




Your Membership Benefits From — THE LIFE EXTENSION FOUNDATION

The Foundation provides its members with "inside" information about life extension research, nutrients, and drugs from medical centers, universities, and companies throughout the world. This includes the latest findings on therapies to slow down and reverse the aging process, and to prevent and treat lethal diseases. The Foundation is the foremost advocate of healthcare freedom in the nation and is in the frontlines of the fight against FDA tyranny. Among the lifesaving information you'll receive when joining The Foundation are:

- How to obtain high-tech anti-aging therapies such as **growth hormone, DHEA, melatonin, thymosin, etc.**
- How to obtain **European "smart" drugs**, including one drug that is *four times more effective* than Hydergine.
- A drug available in the United States that has extended the lifespan of animals to the **human equivalent of 160 years.**
- Exciting new therapies to treat "**terminal**" victims of cancer, Alzheimer's Disease, and other incurable diseases.
- High power nutrients that boost immune function, protect against cancer and heart attacks, and shrink benign enlarged prostate glands.

As a member of The Foundation, you'll receive *at least 80 newsletters a year* in our unique **LIFE EXTENSION MAGAZINE** (which also includes special reports and ads). You'll receive frequent issues of *Life Extension Report, Life Extension Update, Life Extension Abstracts, The FDA Raid Report, Quality Control News, Fitness For Longevity, AGE NEWS* (the American Aging Association's newsletter) and special *Life Extension Hotline Bulletins* whenever we have news that we simply cannot wait to send you! In addition to this priceless information, you'll also receive:



1. **THE DIRECTORY OF INNOVATIVE DOCTORS.** A nationwide directory of doctors knowledgeable about advanced life extension therapies who may be willing to prescribe them for you.
2. **THE PHYSICIAN'S GUIDE TO LIFE EXTENSION DRUGS.** The first book to provide American doctors with information about safe and effective "unapproved drugs". The book is referenced to enable lay persons to find therapies for specific purposes.
3. **THE DIRECTORY OF INNOVATIVE MEDICAL CLINICS.** "Incurable" diseases may have effective therapies that the FDA has not yet "approved". This unique book tells you about innovative medical clinics where the victims of cancer, Alzheimer's Disease, etc., are benefiting from such therapies.
4. **DISCOUNTS OF 25% to 60% on all your vitamin purchases,** including unique, advanced multi-nutrient formulas which will protect you against lethal diseases, increase your energy level, improve your intelligence, boost your sex drive, help you sleep better, and make you feel better. Special sales throughout the year will enable you to enjoy **even higher discounts** on these extraordinary products.
5. Direct access to "unapproved" life extension therapies such as **acetyl-l-carnitine, melatonin, and pygeum/saw palmetto herbs** at lower than wholesale prices through your exclusive membership in the only **LIFE EXTENSION BUYERS' CLUB** in the world.

The Life Extension Foundation is the **only** organization that tells you how to obtain the most advanced life extension therapies in the world long **before** they are "approved" by the FDA. For only \$75, you'll be among the **first** to find out about remarkable therapies that will extend and enhance your life!

To Join
The Foundation
call:

1-800-841-5433

or use the coupon provided.

✂ Enclosed is my \$75 membership. Please enroll me in your life extension program, which includes 80 newsletters a year, directories of innovative doctors, drugs, and clinics, and super discounts on all my life extension product purchases!

Name _____

Address _____

City _____ St _____ Zip _____

Credit Card # _____ Exp Date _____

Mail to: **Life Extension Foundation • P.O. Box 229120 • Hollywood, FL 33022-9120**


Neuroscience Pioneers

PROFILE

| | |
|--------------------------------------|-----------|
| Building Brains | 18 |
| Intelligent Assistants | 23 |
| DNA Computers | 26 |
| Memory and its Improvement | 27 |
| Imaging the Brain | 29 |
| Fuzzy Logic & Neural Nets | 31 |



THE TRANSHUMAN TASTE: REVIEWS




Better Sex Through Chemistry 54

NANO: The Emerging Science of Nanotechnology 56




River Out of Eden: A Darwinian View of Life 57

Dr. Walford's Interactive Diet Planner 59



Permutation City 60



Intelligent Information Filters and Enhanced Reality

Sasha Chislenko

13

From personalizing the way information is presented to you, to virtual overlays on your “natural” perception of the world. In the coming years Enhanced Reality will allow us to boost our senses and customize our window on the world. VR will never be the same again!

Complacency & Conservation

33

Julian L. Simon

The author of *The Ultimate Resource* responds to criticisms of his ideas on complacency about the environment. Simon refutes the allegations and explains the need for trade-offs.

Idea Futures on the Web

35

Duane Hewitt

A betting market on ideas, especially scientific and technological questions, has been implemented on the World Wide Web. We will carry the latest market prices in a regular Idea Futures column.

Smart Contracts

Building Blocks for Digital Free Markets

Nick Szabo

50

What parts of our hard won legal tradition will still be valuable in the cyberspace era? What is the best way to apply these common law principles to the design of our on-line relationships? Smart Contracts discusses principles of observability, verifiability, privity, and enforceability.

D E P A R T M E N T S

06 **Editorial**

07 **Letters: Transmissions**

9 **Advances**

PGPfone 1.0 Beta

Space

Nanotechnology

Neuroscience

Internet

Computers

Economy

Population

39 **Extropy Institute**

42 **Bio-Enhancement Update:**

Melatonin, pt 2: Dreams like you've never dreamed

Dr. Ray Sahelian

49 **ENIGMA:**

Murder at the Liar's Club

M.J.P. Wolf

50 **Mindsurfing:**

How the Internet is Like a Baby's Head

Yow

66 **Back issues**

DATE: October 24 1995

Advancing

by Max More

“The Extropian view continues to spread rapidly and unstoppably.”

Extropians are those who consistently espouse ideas and values directed at overcoming limits to human possibility, both individually and socially. The Extropian view (or views) of life make sense of the world and our place in it. This view coherently connects together ideas, technologies, and practices into a powerful way of understanding and acting.

The Extropian view continues to spread rapidly and unstoppably. Even many people not directly involved in explicitly Extropian activities are appreciating the scope and utility of this way of thinking. Often this understanding is only partial. We can see an example of this in the current guest editorial of the *SPIN* magazine special issue on the future by VR pioneer Jaron Lanier. Part of what Lanier says is this:

“The new divide is between what I’ll call the Extropians and Stewards. A Steward is somebody who wants to manage the world as a precious resource, and an Extropian is someone who wants to let some big, impartial evolution-like process run wild with it... Extropians don’t believe in any one “natural” order and don’t care if it is turned upside down. Extropianism is in my dreams because it is creative and unbounded, and yet it gives me the creeps.”

The editorial does give some idea of what Extropians are about, but also makes us appear more cavalier than is warranted. I would point out that Extropians do not see the issue as maximum growth at the cost of



Leigh Christian

environmental damage vs. heavy regulation. Rather we favor the use of market mechanisms (such as the recently implemented pollution shares) to rationally balance people’s preferences for environmental quality and other goods. By setting up incentives intelligently, and encouraging rapid technological progress, we expect to achieve fast growth and maintain a high degree of environmental quality.

In this issue, economist Julian Simon addresses exactly this kind of misunderstanding in “Complacency and Conservation.”

This issue also sees the introduction of a new Advances feature. Advances provides summaries of technological and economic progress. We encourage readers to contribute relevant items and to let us know what you’d like to see covered in this section.

This record size issue contains the usual range of writing, covering everything from the next step beyond VR and smart contracts for business on the Net to chemicals to improve sexual function. Enjoy!

Upward and Outward!

Max More, Ph.D.



EXTROPY — a measure of intelligence, information, energy, life, experience, diversity, opportunity and growth. Extropians are those who consciously seek 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]

TRANSMISSIONS

Send comments to maxmore@primenet.com or to:

The Editor, *Extropy*, 13428 Maxella Avenue, #273, Marina Del Rey, CA 90292

Digital Checking

I want to address a few of the points made by Lawrence White in the article in *Extropy* #15 on digital cash.

I had one quibble with White's description of currency and cash cards. He seemed to overlook the salient feature of electronic cash, that once withdrawn it gets "re-blinded" so that it is not recognizable to the bank. So this is fundamentally different from electronic checking, where the payor knows the identity of the payee. This is not very significant to what I mostly want to talk about, though.

Robin [Hanson] had asked about White's challenge to the dogma that cryptography in general and digital cash in particular will lead to collapse of governments and the advent of a libertarian/anarchist state. This view has been advanced most forcefully by former [Extropians e-mail] list member Tim May, who refers to this scenario as "crypto-anarchy".

Actually in reading White's article I don't find very much to challenge or address this view. After the technical discussion White argues that the new technologies will make it a lot easier for people to get overseas bank accounts. I had echoed this point in my earlier message, probably dimly remembering White's argument.

White then asks two questions: will this hurt the Fed's control of the money supply (of dollars), and will people switch to using some new currency in place of dollars. He thinks neither is likely.

Two things about these issues stand out: one is their US-centric nature, which is in contrast to the international net. Why ask about the Fed's control of dollars and not the British government's control of pounds, etc.? White seems fixed on a world in which people have well-defined currencies that they "should be" using, but my guess is that the boundaries between currency usefulness are likely to blur for more and more people. Secondly, even if White is right, it doesn't really address the larger issue of crypto anarchy; people may stick with dollars, but if they can exchange them anonymously for

enough of their personal finances, they may be able to evade taxes and thus starve the government of its revenues. We went over this already, but White doesn't touch on it at all. His issues are more narrowly focused.

As far as the money-supply issue, White says that the Fed can simply change the reserve requirements to keep control of money supply, as it does today. However he just finished arguing that more overseas bank accounts will exist, hence presumably there are a lot more "Caribbean dollars" outside the control of the Fed. The Fed still will wield a lot of leverage, but presumably the internationalization of dollars that he predicts will make it less powerful. The Eurodollar market which has become so significant in recent years may increase by orders of magnitude once every Uncle Joe has his own overseas bank account with which he communicates with perfect anonymity and secrecy.

As far as the other point, that people won't stop using dollars (or yen in Japan, pounds in England, etc.) I think as I said above that internationalization will hammer at this assumption. If most of the things on the Net that I want to buy are in Japan it may be that I will want to get paid in yen rather than dollars. The net will allow more people to work farther from where they live, including overseas. The company I work for is considering farming work out to India. We can use email to send and receive the programs they will write for us. They don't charge very much.

White's argument seems hardly more than saying that people will do what they have always done. But the net will offer an opportunity for new financial instruments to get a competitive edge. It is like a new ecological niche, and it may be that the stodgy old currencies of the past won't be the strongest to colonize.

Another thing I expect to happen is a reduction of transaction costs in switching currencies. With the market volume increasing there should be economies of scale, a more competitive situation. So it may not matter that much what currency you get paid in online, you can easily switch it to the local currency of whatever

you want to buy.

So, overall, I did not think White's analysis was very useful. He focused on some narrow issues which might be of interest to the professional economist but did not seem to look at the bigger picture, or to consider the larger scale changes which the net may cause.

Hal Finney

hfinney@shell.portal.com

Reply by Larry White

Hal Finney is right that my article didn't address the proposition that "cryptography in general and [encrypted payments] in particular will lead to the collapse of governments and the advent of a libertarian/anarchist state." As he noted, my focus was narrower. Let me broaden it a bit here.

First, as to the fiscal implications: I agree with the mild proposition, cited by Robin Hanson, that encrypted digital checking will *weaken* governments (not necessarily to the point of collapse) by undermining their ability to levy some kinds of taxes. An increase in the convenience of using anonymous and untraceable ways of transacting will clearly make taxes on more kinds of transactions harder to collect. Casual income (e.g. from babysitting), received in cash, is in practice untaxable already today. The empirical question, about which we can only speculate, is *how big* a difference it will make when the anonymity and untraceability of cash payments becomes available for check-like payments.

I doubt that it will make a big difference to sales tax collection, because the option of cash payments already exists today in face-to-face retail transactions. (Mail-order transactions are already sales-tax-free, at least in the United States.) Retailers very seldom (at least in my own experience) offer a sales-tax-free price if one will pay in cash. Presumably they fear a sting operation by tax agents posing as customers, a fear that would remain even with a customer offering payment by encrypted deposit transfer. The same con-

siderations suggest that neither will there be a big difference to payroll taxes paid by regular employers.

Encrypted desktop funds transfer may well make a big difference to the taxability of gambling, by allowing gambling via the Internet (see the article in the October '95 issue of *Wired*). On-line competition from offshore lotteries and virtual casinos will undercut the taxability of local lotteries (today state-run monopoly lotteries tax their customers via lousy odds) and physical casinos.

I'm having trouble, though, thinking of many other transactions where nowadays the parties acquiesce to paying taxes only because it is too much of a hassle to make the deal in untraceable cash and they are afraid to evade taxes when paying by check or credit card payment. If I am overlooking some major sets of transactions for which this situation obtains, I hope someone will point them out.

Second, as to the monetary implications: Hal Finney wonders why I "ask about the Fed's control of dollars and not the British government's control of pounds, etc." Theoretically, these are the same question. The dollar was just the example I chose.

My point about reserve requirements was that the Federal Reserve (to use the U.S. example again) could, if it wanted to, easily apply them to smart card currency balances. The total of such balances will appear on the balance sheets that banks report to the Fed each week. But reserve requirements on bank-issued currency are *not* needed to keep control of the quantity of money; neither are reserve requirements on bank deposits. Reserve requirements are just a tax.

By making it easier for domestic citizens to move deposits and currency balances to privacy-respecting offshore banks that don't report to the Fed, I argued, desktop electronic funds transfer will make the reserve requirement tax, and other taxes on banks, easier to evade. The Fed will find it more difficult to impose taxes and regulations on banking, and in that sense its power will decline. But the Fed won't find it harder to control the global aggregate of deposit and currency dollars, unless the proportions between onshore and offshore dollars becomes harder to predict. I don't see any reason to think that this proportion will be any harder to predict than it is today, when offshore banking is already huge, and when about half of all Federal Reserve notes circulate outside the United States.

Finney's point is certainly well-taken that desktop EFT may, by providing cheaper access to low-commission interbank exchange rates for switching between currencies, make ordinary moneyholders more sensitive (as multinational firms already are) to the strength of their local currency. The more readily moneyusers will dump a weakening currency (with a rising inflation rate) in favor of a substitute strong (low-inflation) currency, the lower the rate of inflation a central bank will find worth having.

My argument for expecting people to persist using their national currencies was not that "people will do what they have always done," but that there is a kind of lock-in to an established monetary standard. People will not substitute out of the local currency for a trivial difference in inflation, though we know from the "dollarization" experiences of high-inflation countries that the public will substitute when the inflation differential gets high enough. Cheaper access to foreign currency lowers the threshold differential (*by how much* is again an empirical question). Central banks will have to be more careful not to vastly exceed the inflation rates of their major trading partners.

Unfortunately, I still see no reason to expect the *de novo* establishment of stateless currencies. It is hard to imagine that any individual would find it his best option to switch to a new currency that none of his trading partners yet uses.

Larry White
lwhite@rigel.econ.uga.edu

Great Idea Seeks Champion

The World Wide Web lets you follow links forward easily, but following them backwards is much harder. Links which could be nearly as easily followed backwards as forwards were a central feature of early "hypertext" visions, such as Xanadu, with which I worked for a while.

We were big fans of criticism as central to the evolution of knowledge, and we hoped that such bi-directional links would, among other things, allow people to easily find criticism of pages. Imagine that people at the home page of the CIA or R.J.Reynolds, or at some high-profile Op-Ed article, could find responding critiques with just one button push! This might have profound implications for public and commercial debate.

Now I know there are various

"backlink" projects working on how to do this in a general/efficient way. But most of these projects leave the ability to follow a link backwards at the discretion of the page linked-to—exactly the wrong incentives for criticism. And in the meantime habits, conventions, and expectations are being formed by actual use of the web, and these now have little place for direct criticism.

Actually, we can get most of the benefit of easily-found criticism from just one small modification to one web browser. That is, to get this feature:

Given that I'm looking at a page (call it page1), I push just one button, and get a page which lists page2s which reference page1.

...you just need your browser to paste the URL of page1 into the query form of a web database, such as Open Text Index, that allows searches on cited URLs.

Of course having such a browser would just be the first step. Fans of criticism should then:

1. Create pages which effectively criticize other pages, and to submit them to this database.
2. Use this browser and its backlinks button, and spread the word about how handy it is.

If this feature gets a good reputation, other browsers and databases should jump on the bandwagon, and we might slowly induce a web culture of criticism, where websurfers get into the habit of looking for criticism when they have doubts about a page.

Yes, there is a lot more one can do with backlinks besides finding criticism. And for pages with many backlinks, we'd like better ways to help folks find the good criticisms. But let's not try to work everything out before we get started—that was Xanadu's mistake.

This idea needs a *champion*, someone who really values criticism, and who will take the time to make it real. I don't have time to be this champion, and neither does Rob Jellinghaus, who also advocates it. Hence this plea, which can be found at:

<http://www.hss.caltech.edu/~hanson/findcritics.html>

Please mention this page to whomever you think might care.

Robin Hanson
hanson@hss.caltech.edu

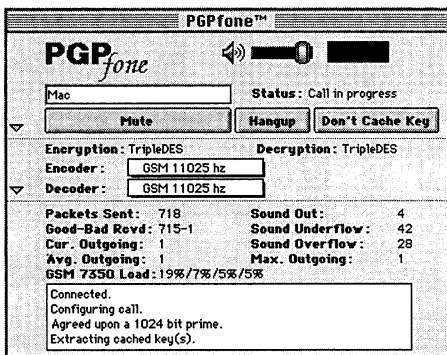
ADVANCES

PGPfone 1.0 Beta

To follow up on the success of his public-key encryption program, PGP, Phil Zimmerman recently released a Beta version of PGPfone, encrypted voice software for Apple Macs. PGPfone runs native on both 68k and PowerMacs, allowing you to use your computer as a secure telephone either over Appletalk or by modem.

Unlike PGP, which uses the RSA and IDEA encryption algorithms, PGPfone uses the Diffie-Hellman public key algorithm to exchange keys, and gives a choice of either Triple-DES or Blowfish to encrypt the conversation. Triple-DES is considered to be one of the most secure encryption algorithms available, but is CPU-intensive. Blowfish has been subject to much less analysis, but is fast and probably of a similar security level.

On startup PGPfone displays a window with a box to enter either the telephone number to call or the name of a Mac on the Appletalk network. When this has been entered, the computer will attempt to open a connection to the PGPfone program running on the other machine, and if successful the recipient will hear a realistic



rendition of an old-style telephone ring and may choose to answer or ignore the call (there is no equivalent to caller-id included).

When the call is answered the computers will spend several seconds exchanging keys. The Diffie-Hellman algorithm is secure against eavesdroppers, and unlike RSA if an encrypted conversation is recorded, even the people who made the call will not be able to decrypt it later. However, an active attacker can break into the key

exchange as it occurs, so PGPfone displays a list of four words based on the exchanged key. If the same list is displayed to both users, then they can be sure that no such break-in has occurred.

Because of the wide variation in performance of different Macs and modems, PGPfone includes a variety of speech compression algorithms, based on the GSM algorithm used for digital cellphones but with data rates from 4410 Hz up to 11025 Hz. The default 5512 Hz setting is tolerable, and the best rate better than a standard telephone, however while speech at the slowest rate can be understood it is difficult to distinguish different speakers. An additional problem is a noticeable propagation delay, taking around one second for speech to be transmitted and decoded.

Two PowerMacs connected over Appletalk will happily run at the fastest rate using Triple-DES, though occasional dropouts occur. Slower Macs, or those using 14.4k or 9.6k modems, will have to select a lower rate, and very slow Macs will have to run in 'half-duplex' mode, with a 'Press to Talk' button similar to CB radio.

For a Beta, the software is very robust, aside from one bug which crashes your machine with some Appletalk setups if you attempt to connect without specifying a machine name. Otherwise it is a worthy counterpart to PGP, and hopefully will become just as widespread.

Windows95/NT and Internet support are promised for the final release. I look forward to both.

PGPfone is available by ftp from net-dist.mit.edu in /pub/PGPfone for US and Canadian users, or from utopia.hacktic.nl in /pub/voice for the rest of the world. Due to US ITAR restrictions you should not download the software from MIT if you are outside the US and Canada!

INTRODUCTION

This is the "beta" test introduction of a new Advances feature for *Ex-tropy*. Long time readers will remember the Intelligence At Work section from the early issues.

Advances will feature short summaries of advances in science and technology. Our focus will be on developments that further our extropic goals of extended life, intensified intelligence, increased freedom, and other ways of overcoming human limits.

Advances will also present economic information, especially as it relates to standards of living and investment opportunities in technological companies. —MM

Direct all contributions of information for this section to Advances Editor Eric Watt Forste: arkuat@pobox.com or send to *Ex-tropy*, Advances, 13428 Maxella Avenue, #273, Marina Del Rey, CA 90292.

SPACE

ULYSSES IS NOW PASSING OVER THE SUN'S NORTH POLAR REGION. Launched in 1990, the Ulysses spacecraft does not take photographs but instead monitors the magnetic fields and the fluxes of cosmic rays and solar wind particles in the greater solar environment. Now that Ulysses has gone over the top its main task has been fulfilled, but scientists are hoping that the mission can be extended at least to the year 2000, when the craft would return to the solar antipodes at a time when the sun would be in the most active phase of its 11-year cycle.

(Eos, 25 July 1995.)

NANOTECHNOLOGY

TWO IMPORTANT AND INDEPENDENT ADVANCES in nanotechnology were announced this summer:

1) Carbon tubes of nanometer diameter have been around for about 4 years, but until now nobody could do much with them. In the current issue of *Nature* it's reported that molten vanadium oxide can form a coating on the carbon tubes, the carbon can then be dissolved away using conventional chemical techniques leaving pure vanadium oxide tubes of nanometer diameter.


Unlike carbon, vanadium oxide is active optically and so can be used as an optical switch in a super fast computer. Also vanadium oxide is a powerful catalyst for many chemical reactions so it might be possible to use the tiny tubes as molds for all sorts of different materials.

2) MOVING ATOMS around with a scanning tunneling microscope is, of course, old hat but it has not been possible to break individual chemical bonds between atoms, until now. In the June 16 issue of *Science* it's shown that if electrons of the correct energy are shot at an atom from the tip of a scanning tunneling microscope the atom will resonate and the resulting vibration will break the chemical bond.

We can now do chemistry, and not just physics on atoms one at a time. The

procedure is somewhat faster than expected and does not require any exotic conditions such as very low temperature. J.W. Lyding, one of the authors of this report, is quoted as saying "We'd like to make small, electronic devices on the nanometer scale".

DENSE ARRAYS OF NANOWIRES can now be made. Huber and colleagues at the Naval Surface Warfare Center in Maryland forced molten metal into 250 nm-diameter channels using injection molding. Most of the wires in the resulting array were conductive. Researchers aim to reduce the size of the channels to 3 nm. Metals used in this study include tin, indium, aluminum, and several semiconductors.



Scanning tunneling microscope

[*Science*, 11 Feb 94]

ATOMICALLY PERFECT fullerene gears have been designed in simulations at NRL with 5-8 sprockets and 290-464 atoms. A fullerene shaft was added to a six-tooth gear. During a simulation, one gear was spun up the other from stationary to 20 gigarev/sec in a quarter of a turn. The overall bind-

ing energies were determined to be more stable than C_{60} . [Novel Forms of Carbon II, 283-288]

CARBYNE RODS with chain lengths over 300 carbon atoms were synthesized by Lagow *et al.* Rod logic nanocomputer proposals by Drexler involve just these kind of rods in their construction. It is therefore encouraging that these rods are stable at room temperature despite the absence of any structure to prevent collisions.

[*Science* 267:362-367, 20 Jan 95]

NEUROSCIENCE

A SILICON DEVICE FOR TRIGGERING ELECTRICAL ACTIVITY IN A NERVE CELL has been constructed, opening possibilities for two-way, non-toxic communication between computer chips and nerve cells. Previous devices for stimulating nerve cells were metallic devices generating ordinary electric currents. Not only do such devices have corrosion-prone electrodes, but their currents create electrochemical byproducts and heat that could damage the nerve cells and themselves.

The silicon device, constructed by researchers at the Max Planck Institute for Biochemistry in Germany, contains a "stimulation spot" that triggers neural activity simply through the rearrangement of electric charge. Insulated by silicon oxide, the stimulation spot has a size (between 10 and 50 microns) matched to that of a leech nerve cell to which

it is connected. A voltage pulse applied to the spot rearranges electric charge on the silicon oxide film and the insulating membrane of the nerve cell, creating a buildup of positive charge in the nerve cell which causes it to fire above a certain threshold. The silicon device is capable of triggering a single nerve cell without affecting other nearby neurons.

The device complements the previously designed "neuron transistor," which receives ionic signals from nerve cells and transcribes them to electronic signals in silicon. "These two worlds join the two worlds of information processing, the silicon world of the computer and the water-world of the brain," says the Max Planck Institute's Peter Fromherz (fromherz@vms.biochem.mpg.de). Developing this device for biomedical applications, such as computer-controlled artificial limbs, is not envisioned at the present moment, as researchers will first need to build and understand devices that interact with connective tissue and other non-neuronal cells in the body.

[Physics News Update, *The American Institute of Physics Bulletin of Physics News* # 236 August 7, 1995 by Phillip F. Schewe and Ben Stein]

SYNTHETIC BRAIN: Building an "artificial brain" by 2001 is the goal of researchers at STR laboratories in Kyoto, Japan. According to ATR computer scientist Hugo de Garis, the CAM-Brain Project intends to create a silicon brain with more than a billion artificial neurons.

Designing and building such a massive device in all its details would not be feasible. What researchers will start with a neural network running on a massively parallel computer. They will use a form of artificial life called cellular automata to make links so that the system builds itself up into an extremely complex system.

In this "evolutionary engineering", as de Garis calls it, the cellular automata, each of which is a computer, sends signals to other cellular automata which spur growth through something like the synaptic connections in biological brains. At the current stage of the project, the connections are growing in a plane, but they will soon start to connect up in three dimensions. De Garis refers to the artificial brain as a type of "Darwin Machine" because of the evolutionary manner in which it develops.

INTERNET

WWW Growth: 1993: 443,931% 1994: 1,713%

Countries connected to the Net: 1993: 137 1994: 159

Number of US invention patents available for searching on the Internet: over 158,000

At current growth rates, estimated time at which everyone on Earth will be on the Internet: 2004

Estimated number of subscribers to online services: 12 million

Dataquest estimate of revenues for online services in

1997: \$3.3 billion

Approximate number of domain name requests, per minute, handled by the InterNIC: 1.2

Estimated number of Internet hosts at the end of the century: 101 million

Past issues and citations to sources can be found at <http://www.openmarket.com/info/internet-index/>.
internet-index-request@OpenMarket.com.

COMPUTERS

PERSONAL COMPUTER sales are expected to reach 100 million per year by the end of the decade.

CYRIX expects to ship in December '95 a fast Pentium-compatible chip. Cyrix's chip reportedly runs 1.2 to 1.3 times faster than an equivalent Pentium. Cyrix currently has about 7% of the 486 chip market. It could hold up to 10% of the Pentium-level market in 1996. Another incentive for Intel to keep accelerating development times.

COMPUTERS AGAINST BLINDNESS: Computer technology may be starting to restore partial sight to some blind people. Dr. Joseph Rizzo and colleagues have developed a tiny computer chip which will float on the retina. This solar-powered computer uses electricity provided by a laser beam emitted from special eyeglasses. The eyeglasses have miniature TV cameras in place of lenses.

The chip consists of two layers. The solar cells are

on the top layer. Underneath is the circuitry. A strip emerges from the layers bearing electrodes that send signals through the retina's nerves to the brain.

The ingenious glasses use optical detectors developed from those used in video cameras. They convert images into coded streams of light pulses conveyed by laser to the chip. Researchers expect the chip to allow only a narrow visual field. However, the practical difference between that and no vision at all will be immense.

and stands for Visa, MasterCard International Inc., and Europay International. These three companies agreed to the standard, specifying both cards and processing terminals, in September '95. Cards using the standard can be used anywhere in the world participating in the program.

Ivan Yim, country manager of Visa for Hong Kong and Macao, said of the Visa Compass Card: "It is the world's first and Hong Kong's first VME compatible program. Throughout



Asia/Pacific, there are a number of issuing banks and non-bank financial institutions who are in the race to launch chip card programs." Partners with Visa in the Compass Card project are Hutchison Whampoa Ltd., and Dao Heng Bank Group's subsidiary OTB Card Co.

The current strips found on the backs of credit cards can hold only a line or two of information. These will gradually be replaced with a chip capable of storing the equivalent of two pages of text.

Visa expects to encourage use of the card through a "frequent buyer" program. Visa Compass cardholders can collect points by spend-

These retinal chip implants are the latest computer based sensory aid. Already in use are electronic cochlear implants which have allowed almost 10% of previously deaf persons receiving them to hear well enough to hold a telephone conversation.

Researcher Dr. Joseph Rizzo is a neurologist and ophthalmologist at Harvard Medical School.

ELECTRONIC MONEY takes another step. Visa International Inc. announced a credit card, to be issued first in Hong Kong, conforming to a new standard for "smart card" computer chips. The standard is called VME,

ing any any of 12 million Visa merchants around the world. The points are turned into cash that can be spend at 300 retailers in Hong Kong.

Mastercard has been trying pilot programs for the cards in a number of places, but pervasive use of the cards isn't expected for a year. Credit cards initially were accepted only reluctantly in Asia according to Visa's Yim. Yim stated that Visa averaged 20 percent growth a year in sales in Hong Kong. Sales last year were up to \$5.2 billion out of \$83 billion in the Asia/Pacific region.

Oct 13, 1995

ECONOMY

THE PRODUCTIVITY PARADOX has puzzled economists and frustrated businesses. The paradox refers to close to static productivity despite a flood of new information technologies. Part of the problem has been the difficulty of taking into account many ways in which increased productivity shows up in an information-based economy. Finally, despite measurement shortcomings, the USA is seeing an acceleration in productivity gains.

