The Information Age

BRIAN HAYES

The Electronic Palimpsest

Digital documents for all occasions: erasable, correctable, reproducible, forgeable

AS A WRITING INSTRUMENT, THE computer is not so much a better pencil as a better eraser. Although it serves well enough to put words on the page, where it really excels is in wiping them out again. Writing with a computer affords you the luxury of changing your mind, again and again, without penalty. The excised word leaves no scar; the page never becomes gray or tattered from rubbing; the margins do not fill up with afterthoughts; there is no tangle of arrows showing how sentences are to be rearranged. When you write on the glass screen, the world need never know how you labored to achieve that easygoing prose style. Indeed, this very paragraph conceals the tortured history of its own composition: you the reader cannot see in the space between the lines how I have revised it, a dozen times or more, until hardly a word of the first draft survives.

When I got my first chance to write with a computer, it was an exhilarating experience. I would insert a word into the middle of a paragraph and marvel as all the following words automatically rearranged themselves to make room, cascading from line to line in a kind of domino effect. Or I would hold down the delete key and suck up whole sentences like spaghetti. Suddenly prose became a kind of clay that never hardens, a medium that one can always reshape yet again.

But if the plasticity of electronic text is a great liberation for the author, it can also license the forger, the plagiarist, the swindler, the impostor; and it is not an unmixed blessing for the scholar, the historian or even the ordinary reader. Words stored in electronic form are in certain ways less secure and less permanent than words on paper. When writing is inscribed in the magnetic domains of a spinning disk, can one trust its integrity? Fifty years from now, will anyone even be able to read it? As more of the world's documents migrate from memo pads and filing cabinets and bookshelves into computer memories, those questions are going to take on considerable importance.

ONE WAY OF EXPLORING THE ISSUES IS to imagine a world without paper, in which all documents are electronic. Such a world is not far off. True, the "paperless office" has so far turned out to be a bad joke, as office paper consumption doubles every four years. But all those pages spewing out of all those laser printers are coming from computers. In many cases the computer files are already the primary versions of the documents, and the printouts are just a means of distribution or archival storage. In the long run, paper will surely be supplanted in those roles as well.

The first thing you notice in a paperless world is that certain awkward situations become even more awkward:

You receive a letter (in the form of a computer file) in which your long-lost sister claims she is being held in a Turkish prison for crimes she didn't commit. Please send her $20,000 to bribe the prosecutor. How do you establish that the letter was written by your sister?

A Washington friend asks you to take a discreet look around someone else's office late one night. In case you get into any trouble, he gives you a letter, stored on a computer disk, that explains the importance of your work to the nation's security. When you are charged with burglary, however, the friend disclaims all knowledge of the letter—and of you. How can you prove the letter is not a forgery? Note that this task is harder than the first one. With the letter from your sister, you need only convince *yourself* of its authenticity; to avoid jail—or at least to take your

friend with you—you must convince a judge and a jury of the letter's provenance.

You pick up a hitchhiker in the desert, and in gratitude for that small kindness he gives you a floppy disk bearing a promise of millions of dollars upon his death. How do you prove the bequest is from Howard Hughes? Again, you must demonstrate to others that the document is genuine. Furthermore, you may well have to show not only that Hughes wrote the note but also that you have not altered it (changing "two million" to "two hundred million," say).

In the world of paper documents the primary tool for settling such controversies is the examination of handwriting. You know the letter is from your sister because you recognize her hand; experts compare your cover letter from the White House or your note from Howard Hughes with specimens known to be authentic. But the bits and bytes of a computerized document are all alike, with none of the idiosyncrasies that might identify individual authorship. Anyone could have typed those letters, on almost any computer.

The introduction of "pen"-based computers, which substitute a stylus for a keyboard, will not solve the problem. On such a machine one might confect an ornate and quite inimitable signature, ending with the most swashbuckling paraph, but a document signed in that way offers only a weak warrant of authenticity. The reason is that such a digitized signature—or any other graphic object—can be copied in an instant with the help of a computer. Give me one "signed" electronic document, and I can forge your name to anything I please. As a matter of fact, the widespread availability of high-resolution scanning and printing equipment raises questions about the security of signatures on paper. There is nothing the modern forger might need that can't be found at the local Kinko's.

