The Political Science of the Internet

Jerrold D. Prothero
Human Interface Technology Laboratory
prothero@hitl.washington.edu
http://www.hitl.washington.edu

May 1996
(Last Update: July 4rth, 1996)

HITL Technical Report Number TR-96-2

Contents

"So we must lay it down that the association which is a state exists not for the purpose of living together but for the sake of noble actions."
[ARIS]

Abstract

Stable trends are changing the foundations of society. These trends are briefly reviewed. Broadly, they consist of information theory replacing physics as the dominant framework; information replacing material goods as the dominant economic sector; and network topology replacing geography as the dominant space of interaction. This paper is concerned with the political and social implications of these trends. While the details of the future are chaotic, hence unpredictable, it is argued that the broad strokes of future social development are driven by exponential curves which have been remarkably stable. The primary conclusion is that geographic organizations, including geographic nation-states, will become much less relevant and that standing online organizations and nation-states will form with a new type of political structure. The implications of this and other conclusions are addressed.

1.0 Introduction

Order-of-magnitude improvements in technology result not simply in the old done more efficiently, but in the creation of entirely new capabilities. Motorized transport was not simply better horses. It afforded fast, cheap transportation. This new creation allowed for new types of human interaction. Aside from applications to warfare, motorized transport led to the growth of the suburbs and of inter-regional trade. To the extent that trade ties groups together, it has played a role in the political unification of Europe.

The combination of high-bandwidth communication, powerful computers and strong encryption also creates a new capability. As with motorized transport, the social implications of the information revolution are more profound than the technology. In Mumford's account [MUM], principles of the political state appear to have developed by 3000 B.C. during the creation of the first city-states (amidst the most dramatic technological explosion prior to modern times). These principles can be taken to be: a strong government with a local monopoly on the use of force; an elaborate legal system; systematic taxation; a stratified society; and the ability to organize large-scale operations. All supported by coercion if necessary. These basic principles have not fundamentally changed since 3000 B.C., although the terms of their application (e.g., monarchy, oligarchy, democracy) have fluctuated. A primary purpose of this essay is to suggest that the information revolution will strike at several of these tenets for the first time in five millenia. While not a new idea, the attempt here is to systematically document the underlying trends and to unfold as carefully as possible their likely political implications.

Nothing dates faster than prediction. Beyond unforeseeable events, one can not know all one should about current engineering, science, economics and politics, much less the many useful historical insights which are available. This is particularly true for an essay such as this one, written briefly in time, length and resources. My justification is that the underlying trends on which this essay is based seem sufficiently stable to support useful prediction. Beyond this, remember the word "essay" comes from the French for "trial" or "attempt", and read what follows in that light.

Finally, the following predictions are what I believe we can plausibly expect. Not all are desirable. The Society for Creative Anachronism [SCA96] is devoted to the recreation of the past not as it was, but as it should have been. However, one must predict not as it should, but as it will be.

2.0 Trends

This section briefly discusses the trends on which this essay is based and why one may reasonably expect them to be stable. The consequences of these trends are discussed in the next section.

2.1 Cheaper Computation

Moore's Law [INTEL96] states that the density of transistors on a chip, and thus price/performance, doubles every 18-24 months. The law has held since the early sixties and is expected to continue for at least the next decade, perhaps accelerating [AU94]. Moore's Law may be an underestimate of future computational power. Parallelism, new synthesis techniques, and new computing technologies (including optical) can only accelerate the trend: they will have an impact only if they succeed in beating the monoprocessor price/performance curve.

In particular, we have currently sequenced less than 1% of the human genome and know little about micromanufacturing. Within less than a decade, we will have sequenced essentially 100% of the genome and have a developing knowledge of how to apply it [SN9654]. Incorporating billions of years of study on how to structure matter at the molecular level cheaply and effectively is likely to result in molecularly-engineered computing engines. One can't predict which application of this technology to computing will prove most cost-effective, but possibilities include:

° Modification of self-assembling molecules like actin to produce cellular automata.

° Neurons grown into useful computational layouts, based on an understanding of biological neural nets.

° Bacteria trained to etch the resist on conventional silicon chips, resulting in huge, regular field-programmable gate-arrays implemented in standard electronic technology, which can then be programmed to emulate the desired CPU architecture.

2.2 Cheaper Communication

Although once lagging, telecommunications bandwidth is now conservatively estimated to be growing at 65% per year, faster than Moore's Law. Individuals are for the most part still in the kilobit-per-second range, but the telephone companies were expected to deliver 2.4 gigabits-per-second to their corporate vendors by the end of 1995, with 10 Gbps and 40 Gbps projected for the following two years [MMG95].

2.3 Generic Computational Engines Will Dominate On Price/Performance

The cost of building a new state-of-the-art fabrication plant for a microprocessor is about a billion U.S. dollars. The plant is considered obsolete after 5 years [INTELFAB]. A large sales volume is required to pay off such an investment, which implies that specialty microprocessors are unlikely to be able to compete on price/performance with generic processors aimed at the widest possible market. Consequently, computers are likely to consist of a single generic microprocessor (on the low-end) or else a combination of generic microprocessors in a single platform or networked together.

2.4 Engineering Will Be Limited By Complexity

Software is limited primarily by our ability to state what it is that we want to have happen, rather than by material properties or physical resource constraints. Consequently, the limiting factor on the quantity, quality, and price of the software we produce is our ability to surmount the complexity of the tasks we describe.

It is sometimes argued that the software crisis is due to a failure of the software industry to adopt a manufacturing approach, with standardized parts and procedures for construction. This misses the point. The software equivalent of manufacturing is not only done, it is done so cheaply and efficiently that the process is usually ignored. It consists of copying the finished software. The right comparison is between software development and the design of the manufacturing process for a new product. As mentioned previously, Intel spends a billion dollars on a new fabrication plant [INTELFAB]. Software development for major projects is not inherently a less complex design problem. In principle, it can be much more complex.

Is it possible for design problems, such as software development, to be solved by assembling standardized parts, as is done for manufacturing? To some extent it is, as the proliferation of object-oriented toolkits testifies. But it is no secret that these toolkits have hardly served to end the software crisis. Design resists standardization because there is a fundamental trade-off between efficiency and flexibility, as Conrad [CONRAD] demonstrates on thermodynamic grounds.

Intuitively, an explanation for the efficiency-flexibility trade-off is as follows. A problem solution can be expressed in terms of a language, in which the components of the solution are the words and the rules for combining the components are the grammar. For simplicity, assume that all problems in some set are equally important (removing this assumption does not change the result). Then a language should try to have an encoding of equal length for all the problems it covers. Imagine using two languages to express the same particular problem. One of these languages is very flexible: it can solve a wide range of problems in addition to the particular one. The other language is specialized to the neighborhood of the particular problem in question. Then the length (and therefore complexity and inefficiency) of the representation in terms of the flexible language must be longer, because the specific problem in question represents a smaller part of the flexible language's total problem space.

Another way of looking at the limitation of toolkits is that doing anything new with a set of tools means plugging them together in ways which have not been done before, and verifying that a new combination works is inherently a difficult problem (NP-hard, unless one believes that correctness conditions for typical assemblies are simpler than the boolean satisfiability problem). Consequently, assembling from standard parts may easily turn out to be more work than writing custom components from scratch.

Standard toolkits achieve flexibility at the expense of being less efficient than a special-case solution for any given problem, both in terms of the resource use of the final product and in terms of the complexity and therefore maintenance problems of the implementation. This is not to say that toolkits are not useful! If the problem space the toolkit was designed for is centered fairly closely and narrowly around the problem one is trying to solve, then the start-up costs of a special-case solution will be much greater than the cost of the extra inefficiency and complexity of a solution built from the toolkit. Toolkits are useful short-cuts for common problems. But to the extent that we need solutions to new or uncommon problems, or a minimum of inefficiency or maintenance problems, toolkits can not solve the software crisis. On the contrary, the software crisis (which we should really call a design crisis) is likely to increasingly extend itself into other branches of engineering as construction techniques become more sophisticated (i.e., more limited by information and by the complexity of this information).

(The same flexibility-efficiency trade-off limits the scalability of software across an order-of-magnitude or more in hardware performance. Each order-of-magnitude improvement in hardware creates new potential solutions and therefore a new environment. Software sufficiently flexible to work at multiple orders-of-magnitude will be less useful at any particular level of hardware performance than software specialized to that particular level of performance.)

Since wealth is tied to solutions to limiting problems, and since solutions to the complexity limit take the form of information, the complexity limit is one of the factors driving the next trend, the relative growth of the information economy (IE) relative to the physical economy (PE). It is perhaps this which has allowed Japan to have the second-largest economy in the world, despite an absence of material resources.

2.5 Continued Relative Growth of the Information Economy

This section discusses the growth of the IE, both absolutely and in comparison to the PE. According to Katz, information services have grown relative to the rest of the economy since 1860, with a sharp acceleration after 1940 [KATZ]. For reasons given below, the information economy will be centered on the Internet. We can think of the IE as having two parts: the telecommunications infrastructure ("medium") and information goods and services ("content"). By content we include everything whose value is primarily intangible, rather than physical. This includes such things as advertising, banking, brokering, consulting, entertainment, insurance and software. Computer hardware could also be added, in that it is only of value to the extent that it supports the above. The economics of information medium and content are discussed separately below.

2.5.1 The Growth of the Telecommunication Infrastructure

One source puts the worth of the U.S. telecommunications industry at $700 billion, one-sixth of the U.S. economy [LW963a]. Looking simply at the growth in connectivity, it seems clear that the part of the economy related to telecommunication infrastructure will grow rapidly. Consider the following.

° From 1994 to 1995, the number of U.S. networks on the Internet increased by 350% to more than 27,000, while the number of non-U.S. networks on the Internet grew by 225% to more than 21,000 [BLUE96A].

