Self and community:

Turkle, Sherry; _The Second Self_, 1984

Turkle, Sherry; _Life on the Screen_, 1995

Ess, Charles; _Culture, Technology, Communication_, 2001

Shirky, Clay: _Here Comes Everybody: The Power of Organizing Without Organizations_, 2008

Preece, Jenny: _Online Communities_, 2000

Mossberger, Karen, et. al.: _Digital Citizenship_, 2008

Montgomery, Kathryn: _Generation Digital_, 2007

Cohoon, J. McGrath, et. al.: _Women and Information Technology_, 2006

American Behavioral Scientist v.45 n.3, 11/2001:
"The Internet in Everyday Life" issue

Privacy:

Agre, Philip, and Marc Rotenberg;
Technology and Privacy: The New Landscape, 1997

Diffie, Whitfield; Susan Landau: _Privacy on the Line_, 2007

Risks:

Neumann, Peter; _Computer-Related Risks_, 1995

Evan, William, and Mark Manion; _Minding the Machines_, 2002

Himma, Kenneth: _Internet Security_, 2007

Socially aware design:

Friedman, Batya; _Human Values and the Design of Computer Technology_, 1997

Johnson, Robert; _User-Centered Technology_, 1998

Berman, David: _Do Good Design_, 2009

Vinck, Dominique: _Everyday Engineering_, 2003

Hacking:

Himanen, Pekka; _The Hacker Ethic_, 2001

Haring, Kristen: _Ham Radio's Technical Culture_, 2007

Legal issues:

Cavazos, Edward, and Gavino Morin; _Cyberspace and the Law_, 1995

Godwin, Mike: _Cyber Rights_, 2003

Lessig, Lawrence: _Code_, 1999

Miller, Steven: _Civilizing Cyberspace_, 1996

Education:

Bromley, Hank, and Michael Apple; _Education/Technology/Power_, 1998

DiSessa, Andrea: _Changing Minds_, 2000

Margolis, Jane, et. al.: _Stuck in the Shallow End_, 2008

Klopfer, Eric: _Augmented Learning_, 2008

Reif, Fred: _Applying Cognitive Science to Education_, 2008

Work, War:

Rochlin, Gene; _Trapped in the Net_, 1997

Intellectual Property:

Stallman, Richard; _Free Software, Free Society_, 2002

Gillespie, Tarleton: _Wired Shut_, 2007

Feller, Joseph, et. al.: _Perspectives on Free and Open Software_, 2005

Games:

Sicart, Miguel: _The Ethics of Computer Games_, 2009

Bogost, Ian: _Persuasive Games_, 2007

Consalvo, Mia: _Cheating_, 2007

Miscellaneous:

Forester, Tom; _Computers in the Human Context_, 1989

Forester, Tom, and Perry Morrison; _Computer Ethics_, 2/e 1995

Brook, James, and Iain Boal; _Resisting the Virtual Life_, 1995

Borgmann, Albert; _Technology and the Character of Contemporary Life_, 1984

Winston, Morton, and Ralph Edelbach; _Society, Ethics, & Technology_, 2000

Johnson, Deborah; Jameson Wetmore: _Technology and Society_, 2009

Felon DNA Registration

The collection of DNA for use in government databases is a widely debated and controversial issue. This delicate balance of personal privacy verses public utility faces challenges in many areas, ranging though military, civilian, and criminal demographics. The last of these groups, convicted felons, is currently a topic of intense political and legislative argument. DNA is an excellent medium for identifying suspects, convicting criminals and even exonerating the innocent. Criminals may wear gloves, to hide their fingerprints, but it is difficult to commit most crimes without leaving behind DNA evidence. Saliva from a cigarette butt, sweat from a baseball cap, hair, and blood are just a few examples of evidence that will carry DNA information. Using DNA databases, therefore, is a largely successful means of solving crimes. There is, invariably, another side to this conflict, as DNA collection constitutes a possible threat to personal privacy. This paper is designed to both highlight the salient information of criminal databases, and to summarize the arguments for and against the collection of felon DNA. As much of the paper is an explanation and advocacy of current database practices, it is logical to begin with some criticisms of criminal DNA cataloging in general, so that the relevant details may be discussed.

Most of the objections to cataloging felon DNA information fall into two broad categories. The first is an argument for the protection of personal privacy. Many criminal defense lawyers and privacy advocates disagree with DNA indexing on the grounds that it is a breach of the Fourth Amendment right to be free of unreasonable search and seizures. Opponents of DNA collection demand that one has the right to be "secure in their person," and suggest that DNA collection violates this right. The second general argument against criminal genetic indexing is that collecting such information has nothing to do with the crime committed, but instead with the anticipation that a future crime may be perpetrated. This seems to be

inconsistent with the notion of criminal rehabilitation. Although far from an exhaustive list of criticisms, these two concerns would necessarily need to be considered in any successful criminal DNA database.

In order to address these two concerns with regard to current DNA indexing, it is first necessary to understand exactly how these databases are implemented and regulated. The basic California database consists of two DNA indices. The Convicted Offender Index contains DNA profiles from individuals who are convicted of crimes that warrant their inclusion (i.e. sex offenses and other violent crimes). The second is the Forensic Index which contains DNA profiles collected at crime scenes and constitute unsolved cases. When a profile in one index, matches another, a known felon is positively identified as a suspect in an unsolved case. This match is called a "cold hit" and generating these correlations is the primary function of the database.

The California DNA database is a subset of CODIS (combined DNA Index System) which is a national network of criminal DNA profiles. CODIS assumes a hierarchical structure beginning at the local level with the LDIS (Local DNA index system). This database is installed in local forensic laboratories that are operated by police departments, sheriff offices, or state police agencies. Profiles in this layer are entered and compared with each other and then passed on to the state level, the SDIS. The State DNA index system is operated by the agency responsible for the state's convicted offender statute, namely the California Department of Justice. As of July, 2001 the CAL-DNA database contained approximately 200,000 profiles. Bringing all the state databases together, is the National DNA index System, which is operated by the Federal Bureau of Investigation. The DNA Identification Act passed in 1994 formalized the FBI's jurisdiction in controlling the national database.

The most important details to keep in mind when summarizing the database structure are the actual contents stored in each DNA profile. Each profile that is entered into the CODIS system consists of: a specimen identifier, the sponsoring laboratory identifier, the name of the personnel responsible for the DNA analysis, and the actual DNA characteristics. The characteristics are not the full DNA content of the individual, but a tiny subset of the genome that allows for differentiation among others. From within the DNA is extracted thirteen Short Tandem Repeats loci that can positively identify one person from another. The probability of two randomly selected, unrelated individuals possessing an identical thirteen-locus DNA metric is one in 1.8×10^{15} for African Americans and one in 3.8×10^{14} for U.S. Caucasians. The CODIS profiles do not include criminal histories, case information, social security numbers, or personal physical profiles.

It is important to stress the fact that the current DNA database system does not contain the entirety of the individual's DNA information, but only a sufficient amount to differentiate between different profiles. This means that detailed hereditary and like studies cannot be performed on the samples in CODIS. Because of this restriction on the amount of DNA information that its profiles contain, CODIS is very similar to the AFIS (Automated Fingerprint Identification System). Both constitute identification tools only, and do not contain more private information that is protected in the Bill of Rights. This analogy is the key response to critics' contentions that criminal DNA collection is a violation of the Fourth Amendment, which reads:

> *The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no warrants shall issue, but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.*

The opponents of DNA analysis classify such collection as seizure and search of personal and privileged information. Advocates of CODIS contend that the database is not such a violation. In 1969 the Fourth Amendment was in question by the Supreme Court in the common law decision of Davis vs. Mississippi. Supreme Court Justice Brennan ruled that traditional fingerprinting in criminal investigation complies with the Fourth Amendment because it "involves none of the probing into an individual's personal life and thoughts that marks an interrogation or search. Furthermore, fingerprinting is an inherently more reliable and effective crime fighting tool " (Imwinkelried, 4). If the profiles in CODIS contain such a small amount of the individuals DNA, then the argument for fingerprinting is applicable to DNA indexing. The American Civil Liberties Union has a logical reservation with this defense. For this argument to be valid, the original biological sample must be destroyed. "It is one thing for the government to permanently store a genetic fingerprint; it is altogether different for the government to permanently retain the biological samples, which can be used for future genetic testing." (Steinhardt, Barry)

Thanks - this is a very useful clarification.

The Supreme Court also ruled that criminals can't expect the same measure of privacy as a non-offending citizen while they are under incarceration or disciplinary probation. The state is then free to reasonably intrude on their privacy by analyzing and storing their partial genetic profile because the individuals have proven themselves criminals. When a suspect is convicted of a criminal offense, his permanent identification becomes of great state interest in order to solve past and future crimes. In 1992, the U.S. Court of Appeals upheld the constitutionality of Virginia's (a pioneer in DNA profiling) legislation to include all felons in their DNA data bank program. It was ruled that the data bank "was neither a violation of an inmate's constitutional

protection against unreasonable search and seizures, nor did it violate the Constitution's ex post facto laws" (Hibbert, 768).

If the cataloging of criminal DNA is not a violation of the Fourth Amendment in itself, then perhaps it is an unreasonable assumption of the state that recidivism will be prevalent in criminal populations. Statistically, this is not an unreasonable assumption – in fact it proves to be quite necessary. Current California law dictates that violent criminals be added to the CAL-DNA database. It has omitted this necessity for many other felonies such as burglary. If convicted burglars, for example, are shown to escalate to more violent, qualifying crimes, the argument of regular recidivism would carry more clout. In fact, analysis of criminal records has determined that most violent crimes are preceded by a non-violent crime, such as burglary. A large number of sex and violent crimes solved by the Virginia DNA Data Bank program came from profiles of non-violent crimes in the past. Virginia's data bank program, as mentioned earlier, contains profiles for all felons. Sixty per cent of all crimes solved in the Virginia data bank correspond to criminals that were found guilty of a property crime. Fifty-two percent of all sex crime cases such as rape, sexual assault, indecent exposure, child molestation, etc, solved in Virginia as of December 2000, were perpetrated by individuals with prior non-violent convictions. Fifty six per cent of the cold hit cases would have gone unsolved, if Virginia had the same policy of limiting qualifying felonies, as dictated in the California DNA collection statute (Migden, AI). Virginia's data provides statistical support that many criminals begin their career with non-violent property crimes, and then escalate to greater and more violent offenses.

Although California has not included burglary in their DNA databases in the past, important research by the California Department of Justice upholds the contention that criminals often begin with non-violent crimes but escalate to more serious offenses. Recently the histories

→ What % of non-violent convicts later commit violent crimes?

of the CAL-DNA database felons and current burglary inmates were reviewed. A large number of violent and sex offenders currently listed in the CAL-DNA database have lesser crimes on their record. Forty-one per cent of the total samples have a burglary conviction in their criminal record, and more importantly, thirty-two per cent have a conviction prior to the offense that qualified them for inclusion in the CAL-DNA data bank. On top of the actual data bank profiles, inmates serving time in California Department of Corrections institutions also show a proclivity to repeated crime. Thirty-three percent of the burglar population have been convicted of a crime that would warrant their inclusion in the CAL-DNA database. Once again, a significant amount, twenty three point three per cent, of the total burglar population have a burglary conviction prior to the more serious qualifying crime. It would be prudent, therefore, to list such criminals in the DNA database as, statistically, they show a tendency towards repeated crime.

It is important to note that these statistics are not meant to imply that perpetrators of lesser crimes always escalate to violent crimes. Nor is the data presented to argue that criminals continue in illegal careers by their own volition. They simply speak to the fact that current penal and criminal populations often tend toward repeated crime, regardless of the reason – whether it individual, institutional, or governmental. On a probabilistic base putting such criminals into the database would prove advantageous to solving past and preventing future crimes.

There is a subtle detail involved in more general DNA indexing that naturally arises from criminal profiling. As police investigate all crimes, such as kidnapping, there is often groups of society that are urged or compelled to give fingerprints to AFIS. School teachers are required to submit fingerprints and institutions such as the National Center for Missing and Exploited Children urge parents to fingerprint their children. It is only logical to assume that DNA indexing starting in a criminal context would eventually come to this point as well. In fact, in

→ Of all the numbers you give this is the only one that helps your argument. And it's the smallest, of course. Fewer than 1/4 of convicted burglars later commit violent crimes. The other 3/4 would be in the database for no reason

→ Is 23.3% "probable cause"? What % of victims of child abuse grow up to be abusers? Should we file their DNA? I don't mean that as rhetorical Q's — for me this is a difficult decision.

1996, the Virginia General Assembly lowered the qualifying age for inclusion in their DNA database to 14. As of that date, any minor older than 14 that committed a specified felony would have their DNA loci taken. Furthermore, in 1998, New York Police Commissioner Howard Safir pushed for the power to collect DNA from every individual that is arrested (Associated Press, ACLU News 1998). This obviously blurs the line between convicted felons and non-convicted suspects. Barry Steinhardt, Associate Director of the ACLU, comments that "while DNA databases may be useful to identify criminals, I am skeptical that we will ward off the temptation to expand their use." As the debate over Criminal DNA proceeds, it is important to remember that other groups besides criminals would likely be directly influenced.

Although the general issue of DNA indexing still evokes great controversy and argument, the profiling of convicted felons is perhaps more universally acceptable. Undoubtedly, there are further objections to felon DNA profiling that have not been discussed. Perhaps the abuse of criminal DNA at the level of the forensic lab could be of concern. Because the lab and the personnel responsible for the analysis are listed in the CODIS profile, however, this type of misuse can be deterred and punished. Even an argument that DNA evidence can easily be used to implicate an innocent person, is no different than the same argument applied to fingerprinting. The fact is that advances in technology have made DNA analysis, which once took hundreds of dollars and weeks for each sample, into a cheap and efficient means to identify criminals. There are objections to such use, but there are also powerful answers to these concerns. Convicted felons do not have the same rights as non criminals, and the use of DNA indexing can drastically impact the conviction and prevention of crime. However, it is of paramount concern to carefully weigh the value of ever more prevalent DNA indexing with the vital issue of personal privacy.

