview, by David Krieger, Pt.1: Creole Physics & the Credit Theory of Identity; Nanocomputers: 21st Century Hypercomputing, by J. Storrs Hall; The Transhuman Taste (Reviews): Two books on Ayn Rand & Objectivism; *Nanosystems*; *Genius*.

#9, Vol.4 No.1 (Summer 1992): The Extropian Principles, 2.0, by Max More; Extropy Institute Launches, by Max More; Persons, Programs, and Uploading Consciousness, by David Ross; Nanotechnology and Faith, by J. Storrs Hall; The Making of a Small World (fiction), by R. Michael Perry; Genetic Algorithms, by Simon! D. Levy; Time Travel and Computing, by Hans Moravec; Futique Neologisms 3; Exercise and Longevity, by Fran Finney; The Transhuman Taste (Reviews): *The Anthropic Cosmological Principle*, *The Blind Watchmaker*, *The Ultimate Resource*, *Population Matters*, *The Resourceful Earth*, *Bionomics*.

#8 Vol.3 No.2 (Winter 1991-92): Idea Futures: Encouraging an Honest Consensus, by Robin Hanson; Dynamic Optimism, by Max More; Neurocomputing 5: Artificial Life, by Simon! D. Levy; Futique Neologisms (futurist lexicon); Extropia: A Home for Our Hopes, by Tom Morrow; Human-Transhuman-Posthuman, by Max More; Reviews of: Stiegler's *David's Sling*, Drexler's *Unbounding the Future*, Platt's *The Silicon Man*; News of scientific advances and movement news; Reviews of zines.

#7 Vol.3 No.1 (Spring 1991): A Memetic Approach to 'Selling' Cryonics, H. Keith Henson & Arel Lucas; Privately Produced Law, Tom Morrow; Order Without Orderers, Max More; Futique Neologisms; Neurocomputing 4: Self-Organization in Artificial Neural Networks, by Simon! D. Levy; Forum on Transhumanism; Reviews of *Smart Pills*, *Surely You're Joking Mr Feynman*, *Great Mambo Chicken and the Transhuman Condition*; and more...

#6 (Summer 1990): Transhumanism: Towards a Futurist Philosophy, by Max More; The Thermodynamics of Death, Michael C. Price; The Opening of the Transhuman Mind, by Mark Plus; The Extropian Principles, by Max More; Neurocomputing Part 3, by Simon! D. Levy; Forum on Arch-Anarchy and Deep Anarchy; Reviews: *Order Out of Chaos*, *The Emperor's New Mind*, *A Neurocomputational Perspective*, *Loompanics Greatest Hits*, *The Machinery of Freedom*; Extropian Resources, and more.

#5 (Winter 1990): Forum: Art and Communication; Leaping the Abyss, by Gregory Benford; Arch-Anarchy, by A; Deep Anarchy, by Max O'Connor; I am a Child, by Fred

Chamberlain; Perceptrons (Neurocomputing 2), by Simon D. Levy; On Competition and Species Loss, by Max O'Connor; A Review of Intoxication, by Rob Michels; Intelligence at Work, by Max O'Connor and Simon D. Levy; Extropian Resources, by Max O'Connor and Tom W. Bell; The Extropian Declaration, by Tom W. Bell and Max O'Connor; Our Enemy, 'The State,' by Max O'Connor and Tom W. Bell.

#4 (Summer 1989): Forum; In Praise of the Devil, by Max O'Connor; Neurocomputing, by Simon D. Levy; Why Monogamy? by Tom. W. Bell; What's Wrong With Death? by Max O'Connor; Reviews: Are You a Transhuman? Postscript to "Morality or Reality" by Max O'Connor; Efficient Aesthetics, by Tom. W. Bell; Intelligence at Work: Advances in Science by Max O'Connor.

#3 (Spring 1989) is out of print.

#2 (Winter 1989): Review of Mind Children, by Max O'Connor; Darwin's Difficulty, by H. Keith Henson and Arel Lucas; A Truly Instant Breakfast, by Steven B. Harris M.D.; Wisdomism, by Tom W. Bell; Nanotechnology News, by Max O'Connor; Weirdness Watch, by Mark E. Potts.

#1 (Fall 1988): A brief overview of extropian philosophy and an introduction to some of the topics we plan to address: AI, Intelligence Increase Technologies, Immortality, Nanotechnology, Spontaneous Orders, Psychochemicals, Extropic Psychology, Morality, Mindfucking, Space Colonization, Libertarian Economics and Politics, Memetics, and Aesthetics; "Morality or Reality," by Max O'Connor.

Continued from page 40

now an Extropian presence in the Metaverse (initiated by Geoff Dale), including virtual offices, meeting places, and apartments. (See the back cover for information on most of these cyberfora.)

What's Coming in 1994?