Nonfarm business productivity growth is growing:

| Productivity annual growth | |
|----------------------------|------|
| 1973 Q3 to 1980 Q1 | 0.6% |
| 1980 Q1 to 1981 Q3 | 0.7% |
| 1981 Q3 to 1990 Q3 | 0.9% |
| 1990 Q3 to 1995 Q1 | 2.2% |

The last five years have seen such an improvement that the USA has the highest productivity and creates more jobs than any other industrial country. All this despite a savings rate kept low by heavy income taxation. Over a few years, this improved growth in productivity generates enormous wealth. Raising productivity by only 0.5% for 10 years produces an additional \$300 billion.

These figures are likely to be revised downward slightly when the government statisticians revise their measures soon. However, these measures fail to account for many improvements in services and proliferation of consumer options. Other accounting choices artificially understate productivity. For instance, although much spending on capital involves purchasing software, this is not counted as investment.

The investment in information technology is finally paying off. An MIT survey of nearly 400 companies suggested a better than 50% return on investment in these technologies. We can explain the productivity paradox in terms of a learning process for new technologies. Information technology, to significantly boost output, requires new ways of working, new organizational structures, and more flexibility. Rigidly hierarchical and inflexible corporate organizations worked well for mass production but do

not suit many of today's industries. Information technology is paying off as corporations learn to shift towards more flexibility—adhocracy, flattened hierarchies, and self-managing teams. This process is just getting under way on a large scale.

As state schools deliver ever-worsening schoolings, ill preparing graduates for the new technologies, businesses are taking up the slack. Over the last decade formal company training has increased by almost half. Helping the process, software is becoming ever more easy to use, allowing relative novices to perform tasks previously demanding advanced skills.

(Data: Bureau of Labor Statistics.)

CPI ADJUSTMENTS: A panel of economists headed by Michael Boskin (previously head of the Council of Economic Advisors) argue that the consumer price index (CPI) has been showing the inflation rate as about one percentage point higher than it should. One consequence: workers are better off than had been thought. It has looked like real wages had declined by 0.2% per year since 1990. However, if the critics of the current CPI measure are right, real wages have advanced 0.8% annually. Hourly compensation, including medical and retirement benefits, has done even better, up 1.3% per year.

If the statistical bias is

corrected, the budget deficit in 2005 should be \$140 billion less than otherwise. This would result from lower expenditure on inflation-adjusted programs such as Social Security.

ANOTHER SIGN that the fear of the Asian monster is outdated: The USA had 63 personal computers per 100 people employed. Japan has 17. Japan's growth in productivity between 1990 and 1994 was lower than that of the U.S., at about 2.4%.

POPULATION

POPULATION GROWTH in the developing nations is slowing rapidly. Many will consider that good news. A case can be made for the contrary view: more people, at least when well educated, means more brains solving problems, and bigger markets with greater specialization.

The economists are turning out to be right, and the doomsaying ecologists wrong: population growth is slowing as less developed countries become wealthier and family planning takes effect. In the 1960s six children were born to women on average in developing countries. That figure has plummeted to 3.5. China's fertility rate is now below the replacement level of 2.1. The rapidly growing income of Hong Kong, Taiwan, Korea, and Singapore has kept their fertility rates low. ☛

Intelligent Information

Filters

and



(C) 1995

Alexander Chislenko

Preface

I started to think seriously about the ideas of augmented perception and personalized views of reality after reading a number of Internet messages containing proposals to introduce language standards for on-line communications. Frequently, people suggest restricting certain forms of expression or polishing the language of the posts to make them less offensive and more generally understandable. While looking forward to the advantages of improved communications, I would like to see them provided by tools that would at the same time make the language mix of the Net more free and diverse.

In this essay, I suggest that active information filtering technologies may help us approach this goal for both textual and multimedia communications. I also pursue this concept further, discussing the introduction of augmented perception and enhanced reality (ER), and share some observations and predictions of the transformations in people's perception of the world and of themselves in the course of the technological progress.

Text Translation and Its Consequences

Many of us are used to having incoming e-mail filtered, decrypted, formatted and shown in our favorite colors and fonts. These techniques can be taken further. Customization of spelling (e.g., American to British or archaic to modern) would be a straightforward process. Relatively simple conversions could also let you see any text with your favorite date and time formats, use metric or British measures, implement obscenity filters, abbreviate or expand acronyms, omit or include technical formulas, personalize synonym selection and punctuation rules, and use alternative numeric systems and alphabets (including phonetic and pictographic). Text could also be digested for a given user, translated to his native language and

even read aloud with his favorite actor's voice.

My friend Gary Bean suggested possible implementation of "cliché translators" that would explicitly convey the meaning of a sentence that is known to the translator, but not necessarily to the reader. For example, the phrase "*that's an interesting idea*" might be translated as "*I have serious reservations about this*". In the reverse operation, words and phrases can be replaced with politically correct euphemisms.

After the recent Communication Decency Act, Robert Carr developed a remarkable "*HexOn Exon*" program that allows the user to convert obscene words in the messages into the names of the senators responsible for this Act, and vice versa. Besides presenting a humorous attempt to bypass the new obscenity censorship, this program demonstrates that allocating both responsibilities and rights for the contents of a message among multiple authoring and filtering agencies may not be easy.

Translation between various dialects and jargons, though difficult, should still take less effort than the translation between different natural languages, since only a part of message semantics has to be processed. Good translation filters would give "linguistic minorities" — speakers of languages ranging from *Pig Latin* to E-

Prime — a chance to practice their own languages while communicating with the rest of the world.

Some jargon filters have already been developed, and you can benefit from them by enjoying reading "Ible-Bay", the *Pig Latin* version of the Bible, or using Dialectic program to convert your English texts to anything from Fudd-speak to Morse code.

Such translation agents would allow rapid linguistic and cultural diversification, to the point where the language you use to communicate with the world could diverge from everybody else's as far as the requirement of general semantic compatibility allows. It is interesting that today's HTML Guide already calls for the "divorce of content from representation", suggesting that you should focus on *what* you want to convey rather than on *how* people will perceive it.

Some of these features will require full-scale future artificial intelligence, such as "sentient translation programs" described by Vernor Vinge in "*A Fire Upon The Deep*". In the meantime, they could be successfully emulated by human agents.

Surprisingly, even translations between different measurement systems can be difficult. For example, your automatic translator might have trouble converting such expressions as "*a few inches away*", "*the temperature will be in the 80s*" or "*a duck with two feet*". A proficient translator might be able to convey the original meaning, but the best approach may be to write the message in a general semantic form which would store the information explicitly, indicating in the examples above where the terms refer to measurements, whether you insist on the usage of the original system, and the intended degree of precision. As long as the language is expressive enough, it is suitable for the task—and this requirement is purely semantic; symbol sets, syntax, grammar and everything else can differ dramatically.

A translation agent would interactively convert natural-language texts to this semantic lingua franca and interpret them back according to a given user profile. It could also reveal additional parts of the document depending on users' interests, competence in a given field, and access privileges.

Currently, we can structure our mental images any way we want so long as we can translate them to a common language. This has led to relatively stable standardized languages and a great variability among minds. Likewise, intelligent software translators could let us make our languages as liberated as our minds and push the communication standards beyond our biological bodies. (It really means just further *exosomatic* expansion of the human *functional* body, but the liberation still goes beyond the traditional human interpretation of "skin-encapsulated" personal identity.)

So will there be more variety... or more standardization? Most likely both, as flexible translation will help integrate knowledge domains currently isolated by linguistic and terminological barriers, and at the same time will protect linguistically adventurous intellectual excursions from the danger of losing contact with the semantic mainland. Intelligent translators could facilitate the development of more comprehensive semantic architectures that would make the global body of knowledge both more diverse and more coherent.

Information may be stored and transmitted in the general semantic form. With time, an increasing number of applications can be expected to use the enriched representation as their native mode of operation. Client translation software will provide an emulation of the traditional world of "natural" human interactions while humans still remain to appreciate it. The semantic richness of the system will gradually drift away from biological brains, just as data storage, transmission and computation have in recent history. Humans will enjoy growing benefits from the system they launched, but at the expense of understanding the increasingly complex "details" of its internal structure, and for a while will keep playing an important role in guiding the flow of events. Later, after the functional entities liberate themselves from the realm of flesh that gave birth to them, the involve-

ment of humans in the evolutionary process will be of little interest to anybody except humans themselves.

Enhanced Multimedia

Similar image transformation techniques can be applied to multimedia messages. Recently, a video system was introduced that allows you to "soften the facial features" of the person on the screen. Advanced real-time video filters could remove wrinkles and pimples from your face or from the faces of your favorite political figures, caricature their opponents, give your mother-in-law a Klingon persona on your video-phone, re-clothe people in your favorite fashion, and replace visual clutter in the background with something tasteful.

It also seems possible to augment human senses with transparent external information pre-processors. For example, if your audio/video filters notice an object of potential interest that fails to differ from its signal environment enough to catch your attention, the filters can amplify or otherwise differentiate (move, flash, change pitch, etc.) the signal momentarily, to give you enough time to focus on the object, but not enough to realize what triggered your attention. In effect, you would instantly see your name in a text or find Waldo in a puzzle as easily as you would notice a source of loud noise or a bright light.

While such filters do not have to be transparent, they may be a way to provide a comfortable "natural" feeling of augmented perception for the next few generations of humans, until the forthcoming integration of technological and neural processing systems makes such kludgy patches obsolete.

Some non-transparent filters can already be found in military applications. Called "target enhancements", they allow military personnel to see the enemy's tanks and radars nicely outlined and blinking.

More advanced filtering techniques could put consistent dynamic edits into the perceived world. **Volume controls** could sharpen your senses by allowing you to adjust the level of the signal or zoom in on small or distant objects. **Calibration tools** could expand the effective spectral range of your perception by changing the frequency of the signal, allowing you to hear ultrasound or perceive X-rays

and radiowaves as visible light. **Conversions** between different types of signals may allow you, for example, to "see" noise as fog while enjoying quiet, or convert radar readings from decelerating pe-



destrians in front of you into images of red brake lights on their backs. **Artificial annotations** to perceived images would add text tags with names and descriptions to chosen objects, append warning labels with skull and crossbones on boxes that emit too much radiation, and surround angry people with red auras (serving as a "cold reading" aid for wanna-be psychics). **Perception utilities** would give you additional information in a familiar way — project clocks, thermometers, weather maps, and your current EKG readings upon [the image of] the wall in front of you, or honk a virtual horn every time a car approaches you from behind. They could also build on existing techniques that present us with recordings of the past and forecasts of the future to help people develop an immersive trans-temporal perception of reality. **"World improvement"** enhancements could paint things in new colors, put smiles on faces, "babify" figures of your incompetent colleagues, change night into day, erase shadows and improve landscapes.

Finally, completely **artificial additions** could project northern lights, meteorites, and supernovas upon your view of the sky, or populate it with flying toasters, virtualize and superimpose on the image of the real world your favorite mythical characters and imaginary companions, and provide other educational and recreational functions.

I would call the resulting image of the world **Enhanced Reality (ER)**.

Structure of Enhanced Reality

One may expect that as long as there are things left to do in the physical world, there will be interest in application of ER technology to improve our interaction with real objects, while Virtual Reality (VR) in its traditional sense of pure simulation can provide us with safe training environments and high-bandwidth fiction. Later, as ER becomes considerably augmented with artificial enhancements, and VR incorporates a large amount of archived and live recordings of the physical world, the distinctions between the two technologies may blur.

Some of the interface enhancements can be made common, temporarily or permanently, for large communities of people. This would allow people to interact with each other using, and referring to, the ER extensions as if they were parts of the real world, thus elevating the ER entities from individual perceptions to parts of shared, if not objective, reality. Some of such enhancements can follow the existing metaphors; for example, a person who has a reputation as a liar, would appear to have a long nose.

Other extensions could be highly individualized. It is already possible, for example, to create personalized traffic signs. Driving by the same place, an interstate truck driver may see a "no go" sign projected on his windshield, while the driver of the car behind him will see a sign saying "Bob's house — next right". More advanced technologies may create personalized interactive illusions that would be loosely based on reality and propelled by real events, but would show the world the way a person wants to see it. The transparency of the illusion would not be important, since people are already quite good at hiding bitter or boring truths behind a veil of pleasant illusions. Many people even believe that their entirely artificial creations (such as music or temples) either "reveal" the truth of the world to them or, in some sense, "are" the truth. Morphing unwashed Marines into singing angels or naked beauties would help people reconcile their dreams with their observations.

Personal illusions should be built with some caution, however. The joy of seeing

the desired color on the traffic light in front of you may not be worth the risk. As a general rule, the more control you want over the environment, the more careful you should be in your choice of filters. However, if the system creating your personal world also takes care of all your real needs, you may feel free to live in any fairy tale you like.

In many cases, ER may provide us with more true-to-life information than our "natural" perception of reality. It could edit out mirages, show us our "real" images in a virtual mirror instead of the mirror images provided by the real mirror, or allow to see into — and through — solid objects. It could also show us many interesting phenomena that human sen-

a multitude of alternative intelligent representation agents.

The implementation of ER extensions would vary depending on the available technology. At the beginning, it could be a computer terminal, later a headset, then a brain implant. The implant can be internal in more than just the physical sense, as it can actually post- and reprocess information supplied by biological sensors and other parts of the brain. The important thing here is not the relative functional position of the extension, but the fact of intentional redesign of perception mechanisms — a prelude to the era of comprehensive conscious self-engineering. The ultimate effects of these processes may appear quite confusing to

Of course, unless you are forced to "wear glasses", you can take them off any time and see the things the way they "are" (i.e., processed only by your biological sensors and filters that had been developed by the blind evolutionary process for jungle conditions and obsolete purposes).

sors cannot perceive directly. Giving us knowledge of these things has been a historical role of science. Merging the obtained knowledge with our sensory perception of the world may be the most important task of Enhanced Reality.

Historical Observations

People have been building artificial symbolic "sur-realities" for quite a while now, though their artifacts (from art to music to fashions to traffic signs) have been mostly based on the physical features of the perceived objects. Shifting some of the imaging workload to the perception software may make human interactions with the world more balanced, flexible, powerful and inexpensive.

With time, a growing proportion of objects of interest to an intelligent observer will be entirely artificial, with no inherent "natural" appearance. Image modification techniques then may be incorporated into integrated object designs that would simultaneously interface with

humans, as emergence of things like personalized reality and fluid distributed identity could undermine their fundamental biological and cultural assumptions regarding the world and the self. The resulting "identity" architectures will form the kernel of *trans-human civilization*.

The advancement of human input processing beyond the skin boundary is not a novel phenomenon. In the audiovisual domain, it started centuries ago with simple optics and hearing aids and has now progressed to all kinds of recording, transmitting and processing machinery. With such development, "live" contacts with the "raw world" data might eventually become rare, and could be considered inefficient, unsafe, and even illegal. This may seem an exaggeration, but this is exactly what has already happened during the last few thousand years to our perception of a more traditional resource — food. Using nothing but one's bare hands, teeth and stomach for obtaining, breaking up, and consuming naturally grown food is quite unpopular in all mod-



...give your mother-in-law a Klingon persona on your videophone.

ern societies for these very reasons. In the visual domain, contacts with objects that have not been intentionally enhanced for one's perception (in other words, looking at real, unmanipulated, unpainted objects without glasses) are still rather frequent for many people, and the process is still gaining momentum, in both usage time and the intensity of the enhancements.

Rapid progress of technological artifacts and still stagnant human body construction create an imperative for continuing gradual migration of all aspects of human functionality beyond the boundaries of the biological body, with human identity becoming increasingly exosomatic (non-biological).

Truth vs. Convenience

Enhanced Reality could bring good news to privacy lovers. If the filters prove sufficiently useful to become an essential part of the [post]human identity architecture, the ability to filter information about your body and other possessions out of the unauthorized observer's view may be implemented as a standard feature of ER client software. In Privacy-Enhanced Re-

ality, you can be effectively invisible.

Of course, unless you are forced to "wear glasses", you can take them off any time and see the things the way they "are" (i.e.,

used for augmenting people's view of their favorite object of observation — themselves. Biological evolution has provided us with a number of important self-sensors, such as physical pain, that supply us with information about the state of our bodies, restrict certain actions and change our emotional states. Nature invented these for pushing our primitive ancestors to taking actions they wouldn't be able to select rationally. Unfortunately, pain is not a very accurate indicator of our bodily problems. Many serious conditions do not produce any pain until it is too late to act. Pain focuses our attention on symptoms of

processed only by your biological sensors and filters that had been developed by the blind evolutionary process for jungle conditions and obsolete purposes).

In my experience, though, people readily abandon the "truth" of implementation details for the convenience of the interface and, as long as the picture looks pleasing, have little interest in peeking into the "source code" or studying the nature of the physical processes they observe—or listening to those who understand them. Most likely, your favorite window into the real world is already not the one with the curtains—it's the one with the controls...

Many people seem already quite comfortable with the thought that their environment might have been purposefully created by somebody smarter than themselves, so the construction of ER shouldn't come to them as a great epistemological shock.

Canonization of chief ER engineers (probably well-deserved) could help these people combine their split concepts of technology and spirituality into the long-sought-after "holistic worldview".

Biofeedback and self-perception

Perception enhancements may also be

the disease rather than causes, and is non-descriptive, uncontrollable, and often counterproductive.

Technological advances may provide us with the informational, restrictive and emotional functions of pain without most of the above handicaps. Indicators of all important/critical/abnormal bodily functions could be put on output devices such as a monitor, watch or even your skin. It is possible to restrain your body slightly when, for example, your blood pressure climbs too high, and to emulate other restrictive effects of pain. It may also be possible to create "artificial symptoms" of some diseases. For example, showing to a patient a dial indicating a spectral divergence of his alpha- and delta- rhythms that may indicate some neurotransmitter deficiency, may not be very useful. It would be much better to give the patient a diagnostic device that is easier to understand and more "natural-looking":

— "Hello, Doctor, my toenails turned green!"


— "Don't worry, it's a typical arti-symptom of the XYZ condition, I'm sending you the pills".

(Actually, a watch may serve a lot better than toenails as a display.)

Sometimes, a direct feedback generating real pain may be implemented for patients who do not feel it when their activities approach dangerous thresholds. For example, a non-removable, variable-strength earclip that would cause increasing pain in your ear when your blood sugar climbs too high may dissuade you from having that extra piece of cake. A similar clip could make a baby cry out for help every

time its EKG readings go bad. A more ethical solution with improved communication could be provided by attaching this clip to the *doctor's* ear. "I feel your pain..."

Similar techniques could be used to connect inputs from external systems to human biological receptors. Wiring exosomatic sensors to our nervous systems may allow us to better feel our environments, and start perceiving our technological extensions as parts of our bodies (which they already are). On the other hand, poor performance of your company could now give you a *real* pain in the neck...



girl next door may be wary that your "re-clothing" filters leave her in Eve's dress. Parents could be suspicious that their clean-looking kids appear to each other as tattooed skin-heads, or replace their obscenity masks with the popular "Beavis and Butt-head" obscenity-enhancement filter. Extreme naturalists will demand that the radiant Microsoft logo and Coca-Cola bottle

Distant Future

Consequent technological advances in ER, biofeedback and other areas will lead to further blurring of demarcation lines between biological and technological systems, bodies and tools, selves and possessions, personalities and environments. These advances will eventually bring to life a world of complex self-engineered interconnected entities that may keep showing emulated "natural" environments to the few remaining [emulations of?] "natural" humans who would never look behind the magic curtain for fear of seeing that crazy functional soup...

Terminological Exercises:

ER ↔ EP → ... → IE → ?! →

You must realize that most ER technologies suggested in this article have little to do with changing *reality* and everything to do with changing our *perception* of it. So at this point it would be better to call these techniques EP, for Enhanced Perception, and reserve the ER term for conceptualizing traditional technologies. The traditional technologies have always been aimed at improvement of human perception of the environment, from digestion of physical objects by the stomach (cooking) to digestion of info-features by the brain (time/clock). Since there is hardly any functional difference in how and at what stage the clock face and other images are added to our view of the world, and as the technologies will increasingly intermix, the more appropriate term may be Enhanced Interface of Self with the Environment — and, as in the case of biofeedback, the Enhanced Interface of Self with Self.

With future waves of structural

Your angry "Klingon" relatives may find unexpected allies among "proboscically enhanced" (a.k.a. big-nosed) people protesting against using their "alternative standard of beauty" as a negative stereotype.

change dissolving the borders between self and environment, the term may generalize into Harmonization of Structural Interrelations. Still later, when interfaces become so smooth and sophisticated that human-based intelligence will hardly be able to tell where a system ends and interface begins, we'd better just call it Improvement of Everything. Immediately after that, we will lose any understanding of what is going on and what constitutes an improvement, and should not try to name things anymore. Not that it would matter much if we did...

Social Implications

We can imagine that progress in human information processing will face some usual social difficulties. Your angry "Klingon" relatives may find unexpected allies among "proboscically enhanced" (a.k.a. big-nosed) people protesting against using their "alternative standard of beauty" as a negative stereotype. The


gracefully crossing their sky should be replaced by sentimental images of the sun and the moon that once occupied their place. *Libertarians* would lobby their governments for the "freedom of impression" laws, while *drug enforcement agencies* may declare that the new perception-altering techniques are just a technological successor of simple chemical drugs, and should be prohibited for not providing an *approved perception* of reality.

My readers often tell me that if any version of Enhanced, Augmented or Annotated Reality gets implemented, it might be abused by people trying to manipulate other people's views and force perceptions upon them. I realize that all human history is filled with people's attempts to trick themselves and others into looking at the world through the wrong glasses, and new powerful technologies may become very dangerous tools if placed in the wrong hands. So adding safeguards to

cont. on p.26

TED BERGER

Building Brains



Professor of Biomedical Engineering and Neurobiology.
Hedco Neurosciences and Engineering.

What area of neuroscience do you work in?

I work on the neurobiological basis of memory and learning: How the brain stores information, how we acquire new associations. Traditionally this problem has been approached by recording from single cells in the brain and seeing how their activity changes during the course of learning.

One particular part of the brain is essential for forming memories—the hippocampus. In many well documented cases where patients have suffered damage to the hippocampus, we find they still retain old memories but lose the ability to lay down new memories. This doesn't affect learning of *abilities*, only learning of fact-based information and associations.

People have tried to find out what kinds of neural representations exist, what does that activity correlate with, and then how does that change. The problem is there are five to ten million neurons in the hippocampus. So how do you learn how a system like that works by looking at the individual elements? There are too many cells. It would be like going into a computer and looking at the voltage at a point on the chip and then trying to figure out how the computer does its job. We just can't do that.

I did that kind of work for 5 or 6 years and pushed it as far as possible. It isn't enough to understand the processes at the system level and how cell activity relates to the memory process itself. What's needed is two things: You really need a mathematical model of individual elements and the whole system, so you can take the data about individual neurons and try to relate that to some structure. Secondly, you've got to have the kind of technology that will allow you to record from many neurons at the same time, and to be able to mimic the computational characteristics of the brain system when it's fully operational. You can describe it as what seems to be a series of parallel circuits. The hippocampus seems to function as a parallel processor. Parallel processors are not the kind of computers that we have on the desktop, so the computational basis for this is not easily available.

So I moved slowly from the area of neurobiology into the

areas of engineering and mathematics and began to collaborate with engineers who had developed modeling methods that were particularly good for capturing the dynamic properties of single cells and the collective dynamics of neural networks. I've also begun to collaborate with other engineers who are capable of designing computer chips that are used as a series of detectors, and we use those as electrodes to implant them into the brain to record many different neurons simultaneously so we can get the same population activity, the same population dynamics that

We're also working with a colleague in photonics where they use light signals to connect analogue VLSI devices. We're going to apply this technology to try to create a 3 dimensional structure which has the same properties of at least part of the brain system.

these brain cells exhibit.

These are analogue devices in the sense that you use conductive points on the chip as a basis for recording analog signals from the brain and then take those off the chip and analyze them in a computer. Assuming we continue to be successful in being able to record the activity of many cells simultaneously, then there's the issue of how do you take that information then mimic the computational characteristics of some part of the brain? So we've been working with some other colleagues to develop analog VLSI chips that have the characteristics of the brain cells that we've studied. We study a single cell and model the properties of that cell, then construct a circuit on the analog VLSI chip that will mimic the properties of that cell. And we've constructed the circuits for many cells and put them all onto one analog VLSI chip. Now we have a chip which essentially has the same characteristics the small population of neurons that we've looked at.

These chips have the exact characteristics of the cells they are learning. Such a chip will allow you to predict the activity of the neurons. Having developed an analog VLSI chip that

NEUROSCIENCE PIONEERS

BY MAX MORE

Even that dinosaur, the government, has declared this to be the Decade of the Brain. The premier science of this century has been physics. We have looked up to physics as the exemplary science. We have tried, sometimes misguidedly, to base other disciplines on the methods of physics. We have used physics for sources of metaphor and metaphysics. The next century will see the overthrow of physics and the installation of biology and the neurosciences as first among disciplines.

For this issue, I decided to go beyond my reading by talking to some researchers in neuroscience. Having just completed my Ph.D. at the University of Southern California, I investigated USC's new interdisciplinary neuroscience department. I was delighted to find a top class institution. The Hedco Neurosciences department was created in the late '80s and is home to an impressive collection of brilliant researchers in neurobiology, cognitive science, linguistics, engineering, and other fields. The range, depth, and innovativeness of this department may well be unmatched anywhere else.

In this series of profiles, I present just a taste of the richness of the research going on at the Hedco Neurosciences and its Neural, Informational, and Behavioral Sciences (NIBS) program. Here we have researchers building synthetic neurons, figuring out cognition and intelligent agents, understanding memory, scanning the brain, and fusing neural networks and fuzzy logic. Enjoy this glimpse of the future.

essentially has on it a small population of neurons, that's essentially equivalent to creating a hardware model of a slice through a 3-dimensional structure. They have so far created a population of nine neurons and have designed one for 100. We've essentially modeled a 2-D plane of a 3-D structure. If you want to mimic how the whole system works you need to do this in three dimensions. So we're also working with a colleague in photonics where they use light signals to connect analogue VLSI devices. They've developed a brand new technology that will allow you to stack analog VLSI chips together and sandwiched in between them is the photonic technology for connecting those VLSI chips. We're going to apply this technology to try to create a 3 dimensional structure which has the same prop-

erties of at least part of the brain system.

There will be 100 neurons in each of the planes. As many of those as we can stack together. Then we can have the 3-dimensional characteristics of part of the brain system. That allows you to study the parallel processing capabilities of a brain in a way which you can't on the kinds of computers that you use right now. We know the properties of those cells have the same properties as parts of the brain. We can begin to ask what are the dynamics of this brain system, and how is it that a network of this kind can be trained to learn something new.

What are your objectives with this research?

There are several objectives to this project. One is to understand how brain systems work. What are the computational char-

acteristics and computational limits of different brain systems? We really need to have the three dimensional structure of that brain system in a model to be able to answer those kinds of questions. There is a second objective: We want to create a hardware device which will function like the parts of the brain. There are three major advantages to a hardware device. One is you can incorporate true parallel processing. The second is speed: you can do very rapid processing. The third is size. So the second objective relates to those three advantages.

If we can mimic the computational characteristics of the brain at a reduced scale using a hardware device then there's no reason why we can't begin to contemplate *replacing* parts of the brain that are damaged, with computer chips that have the same properties and can be connected to the rest of the brain through the specially designed interface electrodes. We can sense the activity within the brain and we can send out signals into the brain. These kinds of sensing probes or signaling probes could be sandwiched on the end of a 3-dimensional structure that could perform the same function as the part of the brain that we want to replace. The replacement parts would be of a similar volume of the parts of the brain they are replacing.

The third major objective is to understand enough about how the brain works to be able to build devices to solve problems in the real world that take advantage of the things the brain does really well. One of the things the brain does really well is to associate arbitrary kinds of objects. There are some real world problems the brain is very good at and does better than any other kind of device. If we can understand what those principles are, then we can build devices that will solve problems in the real world. We've designed a device that could function as a wireless duplication receiver based on some of the first principles that we've understood about the hippocampus. That may have an application in cellular phone technology.

How far off do you think any kind of neuroprosthesis will be?

It actually depends on the parts of the brain we're talking about. There may be

lower level functions—spinal cord functions, motor systems that control the limbs—that may be possible within ten years. For replacing the kind of higher cognitive functions that involve learning and memory, that would be more like ten to twenty year range. Replacing a damaged point in the spinal cord may be possible in ten years. The tissue above and below the point of damage is functioning normally. If you can sense the activity of all the cables that are on the brain side, and you can drive the activity of all the cables that are on the lower spinal cord side, and you have a set of chips which performs the correct connection and the correct transformational activity from the brain to the spinal cord, then why not?

How do inorganic chips connect to biological neurons?

There has been a lot of research identifying the kinds of conditions that will allow neurons to attach to the electronic current sensing part of the chip. Under the right conditions cells will attach to the metal surface and stay attached. Whenever the cell exhibits electrical activity then the underlying circuit will detect that activity and transmit it elsewhere. Or, in just the reverse way, you could actually supply current to the chip and that can drive the activity within the cell. So there's a way of interfacing neurons and chips.

Personally I think it would be extremely interesting to find out how much we could enhance human brain function. I think it should be tried.

The additional problem is how to get the neurons that are interfaced with the chips to connect with the rest of the brain directly. That's a problem where there are a lot of unknowns. But we do know two essential principles: One is, nerve cells connect themselves up together. There are growth factors and a lot of other things that guide the connections from one neuron to another. They may not find the *right* pairs, but they do find each other. Secondly, it turns out that during development these connections are formed between cells that are active simultaneously.

It involves part of the same process that's used in the brain of the adult animal to store information. The strengths of connection between neurons in the hippocampus and in other parts of the brain changes as a function of activity. If two cells are active at the same time then the synaptic connection between them is strengthened. There are other conditions under which those connections can

If we can mimic the computational characteristics of the brain at a reduced scale using a hardware device then there's no reason why we can't begin to contemplate *replacing* damaged parts of the brain with computer chips.

weaken.

We now understand a great deal about the principles for how connections between neurons are strengthened, and it's primarily on the basis of activity. So if a cell has been grown to the surface of a chip and we put this chip into the brain, and we want to connect up correctly the cells that are on the interface one of the ways to do that is to drive the activity of the cells. We can control that and as a result control in part how these cells wire themselves up to the rest of the brain. Although that will be a very difficult problem we can see the

beginnings of how to approach the problem.

How does this approach differ from things like NetTalk (a neural network for recognizing speech)?

