
Internet Dreams
Archetypes, Myths, and Metaphors

Mark Stefik

The MIT Press
Cambridge, Massachusetts
London, England

Third printing, 2001

© 1996 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was set in Sabon by Asco Trade Typesetting Ltd., Hong Kong and was printed and bound in the United States of America.

Images of the metaphors are by Eric P. Stefik.

Library of Congress Cataloging-in-Publication Data

Internet dreams : archetypes, myths, and metaphors / Mark Stefik.

p. cm.

Includes bibliographical references (p.) and index.

ISBN 0-262-19373-6 (HB), 0-262-69202-3 (PB)

1. Information superhighway—United States. 2. Internet (Computer network)—United States. I. Stefik, Mark.

ZA3250.U6I58 1996

303.48'33—dc20

96-28249

CIP

Letting Loose the Light: Igniting Commerce in Electronic Publication

Mark Stefik

Connections

In "The Digital Library Project: The World of Knowbots" in Part 1, Robert Kahn and Vinton Cerf ask, "If a thousand books are combined on a single CD-ROM and the acquirer of the CD-ROM only intends to read one of them, what sort of royalty arrangement is appropriate to compensate the copyright owners? How would compensation be extended for cases in which electronic copies are provided to users?" Their questions show how, in 1988, issues about copyright protection and payment for using information arose in the context of early CD-ROM distribution.

By 1994 copyright issues not only had not been settled, they were coming to a boil. Laura Fillmore's effort to build a successful publishing business on the Internet reveals the limitations of what was practical in May of 1994. Although digital works were being sold on the Internet, provisions for commerce were primitive. Furthermore, the ease of copying digital works had led many people to believe that digital information should be free. Fast access to the network had made trading programs or other data as easy as mixing songs on audio tape. In short, it had become much simpler for network users to infringe copyright than to uphold it.

This is the context for the oft-quoted statement by John Perry Barlow of the Electronic Freedom Foundation, "Copyright is dead." Advocates of free information argue that because you don't lose the original when you make a copy of a digital work, there should be no charge for copying information. The conventional wisdom among publishers in late-1994, when this article was written, was that digital containers for software

were inherently leaky vessels and that no viable solution would ever be found. The article suggests, however, a way to sustain commerce for those who want to sell information on the network.

Throughout the time I've been groping around cyberspace, an immense, unsolved conundrum has remained at the root of nearly every legal, ethical, governmental, and social vexation to be found in the Virtual World. I refer to the problem of digitized property. The enigma is this: If our property can be infinitely reproduced and instantaneously distributed all over the planet without cost, without our knowledge, without its even leaving our possession, how can we protect it? How are we going to get paid for the work we do with our minds? And, if we can't get paid, what will assure the continued creation and distribution of such work?

John Perry Barlow, "The Economy of Ideas"

No problemo.

T-101 (Arnold Schwarzenegger) in *Terminator 2*

It all depends on whether you really understand the idea of trusted systems. If you don't understand them, then this whole approach to commerce and digital publishing is utterly unthinkable. If you do understand them, then it all follows easily.

Ralph Merkle

Across many cultures, knowledge and inner knowing are described as light. *Letting loose the light* refers to spreading knowledge in the world, typically in written form. Consistent with this metaphor, the period in the eighteenth century characterized by a burst of writings in philosophy and science is called the *Enlightenment*. In the present century the metaphor of knowledge as light is both poetic and physically realized. Books, pictures, movies, musical performances, and other works can be conveniently represented digitally. With fiber optics, digital works are actually transmitted by the shining and pulsing of light.

The digital representation of works and their nearly instantaneous transmission has profound consequences for commercial publishing. Three of the fundamental economic factors affecting the publishing industry—printing costs, inventory costs, and transportation costs—can be drastically reduced. Digital works can be copied at minuscule costs, stored in almost no space, and transported instantly anywhere in the world.

This portability opens up visions of a greater information age. For libraries, universal access to the world's written knowledge is a centuries-old vision. Today many libraries have electronic catalogs accessible to anyone with a computer. Articles can be delivered to anyone with a fax machine. In the technophile's idealized vision, books and magazines need never be printed on paper at all; any digital work could be made available to anyone, anytime, anywhere in the world.

However, the dream of universal digital access to high-quality works dangles just beyond reach. Such works are not usually available, because of publishers' concerns that uncompensated copying will infringe and erode their ability to make a living. History suggests that this problem will not go away. Publishing thrives only when it is profitable, and profitability depends on limiting uncompensated copying.

The conventional wisdom—based on the way computers are used for word processing, electronic mail, and computer networking—is that copying digital works is easy and, therefore, inevitable. There appears to be a clear, inherent conflict between representing works digitally and honoring the commercial and intellectual property interests of creators and publishers. Fortunately, computers need not be blind instruments of copyright infringement. Properly designed digital systems can be more powerful and flexible instruments of trade in publications than any other medium. The seeming conflict between digital publishing and commerce is merely a consequence of the way computer systems have been designed to date.

The technological means for commerce in digital works are now at hand. New and unconventional when compared with today's uses of computers, these means will enable us to buy, sell, and lend digital works much as we now buy, sell, and lend printed books and other publications. They will change the way digital works are purchased and delivered and will give consumers access to all sorts of works at any time of the day—though not necessarily for free. Consumers will be able to sample works, borrow them, rent them for nominal fees, and make copies for friends. Creative people will be able to circulate their works to networks of friends while earning a reliable living from people who make copies of them. This technological system will affect everything from digital books to digital television, from digital music to digital video

games. It will radically change our concepts of digital libraries, digital bookstores, digital music stores, digital newspapers, and digital television stations. Moreover, any competent technological company will be able to implement the required systems.

Here is a road map to this new land. First, we discuss the history of copyright law and the reasons for the widespread, but incorrect, belief that works represented digitally will be copied without permission. We then describe the technological innovations that can enable and support commerce in digital publishing. Finally, we introduce the institutional and business challenges that lie ahead. What we require to overcome them is the wit, will, and means to create institutions that provide the necessary security, convenience, vision, and longevity.

The Origin and Rationale of Copyright

It is harder to be honest than to cheat when copying digital works on general-purpose computers. The license printed on the package of most purchased computer software authorizes a buyer to load the software into one computer and use it there. Getting another legal copy for a friend involves driving to the computer store and buying it. It is much easier, faster, and cheaper to simply load the same software into another computer. Such copying is so private and easy to do that most people do it without thinking, and without guilt.

Unauthorized copying on computers is not, of course, limited to purchased software. With a few keystrokes, it is often possible to copy a paragraph, an article, a book, or a life's work without compensating its creators or publishers. Nor are unauthorized copying and use new phenomena. Anyone who ignores the FBI warning message on video tapes to make copies for friends infringes a copyright, as do people who copy compact discs onto cassettes. As a practical matter, it has not been feasible to enforce the copyright law in these cases. There are simply too many people with recording devices to make rigorous enforcement practical or cost-effective.

It is widely believed that there is no viable technical solution to this problem for digital information. John Perry Barlow, a prominent spokesperson in the computer industry, says that the idea of patents and copy-

rights needs to be rethought in the digital age. Information, he argues, cannot be contained or owned. It wants to be free. Cyberspace is the new frontier, and its leaders and pioneers are today's radical thinkers about freedom of information. Barlow suggests, in fact, that we abandon all notions of intellectual property and market regulation. This solution was tried at least once, and it didn't work. Apparently, for high-quality works to spread in the world people need to be able to make a living from creating and distributing them.