As announced last issue, *Extropy* will be published quarterly instead of semi-annually to accommodate the swelling tide of Extropian writing. *Exponent* also is now appearing more often so that every month sees the publication either of *Exponent* or *Extropy*. We expect to see *Extropy's* circulation surpass 4,000 and head for 5,000 this year (circulation was only 750 as recently as the summer of '92). We intend to repeat the 160% growth in ExI membership achieved in '93, in part by sending information to thousands of persons on promising mailing lists.

I've already mentioned the forthcoming Extro 1 conference on Transhumanist Thought. This will be accompanied by publication of a substantial Proceedings volume. We will be fostering the growth of more local discussion groups and international chapters of ExI, and we will publish a new introductory booklet about the Institute's purpose and activities. We also look forward to the continued development of the Extropians e-mail list software and culture.

We hope you will join us as an active participant in the Extropian movement!

Forthcoming in *Extropy*

Part Two of A Conversation with Dave Ross

Neural-Computer Interfacing

Two Questions for Extropianism, by Charles Platt;
response by Max More

Neurocomputing: Sequential Networks

Utility Fog (nanotech)

Extropian Values and Beauty

Liliputian Uploads

Transhuman and Posthuman Sexuality

Longevity Diets: Vegetarianism, Veganism, and
High/Low (Walford).

EXTROPY #13 (vol.6, no.2) on sale April 15 1994.

The Transhuman Taste

REVIEWS OF EXTROPIAN INTEREST

Fuzzy Thinking: The New Science of Fuzzy Logic

By Bart Kosko

Hyperion, New York, 1993.
318 pages. ISBN: 1-56282-839-8 \$24.95

Reviewed by Max More

2700 v.1

Seeking to continually improve the power and accuracy of our thinking, we Extropians ceaselessly search for new tools for thought. Fuzzy logic and, more broadly, fuzzy thinking — despite its initially off-putting nomenclature — adds to our intellectual arsenal of stupidity-destroying weapons. *Fuzzy Thinking* goes deeper than McNeill and Freiburger's recent groundbreaking book on the topic, as is to be expected since Dr. Kosko may fairly be described as the leading fuzzy proponent. Kosko teaches classes on neural networks and fuzzy systems at USC, and has written two technical books in the area (*Neural Networks and Fuzzy Systems*, and *Neural Networks for Signal Processing*). *Fuzzy Thinking* is of further interest due to Kosko's explicit advocacy of libertarian values and physical immortality and cryonics, and his discussion of the prospect of nanotechnology and the shape of the future. I first became aware of Kosko in 1988 when he gave a guest lecture to a cognitive science class in the philosophy department at USC (where he did undergraduate work) — a class attended by Tom Morrow and myself.

First introducing the reader to the Fuzzy Principle (Everything is a Matter of Degree), *Fuzzy Thinking* proceeds historically, with sections on The Fuzzy Past, The Fuzzy Present, and The Fuzzy Future. Though much fuzzy theory is the product of the last three decades, it originated in the work of philosophers such as Charles Pierce, Bertrand Russell, and Max Black. Heisenberg's uncertainty principle prompted logicians such as Jan Lukasiewicz in the 1920s and '30s to develop multivalued logic. The invention of fuzzy set theory — a crucial advance — was made by philosopher Max Black in his 1937 paper "Vagueness: An Exercise in Logical Analysis", published in *Philosophy of Science*. Unfortunately Black's untraditional analysis was ignored by other philosophers and scientists, otherwise the field might now be termed 'vague logic'. The name 'fuzzy logic' was bestowed by fuzzy pioneer Lofti Zadeh in 1962. Bart Kosko himself two offers definitions of 'fuzzy logic':

The first meaning is multivalued or 'vague' logic. Everything is a matter of

degree including truth and set membership. This dates back to the turn of the century. The second meaning is reasoning with fuzzy sets or with sets of fuzzy rules. [292]

The choice of the term 'fuzzy' strikes me as an unfortunate strategic error. The connotation of imprecision and wooliness, though unmerited, will stiffen opposition to fuzzy logic's acceptance by the scientific and rationalist communities. This nomenclature probably contributed to the distress evident in this comment by Professor William Kahan: "Fuzzy theory is wrong, wrong, and pernicious. What we need is more logical thinking, not less. The danger of fuzzy logic is that it will encourage the sort of imprecise thinking that has brought us so much trouble. Fuzzy logic is the cocaine of science." I would be tempted to reply that many scientists need some kind of stimulant to their thinking.

Compounding the obstacle raised by the field's name is Kosko's linking of it with Eastern philosophy. While Eastern philosophy indeed has — to a minor extent — anticipated the field, emphasizing the connection will repel many Occidental-centric rationalists. Prior to actually reading the book they are likely to see it as another ludicrous comparison in the tradition of Fritjof Capra's linkage of Eastern mysticism to quantum physics in *The Tao of Physics* (a book I cannot recommend too lowly).