° The number of commercial domains on the Internet grew by 61% in the first quarter of 1996, to 276,000 [WEBCOM].

° Starting from zero in 1993, there were 27,000 World-Wide Web (WWW) sites as of May 1996, a number which is doubling every 50 days [LW963b].

° The number of U.S. households with Internet access is expected to triple to 35 million (about a third of the total) by 2000 [R960429].

° According to Matrix Information and Directory Services (MIDS), the number of people on the Internet has been growing at about 100% per year since 1988 [MIDS9510]. As of October, 1995, MIDS estimated the Internet population as consisting of 39 million users of mail and news; 26.4 million users who can access information via interactive TCP/IP services such as Telnet or WWW; and 16.9 million users who can provide information via TCP/IP services. Estimates for the number of people who will have access to the Internet by the end of the century vary between 170 million and 700 million [BOXX].

These trends are reflected in the sales growth of Cisco, the leading computer networking company, by 93% from the corresponding quarter of 1995 to 1996 [NEWS051196].

All while the U.S. economy as a whole is growing at 5% or less in real terms. And all prior to the widespread use of secure mechanisms for monetary transactions on the Internet.

2.5.2 The Growth of Information-Related Goods and Services

The previous subsection covered strictly the medium of the IE. If one includes content, the growth of the IE, both absolutely and relative to the PE, becomes even more dramatic.

The Computer & Business Equipment Manufacturer's Association (now the Information Technologies Industry Council) published in 1992 statistics on the historical and projected growth of the information industry [CBEMA]. While this source is now somewhat dated, it suffices to show the strength and stability of the growth trends. The U.S. information industry has grown faster than the economy as a whole every year since 1971 except for 1976. The average annual growth for the information industry from 1971-1981 was 13.7%, compared to 10.7% for the economy as a whole. From 1981-1991, the comparable numbers were 9.0% and 6.5%. The CBEMA estimated the information industry at $542 billion in 1991, and projected it to $1,273 billion in 2002.

Looking simply at software goods and services (SGS), the purest of information technologies, the average annual growth rate from 1971-1981 was 21.5%, and 17.7% from 1981-1991. SGS stood at $107 billion, or 5% of the U.S. GNP by 1991 [CBEMA]. The estimated annual growth for 1992-2002 was 12.7%, to $401 billion. Looking at recent computer hardware sales in the U.S., the $125 billion dollar industry is currently growing at 19% per year, after growing 25% the previous year [NEWS051396]. Manufacturing automation is part of the trend towards the relative growth of the IE over the PE. Automation has the effect both of making physical objects cheaper and of requiring an investment in the information needed to design and operate the new manufacturing processes.

In addition to the IE consistently outgrowing the PE, another important trend is that the IE is increasingly centered on the Internet. In an interesting article on software trends, McNamee [MCNAM] points out that the value of the traditional mainstays of the software industry, word processors and spreadsheets, have been dropping steadily. Word processors began as special-purpose units valued at $7,000-$10,000. By the early 1990's, they had become part of $100-$150 application packages. Meanwhile, the growth of Internet connectivity has created a need for an entirely new set of tools, centered around communication.

The Internet has many advantages over other means of information exchange, such as print, telephone, radio or television. These advantages (which include speed, price, universal accessibility, search capabilities and two-way communication) seem sufficiently compelling as to suggest that the vast majority of information-related economic activity in the developed world will migrate to the Internet. For instance, for the price of a single full-page ad in a newspaper or trade journal, a business can create a Web page which will reach a worldwide audience for a year [BUSIQ1]. Furthermore, people with an interest in the business's service will be able to quickly find this page from anywhere in the world by using net search engines.

It should also be pointed out that the Internet is swallowing up other modes of communication. Not only books, but also magazines are now available online (e.g., [PENNWELL]). Static and moving pictures are already distributable over the Internet. Internet audio links will be provided by Netscape as of the autumn of 1996 [AUDIO]. ("There could be 10 million people out there with this capability within a year or two," Andreessen said. "The implications for the telephony business are interesting.")

In addition to immigrant companies using the Internet as a new market for pre-existing products, there are also a growing set of "native" companies which could only exist on the Internet. These include companies vending browers, encryption, search engines, Internet advertising, digital cash, Internet consulting services, news filters, databases, etc.

In the early days of the computer industry, most of the value was in hardware, not software. This has since reversed. It is suggested that the same will hold generally for the physical versus information economies. Which implies that most of the world's wealth will take the form of information stored and traded on the Internet.

Despite the high growth rate, for the short term the Internet will be a comparatively modest percentage of the complete economy. Hambrecht & Quist, an investment banker for emerging growth companies, estimates that the Internet market will grow from $117 million in 1995 to $23.2 billion in 2000 [BOXX]. Zona Research's projection is from $5.3 billion in 1995 to $41 billion in 1999. Forrester Research estimates that online commerce will account for $31.6 billion by the year 2000, with $6.9 billion from retail sales. In addition, $15 billion of the $300 billion electronic data interchange (EDI) trades will move to the Internet by the millennium [ROSE]. A list of Internet-related market segments and companies is provided by Global Villager [VIL].

2.6 Universally-Available Unbreakable Encryption

Cryptology is the science of the secure transmission and storage of information by transforming it into an encrypted form that cannot be read by third parties [ASCOM]. While formerly used primarily by the military or diplomatic services, the growth in the amount of information transmitted over potentially insecure telecommunications channels, or stored in potentially insecure computers, has increased the necessity of encryption both for individuals and for businesses. Simultaneously, for the first time in history, the new computational power has made fast encryption available to businesses and individuals which it is impossible for competitors or governments to break.

A background on encryption in general, and the PGP algorithm in particular, is available from [EFH]. The idea behind classical single-key encryption is to transform information using a key, in such a way that the transformation can only be reversed with a knowledge of that key. Making the key longer does not greatly increase the time to encrypt a message, whereas it increases exponentially the time to guess the key. Therefore by picking sufficiently long keys one can make guessing the key arbitrarily difficult, beyond the capabilities of any third party. The drawback to single-key encryption is that both parties have to have a copy of the key, which poses the problem of how to communicate it securely, particularly if one wants to communicate on short notice with persons one has not previously met.

An interesting development in the 1970's was public key cryptography (PKC), due to Whitfield Diffie and Martin Helman [DH]. In PKC, one distributes a public key which allows a message to be encrypted, but a private key is kept secure. Only the owner of the private key can decode messages encrypted with the public key. PKC avoids the problem of securely communicating keys. PKC also solves the authentication problem. The owner of a private key can prove that a message came from him or her by encrypting a message with the private key, which can be decrypted by the corresponding public key. PKC algorithms are widely available.

For those concerned with traffic analysis, PKC can also be used to make the source and destination of a transaction secure from those watching messages going by on the Internet [FRIEDMAN].

One method is to send a message through an anonymous mailing service, with the destination address encrypted using the mailing service's public key. Anyone watching the traffic from the mailing service will be able to find all the sources and all the destinations, but will have no way of connecting a given source to a given destination. If the integrity of a given mailing service is in question, any number of them in sequence (possibly in different countries) can be used. If any is secure, or not cooperating with the rest, the source and destination of the transaction are secure. If the source and destination have only the version of each other's address encrypted with the mailing service's public key, they can perform transactions without knowing each other's real identities. A more sophisticated scheme, which does not require trusting a mailing service, is also possible [CHAUMDC].

PKC suffers from two drawbacks. The first is that it is inefficient compared to single-key encryption, since PKC requires complex operations, such as the multiplication of large numbers. This problem can be finessed by using PKC to transmit a regular key created on the fly for a particular communication, then using single key encryption. The PGP algorithm uses this hybrid approach [EFH].

The second and more fundamental problem with PKC is that, in the long run, it may not be secure. Beneath PKC is the belief that there are invertible mathematical operations, "trapdoor functions", which are much harder to do in one direction than the other. Factoring numbers, for instance, is at least superficially an exponentially more difficult problem than multiplying primes together. By using longer and longer primes, one does not increase substantially the time to multiply the primes together, but one appears to make it arbitrarily difficult to factor the resulting product. It is an open question, however, whether trapdoor functions really exist. In 1994, Shor demonstrated that numbers can be factored efficiently by a computations done at the quantum level [BRAUN]. While quantum computers are not immediately feasible, they are not inherently impossible, either. Fundamentally, the existence of trapdoor functions appears to be incompatible with the linearity of quantum mechanics.

What quantum mechanics takes away with one hand it gives with the other. Techniques for securely distributing keys using quantum uncertainly have been proposed [BRASSARD]. However, it is far from clear that such techniques will be sufficiently feasible and inexpensive for routine use.

PKC removes the difficulty of communicating a key securely. But even if PKC is broken, it will only reduce the convenience of secure cryptography, not make it unavailable to the general public. A preliminary step will be necessary in which keys are exchanged. Using conventional cryptography, one needs a conventional key (or one-time pad, another approach [RANUM]) for each trusted entity one wants to directly communicate with, and it is desirable to have physical contact to get the first key. PKC lets one bootstrap without that first physical contact, which is nice but not critical: there's nothing unreasonable about visiting the local offices of two or three major local network providers to download keys from them. If one trusts them, one can then download more keys via them in future. Given a few dozen hubs that most people have accounts with, and which are interlinked, any pair of people can communicate privately by both dialing into their respective hubs and then using the lines between them, which would likewise be encrypted. Each pair of hubs can physically exchange keys once per week, and shift to a new key every hour (or minute or millisecond) for the logical link joining them. Any pair of people who don't trust any of the hubs can physically meet and exchange their own cryptographic keys, and add their own encryption to that used by by hubs.

The benefits of secure data transmission to individuals and organizations are obvious. The dangers of a situation in which the government can not monitor information, even in the case of clear wrongdoing, are summarized by Denning [DENNING].