Bibliography[1]

Alfano, William. "Should Burglars Be Included in the Cal-DNA Data Bank Program?" California Department of Justice, DNA Lab. (2001)

Associated Press. "NYC Police Want DNA Samples from All Arrestees." (http://www.aclu.org/news/w121598a.html).

Couffman, William. "Florida DNA Data Bank." Florida Department of Forensics. (2001)

Cornell Law School. "Bill of Rights." United States Constitution. (2001). (www.law.cornell.edu/constitution/constitution.billofrights.html#amendmentiv).

Deering's Penal Code: Including Penal Provisions of Other Codes. Lexis Publishing, San Francisco (2001): 167. California Penal Codes § 296

"DNA Typing in Action: Data basing in the Commonwealth of Virginia." Profiles in DNA Vol. 3, No. 1 (2000). (www.promega.com/profiles/301/301-03/)

Edwards, Bob. "The Use of DNA." NPR Online. 8 March 2001. (http://search.npr.org/cf/cmn/cmnpd01fm.cfm?PrgDate=03%2F08%2F2001&PrgID=3)

Goldberg, Gary. "DNA Data Banks Giving Police A Powerful Weapon, and Critics." NY Times. 19 Feb. 1998.

Hibbert, Michelle. "DNA Data Banks: Law Enforcement's Greatest Surveillance Tool." Wake Forest Law Review 34 (Fall 1999): 767-825.

Holt, Cydne and Clinton Stauffer. "Practical Applications of Genotypic Surveys for Forensic STR Testing." Forensic Science International. Vol. 112 (2000): 91-109.

Imwinkelried, Edward, David Kaye. "Forensic DNA Typing: Selected Legal Issues." (2000) (www.law.asu.edu/kaye/pubs/dan/ncfdna-report2-000202.htm)

Migden, Carol. "A.B. 673: Forensic Identification." Assembly California Legislature. California Department of Justice. (2001).

Migden, Carol. "A.B. 673: Forensic Identification, Attachment I." Assembly California Legislature. California Department of Justice. (2001).

Newcombe, Tod. "Slow Spiral: State DNA Lag Databases." Government Technology. April 2000. (www.govtech.net/publications/crimetech/Apr00/CTEDNA.html)

Schiermeier, Lisa. "Virginia's Data Bank Hits". Virginia Forensic Department. (2000)

Schoenberg, Tom. "DNA Dragnet." Legal Times. 10 May 1999.

Smith, Alling, Lane. "DNA Legislation and News." 16 Feb. 2001 (www.ncsl.org/public/siteleg.htm).

Smith, Alling, Lane. "DNA Legislation and News." 2 Mar. 2001 (www.ncsl.org/public/siteleg.htm).

[1] Some of these sources were not directly cited, but found relevant and helpful in the writing of this paper.

Snyder, Howard H.. "Sexual Assault of Young Children as Reported to Law Enforcement: Victim, Incident, and Offender Characteristics." United States Department of Justice. Bureau of Justice Statistics (2000).

Steinhardt, Barry. "Testimony of Barry Steinhardt Associate Director of the American Civil Liberties Union Before the House Judiciary Committee Subcommittee on Crime." March 23, 2000 (http://www.aclu.org/congress/l032300a.html)

This is among the best papers in the class. It's carefully researched, and presents balanced arguments while still reaching a serious conclusion. Thanks!

P.S. – I wonder when we'll start hearing about criminals deliberately dropping other people's hair at crime scenes! Someone should look into the security of floor sweepings in prison barber shops.

"Free" for All — Two Systems for Privacy and Anonymity on the Internet

Some features of communication on the Internet might seem to be determined just by the intrinsic features of the medium. For instance, it might be socially valuable for people to be able to send or receive information privately, but since every packet of data on the Internet includes a source and a destination address, it isn't clear that anonymous communication is even possible. In fact, several projects have recently been developed that attempt to overlay a privacy-preserving communications system on top of the existing Internet. Two of the most visible of these projects, despite their similar names of the "Freedom Network" and the "Freenet" project, are actually quite different. The Freedom Network was a subscription-funded commercial network that attempted to provide privacy protections for common tasks like web browsing and email. The Freenet project is a volunteer-based open source attempt to create a censorship-proof anonymous information sharing network. Because of their differing goals, ideologies, and social contexts, the two systems work differently in many ways, but they also have interesting similarities. Among the most interesting, both systems confront the problem that a completely name-less system isn't really usable — individuals and documents still need some sorts of names if they are to be accessed.

Examining the architecture on which the Internet is based, it's clear that questions of privacy and anonymity weren't very high in the minds of protocol designers. Since the Internet was first developed in a relatively trusted research environment, it was designed neither to facilitate anonymous communication nor to track every possible communication, though moves in both directions have occurred later. In the interests of simplicity and efficiency, a basic packet of information traveling across the net records only the network locations of the computer originating it and the one to which it is destined, and no other globally identifiable information. Historically, this information has been enough to keep users from being substantially anonymous, but not enough to make tracking them easy. When multi-user systems were more common, users could easily blend in with others using the same machine; the more recent trend towards single-user computers has been balanced by increased use of dynamic address assignment schemes under which a machine's address changes over time. In addition, there's never been a strong mechanism to verify the integrity of source addresses — they can be forged by a machine with sufficient network access. The higher-level protocols of the Internet tend to add additional identification and authentication information, though these are usually optional in the sense that they're required only as part of accessing a particular service, not connected with every communication.

As the Internet became commercialized and accessible to a larger community, questions of trust began to require more attention. As luck would have it, this need became apparent around the same time a new tool was becoming available to address it, namely the private-sector development of cryptography, and especially the invention of public-key cryptography. While generally speaking it doesn't require much extra technical sophistication to reveal or collect extra information, moving in the opposite direction to conceal or control the distribution of information is more difficult, and public-key cryptography provided the needed mechanism. The first influential use of cryptographic privacy protection, which helped inspired most later systems, was 'anonymous remailers.' Anonymous remailers represented a progression of increasingly sophisticated mechanisms to conceal the identity of the source of one or a sequence of email messages. As the name implies, the basic mechanism is that a remailer is an intermediary in the transmission of a message, receiving it from the original sender, removing identifying information, and passing it on to the intended recipient. The main challenges of designing a remailer system arise in making it possible for the recipient to reply to the author without compromising his or her anonymity, and in protecting the integrity of the whole system even if some parts are compromised. Remailer developers hit on two main techniques to solve these problems: using public-key encryption to prevent parts of messages from being read by anyone other than their intended recipient, and creating a network of remailer computers, so that a message can be routed through several machines, and its anonymity will be breached only if all these machines cooperate. Though remailers solved many of the technical challenges of anonymizing email, they've never been as easy to use as conventional email systems, and haven't been widely popular.

The design of the Freedom Network, as described in [Goldberg99], [Goldberg00], and [Bouch-

er00] can be seen as an application of the basic principles behind anonymous remailers to Internet communication using any protocol. The most important part of the Freedom Network was a geographically distributed collection of proxy server computers on the Internet, each of which was able to decrypt, encrypt, and forward individual packets of information in the same way remailers process individual mail messages. To communicate across the network, a computer picked a series of servers in the network, and created packets that contained the addresses of these machines, but encrypted in such a way that the only part of the packet any computer in the chain could read is the part that directs it where to forward the packet on to next. A special case was the last server in the chain, which was responsible for sending the message on to its destination somewhere on the regular Internet. Reply packets used a server sequence encrypted in the opposite direction to find the route back to the originating host. As described in the Zero-Knowledge, Inc., papers referenced above, the Freedom Network was a commercially deployed system, in which the proxy servers of the Freedom Network were operated by ISPs paid by Zero-Knowledge, and the software and privileges to connect to the network were sold by Zero-Knowledge. The complete system included some ISP-like services (such as email) provided by Zero-Knowledge, and a complicated system to ensure that users could pay to create a virtual identity (called a nym) without allowing that identity to be linked with the credit card used to pay for it [Russell00]. For a time the Freedom Network was up and running and selling identities, and operating as designed, though it would appear it was not successful in the sense of a business plan, as Zero-Knowledge recently deemphasized and then discontinued the service in favor of more conventional single-computer and 'enterprise' privacy software [Jesdanun].

Though the architecture of the Freenet project involves some of the same basic concepts of distribution over a network and use of encryption as the Freedom Network, it also includes some more radical changes in the way the network stores information compared to the regular Internet. While the Freedom Network is fundamentally about providing anonymity for a transient, point to point communication of information in a single packet, Freenet treats the network as a more abstract information storage facility, more like the World Wide Web in particular than the whole Internet.

Unlike the web, however, content in Freenet isn't identified as being located on a particular computer; instead, a file is identified simply by a hashed code of its contents, and at any time it may reside on one, more than one, or none of the computers at once. Information can be inserted into one 'node' of the network, and it then moves to other machines in the network as it is requested. When a user at one computer asks for a file, if it isn't present there the request is passed on to an adjacent computer in the network, and so on until it is found. Once the information has been located, all the machines in the path keep a copy of it, so that it can be accessed more quickly the next time it is requested. Conversely, if a piece of information is not requested, it is gradually dropped from machines to make room for more popular files, until potentially it can be lost from the network completely. The basic principles of Freenet were developed by Ian Clarke as an academic project [Clarke99], and it has since become established as a free software project being developed by a group of programmers over the Internet. More recent iterations have made increasing use of encryption, so that the computers on which information is stored can't easily determine what information they contain.

Before getting to more substantive comparisons between the two systems, it is perhaps interesting to consider the implications of both projects' choice of the word 'freedom' or its adjective form 'free' to describe their software. In both cases, the word choice can probably be described as a marketing decision — though only Zero-Knowledge is literally associated with trying to entice purchases, Freenet too has reason to be concerned about the impression it conveys to readers of its web site. Freenet needs to attract users, since the network needs to have an adequate number of computers on which to store data, but it also must attract developers, who must be willing to devote their talent to the project for no direct compensation. This marketing effort is also similar for both projects in that they are both taking advantage of the virtuous connotations of 'freedom' to deflect possible criticisms of their work. The Freenet project explicitly invokes freedom of speech as its fundamental virtue [ClarkeUD], but the Freedom Network, though not as explicit, also marketed its software as providing freedom, say from unwanted information gathering. By comparison, the connotations of anonymity are at best morally neutral: though many would agree that anonymity has socially positive

uses, it can also draw up less positive feelings about secrecy and unaccountability.

On a more substantive point, it's clear that much of the difference between the Freedom Network and Freenet has to do with the differing goals and ideologies that motivate the two projects. To the extent that the Freedom Network was a commercial project, one could claim that its main motivation was simply to turn a profit for the company that produced it, but more usefully we can speak of a goal in terms of what the system was to offer its users. Though the low-level mechanisms of the Freedom Network allowed for anonymous communication, the service it provided from the user's point of view was privacy, achieved by associating information not with the user's real name, but with a persistent pseudonym. In most cases, users would keep using the Internet much as they would without the Freedom Network, but secure in the knowledge that their personal information would not be misused. One consequence of this choice is that just as most users of the web don't publish significant information of their own, the Freedom Network didn't provide a facility for pseudonymous web sites. Of course, the decision not to provide this feature was certainly also motivated by business considerations such as the cost of serving popular information, and the possibility for offensive materials generating complaints (Goldberg's thesis [Goldberg00] contains a discussion of how a Freedom-like network could be used to publish pseudonymously, but as far as I know this never became part of the fielded system). For the Freenet project, on the other hand, the ideology of free speech is of paramount importance, and this necessarily involves a sort of 'publishing.' The demand-based data replication scheme used in Freenet addresses the popularity issue, by automatically 'mirroring' (making duplicate copies available of) popular pieces of information. This is in fact an important part of giving any author the possibility of being widely heard, as the free software community has learned from the so-called 'Slashdot-effect' — if a suddenly popular document on the regular Web is hosted by a computer of only average capacity, that computer will often fail or at least substantially slow under the load, effectively making the information unavailable.

Many of the important issues to consider in evaluating a software system relate to the ways it might be misused or abused, but we should be careful here to distinguish between several types of

such problems. One problem occurs when a user of a system manages to subvert the mechanism of a system to do something that should be impossible; in security terminology, this is an 'attack.' A second kind of problem occurs when a user uses a system to do something other that what it was intended to do; we might call this 'abuse.' Finally, a user might use a system in a way that's consistent with the system's rules of operation, but which is undesirable in some broader legal or moral sense; this can be called 'misuse.'

Both the Freedom Network and Freenet are vulnerable to all three kinds of problems outlined above, though the designers of the Freedom Network appear to have taken these issues more seriously. The question of the possibility of attacks is squarely an issue to be dealt with in the design and construction of a system — in a perfect system, every attack would be recognized and contemplated before the software was released, and any attacks that could be defended against would be. Obviously this is an unrealistic goal, but it can be approximated by a careful design and review process. Zero-Knowledge is to be commended for the extent to which it has tried to follow the model set by other security-critical software, in which a system is carefully reviewed by experts not involved with its development, and the results made publicly available for further examination (this is done in [Adam01], a surprisingly honest document to be linked from a company's home page). Some of the same sort of peer-scrutiny is implicit in any project like Freenet that has freely available source code, but there isn't evidence of such a systematic review. Arguably, Freenet is so novel and complex, and still in an early enough stage of its development, that its users should know better than to expect a completely secure system, but its developers do at least appear honest about the flaws they know of.

On the question of abuse, the legal status of the Freedom Network operators forces them to be more responsible than the largely unaccountable developers behind Freenet. In fact, since the Freedom Network provided many of the same services as an ISP to its customers, it was led to adopt the same model of abuse management as a traditional provider would, except without benefit of some standard tools like comprehensive logging (these efforts are described in [Bratzer01]). For instance, an anonymous account would have been a very appealing place from which to send spam, except

that the number of message sent per day was limited to a reasonable number. A more complicated problem occurred when a barrage of spam was sent from somewhere else on the Internet, but using a Freedom Network pseudonym to collect responses; in these cases, it was necessary to investigate further to determine if the mail was indeed connected. In its current form of a distributed network, it's not clear how Freenet could ever have an 'abuse department' in the same way Zero-Knowledge did. The operators of individual nodes in the network might have some ability to control the way their resources were used, but the only scalable solution is to build abuse-resistance into the protocol itself. Some design work has been done securing the protocol against obvious techniques such as trying to insert a large number of random messages [Clarke01], but even if they were all implemented it's not clear this would suffice; malicious users can exhibit surprising creativity.

