

Authenticity in a Digital Environment

May 2000

Council on Library and Information Resources
Washington, D.C.

ISBN 1-887334-77-7

Published by:

Council on Library and Information Resources
1755 Massachusetts Avenue, NW, Suite 500
Washington, DC 20036

Web site at <http://www.clir.org>

Additional copies are available for \$20.00 from the address noted above. Orders must be prepaid, with checks made payable to the Council on Library and Information Resources.



The paper in this publication meets the minimum requirements of the American National Standard for Information Sciences—Permanence of Paper for Printed Library Materials ANSI Z39.48-1984.

Copyright 2000 by the Council on Library and Information Resources. No part of this publication may be reproduced or transcribed in any form without permission of the publisher. Requests for reproduction should be submitted to the Director of Communications at the Council on Library and Information Resources.

About the Authors

Charles T. Cullen is president and librarian of the Newberry Library in Chicago. A legal historian, Mr. Cullen earned his Ph.D. degree from the University of Virginia. He has taught at Princeton University, the College of William and Mary, and Averett College. While at the College of William and Mary, he worked on the Papers of John Marshall, and became editor of that project in 1976. At Princeton, he was editor of the Papers of Thomas Jefferson.

A charter member of the Association for Documentary Editing, Mr. Cullen served as its president in 1982–1983. He currently represents that organization on the National Historical Publications and Records Commission. He has served on the New Jersey Historical Commission, and was vice-chairman of the board of the Founding Fathers Papers, Inc., from 1986 until 1992. He is currently vice-president of the board of the Modern Poetry Association, and chairs the board of the Heartland Literary Society.

Mr. Cullen has published or contributed to 30 books and articles. He has lectured widely on subjects relating to the Age of Jefferson, the scholarly use of computers, and the role of humanities research libraries. His scholarly computing expertise was recognized by the Association for Documentary Editing, which in 1987 awarded him its Distinguished Service Award “for outstanding contributions to the field of documentary editing through the use of computers.”

Peter B. Hirtle is co-director of the Cornell Institute for Digital Collections (CIDC). CIDC is responsible for developing digital resources, supporting their use campus-wide, and conducting applied research that advances the production and utility of such resources. Mr. Hirtle is associate editor of *D-Lib Magazine*, a monthly magazine about innovation and research in digital libraries.

Prior to his arrival at Cornell, Mr. Hirtle worked at the National Archives and Records Administration (NARA), first for the technology research staff (where he helped complete its most recent digital imaging report), and then as coordinator of electronic public access for the agency. He has also served as curator of modern manuscripts at the National Library of Medicine. Hirtle has a master’s degree in History and an M.L.S. with a concentration in archival science. He has served on several of the units sponsored by the Society of American Archivists, most recently as a member of the executive committee of its governing council, and on the Commission on Preservation and Access/Research Library Group’s Task Force on Digital Archiving. He currently serves on the Research Libraries Group/Digital Library Federation Task Force on Long-Term Retention, and is a member of the National Initiative for a Networked Cultural Heritage’s Working Group on Best Practices in Networking Cultural Heritage.

David Levy is an independent consultant in the areas of documents, digital libraries, and publishing. Between 1984 and 1999, he was employed as a researcher by the Xerox Palo Alto Research Center. The

focus of his work over the past decade has been the nature of documents and the tools and practices through which they are created and used. His research interests include digital libraries; the reuse of documents; document design, structure, and standards; work practice studies; the combined use of paper and digital media; and the relation between technology and the character of modern life.

Mr. Levy holds a Ph.D. degree in computer science from Stanford University and a Diploma in calligraphy and bookbinding from the Roehampton Institute, London. He is currently a member of a National Research Council commission charged with advising the Library of Congress on its technology strategy. He is completing a book, *Scrolling Forward: Making Sense of Documents in a Digital Age*, which will be published by Arcade.

Clifford A. Lynch has been executive director of the Coalition for Networked Information (CNI) since 1997. CNI, jointly sponsored by the Association of Research Libraries and EDUCAUSE, includes about 200 member organizations concerned with the use of information technology and networked information to enhance scholarship and intellectual productivity. Before joining CNI, Lynch spent 18 years at the University of California Office of the President. For the last 10 of those years, he was director of library automation. In that post, he managed the MELVYL information system and the inter-campus internet for the University. Mr. Lynch, who holds a Ph.D. degree in computer science from the University of California, Berkeley, is an adjunct professor at Berkeley's School of Information Management and Systems.

He is a past president of the American Society for Information Science and a fellow of the American Association for the Advancement of Science. He currently serves on the Internet 2 Applications Council and the National Research Council Committee on Intellectual Property in the Emerging Information Infrastructure.

Jeff Rothenberg is a senior computer scientist at The RAND Corporation in Santa Monica, California. He has a background in artificial intelligence and modeling theory. His research has included developing new modeling methodologies, studying the effects of information technology on humanities research, and investigating information-technology policy issues. He has been researching the problem of digital longevity since 1992, when he coauthored a prize-winning paper in *The American Archivist*. He has since explored the dimensions of the problem with archivists, librarians, and others in the United States and Europe. He published a widely cited article on the subject in *Scientific American* in 1995. He also appeared in the documentary film "Into the Future," which was produced by the Council on Library and Information Resources in 1998. Mr. Rothenberg recently completed a project for the Dutch National Archives and Ministry of the Interior that recommended a strategy for the long-term preservation of digital records in The Netherlands. He is currently working with the Dutch Royal Library on related issues.

Contents

Introduction	vi
Authentication of Digital Objects: Lessons from a Historian's Research, by Charles T. Cullen	1
Archival Authenticity in a Digital Age, by Peter B. Hirtle	8
Where's Waldo? Reflections on Copies and Authenticity in a Digital Environment, by David M. Levy	24
Authenticity and Integrity in the Digital Environment: An Exploratory Analysis of the Central Role of Trust, by Clifford Lynch	32
Preserving Authentic Digital Information, by Jeff Rothenberg	51
Authenticity in Perspective, by Abby Smith	69
Appendix: Conference Participants	76

Introduction

What is an authentic digital object? On January 24, 2000, the Council on Library and Information Resources (CLIR) convened a group of experts from different domains of the information resources community to address this question. To prepare for a fruitful discussion, we asked five individuals to write position papers that identify the attributes that define authentic digital data over time. These papers, together with a brief reflection on the major outcomes of the workshop, are presented here.

Our goal for this project was modest: to begin a discussion among different communities that have a stake in the authenticity of digital information. Less modestly, we also hoped to create a common understanding of key concepts surrounding authenticity and of the terms various communities use to articulate them.

“Authenticity” in recorded information connotes precise, yet disparate, things in different contexts and communities. It can mean being original but also being faithful to an original; it can mean uncorrupted but also of clear and known provenance, “corrupt” or not. The word has specific meaning to an archivist and equally specific but different meaning to a rare book librarian, just as there are different criteria for assessing authenticity for published and unpublished materials. In each context, however, the concept of authenticity has profound implications for the task of cataloging and describing an item. It has equally profound ramifications for preservation by setting the parameters of what is preserved and, consequently, by what technique or series of techniques.

Behind any definition of authenticity lie assumptions about the meaning and significance of content, fixity, consistency of reference, provenance, and context. The complexities of these concepts and their consequences for digital objects were explored in *Preserving Digital Information: Report of the Task Force on Archiving of Digital Information*, published by the Commission on Preservation and Access in 1996. There is no universally agreed-upon mandate about what must be preserved and for what purpose. For example, an archivist will emphasize the specifications of a record that bears evidence; a librarian will focus on the content, knowing that it could serve multiple purposes over time. That being the case, there may be many ways to describe an item being preserved and what aspects of that item must be documented to ensure its authenticity and its ability to serve its intended use over time. For certain purposes, some argue, migration may suit the preservation needs of a digital object. For those objects most valued as executable programs, others argue, emulation is pref-

erable. Beyond the technical options undergirding metadata and preservation decisions, numerous nontechnical questions beg to be asked. The issue of authenticity must be resolved before humanists and scientists can feel confident in creating and relying upon digital information.

Creating a common understanding about the multiple meanings and significance of authenticity is critical in the digital environment, in which information resources exist in many formats yet are interactive. From peer-reviewed journal articles to unpublished e-mail correspondence, these resources are integrated; they can interact and be modified in a networked environment. We wanted to know whether the distinctions that have proved to be helpful heuristic devices in the analog world, such as edition or version, document or record, could help us define a discrete piece of digital information. Can we define the distinct attributes of an information resource that would set the parameters for preservation and mandate specific metadata elements, among other important criteria?

We charged the five writers—an archivist, a digital library expert, a documentary editor and special collections librarian, an expert on document theory, and a computer scientist—to address one essential question: What is an authentic digital object and what are the core attributes that, if missing, would render the object something other than what it purports to be? We asked each to address this question from the perspective he found most congenial. We emphasized our interest in the essential elements that define a digital object and guarantee its integrity, but left the writers free to grapple with that question as they saw fit.

In considering this central issue, we asked that they think about the following:

- If all information—textual, numeric, audio, and visual—exists as a bit stream, what does that imply for the concept of format and its role as an attribute essential to the object?
- Does the concept of an original have meaning in the digital environment?
- What role does provenance play in establishing the authenticity of a digital object?
- What implications for authenticity, if any, are there in the fact that digital objects are contingent on software, hardware, network, and other dependencies?

These are some of the issues that we anticipated would arise in the course of the workshop.

In thinking of which communities to include in the workshop discussion, CLIR sought expertise from the major stakeholders in these issues: librarians, archivists, publishers, document historians, technologists, humanists, and social scientists. Because so many concepts of authenticity derive directly from experience with analog information, we called upon experts in the traditional technologies, such as printing and film, to elucidate key concepts and techniques

for defining and securing authenticity of information bound to a physical medium.

The authors were given time to revise their papers in light of the discussion and any comments they received from the participants. Some chose to revise their papers, and others did not. The task of writing a position paper on this complex subject (a paper that we limited in size but not scope) was quite difficult. Each writer took a different approach to the subject, and the papers differ greatly one from another. This seeming disparity proved a boon to the discussions. During that time, each writer had a chance to “unpack” the various nuances of thought that the papers held in short form only, and participants were confronted with the diverse ways that such common words as *copy*, *original*, *reliable*, or *object* are used. Much of the substance of the discussion is included in the concluding essay.

As one participant remarked, authenticity is a subject we have avoided talking about, primarily because the issues it raises appear so intractable. We are deeply grateful to Messrs. Cullen, Hirtle, Levy, Lynch, and Rothenberg for agreeing to form the advance party as we ventured into terra incognita. They were willing not only to think deeply about a vexing issue but also to commit their thoughts to writing and to careful scrutiny by others. Their papers, together with the oral summaries they delivered at the meeting, marked out several different trails to follow, each of which opened onto ever-larger vistas—some breathtaking, some daunting. We are also grateful to the participants, many of whom came from very distant places. Their thoughtful preparation and frank discussion confirmed our sense that authenticity is important for many communities, and that they are ready to engage the issue.

Abby Smith
Director of Programs

Authentication of Digital Objects: Lessons from a Historian's Research

by Charles T. Cullen

The issues stemming from authenticating digital objects are quite similar, and in some cases identical, to those relating to holographs or printed books. Everyone dealing with important material in any form should approach it with a bit of skepticism, but scholars especially need to question what it is they are using. In other words, they need to authenticate all documentation they use in the processes of learning and of creating new scholarship. An authentic object is one whose integrity is intact—one that is and can be proven or accepted to be what its owners say it is. It matters little whether the object is handwritten, printed, or in digital form.

Over time, we have established various measures of authenticity for analog forms that we trust almost without question. Our trust is, however, much greater for printed books than for handwritten objects. In fact, handwritten objects raise many of the same questions of authenticity as digital objects do. The difference is that in the case of the former, the answers may be more easily found. Take Thomas Jefferson's manuscript "Report on the Navigation of the Mississippi," for example. Could he have written it? Is it his handwriting? Is the paper watermarked, and from the appropriate time period? Is the ink contemporary? Do other copies of the manuscript exist? Has its recipient or any other contemporary endorsed it? Is there other internal evidence? Who has described it for us? Has it been identified by a trusted third party?

Is a book authentic? Who published it, and who wrote it? Can they be trusted (are they worthy of one's research time)? Is the rare book what it purports to be? Is the manuscript correspondence actually by the person to whom it is attributed, and is its date accurate? These questions are now being asked more openly of objects that originate in digital form because we have not yet adopted practices or standards for providing ready answers to them. When objects are

presented digitally, deciding what is required to authenticate them may be informed from past practices with non-digital objects.

Two experiences with paper objects inform my views of this subject. The first is a multi-page autograph document that lies in the John Marshall Papers at the Virginia State Library. It is labeled in the hand that wrote the entire piece, "John Marshall's Notes on Evidence in Commonwealth v. Randolph, 1796." Although the title itself might raise some question about who penned it (How often does an author—even an eighteenth-century author—use his own name in a title of one of his documents?), this document has been used for decades for the source of historical articles and at least one full-length book on the investigation of Richard Randolph for murder. Randolph, a member of the famed Randolph family of Virginia, was related to Marshall and to Thomas Jefferson and to many other members of Virginia's "first families."

Examination of the writing by those familiar with John Marshall's hand, however, quickly reveals that he did not pen this document. Knowing who did write it is important, but does not help make it more authentic as a Marshall document. The possibility that someone in possession of Marshall's holograph could have copied the document raises new questions, not the least of which assigns significant importance to the value of the original of a document, regardless of its form. Internal evidence, obtained by a close reading of this document, reveals that it might be a partial transcript of a hearing in Cumberland County Court where witnesses are questioned by attorneys, and it has been used by historians as a partial record of Randolph's "trial." But Marshall's name never appears as one of the questioners, and a knowledge of Virginia law at the time would reveal that whatever was taking place could not be an actual trial, because white men could not be tried for felonies at the county court level during that period. In short, efforts to authenticate this document raise more questions than they answer. At the least, such efforts reveal that the document may not be what many had long thought it to be, and that it may not be even what its title says it is.

This example is somewhat esoteric, to be sure, because it is unlikely that one would use a similar digital document without asking the questions that eventually were asked of the document attributed to Marshall. But the questions asked of it suggest attributes that must be held by a holograph as well as by a digital object that is to be regarded as authentic. Is it the author's work or a copy? Is it what its title purports it to be? What tests can be applied to answer these questions convincingly?

A second example is more to the point. In the collection of Thomas Jefferson's papers at the Library of Congress is a document that appears to be a list of letters written and received between 1791 and 1793, a period of time during which Jefferson was Secretary of State. An examination of the handwriting reveals that it is most likely Jefferson's, but the list is unlike his other journals of letters sent and received. A close look at the original document suggests that it was written in only one or two sittings (the ink changes only once or

twice), rather than over three years. The most significant evidence relating to this document's authenticity lies in the paper itself. Holding it before a light source reveals a watermark that indicates the paper was manufactured in 1804. The document, therefore, could not be an authentic 1791–1793 document.

Almost all these tests can be applied to digital objects, and they need to be. But because digital objects bear less evidence of authorship, provenance, originality, and other commonly accepted attributes than do analog objects, the former are subject to additional suspicion. Tests must be devised and administered to authenticate them.

In many cases, problems of authentication arising from objects that originate as digits are obvious. In trying to find solutions to those problems, however, we must carefully test all suggestions to ensure that they do not themselves open new issues that may be inherent in this medium. The problems of preserving digital objects have received more attention than have questions of authentication (people, I suppose, are less worried about authenticity than about preservation). But why preserve what is not authentic? Might the preservation of a digital object imply an endorsement of authenticity, even if nothing else is done to it? More than one archivist has stated that the only sure means of preserving a digital object is to save a printed copy. Concerns with format codes, migrations from version to version, dependence on hardware—would all be solved by printing a copy (or many copies) and putting it (them) in a safe place. Do that to a digital object before confronting the questions of authenticity, and all that is valuable may be lost. Converting a digital object from one program to another, or migrating it from version to version, could present problems of authenticity that may or may not be solved by careful attention to provenance.

A digital object must be authenticated at the time of its creation by a means that will convey a high degree of confidence to all users, including subsequent use by the originator. Clifford Lynch wrote an interesting and convincing article on the integrity of digital information, published in the December 1994 issue of the *Journal of the American Society of Information Science*. He seems to assume, from traditional experience perhaps, that readers will be responsible for authenticating copies being used on the basis of cataloging data to which they must be alert. Retrieving electronic files by title, for example, might lead one to a revised work, different from the original. The reader must exercise caution, Lynch writes, and be ready to detect signs of alteration. "The expectation should be that violations of integrity cannot be trivially accomplished," he says. Accepting this in the world of printed objects is relatively easy. It is much more difficult in the realm of electronic digital information.

Andy Hopper of Cambridge University suggests an authentication strategy that is worthy of consideration, if not adoption. In his system, the concept of a trusted third party is borrowed from the print world. According to this concept, trusted librarians help authenticate their print holdings through recognized acquisition pro-

cesses, accepted cataloging procedures, and careful stewardship of their collections, especially those in manuscript form. If a special collection librarian tells us, either directly or by means of a catalog card, that the book in hand is one of two extant copies of Ariosto's *Orlando Furioso* printed on vellum in Venice in 1542, and that it was prepared for the dauphin of France, the library's and the librarian's reputation go a long way toward instilling some degree of confidence that the document is indeed authentic. Moreover, all of this information may be checked. If another librarian delivers to a reader a box of letters cataloged as Ernest Hemingway's, authentication is assumed until internal or physical evidence suggests someone has made a mistake. Knowing that the materials—hard-copy objects—have gone through a process of description and identification, if not authentication, conveys a sense of trust that they are authentic, at least until proved otherwise. Some of the problems of description that help authenticate printed special collection objects have similar, if not identical, examples in the digital world. Take one final example as evidence: in the Newberry Library's special collections is a printed copy of the classic book on rhetoric in Renaissance England, *Arte of Rhetorique* by Thomas Wilson (1525–1581). This particular copy is identified as having belonged to Elizabeth I as part of her royal library, and it is authenticated as such by its original binding, which bears the mark of the royal arms. Book historians know that until the time of James I, the royal arms were put only on the books within the monarch's own library. (After 1603, King James allowed them to be placed on books bound for other members of court.) Elizabeth's coat of arms on the binding of this copy of Wilson's *Rhetorique* therefore marks it as authentic, as long as external evidence does not dispute it. (If it could be shown that the binding was not sixteenth century, for example, or that it resembled the work of a sixteenth-century forger, the authenticity might be questioned).

Some accepted system of similar assumption of authentication needs to exist in the digital world, but it is more difficult to achieve because digital material is more changeable, accidentally or deliberately. Andy Hopper and others suggest that some means of marking digital objects could help solve many of the problems of authentication. Hopper argues that libraries might serve as authenticators by marking digital objects by some means that would remove doubt as to their characteristics at time of origin. A method must be developed whereby a trusted third party, ideally a trusted librarian, would put a marker on a digital document—a marker that could not be predicted or devised (guessed)—that would mark the document's time and date. The marker might be a number based on sonic rays at various times during the day, a number large enough to prevent guessing (Hopper suggests 100 digits). A professor writing a paper could send the document to the librarian to be marked, and it could then be returned to and held by the author. In the future, the object could be authenticated by its marker, regardless of who held it. Any change in the document would remove the marker. This procedure would be used by librarians who receive digital objects from donors. The

marker would ensure that digital objects are as authentic as analog objects at time of cataloging.

Despite its science-fiction flavor, such a method seems to meet accepted tests of authenticity. A trusted third party can claim nothing more about an object, analog or digital, than what can be cataloged, and that information derives largely from physical evidence. Identifying an object in a catalog record or a collection description puts a marker on it that most of us use as the first step in the process of authentication. Is the document what it purports to be or what its owner claims it to be? Scholars often require means to test the cataloger, and the physical attributes of analog objects offer more opportunities to do such testing than do those of digital objects. Handwriting, publishing history, bindings, watermarks, inks, and various forms of internal evidence provide answers to questions of authenticity in analog objects that are lacking in digital objects. Digital objects have attributes that can be used to help with authentication, but none is sufficiently trustworthy or stable to be acceptable unless a workable system of certain marking can be devised.

Certifying that a digital object is the product of its author is difficult when the object originates in electronic form. Without a deliberate and distinctive marking caused by the author that could not be guessed by another or altered by anyone, it seems impossible to authenticate an electronic document beyond doubt. Authors of files or images must take steps to establish authorship of their work; if not, our only option is to accept the assertions of others. Electronic files left behind by someone who has not taken action to establish authorship are subject to suspicion if authorship is asserted by anyone else at the time of "cataloging." This leaves us where we have been all along—at the mercy of catalogers. But, in the case of a digital object, we are actually worse off than we would be if we were dealing with an analog object. This is because we lack the physical evidence provided by analog objects—evidence that offers the means to test the cataloger. This ability to test both reassures the user and helps keep the cataloger honest. I find no corollary in the digital object realm.