DIGITAL DOCUMENTS *CAN* BE SIGNED, however; what is needed is not a digitized signature but a truly digital one. A technique for creating such signatures was proposed in 1976 by Whitfield Diffie and Martin E. Hellman, both then at Stanford University, as part of their ingenious public-key cryptosystem; the idea was refined a few years later by Ronald L. Rivest, Adi Shamir and Leonard Adleman (a triumvirate known as RSA), all then at MIT. In the RSA cryptosystem each user has two keys, one of which is made public and the other is held in secret. A message encrypted with the public key can be decrypted with the private key, and vice versa.

When the system is used for secrecy, a message is encrypted with the recipient's public key (which anyone can look up in a directory); then only the recipient can decrypt the message with the corresponding private key. A simple variation on the protocol yields highly secure digital signatures. To sign a document, you encrypt a copy of it with your private key. Anyone can then verify the signature by decrypting it with your public key. The mechanism works like Cinderella's slipper: whoever owns the key that fits must be the author of the document.

A further refinement has since been added to the digital-signature protocol. If you encrypt an entire document in order to sign it, the signature is as large as the document itself. To reduce the bulk, and at the same time avoid a subtle weakness lurking in the original scheme, the document is collapsed to a "digest" of just 160 bits, and then only the digest is signed. The digesting is done in such a way that even the slightest change to the document is almost certain to yield an entirely different digest. The recipient verifies the signature by applying the digesting algorithm to the document and decrypting the digest with the sender's public key; the result should match the supplied signature.

The National Institute of Standards and Technology is at work on a digital-signature standard based on the public-key principle. The standard-setting process has been going on for years, buffeted by much controversy, but it now appears to be nearing a conclusion.

THE DECRYPTING *mechanism works like Cinderella's slipper: whoever owns the key that fits must be the author of the document.*

Digital signatures would probably deal quite well with the three situations described above. If signatures accompanying the letter from the Turkish prison and the note from Howard Hughes could be decrypted with the appropriate public keys, that would count as strong evidence for the documents' authenticity. Similarly, your Washington friend would have a hard time disowning a letter that had been signed by means of encryption with his own private key. The signatures also protect against after-the-fact tampering. There is no way you could have exaggerated Hughes's generosity in the signed bequest without knowing his private key.

ALTHOUGH DIGITAL SIGNATURES ARE dashed clever, they fall short of solving all the problems of the paperless society.

It might seem at first that digital signatures would provide all the security necessary for a system of electronic checks. To pay your rent, you would merely type a note on the computer, or fill in a template, stating the date, the amount and the payee and identifying your bank and your account number; then you would sign the check with your private key and send it off by electronic mail; the recipient would sign it as an endorsement and mail it on to his own bank. You would be confident that your unscrupulous landlord could not alter the amount, because the bank would detect the change when it verified your signature. Unfortunately, you would remain vulnerable to a cruder kind of fraud. The computer is not only a good eraser but also a flawless copier, and your landlord could simply duplicate your check (along with its signature) and deposit multiple copies, all of which would appear equally authentic. Rivest, Shamir and Adleman suggest including a unique serial number in each signed check and requiring banks to accept only one check with a given serial number. But that puts the onus of vigilance and record keeping on the banks, which may be reluctant to accept it.

Another problem arises when readers must verify not only the authorship and the integrity of a document but also its time of composition. In your computerized laboratory notebook you record the discovery of a new comet or a new virus, and you apply a digital signature to the entry. Later, a rival challenges your claim to priority. Naturally you included the date in your signed notebook entry, but your opponent is not much impressed by that evidence. He points out that the digital signature prevents others from tampering with the document, but since you know your own secret key, you could alter the date—or alter other parts of the document—at any time, then resign it. In other words, the signature proves that you wrote the notebook entry, but it cannot establish when you wrote it.

With paper documents there are at least two ways of dealing with the problem. First, important documents are witnessed as well as signed, and the witnesses can later be called to attest to the dating of the material. That practice can be adopted just as easily with digital signatures. Second, laboratory notes are generally kept in a bound volume with numbered pages, so that sheets cannot be inserted or removed. When disputes arise, the notebook will be credible evidence only if it can be read as a complete, continuous and contemporaneous record of laboratory activity. The pages must be filled up in sequence, without leaving gaps where back-dated entries might be inserted. In the recent controversy over the work of Thereza Imanishi-Kari of Tufts University, the U.S. Secret Service was asked to examine certain notebooks in an attempt to verify the chronology of the entries.