Barlow's arguments are reminiscent of the intense debates about intellectual property that took place in France during the French Revolution. Like Barlow, revolutionaries argued that ideas cannot be owned and should not be regulated. During the revolution, many writers and underground publishers emerged as civic heroes of public enlightenment by arguing against tyranny and for freedom of the press. The revolution of the mind, they said, required the dismantling of the laws and institutions governing authorship, printing, publishing, and bookselling. Absolutely free communication was one of the most precious rights of man. All citizens should be able to speak, to write, and crucially, to print freely. According to this philosophical ideal, people had a will to know and should be allowed to read and learn from anything they liked. The wide availability of books and the right to publish were seen as keys to this spread of knowledge.

In 1789, the revolutionary government wholly deregulated the press, believing that the works of the great writers of the Enlightenment would thus be made universally and cheaply available. The writers and publishers certainly never expected what actually happened. Instead of works of enlightenment, the presses turned out mostly seditious pamphlets and pornography. Printers also competed with each other to bring out cheap editions of books others had spent money developing. So little money could be made producing the good books that quality declined; most editions were abridged and contained many errors. Publisher after publisher went into bankruptcy and then out of business. The disastrous nature of an unregulated press, largely unanticipated in the heat of the revolution, became blatantly obvious as the publishing industry fell into shambles. The same leaders who had clamored for the freeing of the presses came belatedly to understand the folly of their action. In the

chaos of the unregulated press; some prominent and popular writers even stopped publishing; because they could not control the printing of their works, they could not make a living by writing.

In 1793 legislation to restore order to publishing was passed. It recognized the rights of authors and grounded the publishing industry in the principles of the marketplace, establishing the author as creator, the book as property, and the reader as an elective consumer. This law reflected a fundamental shift in the Enlightenment perspective, which now saw that the widespread creation and publication of creative works was better served when the authors could own the products of their minds. This history of the treatment of intellectual property in France is discussed by Carla Hesse in her book *Publishing and Cultural Politics in Revolutionary Paris, 1789–1810*.

Today, most people see the infringement of copyright on digital systems as unavoidable. In the remainder of this section, we describe the assumptions about computer design behind this belief and argue that we need to go beyond conventional ways of thinking to solve the problem.

Three main factors currently inhibit the development of digital publishing: (1) the absence of high-contrast, low-power, cheap flat-panel displays; (2) lack of an inexpensive and reliable way of handling money digitally; and (3) the need for a widely accepted means of accounting for the use and copying of digital works. Improvements in technology will almost certainly solve the display problem in the next five to ten years. Most people see such displays as crucial to making electronic books and newspapers portable. They matter less, however, in applications for which desktop displays are satisfactory or where displays are not necessary—such as in transmitting musical works. The second factor—methods of handling money digitally, in the form of checks, credit cards, or anonymous cash—has recently become the subject of much field experimentation. Our focus is on the third problem, techniques for commerce in what we call *digital property* rights or *usage rights*, a generalization of the idea of copyright that delineates several kinds of rights besides copying.

Some publishers see illicit copying as too big a business risk and do not publish in digital form at all. Digital newspapers often leave out impor-

tant and high-value content such as the pictures or graphics, and consumers of these lower-quality papers are unwilling to pay much for them. The perception of low quality leads to a chicken-and-egg problem in which the publishers make little money and consumers have few choices. Ironically, publishers of works that need periodic upgrading, such as computer software, have found that some leakage increases their customer base, even though it is often reported that there are more unauthorized copies of a program in use than authorized ones. Software publishers have decided that the revenue losses of illegal copying are affordable, although they lead to unfair billing. Software publishers charge all users the same price, regardless of the use to which they put the program, arguably overbilling people who use the work infrequently.

As computers and computer networks have proliferated, the need for a better approach to protecting digital works has become more widely appreciated. Moreover, as new kinds of works—such as music, video, and multimedia works that mix these forms—are now available digitally, people from different industries are searching for solutions. Given this wide acknowledgment of the need, why have solutions seemed so elusive? Apparently, we are stuck in a rut, assuming that things must be done the way they have always been done with electronic mail, word processing, and other current applications.

Conventionally, we use general-purpose computers with general-purpose operating systems and general-purpose programs. The computer industry, grounded on the premise that computers can do anything that can be programmed in software, produces a wide range of programs—word processors, spreadsheets, databases, calendars, graphics programs, and computer games. Manufacturers accept no liability when someone uses a computer to copy a copyrighted file. After all, one company builds the computer, another writes the software that does the copying, and both hardware and software are intended for general purposes—that is, any purpose the user wants to put them to. The manufacturer wants no responsibility for someone who uses the computer in a way that just happens to infringe a copyright, nor does the software publisher. The perpetrator is the consumer, who finds it easier to make an unauthorized copy than to be strictly honest.

Stuck within this framework the community of computer users protests against any attempt to regulate the copying of digital property. If we continue to accept this framework, with all of its assumptions, no party will be motivated or empowered to break the cycle and no effective way to protect digital property will be developed. At present, without enforceable property rights, the writers of words, interactive games, and songs often are not compensated for their work. And without their works the world is a darker, poorer place. Honoring their creative work in the digital systems of tomorrow requires us to challenge the design assumptions of the systems we use today.

A New Design for Digital Publishing

The technical core of the approach we propose is based on two ideas: (1) that digital works can be bought and sold among trusted systems, and (2) that works have attached usage rights that specify what can be done with them and what it costs to exercise those rights.

Trusted Systems

The term *trusted system* refers to computers that can be relied on to do certain things. For example, suppose that a creator or publisher forbids all copying of a particular digital work. A trusted system in this context would reliably and infallibly carry out that stipulation; no amount of shouting or coaxing would coerce it to copy the work. The trusted system might be very polite, but ultimately it would always refuse to make an unauthorized copy. Similarly, suppose that a trusted system *could* copy a work but only if it reliably records a set fee to be paid when it has done so. A trusted system would always record the fee whenever the work was copied. If the copying process is interrupted part way through, the trusted system would follow a standard policy; for example, it might delete the partial copy, record no fee, and note that a copying attempt was begun but not completed. Again, no amount of coaxing would change its behavior. It could always be counted on to follow the rules of the trust.

A common but false analogy claiming to show why digital works cannot be protected in computers is that of genies and bottles. In this

analogy, a valued digital work corresponds to a genie and the bottle is a place to store it. When a digital work is sent to a computer, for example, it may be sent in coded form, so that even if the transmission is intercepted it is useless to a wiretapper. Once people have a legal copy of a digital work, however, they can make more copies of it. Since they have a key, they can just decode the work and make copies of it. Alternatively, they can copy the coded version and give away copies of the key. Once the content genie is out of the bottle, according to this scenario, you can't put it back in and unauthorized digital copies are sure to circulate. This is the problem trusted systems can fix.

Trusted systems speak a communications protocol with other trusted systems and will not transmit information to any system not recognized as another trusted system. This strategy ensures that copies of digital works are either inside trusted systems or they are encrypted. When they are inside trusted systems, usage is controlled. When they are outside trusted systems, usage is practically impossible without breaking the code. The important issue, however, is not just protection and containment. The greater good is not served by simply limiting the flow of information. It is served by supporting and encouraging a lively trade in information. *Rather than just confining genies to specific bottles, we want to encourage them to travel between bottles under rules of commerce.*

A very concrete question about such a system is "Why couldn't I just copy a file onto a diskette and give that away?" Unless there is permission to do so, a trusted system would never copy a work to a diskette or anywhere else. Even if permission to copy a work is given, a trusted system would not make a copy on a diskette, because a diskette is not a trusted system. Nor are magnetic tapes, compact discs, or, even, the disk drives of trusted systems. Trusted systems contain computers, have internal protected storage, and communicate by protocol. From a user's point of view, the trusted system is the storage device. Trusted systems only make copies of a digital work on themselves or on other trusted systems. Putting an unencrypted copy on a diskette is letting the genie out of the bottle onto an unprotected medium that can be accessed by a general-purpose computer that does not honor usage rights.