Perhaps the most obvious danger, of these or any other anonymous communication systems, is that they'd be used to transmit material that is illegal or otherwise undesirable, and that because of the anonymity involved we'd be powerless to stop it. This fear can be made concrete by simply considering all the kinds of information that the government or civil courts currently try to control: entertainment works protected by copyright, computer programs protected by copyright, trade secrets, including those protected by non-disclosure agreements, sexually explicit text or pictures (especially child pornography), libel, instructions for bomb-making, instructions for producing illegal drugs, the communications of criminals (including organized crime), the communications of terrorists (foreign or domestic), information that was gathered with an expectation of privacy, or fraudulent advertising claims, to list some of the more commonly-cited examples. Inasmuch as these systems themselves don't know or care what information they're transferring, there's no important technical difference between all of these potential uses, just differences of how likely a threat they seem, and how dangerous they're considered in a particular political climate or from the point of view of some interest group. It's also important to point out that the difference between these illegal examples and the more general uses of anonymity is often rather subtle. Though we'd like to keep some information secret simply because it would be embarrassing or inconvenient to have associated

with us, many other things one would like to keep private exactly because they would be disapproved of by some larger part of society, even if this disapproval doesn't reach the level of illegality.

It's not clear to what extent the Freedom Network may have been used for illegal activities, both because Zero-Knowledge would be understandably reluctant to release such information, and because those involved would likely take steps to keep their communications secret. Because as was previously mentioned the Freedom Network doesn't provide publishing capabilities, most such uses would have occurred in private media such as email, making them more difficult to track. Though the Freedom Network's terms of service prohibited illegal activity [Bratzer01, 10], it was effectively only able to enforce this rule with regards to public statements like Usenet postings; the capabilities to track other information were explicitly limited, even to Zero-Knowledge, so that theoretically they couldn't be abused by the company or coerced by legal action. Because of the anonymity of individual users, of course, the largest penalty that can be applied is to discontinue the offending nym; but since these must be paid for, and may have an established reputation, this is not a completely empty threat.

Because of the comparatively public nature of Freenet, it's easier to ascertain exactly how it's being used. Obviously there's no centralized statistics system keeping track of everything on Freenet, but it isn't hard to take a representative sample. Probably few were surprised by the results of [Orwant00], which showed the contents of Freenet to include 16% pornography, 22% texts about drugs, 17% copyright-infringing audio, and 2.4% proprietary software. Though some of this distribution, especially in its particulars, is influenced by the tastes of a few early-adopters of the system, the large proportion of contraband material is easy to explain — since Freenet is still quite cumbersome to use compared to the usual World Wide Web, it contains mainly content the demand for which can't be satisfied by that better known medium. For instance, there isn't much explicitly political writing available on Freenet, not necessarily because users wouldn't be interested in it, but because that content is not censored, and therefore fairly easily available, through conventional channels on the Web for its mainly US and western European users.

Of course, Freenet's proponents wouldn't necessarily consider the large proportion of illegal traffic a condemnation of their system, which reflects a particular ethical position. Though most Americans would likely assent to 'freedom of speech' being a desirable social principle, Freenet implicitly endorses (and its creator explicitly endorses, see [ClarkeUD]) a stronger statement that in every context freedom of speech should be a guiding principle. Though he also gives pragmatic arguments about the utility of free speech in various situations, not to mention some of the favorite arguments of Napster apologists, it's pretty clear that Clarke's personal views fall in the more absolutist camp. If we assume for the sake of argument that Freenet can live up to all of its claims, we're left with a thorny question that we might ask about any particular ethical principle: is it so universally valid that we would be well off to embed it directly in the infrastructure of our society? For some virtues, in some political ideologies, the answer is yes, but it's sure to be an issue that will continue to be contested in the future.

Since at least the beginnings of the wide popularization of the Internet in the early and mid-1990s, prognosticators have been extrapolating ~~developing in~~ how we can communicate over the Internet to envision a future in which the free flow of any sort of information is unstoppable. The argument is, generally, that there's a natural tendency of information technology to enhance rather than restrict the flow of information, and that the dissemination of information and technology is a one-way process — once a particular idea has been publicized, or a new kind of software distributed, it can't be taken back. It is of course an empirical question whether this will in fact be an unstoppable trend, but it seems to be playing out in many ways on the Internet today. Though there isn't an absolute rule that says information will always become more widely available, the workings of a network make it much easier to spread information than to keep it under wraps, so a small group of people can often disseminate information even if powerful entities like governments or multinational corporations want to control it. Though neither Freenet nor the Freedom Network has yet had very much success, under this theory the idea is out of the bag. Freenet can continue to grow, and though the Freedom Network has ceased to exist as a network for the moment, the technology is available

for Zero-Knowledge or anyone else to try the idea again if conditions become more favorable. At least as long as some fraction of the world's knowledgeable computer users believe a future of 'free' information is inevitable, it will continue to become closer and closer to reality.

It is pointless to discuss the theoretical societal implications of a technology, though, if it is never used by anyone other than its creators. Both Freenet and the Freedom Network have had trouble attracting as many users as they had hoped, failing to break out of small groups of technically savvy and/or politically wary early adopters. For the Freedom Network, this levelling off was what marked the end of its viability as a business model, and therefore its shutdown. Freenet is still up and running, but also without capturing a large user base among general Web users. It recently experienced a increase in interest as it was considered a Napster replacement, but it hasn't gained a dominant place in this arena, likely because of a still cumbersome interface for configuration and use. Because most Napster users aren't especially concerned about their taste in music getting out, the only necessary guarantee is that those publishing music be unattractive lawsuit targets, which so far can apparently be guaranteed by less technically sophisticated measures. In fact, these systems would appear to furnish further examples in a pattern of security-enhancing software, in which the only systems that become widely used are those that require virtually no extra effort on the part of new users. Email encryption programs like PGP get the same sort of usage — most users find them too cumbersome to use except when sending particularly sensitive information, so they don't get the broad use that would make them commonplace and the protection they provide standard.

Though Freenet and the Freedom Network have some technical similarities, and they both provide a mechanism for private or anonymous communication, they are also significantly different in important ways. Many of these differences can be tied to the fact that the Freedom Network was developed as a commercial product, while Freenet was not, while others are driven by the different ideologies of the system creators. Whatever the successes or failures of these particular systems, though, it's clear that some of the issues they raise will continue to be important as we consider what communications will look like in the future.

References

Back, Adam et al. "Freedom 2.1 Security Issues and Analysis" unpublished white paper, Zero-Knowledge Systems Inc, 2001. http://www.freedom.net/products/whitepapers/index.html

Boucher, Philippe et al. "Freedom System 2.0 Architecture" unpublished white paper, Zero-Knowledge Systems Inc, 2000. http://www.freedom.net/products/whitepapers/index.html

Bratzer, David and Andrew Elkin. "Freedom 2.2 Abuse Issues and Analysis" unpublished white paper, Zero-Knowledge Systems Inc, 2001. http://www.freedom.net/products/whitepapers/index.html

Clarke, Ian. "A Distributed Decentralised Information Storage and Retrieval System" unpublished report, Division of Informatics, University of Edinburgh, 1999. http://freenetproject.org/freenet.pdf

Clarke, Ian. "The Philosophy behind Freenet". Undated web page. http://freenetproject.org/index.php?page=philosophy

Clarke, Ian et al. "Freenet: A Distributed Anonymous Information Storage and Retrieval System" in *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, H. Federrath, ed. Springer, 2001. http://freenetproject.org/index.php?page=icsi-revised

Froomkin, A. Michael. "Flood Control on the Information Ocean: Living With Anonymity, Digital Cash, and Distributed Databases" *University of Pittsburgh Journal of Law and Commerce*, 1996, p. 395. http://www.law.miami.edu/~froomkin/articles/ocean.htm

Goldberg, Ian. *A Pseudonymous Communications Infrastructure for the Internet*. Ph.D thesis, Computer Science Division, University of California at Berkeley, 2000. http://www.isaac.cs.berkeley.edu/~iang/thesis.html

Goldberg, Ian and Adam Shostack. "Freedom Network 1.0 Architecture" unpublished white paper, Zero-Knowledge Systems Inc, 1999. http://www.freedom.net/products/whitepapers/index.html

Jesdanun, Anick. "Online anonymity tool to shut down" *Associated Press*, October 4th, 2001. http://www.canoe.ca/CNEWSTechNews0110/04_zero-ap.html

Kleiner, Kurt. "Free speech, liberty, pornography" *New Scientist*, March, 2001. http://www.newscientist.com/features/features.jsp?id=ns22812, mirrored at http://www.efc.ca/pages/media/2001/2001-03-10-a-newscientist.html

Korman, Richard. "Free Radical: Ian Clarke Has Big Plans for the Internet" *OpenP2P*, O'Reilly & Associates, November 14th, 2000. http://www.openp2p.com/pub/a/p2p/2000/11/14/ian.html

Orwant, Jon. "What's on Freenet?" *OpenP2P*, O'Reilly & Associates, November 21st, 2000. http://www.openp2p.com/pub/a/p2p/2000/11/21/freenetcontent.html

Samuels, Russell and Ed Hawco. "Untraceable Nym Creation on the Freedom 2.0 Network" unpublished white paper, Zero-Knowledge Systems Inc, 2000. http://www.freedom.net/products/whitepapers/index.html

Interesting and off the beaten track. Thanks.

Facing the Issues

Advances in computing now allow computers to identify a person through unique physical or biological features. One of the most promising and controversial technologies is facial recognition technology. When local police at the Super Bowl deployed the technology in January, its use stirred great controversy.[1] While facial recognition technology is a tool with numerous useful applications, the deployment of these systems in public gives cause for concern.

Facial features are an example of biometrics – measurable physiological and behavioral features that can be used to identify an individual.[2] Other examples of biometrics include fingerprinting, iris and retinal scanning, voice patterns, and of course, DNA. The use of fingerprinting for verification/identification purposes is a long established practice, with the use of DNA analysis rising dramatically in the last decade.

While facial recognition is one of many biometric technologies, the methods and scope of its potential applications set it apart. Facial recognition can be used in verification (one to one matching), identification (one to many matching), and surveillance. Companies have developed systems that allow cameras linked to computers to scan faces and compare them to photographs stored in databases. They market the systems for use in law enforcement, security, and commercial businesses.

[1] http://www.washtech.com/news/regulation/11586-1.html Online article about use of facial recognition technology at the Super Bowl. Robert O'Harrow Jr., August 8th, 2001.
[2] http://homepage.ntlworld.com/avanti/ Web site that introduces basics of biometrics

Companies cite the advantages of biometric face recognition over traditional verification methods. Passwords and PINs are easily transferable from one person to another, so they don't actually verify that the individual who presented the token is the legitimate user. A biometric such as a person's distinct facial features, however, is sufficiently unique and very difficult to transfer or counterfeit. The options of expensive plastic surgery or makeup disguises do exist, but their effectiveness against facial recognition depends on which parts of the face are encoded by the underlying algorithm. The automation of facial recognition also streamlines the process of verification, which leads to savings in time, manpower, and money.

Despite the marketing brochures, biometric facial recognition is still being perfected. In its current form, the systems need a full frontal view of the face, especially both eyes, and adequate lighting in order to maximize the chances of a match. Despite certain limitations, it is already finding widespread uses. The earliest adopters have been law enforcement agencies and casinos.[3] Police officials have used the systems to scan for suspected criminals in public areas and plan to use it to manage mug shots in a crime database.[4] The DMV has used systems to detect duplicate or fraudulent license registrations, while businesses have used it to prevent fraud. In particular, casinos use the system to guard against cheaters, but there are plans to use it as a way to identify the high rollers as well.[5]

[3] http://wbbm.cbsnow.com/now/story/0,1597,274604-240,00.shtml Online article about the use of facial recognition technology in casinos. Associated Press, February 26th, 2001.
[4] http://biz.yahoo.com/bw/010814/142123.html Online article includes statement made by Visionics Corp. about getting the contract from Minnesota.
[5] http://www.uniontrib.com/news/metro/20010717-9999_1n17cameras.html Online article about casinos in San Diego area using facial recognition technology. Chet Barfield, July 17th, 2001.

And therein lies the controversy. Do high rollers want to be identified easily? While the idea of recognizing an individual by their facial features is not new, the use of computers to make it an automated process opens up new possibilities that have to be dealt with. Specifically, computerized facial recognition makes it so much easier to identify people. With the ability to do perform a task faster with computers, there exists the worry that this time it may cost people more of their right to privacy.

The constitutional "right to privacy," although not explicitly stated, is inferred from the Bill of Rights and Supreme Court rulings. The First Amendment guarantees the freedom of expression and association, while the Fourth Amendment guards against "unreasonable searches and seizures."[6]

Originally, the Fourth Amendment was invoked to only protect personal property against such searches and seizures. However, the Supreme Court eventually ruled that the Amendment "protects people, not places," stating that "the principal object of the Fourth Amendment is the protection of privacy rather than property."[7]

In addition, the Supreme Court has ruled that the Due Process Clause of the Fourteenth Amendment protects "personal decisions relating to marriage, procreation, contraception, family relationship, child rearing, and education."[8] Therefore, the constitutional "right to privacy" consists of not only *physical* privacy, but also the notion that "citizens should be able to control certain information about themselves and to make decisions free of government compulsion."[9]

[6] *Bill of Rights* (1787)
[7] *Katz vs. United States*, 389 U.S. 347 (1967)
[8] *Planned Parenthood of Southeastern Pennsylvania v. Casey*, 505 U.S. 833, 851 (1992).
[9] http://www.rand.org/publications/IP/IP209/IP209.pdf "Super Bowl Surveillance: Facing Up to Biometrics," John D. Woodward, Jr. RAND 2001

The collection and dissemination of personal information about an individual, by the government, businesses, or other organizations, has always worried those concerned with privacy rights. The Supreme Court itself noted that it was "not unaware of the threat to privacy implicit in the accumulation of vast amounts of personal information in the computerized data banks or other massive government files."[10]

With facial recognition systems, not only will more information be accumulated about a person, the nature of that information is much more personal than most other personal information. While a person can change their name, move to a new address, and get a new phone number, it's highly difficult to change one's facial features. Face recognition systems are engineered to be "resistant to lighting, skin tone, facial hair, hair styles, eyeglasses, [facial] expression, and pose."[11]

Another one of the concerns about facial recognition is that it doesn't require an individual's explicit consent to record the information. In most cases, people don't even know that they are being scanned by these systems. It's one thing to write down one's name, birth date, address, phone number, and e-mail address to apply for a credit card, bank order, or driver's license, and another to have that personal information recorded about a person without his knowledge or consent. Facial recognition systems have the ability to detect a person's face and scan their facial features at a distance without their knowledge.

