The Bit and the Pendulum

Balancing the Interests of Stakeholders in Digital Publishing

Mark Stefik and Alex Silverman*
Xerox Palo Alto Research Center
Palo Alto, California 94304

 The authors would like to thank Jonathan Zaremski, Robert Shafter, Margaret Jane Radin, Richard Domingo, Stuart Card, and Harvey Brownrout for their valuable comments on earlier drafts of this article. The views and opinions expressed here are those of the authors, not necessarily those of Xerox Corporation.

The Pendulum Swings

Personal computers and computer networks have the potential to become an ideal basis for digital publishing. But the potential for digital publishing remains just that—a potential. The market for digital works remains nascent, because the medium has failed so far to balance the interests of important stakeholders. Computers and the digital medium are sometimes seen as the root of this problem (Barlow, 1994). In this article we explore how computers designed as trusted systems could bring things more into balance.

By *digital publishing*, we mean the on-line sale and distribution of digital works. A digital work can be anything in digital form: an article, a book, a program, or any multimedia combination involving programming, music, text, and video. The advantages of the digital media include nearly instant distribution, low production costs, and the convenience of 24-hour automated shopping.

When personal computers and desktop publishing first appeared in the early 1980s, many publishers saw digital publishing as being too risky. Although numerous factors influenced publishers' judgments in particular cases, the dominant and recurring factor was the fear of widespread unauthorized copying. Realistically concerned about loss of control over their intellectual assets, many publishers avoided the digital medium. From the publishers' perspective, the pendulum representing the balance of power between creators and consumers had swung too far towards consumers.

In the late 1990s, trusted systems began to appear from several vendors, including Folio, IBM, Intertrust, Xerox, and Wave Systems (Kahin and Arms, 1996). Trusted systems vary in their hardware and software security arrangements, but in general, they automatically enforce terms and conditions under which digital works can be used. For example, rights can expire after a period of time. Different people can pay different fees for using a work, depending on digital licenses for membership in such groups as affiliated book clubs. Trusted systems differentiate between different uses such as making a digital copy, rendering a work on a screen, printing a work on a color printer, or extracting a portion of a work for inclusion in a new work. When asked to perform an operation not licensed by a work's specific terms and conditions, a trusted system refuses to carry it out. So dramatically do trusted systems alter the balance of power between publishers and consumers that some observers have suggested that the pendulum has now swung too far towards publishers.

Copyright and Paper Publishing

Much of copyright practice and law derives from several centuries of experience with publishing on paper. Copyright law grants certain exclusive rights to rights holders—authors and publishers—in order to promote the creation and distribution of useful works. Among other things, authors have the exclusive right to reproduce copies of their works. This protects the interests of authors and publishers, who invest heavily in the development of works, as well as in printing, warehousing, and distribution. The law forbids a

second publisher from undermining the market by selling competing copies. Copyright law also addresses the public's interests, for example, through provisions such as the first sale doctrine and fair use.

Copyright law, by itself, does not prevent unauthorized copying. However, where perfect enforcement of copyright law has been impractical, the technology and economics of paper publishing have helped to keep infringement in check. Although photocopying makes it easy for an individual to make a single copy of a work for personal use, it generally does not promote large-scale copying and distribution by individuals. The offset presses and finishing equipment used by publishers make it difficult for individuals with photocopiers to match the quality and mass production economics of book and magazine production. Furthermore, it is usually too costly for individuals not in business to make and distribute thousands of paper copies. Thus publishers in the paper medium are protected by two practical means. Large-scale infringement by well-funded rogue publishers is addressed by legal remedies; large-scale infringement by individuals with photocopiers is inhibited by the relative economics of one-at-a-time versus mass production. Together, these forces have created a point of balance between publishers and consumers for paper-based publishing.

Copyright and Personal Computers

The balance of power from the paper medium does not directly translate over to the digital medium. Even though digital works can be expensive to develop, copying is essentially free. Using a personal computer, a consumer can copy a digital work as cheaply and with the same quality as a publisher. Furthermore, during the PC revolution, community support for copyright has been crucially different for paper and digital media. Whereas community support for copyright in paper media is strong and well established, community support for copyright in digital media is weaker, or even absent.

Many key institutions have long been active in advocating copyrights for works on paper. In the United States, these have included the Library of Congress, the Copyright Clearance Center, and the American Association of Publishers. Worldwide, although there are international differences in copyright law, an international body (WIPO, the World Intellectual Property Organization) has been active in promoting treaties and standards to harmonize national laws. Overall, copyright has worked well enough to support the established book, magazine, and newspaper publishing industries.

As paper publishers have begun to consider digital publishing, however, they have faced resistance from many people in the computer community who deeply believe that computer software and information ought to be free. Although this attitude may seem contrarian, it has grown naturally out of the environment and prior experience of the computer community, as a brief historical sketch will explain.

Before the rise of personal computers in the 1980s, the use of computers for digital publishing was mostly the province of the academics and scientists, who used the same computers for publishing as for their scientific work. Publishing in the academic community is a special case, in that free sharing of results is

rooted in a philosophy favoring competition of ideas and the open search for knowledge. Academic journals are also unusual in that the structure of payment for works is the exact opposite of the commercial publishing case: Authors for whom scholarly citation and academic reputation is important for career advancement often pay to be published through page charges, rather than being paid for their writings.

Software in the computer-using academic community was created by and freely exchanged among many groups of developers. Having one's software widely used was a way to build reputation and to establish one's place in the community. This was a natural extension for a community in which being read was key to advancing an academic career. Free exchange of programs, starting in the 1960s, accelerated the development of the field of computer science, a field directly concerned with the creation and study of the algorithms. Thus, for computer science, sharing of software was a natural part of the academic processes for sharing results, community building, and peer review.

The PC community inherited from the academic community the values supporting free exchange of information. When the PC community first emerged as a hobbyist fringe of academic computer science, these values served it well and hobbyists built their computers and traded programs. However, as the personal computer culture matured, the underlying assumptions supporting free sharing no longer applied. Few computer users created software or built on each other's software; instead, they bought software from publishers. Publishers did not rely on academic grants for support; they had to make a living selling their software.

To keep prices down, software publishers amortized the cost of support and production over their user base. Software publishers tried several measures to cope with unauthorized copying. In particular, various copy protection schemes were tried, such as rigging computer disks in various ways to make them difficult to copy. However, even legitimate customers found the protection approaches too inconvenient and they were eventually dropped. Ultimately, very large publishers learned that copying could be good for business. People would get trained in using pirated copies of their software and later would buy legitimate copies when new versions were released. Indeed, this process led to a market dynamic that helped big publishers to dominate and marginalize smaller ones.

Finally, software publishers operated in a legal regime in which the strength of safeguards against copying of digital works appeared to be, at best, uncertain. Courts pondered the extent to which copyright ought to protect the structure, sequence, and organization of computer programs (*Computer Associates v. Altai*[1]). Certain well-known computer copyright cases, such as Apple's dispute with Microsoft over the Macintosh graphical user interface (*Apple v. Microsoft*[2]), took many years to resolve. Meanwhile in other electronic venues, such as videocassette recorders, court challenges suggested that certain kinds of copying were

---

[1] *Computer Associates International, Inc. v. Altai, Inc.*, 982 F.2d 693 (2nd Cir. 1992).
[2] *Apple Computer, Inc. v. Microsoft Corp.,* 35 F.3d 1435 (9th Cir. 1994).

permissible under the grounds of fair use (*Sony v. Universal Studios*[3]). When the early satellite broadcasts of television were widely compromised by unauthorized descrambling devices, Congress refused to pass laws to regulate copying until the broadcasters used adequate technological measures to thwart copying.