There is an important issue about the perception of trusted systems. One way of looking at them is to say that trusted systems presume that the consumer is dishonest. This perception is unfortunate, and perhaps incorrect, but nonetheless real. Unless trusted systems offer consumers real advantages they will probably view them as nuisances that complicate our lives. A more favorable way to look at trusted systems is to compare them to vending machines. They make it possible to order digital works any time of the day and get immediate delivery. Faster than a telephone-order pizza, a digital work can be delivered immediately over the same telephone line it was ordered over.

In summary, the first key to commerce in digital works is to use trusted systems. We have spoken of these systems as computers, but they are not limited to devices like personal computers and need not seem like computers at all. They could be personal entertainment devices for playing music, video game devices, laptop reading devices, personal computers, devices for playing digital movies at home, credit-card-sized devices that fit in your pocket, or whatever. In the following discussion we refer to these trusted systems as *repositories*, an architectural plan that can have different embodiments. Repositories communicate digitally with other repositories and not with anything else. In contrast to such current passive media as compact discs, repositories have no externally defined limits on storage capacity; so successive generations of repositories could increase in capacity while remaining completely compatible with earlier systems. Digital works would be communicated between repositories using secure coded protocols. Repositories would read the rules that apply to a given digital work and follow them. This brings us to the next issue: How do repositories know what the rules are?

Attached Usage Rights

We start with an analogy. When we go to a store to buy a shirt, there are various tags attached to it. One kind of tag is a price tag. If we want to buy the shirt, we must pay the amount on the tag. Another tag gives cleaning instructions: for example, wash by hand in cold water or dry clean only. Still another tag might say something about the style of the shirt or the history of the shirt company.

This is roughly the idea of *usage rights* on digital works. Digital works would come with tags on them. The tags—put there by the creators, publishers, and distributors—would describe the usage rights for the digital work: what can be done with it and what it costs.

There are some important differences from the shirt's tags. The first is that the tags are digital and intended to be read and used by the repository itself, although consumers can also read the tags through the repository's user interface. They are written in a machine-readable language and give the repository the rules for using the work; they are an electronic contract enforced by the repository. Another difference is that the tags are not removable. Finally, there can be tags attached to different parts of a work. For a shirt, it is as if there were tags on the pockets, tags on the buttons, tags on the collar, tags on the sleeves, and so on. Each tag would grant rights to that part, and different rights could pertain to different parts of a work. For example, a digital newspaper might have certain rights on local stories, others on photographs or wireline stories or advertisements, and so on.

Suppose that the digital work is a piece of music. A statement describing a right might say the following:

This digital work can be played on a player of type Musica-13B. This right is valid from February 14, 1995 to February 14, 1996. The repository must have a security level of three. No other authorizations are needed. The fee for exercising this right is one cent per minute with a minimum of five cents in the first hour. Usage fees are paid to account 1997-200-567131.

Of course, such an internal statement would not be in English, although it should be in a well-defined computer language. Here is an example of a machine-readable statement in a usage rights language:

Right Code: Play Player: Musica-13B
 Copy Count: 1
 Time-Spec: From 95/02/14 Until: 96/02/14
 Access-Spec: Security-Level: 3
 Fee-Spec: Fee: Metered \$0.01 per 0:1:0
 Min: \$0.05 per 0/1/0
 Account: 1997-200-567131

Computer languages are more precise than natural languages and have formal grammars and semantics that define how to interpret each phrase in the language. Computer languages are not at all poetic, but they are much less ambiguous, if less expressive than natural languages. Because the sentences of a digital property language are parts of potential contracts between the creators of digital works and consumers, clarity and simplicity are exactly what we want. Interpreting a usage rights language is quite simple. In level of difficulty, it is more like reading bar codes from packages at the supermarket checkout than it is like reading and understanding an English sentence in a story.

A digital property language needs to define several different kinds of rights, mainly those concerned with how the work can be transported, how it can be rendered, and whether it can be used in derivative works. Other, special rights relate to making and restoring backup copies to protect against hardware failure. The easiest way to understand usage rights is to consider some examples.

Transferring Digital Works When we copy files for friends on a general-purpose computer, we increase the number of copies of a digital work, fail to compensate the work's creator, and infringe the copyright. A repository, in contrast, never infringes copyright.

Our first scenario illustrates how copy and transfer rights would work in a repository system. Suppose that Morgan buys a copy of a digital book, perhaps at the book kiosk at the supermarket. To do so, he exercises a right to copy the book and pays a fee; copying the book records a transaction between the seller's repository and (say) a card-sized repository that Morgan carries with him. Alternatively, he could buy a copy of the digital book from home by telephone. In either case, the digital book is delivered electronically by a communications protocol between the vendor's repository and Morgan's repository. At the end of the transaction, Morgan has spent some money, has a copy of the digital book in his repository, and can now read it on a reader. The book arrives with all its usage rights intact.

Now suppose that, when Morgan finishes reading the book, his friend Andy asks to borrow it. They plug their repositories together, and Morgan exercises a *transfer right* to move the digital book to Andy's reposi-

tory. With paper books, once we have bought a book we can give it away or dispose of it in any way we please, and the same right could apply to Morgan's digital book. At the end of the transfer transaction, the digital book resides on Andys repository and not on Morgan's, and no money has been exchanged. Andy can now read the book, but Morgan cannot. The crucial point is that the transfer transaction preserves the number of copies of the digital book.

We now consider a scenario involving a loan right. Again Morgan has a digital book that his friend Ryan wants to borrow for a week. They plug their repositories together and Morgan exercises a *loan right*. Again, while the digital book is loaned out, Morgan cannot use it. Suppose, however, that Ryan goes off on vacation and, while he is playing volleyball on a beach thousands of miles away, the week's loan period runs out. He has completely forgotten the book. Because both repositories have clocks in them, Ryan's repository deactivates its copy when the week is up. Meanwhile, Morgan's repository also notices that the loan time is up and marks its temporarily deactivated copy as usable again. Without any action by either person, or even any communication between their repositories, the digital book has been returned automatically. If Ryan still wants to access it later, he could pay a nominal fee to rent the work or to make his own copy. The point of both scenarios is that the repositories follow rules, which in this case mimic and improve on the rules of loaning for paper books. The ability to return loaned materials automatically would probably be widely used in digital libraries.

Rendering Digital Works To read a digital book you have to be able to see it; to listen to digital music you have to be able to hear it; to enjoy a digital video game, you have to be able to see and hear it. We use the term *render* to mean the processing of a digital work so that it can be experienced. Like copying, transferring, and loaning, rendering is controlled by usage rights.

We distinguish two forms of rendering: playing and printing. When we play a digital work we send it to another person through some kind of transducer so that he or she can experience it. The term *play*, usually employed in phrases like *playing music* or *playing a movie*, is also used to

denote displaying part of a book, running a computer program, or running an interactive video game. The term *print* in the digital context means to make a copy of the work on media outside usage rights control, either on paper or by writing a file to an external storage device.