Of course, even though facial features are highly personal and unique in nature, they are also universally exposed. However uncomfortable people are about having their faces scanned and compared to other faces in a database, people don't have a "reasonable

[10] *Whalen v. Roe*, 429 U.S. 589, 605 (1977).
[11] http://www.visionics.com Web site of Visionics Corp., a tech company offering biometric verification/identification systems.

expectation of privacy" with respect to their faces in public unless they take steps to conceal them.[12]

In that case, *United States vs. Dionisio*, the defendant was ordered to give a voice sample to the grand jury for identification purposes. Dionisio argued that the order constituted an "unreasonable seizure" on his person, thereby violating the Fourth Amendment. However, the Supreme Court wrote in its ruling, "Like a man's facial characteristics, or handwriting, his voice is repeatedly produced for others to hear. No person can have a reasonable expectation that others will not know the sound of his voice, any more than he can reasonably expect that his face will be a mystery to the world."[13] In addition, the face recognition systems currently discard all the digital "faceprints" that don't register a match in the computerized database, meaning no personal information is retained, disclosed, or transferred.

Civil libertarians argue that there are currently no restrictions on actually storing someone's facial features in a database instead of deleting it. Storage capacity and cost isn't an obstacle at all. A faceprint for a given person's face can be less than 100 bytes.[14] Adding 100 bytes to each profile in an existing database is very cheap in terms of space and cost. If the database has 250 million profiles, that would be adding the equivalent of 25 gigabytes to the existing database, which is smaller than the capacity of an average hard drive. Storing six billion faceprints, one for every person on Earth, would only result in an additional 600 gigabytes.

[12] *United States v. Dionisio*, 410 U.S. 1, 14 (1973).
[13] IBID, p. 18
[14] http://www.visionics.com The company web site claims that a "faceprint" can actually be stored in 84 bytes.

In reality, computerized databases storing people's faces already exist at places like schools, the DMV, and the credit card company. Faceprints can be easily generated from these images. The greatest concern for privacy advocates is who will have access to this information and how they will be used. Civil libertarians' ultimate fear is that the technology will evolve into a surveillance tool that will be abused by governments, businesses, and other groups.

Automatic face recognition could be used to provide tracking and surveillance on individuals wherever they went. Moreover, different systems can be networked to share information on a particular person, allowing the different organizations to construct a very detailed profile of that individual. For example, in light of the events of September 11th, some airports around the country are deploying facial recognition systems. Prior to that, it was already being deployed in public municipalities around the U.S. If people thought that their movements and actions can be tracked and recorded, they will be less likely to freely express themselves and engage in activities that oppose powerful interests. Critics say the effect is tantamount to social coercion and control.[15]

The vital question is, therefore, whether facial recognition systems deployed in *public* violate an individual's right to privacy under the Fourth Amendment. The Supreme Court has mostly followed a two-part test formulated by Justice John Marshall Harlan II, given in *Katz vs. United States* back in 1967. Justice Harlan stated the right to privacy has a "twofold requirement, first that a person have exhibited an actual (subjective) expectation of privacy and, second, that the expectation be one that society is

[15] http://dlis.gseis.ucla.edu/people/pagre/bar-code.html "Your Face Is Not a Bar Code: Arguments Against Automatic Face Recognition in Public Places," Phillip E. Agre, last updated October 30, 2001.

prepared to recognize as 'reasonable.'"[16] In Katz, the defendant was bugged electronically when making a phone call from a public phone booth without a warrant. Justice Harlan interpreted the act of entering a phone booth and paying for a phone call as satisfying both parts of the test for reasonable expectation of privacy. Therefore, the conviction based on the electronic surveillance was overturned.

When a person walks out in public, however, unless he has taken steps to conceal his face from view, the individual can't expect others not to look at his face. If that's true, then an individual can't expect not to be *recognized* by someone who sees his face. Under these guidelines, facial recognition technology doesn't seem to violate the Fourth Amendment. However, society does need to be concerned about the vast accumulation of personal information by governments, businesses, and how facial recognition systems add an intensely personal and specific piece of information to that collection.

In its current form, facial recognition technology still faces numerous technical obstacles but performs adequately for its current uses. As the technology is refined and becomes more powerful, the social implications increase dramatically. Properly employed, it will provide a dependable, streamlined, and cost effective alternative to traditional practices of identification and verification. Without careful oversight, however, it can be abused to infringe even more on the privacy of citizens, even to the point of turning society into an Orwellian surveillance state. Therefore, it is important that guidelines and regulations be established to safeguard society against the abuse of this new technology.

Good - you did a great job of filling in the gaps in the first draft.

[16] *Katz vs. United States*, 389 U.S. 347 (1967)

The Economics of Information:
Physical & Digital

INTRODUCTION

The very *nature* of information presents a dilemma for both consumers and suppliers of information goods. On the demand side, consumers are uncertain about the utility of information (its worth) because it is difficult to determine its value until one has it. But they cannot have information until they have paid for it. Yet they cannot possibly know how much to pay for information until they have determined its utility by having it.

Suppliers have an equally formidable problem. Information is extremely costly for them to produce; yet it is relatively inexpensive to transmit. They, therefore, have a difficult task of recovering their investment (e.g., time and research & development) through the sale of information. This is because the first consumer of an information good instantly becomes a competitor of the original supplier, resulting in an iterative process that eventually drives down the price of the good to its (very low) cost of distribution.[1]

This suggests that the market, barring some form of governmental regulation, will result in an undersupply of information; producers, a priori, will recognize that they cannot recover the high sunk costs of production, and consumers, a posteriori, will desire to become free riders[2] for information. Recognizing this disjoint, the Government has developed intellectual property rights—or laws creating property rights for ideas.

This paper, in Section I, will focus on Copyright protection, which grants ownership rights to authors, artists, and composers. The tradeoff between the incentive for the creation of

[1] The fact that producers have difficulty selling information for more than a fraction of its value is called the problem of non-appropriability—that is, they have difficulty appropriating the idea's value. Appropriation, colloquially, may have come to mean "grabbing something that was already property, i.e., theft—but here, appropriation is to take possession of it, i.e., turn it into property.

[2] Free riders are those consumers who may enjoy a product without paying their "fair" share.

ideas and the incentive for their dissemination will be discussed here. Section II will explain the fundamental differences between physical property and digital property, highlighting the changes in the economics of digital information. Section III questions whether legal protection or technical protection of intellectual property in the digital age will be more efficient. Section IV asks whether protection for the appropriation of monetary value is the only objective of the suppliers of information, and the Conclusion attempts to sum up the main points of the paper.

I. COPYRIGHT PROTECTION

Because an unregulated market will produce a sub optimal amount of information in the economy, and since this, in turn, threatens total economic welfare, the Government has developed three areas of Intellectual Property Law that confronts the tradeoff between the creation and dissemination of ideas. Due to limited scope of this paper, however, only Copyright protection will be discussed here, although the other two fields could be included.[3]

Granting exclusive rights to the creator of an idea, upon demonstration of original expression, under copyright law, allows one theoretically to appropriate much of its social value. But while this may create an incentive for efficient innovation, there is a concomitant social cost. Because the owner of an idea can exclude others from its use—without sale or license—this creates a problem for the dissemination and application of the information.[4] The breadth and duration of a Copyright, therefore, affect the social benefits and the social costs of appropriation, as required for efficiency.

[3] The other two areas of Intellectual Property law are: the patent system, which establishes ownership rights to inventions and other technical improvements; and the trademark system, which establishes property in distinctive commercial marks or symbols. The principal economic justifications for granting property rights to trademarks are that they protect consumers against fraud and create an incentive for producers to supply goods of high quality.

[4] Although all information is *not* similar, the point here is that if owners of ideas do *perfectly* exclude others from their use then the result is duplicative research (since ideas have to be re-researched and re-created)—instead of applicative research, which would apply the ideas already created.

Holding this footnote doesn't answer my objection—in fact, this makes it worse, since digital media have affected research dissemination much less than entertainment.

A. Breadth

The breadth of a Copyright concerns the uses to which the material can be put without authorization. A broader scope forbids any unauthorized use, while a narrow scope permits some unauthorized uses. These uses tend to fall under fair-use exclusions, which vary depending on the medium of information.[5]

B. Duration

The duration of a Copyright in the United States now stands as the creator's life plus 70 years. The optimal duration of a Copyright concerns the problem of tracing costs, which accrue from "tracing" an owner and obtaining permission for the use of copyrighted material. But the limited duration of a copyright tends to ameliorate these costs.

It is clear from this discussion that a copyright of wider scope and longer duration will strengthen the incentives for innovation, while at the same time weaken those for dissemination and application.

II. PHYSICAL PROPERTY VERSUS DIGITAL PROPERTY

Although Copyright protection provides a legal framework from which owners of ideas—either embodied in a physical product or something more intangible—can appropriate its value, there are a number of assumptions that must be relaxed within the context of digital information, exposing again the problem of nonappropriability. Thus, while Copyright protection of

[5] For example, recording over-the-air copyrighted television programs on a videocassette recorder is fair use when done for "time-shifting" purposes, but not necessarily for purposes of "archiving". There is a fine line that divides fair and unfair unauthorized copying.

information in the real world has largely been successful, Intellectual Property rights in the digital age have proved insufficient at best. This disparity exists because there are three fundamental differences between physical property and digital property; that is, digital information changes the economics of reproduction, distribution, and publication.[6]

A. The Economics of Reproduction

Information that is protected under Copyright law allows the owner of an idea to charge a price above its marginal cost (the extra cost associated with a one-unit increase in output)—a situation similar to a monopoly, in that prices are higher and output is lower than in a model of perfect competition. An owner wields monopoly power—at a cost to society—but it theoretically gives these suppliers the incentive to innovate in the first place.[7]

Copyright protection has traditionally worked well with information that is embodied in physical property because of the natural barriers to reproduction—that is, the high cost and the decrease in quality from copying a work.[8] But digital information radically reduces the difficulty and cost of reproduction, without any decrease in quality. Most infringement can be accomplished in *private*—violating the law seems trivial—with technology that is not necessarily beyond the resources of *large* numbers of people. Further, copies produce exact replicas of the original, with minimal loss of quality even from *copy to copy*.

not obvious why these are equivalent /

B. The Economics of Distribution & Publication

[6] Samuelson, Pamela and Davis, Randall. *The Digital Dilemma: A Perspective on Intellectual Property in the Information Age*, 7.
[7] Lord Macaulay said that copyright is "a tax on readers for the purpose of giving a bounty to writers." Macaulay, Thomas B. "Speeches on Copyright", 25.
[8] Cost here includes the difficulty of access to technologies that enable infringement. The lower quality is due to successive generations of copies in analog media.

this, too, is more about entertainment than about research, where a 2024 photocopy is often good enough.

The essence of digital information and the concomitant ease of its reproduction have removed the physical barriers to replication. Digital information has further changed the method of distribution: the dominant form of transmission of information goods has hitherto been the sale of a physical copy of a work. Today, digital information is primarily licensed, transferring not complete ownership rights to an idea, but rather allowing a limited transfer of rights on specific terms and conditions. Because the sale of a good involves physically transferring the product from the owner to the consumer, the transaction thus leaves the original supplier without the good. Ownership of physical goods, therefore, exhibits the characteristic of excludability.[9] Licensed digital information, on the other hand, leaves the *owner*, in addition to the consumer, with a copy of the work.

This was always true.

Whereas "physical" information has relied mostly on distribution through vendors like book, music, or video stores, digital information can now be distributed almost instantaneously to anyone with a computer connected to the Internet. These vast computer networks have made distribution of, and access to, digital information cheap and relatively easy.

The World Wide Web has also allowed anyone to be a publisher of information, vastly increasing the publication of material. While there may be a greater supply of ideas, this has caused radical changes in the publication industry.

III. LEGAL PROTECTION OR TECHNICAL PROTECTION IN THE DIGITAL AGE

The differences between physical and digital property have important implications regarding how information will be protected. Protection of information is significant for both owners, who will want to appropriate the value of their ideas, and for consumers, who can be assured of the

[9] The distinction here is between information bound to a physical good, and an idea itself—which may or may not be excludable.

information's authenticity. Suppliers may rely on Copyright as the dominant law of the digital age, or they may instead invest in technical protection, as a security against any sort of unauthorized infringement.[10]

A. Copyright As Legal Protection

The notion that most people will pay some fee for information from a few suppliers is represented as the "Celestial Jukebox" model, based on the old-fashioned music devices that play songs after receipt of some money. In this model, most information will be conglomerated in a few—recognizable and highly visible—large suppliers to which users would pay royalties for the right to utilize some form of digital information. The Copyright system would have a large role in the digital age, as Internet traffic would be highly regulated. Enforcement of infringement would make unauthorized information difficult to locate.[11] As Lance Rose argues in his article first published in *Wired* magazine, "It is irrelevant whether any given infringement goes unpunished—as long as it is kept outside the public marketplace."[12] That is, Copyright law succeeds at maintaining public markets for copyrighted products.

B. Technical Protection

While the previous model relies heavily on the Copyright system to be successful, the model of "digital libertarianism" assumes little room for Copyright law, as technical protection will make legal protection unnecessary. In this model, cheap encryption by individual suppliers of information will be a more efficient protection of intellectual property.