Comparing computers to leaky bottles, John Perry Barlow (Barlow, 1994) argued that once a work is digital, it can inevitably be copied. All these events led to the hyperbole that copyright had failed.

Copyright and Trusted Systems

Beginning in the 1990s, it was realized that computers could become part of the solution to the copyright problem that they were said to cause. The key was the development of trusted systems technology.

There are two main ideas behind trusted systems: that the terms and conditions governing the authorized use of a digital work can be expressed in a computer-interpretable language, and that computers and software can be designed to enforce those terms and conditions (Stefik, 1997a,b). An example of a rights language is Xerox's DPRL (Digital Property Rights Language).

Digital rights cluster into several categories. Transport rights include rights to copy, transfer, or loan a work. Render rights include playing and printing. Derivative work rights govern extracting portions of a work, controlled editing of changes to it, and embedding of the portion in other works. Other rights govern the making and restoring of backup copies. With trusted systems, a publisher can assign rights to a digital work. Each right can specify fees that must be paid to exercise the right. Each right can specify access conditions that govern who can exercise the right.

Trusted systems enforce the terms and conditions. They also exchange copies of the work only with systems that can prove themselves trusted via challenge-response protocols (Summers, 1997). In exchanging digital works, trusted systems form a closed network of computers that exclude non-trusted systems and collectively support use of digital works under established rules of commerce. When digital works are sent between trusted systems, the works are encrypted. When digital works are rendered—by printing them on paper, displaying them on monitors, or playing them on speakers—the rendering process can embed machine-readable watermark data in the signal to make it easier to trace the source of any external copying of the works.

In general, the higher the security of a trusted system, the higher its cost. High-security trusted systems can detect any physical tampering, set off alarms, and erase secret key information inside. Intermediate security trusted systems have more modest physical, encryption, and programmatic defenses. Using challenge-response protocols, trusted systems have the capability to recognize other trusted systems and to determine their security levels. For any particular work, publishers can specify the security level required by a trusted system that can receive it. An expensive industry report might require an expensive and secure corporate

[3] *Sony Corp. of America v. Universal City Studios, Inc.,* 464 U.S. 417 (1984).

trusted system with advanced security measures. A digital newspaper for wide distribution and subsidized by advertisements might require only modest security measures for home computers.

Trusted Systems and the Balance of Interests

There are many stakeholders in digital publishing. Beyond the government itself, U.S. copyright law focuses on two parties or categories of people: rights holders (that is, the authors and publishers who hold the copyrights) and the public. However, trusted systems delegate enforcement and control to computers. One of the effects of this delegation is that it introduces third parties to the arrangement, including distributors, trusted system vendors, financial clearing houses, and multiple governments. This complicates the balance of interests, in that it introduces more parties whose interests need to be considered.

The use of trusted systems to enforce terms and conditions provides a much finer grain of control than copyright law, and moves the legal basis of protection in the direction of contracts and licenses. The finer grain of control includes distinctions between different kinds of usage rights such as copying, loaning, printing, displaying, backup, and so on. It also includes provisions for identifying specific users, specific kinds of devices for rendering, and fees for uses. Further, trusted systems provide a finer grain of control opens in that it becomes possible for rights holders to monitor and negotiate over transactions in copyrighted works in situations where, in the past, such monitoring and negotiation would have been impractical, if not impossible.

Below we consider the sometimes competing interests of the stakeholders, together with technological and institutional implications for digital publishing and trusted systems. In particular, we contrast copyright-based and contract-based protection of digital works and examine how issues such as fair use, liability, privacy rights, and national borders play out in the context of a trusted systems regime.

Copyright Law

Excellent summaries of U.S. copyright law are available elsewhere (Schlachter, 1997; Dept. of Commerce, 1995) and we will not recapitulate them here. However, it is useful to explain some key aspects of copyright law in order to understand stakeholder interests and the design of trusted systems.

Like other forms of legal protection for intellectual property including patent and trademark law, copyright law encourages the creation of intellectual property by granting certain exclusive rights to its creators. Article 1, section 8 of the U.S. Constitution authorizes Congress to create legislation "to promote the Progress of Science and useful Arts, by securing for limited Times to Authors ... the exclusive Right to their respective Writings." Over the years, Congress has enacted a series of copyright laws, beginning with the Act of May 31, 1790. The most recent major overhauls of the copyright law have been the Copyright Act of 1909 and the Copyright Act of 1976.

Through copyright law, the federal government grants to a copyright owner certain exclusive rights in his or her copyrighted work. These include the right to reproduce the work in copies, to prepare derivative works

based on the copyrighted work, and to distribute copies to the public by sale or other transfer of ownership, or by rental, lease, or lending. For certain kinds of works, such as literary, musical, and dramatic works, exclusive rights to perform the work publicly and to display the work are granted as well. The rights last for relatively long periods of time. For works created after January 1, 1978, copyrights cease 50 years after the death of the author or, in the case of works made for hire, the earlier of 75 years from the date of publication or 100 years from the date of creation. (Legislation now pending in Congress would lengthen the copyright term to run until 70 years after the author's death or, in the case of works made for hire, the earlier of 95 years from the date of publication or 120 years from the date of creation.) Although the duration of copyright is long, it is finite, and once it is over, the copyright owner's exclusive rights end. Thereafter the work falls into the public domain, and anyone may copy it freely.

Copyright law represents an attempt to strike a balance among the competing interests of various stakeholders. Most obviously, there is the balance between the interests of the rights holders and the public. Copyright law addresses this balance by limiting the exclusive rights that it provides. The legal protection afforded by copyright is limited in time, as described above. It is also limited in scope. For example, the Supreme Court has held that copyright protection does not extend to the "sweat of the brow" invested by the work's creator, but only to the original contributions of authorship. Thus an alphabetical white pages telephone directory may lack sufficient originality to be protected by copyright, even if the compilation of its content required considerable effort (*Feist v. Rural Telephone*[4]; Samuelson, 1997). Further, copyright protection extends to expression but not to ideas. Thus in the area of computer software, courts have held that although the code of a computer program can be (narrowly) protected by copyright, the program's functionality can be protected, if at all, only by patent or trade secret law, not by copyright (*Sega v. Accolade*[5]*; Atari v. Nintendo*[6]).

Further in striking the balance between rights holders and the public, copyright law provides that the fair use of a copyrighted work for purposes such as criticism, comment, new reporting, teaching, scholarship, or research, are not infringements of copyright. We will consider fair use at length in a later section. Additionally, provisions of the law (sections 108 to 120 of the Copyright Act of 1976) establish a framework—some might say, a patchwork—of more specific rights limitations, scope restrictions, and licensing arrangements. Many of these provisions are designed to address the concerns of particular interest groups, such as religious organizations, small businesses, lending libraries, blind and handicapped persons, cable television stations, and noncommercial broadcasters.