Our objective is to create a hardware model of the function of the hippocampus. It's situated in a part of the brain where after the rest of the sensory systems break down the incoming signal, determine what the features are and integrate all those features that identify a face—it's that information that goes into the hip-

pocampus. All the feature analysis has been completed. That information is processed in some way, along with, for example, the auditory sounds that the creature made, so that the features (which have already been identified as a face) and the auditory signals which identify how your name sounds, those two things get fed into the hippocampus and they're associated in some way and then sent back

to the cortical regions that do the sensory analysis, and they're stored there. This signal transformation process and the association process is done in some way that allows this human to learn this new information and to insert it into long term memory without disrupting all the other long term memories. The associations formed by the hippocampus allow each of our databases to be updated without destroying the existing databases and make retrieval of that data optimal. It's that function that we're trying to emulate. Not just learning new information or identifying speech patterns, but how to take that new item that's learned and insert it into a database so that it has the correct associations with all the other things the person has learned.

This is a unique effort. People with very different backgrounds have agreed to work together. I'm one of five people. Without the different backgrounds the problems couldn't be solved. The Hedco Neurosciences program is unique. The purpose is to get neurobiologists, biologists, computer scientists, engineers, psychologists, all in the same building. All the members of the team are willing to be members of a team. It turns out that there are an awful lot of problems that we have with the neurobiology that there are already solutions for in the field of engineering. We just don't know about them. So breaking down the disciplinary areas is extremely important.

How long will it be until we have a real artificial brain of human level intelligence?

What's most important in answering that question is that in the last couple of years we have reached a point both in understanding the neurobiology of the brain and understanding the fundamental principles of engineering and computer science where we can entertain that question; where's it's actually reasonable to ask how long do you think it will be. Just five or ten years ago, that would be seen as a science fiction question. But now it is feasible to start thinking about things like

It seems to me that we should be able to upload into a human the correct series of input signals at the right places at the right times; we'll be able to build into the memory banks new associations that we haven't in fact experienced.

that, just as its feasible to start taking about replacing parts of the brain. I don't have the faintest idea. I could say 50 years from now. I don't think it's unreasonable to think in those terms.

Much of the population is uneasy with the idea of replacing parts of the brain. They believe that there's something up here that's outside physics and chemistry. If we replace parts of the brain with things we've built, then aren't we just a machine of some kind? How do feel about that and how do you think people ought to feel about it? Should we be losing a sense of being special, or should we just realize that we're the most magnificent machines around?

To understand how the brain works and to approach the problem of understanding cognitive behavior scientifically it's imperative to treat the brain as a kind of machine, that this system can be reduced to a set of parts and they have relations to one another. When those relations are allowed to exist there's a dynamic that unfolds that explains the global behavior of the system. When you're trying to explain the complex thought processes that we engage in, when you have many elements and the dynamics of those elements are complex, it becomes a much harder problem. But nonetheless, the es-

sential tenet of the scientific approach is that everything can be treated as a machine and broken down into its component parts.

People will become more comfortable with that—with the consequences of that approach, such as with the consequence of putting computers into the brain. Once they experience the benefits... Everyone has a problem with their mother having Alzheimer's disease or their child having epilepsy. Any solution is a good solution to such problems. We not talking about changing the entire function of the

individual but simply replacing the part that used to be there, replacing the function that used to be there. Then I think a lot of that resistance will melt away.

What if we do start talking about altering and enhancing our capabilities? Without having to sit down out a computer, what if I can work out complex equations and trajectories, what if I can do all those things in my head? What if I can work out complex strategies and patterns that would normally have to do on a very powerful computer — do you think people will be more upset about those?

Yes, I think there are going to be great social debates over whether we should enhance the capabilities of the brain, and do anything other than replace non-functioning parts of the brain. As soon as the work began to move towards enhancing function of the brain I think that would cause a great deal of social debate. Personally, I think it would be extremely interesting to find out how much we could enhance human brain function. I think it should be tried. Now we're talking way downstream, but if we can replace brain function, well why not try to find out how well we can enhance brain function. I think it would be incredibly interesting to

try. I would love to be able to remember things that I forget. How many times has it been that you wished you were in a certain cognitive or emotional state, and you can't be in that state. It would be an incredible advantage to be able to have that choice.

What do you think of the idea of uploading the contents of a brain into a computer?


I think it's far more likely that these technological developments will allow the uploading into the brain of information from the computer. If we can understand how it is that certain signals input into a system, how that neural representation is transformed and how it's associated with other representations, then it seems to me that we should be able to upload into a human the correct series of input signals at the right places at the right times; we'll be able to build into the memory banks new associations that we haven't in fact experienced. That's because we've identified a very discrete part of the brain that's important for laying down new memories. But we don't know where the memories are stored. We think we know. We think they're stored in the neocortex—phylogenetically the newest part of the brain. But exactly where and how it's stored no one really has much idea.

Do you personally look forward to having some neurons replaced, some functions augmented?

I would love to. That would be extremely interesting to me. It's a challenge in sense of being an entirely different dimension of testing how well you've understood the system. Replacing functions is one thing, but when you're trying to enhance brain function, that's potentially different problem. We might not know how to change properties of the brain so you enhance the functions you wanted without disrupting other functions. That's a very hard problem and a very interesting one.

So if it takes another 100 years or 150 years, would you want to stick around for that?

Oh, you better believe it! Without a doubt. I'd love to.



Skeptic Magazine

Provocative & Informative!

SUBSCRIBE NOW AND SAVE \$10 OFF THE REGULAR SUBSCRIPTION PRICE, PLUS GET A FREE BACK ISSUE OF YOUR CHOICE!!!

"One of the...top 10 best new publications!"

— Library Journal

"Skeptic kicks ass."

— Penn Jillette of "Penn & Teller"

"...the best journal in the field."

"Let me congratulate you on the superb design, layout, typography, and intellectual content that has already made Skeptic the best journal in the field. I wish you the very best for the success of your important magazine and its even more vital mission."

— Stephen
J a y
G o u l d

Are you interested in fringe and revolutionary science? Are you intellectually curious about the paranormal? Subscribe to *Skeptic* magazine and learn what the world's leading scientists and experts have to say. Read investigative reports on pseudoscience. Succinct summaries of both sides of contemporary controversies and guides for further investigation. *Skeptic* Magazine devotes more space to opposing views than any other forum for reader's responses. *Skeptic* also provides reviews of noteworthy books and articles of related interest. *Skeptic* is a must read for curious minds.

The Skeptics Society is a member-supported organization of scholars, scientists, historians, magicians, and the intellectually curious, that sponsors a lecture series at Caltech and publishes the quarterly magazine *Skeptic*, for the purpose of promoting science and critical thinking, and disseminating information on pseudoscience, pseudo-history, the paranormal, superstition, magic, fringe claims and groups, revolutionary science, frontier-science, and the history of science and pseudoscience, in articles, essays, reviews, and letters. Transcripts of Caltech lectures are reprinted in *Skeptic*, and videotapes of the lectures are available through the magazine.

BACK ISSUES: \$6.00 each

V.1, #1-Tribute to Isaac Asimov
By Steve Allen, Harlan Ellison, Martin Gardner; The "Amazing" Randi on the Paranormal; A Skeptical Manifesto; Triumph of Science.

V.1, #2-Can Science Cheat Death?
Cryonics: Reaching for Tomorrow; Q & A About Cryonics; Can Science Cheat Death?; Society of Persons Apparently Dead; Abuse of Statistics.

V.1, #3-Revolution in Evolution?
Punctuated Equilibrium by Gould; Punc Eq at 20; Historic 1842 Darwin Sketch on Evolution; The Darwinian Culture; Edgar Cayce investigation.

V.1, #4-Heretics, Scientists & Resistance to New Ideas
Spirits, Witches & Science; Edgar Cayce Foundation Response; Can Religion be Rational?; Today's Heretic, Tomorrow's Saint; Ideological Immune System; Copernicus as a Heretic.

V.2, #1-Genius: Myths & Reality
By Steve Allen, Paul MacCreedy, James Randi, Marilyn Vos Savant; Gleich on Feynman; Gene Roddenberry; The Amadeus Myth; Noah's Ark; Who & What Mattered in History?

V.2, #2-Science, Religion, Cults
Waco, Jonestown, and All that Madness; Millennium Watch; The Jesus Cults by Steve Allen; The Resurrection Myth; Was Christianity a Cult?; Realism and Religion by Milton Rothman; The Unlikeliest Cult; Mathemagics.

V.2, #3-Fad Psychology
False Memory Syndrome; Facilitated Communication; Diagnoses are not Diseases by Szasz; William Reich; Sex, Brains, & Hands by Halpern; Pseudoscience/Psychiatry by Tavis; How Thinking Goes Wrong; Flood Myths; The Ark Hoax; Anti-Science; Millennium; Mathemagics.

V.2, #4-Pseudohistory: Holocaust Revision & Afrocentrism
Holocaust Revisionism; Free Speech & Neo-Nazis; Denying the Holocaust review; Creationism & Racism; Goodby Columbus?; Were the Egyptians & Greeks Black?; Beyond Race; Noah's Ark Hoax Update; Randi on the Millennium.

V.3, #1-Pseudomedicine
Homeopathy; Therapeutic Touch; Alternative Healthcare; Randi on Liquefying "Blood"; Skeptics & Religion; Update-FMS; Leftist Science & Skeptical Rhetoric; Mattoon Mass Hysteria.

V.3, #2-The AIDS Heresies
The AIDS Heretics; Charles Murray Interview; Astrology Brouhaha; Extraordinary Claims/Extraordinary Evidence?; Wallace, Darwin/Priority Dispute; Culture Wars; Debunking Hal Lindsay; Should the Bible be Taken Literally? Why Did God Make Rice Cakes? Star Trek.

V.3, #3-Race, IQ, & Intelligence
Skewed Logic of *The Bell Curve* by Halpern; A Place In The Sun by Tavis; Robert Sternberg interview; Defense of *The Bell Curve* by Sarich. Scientology v. Internet; "Curse" of King Tut; Moral Panics; FMS Confession; Eat a Light Bulb.



YES! I want to join the Skeptics Society and subscribe to *Skeptic*, a quarterly magazine, at the special limited time offer price of \$25. I will **SAVE \$10** off the regular price of \$35. The **FREE BACK ISSUE** that I want with my subscription is Vol. _____ # _____.



YES! I want to order additional back issues, videos, or books using the order form on the back of this page.

CHARGE IT! FAX your credit card order to 818/794-1301. To order by mail send a check made out to Skeptics Society, or your credit card number to: Skeptics Society, P.O. Box 338, Altadena, CA 91001. For information phone 818/794-3119. E-mail us at skepticmag@aol.com

NAME _____
 ADDRESS _____
 CITY _____ STATE _____ ZIP _____
 PHONE _____
 Charge my:
 MASTERCARD VISA EXPIRATION DATE _____
 Month _____ Year _____
 Account No. _____

Signature as name appears on credit card

Special price *Skeptic* Subscription/Membership at \$25 for one year, (includes one free back issue, with free shipping.).....

Special price *Skeptic* Subscription/Membership at \$45 for two years. (includes one free back issue, with free shipping.).....

Foreign Subscription at \$55 for one year, (includes one free back issue, with free shipping.).....

Foreign Subscription at \$90 for two years (includes one free back issue, with free shipping.).....

Total for additional *Skeptic* back issues, plus books and videos from the back of this page. (Please don't forget to include shipping and handling).....

TOTAL ENCLOSED _____

Subscribers—Moving? Please notify us immediately. The post office does not forward magazines!

MICHAEL ARBIB

Intelligent Assistants

Michael Arbib's two primary appointments are in Computer Science and Neurobiology at the University of Southern California.

What has your work focused on in the past?

From undergraduate days I've been interested in what was then called cybernetics: the attempt to draw parallels between brains and machines. I've stuck with that ever since. Trying to understand how to build intelligent machines and trying to use new concepts to probe how the brain really works. My concern with philosophy of mind has really been a corollary of that, trying to see to what extent my view of the brain as a highly unusual computer, very different from anything we've built so far, could factor into an analysis of mind.

"So you've been working primarily on the computational aspects rather than the neurobiology?"

I'm very interested in the question of how we use vision to structure our environment on the basis of which we can act. Of course that has a corollary in understanding the structure of memory so that experience in the world shapes how we behave

MIT in particular and many people in mainstream AI [artificial intelligence] in general had become blinded to what we in the brain or neural nets community knew about and then they adopted it.

in the future. So that work has really gone in two ways. One way is in working very closely with experimentalists, some doing neurophysiology on frogs to look at how creatures might go for basic survival things, others working with rats in terms of how they learn a complex spatial environment, other studies with monkeys, looking at eye, arm, and hand movements

and trying to tie that down to detailed functioning of the brain. Finally with humans, with behavioral experiments and using brain imaging to capture the changing activity of the brain.

So, this provides the inspiration for a lot of the computational models, and then the predictions of the models factor back in to the design of new experiments. The other side of this is taking some these ideas about an animal or human can find its way around the world and applying that to the design of robots, both hard-arm robots and to mobile robots who have to locomate their way around a complex environment.


What is the current state of robotics research? Hans Moravec spoke at the EXTRO² conference about a robot-driven truck...

The approach that the CM [Carnegie-Mellon] group took was to take a van and drive it on roads and recognize where the

side of the road was, so the vehicle could drive along the side of the road. They haven't yet really looked at what happens when there's traffic or obstacles in the road. But we're allowing the environment to be more complex. Our study of frogs has got us to think about competing strategies and trade offs: there's somewhere you want to go and there's something you

want to escape—do I go down that road where's there's good food, but there's a risk of being attacked?

Similarly with the hand-eye stuff, the



state of the art tends to be you tell the robot to do some stereotypical task. We've been looking at monkeys and humans to see how you can continually monitor the environment using vision to restructure the instructions to the robot, then going even further using adaptive neural nets so that we don't have to explicitly understand the problem. We can have the robot learn by being put in a variety of situations and it will improve its performance over time.

Are processors fast enough to respond to the real world right now?

In a lot of applications. The vehicle you mentioned uses a neural net to learn to recognize different fragments of a road scene captured by the camera and therefore learn what is the appropriate way to steer. In that case it is fast enough. The power of workstations has increased so much in the last few years that for many applications you can get away with these. But it's true that they have limitations. At USC we have Berger's group, and we

have groups working on using VLSI electronic chips to make parallel neural networks, and also a group working on optical methods.

The idea is that as we build more powerful algorithms in neural nets we need more computational power. In some of my work in modeling the brain of a monkey we may use tens or hundreds of thousands of neurons, and some of those neurons may be automatically changing their connections in a learning procedure. And then there's no way with such a complex network we can get real time performance with an ordinary computer. In the long run we think we'll need a network of that size.

Are these networks being designed to be very similar to the human brain, or having neural networks that will be effective regardless of how closely they mirror the brain?

I suspect that there's always going to be a tradeoff between that work that's trying to understand the brain and trying to build in the details of the biology and chemistry of the brain and those that are seeking technological efficacy. If you want to add numbers, you can have a special purpose chip that can do it effortlessly and a million times faster. For technology the migration path is not that you study the brain and then imitate it, rather you study the brain to get ideas and then you use those ideas to get technological efficacy. For example, one approach to putting neural networks on a chip is not that you represent all the connections that are present in the brain, but you use time-sharing so that particular processes can be used to do computations for many different neurons. In electronics the speed is thousands or millions of times faster than that of the individual neurons. On the other hand we have perhaps a 100,000 connections for every cell and it's just impossible in near future technology to put that many connections into an electronic chip. There you trade off the speed of the computing against the time-sharing of the communication links. You have a technological answer that's inspired by the brain but exploits the nature of the electronic medium.

Do you have an opinion on Marvin Minsky's

Society of Mind thesis?

Well it's funny because Marvin in about 1975 at an AI meeting gave a very staunch talk on why intelligence must be serial and why any consideration of distributed computing is totally inappropriate. And then about two years later he published *Society of Mind* which was very consistent with what those of us had stuck with the brain throughout and tried to build our models of intelligence on that basis agreed with. So I have always had, from both from modeling the brain and from schema theory, the attempt to understand the interaction of different units of the brain. I have always had a view that is similar in spirit at least to Minsky's society of mind.

Rodney Brooks a few years later put out his subsumption architecture, which again was amusing in the sense that it was someone in the heart of MIT's AI group finding religion as it were, talking as if this idea of having different layers of

Penrose does this incredible leap which says we really need some breakthrough in quantum gravity in order to solve the problem. [But] None of those subtleties bring us up against the need to invoke a new physics.

computational interaction with each other was a new departure. But it was rather that MIT in particular and many people in mainstream AI in general had become blinded to what we in the brain or neural nets community knew about and then they adopted it. Now the lines are very much blurred. A lot of people in AI, a lot of people in autonomous robots, in animal behavior, in brain modeling, subscribe to these ideas that overlap what Minsky, Brooks, myself, and others espouse.

Do you have views on positions by Penrose or Searle?

I've written criticisms of both. Searle's writing reminds me of a great story about Norbert Wiener who, many years ago, was said to have proved a very big mathematical result called the Reimann hypothesis, and all the mathematicians came flocking to MIT to hear this. Wiener started proving this on the board, and then finally starting writing less on the board


and puffing more on his evil black cheroot and finally said, "It's no good, it's no good. I've proved too much. I've proved there are no prime numbers." I think if you look at Searle's argument, because he goes through these intermediate steps and yet believes at the end that the brain is a special kind of machine, I think he's proved too much. I think he's proved that we can't think, that we're not intelligent. Because he doesn't have any subtle appreciation of how accumulation of complexity can yield a difference in kind.

Now Penrose is quite different. Penrose is quite bizarre. I analyzed his second book in immense detail in the London *Telegraph*. I can summarize his argument like this. He first says machines can't think because Godel showed that no matter how good a set of axioms you put down for arithmetic, there would be truths for arithmetic that wouldn't be theorems. But that totally ignores that we do not

operate *our* intelligence in terms of fault-free inference. We're continually making analogies, we're continually making mistakes, and we're continually learning from our mistakes. So the consistent axiom system model just has nothing to do with the way the brain is. Although Penrose in his second book has discussed that, it's not satisfying. He insists that *he* has shown that machines can't think.

Then he does this incredible leap which says we really need some breakthrough in quantum gravity in order to solve the problem. The way in which this breakthrough is going to solve the problem is that there are things called microtubules in nerve cells and these are so fine they could exhibit some weird quantum gravity effects that would break the Godel barrier and allow the brain to be intelligent. There's so much work on microtubules which show the supporting structure and transfer of chemicals through the cell. There is nothing in neuroscience,

except for a few people very much on the fringe, which says there is a gap. The problem about the brain is that it has about a hundred billion components with on the



Hedco Neurosciences, University of Southern California

average of ten thousand connections. We're continually learning more about the molecular biology of the cell. Not only do we have more elements to work with, each of those elements is revealing more and more subtlety. None of those subtleties bring us up against the need to invoke a new physics.

How is all this growing information about the brain to be accessed and used effectively?

I spearhead the USC Brain Project. This brings together many different students and faculty to build databases about the brain, to provide new tools for visualizing the brain. Then to provide simulation tools so we can try to capture what we know in models of increasing complexity in such a way that it becomes *relatively* easy to go back and forth between huge

arrays of data stored in the database and models which require immense amounts of computation to understand how things hang together. I believe this, rather than a breakthrough in physics, is the way we're going to improve on our current understanding of the brain. The database consists of several parts: How do we get different data all the way from the molecules up to human activity? Multimedia... We may need recordings of sounds that cells make while they're active, we may need pictures of different parts of the brain, we may need atlases. Finally, the simulation component.

Are you willing to speculate how long it will be until we have human-level machine intelligence?

I'm not sure what you mean by human-

Once you put it in the silicon or whatever, as distinct from in a body with all these physiological markers of emotion, then I'm not sure to what extent you'll want to see it as an emotion rather than as hierarchical priority setting.

level AI. If you mean as good at music as Beethoven, as good at physics as Einstein, that's rather hard. If you mean as good at physics as you are, as good at musical composition as I am, then it's not that hard. I suspect that we will not try to replicate general human intelligence, let alone genius, but rather we'll try to build conversational interfaces so that people can sample repositories of information

using the sort of conversation they would use in asking a human expert, with having themselves to be expert programmers or computer users. We already see that it can be done in very specialized domains quite effectively.

I have no problem with saying that in 20 years you will be able to talk to your computer about all your day to day assignments as if you were talking to a human assistant. The conversation may be similarly broad ranging. Notice that that answer includes two different things: One is the advance in the science of intelligence or understanding of the brain and the other is in miniaturization. So you'll probably have a range of assistants, just like you might talk to one person about politics, someone else about the weather, and someone else about good restaurants. In the same way, you might not have all the data in your own personal computer, but

with the advances in networking you might well plug your computer in, in terms of what you want.

Are you confident that it will be possible to build intelligent machines without emotions? Or will we need to build in something like emotions to give them motivation, to give them interest in problems?

*is thinking
for yourself
a crime?
we don't think so...*

"Badly needed"
--Robert Anton Wilson, Author, *Prometheus Rising*

"Separates the wheat from the chaff"
--Joan Kennedy Taylor, Vice President, Feminists for Free Expression

"Exciting resource"
--Peter Breggin, MD, Author, *Toxic Psychiatry*

Tools to increase your critical thinking skills:

- ⌘ **Independent Thinking Review** (quarterly journal, \$18/4 issues)
- ⌘ **Book Catalog** (critical thinking/sensible psychology/questioning authorities/springing belief traps)
- ⌘ **Children's Catalog** (critical thinking books for kids)

For samples, send \$1 to:
Resources for
Independent Thinking
484 Lake Park Avenue, N^o 24X
Oakland, CA 94610 • (510) 601-9450

RICHARD THOMPSON

Memory and its improvement


Richard Thompson, Ph.D. is Keck Professor of Psychology and Biological Sciences at the University of Southern California and Director of the Program in Neural, Informational and Behavioral Sciences at USC. Previous positions include Professor of Psychology at Stanford and Harvard Universities. He is President of the Western Psychology Association and President of the American Psychological Society.

What area of neuroscience has your work focused on?

My particular interest has been for a long time brain substrates of learning and memory, specifically how it is that the brain can code, store, and retrieve memories. No one knows the answer to this question for any kind of memory. What we have learned over the last couple of generations is that there are many different kinds and forms of memory. My own view has always been that whether or not a memory is localized to a particular place or is distributed to a number of places, we won't be able to analyze the mechanisms of memory storage at a molecular and

be distributed widely throughout the hippocampus. The hippocampus is critically important in *declarative* memory. We've been interested in both *declarative* and *procedural* memory. Even in elementary forms of learning like classical conditioning there are massive increases in neuronal activity in the hippocampus, so the hippocampus is coding this form of memory but in a different way. The basic paradigm we use, classical conditioning, is a procedural kind of learning. Simple forms of learning have the same properties in animals and humans.

What we find in mammals is that the essential memory traces for this kind of learning are stored in the cerebellum. This is the first time that anyone has been able to localize a particular form of memory storage to the cerebellum. By "essential" I mean that you can lesion appropriate regions of the cerebellum, and completely prevent the learning and completely and permanently abolish memory of the learned response without interfering with the ability to make the response. What's changed is that the animal can no longer learn to associate that response to any



that learns about the *significance* of what's going on. It learns that the tone means that there's something bad going to be happening to the eye and that it had better do something about it. These are very different kinds of learning situations.

There are still other memory systems. One is a learned fear system. Yet another forebrain structure, the amygdala is critical for the learning of fear. We're interested in all of these, ultimately we're interested in human memory in all its guises. That's why the work on the simple learned response has been so satisfying. All of the work we've done on the rabbit has been replicated in humans.

More generally, as Director of the Neuroscience program here, we are interested in understanding the brain in all its facets. How the human brain works in the ultimate goal. A unique aspect of our program compared to many others is the emphasis on mathematical, computational and cognitive aspects of neuroscience. Like most other programs we have strong groups in molecular and cellular and systems level neuroscience, but here we also have very strong groups in computer science—the kind of work that Michael Arbib and Christopher von der Malsburg do in making models of neural systems; we even have people like Armand Tanguay, an engineer who's developing optical com-

We know the basic circuitry of the hippocampus, how it works, how we can produce long term potentiation. We can build that into a chip.

cellular level until we know where the memories are stored. They don't have to be stored in one place but we need to identify what parts of the system are involved in the memory storage prior to analyzing the mechanisms.

Has that been narrowed down at all? Such as to the hippocampus?

The hippocampus is critically important. If any structure is distributed, it's the hippocampus—distributed in the sense that information projected into the hippocampus is not represented by stimulus modality—you don't find cells in the hippocampus that respond to visual or auditory stimuli. Instead information tends to

neutral warning experience. The eyeblink is a form of defensive reflex. A tone sounds and the animal learns to close the eye and avoid the air blast. However, once the lesion is made the organism can no longer learn to make that adaptive response. We've been doing a lot of work recording neural activity, looking at chemical changes, electrical stimulation, to localize where these memories are formed in the cerebellum.

When animals or humans learn these simple responses, there is a system in the brain that learns what to do about it, which in this case is the cerebellum—a procedural system. Then there's another system, a hippocampal cortical system

Memory and its improvement

puting systems. In this building we have groups ranging everywhere from molecular neurobiology and cellular neurobiology through systems, cognitive and computational approaches to the study of the brain, and people even in areas like linguistics.

A significant part of the interest we have in our program relate to your field—philosophy, because we're terribly interested in the fundamental problem of all time: What is the relation between brain and mind? It's still a wide open question. I don't mean we haven't made progress. I think it's safe to say that whatever the

What about vasopressin?

Yes, that's interesting. We published a study looking at vasopressin's effect on transmission. It enhances synaptic transmission. The animal literature says vasopressin does enhance memory, as do some of the hormones like norepinephrine. I say these are significant, but they're not huge effects though statistically the animals are significantly better than without it. I doubt that we can do much better. Perhaps we can. Some people are developing drugs to modulate a certain receptor in the hippocampus. There's animal evidence that those drugs may enhance

cal stimulation. So if you have single axons on single connectors on a transistor in the right circumstances, you can activate individual neurons. I say easily but it's not practically possible to do it now but it's certainly conceptually possible. And of course you get into the possibility that we could plug a brain right into a huge computer.

I see nothing conceptually different between a complex computer and a human mind. It's materialist view, but so far no one has come up with any reason to doubt it... I see nothing in principle impossible about building a computer that has all the properties of a human brain, plus a better memory!

I see nothing conceptually different between a complex computer and a human mind. I see nothing in principle impossible about building a computer that has all the properties of a human brain, plus a better memory.

mind is, it's an emergent property of the brain. Without the brain there is no mind. Beyond that—what that means—we don't know.

The fundamental problem from a scientific point of view is to figure out methods of measuring the mind independent of behavior. The two approaches right now are to study behavior and the other is to record neural activity. We don't have any

memory performance as well.

There's another approach that I think may eventually be possible. Not a drug approach but rather the approach of hooking up a brain to a computer. Ted Berger, Armand Tanguay, and I planned to do some research in that area, ultimately trying to attach computer chips to nerve cells. It's been done a little bit with invertebrate nerve cells which are much stur-

Are you familiar with the practice of biostasis, or cryonics? Do you have a view about the possibility of frozen persons being revivable in the future?

You can freeze sperm right? They are perfectly viable after X years. There are a lot of technical problems, such as with crystallization, for example. But I don't see anything in principle impossible about it... We do think that long term memories ultimately must have a structural basis in terms of changes in actual structures of synapses.

What do you see as the most exciting and promising research areas over the next 20 years?

There are two areas: the ends of neuroscience. One end is the molecular, cellular aspects which get down to the changes in the genome. It's increasingly clear that if you're going to have structural changes in synapses to code memories, that requires changes in gene expression and that involves the whole cascade of molecular machinery inside the nerve cell which produces the structural substrate of these long term memories.

The other end is computational cognitive neuroscience which involves everything from PET scanning and brain imaging to developing computer models of the brain to eventually developing pieces of brain-like chips that can be stuck in brains. Those are the two areas in which I think the next level of advance will be made.

I don't see anything in principle impossible about it [reviving frozen brains]

other independent method for measuring mind. So it's a still a problem because the introspectionist method is very unreliable.

Do you think there are avenues towards eventually being able to improve human memory beyond the healthy norm?

There are just a few drugs that do improve memory performance significantly beyond normal performance. Two of those drugs, neither of which is recommended, are amphetamine and nicotine. Both do produce significant improvement in complex memory like declarative memory—both reasoning and memory, and perhaps attention.

dier than mammalian nerve cells. You can grow a few invertebrate neurons in a dish and they're pretty tough, and you can grow them on a transistor grid and actually make functional connections between the neurons. This is purely speculation. We know the basic circuitry of the hippocampus, how it works, how we can produce long term potentiation. We can build that into a chip. What if someone's hippocampus is damaged to the point where they are having memory problems. Couldn't we eventually develop a chip-like hippocampus that we can implant into the human brain to substitute for the damaged tissue. It's not inconceivable because you can activate nerve cells very easily by electri-

RICHARD LEAHY

Imaging the Brain

USC, Professor, Electrical Engineering—Systems.

You work on imaging brain function. What methods are used for imaging the brain?

There are two basic types of imaging systems: *Anatomical imaging* systems, and *functional imaging* systems. X-ray, CT, CAT scans, are anatomical imaging systems, which produce images of soft tissue and bone, and won't tell you what the body is *doing*. You can scan a cadaver or a live person and they look much the same. Standard magnetic resonance imaging (MRI) is similar in that it produces images of anatomy. It images hydrogen nuclei and the molecules they are connected to.

The more interesting modalities are the functional ones. Of those the best is

Magnetoencephalograph (MEG)



probably PET scanning (Positron Emission Tomography). You inject into the subject a positron emitting isotope of something you're interested in. If you want to look at brain metabolism, you inject an analog of glucose that has a fluorine-emitting positron isotope in it; if you want to look at blood flow, which is used for activation, looking at changes in brain states as you're doing different cognitive tasks, then you use a positron isotope of oxygen in water, or you can look at neurotransmitters, neuroreceptors by choosing very specific agents.

This gives you an image in real-time?