[10] From: Cooter, Robert and Ulen, Thomas. *Law & Economics*, 136.
[11] Readers may be concerned with the amount of faith placed in the Copyright system in this model.
[12] Rose, Lance. *The World Wide Web and Copyright Law* from: Ermann, M. David et al. *Computers, Ethics, and Society*, 222.

As technology improves, there will be better and more effective ways of protecting information. Some limitations of digital information may be allowing consumers only to view text-based information products—such as books, magazines, newspapers, research journals, manuscripts, and so on—without an option to either print or save the material. Other digital information may be provided on only a time- or audience-limited basis, further limiting its ability to be shared among consumers.

IV. OBJECTIVES OTHER THAN PROTECTION TO APPROPRIATE VALUE[13]

Whether Copyright, technical protection, or some combination of the two becomes the dominant law of the digital age may not be a point of contention for some suppliers of information. Those in this category may be interested in objectives other than *protection* as a means to appropriate the value of information.

Some owners of digital information may distribute certain products free, in order to obtain an indirect benefit in a related market (such as free web browser software to command market power in web server software[14]); others may want to distribute intellectual property in order to build community (such as the operating system, Linux); and still others may want to keep digital information private (such as trade secrets[15]).

Protection may also not be important to those suppliers of information who rely on a "Ransom Model" to make a sufficient return on their investment in ideas.[16] Instead of flatly giving away their intellectual property, owners use a portion of their product as a teaser to attract

[13] From: Samuelson, Pamela and Davis, Randall. *The Digital Dilemma: A Perspective on Intellectual Property in the Information Age*, 14.
[14] Hoping to gain a return to investment, these suppliers, of course, will want protection of the related intellectual property.
[15] A well-known example is the formula for Coca-Cola.
[16] Varian, Hal. *Internet Changes the Economics of Information Industries.*

donations from consumers, which allows the owners, after receiving a sufficient amount of money to cover costs, to release the material in its entirety. Although the free rider problem may be present here[17], those consumers who contribute may receive related material—such as T-shirts, autographs, or concert tickets (depending on the product)—as additional enticement beyond pure interest.

CONCLUSION

As Pamela Samuelson, a well-known Berkeley Professor of the School of Information Management and Systems, explains, the new information infrastructure has the potential to be either an "information leveler"—providing information to all those who had little or no prior access—or an "information stratifier"—deepening further the divide between those with and without information. This is because the supply of information digitally (instead of primarily through physical goods) has changed the economics of reproduction, distribution, and publication. But the nature of information presents a problem of incentive for the suppliers of digital information: the high cost of investment in the creation of an idea cannot, in an unregulated market, be covered by the correspondingly low price that must be offered (due to the relatively cheap cost of transmission). In response to a potential reduction in economic welfare by an undersupply of information, Intellectual Property rights—Copyright here—and/or technical mechanisms, therefore, provide an opportunity for owners to appropriate the value of their information through protection. But if owners of information are motivated by objectives other than those that are pecuniary, then protection may have a much smaller role for information on the Internet.

Good paper.
of course the trouble with folks like Samuelson and Varian is that they leave out a lot of the political reality—such as the fact that huge corporations own the legislatures. You might want to see if you can find any Marxist economists who study IP issues.

[17] See note 1 above.

USA PATRIOT Act vs. Encryption

Under pressure to respond to terrorist attacks on the World Trade Center and the pentagon, Congress passed the Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism (USA PATRIOT) Act of 2001 upon a 96-1 vote in the Senate on October twelfth, 2001. Included in the bill were many additions to police power meant to prevent the use of the computers and the internet for terrorist purposes. Some are concerned that the new government powers are too expansive and will prevent people from protecting their privacy through encryption.

In general, those who support increased police powers do so because they believe that additional criminals, particularly kidnappers, will be caught and the nation will be more secure against espionage with the additional powers. The price of decreased privacy is worth these benefits. However, as lately as 1995, the FBI had not encountered a single encrypted voice communication. Wiretaps were used only in two to three kidnapping cases a year on average (between 1968 and 1993). No wiretaps were used in terrorism cases between 1988 and 1994. (7) The opposition argues that this government intrusion goes too far, even going as far as to say that "privacy is at the very soul of being human" (7) and that this legislation sacrifices too much privacy. The fourth amendment of the Constitution in particular guarantees freedom from unreasonable search and seizure.

Background on Public Key Encryption

Public key encryption is a method for two parties to communicate securely through an insecure channel. For instance, if two people, Alice and Bob, wanted to communicate privately, over a line which someone could be listening, they could employ such a system. If Alice wants to send a message to Bob, Bob must first produce two keys, his private key and his public key. He shows his public key to anyone who wants to see it and keeps the private key to himself. Alice then uses the public key to scramble the message she wants to send to him. When Bob receives the message, he can use his private key to unscramble the message. Without Bob's private key, the original message cannot be retrieved, and since only the scrambled message is transmitted over the line, even if someone else received the scrambled message, only Bob can unscramble it.

Public key encryption has been widely applied to secure a variety of messages such as e-mail, bank transactions, and military communication. Several activities such as e-commerce and web-based university class registration would be unable to exist without the security guarantees of cryptography.

Encryption over the internet differs from previous means of long distance communication in that the government is unable to easily access the content of the message. Mail can be opened and telephone wires can be tapped, but a person's private key is not as easy to acquire without accessing their computer.
This has created a desire to limit the ability of encryption to transmit any message securely.

Some attempts have been made in the past that would protect privacy and (in a seeming self-contradiction) allow the government to access messages without allowing wider access to the protected messages.

The first attempt was to limit the export of programs that could generate large keys. With only small keys, the government (or others) could potentially calculate someone's private key from their public key. The algorithms necessary to implement a public key cryptographic system with keys of any size have been widely published in journals, textbooks, and the internet. Restricting export as a means of control has been called "a practice about as pragmatic as restricting the export of wind." (1) Hobbyists in foreign countries quickly created their own cryptographic programs unimpeded by the export controls. In addition, in 1999, the ninth circuit court ruled that the government could not regulate the export of these algorithms because they are a form of protected free speech. (2)

A second attempt was made in the form of a key escrow system. Under this system, a trusted third party, such as the government, would be entrusted with the private key. Although this system would give government the access it desires, it has critical weaknesses. The first and most obvious is that the private key would no longer be private if it were revealed to a third party. The third party would be an obvious target for attack since infiltrating the trusted third party would get the intruder access to all of the messages sent to those people who had trusted this third party, so the security of the entire system is now dependent on the third party, instead of the people who received the messages. The second is that using this cryptography system does not preclude the use of the original system in conjunction with the escrow system which would make an attempt by the trusted third party to unscramble the message just as difficult as attempting to unscramble the message without the key escrow system. (3)

Despite several attempts to limit encryption with these methods and various variations on these themes, none have been both popular and effective because in each case, the security provided by the new system is weaker than that of the old and requires at least the same amount of effort on the part of the user.

The USA PATRIOT Bill

Recent anti-terrorism legislation has given law enforcement agencies many new powers in an effort to prevent crime, especially terrorist crimes. However, these new abilities raise concerns about privacy and security of citizens.

Intercepting Internet Traffic

Part of section 216 of the USA PATRIOT act reads:

"the court shall enter an ex parte order authorizing the installation and use of a pen register or trap and trace device anywhere within the United States, if the court finds that the attorney for the Government has certified to the court that the information likely to be obtained by such installation and use is relevant to an ongoing criminal investigation. The order, upon service of that order, shall apply to any person or entity providing wire or electronic communication service in the United States whose assistance may facilitate the execution of the order. Whenever such an order is served on any person or entity not specifically named in the order, upon request of such person or entity, the attorney for the

Government or law enforcement or investigative officer that is serving the order shall provide written or electronic certification that the order applies to the person or entity being served." (4)

The ability granted by this paragraph is to read the internet traffic of anyone "relevant" to an investigation without having to disclose what is discovered. In addition, law enforcement officials decide who is "relevant." Those whose information will be captured need not be the subject of investigation themselves. It is easy to see how this would be convenient for law enforcement officials wishing to monitor the internet use of a target and his or her associates. Without this law, the permission of the court would be required to intercept this traffic. Now, law enforcement officials no longer need to deal with the court in order to do this. This could potentially expedite investigations and allow more communications to be intercepted, leading to the conviction of more criminals.

Balanced against that are the concerns about the privacy of individuals. The term "relevant" is so broad that it could apply to almost anyone at any time. Since the traffic of any "relevant" person may be captured without notice or disclosure, one does not know whether or not the government is watching. Under the assumption that ones computer communication is being observed by a third party, encryption offers a solution to avoid the compromising ones privacy. In practice, for most people, any such observation will not lead to discovery of criminal actions. Assuming that criminals using computers to commit crimes already use readily available cryptography software to discuss criminal actions, this expanded interception would not assist the government in uncovering criminal activity.

Secret Searches

The USA PATRIOT legislation allows law enforcement to use warrants in a new way. A part of the section 213 of the act reads: (5)

DELAY- With respect to the issuance of any warrant or court order under this section, or any other rule of law, to search for and seize any property or material that constitutes evidence of a criminal offense in violation of the laws of the United States, any notice required, or that may be required, to be given may be
delayed if—

(1) the court finds reasonable cause to believe that providing immediate notification of the execution of the warrant may have an adverse result (as defined in section 2705);

(2) the warrant prohibits the seizure of any tangible property, any wire or electronic communication (as defined in section 2510), or, except as expressly provided in chapter 121, any stored wire or electronic information, except where the court finds reasonable necessity for the seizure; and

(3) the warrant provides for the giving of such notice within a reasonable period of its execution, which period may thereafter be extended by the court for good cause shown.

In short, law enforcement officials no longer have to give notice to the person or owner of the property to be searched or seized in advance of such a search or seizure. For instance, the police can enter a person9s home and examine the contents of their hard drive without their knowledge and without needing to immediately present a warrant. Besides securing communication, encryption can also be used to protect information stored on hard drives. Encrypting hard drives can protect a person's data from this kind of examination. However, this is not a complete solution since the person9s password can be acquired by bugging the computer they use to access the data. This approach was recently publicized in the case against Nicodemo Scarfo, who is accused of loan-sharking and illegal gambling and questions of its legality are working their way through the courts. (6)

The use of this power can successfully bypass protection provided by cryptography. By accessing the computer of the suspect, the authorities can access the suspects private key which will then allow them to quickly decrypt intercepted messages. Even if the computer is password protected as in the Scarfo case above, eventually the suspect will try to use the computer at which point the password must be entered and can be detected by a bug. However, this method is the most invasive because it requires authorities to install a bug in physical proximity to the computer. Since the legality of this method before the law was passed was undecided, it is unknown if this provision provides the police with any additional power in this case.

Sharing information

A part of section 203 of the USA PATRIOT bill reads:

Any investigative or law enforcement officer, or attorney for the Government, who by any means authorized by this chapter, has obtained knowledge of the contents of any wire, oral, or electronic communication, or evidence derived therefrom, may disclose such contents to any other Federal law enforcement, intelligence,protective, immigration, national defense, or national security official to the extent that such contents include foreign intelligence or counterintelligence (as defined in section 3 of the National Security Act of 1947 (50 U.S.C. 401a)), or foreign intelligence information (as defined in subsection (19) of section 2510 of this title), to assist the official who is to receive that information in the performance of his official duties. Any Federal official who receives information pursuant to this provision may use that information only as necessary in the conduct of that person's official duties subject to any limitations on the unauthorized disclosure of such information.".

The government hopes that sharing information gained through surveillance between federal agencies will make it easier to solve crimes and when suspects are in custody, the case against them can be made stronger by including more evidence in the possession of other parts of the federal government.

Similarly to the first two attempts to address criminal activity, encryption is able to protect the communication of those who use it. If incriminating communications are made using encryption, the government will be unable to unscramble the message, so no information useful to the prosecution is gained. However, information of no immediate

use acquired by wiretap and other means can be shared all over the federal government, destroying the privacy of that information.

Conclusion

Most of the USA PATRIOT bill does not overcome the security provided by encryption technology but allows government officials to intrude on the privacy of private communications. Although the bill is aimed at catching criminals through their electronic communications, simple encryption of these messages would thwart casual attempts to unscramble them. Even though the bill may help catch those criminals who do not use encryption, its provisions critically undermine the privacy of the communications of all citizens and the security of their homes.

Because of this almost one-sided comparison of the benefits of additional security to the price of privacy, the position of security is indefensible as an argument supporting this bill. Without the force of serious national security and crime prevention arguments, this bill has no legs to stand on.

Bibliography

1. http://www.eff.org/Privacy/Key_escrow/decrypting_puzzle_palace.article 11/5/01
2. http://www.epic.org/crypto/export_controls/bernstein_decision_9_cir.html Bernstein v. USDOJ (9th Cir. May 6, 1999)
3. http://www.eff.org/Privacy/Key_escrow/960724_isoc_crypto.statement IAB and IESG statement on cryptographic technology and the Internet July 24, 1996
4. http://www.politechbot.com/docs/usa.act.final.102401.html Final text of the USA anti-terrorism bill
5. http://www.aclu.org/congress/1110101a.html November 1, 2001:ACLU Legislative Analysis on USA PATRIOT Act
6. http://www.wired.com/news/politics/0,1283,46329,00.html Scarfo:Feds Plead for Secrecy
7. Privacy on the Line: The Politics of Wiretapping and Encryption, Diffie and Landau, 1999 MIT Press

Good – you've acknowledged some weaknesses in the argument, which strengthens your effectiveness.

But I still don't quite believe it. Most criminals, I think, don't use encryption – they lack the technological savvy. So USA-PATRIOT might not catch Osama bin Laden, but it might catch the next Timothy McVeigh – and meanwhile it'll probably catch bunches of non-political minor criminals.

So I don't find "it won't work" a compelling argument. What convinces me to oppose such laws is that they will work all too well, bringing us closer to the virtually crime-free Singapore.