The history of copyright law has been, in part, a history of changes in the law striving to keep pace with changes in technology and media. In 1790, the copyright law governed books and charts; today, it governs not just writings on paper, but any original works of authorship fixed in a tangible medium of expression,

[4] *Feist Publications v. Rural Telephone Service,* 499 U.S. 340 (1991).
[5] *Sega Enterprises, Ltd. v. Accolade, Inc.,* 977 F.2d 1510 (9th Cir. 1993).

such as motion pictures, architectural works, sculptural works, and sound recordings. Nevertheless, the major revisions of the copyright law are few and far between (almost 70 years between the 1909 and 1976 Acts), and even the lesser changes through statutory amendment, regulation, or case law can take years to accomplish and so can lag behind the latest technological developments. In short, the law tends to fall behind the technology.

This has become especially apparent in the era of modern digital media. Digital technologies change at an unprecedented rate. Moreover, digital media blur the boundaries between traditionally separate categories of work, so that statutory provisions intended to apply to a particular type of work or serve a particular set of interest groups can rapidly and unexpectedly become applicable elsewhere. Digital media can create confusion over what type of a work something is. For instance, is a Java program that causes display of an animation on a Web page an audiovisual work or a computer program, or both, or something entirely new? Or consider that in packet-switched computer networks, such as the Internet, the content-bearing information packets sent between computers can travel on cables or over the airwaves. The choice of transmission path depends on routing conditions in a manner that is unknown or even unknowable to the users. Should provisions of the copyright law designed for cable television or for broadcast transmissions apply here? Does it make sense to define a single concept of "transmission" in the digital media and, if the answer is less than clear, should the law provide a new exclusive right associated with transmission, as some have proposed (Dept. of Commerce, 1995)? Yet another gray area: As an information packet travels across the Internet, are the bits in the packet sufficiently "fixed in a tangible medium of expression" so that the packet itself may constitute a work subject to copyright? And so forth.

Digital media can also create confusion about who is the creator or owner of a work. For example, digital sampling is a technique in digital audio processing that makes it possible to take sounds (possibly from copyrighted works), to process them in various ways, and to include them in other derived works. The source or sources of the sampled works may or may not be recognizable in their processed form. In particular, the samples can be very short, perhaps a single drumbeat or a single note of music. In the absence of trusted systems, determining the original source from which such a short sample was taken can be difficult or impossible. Even if the source can be determined and is found to be someone else's copyrighted work, it is not always clear whether the person doing the sampling should be required to pay a copyright royalty to the owner of the earlier work, or whether and in what circumstances such borrowing of short segments may constitute fair use. So digital sampling—or, more aptly, the lack of fine-grained control over digital samples in the absence of trusted systems—has posed difficult issues for the recording industry and recording artists alike. These issues have become even more complicated now that samples can be sent and traded across the Internet among thousands of musicians worldwide.

---

[6] *Sega Enterprises, Ltd. v. Accolade, Inc.,* 977 F.2d 1510 (9th Cir. 1993).

Without trusted systems, effective enforcement of copyright in the digital medium can be nearly impossible. Like a proverbial sieve with thousands of little holes, it is too hard and expensive to find all the little infringement leaks of isolated individuals making copies. Furthermore, living with the leaks has its own deep risks. By publishing without copyright enforcement in a community that routinely makes unauthorized copies, rights holders risk that, over time, such copying could become established as common practice or even sanctioned by courts as fair use.

In sum, the move toward digital media poses challenges for copyright law and creates uncertainty for rights holders, especially for would-be publishers in the new media. The uncertainty tends to hamper the adoption of the new media and to discourage publishers from publishing in it. The impracticalities of enforcing copyright on untrusted, networked systems, the gray areas of legal interpretation for digital works, the lack of fine-grained control in copyright law, and the risk of an emerging legal claim of fair use for digital copying all motivate would-be authors and publishers in the digital medium to find means other than copyright law for protecting their interests.

Contract Law and Digital Contracts

In a representative scenario of digital publication on a trusted system, an author begins by creating a digital work. When the work is ready, the author finds a publisher (possibly himself) to develop the work further and to sell it. The publisher develops a set of terms and conditions for use of the work. Using a rights management language like DPRL, the publisher specifies the time period over which the rights apply. He determines what rights to include, for example, whether printing is allowed, whether the work can be loaned out for free, whether there is a special discount for members of a particular book club. He may assign different fees for different rights. For example, he may decide either to disallow creation of derivative works or to encourage creation of such works as a source of further revenue. He may mandate that the reader of the digital work must have proof in the form of a proper digital certificate that he is over 18. In DPRL (Stefik, 1997a), each right specification statement includes a type of right, a time specification governing when the right is valid, an access specification governing any special licenses required to exercise the right, and a fee specification governing billing. Using a trusted system, these rights are associated with the digital work, either by bundling them together in an encrypted file or by assigning the work a unique digital identifier and by registering the work and its rights in an on-line database.

Why would a publisher or a consumer want to have a specific and detailed agreement about use? The alternative, based on copyright as used for most printed works, is to have a single fee to purchase a work and then general legal standards about how works can be used. In the previous section we considered motivations for publishers to use specialized terms and conditions for their digital works. Specialized rules also have potential economic advantages for consumers. In the established software market, a software license is typically purchased for a fixed fee. This means that a user who expects to make little use of a software product must pay the same fee as someone who would use it for many hours a day. In some

markets, this situation is bad for both publishers and consumers because many low-usage consumers will decide not to purchase the software at all. Trusted systems offer the possibility of differential pricing and "metered use" in which the amount that someone pays to use software depends on how much they use it. One way to look at metered use is that it allows "renting" software, where the rental terms can be flexible enough to provide for decreasing costs or caps with increased volume of use.

Another example of mutual economic advantages concerns the first sale doctrine. When consumers buy a paper book, they receive and own the copy of the book. When they are done with the book, they are free to give it to a friend or to sell the book to someone else. The first sale doctrine from copyright law guarantees these rights. In the DPRL language, the analogous usage right is called a *transfer right*. When one trusted system transfers a digital work to a second trusted system, the copy on the first trusted system is deleted or deactivated so that it can no longer be used. Analogous to handing a book to a friend, a transfer operation preserves the number of usable copies of the work. Analogous to the first sale doctrine, the terms and conditions on a digital work could allow it to be transferred at no charge.

A free transfer right is exactly what a consumer might want if he or she were buying the digital work for a friend, or intended to share the work serially with others. On the other hand, from a publisher's perspective, a free transfer right is a threat to future sales. If each person who reads a copy of a digital work needs to buy their own copy, the publisher would sell more copies. A publisher could offer two different combinations of rights with a work. In one combination, the consumer pays the "standard" amount for a work, and can transfer the work without fee just as with the first sale doctrine. In another combination, the consumer gets a discount for a non-transferable work or must pay a fee to transfer it. This discounted purchase might be preferred by a consumer who buys the work for his personal use and who does not anticipate giving it away. Arguably, first sale doctrine is grounded in experience with paper-based works and the copies were treated as physical objects, independent of their creative content. Like tools or food, such physical objects could be resold at the owner's convenience. Enforcement of a law to prevent resale or giving of books would be difficult in any case, so the first sale doctrine makes sense for paper-based works. For digital works and trusted systems, these considerations are less relevant. The publisher and the consumer are free to enter into an agreement that each sees as economically advantageous.