No, one of the main limitations of positron tomography is that the resolution is slow; it's from ten seconds to minutes or hours. What you're detecting is that the positron is emitted from the isotope, it annihilates with the electron and produces a pair of high energy photons that gets detected by a set of scintillation crystals surrounding the patient—it uses antimatter to image what's happening in the brain. By detecting the pair of photons that are produced you can tell with some uncertainty where the positron was emitted from, which is where the molecule it came from is located. So it produces this image of positron emissions and those are a direct correlate of the spatial distribution of whatever it was you introduced into the body. If you want to look at glucose metabolism in the brain, the areas where there's the highest emission of positrons is where the largest amount of glucose was metabolized.

With glucose metabolism the person does the task and then they put them in the scanner. It gets partially metabolized and gets trapped in the brain so then you put them in the

scanner 20 minutes afterwards. It's not very specific. It detects broad areas of activation across the brain. You see differences in schizophrenics and normals, or people with Alzheimer's disease. With Alzheimer's you see a general decrease in activity; with more specific diseases such as Parkinson's disease you see one particular part of the brain that's lacking activity.

Don't you need only very small amounts of antimatter to combine with matter to cause large explosions. Is that a problem?

You do generate quite a lot of energy. Antimatter is the opposite of matter, so the antimatter particle for an electron is a positron. You put the two together and they annihilate and produce quite a lot of energy. I don't think your head's going to explode!

It takes some work to generate a positron isotope. You have to have a cyclotron nearby to generate these particles, and they have very short half-lives. It's an extremely expensive modality because, in addition to the PET scanner which costs a couple of million dollars, you need a five million dollar cyclotron.

There's been a lot of recent interest in functional MR. That looks at local changes in blood flow in the brain. If you do a simple experiment like flash lights in front of the patient who is lying on a scanner, and you take an image then you take an image without the lights, then when you're flashing the lights there will be increased neural activity in the visual cortex. That increased activity requires replenishment of the cells nearby so that blood flow to that area increases. The inflow of oxygenated blood brings paramagnetic particles with it, and that changes the local magnetic properties of the brain in that area which changes the image. So you take the difference between the activated image and the rest image; you see little white spots where the brain has been activated. That's exciting because everybody has MR scanners but not so many have PET scanners. You don't have to

inject radioactive material into people to find out what they are thinking about.

What prospects are there for doing real time scanning?

With positron tomography it's impossible. You're imaging individual photons or pairs of photons, and one photon tells you next to nothing. You have to collect on the order of at least a couple of hundred thousand for a two-dimensional image. There's a physical limit to the rate at which the machine can count. It's not the radiation dose that limits, it's the machine itself.

For magnetic resonance imaging, where you can look at local blood flow, you don't have that limit. You can do fast scanning. With the newer systems you can collect 20 frames per second. But you're looking at changes in blood flow and that has its own time constant, on the order of not more than a second. So it's the physiological process that's limiting you there.

There's only one modality that allows you look at the brain at the rate at which it functions, and that's magnetoencephalography (MEG), which is what I work on. This uses the magnetic fields of the brain. You have an array of magnetometers around the brain. They consist of a pair of coils with a Josephson Junction on the end. You need a pair of coils because if you just had a single coil all you'd measure would be the Earth's magnetic field, which is massive—about seven orders of magnitude than any field that comes from your head. You monitor around a thousand samples a second in each of these magnetometers. The biggest system that has been built so far has 122 magnetometers.

The sources are assumed to be the pyramidal cells in the cortex, specifically the current flow in the dendrites. You have synchronous or near synchronous activation of several thousand neurons together. This produces magnetic flux that cuts each of these coils and that produces a signal that you measure at the outputs for these devices. You use Maxwell's equations to give you the relationship between the source and the measurements.

You can combine this with EEG. You can put electrodes on the scalp and simultaneously measure the EEG signal. You need to average about a hundred of these for a typical stimulus, like a tone played in the ear. At most you put on 128 electrodes. They take a long time to put

on! The nice thing about MEG is that you just sit with your head inside a gantry.

Do you see new technologies coming along, or refinements of the current ones for improving scan resolution?

Yes. The resolution of functional MR studies will probably improve. It's still improving. There are no well worked out theoretical limits. In positron tomography you're fundamentally limited by the fact

Everybody has MR scanners but not so many have PET scanners. You don't have to inject radioactive material into people to find out what they are thinking about.

that what you're measuring is the point of the annihilation of the positron and electron. The positron gets emitted from the nucleus and travels up to a couple of millimeters before it annihilates with an electron. So, no matter what you do, you can never get resolution beyond that couple of millimeters.

MR doesn't have that fundamental limitation. There are some systems built to scan little animals. They have high field magnets with a small bore that can give you ten-of-microns resolution. In principle you can get that same tens-of-microns resolution in a human scanner, but you might get such huge fields that it's not feasible. A lot of it's the engineering to build machines with high enough and clean fields with nice linear gradients.

With current systems, you're looking at activation of around ten thousand neurons. The most sensitive techniques use a positron isotope of carbon on some relation of dopamine. That can show you very low levels of chemical activity. You can do that because with positron tomography each event is a single molecule or a single nucleus giving off a positron. You can do that with a couple of hundred thousand molecules and get in image in two dimensions. That's many orders of magnitude finer than you're looking at with magnetic resonance. You're fundamentally limited by noise considerations.

It depends how invasive you want to get. You can record a single neuron if you want to stick a micron probe in there. I've spend a lot of time trying to get high spatial resolution. My gut feeling is that there's got to be some undiscovered method because, until it came along, MR was unknown and in retrospect, given

that it had been used in chemistry for about 30 years before they figured out you could image with it. In hindsight it's obvious. There may well be some other technique out there that hasn't been considered.

One of the limitations is whether you're dealing with a passive or an active system. One of the problems of MEG is that it's entirely passive: it's just measuring the fields that you produce outside

your head through what's happening in this very complex three-dimensional volume. PET and MR get around that by selecting a specific part of the brain and encoding something on to that so you're only looking at that part of the brain. Those active techniques are able to do something that you can't do with a passive system. With MR you can only look at a certain part of the brain at a time, though that's improving. With PET you look at the whole brain, but the resolution is limited. In both cases you're looking at analogs of neural activity rather than the activity itself. You're looking at glucose metabolism, which occurs over a timescale of minutes, or blood flow over a timescale of seconds. What you want to look at activity on the millisecond timescale. MEG is the only modality that but it doesn't let you image. Maybe the future is to combine the relative attributes of both, so you get good spatial resolution from one and good temporal resolution from the other.

What if you had a brain in suspension, so that you can take as much time as you like to scan it. Can you then get a much higher resolution scan?

Yes. You'd have to slice it up and put it under a microscope. It would take up a lot of space. The closest thing to that for the whole body is this Adam and Eve project. They took a male and female cadaver, microtomed them and frozen them, enclosed them in cellulite. They chopped off half a millimeter at a time, photographed it, and scanned that into a computer to make a 3D image of the body. The results are on the Web. It's gigabytes of data.*


*[This is the Visible Human Project. See: <http://www.nlm.nih.gov> —ed.]

BART KOSKO: Fuzzy Logic & Neural Nets

Prof. Electrical Engineering, USC

Q: What is the focus of your current research?

A: My research focuses on how nonlinear systems learn. Most math models in science are linear models or linearized models even though no one has found a truly linear process in nature. Nonlinear math models have more power and accuracy but have few closed-form properties or guarantees of stability or convergence.



The truth is the space shuttle runs as much on computer simulation as it does on formal models of space mechanics.

Neural and fuzzy systems are tools that let us model nonlinear systems without having to guess at their exact mathematical form. No one knows the equations that govern single-lane platoons of smart cars that drive 70 mph over curves and bumps. Neural and fuzzy systems can approximate these systems with finer and finer accuracy as test data and expert advice come in. They learn in the sense that the test data tunes their parameters. The data may change the synaptic weights in a neural network or change the size or structure of a rule in a fuzzy approximator.

I am working with students and colleagues on applying these tools to a wide range of problems in signal processing, control, and multimedia systems. These applications range from image compression and motion estimation in film to

smart-car control and the design of neural-fuzzy intelligent agents. Meanwhile I always keep working on a few purely abstract math problems but like to keep them private until the theorems fall in line.

Q: In what ways do the kinds of neural networks you use differ from actual biological brain function, and in what ways are they the same? What practical significance is there to such similarities and differences?

A: There are four broad types of neural

There is little evidence for supervised learning in flesh. That would be like you not making a move until you saw what effect you had on the gross domestic product or the interest rate of your country.

models based on how they learn and the structure of the synaptic connection topology. Learning is **unsupervised** or **supervised**. The topology is **feedforward** or **feedback** (with closed loops).

Biological neural nets are largely **unsupervised feedback** neural networks. These tend to occur deeper in the cortex and away from direct sensory input. I developed a large family of unsupervised feedback models in this area called RABAMS or random adaptive bidirectional associative memories. These models can learn new patterns while they display and recall old ones. That means their synapses and neurons can both change at the same time though on different time scales. Synaptic changes are slow and neuronal changes are fast. Real nerve nets certainly have this property. RABAMS can also operate in the presence of a great deal of noise or "unmodeled effects" and they give back as a special case many popular feedback neural models: Hopfield circuits, adaptive resonance theory (ART) nets, simulated annealing or "genetic algorithm" nets, and others. The trouble is no one has figured out how to get such feedback neural nets to do anything of real practical interest. A decade ago researchers put forth some tantalizing toy prob-

lems but they were just that and researched interest soon waned. Someday research interest will surely return to this no-man's land of nonlinear feedback dynamical systems.

The most popular neural models are **supervised and feedforward**. A godlike teacher has to tell each neuron or synapses how well it helped or hurt some global payoff measure of error or cost. Signals flow from left to right in the feedforward net. So the teacher passes the error data backward from right to left to adjust all the parameters and may have to do this hundreds of thousands of times for a data set. These nets have real power in terms of

learning complex boundaries between cancerous and non-cancerous pap smears or between bomblike and non-bomblike x-ray scans. But they cannot explain themselves. And like real brains they forget some of what they have learned each time they learn a new pattern and you never know which patterns they have forgotten.

My own work does not use a neural architecture for feedforward supervised learning. I use a **feedforward fuzzy system** instead. It contains a set of fuzzy if-then rules instead of a web of neurons and synapses. I have derived like supervised learning laws to tune the rules for the fuzzy sets that make up the rules. The fuzzy system learns a mapping from input to output as does the neural net. Both are in theory universal approximators. But the rules act as units of compression and are modular. You can "open" the fuzzy black box and take out rules or put them in and test to see which ones are most or least important. But this can come at a high price of too many rules. A neural net might get by with far fewer neurons for the same problem. You don't know in advance.

The next most popular neural nets are **unsupervised feedforward nets**. These nets strike a nice balance between

mathematical tractability and biological accuracy. Much of early vision uses feedforward layers of competing neurons. Each neuron competes with its neighbors for the incoming activation. The more it gets the more it increases its own firing strength and decreases that of its neighbors. The net acts much like a board full of lightbulbs that light up to varying degrees and then one light bulb wins and turns on 100% and the other turn off 100%. Competitive learning means that only the synapses of the winning neuron change to encode the incoming pattern. The new pattern does not affect the other neurons or their synapses and so they still remember what they have learned in the past. The winner forgets its old pattern in favor of the new one. This is just "blind clustering" at the math level as are most unsupervised learning schemes.

Here is one way to think about the difference between unsupervised and supervised learning. Consider the Frankenstein monster coming to life at an international airport such as LAX. The monster hears many voices in many languages. The first thing it does is associate like with like. It clusters English with English and Spanish with Spanish and so on. The more speech it hears the better it gets at forming the pattern clusters. The striking thing is that no one tells the monster how to do it. The monster just matches features in some way.

Now suppose Dr. Frankenstein shows up. He can tell the monster whether a given speech sample is English or not or Spanish or not and so forth. He supervised the learning and compares the real pattern class to what the monster says or thinks it is to form the error signal. Frankenstein punishes misclassifications and rewards (or does not punish) correct classifications. Such supervised learning is powerful but good teachers are hard to come by.

There is little evidence for supervised learning in flesh. That is why most neural algorithms have only metaphorical connection to real mammalian brains. Supervised learning would require the body to somehow compute an error signal each second and feed that back to perhaps trillions of learning units. It may happen on a small scale in matters of coordination but certainly not at the large-scale brain level. That would be like you not making a move until you saw what effect you had on the gross domestic product or the interest rate of your country.

I have worked with unsupervised nets

both with neural systems and with fuzzy systems. In 1985 I introduced something called differential Hebbian learning in contrast to correlation or Hebbian learning. The old idea of Donald Hebb from 1949 (or Friedrich Hayek from a year or two before in his book *The Sensory Order*) is that the synapses between neuron A and neuron B grows or decays according to the joint activity at A and B. Most mathematicians have understood this to mean that you multiply A and B to get the learning product AB.

I had a problem with that. I started out modeling the link from A to B not as a synapse but as a causal link in a semantic network or cognitive map. This puts the learning question in the context of causal induction in philosophy. Then Donald

Fuzzy rules or concepts act like chunks of animation in the virtual world. The chunk size controls the VR's conceptual granularity. Again neural systems can help figure out some of these fuzzy chunks and tune them to suit each user.

Hebb looks like David Hume who said that causality was just a "constant conjunction of events" or AB. I thought John Stuart Mill had gotten it right in his *Logic* in the 1840s with his notion of "concomitant variation." You don't infer a causal link between my arm and the light being on just because both are present. Rather you tend to infer cause and effect if the light goes off and on as my hand moves up and down. So I replaced the Hebb product AB with the product of time changes.

Only much later did I explore the effect of this at the neural level. Meanwhile other researchers did and have put forth some interesting evidence for it. Gluck and Parker at Stanford showed that if you modeled neural signals as pulses then the differential Hebbian model gives back a simple form. So the synapse does not have to compute a change. The presence or absence of the arriving neural pulse shows whether that change is positive or negative. In engineering we call this a form of delta-modulation. Graeme Mitchison of Cambridge has also argued for some form of this "differential synapse" at the biological level.

In 1988 I realized I could apply the new differential idea to competitive learning as well. That led to the unsupervised scheme that I call differential competitive learning or DCL. The idea is that you

don't just want the winning neuron to learn in a neural slug out. But you want to weight it by its rate of learning. So the synapses change or learn only if the competing neuron changes its win-loss status. This gives a type of pseudo-error signal to the learning process. The plus or minus value of the win rate acts like Dr. Frankenstein saying "Yes" or "No." We benchmarked the DCL scheme against both unsupervised and supervised competitive learning and found that it always held its own and often did better than its supervised rivals—even though it used less information. I know of no evidence for DCL at the biological level. Such is the tradeoff between engineering utility and biological accuracy.

In my work I have used unsupervised

competitive learning schemes (including DCL) to grow the first set of fuzzy rules in a fuzzy system. The idea is that each competing neuron gives rise to a competing list or vector of synapses. In the big state space in the sky this vector forms a point. The system learns when the synapses change and thus when the point moves. But it tends to move in fits and starts and jumps and leaps. So there is an error ball or "covariance ellipsoid" about each synaptic vector. These balls define the big fuzzy subsets or patches of the input-output state space that in turn define a fuzzy if-then rule of the form: "If the input is the fuzzy set A then the output is the fuzzy set B." Noisy or sparse training data leads to a big error ball and thus a large and uncertain fuzzy rule. More accurate data tends to give a smaller error ball and fuzzy rule patch.

My students and I have long since found that we can combine unsupervised and supervised learning to improve a fuzzy or neural system. First the unsupervised learning looks at the stream of sample data and clusters it to form the first set of rules. This is a hard task and often there is no expert to do it. We had this experience when we searched for rules to control the throttle and braking subsystems of

continued on p.40

Complacency and Conservation

Julian L. Simon

Prof. of Business Administration, Univ. of Maryland at College Park

Editor's note: This essay, some of which is excerpted from Dr. Simon's forthcoming book, is a response to posting on the Extropians e-mail list.

Concerning complacency: Of course progress does not come about automatically. And my message certainly is not one of complacency. In this I agree with the doomsayers—that our world needs the best efforts of all humanity to improve our lot. I part company with them in that they expect us to come to a bad end despite the efforts we make, whereas I expect a continuation of humanity's history of successful efforts. And I believe that their message is self-fulfilling, because if you expect your efforts to fail because of inexorable natural limits, then you are likely to feel resigned, and therefore to literally resign. But if you recognize the possibility—in fact the probability—of success, you can tap large reservoirs of energy and enthusiasm.

Adding more people causes problems, but people are also the means to solve these problems. The main fuel to speed the world's progress is our stock of knowledge, and the brakes are our lack of imagination, and of unsound social regulations of these activities. The ultimate resource is people—especially skilled, spirited, and hopeful young people endowed with liberty—who will exert their wills and imaginations for their own benefit, and so inevitably they will benefit not only themselves but the rest of us as well. Thank you.

Concerning the preservation of diversity, including human diversity, which Paul Goetz wrote about, plus some more general discussion of the issue, from *Ultimate Resource II* (out next spring, if lucky):

SHOULD WE CONSERVE RESOURCES FOR OTHERS' SAKES? WHAT KINDS OF RESOURCES NEED CONSERVATION?

I finally got to ask Marla Maples a question. It was at a frenetic press conference where the 26-year old actress, having pocketed a cool \$600,000 for endorsing No Excuses jeans, was pirouetting for the horde of photographers in a skintight pair. Was she, I inquired, simply exploiting her notoriety as the Alleged Other Woman? Au contraire, she said, this was part of her new campaign to save the environment. When pressed for specifics, Maples said breathlessly, "I love the ocean". (Howard Kurtz in *The Washington Post Magazine*, Aug 19, 1990, p. 37).

Should we try to conserve our resources? It depends. Should we try to avoid all waste? Certainly not. Are the Sierra Club, Friends of the Earth, and other conservationist groups barking up the wrong tree? Yes and no.

This is a topic so apparently "simple" and "common-sensical" that adults delight in instructing children in it: "Environmental singer Billy B. sings about recycling on the stage at Wolf Trap." Children are told to "Cover both sides of every sheet of paper you use...." (Yes, Einstein did that, but it is disastrous advice for any office-worker at today's paper prices.) They are instructed to "encourage your family to take part in your community's recycling program," implying that the author of the article is prepared to have children induce guilt in parents she does not know, with needs she cannot discern, for the sake of her own values. That is the sort of social relationships that recycling programs engender. More about this in the next chapter.

The kids get the message—too well. A Wooster, Ohio seventh-grader writes to the newspaper: "On Earth Day we think people should restrain from using aerosol cans [which presumably pollute the atmosphere] and disposable diapers, and they should recycle everything they can."

We can clarify conservation issues by distinguishing among the following: (1) Unique resources, which are one of a kind or close to it, and which we value for aesthetic purposes; examples include the Mona Lisa, an Arthur Rubenstein concert or a Michael Jordan basketball game, and some species of animals. (2) One-of-a-kind resources that we value as historical artifacts; examples include the original U.S. Declaration of Independence and the Dead Sea Scrolls, Abraham Lincoln's first log cabin (if it exists), and perhaps the Mona Lisa. (3) Resources that can be reproduced or recycled or substituted for, and that we value for their material uses; examples include wood pulp, trees, copper, oil, and food. Categories 1 and 2 are truly "non-renewable" resources, but contrary to common belief, category 3 resources (including oil) are all renewable.

This essay deals mainly with resources in category 3, those we value primarily for their uses. These are the resources whose quantities we can positively influence. That is, these are the resources for which we can calculate whether it is cheaper to conserve them for future use, or use them now and obtain the services that they provide us in some other way in the future. The benefits we get from the resources in the other categories—the Mona Lisa or Lincoln's log cabin—cannot be adequately replaced, and hence the economist cannot determine whether conservation is economically worthwhile. The value of a Mona Lisa or a disappearing breed of snail must be what we as a society collec-

tively decide is the appropriate value, a decision upon which market prices may or may not shed some light.

Conservation of resources and pollution often are opposite sides of the same coin. For example, waste newspapers are a pollution, but recycling them reduces the number of trees that are planted and grown.

The costs and scarcities of resources in category 3—mainly energy and extractive materials—are likely to decline continuously in the future, according to the analyses in chapters 1-3. But this chapter asks a different question: whether as individuals and as a society we should try to use less of these materials than we are willing to pay for. That is, should we make social efforts to refrain from using these natural resources, and hence treat them differently from the consumption of pencils, haircuts, and Hula-Hoops for reasons *other than* their costs? The broad answer is that, apart from considerations of national security and international bargaining power, there is no economic rationale for special efforts to avoid using the resources.

Conservationists perform a valuable service when they alert us to dangers threatening humanity's unique treasures, and when they remind us of the values of these treasures to ourselves and to coming generations. But when they move from this role, and suggest that government should intervene to conserve pulp trees or deer beyond what individuals are willing to pay to set aside the trees or the deer's habitat, they are either expressing their own personal aesthetic tastes and religious values, or else they are talking misguided nonsense. (When the Conservation Trust in Great Britain puts "Re-Use Paper, Save Trees" on an envelope, it is simply talking trash; the paper comes from trees that are planted in order to make envelopes.) And when some famous conservationist tells us that there should be fewer people so that it is easier for him or her to find a deserted stretch of beach or mountain range or forest, she or he is simply saying "gimme"—that is, "I enjoy it, and I don't want to share it." (In chapter 29, we shall see how population growth paradoxically leads to more wilderness, however, rather than less.)

Thinking straight about conservation issues is particularly difficult because we must do what we human beings desperately resist doing: Face up to the fact that we cannot have it both ways. We cannot both eat the pie and continue to look at it with pleasure. Grappling with such tradeoffs is the essence of

microeconomic theory. For example, it is obvious that having the wilderness be pristine and not hearing other human voices when one is there means denying the same experience to others. Many who in principle would like others to be able to have that experience, as well as themselves, do not face up to that inevitability.¹

Anthropologists lament the arrival of civilization to the Yononami Indians of Brazil. But the anthropologists also seek the health and cultural benefits of civilization on behalf of that group. Whichever way it goes they will feel regrets, and both cannot be the case. Or, Jews in Israel yearn for the ingathering of Jews from the Diaspora into Israel, where most immigrants live a healthier life than before. But when the Jews of Yemen leave their ancient home, Israeli Jews lament the passing of the 2500 year old Yemenite community, despite the present miserable state of that community.

It is natural to want things both ways. When an economist uses quintessential economic thinking to point out that we must accept the necessity for a tradeoff and that we cannot usually have our cake and eat it too, the argument is met with denial of any such necessity—say, denial that reserving a forest for spotted owls means fewer jobs and lost income—or with charges that harvesting wood is an "obscenity". This makes it very difficult to think straight about conservation issues.

One possible reason why some people refuse to accept the economist's stricture that trade-offs are necessary is that the economist's motivations somehow are not considered noble. And indeed, the economist does try to focus on matters other than motivations. As Murray Weidenbaum wisely notes, economists "care more about results than intentions." If we can succeed in focusing others' attention on results rather than intentions, too, we will achieve results that people will like better than they will otherwise obtain.

It is useful perspective to go back and re-read the classics of the conservation movement in the United States in the early years of the 20th century—for example, the great 1910 book by Charles van Hise. There one finds all the themes being sounded today, and expressed very well. There is one great difference between that literature and the present writings, however. In van Hise's day people believed as follows: "What is the purpose of conservation? It is for man." As another chapter discusses, humankind's welfare is no longer the only—or even the main—goal for many conservationists. ♻

Extropians may have previously encountered the Idea Futures (IF) concept either in the original paper by Robin Hanson [*Extropy* #8, 3:2, Winter 1991-2] or our paper and presentation for EXTRO². Idea Futures allows participants to make bets upon the likelihood of future events, called **claims**. For example, the purchase of a share in a claim at a price of 38 indicates the buyer believes there is a 38% or more probability of it coming true. The Idea Futures marketplace is basically a clearinghouse where individuals place bets upon various issues or even create markets in issues that interest them. The claim prices fluctuate depending upon the most current information related to the event.

The WWW Idea Futures market opened in September of 1994 with just ten players; over the past year the market has grown to over 1500 participants and gained international recognition by winning The Golden Nica Award of the *Prix Ars Electronica* and more recently scoring in the top 5% in the Point Survey of Web sites. It has also received media coverage in Japan and Europe, and in *Wired* and the *Los Angeles Daily News*. It can be found on the Web at <http://www.ideafutures.com/> Make sure to add it to your bookmarks file or hotlist.

How might an Idea Futures market change things?

Originally IF was envisioned as an alternative method to generate funding for scientific research. A real life example where IF could have been useful involves some controversy regarding the cause of stomach ulcers. Dr. Barry Marshall discovered a bacterium that he believed to be associated with ulcers in 1982. He had believed that a treatment with a standard antimicrobial regimen would be developed within a few years. However, this has not occurred. The association of the bacteria with ulcers has faced an uphill battle and the researcher even dosed himself with the bacterium to induce an ulcer and make his point. He was driven to this course of action by the refusal of the medical community to accept his research. As late as 1993 antacids used to treat ulcers without curing them were among the top selling drugs in the United States. Therefore there was a huge industry focused around treating the symptoms rather than curing the disease. Now after over thirteen years an antimicrobial treatment is finally being developed.

If an Idea Futures market had been available at that time the following might have occurred. The researcher could have initiated, upon his discovery, an IF claim as follows:

"By January 1st, 1987 it will have been demonstrated that ulcers are caused by a bacterium and can be treated with antibiotics. "

He could then have seeded the market with \$2000 from his savings and put in bids that he would take bets at a 65% probability of this claim being judged true. The creation of the claim and the odds being offered are published in some journals that follow the Idea Futures market and in some medical journals

that track relevant claims. Some pharmaceutical companies with anti-ulcer medications coming through the approval pipeline and some biased skepticism drive the odds down to 0 by overwhelming the initial investment of the researcher. However, at this point the researcher has convinced a venture capitalist to back him and \$50,000 is pumped into the market and the researcher is given some capital to continue his work. With this infusion of money the market odds rise to 45%. The companies decide to allocate enough money to push the odds down to 20% and they also allocate some funding for research to debunk the claims of the lone researcher.

The market's capitalization creeps over the \$200,000 threshold and starts to draw the attention of professional speculators.

Idea Futures on the Web

Duane Hewitt

The initial findings of the pharmaceutical companies contradict those of the researcher and the release of these results to the press results in the odds being driven down to around 10%. The lone researcher's work has progressed and his results have been confirmed. He takes out a second mortgage on his house and puts \$40,000 into the market. His actions do not go unnoticed as one professional speculator is having him tailed. Upon discovering the researcher's investment the speculator deploys his resources and those of some of his clients in the market and then sells the information to other speculators. The market price climbs inexorably from 15% to around 80% with sudden surges in the price as people cancel booked orders until they can find out whether the information that the market is moving is sound. A press release of his latest results as well as that of some

independent researchers causes the market to move to 95% and he sells out all the shares that he has accumulated in the claim for over \$400,000. He then uses this money to buy out his VC partner and to gear up his company for the delivery of this treatment. This capitalist fairy tale involves those who are correct being rewarded and those who are incorrect being penalized.

What is going on in the WWW Idea Futures market?

As in all markets there are buyers, sellers, commodities, prices and a currency. The buyers and sellers in the WWW Idea Futures market are the registered participants. (Free registration) Currently we are trading a proprietary currency called credibills that has no actual value outside the game. With this currency players may buy or sell the commodities in this market which are "coupons" in claims. Coupons are assets that have a value (price) that depends on the market consensus of the probability of the claim being judged. The value is quoted on the Transaction Form as credicents (hundredths of a credibill). This value varies between 0 and 100 and can be read simply as the estimated probability of the claim coming true. Players may purchase a coupon that may either be a YES (the stated claim will come true) or a NO (the stated claim will not come true).

On the Table there are **Bid** and **Ask** prices for several claims from the IF market. Let us imagine that I am a firm believer that Cryonics is going to be a major growth industry in the next five years. I could then buy YES coupons in *Cryo* in anticipation of this growth meeting the criteria for the claim being judged true. One note of caution: the claim wording often is very important when estimating the probability of the claim being judged true or false. Therefore it is strongly recommended that the **Long Description** (separate Web page) be read carefully. I would have to pay the current **Ask** price if I wanted to make the transaction immediately. Therefore I would have to pay 0.14 credibills per *Cryo* YES coupon that I wished to purchase. If I wished to buy 100 YES coupons it would then cost me \$14 (in credibills).

On the other hand, if I was skeptical about Cryonics catching on in the next

five years I could sell YES coupons short in the IF market. I would do so by selling YES coupons at the **Bid** price and since I don't have any YES coupons to begin with what I will end up with is a negative number of YES coupons held which is the equivalent of holding NO coupons. (At this point I would strongly recommend the IF Web documentation and the IF tutorial by Ken Fishkin to further elucidate these concepts.) Another important thing to realize in dealing with NO coupons (negative YES coupons) is that the following formula applies:

$$(\text{Ask Price of NO Coupon}/100) + (\text{Bid Price of YES coupon}/100) = \$1(\text{credibill})$$

From this you can derive:

$$(\text{Ask Price of NO coupon}/100) = \$1 - (\text{Bid Price of YES coupon}/100)$$

In order to buy 100 *Cryo* NO coupons I would have to spend $(\$1 - 11/100) = \0.89 per coupon which ends up being 89 credibills. This brings to light one feature of the market—*leverage*. This allows people with unpopular or minority opinions a disproportionate effect on the market. In the *Cryo* claim we see that it takes \$14 to purchase 100 YES coupons but requires \$89 to purchase 100 NO coupons.

Idea Futures is a zero sum game in that the server processes the transactions for two parties and does not take any of the risk inherent to the wager. The winnings of one individual correspond to the losses of (an)other individual(s). Therefore the Idea Futures market can serve as a testing place for futurists, pundits, psychics and consultants in order to quantify their ability to make accurate predictions. This is currently reflected as the **Score** in the WWW market which is public.

Where do the prices come from?

In a market there are often people who are unwilling to pay more than a certain price for a commodity. Therefore there are often periods in which no activity occurs because nobody is willing to trade at the current price. In IF we have created a mechanism by which players can "book" orders at a specified price. What happens is that the individual places an order for a specific amount of

coupons at a specific price. This order is not executed until someone else is willing to make the complementary trade. Therefore every trade in IF involves one person who is logged on and completing someone's booked order. The book is public and therefore you can see who is willing to buy coupons and at what price. This is a way to gauge the depth of the market in the claim.

What do the prices mean?