The concern over authenticating digital copies of analog objects is almost as important as that relating to objects that originate in digital form. Scholars are keenly interested in having access to documentary evidence in digital form, and librarians have begun to consider digitization a desirable means of preservation, in spite of the recognized problems inherent in it. Those who hope to use this material, once it has been digitized, must be able to rely on its authenticity, just as they have become accustomed to do in all the forms currently available. Documentary editors, as well as librarians, have new responsibilities as they publish and provide access to their materials in digital form with all the value they have added intact. The work of documentary editors offers some insight into the questions raised over authenticity of digital objects, especially those that derive from analog or holographic objects.

The first task of a documentary editor who is working on an edition of a subject's papers is to locate all the objects that have ever ex-

isted as part of the corpus, incoming as well as outgoing. This sometimes requires reliance on copies of papers that evidence suggests once existed in original form but which have not been found. Once the collection is organized, each item must be dealt with separately. That is the first stage for authentication tests, starting with the question of whether the item is what it appears to be. All the available physical attributes assist in answering these questions, but sometimes only internal evidence leads to a final answer (as in the Marshall and Jefferson examples described earlier). The editor is obliged to share these findings with readers and to describe the item in such a way that few, if any, questions remain about the document as object. Not unimportant in this description is all available information about other copies of an item, be they photocopies, carbons, letterpress, polygraph, drafts, or additional holographs. Knowing as much as the editor about all copies is the only sure way for other readers to test the "cataloger's" description, and only by having this information available can a reader have full confidence that all questions of authenticity have been asked. In preparing digital files of historic documents, editors begin their publication by attaching a full document description to a transcript. This is the scholar's seal of authenticity, as it were, or at least as much of a seal as a scholarly editor can provide.

Preparing a digital transcript of a historic object introduces new problems to the issue of authentication. How do we know the transcription is accurate and that it is exactly what the editor prepared originally? The method of providing access to journal articles adopted by JSTOR may offer the best answer for authenticating modern digital transcripts of manuscripts or printed material that originated in analog form. They provide the user with a digital transcription of the text, which is fully searchable and otherwise subject to all the vagaries of digital files. They also provide an image of the original text. If both copies could carry some form of marking that could not be manipulated, the problem of authentication would be solved. This system should work quite well for documentary editors and the readers of their digital publications. Providing an image of the document that is transcribed would be an important improvement over present forms of presentation, because it would permit easy verification of transcriptions. Inaccurate transcriptions are the downfall of documentary editors (as they should be), and mistakes often go undetected. The reader, who may have a high level of confidence in the scholarly work of the editor, is left to assume that the transcription is accurate and authentic. Having a means of testing this assumption would be a great improvement.

Related problems that arise from considerations of authenticity seem to offer little to assist us in answering the primary question. Creating a digital file, and even marking it in such a way that will ensure authenticating it as my own, will mean little if the file itself cannot be read at any point in the future. If the file cannot be read, it cannot be authenticated as mine. (It would be even more maddening if the file could be authenticated but not read.) The same can be said

for provenance. If a file can be marked in such a way that its authenticity is assured, issues of its subsequent provenance might not matter in questioning its authenticity. But if a file cannot be read, its provenance will mean little, even if it can be tracked over a long period. Without a marker of authenticity, provenance of a digital object would be of limited use in establishing authenticity. It would help test the cataloger, but the current technology would render uncertain any assertions of authenticity. The instability of software alone would introduce questions that would challenge any claims of authenticity suggested by a trusted provenance.

Paul Conway (1999) says the existence of digital objects moves challenges of preservation from guaranteeing the physical integrity of objects to assuring their intellectual integrity, including their authenticity. He adds that librarians can control this by “authenticating access procedures and documenting successive modifications” to digital files. Authenticating access procedures may affect provenance more than the integrity of the digital object itself, but it would be difficult to guarantee authentication with only this control. It seems that, in this argument, the alteration of an original record is acceptable as long as it is documented. Acceptance of changes with documentation is unreasonable over time and places unnecessary burdens on users. In this case, as in others, preservation without authentication results in a loss of intellectual integrity.

We are not close to having a means of marking digital documents that cannot be challenged—a means that would establish authenticity. Absent such a technique, we are left to consider what other attributes, if any, might approach the establishment of authenticity. Few suggest any high degree of confidence that would come close to what we have for analog materials, but consideration of the problem raises some issues that relate to other concepts that bear on the problem. How confident can one be when an object whose authentication is crucial depends on electricity for its existence? Surely there are higher degrees of confidence in some cases than in others, but something more than provenance or traditional testing methods established for analog objects is needed. I believe it is easier to describe the characteristics of an authentic digital object than to support the authentication beyond a reasonable doubt. My definition is conditional; it depends on an object’s capability of being proved to be authentic. Establishing a method of authentication of digital objects that would be unconditional may be possible. At the least, we must agree on some means of testing the authentication of digital objects. The consequences of not doing so are dire.

REFERENCE

Conway, Paul. 1999. *The Relevance of Preservation in a Digital World*. Technical Leaflet, Section 5, Leaflet 5, p. 8. Andover Mass.: Northeast Document Center. Available from <http://www.nedcc.org/plam3/tleaf55.htm>.

Archival Authenticity in a Digital Age

by Peter B. Hirtle

Archival Authenticity: An Example

Downtown Baltimore is a vibrant, dynamic place filled with new office towers and hotels that rise above shops, plazas, and museums. At the heart of Baltimore is the Inner Harbor, an area that is crowded year-round with residents and tourists who are sightseeing, dining, shopping, or watching baseball at nearby Camden Yards. Over the past two decades, the Inner Harbor has become the living center of a revitalized downtown.

The defining feature of Baltimore's Inner Harbor, unlike that of so many American cities, is not a glass structure, a shining space needle, or a distinctive sculpture. The Harbor is marked instead by the sturdy masts and graceful spars of the USS *Constellation*, a historic wooden-hulled naval vessel permanently moored there.

The famous ship arrived in Baltimore in 1955, and for the next 35 years, the city celebrated its frigate, taking pride in the illustrious history of a ship that had been built in Baltimore in 1797 as a sister ship to the equally famous USS *Constitution* anchored in Boston. The story of the *Constellation* took a different turn in 1991, however, with the publication of *Fouled Anchors: The Constellation Question Answered*, a report by Dana Wegner, the chief of ship models at the U.S. Navy's David W. Taylor Research Center in Carderock, Maryland. Rumors had circulated for half a century that the *Constellation* was not what its promoters claimed it to be, and Wegner's report confirmed them. Investigators from the Navy discovered that the supposed Revolutionary War-era frigate in Baltimore Harbor was actually a Civil War era sloop that had been built in Norfolk, Virginia, in 1854. All it shared with the frigate built in Baltimore in the eighteenth century was its name. It resembled a Revolutionary War-era frigate because during early renovations, some of the ship's admirers had "restored" the *Constellation* to appear to be almost 60 years older than it was; for

example, they added a second gun deck and made other alterations. For most of its tenure in Baltimore, the *Constellation* was living a lie (Wegner 1991; LeDuc 1999).

Many themes are at work in the story of the true identity of the *Constellation*. Early citizens of Baltimore, for example, seemed to have a stronger need to connect to the Revolutionary War than to the Civil War. They may have felt that “older is better,” and that the ship would be of greatest interest if it was thought to have a Baltimore connection (i.e., if it had been built there). Nonetheless, their distortion of history came at the expense of the *Constellation*’s very interesting own history. It was, for example, the last and largest all sail-powered sloop commissioned by the U.S. Navy, and while it did not engage in a famous sea battle, as did its predecessor, it did work to interdict the slave trade during the mid-1800s.

The most interesting themes in the *Constellation* story, however, revolve around the issue of authenticity—not the authenticity of the ship itself, but rather the authenticity of the *documentation* about the ship. For it was not just the appearance of the ship that was “forged,” but also the written record concerning the ship.

Some of the changes to the written record may not have been an intentional effort at deceit. Between 1854 and 1908, for example, the annual reports of the Navy listed the ship as having been built in Norfolk in 1854; however, from 1909 onward, the reports listed the ship as having been built in Baltimore in 1797. Was this an intentional effort to deceive or an honest effort to correct what naval officers may have thought was a past mistake? Wegner could not determine the answer.

In the 1950s, however, documents began to appear that Federal Bureau of Investigation (FBI) investigators later determined were forged. One document, allegedly written in 1918, was found to have been written with a typewriter made after 1946. Some of the forged documents in the possession of researchers bore forged stamps indicating that they were copies of records found in the National Archives. Other forged documents were inserted into historical files at the National Archives and at the Franklin Roosevelt Presidential Library, where they were subsequently “found” by researchers.

The need to alter the archival written record to conform to a particular historical interpretation speaks to the power of archives to authenticate. At rest in Baltimore Harbor was a physical artifact, a wooden ship, measuring over 180 feet long and weighing several hundred tons. The existence of the artifact per se, however, was not enough to establish its authenticity. To confirm beyond doubt the nature and history of the *Constellation*, both supporters and critics of the “*Constellation* as frigate” theory turned to a few sheets of paper housed in a few archives.

What characteristics of traditional analog archives give them the power to authenticate? And how can this power be maintained in the digital world, both for archives and for other cultural heritage repositories in general?

The Nature of Archives

To understand why users turn to and trust information found in analog archives, it is necessary to understand the nature of archives. In the vernacular, the word *archives* has come to mean anything that is old or established, be it collections of old movies (such as the Pacific Film Archive), a journal that publishes what the editors hope will be papers of enduring value (for example *Virchows Archiv*, the official journal of the European Society of Pathology), or even rock-and-roll oldies on cable television (in the VH1 Archives) (Maher 1997). Even information professionals have not been loath to extend the definition of archives beyond that found in the American Library Association (ALA) Glossary or other official lexicons when they speak of “digital archiving,” a generic term for the preservation of electronic information.

While archivists often inherit responsibility for old things, a collection of historic documents or artifacts, in and of itself, does not make an archives. A true archives is a contextually based organic body of evidence, not a collection of miscellaneous information. A manual written by Dutch archivists almost a century ago codified existing German and French archival theory and developed a modern basis for archives. According to these authors, archives are “the whole of the written documents, drawings and printed matter, officially received or produced by an administrative body or one of its officials . . .” (Muller, Feith, and Fruin 1968). This definition has been adopted in one form or another by most of Western society.

Found within this definition are the essential elements that define an archives and are the source of much of its power to authenticate. First, archives consist of documents. For the Dutch, these documents had to be written or printed; modern archivists extended the definition to include multimedia records, including sound recordings and motion pictures. More recently still, archivists (and the courts) have added electronic records to the definition of documents. A recent court case even argued (unsuccessfully) that “cookies,” the small transactional files created by many Web browsers when surfing the Internet, were government records when found on a computer used by a government official; others have argued that voice-mail messages are documents (Welch 1998). In short, archives consist of documents, regardless of their form.¹

The documents constituting a formal archives are further distinguished by the fact that they have to have been officially produced or received by an administrative body. Such documents become records. According to the most recent glossary of archival terms, published by the Society of American Archivists, a record is a “document created or received and maintained by an agency, organization, or individual in pursuance of legal obligations or in the transaction of business” (Bellardo and Bellardo 1992). When someone requests a Social Security card, when a business reports its revenues for tax

¹ Of course, the question of what constitutes a “document” can be problematic (Buckland 1997).

purposes, or when President Clinton issues a proclamation, documents are created. These documents are records because the agencies or officials involved in each transaction are fulfilling legal obligations as they conduct their business. Similarly, when a faculty committee approves tenure for an assistant professor, or when an organization issues an invitation to a meeting, a record is created.

Note that under this definition, the archivist is not concerned about the value, accuracy, or utility of the content of the record. A document may contain lies, errors, falsehoods, or oversights—but still be evidence of action by an agency. Nor does a record have to be particularly interesting or important, or even something that anyone would ever want to consult again. Pure archival interest in records depends not on their informational content, but on the *evidence* they provide of government or business activity. As the Australian archivist Glenda Acland has noted, the “pivot of archival science is evidence, not information” (Acland 1992).

For a time, the essence of records as evidence slipped from center of the archival vision. Ironically, the challenges inherent in dealing with the most modern of records—electronic records—forced creative archivists to reinvestigate basic archival principles. Perhaps the most notable of these individuals is David Bearman, author of many publications on electronic records. His collection of essays on *Electronic Evidence: Strategies for Managing Records in Contemporary Organizations* is particularly noteworthy (Bearman 1994). Similar analysis has been conducted by the Australians Sue McKemish, Frank Upward (McKemish and Upward 1993), and Glenda Acland, and by the archival educators Luciana Duranti in Canada (Duranti 1998) and Margaret Hedstrom in the United States (Hedstrom 1995). All these authors have concluded to some extent that one can deal effectively with electronic records only if one returns to the first principles of archival theory, including the importance of records as evidence.

Records as evidence provide internal accountability for an agency and make it possible for the agency to determine what it has done in the past. More important, archives—when they contain records that can serve as evidence—can force leaders and institutions to be accountable for their actions. Government archives that contain evidence of the actions of the government can ensure that the rights of individual citizens are protected.² They can also provide evidence of when, where, and why the Navy might build and name a new ship.

Records preserved as evidence may also be interesting because of their informational content. For example, census records retained in an archives because of the evidence they provide about the activity of the Census Bureau, may be of great interest to genealogists. To

² These two themes—the ability of archives to hold public officials accountable and to protect the rights of individual citizens—form the basis of the new mission statement of the National Archives and Records Administration, i.e., “to ensure ready access to essential evidence [and note the emphasis on *evidence*] . . . that documents the rights of American citizens, [and] the actions of federal officials . . .”

many archivists, however, the fact that the Census Bureau creates census returns in the course of conducting its legally mandated business—not the information contained in the record—is of paramount importance.³

At the heart of an archives, therefore, are records that are created by an agency or organization in the course of its business and that serve as evidence of the actions of that agency or organization. The agency or organization maintains those records for its business purposes. At the point when the records are no longer of immediate value to the organization, it may elect to transfer its records to an archives. The archives become responsible for maintaining the evidentiary nature of the materials after the records have left the control of the agency that created them.

One way in which archivists working with analog records have sought to ensure the enduring value of archives as evidence is through the maintenance of an unbroken provenance for the records. Archivists need to be able to assert, often in court, that the records in their custody were actually created by the agency specified. Furthermore, the archivist must be able to assert that the records have been in the custody only of the agency or the archives. In an analog environment, the legal and physical transfer of the documents from the agency to the archives ensures an unbroken chain of custody.

Archives truly exist only when there is an unbroken chain of custody from the creating agency to the archives. For a government archives, the transfer of custody is best accomplished as a matter of law. As Margaret Cross Norton, a pioneer theorist of American archives, noted:

We must disabuse ourselves of the concept that the acquisition by the state historical society of a few historical records . . . automatically transforms the curator of manuscripts into an archivist . . . An archives department is the government agency charged with the duty of planning and supervising the preservation of all those records of the business transactions of its government required by law or other legal implication to be preserved indefinitely (Mitchell 1975).

In a nongovernmental agency, policy can take the place of law if the policy identifies what records of business transactions need to be preserved indefinitely. Either law or policy, however, should govern the transfer of records to an archives.

Why is the authorized transfer of a complete set of records to an archives with an unbroken chain of custody important? First, it helps maintain the evidentiary value of the records. An archivist can be called upon to testify in court about the nature of the records in his or her custody. That archivist would not be expected to testify as to

³ While most archivists would agree with the definition of a record as presented in this paper, there are strong differences about what criteria should be used in the appraisal of records for retention or possible destruction. Some archivists argue that only the evidentiary value of the records should be taken into account, others argue that sociocultural requirements, including the need to establish memory, should be considered (Cook 1997; Cox 1994; Cox 1996).

the accuracy of the contents of the records. However, he or she should be able to assert that on the day when the records left the custody of the originating agency or organization, a particular document was included as part of the records.

Equally important as unbroken custody in establishing the integrity of records is the completeness of the documents. Only records that are complete can ensure accountability and protect personal rights. As soon as records become incomplete, their authority is called into question. For example, when information is missing in a record, we do not know if it is because the information was never created or because it has been discarded. Individual records must be complete; they must contain all the information they had when they were created. They must also maintain their original structure and context.

In addition to each individual record being complete, it is also necessary that the record series in which the record is created be complete. Because records gain meaning from their context, it is important to know the nature of other records. Take the example of a case file. A case file is a record relating to one person as he or she interacts with a government agency. It might be an application for food stamps, an assessment of eligibility for veterans' benefits, or a request for a reproduction of a photograph in an archives. By itself, a case file can tell the user a great deal, but it does not reveal whether the individual in question was treated differently from other people in the same situation. To understand a single record in context, one needs the whole series. There may be references from the case file to other records in the same series. Whenever possible, therefore, archivists seek to preserve entire series.

This does not mean that archivists never throw anything away. The normal archival principle is to save only 2 to 4 percent of an organization's records. What archivists try to avoid, however, is assessing individual records or parts of records. One either keeps the entire record or discards the entire record. Similarly, the normal presumption is that one either keeps or discards an entire series of similar records (though there may be times when the bulk of the records makes this impossible).

Hilary Jenkinson, a leading archival theoretician, neatly summed up the importance of both the legal basis for the transfer of records to an archives and the need for completeness within the record series and the individual records. He noted the importance of authenticity to archives and defined it as the principle that archives are "preserved in official custody . . . and free from suspicion of having been tampered with" (Jenkinson 1965). According to Jenkinson, the archivist's primary task is "to hand on the documents as nearly as possible in the state in which he received them, without adding or taking away, physically or morally, anything: to preserve unviolated, without the possibility of suspicion, every element in them, every quality they possessed when they came to him" (Jenkinson 1984).

Archivists have a responsibility to ensure the integrity of the documents even after they are legally transferred to a repository. In

an analog environment, this is done by a number of mechanisms. Users of archives, for example, normally must work under the supervision of an archival staff member. The users are instructed to maintain the order of records as they are found and are cautioned against adding material to or removing it from the file. In some cases, especially when documents are known to be of great economic value, an archival staff member may count the documents delivered to and then returned by a researcher. (Normally, however, the volume of material in an archives works against any sort of item control.)

The example of the *Constellation* illustrates both the promise and the dangers associated with the evidentiary power of traditional archives. Some of the forged documents that seemingly proved that the ship in the Baltimore harbor had been built in 1797 were found among the records of the U.S. Navy located in the National Archives and Records Administration. Transfer of the records presumably took place under the legal authority of the Federal Records Act, and an unbroken chain of custody had been established. Users of the records, therefore, could assume that any documents found in the record series had been created and maintained by the Navy until they were transferred to the National Archives. The National Archives then maintained the records as they were received from the Navy. The powerful presumption must be that documents found in the Navy files in the Archives are an accurate reflection of the Navy's files at the time of the transfer. Regardless of the content of the records, the organizational context alone would be enough to argue for their authenticity.

We now know that in the case of the *Constellation*, it was wrong to presume that all of the documents in the Navy files, as they were found in archives, were authentic. Archivists had sought to preserve the records in the context of the office that had created them and they had accessioned a complete series into the archives. Normally, this would be enough to ensure the authenticity of the records. In this case, however, it was also necessary to turn away from the context of creation of the record and to examine the individual record itself.

When Wegner, assisted by forensic document examiners at the FBI, examined the problematic documents, he found a number of elements within the documents that led him to question their authenticity. Since most of the documents were copies, it was not possible to test inks and papers. On the basis of the typeface on some of the documents, however, the FBI could determine that the documents had been typed on typewriters that did not come into existence until 30 years after the documents had supposedly been created. Other documents were undated and unsigned, raising questions about their authenticity. In yet another instance, the investigators noticed 14 spelling and typographical errors in a simple document. The investigators knew that the office from which this document supposedly originated had strict requirements for accuracy; the suspect document could not have originated in an office that enforced those requirements.

Without realizing it, the investigators had used one of the oldest archival sciences to test the authenticity of the documents: the sci-

ence of diplomatics. Diplomatics is a body of concepts and methods, originally developed in the seventeenth and eighteenth centuries, “for the purpose of proving the reliability and authenticity of documents.” Over time it has evolved into “a very sophisticated system of ideas about the nature of records, their genesis and composition, their relationships with the actions and persons connected to them, and with their organizational, social, and legal context” (Duranti and Eastwood 1995, quoted in Duranti and MacNeil 1996). Perhaps because diplomatics emerged from the need to understand and authenticate medieval charters, patents, and other legal documents, American archivists knew little about the field until quite recently. In addition, the primary problem facing American archivists for most of this century has not been to understand individual documents but rather to deal with the flood of documents on paper and in other formats generated by a bureaucratic, paper-intensive society.