In many ways, a set of terms and conditions in DPRL is much like a contract or license agreement for using a digital work. For convenience here, we will call such a set of terms and conditions a *digital contract*. However, it should be remembered that a digital contract differs from an ordinary contract in crucial ways. Notably, in an ordinary contract between people, compliance is not automatic and is the responsibility of the agreeing parties. There may be provisions for monitoring and checking compliance with the terms and conditions, but the responsibility for acting in accordance with the terms falls on the parties, and enforcement of the contract is ultimately the province of the courts. In contrast, with trusted systems, a substantial part of the enforcement of a digital contract is carried out by the trusted systems themselves. In the short term, at least, the consumer does not have an option to disregard a digital contract, for example, to

make infringing copies of a digital work. A trusted system will refuse to exercise a right that is not sanctioned by the digital contract. Over the longer term, it may be possible for consumers or consumer advocacy groups to negotiate with publishers to obtain different terms and conditions in the digital contracts, but even then, the new digital contracts will be subject to automatic enforcement by trusted systems.

Contract law is a complex subject that we will not try to summarize here. We note that as of this writing, there is an ongoing, controversial effort to add new provisions concerning "licenses of information and software contracts" (UCC Article 2B, section 2B-103(a), Draft of May 5, 1997) to the Uniform Commercial Code, which is the body of statutory law that governs commercial contracts in most U.S. states. If, eventually, these proposed new provisions are adopted into law, they could have significant implications for digital publishing in general and for trusted systems in particular. Even so, it is worthwhile to point out some of the basic provisions of contract law in order to understand stakeholder interests and the design of trusted systems.

Contracts are agreements entered into by two or more parties. In a representative case, the parties negotiate and come to an agreement on terms and conditions under which each party provides economic value to the other. This bargain for a mutual exchange of value is an important part of what makes an agreement a contract. (Technically, a contract includes one or more promises backed by what is legally referred to as "valid consideration.") Typically, the terms of the parties' agreement are set out in a written document, and the parties formalize their agreement by signing and dating the document. In cases warranting the extra care, the contract document may also be notarized by a registered third party who checks the identity papers of the parties and who may also take thumb prints or require other forms of personal identification from the parties.

A contract is backed by the force of law. Generally, each party can enforce the contract through the courts in the event of another party's failure to comply with the agreed-upon terms and conditions. However, there are various circumstances under which the terms and conditions of a contract are not legally enforceable. The courts have developed various doctrines (unconscionability, fraud, illegality, contracts of adhesion, contracts void as against public policy, etc.) to prevent enforcement of certain contracts that are deemed unfair or improper. Also, the provisions of the Constitution and the provisions of certain statutory laws may preempt some contracts. In particular, the copyright law contains a preemption provision (section 301 of the 1976 Copyright Act) that may, in some cases, render certain contracts unenforceable.

Here is an illustrative example of an unenforceable contract. It is not unusual for landlords to offer standard rental agreements. These documents may consist of several pages of formal "legalese." Suppose that a tenant signs a document without reading it carefully. He fails to notice a clause in small print saying that if he eats mushrooms on Tuesdays he must pay an additional $1000 in rent. The landlord throws a party on

the next Tuesday and slyly offers the tenant mushrooms. Such clauses are outside the normal scope of what would be found in a rental agreement, and would very likely be held unenforceable by a court.

Courts provide checks and balances in contract law by deciding what contracts to enforce and how to interpret the terms and conditions of those contracts. With properly designed trusted systems, many of these checks and balances can be made available automatically. Consider a digital publishing scenario again. The author has finished the work and the publisher assigns terms and conditions. Just as there can be conventional (so-called "boilerplate") language used in putting together an agreement, there can be digital boilerplate in the form of templates and default conditions in setting up a digital contract. Suppose that the publisher has included some very unusual terms and conditions in the agreement. When the consumer's trusted system is in communication with the publisher's trusted system, it can first retrieve the terms and conditions of the digital contract. It shows these to the consumer. Before the consumer accepts receipt of the digital work, a program can check for and highlight unusual conditions in the digital contract. Because rights management languages like DPRL are simple and formal languages with limited complexity, simple grammar and predetermined meanings, this checking is straightforward for a computer. In particular, the contract checker can look for unusual or high fees on certain rights, unrealistic expiration dates, or any other requirement that is outside of the usual practice. (As a somewhat bizarre example, consider a digital work that the consumer can copy for free but, surprisingly and inconveniently, costs $10 to delete.) The consumer is given an opportunity to agree to the terms and accept delivery or to refuse the terms and not take delivery of the work. If the consumer agrees, his trusted system can digitally sign a form marking his agreement to the contract. This signing can be digitally notarized by a third party (a "digital notary") known to both parties.

The sequence of events in this example illustrates several checks and balances in the process. Both the publisher and the consumer can use computational aids to check the normalcy and appropriateness of the contract. More than being a labor-saving or time-saving procedure, this approach also helps to compensate for the somewhat less tangible nature of information inside computers. It gives increased confidence to both parties that the terms and conditions used by the trusted systems will be reasonable.

It is helpful to think of a digital contract as encompassing several distinct legal contracts. There is the contract for access to the copyrighted work itself. Further, there is a contract for the service of delivering digital data to the consumer, irrespective of whether that data is or can be copyrighted. For example, if the publisher provides an uncopyrightable database (*ProCD v. Zeidenberg*[7]) or telephone white pages directory (*Feist*) to the consumer via a trusted system, the publisher can fairly charge for this service, even though the consumer could, in principle, get the uncopyrightable data elsewhere or put together the database himself. Similarly, the digital publisher could charge the consumer for a copy of the complete works of Shakespeare even though that is in the public domain, just as print publishers can charge for printed copies of

[7] *ProCD, Inc. v. Zeidenberg,* 86 F.3d 1447 (7th Cir. 1996).

Shakespeare's works. Put another way, like the print publisher, the digital publisher has made life more convenient for the consumer, and the consumer pays the publisher for this convenience. What the publisher of an uncopyrightable or public domain work cannot do is to prevent another publisher from offering consumers the same or a comparable work, as would be the case if the work were copyrighted. Finally, there is a third contract implicit in the digital contract, namely, the service agreement by which the consumer is entitled to access the network of trusted systems in the first place. This agreement may be arranged between the consumer and the publisher, or between the consumer and one or more network service providers who may or may not be affiliated with the publisher.