The concept of usage rights allows great flexibility in marketing digital works. Today, when you buy a compact disc at the music store, you pay for the copy and play it for free. The same is true for a book. You buy the book and read it as often as you want; generally, you aren't supposed to make copies of it, but you can give it away. By contrast, keeping digital works in repositories would provide more flexibility.

Suppose for example that Andrea's mother is at the music store but does not know exactly what music her teenage daughter wants to buy. She transfers a selection of music to her own repository, choosing collections by half a dozen bands that she knows Andrea likes. At this point she does not need to pay anything for the right to make the copies. When she gets home, she transfers the music to Andrea's repository for her to check out at her leisure. Like all repositories, Andrea's home repository has a built-in credit server that transfers funds electronically. She can listen to short demonstration samples for free or listen to pieces she selects for twenty-five cents an hour; or she can pay for five years of unlimited playing for \$10. Thus, Andrea exercises a pay-for-play right rather than a pay-for-copying right. If pay for play has an infinite term, there are no fees for playing the music as often as she likes. The terms and alternatives for usage rights and usage fees would be set by the music's creators and distributors. What Andrea gains by this arrangement is flexibility and the convenience of trying out different digital works. She might even be able to use technology similar to a cellular telephone to order music and download it from the music store to her car repository.

For books, however, the idea of pay for play may not seem very useful if the typical book we have in mind is the paperback novel. As it is already cheap, it does not seem worthwhile to charge for the time needed to read it. Besides, why should slow readers pay more than fast readers? On the other hand, consider large, expensive reference works like encyclopedias. People do not casually pay out hundreds of dollars for these works; nor do they usually read them from beginning to end. Paying a

small fee for each hour of actual use may make it feasible to bring high quality digital encyclopedias into many households that could not otherwise afford them. As in the music example, purchasers could decide whether to pay by the hour or for large or infinite blocks of usage time.

As yet another example, consider the digital newspaper, which could, in principle, be delivered in several ways: bought at the corner newsstand, downloaded over the telephone, or broadcast by a digital radio station. It could be available though pay-for-play usage rights or by monthly subscription. Suppose, however, that a particular newspaper is reluctant to allow its customers to make paper copies of the newspaper, even for a fee, for fear that some enterprising person will print up enough copies for the entire neighborhood and cut into circulation. To prevent this occurrence, usage rights could be designated so that people trying to print the newspaper would discover they have no printing rights.

Suppose, again, that a month later customers find that they can print the newspaper's old news without hindrance but not its new stories. In that case, the usage right for printing would be dated to prohibit printing until a month had passed since publication, perhaps because the publisher figures that printouts of old news are more like advertisements for the newspaper than threats to circulation. Variations in what the publishers might allow or encourage are virtually endless. The rights granted on stories or photos or whatever could even become a basis for competition among digital newspaper publishers.

Making Derivative Works Distributors add value to products by advertising and selecting works and presenting them to consumers, making a living by performing these functions and requiring compensation for them. Today bookstores and music stores operate on a per-copy basis, charging for "hard copies" of books or compact discs. If pay for play became popular for digital works, how would distributors make money? What is needed, of course, is a mechanism for paying distributors when a consumer chooses to pay for play. This mechanism, called a *shell*, would enable distributors to modify rights and add new usage fees.

We can understand shells as analogous to gift boxes of different sizes. A common and amusing trick is to put a present in a small box, wrap it, and put that box in a bigger box, wrapping that box and putting it in a

bigger box, and so on, perhaps attaching a gift card to each package. In digital work, the boxes would correspond to digital containers (shells) and the gift tags would correspond to attached usage rights. Putting a digital work into an empty shell, or into a shell containing another digital work, is called *embedding*, and it is controlled by an *embed usage right*.

Consider the following case. Nick has written a novel and offers it in digital form. He determines what he wants to charge per copy and attaches usage rights to it, specifying the fee to be paid into his account every time a copy is made. His publisher agrees to publish the book and puts another shell around it directing that when the work is copied, an additional fee is to be paid to the publisher. The publisher may help the author improve the book in various ways and spend money advertising it. Finally, a bookstore puts an additional shell around the publisher's shell, directing that when the kiosk makes a copy of the book, a fee should be paid to the bookstore. Thus when a customer buys a copy of the work, he or she pays, automatically, a fee to the bookstore, to the publisher, and to Nick. This works because the accounting system follows the instructions embedded in every shell of the copied work.

Or consider Paige, a college professor at a business school that bases courses on case studies collected into a reader. Paige chooses the cases of interest for her course as digital works with attached rights. If an interesting case has an extraction right, she can remove a copy of it from its digital source; if it has qualified editing rights, she can make certain kinds of changes to it; if it has embed rights, Paige can add it to her own collection. At each stage, the continuing rights to the work are controlled by its creator's specifications. Paige can put all the works she has collected into a shell and add her own usage right specifications to the shell. When a student buys a copy of her course reader, fees are paid to the creators of each case study and to Paige herself.

It is interesting to compare this process of controlled reuse to the existing practice in which one author requests reprint permission for an article by another author. The process of granting and obtaining permissions is tedious and time consuming and is often assigned to editorial assistants by publishers who do not expect to earn much from reprint requests. This approach assumes that they are willing to agree to the most usual rights, fees, and conditions for reprinting a work. By lowering

the perceived hassle and cost of reuse, usage rights may trigger a substantial increase in the commercial reuse of works.

Consumer-based Distribution One of the most radical possibilities for distributing digital works is *consumer-based distribution*, sometimes called super distribution. With ordinary media, consumer copying and sale of works is considered a problem, because creators and publishers receive no compensation for such copying. In contrast, consider what would happen in our earlier scenario if Morgan, instead of giving or lending his digital book, makes a copy of it for his friend Andy. The repositories would record the transaction and bill Morgan or Andy for the new copy. Depending on how the shells have been set up, fees would be collected for the store where Morgan originally bought the work, for the distributor, for the publisher, and for the creator, even though none of them are present at the transaction. Every consumer would become a potential salesperson, a word-of-mouth sales channel.

This possible future is in radical contrast to the problem foreseen by Barlow in the quotation at the beginning of this piece. Digital property can be anywhere on the planet without the knowledge of its creators and still make money for them whenever it is used or copied by a repository.

Licenses and Tickets In a trusted system, licenses and tickets would be special kinds of digital works that play direct roles in commerce. *Licenses* would be digital certificates that enable someone to exercise certain usage rights. Think of them as similar to driver's licenses or identification cards that authorize someone to drive a car or enter a restricted area. A digital license would let someone exercise certain rights, such as copying or printing a particular work. When a consumer asks to use a licensed work, an authorization server or "digital authority"—a program on a repository—would check his or her digital license. *Digital tickets*, a kind of coupon offered by publishers for prepaid uses or discounts, would enable a possessor to exercise a right exactly once. Think of digital tickets as comparable to movie or train tickets; once you have entered the theater or boarded the train, the ticket is punched and cannot be used again. A digital ticket is punched by a digital ticket agent that is a program on a repository.

When an author creates a work and specifies its usage rights, he or she can require buyers to have particular licenses or tickets to exercise certain rights. These digital licenses and tickets would be essentially impossible to forge; would let consumers exercise usage rights in living rooms, school dormitories, or anywhere else; and would assure authors that fees will be collected and that the tickets and licenses specified will be required.