CS 195

"The Effect of Supercomputing

in Arms Control Policy"

As nuclear stockpiles left in the aftermath of the Cold War forced world governments to reevaluate both foreign and domestic policies, the world itself seemed to stand still. Frozen by the dream of a life without the inevitability of nuclear war, initiatives such as the Comprehensive Test Ban Treaty (CTBT) were created to protect against proliferation and nuclear explosion. Created and not signed, signed and then never ratified. And so never made into reality, for the world is still frozen in this dream. And yet while the United States advocates the CTBT as a proponent of foreign policy, it still undermines its purpose domestically. The past decade has seen the emergence of the Stockpile Stewardship Program and the Accelerated Strategic Computing Initiative (ASCI), operated under the Department of Energy to ensure the reliability of the nuclear arsenal. Supercomputers funded by the ASCI and built by domestic companies will use simulations to ensure the capability of US nuclear weapons without the need for physical testing. And therefore allow for a secure nuclear arsenal within the guidelines of the CTBT, which does not prohibit computer modeling. However, these supercomputers are unique to the United States alone and not within the capabilities of other nuclear nations. Thereby making the CTBT, in the light of these computing advancements, not a step toward a non-nuclear world, as it was intended, but a liability for world governments. Effectively, ratifying the CTBT would allow these nations to keep weapons without being able to test them, and subsequently without a nuclear deterrent. For what is the reliability of having guns, left for years, without knowing if they could still even fire a bullet? In order to make the CTBT a viable policy option, a world policy, the issue of supercomputing must be addressed. The scientists behind these machines can play a major role in this process, by supporting the need for the CTBT to address simulated testing and prompting a discourse between nations on this issue; by eliminating the black and white view of the CTBT, in order to resolve the gray area of simulated testing. For the past decade has shown that the creation of the ASCI, the debate over the CTBT, and the emergence of the supercomputing alternative have been inextricably linked. And at the center of this history, is the scientist.

Within the past decade, the ratification of the CTBT has been a precarious balance between security and a world without nuclear weapons. In this sense immediately ironic, for governments look to the very weapons they seek to eliminate for protection. And the United States stands at the frontline of this apparent contradiction, holding one hand in a gesture of non-proliferation and disarmament while trying to maintain its nuclear deterrent. For "the robustness and reliability of America's nuclear deterrent remain crucial to security and stability: America and its allies need to be confident that the weapons that remain will go bang if needed, and only if needed". [1] And so on July 3rd, 1993 President Clinton announced in a radio address a proposal to negotiate the CTBT and a worldwide nuclear moratorium.[2] However, October 5th, 1993 changed dramatically the nature of this "proposal" when China conducted its first nuclear test since the address.[3] the "robustness and reliability" of nuclear weapons, it seemed, was not a concern to the United States and its allies alone.

On January 30th, 1995 the President launched another initiative, extending the nuclear moratorium and reinterpreting the United States stance on the CTBT, retracting the special "right to withdraw" privilege that the US had been advocating for itself, a right that allowed the President to withdraw from the CTBT for a span of 10 years after its signing. During this concession, the American people are reassured by APNSA (Assistant to the President for National Security Affairs) Lake that "the President considers the maintenance of a safe and reliable nuclear stockpile to be a supreme national interest of the United States"[4]. And it was true, 1995 marked the creation of the "Stockpile Stewardship Program"[4], a $4.5 billion dollar per year effort to maintain the US supreme national interest. Its goal to preserve the functionality of the nuclear stockpile without physical testing, in essence, to maintain the nuclear alternative. From its budget, $1 billion was dedicated to ASCI, the Accelerated Strategic Computer Initiative.[5] The three national defense laboratories, SNL (Sandia National Laboratory), LANL (Los Alamos National Laboratory), and LLNL (Lawrence

[1] The Economist. May 24th, 2001. www.clw.org/pub/clw/coalition/briefv5n13.html
[2-3] The White House, September 22nd, 1997, www.usatoday.com/life/cyber/tech/review/crh273.htm

[4] U.S. Department of State, October 8th, 1999.
www.state.gov/www/global/arms/factsheets/wmd/nuclear/ctbt/fs_991008_stockpile.html
[5] "Cray Research-Silicon Graphics wins DOE award
for world's most powerful supercomputer" www.lanl.gov/orgs/pa/News/101196.fulltext.html

Livermore National Laboratory) united scientists under ASCI grants to develop a new wave of supercomputers. The goal was to create a 100 teraflop machine capable of testing nuclear weapons accurately.[6]

Subsequently, on August 11th, 1995 President Clinton announced US support for a true zero yield CTBT, a proposal that would ban any nuclear weapons test and any other nuclear explosion.[7] In 1996 the "ASCI Option Red Supercomputer" is built by Intel under ASCI funding, capable of 1.06 teraflops. And within the same year, a $110.5 million grant is awarded to Cray/SGI to build a 3 teraflop computer called "Blue Mountain" in LANL.[8] And so as the United States advocates a safer world, one with a ratified, zero yield CTBT, it is creating a supercomputing alternative to maintain the integrity of its nuclear arsenal after the treaty.

> "The Comprehensive Test Ban Treaty is a giant step toward a safer, more peaceful world. We also need to ensure the safety and reliability of a reduced U.S. nuclear stockpile," said President Clinton. "This agreement will provide the Los Alamos National Laboratory in New Mexico with the world's most powerful supercomputer - a computer that will provide a reliable substitute to the underground testing we have worked so hard to ban"[9]

IBM also joins the effort in 1996, the ASCI gives $94 million to the company to build the "ASCI Blue Pacific" at LLNL.[10] While both supercomputers are under construction, President Clinton sets a precedent to the world by becoming the first world leader to sign the CTBT on September 24th, 1996. By 1997 the power of the American supercomputer is unmistakable, IBM's "Deep Blue" wins a six game rematch against the World Champion of Chess, Garry Kasparov. By 2000 the "ASCI White" is completed by IBM and delivered to LLNL at a $110 million dollar cost, and it far exceeds its original expectations, surpassing Moore's Law by 2.3 teraflops. Capable of 12.3 teraflops, the

[6] "Three-Day Calculations Set Engineering Analysis Milestone" May 8th
2000.www.sgi.com/newsroom/press_releases/2000/may/blue_mountain.html
[7] The White House, September 22nd, 1997, www.usatoday.com/life/cyber/tech/review/crh273.htm
[8-9] "Cray Research-Silicon Graphics wins DOE award
for world's most powerful supercomputer" www.lanl.gov/orgs/pa/News/101196.fulltext.html

[10] www.epcc.ed.ac.uk/direct/newsletter5/node33.html

ASCI White is already 1,000 more powerful than Deep Blue.[11] In 2000, Compaq is given a $200 million dollar contract to produce the "ASCI Q", the next line of ASCI supercomputers. Compaq promises a 30+ teraflop machine to be completed in 2002.[12] But not to be ousted, IBM announces its plan for "Blue Gene", a supercomputer capable of a 1000 teraflops.[13] The promise of technology and its advancement continues unhindered as the past decade has proved, but what of the promise of a world with a CTBT?

Understanding the Scientist's Role in Policymaking

"A model democracy would not relieve the scientists of the burden of responsibility; man is responsible for his actions, particularly if he belongs to the few who have the ability to assess, better than anybody else, the consequences of discoveries and the implications of their applications"[14].

The history of the ASCI program has united government and private companies in a series of coalitions between think tanks, private resources, and of course, federal money. The scientist remains crucial to this effort, and yet produces technology without being a part of the politics of its use. While the research and the creation of these machines continues successfully, the influence of these computers in arms policy goes unheeded. The scientist is relieved of the burden of responsibility, following the example of Edward Teller, who "when defending the continuation of nuclear tests, said that it is the duty of scientists to find out about all the potentials of a new discovery, but it is the duty of the politicians, the people's representatives to decide what to apply and what to leave out"[15]. In terms of the CTBT, this approach will no longer suffice. As the environment in which the treaty was originally created has already changed dramatically by the invention of the supercomputing alternative. This new alternative must be addressed and the CTBT reevaluated.

[11] ww.cbsnews.com/now/story/0,1597,210684-412,00.shtml
[12] "U.S. Department of Energy Selects Compaq to Build World's Fastest and Most Powerful Supercomputer"www.compaq.com/newsroom/pr/2000/pr2000082202.html
[13] ww.ecommercetimes.com/perl/story/4104.html
[14] Supek, Ivan and Malecki, Ignacy. Scientists, the Arms Race and Disarmament: "Scientists in the Contemporary World". Taylor & Francis Ltd, London 1982. p179.
[15] Supek, Ivan and Malecki, Ignacy. Scientists, the Arms Race and Disarmament: "Scientists in the Contemporary World". Taylor & Francis Ltd, London 1982. p179.

The scientist's role, therefore, is to propel this discussion. For while creating supercomputers is not unethical, they cannot sit idly while politicians use their inventions to circumvent the CTBT. That supercomputers would give the United States a nuclear alternative that no other nation in this treaty could maintain. If the CTBT is to be a sign of world trust and of world security, then it must be ratified with these intentions in mind. And, indeed, the scientist can play a very important role in ensuring that this goal is achieved. Jeffrey W. Knopf writes:

> "By pooling their resources and coordinating their activities with elites who share some of their policy goals, citizens' groups can stimulate action by and enhance the capabilities of like-minded political elites. . . leverage is possible through this process because the president needs the support of a majority Congress for certain elements of his program that are related to arms control policy, such as weapons' appropriations and appointments of negotiators, and of course two-thirds of the Senate it he wants a treaty ratified. For this reason, the president will normally also be concerned about the stance of other elites who might affect congressional opinion, such as nuclear scientists. Because presidents need a winning coalition in Congress and value an elite consensus behind their policies, there is a route for activist influence even if arms control is not a major concern with most voters"[16].

The scientist, by exerting influence among and apart of citizens' groups, can motivate a reevaluation of the CTBT. They form the crucial elite opinion that can alter the political environment and arms control policies as active, important components of the process. For as a "model democracy would not relieve the scientists of the burden of responsibility", the scientist should also work to make the democracy model.

By supporting the need for the CTBT to address simulated testing, the United States can set a precedent to the world. That the CTBT is still a proponent of trust and security, that it is still the dream that nations shared twenty years ago. And a continuation of the "Russell-Einstein Manifesto", as "there lies before us, if we choose, continual progress in happiness, knowledge, and wisdom. Shall

[16] Knopf, Jeffrey W. Domestic Society and International Cooperation: The Impact of Protest on US Arms Control Policy. Cambridge University Press, Cambridge 1998. p60-61.

we, instead, choose death, because we cannot forget our quarrels? We appeal as human beings to human beings: Remember your humanity, and forget the rest"[17]. To remember that the CTBT is still a proponent of peace, and to acknowledge that simulated testing may derail it. To act within the United States to bring this issue to a resolution, and to make the CTBT not just a signed document but one also that is sworn to and ratified by.

So, your position is that the CTBT should prohibit the use of supercomputers for weapon simulation? If that actually happened, there wouldn't be any supercomputers — the funding would disappear. ∴ no jobs for supercomputer researchers. So you're asking a lot from these scientists.

That doesn't mean you're wrong. But it makes the ethical decision-making harder than you seem to suggest.

[17] ww.pugwash.org/about/manifesto.htm

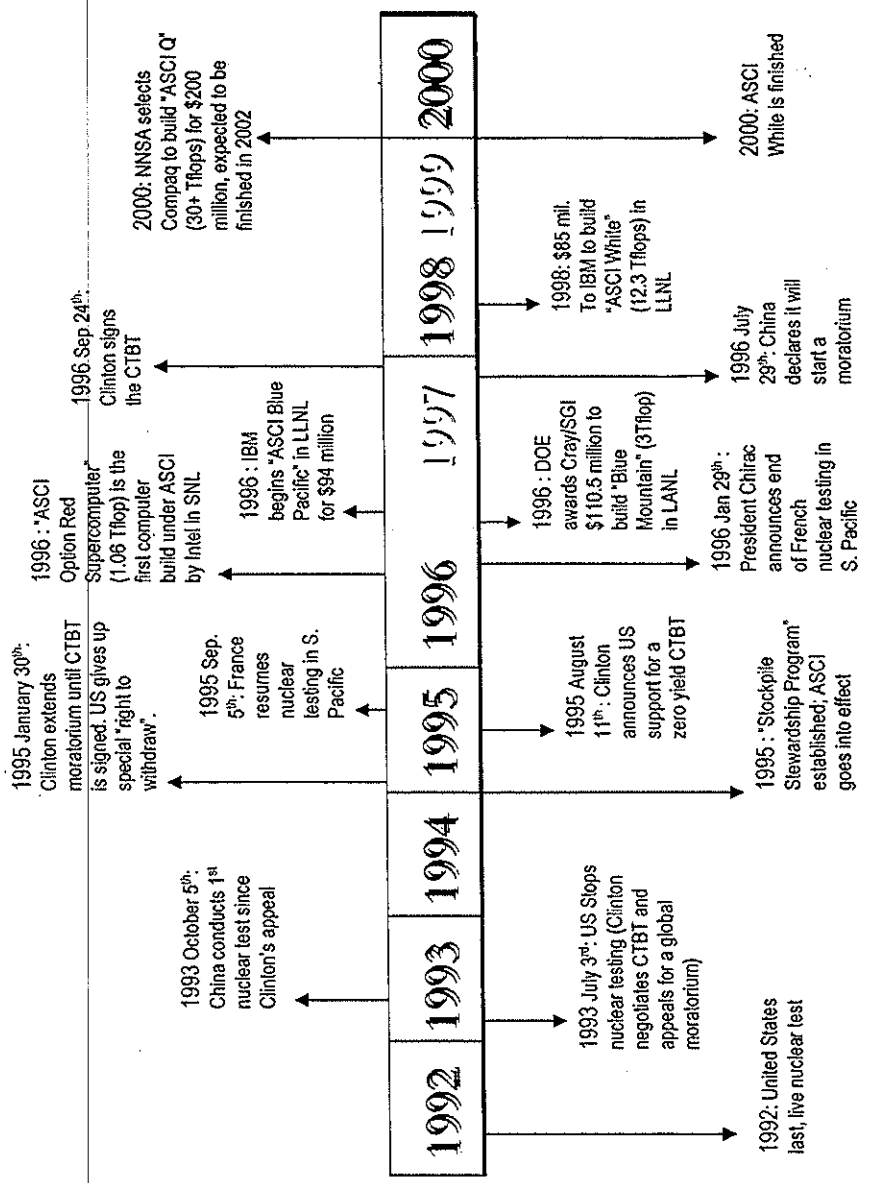

It Is Time To Think About Dynabook Again

In 1977, Alan Kay, the pioneer of today's graphical user interface (GUI) based personal computers, introduced what he called "Dynabook". The Dynabook was a conceptual portable personal computer which everyone can use as easy as a paper and pen. The Dynabook consisted of a flat screen touch panel display, a keyboard, GUI, a wireless networking capacity and so on. It made a great impact on all computer scientists.[1]

Today, we all know that personal computers are a very convenient tool. Using a computer, you can check the latest news on the internet, you can chat with your friends or you can type your papers without a white-out. This is, however, true only if you know how to use them. Of course many people know how to use computers. Most high schools have computer classes, businessmen use computers for their paperwork and college students download their favorite music from the internet. Still, there is a group of people who have not gotten any opportunities to learn how to use computers, namely senior citizens, even though there are many applications of computers useful for their daily lives.[2]

One of the reasons why many senior citizens do not use a computer is that some of them think that they do not need a computer in their daily life at all. They can write a letter by hand without using word processing software. They can call or fax to someone instead of sending e-mail. They can see the news on TV or on papers. They can go shopping on foot. For these people, computers might not help their daily live very much. Though there are still many computer applications that help their daily life such as monitoring senior citizens who have a chronic illness by a computer so that doctors will be able to cure them more properly.