"Crypto anarchy" makes many types of crime very difficult to trace, including tax evasion, money laundering, political and corporate espionage and contract murder. These dangers are not lost on governments. Encryption is illegal in some countries, including Iran and Iraq [PGPFAQ]. The U.S. restricts the export of encryption software, which is classified as munitions. It attempted to make standard the Clipper chip (based on an undisclosed encryption algorithm), with two federal organizations keeping copies of the separate keys needed to decrypt any message encrypted by a given chip.

The attempt by the U.S. or any other government to restrict strong encryption is very unlikely to succeed for two reasons. The first is the opposition of the business community, which is faced with restrictions not imposed on international competitors [TESEC].

The second is the inherent difficulty of preventing people from using the strong encryption algorithms (if necessary, in addition to the Clipper or similar chips) which are easily available and easily implemented. Even if a government could somehow scan all Internet traffic and impose fines for cases of strong encryption, and could survive the political backlash, it is not difficult to hide messages so well that it is impossible to tell that a message was sent. A good encryption algorithm removes all obvious patterns from the message, which leaves the encrypted message looking like noise. One could, for instance, encode a message into the least-significant (and most noisy) bits of the images which are already widely distributed over the Internet. The message would not noticeably alter the image. Thus, one could not even prove that a message was sent at all, much less that restricted encryption methods were used.

2.7 Secure Online Monetary Transactions Will Be Available

The economic development of the Internet has been limited by the difficulty of making secure payments. Both strong economic pressure and technical feasibility ensure that this situation will soon end. In Feb. 1996, MasterCard [MC, MEP] announced that it had developed a joint standard with Visa for secure credit card transactions on the Internet. The standard is backed by GTE, IBM, Microsoft, Netscape Communications Corp., SAIC, Terisa Systems and Verisign. America Online, the largest online service in the U.S., recently announced an agreement with other providers to bring secure electronic shopping to the Internet [AOLIC].

Aside from credit cards, another approach to online transactions is digital cash, in which a given number refers to a fixed amount of money. Such a scheme is available, for instance, from NetCash [NETC]. In addition, DigiCash has aroused considerable interest in its implementation of digital cash, in which customers are able to make electronic purchases while retaining their anonymity both from the bank and the seller [DC, CHAUMEP]. This is built on top of PKC.

As with strong encryption, there are both benefits and dangers to anonymous and untraceable electronic purchases. The issues are discussed by Rose [ROSE]. Like strong encryption, it appears inevitable that anonymous purchases, at least for intangible items, will be easily available regardless of governmental attempts at regulation. The reason is that such purchases can always be made in some international jurisdiction where anonymous cash is allowed, or else effective anonymity can be achieved by indirecting through one or more discreet international banks. And of course, if anonymous cash is made available for small purchases it is easily available for large purchases, simply by breaking the large purchase into small payments.

Whether anonymous or not, digital cash has an important security advantage over credit card transactions, which is that the amount of money which the seller can extract is inherently limited. We may soon enter an age in which individuals often make thousands or millions of purchases in a day, most for fractions of a cent. For instance, one might have a program searching the Web for information and arranging it into some structure. The sites being searched might charge a small amount of money per query. There are obvious dangers to giving out a credit card number that many times.

A special issue of the Journal of Computer-Mediated Communication has been devoted to electronic commerce [JCMC13]. Crede [CREDE] makes an interesting comparison between the development of the medieval fair in 11th century Europe and Internet commerce today. Both introduce(d) new modes of economic interaction which require(d) extensive institutional changes. Many current banking practices date to the 11^{th} century.

What should be the economic value underpinning a currency? Money began from barter, in which the “currency” was specific goods. Coins made of precious metals such as gold developed as a means of referring to goods: so much gold was worth so much merchandise. The world economy at present operates on an interlinked set of fiat currencies, whose values are dependent on the actions of the governments backing the currencies. Aside from the problematical aspects of having wealth tied to the good will and responsibility of governments, the present system introduces economic uncertainty through the fluctuations between national currencies. Uncertainty always reduces efficiency (because it requires additional flexibility). There are therefore market pressures in the direction of developing an international currency not under the direct control of governments. An interesting possibility is that a de facto currency will arise based on an average value of a set of stock markets, possibly factoring in corporate debt. This can be thought of as a generalization of the gold standard: the gold standard is based on the value of a single commodity on the stock market.

A general resource on electronic banking issues is [EBRC]. A reference on performing small electronic transactions securely and cost-effectively is Manasse [MANA].

2.8 The Cost of Censoring the Internet Will Be Prohibitive

"King Canute had as much success commanding the tides to retreat as a national government will have regulating cyberspace."
[WISE]

"The Net treats censorship as damage and routes around it."
[GIL]

As this is written in May 1996, the Paris headquarters of the Internet providers FranceNet and WorldNet have recently been raided for providing access to Usenet (an international, participatory public news forum which, along with everything else, contains pornography) [NEWSFR]. The U.S. government has recently passed the Telecommunication Reform Act, attempting to censor the distribution of indecent material on the Internet. The Act, currently being challenged on Constitutional grounds in U.S. courts, has aroused concern in the Internet community that governments will attempt to regulate the free flow of thought on the Internet, as governments sometimes do for other media. One indicator of this concern is Barlow’s widely-cited (and sited) “Cyberspace Independence Declaration” [CID]. Opposition to governmental encroachments on the liberties of the Internet community sparked the founding of the Electronic Frontier Foundation in 1990 [EFF].

Despite the current struggle, effective attempts to regulate content on the Internet appear foredoomed by their difficulty and by the economic dislocation they would engender if successful. Even if one country attempts to limit content, it can easily move to another country which is just as easy to access from anywhere in the world. Canada discovered this when Justice Kovacs attempted to limit coverage of the Homolka case in 1993 [WISE]. In addition, encryption allows information to be transmitted between two or more people without being intercepted. Encryption of source and destination addresses, as described above, adds to the regulatory headaches. The only way to regulate information flow on the Internet would be to severely limit access. But as the world economy comes to be centered on the Internet, limiting access can only be done at the cost of considerable damage to the economy and therefore the government attempting the regulation. A “North Korea” model of control is possible, but only at the expense of rendering the country which chooses it economically marginal. In short, governmental attempts to censor the Internet are insignificant because governments which succeed in censoring the Internet will render themselves insignificant.

This is not to say that individuals will not be able to control the information they receive. As will be discussed below, various types of information filters are likely to be a major industry on the Internet. What is not economically viable is to limit is the ability of two or more persons to voluntarily communicate information.

2.9 The End of Copyright Law

Property is a monopoly over the use of some resource. In the case of material resources, the monopoly develops fairly naturally: the resource has well-defined boundaries and access to it can be controlled. By controlling access, those who invest in the improvement of the resource receive a return from its increased value.

Information does not have well-defined boundaries: consequently, the link between investment and return is less inherent. Intellectual property right law, which was developed to re-create this link, takes four forms: trade secrets, trademark, patent and copyright. The last two have become increasingly problematical in the information age.

Patent law was intended to protect rare, difficult and distinct insights for a time which was fairly short compared to the pace of innovation (less than 20 years). This is a poor match for the information industries, which boast continual innovation, much of it straight-forward, in which ideas blend seamlessly into one another, and in which even a couple of years is a very long time. Copyright law was intended to protect the fixing of ideas into a particular form. It has become more difficult to enforce as copying has become easier. The combination of easier copying and the growth of the IE will result in increasing tension in the short-term, of which the present tussle between China and the U.S. is one example.

A technological fix might be maintainable for software. Agreements between software and hardware manufacturers would allow software to check for serial numbers indicating that the software was running on a particular machine, for which the software was purchased. For other forms of information, however (such as music, films and books), with which observing the information is sufficient to extract its essence, it is hard to see how the current copyright system can survive the Internet. Countries which attempt to enforce copyright law will have even less success than they have had enforcing drug laws, since drugs (unlike information) have a physical representation which is detectable. Nations attempting to enforce copyright law will impose a cost on their economies which is not borne by nations with more lax enforcement. This will tend to erode not only the fact, but also the form of copyright law.

Even for software, assuming hardware serial numbers, the future of copyright as a basis for wealth appears gloomy. As will be discussed below, one of the features of the Internet is that it allows talent to be assembled quickly and efficiently from around the world. This allows volunteer efforts, producing freely-available software, to match the personnel, talent and resources of major commercial efforts. Programs such as Linux and GNU Emacs are as good or better than anything commercially available under copyright law. Programmers have an incentive to contribute to free software projects both from a sense of community and because it gives them access to software tools which are otherwise unavailable to individuals. As the amount of talent on the Internet grows, it can only favor the ability of free software projects to out-perform commercial operations on crucial software problems of interest to many volunteers.

Attempts to control standards, or to copyright or patent the “look and feel” of interfaces, is in a slightly different category from the above. As with music, the essential information is easily available to copy. But unlike music copied for private use, standards and interfaces tend to be only useful if made generally available, which makes the legal violation (if there is one) detectable. Even here, however, copyright law is unlikely to hold up. Firstly, customers have developed a distaste for proprietary standards, which lock them into one or a small group of vendors, reducing competition and therefore price/performance, and sometimes leaving the customer with orphaned technology if the vendor goes out of business. Secondly, vendors themselves are wary of betting on an attempt to establish a proprietary standard, since if they lose the attempt will have put them behind their competitors. These factors played a role, for instance, in the development of the Posix standard for Unix. Thirdly, software is very malable, meaning that it can easily be configured so that it does not ship in violation of any “look and feel” laws, but can be configured by end-users so that it does. Fourthly, software which violates “look and feel” laws in one country can easily be downloaded from another country.

Similar problems apply to software patents. Patents related to material goods have a better chance of holding up, to the extent that they refer to more easily traceable physical objects and are less ambiguous than claims to malable software.

