Tragic loss or good riddance?
The impending demise of traditional scholarly journals

FULL VERSION

Andrew M. Odlyzko
AT&T Bell Laboratories
amo@research.att.com

October 24, 1994

Table of contents:

[1]To be published in *Intern. J. Human-Computer Studies* (formerly *Intern. J. Man-Machine Studies*) and reprinted in *Electronic Publishing Confronts Academia: The Agenda for the Year 2000*, Robin P. Peek and Gregory B. Newby, eds., MIT Press/ASIS monograph, MIT Press, 1995. Condensed version to appear in *Notices Amer. Math. Soc.*, Jan. 1995.

1. Introduction

Traditional printed journals are a familiar and comfortable aspect of scholarly work. They have been the primary means of communicating research results, and as such have performed an invaluable service. However, they are an awkward artifact, although a highly developed one, of the print technology that was the only means available over the last few centuries for large-scale communication. The growth of the scholarly literature, together with the rapidly increasing power and availability of electronic technology, are creating tremendous pressures for change. The purpose of this article is to give a broad picture of these pressures and to argue that the coming changes may be abrupt.

It is often thought that changes will be incremental, with perhaps a few electronic journals appearing and further use of email, ftp, etc. My guess is that change will be far more drastic. Traditional scholarly journals will likely disappear within 10 to 20 years. and the electronic alteratives will be different from current periodicals, even though they may carry the same titles. There are obvious dangers in discontinuous change away from a system that has served the scholarly community well [Quinn]. However, I am convinced that future systems of communication will be much better than the traditional journals. Although the transition may be painful, there is the promise of a substantial increase in the effectiveness of scholarly work. Publications delays will disappear, and reliability of the literature will increase with opportunities to add comments to papers and attach references to later works that cite them. This promise of improved communication is especially likely to be realized if we are aware of the issues, and plan the evolution away from the present system as early as possible. In any event, we do not have much choice since drastic change is inevitable no matter what our preferences are.

Predictions and comments in this article apply to most scholarly disciplines. However, I will write primarily about mathematics, since I am most familiar with that field and the data that I have is clearest for it. Different areas have different needs and cultures and are likely to follow somewhat different paths in the evolution of their communications.

The rest of this section is devoted to an outline of the paper. The impending changes in scholarly publications are caused by the confluence of two trends. One is the growth in the size of the scholarly literature, the other is the growth of electronic technology.

Section 2 discusses the first factor that is forcing a change in traditional publishing. The number of scientific papers published annually has been doubling every 10-15 years for the last two centuries. Growth has stopped recently, but this is likely to be a temporary pause of the kind that have occurred before.

The exponential growth in scholarly publishing has interesting implications. When the number of articles per year doubles in 10 years, as it did for several decades after World War 2 in many fields, about half of the entire literature will have been produced during the preceding 10 years. Even if growth then stops suddenly, the total number of articles will double in another 20 years. Since library costs are already high, there is bound to be pressure to change the system of scholarly publications.

Section 3 describes one of the electronic information dissemination systems that are becoming widespread. Their costs are negligigble compared to those of traditional print journals, and they provide most of the services that scholars have come to expect.

Sections 4 and 5 discuss the technological trends that are making new methods of publication possible. Processor and transmission speeds are increasing at rates far higher than the growth rates of scholarly literature. This will make a change to electronic publishing feasible in the next few years.

Section 6 describes some of the recently established electronic journals. Most are operated by scholars and do not charge fees for access, and thus solve the cost and availability problem of print journals.

Section 7 argues that although there have been many previous predictions of dramatic changes in scholarly publishing, this time it will really be different, and the change is imminent. Technology is about to provide tools for a far superior information dissemination system, and the problems with the existing one are severe enough that a change will be unavoidable.

Section 8 is devoted to a description of future electronic publications. It is not simply that the escalating costs of print journals will force scholars to move away from paper publications. There are also novel features that can be implemented on the networks but not in print that are already pulling scholars towards electronic communication. First some of the present network communication methods are surveyed, and then projections are presented of what future systems will be like.

Section 9 examines in detail the costs of the current system of print journals, how it might change, and how this change will affect publishers and librarians. There are many uncertainties about the precise shape of future systems, but costs will have to be reduced dramatically. This will force publishers to shrink. Libraries will also be affected. It is possible that in most fields, such as mathematics, a few dozen scholars and librarians at a central organization might be able to provide electronically all the services that a thousand reference librarians and their assistants now do.

2. Growth of literature

The impending changes in scholarly publications are caused by the confluence of two trends. Both trends have exponential growth rates (where exponential refers to the mathematical meaning of this word, not the current journalistic usage). One trend is in the size of the scholarly literature, which is causing a crisis in the traditional approach to dissemination of research results. The other trend is the increasing power and availability of electronic technology. In this section we look at the growth in publications.

Not all scholars perceive that there is a genuine crisis in the present system of journal publications. However, librarians and publishers are well aware of it. New journals are springing up all the time, yet libraries are not buying them, and are even dropping subscriptions to old journals. The blame for this crisis is usually ascribed either to greedy publishers or short-sighted administrators. The basic underlying problem, though, is the exponential growth in the scholarly literature. The number of scientific papers published annually has been doubling every 10-15 years for the last two centuries [Price]. In many fields, there was an acceleration in the rate of growth after World War 2 to a doubling every 10 years or less. For example, in geosciences the number of publications was doubling about every 8 years over an extended period of time [Hall]. In other fields growth was slower, and in astronomy the doubling period seems to have been closer to 18 years [DavoustS]. For a more comprehensive view of recent developments we can look at *Chemical Abstracts* (*Chem. Abs.*). This review journal deals with large areas of biology, physics, and related fields as well as with chemistry. The number of abstracts it publishes was doubling about every decade from 1945 to about 1980, when it reached 475,000, and has grown more slowly since then, to a level of about 550,000 per year.

Mathematics is an old discipline, with an extensive literature accumulated over the centuries. However, what is seldom realized is that, just as with other areas, the bulk of this literature is young. In 1870 there were only about 840 papers published in mathematics [GroetschelLS]. Today, about 50,000 papers are published annually. (If you ever wonder why we no longer have mathematicians like Hilbert and Poincare, who had a comprehensive understanding of mathematics, these figures are surely a large part of the story.) The increase from 840 to 50,000 over 124 years corresponds to a doubling in the rate of publication about every 20 years. However, this growth was not even, and a more careful look at the statistics shows that during the post-World War 2 era, the number of papers published has been doubling about every 10 years [MR]. Growth has stopped recently. According to Jane Kister of *Mathematical Reviews* (*Math. Rev.*), the number of items entered into the *Math. Rev.* data base reached about 57,000 per year in 1990, and has remained at that level since then (private communication). *Math. Rev.* has stayed at approximately 47,000 reviews per year only by limiting its coverage.

The exponential growth in mathematical publishing has interesting implications. Adding up the numbers in [MR] or simply extrapolating from the current figure of about 50,000 papers per year and a doubling every 10 years, we come to the conclusion that about 1,000,000 mathematical papers have ever been published. What is much more surprising to most people (but is a simple consequence of the geometric growth rate) is that almost half of them have been published in the last 10 years. Even if the rate of publication were to stay at 50,000 papers per year, the size of the mathematical literature would double in another 20 years.

Similar exponential growth rates can be seen in other indicators of the size of the mathematical research enterprise. The American Mathematical Society (AMS) had 1,926 members in 1930, 6,725 in 1960, and 25,623 in 1990, a doubling roughly every 16 years. (The number of papers published has been doubling faster, almost every 10 years, and the share of the papers written by mathematicians in North America has not dropped too much recently, and went up substantially in the '30's and '40's. Does this mean that mathematicians have become more productive, possibly because their jobs involve less teaching and more research, or is it just that they publish more, or are they publishing shorter papers? These are intriguing questions that invite further study.) The Division of Mathematical Sciences of the NSF has seen its budget grow from about $13 M in 1971 to about $73 M in 1991, which is about a doubling in inflation-adjusted dollars. Awards (which are roughly proportional to the number of researchers supported) went from 537 to 1424, an increase by a factor of 2.65. The complaints one hears about reduced NSF support, just like those about reduced library budgets, are not so much about reductions in absolute or even inflation-adjusted budgets, but about their growth rates not keeping up with the number and output of researchers.

Exponential growth of the scholarly community cannot continue for long. The data from *Chem. Abs.* and *Math. Rev.* cited above show that there has already been substantial slowdown in the growth of publications. There are clear signs (evident in the current job market and the projections one reads) that the number of jobs in North America is not likely to grow at all, or at least will not grow anywhere near as fast as it used to. However, there have been similar slowdowns before, and there is scope for continuing exponential growth in the literature. Currently most scholars are educated and work in Europe, North America, and Japan. Continued rapid economic growth and better education in countries such as China and India will enlarge our ranks. With world population growing towards 10 billion in the next three or four decades, we might find ourselves in the second half of the 21st century with 10 times as many researchers as we have today. Could your library conceivably cope with a 10-fold increase in the literature? Repeating the point made earlier, even if we don't project more than 50 years into the future, and if we somehow manage to stop the growth in the quantity of scholarly publications, we will still double the size of the existing literature in the next 20 years. Can your library cope with that?

Scholarly publishing has some features that sharply differentiate it from the popular fiction or biography markets, and make rapid growth difficult to cope with. Research papers are written by specialists for specialists. Various estimates have been made of how many readers a scholarly paper usually attracts. These estimates are not too trustworthy, since there is difficulty in collecting data, and also in defining what it means to read a paper. Typical estimates for the number of serious readers in technical fields are under 20. (One version of an old joke has an editor tell a referee, "You are probably the second and last person to read this paper," to which the referee replies, "Second? Are you sure the author has read it?") To check whether this number is reasonable, ask yourself how many papers you have read in the last year. Great value is also derived by scholars from less thorough reading of articles, but even there, the number of times that an article is scanned is not very high. (See Section 9.4 and [KingMR] for some estimates of this number.) Whatever the correct average is for either careful study or scanning of papers, it is small, and is unlikely to grow. This is a consequence of a simple arithmetic relationship. If a scholar reads x papers and writes y papers per year, and those numbers do not change much, then the average paper will be read by x/y scholars, no matter how large the total community is. (This argument is not rigorous, since it assumes a steady state. With growth in the ranks of scholars, many scholars who do read papers but do not write any, and papers being read many years after they are published the numbers can be tricky to evaluate, but the general conclusion seems right. All the numbers in this article are back-of-the-envelope type estimates. Given the exponential growth rates we are dealing with, more precise figures would not change the implications of the message.) This would change if we could attract more readers from outside our areas. However, given the increasing specialization that applies in all fields, this is unlikely. Interdisciplinary interactions are occurring with increasing frequency, but they almost always involve specialists, and this does not affect the trend towards a narrower focus of research, which precludes attracting a general lay audience. The trade press is different. If the potential audience doubles (either because of natural increase in population, or because readers in other countries start buying translations), a popular writer like Le Carré can expect to sell twice as many copies of his books. (There will be further discussion on the differences between scholarly and trade presses later. See also [Harnad2].)