The idea that a digital contract includes multiple legal contracts provides a coherent rationale for why digital contracts ought to be enforceable, even as to uncopyrightable works. For example, suppose that a publisher provides a public domain work, such as the complete works of Shakespeare, to the consumer via a trusted system. However, the digital contract for this work prohibits the consumer from copying or further transferring the contents of the work, at least not in digital form. The consumer is unhappy about this. He knows that the work is not protected by copyright and, when the bill arrives from the publisher, he refuses to pay, or else sues to get his money back. In court, the consumer argues that for the publisher to charge for what is no longer protected by copyright is in violation of the policy of copyright to establish limited-term monopolies for authors. Therefore, says the consumer, the digital contract should be preempted by the Copyright Act and should be held unenforceable. (The consumer might also argue that, because the publisher accepts the consumer's money while providing in return only a public domain work that ought to be available for free, the agreement fails for lack of consideration). The publisher responds that what is being sold here isn't the work, but rather the service of delivering the work. The publisher says, in effect, "Consumer, by dealing with me, you save time and energy and money over other delivery mechanisms such as conventional bookstores. But if I, as vendor, want to continue successfully to provide this service to others, then I am entitled to collect revenue at every transaction, not just the first one. Therefore, I can legitimately prevent you by this digital contract from transferring the copy of the work I just sold you." We think that the publisher has the better argument here. The consumer can pay the publisher for the right to print out the contents of the book, and can then copy the contents, for example, by hand or by scanning with an untrusted optical scanner. Also, other publishers can produce similar books containing identical texts, and a not-for-profit library could make these texts available for free. In short, the publisher has not overstepped the bounds of copyright.

Another point of possible concern with a digital contract is the extent to which a user is realistically in a position to negotiate the terms of the contract. In court cases concerning the viability of shrinkwrap licenses (*ProCD*; O'Rourke, 1997), one of the legal arguments used to challenge the validity of the license is that a publisher has an advantaged position of power and leaves the user with only a "take it or leave it" proposition. In this situation, many consumers do not bother to read the shrinkwrap license. In the case of trusted systems, it may be important that a consumer agent could be called upon to highlight terms and

conditions likely to be unacceptable. In principle, one of the options when such terms are found is for the trusted systems to open a channel for negotiation and possible change of the terms. It is worth noting, however, that one of the main advantages of digital publishing is the possibility of fully automated systems providing 24-hour shopping convenience. In that setting, one might not expect to negotiate the terms of purchase for a mass-market digital work any more than one would expect to negotiate the price of buying a best-seller paperback at a convenience store in the middle of the night. The consumer would simply have to either accept the terms as they stand, or postpone her purchase until such time as a human agent became available to negotiate the terms.

Trusted Systems and Fair Use

In addition to specific statutory exceptions to the exclusive rights provided to rights holders, section 107 of the Copyright Act of 1976 sets forth four factors to be considered in determining whether an otherwise-infringing use of a copyrighted work is to be considered a "fair use":

1. the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes;
2. the nature of the copyrighted work;
3. the amount and substantiality of the portion of the work used in relation to the copyrighted work as a whole; and
4. the effect of the use upon the potential market for or value of the copyrighted work.

Technically, fair use is a defense that can be raised in the event that a copyright owner challenges a use of a copyrighted work. In a representative case, the fair use defense works as follows: A copyright owner publishes a copyrighted work. Without seeking permission, a second party obtains a copy of the work and incorporates portions of it in a new work. The copyright owner objects and takes the second party to court. In court, the second party argues that the use is a fair use and should be permitted notwithstanding the copyright. For example, the second party might argue that the use was made for purposes of legitimate commentary or satire. The copyright owner argues that the use is not a fair use and should be prohibited. The court is supposed to decide on the basis of all four factors whether the use is a fair use. The outcome depends on the particular facts of the case. In practice, the fourth factor, about undermining the market for the work, is often considered the most important.

One of the concerns that some have raised about trusted systems is that they might preclude consumers' access to works on a fair use basis. For example, if a consumer tries to extract a portion of a digital work using an editor in a trusted system, the system will prevent the operation. Therefore, the consumer never even has an opportunity to commit the act of copying that would occasion the fair use defense. This is not to say that a trusted system would preclude a user from manually retyping portions of a copyrighted book or from digitally recording excerpts from a copyrighted audio or video work using a computer's microphone

or digital camera. Rather, it means that the operations of cutting and pasting directly from a rights-protected digital work with an editor would be automatically disabled if the consumer of the digital work did not have appropriate derivative work rights. For trusted systems, this example shows how the pendulum has swung towards giving more power to authors and publishers.

Arguments about fair use for digital works sometimes tacitly (and incorrectly) assume that risks in the digital medium are similar to risks in the paper medium. However, in the paper medium, as discussed earlier, it is unlikely that a consumer will make and distribute thousands of infringing copies of a digital work. In the digital medium, an infringer can copy and mail a thousand copies of a digital work with no expense and a single keystroke. Restated, in granting unrestricted access on the fair use principle to each and every anonymous user, a publisher would routinely risk loss of all content assets.

The other side of the coin is that fair use serves the public interest, in particular by helping to ensure freedom of speech. Parody, academic and social criticism, satire, and other forms of speech that rely on the ability to quote from, paraphrase, and modify portions of other people's copyrighted works are essential ingredients in the mix of speech that contributes to a democratic society. Fair use helps to ensure that such borrowing from others' works is possible. The holder of a copyright may be very much against letting someone else use even portions of their work if the purpose is to critique or lampoon that work. Without fair use, rights holders could effectively quash criticism by preventing the critics from publishing. The importance of this public interest aspect of fair use is amplified in the digital medium, because the ease of wide-area communication in the medium promotes the ability of everyone in the society to engage in such speech and to be heard by all. Ultimately this benefits society as a whole.

Is there a way in the digital medium to balance the risks and benefits for publishers and consumers around fair use? One approach adopts the view that pricing and market forces will, in many instances, make the fair use issue moot in the digital medium. Currently, fair use is a binary decision: Either the use is a fair use or it is not. Thus, whenever the law dictates that a particular use is a fair use, rights holders must necessarily forfeit income for that use; if not, consumers are obliged pay for that use. This creates high stakes for cases bearing on fair use. The issue is especially awkward in cases where the cost of being honest greatly exceeds the expected fee for using the work (Stefik, 1997b). In contrast, with trusted systems, fees for use of copyright works in the digital medium can be either large or small and can be collected automatically. Market forces can prevail so that in many cases where there might otherwise be fair use, disputes will be resolved as fees reach an appropriate level.

An example illustrates. Suppose that a consumer wants to copy a short portion of a song from an audio CD album owned by her friend. Today, the consumer makes a tape recording for which the record company and the recording artist are not compensated. The consumer will probably never be found out by the record company but, even if she were, she would argue that copying only a short part of the song was fair use. The company would likely disagree. The point is that the situation is all-or-nothing: Either the consumer has to

purchase the entire CD just to get a few seconds' worth of audio, or else the consumer pays nothing (since her use is a fair use) but the cumulative losses to the record company over thousands of such consumers are substantial. With trusted systems, though, a middle ground is possible. The fine-grained control offered by trusted systems make it possible for the consumer to purchase exactly the portion of the audio CD that she wants, and for the record company to charge and collect a fair price for that audio. In other words, to the extent that fair use is a response of the copyright law to market failure (Merges, 1997), well-designed trusted systems can help to correct that market failure, so that the fair use issue simply disappears.