The end of copyright law makes general economic sense. In bulk, the cost of an item should be related to the marginal cost of making one more of the item. To charge more than this limits its distribution, which limits the wealth of the society as a whole (as opposed to simply the vendor). For information, the marginal cost happens to be essentially zero.

Wealth flows to solutions to limiting problems; copying was once a reasonable place to assign intellectual wealth because copying, prior to photocopiers and especially computers and networks, was a limiting problem. This is no longer true. If we are to predict how intellectual wealth will be assigned in the future, we should consider what the limiting problems will be. Here is a by no means perfect taxonomy. Firstly, solutions to specific problems. New problems will constantly arise. Solving them will require hiring the relevant skills. Secondly, immediate access to timely information, such as sporting events or the latest business news or commentary. Thirdly, access to unique sites, such as: online social or business communities; large or specialized databases; specialized display or computational environments such as entertainment centers; or the opportunity to be present at a live performance. Barlow mentions that the Grateful Dead did not charge for copies of their music from the seventies on, and benefited in terms of increased attention and therefore attendance at their concerts [BARLOW].

A possible model is that instead of individuals buying an information service, groups will buy joint, unlimited access to the service, which may then be specialized to the interests of that group. By charging a group rate, it becomes feasible to provide the service without enforcing a per copy fee. And by specializing the service to the group, there is less of a problem with the information being used by those who have not paid for it. Perhaps an early example of this is the recent agreement between CompuServe and Time Warner. CompuServe will pay Time Warner a licensing fee for its members to access Time Warner’s Pathfinder [IMWLF].

The Electronic Freedom Foundation has a resource page for online intellectual property issues [EFFIP].

2.10 Increasing Gap Between Rich and Poor

The industrial revolution greatly reduced the need for the employment of human strength, but created a more-than-compensatory demand for human intelligence to guide the new machinery. The information revolution greatly reduces the need for humans to guide machinery, without creating a corresponding number of new employment opportunities. The consequence is a sharp divide into a majority whose skills are of declining economic relevance, and a smaller group with the fortune or training to stay on top of the system. Those whose skills are relevant to the IE can reach a much wider market through the new technology, which allows information to be amplified. Their earnings rise proportionately. And because they can reach a wider market, they further decrease the need for general employment. The median real income in the U.S. has declined every year since 1989, combined with a sharp rise in wages and salaries earned at the very top of the income scale [EVANS].

Looking at this trend on a worldwide scale, it is worth pointing out that the last several hundred years were unusual in that the growth of the need for labor outpaced the growth of the population. In the 19th century, the GNP of countries such as the U.S., Canada and Australia benefited from immigration, which served to fill available farmland and manufacturing positions. Today, according to the WorldWatch Institute, the climbing worldwide refugee population is equal to that of Canada [WWI]. There appears to be no country which believes that

accepting unskilled refugees will aid that country economically, although refugees are still accepted on humanitarian grounds. The situation is made more critical by the declining grain production per person, increasing grain prices and the continuing growth in the population, currently at the rate of 87 million people per year [WWI]. On a worldwide scale, therefore, there is an increasing divide between a relatively small group which benefits from the new economy and a majority tending towards the abyss.

2.11 The Defense Will Hold

"At every level of life one trades mobility for security,
or in reverse, immobility for adventure."
[MUM5]

The U.S. Computer Emergency Response Team (CERT) was established in 1988 [CERTPR] following the release of the Internet worm which shut down much of the Internet for several days [WORM]. In 1994, CERT issued 15 advisories and was contacted more than 33,000 times to request security-related information affecting more than 40,000 sites [BLUE96B]. A recent penetration study of 8932 computers by the Defense Information Systems Agency showed that 88% of the computers could be successfully attacked [DENNING]. (This depends, of course, on the definition of a successful attack. For instance, virtually all computers are vulnerable to having their mailboxes overloaded.) The security problems on the Internet can reasonably be expected to become worse before they get better, given the evolution of new tools with unforeseen holes, the quickly growing Internet population not always well-versed in security measures, and the increasing amount of economic activity and therefore tempting targets on the Internet. Electronic assault is a common topic in futuristic fiction.

Nevertheless, it is suggested that security is not a serious long-term problem for the Internet. The reason is that absolute safety is available by disconnecting, and beyond that people can choose the convenience/safety trade-off which suits their needs. The software industry is useful, despite our inability to completely remove bugs, because the serious problems can be fixed and the minor problems can be lived with. In the same way, the serious security holes will be patched and the minor ones are survivable.

2.12. No New Physics Which Contradicts Existing Physics

While it is common to speak of revolutions in 20th century physics, new theories have clarified the boundaries of previous knowledge rather than overthrown it. Newtonian mechanics is still used to calculate comet and spacecraft trajectories. Newton's theory is corrected by quantum mechanics for very short distances, and by general relativity for high speeds and strong gravitational fields: but for the range of phenomena Newton was concerned with his theory remains an adequate description.

While we do not have a complete description of physics, we know enough that it has become prohibitively expensive to perform experiments to which we can not, in principle, predict the outcome ahead of time. (Witness the recent cancellation of the Superconducting Super Collider in the U.S. for budgetary reasons.)

It is to be expected that new formulations will arise which will extend current theories into physical domains which are not currently understood. But it is also to be expected that these formulations will not fundamentally contradict existing theories.

In particular, we assume that the speed-of-light will remain an upper bound on the speed of communication. This may turn out to be false, of course, but it is a quite reasonable assumption.

3.0 Consequences

The previous section was intended to document trends which can be studied fairly well in isolation. This section examines the interactions between these trends.

3.1 The Rise of Virtual Organizations

"Of them alone it may be said that they possess a thing almost as soon as they have begun to desire it, so quickly with them does action follow upon decision."
[THUC]

An organization is a set of people united for common purposes. To be unified they must communicate with each other. Traditionally, the speed and quality of communication has been heavily constrained by distance: hence a requirement for organizations to put all critical resources in a single location. The increasing availability of high-bandwidth communication; of cheap computation to process and encrypt the communication; and of sophisticated interfaces for teleconferencing is removing the geographical constraint on communication. Also, as the economy becomes increasingly information-based, the importance of transporting physical goods is reduced, which lessens another geographic constraint.

In addition to the reduction in the geographic constraint on the structure of organizations, non-geographic constraints are becoming stronger. These are primarily the increasing speed of events and the increasing sophistication of the response required. Both may require assembling specialized talent and resources from around the world, possibly on short notice, possibly for only a limited amount of time. To the extent that organizations tend to rely only on talent and resources available at one location they will pay an increasingly unacceptable cost in terms of the time, quality and expense of solutions produced. Staying local forces one to either have more infrastructure than one will usually need, or else to accept inferior solutions. In general, one should expect the size of organizations to decline, as it is not necessary to retain permanently what can be acquired efficiently as needed from outside.

The combination of weakening geographic constraints and strengthening non-geographic constraints points to a new type of organization, organized solely around the nature of particular problems, rather than around geographic similarity. “Virtual organizations” which have essentially no existence except on the Internet, assembling talent and resources as needed from diverse geographic locations. To the extent that information comes to dominate physical considerations, retaining a strong geographic basis for an organization imposes costs but adds few benefits. Virtual organizations therefore are not only possible but should quickly become the rule.

Aside from their value as a geographic bridge, virtual organizations also have a valuable role to play as information filters. Nothing is more valuable in the information age than a high signal-to-noise ratio, and nothing is more endangered by the ease of distributing messages on the Internet. By restricting access or posting privileges, virtual organizations can provide a much higher signal-to-noise ratio than the Internet as a whole, where “signal” is defined in terms of the interests an of the members. (Moderated newsgroups are an early form of this idea, as are the academic muds, which often restrict access [WILSON].)

The above is true of organizations generally. The rest of this section is devoted to economic organizations (corporations). Virtual corporations face difficult organizational problems for both resources and people. For informational resources, there is no fundamental organizational difficulty; organizing physical resources remotely is more interesting. This problem is receiving considerable attention in the manufacturing community. An indication of this attention is the Agile Infrastructure for Manufacturing Systems (AIMS) project in the U.S., which is intended to allow major manufacturing facilities at different sites and companies to be coordinated as needed over the Internet. “When fully developed, AIMS will provide access to a national network of agile manufacturing services that any company can utilize as seamless extensions of their own internal production capabilities.” [PARKEA]

Virtual corporations also face intriguing social organization problems. For instance, because virtual corporations can be created quickly and require few fixed assets, they have sufficiently low overhead that one can imagine creating a corporation for a specific project, then dissolving the corporation after the project is finished. Since these “instant companies” will incur less costs than companies which stay in existence between contracts, we can expect a high percentage of short-lived companies. This poses questions such as the following. How will these companies be called into existence? Presumably some form of online market will be necessary, in which job requirements are posted on the buyer side, qualifications are posted on the seller side, and either automated or human match-makers suggest the outlines of a possible company. Will buyers trust a company with no track record? An open question, but the track records of the technical and managerial talent may be sufficient. Will there be a pension plan, etc.? If so, who will administer it?

There are other questions about virtual organizations which will be dealt with in subsequent sections. For instance, Whose laws will govern an organization with no fixed address? Can the laws be enforced? Are the management principles for virtual organizations the same as for traditional organizations?

3.2 The Decline of the Geographic Governments

Nations (as with other organizations) have formed along geographic lines, for reasons of communication, transportation and consequent ease of attack and defense. One problem facing the geographic government will be the declining importance of regional geography.

As issues become increasingly non-geographic (such as the IE) or pan-geographic (such as environmental decay or pollution), the relevance of geographic political structures decreases. But this is not the main problem facing the geographic government. The main problem is the decline of its tax base, due to the internationalization of commerce and the ease of hiding resources.