The Eastern connection enters by way of the "Aristotle vs. the Buddha" theme. Aristotle invented Western bivalent logic, and represents that tradition. According to Aristotelian bivalence, everything is either A or not-A — either grass is green or is not green. Nothing can both have an attribute and not have it at the same time. As Kosko argues, this view has rarely been questioned and is built into the thinking of practically all Western scientists and philosophers. (There are exceptions: Philosopher Derek Parfit, in *Reasons and Persons*, bases his personal identity theory on a relation of psychological connectedness that holds to varying degrees.) The Buddha is taken to represent the idea, never developed formally in the East, that truth is a matter of

degree, that a thing can be A and not-A to varying degrees. A man may be tall and not-tall to a certain extent. The fuzzy view allows that not only is the part contained in the whole, but the whole is contained in the part — to some degree ranging from 0% to 100%. These conflicting views are revealed by these two authoritative statements:

Everything must either be or not be, whether in the present or the future. (Aristotle, *De Interpretatione*)

I have not explained that the world is eternal or not eternal. I have not explained that the world is finite or not finite.

(The Buddha, *Majjhima-Nikaya*)


Subsethood appears able to take the place of probability in explaining phenomena, and avoids the need to see probability as a metaphysical rather than an epistemic notion (i.e., avoids seeing things as having a probability, rather than probability being a statement about our knowledge of things). Kosko shows that the degree of subsethood equals the conditional probability of traditional theory. "What is the probability of success? The degree to which all trials are successful, the degree to which the set of successful trials contains the set of all trials... In general the probability of a set or event A equals the degree to which the part A contains the 'sample space' X."

Randomness and metaphysical probability, Kosko argues, can be dispensed with.

And Einstein looks right again: God need not play dice. The universe is not random. You can take it one step deeper and get rid of the "randomness." The universe is deterministic but gray. Chaos theory had already gotten the determination part right. Fuzzy theory now confirmed that and [showed] that all things were matters of degree too." [63]

Some care should be exercised in interpreting a statement such as "the universe is deterministic but gray." At the quantum level, entities may truly be fuzzy themselves. Here, fuzzy logic can be applied directly, as has already been done (for instance by Hilary Putnam in his discussion of quantum logic as applied to the Measurement Problem. At the macro level things themselves are determinate and unfuzzy. Individual objects have a determinate mass, number of atoms, positions, and so on (unless they are things like clouds...). Fuzziness appears when we consider attributes or qualities of things, and when we consider concepts of things-of-a-type. In the case of attributes ("This apple is red") our descriptions are fuzzily true or false (the apple may be mostly red but also partly green). In the case of concepts of types of things or actions (including scientific concepts), fuzziness is revealed when we realize there are many borderline cases. Our concepts are fuzzy sets. I like one of Kosko's illustrations of the fuzziness of an action: "Touch your mother's toe. Is that incest or not incest? Touch her ankle, her shin, her knee. Is that incest? And so on up."

Having introduced the Fuzzy Principle, in the Fuzzy Present, Kosko explains that fuzzy logic is reasoning with fuzzy sets. This section begins by arguing that mathematics is just a limiting case of fuzzy set theory, then develops the idea of fuzzy entropy as a measure of fuzziness (measured in fuzzy units or *fits* as distinguished from binary units or *bits*). This section of the book contains plenty of solid material, such as an accessible presentation of the Fuzzy Approximation Theorem (FAT), Fuzzy Associative Memory, adaptive fuzzy systems, and fuzzy cognitive maps. Kosko demonstrates that fuzzy logic, rather



than being merely a fascinating idea, has immediate practical applications, providing a list of patented applications such as fuzzy systems that control anti-lock brakes, eliminate hand-jitter from cameras, regulate the mixing of chemicals, and one that stabilizes a helicopter in flight when it loses one of its rotor blades — a feat unmatched by any human operator or math model.

"Chapter 11: Adaptive Fuzzy Systems" clearly explains the tight connection between fuzzy systems and neural networks. Anyone familiar with neurocomputing will immediately see how fuzzy rules (or principles) and the Fuzzy Approximation Theorem are interwoven with connectionist systems. Fuzzy reasoning already has increased machine IQ, allowing computers to adapt, learn, and recognize far more effectively than allowed by traditional rigid, rule-based, symbol-manipulating AI approaches (what Fodor calls GOF AI — Good Old-Fashioned AI). Fuzzy Thinking (or Kosko's textbook, *Neural Networks and Fuzzy Systems*) combine with neurophilosopher Paul Churchland's brilliant *A Neurocomputational Perspective* to illuminate the power and promise of this approach to machine intelligence and neuroscience.