Scholarly publishing, because of its nature, cannot benefit from the economies of scale that the trade press enjoys. As our numbers grow, we tend to work in narrower specialties, so that the audience for our results stays constant. Further, the centers in which we work (typically university departments) do not grow much, so we get increasingly dispersed. On the other hand, in principle we need access to all the literature. This leads to a crisis, in which we cannot afford to pay for all the paper journals we need in our work. Libraries are coping with the increasing volume and cost of journals by cutting back on subscriptions, which only serves to increase the costs to the remaining subscribers. Many mathematics journals have been suffering from drops in the number of subscriptions of about 4% per year, while the total number of journals grows.

Scholarly publishing would be facing a minor inconvenience and not a crisis if the scale of this enterprise were small enough. If a university department were paying $5,000 per year for journals, it could deal with several decades of doubling in size and cost of the subscriptions before anything drastic had to be done. However, good mathematics libraries spend well over $100,000 per year just for journal subscriptions, and the best ones spend close to $200,000 (with those at the top of the range obtaining also some of the most relevant physics, computer science, and engineering literature). To obtain the true cost of the library, we have to include the cost of staff and space. According to data for 1993 compiled by the Association of Research Libraries (ARL), and available on-line on the World Wide Web at URL http://arl.cni.org, costs of all acquisitions are usually about one third of the total cost of a research library. Hence a mathematics library that spends $150,000 on books and journals per year probably costs close to $500,000 per year to run. Totals for all departments at a university are often staggering. According to figures from ARL, Harvard University spends $58 million per year on all its libraries. Harvard stands out for the size and expense of its libraries, but other institutions also have large budgets, with (to cite just a few examples) Stanford at $36 M, University of Michigan at $28 M, Princeton at $21 M, MIT at $12 M, and Dartmouth at $10 M. Budgets that large are not easy to increase, and are likely to be scrutinized for possible reductions.

3. A brave new world

Modern technology is making possible methods of information dissemination that are dramatically cheaper than traditional journals. An example can be seen in the system that my colleagues in the Mathematical Sciences Research Center of Bell Labs and I have started to use recently. My recent preprints (including this essay) can be accessed through Mosaic at URL ftp://netlib.att.com/netlib/att/math/odlyzko/index.html.Z. (Preprints of some older, already published papers are also available there, but may have to be removed if publishers complain.) For those without access to Mosaic, ftp access is available on machine netlib.att.com. After logging in as "anonymous" and giving the full email address as password, all the user has to do is give the commands

```
cd netlib/att/math/odlyzko
binary
get index.Z
```

to obtain a copy of the (compressed) index file, which describes what preprints are available. Alternatively, those preferring email or simply not having ftp available can send the message

```
send index from att/math/odlyzko
```

to netlib@research.att.com, and the index file will arrive via return mail, with instructions for retrieving individual papers. (For papers of my colleague Neil Sloane, use the same commands as above, but with "odlyzko" replaced by "sloane," and so on.)

The system described above provides access for all the 20 million users of the Net (as the Internet and various other networks are called) to the preprints that my colleagues and I write. This access is free and available around the clock. Further, this access is easy, and can accommodate users with various levels of connectivity and terminal equipment. What is most remarkable about it, though, is that it places a minimal burden on the author. All I need to do (once a paper has been typeset in LaTeX, say) is to give the commands

```
latex analytic
dvips analytic.dvi > /usr/math/odlyzko/analytic.comp.ps
```

and edit the file /usr/math/odlyzko/index by adding to it the lines

```
file   att/math/odlyzko/analytic.comp.ps
title  Analytic Computations in Number Theory
by     Andrew M. Odlyzko
#      to appear in "Mathematics of Computation 1943-1993,"
       W. Gautschi, ed., Amer. Math. Soc., Proc. Symp. Appl. Math., 1994.
```

Everything else is done automatically by this public domain system, which was written by my colleague Eric Grosse. (In practice there is a bit more work, since I also make the source files available in the src directory.) The plummeting costs of storage mean that I do not have to worry about sizes of data sets. I therefore make available both PostScript files, for portability, and the TeX source files, to make text searches easier.

The only time-consuming part in using Grosse's system is the typesetting of the paper, but that would be done in any case. The extra effort needed to make the preprint available is a matter of a minute or two. This is a dramatic change compared to the situation of even a few years ago, and certainly to that of a few decades ago, when the only way for a scholar to communicate with a wide audience was to go through the slow and expensive process of publishing in a conventional journal. Now it is possible to reach a much broader audience with just a few keystrokes.

Similar systems are becoming widely used, and are likely to spread. There is a compelling logic to them, as there is to the uniform email addressing convention that is becoming common, in which if Gauss were active at Celestial University, he might be reachable as gauss@math.celestial.edu. Mathematics departments are beginning to set up preprint archives at addresses of the form ftp.math.celestial.edu. Existence of such publicly accessible preprint archives is a great boon to scholars, but it is extremely subversive of journal publications. If I can get a preprint of a published paper for free, why should I (or my library) pay for the journal?

Preprints are treated differently from refereed journal publications. However, the system described above can be used just as easily to run an electronic journal. If I am an editor, all I need to do after receipt of a revised paper that has been refereed is to follow a procedure analogous to that described earlier, simply placing the paper in a directory designated for the journal. This is again in dramatic contrast to the traditional system, where the publisher provided the extensive range of skills and facilities (typesetting, copy editing, printing, distribution, etc.) that were needed to operate a journal. An electronic system can be run by scholars alone. The skills involved are of even higher caliber than those in the traditional setting, but they are now embodied in the hardware and software that is available at low or no cost.

The facilities described in this section are examples of what is becoming common. The next two sections consider in detail developments in modern technology that are making this possible.

4. Hardware improvements

A doubling of scholarly papers published each decade corresponds to an exponential growth rate of about 7% per year. This is fast, but nowhere near as fast as the rate of growth in information processing and transmission. Microprocessors are currently doubling in speed every 18 months, corresponding to a growth rate of 60% per year. Similarly dramatic growth figures are valid for information storage and transmission. For example, the costs of the NSF-supported backbone of the Internet increased by 68% during the period 1988-91, but the packets transmitted went up by a factor of 128 [MacKieV]. The point of citing these figures and those below is that advances in technology have made it possible to transform scholarly publishing in ways that were impossible even a couple of years ago.

Recall that about 50,000 mathematical papers are published each year. If they were all typeset in TeX, then at a rough average of 50,000 bytes per paper, they would require 2.5 GB of storage. If we include the fewer than 1,000 mathematics books (not counting college textbooks, say) that are published each year, we find that the annual output of the mathematical community requires under 4 GB of storage. (This assumes everything is in TeX. If we use dvi files, the requirement about doubles, while with PostScript it about quintuples. On the other hand, compression can reduce storage requirements by a factor of 2 or 3 for TeX source code, and 5 for PostScript. Given the rapid growth in the availability and cost of storage system, factors of 2 or 5 mean only a difference of a couple of years.) These numbers, which looked daunting a few years ago, are now trivial.

For comparison, the EDGAR database of financial reports receives about 30 GB of data per year. (It is available through several commercial vendors and is now being made available for free on the Internet by the New York University School of Business Administration with the support of a grant from NSF.) The Canadian Meteorological Centre receives 3 GB of data per day.

We can now buy a 9 GB magnetic disk for about $3,000. For archival storage of papers, though, we can use other technologies, such as optical disks. A disk with a 7 GB capacity that can be written once costs $200–300. (The equipment for writing data on it is still expensive, and costs $20,000 - 40,000, but it can be shared by many individuals and even departments). Digital magnetic tape is even cheaper. The standard CD-ROM disks, with about 0.7 GB of storage capacity, cost a few dollars each to produce (with the information written on it) in runs of a few thousand. Digital tapes with 250 GB capacities are expected to become available soon. Thus the electronic storage capacity needed for dissemination of research results in mathematics is trivial with today's technology.

We conclude that is it already possible to store all the current mathematical publications at an annual cost much less than that of the subscription to a single journal. What about the papers published over the preceding centuries? Since there are 1,000,000 of them, it would require about 50 GB to store them if they were all in TEX. It appears unlikely (but not impossible) that anyone will undertake the project of converting them all into TEX (or any other modern electronic format). What is more likely is that optical character recognition will eventually be applied to the texts, since this will enable rapid computerized text searches, while the equations will be stored as bitmaps. To provide reliable access to the text, whole papers will have to be available as full bitmaps. This dramatically increases the storage requirements. However, even then, they are not prohibitive. With current fax standards (which produce copies that are not pleasant to read, but are usable), a page of mathematical text requires 30-50 KB. Therefore all the 1,000,000 mathematical papers can be stored in less than 1,000 GB. This is large, but it is still less than 150 of the current large optical disks. For comparison, Wal-Mart has a database of over 1,000 GB that is stored on magnetic disks, and is processed intensively all the time. The credit card transaction records maintained by American Express come to 220 GB. The DIALOG set of databases contains about 3,500 GB of data.

The storage requirements of scholarly literature are likely to become even more ordinary in the near future. Cable TV and phone companies are getting ready to deliver movie-on-demand to the home. This will require tremendous storage capacity, and is bound to stimulate development of much denser and cheaper media. There is not much progress in optical disks, but magnetic hard disks are becoming both larger and cheaper, and there are many promising new technologies on the horizon, such as optical tape that might be able to store over 1,000 GB on a single unit. We can start thinking of 1,000 GB storage devices for personal computers becoming available in a decade or so. This means that any one will be able to have all the mathematical literature available right on the desktop. Librarians and scholars often express fears about availability and durability of data in the electronic world. However, once we can put all the mathematical papers on a single $50 tape, every professional mathematician, every graduate student, and every local public library will be able to have one. Instead of 1,000 libraries in the world having a copy of a particular journal, we will have 100,000 libraries and individuals owning it.

Even before systems able to store all the mathematical literature become available to individuals at low prices, though, larger institutions, such as university mathematics departments, will be able to make such systems available to their members. This ability will mean a dramatic change in the way we operate. For example, if you can call up any paper on your screen, and after deciding that it looks interesting, print it out on the laser printer on your desktop, will you need your university's library?