Fortunately, in 1989 an Italian archivist teaching in Canada introduced North American archivists to the primary concepts of diplomatics through a series of six articles published in the Canadian journal *Archivaria* (Duranti 1998). In these articles and in her later work on reliability and integrity, Duranti expands on the interrelationship between the form, structure, and authorship of documents. The form of a record and the procedure for its creation, she asserts, determine the reliability of the record. A record is more likely to be reliable when its form is complete than when it is incomplete. While documents can require many elements, the two most commonly required elements of form are the date and an element, usually a signature, that assigns responsibility to a person for the content of the record (Duranti 1995).

Diplomatics also provides a mechanism for evaluating the authenticity of copies. Why is an original more reliable as evidence than a copy? It is because the original has the maximum degree of completeness and a higher degree of control in the procedure of creation of the document. Creating a copy always introduces the possibility for variation or change from the original.

On the other hand, there are times when a copy may be more reliable than an original. For example, a contract for the sale of the house that is copied into the deed books of a village government may be more reliable than the original, because a third, impartial, authority can attest to the agreement of the parties represented in the contract. Archives have a long tradition of producing authentic copies, i.e., copies that have not been subject to manipulation, substitution, or falsification after the completion of the process that created the original record. Such copies often entail a change in format (for example, from paper to microfilm) and require that procedures be in place to ensure the authenticity of the resultant copies. If the latter condition is met, archivists willingly discard the originals.

An archivist could use the principles of diplomatics to judge the reliability and the authenticity of the individual documents in the *Constellation* case. For example, questioned documents that lacked a date or a signature would fail the fundamental test for reliability. The

document filled with misspellings and typographical errors would also fail. The form of a document that does not follow the documentary conventions of the creating office is suspect; the document itself may be unreliable.

In summary, traditional archival theory has developed two approaches for ensuring the authenticity of the document. The first approach, the basis for most American archives, seeks to understand and control the context in which records are created. Records that are generated in an agency, transferred by law or policy to an archival agency through an unbroken change of custody, and maintained complete and inviolate by that archival agency are presumed to be authentic. The second approach, as exemplified in the works of Durrant, focuses on the individual record: its form and the circumstances of its creation. Together, these two approaches are used to ensure the authenticity of records in the analog world.

Archival Authenticity in a Digital World

The archival profession has established a theoretical base to justify the assertion of authenticity when dealing with analog records. But will the principles that have worked so well in the analog environment transfer to the new digital world? Wendy Duff has noted, "As records migrate from a stable paper reality to an intangible electronic existence, their physical attributes, vital for establishing the authenticity and reliability of the evidence they contain, are threatened" (Duff 1996). The ease with which records in electronic form can be created, transferred, and modified only heightens the importance of maintaining their integrity. The central question facing all archivists, therefore, is how to ensure the authenticity of records in digital form. Can the traditional archival methodologies developed for analog records be used for digital records? Or must new methodologies and techniques be developed to ensure that the archival records remain authentic over time?

A number of important initiatives are under way to explore how the integrity of records can be preserved in a digital environment. None of the strategies has yet become widely accepted, primarily because they have not been tested in the field. As Philip Bantin has concluded, "In short, there are no clear-cut answers available yet, but there are plenty of very good ideas and emerging strategies out there" (Bantin 1999). Two of the more promising approaches can be summarized here.

The University of Pittsburgh Functional Requirements for Evidence in Recordkeeping Project

The University of Pittsburgh conducted one of the first and most extensive research projects that sought to identify the functional requirements for the preservation of electronic evidence. Its project, the "Functional Requirements for Evidence in Recordkeeping," consisted of three main components. First, the project identified the func-

tional requirements for recordkeeping in a variety of communities. The project recognized that groups other than archivists (e.g., the legal, medical, and business communities) also had need for authentic, reliable records. Laws, standards, customs, and the best practices of each community contain the justifications for record keeping. To ensure that electronic records meet the needs of those communities (i.e., that they become what the project identified as “business acceptable communications”), one must identify the requirements for recordkeeping in each community and then establish metadata that meet those requirements. The project did this by establishing the recordkeeping requirements and practices of organizations—the “literary warrant” (Duff 1996; Bearman 1996).

Using the requirements necessary for literary warrant, the project then produced a general specification of the attributes of evidentiality. The specification consists of 13 properties that are categorized into three groups. The first group requires a *conscientious organization* that complies with legal and administrative requirements for recordkeeping. The second group specifies the requirements for *accountable recordkeeping systems*, including policies, assigned responsibility, and formal methodologies for their management and accurate and complete documentation. The Pittsburgh system presupposes that accountable recordkeeping systems are used at all times in the normal course of business. The third group defines the requirements that relate to the *record itself*, specifically how the record is created or captured, how it is maintained, and what is necessary for the record to be used.

In addition to developing the general specification of the requirements for evidentiality, the Pittsburgh project developed a set of production rules to express formally each functional requirement. David Bearman, a consultant on the project, has turned the production rules and general analysis into a set of metadata requirements. The goal is to be able to create records that are encapsulated metadata objects: content in an envelope of metadata that ensures the authenticity, integrity, reliability, and usability of the content.

Implicit in the Pittsburgh approach is the assumption that “recordness” and “evidentiality” (the elements that determine the trustworthiness of records in business and legal settings) can be maintained in an electronic system only if the requisite functionality is built into the record system from the start. Several efforts have been made to implement the Pittsburgh model, most notably in projects under way at Indiana University, a Swedish pharmaceutical company, and the City of Philadelphia, but there is no consensus whether the Pittsburgh project has identified the true functional requirements for authenticity. Some worry that the Pittsburgh model may be too complex, and hence too costly, to implement. Furthermore, it presupposes radical changes in how documents are generated. For example, if one wishes to write a report, one currently opens a word processing package and begins writing. The Pittsburgh system seems to propose that in the future one would open instead a report-writing module. The module would “know” who you are, what your author-

ity for writing the report is, and in what format you are writing the report. The software would automatically encapsulate each draft of the report with this management information. While highly desirable or even mandatory, to ensure the authenticity of the electronic file, such an approach does not reflect how people currently use software.

University of British Columbia Preservation of the Integrity of Electronic Records and InterPARES Projects

Two projects at the University of British Columbia (UBC) are investigating the integrity of digital information over time. The first project, "Preservation of the Integrity of Electronic Records," sought to identify the best methods for preserving the reliability and authenticity of electronic records over time. The UBC analysis determined that generic information systems designed to collect, process, store, and disseminate information lack some of the functionality needed to produce, maintain, and preserve reliable electronic records. For example, most current systems do not adequately relate the content of records to business transactions. They also lack sufficient metadata to monitor the creation and maintenance of records in a way that ensures they will be both reliable and understandable when retrieved in the future. The project concluded that reliability and authenticity of electronic records are best ensured when procedural rules for record-keeping are embedded into the overall records system. This finding is similar to that of the Pittsburgh project, which expressed an interest in building into systems the automatic capture of the metadata it has determined are needed to ensure the recordness of the data (Duranti and MacNeil 1996; Hedstrom 1996).

In other ways, however, the UBC project was fundamentally different from the Pittsburgh project (Duranti and MacNeil 1996; Bantin 1999; Marsden 1997). For example, the analysis of the requirements for recordkeeping in the two projects differed greatly. The Pittsburgh project based its analysis on literary warrant, whereas the UBC project's analysis was based on diplomatics and archival theory.

In part because of the difference in starting points, the two projects reached fundamentally different conclusions in some areas. One of the most striking differences relates to the role of the archives in ensuring authenticity. The Pittsburgh project did not assume that an archives is needed to ensure the preservation and authentication of records. In the Pittsburgh system, it is the metadata, not the custodial agency, that determine the authenticity of records. Records can, and in most cases should, remain in the custody of the agency that created them. As one of the Pittsburgh project members has argued, "Archivists cannot afford—politically, professionally, economically, or culturally—to acquire records except as a last resort . . . Indeed, the evidence indicates that acquisition of records and the maintenance of the archives as a repository gets in the way of achieving archival objectives and that this dysfunction will increase dramatically with the spread of electronic communications" (Bearman 1991). The

UBC project, in contrast, placed archives at the heart of the authentication system for electronic records, in a fashion similar to the role played by archives in protecting and authenticating paper records. This project concluded that “the routine transfer of records to a neutral third party, that is, to a competent archival body, invested with the exclusive authority and capacity for the indefinite preservation of inactive records, is an essential requirement for ensuring their authenticity over time” (Duranti and MacNeil 1996).

The “Preservation of the Integrity of Electronic Records” project at UBC sought to establish a theoretical framework based in traditional archival principles for the authentication of digital information. A follow-on project is now seeking to put some of these principles into action. The InterPARES (for “**I**nternational **R**esearch on **P**ermanent **A**uthentic **R**ecords in **E**lectronic **S**ystems”) project is an international collaboration spearheaded by UBC. Its goal is to use the tools of archival science and diplomatics to develop the theoretical and methodological knowledge essential to the permanent preservation of inactive electronically generated records. It will then formulate model strategies, policies, and standards capable of ensuring the preservation of those records. The InterPARES project has generated great interest in the archival community, in part because it is based on familiar principles and practices. The community eagerly awaits reports of its findings.

Conclusion

It is not possible at this early stage to say whether Pittsburgh or UBC has the better approach for ensuring the authenticity of records. Both approaches need to be tested in the field (Bantin 1999). As Margaret Hedstrom has noted, “What we lack is an evaluation of the usefulness of these findings from the perspective of organizations that are responsible in some way for preserving and providing access to electronic records. We need assessments from the administrators of archival and records management programs about the feasibility of putting the proposed policies, and models into practice. We need reactions from people outside the archival community especially where related research and projects are being conducted” (Hedstrom 1996).

In the interim, however, it is easy to speculate that some combination of the Pittsburgh and UBC approaches will come to dominate. The Pittsburgh project’s basis in the actual documentary requirements of different communities is very appealing, and the project’s desire to include administrative metadata from the very moment of creation is highly desirable.

On the other hand, it is unlikely that all information of interest to future users of records systems will be found in records creation management systems fully compliant with the Pittsburgh metadata. Scholars will be willing to access, use, and evaluate the information found in the electronic files, regardless of whether the actual data convey the true quality of “recordness.” An archival purist might in-

sist that if information is not stored in a record keeping system, then the information cannot be a record and therefore should not be part of the archival record. In reality, however, our repositories are filled with interesting information that may not meet the formal definition of "record" or may not have been created with a record keeping system in mind.

A good example of how material that is not formally a record can be valuable to the researcher is the famed PROFS case (Bearman 1993). PROFS refers to a proprietary IBM communication system used in the White House under Presidents Ford and Reagan. Because they were system back-up tapes, the PROFS tapes lacked even the rudiments of record keeping functionality. Nevertheless, a consortium of historical groups sued for the release of the tapes. In the absence of controlled records, the information on the back-up tapes was the best the researchers could find. For researchers, the value of the tapes was great because they were still held by the agency and were surprisingly complete. However, even if only selections of the e-mail messages had survived and were located only in nongovernmental repositories, researchers would still try to use them, even though their authenticity was more questionable.

In short, social mechanisms of control promise to be the fundamental basis for the establishment of digital authenticity. It would be desirable if all digital information consisted of true records created in a system that encapsulates with the record the information needed to maintain the evidential value of the records. For most digital information, however, the fact that it is in an archives, an unbiased third party, will have to suffice. As with the paper records used in the *Constellation* example, the fact that digital information is found within a trusted repository may become the base upon which all further assessments of authenticity build.

Even if the physical presence of digital data in a trusted repository is the basis for future assessments of authenticity, archivists will still need to associate with those digital documents metadata that researchers can use to understand and assess digital information. We need self-conscious documentation by the creators and preservers of digital representations that details the methods employed in making and maintaining the representations. We also need to know what researchers need to know about the transformation from analog to digital format, as well as about any transformations that may occur as digital data are preserved. To determine the latter, we need to understand the "digital literacy" that future researchers will need "to assess digital information, identify known artifacts introduced by particular processes, and correctly identify as yet unknown sources of distortion" (Bearman and Trant 1998). Only by understanding the interactions between researcher and document and records and repositories will we be able to convey into the future the trust mechanisms of the paper world.

REFERENCES

- Acland, Glenda. 1992. Managing the Record Rather than the Relic. *Archives and Manuscripts* 20(1):57-63.
- Bantin, Philip. 1999. Strategies for Managing Electronic Records: A New Archival Paradigm? An Affirmation of our Archival Traditions? *Archival Issues*.
- Bearman, David. 1991. An Indefensible Bastion: Archives Repositories in the Electronic Age. In *Archival Management of Electronic Records*, edited by David Bearman. Archives and Museum Informatics Technical Report #13. Pittsburgh: Archives and Museum Informatics.
- Bearman, David. 1993. The Implications of *Armstrong v. the Executive Office of the President* for the Archival Management of Electronic Records. *American Archivist* 56(4):674-89.
- Bearman, David. 1994. *Electronic Evidence: Strategies for Managing Records in Contemporary Organizations*. Pittsburgh: Archives and Museum Informatics.
- Bearman, David. 1996. Virtual Archives. Paper presented at the International Congress of Archives meeting in Beijing, China, September 1996. Available from <http://www.lis.pitt.edu/~nhprc/prog6.html>.
- Bearman, David and Jennifer Trant. 1998. Authenticity of Digital Resources: Towards a Statement of Requirements in the Research Process. *D-Lib Magazine*. Available from www.dlib.org/dlib/june98/06/bearman.html.
- Bellardo, Lewis J., and Lynn Lady Bellardo, compilers. 1992. *A Glossary for Archivists, Manuscript Curators, and Records Managers*. Chicago: Society of American Archivists.
- Buckland, Michael. 1997. What is a "Document"? *Journal of the American Society for Information Science* 48(9):804-9.
- Cook, Terry. 1997. What Is Past Is Prologue: A History of Archival Ideas Since 1898, and the Future Paradigm Shift. *Archivaria* 43:17-63.
- Cox, Richard. 1994. The Record: Is It Evolving? *Records & Retrieval Report* 10(3):1-16.
- Cox, Richard. 1996. The Record in the Information Age: A Progress Report on Research. *Records & Retrieval Report* 12(1):1-16.
- Duff, Wendy. 1996. Ensuring the Preservation of Reliable Evidence: A Research Project Funded By the NHPRC. *Archivaria* 42:28-45.

- Duranti, Luciana. 1995. Reliability and Authenticity: The Concepts and Their Implications. *Archivaria* 39:5-10.
- Duranti, Luciana. 1998. *Diplomatics: New Uses for an Old Science*. Lanham, Md.: Scarecrow Press.
- Duranti, Luciana, and Terry Eastwood. 1995. Protecting Electronic Evidence: A Progress Report. *Archivi & Computer* 5(3):213-50.
- Duranti, Luciana, and Heather MacNeil. 1996. The Protection of the Integrity of Electronic Records: An Overview of the UBC-MAS Research Project. *Archivaria* 42:46-67.
- Hedstrom, Margaret. 1996. Electronic Records Research Issues: A Summary of Recent Research. Prepared for the Invitational Conference on Electronic Records, June 28 and 29, 1996. Available from <http://www.si.umich.edu/e-recs/Research/NHPRCSum.html>.
- Hedstrom, Margaret. 1995. Electronic Archives: Integrity and Access in the Network Environment. *American Archivist* 58(3):312-324.
- Jenkinson, Hilary. 1965. *A Manual of Archive Administration*. London: Percy Lund, Humphries & Co. Ltd.
- Jenkinson, Hilary. 1984. Reflections of an Archivist. In *A Modern Archives Reader*, edited by Maygene F. Daniels and Timothy Walch. Washington, D.C.: National Archives and Records Service.
- LeDuc, Daniel. 3 July 1999. Shored Up and Shipshape: Refurbished USS *Constellation* Returns to Inner Harbor, Minus a Myth. *The Washington Post*.
- Maher, William. Society and Archives. 1997. Incoming presidential address to the Society of American Archivists. Available from <http://www.archivists.org/governance/presidential/maher-1.html>.
- Marsden, Paul. 1997. When is the Future? Comparative Notes on the Electronic Record-Keeping Projects of the University of Pittsburgh and the University of British Columbia. *Archivaria* 43:158-73.
- McKemmish, Sue, and Frank Upward, eds. 1993. *Archival Documents: Providing Accountability through Recordkeeping*. Melbourne: Ancora Press.
- Mitchell, Thornton W. 1975. *Norton on Archives: The Writings of Margaret Cross Norton on Archival and Records Management*. Carbondale, Ill.: Southern Illinois University Press.

Muller, S., J. A. Feith, and R. Fruin. 1968. *Manual for the Arrangement and Description of Archives*, translated by Arthur H. Leavitt. New York: H.W. Wilson.

Wegner, Dana M. 1991. *Fouled Anchors: The Constellation Question Answered*. Bethesda, Md.: David Taylor Research Center.

Welch, Matt. 15 October 1998. The Cookie Monster of Putnam Pit. *Salon Magazine*. Available from <http://www.salon.com/21st/feature/1998/10/15feature.html>.

Where's Waldo? Reflections on Copies and Authenticity in a Digital Environment

by David M. Levy

Introduction

You have probably seen the “Where’s Waldo?” children’s books. Each double-page spread contains drawings of hundreds of cartoon figures. Your job is to find Waldo, a character who is always dressed in a red-and-white striped woolen cap and shirt and is wearing glasses. Often there are characters who look a lot like him, but if you look closely you can see that some detail or other is wrong (e.g., it is a woman, the cap is solid red). In other words, only one of the figures on the page is the *real* Waldo; the rest are impostors, look-alikes, or close matches. “Pay attention,” these drawings seem to say, “Appearances can be deceiving.”

Waldo presents the problem of authenticity in graphical form. Although a number of the cartoon figures *seem* to be Waldo, only one *is* the authentic Waldo. Being authentic in this case means being who or what you seem or claim to be. In Waldo’s case, there can only be one right answer, since we are talking about a unique individual. But in other cases, there may be more than one right answer. This happens when we are concerned with, say, group membership (being a medical doctor) or with types (being a 1956 Chevy). It is only because we live in a world of multiplicity—where several people or things may appear to be the same—that duplicity is possible. Judgments of authenticity, as I understand it, allow us to navigate through a world by distinguishing genuine multiplicity from duplicity.

In the realm of written forms—in the world of paper and other tangible media—we have, over the centuries, developed elaborate procedures for identifying authentic documents and for ferreting out impostors. In the digital realm, we have barely begun to do this, and there are many technical and social challenges to be met. One challenge comes from the fact that the digital realm produces copies on an unprecedented scale. It is a realm in which, as far as I can tell,

there are no originals (only copies—lots and lots of them) and no enduring objects (at least not yet). This makes assessing authenticity a challenge.

What Are Documents?

I use the word *document* where others might use *text*, *record*, *information-bearing artifact*, or *written form*. “Document” is a cover term for a large group of artifacts, including textual materials, whether handwritten or mechanically realized; graphics and photographs; and audiovisual presentations. But by what criterion do all these things fit into a single, coherent category?

I have come to understand documents by analogy with human beings. Documents are surrogates for people. They are bits of the material world (stone, clay, wood pulp, and now silicon) that we create to speak for us and take on jobs for us. A receipt bears witness to and thereby validates a financial transaction; a restaurant menu speaks for the establishment, the restaurant; a novel tells a story; a political flyer speaks for a candidate or political organization; and so on.¹ By saying that documents “speak,” I do not mean to limit them to textual or verbal materials. Pictures, drawings, diagrams, moving images, and other conventional forms of communication also speak in the metaphorical way in which I am using the term: they communicate, they tell us things about the world. And when I say documents “take on jobs,” I am referring to the way we tailor their form and content to particular tasks and contexts. Genre (whether a receipt, a menu, a novel or a flyer) is, in effect, the clothing of conventional content to do particular tasks in the world (to witness a financial transaction, recite the dishes available and their prices, etc.) (Levy 1999).

For a document, speaking per se is not enough. It also must be able to speak *reliably*. We depend on documents to carry messages through space and time. In many cases, this reliability is achieved through fixity: letterforms inked on paper can survive for long periods of time. But with newer media, such as video, this reliability is achieved not by fixity but by repeatability. The moving images on a video screen are by their very nature transient. I will never be able to see those very images again. But I can play the tape repeatedly, each time seeing a performance that, for all practical purposes, is “the same as” the one I saw the first time.

If documents are meant to be reliable surrogates for human beings, then it makes perfect sense that we would be critically concerned with their authenticity. Steven Shapin (1994), a sociologist, argues that human social order—that human life itself—is fundamentally based on trust, i.e., on our ability to rely on one another.

¹ There are great complexities and ambiguities regarding who is speaking in or through a document. In literature, for example, distinctions have been made between the narrator, the implied author, the “real” author, etc. Such complexities and ambiguities also exist, however, when a human being is speaking.

“How could coordinated activity of any kind be possible if people could not rely upon others’ undertakings? No goods would be handed over without payment, and no payment without goods in hand. There would be no point in keeping engagements, nor any reason to make engagements with people who could not be expected to honor their commitments,” he writes. Much as we rely on one another, we also have come to rely on documents in the making and maintaining of a shared, stable, social order. So it is no accident that words such as *trust*, *reliability*, and *truthfulness*, which are fundamentally social, would apply to documents as much as to people. It is likewise no accident that documents, as surrogates for us, would be accountable in the same terms.