On the table the price for *Cash* is a **Bid** of **88** and an **Ask** of **89**. Therefore there is a market consensus that there about an 88.5% chance that an electronic currency will be used for transactions before a specific date in 1997. The fact that the spread is only 1% indicates that there is a fairly strong consensus that the probability belongs in that range. Several other claims in the table share the characteristic of a very small spread. (*Cr56, IDEA, Immo, Moon, Spce, SSTO, Stew, surg*) This is very surprising for some of these claims because of the long time frames involved and the unpredictable nature of the advances required to make them possible. (eg . *Immo, Moon, Stew*) Some of these claims have larger spreads which indicates that the market estimates that they are more uncertain. (*Canc, MLAW*)

Some Observations from the market

Some interesting information can be gleaned from the market dynamics such as the estimated probability of the event (price), the uncertainty about the event's probability (the spread) and the amount of interest in the issue (capitalization). The price is fairly straightforward in that it represents the estimated probability of the event occurring (as a percentage) and serves as the value of the coupons on the market. The *spread*, which is the difference between the **Bid** and **Ask** prices indicates how narrowly people are willing to pin down the probability of the claim coming true. The spread may be large because of uncertainty as to the exact probability of the claim coming true. This may be due to unclear claim wording or to uncertainties about the actual event. The *capitalization* indicates the credibill value of the outstanding coupons in the claim. It is a measure of the interest in this claim. ☼

Some Idea Futures Claims of Interest to Extropians

| Symbol | Bid | Ask | Short Description | Change | Capitalization |
|--------------|-----|-----|-----------------------------------|--------|----------------|
| <i>\$vIF</i> | 84 | 86 | Real \$ Ver. of IF (<1.1.2000) | +4 | \$1974 |
| <i>Cryo</i> | 8 | 11 | Cryonics catches on by 2000 | 0 | \$428 |
| <i>Canc</i> | 34 | 36 | Cancer Cured by 2010 | +4 | \$2371 |
| <i>Cr56</i> | 82 | 85 | 56 bit cipher cracked <1998 | +12 | \$209 |
| <i>Cybo</i> | 36 | 38 | Cyborgs by 2035/12/31 | +1 | \$1533 |
| <i>Cash</i> | 88 | 89 | E-cash implemented before '97 | +4 | \$4589 |
| <i>IDEA</i> | 16 | 17 | Idea Cipher cracked by 2000 | N/A | \$245 |
| <i>Immo</i> | 20 | 24 | Immortality by 2050 | +3 | \$1420 |
| <i>MLAW</i> | 12 | 18 | Minimalist Legal System Wins | 0 | \$395 |
| <i>Moon</i> | 58 | 59 | Moonbase by 2025 | +2 | \$2003 |
| <i>Spce</i> | 58 | 59 | Private Space Exploitation | +1 | \$1194 |
| <i>SSTO</i> | 62 | 65 | Single Stage to Orbit | +1 | \$1490 |
| <i>Stew</i> | 20 | 21 | Nanotech Stewart Platform | +1 | \$1149 |
| <i>surg</i> | 68 | 71 | Unnatural Cosmetic Surgery | -7 | \$983 |

Prices are as of October 1st, 1995

Symbol:

This is the four letter trading symbol that represents the claim. It is similar to the use of trading symbols in normal stock markets.

Bid:

This is the price (in credicents) that someone is willing to pay for YES coupons in the claim specified.

Ask:

This is the price (in credicents) that someone is willing to sell YES coupons in the claim specified.

Change:

This is the change in the last price since last reported . Since this is the first report I am going to report the change in price since August 8th, 1995.

Capitalization:

This represents the number of outstanding coupon pairs on a claim and is a measure of the interest and amount of disagreement there is on the claim.

Robin Hanson, originator of the Idea Futures idea has several published papers on the topic, including those in *Extropy* #8, *Wired* (Sept. 1995), and *The Journal of Social Epistemology*. You can find a summary of idea futures and pointers to Web versions of these papers at: <http://www.hss.caltech.edu/~hanson/ideafutures.html>

Calling All Individualists!

Liberty is the provocative magazine that celebrates libertarian thinking and individualist living. Neither left-wing nor right-wing, *Liberty* never compromises in its defense of individual freedom and self-reliance:

✓ *Liberty* is not afraid to distinguish its views from conservatism's traditionalism and hostility to civil liberties, as well as from modern liberalism's contempt for the spontaneity of the free market. *Liberty* publishes thought-provoking reviews and essays on the nature of freedom's ideological opponents, and the on cultural and psychological sources of anti-libertarian ideas.

✓ *Liberty* is almost alone in combining concern for the environment with strong opposition to Big Government and intrusive regulation. *Liberty* debunks environmentalist phantasies and anti-environmentalist scare stories, and attacks porkbarrelling by Greens and industrialists alike.

✓ *Liberty* is never dry or dull; its writers love fun as much as they love freedom.

✓ *Liberty's* writers refuse to conform to one ideological mold, and debate each other almost as fervently as they do the statist.

In the upcoming November issue,

Free with your subscription! Liberty via Encryption

New encryption technologies are going to revolutionize the world by making absolute privacy possible for the very first time. In "Liberty via Encryption," innovative economist David Friedman (author of *The Machinery of Freedom*) explores the encrypted world of the twenty-first century: a world of surveillance-proof electronic money, tremendously decentralized power — and a new threat of cyberspace crime.

Friedman explains why the government will be unable to stop the new technologies, describes how encryption will benefit the cause of freedom, and suggests some ways to deal with the problems the new world will bring. You won't want to miss this riveting tape. And it's yours — free — with your subscription to *Liberty*.

Liberty

November 1995 Vol. 9, No. 2 \$4.00

Hiroshima: Paeans to Peace Are Not Enough

Green Light To Kill: The FBI at Ruby Ridge

by R.W. Bradford

How to Launder Money

by J. Orlin Grabbe

Inside Bosnia Today

by Bryan Alexander

The Return of H.L. Mencken

by Max Schulz

Dope Dealing in Chicago

by Matt Asher

"Where *Liberty* lives, there is my homeland." —Algernon Sidney

Liberty remains true to form:

- "The Randy Weaver Show" — R.W. Bradford tunes in to the prime-time police state;

- "A Warm Day in Peace Park" — Michael Oakes revisits Hiroshima on the semicentennial of the atomic bomb, and finds too little rebellion and too much peace;

- "The Money Laundromat" — J. Orlin Grabbe investigates the seamy world of money-laundering, surveillance, and scandal;

- "Undercover Economist" — Matt Asher descends into the world of the drug trade, and examines the economics of the Chicago

marijuana market;

- "High Noon for the Feds?" — In the wake of the Lopez decision, David Owsiany asks whether federal authority has reached its peak;

- "A Week in Bosnia and Points West" — Bryan Alexander tells tales of the Bosnia the mass media misses: of armies of drunks, competing police forces, Satanist dissidents, and governments no one is certain exist;

- "Triumph of the Image" — Richard Kostelanetz reviews the career of the late critic, Edmund Wilson.

Plus the outrageous cartoons of James Gill, Baloo, and John Bergstrom; the wit of "Reflections"; the wonderful weirdness of "Terra Incognita"; fiction; book and movie reviews; and more!

Money-Back, Double Guarantee!

We're so confident that you'll like *Liberty*, we'll make your subscription risk-free. Subscribers are protected by our money-back double guarantee detailed in the box below.

Act Today!

Liberty offers you the best in libertarian thinking and writing. So don't hesitate. Call us at

1-800-854-6991

With our guarantee, you have nothing to lose. And you have the fruits of *Liberty* to gain!

Money-Back Double Guarantee

1. We guarantee a full refund with no questions upon your request after you have received your first issue.
 2. At any time during your subscription, we will guarantee a full pro-rated refund for any unmailed issues.
- In any event, your free bonus and first issue are yours to keep with our compliments.

R.W. Bradford
Publisher

Yes! I'd like to try *Liberty*. Please enter my subscription as follows:

- Twelve issues: only \$35.00, two full years plus "Liberty via Encryption"
- Six issues: \$19.50, one full year plus a free copy of "Liberty via Encryption" foreign orders add \$5.00 per year
- My check (or money order) is enclosed
- Charge my: MasterCard Visa expires: _____

Account # _____


Signature _____

Name _____

Address _____

City, State, Zip _____

Send to: *Liberty*, Dept. EX1, P.O. Box 1181, Port Townsend, WA 98368



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 in 1988 by Max More and Tom Morrow). We created ExI in order to provide a structure and network that would facilitate the

spread and evolution of extropic ideas, values, and culture.

Awareness of extropian ideas and their significance has spread rapidly. The current (November '95) issue of SPIN magazine (a special issue on the future) has a guest editorial by VR pioneer Jaron Lanier. Lanier says:

"I think I now have a sense of what the drama of the next 50 years is going to be like. There will be a struggle between two competing ways of thinking about the world, and this struggle will replace the old Left/Right struggle that defined the twentieth century. The new divide is between what I'll call the Extropians and Stewards... Stewards speak a language of what's already here, like human beings and rocks, while Extropians believe that everything here is going to be replaced by new, evolving things anyway."

Lanier is among those recognizing the obsolescence of old worldviews and the central place Extropian perspectives will have in shaping the ideas and practices of the future. Extropy Institute is the foremost organization developing and disseminating these ideas.

This organizational mission encompasses 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 "posthumans" 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.

ExI pursues these *transhumanist* goals in several ways. Complementing our primary publication, *Extropy*, is our members' newsletter, *Exponent*. *Exponent* carries shorter articles, membership information such as forthcoming meetings, reports on progress of projects, new media coverage, and discussion of organization questions.

We hold a variety of meetings, including Idea Forum discussion meetings and dinner gatherings, lunch meetings, and impromptu celebrations and outings with extropic themes. As membership grows, local events across the country and abroad are taking place. Spring '94 saw an important new development: EXTRO¹ heralded the start of a series of annual conferences where ideas can be explored in depth, and bounced off persons of many different specialties and perspectives. (The main talks from each session can found in the Proceedings volume. SEE BACK COVER)

As befits a transhumanist, high-tech subculture, supplementing printed publications and physical meetings is the on-line Extropian virtual community. The Extropian cyber-community continues to expand, encompassing the main Extropians e-mail list (now in its 4th year), the ExI Essay list, five local e-mail lists for arranging meetings, parties, and

other joint activities, a newsgroup, and an Extropian presence on the World Wide Web. To join the main Extropians e-mail list, send a request to: extropians-request@extropy.org Eric Watt Forste maintains a Web site with plenty of Extropian information: <http://www.c2.org/~arkuat/>




If this issue is your first real contact with extropian ideas, the short version of The Extropian Principles (p.40) will help clarify our shared values and goals. (The full text appeared in *Extropy* #11) The Principles is intended not as a detailed statement or final word on any topic, but as a codification of some of our shared values and attitudes.

Our second conference, EXTRO² brought together 133 extropically-oriented people for a weekend of intellectual stimulation, social interaction, and pure pleasure. Prof. Marvin Minsky gave the keynote address and the program was filled out by 12 other sessions. The conference Proceedings and audio tapes can be ordered. (See back cover.) Be sure to plan ahead for EXTRO³ in July 1996.

RECENT EVENTS

In addition to the conference, recent months have seen discussion meetings on topics such as "How well can the future be pre-

dicted". Extropians have gathered at several social events and have been



Clockwise from top: Fiorella Terenzi & Roy Walford at EXTRO²; Tanya Jones, Ralph Whelan, Rob Michels, & Abe Heward outside conference hotel; Nancie Clark, Mike, Regina Pancake, Keith at Siggraph; Eric Messick, Max More, Ray Sahelian, Marvin Minsky, Steve, at pre-conference reception; Ralph Merkle, Nick Szabo, David May, Max More, Hara Ra, Fred Moulton, Peter Voss, Peter McCluskey, Mark Miller at discussion meeting.

ExI Directors

Max More, President, Editor of *Extropy*. maxmore@primenet.com
Tom Morrow, Vice President. t0morrow@aol.com
Tanya Jones, Treasurer. tanya@alcor.org
Ralph Whelan, Secretary. ralph@alcor.org
Eric Watt Forste. arkuat@pobox.com
Russell E. Whitaker. russw@netcom.com

Council of Advisors

Jamie Dinkelacker, Ph.D.
Prof. Bart Kosko, USC
Prof. Marvin Minsky, MIT
Sharon Presley, Ph.D., Resources for Independent Thinking

EXTROPY INSTITUTE

13428 MAXELLA AVENUE, #273

MARINA DEL REY, CA 90292

310-398-0375 EXI-INFO@EXTROPY.ORG

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

EXTROPIAN PRINCIPLES v.2.5

(Full version in *Extropy* #11)

"This is a philosophy of boundless expansion, of upward- and outwardness, of fantastic superabundance. It's a doctrine of self-transformation, of extremely advanced technology, and of dedicated, immovable optimism. Most of all, it's a philosophy of freedom from limitations of any kind."

Ed Regis, "Meet the Extropians", *Wired*, October 1994.

"Extropians remain die-hard rationalists, resistant to revealed truth of any kind, even if it's the truth of their own predictions... The Extropians' is a libertarianism of rare sophistication... [their] vision could turned out to be our best guide through the strange eons to follow."

Village Voice, December 1994.

found at events such as Siggraph '95 in Los Angeles, the opening of Resources for Independent Thinking (run by ExI Advisor Dr. Sharon Presley), and the Bionomics Conference in October.

We are fostering the growth of more local discussion groups and international chapters of ExI, and we will continue to develop our network of communication and action. We look forward to the continued development of the Extropians cyberculture, especially on the World Wide Web. We will continue to build cooperation with other organizations for shared goals and make contact with more scientists, technologists, philosophers, and artists to strengthen our network.

We hope you will join us as an active participant in the Extropian movement. (See p.2 for membership information.) Help shape the future!

Upward and Outward!

Max More
President



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.

B.E.S.T. D.O. I.T. S.O.

KOSKO from p.32

single-lane platoons of smart cars. Then supervised learning can tune the first set of rules. This works better than using either unsupervised or supervised learning itself.

The fourth and last type of neural model is a **supervised feedback net**. Researchers have played with these models since the late 1980s but no one has put forth an efficient algorithm or really shown anything but very small-scale examples. Supervised learning requires hundreds of thousands of training cycles and those can easily destabilize a feedback network. But these nets do hold out the promise of learning more complex time-varying patterns than standard feedforward nets can learn. You have to somehow tie the time and rhythmic structure of a musical tune to the swirling cycles of neuronal activity in the net. No one has yet worked out that dynamic balance.

Q: What future applications of fuzzy logic and neural networks do you most look forward to?

The most promising applications lie in communications and multimedia. Both these fields work with much faster data rates and heavier data loads than do many problems of control and even of signal processing.

We have already designed fuzzy systems to "randomly" spread and despread wireless signals over wide stretches of the frequency spectrum. Other fuzzy systems can help route calls in local area networks or help pick time or frequency slots to multiplex the calls together.

Multimedia and virtual reality problems offer a still greater challenge for realtime intelligent decision making. We want to design neural-fuzzy agents that learn your preference maps and then search through databases to fetch or browse on your behalf. Preference maps are sheets with hills and valleys. The higher the hill the more you like the object. The lower the valley the more you dislike it. Neural fuzzy agents try to place rules at the highest peaks and lowest valleys and then fill in from there. The sheet changes slightly each time you sample new object or re-sample old ones.

It takes a full-blown math model to capture a "real" virtual reality. In some sense math models are virtual realities. No one knows those equations and no computer would have enough time or power to process them even if you did know them. Feedback fuzzy systems offer one way to deal with the VR problem. Fuzzy rules or concepts act like chunks of animation in the virtual world. The chunk size controls the VR's conceptual granularity. Again neural systems can help figure out some of these fuzzy chunks and tune them to suit each user.

The trouble is we have almost no knowledge of how such swirling feedback fuzzy systems work. So this should remain an open and active research area for some time.

These nets have real power in terms of learning complex boundaries between cancerous and non-cancerous pap smears or between bomblike and non-bomblike x-ray scans. But they cannot explain themselves.



MINDSURFING




Subject: Why the Internet is like a Baby's Head

A quick look at the many threats facing the global network

Date: September 24 1995

From: Yow <yow@primenet.com>

Fontanels are the soft, squishy spots on a baby's head. The Internet (a newborn, after all) sports a number of unnerving soft spots:



1. Bad speech. Network traffic routinely contains obscenity, hate speech, vulgarity, heresy, bomb-making instructions [1] and LSD recipes. Not to mention malt beverage websites, make-money-quick and phone sex Usenet postings, and unsolicited e-mail on how to meet Russian women.

2. Bad acts. The network is sometimes used to facilitate crimes such as child molestation, terrorism, fraud, and software piracy. On a percentage basis, however, true bad acts are extremely rare.

3. Bad legislation. In reaction to a perceived crisis of bad speech and bad acts on the network, our presumably well-meaning lawmakers rush to create legislative "solutions" which if enacted would inhibit or even criminalize on-line speech [2]. Witness the universally reviled Exon Communications Decency Act [3], or the farcical Grassley "cyberporn" hearings based on the widely discredited "Rimm" on-line pornography study [4]. (In each instance, countless goofy Internet petitions follow.)

4. Bad lawsuits. Access providers, content providers and individual users are routinely sued for copyright infringement and defamation.

5. Bad policies. In fear both of legislators and lawsuits, major access providers enact and enforce extremely restrictive and paternalistic acceptable use policies, including "vulgarity guidelines" and so-called "guides", company *capos* who monitor chat rooms for bad speech [5].

6. Bad seizures. In connection with criminal cases and private litigation such as copyright actions, law enforcement raids on homes and businesses, seizing CPUs and floppies. In some cases, the seizures are carried out by the high-tech equivalent of Pinkertons, rather than by government law enforcement agents.

7. Bad privacy. Personal privacy is eroded by government databasing, the direct marketing in-

dustry, employer monitoring and the information black market. Anonymous remailers are compromised. The patron saint of guerrilla cryptographers, Phillip Zimmerman, is hassled relentlessly. Everything you do on-line may be monitored by your service provider, and although federal law prohibits disclosure of the content of your communications, nothing restricts your service provider from disclosing what on-line information you consume (think about that the next time you're lurking in alt.binaries.pictures.erotica.oral) or what products you buy on-line.

All those soft spots! What will become of our poor baby?

As it turns out, fontanels really aren't as delicate as they look. Covered by flexible, tough membrane, they first allow a baby's skull to compress during passage through the birth canal, and in infancy permit rapid expansion of the brain. Baby heads don't squish like melons, even if they look like they might.

Perhaps the Internet is also hardier than it may first appear, for several reasons. First, there is a common interest shared by the on-line pornophiles, racists, anarchists, body piercing fanatics, neo-Nazis, vulgarians, fundamentalist Christians, heretics, pagans, psychopharmacologists, and many others with incredibly diverse and often extreme viewpoints: a love for the Internet, and, by implication, for the ability to express oneself freely and consume information freely that only the Internet makes possible. Humanity's shared overriding concern for the well-being of the network may someday give rise to an Internet political party or even a larger movement that transcends political and national boundaries.

(Say what you want about Newt Gingrich in any other context, his personal decision to oppose the Communications Decency Act in June 1995, after its 84-16 approval in the Senate, probably did more for the future of on-line speech than any allegedly liberal politician did or ever will do. Don't look to California Senators Boxer and Feinstein for leadership, for example... they voted for the CDA. [6])

Second, as a practical matter, the Internet (and its powerful ancillary tools such as PGP

Cont. on p.64

Bio-Enhancement Update

Ray Sahelian, M.D.




Photo by Nancie Clark

The melatonin article from the last issue was quite timely. Over the last few weeks we have seen an unprecedented coverage of melatonin in the news including hundreds of articles in magazines and newspapers, along with local and national radio and television segments. The August 7th

What about negative feedback?

Dr. Rioux continues: "Is it possible that in taking the supplement long-term one may inadvertently send a negative feedback to the pineal gland so that it would produce

MELATONIN part 2

Dreams like you've never dreamed

Melatonin is a tryptamine, chemically similar to DMT and other potent hallucinogens. Many users of melatonin have noticed that their dreams are incredibly vivid. Is melatonin our "dream molecule"?

Newsweek article was mostly responsible for accelerating the melatonin mania. It quotes me as saying "This will eventually make prescription sleeping pills all but obsolete."

Quite a few readers contacted me with questions regarding the use of this supplement. Dr. Pierre Rioux from Minot, North Dakota, asks:

How much melatonin is absorbed?

When melatonin is swallowed in tablet or capsule form, it is easily absorbed from the stomach or intestines and subsequently is metabolized by the liver. Only a portion of the ingested melatonin eventually makes it to the bloodstream. In one study volunteers were given 80 mg of melatonin and then had their blood levels checked. There was a 300 fold variation in the amount of melatonin present in their blood. This shows the uniqueness of each person's absorption and metabolism. (Waldhauser F, Sleep laboratory investigations on hypnotic properties of melatonin. *Psychopharmacology* 100:222-6, 1990.)


The sublingual form of melatonin is more consistently effective since the absorbed melatonin from the oral mucosa goes directly into the bloodstream without having to initially be metabolized by the liver. A lower dose from a sublingual is often as effective as a higher dose from a pill.

even less melatonin, resulting in a person then becoming completely dependent on exogenous melatonin?"

A reliable way to find the answer to this question is by giving humans melatonin supplements for a long period, such as a year or two and then remove the pineal gland to analyze it under the microscope and see how healthy it is. This study, for obvious reasons, is not likely to be done. Another way of finding out is again supplementing for a long period, stopping the melatonin, and testing the amount of melatonin present in the blood stream the subsequent nights. To my knowledge, this study has also not been done.

However, we do have plenty of anecdotal evidence that does not indicate any clinical problems with discontinuation of chronic melatonin use. A number of my patients and people I have surveyed who stopped regular use of melatonin after a period of months subsequently kept sleeping well without any signs of withdrawal. The abrupt discontinuation of most pharmaceutical sleeping pills does lead to withdrawal.

Rats injected with melatonin every day showed reduction in pineal calcium concentrations and appeared to have healthier glands. Is it even possible that taking exogenous melatonin could keep our pineal glands in tip top shape? (Puigdevall V, Hypothalamic and pineal concentrations of magnesium and calcium after a sustained



A CHILD'S MIND

KNOWS

NO BOUNDARIES

AND

SETS

.....EXPAND YOUR MIND'S LIMITS

NO LIMITS

COGNITIVE ENHANCEMENT

- ACETYL-L-CARNITINE
60 caps 500 mg. \$49.95
- MEMORY BOOST
180 caps \$17.95
- GINKGO BILOBA
90 caps 60 mg. \$17.95
- EXTENSION I.Q.
90 caps \$15.95
- MELATONIN
60 caps 3 mg. \$9.00

VITAMIN RESEARCH PRODUCTS, INC.

1-800-877-2447

<http://www.connectus.com/~vrp/>

25% DISCOUNT ON ALL FIRST ORDERS, CALL FOR A FREE CATALOG.

VITAMIN RESEARCH PRODUCTS, INC. • 3579 HWY 50 E, CARSON CITY, NV 89701 • (702) 884-1300

administration of melatonin. *J Endocrinol Invest* 16 (Suppl 1 to no. 8):253, 1993.)

wooden plank that seemed to support a corpse. It passed silently, no one besides the corpse was on board. Strange...

The actual chemical name of melatonin is 5-methoxy-N-acetyl-tryptamine. Tryptamines are compounds found abundantly in certain hallucinogenic tropical plants. In South America, amazonian natives use a form of tryptamine called DMT, (N,N-dimethyl-tryptamine), to induce an intense hallucinogenic experience. Some amazonian shamans use DMT and related hallucinogens for healing purposes and spiritual connections.

Vivid Dreams

Gary Miracle, a chemistry student from the University of Wisconsin, points out that melatonin is a tryptamine, chemically similar to DMT (dimethyl-tryptamine) and other potent hallucinogens. Many users of melatonin have noticed that their dreams are more memorable. Is there a connection?

I was riding a gigantic elephant through a thick jungle in India with vines crisscrossing from tree to tree above my head. We soon reached a clearing where there appeared a pond. The elephant continued on its steady march, its thick feet and legs wading through the shallow waters. There was something unusual about the bottom of the pond. It glittered... the glitter of priceless gems—of sapphires, rubies, and diamonds. A canoe filled with cut flowers slowly passed to my right. It had four posts holding a

Vivid dreams are common with melatonin use. After experiencing more intense dreams myself, and finding out through my surveys that others had experienced the same, I became extremely curious. Throughout ages humankind has been seeking to find out about the nature of dreams and what causes them. Various metaphysical explanations have been proposed. Could the answer simply be melatonin? I strove to learn more about the chemistry of melatonin. This is what I've learned and wish to propose as an explanation for the chemicals responsible for dreams.

When melatonin is metabolized in the pineal gland, it is converted to 5-methoxy-tryptamine. When rats are given melatonin or 5-methoxy-tryptamine, they spend more time in REM sleep, the stage of sleep associated with dreams. It makes

I found that slightly over half of respondents had experienced vivid dreams. The higher the melatonin dose, the more likely the possibility.

LifeLink

445 Lierly Lane
Arroyo Grande, CA 93420

either or

Melatonin: (mel-a-toe-nin) n. 1. The principal hormone produced by the pineal gland. 2. Scientists say melatonin can boost the immune system, prevent cancer, ease insomnia, protect cells from free radical damage, reset the body's aging clock and extend life.

3mg x 120 tablets.....\$13.00

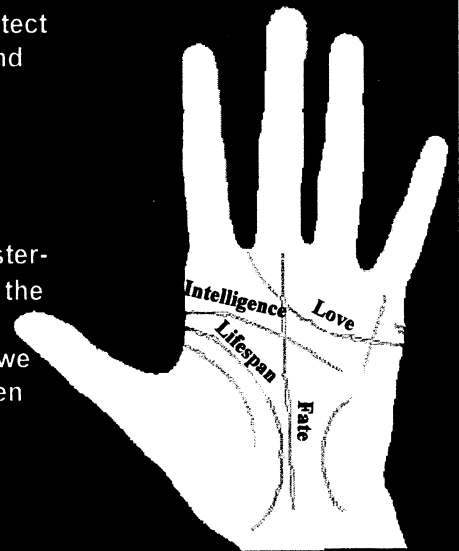
3mg x 240 tablets.....\$25.00

DHEA: dehydroepiandrosterone (dee-hi-dro-epp-ee-an-dro-ster-own) n. 1. A naturally occurring steroid precursor produced by the adrenal gland and found in plasma levels and brain tissue. 2. Production rises dramatically during puberty but wanes as we age. 3. Is used by the body to make the sex hormones estrogen and testosterone. 4. Increases our ability to cope with stress, decreases joint pain and improves mood.

25mg x 90 capsules.....\$22.00

25mg x 180 capsules.....\$40.00

25mg x 360 capsules.....\$75.00



Extropy subscribers receive a 20% discount

Orders must be prepaid (\$5.00 S & H).

Phone: 805.473.1389, Fax: 805.473.2803, e-mail: delano@dax.win.net.

sense that melatonin, 5-methoxytryptamine or related tryptamines are involved with dreams. Are there more chemicals involved? We don't know all the answers yet.

We've always thought that hallucinogens were substances foreign to the brain. It is interesting to note that we can synthesize natural hallucinogens during sleep; they are a normal part of our brain chemistry. Moreover, our brain does not seem to develop a tolerance to these tryptamine hallucinogens since we dream every night, whether we remember the dreams or not.

In order to find out whether anyone else had come up with the proposition that tryptamines were involved with dreams, I placed a question on the internet. Two days later I got a message from J.C. Callaway, Ph.D., from the University of Kuopio, Finland. He referred me to an article he had published in *Medical Hypothesis* titled "A proposed mechanism for the visions of dream sleep." (26:119-124, 1988.) Dr. Callaway suggests tryptamines and related compounds are the chemicals responsible for dreams and his article gives a detailed explanation.

Throughout my surveys I found that slightly over half of respondents had experienced vivid dreams. The higher the melatonin dose, the more likely the possibility. David, a 35 year old, says, "I've been taking melatonin 3 mg for about five months—not every night, but perhaps 4 nights a week. I tend not to dream. Recently I tried 6 mg and noticed that I do have dreams on this dose."

In my surveys, about 8% of people taking melatonin reported eventually having had a very bad dream, or nightmare. Melatonin accentuates dreams. If an individual's dreams are usually pleasant, they will continue being pleasant, but more intensely. If dreams are normally unpleasant, then it is possible they will be more unpleasant with melatonin use. Pam, a 21 year old student writes, "I took 3 mg of melatonin. It worked quite fast and I slept fairly soundly except after an awful nightmare. From then on my sleep was restless and I kept waking up. This was the second time I had taken melatonin. The first time I had one particularly vivid dream, and it was okay."

In order to reduce the likelihood of unpleasant dreams, develop a positive attitude and good self-esteem. Steer away

from negative brain input such as violent TV shows or horror movies. Nurture a loving connection with people, nature, and animals. For a detailed approach to mood improvement and enhancing quality of life, thus more pleasant dreams, please refer to my book *Be Happier Starting Now*.

If you need to take melatonin for a sleep disorder but do not wish to have vivid dreams or nightmares, take as little a dose as necessary. Low doses, such as 0.5 mg, are less likely to intensify dreams.

Whether you believe that dreams have meaning or are due to random firing of neurons, or both (sleep researchers are still not in total agreement on this issue), it can be fun to recall these vivid dreams and write them down in a diary soon upon awakening. Even if you can't find meaning in them, they can always be interesting conversational pieces over the breakfast table.

Don't Blame the Burrito.

Every evening, after a long day in the office, I take a few mile walk by the ocean and streets of Marina Del Rey and Venice to release unused physical energy. On one of these recent walks, while passing by a Mexican restaurant, I got an incredible craving for a vegetarian burrito. After quickly downing one, I ordered a second, and since my mouth was dry, I went against my custom and ordered a small cola. I hadn't drank a coke for a long time and I truly enjoyed the slow sips. Against my better judgment I ordered a refill.

That night I went to bed at my usual hour, midnight. Twenty minutes later, I was still awake. I got out of bed and turned on the TV. Not much was on except some Psychic Network enticements promising an accurate prediction of my upcoming love life and an infomercial that guaranteed an income of \$8,000 a month from simply placing small classified ads in the backs of newspapers. Thinking to myself, "There's got to be a better way I can nurture my neurons," I picked up a recent article from *Neuropsychopharmacology*.