Although the arguments used here are repetitive, it is important to stress that technological progress has pushed the state of what is available with routine off-the-shelf systems far ahead of what is required for scholarly publishing. This makes many arguments against electronics invalid. For example, there is the famous story of warehouses full of tapes of US census data that have deteriorated so that they cannot be read. The argument is that if census data cannot be preserved in digital format, how can one possibly expect to preserve scholarly literature? However, the difficulty with the census tapes was that they were the product of an earlier low-density storage technology that required warehouses for what could today be accommodated on a handful of optical disks. During my early experiences in computing, on a mainframe in the mid-70s, disk storage was expensive, so I could not even afford to store on disk all my programs, much less data, and had to toss away or archive most of what I had. Given the slow turnaround in retrieving things from archives, I tended to avoid them, and when we moved to new machines, I would sometimes not bother recovering the archived material. These days everything goes much more smoothly. Recently it was announced that an old machine that used to be my home was going to be retired. I had been a big user of it, and had accumulated around 40 MB of files, which was a lot 10 years ago. However, now I am one of a handful of users sharing a 3 GB disk, so all I did was to issue a one-line command to copy all the files from the old machine, including the useless executables, into a new directory on my current machine. I am convinced that in the future things will be handled similarly. Yes, if you are an astrophysicist who is getting terabytes of data from the Hubble space telescope, and you build some kludgy one-of-a-kind holographic memory that happens to give you the factor of 5 improvement in capacity over what's on the market today, you may have trouble maintaining your data. However, if you stay well behind the technological frontier, then every 5 years you will move to a file system with 10 times the capacity of the old one, and you will simply devote 10% of it to storage of files from your old machine.

The argument about scale of scholarly publishing suggests that some technologies will not suffice. CD-ROMs are rapidly growing in popularity, and seem set to become the preferred medium of electronic publishing for general literature. For example, in 1993, more CD-ROM encyclopedias were sold in the US than paper ones. However, CD-ROMs are not ideal for scholarly publishing because of their size. At 0.7 GB capacity, they are not capable of storing even the complete set of reviews of the mathematical literature. It is true that a CD-ROM might hold almost all the papers that an individual mathematician will work with, but only if one can decide beforehand what those papers are going to be. There is simply not enough capacity to put in all that might possibly be needed, as could be done if the capacity were 70 GB, and not 0.7 GB. (Multiple CD-ROMs can be an answer, but not a convenient one.) Therefore I suspect that mathematicians and other scholars will depend on network connections more than on CD-ROMs for access to information, even though this technology will be useful for distribution of some data. CD-ROMs may even act as a distraction, with great effort devoted to trying to squeeze data onto them. As in the case of fax vs. electronic mail, this may develop into a case where the "good enough" may delay the arrival of "best."

Optical disks can be shipped by parcel post at low cost. However, it is desirable to have faster communication links. Information transmission is a barrier right now. Anyone who has to download files at 2,400 baud (roughly bits per second, bps) can attest how annoyingly slow that is. However, communication speeds are increasing rapidly. Most departments have their machines on Ethernet networks, which operate at almost 10 Mbs (millions of bits per second). Further, almost all universities now have access to the Internet, which was not the case even a couple of years ago. The Internet backbone operates at 45 Mbs, and prototypes of much faster systems are already in operations. Movies-on-demand will mean wide availability of networks with speed in the hundreds of megabits per second. If your local suppliers can get you the movie of your choice at the time of your choice for under $10 (as they will have to, in order for the system to be economic), then sending over the 50 MB of research papers in your specialty for the last year will cost pennies. Scholars might not like to depend on systems that owe their existence to the demand for X-rated movies, but they will use them when they become available.

There is concern in the scientific community that with the withdrawal of NSF support for the Internet, the electronic data highway will charge tolls that will be prohibitive for university researchers. That concern is misplaced. In the first place, the NSF subsidy is only \$20 M per year [MacKieV], and covers around 10% of the cost of running the Internet. Compared to the 20 M users of the Internet, that is not much. Further, this concern is based on a static view of the world and does not take into account the increasing power and decreasing cost of technology. Yes, the commercial data highway will have tolls. Yes, the networks that the cable TV companies build may be structured primarily for the broadcast delivery of movies to the home, and may not have the full communication capabilities that scientists require. The point is that these networks are going to be large enough to justify tremendous development efforts that will drive down the costs of all communication technologies. The tolls on the commercial data highway will have to be low enough to allow transmission of movies. Therefore the cost of transmitting mathematics will be trivial. Even if the commercial data highway is not structured in an ideal way for science, or is a commercial flop, universities will be able to use the technology that is developed and build their own cheap networks. Just as the recent progress in computers is caused much more by demand from businesses running spreadsheets on PCs than from scientists modeling global warming on supercomputers, so the progress in communications will come primarily from commercial uses of the data highway. Some scholars (such as those getting atmospheric data from satellites) will have to worry about the costs of communications, since they will be transmitting giga- and tera-bytes of data. The rest of us can sit back and go along for the ride provided almost for free by commercial demand.

Stevan Harnad (private communication) notes that even now, non-scientific and non-educational uses of the Internet are a significant fraction of the traffic. Some 1993 statistics for net-news (the decentralized system of discussion groups) collected by Brian Reid of DEC show that the largest volume was generated by alt.binaries.pictures.erotica, which has an estimated 240,000 readers, and generated 31.4 MB of postings during one sampling period, even though it is received by only about half the sites. The second largest load was 16.3 MB from alt.binaries.pictures.misc, which has 150,000 readers. The third highest was 14.1 MB from bionet.molbio.genbank.updates, which has 24,000 readers. Sci.math, which has an estimated 120,000 readers (the highest of any of the sci.* groups, but only 39-th highest among all the groups) generated only 3.6 MB. Even alt.politics.clinton generated 7.1 MB during the sampling period! This only reinforces the conclusion drawn earlier that most of scientific communication can be accommodated in a tiny fraction of the capacity of future networks. (The gene bank data, like the atmospheric data mentioned before, is a special case.)

The arguments above do not mean that the existing networks are already sufficient for scholarly communication. Even though the traffic over the Internet backbone has been doubling every year for the last three years, it still comes to only between 1 and 3 MB per person per month (depending on how many users there really are). If email and ftp of papers were all that was going to be used, then 3 MB per month would be plenty. With the growth of multimedia traffic, though, we might soon run into a crisis situation that will require a substantial increase of the network, and might require imposition of tolls. Still, this is likely to be only a momentary problem, since technology is progressing rapidly, and faster networks are being developed.

5. Software improvements

Not only have information storage and transmission capacities grown, but the software has become much easier to use. Computerized typesetting systems have become so common that it is rare to encounter a manuscript typed on an ordinary typewriter. Moreover, scholars are increasingly doing their own typesetting. This trend is partially due to cutbacks in secretarial support, but is caused primarily by scholars preferring the greater control and faster execution that they can obtain by doing their own typesetting.

Two centuries ago there was a huge gap between what a scholar could do and what the publishers provided. A printed paper was far superior in legibility to hand-written copies of the preprint, and it was cheaper to produce than hiring scribes to make hundreds of copies. Today the cost advantage of publishers is gone, as it is far cheaper to send out electronic versions of a paper than to have it printed in a journal. The quality advantage of journals still exists, but it is rapidly eroding. Systems such as TeX, widely used by scholars, are being increasingly adopted by publishers in response to the economic pressure to lower costs. (It was the switch to a poor electronic typesetting system by Don Knuth's publisher that prompted him to invent TeX in the first place.) Furthermore, these systems, because of their wide use, are improving so that they are almost as good as the best conventional typesetting systems. The software improvement is aided by increasing quality of laser printers (with 600 bpi printers replacing 300 bpi ones, and 1200 bpi printers likely to become common soon). Authors of papers and books with numerous figures are frequently told by publishers that it would not have been economically feasible to publish their manuscripts with old systems. Most journal papers still have the advantage of better editing, but even that advantage is eroding with the development of better software and better training of scholars. Further, many authors complain that the advantages of professional manuscript preparation are often vitiated by the mistakes that this process introduces.

The progress in communications is even more striking than in preparation of manuscripts. At the beginning of the electronic era most of us had email. Then came anonymous ftp. Together with the network connections that are now almost universal, they provide powerful tools for communication. However, this is only the beginning. We now have a plethora of new tools, such as Archie, Gopher, and WAIS, which allow easy access to the huge stores of information on the Internet without having to learn arcane details of how it functions. In the next few years these tools are likely to evolve much further, and become widely used in the day-to-day life by most scholars. The one software tool that seems likely to have the greatest impact is Mosaic. It is a browser, a tool for reading documents in the World Wide Web, a system of distributed hypermedia databases. The usage of Mosaic is growing explosively, as it seems to have finally reached the level of user-friendliness that the wide public can live with. (This seems to be another case of the discontinuous change that is common with new technologies. As some threshold of either price or convenience is passed, there is a snowball effect. We have seen this with fax machines and CDs, and I argue that we are likely to see this with electronic publishing of scholarly information.) Mosaic is often touted as the "killer app" for the Internet. With its "point-and-click" feature, it enables users to move from file to file, and from machine to machine. It's as if a visitor to a library, on encountering a reference in a book, could press a button and be magically transported to the shelf containing the referenced journal article. Many of the existing and planned electronic journals are adapting to Mosaic as their main access tool. I will not try to describe it, as that has been done extremely well in many places.

All I will say is that I am an enthusiastic user of it, and I urge any reader who has not seen it to arrange as soon as possible to try it out.

While Mosaic is a great tool, it does have its limitations. After a prolonged period of "surfing the Infobahn," users begin to express their dissatisfaction with the lack of structure in the databases and difficulty in quickly reaching what is needed. Mosaic may eventually be replaced by a better tool, just as VisiCalc was the first spreadsheet and helped launch the PC revolution, yet was completely displaced within a few years by better products such as Lotus 1-2-3. It might also be that Mosaic will evolve. Many software tools can be integrated easily with Mosaic. For example, RightPages(TM) is an experimental Bell Labs system that gives the user access to about 60 technical journals from about 10 publishers [Story]. After selecting an issue of a journal involved in the experiment, the user is shown the table of contents, and by clicking on an item, he or she is shown the first page of the article. If it looks interesting, a click on the appropriate icon orders a paper copy for delivery to the user. (At the moment only the first page is scanned, but this is not a fundamental limitation, and is caused by the lack of resources in a small-scale experiment to do more.) The RightPages system is used inside Bell Labs, and also at the University of California in San Francisco, in the Red Sage (TM) project. Its advantage is that it is better adapted for presenting journal papers than is raw Mosaic.

In addition to improvements to or on Mosaic, we will soon have automated personalized "agents" that will perform data searches automatically. Eventually such agents will have the ability to interact among themselves, specialize and become a secondary level of information.