What Is a Copy?

I worked for Xerox for a number of years, so it should hardly be surprising if some of my thinking and my examples come from the world of photocopying. In that world, “to make a copy” means to put one or more pieces of paper on the photocopier platen or in the RDH (recirculating document handler) and push the Big Green Button. What comes out at the other end of the machine is a “copy.” In this context, a copy is something that is the result of a *process* of copying. It says nothing about whether the result is a good copy or a bad copy, or whether or not it is useful.

But there is a second notion of copy, which has more to do with the *product* than the process. To be a copy in this sense is to stand in a certain relation to an original, that is, to its origin. To be a copy in this sense is to be faithful to the original. The definition of “faithful,” however, depends on the circumstances in which the copy is being made and on the uses to which it will be put. The context of use, in other words, determines which properties of the original must be preserved in the copy. Does it matter that I have just made a photocopy of a signed will? It depends on what I intend to do with it. If it is for informational purposes (to show you what my will says), then it is an adequate copy; for some legal purposes, however, it won’t do.

The point is, a document can be *identical* only with itself, if “identical” is taken to mean “the same in every respect.” When we say that something is “the same,” we generally mean one of two things. We either mean that it is “the very same” thing (as in “This is the same car I drove yesterday”) or that it is “of the same type” as something else (“I read that same book last year”). It is this second notion of sameness—sameness of type, sameness in virtue of sharing certain properties—that is at issue in copying (Levy 1992).

Even an extremely high-fidelity copy will be different from the original in innumerable ways, because to copy is to *transform*. The copy will be on a different piece of paper that has its own unique properties. The process of photocopying will make letterforms thicker or thinner than those on the original, and will make images lighter or darker; it will add noise or remove it; it will change tones, shapes, aspect ratios, and so on. Differences will always be introduced in

copying; the trick is to regulate the process sufficiently so that the resulting differences are of little or no consequence and that the properties of greatest consequence are shared. Determinations of which properties matter are made in the context of purpose and use.

Copying Without an Original

I have presented a simple and straightforward notion of copying. Although I have used the photocopier to illustrate how it works, this notion is not dependent on any particular technology. Making a copy by hand embodies the same idea. Moreover, although I have talked about making a single copy, one can obviously make multiple copies of an original—an indefinite number, in fact. It is common for someone to create a “master” document and to produce any number of copies from it. What is crucial in this scheme is that there is an *original* from which the copies are made.

But there is another scheme—one that does not require an original. It is a manufacturing technique, a means of producing a large number of artifacts from a single source. If you want to make coins, for example, you can create a mold and pour molten metal into it to cast the coins. This is also the way the printing press works. You create a set of printing plates that are used to produce inked pieces of paper.

The reason I say there is no “original” in this technique is that the *source*² from which the copies are made (the mold or the printing plate) is a very different kind of thing than the copies.³ You cannot spend the mold (although you may be able to mint more coins); you would not normally choose to read the text on the printing plate. This means that the word *copy* is being used in a somewhat different sense. It perhaps harks back to the root meaning of the word (copious, plentiful). But there is another sense in which the artifacts produced in this way are copies: They are copies *of one another*. Indeed, to a large extent, the purpose of this technique is to manufacture a set of “identical” artifacts—artifacts that are all “the same,” that is, of the same type. These artifacts are identical in the sense that they are interchangeable with one another for certain purposes.

The examples I have given so far involve the production of enduring physical artifacts, or things. But this method of copying from a source also works for producing activities or events, which by their very nature are transient. Consider the case of a play, where a script (the source) serves as the basis for a number of performances (the copies) or an audio or videotape (the source), which leads to the realization of sounds or visual images, or both.

² I will use the word *source* to designate the thing from which copies are made in this method, and the word *original* when I mean something that is of the same kind as the copies.

³ I do not mean to suggest that there can *never* be an original that is used to guide the making of the source. I may print an edition of *Leaves of Grass*, taking the text from the 1891 edition. In this case, some actual printed copy of the 1891 edition is my original. Nevertheless, the production of my new edition is mediated by the printing plates I have created, and these plates are *not* an original.

In none of these cases, however, is the source ever enough. Manufacturing the intended artifacts also requires a complex of skills, know-how, and, often, technical equipment. The mold for coins is useless without the right metals and the skill to do casting; a printing plate is useless without a printing press and knowledge of how to use it; the script needs a cast of actors; and the videotape needs a video player. In each case, the quality of the product or the performance depends on a skillful and properly executed process of production. The source, in other words, does not and cannot fully specify the properties of the things it is used to make. There is a division of responsibility between the source and the environment in which it operates.

It is worth comparing print with analog audio or video recording before talking about the digital case. In the case of printing, the source is used to produce a definite number of copies, an edition. Each copy in an edition is a stable physical object whose existence is independent of the source. But in the case of the recording, when the tape is defined as the source, there is no notion of a definite number of copies (e.g., replayed performances); rather, once you have the tape and an appropriate player, you can produce a (relatively) unlimited number of copies, or performances. Moreover, unlike the products of print, the copies are completely dependent on the source for their existence. Should the tape be damaged or lost, there will be no more performances. This gives the source a greater importance in the case of recordings. You *have* to preserve it if you want copies in the future. (And, of course, you have to preserve the player, which is the means of making copies from the source.) In the case of printing, by contrast, once the source has done its work, it is no longer needed. (Indeed, the advantage of movable type is that it can be reused, i.e., the elements of the source can be recycled.)

Digital Documents

Like printed documents and recorded audio and video performances, digital documents are founded on a distinction between a source and the copies produced from it. The source is a digital representation of some kind, a collection of bits. The copies are the sensible impressions or manifestations—text, graphics, sound, whatever—that appear on paper, on the screen, and in the airwaves. Getting from the source to the copy requires a complex combination of technical and social environment, including an elaborate configuration of hardware and software.

In one sense, digital technologies are very much modeled on the printing press. They allow users to create what amount to digital printing plates from which they can “print” an arbitrary number of copies. The relation with traditional print is particularly strong when the copies produced are textual and graphical in nature, as is so much of the material on the Web today. But digital documents, even those with textual content, share significant features with analog audio and video recordings as well. With audio and video, we tend to

think of the source (in this case, the audio or videotape) as more permanent than the copies produced from it (the performances), which are inherently transient. Currently, we seem to be importing this same hierarchy of permanence into the digital domain. We think of the digital source (such as a Microsoft Word file) as more permanent than the text and images that appear on the screen. This makes sense, because we know how to “save” the file. When we have done so, it will typically survive on a hard drive or a floppy despite power loss, whereas the screen image cannot. But as we adopt this way of thinking, we are also coming to treat paper copies (analogous to screen images) as more transient than the source file. We often print out a paper copy to read and then toss it away, confident that we will be able to print out another as long as we have the file. But the truth is, at least for the moment, that paper has a better chance of survival than a digital source.

Indeed, digital entities are generally less stable than their counterparts on paper and other tangible media, and digital production tends to yield much greater variability of product than analog production does. In the case of print, once we have the plate and a press, the amount of variability is limited. Even more so is this the case with an analog recording: once we have the tape and an appropriate player, the amount of variability in performances is typically fairly well constrained. The differences generally are limited to minor variations in quality. For digital copies, however, there is likely to be a much greater range of variability. Some of the variability is intentional and it is a great strength of the technology. We can easily edit digital documents and quickly produce variants. Some variability is unintended and is an unresolved problem: digital copies are extremely sensitive to the technical environment, to the point that features we would like to preserve in subsequent copies may be hard (or impossible) to maintain. Displaying the file on a different computer may lead to font substitutions, different line breaks, and so on. These same sorts of variability may even occur on the same computer if, in the interim, the environment has changed in some crucial way.⁴ Consequently, two different viewings of the “same” source may differ in important ways—they may not be “the same.”

Under such circumstances of radical variability, there does not appear to be anything like a stable document or object. Over time, the digital source may move from server to server. The version that ends up on your local computer may have been copied from a server and will likely have undergone further transformation; for example, your local browser or editor may generate other local, and possibly partial, digital sources in the process of creating something you can actually see. What you do see at any given moment will be the product both of the local digital source and of the complex technical environment (hardware and software), which is itself changing in complex and unpredictable ways. The digital source, the perceptible

⁴ As sound and motion are digitally recorded, issues of uncontrolled variability will increasingly arise here, too.

copies, and the environment are all undergoing change in ways that no one yet knows how to control.

Authenticity in a Digital Environment

Assessments of authenticity in the world of paper and other stable, physical media rely heavily on the existence of enduring physical objects. If you want to determine whether the document in front of you is the unique individual it purports to be (someone's last will and testament, for example), you can try to determine its history. But you can do this only because it *has* a history, an extended existence in time. If you want to determine the authenticity of something that is one of many (a member of an edition, for example) you can compare it with another copy, a reference copy. And even where the thing in question is transient (such as the performance of a play), you still may be able to make use of a stable reference object (such as the script). In all these cases, either the object in question or a reference object has an enduring, physical existence that helps ground the determination of authenticity.⁵

What happens in the digital case if there are no stable, enduring digital objects? One possibility is that we will find a way to *create* them. In one current view, objects are at least in part socially constructed; they are bounded and stabilized through social interaction (Smith 1996). Literary works (e.g., *Hamlet*) are a clear example of this. Although we cannot really say what works are, we have nonetheless created a cultural mechanism (copyright and the courts) to help us decide where the boundaries between works lie. Here there can be no question of ultimate, natural answers—only social answers based on law and politics. In the digital domain, I see Jeff Rothenberg's proposal (in this collection) to stabilize digital environments through emulation as one attempt to create stable digital objects. (I am not sure it is a workable solution, but that is another matter.)

Without the security of stable digital objects, what might we do? One possibility would be to maintain audit trails, indicating the series of transformations that has brought a particular document to the desktop. Such a trail (akin to an object's provenance) could conceivably lead back to the creation of the initial document or, at least, back to a version that we had independent reasons to trust as authentic. Having such an audit trail (and trusting it) would allow us to decide whether any of the transformations performed had violated the document's claimed authenticity. A second possibility would ignore the history of transformations and would instead specify what properties the document in question would have to have to be authentic. This would be akin to using a script or a score to ascertain the authenticity of a performance.

⁵ How do we know whether to trust the authenticity of reference objects? The whole process recurses. I agree with Clifford Lynch, who suggested in his presentation at this workshop that the process is ultimately grounded in our trust of others. The "buck stops" when we accept someone's (or some institution's) claim that some object in the chain of reasoning is authentic.

Conclusion

I have no conclusion other than this: Understanding what we *want* to accomplish, and what we *can* accomplish, with regard to authenticity in the digital realm will take considerable effort. If nothing else, this workshop has convinced me of the cultural importance, as well as the difficulty, of the work that lies ahead.

REFERENCES

Levy, D. M. 1992. What Do You See and What Do You Get? Document Identity and Electronic Media. In *Screening Words: User Interfaces for Text; Proceedings of the Eighth Annual Conference of the UW Centre for the New OED and Text Research*. Waterloo, Ontario: University of Waterloo Centre for the New Oxford English Dictionary and Text Research.

Levy, D. M. 1999. The Universe Is Expanding: Reflections on the Social (and Cosmic) Significance of Documents in a Digital Age. *Bulletin of the American Society for Information Science* 25(4):17-20.

Shapin, S. 1994. *A Social History of Truth: Civility and Science in Seventeenth-Century England*. Chicago: The University of Chicago Press.

Smith, B. C. 1996. *On the Origin of Objects*. Boston: MIT Press.

Authenticity and Integrity in the Digital Environment: An Exploratory Analysis of the Central Role of Trust

by Clifford Lynch

Introduction

This paper seeks to illuminate several issues surrounding the ideas of authenticity, integrity, and provenance in the networked information environment. Its perspective is pragmatic and computational, rather than philosophical. Authenticity and integrity are in fact deep and controversial philosophical ideas that are linked in complex ways to our conceptual views of documents and artifacts and their legal, social, cultural, and historical contexts and roles. (See Bearman and Trant [1998] for an excellent introduction to these issues.)

In the digital environment, as Larry Lessig (1999) has recently emphasized, computer code is operationalizing and codifying ideas and principles that, historically, have been fuzzy or subjective, or that have been based on situational legal or social constructs. Authenticity and integrity are two of the key arenas where computational technology connects with philosophy and social constructs. One goal of this paper is to help distinguish between what can be done in code and what must be left for human and social judgment in areas related to authenticity and integrity.

This paper has been modestly revised based on discussion at the workshop and a reading of the other papers presented there. All of the papers, but particularly those of David Levy and Peter Hirtle, raise important issues that are relevant to the topic of this article. From Hirtle's paper, I had the opportunity to learn something of the science of diplomatics, and at the workshop, I had the opportunity to learn much more from Luciana Duranti. Her book, *Diplomatics: New Uses for an Old Science* (1998), offers valuable and fresh insights on the topics discussed here. These other works provide important additional viewpoints that are not fully integrated into this paper and I urge the reader to explore them. My thanks also to the participants in the Buckland/Lynch Friday Seminar at the School of Information Management and Systems at the University of California, Berkeley, for their comments on an earlier version of this paper.

Gustavus Simmons wrote a paper in the 1980s with the memorable title “Secure Communications in the Presence of Pervasive Deceit.” The contents of the paper are not relevant here, but the phrase “pervasive deceit” has stuck in my mind because I believe it perfectly captures the concerns and fears that many people are voicing about information on the Internet. There seems to be a sense that digital information needs to be held to a higher standard for authenticity and integrity than has printed information. In other words, many people feel that in an environment characterized by pervasive deceit, it will be necessary to provide verifiable proof for claims related to authorship and integrity that would usually be taken at face value in the physical world. For example, although forgeries are always a concern in the art world, one seldom hears concerns about (apparently) mass-produced physical goods—books, journal issues, audio CDs—being undetected and undetectable fakes.¹

This distrust of the immaterial world of digital information has forced us to closely and rigorously examine definitions of authenticity and integrity—definitions that we have historically been rather glib about—using the requirements for verifiable proofs as a benchmark. As this paper will demonstrate, authenticity and integrity, when held to this standard, are elusive properties. It is much easier to devise abstract definitions than testable ones. When we try to define integrity and authenticity with precision and rigor, the definitions recurse into a wilderness of mirrors, of questions about trust and identity in the networked information world.

While there is widespread distrust of the digital environment, there also seems to be considerable faith and optimism about the potential for information technology to address concerns about authenticity and integrity. Those unfamiliar with the details of cryptographic technology assume the magical arsenal of this technology has solved the problems of certifying authorship and integrity. Moreover, there seems to be an assumption that the solutions are not deployed yet because of some perverse reluctance to implement the necessary tools and infrastructure.² This paper will take a critical view of these

¹ Confusingly, however, we have the appearance of perfect forgeries (at least in terms of content; the packaging is often substandard) of digital goods in the form of pirate audio CDs, DVDs, and software CD-ROMs. In these cases, the purpose is not usually intellectual fraud so much as commercial fraud through piracy. One might argue that these copies have integrity (they are, after all, bitwise equivalent); however, their authenticity is dubious, or at least needs to be proved by comparison with copies that have a provenance that can be documented. Another case that bears consideration and helps refine our thinking is the bootleg or “gray-market” recording—perhaps an audio CD of a live performance of a well-known band, released without the authorization of the performers and not on their usual record label. This does not stop the recording from being authentic and accurate, albeit unauthorized. The performers may or may not be willing to vouch for the authenticity of the recording; alternatively, one may have to rely on the evidence of the content (i.e., nobody else sounds like that) and, possibly, metadata provided by a third party that potentially has its own provenance.

² It would be useful to better understand why there has not been a greater effort to deploy these capabilities, even though they have substantial limitations. Contributing factors undoubtedly include export controls and other government regulations on cryptography, both in the United States and elsewhere; legal and

cryptographic technologies. It will try to distinguish between the problems that cryptographic technologies can and cannot solve and how they relate to the development of infrastructure services. There seems to have been surprisingly little examination of these questions; this is itself surprising.

Before attempting to define integrity or authenticity, it is worth trying to gain an intuitive sense of how the digital environment differs from the physical world of information-bearing artifacts (“meatspace,” as some now call it). The archetypal situation is this: We have an object and a collection of assertions about it. The assertions may be internal, as in a claim of authorship or date and place of publication on the title page of a book, or external, represented in metadata that accompany the object, perhaps provided by third parties. We want to ask questions about the integrity of the object: Has the object been changed since its creation, and, if so, has this altered the fundamental essence of the object? (This can include asking these questions about accompanying assertions, either embedded in the object or embodied in accompanying metadata). Further, we want to ask questions about the authenticity of the object: If its integrity is intact, are the assertions that cluster around the object (including those embedded within it, if any) true or false?

How do we begin to answer these questions in meatspace? There are only a few fundamental approaches.

- We examine the provenance of the object (for example, the documentation of the chain of custody) and the extent to which we trust and believe this documentation as well as the extent to which we trust the custodians themselves.
- We perform a forensic and diplomatic examination of the object (both its content and its artifactual form) to ensure that its characteristics and content are consistent with the claims made about it and the record of its provenance.
- We rely on signatures and seals that are attached to the object or the claims that come with it, or both, and evaluate their forensics and diplomatics and their consistency with claims and provenance.
- For mass-produced and distributed (i.e., published) objects, we compare the object in hand with other versions (copies) of the object that may be available (which, in turn, means also assessing the integrity and provenance of these other versions or copies).

liability issues involved in an infrastructure that addresses authentication and identity; and social and cultural concerns about privacy, accountability, and related topics. Patent issues are a particular problem. It is hard to develop infrastructure, widely deployed standards, and critical mass when key elements are tied up by patents. With the recent insane proliferation of patents on software methods, algorithms, business models, and the like, uncertainty about patent issues is also a serious barrier to deployment. All of these have been well covered in the literature and the press. What has been less well examined is the lack of clear, well-established economic models to support systems of authentication and integrity management. To put it bluntly, it is not clear who is willing to pay for the substantial development, deployment, and operation of such a system. While many people say they are worried about authenticity and integrity in a digital environment, it is not clear that they are willing to pay the increased costs to effectively address these concerns.

In the digital environment, there are few forensics or diplomatics,³ other than the forensics and diplomatics of content itself. We cannot evaluate inks, papers, binding technology, and similar physical characteristics.⁴ We can note, just as with a physical work, that an essay allegedly written in 1997 that makes detailed references to events and publications from 1999 is either remarkably prescient or incorrectly dated. There are limited forensics of availability, and they mainly provide negative information. For example, if a document claims to have been written in 1998 and we have copies of it that were deposited on various servers in 1997 (and we trust the claims of the servers that the material was in fact deposited in 1997), we can build a case that it was first distributed *no later than 1997*, regardless of the date contained in the object. Nevertheless, this does not tell us when the document was written.

The fundamental concept of publication in the digital environment—the dissemination of a large number of copies to arbitrary interested parties that are subsequently autonomously managed and maintained—has come under great stress from numerous factors in the networked information environment. These factors include, for example, the move from sale to licensing, limited distribution, making copies public for viewing without giving viewers permission to maintain the copies, and technical protection systems (National Research Council 2000). While the basic principle of broad distribution and subsequent autonomous management of copies remains valid and useful as a base of evidence against which to test the authenticity of documents in question, the availability of relevant and trustworthy copies may be limited in the digital environment, and assessing the copies is likely to be more difficult. Moreover, the forensics and diplomatics of evaluating seals and signatures, and documentation of provenance, become much more formal and computational. It is difficult to say whether digital seals and signatures are more or less compelling in the digital world than in the analog world, but their characters unquestionably change. Finally, provenance and chains of custody in the digital world begin to reflect our evaluation of archives and custodians as implementers and operators of “trusted systems” that enforce the integrity and provenance records of objects entrusted to them.

At some level, authenticity and integrity are mechanical characteristics of digital objects; they do not speak to deeper questions of

³ It is worth carefully examining the forensic clues available when evaluating a digital object as an artifact. Today, many of them seem trivial, but as our history with digital technology grows longer, understanding them will likely become a specialized body of expertise. Examples include character codes, file formats, and formats of embedded fonts, all of which can help at least place the earliest time that a digital object could be created, and perhaps even provide evidence to argue that it was unlikely to have been created after a certain time. For an object that has undergone format conversions over time as part of its preservation, these forensic clues help only in the evaluation of the record of provenance.

⁴ For digital objects created by digitizing physical artifacts, if we can identify and obtain access to the source physical artifact, we can apply well-established forensic and diplomatic analysis practices to the source object.

whether the contents of a digital document are accurate or truthful when judged objectively. An authentic document may faithfully transmit complete falsehoods. There is a hierarchy of assessment in operation: forensics, diplomatics, intellectual analyses of consistency and plausibility, and evaluations of truthfulness and accuracy. Our concern here is with the lower levels of this hierarchy (i.e., forensics and diplomatics as they are reconceived in the digital environment) but we must recognize that conclusive evaluations at the higher levels may also provide evidence that is relevant to lower-level assessment.