The final section of the book, *The Fuzzy Future*, applies fuzzy thinking to issues of life and death, the "social contract", why the universe exists, and our future. In *Life and Death*, Kosko analyzes the fuzzy nature of life and death, illustrating the latter by introducing many readers to cryonics and nanotechnology. Kosko boldly comes out and states that he himself has made arrangements for being cryonically suspended. "Chapter 15: Man and God", being crammed with several fascinating and stimulating discussions, left me hungry for more. Included here are speculations

as to how mathematics may *require* that the universe exist, the effects on us of widely used smart drugs and smart weapons, and the prospects of advanced machine intelligence and indefinite lifespans.

I would not have fulfilled my duty as a reviewer if I failed to pick out a few errors that detract from the overall excellence of this book. Kosko offhandedly claims that "Most modern philosophers are behaviorists" (82) — an assertion that I've heard from others but which clashes with my experience. Though true in the '50s and '60s, philosophy has moved on, most contemporary philosophers esposing functionalism (a view which developed out of behaviorism but which allows room for the causal power of internal cognitive states in addition to sensory input and behavioral output).

I also take issue with Kosko's equation of a coherence theory of truth with logical truth: "Coherent truth is empty if achieved, and self-contradictory if not achieved..." Certainly this is correct if coherence is restricted to the domain of logical and mathematical truths (excluding applied geometry), but coherence theories apply also to factual and moral questions. A coherence theory holds that the justifiability of a belief depends on the other beliefs one holds, and denies that there is any special set of beliefs (such as those based on sense perception) that are certain and that provide foundations for other beliefs.

I wish to emphasize that *Fuzzy Thinking* is not a dry, if trenchant, study. The book is effectively balanced by absorbing personal anecdotes and vignettes of leading figures in fuzzy research, such as Professor Takeshi Yamakawa — a fifth dan black belt in Shotokan karate, fuzzy chip designer, and leader of the Fuzzy Logic Systems Institute in south Japan.

Although irrelevant to the value of the ideas, Extropian readers will appreciate Kosko's personality as it emerges throughout the book (and as reported in McNeill and Freiberger's earlier book). Clearly Kosko is extropian: Highly intelligent, self-confident, physically fit, and an iconoclast. Far from being the vague, soft-headed, lazy person that might be connoted by the word 'fuzzy', Kosko comes across as thoroughly undecadent — hardworking, innovative, and disciplined.

Dr. Kosko's exceptional ideas provide powerful new tools for thought, to the delight of anyone committed to intensifying their intelligence and enhancing their rationality. Lucid, incisive, and startling, *Fuzzy Thinking* should spur a productive re-evaluation of methodology in numerous fields of science, and illuminate some of the fundamental questions about life that we all ponder. Expect to see widespread and virulent opposition to fuzzy logic from those who choose tradition over transformation and agreement over advancement.

The Children's Machine: ***Rethinking School in the Age of the Computer***

by Seymour Papert

Published by Basic Books, a division of Harper Collins
Publishers, Inc. New York, 1993.
230 pages; ISBN: 0-465-01830-0

Reviewed by Harry S. Hawk

Seymour Papert is one of the more well known faculty members of the MIT Media Lab. He is also known for his work with Logo and Lego Logo, the computer language/learning environment. His latest book, "The Children's Machine," focuses on many aspects of society that are of interest to Extropians. This includes the evolution of technology, and the process of learning. He writes, "this book focuses on... How does the relationship between children and computers effect learning," adding, "Understanding this relationship will be crucial to our ability to shape the future." He is optimistic because he sees "synergy between the technological revolution and the epistemological changes in how we think about knowledge."

Michael Rothschild in *Bionomics*, puts forth the following idea: Faster modes of learning lead directly to economic efficiencies and profits. Papert acknowledges this theme when he writes, "the Japanese success is exactly the ability responsible for America's past success — the willingness to learn. Complainers would do well to relearn from the Japanese the skill of learning, at which America was once the world's champion." Papert has made his career in learning about how young children learn.

Papert makes an insightful point about how our society has been changed by technology. He uses an example of an operating room, and the technology found in that operating room, and how much that technology has changed in the last 200 years. Then he compares the technology and other aspects of school rooms. He feels, and I agree, the technology has changed very little in the same time period. He feels that this a major problem facing America.

Papert asks, "Why, through a period when so much human activity has been revolutionized, have we not seen comparable change in the ways we help our children learn?" Extropians can easily answer this question, "government control of our schools." Papert might lump Extropians in the group he calls, "the Yearners, who respond by citing impedi-

ments to change in education such as costs, politics, the immense power of the vested interests of school bureaucrats..." He has found another group, "the Schoolers [who] are taken aback by the suggestion for megachange [in schooling]." Papert who is clearly a Yearner wants to overcome the difficulties associated with bringing change to our schools.