"Caffeine reduces Low-Frequency Delta Activity in the Human Sleep EEG," was the title of the article by Landolt and colleagues from the Institute of Pharmacology, University of Zurich, Switzerland. How appropriate! Actually, "Caffeine Stops Entry Into Stage I Sleep," might have been a more accurate title for

my case.

There are four major stages of sleep defined as stages I, II, III, and IV. Each stage is progressively deeper, with IV being the deepest, most restorative stage. In the above study, eight young males paid to participate in a study—probably poor college students who chose this time to spend a sleepless night in the laboratory to pay for needed pencils and paper rather than donating their last remaining pint of blood to the blood bank—were given a capsule of 100 mg of caffeine at bedtime. A regular cup of caffeine on average contains anywhere between 80 to 140 mg of caffeine. They were then thoroughly monitored including a detailed EEG (electroencephalogram).

When levels of caffeine were checked in their saliva an hour later, a concentration of 7.5 micromol/L was found which leads to a concentration in the cerebrospinal fluid of approximately 6 micromol/L. At this concentration, caffeine blocks A1 and A2 adenosine receptors.

Caffeine has long been known to be a stimulant that inhibits the onset of sleep. This study also confirmed that Delta sleep, an EEG recorded pattern that occurs during stages 3 and 4, was suppressed. The researchers propose that adenosine receptors may be involved in the regulation of sleep.


Seven hours after the administration of caffeine, the amount present in saliva had dropped to 3.5 micromol/L. But what was really interesting was that the following night, the volunteers, even though they did not receive any additional caffeine pills, took longer to fall asleep. Delta sleep was also reduced. The researchers speculate that there may have been some residual effects from the previous night's caffeine administration.

Reading this article confirmed some of my suspicions. Quite a number of my patients over the years have informed me that they thought drinking coffee, even in the morning or early afternoon, was interfering with their sleep. I have noticed this in myself, too. I'm very sensitive to stimulants. For instance, even two mg of deprenyl taken in the morning can sometimes lead to insomnia. Keeping this in mind, I generally recommend those who take stimulants of any kind (tyrosine, phenylalanine, DMAE, ginseng, ginkgo, diet pills, ephedrine, etc.), and are prone

cont. on p.65

The phenomenal best-seller, as mentioned in *Newsweek* and *U.S. News and World Report*.

- Sleep like a baby
- Improve your mood
- Have more energy
- See vivid dreams
- Prevent jet lag
- And possibly live longer.



Melatonin, a natural hormone made by our pineal gland each night to help us sleep, is now available in health food stores, and where supplements are sold. Many claim it can improve our immune system, treat cancer, be a powerful antioxidant, and reverse the aging process. Is melatonin this miracle medicine?

"A splendid and well-written source of information on all aspects of melatonin use, including the all-important precautions. Anyone thinking of taking melatonin supplements ought to read it. Excellent."

Durk Pearson and Sandy Shaw, authors of *Life Extension: A Practical Scientific Approach*.

"Well referenced, concise, and readable, Melatonin: Nature's Sleeping Pill covers all aspects of this remarkable substance. I highly recommend Dr. Sahelian's book."

Julian Whitaker, MD., Editor of *Health and Healing* Newsletter.

ISBN: 0-9639755-7-9

For your autographed copy of *Melatonin: Nature's Sleeping Pill*, send a \$16.95 check or money order (includes shipping; California residents add \$1.05 tax, for a total of \$18) to the above address.

To order both books, send a \$29.95 check or money order (includes shipping; California residents add \$2.45 tax, for a total of \$32.40). Be sure to include your name and mailing address. Orders are shipped promptly. Residents outside of U.S. add \$6 postage for one book, or \$8 for two books.

Or, have your credit card ready and call 1-310-821-2409 to order either or both books now. Orders are shipped the same day. Melatonin is also available or can be ordered from most bookstores or health food stores.

"A self-actualization book different from any other of its genre. In addressing one of the most universal of human quests—happiness—it promotes a humanist approach. Spirituality is explained as a sense of connection associated with brain chemicals, such as serotonin and endorphins. The mind-body connection is explained in terms of interactions of the neural, hormonal, and immune systems."

—*Skeptical Inquirer*, July/August 1995


"Our lives can be filled to overflowing with joyful, peaceful, vigorous happiness, if only we act with self-control informed by wisdom. Dr. Ray Sahelian's delightful, well-rounded book lyrically delves into that wisdom — wisdom that we often forget in the rush of daily life. This book is gentle and soothing; we sense the author's strong sense of benevolence. Be Happier flows remarkably smoothly, informs broadly, and encourages us to renew our appreciation of the preciousness and glory of our lives."

—Max More, President, Extropy Institute.

ISBN 0-9639755-6-0

For your autographed copy of *Be Happier*, send a \$15.00 check or money order (includes shipping; California residents add 99 cents tax for a total of \$15.99) to:


Be Happier Press
PO Box 12619ex, Marina Del Rey, CA 90295



NATURAL SLEEP

The New Sleep-Enhancing Melatonin Formula

From The Life Extension Foundation



NATURAL SLEEP™ is now being sold on national TV and will soon be available in pharmacies at **\$19.95** for 60 capsules. Members of **The Life Extension Foundation** can obtain this unique sleep formula for **only \$14.96 per bottle**. If a member buys four bottles of **NATURAL SLEEP™** the price is just **\$10.87 a bottle**.

To order your supply of **NATURAL SLEEP™ with Melatonin Plus™** just call:

1-800-544-4440

Or order by mail from: **PROLONGEVITY, LTD.,**
10 Alden Road, Unit 6,
Markham, Ontario L3R2S1, Canada

The *natural* way of enjoying a good night's sleep *without* side effects is to take **melatonin**, the body's own sleep-enhancing agent that induces drowsiness and enables you to enjoy deep, restful sleep *every* night!

Good sleep will do more than make you feel better. Studies show that insomnia is the primary cause of lost productivity in the work place. The bottom line is that you need **NATURAL SLEEP™ with Melatonin Plus™** to help you feel *and* work better!

Young people secrete **melatonin** in abundance after it gets dark to enable them to enter deep, restful sleep easily. Once sleep occurs, **melatonin** enables them to enjoy more of the deep sleep they need to feel fully rested and refreshed in the morning.

As we grow older, however, our natural production of **melatonin** begins to decline, which can lead to insomnia in people who never had sleep problems before. The problem worsens the older we get, as our **melatonin** production continues to decline.

Scientific studies have documented that **melatonin deficiency** is the primary reason people over 40 have difficulty in enjoying a good night's sleep.

Tens of thousands of insomniacs have benefitted by taking **melatonin** at bedtime. However, there are other nutrients that also promote good sleep to *enhance* the action of **melatonin**. **The Life Extension Foundation** has put together a new formula—**NATURAL SLEEP™ with Melatonin Plus™** that combines the finest **melatonin (99.8% pure)** with additional nutrients to help promote safe, natural sleep.


Here are the ingredients of **NATURAL SLEEP™ with Melatonin Plus™**:

- **Melatonin**...the body's own natural sleep agent improves the *quality* of sleep by facilitating easier entry into the super-restful deep sleep required for optimal recuperation from fatigue. Each capsule of **NATURAL SLEEP™** contains **3 mg of pharmaceutical quality melatonin** certified **almost 100% pure!** (Most *commercial* melatonin products are only 98% pure).
- **Pyridoxal 5-phosphate**...converts the dietary amino acid tryptophan into serotonin, which is required to achieve high quality sleep. Pyridoxal 5-Phosphate converts tryptophan into **serotonin**, some of which is then converted into **melatonin**. By drinking a glass of milk or eating a slice of turkey an hour before bedtime, you consume tryptophan which can then be converted into sleep-inducing **serotonin** with the help of **Pyridoxal-5-phosphate**—a non stimulating form of **vitamin B-6**.
- **Chromium polynicotinate and chromium picolinate**...lowers serum glucose. Excess blood glucose (sugar) levels can inhibit the ability to fall and stay asleep.
- **Vitamin B-12**...normalizes the body's own natural circadian rhythms, which enable you to enter deep, restful sleep effortlessly, without stress or tension.
- **Niacinamide ascorbate**
- **Magnesium chloride**
- **Calcium citrate**
- **Inositol**.....induce the state of relaxation needed to enter sleep.

MURDER AT THE LIAR'S CLUB

by Mark.J.P. Wolf

There's been a murder at the Liar's Club in Merm, and as interim sheriff, you are called upon to investigate. Your deputy, Ned, did some preliminary interrogation, but has given up in frustration; members of the Liar's Club *always* lie, and there are at least three members, and at least three non-members (truth-tellers), among the seven people involved (the six suspects and the murdered man). Everyone involved knows who is a member and who isn't, and even who the killer is, but no one wants to get involved, for fear of becoming the next victim. Ned's first question to the six suspects was "Are you a member?", to which each person answered "No." Finally in desperation, Ned took a statement from each suspect, leaving you to sort out the clues. Along with the statements is a written statement from Hart, the murder victim, which he wrote as he was dying. In his note, he accuses Archer of killing him; however, Hart may be a member, and thus a liar. Below is a summary of the six suspects statements, and Hart's note which he wrote as he died. The case is definitely *not* suicide, and it is up to you to determine from the seven statements who is a member and who isn't, and who the killer is.



Archer: "Davis is a member, and he would agree that Flint would agree that Hart and Edgar are both members."

Brown: "Hart would agree that Clark would agree that Archer and Davis are both non-members."

Clark: "Davis and Brown would agree that Edgar would agree that the murderer is not a member."

Davis: "Hart's membership status is the same as the murderer's."

Edgar: "Clark would agree that Brown and Hart would agree that Flint is a member."

Flint: "Brown would agree that Archer and Clark do not have the same membership status."

Hart: "Archer killed me!"

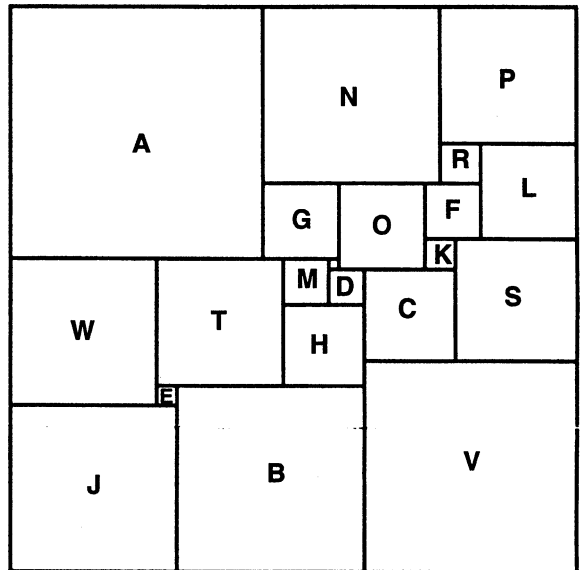
(Answers to appear in the next issue.)

See next page for answer to last issue's puzzle.

Congratulations to Sasha Chislenko for being the first to solve last issue's Enigma.

SQUARED DEAL — SOLUTION

| Republic | Side length (miles) | Land Area (square miles) |
|----------|---------------------|--------------------------|
| Zevo | 2 | 4 |
| Elga | 4 | 16 |
| Koid | 6 | 36 |
| Dorra | 7 | 49 |
| Rudra | 8 | 64 |
| Merm | 9 | 81 |
| Froll | 11 | 121 |
| Gom | 15 | 225 |
| Hort | 16 | 256 |
| Okell | 17 | 289 |
| Curro | 18 | 324 |
| Lenif | 19 | 361 |
| Slome | 24 | 576 |
| Telka | 25 | 625 |
| Phydra | 27 | 729 |
| Wintz | 29 | 841 |
| Jarp | 33 | 1089 |
| Nurin | 35 | 1225 |
| Bolta | 37 | 1369 |
| Voth | 42 | 1764 |
| Alto | 50 | 2500 |
| The USSR | 112 | 12544 |



two squares; the smallest square which is the sum of two squares is 100, so Telka's length must be at least 10; Telka's length could be 10, 13, 15 (which can be eliminated because Gom is 15), 17, 20, or 25 miles wide. Numbers over 25 are too big, since Alto, which is four times the area of Telka, can be no bigger than 50. Thus, Alto must be 20, 26, 34, 40, or 50 miles wide. This also limits the possibilities for the Slome & Dorra combination to {5,12}, {6, 8}, {7, 24}, {8, 15}, and {12, 16}. But, by clue C., 5 is not a possibility for square width, and Gom is 15, so Slome & Dorra are either {6, 8}, {7, 24}, or {12, 16}. Telka is either 10, 20, or 25 miles wide, and Alto is either 20, 40, or 50 miles wide.

By clue G., we know Froll is larger than at least two other republics, so it must be 6 or more miles wide. Since Froll cannot be 4, Voth cannot be 41, ruling out one of the possibilities. Voth must be either 42 or 50. Given the possible values for Dorra (8, 16, and 24), we see that one of remaining possibilities, {Voth=50, Gom=15, Bolta + Dorra + Froll = 25 + 27 + 39}, does not work either, because Dorra can be none of those three values. Therefore Voth must be 42 miles wide, making Bolta= 37 miles wide, Froll= 11 miles wide, and Dorra= 7 miles wide. Taking clue F. into account, the only value possible for Slome which is a square when added to 49, is 24, so Slome= 24 miles wide, Telka= 25 miles wide, and Alto= 50 miles wide.

From the preceding steps, we know the values and names of nine of the republics

Quadra, the capital city, lies in the republic of Merm.

The first step in solving the problem is to determine the sizes of the republics, and clue A. helps in finding the range of sizes. Since there are 21 squares of unequal size, the second largest square would have to be at least 20 miles wide; thus the maximum size of the largest square would be limited to 92 miles wide (otherwise they could not fit within the USSR). From clue A., the largest square must be either 75 or 50 miles wide (25 is too small a size for the largest square, since the total for the 21 unequal squares comes to $1122 = 12544$ square miles). The smallest territory, then is either 2 or 3 miles wide.

When clue B. is taken into account with clue A., the size of the second largest square must be either 28, 35, 42, or 49 miles wide, and the third smallest square must be either 4, 5, 6, or 7 miles wide, as these are the only combinations possible which have values that fit inside the bounds set in the two clues, and the USSR size of 112 miles square. If the largest square is 75, then only the {35, 5} combination will work for clue B., since the room left in the 112 x 112 square will not fit squares bigger than 37 alongside a 75 mile wide

square. If the 35 mile one is used, however, there will be a 2-mile gap, and since 3 is the smallest square, this combination will not work. Thus, the largest square is not 75, but 50, and the smallest square is 2 miles wide.

Using the rest of clue A., and knowing 50 to be the largest square, and 2 to be the smallest, we can determine the six smallest squares to be 2, 4, 6, 7, 8, and 9 (two combinations of six squares add up to 250, but only one satisfies Clue B.). Also, since we know 6 to be the third smallest, by clue B. that the second largest square is 42 miles wide.

Clue E. reveals that Nurin is 1225 square miles, and that Gom is 225 square miles. Since Bolta is larger than Nurin, and Gom is 225, Voth must be more than 1450 square miles, and thus must be either 40, 41, 42, or 50 miles wide (taking clue C. into account). The value 40, however, does not work, so Voth must be either 41, 42, or 50 miles wide, since combinations of squares are possible for all three values. Since there are only four squares larger than 225 by less than 150, we know (by the rest of clue E.) that 16, 17, 18, and 19 are side-lengths of four other republics.

Clue F. reveals Telka as the sum of

Cont. on p.64

Smart Contracts

Copyright 1995 by
Nick Szabo

Introduction

The contract, a set of promises agreed to in a “meeting of the minds”, is the traditional way to formalize a relationship. While contracts are primarily used in business relationships (the focus of this article), they can also involve personal relationships such as marriages. Contracts are also important in politics, not only because of “social contract” theories but also because contract enforcement has traditionally been considered a basic function of capitalist governments.

Whether enforced by a government, or otherwise, the contract is the basic building block of a free market economy. Over many centuries of cultural evolution has emerged both the concept of contract and principles related to it, encoded into common law. Algorithmic information theory [Li & Vitanyi] suggests that such evolved structures are often logically deep — they are prohibitively costly to recompute. If we started from scratch, using reason and experience, it could take many centuries to redevelop sophisticated memes like property rights that make the modern free market work [Hayek].

The success of the common law of contracts, combined with the high cost of replacing it, makes it worthwhile to both preserve and to make use of these principles where appropriate. Yet, the digital revolution is radically changing the kinds of relationships we can have. What parts of our hard-won legal tradition will still be valuable in the cyberspace era? What is the best way to apply these common law principles to the design of our on-line relationships?

Computers make possible the running of algorithms heretofore prohibitively costly, and networks the quicker transmission of larger and more sophisticated messages. Furthermore, computer scientists and cryptographers have recently discovered many new and quite interesting algorithms. Combining these messages and algorithms makes possible a wide variety of new protocols.

New institutions, and new ways to formalize the relationships that make up these institutions, are now made possible by the digital revolution. I call these new

contracts “smart”, because they are far more functional than their inanimate paper-based ancestors. No use of “artificial intelligence” is implied. A smart contract is a set of promises, specified in digital form, including protocols within which the parties perform on the other promises.

From common law, economic theory, and contractual conditions often found in practice, we can distill four basic objectives of contract design. The first of these is *observability*, the ability of the principals to observe each others’ performance of the contract, or to prove their performance to other principals. The field of accounting is, roughly speaking, primarily concerned with making contracts an organization is involved in more observable.

The second objective *verifiability*, the ability of a principal to prove to an arbitrator that a contract has been performed or breached, or the ability of the arbitrator to find this out by other means. The disciplines of auditing and investigation roughly correspond with verification of contract performance. Observability and verifiability can also include the ability to differentiate between intentional violations of the contract and good faith errors.

The third objective of contract design is *privity*, the principle that knowledge and control over the contents and performance of a contract should be distributed among parties only as much as is necessary for the performance of that contract. This is a generalization of the common law principle of contract privity, which states that third parties, other than the designated arbitrators and intermediaries, should have no say in the enforcement of a contract. Generalized privity goes beyond this to formalize the common claim, “it’s none of your business”. Attacks against privity are epitomized by third parties Eve the eavesdropper, a passive observer of contents or performance, and malicious Mallet, who actively interferes with performance or steals service. Under this model privacy and confidentiality, or protecting the value of information about a contract, its parties, and its performance from Eve, is subsumed un-

Building

Blocks

for

Digital

Free

Markets

der privity, as are property rights. The field of security (especially, for smart contracts, computer and network security), roughly corresponds to the goal of privity.

The fourth objective is *enforceability*, and at the same time minimizing the need for enforcement. Improved verifiability often also helps meet this fourth objective. Reputation, built-in incentives, “self-enforcing” protocols, and verifiability can all play a strong part in meeting the fourth objective. Computer and network security also can contribute greatly to making smart contracts self-enforcing.

Smart contracts often involve trusted third parties, exemplified by an intermediary, who is involved in the performance, and an arbitrator, who is invoked to resolve disputes arising out of performance (or lack thereof). Privity implies that we want to minimize vulnerability to third parties. Verifiability and observability often require that we invoke them. A mediator must be trusted with some of the contents and/or performance of the con-

tract. An arbitrator must be trusted with some of the contents, and some of the history of performance, and to resolve disputes and invoke penalties fairly. In smart contract design we want to get the most out of intermediaries and arbitrators, while minimizing exposure to them. One common outcome is that confidentiality is violated only in case of dispute.

Today's forerunners of smart contracts include POS (Point of Sale) terminals and cards, EDI (Electronic Data Interchange, used for ordering and other transactions between large corporations), the SWIFT, ACH, and FedWire networks for transferring and clearing payments between banks, and allocation of public network bandwidth via automated auctions [Agorics]. These implement commercial security models, but with little heed paid to the contractual needs and obligations of the parties.

Multinational Small Business

Currently the vast majority of businesses on the planet are small, but the vast bulk of multinational business is conducted by large corporations. In the future the size distribution of multinationals will approach that of local business. The phase change between these states may be quite rapid as telecomm and transport costs pass through a "melting point", creating a wide variety of new multinational small businesses, and industries to support those businesses.

Barriers to multinational small business include

- * legal: prohibitively complex maze of jurisdictions
- * language/culture*
- * telecomm costs
- * transport costs

Both transport and telecomm costs have become unprecedentedly low, and with fiber optics telecomm costs will drop orders of magnitude further. The mass media breaking down language and cultural barriers. This leaves legal barriers the most severe cost of doing business across many jurisdictions. Smart contracts can cut through this Gordian knot of jurisdictions. Where smart contracts can increase privity, they can decrease vulnerability to capricious jurisdictions. Where smart contracts can increase observability or verifiability, they can decrease dependence on these obscure local legal codes and enforcement traditions.

The consequences of smart contract design on contract law and economics, and on strategic contract drafting, (and vice versa), have been little explored. As well, I suspect the possibilities for greatly reducing the transaction costs of execut-

ing some kinds of contracts, and the opportunities for creating new kinds of businesses and social institutions based on smart contracts, are vast but little explored. The "cyberpunks"[4] have explored the political impact of some of the new protocol building blocks. The field of Electronic Data Interchange (EDI), in which elements of traditional business transactions (invoices, receipts, etc.) are exchanged electronically, sometimes including encryption and digital signature capabilities, can be viewed as a primitive forerunner to smart contracts. Indeed those business forms can provide good starting points and channel markers for smart contract designers.

Observability & Hidden Actions

One important task of smart contracts, that has been largely overlooked by traditional EDI, is critical to "the meeting of the minds" that is at the heart of a contract: communicating the semantics of the protocols to the parties involved. There is ample opportunity in smart contracts for "smart fine print": actions taken by the software hidden from a party to the transaction. For example, grocery store POS machines don't tell customers whether or not their names are being linked to their purchases in a database. The clerks don't even know, and they've processed thousands of such transactions under their noses. Thus, via hidden action of the software, the customer is giving away information they might consider valuable or confidential, but the contract has been drafted, and transaction has been designed, in such a way as to hide those important parts of that transaction from the customer.

To properly communicate transaction semantics, we need good visual metaphors for the elements of the contract. These would hide the details of the protocol without surrendering control over the knowledge and execution of contract terms. A primitive but good example is provided by the Secure Mosaic software from CommerceNet. Encryption is shown by putting the document in an envelope, and a digital signature by affixing a seal onto the document or envelope. On the other hand, Mosaic servers log connections, and sometimes even transactions, without warning users — classic hidden actions.

Cryptographic Building Blocks

Protocols based on mathematics, called *cryptographic protocols*, are the basic building blocks that implement the im-

proved tradeoffs between observability, verifiability, privity, and enforceability in smart contracts. Contrary to the common wisdom, obscurity is often critical to security. Cryptographic protocols are built around foci of obscurity called *keys*. A key's immense unknown randomness allows the rest of the system to be simple and public. The obscurity of a large random number, so vast that a lucky guess is unlikely in the all quantum events of a trillion universe lifetimes, is the foundation upon which cryptographic protocols, and in turn smart contracts, are built.

A wide variety of new cryptographic protocols have emerged in recent years. The most traditional kind of cryptography is *secret key* cryptography, in which Alice and Bob (our exemplar parties to a smart contract) use a single shared, prearranged key to encrypt messages between them. A fundamental problem we will see throughout these protocols is the need to keep keys secret, and *public key* cryptography helps solve this. In this technique, Alice generates two keys, called the private and public keys. She keeps the private key secret and well protected, and publishes the public key. When Bob wishes to send a message to Alice, he encrypts a message with her public key, sends the encrypted message, and she decrypts the message with her private key. The private key provides a "trapdoor" that allows Alice to compute an easy inverse of the encryption function that used the public key. The public key provides no clue as to what the private key is, even though they are mathematically related.

Public key cryptography also makes possible a wide variety of *digital signatures*. These prove that a piece of data (hereafter referred to as just an "object") was in active contact with the private key corresponding to the signature: the object was actively "signed" with that key. The digital signature probably should have been called a "digital stamp" or "digital seal" since its function resembles more those methods than an autograph. In particular, it is not biometric like an autograph, although incorporation of a typed-in password as part of the private key used to sign can sometimes substitute for an autograph. In many Asian countries, a hand-carved wooden block, called a "chop", is often used instead of autographs. Every chop is unique, and because of the unique carving and woodgrain cannot be copied. A digital signature is similar to the chop, since every newly generated key is unique, but it is trivial to copy the key if obtained from the holder. A digital signature relies on the assumption that the holder will keep the private

key secret.

A *blind signature* is a digital signature and secret-key encryption protocol that together have the mathematical property of commutativity, so that they can be

Eve is prevented by the Russian-doll encryption of the message by the sender with the public keys of each mix operator in the chain, and the mixing of messages by each operator, so that panoptic wiretapper

Eve loses track of the messages. For the sender/recipient pair to remain confidential, only 1 out of N of the operators needs to be trusted with their local traffic information.

have a PGP-compatible board. This is economical for central sites, but may be less practical for normal users. Besides better security, it has the added advantage that hardware speeds up the public key computations.

If Mallet's capability is to physically seize the machine, a weaker form of key protection will suffice. The trick is to hold the keys in volatile memory. This makes the PC proof from physical attacks — all that needed to destroy the keys is to turn off the PC. If the key backups can be hidden in a different, secure physical location, this allows the user of this PC to encrypt large amounts of data both on the PC itself and on public computer networks, without fear that physical attack against the PC will compromise that data. The data is still vulnerable to a "rubber hose attack" where the owner is coerced into revealing the hidden keys. Protection against rubber hose attacks might require

A nice side effect is a refreshing new view of key distribution that substitutes contractual relations as the basis for key distribution protocols, for tired models based on politics within military organizations.

stripped in reverse of the order they were applied. The effect is that Bob "signs" an object, for which he can verify its general form, but cannot see its specific content. Typically the key of the signature defines the meaning of the signed object, rather than the contents of the object signed, so that Bob doesn't sign a blank check. Blind signatures used in digital bearer instruments, where Bob is the clearing agent, and Chaumian credentials, where Bob is the credential issuer.

Secret sharing is a method of splitting a key (and thus control over any object encrypted with that key) into N parts, of which only M are needed to recreate the key, but less than M of the parts provide no information about the key. Secret sharing is a potent tool for distributing control over objects between principals.

The *zero-knowledge interactive proof* can be used to prove that a party possesses a private key, corresponding to a well-known public key (and by weaker implication, that otherwise normally functioning parties who have an incentive to respond properly to the challenge, but fail to do so, do not possess the key), without revealing any information about that key to the challenger or any eavesdroppers. ZKIPs are currently used for authentication, and in smart weapons for Identification Friend or Foe (IFF).

Information about who is talking to whom, such as can be found on telephone bills, can be quite valuable even without records of the actual content. Confidential messaging is necessary for some of the privacy features of Chaumian credentials and bearer securities to be strongly implemented on an actual network. To provide this traffic confidentiality, a *digital mix* can allow parties to communicate across a network without revealing their partners to network providers or the outside world. In a mix, traffic analysis by

tion, although Eve can sometimes gather statistics over large numbers of messages between the same partners to eventually guess who is talking to whom. The communicating parties can also be mutually anonymous, and with normal encryption need trust no other parties with the content of messages. The "Mixmaster" software on the Internet implements most of the features of a digital mix [Mixmaster].

Protection of Keys

So far, we've assumed parties like Alice and Bob are monolithic. But in the world of smart contracts, they will use computer-based software agents and smart cards to do their electronic bidding. Keys are not necessarily tied to identities, and the task of doing such binding turns out to be more difficult than at first glance. Once keys are bound, they need to be well protected, but wide area network connections are notoriously susceptible to hacking.

If we assume that the attacker has the ability to intercept and redirect any messages in the network protocol, as is the case on wide area networks such as the Internet, then we must also assume, for practically all commercial operating systems, that they would also be able to invade client if not merchant computers and find any keys lying on the disk.

There's no completely satisfactory solution to end point operations security from network-based attacks, but here's a strategy for practically defanging this problem for public-key based systems:

All public key operation are done inside an unreadable hardware board on a machine with a very narrow serial-line connection (i.e., it carries only a simple single-use protocol with well-verified security) to a dedicated firewall. Such a board is available, for example, from Kryptor, and I believe Viacrypt may also



some form of Shamir secret sharing which splits the keys between diverse physical sites.

How does Alice know she has Bob's key? Who, indeed, can be the parties to a smart contract? Can they be defined just by their keys? Do we need biometrics (such as autographs, typed-in passwords, retina scans, etc.)?

The public key cryptography software package "Pretty Good Privacy" (PGP)[15] uses a model called "the web of trust". Alice chooses *introducers* whom she trusts to properly identify the map between other people and their public keys. PGP takes it from there, automatically validating any other keys that have been signed by Alice's designated introducers.

There are two entirely separate criteria PGP uses to judge a public key's usefulness:

1) Does the key actually belong to whom it appears to belong? In other words, has it been certified with a trusted signature? 2) Does it belong to an introducer,

someone you can trust to certify other keys?

Having been told by Alice the answer to the second question, PGP can calculate the answer to the first question for the public keys Alice has collected.

Keys that have been certified by a trusted introducer are deemed valid by PGP. The keys belonging to trusted introducers must themselves be certified either by you or by other trusted introducers. This "transitivity" introduces an implicit third criterion:

3) Does the key belong to someone you can trust to introduce other introducers?

PGP confuses this with criterion (2). It is not clear that any single person has enough judgment to properly undertake task (3), nor has a reasonable institution been proposed that will do so. This is one of the unsolved problems in smart contracts.

PGP also can be given trust ratings and programmed to compute a weighted score of validity — for example, two marginally trusted signatures might be considered as credible as one fully trusted signature.

Any keys in Alice's secret key ring

swears to the truth of some affidavit before a notary or some other officer entitled to take oaths. This does not require the notary to know who the affiant is. The second act is when someone "acknowledges" before a notary that he has executed a document as "his own act and deed." This second act requires the notary to know the person making the acknowledgment. Thus, for example, the form of an acknowledgment can go something like this:

On this ____ day of _____, 19__, personally appeared before me _____, known to me and known to me to be the person who signed the foregoing instrument, and acknowledged that he executed the same as his own act and deed.


In the first type of act of notarization, the notary merely certifies that the affiant swore the statement was true. In the second type the notary actually vouches that the person making the acknowledgment was who he claims to be.