Not all senior citizens think that they do not need a computer. In fact, some of them are eager to know how to use computers.[3] They are the people that computer scientists have to consider more carefully because most of the factors that make computers harder to learn for them can be eliminated easily with a little consideration of computer designs.

The first thing we look at is the concepts of personal computers. Personal computes have become much easier to use since the GUI was implemented than before. The very first GUI-based personal computer was Alto, developed by the team led by Alan Kay in Xerox's Palo Alto Research Center (Xerox PARC).[4] Based on the concept of Alto, in 1984, Apple Computer introduced Macintosh, a fully GUI-based personal computer.[5] The concept of GUI these computers brought was very clear. The monitor is your desktop. The cursor is your hand on the screen and you move it with moving a one-button mouse. You can open different applications on your desktop just like you open your books or put your typewriter on your desk. You trash unnecessary files into a trashcan. This idea was successfully accepted by a huge number of computer users and GUI-based computers became the mainstream of the personal computer market.

However, after 1985, this simple concept became more complicated. Microsoft introduced the first version of Windows in 1985.[6] Windows's GUI was very similar to Macintosh's but there are many differences between them that confuse people. The

desktop is now a wall and you put your favorite wallpaper on it. A mouse now has two buttons and instead of double-clicking, you may be asked to "right-click". Folders are called directories. These small conceptual differences among various systems are now mixed and you put your favorite wallpaper on your desktop today. It sounds a little ridiculous to argue about these small things; however, these small things really make the concept of computers hard to understand for people who do not have a concept of computers at all, especially for senior citizens. This problem is very hard to solve unless we redefine all the metaphors. 7

Hardware designs also have to be considered carefully. Personal computers consist of lots of quite complicated hardware components. This is probably the first difficulty beginners face. Many senior citizens that first see a computer always ask you where the power switch is. Once they turn on the computer, they now try to turn off the computer with the same switch and often they are told not to do that or they will break the computer. This is clearly a hardware design problem. The power switch should be the switch to turn on and off the computer. If they hear that they may break the computer, they will certainly be afraid of computers. This problem is clearly easy to fix, the operating system only has to handle the signal from the power switch correctly. This is a good example of computer designers' lack of consideration.

Input devices are very user unfriendly in various ways. For example, a mouse is a very bad human interface. Beginners often place the tail of the mouse on his side. In other words, they hold the mouse upside down. They are then told that it is wrong. This also makes them afraid of computers. Then, once they hold the mouse correctly, they are told to move it and click the button. This procedure is actually very difficult for senior

What happens when the OS freezes? You still need a hardware power-off switch. What you mean is that the big, visible power switch should be software-controlled, as on most laptops.

citizens. Some of them cannot move the pointer straight or cannot click the button without moving it. There are many other difficult actions they have to master such as double-clicking, pressing the right-button of a mouse, dragging or moving mouse with holding a key on a keyboard. These actions are way far from the actual actions you take on your actual desktop such as moving your books or placing a typewriter.

A keyboard is also a bad input device. In fact, this is true for not only senior citizens, but also all beginners. There are more than one hundred buttons on an ordinary keyboard. The alphabets are usually located in non-alphabetical order, for example both of the qwerty layout and the dvorak layout are not an alphabetical order layout. Other keys have weird names such as "ctrl" or "alt". You also often have to press more than two keys at the same time sometimes. Windows or MacOS very often requires you to press 3 keys at the same time to reboot the system. The usability of a keyboard is a very serious problem for people who do not speak English, for example, Japanese. Japanese have more than fifty letters. Each letter has two forms just like English alphabets have upper and lower case. Also, Japanese uses hundreds of thousands of Chinese characters at the same time. The way they type Japanese is that they first type everything with Japanese letters and then convert the letters into Chinese characters if needed. In addition, there are two ways to type Japanese letters. One is that all Japanese characters are mapped into one key just like English letters are mapped into one key. However, in this way, they have to memorize where all the fifty characters are located. The other way of typing Japanese is that they first type the English characters printed on the keyboard and according to the sound, computer converts the English letters into Japanese letters. This

This is a good thing! The idea is that rebooting is an unusual event, and it should be hard to do by accident.

I don't think you got the point of my comment about using keyboard buttons with the mouse. This isn't just one more problem to add to the list. It's an example of a profound, general problem: There's a limit to how simple you can make the UI, because people want to do complicated things! So you need the idea of layers of UI—simple things simple, complicated things possible.

sounds a very good solution, however, for some Japanese senior citizens, this is impossible because they do not know anything about English alphabets.

For these reasons, many universities and companies are researching about the input devices in Japan. These researches are very demanded since Japan is one of the most aged societies. For your information, the ratio of the population older than 65 years old to the population of Japan is 17.9% in 2001 (July 2001). This is higher than Italy (17.7%, Jan.1999), Sweden (17.3%, Dec. 1999) or the United States (12.9%, July 1999).[8] There are mainly two approaches to solve the input device problems. One is a touch panel and the other is a voice-control system.

With a touch panel system, you can avoid the nightmares of the uncontrollable mouse and the unfriendly keyboard. You can click buttons on screen with your fingers, or you can actually write characters with a pen without worrying about the character conversions. Casio first introduced their first pen-based PDA in 1983 in Japan.[9] Since then, the handwriting recognition systems for Japanese have been improved dramatically. However, even though this technology is already available, the touch panel based personal computers are not popular yet because there is one very serious problem. You cannot feel the feedback from touching the screen. This is a quite big disadvantage, especially for senior citizens. Since the virtual buttons are a new concept for them, they do not think that they pressed the button without feeling it. However, NTT Docomo finally solved this problem in 2001.[10] They introduced a "clickable" touch panel. Its idea is very simple. When a user touches the panel, it vibrates the whole screen once. Since you always push one button at once with one finger, only the finger touching the panel feels the vibration and you feel the "click". With the handwriting recognition and the clickable touch panel, the notorious mouse and keyboard problem may be solved soon. Alan Kay has claimed recently that the handwriting recognition system cannot be an alternative of a keyboard[11], however, I believe that a handwriting recognition system is more useful for non-English speaking people then a keyboard.

The voice control system is a remarkable technology also. This may change the whole concept of personal computers. The idea of the voice control system is that you ask the computer to do something and the computer does the job for you. The most famous voice recognition system today is IBM's ViaVoice. It supports ten languages today including Japanese, Chinese, French and so on.[12] People say that it recognizes natural conversation quite well, however, this system has a problem as well. It is not good for inputting commands such as resizing a window and moving it to the location (50,100). Sharp has developed a solution for these voice control system recently. It is called CG assistant.[13] The software displays a human-like CG character on the screen and asks you what to do. Once you tell it what to do, it does the job for you. With the system, you do not have to say "d-e-l space colon ..." but you can just ask the character to delete something. No one doubts that this technology is clearly useful for senior citizens.[14]

There are still many small computer design problems. Text characters on the screen are too small or the display is too bright and so on. However, these problems have already been solved partially. Many operating systems have an option to make the text characters bigger. Even though most operating systems still use the same GUI layout as the one for smaller fonts so that it does not solve the problem well, some operating systems have a special GUI layout for larger fonts and it solves this problem well. Many

companies have developed EL displays or paper displays that are more gentle for your eyes than CRT or TFT LCD screens.

The jargons and unreadable manuals are also a big problem. However this problem has already recognized and many companies are now researching on this problem. For example, Sony has considered changing all fancy icons in their manuals that make no sense, or replacing the jargons such as "wizard" or "web".[15] Also, many Japanese companies are trying to replace all English-based words to normal Japanese words.

Today, all of these technologies are available, and we can easily make the Dynabook-like computers. Still, there is one more big step we have take and it is the most difficult one, changing the society.[16] The most important thing to do now is to care about senior citizens more and make a society in which everyone supports senior citizens. Today, computer companies are only focusing on inventing new technologies, making faster computers, designing fully featured complicated systems. Of course, these are clearly necessary for computer industry to grow. However, I believe that we need to split the computer industry into two different groups. One focuses on inventing new technologies just as we have done, and the other one focuses on people who use computers. Computers used as a "computer" such as enterprise servers or office workstations must be treated differently from computers used as a "tool" for people's daily life like TV sets or telephones. If we care about "people" in this way and spend some of our time and money, the real Dynabook will come true in very near future.

Interesting idea, but not new — TVs already have computers inside, as do some telephones, and they don't use the desktop metaphor. But the Dynabook idea is about giving people general computing power, not just smart appliances.

Notes:

1. Gasch, Scott "Alan Kay" <http://ei.cs.vt.edu/~history/GASCH.KAY.HTML>
   Rampersad, Brian "Key Things To Remember" <http://www.sheridanc.on.ca/~randy/design.dir/software.dir/key.htm> Apr. 1997
2. Nagashima, Hiromi "Technical ends and the computer usage of senior citizens" <http://buri.sfc.keio.ac.jp/access/research/rep96/eishima.html>
3. "Learning Computers" Asahi Shinbun 6 Apr. 2001 <http://www.asahi.com/tech/feature/K2001040602164.html>
   "Computer Schools – they want to send email to their grandchildren" Asahi Shinbun 20 Apr. 2001 <http://www.asahi.com/tech/feature/K2001042001110.html>
4. Xerox PARC <http://www.parc.xerox.com>
5. Apple Computer Inc. <http://www.apple.com>
6. "Windows98: A History of Windows" PC Magazine <http://www.zdnet.co.jp/magazine/pcmag/9806/windows98/history.html>
7. "Thinking About User Interface" Isys Information Architects Inc. <http://www.iarchitect.com/mshame.htm>
8. "Population of aged people" Statistics Bureau of Japan 2001 <http://www.stat.go.jp/data/guide/5-3-1.htm>
9. Casio Corporation <http://www.casio.co.jp>
10. VRSJ Newsletter Vol.6, No.4 <http://www2.vls.gifu-u.ac.jp/vrsj/nl/msg00012.html>
    NTT Docomo Corporation <http://www.nttdocomo.co.jp>
11. Sellers, Dennis "Kay comments on agents. Dynabook" Macworld <http://maccentral.macworld.com/news/0106/19.kay.shtml> 19 June, 2001
12. IBM Corporation <http://www.ibm.com>
13. "CG model helps web searching. Sharp's experimental model" Asahi Shinbun 7 Sep. 2001 <http://www.asahi.com/tech/feature/K2001051807813.html>
    Sharp Corporation <http://www.sharp.co.jp>
14. "IT for disabled people" Asahi Shinbun 5 Oct. 2001 <http://www.asahi.com/tech/feature/K2001100400065.html>
15. "Manual – Confusing terminologies" Asahi Shinbun 18 May 2001 <http://www.asahi.com/tech/feature/K2001051807813.html>
    Sony Corporation of Japan <http://www.sony.co.jp>
16. "Braille, Voice, Larger Fonts"

Mainichi Shinbun 1 Aug. 2001
<http://www.kaigo-fukushi.com/seikatsu/200108/seikatsu2001080101.html>

good – well researched and presented.

Can You Handle This: Denial of Service Attacks

Without a doubt, the term *Denial of Service* (DoS) is slowly finding its place in our modern day lexicon of illicit activity on the Internet, right beside its more glamorous counterparts such as *hacker* and *virus*. This should be no surprise to anyone after the more than extensive (possibly overblown) media coverage of DoS attacks against corporations like Yahoo.com, Amazon.com, Buy.com (attacked a mere hour after their initial public stock offering), ZDNet.com, E-Trade.com, eBay.com, and CNN.com perpetrated by the 15 year old Canadian hacker Mafiaboy in February of 2000. Even Microsoft, the overly confident software behemoth, suffered a crippling attack against its DNS server that rendered its Hotmail network inaccessible for over two hours (I know because I kept trying to check my e-mail in vain during that time). What people don't seem to realize is that DoS has actually been around for a while. Only recently has it gained a great deal of attention, a trend that will likely continue into the future.

What is DoS?