Different areas of publishing have different crucial problems for which digital licenses and digital tickets could offer solutions. In some industries, there is an advantage in ensuring that only authorized distributors can sell digital works. The author of a computer game, for example, may want his or her game sold only through distributors licensed to advertise, promote, and demonstrate it. If a dealer without that license tries to copy and sell the game, the repository would refuse the transaction. By determining who gets distributor's licenses and arranging that licenses cannot be transferred between repositories without authorization, the author could maintain control over distribution. Such digital licenses could also restrict distribution rights to a certain time period.

In the music and video industries, companies try to limit the playing of their recordings to home use. While current technology provides no effective way to enforce this provision, trusted entertainment systems using licenses could distinguish between equipment for home use and equipment for theater or broadcast use. Equipment for playing digital recordings would come with different licenses and fee schedules for home and public use. Alternatively, radio transmitters and receivers might all be linked to repositories; a station might broadcast a work but require each listener to pay a nominal fee for receiving it. Trusted systems could provide a basis for many different kinds of relationships between creators, broadcasters, and the public than are feasible today.

In the book publishing industry, it is common to offer big discounts on books that remain unsold after a certain period of time. Imagine a usage right that allows a copy of a digital book to be made for nothing in exchange for a certain ticket. Later, a book club might offer these tickets for sale and let a holder get any three digital books in exchange for three tickets. The price of a digital book could thus be determined later by the price of the ticket, which could vary according to demand.

In the computer software industry, it is common to release new versions of software to fix bugs in earlier versions and to offer the upgraded software for free or for low prices to purchasers of the original version. The problem lies in making sure that the upgraded versions are not given to people who did not purchase the original software. Digital tickets would offer a solution by allowing vendors to bundle upgrade tickets with the original works. Consumers could use any copy of the upgrade versions of the software to make their own new copy by simply using their tickets; this approach lets consumers upgrade their software without going back to the dealer. Because digital tickets get punched when they are used, the approach would allow exactly one upgrade per original version.

Because metered and per-copy fee arrangements can make usage budgets unpredictable for organizations, they may prefer *site licenses*. Such licenses would grant members of an organization the right to use a digital work subject to restrictions. For example, a site license would generally preclude making copies for use by people not in the organization; or it might distinguish between different kinds of uses in different departments. It might also limit the number of people who can use a work at the same time, perhaps leaving its administration to someone in the organization. The supplier of the digital work would specify in the usage rights that the organization's site license be recorded on a repository before the work could be used. If different departments require different regulated uses, then each department would have its own specialized license. To monitor such global constraints as the number of copies of a work simultaneously in use, a digital license could instruct the repository's authorization server to communicate periodically with a site authorization server that registers and counts users.

Licenses and tickets could be established for diverse categories and purposes, including social purposes. For example, a charity or governmental organization could issue certificates to low-income people or inner-city youth. Socially conscious publishers could then offer discounts or limited free use of certain digital works to people holding such certificates. From a profit perspective, they may assume that offering the certificates will increase the potential base of customers by contributing to public literacy and that, in any case, these consumers would not buy the

work. The same digital-license mechanism could provide special rights to certified librarians, researchers, and teachers; and certificates could be dispensed at libraries to allow readers to browse works for limited periods.

One area that has been much discussed recently is control of access to works on computer networks. In the present free-wheeling environment of computer networks, interested people can set up discussion groups on explicit and adults-only topics. The same computer network that offers on-line museums and information sources for kids may also offer on-line pornography. Controlling access to such materials requires an approach that balances social interests in free speech and commerce with community interests and responsibilities regarding adult materials.

In public settings off the network, the issue of adult material usually arises in reference to magazines and videos. It is common practice for dealers to display adult magazines so that passersby need not see suggestive covers. Video distributors generally follow the movie rating system: G for general audience, PG for parental guidance advised, R for restricted, and X for adult movies. Both these approaches have analogs to digital licenses. Thus, works can come with ratings established by appropriate institutions or community organizations. For example, a G-rated movie would require no license, whereas viewing a digital movie rated PG-13 would need either an identity certificate specifying that the consumer is over 13 or a permission ticket issued by his or her parent.

The foregoing examples show how central digital tickets and licenses are to the usage rights approach and how one overall infrastructure could serve a wide range of social and commercial purposes.

Foundations of Trust in Repositories

Next we explore the question, what is it that we want to trust about repositories and what is the basis for such trust? In general, a trusted system is one that can be relied upon to take responsibility for a given operation or set of operations. In the case of digital works on repositories, the requirement for trust is that the repositories follow—at all times and in every instance—the rules about how digital works are used. They must be accountable for all uses of the works and for the fees charged for those uses. Responsibility is, fundamentally, an issue of *integrity*. For

repositories that integrity has three parts: *physical integrity* refers to the soundness of the physical device itself; *communications integrity* means roughly that repositories cannot be easily fooled by telling them lies; and *behavioral integrity* means that repositories will exercise their functions exactly as they are supposed to, 100 percent of the time.

Physical integrity applies both to the repositories and to the digital works they protect. One threat to a repository is that someone will pry open the case and gain access to what is stored inside. In a trusted system different repositories could have different levels of security. A repository that can be compromised with a power drill and a screwdriver would have a low security level. A somewhat higher level of physical security would be a system with sensors that enable it to detect a threat and erase certain key data. A still higher level of security, suitable perhaps for a real life James Bond, might be a system that self-destructs when it detects that it is under threat, perhaps setting off alarms and telephoning for help.

Even at its lowest level, security for a repository would be much higher than the security for such passive media as videotapes, compact discs, or computer diskettes. These media record their information out in the open where it can be accessed by any general-purpose reading device; they cannot detect intrusion or take any kind of protective or evasive action. Repositories, on the other hand, would never present data to any device that fails to establish itself as a bona fide, qualified repository.

The second kind of system integrity, communications integrity, would ensure that repositories could not be easily fooled by being connected to illegitimate computer systems masquerading as legitimate repositories. When repositories connect with each other, they would go through a registration process identifying themselves to each other and establishing their bona fides. Imagine, for example, two secret agents unknown to each other who meet by arrangement. What's the secret word? How do I know you are who you say you are? Are we sure that nobody else is listening? These are the kinds of concerns two repositories would have when they are connected. They would put each other through a series of tests intended to weed out impostors and protect the works with which they are trusted. Only when registration succeeds would they establish a trusted session.

A few words about how this works are in order. At the heart of the registration process is a security concept called *public key encryption*, a

well-known and much-studied system for secure and secret communication. In this approach, each repository is given a private key or code, which it keeps secret, and a public key. In a trusted system, these keys would be given to a repository when it is manufactured and would be certified by a master repository known to be highly secure. (One of the principal requirements for the system architecture—discussed in the last section of this paper—is an institution that can control and safeguard the master repositories.) Communications integrity is ensured because all communications among repositories are in codes that are extremely difficult to break. In addition, provisions to detect attempts to tamper with communication and to isolate repositories identified as compromised can be built into the protocols.

Finally, we come to behavioral integrity. Even if the repositories have not been physically compromised and can prove their identities to each other, how do we know they will work properly? In the secret agent analogy, how do we know that the other agent hasn't been compromised or turned? Because the behavior of computers is determined by their programming, the programs used in repositories must be thoroughly tested and certified, which is a lot of work. What makes the certification task easier for repositories than for computers in general is that repositories would have limited functions. They would need to carry out a limited number of very specific operations relating to usage rights, protocol handling, and accounting. Furthermore, there would be procedures that guarantee that all installed software is inviolable to tampering or modification. Finally, even if a repository were compromised, it would need to identify itself for any transaction with a certificate from a master repository; other repositories, given the identity of the compromised repository, would refuse to carry out further transactions.