Here is another example of the microtransactions approach to fair use. Today, a person who buys a computer game or other software program for home use expects to be able to share the software with other family members. In a household with several computers, this sharing may be accomplished by making multiple copies of the software, one for each computer. Such copying today is typically in violation of the shrinkwrap license that accompanies the software. Nonetheless, the family might assert that such copying is fair use, because in practice, only one copy of the software gets used at a time, so that in effect the household members are just sharing a single copy of the software. In a trusted systems regime, would software publishers prevent this practice by requiring that each individual family member pay a separate fee for use of their particular copy of the software? Such pricing practices would substantially increase the family's costs, with no apparent benefit to the family. The answer is, not necessarily. Software publishers could instead keep prices as they are today, even with trusted systems. For example, a publisher could sell software intended for household use with built-in digital contract provisions designed to make it easy to share the software among family members. These provisions would allow a certain number of copies to be made for free, but the copies would be restricted for use only by the family and only on household computers. If a family member attempted to provide one of their "free" copies to someone not in the family, or to use their copy on a trusted system at work, they would be prevented from doing so unless an additional fee was paid.

In yet another instance of the microtransactions approach to fair use, consider how the now-common practice of making a photocopy of a single page of a book owned by someone else might change under trusted systems. Today, people typically make such copies without publishers' knowledge. So long as the copying is fairly minimal, it is popularly considered (rightly or wrongly) as fair use. With trusted systems, a publisher could discourage unauthorized copying and, at the same time, benefit consumers. For example, the publisher could charge less for the right to view a page of a digital book on a display screen than for the right to make a print out a hardcopy of the page. This would benefit the reader who just wanted to look at the one page rather than the whole book. The charge to the reader for viewing, but not printing, the page might be made cheaper than the cost to the reader of making a traditional photocopy. As another example, the publisher could charge more for the right to print the page in a form containing a machine-readable watermark than to print the page without the watermark. The watermarked version of the printout would be less prone to unauthorized copying, because trusted digital photocopy systems would detect the watermark

and charge for the privilege of copying. Although consumers could still make copies on older photocopiers, the digital watermark would inhibit at least some unauthorized copying.

A second approach to fair use, one not grounded in faith in market forces or microtransactions, is to institute fair use licenses (Stefik, 1997b). In a representative scenario, Joe, a consumer, applies for a fair use license, in the way he might apply for a driver's or a radio operator's license. To earn the fair use license, Joe studies the rules of fair use and must pass an examination by an appropriate organization that we call the DPT (for Digital Property Trust). The DPT certifies Joe's identity and issues him a physical certificate as well as a personalized digital license for use on Joe's trusted system. A publisher publishes a digital work. In publishing a work, a publisher assigns privileged rights that can be exercised by a fair use licensee. The publisher also declares an insurance limit for the expected commercial value of his rights to the digital work. When Joe exercises a copy transaction to obtain a copy of the digital work, the transaction fee includes an additional small amount, which is a share of the insurance premium on the digital work. Given his fair use license, Joe can exercise privileged rights on the digital work not normally available. His privileged uses are perhaps open to monitoring, logging, and reporting, subject to appropriate considerations for his privacy. Suppose that Joe now sends thousands of usable copies of the digital work to a mailing list on the Internet. The publisher takes Joe to court, claiming damages beyond Joe's ability to pay. The court takes into account the digital contract and the four factors of fair use. If the court finds in favor of the publisher, the fair use insurance pays at least part of the damages. If the court finds in favor of Joe, the fair use insurance pays for some or all of Joe's court costs and attorney fees.

The main point of this example is to illustrate the different risks and interests surrounding fair use in the digital medium. Fair use is treated here as a licensed privilege analogous to a driver's license, rather than as a legal defense. From a legal perspective, this is a substantial reframing of fair use intended to take into account the greater risks of misappropriation in the digital arena. The example implicitly raises several interesting policy issues:

- Does Joe pay for his own license? If fair use is seen as essentially related to free expression, then fair use licenses could be free, subsidized in some way by taxes or by the publishing industry. If fair use is like a driver's license, the potential licensee pays for it, but the cost is modest.

- What rights does a fair use license grant? In the example above, the publisher decides what additional rights go with a fair use license. In an alternate scenario, there is a standard set of rights, possibly defaulting to zero fee versions of all possible rights. Fair use insurance covers the financial risks to the publisher if the work is turned loose on the Internet.

- Who pays for the insurance? In the example above, a per-transaction fee levied on all consumers of the digital work can pay the cost of insurance. This has the advantage of natural scaling with the popularity of the work. Alternatively, the publisher can pay for insurance costs, although that may amount to almost the same thing as a per-transaction fee if the costs are passed along to consumers.

For self-publishers, there would be an advantage of per-transaction-based collection of insurance premiums, because the publisher could avoid up-front payment of the premium. In still another alternative, the fair use licensee can pay insurance premiums.

- Who does fair use protect? In the example above, fair use is intended to protect the publisher against losses. Should there also be a kind of liability insurance or fair use bonding for consumers to guard against claims for damages by publishers that their use was not fair use? Should there be deductibles on such insurance?

- Should fair use actions be monitored? One position is that this violates privacy. Another position is that this makes it more feasible to detect privileged violations of a fair use license. An intermediate position is to log and encrypt actions by the fair use licensee and to make these records available to law enforcement parties only when there is appropriate reason to believe that the law has been broken.

In summary, the risks to a publisher in the digital medium for unencumbered fair use are much greater than they are in paper. In this section we considered two approaches to preserving the spirit of fair use in the digital medium. One approach relies on market forces and microtransactions to make most uses of a digital work very inexpensive. Another approach, more attuned to the possibility of zero-cost fair use (and perhaps more relevant to the safeguarding of free speech), balances the risks and benefits by institutionalizing fair use as a licensed privilege with fiduciary responsibility backed by insurance.

A recurring theme in discussions of fair use and trusted systems is the fear that publishers will tend to use trusted systems to take advantage of consumers, as by unfair pricing, and that consumers will be unable to mobilize successfully to combat this trend. We believe, however, that trusted systems must serve everyone's interests, or they won't serve at all. Publishers and consumers alike will be better served if publishers use trusted systems in a way that recognizes and responds to legitimate consumer expectations, for example, by creating digital contracts that comport with traditional notions of fair use. Publishers can either choose to self-regulate, or else they risk being regulated by outside forces: the legal system, the marketplace, and public opinion. Those publishers who fail to consider consumers' interests may find themselves under attack by free speech and civil rights organizations, consumer advocacy groups and boycotts, commentators in the media, and public outcry. Effective regulation may then emerge from legislative action. Also, if digital publishers get too greedy, market forces will tend to push back. Other digital publishers will offer better deals. Or, consumers will simply prefer works published in more traditional forms to digitally published works, so that the market for digital publishing remains limited, to everyone's detriment. Accordingly, publishers who want to promote the growth of the market for digital publications will consider interests other than their own, and that will include making allowances for fair use. At the same time, our very notion of what kinds of practices constitute fair use will evolve as well, as trusted systems make it

easier for consumers to respect publishers' copyrights than to risk infringement and rely on the fair use defense.

Trusted Systems and Liability for Security Failures

Copyright law and the previous scenarios of using digital works on trusted systems focus mainly on two parties whose interests figure into the balance: rights holders and the public. However, with trusted systems, vendors for trusted systems also have a role in the transactions. This role arises because trusted systems have ongoing responsibilities for handling digital works securely and for enforcing digital contracts.