Exploring Definitions and Defining Terms: Digital Objects, Integrity, and Authenticity

The Nature of Digital Information Objects

Before we can discuss integrity and authenticity, we must examine the objects to which we apply these characterizations.

Most commonly, computer scientists are concerned with digital objects that are defined as a set of sequences of bits. One can then ask computationally based questions about whether one has the correct set of sequences of bits, such as whether the digital object in one's possession is the same as that which some entity published under a specific identifier at a specific point in time. However, this is a simplistic notion. There are additional factors to consider.

Bits are not directly apprehended by the human sensory apparatus—they are never truly artifacts. Instead, they are rendered, executed, performed, and presented to people by hardware and software systems that interpret them. The question is how sophisticated these environmental hardware and software systems are and how integral they are to the understanding of the bits. In some cases, the focus is purely on the bits: numeric data files, or sensor outputs, for example, that are manipulated by computational or visualization programs. Documentary objects are characterized primarily by their bits (think of simple ASCII text), but the craft of publishing begins to make a sensory presentation of this collection of bits—to turn content into experience. Text, marked up in HTML and displayed through a Web browser, takes on a sensory dimension; the words that make up the text being rendered no longer tell the whole story. Digital objects that are performed—music, video, images that are rendered on screen—incorporate a stronger sensory component. Issues of interaction with the human sensory system—psychoacoustics, quality of reproduction, visual artifacts, and the like—become more important. The bits may be the same across space and time, but because of differences in the hardware and software used by recipients, the experience of viewing them may vary substantially. This raises questions about how to define and measure authenticity and integrity. In the most extreme case, we have objects that are rendered experientially—video games, virtual reality walk-throughs, and similar interactive works—where the focus shifts from the bits that constitute the digital

object to the behavior of the rendering system, or at least to the interaction between the digital object and the rendering system.

Thus, we might think about a hierarchy of digital objects that could be expressed as follows:

(Interactive) experiential works

Sensory presentations

Documents

Data

As we move up the hierarchy, from data to experiential works, the questions about the integrity and authenticity of the digital objects become more complex and perhaps more subjective; they address experience rather than documentary content (Lynch 2000). This paper will focus on the lower part of the digital object hierarchy. The upper part is poorly understood and today is addressed only in a limited way; for example, through discussions about emulation as a preservation strategy (Rothenberg 1999, 1995). It seems conceivable that one could extend some of the observations and assertions discussed later in this paper to the more experiential works by performing computations on the output of the renderings rather than on the objects themselves. However, this approach is fraught with problems involving canonical representations of the user interface (which, in the most complex cases, involves interaction and not just presentation) and agreeing on what constitutes the authentic experience of the work.

In meatspace, we cheerfully extend the notion of authenticity to much more than objects—in fact, we explicitly apply it to the experiential sphere, speaking of an “authentic” performance of a baroque concerto or an “authentic” Hawaiian luau. To the extent that we can make the extension and expansion of the use of authenticity as a characteristic precise within the framework and terminology of this paper, these statements seem to parallel statements about integrity of what in the digital environment could be viewed as experiential works, or performance.

Even as we struggle with definitions and tests of integrity and authenticity for intellectual works in the digital environment, we are seeing new classes of digital objects—for example, e-cash and digital bearer bonds—that explicitly involve and rely upon stylized and precise manipulation of provenance, authenticity, identity and anonymity, and integrity within a specific trust framework and infrastructure. While these fit somewhere between data and documents in the digital object hierarchy, they are interesting because they derive their meaning and significance from their explicit interaction with frameworks of integrity, authenticity, provenance, and trust.

Canonicalization and (Computational) Essence

Often, we seek to discuss the *essence* of a work rather than the exact set of sequences of bits that may represent it in a specific context; we are concerned with integrity and authenticity as they apply to this essence, rather than to the literal bits. Discussions of essence become more problematic as we move up the digital object hierarchy. How-

ever, even at the lower levels of data and documents, we encounter a troublesome imprecision that is a barrier to making definitions operational computationally when we move beyond the literal definition of precisely equivalent sets of sequences of bits. Those approaching the question from a literary or documentary perspective cast the issue in a palette of grays: there are series (not necessarily a strict hierarchy; at best a partial ordering) of intellectual abstractions of a document that capture its essence at various levels, and the key problem is whether this abstract essence is retained. The abstraction may involve words, layout, typography, or even the feel of the pages. Are hardcover and paperback editions of a book equivalent? Does equivalence depend on whether the pagination is identical? Elsewhere, I have proposed *canonicalization* as a method of making such abstractions precise (Lynch 1999). The fundamental point of canonicalization as an organizing principle is that it defines *computational algorithms* (called “canonicalizations”) that can be used to extract the “essence” of documents according to various definitions of what constitutes that essence. If we have such computational procedures for extracting the essence of digital objects, we can then compare digital objects through the prism of that definition of essence. We can also make assertions that involve abstract representations of this essence, rather than more specific (and presumably haphazard) representations that incorporate extraneous characteristics.

The hard problem, of course, is precisely defining and achieving a consensus about the right canonicalization algorithm, or algorithms, for a given context.

Integrity

When we say that a digital object has “integrity,” we mean that it has not been corrupted over time or in transit; in other words, that we have in hand the same set of sequences of bits that came into existence when the object was created. The introduction of appropriate canonicalization algorithms allows us to consider the integrity of various abstractions of the object, rather than of the literal bits that make it up, and to operationalize this discussion of abstractions into equality of sets of sequences of bits produced by the canonicalization algorithm.

When we seek to test the integrity of an object, however, we encounter paradoxes and puzzles. One way to test integrity is to compare the object in hand with a copy that is known to be “true.”⁵ Yet, if we have a secure channel to a known true copy, we can simply

⁵ As soon as we begin to speak of copies, however, we need to be very careful. Unless we know the location of the copy through some external (contextual) information, we run the risk of confusing authenticity and integrity. For example, if we have an object that includes a claim that “the identifier of this object is N” and we simply go looking for copies of objects with identifier N on a server that we trust, and then securely compare the object in hand with one of these copies, what we have really done is simply to trust the server to make statements about the assignment of the identifier N and then confirmed we had an accurate copy of the object with that identifier in hand. The key difference is between trusting the server to keep a true copy of an object in a known place and trusting the server to vouch for the assignment of an identifier to an object.

take a duplicate of the known true copy. We do not need to worry about the accuracy of the copy in hand, unless the point of the exercise is to ensure that the copy in hand is correct—for example, to detect an attempt at fraud, rather than to be sure that we have a correct copy. These are subtly different questions.⁶

If we do not have secure access to an independently maintained, known true copy of the object (or at least a digest surrogate), then our testing of integrity is limited to internal consistency checking. If the object is accompanied by an authenticated (“digitally signed”) digest, we can check whether the object is consistent with the digest (and thus whether its integrity has been maintained) by recomputing the digest from the object in hand and then comparing it with the authenticated digest. But our confidence in the integrity of the object is only as good as our confidence in the authenticity and integrity of the digest. We have only changed the locus of the question to say that *if* the digest is authentic and accurate, then we can trust the integrity of the object. Verifying integrity is no different from verifying the authenticity of a claim that “the correct message digest for this object is M” without assigning a name to the object. The linkage between claim and object is done by association and context—by keeping the claim bound with the object, perhaps within the scope of a trusted processing system such as an object repository.

In the digital environment, we also commonly encounter the issue of what might be termed “situational” integrity, i.e., the integrity of derivative works. Consider questions such as “Is this an accurate transcript?”, “Is this a correct translation?”, or “Is this the best possible version given a specific set of constraints on display capability?” Here we are raising a pair of questions: one about the integrity of a base object, and another about the correctness of a computation or other transformation applied to the object. (To be comprehensive, we must also consider the integrity of the result of the computation or transformation after it has been produced). This usually boils down to trust in the source or provider of the computation or transformation, and thus to a question of authentication of source or of validity, integrity, and correctness of code.

Authenticity

Validating authenticity entails verifying claims that are associated with an object—in effect, verifying that an object is indeed what it

⁶ One thing that we can do with cryptographic technology—specifically, digest algorithms—is to test whether two copies of an object are identical without actually exchanging the object. This is important in contexts where economics and intellectual property come into play. For example, a publisher that is offering copies of a digital document for license can also offer a verification service, where the holder of a copy of a digital object can verify its integrity without having to purchase access to a new copy. Or, two institutions, each of which holds a copy of a digital object but does not have to rights to share it with another institution, can verify that they hold the same object. Digest algorithms are also useful for efficiency purposes, because they avoid the need to transmit copies of what may be very large objects in order to test integrity. We should note that digest algorithms are *probabilistic* statements, however; the algorithms are designed to make it very unlikely that two different objects (particularly two similar but distinct documents) will have the same digest.

claims to be, or what it is claimed to be (by external metadata). For example, an object may claim to be created on a given date, to be authored by a specific person, or to be the object that corresponds with a name or identifier assigned by some organization. Some claims may be more mechanistic and indirect than others. For example, a claim that “This object was deposited in a given repository by an entity holding this public/private key pair at this time” might be used as evidence to support authorship or precedence in discovery. Typically, claims are linked to an object in such a way that they include, at least implicitly, a verification of integrity of the object about which claims are made. Rather than simply speaking of the (implied) object accompanying the claim (under the assumption that the correct object will be kept with the claims, and that the object management environment will ensure the integrity of the object) one may include a message digest (and any necessary information about canonicalization algorithms to be applied prior to computing the digest) as part of the metadata assertion that embodies the claim.

It is important to note that tests of authenticity deal only with specific claims (for example, “did X author this document?”) and not with open-ended inquiry (“Who wrote it?”). Validating the authenticity of an object is more limited than is an open-ended inquiry into its nature and provenance.

There are two basic strategies for testing a claim. The first is to believe the claim because we can verify its integrity and authenticate its source, and because we choose to trust the source. In other words, we validate the claim that “A is the author of the object with digest X” by first verifying the integrity of the object relative to the claim (that it has digest X), and then by checking that the claim is authenticated (i.e., digitally signed) by a trusted entity (T). The heart of the problem is ensuring that we are certain who T really is, and that T really makes or warrants the claim. The second strategy is what we might call “independent verification” of the claim. For example, if there is a national author registry that we trust, we might verify that the data in the author registry are consistent with the claim of authorship. In both cases, however, validating a claim that is associated with an object ultimately means nothing more or less than making the decision to trust some entity that makes or warrants the claim.

Several final points about authenticity merit attention. First, trust in the maker or warrantor of a claim is not necessarily binary; in the real world, we deal with levels of confidence or degrees of trust. Second, many claims may accompany an object; in evaluating different claims, we may assign them differing degrees of confidence or trust. Thus, it does not necessarily make sense to speak about checking the authenticity of an object as if it were a simple true-or-false test—a computation that produces a one or a zero. It may be more constructive to think about checking authenticity as a process of examining and assigning confidence to a collection of claims. Finally, claims may be interdependent. For example, an object may be accompanied by claims that “This is the object with identifier N,” and “The object with identifier N was authored by A” (the second claim, of course, is

independent of the document itself, in some sense). Perhaps more interesting, in an archival context, would be claims that “This object was derived from the object with message digest M by a specific re-formatting process” and “The object with message digest M was authored by A.” (See Lynch 1999 for a more detailed discussion of this case.)

Comparing Integrity and Authenticity

It is an interesting, and possibly surprising, conclusion that in the digital environment, tests of integrity can be viewed as just special cases and byproducts of evaluations of authenticity. Part of this comes from the perspective of the environment of “pervasive deceit” and the idea that checking integrity of an object means comparing it with some precisely identified and rigorously vetted “original version” or “authoritative copy.” In fact, much of the checking for integrity in the physical world is not about ferreting out pervasive deceit and malice, but rather about accepting artifacts for roughly what they seem to be on face value and then looking for evidence of damage or corruption (i.e., torn-out pages or redacted text). For this kind of integrity checking, a message digest that accompanies a digital object as metadata serves as an effective mechanism to ensure that the object has not been damaged or corrupted. This is true even if the message digest is not supported by an elaborate signature chain and trust assessment, but only by a general level of confidence in the computational context in which the objects are being stored and transmitted. In the digital environment, there is a tendency to downplay the need for this kind of integrity checking in favor of stronger measures that combine authenticity claims with integrity checks.

The Role of Copies

David Levy argues that all digital objects are copies; this echoes the findings of the National Research Council Committee on Intellectual Property in the Emerging Information Infrastructure that use—reading, for example—implies the making of copies (National Research Council 2000). If we accept this view, authenticity can be viewed as an assessment that we make about something in the present—something that we have in hand—relative to claims about the past (predecessor copies). The persistent question is whether a given object X has the same properties as object Y. There is no “original.” This is particularly relevant when we are dealing with dynamic objects such as databases, where an economy of copies is meaningless. In such cases, there is no question of authenticity through comparison with other copies; there is only trust or lack of trust in the location and delivery processes and, perhaps, in the archival custodial chain.

Provenance

The term *provenance* comes up often in discussions of authenticity and integrity. Provenance, broadly speaking, is documentation about the origin, characteristics, and history of an object; its chain of custody; and its relationship to other objects. The final point is particularly

important. There are two ways to think about a digital object that is created by changing the format of an older object that has been validated according to some specific canonicalization algorithm. We might think about a single object the provenance of which includes a particular transformation, or we might think about multiple objects that are related through provenance documentation. Thus, provenance is not simply metadata about an object—it can also be metadata that describe the relationships between objects. Because provenance also includes claims about objects, it is part of the authentication and trust infrastructures and frameworks.

I do not believe that we have a clear understanding of (and surely not consensus about) where provenance data should be maintained in the digital environment, or by what agencies. Indeed, it is not clear to what extent the record of provenance exists independently and permanently, as opposed to being assembled when needed from various pools of metadata that may be maintained by various systems in association with the digital objects that they manage. We also lack well-developed metadata element sets and interchange structures for documenting provenance. It seems possible that the Dublin Core, augmented by semantics for signing metadata assertions, might form a foundation for this, although attributes such as relationship would need to be extended to allow for very precise vocabularies to describe algorithmically based derivations of objects from other objects (or transformations of objects). We would probably also need to incorporate metadata assertions that allow an entity to record claims such as “Object X is equivalent to object Y under canonicalization C.”

Watermarks, Authenticity, and Integrity

In the most general sense, watermarking can be viewed as an attempt to ensure that a set of claims is inseparably bound to a digital object and thus can be assumed to travel with the object; one does not have to trust transport and storage systems to correctly perform this function. The most common use of watermarks today is to help protect intellectual property by attaching a copyright claim (and possibly an object-specific serial number to allow tracing of individual copies) to an object. Software exists to scan public Web sites for objects that contain watermarks and to notify the rights holders about where these objects have been found. A serial number, if present, helps the rights holder not only identify the presence of a possibly illegal copy but also determine where it came from. Various trusted system-based architectures for the control of copyrighted works have also been proposed that use watermarking (for example, the Secure Digital Music Initiative [2000]). The idea is that devices will refuse to play, print, or otherwise process digital objects if the appropriate watermarks are not present.⁷ The desirable properties of watermarks include being very hard to remove computationally (at least without knowledge of the private key as well as the algorithm used to generate the watermark) and being resilient under various alterations that

may be applied to the watermarked file (lossy compression, for example, or image cropping). The development of effective watermarking systems is currently a very active area of research.⁸

From the perspective of authenticity and integrity, watermarks present several problems. First, they deliberately and systematically corrupt the objects to which they are applied, in much the same way that techniques such as lossy compression do. Fingerprints (individualized watermarks) are particularly bad in this regard since they defeat comparisons among copies as a way of establishing authenticity—indeed this is exactly what they are designed to do, to make each copy unique and traceable. Applying a watermark to a digital object means changing bits within the object, but in such a way that they change the perception of the object only slightly. Thus, finding and verifying a watermark in a digital object give us only weak evidence of its integrity. In fact, the very presence of the watermark means that integrity has been compromised at some level, unless we are willing to accept the watermarked version of the object as the actual authoritative one—an image or sound recording that includes some data that allegedly does not much change our perception of the object. If a watermark can easily be stripped out of an object (a bad watermark design, but perhaps characteristic of watermarking systems that try to minimize corruption), then the absence of such a watermark does not tell us much about the possible corruption of other parts of the object.

A second problem is that some watermarking systems do not emphasize preventing the creation of fake watermarks; they are concerned primarily with the preservation of legitimate watermarks as evidence of ownership or status of the watermarked object. To use watermarking to address authenticity issues, it seems likely that one would need to use it simply as a means of embedding a claim in an object, under the assumption that the claim would then have to be separately verifiable (for example, by being digitally signed).

To summarize: If one obtains a digital object that contains a watermark, particularly if that watermark contains separately verifiable claims, it can provide useful evidence about the provenance and characteristics of the object, including good reasons to assume that it is a systematically and deliberately corrupted version of a predeces-

⁷ This is not a universally accepted definition of a digital watermark. The term is also used to refer to other things, such as modifications to images that allow them to be viewed on-screen with only moderate degradation but that produce very visible and unsightly artifacts when the image is printed. The description here characterizes what I believe to be the most commonly used definition of the technology. Sometimes “watermark” is reserved for a “universal” encoding hidden in all copies of a digital object that are distributed by a given source (for example, containing an object identifier) and the term “fingerprint” is reserved for watermarks that are copy-specific, that is personalized to given recipients (containing a serial number or the recipient’s identifier). The fingerprint individualizes an object to a version associated with a specific recipient.

⁸ See, for example, the proceedings of the series of conferences on Information Hiding (Anderson 1996, Aucsmith 1998, Pfitzmann 2000). See also proceedings from the first, second, and third international conferences on financial cryptography (Hirschfeld 1997, Hirschfeld 1998, Franklin 1999).

sor digital object that one may or may not have access to or be able to locate. The watermark may have some value in forensic examination of digital objects, but it does not seem to be a good tool for the *management* of digital objects within a controlled environment such as an archive or repository system that is concerned with object integrity. It seems more appropriate to require that the environment take responsibility for maintaining linkages and associations between metadata (claims) and the objects themselves. Watermarks are more appropriate for an uncontrolled public distribution environment where integrity is just one variable in a complex set of trade-offs about the management and protection of content.

Semantics of Digital Signatures

One serious shortcoming of current cryptographic technology has to do with the semantics of digital signatures—or, more precisely, the lack thereof. In fairness, many cryptographers are not concerned with replicating the higher levels of semantics that accompany the use of signatures in the physical world. They regard these issues as the responsibility of an applications environment that uses digital signatures as a tool or supporting mechanism. But wherever we assign responsibility for establishing a system of semantics, the need for such semantics is very real, and I believe that many people outside the cryptographic community have been misled by their assumptions about the word *signature*. They do not understand that the semantics problem is still largely unaddressed.

At its core, a digital signature is a mechanical, computational process. Some entity in possession of a public/private key pair was willing to perform a computation on a set of data using this key pair, which permits someone who knows the public key of the key pair to verify that the data were known to and computed upon by an entity that held the key pair. A digital signature amounts to nothing more than this. Notice that any digital data can be signed—not just documents or their digests, but also assertions about documents. The interface between digital signature processing and documents is extremely complex, questions about the semantics of signatures aside. The reader is invited to explore the work of the joint Worldwide Web Consortium/Internet Engineering Task Force on digital signatures for XML documents (1998) to get a sense of how issues such as canonicalization come into play here.

The use of digital signatures in conjunction with a public key infrastructure (PKI) offers a little more.⁹ People can choose to trust the procedures of a PKI to do the following kinds of things:

- To verify, according to published policies, a user's right to an "identity" and to subsequently document the binding between that identity and a public/private key pair. Verification policies

⁹ See, for example, Ford and Baum 1997; Feghhi, Geghhi, and Williams 1999.

vary widely, from taking someone's word in an e-mail message to demanding witnesses, extensive documentation such as passports and birth certificates, personal interviews, and other proof. In essence, one can trust the PKI service to provide the public key that corresponds to an identity. The identity can be either a name ("John Smith") or a role ("Chief Financial Officer of X Corporation"). Attributes can also be bound to the identity.

- To provide a means for determining when a key pair/identity binding has been compromised, expired, or revoked and should no longer be considered valid.

Compare this mechanistic view of signatures with the rich and diverse semantics of signatures in the real world. A signature might mean that the signer

- authored the document;
- witnessed the document and other signatures on it;
- believes that the document is correct;
- has seen, or received, the document;
- approves the actions proposed in the document; or
- agrees to the document.

There are questions not only about the meaning of signatures but also about their scope. In some situations, for example, documents are signed or initialed on every page; in others, a signature witnesses only another signature, not the entire document. Questions of scope become complex in a digital world, particularly as signed objects undergo transformations over time (because of reformatting, for example). Considerable research is needed in these areas.