In summary, the physical integrity, communications integrity, and behavioral integrity of a repository are the foundations of a trusted system. A repository is designed with the ability to detect tampering and communications errors and to ensure certified behavior. These characteristics could be achieved by computers. Until now, however, the value of these attributes has not been appreciated, so that elements of the technology are not widely available and computers are designed without these goals in mind.

The Accountant Inside When someone makes a copy of a digital book and the repositories record the charges, what's to say that the author or publisher will ever get paid? Suppose, for example, that the repository is a credit-card-sized device. Why not just throw it away when a large enough bill accumulates on it? Also, if repositories contain the equivalent of money, what will prevent theft of that money?

There are many ways to approach these questions. During 1994 several groups conducted field trials of the national computer networks that operate digital checking accounts, cash, and network credit cards. These trials provide a base of experience relevant to usage rights.

Repositories would have a substantial advantage over such systems because they would record a transaction without requiring an immediate telephone or computer-network connection to a financial clearinghouse. This is especially relevant for billing tiny amounts, microtransactions, for which the cost of a telephone call would dwarf the amount of the payment. Monthly connections to a clearinghouse would be more convenient. A repository could be connected to a clearinghouse in any of several ways: by plugging it into a bank teller machine or a special telephone or connecting it to a computer network through a personal computer. In such an operation, the repository would open a channel to a clearinghouse with a registration protocol similar to the ones used when two repositories connect to each other. During the monthly session, all transactions could be reported in a single communication and a fresh credit limit for the repository could be set.

Some security measures would be necessary. First, a personal identification code similar to those used at automatic tellers would be required before charges could be accrued. Second, all transactions would be reported by both parties; this would not prevent cheating but would require two people to collude in losing their repositories. Third, the repository itself, which is likely to be more expensive than a credit card, would require insurance to replace it. Finally, as happens with credit cards today, anyone who regularly reports the loss or theft of a repository would eventually have difficulty getting a replacement or would be forced to use a repository with more stringent security arrangements.

Implications and Institutional Challenges

Stakeholders in digital publication will ultimately come from many industries, walks of life, and parts of the world. For a variety of reasons, they are now accustomed to different versions of copyright law and have different conventions of what constitutes fair use of copyrighted material. As more and more kinds of publication go digital, people's different expectations are likely to create a tug of war. Digital property rights may provide a way to manage this evolving situation. Because copyrights traditionally last a long time, we need to create and shape institutions governing usage rights that will serve us well for a long time.

Usage Rights and Copyright Law

Copyright law is not static. Over time, there have been various changes. Notable copyright reforms were made as recently as 1976, and several more are currently under active discussion. Defining a few terms will bring us to the nub of many copyright issues and show how digital usage rights are relevant to them.

Copyright law has provisions for what is called *fair use*, the amount and kind of quotations from a copyrighted work that can be made without permission. In general, creators and publishers have an interest in limiting free use of their own materials under fair use so that they can require extra fees for particular uses. Consumers, librarians, and scholars, on the other hand, have an interest in ensuring that certain uses are unencumbered.

Here are some examples. A person who buys a copy of a work is permitted to use it in a variety of ways. Although making copies of the work to sell or give away is generally not legal, some copying is allowed—at least for certain kinds of works and in some contexts. For example, it is usually all right to make a copy for personal use and to quote passages from a book in a book review or scholarly work, as long as the original source is cited. As new media have become important, similar issues have arisen about fair use in quoting from recorded music and movies.

Sometimes more than one kind of use is distinguished for a work. An easily understood example is that of a play script. The usual copyright provisions govern the copying of a script. The right to create a stage production from the script and to perform it publicly for profit is not,

however, covered by fair use. A public performance requires a different kind of right, a *performance right*. Similar issues arise for musical scores and recorded music. When we buy a compact disc or cassette tape of music, fair use includes playing it at home, whereas a radio transmission constitutes a public performance. Because performance rights are not included in the purchase, radio stations are supposed to pay for putting such works on the air. The small print on a compact disc generally includes the phrase "all rights reserved," the publisher's way of claiming all rights not explicitly granted.

The fair use doctrine arises when possessing a copy of a work gives a person the potential to use it in ways its creator thinks are unfair. In such cases, usage rights could provide for different kinds of uses and fees, distinguishing between copying rights, loan rights, transfer rights, play rights, broadcast rights, print rights, extract rights, embed rights, editing rights, and several others. On a trusted system, these specific rights could be granted by requiring particular digital licenses or tickets and different fees. Some kinds of digital players, for example, might have built-in licenses and codes that determine whether they can be connected to public broadcast systems or only to home systems.

In many fair use scenarios, the gap between what is fair use and what is infringement is exacerbated by the fact that there is nothing in between. Fair use costs nothing. Other uses cost fees; and even when the fees are small, the cost and bother of obtaining and accounting for them is high. The digital property rights approach, however, could permit transactions for even nominal amounts of money, changing the confrontational issue of fee versus free to a practical issue of "how much?"

Usage rights could also make it possible to grant rights to designated categories of users for social reasons. Digital certificates could be made available to librarians, library users, teachers, students, impoverished people, and so on. Some interest groups are already drafting position papers about fair use in the electronic age, spelling out certain rights that they believe should incur no fees, such as the right to print temporary paper copies for personal use. Curiously, the concept of usage rights reverses our conventional assumptions about making copies; printing a paper copy, we can now recognize, is moving digital content out of repository control. In a usage rights context, paper copies would have greater potential for unauthorized copying than digital copies would.

In summary, the usage-rights approach provides repository-mediated contracts to govern the various uses of digital works. It creates specific language for common kinds of uses and their fees. It is a tool never imagined by the creators of copyright law, or by those who believe laws governing intellectual property cannot be enforced. What copyright law protects is the expression of ideas. As John Perry Barlow put it, "The point at which this franchise was imposed was that moment when the 'word became flesh' by departing the mind of its originator and entering some physical object. . . . Protecting physical expression had the force of convenience on its side. Copyright worked well because, Gutenberg notwithstanding, it was hard to make a book. . . . Unlike unbounded words or images, books had material surfaces to which one could attach copyright notices, publishers marques, and price tags (*Wired*, March 1994).

The ability to attach appropriate tags proclaiming the rights and fees for different uses is exactly what repositories do for digital works. Such tags, permanently attached and honored by the trusted systems, would enable us to experience the works. With these tags, the basic concepts of copyright law seem to work just as they are. In this way, an unconventional redesign of computer operation can preserve and even improve the now conventional social contract between those who create and those who consume works of the mind.

What Repositories Can and Can't Do

Repositories cannot, of course, prevent all unauthorized copying of digital works. No technology can keep someone from reading a digital book, then laboriously typing the words verbatim at a keyboard. A plagiarizer could even use technological aids, for example, a television camera aimed at a display and feeding its output into a computer equipped with an optical text reader. Digitally published music could be played through a repository's speaker and recorded through a microphone. Recordings of interactive works such as video games are of little use; still, if a work can be experienced by the human senses, it can be recorded. Trusted systems would simply inhibit the unauthorized making of perfect digital copies.