The Children's Machine is Papert's journey through the education landscape; it is his call for action for megachange in our education system. While Extropians might prefer a frontal attack on governmental control and entrenched bureaucrats through voucher systems or tax credits, they should find Papert's work quite rewarding, even though he doesn't call for any of these measures. He is seeking to radically alter school curricula through technology. The book is divided into 10 chapters: Yearners and Schools; Personal Thinking; School: Change and Resistance to Change; Teachers; A Word for Learning; An Anthology of Learning Stories; Instructionism versus Constructionism; Computerists; Cybernetics; and What Can be Done? What is interesting about Papert's ideas is that those who are interested in them don't have to wait for the government to implement them. We can use them to teach our children and perhaps ourselves. However, Papert is a strong supporter of pluralist, democratic public education and is concerned that the changes he is proposing might "first enhance the lives of the children of the wealthy and powerful." He hopes the changes will first occur in the public school system. Given that private school students often seem better motivated with the traditional approaches, a public school implementation might have the most impact. However, I feel that Papert's ideas are too good to wait until the public school sector embraces them.

Papert coins a new word — *mathetic* — and then uses this word throughout the book. He was looking for a word to denote learning as heuristics denotes problem solving. He revisited a "family of Greek words related to learning... ma[king] restitution for a semantic theft perpetrated by [those]... who stole the

word mathematics...*mathematikos* meant 'disposed to learn,' *mathema* was 'a lesson,' and *menthanien* was the verb 'to learn.' "

A central theme of mathetic based education is constructing environments where students are self-motivated to learn because of their interactions with an environment. Certainly that is how many of us learn. For example, I observed Perry Metzger using PERL, and I wanted to play with it too. In playing with PERL I learned about not only programming, but how to run a mailing list, and about other aspects of the Internet. It is unlikely I would have acquired so much knowledge so quickly in a formal learning environment. Indeed Papert rejects the formal; Papert disputes the view that there is a single method of teaching a single subject. That makes sense on several levels. First, given the memetic differences among individuals it makes sense that individuals will behave as individuals. That the most effective learning mode for them will be a learning environment that focuses (or rather allows them to focus) on their own interests. That is the essence of Papert's view.

Papert introduces John Dewey who, "began his campaign for a more active and self-directed style of learning in schools over a hundred years ago... Dewey remains a hero to those who believe in a twentieth-century vision of a child as a person with the right to intellectual self-determination." He likes to use video games as one example of an "ideal." He says, "School would have parents... believe that children love them [video games] and dislike homework because the first is easy and the second is hard. In reality, the reverse is more often true... These toys, by empowering children to test out ideas about working within prefixed rules and structures in a way few other toys are capable of doing, have proved capable of teaching students about the possibilities and drawbacks of a newly presented system in ways many adults should envy. Video games teach ... [us] that some forms of learning are fast-paced, immensely compelling, and rewarding. By comparison school strikes many young people as slow, boring and frankly out of touch."

As individualists, I think many if not most Extropian-minded persons can relate to the idea of intellectual self-determination. However, based on my conversations with many Extropians, I am unsure how many would endorse this idea for children (biological, mind, or other...). I have certainly heard many call for a return to some sort of classical education, you know, "Education the way it used to be when it was good!; force children to learn Latin, etc." Papert would strongly disagree; he wants to turn education upside down — and I agree with him. He writes, "there can be little doubt that a child treated with respect and encouragement rather than threatened with rejection and punishment will fare better..."

Papert is trying to create such environments for learning. In Lego Logo and video games there is no right and wrong¹, there are only techniques to be learned and rules to be discovered rather than dogma being preached and reinforced by teachers. Papert writes, "I have always yearned for ways of learning in which children act as creators rather than consumer of knowledge... There is a family resemblance ... between the vision of learning I am presenting here and certain philosophical principles expressed in the diverse forms of innovations that go under such names as progressive or open or child-centered or constructivist or radical education."

A critical issue is how do we teach subjects that many consider highly formal, like mathematics. Papert states,

My goal became to create an environment in which children could learn algebra and geometry... in ways more like the

informal learning of the unschooled toddler or the exceptional child than the educational process followed in schools... [For example] every preschool child has amassed on his or her own special mathematical knowledge about quantities², about space, about the reliability of various reasoning processes, elements that will be useful later in the math class... The central problem for math education is to find ways to draw on this vast experience... [to] construct microworlds in which children pursue mathematical activity... giving children the opportunity to learn and use mathematics in a nonformalized way of knowing [that] encourages rather than inhibits the eventual adoption of a formalized way.³

Papert is clear that these nonformalized methods can't be tricks: "the point of developing nonformalized ways of knowing... is entirely subverted if these are conceived as... a trick to lure children into formalized instruction. They have to be valued for themselves and genuinely useful to the learner in and of themselves." He provides an example of a class programming with turtle geometry⁴ African textile designs. All of what they learned can be used to build a formal knowledge of geometry *but* the purpose in which they used the turtle geometry for was to simply learn about African textile designs. Seymour says, "Geometry is not there for being learned. It is there for being used." He finds that since different students learn differently they must be 'taught' differently; some will thrive with rigid formalism while others will thrive when allowed to follow their educational 'muse.' I

don't think anything could be more extropic that the acknowledgment of the uniqueness of an individual. Papert describes a few of the programming projects undertaken with Lego Logo and the results are very interesting. Students learning math and engineering skills while they are building "trucks, robots, and houses."