The IE has no crucial link to geography. As the IE grows relative to the PE, the ease with which wealth can be transferred also increases. This reduces the ability of any geographic government to extract wealth from the economy. To take the extreme case, virtual corporations have no well-defined geographic existence at all: they can define themselves to be in whatever jurisdiction is most convenient, or to be in no jurisdiction. The IE is not only fluid, it can also make itself invisible. If an individual residing in one country works for a company based in a second, on a contract in a third, and deposits payments in a fourth country (if the money needs to be given a geographic location at all), and if all of the transactions are done using strong encryption, how, realistically, is any taxing agency going to find out about it? Particularly if the money is spent using anonymous transactions. Furthermore, the IE not only makes wealth more mobile, it also makes individuals more mobile. Those who make their living off the Internet can do so from anywhere where they can connect to the Internet: with satellite links, this means the entire globe. Individual mobility will tend not only to increase the number of tax refugees, but also to make it easier to spend wealth without tipping off tax agencies in one's home country. (The country where the wealth is being spent is unlikely to reduce its own cash flow by reporting expenditures to the spender's home country.) Even before the Internet takes hold, tax evasion is already significant. Currently in Brazil, it is estimated that 50% of taxes go unpaid [BRAZ].

Geographic governments can never become entirely irrelevant, both because the physical economy will never disappear and because problems such as physical crime, housing, sewage, roads, etc., are essentially regional. But the nature of their funding will change. Geographic governments can only tax what they can find: this would appear to mean sales taxes on physical goods and a poll tax. As argued above, those controlling most of the wealth are going to be very mobile: consequently, in order to retain their income, governments will have to compete on the quality of their services. If one thinks of a poll tax as a rent, one has the ironical prediction of geographic governments becoming essentially landlords.

Those who like the thought of geographic governments, from nation-states down to local government, reduced to landlords competing on quality of service should superimpose this trend on the centralization of wealth, which the inability to tax IE income will accelerate. Politics is a form of symbolic warfare, which saves the expense and uncertainty of real warfare. For as long as the political system (counting both the official form and unofficial corruption) produces results which reflect the relative power of the participants, no one can expect to gain by attempting to alter the system. But a political system can not survive if it does not reflect the underlying power distribution. The systematic centralization of wealth is not compatible with anything resembling a democracy of the entire population. (It is compatible with a democracy having citizenship limited to significant property owners, as was practiced by the ancient Greeks.) It therefore seems likely that a significant part of the population will be rendered increasingly marginal both economically and politically. Those not covered by the social contract have little reason to obey it. Which increases the tendency towards free-lance wealth redistribution. Which in turn causes the wealthy to more thoroughly separate themselves into geographically distinct areas, eventually with a separate political system.

Of course, it is possible that a more equitable distribution of wealth than projected here will arise, which would avoid considerable unpleasantness. Certainly, with a reduced population and roughly even distribution of wealth, the material basis for the good life would be generally available. The main limitation to human happiness is humans: but this is a significant limitation. Perhaps the best historical analogy to the early part of the information age is the early part of the modern age, when the towns arose in Western Europe and separated themselves from the feudal system. “For the burgesses the country population existed only to be exploited. Far from allowing it to enjoy their franchises, they always obstinately refused it all share in them. Nothing could be further removed from the spirit of modern democracy than the exclusiveness with which the medieval towns continued to defend their privileges, even, and indeed above all, in those periods when they were governed by the crafts.” [PIRE]

3.3 Government on the Internet

“The objections which have been brought against a standing army,
and they are many and weighty, and deserve to prevail,
may also at last be brought against a standing government...
the government itself, which is only the mode which the
people have chosen to execute their will, is equally liable to
be abused and perverted before the people can act through it.”
[THOR]

A government is a type of organization. By a virtual government would be meant an organization with no fixed geography which administers the activities of some group of people, particularly their activities on the Internet. The decline of the geographic governments raises the possibility that the vacuum will be filled by a corresponding emergence of virtual governments.

Any government, geographic or virtual, must have some means of influencing events in its jurisidiction: this is what it means to govern. Other classifications are possible, but it seems natural to view influence as arising in one of three ways: control over force; control over resources; or control over the creation and guidance of the dominant modes of thought. (A mnemonic might be military, money and mind.)

The three types of influence differ in terms of their relative local (or short-term) and global (or long-term) effectiveness. Force is the easiest to apply locally but the hardest to maintain globally. The simple application of force produces discontent which manifests itself wherever the force is absent. If we consider control over cultural or religious instruction as one example of thought influence, we have the opposite extreme from a monopoly over force. Cultural education may take an entire childhood to instill in each individual, but the result is very often an adult willing to die to defend a particular group. Control over resources falls between the other two categories.

Power, fungible as energy, may of course be bought in one coin and spent in another: for instance, an army can seize resources; resources can buy control over the means of communication or instruction, and thus eventually the culture; and the culture can inspire the raising of an army.

In considering what forms of virtual government are possible, if any, one must begin by examining what forms of influence are possible on the Internet. Each of the above three categories will be considered separately. Finally, the three individual threads will be brought back together with some thoughts about their implications.

3.3.1 The Control of Force

A local monopoly on the use of force is the most fundamental characteristic of geographic governments. (A cynic might define a geographic nation as merely an army which has sprouted a bureaucracy.) The same route, however, does not appear to be open to the formation of virtual governments. The Internet per se does not lend itself to physical coercion. One could imagine a virtual government contracting to a geographic government to enforce its rulings, just as the Spanish Inquisition turned its difficult cases over to the “secular arm”, were it not for the ease with which individuals can hide their physical location and place themselves in a geographic jurisdiction amenable to their interests.

One could generalize the definition of force somewhat to include electronic assault on information assets. But as argued previously, it is likely that sites will be able to defend themselves adequately against electronic assault.

3.3.2 The Control of Resources

The second method of achieving influence is through the control of resources. Corporations are proto-governments, although their scope of interest is usually more narrow than the geographic governments. Trade organizations which are able to maintain a near-monopoly for extended periods of time may become more influential than geographic governments: witness the Hanseatic League. Possible physical and informational resource monopolies are considered below.

For purposes of controlling activities on the Internet, a monopoly on a physical resource would have to be closely related to the Internet. There appear to be only two possible categories: communication hardware and computational hardware. A monopoly on communication hardware would imply the ability to control access to the Internet. Given the decentralized nature of the Internet, the ease of adding new connections and the ease of connecting from different locations, a communication monopoly is unlikely to develop.

A computational monopoly could imply not an ability to restrict the availability of computers (which is obviously unlikely), but rather control over a computational resource so vastly more powerful than what is generally available that access to it becomes an economic or social necessity. This is unlikely to occur because of the dominance of generic microprocessors. Because large computers will be built out of generic parts, the price/performance of cheap computers improves faster than for expensive computers. Consequently, a group of individuals, linked together, can match the computational power of large organizations. This fact has decentralizing implications, both economically and politically [MCCANN]. (In this light, it is interesting that Hewlett-Packard recently unplugged its last mainframe, to considerable fanfare [HPPLUG].)

As a virtual government would exist in the information bath of the Internet, information monopolies would be the most natural thing for it to exploit, were it not that information does not have the firm boundaries which lend themselves to monopolies. Possible monopolies could be based on either the form or the content of information.

A monopoly on form might be created by control of a standard, or the “look and feel” of an interface. This has been a source of considerable wealth for software companies in the past. For reasons discussed in conjunction with copyright law, however, this form of monopoly is unlikely to be sustainable in the future.

A monopoly on content might be built around information which is both crucial and expensive to generate. Perhaps the most important, because of its central role, would be a monopoly on the mechanisms for online monetary transactions. If online transactions are done using credit cards, vast resources might be necessary to carry loans between the time of purchase and the time of payment. One might argue that this will result in only a small number of online banks, with very considerable influence. However, the speed with which online transactions can be resolved reduces the need for deep pockets, as does digital cash. Furthermore, the ease of setting up shop on the Internet would allow smaller banks to enter the market, limiting their scope to relatively small transactions.

One content monopoly which is viable is the limitation of access to the information on a particular site. This is a purely local monopoly, of course, since each site can set its policies independently. The importance of this kind of monopoly depends on the

importance of the information at a particular site. Is it possible for a given site to have information so important, and so difficult to replicate, that access to it becomes a necessity? If the information could be easily generated by standard toolkits, it obviously can not be so difficult to replicate as to constitute an important monopoly. However, as was argued in conjunction with the complexity limit for engineering, there are tasks which can not be efficiently accomplished using simply standard toolkits. This implies that useful information can be generated which would be hard to duplicate. Does this constitute a potentially important monopoly?

The answer appears to be no. The reason that generic toolkits can not solve all problems effectively is that they are too flexible, which makes them less efficient and more complex than special-purpose solutions for particular problems. So improvements on easily-available solutions arise through specialization. But as something becomes specialized, it necessarily becomes easier to do without, by working around it, and therefore unsuitable as the basis for an important monopoly.

3.3.3 The Control of Thought

"The dictum that truth always triumphs over persecution is one of those pleasant falsehoods which men repeat after one another until they pass into commonplaces, but which all experience refutes."
[MILL]

Our culture molds our thoughts. This is apparent in the frequency with which people raised in a particular culture adopt the religion, politics and general outlook of that group, through many generations. This mallability of human thought serves its purpose, both in passing on useful truths and in maintaining group cohesion. The same mallability can and is used by governments to bolster particular political arrangements, for better or worse. Consequently, the ability to guide the dominant flow of thought is a potential basis for governmental power.

The Internet, which gives the power of the printing press to everyone, reduces the ability of any central group to control the flow of thought. This becomes particularly true if, as suggested below, there is a trend towards decentralized and ad hoc education. It was previously argued that the Internet is not censorable, in which case the ability to guide the flow of thought through distribution limitations will not be possible in the future. Of course, this does not imply that thought will not be influential in the future, only that centrally-applied distribution limitations will not be effective. Ideas can still thrive based on their inherent quality.