Two things could happen when performances of digital works are re-recorded outside a repository. The first is a loss of fidelity, familiar to anyone who has made cassette-tape copies from a compact disc. The first-generation copy of a digital work re-recorded outside a repository

would not be perfect, although it may be very good. Subsequent generations of uncontrolled copies would be as good as the first-generation copy. The second thing that may occur is the copying of identification information hidden on an original version. This information can be made invisible and inaudible to the human senses but detectable by special equipment. In the event that unauthorized copying is frequent enough to justify intervention, such identification information could be used to trace an illegal copy back to the repository where it was made.

In summary, repositories—like all technology—are imperfect. Yet they can make it much easier to be honest and easier for creators to make a living. Moreover, we would expect them to have other social effects in the long term. More digital works would be created because more potential authors would see a possibility of making a living creating them. That is the main effect—letting loose the light.

There are other predictable effects. With usage rights, electronic distribution of digital works would be much easier than it is now. Using computer networks and the telephone system, consumers could take immediate delivery of digital works from distributors located at great distances. Many distributors of digital works, such as digital bookstores, digital libraries, and digital music stores would serve much larger geographic areas. Consumers would be able to choose from a wide range of potential distributors and to access or purchase quality digital works at a moment's notice and at any time of the day. Distributors who add little value to the chain between producers and consumers would probably be squeezed out.

What kinds of works will first be distributed on repositories? The answer to this question, and many of the factors bearing on it, are still unknown. Repositories will be adopted first in areas where they can solve a pressing problem. In the music industry, publishers have lobbied successfully against the manufacturing of devices that can make digital recordings. There may be an opportunity in this area if the trusted systems developed give advantages to both consumers and publishers. It will probably also be easy to introduce repositories, where other technological changes are leading to system changes. In television, for example, repositories could become more viable when digital television offers fidelity that cannot be captured by the today's videotape formats.

The adoption of usage rights would remove many of the barriers to self-publishing and induce more creators of digital works to self-publish, offering their works for sale on computer networks and reaping the benefits of consumer-based distribution. Self-publishing would not, however, eliminate the role of digital publishers. Publishers often improve the quality of the works they publish and provide brand names recognized by the consumer as indicative of quality and style. The continued need for this kind of quality assurance will give publishers of digital works an enduring role. In fact, new publishers will probably appear. People who review works may put together collections, providing a service similar to brand-name recognition by appraising the works they offer. The net effect of these opportunities for new publishers could be the broadening of influences on popular taste.

The Digital Melting Pot

The metaphor of the melting pot has long been used to describe the culture of the United States created by blending the traditions of people from all over the world. Digital publishing, we believe, is creating a melting pot of genres. The term *multimedia* refers to the mixing of multiple kinds of media—books, newspapers, musical recordings, videos, video games, and computer software—in a single production. But the blending of forms, unlike the blending of water and oil, is creating new forms: for example, interactive movies and travel guides and annotated presentations of plays that include scripts, multiple performances, and reviews, all in a single hyperlinked work.

Some currently distinct genres may be evolving toward similar digital interactive forms. Ultimately, the form of the digital news program and the digital newspaper may be the same. Both may become a digital work broadcast in the air or by cable several times a day and laid out on multiple electronic pages. The pages will contain short film clips of anchors giving the news and, perhaps, animated advertisements and infomercials. Both versions may be interactive and playable on color screens with high-fidelity sound. Today's newspapers and television news programs are the forerunners of these future interactive news forms.

As publication forms blend, what will become of the provisions of copyright law and fair use, which today have different provisions for

newspapers, videos, and computer programs? It is likely that some of today's legal distinctions will not be sustainable in the new digital forms.

In the Fall of 1993 Bruce Lehman, head of the U.S. Patent Office, conducted a public hearing on intellectual property and digital networks. The hearing was called to discuss these issues in the context of the national information infrastructure, the so-called Information Superhighway touted in the press. At this Washington meeting, representatives of the cable industry sat down the aisle from small-town librarians and civil libertarians. Representatives from the music industry rubbed shoulders with people from the computer hardware and software industries. All were aware that their own businesses and institutions were becoming more involved in digital publication and that new media were creating forms and genres that confound present definitions of fair use.

The representatives found that they had very different assumptions about the appropriate means for protecting intellectual property. The music industry uses statistical sampling of radio station broadcasts to check that royalties are properly paid, while the computer industry has no effective organization to check on copyright infringements. Paper-based publishing mostly uses the copyright clearance center, which expedites payment of copying fees. The representatives were well versed in the issues pertaining to their own businesses and recognized that the conditions for the other businesses were different. They knew that as the various media merge into new forms new ways of doing business will emerge. Whatever new rules and ways of business develop, they want to ensure that their own industry thrives.

Planning for the Generations

The laws governing the length of a copyright have changed several times. In 1978 the copyright for a work was established as the lifetime of the creator plus fifty years. Thus, the total period of a copyright may easily amount to a hundred years. In our fast-paced society, we do not often design institutions to last for hundreds of years.

Yet the long view is not unheard of. When we set aside parklands in the public trust, we are planning for the long term. Such consciousness of time is akin to that expressed by Native Americans when they speak of planning for seven generations: three past generations of parents, grand-

parents, and great-grandparents; the generation that makes the decision; and three future generations of children, grandchildren, and great-grandchildren. Adapting this perspective could help us create an institution that has a lasting value for humankind.

There are many stakeholders in digital publishing, including consumers, authors, publishers, distributors, platform vendors, financial institutions, and governments. The approach based on usage rights presumes that all will share a digital property language, compatible platforms from multiple vendors, and broad general agreement about what rights mean. The details of the approach and the particular kinds of rights defined will evolve over time.

The security needs of the approach assume the existence of an authoritative institution that issues digital certificates warranting that particular platforms and software uphold and enforce the concept of usage rights. We call this institution the Digital Property Trust (DPT). Although some initial seed funding would be needed, the DPT could eventually fund its activities from a small tithe on commercial repository transactions or by renewable licenses on platforms and software.

Although planning for a DPT is tentative at this time, a few observations about its role and structure are appropriate. At present, there are many different kinds of social and international structures—government, banking, political, and standards organizations, among others—with their own bases of authority. In the future these organizations should be able to establish and publish their own digital certificates, publishers issuing their own kinds of digital licenses and authors and distributors setting whatever usage rights and fees they please. The role of the DPT would be to promote widespread commerce in digital works. To this end, it would certify and maintain the security of trusted systems and establish a common operational terminology of usage rights to meet social and economic needs.

Achieving this long-term vision for the DPT requires the wisdom and organization to appropriately balance interest groups. The DPT would need to be evenhanded in its treatment of hardware providers, software designers, and publishers. Because the certification process for hardware and software would require DPT to have detailed knowledge of exactly how products work, maintaining evenhandedness may be challenging. At

the same time, the DPT would need the representation and support of the most powerful vendors and publishers for its decisions to carry the weight of authority.

Establishing Trusted Systems

Today's computers and installed software are not programmed to honor usage rights. This fact raises a key question: Assuming that the usage rights language is appropriate and the DPT can be established, how do we go from a world in which most systems are not trustworthy to a world populated by trusted systems? The realistic answer is that the world will not change suddenly. Rather, we need incremental approaches to establish trusted systems widely.

One incremental approach would distinguish between individual and organizational repositories, starting with institutions dealing in documents of high value and limited distribution. These might be bookstores that print documents on demand or legal offices that provide rapid access to thousands of scanned documents. Beginning with organizations like these could demonstrate the viability of usage rights without requiring tamperproof systems for authentication, authorization, and accounting.

Another incremental approach would focus on digital works in a particular market niche. For example, rather than starting with computers, it may be easier to begin with personal entertainment systems for music or video games, for these systems do not need compatibility with general-purpose application programs and operating systems.