A fiducial responsibility creates potential liability for platform vendors. Consider the case where the security arrangements for a trusted system fail and, with no user action or intention, a document is released. Is the platform vendor liable? Consider another case where an individual purchases a digital work through a trusted system built by vendor A. That trusted system is later used to transfer the work to a trusted system built by vendor B. Subsequently, the work is transferred to a trusted system built by vendor C, which fails in some way and releases the work onto the Net. Are vendors A and B liable because their trusted systems gave a copy of the work to vendor C's system, which proved not to be trustworthy?

The following scenario shows one way that the risks and liabilities for the failure of trusted systems might be handled. Vendor A builds and sells trusted systems, involving hardware and software. Before bringing the system to market, vendor A takes the system to an independent testing organization, the DPT (or Digital Property Trust). The DPT tests the system, gives it a security rating and issues signed digital certificates used by the trusted system in its authentication protocols. These challenge-response protocols and digital certificates make it possible for other trusted systems to determine the identity and security level of this model of trusted systems made by vendor A. They also make it possible to register all transactions and, in particular, to keep track of which trusted systems have handled which documents.

Continuing the scenario, a publisher assigns rights to a digital work, declares an insurance limit, declares the required security level of trusted systems that can receive the work, and offers the work on the Net. Using a trusted system by vendor A, consumer Joe buys a copy of the digital work. Later, Joe transfers the work to a friend's system built by vendor B. Subsequently, the work is transferred to a system built by vendor C, which fails in some way (or, perhaps, is tampered with by an intruder) so that the work is turned loose on the network. As part of an industry coalition, all the trusted system vendors take measures to isolate the damage. Document insurance evaluates and pays for the publisher's losses.

The example raises several policy issues. Who is liable if a trusted system improperly releases a document because of a design failure? Who is liable if the security of a trusted system has been undermined by tampering or by a computer virus? Who is liable if a computer with outdated security measures participates in a transaction? Is it reasonable for publishers to take the entire risk according to informed consent about

the nature and limitations of the trusted systems? What are the prudent and appropriate actions for vendors when a model of their system is apparently compromised? Should periodic testing and upgrading of security be mandatory for trusted systems? Should the insurance rates be higher for works on trusted systems of low security than for systems of high security? Over time, security requirements are likely to increase and system failures are inevitable. In general, stakeholders in digital publishing will be served best if prompt, concerted, and coordinated action can be taken by responsible parties to contain damage and preserve business as usual. The process of determining a failure mode can be complex. If the process of failure diagnosis and recovery is highly adversarial, vendors of trusted systems may have difficulty cooperating and sharing information in a way that facilitates quickly containing security problems. By creating an industry insurance pool and standards for cooperation, publishers and platform vendors can spread their risks.

In summary, the role that a trusted system plays in enforcing usage makes vendors of trusted systems a party to the system for honoring intellectual property rights. This example illustrates an approach that combines security technology with institutional arrangements. This approach is intended to create a business-as-usual marketplace in digital works in which the risks are amortized by insurance, concerted and coordinated action by vendors is prompt, and the compliance and security of trusted systems is determined by an independent organization.

Trusted Systems and Privacy

Much of the "trust" in trusted systems may seem to be about protecting the economic interests of authors and publishers. However, when automatic systems are used to mediate and bill for what we read and how much time we spend reading it, there is a potential for misuse of the information to violate privacy (Cohen, 1996).

For example, if a person browses some digital magazines on fishing or any other topic, is it appropriate if this interest is noticed by computer systems across the country leading to a deluge of special offers on fishing gear and fishing magazines? If a person with a heart condition develops an interest in open sea fishing, is it appropriate for his insurance rates to go up on the assumption of greater risks due to a heart attack at sea?

In order to do accurate billing, trusted systems must keep billing logs of a consumer's use of various digital works. Privacy issues arise when this billing information is misused by combining it with other data and reporting to third parties.

One way to address the privacy issue with trusted systems is to extend the notion of compliance and trustworthiness to other systems in the billing chain. For example, billing systems could be tested for compliance with policies for accuracy, data reporting, and appropriate data aggregation. These policies would balance the interest of publishers of digital works in knowing their readers with the potential interest

of consumers for privacy. Determining appropriate policies and standards in this area is very much an open issue. One consequence of introducing privacy into the overall picture of protection of intellectual property is that it introduces financial clearinghouses into the set of participants with interests to balance.

An additional set of privacy issues concerns how to shield trusted systems and their users from inappropriate government intrusion or surveillance. For example, pursuant to a criminal or other official investigation, government investigators might want to access the private records contained in trusted systems in order to monitor an individual's reading habits or an organization's publications and subscribers. Depending on the surrounding facts and circumstances, such access could be a legitimate use of authority or a violation of civil liberties. Presumably, the protections of the Fourth Amendment against unreasonable search and seizure would apply to trusted systems. Whether present-day laws, such as statutes prohibiting unauthorized wiretapping, would provide adequate safeguards for users of trusted systems while appropriately taking legitimate government interests into account, or whether additional protective legislation specifically tailored to trusted systems would be needed, is an important subject that we must leave for another day.

Governments as Stakeholders in Digital Publishing

As pundits have observed, international borders are speed bumps on the information superhighway. The rapidity of digital communications and the possibility for information goods to travel across national and other boundaries raises issues of interest to governments that are relevant to the design of trusted systems. Regarding commerce and boundaries, national governments have traditional interests in import and export controls, national security, and taxation.

With regard to taxation, the automatic billing capabilities of trusted systems will almost certainly attract the interest of taxing authorities unless as President Clinton has recently suggested, the Internet is to be treated as a "free trade zone." In principle, billing for taxes can be automatic if trusted systems can be kept abreast of changes in the tax laws.

More problematic is the issue of determining boundaries in cyberspace. To a large extent, one computer looks like another in cyberspace and location is not easy to determine reliably. Using intermediate agents, for example, it is possible to disguise the ultimate destination of a digital work. Furthermore, with portable computers, the actual physical location can change continuously or frequently.

Consider the scenario where a French citizen carries a laptop computer into the United States and downloads software from a U.S.-based software publisher. Has he exported it yet? Does he export it when he subsequently carries the laptop through customs? What if a U.S. citizen is in France with a laptop computer that he intends to return to the United States, logs onto the Net, and downloads some software. Has he exported or imported software? Is he subject to French taxation? When he returns to the U.S. with

the computer, has he exported the software twice or not at all? The most direct answers to these questions based on current law may or may not make much sense in cyberspace.

One approach to reframing import and export issues is to register computers in a way analogous to the way we now register ships or have national embassies in foreign lands (Stefik, 1997b). In this approach, export occurs when someone transfers software from a U.S. registered computer onto a French registered computer, independently of the physical locations of the two computers. Trusted systems could carry digital certificates that would authenticate their nation of registry independent of their physical locations.

To summarize, although copyright law tends to be a national concern (modulated by the Berne convention), the digital medium intimately connects computers that reside in different nations. The possibility of automatic taxation on electronic commerce is likely to become a substantial interest for all nations. This section has sketched an approach to rethinking the mapping of computers to countries through a process of assigning computers to nations of registry.

Conclusion

With trusted systems, copyright is potentially alive and well in the digital era, as is the balance of interests that copyright represents.