He talks about Schoolers having immunological reactions to his methods; I can understand that. The meta point about Papert's call for megachange is that not only does he want to challenge how lessons are taught, but

A central theme of mathetic based education is constructing environments where students are self-motivated to learn because of their interactions with an environment.

he also wants to change the subjects. He is looking for a big "departure from the [current] curriculum." He wants us to question, "not only how a school teaches but what as well." Papert is also a fan of what I call Integrationism — that many things are related and do not have to be divided and taught separately⁵.

Surprisingly, by the end of the book, he does come out strongly against large bureaucrat school administrations. He also compares national standardized testing and the curriculum to support it to GOSPLAN in Russia under the communists. In fact he compares all hard inflexible curricula to GOSPLAN. What he calls for, ultimately, is a learning environment that is almost like a market for ideas and assignments. He is saying that in the right environment children will order themselves and their learning spontaneously. That certainly sounds extropian. It should be noted that he is not proposing, directly, such an educational system for high school students but rather for grammar school students. He is saying before we teach children to read and count by rote, we should first put them in an environment where they learn how to learn and more importantly learn how to think and reason with logic. He wants teachers who are not specialized technicians but rather facilitators of knowledge; he wants students to bring and retain their innate desire for fun⁶ into the class room and harness it to turn them into world class learners.

I highly recommend this book. Any Extropian who is interested in either improving educational instruction and is interested the subject of how children learn will find it especially useful.

[Notes on next page]

- EXPANSION TOOLS -

Mind Matters - how thought becomes reality

ISBN 1-884318-00-2 Book. Suggested donation \$7 USA, \$9 (US) outside USA.

What we are, how we do what we do, and what happens. Compiled fundamentals and axioms for better understanding. Intro to Manifestational Chaos Theory. Re-defined: creator, individual, randomness, emotion, ethics, etc.

The Effectiveness Profile

ISBN 1-884318-02-9 (3½") ISBN 1-884318-03-7 (5¼") PC-DOS software. \$33 suggested donation.

Computer-assisted diagnostic and progress-check tool. Enter responses to 100 multiple-choice propositions. Program prints a graph of relationships between 20 Attributes in 6 categories.

Decisions: perspective and process in ethics

ISBN 1-884318-01-0 Book. Suggested donation \$11 USA, \$14 (US) outside USA.

Theory and application for improving conditions. Step-by-step procedures that integrate management by goals, principles, performance statistics and much more. Un-militaristic, non-adversarial, respect-based, affirmative.

BlueFlame Impressions PO Box 390696 Mountain View CA 94039

¹ Right and Wrong here means that there is not one true path - no educational dogma; in Super Mario Bros. there are many paths to success; a teacher cannot "put down" a student because they have solved the puzzles in a non-standard way.

² A favorite example of this, for me, is the number of "inner city" youth who once involved in some way with the drug trade, master a basic understanding of metric measurements, and are able to convert between pounds, ounces, and grams. Clearly, if placed in a formal learning environment these individuals would not be able to learn the same material with out a great increase of effort.

³ We see this in Vinge's "The Ungoverned" where Paul observes Willie playing a video game that requires the player to understand gravitational influences of planets, etc. and other aspects of physics.

⁴ part of the Logo environment

⁵ An example of this is Extropianism itself, which brings together many areas of interest that were previously not seen as interrelated. Learning about one aspect of Extropianism can help you learn about others. That is Papert's point - you don't have to formally study math to learn math.

⁶ If any one is interested, I would be happy to discuss how I feel a K to 12 school could educate and help its students to learn by *only* focusing on Baseball.

classifieds

Cross verbal swords in **The (Libertarian) Connection**, open-forum magazine since 1968. Subscribers may insert one page/issue free, unedited. Lots of stimulating discussion. Eight issues (one year) \$20. Strauss, Box 3343X, Fairfax, VA 22038.

Open relationships, group marriage, intimate friendship. Interested? Discreetly contact others. SASE for details. Current issue \$4.00. Touchpoint, P.O. Box 408-EX, Chloride, AZ 86431

Dynamic Optimists – NOT!

"What can be palpably absurd than the prospect held out of locomotives travelling twice as fast as stagecoaches?"
The Quarterly Review, 1825.