Digital signatures alone can neither differentiate among the possible semantics outlined earlier, nor provide direct evidence of any one of them. In other words, there is no reasonable "default" meaning that can be given to a signature computation. Such signatures can tell us that a set of bits has been computed upon, and, in conjunction with a PKI, they can tell us who performed that computation. We clearly need a mechanism for expressing semantics of signatures that can be used in conjunction with the actual computational signature mechanism—a vocabulary for expressing the meaning of a signature in relationship to a digital object (or, in fact, a set of digital objects that might include other signed assertions).

One can imagine defining such a vocabulary and interchange syntax for the management and preservation of digital objects—for a community of archives and cultural heritage organizations, for example. But there is another problem that has not been well explored, to my knowledge. It is likely that we will see the development of one or more "public" vocabularies for commerce and contracting, and perhaps additional ones for the registry and management of intellectual property. These vocabularies might vary among nations, or even among states in a nation such as the United States, where much con-

tracting is governed by state law.¹⁰ In addition, we will almost certainly see the development of organization-specific “internal” vocabularies in support of institutional processes. Many of the initial claims about objects will likely be expressed in one of these other vocabularies rather than the vocabularies of the cultural heritage communities; consequently, we will face complex problems of mapping and interpreting vocabularies. We will also face the problems of trying to interpret vocabularies that may belong to organizations that no longer exist or vocabularies in which usage has changed over time, perhaps in poorly documented ways.

The Roles of Identity and Trust

Virtually all determination of authenticity or integrity in the digital environment ultimately depends on trust. We verify the source of claims about digital objects or, more generally, claims about sets of digital objects and other claims, and, on the basis of that source, assign a level of belief or trust to the claims. As a second, more intellectual form of analysis, we can consider the consistency of claims, and then further consider these claims in light of other contextual knowledge and common sense. For example, an object that claims to have been authored in 2003 by someone who died in 2001 would reasonably raise questions, even if all of the signatures verify. We can draw precious few conclusions from objects standing alone, except by applying this kind of broader intellectual analysis. As we have seen, ensuring the validity of linkages between claims and the objects about which those claims make assertions is an important question. The question becomes even more difficult when we recognize that both objects and sets of claims evolve independently and at different rates, because of maintenance processes such as reformatting or the expiration of key pairs and the issuance of new ones.

Ultimately, trust plays a central role, yet it is elusive. Signatures can allow us to trust a claim if we trust the holder of a key pair, and a public key infrastructure can allow us to know the identity (name) of the holder of a key pair if we trust the operator of the PKI. If we know the name of the entity we trust, we can thus use the PKI to determine its public key and use that to verify signatures that the entity has made. We can establish the link between identity and keys directly (we can directly obtain, through some secure method, the public key from a trusted entity) or through informal intermediaries (we can securely obtain the key from someone we know and trust, as is done in the Pretty Good Privacy [PGP] system) (Zimmermann 1995).

It is important to recognize that trust is not necessarily an absolute, but often a subjective probability that we assign case by case.

¹⁰ In the United States, some of this is likely to be determined by how quickly federal law regarding digital signatures is established and by the extent to which federal law preempts developing state laws. Changes to the Uniform Commercial Code will likely play a role. See <http://washofc.epic.org/crypto/dss/> for information on a variety of material on current legislative and standards developments related to digital signatures.

The probability of trustworthiness may be higher for some PKIs than for others, because of their policies for establishing identity. Moreover, we may establish higher levels of trust based on identities that we have directly confirmed ourselves than on those confirmed by others. Considerable research is being done on methods that people could use to define rules about how they assign trust and belief. These rules can drive computations for a calculus of trust in evaluating claims within the context of a set of known keys and identities and PKI services that maintain identities. An interesting question, which I do not think we are close to being able to answer, is whether there will be a community consensus on trust assignment rules within the cultural heritage community, or whether we will see many, wildly differing, choices about when to establish trust.

We also need an extensive inquiry into the nature of identity in the digital world as it relates to authenticity questions such as claims of authorship. Consider just a few points here. Identity in the digital world means that someone has agreed to trust an association between a name and a key pair, because he or she has directly verified it or trusts an intermediary, such as a PKI, that records such an association. Control of an identity, however, can be mechanically transferred or shared by the simple act of the owner of a key pair sharing that key pair with some other entity. We have to trust not only the identity but also the behavior of the owner of that identity.

If we are to trust a claim of authorship, whom do we expect to sign it? The author? The publisher? A registry such as the copyright office, which would more likely sign a claim stating that the author has registered the object and claimed authorship?

Identity is more than simply a name. We frequently find anonymous or pseudonymous authorship; how are these identities created and named? We have works of corporate authorship, including the notion of “official” works that are created through deliberate corporate acts and that represent policy or statements with legal implications. In this case, the signatory may be someone with a specific role or office within a corporation (an officer of the corporation or the corporate secretary, for example). These may be very volatile in an era of endless mergers and acquisitions, as well as occasional bankruptcies. Finally, we have various ad-hoc groups that come together to author works; these groups may be unwilling or unable to create digital identities within the trust and identity infrastructure (consider, for example, artistic, revolutionary, or terrorist manifestos).

We know little about how identity management systems operate over very long periods. Imagine a digital object that is released from an archive in 2100 for the first time—an object that had been sealed since its deposit in 2000. A group of experts is trying to assess the claims associated with the object. One scenario is that all claims were verified upon deposit, and the archive has recorded that verification; the experts then trust the archive to have correctly maintained the object since its deposit and to have appropriately verified the claims. A second scenario is that the group of experts chooses to re-verify the claims. This may take them into an elaborate exploration of the his-

torical evolution of policies of certificate authorities and public key infrastructure operators that have long since vanished, of histories of key assignment and expiration, and perhaps even of the evolution of our understanding of the vulnerabilities of cryptographic algorithms themselves. This suggests that our ability to manage and understand authenticity and integrity over long periods of time will require us to manage and preserve documentation about the evolution of the trust and identity management infrastructure that supports the assertions and evaluation of authenticity and integrity. This, in turn, raises the concern that relying on services and infrastructure that are being established primarily to support relatively short-term commercial activities may be problematic. At a minimum, it suggests that we may need to begin a discussion about the archival requirements for such services if they are to support the long-term management of our cultural and intellectual heritage.

Authorship is just one example of the difficulties involved in “literary” signature semantics. Consider the problem of assigning publication dates as another example. Every publisher has different standards and thus different semantics.

Conclusions

In an attempt to explore the central roles of trust and identity in addressing authenticity and integrity for digital objects, this paper points to a wide-ranging series of research questions. It identifies the need to begin considering standardization efforts in areas such as signing metadata claims and the semantics of digital signatures to support authenticity and integrity.

But a set of more basic issues about infrastructure development and large-scale deployment also needs to be carefully considered. A great deal of technology and infrastructure now being deployed will be useful in managing integrity and authenticity over time. However, these developments are being driven by commercial requirements with short time horizons in areas such as authentication, electronic commerce, electronic contracting, and management and control of digital intellectual property. The good news is that there is a huge economic base in these areas that will underwrite the development of infrastructure and drive deployment. To the extent that we can share this work to manage cultural and intellectual heritage, we need to worry only about how to pay to use it for these applications, not about how to underwrite its development. Even there, however, we need to think about who will pay to establish the necessary identities and key pairs and to apply them to create the appropriate claims that will accompany digital objects. The less-good news is that we need to be sure that the infrastructure and deployed technology base actually meet the needs of very long-term management of digital objects. To take one example, knowing the authorship of a work is still important, even after all the rights to the work have entered the public domain. It is essential that institutions concerned with the management and preservation of cultural and intellectual heritage engage,

participate in, and continue to critically analyze the development of the evolving systems for implementing trust, identity, and attribution in the digital environment.

REFERENCES

- Anderson, Ross, ed. 1996. Information Hiding: First International Workshop, Cambridge, U.K., May 30–June 1, 1996, proceedings. *Lecture Notes in Computer Science*, vol. 1174. Berlin and New York: Springer.
- Aucsmith, David, ed. 1998. Information Hiding: Second International Workshop, Portland, Oregon, U.S.A., April 14–17 1998, proceedings. *Lecture Notes in Computer Science*, vol. 1525. Berlin and New York: Springer.
- Bearman, David, and Jennifer Trant. 1998. Authenticity of Digital Resources: Towards a Statement of Requirements in the Research Process, *D-Lib Magazine* (June). Available from <http://www.dlib.org/dlib/june98/06bearman.html>.
- Duranti, Luciana. 1998. *Diplomatics: New Uses for an Old Science*. Lanham, Md.: Scarecrow Press.
- Hirschfeld, Rafael, ed. 1997. Financial Cryptography: First International Conference, Anguilla, British West Indies, February 24–28, 1997, proceedings. *Lecture Notes in Computer Science*, vol. 1318. Berlin and New York: Springer.
- Hirschfeld, Rafael, ed. 1998. Financial Cryptography: Second International Conference, Anguilla, British West Indies, February 23–25, 1988, proceedings. *Lecture Notes in Computer Science*, vol. 1465. Berlin and New York: Springer.
- Feghhi, Jalal, Jalil Geghhi, and Peter Williams. 1999. *Digital Certificates: Applied Internet Security*. Reading, Mass.: Addison Wesley.
- Ford, Warwick, and Michael S. Baum. 1997. *Secure Electronic Commerce: Building the Infrastructure for Digital Signatures and Encryption*. Upper Saddle River, N.J.: Prentice Hall.
- Franklin, Matthew, ed. 1999. Financial Cryptography: Third International Conference, Anguilla, British West Indies, February 22–25, 1999, proceedings. *Lecture Notes in Computer Science*, vol. 1648. Berlin and New York: Springer.
- Lessig, Lawrence. 1999. *Code and Other Laws of Cyberspace*. New York: Basic Books.

Lynch, Clifford. 2000. "Experiential Documents and the Technologies of Remembrance," in *I in the Sky: Visions of the Information Future*, edited by Alison Scammell. London: Library Association Publishing.

Lynch, Clifford. 1999. Canonicalization: A Fundamental Tool to Facilitate Preservation and Management of Digital Information, *D-Lib Magazine* 5(9) (September). Available from <http://www.dlib.org/dlib/september99/09lynch.html>.

National Research Council. 2000. *The Digital Dilemma: Intellectual Property in the Information Infrastructure*. Washington, D.C.: National Academy Press.

Pfitzmann, Andreas, ed. 2000. Information Hiding: Third International Workshop, Dresden, Germany, September 29–October 1, 1999, proceedings. *Lecture Notes in Computer Science*, vol. 1768. Berlin and New York: Springer.

Rothenberg, Jeff. 1999. *Avoiding Technological Quicksand: Finding a Viable Technical Foundation for Digital Preservation*. Washington, D.C. Council on Library and Information Resources. Available from <http://www.clir.org>.

Rothenberg, Jeff. 1995. Ensuring the Longevity of Digital Documents. *Scientific American* 272(1):24-9.

Secure Digital Music Initiative. 2000. Available from <http://www.sdmi.org>.

Worldwide Web Consortium/Internet Engineering Task Force on Digital Signatures for XML Documents. 1998. Digital Signature Initiative. Available from <http://www.w3.org/DSig>.

Zimmerman, Philip R. 1995. *The Official PGP User's Guide*. Cambridge, Mass.: MIT Press.

Preserving Authentic Digital Information

by Jeff Rothenberg

Introduction

This paper argues that to better understand what is required to meaningfully preserve digital information, we should attempt to create a foundation for the concept of the authenticity of informational entities that transcends the multiple disciplines in which this concept arises. Whenever informational entities are used and for whatever purpose, their suitability relies on their authenticity. Yet archivists, librarians, museum curators, historians, scholars, and researchers in various fields define authenticity in distinct, though often overlapping, ways. They combine legal, ethical, historical, and artistic perspectives such as the desire to provide accountability, the desire to ensure proper attribution, or the desire to recreate, contextualize, or interpret the original meaning, function, impact, effect, or aesthetic character of an artifact. Each discipline may have its own explicit definition of authenticity; however, in interdisciplinary discussions of authenticity, the dependence of a given definition on its discipline is often manifested only implicitly.

The technological issues surrounding the preservation of digital informational entities interact with authenticity in novel and profound ways. We are far more likely to achieve meaningful insights into the implications of these interactions if we develop a unified, coherent, discipline-transcendent view of authenticity. Such a view would

- improve communication across disciplines;
- provide a better basis for understanding what preservation requirements are implied by the need for authenticity; and
- facilitate the development of common preservation strategies that would work for as many different disciplines as possible and thereby effect technological economies of scale.

Developing a preservation strategy that economically transcends disciplines would free preservationists from the need for discipline-specific definitions of authenticity. In this paper, I will suggest that there is at least one preservation strategy, based on the notion of a *digital-original*, that makes the details of how we define authenticity all but irrelevant from the perspective of preservation. However, to derive this conclusion, it is necessary to examine authenticity in some depth.

Although a discipline-transcendent view of authenticity would be the ideal, it may turn out to be impractical. If so, we may need to settle for a multidisciplinary perspective. This means establishing either a unified concept of authenticity as it is used in a subgroup of disciplines (such as archives, libraries, and museums) or a set of variant concepts of authenticity, each of which addresses the specific needs of a different discipline yet retains as much in common with the other concepts as possible.

Basic Definitions

The term *informational entity*, as used here, refers to an entity whose purpose or role is informational. By definition, any informational entity is entirely characterized by information, which may include contextual and descriptive information as well as the core entity. Examples of informational entities include digital books, records, multimedia objects, Web pages, e-mail messages, audio or video material, and works of art, whether they are “born digital” or digitized from analog forms.

It is not easy for computer scientists to agree on a definition of the word *digital*.¹ In the current context, it generally denotes any means of representing sequences of discrete symbolic values—each value having two or more unambiguously distinguishable states—so that these sequences can, at least in principle, be accessed, manipulated, copied, stored, and transmitted entirely by mechanical means with a high degree of reliability (Rothenberg 1999). Digital informational entities are defined in the next section.

The term *authenticity* is even harder to define, but the term is used here in its broadest sense. Its meaning is not restricted to authentication, as in verifying authorship, but is intended to include issues of integrity, completeness, correctness, validity, faithfulness to an original, meaningfulness, and suitability for an intended purpose. I leave to specialists in various scholarly disciplines the task of elaborating the dimensions of authenticity in those disciplines. The focus of this paper is the interplay between those dimensions and the technological issues involved in preserving digital informational entities. The dimensions of authenticity have a profound effect on the technical requirements of any preservation scheme, digital or otherwise.

¹ One could argue that if the key terms of any discipline are not susceptible to multiple interpretations and endless analysis, then that discipline has little depth.

The remainder of this paper discusses the importance of understanding authenticity as a prerequisite to defining meaningful digital preservation.

Digital Informational Entities are Executable Programs

The distinguishing characteristic of a digital informational entity is that it is essentially a program that must be *interpreted* to be made intelligible to a human: it cannot simply be held up to the light to be read. A program is a sequence of commands in a formal language that is intended to be read by an interpreter that understands that language.² An interpreter is a process that knows how to perform the commands specified in the formal language in which the program is written. Even a simple text document consisting of a stream of ASCII character codes is a program, i.e., it is a sequence of commands in a formal language (each command specifying a character to be rendered) that must be interpreted before it can be read by a human. More elaborate digital formats, such as distributed, hypermedia documents, may—in addition to requiring interpretation for navigation and rendering—embed macros, scripts, animation processes, or other active components, any of which may require arbitrarily complex interpretation.

Some programs are interpreted directly by hardware (for example, a printer may render ASCII characters from their codes), but the interpreters of most digital informational entities are software (i.e., application programs). Any software interpreter must itself be interpreted by another hardware or software interpreter, but any sequence of software interpretations must ultimately result in some lowest level (“machine language”) expression that is interpreted (“executed”) by hardware.

It follows that it is not sufficient to save the bit stream of a digital informational entity without also saving the intended interpreter of that bit stream. Doing so would be analogous to saving hieroglyphics without saving a Rosetta Stone.³

In light of this discussion, it is useful to define a digital informational entity as consisting of a single, composite bit stream⁴ that includes the following:

- the bit stream representing the *core content* of the entity (that is, the encoding of a document, data, or a record), including all structural

² Many programs are compiled or translated into some simpler formal language first, but the result must still ultimately be interpreted. The distinction between compilation and interpretation will therefore be ignored here.

³ Although this analogy is suggestive, it is simplistic, since the interpreter of a digital informational entity is itself usually an executable application program, not simply another document.

⁴ Any number of component bit streams can be represented as a single, composite bit stream.

information required to constitute the entity from its various components, wherever and however they may be represented;

- component bit streams representing all necessary contextual or ancillary information or metadata needed to make the entity meaningful and usable; and
- one or more component bit streams representing a perpetually executable interpreter capable of rendering the core content of the entity from its bit stream, in the manner intended.⁵

If we define a digital informational entity in this way, as including both any necessary contextual information and any required interpreter, we can see that preserving such an entity requires preserving all of these components.⁶ Given this definition, one of the key technical issues in preserving digital informational entities becomes how to devise mechanisms for ensuring that interpreters can be made perpetually executable.

Preservation Implies Meaningful Usability

The relationship between digital preservation and authenticity stems from the fact that meaningful preservation implies the usability of that which is preserved. That is, the goal of preservation is to allow future users to retrieve, access, decipher, view, interpret, understand, and experience documents, data, and records in meaningful and valid (that is, authentic) ways. An informational entity that is “preserved” without being usable in a meaningful and valid way has not been meaningfully preserved, i.e., has not been preserved at all.

As a growing proportion of the informational entities that we create and use become digital, it has become increasingly clear that we do not have effective mechanisms for preserving digital entities. As I have summarized this problem elsewhere: “There is as yet no viable long-term strategy to ensure that digital information will be readable in the future. Digital documents are vulnerable to loss via the decay and obsolescence of the media on which they are stored, and they become inaccessible and unreadable when the software needed to interpret them, or the hardware on which that software runs, becomes obsolete and is lost” (Rothenberg 1999).

The difficulty of defining a viable digital preservation strategy is partly the result of our failing to understand and appreciate the authenticity issues surrounding digital informational entities and the implications of these issues for potential technical solutions to the digital preservation problem. The following argues that the impact of authenticity on preservation is manifested in terms of usability,

⁵ The word *rendering* is used here as a generalization of its use in computer graphics, namely, the process of turning a data stream into something a human can see, hear, or otherwise experience.

⁶ Metadata and interpreter bit streams can be shared among many digital informational entities. Although they must be logical components of each such entity, they need not be redundantly represented.

namely that a preserved informational entity can serve its intended or required uses if and only if it is preserved authentically.

For traditional, analog informational entities, the connection between preservation and usability is obvious. If a paper document is “preserved” in such a way that the ink on its pages fades into illegibility, it probably has not been meaningfully preserved. Yet even in the traditional realm, it is at least implicitly recognized that informational entities have a number of distinct attributes that may be preserved differentially. For example, stone tablets bearing hieroglyphics that were physically preserved before the discovery of the Rosetta Stone were nevertheless unreadable because the ability to read the language of their text had been lost. Similarly, although the original Declaration of Independence has been preserved, most of its signatures have faded into illegibility. Many statues, frescos, tapestries, illuminated manuscripts, and similar works are preserved except for the fact that their pigments have faded, often beyond recognition. Although it is not always possible to fully preserve an informational entity, it may be worth preserving whichever attributes *can* be preserved if doing so enables the entity to be used in a meaningful way. In other words, if preserving certain attributes of an informational entity may allow it to fulfill some desired future use, then we are likely to consider those attributes worth preserving and to consider that we have at least partially preserved the entity by preserving those attributes. Generalizing from this, the meaningful preservation of any informational entity is ultimately defined in terms of which of its attributes can and must be preserved to ensure that it will fulfill its future use, whether originally intended, subsequently expected, or unanticipated.

Deciding which attributes of traditional informational entities to preserve involves little discretion. Because a traditional informational entity is a physical artifact, saving it in its entirety preserves (to the extent possible) all aspects of the entity that are inherent in its physical being, which is to say all of its attributes. Decisions may still have to be made, for example, about what technological measures should be used to attempt to preserve attributes such as color. For the most part, however, saving any aspect of a traditional information entity saves every aspect, because all of its aspects are embodied in its physicality.

For digital informational entities, the situation is quite different. There is no accepted definition of digital preservation that ensures saving all aspects of such entities. By choosing a particular digital preservation method, we determine which aspects of such entities will be preserved and which ones will be sacrificed. We can save the physical artifact that corresponds to a traditional informational entity in its entirety; however, there is no equivalent option for a digital entity.⁷ The choice of any particular digital preservation technology

⁷ In particular, saving the bit stream corresponding to the core content of such an entity is insufficient without saving some way of interpreting that bit stream, for example, by saving appropriate software (another bit stream) in a way that enables running that software in the future, despite the fact that it, and the hardware on which it was designed to run, may be obsolete.

therefore has inescapable implications for what will and will not be preserved. In the digital case (so far, at least), we must choose what to lose (Rothenberg and Bikson 1999).