A third avenue might be to upgrade existing computer systems by adding appropriate software and hardware. Even though such upgrades do not generally lead to high levels of security, the approach could begin with a large base of systems at low levels of security and provide incentives to upgrade them to more secure systems for works of greater sensitivity or value.

Whichever approach is taken, in the long run repositories offer advantages to both publishers and consumers. Consumers may find they have ready, quick, and cheap access to all kinds of digital works that can be delivered by telephone or computer network any time of the day. Creators may find that consumer-based distribution is a large new distribution channel; anyone who buys a digital work will be able to make a

copy and sell it, automatically routing compensation to its creators. The simple provisions for extracting, editing, and embedding small portions of digital work open doors to creative sampling and reuse of multimedia materials. As creators and publishers learn that safety and wider markets are possible in digital publishing, they may bring about a flowering of new works and old works in new digital form.

Repositories provide a way to let loose the light for present and future generations. There are many institutional challenges, and stakeholders need to work together to bring about the necessary changes. The popular adoption of repositories will start small. In what publishing niches will it begin, providing the sparks to ignite the bonfire of publishing?

A Glossary of Terms

Credit server A secure program and database in a repository to keep track of fees owed for the use of digital works. Typically, a credit server would have a credit limit and would need to be regularly connected to a financial clearinghouse to transfer funds to pay the bills. Like a bank automatic teller machine, a credit server would typically require the user to enter a personal identification code before using it.

Digital certificate A digital document attesting to the truth of something. Each repository would have a digital certificate certifying it as a trusted system and identifying its public key and security level. In general, digital certificates could not be transferred between repositories except by specially authorized repositories. Digital certificates are encrypted in the private key of a master repository, making them difficult to forge and providing a simple means of testing their authenticity—decrypting them by using the public key.

Digital property right or Usage right A specification of a contract to use a digital work in a certain way. Such rights would fall into several categories. **Transport rights**, for example, would include the right to copy, transfer, or loan a work; **render rights** would include playing and printing it; **derivative rights** would include the right to extract, embed, and edit it. Usage rights would be represented in a formal language that can be precisely interpreted by repositories. They could also be displayed to a consumer in a variety of simple ways appropriate for the situation.

Digital property rights transaction or Usage rights transaction A series of actions treated as a unit. For example, in electronic banking, a transfer of money from one account to another is a transaction that credits one

account with the money at the same time that it debits the other account. Each digital property right defines a particular transaction that can be carried out. For example, a usage right to copy a digital work would cause a new copy to be made and a credit server to make a record of the required payment.

Digital Property Trust (DPT) An organization that ensures the health of digital publishing and promotes a lively international commerce in digital works. In conjunction with consumers, publishers, creators, and platform vendors, it would set the standard for the evolving digital property language and issue digital certificates to conforming platforms. It would also maintain the master repositories and perhaps ensure security and financial transactions.

Digital license A digital certificate identifying the bearer repository as licensed and thereby authorized for carrying out certain rights. For example, certain digital works might require that copies of a work could only be made and sold when the repository contains a particular digital distributor's license.

Digital ticket A digital certificate or coupon that can be used once to authorize a particular transaction. For example, a digital work may include an upgrade ticket that authorizes the user to replace the digital work with an updated version. When digital tickets are used, they would be punched by a digital ticket agent and could not be used again.

Digital work Any work that can be represented in digital form; for example, a document like a book, magazine, or newspaper. It could also be a recording of music, a movie, a computer game, or any computer program. Sometimes the word **software** is used in a general sense to mean a digital work.

Encryption A process of encoding a digital work by a secret code to render it unusable by anyone without the code. Decoding a work to restore it to usable form is called **decryption**. The preferred method of encoding is public-key encryption, in which there are two keys, a public key and a private key. When the private key is used to encrypt a work, the public key can be used to decrypt it.

Master repository A very high-level repository with the highest security level. Master repositories would be kept by the Digital Property Trust (DPT) and would be used to issue certificates. The public keys of the master repositories would be assigned to all trusted systems when they are manufactured, enabling them to exchange and identify authentic digital certificates in their transactions with other repositories.

Repository Any trusted system used for storing and playing digital works. For example, repositories could be portable entertainment de-

vices, laptop readers, personal computers, credit-card-sized devices, or mechanisms that fit into home entertainment equipment for controlling digital television or music. Repositories would store digital works, together with their usage rights, and include credit servers for keeping track of fees for use.

Security level Different degrees of physical security—ranging from low security to very high security—for protecting digital works against unauthorized use. Repositories for handling extremely valuable works need greater security than those for ordinary and portable use.

Shell A kind of digital container for storing digital works in the filing system of a repository. They could contain both digital works and other shells. Tags specifying usage rights and fees would be attached to each work and its shell. When someone asks a repository to use a digital work, the repository would check the rights and fees recorded on the tags.

Trusted system A system that can be relied on to follow certain rules at all times. In the context of digital works, a repository would be a trusted system that governs the uses and fees for digital works. All digital works would have tags describing uses and fees in a usage rights language; trusted systems would carry out the instructions on these tags infallibly.

Usage right See Digital property right.

Reflections

The network for digital publishing is still being invented. As its creators we will need much wisdom to guide that invention. The hundreds—if not thousands—of traditional myths about creation have tried to answer questions such as: How did the world come to be? Why are humans what they are? Where did they come from? Generally, these stories are intended to explain the actions of the gods to mortals. According to them, we were formed from the thoughts of the gods, arose out of chaos, fell from grace, or were fashioned from mud along with the animals. The myths embody stories of people searching for a land in which to live and stories about how they should live in it to keep the world in balance.

But there are few stories instructing us about how to act as able and responsible creators. Perhaps creation myths are the wrong places to look for guidance; instead we should pay attention to stories about how people in the past worked together as communities. “Letting Loose the

Light draws on those stories, on our understanding of the marketplace and on the history of intellectual property during the French Revolution.

In the past few years, several ways of governing the use of information on CD-ROMs and networks have come into use. In one version, purchasers call a telephone number to pay for keys that unlock specific software distributed on a CD-ROM. Different keys unlock different software. Because there are different versions of the CD-ROM discs, two people buying them are unlikely to be able to share keys. This approach does not, however, control subsequent redistribution of software if people have sufficient storage to copy it from a CD-ROM. Other approaches requiring special hardware and software have also been offered to meter the use of software on a CD-ROM. Efforts to control software with license servers are now widely used in organizations with closed local computer networks. They allow programs to govern the number of simultaneous users of proprietary software or data bases, making it possible, for example, for any five people of an organization of a hundred members to use a database at the same time.

Security and convenience are the key issues for approaches to digital publishing; security is an issue for publishers, and convenience is an issue for users. These concerns are often seen to be in tension. In the simple and convenient low-security approaches, users can load metering software onto existing personal computers; unfortunately, they could also load pirateware, software that defeats security measures. In many proposed systems, pirateware posted on a computer bulletin board could be used to compromise every site on a network. Fear of such attacks makes publishers reluctant to publish their works on such a network. In other approaches, works to be used must be distributed by the equipment manufacturer, which also acts as the accounting and billing company. As most publishers view control over marketing and distribution as a key to their success, they are often unwilling to engage in such arrangements. In contrast with such approaches, the digital property rights approach described in this article separates the publishing business from the financial and platform businesses. It will also require a good deal of coordination among these business sectors to make such trusted systems available.