A LABORATORY RECORD KEPT ON A computer is more like a loose-leaf notebook than like a bound volume. New entries can be inserted anywhere in the se-

quence, or existing entries can be moved around; dates can be misstated or changed after the fact. But a solution is at hand. Inspired by the Imanishi-Kari case, Stuart A. Haber and W. Scott Stornetta of Bell Communications Research (Bellcore), in Piscataway, New Jersey, have devised a time-stamping service for electronic documents. The scheme is conceptually similar to public-key cryptography. When your notebook needs to be validated, you submit a digest to the time-stamping computer, which returns a "certificate" that encodes the time of receipt and other information.

But what if the time-stamping service itself cannot be trusted? For example, someone might tamper with the time-stamping computer, perhaps resetting the clock for long enough to create a fraudulent certificate, then restoring the clock to the correct current value. As protection against such deceptions, each certificate is combined with others issued at about the same time in a treelike structure; the single certificate at the root of the tree, whose value depends on all the individual certificates, is publicly posted. During preliminary trials of the service, documents are being time-stamped in weekly batches, and the root certificates are being published every Sunday in the Public Notices section of *The New York Times*. A certificate for any document stamped in the past week can be verified by rederiving the published root value. (A question that remains for the future is what will happen when *The New York Times* is published electronically instead of on paper.)

Even documents that no one would dream of having time-stamped or witnessed sometimes come under scrutiny. For example, George Bush has made public some of his private diaries in an attempt to establish what he didn't know (and when he didn't know it) about the sale of arms to Iran during his vice-presidency. Suppose Bush had kept his diary on a computer instead of on paper: he would have had great difficulty convincing his critics that no entries had been altered, deleted or back-dated.

Of course another major cache of documentary evidence in the Iran-Contra affair was in electronic form: the electronic-mail messages of Oliver North. Curiously, those messages were accepted as authentic and unaltered precisely because North had deleted them (or rather had tried to delete them) from the disk memory of his computer. They were recovered through a sector-by-sector examination of the contents of the disk. If North had instead copied all the files onto floppy disks and voluntarily handed them over to congressional investigators, the messages would surely have been viewed with greater skepticism. (The hard-core conspiracy theorist knows that the supposed deletion and subsequent recovery of the messages was all a carefully staged means of increasing the credibility of concocted evidence. That some of the recovered messages were incriminating counts for nothing, apart from demonstrating that the real messages must have been much worse.)

THE HANDLING OF LEGAL DOCUMENTS is certainly not the only domain in which a conversion to electronic storage and transmission will change the nature of writing. Even personal correspondence is affected. For example, consider the art of the deft postscript. At the end of a chatty letter home, below the signature, you add, "P.S. I've just heard from Stockholm. Good news." Now, it may be that word of your Nobel prize reached you in the moments after the letter was finished but before it was sealed, but it's also possible that you turned the announcement into a casual afterthought purely for rhetorical effect. With a letter on paper, the recipient could never be quite sure. But with an electronic letter, "P.S." is almost certainly an artifice. After all, with a word processor it is no more trouble to add a sentence at the beginning than at the end.

Other rhetorical devices also lose a bit of their impact. In a letter on paper you might write, "Say hi to ~~dreary~~ dear old Dad," where the strike-through is very much a part of the joke. With a computer, since any mistake can be silently and invisibly corrected, the same trick seems more contrived, less spontaneous.

When a manuscript is being prepared for publication, the kind of invisible mending made possible by computers is often a handicap. Traditionally an editor would return to the author a marked-up copy of the original manuscript, showing all the proposed changes and corrections. When the editing is done with a computer, that record of alterations generally disappears. In fact, software solutions are available for that problem; they are just not widely used. Many word-processing programs offer a "red-lining" mode, which displays insertions and deletions explicitly (though seldom as clearly as they can be with a red pencil). There are also special-purpose programs for annotating text, and the contributions of multiple editors are identified by color.

SUCH TOOLS MAY CAPTURE AN EDITOR'S changes, but what about the author's transformations of the work during its composition? Few writers have the patience to document every stage in the creation of a novel or a poem (much less a love letter or a business memorandum). Indeed, some authors would cite as an advantage of computerized writing the end of old drafts and scribbled notes; all that remains of those scraps is now seamlessly integrated into the final text. From the scholar's point of view, however, a valuable source of information is being lost.