This situation is complicated by the fact that we currently have no definitive taxonomies either of the attributes of digital informational entities or of the uses to which they may be put in the future. Traditional informational entities have been around long enough that we can feel some confidence in understanding their attributes, as well as the ways in which we use them. Anyone who claims to have a corresponding understanding of digital informational entities lacks imagination. Society has barely begun to tap the rich lode of digital capabilities waiting to be mined. The attributes of future digital informational entities, the functional capabilities that these attributes will enable, and the uses to which they may be put defy prediction.

Strategies for Defining Authenticity

It is instructive to consider several strategies that can be used to define authenticity. Each strategy may lead to a number of different ways of defining the concept and may, in turn, involve a number of alternative tactics that enable its implementation.

One strategy is to focus on the originality of an informational entity, that is, on whether it is unaltered from its original state. This strategy works reasonably well for traditional, physical informational entities but is problematic for digital informational entities. The *originality strategy* can be implemented by means of several tactics. One such tactic is to focus on the *intrinsic properties* of an informational entity by providing criteria for whether each property is present in its proper, original form. For example, one can demand that the paper and ink of a traditional document be original and devise chemical, radiological, or other tests of these physical properties.⁸

A second tactic for implementing the originality strategy is to focus on the *process* by which an entity is saved, relying on its provenance or history of custodianship to warrant that the entity has not been modified, replaced, or corrupted and must therefore be original. For example, from an archival perspective, a record is an informational artifact that provides evidence of some event or decision that was performed as part of the function of some organization or agency. The form and content of the record convey this evidence, but the legitimacy of the evidence rests on being able to prove that the record is what it purports to be and has not been altered or corrupted in such a way as to invalidate its evidential meaning. The archival principle of provenance seeks to establish the authenticity of archival records by providing evidence of their origin, authorship, and context of generation, and then by proving that the records have been

⁸ New criteria based on newly recognized properties of informational entities may be added over time, as is the case when evaluating radiological properties of artifacts whose origins predate the discovery of radioactivity.

maintained by an unbroken chain of custodianship in which they have not been corrupted.

Relying on this tactic to ensure the authenticity of records involves two conditions: first, that an unbroken chain of custodianship has been maintained; and second, that no inappropriate modifications have been made to the records during that custodianship. The first of these conditions is only a way of supplying indirect evidence for the second, which is the one that really matters. An unbroken chain of custodianship does not in itself prove that records have not been corrupted, whereas if we could prove that records had not been corrupted, there would be no logical need to establish that custodianship had been maintained. However, since it is difficult to obtain direct proof that records have not been corrupted, evidence of an unbroken chain of custodianship serves, at least for traditional records, as a surrogate for such proof.

Intrinsic properties of the entity may be completely ignored using this tactic, which relies on the authenticity of documentation of the process by which the entity has been preserved as a surrogate for the intrinsic authenticity of the entity. This has a somewhat recursive aspect, since the authenticity of this documentation must in turn be established; however, in many cases, this is easier than establishing the authenticity of the entity itself.

Alternatively, an *intrinsic properties strategy* can be based solely on the intrinsic properties tactic discussed above. This involves identifying certain properties of an informational entity that define authenticity, regardless of whether they imply the originality of the entity. For example, one might define an authentic impressionistic painting as one that conforms to the style and methods of Impressionism, regardless of when it was painted or by whom. A less controversial example might be a jade artifact that is considered "authentic" merely by virtue of being truly composed of jade.⁹ Whether this strategy is viable for a given discipline depends on whether the demands that the discipline places on informational entities can be met by ensuring that certain properties of those entities meet specified criteria, regardless of their origin.

Although there are undoubtedly other strategies, the final one I will consider here is to define authenticity in terms of whether an informational entity is suitable for some purpose. This *suitability strategy* would use various tactics to specify and test whether an informational entity fulfills a given range of purposes or uses. This may be logically independent of whether the entity is original. Similarly, although the suitability of an entity for some purpose is presumably related to whether certain of its properties meet prescribed criteria, under this strategy both the specific properties involved and the criteria for their presence are derived entirely from the purpose that the entity is to serve. Since a given purpose may be satisfiable by

⁹ Here authenticity refers to a specific attribute of an entity (i.e., its chemical composition) rather than to the entity as a whole that is of concern for preservation purposes.

means of a number of different properties of an entity, the functional orientation of this strategy makes it both less demanding and more meaningful than the alternatives.¹⁰ The range of uses that an entity must satisfy to be considered authentic under this strategy may be anticipated in advance or allowed to evolve over time.¹¹

Authenticity as Suitability for a Purpose

In the context of preservation, authenticity is inherently related to time. A piece of jade may be authentic, irrespective of its origin or provenance; however, a specific preserved jade artifact has additional requirements for being authentic in the historical sense.¹² The alternative strategies and tactics presented above for defining authenticity suggest the range of meanings that may be attributed to the concept, but all of these imply the retention of some essential properties or functional capabilities over time.¹³

Authenticity seems inextricably bound to the notion of suitability for a purpose. A possible exception is the case where originality per se serves as the criterion for authenticity. Such is the case, for example, for venerated artifacts such as the Declaration of Independence. Even if such an entity ultimately becomes unsuitable for its normal purpose (for example, if it becomes unreadable), it continues to serve some purpose—in this example, veneration. In all cases, therefore, authenticity implies some future purpose or use, such as the ability to obtain factual information, prove legal accountability, derive aesthetic appreciation, or support veneration.

While recognizing that it is likely to be a contentious position, I will assume in the remainder of this paper that the authenticity of preserved informational entities in any domain is ultimately bound to their suitability for specific purposes that are of interest within that domain.

At any point in time, it is generally considered preferable to be able to articulate a relatively stable, a priori set of principles for any discipline. For this reason, a posteriori criteria for authenticity may

¹⁰ Although it is tempting to consider the suitability of an informational entity to be constrained only by technical factors, legal, social and economic factors often override technical considerations. For example, the suitability of an informational entity for a given purpose may be facilitated or impeded by factors such as the way it is controlled and made available to potential users. Therefore, if it is to serve as a criterion for authenticity, suitability must be understood to mean the *potential* suitability of an entity for some purpose, i.e., that which can be realized in the absence of arbitrary external constraints.

¹¹ Because the strategy potentially leads to dynamic, evolving definitions of authenticity, it has a decidedly a posteriori flavor, which may be inescapable.

¹² In the remainder of this paper, *authenticity* will be used exclusively in the context of preservation.

¹³ Whereas the originality strategy entails no explicit property or capability conditions (though some tactics for evaluating originality may rely on such conditions), it nevertheless implicitly assumes that simply by virtue of being original, an entity will retain as many of its properties and capabilities as possible.

generate a degree of intellectual anxiety among theoreticians. Some archivists, for example, argue that archival theory specifies a precise, fixed set of suitability requirements for authentically preserved records, namely that future users should be able to understand the roles that the records played in the business processes of the organizations that generated and used them, and that users should be able to continue to use the records in any future business processes that may require them (e.g., for determining past accountability). Similarly, some libraries of deposit may require, to the extent possible, that future users be able to see and use authentically preserved publications exactly as their original audiences did. On the other hand, a data warehouse might require that authentic preservation allow future users to explore implicit relationships in data that the original users were unable to see or define.

In different ways, all these examples attempt to allow for unanticipated future uses of preserved informational entities. They also reveal a tension between the desire to articulate fixed, a priori criteria for authenticity and the need to define criteria that are general enough to satisfy unanticipated future needs. This suggests that we distinguish between a priori suitability criteria, which specify in advance the full range of uses that authentically preserved informational entities must support, and a posteriori suitability criteria, which require such entities to support unanticipated future uses. The a priori approach will work only in a discipline that carefully articulates its preservation mandate and successfully (for all time) proscribes any attempt to expand that mandate retroactively.¹⁴ In contrast, an evolutionary, a posteriori approach to defining suitability criteria should be adopted by disciplines that are less confident of their ability to ward off all future attempts to expand their suitability requirements or those whose preservation mandates are intentionally dynamic and designed to adapt to future user needs and demands as they arise.

Authenticity Principles and Criteria

Because it is so difficult to define authenticity abstractly, it is useful to try to develop *authenticity principles* for various domains or disciplines that will make it possible to define authenticity in functional terms. An authenticity principle encapsulates the overall intent of authentic preservation from a given legal, ethical, historical, artistic, or other perspective—for example, to assess accountability or to recreate the original function, impact, or effect of preserved entities. Ideally, an authenticity principle should be a succinct, functional statement of what constitutes authentic preservation from a specific, stated perspective. Requiring that these principles be stated functionally allows them to be used in verifying whether a given preserva-

¹⁴ Non-retroactive expansion can be accommodated by revising the corresponding suitability criteria for all informational entities to be preserved henceforth.

tion approach satisfies a given principle. For example, one possible archival authenticity principle was proposed above, namely, to enable future users to understand the roles that preserved records played in the business processes of the organizations that generated and used them, and to continue to use those records in future business processes that may require them. Alternative authenticity principles might be proposed for archives as well as for other disciplines. It would be desirable to devise a relatively small number of alternative authenticity principles that collectively capture the perspectives of most disciplines concerned with the preservation of informational entities.

Next, from each authenticity principle, it is useful to derive a set of *authenticity criteria* to serve both as generators for specific preservation requirements and as conceptual and practical tests of the success of specific preservation techniques. For example, to implement the authenticity principle described previously, authenticity criteria would be derived that specify which aspects of records and their context must be preserved to satisfy that principle. These criteria would then provide a basis for developing preservation requirements, such as the need to retain metadata describing provenance, as well as tests of whether and how well alternative preservation techniques satisfy those requirements.

The a priori/a posteriori dichotomy mentioned previously arises again in connection with authenticity principles. From a theoretical perspective, it is more attractive to derive such principles a priori, without the need to consider any future, unanticipated uses to which informational entities may be put. If authenticity principles are derived a posteriori, then they may evolve in unexpected ways as unanticipated uses arise. This situation is unappealing to many disciplines. In either case, if authenticity is logically determined by suitability for some purpose, then an authenticity principle for a given domain will generally be derived, explicitly or implicitly, from the expected range of uses of informational entities within that domain. It may, therefore, be helpful to discuss ways of characterizing such expected ranges of use before returning to the subject of authenticity principles and criteria.

Describing Expected Ranges of Use of Preserved Informational Entities

If expected use is to serve as a basis from which to derive authenticity criteria for a given discipline or organization, then it is important to describe the range of expected uses of informational entities that is relevant to that discipline or organization. This description should consist of a set of premises, constraints, and expectations for how particular kinds of informational entities are likely to be used. It should include the ways in which entities may be initially generated or captured (in digital form, for digital informational entities). It should include the ways in which they may be annotated, amended, revised, organized, and structured into collections or series; pub-

lished or disseminated; managed; and administered. It should describe how the informational entities will be accessed and used, whether by the organization that generates them or by organizations or individuals who wish to use them in the future for informational, historical, legal, cultural, aesthetic, or other purposes. The description should also include any legal mandates or other exogenous requirements for preservation, access, or management throughout the life of the entities, and it should ideally include estimates of the expected relative and absolute frequencies of each type of access, manipulation, and use.¹⁵ Additional aspects of a given range of expected uses may be added as appropriate.

Any attempt to enunciate comprehensive descriptions of ranges of expected uses of this kind for digital informational entities—especially in the near future before much experience with such entities has been accumulated—will necessarily be speculative. In all likelihood, it will be over-constrained in some aspects and under-constrained in others. Yet, it is important to try, however tentative the results, if suitability is to serve as a basis for deriving authenticity criteria.

Deriving Authenticity Principles from Expected Ranges of Use

The purpose of describing an expected range of use for informational entities is to provide a basis from which to derive a specific authenticity principle. Any authenticity principle is an ideal and may not be fully achievable under a particular set of technological and pragmatic constraints. Nevertheless, stating an authenticity principle defines a set of criteria to which any preservation approach must aspire.

Different ranges of expected use may result in different authenticity principles. One extreme is that a given range of expected uses might imply the need for a digital informational entity to retain as much as possible of the function, form, appearance, look, and feel that the entity presented to its author. Such a need might exist, for example, if future researchers wish to evaluate the range of alternatives that were available to the author and, thereby, the degree to which the resulting form of the entity may have been determined by constraint versus choice or chance.

A different range of expected uses might imply the need for a digital informational entity to retain the function, form, appearance,

¹⁵ Future patterns of access for digital records may be quite different from historical or current patterns of access for traditional records, making it difficult to obtain meaningful information of this kind in the near future. Nevertheless, any preservation strategy is likely to depend at least to some extent on assumptions about such access patterns. The library community has performed considerable user research on the design of online public catalogs that may be helpful in this endeavor. For example, see M. Ongerling, *Evaluation of the Dutch national OPAC: the userfriendliness of PC3*, Leiden 1992; *Common approaches to a user interface for CD-ROM—Survey of user reactions to three national bibliographies on CD-ROM*, British Library and The Royal Library, Denmark, Copenhagen, April, 1992; V. Laursen and A. Salomonsen, *National Bibliographies on CD-ROM: Definition of User-dialogues Documentation of Criteria Used*, The Royal Library, Denmark, Copenhagen, March 1991.

look, and feel that it presented to its original intended audience or readership. This would enable future researchers to reconstruct the range of insights or inferences that the original users would have been able to draw from the entity. Whereas retaining all the capabilities that authors would have had in creating a digital informational entity requires preserving the ability to modify and reformat that entity using whatever tools were available at the time, retaining the capabilities of readers merely requires preserving the ability to display, or render, the entity as it would have been seen originally.

Finally, a given range of expected uses may delineate precise and constrained capabilities that future users are to be given in accessing a given set of digital informational entities, regardless of the capabilities that the original authors or readers of those entities may have had. Such delineated capabilities might range from simple extraction of content to more elaborate viewing, rendering, or analysis, without considering the capabilities of original authors or readers. As in the data warehouse example cited previously, it might be important to enable future users to draw new inferences from old data, using tools that may not have been available to the data's original users.

As these examples suggest, it is possible to identify alternative authenticity principles that levy different demands against preservation. For example, the following sequence of decreasingly stringent principles is stated in terms of the relationship between a preserved digital informational entity and its original instantiation:

- same for all intents and purposes
- same functionality and relationships to other informational entities
- same "look and feel"
- same content (for any definition of the term)
- same description¹⁶

An authenticity principle must also specify requirements for the preservation of certain *meta-attributes*, such as authentication and privacy or security. For example, although a signature (whether digital or otherwise) in a record may normally be of no further interest once the record has been accepted into a recordkeeping system—whose custodianship thereafter substitutes its own authentication for that of the original—the original signature in a digital informational entity may on occasion be of historical, cultural, or technical interest, making it worth preserving as part of the "content" of the entity, as opposed to an active aspect of its authentication. Similarly, although the privacy and security capabilities of whatever system is used to preserve an informational entity may be sufficient to ensure the privacy and security of the entity, there may be cases in which the original privacy or security scheme of a digital informational entity may be of interest in its own right. An authenticity principle should determine

¹⁶ This requires preserving only a description of the entity (i.e., metadata). The entity itself can in effect be discarded if this principle is chosen.

a complete, albeit abstract, specification of all such aspects of a digital informational entity that must be preserved.

Since an authenticity principle encapsulates the preservation implications of a range of expected uses, it should always be derived from a specific range of this sort. Simply inventing an authenticity principle, rather than deriving it in this way, is methodologically unsound. The range of expected uses grounds the authenticity principle in reality and allows its derivation to be validated or questioned. Nevertheless, as discussed previously, since the range of expected uses for digital informational entities is speculative, the formal derivation of an authenticity principle may remain problematic for some time.

Different types of digital informational entities that fall under a given authenticity principle (within a given domain of use) may have different specific authenticity criteria. For example, authenticity criteria for databases or compound multimedia entities may differ from those for simple textual entities. Furthermore, digital informational entities may embody various behavioral attributes that may, in some cases, be important to retain. In particular, these entities may exhibit dynamic or interactive behavior that is an essential aspect of their content, they may include active (possibly dynamic) linkages to other entities, and they may possess a distinctive look and feel that affects their interpretation. To preserve such digital entities, specific authenticity criteria must be developed to ensure that the entities retain their original behavior, as well as their appearance, content, structure, and context.

Originality Revisited

As discussed earlier, the authenticity of traditional informational entities is often implicitly identified with ensuring that original entities have been retained. Both the notion of custodianship and the other component concepts of the archival principle of provenance (such as *le respect des fonds* and *le respect de l'ordre intérieur*) focus on the sanctity of the original (Horsman 1994). Although it may not be realistic to retain every aspect of an original entity, the intent is to retain all of its meaningful and relevant aspects.

Beyond the appropriate respect for the original, there is often a deeper fascination, sometimes called a fetish, for the original when it is valued as a historical or quasi-religious artifact. While fetishism may be understandable, its legitimacy as a motivator for preservation seems questionable. Moreover, fetishism notwithstanding, the main motivation for preserving original informational entities is the presumption that an original entity retains the maximum possible degree of authenticity. Though this may at first glance appear to be tautological, the tautology applies only to traditional, physical informational entities.

Retaining an original physical artifact without modifying it in any way would seem almost by definition to imply its authenticity. However, it is generally impossible to guarantee that a physical arti-

fact can be retained without changing in any way (for example, by aging). Therefore, a more realistic statement would be that retaining an original *without modifying it in any way that is meaningful and relevant* (from some appropriate perspective) implies its authenticity. The archival emphasis on custodianship and provenance is at least partly a tactic for ensuring the retention of original records to maximize the likelihood of retaining their meaningful and relevant aspects, thereby ensuring their authenticity. Tautologically, an unmodified original is as authentic as a traditional, physical informational entity can be.

If we consider informational entities as abstractions rather than as physical artifacts, however, this tautology disappears. Although the informational aspects of such an entity may be represented in some particular physical form, they are logically independent of that representation, just as the Pythagorean Formula is independent of any particular physical embodiment or expression of that formula. An informational entity can be thought of as having a number of attributes, some of which are relevant and meaningful from a given perspective and some of which are not. For example, it might be relevant from one perspective that a given document was written on parchment but irrelevant that it was signed in red ink; from a different perspective, it might be relevant that it was signed in red yet irrelevant that it was written on parchment. The specific set of attributes of a given informational entity that is relevant and meaningful from one perspective may be difficult to define precisely. The full range of all such attributes that might be relevant from all possible perspectives may be open-ended. In all cases, however, some set of relevant logical attributes must exist, whether or not we can list them.

This implies that retaining the original physical artifact that represents an informational entity is at most sufficient (in the case of a traditional informational entity) but is never logically necessary to ensure its authenticity. If the relevant and meaningful attributes of the entity were retained independently of its original physical embodiment, they would by definition serve the same purpose as the original. Furthermore, since it is impossible to retain all attributes of a physical artifact in the real world because of aging, retaining the original physical artifact for an informational entity may not be sufficient, since it may lose attributes that are relevant and meaningful for a given purpose. (For example, the color of a signature may fade beyond recognition.) Retaining an original physical artifact is therefore neither necessary nor sufficient to ensure the authenticity of an informational entity.

Digital Informational Entities and the Concept of an Original

The preceding argument applies a fortiori to digital informational entities. It is well accepted that the physical storage media that hold digital entities have regrettably short lifetimes, especially when obso-

lescence is taken into account. Preserving these physical storage media as a way of retaining the informational entities they hold is not a viable option. Rather, it is almost universally acknowledged that meaningful retention of such entities requires copying them onto new media as old media become physically unreadable or otherwise inaccessible.

Fortunately, the nature of digital information makes this process far less problematic than it would be for traditional informational entities. For one thing, digital information is completely characterized by simple sequences of symbols (zero and one bits in the common, binary case). All of the information in a digital informational entity lies in its bit stream (if, as argued earlier, this is taken to include all necessary context, interpreter software, etc.). Although this bit stream may be stored on many different kinds of recording media, the digital entity itself is independent of the medium on which it is stored. One of the most fundamental aspects of digital entities is that they can be stored in program memory; on a removable disk, hard disk, CD-ROM; or on any future storage medium that preserves bit streams, without affecting the entities themselves.¹⁷

One unique aspect of digital information is that it can be copied perfectly and that the perfection of a copy can be verified without human effort or intervention. This means that, at least in principle, copying digital informational entities to new media can be relied upon to result in no loss of information. (In practice, perfection cannot be guaranteed, but increasingly strong assurances of perfection can be attained at relatively affordable cost.)

The combination of these two facts—that digital informational entities consist entirely of bit streams, and that bit streams can be copied perfectly onto new media—makes such entities logically independent of the physical media on which they happen to be stored. This is fortunate since, as pointed out above, it is not feasible to save the original physical storage artifact (e.g., disk and tape) that contains a digital informational entity.