The widespread availability of information in no way guarantees that the truth, whatever it might be in a particular domain, will be immediately understood: there are always plausible alternatives. But as Mill points out, truth has the advantage that it will continue to arise [MILL], which means that without distribution limitations truth will tend to emerge over the noise through sheer persistence. (An interesting analogy comes from physics. There are only two forces significant for distances greater than atomic size, gravity and the electromagnetic force (EM). EM is much stronger: so much so that a simple magnet can counteract the gravitational force of the entire Earth. But the EM force comes in two signs, which tend to cancel each other out, whereas gravity is always additive. Consequently, through simple persistence, gravity becomes the dominant force on the galactic scale. An analogy is also possible in terms of electron drift velocities in a wire.)

3.3.4 Virtual Government

The above three sections suggested that the only coercive power available on the Internet is the ability to restrict access to particular sites, which is of only local importance, and that the only other form of long-term influence available is the potential for ideas develop a following through their inherent quality. This is certainly much less power than has been available to geographic governments in the past. Is it sufficient to form the basis for virtual governments? If so, what might such a government look like?

To address this question, one has to inquire what the useful purpose of government is, if any. While geographic governments could be created and maintained through force-of-arms, virtual governments can not. Virtual governments will only come into existence if they solve a legitimate problem which would cause people to voluntarily create and adhere to such an institution, and if the available means of influence are sufficient to make such a government effective. As an example of the limitations, virtual governments are unlikely to arise to interpret and enforce contracts on the Internet, because virtual governments have no more power to enforce contracts between arbitrary people on the Internet than do individuals. Consequently, two parties mutually wanting to settle a dispute are better off hiring a free-lance mediator. (This might work, for instance, by having each party deposit a sum equal to the maximum possible judgement prior to proceedings.)

In general, there is a conflict between local (or short-term) and global (or long-term) interests. To take the simplest example, theft is to the benefit of the thief, but allowing theft is not in the global interest of the society as a whole. Similarly, pollution and ecological destruction serve short-term interests of convenience, but the long-term costs are higher. It is an unfortunate fact that policies with the greatest long-term benefit are not in general the policies with the greatest short-term gain; and in fact, policies with good long-term benefits may not even be able to survive short-term pressures. Long-term investments, which make funds unavailable in the short-term, fall into this category. An example of local-global conflict from the banking industry is the keeping of reserves. Keeping reserves guarantees the long-term solvency of the banking system in the case of an economic downturn; but in the short-term, banks which do not keep adequate reserves gain a significant competitive advantage, because they can increase their return. Similarly, minimizing inventory is in the short-term interests of companies, but results in the economy as a whole becoming more vulnerable to a break-down in the transportation system.

It appears the useful function of government is precisely to mediate between local and global interests (or to internalize external costs, in economic terms). Good laws serve the function of aligning the private with the public interest, through reward or punishment. For the ancient Egyptians, law assembled the private resources needed to develop irrigation, at least theoretically for the general benefit. Today, legislation covering areas from violent crime to anti-trust to regulation of food production at least theoretically serves to make it in the interests of private groups to serve the public good. So whether governments will form on the Internet seems to depend on whether there are cases in which local and global interests will not be aligned, and in which the available means of influence are sufficient to produce alignment. There appear to be two cases in which this condition is satisfied: one is the creation of standards, and the other is the management of large projects.

As suggested above, in conjunction with copyright law, there is a desire by both purchasers and (to a lesser extent) vendors for open technology standards, which are generally accepted and which any company is free to use. Standards are an example of a local-global mediation, in that local interests in particular solutions are smoothed into a general agreement which is probably ideal for no one, but which serves the global community much better than a haze of competing implementations. It is interesting to note that Internet standards have been developing smoothly through international discussions, open to all interested parties, conducted through the Internet Engineering Task Force [IETF]. The resulting protocols are generally accepted, not because the IETF has any coercive power, but rather because the protocols represent the considered judgement of the best available talent and it is in the interests of the community that they be adopted. The IETF may be viewed as an early type of Internet government. It fulfills the governmental role of smoothing local and global interests, but does so in a way suited to the information age: through community consensus and the quality of ideas presented. ("We reject kings, presidents, and voting. We believe in rough consensus and running code."— IETF Credo [CREDO])

Perhaps more traditional is Internet government which arises through the need to manage large projects. To the extent that large projects can be built incrementally from available parts, management is not greatly needed. Indeed, it is reasonable to expect software programs to appear which are very large in the sense of referring to a huge body of code, but which are structured primarily as references into software modules scattered around the Internet. Such huge programs could be created on the fly and erased afterwards; virtual programs flitting into and out of existence. But as mentioned previously, it seems the flexibility-efficiency trade-off implies that there will be important projects which can not be built simply by a series of local increments to existing tools; a global design and substantial up-front effort will be necessary. What these global projects will be, precisely, is hard to predict, since these projects will be phrased in terms of the needs of a community which is only now coming into existence. Possibly creating particular types of virtual environments, for business or recreation; possibly the creation of specialized algorithms; perhaps something to do with mathematics. (Mathematics, which amounts to the formal study of patterns, can only become more important in the information age.)

Regardless of their precise nature, the existence of important tasks not solvable without a substantial up-front effort creates a local-global smoothing problem. Individuals efforts will have to be coordinated as part of a common design, and individuals will have to be paid prior to the generation of income from some joint resource. Assuming that the economic demand is sufficient to justify such efforts, how can they be managed in the absence of the ability to enforce contracts among participants? The fundamental answer seems to be that one has to make very good use of the one available coercive force on the Internet, the ability to restrict access. Restricting access has already been mentioned as a means to improve the signal-to-noise ratio. It can also be used as a means to slowly admit a new member into the inner workings of a group, allowing access only in proportion to demonstrated ability and trustworthiness. Gradual access can serve to keep an individual's required investment in a group (and the future reward for remaining in the group) greater than the benefits of defection.

As with the IETF example, it may seem strange to refer to the above as a governmental structure. It may appear more like a corporation. There are two answers to this. The first is that corporations really are small-scale governments; we do not usually think of them in this light because, in theory at least, they fall under the jurisdiction of geographic governments. In the case of online projects, there is no more powerful government which can impose on them; consequently, their governmental function becomes more interesting. Secondly, online projects differ from corporations in that the lack of coercive power forces them to function primarily on internal trust and goodwill: consequently, online projects are likely to consist not simply of professional collegues, but of friends whose lives become significantly intertwined. This combination of economic and social factors is more in keeping with what we usually think of as government, a point which will be returned to in the conclusion.

3.5 Education

Education is worth discussing separately because it is a limiting problem. The rate at which humans learn is by far the slowest process of interest to the information economy. Furthermore, with the declining economic importance of physical resources, and the quick obsolescence of information resources, developing intellectual skills may become the most stable capital investment one can make. It is likely that we will see dramatic changes in the educational system for two reasons: the reduced importance of geography and changing intellectual requirements.

Schools, as with other organizations, have been formed along geographic lines largely for reasons of communication bandwidth. While proximity is useful for education, restricting education to a single location imposes the usual costs. It is impossible to have locally all the resources which are available globally; consequently, education gains by making both information and instructors available over the Internet. To some extent, we already see this in the addition of networked computers to the traditional school format. But the logical extension, particularly for older students who can guide their own education more effectively, is that the geographical schools will tend to evaporate, possibly under the guise of "home schooling".

Under the current system, students who chose a particular school are restricted to course offerings from that school for a number of years. As it is difficult to switch schools, the educational institution has a limited monopoly over that student's educational choices, with the usual monopolistic problem of limited consumer control over quality and price. The Internet allows and encourages educational institutions, particularly those aimed at older students, to compete for students on a course-by-course basis. This will create the kind of competitive environment with which, abstractly, academics are enthused. A prototype is the Global Network Academy [GNA].

In addition to the loosening of the geographic stranglehold, the other factor diffusing through the educational system is the change in what students will need to be able to do. The intellectual challenge facing the workforce of the 21st century has been investigated by Hunt [HUNT]. The need for humans to do physical or perceptual labor is being replaced by a need to understand intellectually complex and changing tasks. "With a few exceptions, human performance skills that are based on sensory pattern recognition, followed by subsequent, specifiable motor actions, are no longer safe from automation. It's just a matter of time.... jobs are safe if the jobholder's task cannot be programmed." (pp. 239-40.) Essentially, humans can not compete with machines on efficiency, and must therefore do so on flexibility. Both the quickly changing environment and the programmability of well-defined tasks favor those who can learn quickly over those with long experience. This implies an approach to education which favors initiative and the art of self-learning over fixed knowledge and prescribed tasks. Again, this points to a decentralized educational system in which students have a much greater role in the course of their instruction.

A decentralized educational system in which students actively guide their own education through many different channels might appear to make it more difficult to evaluate student performance. One might acknowledge this point, but say that it is inevitable. As we come to require effective response to novel conditions as the hallmark of education, we lose the ability to evaluate through stereotyped tests. Instead, evaluation becomes a more general process of examining the projects students have chosen to undertake and with what success. Graduate education already works, for the most part, on this model.

To the extent that more specific measures are needed, the solution may be to completely separate instruction from evaluation: students who want to demonstrate competence in some area could approach an institution which does nothing except student evaluations. The separation between instruction and evaluation has the added benefit that students who so choose could acquire a degree in much less than the normal amount of time, which provides an incentive for exceptional students to excel. (The external evaluations need not be simply standardized tests, as is usually the case now. In principle, external agency can grade essays as easily as the current school system does, and probably with less variance.)