Take William Wordsworth's long autobiographical poem *The Prelude*. Fragments of the poem are known from as early as 1798; several versions were composed between 1799 and 1804; Wordsworth made sporadic revisions until 1839; various further emendations were introduced by others before a new edition was published soon after the poet's death in 1850. Dozens of manuscript sources survive, and they have enabled scholars to reconstruct the poem's compositional history in detail. There is no consensus among modern readers that the final state of the poem is the best; indeed, the 1805 version has many partisans. Yet

if Wordsworth had had a PC, the history of the poem would probably be lost.

Lord Byron is another intriguing case. He presented himself to the world as an aristocrat of letters whose verses were casual, offhand productions, which he would not deign to correct or revise. Had he been writing with a computer (I imagine him toting the latest laptop model across the Alps into Italy), he might have gotten away with that fib. But recently published facsimiles of his manuscripts show just how labored his process of composition was. A reviewer describes the manuscripts as "bristling with added stanzas, overwritten crosswise, with false starts, impatient deletions, emendations, and adjustments of rhymes." A modern Byron can readily conceal all signs of such unseemly labor, and as a result readers of the next century will likely find that manuscript sources for authors of the 1990s are rather scanty. Information that might well have been preserved on paper is being lost on disk.

The loss is not inevitable, and the cause is not really technological. As a matter of fact, keeping a complete archive of a life's work is surely easier with a computer than it is with a filing cabinet. One approach—one of many—is the WORM drive: a disk memory that can be written on and read from but never erased. (WORM stands for "write once, read many.") WORM drives have ample storage capacity; a single disk would hold all the versions of all the works of a Wordsworth or a Byron, along with all his journals and correspondence. The trouble is, adopting such a device amounts to a declaration that one's every word is worth preserving—which is even more obnoxious than the pose of the poet who claims he never cancels a line.

STILL, SOMEWHERE IN AMERICA TODAY there must be a writer of merit who is either meticulously or absentmindedly saving her complete oeuvre as a series of computer files. Fifty years from now some lucky scholar will sit down in a library carrel to unseal the treasure. There they'll be, packed in a cardboard carton: 600 eight-inch floppy disks from a Radio Shack TRS-80 Model 1. What is the probability anyone will be able to read them? Even supposing the information encoded on the disks has survived, where would one go in the year 2043 to find a working TRS-80? And a copy of the Electric Pencil, the word-processing program the author used to create her works?

It is curious that archival longevity seems to be the last thing anyone worries about when choosing computer hardware and software. In picking a word processor for my own use, for example, I have focused mainly on ease and speed of editing, and the elegance of the on-screen display; I've thought very little about how I will read my own files a few decades from now, when I will have gone on to another computer, another word processor, another disk format (if indeed the very notion of "word processor" or "disk format" is still meaningful). I should know better. I've changed computers five times in ten years. Every few months I need to resurrect a document from some long-gone system, and I spend an exasperating hour puzzling over cryptic formatting commands that were once intimately familiar. What does "@|" mean again? And "«HYØ»"?

I'm not the only one with such a narrowly constrained time horizon. The computer I'm writing on at this moment thinks the world will end in 2040.

Both buyers and sellers of software pay a good deal of attention to the transfer of information between different programs and computer systems, but the emphasis is on synchronic rather than diachronic transfers. We worry about how to move a WordPerfect file on an IBM PC onto an Apple Macintosh equipped with Microsoft Word; we don't pause to ask how our descendants will read any of those files in a century or two, when WordPerfect, IBM, Apple Computer and even Microsoft are only dim memories.

GIVEN ALL THE DRAWBACKS AND DISadvantages of electronic documents, why not just stick with paper? The best way of answering that question is to look back on the one other occasion in human history when a writing medium was replaced. To societies accustomed to writing on stone or clay, paper must have seemed terribly ephemeral stuff, vulnerable to fire and water, with inscribed marks that all too easily smudged or bleached away. And yet paper prevailed. Moses' tablets were stone, but the story of Moses was told on paper. The economic incentives were just too powerful to be ignored: with paper, information became far cheaper to record, to store and to transport. Exactly the same considerations argue that a transition to paperless, electronic writing is now inevitable.

In any case, eternity is too much to ask of any storage medium. Libraries are full of disintegrating paper books; graveyards are full of stone tablets eroded to illegibility; even languages die. Perhaps the best advice, if you must write for the ages, is this: Write very well. In the centuries to come no one will be reading your verses or your novels because they are stored as WordStar files on 1.2 megabyte floppy disks; but maybe someone will preserve the equipment needed to decipher those files and disks if that's the only way to read your deathless works. •

BRIAN HAYES is a contributing editor of THE SCIENCES.