"Landing and moving around the moon offers so many serious problems for human beings that it may take science another 200 years to lick them." *Science Digest*, August 1948.

"The energy produced by the breaking down of the atom is a very poor kind of thing. Anyone who expects a source of power from the transformations of these atoms is talking moonshine." Ernest Rutherford, physicist, ca. 1930.

"There is no plea which will justify the use of high tension and alternating currents, either in a scientific or a commercial sense." Thomas A. Edison, 1889.

"While theoretically and technically television may be feasible, commercially and financially I consider it an impossibility, a development of which we need waste little time dreaming." Lee De Forest, "Father of the Radio", 1926.

Production information

Extropy #12 was produced on a Gateway 486 DX2 50 with 8Mb of RAM, a 630Mb hard disk, 17" NEC 5FG monitor powered by a #9GXE video accelerator with 2Mb of memory, using Pagemaker 5.0 for Windows and Word for Windows 6.0. The proofs were printed at 600dpi on an HP Laserjet 4 with 6Mb of RAM. Layout by Max More.

This issue was printed on a web press by Canyon Printing, Inc., Anaheim, California.

CONTRIBUTORS

Nanc Clark: Nanc received her Master of Fine Arts at the Accademie de Belle Arti in Italy. She seeks ways of applying artistic vision to futurist ideas.

Jamie Dinkelacker: Dr. Jamie Dinkelacker focuses on information services, organizational development, communication technologies, and the marketing of related innovations. He holds a Ph.D. in communication research from Michigan State University, and an MS in communication research and a BS in management from Rensselaer Polytechnic Institute. He serves on the Board of Advisors of both Foresight Institute and Extropy Institute. jamie@netcom.com

Eric Drexler: Eric Drexler extracted a Ph.D. in Molecular Nanotechnology from M.I.T. He wrote *Engines of Creation* (Doubleday), co-authored *Unbounding the Future* (Morrow), and *Nanosystems: Molecular Machinery, Manufacturing, and Computation* (Wiley Interscience) — named the outstanding computer science book of 1992 by the Association of American Publishers. He began studies of molecular nanotechnology in 1977. drexler@netcom.com

Bill Eichman: Bill Eichman, 37, is an independent scholar, philosopher, author, and educator, and the owner of several businesses including a building contracting company, a computer consulting firm, an Internet BBS, and a Management company. One of his special interests is the building of community structures that support self-realizing and self-developing humans. He lives near Penn State University in central PA on a small country homestead with his beloved partner Marisa. wce@hogbbs.scol.pa.us

Hal Finney: Hal received a BS in Engineering from Caltech. He developed several commercial video games for Mattel, then went on to head the Operating Systems Group at Ametek, a major manufacturer of parallel computers. He co-developed Neuralyst, a commercial neural-network simulation package. Hal played a major part in the development of version 2.0 of Phil Zimmermann's encryption program, PGP. hal@alumni.caltech.edu

Robin Hanson: Robin Hanson got an M.S. Physics, and M.A. in the Conceptual Foundations of Science from U. Chicago in 1984. In 1993, Robin became a grad student in social science at Caltech. In between, Robin researched artificial intelligence and Bayesian statistics for NASA, and for himself researched alternative institutions for hypertext publishing, and for buying research, health care, law enforcement, and much more. hanson@hss.caltech.edu

Harry S. Hawk: Manager of Computer Services for Warwick, Baker, & Fiori Advertising. He holds a masters degree in Interactive Telecommunications from New York University. He is Electronic Communications Officer for Extropy Institute and manages the Internet-based Extropians e-mail list. habs@panix.com

David Krieger: Dave Krieger is Director of Publications for Agorics, Inc., a platform software and consulting firm in Silicon Valley. He is *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: Simon! D. Levy is an ExI Director and the editor of ExI's newsletter *Exponent*. He holds a master's degree in linguistics from the University of Connecticut and earns his living writing Macintosh and DECwindows Motif applications at Haskins Laboratories, in New Haven, CT. Simon!'s research interests include neural networks, genetic algorithms, and artificial life. He has written articles about these topics for this magazine. He plans to return to graduate school next fall, to work on a Ph.D. dissertation investigating spontaneous linguistic order in a synthetic environment. levy@yalehask.bitnet

Mark Miller: At Datapoint Mark built the first commercial distributed windows system. He was co-architect of the Xanadu distributed hypermedia server. At Xerox PARC, Mark co-authored (with Eric Drexler) the agoric open-systems papers on market-based computation and market-oriented programming. Currently he is co-director of the Agorics Project at GMU, Chief Technical Officer of Agoric Enterprises, Inc. in Fairfax, VA, and a founder of Agorics, Inc. in Los Altos, CA. mmiller@netcom.com

Max More: Editor and co-founder of *Extropy* and President 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, USC. more@extropy.org