Systems like RightPages are meant to work with the current paper journals and memoranda. They are supposed to compensate for some of deficiencies of the present system, by allowing users to sift through the increasing amounts of information that are becoming available, and also by giving them access to information they cannot get in their local libraries. However, while they do ameliorate the crisis in scientific publishing, they also contribute to it. I can now use the Internet to find out what books and journals the Harvard libraries have. Suppose I could also look on my screen at the articles in the latest issues of the journals that Harvard has received recently, and order photocopies (or even electronically digitized scans that can be reproduced right away on my laser printer) of the articles that interest me. Would there be any incentive to pressure my local library to order that journal? Would there be any reason to have more than one copy of the journal in the country (or the world)? If we do have only one copy, how is it going to be paid for? Thus the arrival of technological solutions to the current crisis in scholarly publishing also threatens to aggravate that crisis.

6. Electronic journals

What I am predicting is that scholarly publishing will soon move to almost exclusively electronic means of information dissemination. This will be caused by the economic push of having to cope with increasing costs of the present system and the attractive pull of the new features that electronic publishing offers. The costs of conventional journals have been mentioned in Section 2 and will be discussed in much greater detail in Section 9. Here we discuss briefly electronic journals.

It is estimated there are currently around 500 regular electronic journals. At most 100 of these journals have rigorous editorial and refereeing standards that are comparable to those of the majority of print scholarly journals. However, while even all the 500 electronic journals come to less than 1% of the print periodicals published in the world, they are growing by around 70% per year. In mathematics, the *Ulam Quarterly* (*UQ*) is the oldest, and has been in operation since 1992. The *Electronic Transactions in Numerical Analysis* (*ETNA*) and the *Electronic Journal of Differential Equations* (*EJDE*) started in 1993, while the *Electronic Journal of Combinatorics* (*ElJC*) and the *New York Journal of Mathematics* (*NYJM*) are starting their operations in 1994. In addition, the *Chicago Journal of Theoretical Computer Science* (*CJTCS*) and the *Journal for Universal Computer Science* (*JUCS*) are scheduled to begin publishing in 1994 and 1995, respectively. (In the interests of full disclosure, it should be mentioned that I am on the editorial boards of *ETNA, ElJC*, and *JUCS*.) These journals are exclusively electronic (although a printed version of UQ is distributed by a commercial publisher for $70 per year). Further, the *Bulletin of the AMS* (*BAMS*) is available electronically in addition to still being printed and sent out to all members of the AMS. The economics of electronic journals will be discussed at length in Section 9 of this article. I will argue that a move towards electronics offers opportunities for dramatic cost savings. Of the journals mentioned above, *CJTCS* is published by the MIT Press, and charges a subscription comparable to other new small journals in the mathematical sciences. All the other journals are currently available for free. (*BAMS, EJDE, ElJC, ETNA*, and *UQ* are available through the AMS e-math machine, URL http://e-math.ams.com/web/index.html, or telnet e-math.ams.org, login e-math, password e-math.) Some either already plan to charge subscription fees in the future or reserve the right to do so. Others, such as *UQ* and *ElJC*, are operated exclusively by their editors with no financial support from subscription fees or subsidies, other than the editors' computers and connections to the Internet, which are paid for by their departments or grants. It seems likely that most future scholarly journals will operate this way, although it is possible that electronic journals with subscription fees might also exist in the scholarly arena (and will surely be the norm in the commercial fields). However, because of the nature of scholarly publishing, I suspect those fees will have to be low. The arguments are presented at length in Section 9.

Electronic journals do not have to be inexpensive, and inexpensive journals do not have to be electronic. However, electronic publication does offer the best opportunity of reducing costs. The economic argument will probably all by itself force a move to electronic journals. There are additional reasons for preferring electronic journals. Some of the most compelling ones involve the interactive potential of the Net, and are discussed in Section 8. They involve novel features that are likely to be of great value to scholars, yet are impossible to implement in print journals. Even if electronic journals did not include any of those features and restricted themselves to the familiar format of print ones, they would still have many advantages. One is easy and general access. Most mathematical journals are available at about 1,000 research libraries around the world. Even for the scholars at those institutions, access to journals requires a physical trip, often to another building, and is restricted to certain hours. Electronic journals will make access available around the clock from the convenience of the scholar's study. It will also make literature searches much easier. For journals without subscription fees, access will be available from anywhere in the world. Electronic publishing will also have a healthy influence on preprint distribution. One point that is made about Ginsparg's theoretical physics preprint service, both by Ginsparg and other users [Ginsparg, Sci1], is that it has made the latest results much more widely available, and diminished the importance of various small "in" groups.

The advantages of easy and almost universal access to electronic publications have a measurable impact on scholars' activities. For example, it is noted in [Ginsparg] that many physicists obtain copies of their colleagues' papers (from the same institution) through Ginsparg's preprint server instead of directly from authors. Similarly, numerical analysts often use the Dongarra-Grosse netlib system to repeatedly obtain the same basic subroutines, since it is easier to obtain them from a centralized and well-organized source than to remember where the previous copy was placed.

Concern is often expressed that electronic publishing will deprive poorer institutions, especially those in the less developed countries, of access to the scholarly literature. The opposite is bound to be true. Few institutions can afford the $21 M per year that Princeton University spends on its libraries. Yet a T1 connection to the Internet (of 1.5 Mbps capacity) costs $20,000–50,000 per year in the US, and would suffice to bring in all the scholarly information that is generated in the world, if only that information were electronic. In other countries connections are more expensive, but even so, less than 1% of what Princeton spends will pay for a satellite earth station of high capacity. Further, while the cost of print journals is going up, that of electronics is going down. Therefore electronic publication is the most promising route for scholars in less developed countries to become full participants in intellectual life.

Once many journals become available electronically, paper copies are likely to disappear. It will be a case of positive feedback operating in a destructive mode. Necessary information will be available electronically, and most of it will have to be accessed electronically, since the local libraries will not be able to provide copies of all relevant journals. Therefore researchers will have less incentive to press for paper journal subscriptions to be maintained, which will lead to diminished circulation, and therefore to higher prices and more pressure from libraries to cut back on subscriptions.

7. Will it really be different this time?

Anyone venturing to predict dramatic changes in scholarly publications has to be mindful of the long history of failed forecasts in this area. Back in 1913 Thomas Edison predicted that movies were going to replace books. A much more carefully considered forecast that also predicted the demise of books was made by Vannevar Bush in 1945 in the article "As we may think," which presented the idea of Memex, a personal data storage device that would contain massive amounts of information. This influential essay is usually regarded as the progenitor of modern hypertext. Since Bush was one of the creators of digital computers, and the article was published in 1945, it is usually thought that it was stimulated by developments in computing. However, the studies in the interesting collection [NyceK] show that Bush started writing early drafts of his article during 1930s, stimulated by the possibilities of dense storage of information on microfilm. These possibilities were also fascinating such thinkers as H. G. Wells, who published the book *World Brain*, motivated by the novel ways to combine information. There were predictions that being able to put entire libraries in cabinet-sized devices in every person's home would lead to a general uplift in the intellectual level of the population. However, microfilm has played a limited role (and one that is usually cordially disliked by scholars).

In more recent times there have been other predictions of how electronic publications were going to sweep the world. Most of the developments forecast in this essay have been described a long time ago. Scholars in library science have been among the pioneers in this area, since they have had to deal with the exponential growth in literature most directly. The books [Licklider, Lancaster] are just two influential examples of the thought that went into this subject. There have even been experimental scholarly information systems with many of the features that this article predicts [Lederberg]. There have also been recent proposals for electronic journals in mathematics that include many advanced features of the kind described later [Loeb]. Yet we still live with and rely on traditional paper journals. Is there any reason to expect that rapid changes are going to take place soon, if they haven't in the last few decades?

A skeptical look at the predictions of a rapid switch to electronic journals has been presented by Schaffner [Schaf1, Schaf2]. She notes various practical obstacles, such as lack of standards for presenting scholarly data in digital formats. She also emphasizes the slow evolution of the print journal, which adapted it well to serve scholarly needs. This evolution owed little to technological developments, and was driven by developments in scholarly culture. Also, while scholars may be intellectually adventurous, they tend to be conservative in their work habits. This conservatism is often reinforced by experiences with new technologies. Much is made of the interactive potential of the Net, and Section 8 of this article will be devoted to this topic. However, Don Knuth, the creator of TEX and an eminent computer scientist, has stopped using email, as he felt it consumed too much of his time. Thus not all the features that are beckoning us to the electronic world are as great as they seem.

Although the arguments against rapid change in scholarly publishing do have some merit, they do not take into account the drastic recent changes in available technology. The traditional journals are refined products of a long evolution. However, the environment they operate in is changing drastically. Many of the features of traditional journals might hinder their survival in the new world. They will not vanish instantly, of course, and can persist in their mode of operation for a few more years, but eventually they will have to either change or die. They still have some time left, since technology is still not quite able to handle scholarly communication needs, and traditions take time to overcome. The dreams of the visionaries of several decades ago who foresaw the dramatic effect that electronics could have on scholarly communication are still not fully realized. The main reason is that it took a long time for technology to provide the tools that made those futuristic dreams possible. Even 20 years ago, computing was largely done in batch mode on mainframes, and the few fortunate enough to have access to time-sharing systems had to content themselves with printing terminals communicating at 300 bits per second. In that environment, electronic publishing and intensive collaborations across oceans were not feasible. With the rapid advance in technology, though, we are just about at the stage where the needs of an electronic publishing system for scholars can be met easily. Moreover, the early predictions (such as those of capacity of storage systems in [Licklider]) often anticipated correctly that it would only be around now that the necessary capability would become available. Thus it is not surprising that no revolution in scholarly publishing has taken place yet.

A major reason for expecting a revolution soon is that we are not dealing with some special system developed just for scholars, as Bush's Memex might have been. Instead, scholarly publishing will be swept along in the general move to the electronic world. We can see this already in the rapid growth in manuscripts that are typeset on computers and in the increasing use of email. We can also see it in some abrupt changes that have taken place in the scholarly arena, especially those associated with the adoption of Ginsparg's electronic preprint distribution system in high energy theoretical physics and a few other fields, which will be discussed in Section 8. Big changes are also occurring in the more general publishing arena, where sales of CD-ROM encyclopedias have surpassed those of paper ones.