As mentioned above, intellectual skills may be the most durable asset available in the information age. As acquiring such skills is a limiting problem, we can expect wealth to flow towards the highly skilled. At least two other factors suggest a bright future for this group of people. One factor is the difficulty of enforcing contracts on the Internet, which forces employers to maintain good relations with workers. The second factor is that as systems become more complex, they become more susceptible to sabotage. Agriculture can be done with slaves, but such an approach is more dangerous with, say, software design. A disgruntled programmer can work inefficiently without being easily detected; or insert bugs, apparently accidentally, which will be very difficult to track down. Consequently, as systems become more complex the degree of voluntary and enthusiastic cooperation required for a project to be successful increases, which again forces employers to treat the highly skilled with considerable respect.

While the status of the intellectually skilled is predicted to rise, it was previously suggested that the average size and life-span of corporations will decrease. If we consider also that large projects were predicted to require stable organizations of people with strong social ties and a gradual admittance of new members which may resemble apprenticeship, the consequence may be the re-emergence of guilds as a dominant force in the world economy.

3.6 The End of Terrestrial Human Civilization

"For I know this thing well in my heart, and my mind knows it:
there will come a day when sacred Ilion shall perish, and Priam,
and the people of Priam of the strong ash spear."
[HOME]

"You fear, you say, that some crueler method of destruction
than that now used will be invented. No. If a fatal invention
were to be made, it would soon be outlawed by international law,
and unanimous agreement among nations would bury the discovery."
[MONT]

"Stop quoting laws to us. We carry swords."
[PLUT]

Powerful tools are indifferent to their application. Combustion engines power ambulances and tanks alike; missiles carry satellites and warheads alike; fission lights and levels cities with equal ease. The demand for better technology and improved health care leads to steadily more powerful tools. There comes a point at which the minimum set of tools made generally available by economic pressures exceeds the maximum set of tools compatible with survival. It would probably be agreed that our civilization could not survive millions of teenagers with their own nuclear weapons. But the coming situation may not be too far from this. Consider some of the possibilities.

Within ten years or so, we will have sequenced the entire human DNA and be moving towards an understanding of biology at the molecular level, an understanding which will be widely available and applied to tasks ranging from chip design to exotic material production to health care. It is apparently only due to evolutionary pressures that diseases which are deadly tend not to be highly contagious (with notable exceptions, such as the bubonic plague which wiped out a third of the European population in the 14th century). Deadly, contagious diseases tend to destroy their host population and therefore themselves. But in principle, it should be possible to engineer diseases deadly as AIDS and contagious as the common cold. It may even be possible to engineer personalized diseases which target individuals or groups with a particular genetic pattern. The availability of the tools for creating such diseases (particularly in conjunction with the predicted emergence of a large, dissatisfied population with limited economic and political rights) may lead, in the words of a traditional curse, to interesting times.

Nor need we limit ourselves to biological havoc. Consider the opportunities created by GPS, encryption and simple explosives. For about $1000 in parts, paid for with anonymous transactions, one could build a flying drone, powered perhaps by solar cells, guided by GPS and satellite messages from an encrypted source, and carrying explosives. Such could target any spot on Earth from any other spot, to an accuracy of centimeters. (The same effect could be achieved less colorfully by remote-controlled robots, which may become common for virtual tourism and telemanipulation tasks.) This approach suits itself to either selective assassination or a widespread attack against infrastructure.

For those preferring heavier ordnances, the increasing flexibility and sophistication of factories should not be neglected. Given the time it takes to build factories and the speed with which they become obsolete, there is a push towards robot factories which can be custom-built out of pre-made parts bought over the Internet [CMU]. It is not difficult to imagine the parts for such factories being bought anonymously and put together to assemble anything from missiles to tanks to small helicopters, all of which can be remotely controlled from an anonymous source. As manufacturing continues to become more sophisticated, it can only become easier to make such weapons.

These are only the most obvious destructive possibilities; doubtless there are many others. For instance, the heaviest ordnance awaits space industrialization, at which point it will be possible to drop asteroids on selected terrestrial targets. The form of civilization we currently have was based upon a limited capacity for destruction: a limitation at first due to technological primitiveness and later to the complexity and expense of the most sophisticated destructive tools. Technological advance will make the capability for mass destruction generally available, as it previously did for computation.

It is possible humanity will be able to dig itself into protective bunkers, or outrun the range of its weapons by spreading through space. But a more likely possibility, as discussed in the next section, is that humans will cease to play a significant role in civilization.

3.7 The Implications of Artificial Intelligence

"The thing thou ask'st is plunged so deep in the night
Of the eternal statute's dark abyss
It is cut off from all created sight.
And to the mortal world take thou back this:
Let men aim lower and no more presume
To encroach on these exalted mysteries."
[DANTE]

"Be not of those
Who think the doomsday is at hand, dear Sir,
Unless they see their path before their nose."
[GOETHE]

"It is possible to believe that all the past is but the beginning
of a beginning, and that all that is and has been is but the twilight
of the dawn. It is possible to believe that all that the human mind

has ever accomplished is but the dream before the awakening...
A day will come, one day in the unending succession
of days, when beings, beings who are now latent in our thoughts
and hidden in our loins, shall stand upon this earth as one stands
upon a footstool, and shall laugh and reach out their hands
amidst the stars."
[WELLS]

What has been done in carbon will be done in silicon. The eventual appearance of artificial intelligence (AI) at the human level is guaranteed by the fact that human-level intelligence has already been shown to be possible. Aside from variations on vitalism, which has a poor track record in science, there is little reason to believe that human-level AI is not possible. Indeed, once AI has reached the human level we can expect it to range far beyond, as it will not suffer from the design constraints which limit humans. AI's can add processing power and functionality as needed; need not forget or grow tired; need not be assembled only from inexpensive and widely available materials; and can learn instantly by copying information. (Humans invented the technique of augmenting slow-changing genetic knowledge with cultural knowledge, at the expense of a long period of education. AI's can have it both ways: quickly changing information which can be learned instantly.)

The question is not whether such beings will arrive, but how soon, and what their implications will be. That the study of AI has so far failed to produce anything resembling human intelligence is in no way surprising. Galileo did not discover the moons of Jupiter until he had a telescope, and we will not develop human-like intelligence at least until we have computers with a computational power comparable to the human brain. The problem is not simply that we do not have the power to run a solution to the AI problem; we are unlikely to find such a solution until we have machines with human-level computational power. Human intelligence is able to solve very general problems, and is therefore likely to involve very inefficient algorithms (by the flexibility-efficiency trade-off), whose properties could not be discovered before the necessary computational power is available.

The first question, therefore, is at what point we will have machines with computational power comparable to the human brain. Only a crude calculation is possible: one should not take the below overly seriously. If we assume the human brain has approximately 10^11 neurons, and each neuron can fire 10 times per second, this implies the brain can produce 10^12 bits of information per second. (Of course, the brain would probably overheat if all neurons fired at the maximum rate, but on the other hand one could argue that the useful computations are done by the synapses, of which there are about a thousand times more than neurons. 10^12 bits/second is probably about as good an estimate of human computational throughput as can be made at this point.)

Currently, machines are in the 100 MIPS range, which corresponds to 10^8 bits/second. One can argue either that an instruction produces 32 bits of results, hence should count as 10-100 neurons, or that it takes 10 instructions by the time one fetches operands and stores the results. Cancelling these two considerations out leaves us about 4 orders of magnitude short of the mammalian brain. It takes about 13 doubling times to cover 4 orders of magnitude. Taking Moore's Law at 2 years, this comes to about 25 years as the time until the typical desktop workstation will be in the class of the mammalian brain (say the year 2020). However, if we assume that leading-edge research supercomputers are 3 orders of magnitude faster than typical workstations, by the above assumptions a few research institutions should have mammalian-level computational power by the end of the century. But time on high-end machines is obviously expensive, and given the difficulty and uncertainty of the AI problem it is unlikely to receive a high priority. So we might pick 2010 as roughly the year in which sufficient computational power will be available to make AI research feasible.

Having machines with computational power comparable to the human brain is necessary to achieve AI, but not sufficient. Cosmology did not appear instantly after the appearance of the telescope. How long will it take after human-level machines become available for AI to unfold? It is impossible to give a precise answer to this question. On the one hand, no one appears to have a convincing description of what the architecture for intelligence might look like. On the other, knowledge of the human brain, which might lead to critical insights on intelligence, is rapidly increasing. (For instance, once we have a better understanding of the human genome, we may be in a position to build intelligences simply by following the instructions. The human genome contains only about 10^9 bits of information, which is not much by today's standards. If we assume that 90% of this is "junk" DNA which codes nothing, and only 10% of the remainder is related to the nervous system, we come to 10^7 bits to code the brain—one megabyte, fitting comfortably on a single floppy. For those who like to believe that intelligence is something which humans have but chimpanzees do not, if we take the figure that 98% of the DNA is common between humans and chimpanzees, then we knock the amount of information needed to code intelligence down by another two orders of magnitude to 10^5 bits, which, for better or worse, is less than this essay.)

The critical question is how difficult of a problem intelligence is. One might think it very hard to achieve if one considers the number of neurons and connections in the brain. But this is apparent complexity: it does not necessarily imply that the underlying design is complex. Indeed, one might argue that the apparent complexity exists precisely to make certain tasks simple, and that if one could express one's theories in terms of these central tasks the design would be seen to be quite simple. At least two points suggest this interpretation, although neither is compelling. The first is the amazing ability of the brain (particularly for young children) to recover from damage by switching functionality to different areas. This suggests that many parts of the brain are running essentially the same algorithm, which can be specialized as needed. The second is the recent suggestion that the cerebral and cerebellar cortices (in keeping with certain other organs) are built out of microscopic repeating units [PROT], which again suggests a general algorithm which is specialized as needed.