Trusted systems can reduce uncertainty about the availability and scope of protection for works falling into the gray zones of copyright. Notably, through trusted systems, fair use can be recast as a market of microtransactions with agreed-upon prices, or as a licensed privilege with fiduciary responsibility backed by insurance. Either model addresses the common good that fair use is meant to protect. Also through trusted systems, it becomes possible to make "digital contracts" to protect those digital works, such as certain kinds of databases, that may fall outside the scope of copyright, again with an appropriate balancing of interests. Publishers can be compensated for the labor and services involved in the preparation and distribution of the digital works, while others who seek to use the uncopyrightable contents may do so.

We have argued that, contrary to what some have suggested, trusted systems do not necessarily swing the pendulum of power too far towards publishers or, for that matter, towards any one interest group. Rather, a well designed trusted systems regime could allow the market for digital publishing to develop and flourish to the benefit of creators, consumers, and the public alike. By the same token, trusted systems are not a panacea. Copyright law will continue to set limits within which the "digital contracts" of trusted systems are made, and will continue to provide a legal basis for thwarting those who attempt to bypass or crack the trusted systems altogether in order to commit large-scale piracy of published works. Contract law will continue to be interpreted by the courts, so that certain provisions of "digital contracts" will be deemed enforceable, others not. Business practices concerning digital publishing will continue to evolve, and consumer advocacy groups will play a role in this regard along with publishers and authors. DPT and other institutions will arise in response to perceived needs, either through government intervention or private initiative.

In the end, trusted systems do not exist in a vacuum, but are complemented by and complementary to the legal, economic, social, and policy frameworks in which they operate. Trusted systems technology, copyright and contract law, market forces, and societal norms (promulgated, in part, through institutions such as DPT), all acting together, can provide the basis for an overall balance. The legitimate and often conflicting interests of all stakeholders in digital publishing—including authors, publishers, consumers, trusted systems vendors, financial clearinghouses, governments, and the public—can be, and should be, fairly represented. The pendulum is still in motion. Without trusted systems, digital publishing is at risk. Through trusted systems and with sound policy choices, we can potentially guide the pendulum to an appropriate point of equilibrium.

References

American Law Institute and National Conference of Commissioners on Uniform State Laws. UNIFORM COMMERCIAL CODE [UCC] ARTICLE 2B—LICENSES. Draft of May 5, 1997, available at <http://www.law.uh.edu/ucc2b/050597/0505_2b.html> (current draft is available from the Uniform Commercial Code Article 2B Revision Home Page, at <http://www.law.uh.edu/ucc2b/>, or from The National Conference of Commissioners on Uniform State Laws, Drafts of Uniform and Model Acts, Official Site, at <http://www.law.upenn.edu/library/ulc/ulc.htm>).

Barlow, John Perry. THE ECONOMY OF IDEAS. *Wired,* March 1994, page 85.

Cohen, Julie E. A RIGHT TO READ ANONYMOUSLY: A CLOSER LOOK AT COPYRIGHT MANAGEMENT IN CYBERSPACE, *Connecticut Law Review*, Volume 28, Number 4, Summer 1996, pp 981-1039.

Elkin-Koren, Niva. COPYRIGHT POLICY AND THE LIMITS OF FREEDOM OF CONTRACT. *Berkeley Technology Law Journal*, 12:1, page 93-113, 1997.

Kahin, Brian and Arms, Kate (Ed), FORUM ON TECHNOLOGY-BASED INTELLECTUAL PROPERTY MANAGEMENT: Electronic Commerce for Content., Volume 2, A Journal of the Interactive Multimedia Association, August 1996.

Lemley, Mark A., INTELLECTUAL PROPERTY AND SHRINKWRAP LICENSES, *Southern California Law Review*, **68**, page 1239, 1995.

Merges, Robert P. THE END OF FRICTION? PROPERTY RIGHTS AND CONTRACT IN THE "NEWTONIAN" WORLD OF ON-LINE COMMERCE. *Berkeley Technology Law Journal* 12:1, pages 115-136, 1997.

O'Rourke, Maureen A. COPYRIGHT PREEMPTION AFTER THE PROCD CASE: A MARKET-BASED APPROACH. *Berkeley Technology Law Journal* 12:1, pages 53-91, 1997.

Samuelson, Pamela. REGULATION OF TECHNOLOGIES TO PROTECT COPYRIGHTED WORKS. *Communications of the ACM*, Vol. 39, No. 7, pages 17-22; July 1996.

Samuelson, Pamela. THE U.S. DIGITAL AGENDA AT WIPO, *Virginia Journal of International Law (forthcoming), 1997.*

Schlachter, Eric. THE INTELLECTUAL PROPERTY RENAISSANCE IN CYBERSPACE: WHY COPYRIGHT LAW COULD BE UNIMPORTANT ON THE INTERNET. *Berkeley Technology Law Journal.* 12:1, pages 15-51, 1997.

Stanford University Libraries. COPYRIGHT AND FAIR USE PAGE. <http://fairuse.stanford.edu/>

Stefik, Mark. LETTING LOOSE THE LIGHT: IGNITING COMMERCE IN ELECTRONIC PUBLICATION. In Stefik, Mark (Ed.), *Internet Dreams: Archetypes, Myths, and Metaphors.* Cambridge, Mass. The MIT Press, November 1996.

Stefik, Mark. TRUSTED SYSTEMS. *Scientific American*. March, pages 78-81, 1997a. (Available on-line as http://www.sciam.com/0397issue/0397stefik.html).

Stefik, Mark. SHIFTING THE POSSIBLE: HOW TRUSTED SYSTEMS AND DIGITAL PROPERTY RIGHTS CHALLENGE US TO RETHINK DIGITAL PUBLISHING. *Berkeley Technology Law Journal.* 12:1, pages 137-159, 1997b.

Summers, Rita C. *Secure Computing: Threats and Safeguards.* New York: McGraw-Hill, 1997.

U.S. DEPARTMENT OF COMMERCE, INFORMATION INFRASTRUCTURE TASK FORCE, INTELLECTUAL PROPERTY AND THE NATIONAL INFORMATION INFRASTRUCTURE: THE REPORT OF THE WORKING GROUP ON INTELLECTUAL PROPERTY (1995), available at <http://www.uspto.gov/web/offices/com/doc/ipnii>.

Biographical Sketches

Mark Stefik is a principal scientist at Xerox PARC. He holds a bachelor's degree in mathematics and a Ph.D. in computer science, both from Stanford University. Current and past research interests include reasoning with constraints, paradigms of programming, as well as applications of AI to problems in molecular genetics, VLSI circuit design, and configuration of computer systems. His current activities include the design of systems for collaborative sensemaking and trusted systems for digital publishing. For several years he taught a course on knowledge systems at Stanford University. His textbook from that course, "Introduction to Knowledge Systems," was published in 1995 by Morgan Kaufmann. His book "Internet Dreams" was published by MIT Press in 1996.

Alexander E. Silverman is Intellectual Property Strategy Liaison at Xerox PARC. He holds a bachelor's degree in physics from Princeton and a law degree from Stanford. Prior to entering the legal profession Alex worked for eight years as a software engineer, specializing in image processing, artificial intelligence, and electronic music. He is a member of the California and Washington State Bar Associations and is registered to practice as a Patent Attorney before the U.S. Patent and Trademark Office.