Even if we accept that change to electronic journals is bound to take place, it is conceivable that it might be on a gradual evolutionary path, with electronic journals slowly gaining on paper ones. The high energy theoretical physicists who now distribute their preprints almost exclusively through Ginsparg's system are still submitting them for publication in traditional journals. Numerical analysts, who have been relying on the Dongarra-Grosse netlib system for distribution of software, are also still publishing their papers conventionally. However, I feel that slow evolution is unlikely. The problem is that the natural development of present preprint distribution systems, described in the next section, is going to make scholarly papers freely available on the Net, so that scholars will be relying on their libraries less and less. They will therefore have less and less incentive to press for paper journal subscriptions to be maintained, which will lead to diminished circulation, and therefore to higher prices and more pressure from libraries to cut back on subscriptions. (Circulations of many mathematics journals have been declining at about 4% per year recently, which has contributed to the price increases.) This phenomenon all by itself can lead to catastrophic declines over a period of a couple of years in print journal circulation. There could also be even more sudden changes at individual libraries. The costs of the traditional system are so high (and are still growing), that they are bound to attract attention of administrators. A library that costs $400,000 per year to maintain, after all, takes the resources that would provide several faculty positions. If nothing is done, then in a decade or so, during the next financial squeeze at a university, a dean might come to a mathematics department and offer a deal: "Either you give up paper journal subscriptions, or you give up one position." Once the hurdle of canceling journal subscriptions is overcome, the next offer is likely be "Either you give up maintaining your old bound journal collection, or you give up a position."

While the transition to electronic publication does appear inevitable, print journals are unlikely to suffer catastrophic declines in their circulation for a few more years. Even though some fields might switch to almost complete reliance on electronic information distribution soon, and do it suddenly (as high energy theoretical physics has done recently), the inertia of a decentralized and heterogeneous scholarly publishing business means that change will take time. Even if the number of electronic journals were doubling every year, it would be over 7 years before it would equal the number of print journals. Thus it seems unlikely that major changes in scholarly publishing will be visible within the next 5 years, at least when measured in the revenues of print publishers. Most papers will continue to be published in the traditional way. However, subscriptions will continue to drop, prices will continue to increase, and the system will be showing more and more signs of stress. At the same time, electronic publications will be developing rapidly, and eventually they will become dominant, most likely between 2000 and 2010.

8. The interactive potential of the Net

Because conventional print journals have been an integral part of scholarly life for so long, their inflexibility is often not appreciated. As one example, consider surveys or bibliographies. They are invariably obsolete even before they are printed, and there are few options for updating them. In contrast, an electronic version can be continually updated by the author. This is only a small example of the novel tools that electronic communication offers to scholars.

Section 9 will discuss the economics of publishing and what role various institutions are likely to have. In this section I describe my vision of what future scholarly publications will be like, with special emphasis on mathematical publications. I start by disputing Frank Quinn's vision [Quinn] of what the present system is and ought to be, and then discuss some novel information dissemination systems that have certain features I expect to find in future systems. Finally I present the basics of the system I expect to emerge over the next decade.

Scholarly journals have evolved during the last three centuries, in the world shaped by Gutenberg's invention of movable type. This invention made possible wide dissemination of scholarly publications. However, because printing, although much cheaper than hand copying, was still expensive, journals were constrained into a format that emphasized brevity. Further, the standards have promoted correctness. Since it took a long time to print and distribute journal issues, and corrections likewise required a long time to disseminate, it was natural to develop a rigorous refereeing standard. (This was not the only reason, of course, but I believe it was an important one in the development of our current editorial practices.) As a result, mathematical literature at least has become reliable, in that mathematicians feel free to use results published in reputable journals in their work, without necessarily verifying the correctness of the proofs on their own. (This correctness of the mathematical literature has increased substantially over the last two centuries. A perusal of Dickson's *History of the Theory of Numbers* shows, for example, that old papers seemed to contain serious mistakes with distressing frequency.) Frank Quinn [Quinn] argues that this feature justifies extreme caution in moving away from paper journals, especially in mathematics, lest we be tempted into "blackboard-style" publishing practices that are common in some fields. In particular, he advocates keeping a strong distinction between informal preprint distribution and the formal refereed publications, even in an electronic format. I agree that mathematicians should strive to preserve and enhance the reliability of mathematical literature. However, I feel that Quinn's concerns are largely misplaced, and might serve to keep mathematicians and other scholars from developing better methods for information dissemination.

The first point that should be made is that electronic publication does not in any way prevent the maintenance of present publishing standards. Electronic journals can follow exactly the same policies (and might even have the same names) as the current paper journals. On the other hand, paper journals are no guarantee against unreliable results, since the practices Quinn deplores are common in some fields, and have been present for a long time. Thus the reliability of literature in any field depends primarily on the publishing norms in that field, and not on the medium. Quinn is right that electronic publication does present increased temptations to careless communication. Computers do promote a less thoughtful style of correspondence. However, that can also be said of the telephone, or even a good postal service. Just read the letters that people used to write in the 18th century. By today's standards, they tended to be literary masterpieces. The difference was that letters took a long time to deliver, and were expensive, so substantial care was taken in writing them. However, nobody is suggesting that the Post Office put a minimum 20-day hold on all letters (even if it sometimes seems they are trying to do it on their own) to promote better writing. In the transition to electronic publishing, we will just have to develop methods of coping with the new trends.

Paper journals serve several purposes in addition to that of providing reliable results. By having a hierarchy of journals of various prestige levels, the present system serves a filtering role, alerting readers to what the most important recent results are, as well as a recognition role, helping in grant and tenure decisions. Here again there is no reason that electronic journals could not provide the same services.

A more serious objection to Quinn's article [Quinn] (and to a large extent also to [JaffeQ]) is that its picture of mathematicians whose main goal is to produce formally correct proofs is unrealistic. (See [JaffeQ2] for some additional comments.) I agree completely with Quinn that it is desirable to have correct proofs. However, it's a mistake to insist on rigor all the time, as this can distract from the main task at hand, which is to produce mathematical understanding. There are many areas of mathematics today that are not completely rigorous (at least so far), with the classification of finite simple groups just one example. This has been true at various times in the past as well. After all, Bishop Berkeley through his "ghosts of departed quantities" jibe about infinitesimals and related arguments did show that 18-th century calculus was not necessarily any more rigorous than theology. On the other, from a historical perspective he lost the argument, since the necessary rigor was eventually supplied. In a similar vein, we speak of the Riemann Mapping Theorem, even though experts agree that Riemann did not have a rigorous proof of it, and that a proper proof was only supplied later. Standards have not improved all that much in the intervening years. The present system does not do a good job of providing the expected reliability, even if by reliability we mean the standards accepted in a given field, and not formal correctness. Even conscientious referees often miss important points. Furthermore, many referees are not all that conscientious. Once a paper appears, there are some additional controls. Sometimes a careful reviewer for *Math. Rev.* will catch a mistake, but that does not happen often. More typically, an error will be pointed out by someone, and then, depending on who and what is involved, there may be an erratum or retraction published, or else a note will be inserted into another paper, or else the mistake will only be known to the experts in the field. If the topic is important, eventually some survey will point out the deficiencies of the paper, but most papers do not get this treatment. Thus we already have situations where published work has to be treated with caution. It is more reliable than that of just about any other field, but it is not as reliable as Quinn's article might lead us to believe.

Lack of complete reliability is only one defect of the current paper journal system. Delays in publication are the one that is best known and most disliked. This is the one area where electronic publication offers a clear advantage. (There is an interesting parallel here to the rise of scholarly journals. They originated in the 17-th century in an attempt to improve on the system of personal letters through which discoveries were being communicated, which in turn developed because the traditional method of communicating through book publications was too slow.) There are others as well. A major one is caused by the emphasis on brevity that was encouraged by an expensive system of limited capacity. Although this is seldom said explicitly, the standard for presentation in a research paper is roughly at the advanced graduate student level. If you write to be understandable to undergraduates, referees and editors will complain that you are wasting space describing basic steps that the reader is supposed to be able to fill in. If you write at a more advanced level, they will complain that the details are too difficult to fill in. The result is exposition that is often hard to follow, especially by non-experts.

Bill Thurston [Thurston], in an article that is largely a rejoinder to that of Jaffe and Quinn [JaffeQ], argues convincingly that formal proofs are fundamental to the correctness of mathematics, but they are a terrible way to convey mathematical understanding. Although this is seldom stated explicitly, implicitly it seems to be well understood by everyone. After all, we do have advanced seminars instead of handing out copies of papers and telling everybody to read them. The reason is that the speaker is expected, by neglecting details, by intonation, and by body language, to convey an impression of the real essence of the argument, which is hard to acquire from the paper. The medium of paper journals and the standards enforced by editors and referees limit what can be done.

While Thurston argues for a more intuitive exposition of mathematical results, Lamport [Lamport] advocates just the opposite. Both Lamport and Thurston feel that the usual standards of mathematical presentation fall far short of the rigor of formal proofs. Lamport feels that our literature would be far more reliable if proofs were more formal. One problem with this suggestion is that it would make proofs even harder to understand. Lamport's solution is to have a hierarchy of proofs, of increasing levels of rigor. However, the current system cannot accommodate this, given the premium placed on brevity.

The ideal system would, as Lamport suggests, have multiple presentations of the results, starting possibly with a video of the author lecturing about them, and going down to a detailed formal proof. Such a system is possible with electronic publishing, given the availability of almost unlimited resources (although it will be a while before video presentations can be included routinely), but cannot be accommodated with our paper journals.

Electronic publishing and electronic communication in general are likely to have a profound influence on how scholarly work is performed, beyond destroying paper journals. It is likely to promote a much more collaborative mode of research. One mode of research that is highly prized in our mythology is that of the individual who goes into a study and emerges a few months or years later with a great result. However, that is not the only mode of operation. In general, team efforts have been increasing, and rapid electronic communication via email and fax has been instrumental in this. Inspection of Math. Rev. shows that the proportion of coauthored papers has increased substantially over the last few decades, and this has also been noted in other disciplines. Given the increasing specialization of researchers, such developments are only natural. Further, collaboration is a congenial mode of operation for many. Laszlo Babai wrote a marvelous article [Babai] that I highly recommend. Its title is a play on words. It is an account of the proof of an important recent result in computational complexity, on the power of interactive proofs. At the same time this article is a description of how the proof was developed through exchanges of TEX files of manuscripts and email messages among a group of about two dozen researchers. There are many such interactions going on, and the electronic superhighway will make them easier and more popular. (It will also create new problems, such as that of assigning proper credit for a work resulting from interaction of dozens of researchers. That is another story, however.)

The following two subsections examine some of the existing methods for information dissemination and their relevance for scholarly publications. All have serious deficiencies, but all have promising features. Subsections 8.3 and 8.4 present my vision of what future systems will be like.