How might AI's compare to human intelligence? Again, only a loose calculation is possible, based on the design limitations for microprocessors. It seems probable that one does not gain significant performance beyond the point at which a signal can cross the microprocessor in one cycle-time. (The brain does plenty of nearest-neighbor computations, for instance in the retina, but it also runs many wires across the brain. At some level, the human computational speed must be constrained by the time it takes information to traverse the brain.) If we assume 10^9 components are necessary to achieve human-level performance, and we can pack them at about 1 nm intervals (about 10 hydrogen atoms wide) in a cube, then it is about 100 nm across the chip. Assuming that signals can propagate at .1 of light-speed through the chip, the maximum clock rate which allows a signal to cross the chip in one cycle-time is about 10^13 Hz. If we assume each component can produce a bit of information per cycle, this comes to 10^9 x 10^13 = 10^22 bits/second of information processing power, which we could round off to 1 billion times the rate of a human, or about a subjective human lifetime per second. (Of course, one could increase this number by adding components to the microprocessor, but one has to consider cooling problems. Perhaps chips will wind up being made of doped diamond?)

Humans would not be even vaguely intellectually competitive with AI's built along these lines. In a competitive environment, organizations which retained more than symbolic human control would quickly be out of business. Computers will not take over the world: humans will give it to them, through economic pressures. It is possible that there will be a transitional phase in which humans augment themselves to become AI's, but the net effect is the same: projections which assume humans will play a significant role in civilization indefinitely appear unrealistic. And why should it be otherwise? They are the children of our minds. Obsolescence is our destiny, not fate.

It is disorienting to visualize a world dominated by such beings. Creatures capable of thinking a billion times faster than us; with fewer of the vulnerabilities which make warfare dangerous; and able to transform themselves, including their personalities, endlessly. Nevertheless, they will not be gods, though they may forget this themselves.

Galileo has an interesting passage in which he alludes to whether it is possible for humans, with finite minds, to understand infinity [GAL]. We can not comprehend infinity, but we can analyze it by studying the properties by which it is constrained. In the same way, we can not comprehend these future beings, but we can examine the constraints which will govern them. The most crippling of these would seem to be the speed-of-light. In 1500, it took about 3 months or approximately 10^6 seconds for information to cross the Atlantic. What is it today? With the rule that light travels a foot per nanosecond, and taking the Atlantic to be 2 x 10^3 miles wide, with 5 x 10^3 feet per mile, we get 10^7 nanoseconds for a signal to cross the Atlantic, or 10^-2 seconds. In other words, in terms of communication time, the Atlantic grew 100 million times narrower between 1500 and 2000. But if we assume AI's will think roughly a billion times faster than us, then measured in terms of subjective experience the Atlantic will be wider in 2100 than it was in 1500. If the critical measure is the ratio of the speed of thought to the time for a signal to cross civilization, then we are currently living in the Golden Age of communication. It can never be repeated.

The geographic constraint therefore re-establishes itself for AI's. This may answer the question of whether there will be a single, worldwide AI or many. There will be many, because communication constraints prevent the degree of coordination which one would associate with a single individual. But that there will be distinct individuals does not compel them to be much different. We are accustomed to individuals having memories and personalities which change only slowly. This need not be true for AI's: they could change their memories and their personalities as often as desired. The entire civilization might chose to run essentially the same software, although it doesn't seem advisable. The degree to which individual differences persist, on the average, may depend on the relative time it takes to compute information as opposed to communicating it. If communication is very cheap then a few "best" solutions may dominate. This is difficult to project.

3.8 The End of Innovation

"The end of all our exploring will be to arrive at where we started, and to know the place for the first time."
-- T.S. Eliot

We have lived in a culture driven by exponential curves for so many generations that it is hard to imagine anything else. And certainly there is no finite limit to the number of ways in which matter can be rearranged, which implies that in principle innovation could continue forever. Nevertheless, one reaches a point of diminishing returns, in which an increasing amount of effort is needed for ever smaller gains. Mines are exhausted; continents are explored; why should the same not happen for engineering and science generally? Perhaps we are already seeing this in particle physics, in which the cost of new experimental information is growing beyond the point at which it affordable.

If we consider what the fundamental components of civilization have been, we might take them to be people, energy, transportation and weapons. (It was the weapons race which has most driven engineering and scientific innovation in the modern age, from the study of ballistics in the Renaissance to electronics in the 20^{th} century.) If we assume that civilization moves out into local and interstellar space, where the resources are, we might project the four components forward as follows: microprocessors; solar cells or fusion plants; solar sails, ion engines, or such; and various no doubt clever types of missiles. All of these technologies appear prone to innovative saturation. There are theoretical limits on how efficiently energy can be processed and materials shaped: as these limits are approached, what will drive innovation?

It is more difficult to say whether pure thought will saturate, since the same thermodynamic efficiency asymptotes may not apply. It has been suggested that AI's will spend their time on some fusion of mathematics, art and music. But even in our time, the arts often give the impression of passing through fashions rather than achieving deeper insights. Perhaps we will see science and engineering fall into an "Aristotelian" solution, in which accepted answers are simply looked up, and the arts fall into a "fashion" solution, in which changes, if not signifying nothing, signify nothing fundamentally new.

We may therefore return to the approximate innovative stasis which was the norm throughout all but the last five thousand years or so of intelligent life on this planet. By some accounts, it was (and may be) a more peaceful and pleasant age.

4.0 Conclusion

In the near term, geography becomes less important for two reasons: the ease of telecommunication and the relative growth of the information economy. The power of geographic governments will decline because the economy will increasingly be outside of their ability to tax. As a fixed geographic base imposes costs, there will tend to arise virtual organizations which exist only as connections on the Internet. Because the communication trend makes it easier to access resources outside a company, companies need less fixed resources; hence the average size of corporations will drop. And because better communication and fewer fixed resources make it convenient to create and dissolve companies, the average lifespan of corporations will also drop.

Simultaneously with the declining power of geographic governments and corporations, the increasing complexity of the information economy will add to the importance of the highly-skilled. Because of the difficulty of enforcing contracts on the Internet, together with the ease with which sophisticated projects can be sabotaged, projects can only succeed in the presence of a high degree of mutual respect and trust. These conditions suggest the evolution of groups of skilled workers which apprentice new members and combine social with professional functions. The consequence may be the re-emergence of powerful guilds. It seems that such guilds would necessarily be pulled together into larger informal associations by tightening connections in the information economy, which require close collaboration. Such associations could not be based on formal laws, which are unenforceable, but rather on consensus and craft, as the IETF is today.

If one considers the prospect of large numbers of people tied together at once economically, socially and politically, the implications are intriguing. It is not of new congregations we speak, but of new nations, of a kind whose properties we can only distantly perceive.

5.0 Acknowledgement

This essay has benefited greatly from the insight and assistance of an anonymous net diety.

6.0 References

[AOLIC] Reuters (1996). *AOL In Deals To Improve Internet Commerce.* May 15th.
http://www.yahoo.com/headlines/960515/compute/stories/aol_1.html

[ARIS] Aristotle. **Politics**, III, ix

[ASCOM] Ascom (1996). *IDEA: The International Data Encryption Algorithm from Ascom*
http://www.ascom.ch/Web/systec/security/idea.htm

[AUDIO] San Francisco Examiner, Nando.net (1996). *Telephone industry moves to ban Internet phone software.*. April 19th.
http://www2.nando.net/newsroom/ntn/info/041996/info5_10951.html

[AU94] Australian Bureau of Transport and Communications Economics Communications (1994). **Futures Project**, Chap. 2.
http://www.dca.gov.au/pubs/comms_project/ch2.htm

[BARLOW] John Perry Barlow. *Selling Wine Without Bottles: The Economy of Mind on the Global Net.*
http://www.eff.org/pub/Intellectual_property/idea_economy.article

[BLUE96A] HPCC (1996) *The Evolution of the Internet.*
http://www.hpcc.gov/blue96/section.2.1.html#EVOLUTION

[BLUE96B] HPCC (1996) *"Technologies for the Information Infrastructure.*
http://www.hpcc.gov/blue96/section.2.4.html

[BOXX] Glo-Boxx (1996). *"The Internet."*
http://www.he.net/~glo-boxx/internet.htm

[BRASSARD] Gilles Brassard (1993). *A Bibliography of Quantum Cryptography*
http://www.iro.umontreal.ca/people/crepeau/Biblio-QC.html

[BRAUN] Samuel L. Braunstein (1995). *Quantum Computation: A Tutorial.*
http://www.ph.unimelb.edu.au/~schmuel/comp/comp.html

[BRAZ] *Brazil's Tax Evasion Problem."* **Offshore** 5(21).
http://euro.net/innovation/Offshore/Offshore_Vol5_21.html

[BUSIQ1] MIWorld (1996). *Can My Business Increase Sales via the Internet?*
http://www.miworld.net/busiq1.html

[CBEMA] **Information Technology Industries Data Book** 1960-2002. Computer & Business Equipment Manufacturers Association, 1992. CBEMA is now the Information Technology Industries Council. Numbers cited from the 7th Annual Computer Industry Almanac, 1994-1995. Brady, New York:NY.

[CERTPR] CERT (1988). *DARPA Establishes Computer Emergency Response Team."*Dec. 13th. ftp://cert.org/pub/CERT_Press_Release_8812

[CHAUMDC] *The Dining Cryptographers Problem: Unconditional Sender and*

Recipient Untraceability. David Chaum in **Journal of Cryptology**, Vol. 1, No. 1, pages 65-75; 1988.

[CHAUMEP] David Chaum (1992). *Achieving Electronic Privacy*. **Scientific American**, Aug., p. 96-101. http://ganges.cs.tcd.ie/mepeirce/Project/Chaum/sciam.html

[CID] John Perry Barlow (1996). *Cyberspace Independence Declaration* http://www.eff.org/pub/Publications/John_Perry_Barlow/barlow_0296.declaration

[CMU] Dick Wilson (1996). *Tiny robots may revolutionize factories.* CNN, May 8th. http://www.cnn.com/TECH/9605/14/t_t/tabletop.factory/

[CONRAD] M. Conrad.