The deeper implication of the logical independence of digital informational entities from the media on which they are stored is that it is meaningless to speak of an original digital entity as if it were a

¹⁷ In some cases, the original storage medium used to hold a digital informational entity may have some significance, just as it may be significant that a traditional document was written on parchment rather than paper. For example, the fact that a digital entity was published on CD-ROM might imply that it was intended to be widely distributed (although the increasing use of CD-ROM as a back-up medium serves as an example of the need for caution in drawing such conclusions from purely technical aspects of digital entities, such as how they are stored). However, even in the rare cases where such physical attributes of digital informational entities are in fact meaningful, that meaning can be captured by metadata. Operational implications of storage media—for example, whether an informational entity would have been randomly and quickly accessible, unalterable, or constrained in various ways, such as by the size or physical format of storage volumes—are similarly best captured by metadata to eliminate dependence on the arcane properties of these quickly obsolescent media. To the extent that operational attributes such as speed of access may have constrained the original functional behavior of a digital informational entity that was stored on a particular medium, these attributes may be relevant to preservation.

unique, identifiable thing. A digital document may be drafted in program memory and saved simultaneously on a variety of storage media during its creation. The finished document may be represented by multiple, identical, equivalent copies, no one of which is any more “original” than any other. Furthermore, copying a digital entity may produce multiple instances of the entity that are logically indistinguishable from each other.¹⁸

Defining Digital-Original Informational Entities

It is meaningless to rely on physical properties of storage media as a basis for distinguishing original digital informational entities. It is likewise meaningless to speak of *an* original digital entity as a single, unique thing. Nevertheless, the concept of an “original” is so pervasive in our culture and jurisprudence that it seems worth trying to salvage some vestige of its traditional meaning. It appears that the true significance (in the preservation context) of an original traditional informational entity is that it has the maximum possible likelihood of retaining all meaningful and relevant aspects of the entity, thereby ensuring its authenticity. By analogy, we therefore define a *digital-original* as any representation of a digital informational entity that has the maximum possible likelihood of retaining all meaningful and relevant aspects of the entity.

This definition does not imply a single, unique digital-original for a given digital informational entity. All equivalent digital representations that share the defining property of having the maximum likelihood of retaining all meaningful and relevant aspects of the entity can equally be considered digital-originals of that entity. This lack of uniqueness implies that a digital-original of a given entity (not just a copy) may occur in multiple collections and contexts. This appears to be an inescapable aspect of digital informational entities and is analogous to the traditional case of a book that is an instance of a given edition: it is *an* original but not *the* original, since no single, unique original exists.¹⁹

It is tempting to try to eliminate the uncertainty implied by the phrase *maximum possible likelihood*, but it is not easy to do so. This uncertainty has two distinct dimensions. First, it is difficult enough to specify precisely which aspects of a particular informational entity are meaningful and relevant for a given purpose, let alone which aspects of any such entity might be meaningful and relevant for any possible purpose. Since we cannot in general enumerate the set of such meaningful, relevant aspects of an informational entity, we cannot guarantee, or even evaluate, their retention. Second, physical and logical constraints may make it impossible to guarantee that any dig-

¹⁸ Even time stamps that purportedly indicate which copy was written first may be an arbitrary result of file synchronization processes, network or device delays, or similar phenomena that have no semantic significance.

¹⁹ Moreover, since there is no digital equivalent to a traditional manuscript, there can be no unique prepublication version of a digital informational entity.

ital-original will be able to retain *all* such aspects, any more than we can guarantee that a physical original will retain all relevant aspects of a traditional informational entity as it ages and wears. The uncertainty in our definition of digital-original therefore seems irreducible; however, its impact is no more damaging than the corresponding uncertainty for physical originals of traditional informational entities.

Although the definition used here does not imply any particular technical approach, the concept appears to have at least one possible implementation, based on emulation (Michelson and Rothenberg 1992; Rothenberg 1995, 1999; Erlandsson 1996; Swade 1998). In any case, any implementation of this approach must ensure that the interpreters of digital informational entities, themselves saved as bit streams, can be made perpetually executable. If this can be achieved, it should enable us to preserve digital-original informational entities that maintain their authenticity across all disciplines, by retaining as many of their attributes as possible.

Conclusion

If a single, uniform technological approach can be devised that authentically preserves all digital-informational entities for the purposes of all disciplines, the resulting economies of scale will yield tremendous benefits. To pave the way for this possibility, I have proposed a foundation for a universal, transdisciplinary concept of authenticity based on the notion of suitability. This foundation allows the specific uses that an entity must fulfill to be considered authentic to vary across disciplines; however, it also provides a common vocabulary for expressing authenticity principles and criteria, as well as a common basis for evaluating the success of any preservation approach.

I have also tried to show that many alternative strategies for determining authenticity ultimately rely on the preservation of relevant, meaningful aspects or attributes of informational entities. By creating digital-original informational entities that have the maximum possible likelihood of retaining all such attributes, we should be able to develop a single preservation strategy that will work across the full spectrum of disciplines, regardless of their individual definitions of authenticity.

REFERENCES

Erlandsson, A. 1996. *Electronic Records Management: A Literature Review*. International Council on Archives' (ICA) Study. Available from <http://www.archives.ca/ica>.

Horsman, P. 1994. Taming the Elephant: An Orthodox Approach to the Principle of Provenance. In *The Principle of Provenance*, edited by Kerstin Abukhanfusa and Jan Sydbeck. Stockholm: Swedish National Archives.

Michelson, A., and J. Rothenberg. 1992. Scholarly Communication and Information Technology: Exploring the Impact of Changes in the Research Process on Archives. *American Archivist* 55(2):236-315.

Rothenberg, J. 1995. Ensuring the Longevity of Digital Documents. *Scientific American*, 272(1):42-7 (international edition, pp. 24-9).

_____. 1999. *Avoiding Technological Quicksand: Finding a Viable Technical Foundation for Digital Preservation: A Report to the Council on Library and Information Resources*. Washington, D.C.: Council on Library and Information Resources. Available from <http://www.clir.org/pubs/reports/rothenberg/pub77.pdf>.

Rothenberg, J., and T. Bikson. 1999. *Carrying Authentic, Understandable and Usable Digital Records Through Time*. RAND-Europe. Available from <http://www.archief.nl/digiduur/final-report.4.pdf>.

Swade, D. 1998. Preserving Software in an Object-Centred Culture. In *History and Electronic Artefacts*, edited by Edward Higgs. Oxford: Clarendon Press.

Authenticity in Perspective

by Abby Smith

This is neither a commentary on the preceding papers nor a summary of the discussions held on January 24. Rather, I will try to give a sense of various views expressed by the presenters, identify the issues raised in light of the subsequent discussions, and highlight the implications of the day's proceedings.

Some Ground Rules

In his seminal work, *Principia Ethica*, the moral philosopher and epistemologist G. E. Moore remarked that, in most complex matters, difficulties and disagreements “are mainly due to a very simple cause: namely to attempt to answer questions, without first discovering precisely *what* question it is which you desire to answer” ([1902] 1988). Conference participants clearly agreed about the question: What is an authentic digital object, and what are the core elements that, if missing, would render that object something other than what it purports to be? The difficulties arose from participants’ legitimate, and perhaps predictable, disagreements about which elements are intrinsic to a digital object and which elements are contingent on context, technologies, encoding schemes and display methods, or other externalities.

As anticipated, communities differ in their understanding of what constitutes intrinsic features of a digital object; these differences mirror their understanding of authenticity of analog objects. After all, the uses of digital and analog information by historians, archivists, publishers, or scientists vary greatly. Most of the workshop participants grounded their thinking about digital objects and their identity in the fitness of these objects for some specified function or purpose, such as a record that bears evidence; a historical source that bears witness to an event, a time, or a life; or data that could produce

a replicable experiment. In other words, what was deemed intrinsic to an object was determined by the purpose for which it was created (or, in the case of archival records, the most narrowly defined of digital objects under discussion, the purpose of bearing evidence about an object's creation and intended use). Regrettably, as Moore pointed out, evidence cannot be adduced for things intrinsic. "From no truth, except themselves alone, can it be inferred that [intrinsic things] are either true or false."

The Key Issues

Perhaps for that reason if no other, neither the presenters nor the workshop participants addressed systematically and directly the question of what an authentic digital object is and what the core attributes are that, if missing, would render the object something other than what it purports to be. However, threaded throughout the discussion were various responses to the other questions raised in the charge.

- *If all information—textual, numeric, audio, and visual—exists as a bit stream, what does that imply for the concept of format and its role as an attribute essential to the object?*

Clifford Lynch proposed a hierarchy of complexity of representation: bit stream, data, documents, interactive objects (i.e., engaging sensory perceptions), and experiential works (e.g., virtual reality). That schema resonated with many of the participants. David Levy pointed out that we might never resolve the paradox of bits being "the stuff" with the fact that bits are inaccessible to our senses and perceptual abilities. This is not how we have dealt with recorded information before. Given that we are just beginning to explore the relationship between humans and computing machines, it is hard to think ahead about how we, as physical creatures, will relate to virtual bits.

In the analog realm, many features of recorded information are an aspect of the object itself and so will not translate into the digital environment. We are generally unaware of how often we use our judgment about the physical integrity of recorded information to *stand in for* a judgment about the integrity of the text. We make instantaneous inferences about a text that we receive with portions blacked out. Evidence provided by the physical object, however, has no counterpart in the digital world; deletions in an electronic text would not be visible to the eye and, consequently, would not raise our suspicion. A digital object has no independent physical manifestation that can accrete information about its fate in this world (such as bookplates, marginalia, coffee mug stains, and so forth). For similarly effective external evidence for a digital object, we must create such things as metadata, which in turn create their own preservation and readability concerns. Once metadata are separated from their object, it is hard to reattach them. It is a fact of life on the Internet that in e-mail correspondence, content will be cut and pasted into

some extraneous document and then widely disseminated without its originating contextual metadata. Cutting and pasting in the analog world, by contrast, leave physical traces that can alert the recipient or reader to the document's provenance.

- *Does the concept of an original have meaning in the digital environment?*

David Levy defined the copying of digital information as a manufacturing process. In effect, a digital file is like a printing plate. Bits may be the source of a document, but they are not and can never be the original. Moreover, there are no unique copies in the digital realm unless they result from a mistake in the manufacturing process, that is, the process of copying the file onto the screen. Jeff Rothenberg argued strongly for the opposite point of view, saying that a *digital-original* is any representation of a digital informational entity that has the maximum possible likelihood of retaining all meaningful and relevant aspects of the entity. This echoed the archival point of view, which suggests that the digital-original is just the same (i.e., works in just the same way) as the analog original. The value of an original is that it is as complete as possible and it is reliable because of the control exercised in its creation. In an archives, the digital-original is simply the first record received.

But this begs the question of what we really mean by "original." In the case of a digital file, we are referring not to an object per se, but to a fixed set of properties that contain information about the digital object and that constitute the digital object itself. Again, this would not make the original unique. One of the difficulties in talking about the issue of "original" is that there is no object fixity in the digital world, as there is in the analog world. As Clifford Lynch helpfully explained, in the analog world, I give you the object and now you have it and I do not. In the digital world, I share with you a file that has the same properties as the file I have—the original, as it were. Now I have it, and you have it, too. But what, precisely, is the "it," the file? It could be characterized as a "fixed set of properties."

All workshop attendees agreed that digital technology obviates the idea of a unique item, because the very act of viewing, say, a digital photograph means creating a copy (on screen). This fact has obvious implications for copyright.

- *What role does provenance play in establishing the authenticity of a digital object?*

The role of provenance is as important in the digital world as in the analog world, if not more so. For the archivists, the role of provenance is well defined. Archives can provide evidence of authenticity by documenting the chain of transmission and custody, and they have internal controls that reduce to an acceptable level the risk of tampering. Within the controlled environment of an archives, the provenance of records is theoretically secure. (Whatever happened to the item before it came to the archives, and whatever happens to it when it leaves, may be another matter altogether.) Archives, of course, deal with limited types of items. They are records—things

created in the order of doing business. The truth value of a record is not what makes it authentic. A record might contain false information but still be authentic as a record.

In the larger context of libraries and beyond, the role of provenance is far more complicated. Archives can serve as a trusted third party only in a relatively controlled environment. Whenever information crosses administrative and technological boundaries, as it does in the more permeable world of publishers and libraries, the role of trusted third parties, while critical for authenticity, is harder to develop and maintain. The partnership between libraries and publishers, a crucial link in the ultimate relationship between author and reader, has evolved slowly, at times painfully, over centuries, and will continue to evolve. Nonetheless, the digital environment will still need trusted third parties to store material, and the libraries and publishers will need to agree on protocols for digital publishing and preservation that work as effectively as have those of the past.

Interestingly, the scholar-participants suggested that technological solutions to the problem will probably emerge that would obviate the need for trusted third parties. Such solutions may include, for example, embedding texts, documents, images, and the like with various warrants (e.g., time stamps, encryption, digital signatures, and watermarks). The technologists replied with skepticism, saying that there is no technological solution that does not itself involve the transfer of trust to a third party. Encryption—for example, public key infrastructure (PKI)—and digital signatures are simply means of transferring risk to a trusted third party. Those technological solutions are as weak or as strong as the trusted third party. To devise technical solutions to what is, in their view, essentially a social challenge is to engender an “arms race” among hackers and their police.

- *What implications for authenticity, if any, are there in the fact that digital objects are contingent on software, hardware, network, and other dependencies?*

Dependencies mean either nothing or everything. What if you have a digital object that you cannot read because you do not have the right software? Jeff Rothenberg argued that you cannot know something is authentic unless you can read it. However, to the archivists, this constituted confusion between meaning and authenticity. In their opinion, you do not need to view the contents of something to say that it is an authentic record. Take, for example, the case of the Rosetta Stone. For centuries, the meaning of the stone was beyond reach, because no one could decipher the codes in which it was written. It was, nonetheless, an authentic record of the time in which it was created. You could not say that it was inauthentic in the eighteenth century and that it became authentic only when—and because—Champollion decoded it.

But publishers, historians, and computer scientists were quick to point out that fixity of a text, much as we take it for granted, is a relatively recent phenomenon. It was the printing press that helped to create the notion of a fixed text. There was little or no fixity of text

before printing, and none exists in unpublished materials such as hand-developed photographs or manuscripts. When a publisher goes to press with an error, he or she feels an obligation to publish an errata sheet. In manuscripts, however, there are no errata sheets; likewise, there need be no such sheets in the digital world. The publisher of a digital resource can simply go in and correct the text. Whether such a publisher chooses to note that change or not is related to his or her sense of obligation to the publication record, not to the truth of the text.

The variability of digital formats is great and will continue to be so. It should not necessarily pose problems to matters of authenticity, depending on how one defines the “fixed set of properties” that constitute the file. After all, difference in display monitors can significantly alter the way things appear to us, even though they display the exact same bit stream. Is the bit stream displayed on my monitor the same document as the one you have, if the bit stream is identical but the appearance it generates in your monitor is different?

Proposed Answers

David Levy proposed that, for purposes of proving something is authentic, we could use the following three methods—all implying a trusted third party for implementation—that answer the question of authenticity by stipulating in reference to what something is authentic.

1. Use of reference object (Does the object match this object?)
2. Metadata (Does the object match this description of an object?)
3. Digital recipe (Could we recreate or reassemble an authentic object using this set of instructions?)

Implications for Preservation

Authenticity, although seldom talked about, is deeply implicated in even the routine decisions we make about preservation. Fortunately, issues of authenticity have seldom been problematical in the print regime, at least for the past century or so. Even the new recording media for audio and video present fewer authenticity issues for professionals than they used to.

In the analog regime, one could not reasonably say that issues of preservation are deeply implicated in authenticity. Any investigation into authenticity per se might, therefore, include preservation, but should not be subsumed by it. It is not clear that this is the case with digital objects. While future discussions of authenticity should be careful to investigate all aspects of the authenticity issue without prejudice, there are nonetheless certain nagging facts about how preservation operates in the digital realm that warrant consideration.

We have known for some time that all of our operating assumptions about selection for preservation are turned on their head by digital technology. Preservation has operated by making choices

about what objects from the past should endure into the future. For collecting institutions, be they libraries, archives, museums, or historical societies, the commitment to preserve is made at the time of acquisition. However, preservation actions—e.g., rehousing into acid-free folders or stabilizing a fragile book—are sometimes taken years after accessioning and, to our shame, years after the physical condition warranted intervention. But there are also materials—those collected “just in time” rather than “just in case”—that may not carry an implied commitment to preserve when they are acquired because someone else has preserved them. In those cases, the collecting institution can make a decision about how valuable the item is likely to be in the future, when an immediate demand for the item no longer exists.

No matter how the preservation selection and action occur, in the analog realm we choose what to preserve well after items have been created, authenticated, and valorized through publishing or, in the case of archives, during appraisal. The item has gone through several processes in which it is selected—from the publisher to the acquisitions specialist to the curator and preservation specialist.

In the digital world, however, the act of selecting for preservation has become a process of constant reselection. We have to intervene continually to keep digital files alive. We cannot put a digital file on a shelf and decide later about preservation intervention. Storage means active intervention. One must refresh data regularly to keep it alive. It is as if suddenly every item in a library—every single book, manuscript leaf, and page of newsprint—demanded preservation action every 18 or 24 months. We do not lose books just because we do not use them, but it is possible to lose digital data just because no one wants access to it within a year or two of its creation. Indeed, many are saying that the preservation of digital data should begin at time of creation. The creator should make all decisions about file format, software and hardware, and even complexity of documentation, in light of the intended longevity of the object. This need to think prospectively about persistence introduces a strong element of intentionality among all actors in the drama of information creation, dissemination, and consumption that has implications for the meaning of authenticity. There is an accidental nature to the evidence borne by physical artifacts that serves to strengthen an item’s claim to authenticity. In one sense of the word, commonly used in scientific laboratories, “artifact” connotes the unintended byproduct of a process, a byproduct usually irrelevant to the outcome. Similarly, among the various physical media on which information is recorded there are byproducts of the recording or printing or manufacturing process that give vital clues to the authenticity of the object precisely because those byproducts were not intended. This is what we lose in the digital environment.

Authenticity in the Perspective of the Future

Fortunately, there are limited circumstances in which authenticity of information is critically important: biomedical data, legal documents,

national security intelligence, proprietary trade and commercial information, and public records. In those cases, information is usually created and managed in controlled environments to reduce the risk of intentional or inadvertent corruption to an acceptable level. But there is much information in digital form that we rely on to be what it appears to be; for example, the historical documents we find on library Web sites, the e-mail messages we receive from colleagues in the course of doing business, photographs that we take on digital cameras, or the online news sources we check for stock quotes. We do not want to live in a world of constant suspicion that what we see is not what actually is. What we, as creators and users of information, need to do is to become digitally literate and to understand better how our machines fulfill the commands we send to them. Specific communities, such as scholars, scientists, and journalists, must decide what information they need to place high trust in and to develop protocols for ensuring the integrity of that information. This means creating appropriate documentation, following standard procedures that leave a transparent trail, and respecting those documents above others. The truth value of most information will always be a matter about which the user must make judgments. These judgments are not guaranteed in the print regime, nor will they be in the digital.

Looking ahead, we can reasonably expect that some digital objects will warrant greater skepticism than their analog counterparts. It took centuries for users of print materials to develop the web of trust that now undergirds our current system of publication, dissemination, and preservation. Publishers, libraries, and readers each have their own responsibilities to keep the filaments of that web strong. Making the transition to a trusted digital environment will require much conscious reexamination of what we take for granted in the print and audiovisual media on which we rely. We can begin by learning more about this new medium of digital information and by clarifying the terms we borrow from the physical world of analog materials to describe the new phenomena of virtual objects.

REFERENCE

Moore, G. E. [1902] 1988. *Principia Ethica*. Reprint. Amherst, N.Y.: Prometheus Books.

APPENDIX I

Conference Participants

Charles T. Cullen

President and Librarian
Newberry Library

Luciana Duranti

Professor
University of British Columbia

Elizabeth Eisenstein

Professor Emeritus
University of Michigan

Peter Givler

Executive Director
Association of American University Presses, Inc.

Rebecca Graham

Research Associate
Digital Library Federation

Peter B. Hirtle

Co-Director
Cornell Institute for Digital Collections

Karen Hunter

Senior Vice President
Elsevier Science

Andrea La Vere

Head of Archive and Asset Management
DreamWorks SKG

David Levy

Consultant

Clifford Lynch

Executive Director
Coalition for Networked Information

Deanna Marcum

President
Council on Library and Information Resources

Jeff Rothenberg

Senior Computer Scientist
The RAND Corporation

Elaine Sloan

Vice President for Information Services and
University Librarian
Columbia University

Abby Smith

Director of Programs
Council on Library and Information Resources

Winston Tabb

Associate Librarian for Library Services
Library of Congress

Ken Thibodeau

Director, Center for Electronic Records
National Archives

Donald Waters

Program Officer, Scholarly Communication
The Andrew W. Mellon Foundation