8.1. Netnews

The Net provides several interesting examples of information dissemination systems. Netnews, the decentralized system of discussion groups, is extremely popular. However, not many serious scholars participate. As has been noted by many, unmoderated discussion groups, such as sci.math, are at the opposite end of the spectrum from the traditional scholarly journals. They have been called a "global graffiti board for trivial pursuit" [Harnad1]. They are full of arrant nonsense, and uninformed commentary of the "As I recall, my high school teacher said that ..." variety. They are also beginning to attract cranks. (I was amazed that sci.math survived for many years without any serious problems with crackpots, although that is unfortunately changing.) Most mathematicians who try sci.math give up in disgust after a few days, since only a tiny percentage of the postings (of which there have been over 80,000 so far) have any real information or interest. I have continued reading it sporadically, more from a sociological interest than anything else. What I have found fascinating is that although there are now cranks posting to it (and, what is worse, generating long discussions), and there is plenty of "flaming," as well as the nonsense alluded to above, there are occasional nuggets of information that show up. Sometimes a well-known researcher like Noam Elkies or Philippe Flajolet will provide a sophisticated solution to a problem that has been posted. What is perhaps even more interesting is that every once in a while some totally unknown person from a medical or agricultural school, say, will post an erudite message, giving a proof, a set of references, or a history of a problem. These are not the people I would think of when choosing referees, yet they clearly have expert knowledge of at least the topic at hand.

Reading sci.math also provides a strong demonstration of the self-correcting nature of science and mathematics. The opposite of Gresham's law operates, in that good proofs tend to drive out the bad ones. For example, every few months, the Monty Hall paradox (with a contestant given a choice of three doors, etc.) crops up again, as a new reader brings it up. There is typically a flurry of a few hundred messages (this is netnews, after all, and there is no centralized control and no synchronization) but after a week or so the discussion dies down, and everyone (or almost everyone, since there are always a few crackpots) is convinced of what the right answer is.

(In these days when we hear constant complaints about lack of public interest in science and mathematics, it is interesting to note that sci.math has an estimated 120,000 readers world wide. This is a large group of people who do have at least some interest in serious mathematics.)

Although the netnews model does have some redeeming features, it is not a solution to the scholarly publishing problem. The information content is far too low. There are real gems out there, such as "This week's finds in mathematical physics" series that John Baez posts, but they tend to get lost in mounds of nonsense. Methods are evolving for screening out unwanted messages, such as the "kill" files that keep messages posted by certain individuals or those responding to postings from such people, from ever being seen by the reader. There are also the thread-following readers that enable users to screen out quickly discussions on uninteresting topics. Even that is not sufficient, though. Even the specialized groups, such as sci.math.numanalysis, which are at a higher level then sci.math by virtue of their greater technical focus, are not adequate, as there is too much discussion of elementary topics. A somewhat better solution is that of moderated discussion groups. Dan Grayson runs the sci.math.research group, which is much more interesting for professional mathematicians, because of the filtering he does. It has had over 2,500 messages posted to it. An interesting (and relevant for later discussions) feature of this group is that Grayson (personal communication) spends only a few minutes per day moderating it, and rejects about one quarter of the submissions. Furthermore, he has not been plagued by crank submissions, and many of the submitters of messages he rejects thank him for pointing out their errors and saving them from embarrassment. However, while sci.math.research is a useful forum for asking technical questions or picking up odd pieces of mathematical gossip or strange results, it is more like a coffee hour conversation in a commons room than a serious publishing venture.

There are other kinds of moderated discussion groups that engage in a form of publication that is clearly useful, but would not fit into the traditional paper journal mode. For example, F. Bookstein (posting to vpiej-l, Pub-EJournals, December 2, 1993) oversees a morphometrics bulletin board. Its main function is to provide technical answers to questions from readers. They are seldom novel, as one goal is to avoid mathematical novelty, and provide references to existing results, or combinations of known results. However, they serve an important function, saving researchers immense efforts by providing needed technical advice.

There are many mailing lists that provide some of the services that Bookstein's bulletin board does. Many mathematicians are familiar with the Maple and Mathematica user group mailing lists. They consist of email messages from users with questions and complaints, and responses to those messages from other users and the paid support staff for those symbolic algebra systems. They do not qualify for traditional paper journal publications. Too many are basic questions about simple features of the system. Most are about bugs in the current release, or system incompatibilities, and nobody would want to make that part of the archival record. However, they are extremely useful, largely because with electronic storage, it is possible to search the great mess of them for the tidbits relevant to one's current needs. Further, they sometimes do veer into deep questions, as when a simple query about solving a system of polynomial equations evolves into a discussion of what is effectively computable with Groebner bases.

All the discussion group formats mentioned above fall easily into the informal communication category. However, they already begin to blur the line that Quinn [Quinn] wishes to preserve between the informal and the formal, reviewed publication. Scholars are increasingly relying on these informal methods. The demarcation line is blurred even further by preprints, discussed below. Quinn regards such blurring as pernicious. However, where Quinn sees danger, I see opportunity.

8.2. Preprint servers and directories

Clear evidence that the present scholarly publication system is not satisfactory is shown by the popularity of preprints. Half a century ago, there were practically no preprints, since there were no technical means for producing them. Journals were the primary means for information dissemination. Today, with xerox machines widely available, preprints are regarded as indispensable. Even Quinn [Quinn], who warns of the dangers of rapid publication, is an advocate of rapid dissemination of results through preprints. Many mathematicians feel that preprints have become the primary method of communicating new results. It is rare for experts in any mathematical subject to learn of a major new development in their area through a journal publication. Usually they either receive preprints as soon as they are written, or they hear of them through the grapevine, and then, when the results seem especially interesting, they request preprints from the author or make copies of their friends' copies. This is also true in many other fields. For example, the article of Cisler in [Grycz] quotes a computer scientist as saying "If I read it in a journal, I'm not in the loop." However, different fields have different practices. For example, in most of chemistry preprints are apparently almost unknown. (Those fields have extremely rapid publications with superficial refereeing, of the kind Quinn [Quinn] warns against.) The extent to which preprints are common in any area is likely to affect significantly the evolution of publications in that field.

Preprint distribution can be done, and increasingly is done, via email. It is much easier to write a shell script or create an alias to send 50 copies electronically than it is to make 50 xerox copies, stick them in envelopes, and address them. Still, this does lead to a messy decentralized system where each author has to decide who is to receive the preprints. There are two possible enhancements to it. Either one of them would be a major advance, and either one could become the main method of disseminating scholarly information in the space of a year or so. Either one would be extremely subversive to the present journal system, and could lead to its demise.

One enhancement to the present preprint distribution system is to use anonymous ftp directories, with possible email and WWW enhancements, of the sort described in Section 3. The other enhancement is to have an automated preprint server. There are several of them already operational in mathematics, such as the one at Duke that covers algebraic geometry. They have not had much influence yet. However, this could change suddenly. An instructive example is provided by the system that was set up by Paul Ginsparg at Los Alamos [Ginsparg, Sci1, Sci2]. In only one year, starting about two years ago, the high energy theoretical physics community switched almost completely to a uniform system of preprint distribution that is run from Ginsparg's workstation (and now also from some other sites that copy the material from his). Apparently nobody in theoretical physics can afford to stay out of this system, as it has become the primary means of information dissemination. Even brief interruptions of service [Sci2] bring immediate heated complaints. This system has already been extended to several other fields, aside from high energy theoretical physics, and the requests to help in setting up the system are a major chore for Ginsparg (private communication). The main point about this system is that it is cheap. The software is available for free, and not much maintenance is required. Physicists submit preprints electronically (in a prescribed version of TeX, and in prescribed format), the system automatically files them and sends out abstracts to lists of subscribers (which is maintained automatically), and then the entire papers can be retrieved through email requests. Recently the system was enhanced by putting it in WWW, so preprints can be browsed using Mosaic.

An important observation about Ginsparg's system is that the transition to it was sudden, at least by the standards of the publishing world. It took under a year from the time Ginsparg wrote his program to the time it became the standard and almost exclusive method for distributing new results in high energy theoretical physics. Ginsparg [Ginsparg] attributes the rapid acceptance of his system to the fact that his area had already switched over to a system of mass mailings of preprints as the primary information dissemination scheme, and regular printed journals were of secondary importance. Mathematics is not at that stage. However, this could change rapidly. The use of email and anonymous ftp for distribution of mathematics preprints is spreading. Other fields have switched or are switching to the use of Ginsparg's system, and I would not find it surprising if mathematicians suddenly did that.

Can Ginsparg's system coexist with paper journals? Theoretical high energy physicists are still submitting their papers to the traditional journals. However, it is not clear how long that will continue. If I have heard of an interesting result that has been published in some journal, and if I can order via email a preprint of the same version (except possibly for the formatting) from an automated preprint server, why should I bother to go to the library to look for the journal? If I don't look at the journal, why should I mind if my library cancels its subscription? If there is anything that is certain, it is that my library will keep coming back each year with lists of proposed journal cancellations, so the pressure to give up more and more subscriptions will continue. One solution would be for publishers to require that preprint copies be deleted from preprint archives once a paper is published. I doubt if this approach can work. It is technologically hard to enforce. How can anyone keep track of all the copies that might have been squirreled away on private disks? Even more important, would scholars tolerate such requirements?

Given the evolution of software, the distinction between a centralized preprint server and a decentralized database will soon be immaterial. As long as preprints are available in electronic form and their authors are interested in disseminating them widely, it will soon be easy to locate and obtain them for anyone on the Net with minimal effort by the authors.

Wide distribution of preprints is supported even by Quinn [Quinn]. However, preprints in most areas have a different status from refereed papers, and Quinn argues strongly for maintaining this distinction. My feeling is that there is no way to do this effectively. The temptation to use preprints stems from the desire for faster dissemination of information. If preprints are to be widely distributed, though, we need to adapt the peer review system to accommodate them. At the moment this is not done in a satisfactory way.

It is possible for a field to rely just on preprints. Ginsparg [Ginsparg] writes that in theoretical high energy physics, they have long been the primary means of communicating research results. Papers are still published in conventional journals, but do not play much role in the development of the field, since they are typically obsolete by the time they appear. Peer review operates, but in an informal way, as several physicists in that area have told me. Rob Pike, a colleague working in operating systems, reports (private communication) that in his area journals have also become irrelevant. Communication is via email and electronic exchange of preprints, typically through anonymous ftp. A recent announcement of reports about a new operating system resulted in over a thousand copies being made via anonymous ftp in just the first week. Peer review operates in this area also, but again journal publications do not play much of a role in it (and probably cannot, since the main product is not embodied in an article, but in the software).

It is easy to argue that the experience of operating systems is not relevant to mathematics. However, the same problems are arising in various areas of mathematics. There have been complaints by some mathematicians that their efforts were not getting proper recognition, since what they were producing was software, whether for solving partial differential equations or doing geometric modeling, and journal publications were not the right way to evaluate their contributions. With electronic publishing, this problem can be overcome.