CERT, a major Internet security-reporting center operated by Carnegie Mellon University, defines Denial of Service as an attack on computing systems and communication networks "in which the primary goal of the attack is to deny the victim(s) access to a particular resource". This definition of DoS is very general and for a good reason, for very many different types of DoS exist. CERT classifies all DoS attacks into three major categories:

- Consumption of scarce, limited, or non renewable (computer) resources
- Destruction or alteration of configuration information
- Physical destruction or alteration of network components

In the first group, hackers take advantage of the fact that computers and networks need certain resources—network bandwidth, memory and disk space, CPU time, data structures, access to other computers, and physical resources, in some cases such as power, cool air, and even water—to accomplish their goals. Taking away any one of these elements can crash or severely limit a computer or network. With regard to the second class of DoS, the performance of computer systems and communication networks are highly dependent on many configuration parameters that describe the environment in which they work. Thus, anyone capable of altering the configuration of computers or network components can either limit performance or bring an entire service to a standstill. Finally, the most barbaric form of Denial of Service is to simply disable or destroy the hardware that provides the attacked service. [1]

As previously explained, DoS comes in a variety of flavors. In practice, however, DoS are most frequently directed against network connectivity. Attacks on network connectivity can further be divided into logic attacks and flooding attacks. Attacks in the first category exploit bugs in software in order to tie up or crash servers. Logic attacks are, in a sense, unavoidable because software is created by humans and thus inherently buggy; however, serious bugs are eventually exposed, though sometimes in the worst way, and patched. Flooding attacks, as the name suggests, consume a network's CPU, memory, or bandwidth resources with the sole purpose to deny legitimate clients access to a server or network resource. Unlike logic attacks, there are no patches for flooding attacks because they are comprised of network traffic that requests the use of legitimate services; the problem is that the traffic is artificial. For example, in a TCP SYN flood attack, a continuous stream of TCP SYN packets is sent to a server; each packet requests a TCP connection between the sender and the server. The server acknowledges each packet, sets aside some of its resources to maintain the connection, and waits for the other side to begin its transmission. Eventually, if enough TCP connections are opened, the server will run out of resources; it will be incapable of servicing other requests for TCP connections. However, in the

case of a TCP SYN flood, there are no legitimate users on the other side of the opened connections because the TCP SYN packets are created with the sole intent to tie up the resources of the server. The important point to realize here is that if too many legitimate users connect to that same server their TCP SYN packets would tie up the server in the same way a TCP SYN flood would. For this reason, flood attacks are difficult to guard against.

What is DDoS?

In order to carry out a successful DoS flood attack, an attacker must be able to generate more traffic than a network is capable of handling. For example, a very crude form of DoS against Internet mail servers might operate by generating more e-mail than the server can handle. However, even the fastest home connection is not capable of generating enough traffic to effectively DoS a large-scale mail server. In order to add more clout to DoS, hackers simply multiply the number of systems involved in the attack. This technique is known as Distributed Denial of Service (DDoS) and possesses the particularly dangerous quality of being able to be carried out by a single individual who posses a minimal amount of computing and financial resources. As the name suggests, DDoS is implemented by mounting several DoS attacks from different physical locations on the Internet. Typically, an attacker will work to compromise several different systems on the Internet in order to install an independent process, or daemon, called a zombie that is capable of carrying out its own DoS attack against a specified target. The process of acquiring zombies can be manual; passive, as in the case of simply posting Trojan horses with DoS capability where people will download them; or, in more sophisticated cases, autonomic. An example of an autonomous DoS process is the case of the infamous IRC (Internet Relay Chat) bots which spread to different IRC servers, seize them, and deploy the program to execute a DoS when given the command to do so from a remote location. [2]

How Serious a Problem is DoS? (Subhead)

Although DDoS has not yet gained the reputation of say viruses, it is quickly becoming very popular. DDoS attacks are rarely carried out against home users, and unlike computer viruses, they do not spread with the purpose to infect as many systems as possible. Instead, the perpetrators of DDoS attacks usually have some definitive target in their scopes. On the other hand, people are starting to be more affected by DDoS because an increase in the number of high-speed and always-on connections has given hackers more targets to compromise for use as zombies that carry out DoS attacks against other specified targets. Furthermore, victims of DDoS attacks are starting to include network services that have a wide range of subscribers, even people with low- to medium-speed connections like modems. As a result, the whole spectrum of Internet users is becoming acquainted with DDoS attacks. ISP's are quickly realizing that the infrastructure that makes up the Internet has no systems in place to combat DDoS. No one could foresee that the Internet would become the perfect breeding ground for DDoS. Because no mechanism exists to stop the problem at its roots, recovering from DDoS has become a network administrator's nightmare.

see end of paper for my comments

The basic principle behind DoS attacks is easier to execute in comparison to other forms of hack attacks such as Man in the Middle (MITM) Attacks or gaining root level access. When executing a DoS attack, the only thing a hacker really has to worry about is not getting caught. DDoS helps with this problem by decentralizing the source of the attack. In addition, hackers know that they can hide their tracks by exploiting the very nature of the Internet Protocol with a technique known as IP spoofing. IP spoofing allows hackers who implement DDoS attacks to hide the source address of their attacks. At the lowest protocol level, IP spoofing involves forging packets to contain an IP address other than the actual IP address of the machine that generated the packet in the source field of the IP packet header. Thus, the machine generating these malformed packets "spoofs" the IP address of the host who actually possesses the IP address. Before IP

What does this mean? A reader who needs DoS explained won't know this either!

But DDoS does require penetration of the intermediaries

spoofing was incorporated into DDoS attacks, it was a very effective technique in gaining trusted host status in the r* service of early Unix systems. These services had very weak authentication systems that amounted to checking the source field of an IP header to confirm that it originated from a trusted host. [3] Of course, a problem that makes this attack less than perfect is that a computer that receives spoofed packets will send its response back to the spoofed address. This, however, is not a problem for DDoS attacks because what is being transmitted is important only in its ability to impair a system's ability to function properly. In other words, the attacker does not care what response it gets from a server; the attacker is only interested in whether the packets sent are capable of denying legitimate clients whatever service the server provides. The purpose of IP spoofing in a situation like this is to prolong the attack by making it more difficult to locate the source of the attack. However, as we shall see later, third parties are using the effects of IP spoofing to study DDoS attacks.

To understand the nature of DoS attacks, it is perhaps useful to obtain an indication of how widespread they are. A recent study has shown surprising statistics on the number of DoS attacks carried out on the Internet. The study, carried out by David Moore, Geoffrey Voelker, and Stefan Savage, is detailed in the report "Inferring Internet Denial-of-Service Activity" and uses a technique known as "backscatter analysis" to estimate the world-wide prevalence of DoS attacks. Backscatter is essentially an indirect consequence of IP spoofing. Spoofed packets enter a network and create conditions that amount to DoS. In the process, the server that is the destination of those packets generates response packets that are directed to the spoofed address. However, if a server exists at the spoofed address, it will transmit some sort of error packet back to the targeted server because it knows that it did not attempt to initiate communication with the server. Because the attackers spoofed source addresses are usually selected in some kind of random process, the response packets are distributed somewhat uniformly across the Internet. It is this inadvertent effect that is called backscatter. A third party could obtain a rough estimate on the number of DoS attacks carried out on the internet by monitoring a randomly distributed sample of IP addresses for replies to spoofed packets. Further, by measuring the rate at which such response packets are generated, it is possible to obtain a lower bound on the intensity of such attacks. Using these methods, the article "Inferring Internet Denial-of-Service Activity" concludes that 12,805 attacks were carried out in a three-week period. Some of these attacks were carried out with an intensity of over 600,000 packets-per-second (pps). It is important to keep in mind that any figure based on backscatter techniques can only provide an estimate on the number of DoS that produce backscatter. [4]

How can DoS be countered?

The question of what can be done to thwart DoS attacks is a good one at this point. At the moment, it is unclear whether DoS attacks will ever be successfully thwarted. DoS attacks come in a variety of shapes and sizes, and as a consequence, the traditional method for hardening systems against DoS is really a hodge-podge of many different security techniques. These techniques, as outlined by the websites of the CERT and SANS (System Administration, Networking, and Security) institutes, boil down to the following:

- Timely application of patches and system updates, especially to potentially exposed machines.
- Deployment of only strictly necessary network services.
- Intrusion detection systems.
- Packet filtering

These aren't really relevant to DoS (as opposed to penetration) attacks, are they?

This section is much improved! You do a good job explaining and discussing the possibilities.

- Address filtering, also known as "egress filtering", of packets leaving the enterprise. This can ensure that packets leaving carry source addresses within the ranges of those sites
- Investing in hot spares, machines that can be placed into action quickly in the event that a similar machine is disabled.

- Investing in more bandwidth to lower your vulnerability to flooding attacks.
- Investing in redundant load-balancing networks and servers. If there are multiple versions of the same Web site operating on different network segments, rogue packets can be distributed evenly amongst them making it more unlikely that any given server will crumble under the weight of an attack.
- Education and communication throughout the community can be extremely helpful. When organizations fail to share information about attacks, this helps give the hacker community an even greater advantage. Systems administrators should participate in industry-wide early warning systems. Information about attacks should be disseminated to vendors and response teams so that it can be applied to the defenses of others.

Unfortunately, these traditional defenses have their weaknesses. The timely application of patches is a war computer security experts have been playing with hackers since day one: it all comes down to who finds the bug first. Usually, patches are made available the moment a weakness has been discovered or exploited; however, keeping up with the massive list of security updates and patches is extremely time consuming, just like constantly monitoring computer systems and networks for potential problems. Many networks simply do not have the time or money to invest in keeping their systems as up to date as they would like to, and it is these networks that are a perfect home for zombies capable of carrying out DDoS attacks. [5]

mixed metaphor

But presumably they can contract this task out to specialists who could provide semi-automated patch distribution.

The main problem with firewalls is that most firewalls are essentially stateless filters. What this means is that they are only able to thwart some DoS attacks by filtering out malformed or unwanted packets. However, as mentioned earlier, most DoS attacks consume server resources by exploiting legitimate features. Firewalls that simply filter malformed packets cannot thwart these types of DoS attacks. Stateful firewalls, on the other hand, are a step up from stateless firewalls because they actually keep the state of connections initiated with the server. Thus, statefull firewalls are capable of monitoring actual traffic patterns and ensuring that the current network traffic is legitimate. However, even the most sophisticated statefull firewalls cannot completely characterize all forms of network traffic, so what happens when a statefull firewall misinterprets a sudden burst of legitimate traffic as a flood attack?

Although if the burst is big enough, the traffic would be unsuccessful anyway! ☺

Egress filtering is an effective technique against DoS. The concept behind egress filtering is pretty simple; a network will not only monitor incoming traffic but also outgoing traffic. For example, it is possible to use egress filtering to filter out spoofed addresses. A local network could have a firewall set up to filter all packets destined for systems outside the local network but with source addresses that do not match those in the local networks domain. While this technique is implemented by many as a "good neighbor" policy, very few sites feel comfortable leaving their security up to outside parties.

Investing in more bandwidth and load balancing is really just a short-term solution to flood attacks. All a hacker has to do is find more zombies to do the work for him.

Traditional defenses against DoS attacks were focused at the downstream level. However, large-scale flood attacks consume bandwidth upstream of the target, making downstream filtering useless. Most websites have some sort of bottleneck between themselves and their service provider.

This is because the pipe that connects the service provider and the site's network is some portion of the total bandwidth the service provider provides to its customers. If a flood attack is big enough to fill the pipe between the ISP and the local network, then filtering on the side of the local network still won't help anyone from another network that is trying to reach the attacked network.

The latest trend in thwarting DoS prevents attack traffic from exiting the service provider's backbone and entering the victim's local network. This type of defense against DoS stops the attack upstream. Although the technology varies among the different companies that offer products, which operate at the ISP level, the basic idea is the same. These products monitor traffic at the ISP level and look for traffic patterns that are the telltale signs of DoS attacks. The products operate on the premise that DoS attacks have a certain signature that distinguishes them from legitimate network traffic. Once a DoS signature is detected, it is traced back to the source and filtered out. The general concept is similar to the way virus-scanning software looks for viruses in computer files, and it is interesting to note that the same problems associated with virus scanning software pop up when discussing the new wave of anti-DoS technology. For example, Asta Network's Vantage series product claims to be able to identify DoS attacks in such a fashion; however, in the product description offered on their website, there is talk about algorithms that minimize "false positives". Thus, it seems this new technology might suffer the same problems that statefull firewalls suffer at the network edges. [6] Is there a chance that this new technology will filter out legitimate traffic by mistake? Maybe, maybe not. But the system is not "fire and forget"; instead, it appears that Asta's anti-DoS solution is more like a sophisticated network administrator's tool. The network administrator has the last say on what gets filtered out, but then how does the network administrator know what's legit and what is not? The technology is relatively new and not exactly proven; at the moment, it has been deployed on the newly created Internet2 backbone to test its effectiveness.

Another problem to consider with this system is how it will perform against the new generation of DoS attacks. If we use the anti virus industry as a base for comparison, then it is unlikely that DoS attacks are going the way of the Dodo anytime soon. Somebody somewhere is going to figure out a DoS technique that is capable of fooling this new technology. Of course, another good point is that just because this new technology does not provide a complete solution does not mean it is worthless. Only time will prove whether the solution is worth the money.

As DoS attacks come in many different shapes and sizes, it is probably safe to conclude that a single one technology is not enough to thwart them all. Even with the introduction of new anti-DoS technology that is deployed at the ISP level, it seems like effective anti-DoS will continue to be a hodge-podge of techniques. Perhaps the best solution to the problem is to make sure that every level of the Internet does what it can to deter DoS. This idea is similar to that of say a neighborhood watch program; if all the members of the Internet community do their part, then the Internet would be a less DoS prone environment. This sort of collaborative solution is necessary because one thing history has proven is that hackers are resilient, and they will come up with newer and better DoS attacks.

References

1. CERT Coordination Center, "Denial of Service Attacks," [Online] Available http://www.cert.org/tech_tips/denial_of_service.html
2. David Moore, Geoffrey M. Voelker, Stefan Savage, "Inferring Internet Denial-of-Service Activity," [Online] Available http://www.caida.org/outreach/papers/backscatter/usenixsecurity01.pdf
3. Marco de Vivo, Gabriela O. de Vivo, Roberto Koeneke, Germinal Isern, "Internet Vulnerabilities Related to TCP/IP and T/TCP," Computer Communication Review, vol.29, (no.1), ACM, Jan. 1999. p.81-5.
4. See 2.
5. Stephen Justin, "The Changing Face of Distributed Denial of Service Mitigation," [Online] http:// www.sans.org/infosecFAQ/threats/face.htm
6. http://www.astanetworks.com/products/how/

Good job. You took a reasonable-sized question and gave a thorough answer.