T.O. Morrow: Tom is Vice President of Extropy Institute, and Law and Politics Editor of *Extropy*. He holds a Masters degree in Philosophy from USC, and a law degree from the University of Chicago. Tom invented the word "extropy," and co-founded the magazine of the same name, in 1988. He currently practices law deep within the belly of the Beast. t0morrow@aol.com

Nick Szabo: Nick Szabo graduated from the University of Washington in 1989 with a Bachelor's in Computer Science. He has worked at Caltech's Jet Propulsion Laboratory, IBM, and Sequent Computer, and currently is a software consultant and writer. szabo@netcom.com

Ralph Whelan: Vice President of Alcor Life Extension Foundation, Editor of *Cryonics* magazine, and a director of Extropy Institute. His main interest is music, which he views as a tremendously precise and complex form of thought/communication. His biggest distraction is computer art/animation. ralph@alcor.org

ExI Audio Tapes, Books, T-Shirt

Everything is Getting Better and Better – I'll Bet On It! by Julian Simon.

Economist Julian Simon uses hard data to counter prevailing gloomy beliefs about the current state of the world and its direction. Practically all measures of human well-being substantiate the Extropian's dynamically optimistic views: Life does tend to improve, though only through the efforts and applied intelligence of free persons. This tape makes an effective introduction to Simon's ideas, and, lent out to your pessimistic friends, will serve as a valuable intellectual catalyst.

\$10.95 (Members \$9.95) EC1 (1-hour audio)

Bionomics On Trial: A Discussion With Michael Rothschild

Rothschild outlines the main contentions of his book *Bionomics: Economy As Ecosystem*, and responds to audience questions. Topics discussed include electronic ecosystems; how bionomics effectively draws 'liberals' into support for free markets; the relation between Austrian/process economics and bionomics; the role of government; how far the economy as ecosystem analogy can be pushed.

\$12.95 (Members: \$11.50) EC2 (80 minutes audio)

Postage:

\$1 per tape. Overseas orders: Surface mail – \$1.50 first tape, \$1.25 each additional tape. Contact ExI for airmail rates.

BOOKS AVAILABLE (domestic postage included):

Maureen Caudill, *In Our Own Image: Building An Artificial Person* \$24.95

K. Eric Drexler, *Nanosystems: Molecular Machinery, Manufacturing, and Computation* \$27.95

Alan Lakein, *How to Get Control of Your Time and Your Life* \$5.95

Marc Steigler, *The Gentle Seduction* \$4.70

Gelemter, *Mirror Worlds* \$24.95

Vernor Vinge, *A Fire Upon the Deep* \$7.49

Forward! Upward! Outward! Into your T-Shirt!

Strike terror into the hearts of gloomy pessimists, life-haters, and statist!

Proudly display your transhumanist colors! Yes, Extropy T-shirts are once again available, currently only in men's eXtra-Large. These shirts, in blue, feature the five-spiral Extropy logo in black and gold up front, and the motto "Forward! Upward! Outward! Into the Future!" on the back.

The cost is \$16 per shirt (\$14 for ExI members).

Check or money order in US dollars drawn on a US bank, payable to "Extropy Institute." Mail your order to: Extropy Institute, Dept. S, 11860 Magnolia Avenue, Suite R, Riverside, CA 92503.

How to Join the ExI Virtual Community

Extropy Institute sponsors, through the endeavor of ExI Electronic Communications Officer Harry S. Hawk, a number of electronic fora for sharing libertarian, life-extensionist, pro-technology and other Extropian ideas with bright, like-minded individuals around the globe.

The most popular service is the Extropians e-mail list, which boils over with lively discussion and debate on numerous topics. To join, send a request to:

extropians-request@extropy.org

When sending your add request, indicate whether you want real time or digest mode. (If unsure, try the digest mode first!) The Extropians list is using the most advanced information-filtering software, allowing you to select which messages (topic, author, etc.) you receive and how you receive them.

There is also an Extropian conference on the Well, one of the longest-running professionally run BBS systems. On the Well, send mail to **habs**.

Another service is the ExI Essay list, for posting longer, more carefully prepared electronic manuscripts. To get on this list, send a message to:

exi-essay-request@gnu.ai.mit.edu

Those ExI-Essay papers containing explicit notices granting permission for redistribution are available by anonymous FTP at lynx.cs.wisc.edu (IP address 128.105.2.216). A list of available essays along with their file names is in the file pub/ExI-Essay/INDEX. Any questions should go to Derek Zahn at derek@cs.wisc.edu.

There are also four "local" lists for announcements and discussions around the San Francisco Bay Area, Boston, Los Angeles, and New York. To join these lists, send messages to:

exi-bay-request@gnu.ai.mit.edu

exi-bos-request@gnu.ai.mit.edu

exi-la-request@gnu.ai.mit.edu

exi-nyc-request@gnu.ai.mit.edu