There are even fields close to mathematics that do not follow the usual procedures for mathematical publications. Theoretical computer science has standards of rigor that are comparable to those in mathematics. There are over a dozen annual meetings, with the original STOC and FOCS conferences the most general and prestigious. Papers (or, to be precise, extended abstracts of about 10 pages) to these conferences are submitted about six months in advance. A program committee then selects around 80 out of 200–300 submissions. Expanded abstracts (which are often full papers) are then published in the proceedings, which are given out at registration to all participants, and can be purchased by anyone from the publishers, ACM and IEEE. The official policy is that since the committee is making decisions on the basis of extended abstracts only, and in any case does not have time to do a thorough refereeing job, the proceedings publications should be treated as preliminary. The final revised versions are supposed to be submitted to refereed journals. However, there have been perennial complaints that this was not being done, and that researchers were leaving the proceedings drafts as the only ones on record. Recently, however, David S. Johnson has compiled a bibliography of STOC and FOCS papers, and it appears that around 60% of them do get published in a polished form in a journal. (Moreover, the 40% that do not get published are not by any means the bottom 40%, but include some of the most influential results. It's just that some authors are negligent and do not carry out their duties to their field.) Although terrible, this 60% figure is considerably better than the folklore would lead one to believe. Still, for working researchers, practically all the information dissemination takes place through early preprint exchanges, and interactions at these conferences. In the words of Joan Feigenbaum, in theoretical computer science "if it didn't happen at a conference, it didn't happen." Moreover, acceptance at these conferences is regarded as more prestigious than journal publications, as one can see from letters of recommendation. Thus here also it is possible to have a healthy field that does not depend on journal publications for information dissemination. The reason computer scientists seem to tolerate the present lack of reliability is that the conference system provides them with rapid evaluation and feedback, and also opportunities for extensive personal interaction. However, their situation is not ideal. There are frequent complaints about errors in the proceedings papers, and there have been a few notorious cases where important claimed results turned out to be wrong. What is lacking here, as with other informal preprint systems mentioned above, is a better peer review system, to provide better assurance of reliability. What we see is the seductive influence of rapid communication through conference proceedings that has led theoretical computer science into the errors that Quinn [Quinn] warns against.

8.3. The publication and reviewing continuum

The popularity of preprints, of netnews groups, and of mailing lists shows that they do fill an important role for scholars. We cannot turn back the tide. On the other hand, there is also a need for reliability in the literature, to enable scholars to build on the accumulated knowledge. One way to resolve the conflict is to follow Quinn's advice [Quinn] and rigidly separate distribution of information, such as preprints, from publication in journals. The former would be done rapidly and informally, while the latter would follow the conventional model of slow and careful refereeing, even with extra delays built in to help uncover problems. I feel a better solution is to have an integrated system that combines the informal netnews-type postings with preprints and electronic journal publication. Stevan Harnad has been advocating just such a solution [Harnad1], and has coined the terms scholarly skywriting and prepublication continuum to denote the process in which scholars merge their informal communications with formal publications. Where I differ from Harnad is in the form of peer review that is likely to take place. Whereas Harnad advocates a conventional form, I feel that a reviewing continuum that matches the publication continuum is more appropriate.

I will describe the system I envisage as if it were operating on a single centralized database machine. However, this is for convenience only, and any working system would almost certainly involve duplicated or different but coordinated systems. I will not deal with the software aspects of this system, which will surely involve hypertext links, so that a click on a reference or comment would instantly bring up a window with that paper or comment in it. My main concern will be with how scholars would contribute to developments of their fields. My basic assumption is that in most scholarly areas, articles will remain the main method of communicating specialized information, although they may be enhanced with multimedia features.

At the bottom level of future systems, anyone could submit a preprint to the system. There would have to be some control on submissions (after all, computers are great at generating garbage, and therefore malicious users could easily exceed the capacity of any storage system), but it could probably be minor. Standards similar to those at the Abstracts of the AMS might be appropriate, so that proofs that the Earth is flat, or that special relativity is a Zionist conspiracy, would be kept out. Discussions of whether Bacon wrote Shakespeare's plays might be accepted (since there are interesting statistical approaches to this question). There would also be digital signatures and digital timestamping, to provide authentication. The precise rules for how the system would function would have to be decided by experimentation. For example, one feature of the system might be that nothing that is ever submitted could be withdrawn. A similar policy is already part of Ginsparg's system, so that papers can be withdrawn, but the withdrawal is noted in the record. This helps enforce quality, since posters submitting poorly prepared papers risk having their errors exposed and publicized for ever. On the other hand, such a rule might be felt to be too inhibiting, and so might not be imposed.

Once a preprint was accepted, it would be available to anyone. Depending on subject classification or keywords, notification of its arrival would be sent to those subscribing to alerting services in the appropriate areas. Comments would be solicited from anyone (subject again to some minor limitations), and would be appended to the original paper. There could be provisions for anonymous comments as well as signed ones. The author would have the opportunity to submit revised versions of the paper in response to the comments (or his/her own further work). All the versions of the papers, as well as all the comments, would remain part of the record. This process could continue indefinitely, even a hundred years after the initial submission. Author X, writing a paper that improves an earlier result $Y(123)$ of author Y, would be encouraged to submit a comment to $Y(123)$ to that effect. Even authors who just reference $Y(123)$ would be encouraged to note that in comments on $Y(123)$. (Software would do much of this automatically.) This way a research paper would be a living document, evolving as new comments and revisions were added. This process by itself would go a long way towards providing trustworthy results. Most important, it would provide immediate feedback to scholars. While the unsolicited comments would require evaluation to be truly useful, and in general would not compare in trustworthiness with formal referee reports, they would be better than no information at all. Scholars would be free to choose their own filters for this corpus of preprints and commentary. For example, some could decide not to trust any unrefereed preprint that had not attracted positive comments from at least three scholars from the Ivy League schools.

Grafted on top of this almost totally uncoordinated and uncontrolled system there would be an editorial and refereeing structure. This would be absolutely necessary to deal with many submissions. While unsolicited comments are likely to be helpful in deciding on the novelty and correctness of many papers, they are unlikely to be sufficient in most cases. What can one do about a poorly written 100–page manuscript, for example? There is need to assure that all the literature that scholars might rely on is subject to a uniform standard of refereeing (at least as far as correctness is concerned), and at the same time control the load on reviewers by minimizing duplicate work. Both tasks are hard to achieve with an uncoordinated randomized system of commentary. A formal review process will be indispensable. That is also Harnad's conclusion [Harnad1, Harnad2]. There would have to be editors who would then arrange for proper peer review. The editors could be appointed by learned societies, or even be self-appointed. (The self-correcting nature of science would take care of the poor ones, I expect. We do have vanity presses even now, and they have not done appreciable damage.) These editors could then use the comments that have accumulated to help them assess the correctness and importance of the results in a submission and to select official referees. (After all, who is better qualified to referee a paper than somebody who had enough interest to look at it and comment knowledgeably on it? It is usually easy to judge someone's knowledge of a subject and thoroughness of reading a manuscript from their comments.) The referee reports and evaluations could be added as comments to the paper, but would be marked as such. That way someone looking for information in homological algebra, say, and who is not familiar with the subject, could set his or her programs to search the database only for papers that have been reviewed by an acknowledged expert or a trusted editorial board. Just as today, there would be survey and expository papers, which could be treated just like all the other ones. (Continuous updating would have obvious advantages for surveys and bibliographies.) As new information accumulated with time, additional reviews of old papers might be solicited as needed, to settle disputes. All the advantages that Quinn claims for our present system in providing reliable literature could be provided by the new one.

Harnad [Harnad1, Harnad2] advocates a hierarchy of groups, with individuals required to pass scrutiny before being allowed to participate in discussions at higher levels. I feel this will be neither necessary nor desirable, especially in mathematics. The best structure might vary from field to field. Harnad is a psychologist, and he may be reacting to the fact that the proverbial people in the street all fancy themselves experts in psychology, and qualified to lecture on the subject to anybody. On the other hand, most people disclaim any expertise in mathematics. Therefore I do not expect that systems in mathematics will be flooded by crank contributions. After all, how many cranks will be interested in discussions of crystalline cohomology, or prehomogeneous vector spaces? There is a need to provide protection against malicious abusers of the system, but beyond that restrictions to specific topics might suffice. In a few cases stronger controls might be needed. For example, any attempted proof of Fermat's Last Theorem might attract many crank or simply uninformed comments that would be just a distraction. That is not likely to be a problem for the bulk of mathematical papers.

Will there be any place for paper in a system such as that sketched above? I suspect it will be limited, if there will be any. In paper publishing nothing can be changed once it has gone to the printer. Therefore this system cannot be adapted to provide the opportunities for continuous updating and correcting that is available with electronic publishing. We rely on this system because it was the only one that was feasible in the past. However, we now have better alternatives, and I expect them to dominate. Perhaps the exceptionally important papers will be collected periodically (after selection by a panel of experts) and printed as a mark of distinction. If that is done, then the proposed system will in many ways resemble the one advocated by Quinn. Quinn suggested a minimum six month delay between submission of a paper and publication. The scenario I am sketching would provide a 10 or 20 year delay, and even greater assurance of reliability! For all practical purposes, though, it is likely that traditional paper journals will cease to matter to scholars.

8.4. A possible scenario

Let me show in detail how I would like the future system to function by using an example of a recent mathematical discovery. This will also offer a way to explore the inadequacies of the present system, and what the difficulties of the current preprint and electronic announcement systems are. One may well object that generalizing from pet cases is not reliable. However, it is often preferable to deal with concrete cases, as opposed to talking in vague terms that are hard to interpret.

The example I will cite is not typical. No example could be typical in the immense area of scholarly inquiry, but this one has several unusual features. I use it precisely for that purpose, because it allows me to specify what I would like to happen in some extreme cases that are not likely to apply most of the time.

On March 19, 1993, Brendan McKay and Stanislaw Radziszowski posted an announcement on sci.math.research that they had proved that $R(4,5) = 25$. One reason for choosing this example is that their result can be explained even to a lay person. What McKay and Radziszowski showed is that if there are 25 people in a room, either there is a subset of 4 of them so that any two know each other, or else there is a subset of 5 of them so that no two know each other. Previously it was known only that this was true if there were at least 28 people in the group. This result falls in the area of Ramsey theory, and a basic theorem in that subject is that for all natural numbers m and n, the Ramsey number $R(m,n)$ exists, so that in any group of at least $R(m,n)$ people, either there will be m mutual acquaintances or else n mutual strangers. What Ramsey theory says is that you cannot have complete disorder, that no matter how you arrange relationships in a group of people, there will be systematic patterns. (There is much more to Ramsey theory than this, as one might expect, but the above is a basic and classical result.)