Problems with Certification

These roles of a notary public are substantially different from the alleged role of hierarchies of "certification authorities"

models suffer from more flaws. A single rooted hierarchy assumes a mythological beast called a universally trusted entity. Hierarchies in general create rigid structures that don't fit the way knowledge about keys and keyholders, and incentives to accurately reflect that knowledge, are distributed among people in the real world. In turn, while PGP distinguishes between "Alice holds a key" and "Alice can be trusted to certify a key", it does not follow that the second claim that Alice can be trusted to validate another issuer. There is little transitivity: seeing Alice's key certified by Bob, whom I know and trust, tells me little about whether Alice's certification of Charlie's key is correct. Seeing Alice certified by Bob as an introducer tells me little about whether Alice can be trusted to certify other introducers. It does tell me that I know to blame Alice if her claim turns out to be wrong; although it's far from clear that Alice has any legal liability.

There is an even more severe flaw when public key is used for digital signatures. Because the claim was only made in the past, both PGP and X.509 allow the end user to plausibly deny that they "signed" a document; and conversely if one's key is surreptitiously stolen, or for other reasons no revocation action is taken, there is no way to prove that one did not "sign" the document digitally "signed" with the stolen key. Finally, there is no widely accepted legal agreement on what is specifically being claimed when one "certifies" a key, nor is there any built-in or widely used mechanism for describing the actual claim one is making. Real world CAs have a nasty habit of



disclaiming liability for their mistakes, or for the misunderstandings that will arise out of the often vague and sloppy, sometimes implicit claims they make. Finally, there are a wide variety of other claims one might make about a key, such as "this key belongs to an office", "this key belongs to a server", "this keyholder has a good credit rating", "use this key to decrypt your new copy of Microsquish Expel", "this key is good for 100 MB of downloads from our web server", etc. which are facilitated by neither X.509 nor the PGP web of trust.

We know too little about the best uses of public-key cryptography to establish such fixed methods with such narrow semantics. This author has suggested a mechanism which modifies the PGP web of trust to create arbitrary certificate with

continued on p.61

are "axiomatically" valid to Alice's PGP program, needing no introducer's signature. PGP also assumes that Alice ultimately trusts herself to certify other keys.

It is believed that PGP causes the emergence of a decentralized fault-tolerant web of confidence for all public keys, but a chain of introduced introducers grows weak very quickly, due to lack of transitivity.

PGP's grass-roots approach contrasts sharply with traditional public key management schemes, such as X.509 and the related Privacy Enhanced Mail (PEM). These standard schemes substitute a hierarchical system of introducers called certification authorities (CAs).

Notary Publics

Two different acts are often called "notarization". The first is simply where one

(CAs) in PEM/X.509, and the "web of trust" in PGP, to "prove identity". In fact the certificates generated by these systems do no such thing. Rather a certificate *proves that a claim was made* by a CA at some time in the past. The (implicit) claim is that a particular key belonged to a particular person at that time. That key is not biometric like an autograph, and can thus be transferred at any time. Furthermore, false claims can be made by a CA about what keys an end-user has held, and the end-user can be stuck with no evidence of CA fraud; nor does the CA have any way of proving that their claim is correct if an end-user challenges it. It is extremely difficult, on the other hand, for notary publics to forge autographs and expect to get away with it often enough to maintain their professional reputations.

Both the PGP web of trust and X.509

The Transhuman Taste REVIEWS OF EXTROPIAN INTEREST

Better Sex Through Chemistry

John Morgenthaler & Dan Joy


Smart Publications, PO Box 4667, Petaluma, CA 94955, 1995.
224pp, softbound, \$14.95 ISBN 0-9627418-2-5

Reviewed by David Jay Brown



Like immortality and the fountain of eternal youth, every culture since the beginning of time has sought out aphrodisiacs and methods for enhancing the sexual experience. To explore this ancient neurochemical connection between sex and drugs with the beacon of modern pharmacology, John Morgenthaler (co-author of *Smart Drugs and Nutrients*) teamed up with Dan Joy (editor of many well-known books on psychedelics), and the result — *Better Sex Through Chemistry*—is a highly readable and practical handbook on how to safely enhance one's sex life through nutrients and a new class of pharmaceuticals they call "prosexual drugs".

Chemical treatments for such age-old problems as impotence, premature ejaculation, loss of interest in sex, as well as difficulty achieving erection, sufficient lubrication, and orgasm, are discussed along with dosages, known contraindications, corresponding scientific studies, and personal anecdotes. Also presented are chemical means to increase one's physical sensitivity and subjective enjoyment of the sexual experience. As is the case with the majority of smart drugs, most of the substances discussed in the book share exceptionally low



that has been described therein as possibly having the activity of an "aphrodisiac". L-Arginine, an amino acid also available in health food stores, has also been shown to increase the frequency of erections by increasing the production of an excitatory neurotransmitter called NO (Nitric Oxide, the only known gaseous neurotransmitter), which is widely recognized as the sole chemical responsible for causing penile erections.

GHB seems to be an aphrodisiac, tranquilizer, hallucinogen, mood elevator, communication enhancer, and highly effective sleep aid all rolled into one.

levels of toxicity. Most are extremely safe, remarkably free of side-effects, non-addictive, and in general improve health.

Yohimbe, for example, is an herb available in most health food stores, which has been shown to reliably produce a greater frequency of erections in men. The active component of the herb—yohimbine—is now actually available by prescription, and is the first and only substance to ever appear in the *Physicians Desk Reference* (PDR, 1994)


I found GHB (precursor to the inhibitory neurotransmitter GABA) the most interesting nutrient discussed. People who have used GHB claim that it lowers their inhibitions, while increasing their tactile sensitivity, the intensity of their orgasms, and their overall interest in and ease of sexual activity. My personal experience with GHB confirmed these properties, and demonstrated the existence of many

more extremely pleasant effects as well. GHB seems to be an aphrodisiac, tranquilizer, hallucinogen, mood elevator, communication enhancer, and highly effective sleep aid all rolled into one. If that weren't enough, this stuff has the secondary effects of increasing dopamine and growth hormone levels, which improve mood and cognitive performance the following day. This extraordinary substance—which is extremely safe if used with the proper precautions—was available in health food stores for years, but the FDA recently banned its sale, due to its use as a “recreational drug”. Although it is illegal to sell GHB, the DEA never scheduled it, so it's still legal to own and use it, if you can find it.

Deprenyl is the pharmaceutical substance which perked up my attention most. Deprenyl—which works primarily on the dopamine (an excitatory neurotransmitter) system in the brain and is often prescribed for the treatment of Parkinson's and Alzheimer's disease—has been shown to increase both sex drive and maximum life span in aging laboratory animals, and many people have verified that at least the elevation in sex drive carries over into the human kingdom. It has been shown to increase maximum life span in laboratory animals by as much as 40%, which is the human equivalent of living to be one hundred and seventy years of age. Interestingly enough, it was those animals that had the highest sex drives which lived the longest.


Better Sex Through Chemistry explores about a dozen other fascinating substances (such as Bromocryptine and L-Dopa) of interest to sexual enthusiasts and biochemical gourmets, along with information on how to obtain them. This assemblage of obscure hard to find information serves as an extremely valuable resource for those of us interested in designing a sexier, smarter, and healthier race of humans. ✨


David Jay Brown is the author of *Brainchild* (New Falcon Publications, 1988), and is co-author of *Mavericks of the Mind* (Crossing Press, 1993). His new book *Voices from the Edge*—which contains extensive interviews with revolutionary thinkers such as Jerry Garcia, Jaron Lanier, Annie Sprinkle, John Allen, and Ram Dass—can be ordered by calling the Crossing Press at: 1-800-777-1048.




“Better Sex Through Chemistry is really a life extension book in disguise. We're using sex to sell life extension.”— says **John Morgenthaler**


Now Morgenthaler is operating (along with **Will Block** of *Living Longer News*) a company called **Life Enhancement Products**. It is a phone/mail order company and carries all these books plus a full line of **prosexual, smart, and life-extending** nutrient products such as:

 **ProSexual Plus™**, an arginine-based drink mix that works by enhancing levels of nitric oxide (NO) in endothelial cells. (On sale for \$27.95)

 **Yohimbe Tonic™**, a standardized herbal extract. Blocks alpha adrenal receptors leading to enhanced libido in men and women and erections in men. (Price: \$19.95)

 **Smart ALEC™**, an acetyl L-carnitine (ALC) based smart nutrient drink. ALC is one of the hottest areas in cognition enhancement research today. This is the product if you want to enhance your memory, mental energy, and creativity. (Sale price: \$29.95)

 Plus many of the **Durk Pearson & Sandy Shaw®** products as well as those of Smart Drug author **Ward Dean, MD** and Russian Scientist **Moris Silber, MD, PhD**.

 Plus all the books in the photo above. Call for a FREE subscription to *Life Enhancement News* (a monthly 16-24 page newsletter on the latest biochemical enhancement technologies).


 Visa, Mastercard, American Express 
 Call **1-800-LIFE-873** Mention Code X1



Nano:

The emerging science of nanotechnology: remaking the world molecule by molecule




by Ed Regis

Little, Brown and Company, 1995
325 pages; ISBN 0-316-73858-1

Reviewed by Rob Michels

bring his ideas to the forefront of science.

He had a difficult time filling out a committee. Marvin Minsky had agreed, had even seemed to be thrilled, to be his supervisor. This made a great deal of sense. The scientific community regarded Minsky to be a bit eccentric. Some regarded him as a bit heretical, if not down right nuts. Case in point: "Nuclear explosions" Minsky once told an audience, "aren't so terrifying, because they're not self replicating. They're just *irritating*." Minsky was already thinking about the most powerful aspect of nanoreplicators—that they were self replicating, or so the theory went. Because of this difference, nanotechnology far surpassed



the emerging science of nanotechnology:
remaking the world — molecule by molecule

Ed Regis
author of *Great Balls of Fire*
and *The Transhuman Condition*

nuclear events both in their usefulness as well as in their destructive capabilities.

Where others thought of nuclear energy as both incredibly productive as well as destructive, Minsky and Drexler saw them as single (albeit incredible) events and hence quite limited in their application. Nanotechnology would make technological *processes* rather than events and that puts nanotechnology in a category of development wholly apart from all of the productive developments of the past.

Minsky had already become familiar with Drexler's ideas, having written the foreword to *Engines of Creation*, and he wanted to see the technical details of that work spelled out in a Ph.D. thesis.

The rest of Drexler's committee was, shall we say, a bit less enthusiastic. One of the problems was that Drexler's work drew on so many disciplines that few, if any, were really qualified to comment on the technical aspects. They took a top down approach to understanding his the-

Darwinian natural selection is nearly obsolete. Self and social selection have become much more important in today's world because they act quickly and, in many cases, just as decisively as natural selection. Not only do these vital changes in our lives mimic Darwinian selection, but they can occur in an afternoon. When I realized my eyesight had deteriorated I had my new glasses "in less than an hour." Had we, as a species, waited for natural selection to improve human eyesight, it would have been an exceptionally long wait. My great great... grandchildren would have been the beneficiaries; or rather I should I say, someone else's great great... grandchildren. In a Darwinian world, the battle for improvement would be between us and other species. But, it's rapidly becoming the case in our world that the battle is between social forces on one hand and autonomous selection on the other. Unfortunately, few people, other than Extropians, seem to think this way.

Nanotechnology, if we can de-

Nanotechnology would make technological *processes* rather than events and that puts nanotechnology in a category of development wholly apart from all the productive developments of the past.

velop it to an extremely sophisticated level, will not only accelerate the process of self selection, but it will give us an enormous number of lives, minds, and bodies from which to choose. Ed Regis, in *Nano*, gives us a book about the recent history of the people behind the development of the theories supporting nanotechnology. The book contains very little technical discussion; it is more of a history of the ideas and biographies of the people who deserve the credit for the

development of nanotechnology.

Regis describes: some of the early meetings of the first nanotechnology study group; the difficulties that Drexler faced at MIT while developing the ideas that we now call nanotechnology; camping trips that he and Chris Peterson took; even how they met, courted, and got married. Regis devotes much of the book to describing Drexler's life and his enthusiasm for this subject, even as an undergraduate.

Regis glorifies the work of Drexler as well as others who became involved in the development of this new technology after Drexler's Ph.D., but at the same time he presents Drexler as a rather down to earth kind of guy who got his kicks at MIT scaring Chris Peterson by climbing tall sculptures on campus without any apparent regard for safety.

Drexler had to fight his way into the position he now occupies in the scientific community. He entered MIT's Department of Electrical Engineering and Computer Science (EECS) in order to have a base department while he sought an interdisciplinary Ph.D. This came after receiving a patent on his designs for a light sail which had been his master's the-

sis. Once accepted to the EECS, he argued that he should not have to take the departmental exams that were required of every EECS degree candidate. He argued that he was not seeking a degree from EECS but rather an interdepartmental degree, so why should he take the exams required of an EECS student? He, after all, had a different agenda; the EECS was just a front and he made no real attempts to make it otherwise. This brought about the first round of battles he had to fight to

sis, looking at this new work through the eyes of people who already had an understanding of a body of knowledge and just had to apply that knowledge to the thesis.

This stands in contrast to people like Ralph Merkle who learned about nanotechnology from the bottom up. Starting with a sophisticated math background, Merkle learned nanotechnology from the basic theories up to the more generalized applications. But such people, other than Minsky, were hard to find and that was nearly the demise of Drexler's thesis. However, having received the respect of Minsky and eventually the acquiescence of the rest of the committee, MIT finally awarded Drexler a Ph.D.

While I find the history of the development of nanotechnology quite interesting, it is far from the most important part of the book, though it does provide the largest part of it. Regis gives us an easily read book that shows how other technologies have also had the manipulation of individual molecules and atoms as their goals. He documents the use of STM's (scanning tunneling microscopes) to write the IBM logo. He takes examples from biology including the study of *E. coli*'s propulsion system which relies on an electric motor envisioned by Drexler to be essential to working nanomachines. From

individual particles which have been isolated for extended periods of time in magnetic fields, to molecular "shuttles" which could move along a "track" of atoms, Regis provides many examples of engineering feats which have manipulated individual atoms and molecules.

The importance of this lies primarily in showing skeptics that this work is already in progress and has achieved some degree of success as well as notoriety. Many people look at the grand claims of nanotechnology and see only science fiction. Regis' book shows that the goals of nanotechnology are being successfully developed piecemeal by scientists in fields that nobody questions.


One of the drawbacks to *Engines of Creation* was that people who were not already sympathetic to many of the ideas presented therein just could not accept the grand conclusions. *Nano* builds up the theories first by showing that the people developing the theories are real people with senses of humor, frustrations, and yes, failed experiments. Then it goes on to bring the successes together into a description of a technology whose claims do not seem so outrageous given the sound foundation that it has in so many other respected and unquestioned fields of study.

Another important part of the book

provides non-technical responses which Drexler has provided to technical objections. These extend from the well known gray goo problem to social and economic problems.

Extropians will not be particularly surprised by this. The rest of the world however, will view many of the goals of nanotechnology as a bit too extreme and implausible. This book shows in a clear way that the goals of nanotechnology are being sought successfully by other disciplines. This will allow supporters of nanotechnology to get their foot in the door of the general populous' minds. We can now quite easily point to many examples of nanomachines or at least other machines which exhibit gears, switches, levers and so on which operate at the nano level.

If you are looking for a technical book, this is not for you. However if you are looking for some light reading with many examples of how nanotechnology is already becoming respectable in other fields (albeit under different names) I recommend this book. You will find out a lot about the people behind the scenes—even those who do not know that they are working toward the goals of nanotechnology.



BREAKTHROUGHS — A TRANSCENTURY
UPDATE

UPCOMING TOPICS: Bodybuilding—the Sport
Women & Eroticism

Flex My Mind Flex My Body

VIDEOTAPES

PART 1 Transhumanist Philosophy

PART 2 Extropy Institute, MI, Uploads

Max More, Robin Hanson, Abe Heward,
William Wiser

PART 3 Indefinite Lifespan, Uploads

PART 4 Alcor Foundation and Cryonics

Max More, Robin Hanson, Chris
Heward, Regina Pancake, Dan Spitzer,
Tim Freeman

**PART 5 Mood Modification and
Identity Choice**

PART 6 The WWW and Virtual Communities

Dr. Ray Sahelian, Max More, Jeannine
Parker, Coco Conn

PART 7 Rational Values 1

PART 8 Rational Values 2

Max More, Peter Voss, Robin Hanson,
Tim Freeman

TURN ON - TUNE IN - TRANSFORM

Each videotape features two Parts.

\$18.50 each donation. Dept. S, Extropy Institute, 13428
Maxella Avenue, #273, Marina Del Rey, CA 90292

HOST/PRODUCER: NANCIE CLARK

River Out of Eden: A Darwinian View of Life

by Richard Dawkins

Basic Books, HarperCollins Publishers, 1995; 172pp.

Reviewed by Greg Erwin



The first thing to note is that this book is part of the “Science Masters” series—a global venture to publish books by leading scientists. If you recognize the names Richard Leakey, Stephen Jay Gould, Daniel Dennett, Jared Diamond, Stephen Pinker, Murray Gell-Mann, you will be anticipating them with joy, for these names are just some of those leading scientists who have written or are slated to write Science Masters books.

There are a mere five sections in Dawkins’ *River*: The Digital River, All Africa and Her Progenies, Do Good by Stealth, God’s Utility Function, and The Replication Bomb.

The river of the title (and the digital river of the first section) is not PishTosh or OshKoshB’Gosh, or any of the other rivers running out of Eden in myths or fairy tales, but is the metaphorical stream of digital information represented by DNA; not just humanity’s, but that of all life. Somehow, sometime, inanimate matter accidentally combined to form something that could replicate itself.

What follows from this occurrence anywhere in the universe, is Darwinian selection and hence the baroque extravaganza that, on this planet, we call life.

From that one event, all the rest follows. With no need for any other help from any outside entity. In fact, if Darwin is right there are no outside entities.

Science shares with religion the claim that it answers deep questions about origins, the nature of life, and the cosmos. But there the resemblance ends. Scientific beliefs are supported by evidence, and they get results. Myths and faiths are not and do not.

Isn’t that refreshing? Science is right. Religion is wrong.

Dawkins covers the digital nature of DNA, mitochondrial Eve, speciation, the design argument. He provides sound answers to many standard questions about

the time available and the rate of mutation, (“there isn’t enough time available for reptiles to change into mammals”); creationist arguments against the possibility of developing an eye (“What good is half an eye, or half a wing?”).

As for the implications, he notes that:

We cannot admit that things might be neither good nor evil, neither cruel nor kind, but simply callous—indifferent to all suffering, lacking all purpose.

However, that is, indeed, the case.

Towards the end, he introduces the term *utility function*, the penultimate chapter being titled “God’s Utility Function.” In economics, the utility function is “that which is maximized.”

[in]... the behavior of a country’s government, you may conclude that what is being maximized is employment and universal welfare. For another country, the utility function may turn out to be the continued power of the president, or the wealth of a particular ruling family, the size of the sultan’s harem,...It isn’t always obvious what individuals, or firms or governments are striving to maximize. But it is probably safe to assume that they are maximizing something.

So, “...imagine that living creatures were made by a Divine Engineer and try to work out, ... what the Engineer was trying to maximize: What was God’s Utility Function?” See, it’s just a metaphor!

Well, it all comes down to DNA survival, nothing more, nothing less, and nothing else. Although we ask the question, why do tragedies happen? Dawkins answers:

...if the universe were just electrons and selfish genes, ... meaningless tragedies ... are exactly what we should expect, along with equally meaningless *good* fortune. Such a universe would be neither evil nor good in intention. It would manifest no intentions of any kind. ... The universe we observe has precisely the

properties we should expect if there is, at bottom, no design, no purpose, no evil and no good, nothing but blind, pitiless indifference....

DNA neither knows nor cares. DNA just is. And we dance to its music.

River Out of Eden is a basic, clear, forthright and thorough explanation of the Darwinian view of life and all of its consequences and implications. The main implication of interest being that if Darwin is right, theistic religion is wrong. Dawkins takes the time to answer the arguments against the various aspects and implications of the Darwinian paradigm and does so successfully. Just because the book is basic, does not make it less valuable; he manages to give a full treatment to the subject. And probably does so better than some others who have taken more time to say less.

Lastly, it is obvious, (he says so) that he has been in touch with others such as Daniel Dennett, while writing the book. It is completely up to date. For those who wish to read further, there is a good bibliography, and many of the books in it were published in 1995 or 1994.

So, the book comes with a thorough recommendation. If you aren’t well-versed in natural selection and evolution, you should be, and this is a very good way to start. If you already are, this is a very handy, short and yet complete reference book to have around. If you are neither, you are not likely to have read this far in the review, and need not be taken into consideration.

In fact, I have read all four of the Science Masters books so far, and can recommend the whole series. In addition to Dawkins there are: Paul Davies, *The Last Three Minutes* (End of the Universe); Richard Leakey, *The Origin of Humankind* (what else?); John D. Barrow, *The Origin of the Universe*; and more to come.

Dr. Walford's Interactive Diet Planner

Reviewed by Max More

I've tried several software programs for monitoring diet. DWIDP (to use its abbreviation) I have found the most useful. This is also the only program created by a world-class gerontologist. Dr. Roy Walford, though most recently familiar as the medic of the original Biosphere 2 team, made his considerable scientific reputation in the field of aging research. (See our interview with Dr. Walford in *Extropy* #14 and #15.) Walford has written several books, including *Maximum Life Span*, *The 120 Year Diet*, and *The Anti-Aging Plan*.

Much of Walford's work has investigated the effects of dietary restriction on life span and age-related disease. His research led him to recommend a High/Low Diet: a diet as high (and balanced) as possible in nutrition while low in calories. Ensuring that you maintain a high level of nutrition at a lowered caloric intake can be tricky. The Diet Planner makes it much easier. Even if you have no interest in lowering your caloric intake, wanting only to optimize your nutrition, the program will be ideal.

I should make a disclaimer here: I have helped out in some peripheral ways

with this program (such as by putting together the Quick Reference Guide). So I'm not an impartial reviewer. However, I was an enthusiastic user of the previous DOS version years before I had any involvement with the creators of this new Windows version. Trust me!


To record the foods you're eating, you create a foodlist. (See illustration 1.) You can enter a food item either in grams, calories, or per serving. Each time you add an item to the foodlist, several nutritional totals (calories, protein, fat, fiber, etc.) will automatically be updated. The nutrient values for thousands of foods (from the Dept. of Agriculture database) are built into the program, so you don't need to type in nutrient amounts for every food. You can add new foods as well as make recipes by saving combinations of foods.

The menus available from the foodlist screen give you access to two ways of

totals for amino acids.

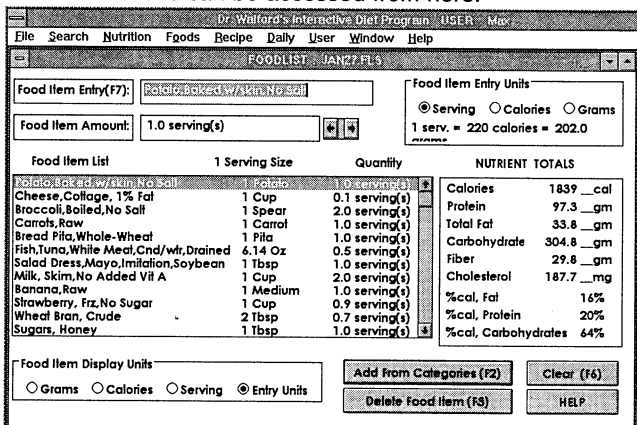
Alternatively, DWIDP can display these totals graphically in the form of a histogram bar chart. (See second illustration.)

You'll find yourself checking your totals frequently as you add foods through the day. Having this information so accessible not only makes it easier to track your nutritional rationality but actually en-



2: Histogram chart of vitamins, minerals, and other nutrients. This chart gives a rapidly graspable display of the same information available in numerical table form.

Foods are entered in this window. Most other frequently used functions can be accessed from here.



1: Foodlist.

displaying detailed information on your nutrient totals. You can choose to View All Nutrients—Table. This gives you figures on the amounts of all vitamins and minerals plus total calories, protein, fat, fiber and others, as a percentage of your daily goal. With one click of the mouse you can switch to a

courages you to improve your diet. Seeing a display of how well I'm doing (or how fair I'm allowing dietary entropy to insinuate) makes me want to push for 100% on each nutrient while keeping the bar representing fat as short as I can.

The nutrient total tables and displays are one way in which you will find yourself becoming painlessly but thoroughly educated about the foods you eat. Another way results from using the search functions. As shown in the third illustration, several options can be selected under the Search drop down menu. In addition to simply minimizing or maximizing nutrients such as saturated fat or vitamin E, the

program allows you to find foods high in some values but low in others. You might want to find foods high in fiber and vitamin A but low in fat.

Search Nutrition
Optimize
Min One
Max One
Min Several
Max Several
Min/Max Several
Repeat

3: Search options

have had enough of), such as meat, nuts, or candy.

DWIDP allows you to save nutrient total information for a range of days in log files. This means as well as seeing how you've done on any particular day, you can survey your nutritional performance over a period of time.

If you want to try caloric restriction, for its many major health and longevity benefits, this program eases the otherwise difficult goal of achieving full nutrition on few calories. Trying to ensure adequate nutrition on a tightly calorie restricted diet is a real pain if you have to look up all those values in a book. Using the program to do it makes it *fun*. If you're unsure of the value of calorie restriction, read Roy Walford's *The 120 Year Diet*. You may decide the cost in effort is worth the extended life span.

The Diet Planner will appeal to several types. For those strongly concerned about longevity and health, the program provides a powerful and efficient means of monitoring nutritional intake. Overlapping this group are those who enjoy keeping personal records, noting figures relating to personal goals, and tracking progress. DWIDP will delight such organized record keepers. The ability to see numerous numbers and graphical displays of nutrients and to maintain daily and longer term totals allows anyone to satisfy, in the healthiest way, their craving for record keeping. ❄

DWIDP requires Windows 3.1. If you are using Win95, call first to see if it has been tested with that OS. (A Mac version may be available in future.) DWIDP costs \$95 + \$3 shipping (and add another \$3 if paying by credit card). California residents add 8.25% sales tax. From: Longbrook Company, 1015 Gayley Avenue, #1215, Los Angeles, CA 90024. 310-392-8208.

We welcome reviews of books, CDs, and software. Contact the Editor if you have want to write a review, or if there is something you think *should* be reviewed in *Extropy*.

Permutation City

by Greg Egan

HarperPrism, 1995; 340pp, \$4.99 pb. ISBN: 0-06-105481-X

Reviewed by Max More

Greg Egan, an Australian writer, has been rapidly making his mark as a fresh, stimulating, intelligent new voice in science fiction. His first novel, *Quarantine*, dealt with the possibilities of neurotechnology more convincingly and absorbing than any I've come across. His latest, *Permutation City*, deals with how we might relate to Copies of ourselves running on a computer. (What we call uploads or infomorphs.)

The story starts grippingly. We follow the experiences of a Copy who discovers that his flesh self has trapped him in a virtual reality in order to conduct experiments. That, at least, *appears* to be the situation, until Egan leads the reader through a series of smartly executed twists. This thoroughly enjoyable initial episode serves as a setup for the main story.

Our next main character, Maria, allows Egan to sketch a virtual world called the Autoverse. The Autoverse compares well with Stephenson's *Metaverse* (from *Snowcrash*). While the *Metaverse* was a well described if standard virtual world, the *Autoverse* is the domain of artificial life and full blown artificial physics and chemistry. The full import of the *Autoverse* only emerges later in the novel.


Rather than give away any of the plot, I'll mention instead another aspect of the book I enjoyed. Our *Autoverse* programmer gets sucked into the intrigue by taking work in order to be able to afford to scan the brain of her terminally ill mother. Maria's mother, a member of the Church of the God Who Makes No Difference, rejects her daughter's plan, arguing that becoming a copy holds no appeal for her and that she is satisfied to die "at the appointed time". Cryonicists who have remonstrated with family or friends will have heard the same peculiar words from those who have grown up to accept death, and will be touched by our protagonist's frustration.

For those familiar with Ayn Rand's *Atlas Shrugged*, one part of the book may give you a technology-tinged sense of *déjà vu*. One of the central characters is making the rounds of uploaded billionaires, giving them detailed arguments for shrugging off society in favor of a virtual world:

She said, "What exactly is he promising these Copies?"

"A refuge. A place where they'll be safe from any kind of backlash—because they won't be connected to the outside world. No telecommunications, nothing to trace. He feeds them a long spiel about the coming dark age, when the unwashed masses will no longer put up with being lorded over by rich immortals—and evil socialist governments will confiscate all the supercomputers for weather control."

Like Egan's previous novel, *Quarantine*, this one was slightly spoiled for me by his preoccupation with a mystical view of quantum mechanics. If you can set this aside, *Permutation City* will reward you with an unusually intelligent and well crafted view of the future. This is a writer to watch. ❄



continued from p.53

a form roughly as follows:

Key about which a claim is being made type of claim, in some standard one-line format (like MIME types) Plain text description of claim Time stamp digitally "signed", Alice's key

In other words, all claims about keys should be *explicit*, readily known from reading the certificate itself, and no kind of claim should be arbitrarily excluded by the mechanism. The legal force of the claim can be based on the text itself, rather than overstated, obscure, and often implicit interpretations about what "certifying" is supposed to mean. Standard kinds of claims will emerge, including perhaps more transitive "good judge of judgment" certificates for which chain-following software can be written, and non-transitive "is-a-person" credentials directly "bound" to traditional notarized identification by a physical notarization protocol that includes both autographs and digital "signatures". More likely, new and more useful kinds of certificates will evolve. These standards should be allowed to emerge out of the wide varieties of possible end uses, much like case law has matured over time, rather than being dictated by our current very inexperienced understanding.

Meanwhile, there is a more practical defense against the man in the middle attack: advertise, early and often. Users of an insecure network can communicate the integrity of a key reasonably well by tying it to a persistent pattern of behavior: for example posts in a persistent style from a persistent e-mail address, persistence of a key unchallenged on a key server, etc. This is the most practical and widely used way by which PGP users gain confidence in public keys, and it does not require certification authorities or introducers. Those who advertise their keys widely, and those who are well known, are more likely to have keys bound to their person.

Virtual Personae

"Identity" is hardly the only thing we might want map to a key. After all, physical keys we use for our house, car, etc. are not necessarily tied to our identity — we can loan them to trusted friends and relatives, make copies of them, etc. Indeed, in cyberspace we might create "virtual personae" to reflect such multi-person relationships, or in contrast to reflect different parts of our personality that we do not want others to link.

A *nym* is an identifier that links only a small amount of related information

about a person, usually that information deemed by the nym holder to be relevant to a particular organization or community. Examples of nyms include electronic bulletin board nicknames, pen names, aliases, and brand names. A nym may gain reputation within its community. For example, a conglomerate may sell a wide variety of brand names, each reputable in its own market niche. With Chaumian credentials, a nym can take advantage of the positive credentials of the holder's other nyms, as provably linked by the is-a-person credential.

A *true name* is an identifier that links many different kinds of information about a person, such as a full birth name or social security number. As in magick, knowing a true name can confer tremendous power to one's enemies. It also can have major economic value among those who cooperate peacefully, as in the use of direct marketing to target product information to those persons most likely to be interested in those particular products.

A *persona* is any persistent pattern of behavior, along with consistently grouped information such as key(s), name(s), network address(es), writing style, and services provided.

A *reputable name* is a nym or true name that has a good reputation, usually because it carries many positive credentials, has a good credit rating, or is otherwise highly regarded. Companies strive to carry reputable brand names, while professionals such as doctors and lawyers strive to have many good personal recommendations of their name. Reputable names can be difficult to transfer between parties, because reputation assumes persistence of behavior, but such transfer can sometimes occur (for example, the sale of brand names between companies).

Constructing Smart Contracts

Blind signatures can be used to construct *digital bearer instruments*, objects identified by a unique key, and issued, cleared, and redeemed by a clearing agent. When an object is transferred, the transferee can request the clearing agent to verify that the key has never before been cleared, and issue a new key. The clearing agent prevents multiple clearing of particular objects, but can be prevented from linking particular objects one or both of the clearing nyms who transferred that object. These instruments come in an "online" variety, cleared during every transfer, and thus both verifiable and observable, and an "offline" variety, which can be transferred without being cleared, but is only

verifiable when finally cleared, by revealing to any the clearing nym of any intermediate holder who transferred the object multiple times (a breach of contract). Privacy from the clearing agent can take the form of transferee-unlinkability, transferer-unlinkability, or "double blinded" where both transferor and transferee are unlinkable by the clearing agent.

Digital cash is the premier example of a digital bearer instrument, in which the clearing agent is a bank. Bearer instrument protocols enable online payment while honoring the characteristics desired of bearer notes, especially unforgeability (via the clearing mechanism) and transfer confidentiality (via blinding).

To implement a full transaction of payment for services, we need more than just the digital cash protocol; we need a protocol that guarantees that service will be rendered if payment is made, and vice versa. Current commercial systems use a wide variety of techniques to accomplish this, such as certified mail, face to face exchange, reliance on credit history and collection agencies to extend credit, etc. Smart contracts have the potential to greatly reduce the fraud and enforcement costs of many commercial transactions.

A *credential* is a claim made by one party about another. A *positive credential* is one the second party would prefer to reveal, such as a degree from a prestigious school, while that party would prefer not to reveal a *negative credential* such as a bad credit rating.

A *Chaumian credential*, named after the prolific cryptographer David Chaum, is a cryptographic protocol for proving one possesses claims made about oneself by other nyms, without revealing linkages between those nyms. It's based around the *is-a-person credential* the true name credential, used to prove the linkage of otherwise unlinkable nyms, and to prevent the transfer of nyms between parties.

Another form of credential is *bearer credential*, a digital bearer instrument where the object is a credential. Here the second party in the claim refers to any bearer — the claim is tied only to the reputable name of issuing organization, not to the nym or true name of the party holding the credential.

Smart Property

We can extend the concept of smart contracts to property. Smart property might be created by embedding smart contracts in physical objects. These embedded protocols would automatically give control of the keys for operating the property to the party who rightfully owns that property,

based on the terms of the contract. For example, a car might be rendered inoperable unless the proper challenge-response protocol is completed with its rightful owner, preventing theft. If a loan was taken out to buy that car, and the owner failed to make payments, the smart contract could automatically invoke a lien, which returns control of the car keys to the bank. This "smart lien" might be much cheaper and more effective than a repo man. Also needed is a protocol to provably remove the lien when the loan has been paid off, as well as hardship and operational exceptions. For example, it would be rude to revoke operation of the car while it's doing 75 down the freeway.

Smart property is software or physical devices with the desired characteristics of ownership embedded into them; for example devices that can be rendered of far less value to parties who lack possession of a key, as demonstrated via a zero knowledge interactive proof.

One method of implementing smart property is through operation necessary data (OND): data necessary to the operation of smart property. For example, a complex, OND can be proprietary firing sequence needed to operate a computerized engine, a CAM file needed to manufacture a specialized part, etc. To avoid theft of service, ZKIP is required to open an encrypted channel to the device. To avoid leaking the OND to Eve, tamper detection combined with a dead-man switch can be used on the device end of the channel.

We might also use ingrained immobilizing or destructive devices to foil attempts to hot-wire smart property.

A *smart lien* is the sharing of a smart property between parties, usually two parties called the owner and the lienholder. This property may be in the proximate possession of the owner or the lienholder, corresponding to the common-law notions of "artisan's lien" and "innkeeper's lien" respectively. Smart liens might be used to secure lines of credit, insurance policies, and many other kinds of contracts that involve smart property.

How can debts be collected? No wise bank will lend to an party that can either be coerced into repaying the debt, or that is more than covered by securely liened collateral plus some conservative function of its reputation for payment in full and on time. For all parties, both credit and liability are closely related, and limited.

The liability of a party is limited by that party's liens and by the ability to deter that party by threatening punishment for violating contracts (i.e., committing

crimes as defined by the contract with the jurisdiction). The potential for other actions an party might take that cause liability, such as damage to others' persons or property, also need to be limited. More on that later.

Many parties, especially new entrants, may lack this reputation capital, and will thus need to be able to share their property with the bank via secure liens. A lien is, in a practical sense, a method of sharing a piece of property between the "owner of record" and a "lienholder", instead of the property having strictly one owner. Liens are used in many large credit transactions, such as auto loans, mortgages, farm loans, etc. They are enforced by the jurisdiction specified in the contract; usually this enforcement is done by the government and subsidized by the taxpayers rather than paid for by the contracting parties. (In fact this usually is the case with contracts and property rights in general, the enforcement clause is an implicit government subsidy). One way to implement a lien without governments is via co-signing with your privately chosen arbitrator (as long as the arbitrator has a good reputation and the contractual right to take appropriate action against you). Smart liens might greatly expand the privity and security of such arrangements.

As is the case today, credit problems will usually be solved by artfully written, menacing dunning letters and dings to one's credit rating long before the lien needs to be invoked. However, the lien needs to be enforceable to make these dunning letters credible over the long run.

What about extending the concept of contract to cover agreement to a prearranged set of tort laws? These tort laws would be defined by contracts between private arbitration and enforcement agencies, while customers would have a choice of jurisdictions in this system of free-market "governments". If these privately practiced law organizations (PPLs for short) bear ultimate responsibility for the criminal activities of their customers, or need to insure lack of defection or future payments on the part of customers, they may in turn ask for liens against their customers, either in with contractual terms allowing arrest of customers under certain conditions (e.g. if they commit acts specified as criminal by the PPL contract) or (more likely for mobile world-traveling and virtual pseudonymous customers) smart liens against liquid assets such as bank accounts and investment portfolios. Smart liens over information, such as digital bearer securities, can be implemented via secret sharing (two or more keys required to unlock the encryption).

Other important areas of liability include consumer liability and property damage (including pollution). There need to be mechanisms so that, for example, pollution damage to others' persons or property can be assessed, and liens should exist so that the polluter can be properly charged and the victims paid. Where pollution is quantifiable, as with SO₂ emissions, markets can be set up to trade emission rights. The PPLs would have liens in place to monitor their customer's emissions and assess fees where emission rights have been exceeded.

Alas, there are some dangers where maximum damage could far surpass any liens. A good rule of thumb here is that if the risk is against a third party, and it cannot be liened or insured against, then PPLs should not allow it to be taken. PPLs that allow their customers to take such risks against non-PPL parties would ruin their credit rating. One example of such a risk is building a nuclear plant for which no insurance company is willing to submit liability coverage. If a plant is safe, presumably one should be able to convince a good insurance company to cover its potential to damage others' property.

A Smart Contract Fable

A long time ago in a galaxy hopefully far, far away...

The battle of the sexes has reached its nadir. The planet Cryptonis bitterly divided between a mob of militant feminists on the continent of Amazonia, and a gang of male chauvinist mafiosi on the continent of Barbaria. The sexes have been segregated accordingly, regardless of their personal preferences. All direct trade and communication between continents faces either stiff extortion by the mobs who rule these lands, or outright ban. The only people free of this tyranny are peaceful eunuchs who inhabit the small islands that dot the planet. These little communities have no prejudices against one sex or the other, and thus communicate and trade freely with both sides.

Ironically, it is the eunuchs who have discovered a form of "vive la difference": some economic preferences are best satisfied by the opposite sex. Thus Alice and Bob often want to trade with each other, but due to their mob masters they can only trade with eunuchs. To satisfy this long repressed demand for illicit virtual intercourse (commercial intercourse, you people), the eunuchs have deployed a whole array of mediation and arbitration services, accompanied by reliable networks and crypto technology: encryption,

One of the techniques of the eunuchs

is *key ownership*. That is to say, ownership via the possession and use of keys, as opposed to the ownership of keys, and in contrast to physical possession, legal dictate, or other means by which one might "own" something. A more precise but clumsier term is "key-based ownership". One might also call this "shadow", "crypto", "mirror", "virtual", "free", or "truly private" ownership, but each of those terms has drawbacks.

We'll study the special case of *key owned corporations*: how our eunuchs create these new corporations, and why, and what "ownership" of such a corporation consists of. This is, alas, only an overview, with many necessary details left out: partly because of space, partly because this storyteller hasn't elaborated these protocols to their full detail yet! Feel free to fill in your own protocols, or critique the scenario described herein.

Alice and Bob like to invest in stocks, preferably with high return and low risk. Holding a portfolio consisting of the same sex exposes one to too much of the same kind of risk, especially if one already "invested" one's career, etc. with that sex. Barbarians could diminish that risk by investing in Amazonia, and vice versa. Alas, this kind of business is either forbidden outright, or exposed to outrageous extortion by the mob. How can the eunuchs play commercial matchmaker?

The simplest way to shadow stock ownership is to create a *synthetic asset* that mimics the real shares financially, but are officially assets of the eunuchs, not of the opposite continent. Synthetic assets are a special kind of derivative that is "synthesized" by combining other securities and derivatives into a more complex instrument. Various ways of creating synthetic variable-income (e.g., stocks) and fixed-income (e.g., bonds) instruments have been implemented[1]. New, fast computers have made it possible for the eunuchs to analyze the cash flows and risk profiles of very complex synthetic assets, a task that would have been far too tedious with their old trusty abaci.

Alice buys synthetic shares of Bob's Big Boy from the eunuchs, Bob buys synthetic shares of Alice's Restaurant, and the eunuchs buy corresponding official shares on each continent. The eunuchs charge Alice and Bob transaction costs plus profit for themselves; Alice and Bob eagerly pay since they can now hedge against the economic failings of their own sex. In effect, the eunuchs have used synthetic assets to get around discriminatory barriers to buying otherwise legal instruments,

Alas, these synthetics still don't allow Alice a vote like the other Big Boy shareholders; the salubrious threat of corporate raids by the opposite sex remains thwarted. It may be possible for the eunuchs to use cryptographic protocols, such as David Chaum's digital bear instruments (for bearer stock certificates) and voting protocols, combined with secret sharing of keys that control the corporation's resources (access to computers, databases, networks, buildings, bank accounts, etc.), to implement a key corporation underneath the official corporation.

The *key shares*, or shadow shares, combine synthetic assets which mimic the official shares financially, with bearer certificates, which can be traded confidentially like digital cash. When it comes time for electing the shadow board of directors, or *key directors*, bearers participate in protocols which combine voting with the splitting of master key(s) between the key directors. M out of N key directors can reconstruct the key. Meanwhile, an official election is held among the official stockholders, but they obtain no keys, only public and mob recognition as the official officers. Finally, the keys of the old key directors automatically expire. There is a solemn, but secret ceremony to mark the transfer of key control to the new directors.

All information security is accomplished through obscurity. A key becomes a focus of that obscurity; its immense unknown randomness allows the rest of the system to be simple and public.

"Possession is 9/10 of the law". Who now controls the corporate assets? The mobsters, with their mastery of violence, still have, if they desire, some control over the physical assets (buildings, machines, etc.), and those informational assets tied to physical assets (such as brains and hardware). But this may be only a crude kind of control, which can threaten destruction and denial of service, but cannot gain fine grained control over the detailed workings of the company.

Assets in software and computer data can be encrypted, the keys distributed amongst the key officers, and/or among trusted large crypto shareholders. Online services and databases can be distributed among several physical locations, and critical services have "hot backup", in which a new computer comes online automatically when the old one goes down (deadman switch), with transactions in progress canceled (the state of the networked computation rolled back to the last transactions completed before the attack). Thus, even denial of online service

attacks become difficult. Theft of online services is hard due to strong authentication (e.g., zero knowledge interactive proofs). With strong encryption theft of information seems to require moles inside the organization; but those moles' lives are made easier by digital mixes.

Access to buildings, machines, etc. can often be restricted by cards requiring the proper key(s). By itself this provides only a temporary barrier against concerted physical attack, and no defense against denial of service. Combined with tamper-indicating technologies and booby traps, key control might be used to render property worthless to would-be robbers and extortionists, thus leaving denial of service as the only attack remaining for the mobsters.

Clearly, there is a jungle of detail to get through when we start delving into just how much control can be maintained by keeping control over a corporation's keys, and which kinds of control remain vulnerable to the ideological and extortionist whims of the mob. The issue of control of the keys is also critical, though assumed in the preceding discussion. Key distribution, revocation, and transfer among key shareholders and key officers need to be elaborated. A nice side effect is a refreshing new view of key distribution that substitutes contractual relations as the basis for key distribution protocols, for tired models based on politics within military organizations.

Similar issues arise in control of other kinds of property; for example with *liens*, parties effectively want to share control over a property, with that control changing upon proof of some event (or lack of it), such as proof of payment. Such "smart liens" will be critical to many kinds of contractual enforcement, such as repayment of loans, in the nonviolent society that the enemies of the mob favor. I look forward to more analysis of these matters; and if you like, to further adventures on our poor repressed planet Crypton!

Smart property may be a ways off, but digital cash and synthetic assets are here today, and more smart contract mechanisms are being designed. So far the design criteria important for automating contract execution have come from disparate fields like economics and cryptography, with little cross-communication: little awareness of the technology on the one hand, and little awareness of its best business uses on the other. The idea of smart contracts is to recognize that these efforts are striving after common objectives, which converge on the concept of smart contracts. ❁

References on p.63

ENIGMA cont. from p.49

lics, and we know 2, 4, 6, 8, 9, 16, 17, 18, 19 to be the values of nine other republics, which we do not know the names of. Adding up the squares of the nine known and nine unknown, we have a total of 9885 square miles. Subtracting their total from the USSR's 12544 square miles, we are left with 2659 square miles, to be divided among the remaining three squares. 2659 is the sum of three squares in two ways; the sum of the squares of 7, 33, and 39, and the sum of the squares of 27, 29, and 33. Since Dorra is known to be 7 miles wide and no squares are equal in size, it must be the second combination of 27, 29, and 33. Used in conjunction with clue D., we know Phydra=27, Curro=18, Merm=9, and Zevo=2.

Now that all the square sizes are known, they can be assembled, much like a jigsaw puzzle, into the larger square of the USSR (albeit a rather counterintuitive one). There is a unique solution, and it is the unique, smallest simplest perfect squared square discovered by mathematician A. J. W. Duijvestijn in 1978 (Thanks, A. J. W.!).

We can identify thirteen of the squares already, and clues G., H., I., and J. allow us to identify the remaining squares, as well as rule out reflections and rotations of the solution. Jarp is in the remaining corner, and is 33 miles wide, it shares a border with Elga, which by clue G., must be the smaller of the two unnamed neighbors, with Wintz the larger one. By clue I., Rudra shares borders with Phydra and Nurin, and so must be the square 8 miles wide; it also borders Lenif, so Lenif must be the square which is 19 miles wide. Clues I. and J. combined place the remaining squares of Koid, Hort, and Okell. Since Quadra, the capital city, is in the center of the USSR, it lies in the republic of Merm, and the map is complete. ☛

MELATONIN cont. from p.45

to insomnia, to take the lowest effective doses and to take them as early in the day as possible so as not to upset precious sleep patterns.

I know better than to consume stimulants in the evening, but knowledge of facts does not guarantee us from occasionally making unwise decisions.

MINDSURFING from p.41

encryption software [7]), may already be so widespread and entrenched as to be unkillable, even if as a political matter such a thing became possible. Further, the intense commercialization of the Internet almost assures the Internet's continued existence simply by reason of corporate self-interest and consumer demand.

The Internet is obviously the largest and most important single thing on the planet. Information and ideas compete in a Darwinian environment in which only the accurate and relevant survive. With this mechanism in place, humanity can collectively heave itself out of the darkness of fear and superstition with astonishing speed. In the year 2525, after your brain is removed from cryonic suspension and processed by nanorobots, exactly what kind of Internet do you expect to be uploaded onto? Think about that—and your baby's head—the next time you're in a voting booth.

An HTML version of this article (with plenty of embedded links) can be found at <http://www.paranoia.com/~ebola/ms/2.html>.

SMART from p.63

References:

- [1] *The New Palgrave: Allocation, Information, and Markets*
- [2] Bruce Schneier, *Applied Cryptography*
- [3] *Crypto* and *Eurocrypt* conference proceedings, 1982-1994.
- [4] "Crypto Rebels", *Wired* #2, also cypherpunks mailing list (mail to majordomo@toad.com with body "subscribe cypherpunks")
- [5] Perry H. Beaumont, *Fixed Income Synthetic Assets*

It's 3 am and I just got my first yawn of the night.

Landolt H, Caffeine reduces low-frequency delta activity in the human sleep EEG. *Neuropsychopharmacology* 12:229-238, 1995. ☛

References

- [1] CDT Counter-Terrorism Page http://www.cdt.org/policy/terrorism/internet_bomb.test.html
- [2] VTW BillWatch Newsletter <http://www.vtw.org/billwatch/>
- [3] EFF Exon Archives http://www.eff.org/pub/Censorship/Comm_Decency_Act/
- [4] Hotwired: The Shoddy Article <http://www.hotwired.com/special/pornscare/godwin.html>
- [5] AOL-Sucks Homepage <http://www.en.com/users/tfinley/>
- [6] Exon Vote Tally gopher://gopher.panix.com:70/00/vtw/exon/senate-response
- [7] PGP Distribution Site (MIT) <http://web.mit.edu/afs/net/mit/jis/www/pgp.html>

Several readers quite correctly pointed out that SLIP and PPP accounts are now generally much cheaper than reported in MindSurf #1 (The tia Transformation). The latest tia version emulates PPP as well as SLIP, and is still a way cool hack for those who only have command-line access.

MindSurfing is a series of articles about the cutting edge of the Internet for the dial-up user. Reader flames and questions may be sent to: yow@primenet.com. ☛

- [6] Tim May, *Cyphernomicon*
- [7] Vernor Vinge, "True Names" (fiction), from *True Names and Other Dangers*
- [8] ??, *Fixed Income Synthetic Assets*
- [9] "The Age of Digital Capitalism", *Forbes*
- [10] Frederick Hayek, *Fatal Conceit: The Errors of Socialism*
- [11] Ming Li & Paul Vitanyi, *An Introduction to Kolmogorov Complexity and Its Applications*
- [12] Agorics white paper.
- [13] Mixmaster (ptr to web site)
- [14] PGP (ptr to web site)
- [15] DigiCash web site

Now reduced print in *The (Libertarian) Connection*, open-forum magazine since 1968. Subscribers may insert two pages/issue free, unedited. Lots of stimulating conversation. Eight issues (year) \$20. 101 South Whiting #700X, Alexandria, VA 22304.

classifieds

I am looking for a publisher for my conservative futurist tracts. Can you help? Call 516-938-3061

CONTRIBUTORS

David J. Brown: Author of *Brainchild* (New Falcon Publications, 1988), and co-author of *Mavericks of the Mind* (Crossing Press, 1993). His new book *Voices from the Edge*—which contains extensive interviews with revolutionary thinkers such as Jerry Garcia, Jaron Lanier, Annie Sprinkle, John Allen, and Ram Dass—can be ordered by calling the Crossing Press at: 1-800-777-1048. dajabr@well.sf.ca.us

Alexander (Sasha) Chislenko: Alexander "Sasha" Chislenko got an M.S. in Mathematics, Computer Science and Education from the University of Leningrad, USSR in 1982. Currently lives in Boston, reading and writing on various futuristic and philosophical subjects, raising a gifted child, dancing, hunting mushrooms, and looking for an occupation in his sphere of interests. sasha1@netcom.com — <http://linux1.uwc.edu/~sasha/home.html>

Nancie Clark: Futurist and media architect, earned graduate honors from Academia Belli Arte (Italy); videos: "Two Women In B&W" (Women In Video) and "T – And Counting" (U.S. Film Festival). Nancie has been an activist in transhumanism for over a decade, the optimism of which is clearly exhibited throughout her work. eonflex@ix.netcom.com

Greg Erwin: Greg Erwin is a former American, now a Canadian citizen. Also a former Methodist, Yogi, Mormon and Unitarian, who has been an atheistic secular humanist and nullifidian since 1988. His electronic magazine, "The Nullifidian" is issued once a month, containing reprints of humanist articles and original items. During the day he works for the Canadian government at the National Parole Board, concealing his secret identity. ai815@Freenet.Carleton.ca

Mark Grant has a Master's Degree in Physics from Oxford, and is currently working as a software engineer, but is intending to become a technomad within five years. He is the author of Privtool, a free PGP-aware mail program for Sun workstations, and is working on a near-future SF RPG with extropian leanings. He owns far too many Italian cars, and maintains the (unofficial) Lancia WWW page. mark@unicorn.com WWW: <http://www.c2.org/~mark/>

Duane Hewitt: Currently completing a MSc. degree in Molecular Biology in the regulation of neuronal specific gene expression at the University of Calgary. Duane is a member of the Deamon Maxwell group, a think-tank focusing on the investigation and implementation of speculative scientific concepts. ddhewitt@acs.ucalgary.ca

Rob Michels: Rob is finishing his Ph.D. in philosophy at the University of North Carolina at Chapel Hill. rmichel@email.unc.edu

Max More: President of Extropy Institute and co-founder of *Extropy*. While studying Philosophy, Politics, and Economics at Oxford University, Max founded the first British cryonics organization, MIZAR (now Alcor-UK). He recently received his Ph.D. in philosophy at the University of Southern California. maxmore@primenet.com

Ray Sahelian, M.D. is a physician certified by the American Board of Family Practice. He obtained a Bachelor of Science degree in nutrition at Drexel University in Pennsylvania and completed his doctoral training at Thomas Jefferson Medical School. In addition to a medical practice, Dr. Sahelian is a frequent contributor to health magazines and has made numerous appearances on radio and television to discuss his personal experiences with melatonin and the current research now being conducted on this interesting health supplement. He is the author of *Be Happier Starting Now* and *Melatonin: Nature's Sleeping Pill*. dr.ray@ix.netcom.com

Julian L. Simon: Professor of Business Administration at University of Maryland at College Park.. Simon is the author many books about the economics of population and resources including *The Ultimate Resource*, and the forthcoming *The State of the World*, as well as other books including *Good Mood: The New Psychology of Overcoming Depression*.

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, Sequent, DigiCash, and currently is contracting for Agorics, Inc. on Internet commerce. szabo@netcom.com

Mark Wolf, Ph.D.: Mark J. P. Wolf recently finished a dissertation on digital technology's effect on art, communication, and cognition. He is now teaching in the Communication Department at Concordia University in Wisconsin. He enjoys recreational mathematics, game design, and writing. mwolf@luther.cuw.edu

Yow: Yow is a pseudonym for Steve Arbuss, an attorney who spends a good part of each day in cyberspace. A partner in the Los Angeles law firm of Pircher, Nichols & Meeks, Arbuss also lectures and writes about the Internet. He can be reached at yow@primenet.com

Production information

Extropy #16 was produced on a Gateway P5-100 with 16Mb of RAM, 1.5Gb hard disk, 17" NEC 5FG monitor powered by a graphics card with 2Mb of video memory, and a quad-speed CD-ROM, using Pagemaker 5.0 for Windows, MS Word 6, Pixar Typestry 1.1, Photoshop 3.0, Ray Dream Designer 3.0, and Aldus Freehand 4.0. Scanned images were input by a Microtek Scanmaker IISP, and processed with Photoshop 3.0. The proofs were printed at 600dpi on an HP Laserjet 4 with 6Mb of RAM, and final output at 2,400dpi on an Agfa typesetter. PageMaker files were sent to the printer on an 88Mb SyQuest cartridge. Complete chaos avoided with the help of Lotus Organizer 1.0. Layout by Max More with Nancie Clark.

This issue was printed by Ripon Community Printers, Ripon, WI. Print run: approx. 5,000

Back Issues

#1-15: \$7 each from
Extropy Institute, 13428 Maxella Avenue, #273
Marina Del Rey, CA 90292

#15 vo.7 no.2 (2nd-3rd Quarter 1995): Future Forecasts by Benford, Bridge, Drexler, FM-2030, Miller, More, Szabo; Introduction to Digital Cash, by Mark Grant; Thoughts on the Economics of "Digital Money", by Lawrence H. White; Hayek's *Denationalisation of Money*, by Max More; Ecological Experiments, Space Habitats, & Long Life: An Interview with Roy Walford, M.D. (Part 2) by David Krieger; Profile: FM-2030; Bio-Enhancement Update: Melatonin, by Dr. Ray Sahelian; Enigma: Squared Deal, by Mark Wolf; Mindsurfing: The Internet Adapter, by Yow; Consciousness: Spontaneous Orders and Selectional Systems, by Reilly Jones; Reviews of *The World of 2044*, *The Theory of Free Banking*, *The Millennial Project*.

#14 vo.7 no.1 (1st Quarter 1995): Utility Fog, Part 2, by J. Storrs Hall; Methuselah's Kitchen: An Interview with Roy Walford, M.D. (Part 1) by David Krieger; Evolutionary Architecture and Extropian Consciousness, by Fred Stitt; Lilliputian Uploads, by Robin Hanson; Bio-Enhancement Update: Prozac and the Next Generation, by Dr. Ray Sahelian; Reviews of *The Physics of Immortality*, *Morals By Agreement*, and *Bright Air, Brilliant Fire: On the Matter of the Mind*.

#13, vol.6 no.2 (3rd Quarter 1994): Boundless Constellations: The Emergence of Celestial Civilization, by Nick Szabo; If Uploads Come First: The Crack of a Future Dawn, by Robin Hanson; Utility Fog, Part One, by J. Storrs Hall; Two Questions for Extropians, by Charles Platt with response by Max More; Souls, Cyberspace, Sins, and Singularity: A Conversation with Dave Ross, Part 2, by Dave Krieger; Neurocomputing 7: Sequential Neural Nets, by Simon D. Levy; Humor: Galactomatic-1000, by Carl Feynman; reviews of *The Origins of Order: Self-Organization and Selection in Evolution*, *Good Mood: The New Psychology of Overcoming Depression*.

#12, vol.6 no.1 (1st Quarter 1994): A Practical Look at Ocean Colonization; The Last Free Place on Earth; Logical Languages: A Path to Posthuman Rationality?; The Open Society and Its Media; God and Man at Yale: A Conversation with David Ross, pt.1; Forum: Nanarchy (automated police and defense systems). Wormhole Warfare; Reviews of *Fuzzy Thinking: The New Science of Fuzzy Logic*, and *The Children's Machine*.

#11, vol.5 no.1 (2nd Half 1993): Uploading Consciousness, by Ralph Merkle; Extropian Principles v.2.5, by Max More; Traversable Wormholes: Some Implications or Contact! A Post-Singularity Phase Change; A Conversation with Mark Miller, Part 2: The Day the Universe Stood Still; "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 1993): Pigs in Cyberspace, by Hans Moravec; Protecting Privacy with Electronic Cash; Technological Self-Transformation, by Max More; Mark Miller interview, Pt.1: Creole Physics & the Credit Theory of Identity; Nanocomputers: 21st Century Hypercomputing; The Transhuman Taste (Reviews): Two books on Ayn Rand & Objectivism; *Nanosystems*; *Genius*.

#9, vol.4 no.1 (Summer 1992): The Extropian Principles, 2.0; Extropy Institute Launches; Persons, Programs, and Uploading Consciousness; Nanotechnology and Faith; The Making of a Small World (fiction); Genetic Algorithms; Time Travel and Computing; Futique Neologisms 3; Exercise and Longevity; 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; Dynamic Optimism; Neurocomputing 5: Artificial Life; Futique Neologisms 2; Extropia: A Home for Our Hopes; Human-Transhuman-Posthuman; reviews of *David's Sling*, *Unbounding the Future*, *The Silicon Man*. (Photocopy of original.)

#7 vol.3 no.1 (Spring 1991): A Memetic Approach to 'Selling' Cryonics; Privately Produced Law; Order Without Orderers; Futique Neologisms; Neurocomputing 4: Self-Organization in Artificial Neural Networks; 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; The Thermodynamics of Death; The Opening of the Transhuman Mind; The Extropian Principles; Neurocomputing Part 3; 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; Arch-Anarchy; Deep Anarchy; I am a Child; Perceptrons (Neurocomputing 2); On Competition and Species Loss; A Review of Intoxication; Intelligence at Work; Extropian Resources; The Extropian Declaration; Our Enemy, 'The State'.

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

#3 (Spring 1989): Forum; Love as a Contractual Relation; Love as a Sharing of Values; Agapeic Love; Sexual Information; Psychedelics and Mind-Expansion.

#2 (Winter 1989): Review of *Mind Children*; Darwin's Difficulty; A Truly Instant Breakfast; Wisdomism; Nanotechnology News; Weirdness Watch.

#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, Immortalism, Nanotechnology, Spontaneous Orders, Psychochemicals, Extropic Psychology, Morality, Mindfucking, Space Colonization, Libertarian Economics and Politics, Memetics, and Aesthetics; "Morality or Reality".