So much for a description of the result. One reason it is unusual is that it attracted attention in the popular press. This then led a lay reader to write a letter to Ann Landers, which she published in her widely syndicated column, asking whether taxpayers' money was not being wasted on such endeavors. This led to a flurry of messages on sci.math, discussing the best ways to respond to the letter and to Ann Landers' request for an explanation of the value of research of the kind McKay and Radziszowski performed. (Unfortunately, while various letters were sent to Ann Landers, none were published in her column.)

How significant is the McKay-Radziszowski result? There is a small part of Ramsey theory devoted to computing values of various Ramsey numbers. However, it is not among the most significant areas of combinatorics. The methods are a combination of mathematics (since raw enumeration of the various possibilities is beyond the ability of any computer that could be constructed in the known universe using the constraints of currently known physical laws) and heavy computation. Judging from the announcement, it appeared that McKay and Radziszowski had extended the state of the art in both mathematics and computing. However, that is about it, and there are no current applications for this result that I know of.

We now take a detour and ask what is worth publishing in mathematics. A certain eminent mathematician said recently that

> "The only thing I care about are the 100 really interesting, and 400 interesting, papers in mathematics published each year. The rest is trash good for the waste-basket, on a par with junk mail."

The McKay-Radziszowski paper (of whose existence I learned by contacting the authors, since none of the announcements on the Net that I had seen mentioned it) does not qualify by most experts' standards for the most interesting 500 papers written in mathematics in 1993. However, I feel that this eminent mathematician is wrong, and that papers such as that of McKay and Radziszowski should be published (as this one will be, in *J. Graph Theory*). The reason is that I, along with many other mathematicians, scientists, and engineers, often find myself looking for some specific result, such as the value of $R(4,5)$. I don't particularly care how hard it was to derive the answer. If it's not a triviality that I should be able to see for myself in a few minutes, I am glad to find it in print, even if I could have done the work better and faster than the author (something that is surely not true here). The point is that I want the answer for some specific reason, and don't want to take a day or a week on a detour from the project I am working on to figure out the answer for myself. (A major strength of the mathematical literature, which Quinn [Quinn] stresses, is precisely that by providing reliable results, it makes this mode of operation possible.) What I want from the Net is the tools that will let me locate such answers quickly, and the editorial/refereeing system to provide me with the assurance that I can trust what I find.

I should stress that I do not have an application for the McKay-Radziszowski result right now. However, I can cite numerous examples where I or my colleagues have used similar results found in the literature. In those cases we did rely heavily on the reliability of the published papers. If I design a code that is supposed to correct certain types of errors, I need to be able to guarantee that it will do so. If disaster strikes, and some product fails because some of the errors that were supposed to be corrected slip through, it won't do to say "Well, I saw this claim on sci.math, by some chap I have never heard of, that ...."

How could the McKay-Radziszowski claim be disproved? One way would be to exhibit a counterexample. If you find a collection of 25 people such that no 4 are mutual acquaintances, and no 5 are mutual strangers, then this will be it. To exhibit such a counterexample, you only need to specify for each of the $25 * 24/2 = 300$ pairs whether they know each other or not. That is a small amount of data. For any specification, there are efficient algorithms that will tell whether your collection has the desired property. How you find the counterexample is immaterial, and it could be something that a grade school pupil came up with. A more likely event is that no counterexample is found, but the McKay-Radziszowski proof is discovered to be faulty, either because the graph theoretic arguments are wrong, or because the computations are wrong.

The reader might well ask why am I taking so much time to explain things about a paper that is of little general interest. I do this to be able to discuss the question of reliability that is attainable with the present system. A hundred years ago, McKay and Radziszowski would have submitted their paper to a journal, and after a peer review a revised version would have appeared in print. For almost everybody, the printed version would have been the first one they would have seen. What assurance would those readers have gotten? If it was a prestigious journal, they would have assumed that at least one competent expert had looked at it and pronounced himself satisfied that it was correct. That is also what readers of current journals get. However, both then and now, there would be no indication of the care taken in the refereeing process. For McKay and Radziszowski, there is both the mathematics and the computing to consider. If one were to see it in print, it would be easy to imagine that the referee or referees checked the graph theoretic arguments that reduce the problem to manageable size. But how carefully was the mathematics checked? There would be no indication of that. What about the computational part? Programs to check problems of this size often take thousands of lines of code. Did the referees check those? It is uncommon for authors to submit program listings with papers. Even if a program listing is available (and according to the McKay-Radziszowski announcement, each author wrote an independent program, and both programs were run with the same outcomes, which I find commendable), how carefully was it scrutinized? Was the program available electronically, so the referees could run it themselves? Did the referees run it? (Seems unlikely, since the announcement mentioned a cluster of workstations that were used, with total run time of over 3 years for a single workstation.)

The aim of the above discussion is to point out that conventional peer review is not too reliable. Moreover, it is not serving the needs of the scholarly community in a timely fashion. In the discussion above, I said that 100 years ago, the first time that almost all people would have seen the paper would have been in the journal. However, we cannot roll back the clock. Researchers will not be denied the interactive potential of the Net. Scholars such as McKay and Radziszowski will keep announcing their results on the Net. However, even without the Net, we would have had problems, caused by the interactive potential of the mail and the telephone. Ever since the invention of xerography, preprints have been proliferating, and have become the dominant method of disseminating information among experts. Now what is the status of a preprint? Just how much can it be trusted? One answer, which seems to be the one advocated by Quinn, is that we should not trust a preprint at all, and should regard it as being in a form of purgatory until it undergoes proper refereeing. However, that is not what is happening. Let me stress that most of what I am saying about publications is descriptive, not prescriptive. I am saying

> "Here is how these fields are developing, and these are the reasons that drive these developments. Let us try to anticipate where this is likely to lead, and whether we can introduce some small modifications on the natural evolution of the system to make it better for scholars."

and not

> "Here is what I think the ideal system would be like, so let us force everybody to adhere to its standards."

What I see is that the current peer review system is not doing its job. That is why I discussed theoretical computer science at some length. There we have a peer review system in operation, but it is not adequate. Similarly, preprints are acquiring the status of official publications. One might say that they should not. But that is contrary to what is becoming accepted practice (a descriptive and not a prescriptive statement again). Is it rare in your field to see preprints cited? It is becoming common in mathematics, physics, computer science, and engineering, the fields I know at least a bit about. (That is not true in all fields, though. In chemistry, for example, they are seldom used.) Preprints are becoming, even if they are not already, the primary information dissemination method. Authors are sending them out to large mailing lists. However, they are not sending out corrections or retractions, except in rare cases. Andrew Wiles posted a message on sci.math in December of 1993, acknowledging that there is a gap in his proof of Fermat's Last Theorem. However, that is rare, and tends to occur only with papers that attract wide interest. To cite a concrete example, I am now handling a paper as editor of a prestigious journal. It is by a person I will call X, and if correct, it is a breakthrough and will undoubtedly be published. However, the two referees that I sent it to quickly spotted a gap in the proof. Revisions have been plagued by problems as well. However, the original manuscript was sent out to many people, and then was passed on to many others by the original recipients. It is still circulating, and I am beginning to see references in other preprints to X's great result. The existence of a gap in the proof is not widely known. This is the reality that we have to cope with, and we need to come up with concrete proposals that will ensure the needed quality control extends to the literature that scholars are working with.

So far I have complained at length about the inadequacies of the present system. Here is what I expect to see in the future. It is not an ideal system, but one that I think is realizable, provides adequate standards of reliability and timeliness, and conforms to the natural evolution of the Net and of scholarly publishing. Let us consider the McKay-Radziszowski case, and how it might be handled in this future system. They make a formal claim to have proved that $R(4,5) = 25$ by submitting it to a lightly moderated newsgroup like sci.math.research. By lightly moderated I do mean something similar to what Dan Grayson is doing now, namely a check that the posting is about mathematics and is not obviously crazy, but nothing more than that, certainly not a verification that the result is new. This could lead to some discussion, with somebody pointing out the result is not new, or figuring out a new application. Everything would still be within the bounds of a weakly moderated discussion group, and most serious researchers would not be paying any attention to it. Those interested in Ramsey theory might have set up their filters to pick up any announcements that mention such keywords, and this might draw them into the discussion on sci.math.research (or a more specialized sublist). Ann Landers' popular column and discussions of whether taxpayers' money should be used to fund the determination of $R(4,5)$ would be kept to other forums.

The next step comes when McKay and Radziszowski produce a paper. I would want them to submit it, together with their computer programs, to an editorial board. (It would be one of many such boards, I expect.) Boards would be appointed in the same way they are now, either by a learned society, or by previous boards that have established a good reputation. This board would then place the preprint and computer programs in a publicly accessible directory (or on an automated preprint server, it does not matter what we call it, since the evolution of the Net will make all these things look the same to users). There would be a link to the earlier sci.math.research posting, so that the readers of sci.math.research could follow the thread to the preprint, and the readers of the preprint service could go back and see when the initial claim was made, and what kinds of discussions it generated. The editors would now invite comments, which would go on the record. Controls on the comments could again be light. There are important questions about the precise rules that I do not think can be prescribed in advance. Would anonymous comments be allowed? Should some provision be made to allow comments to be withdrawn, or at least pushed into the background? (For example, if I say that a particular result reminds me of something I have seen in a German journal of the 1920s, and somebody says that the result in question was published by Emmy Noether in *J. reine angew. Math.* in 1924, it's clear whose comment should carry more weight.) However, these comments could come from a variety of people. For the McKay-Radziszowski paper, even college students could contribute. ("I took their programs and ran them in the spare time on our workstation cluster, and got the same results they did", or else "I could not get their program to compile on my machine, and the compiler error messages suggest there may be a problem with their use of the following data structures.") Usually, of course, the papers would be beyond reach of college students, and even the graph theory in McKay's and Radziszowski's paper would require moderate expertise. However, such expertise is not all that rare, and with some encouragement one could expect a nonnegligible fraction of papers to attract serious scrutiny. After all, we do produce many Ph.D.s who then go into jobs that do not involve research on their part, but who often still have extensive expertise in their fields and often the inclination to follow new developments. It is true that most papers attract almost no attention even when they are published in journals. However, the potentially significant ones are much more likely to attract expert readers, and that is after all what we need.

The editors would at some time ask for official referees. They might do it right when the preprint is submitted, or they might wait 6 months to see what attention and comments the paper attracts. I am not sure what the best policy would be, and it might vary from paper to paper. The referee reports (possibly after some editing) would be added to the record as the process went on, so that in the case of X that I cited above, I would make public edited comments on the first version, and enclose a note that X is working on a revision. Finally, after the referees became convinced of the correctness and novelty of the work, this would be noted, and some mark to designate the degree of confidence the reviewers have in the correctness of the work would be made. Some semi-quantitative evaluation of the significance might also be provided. (Since space would not be a limit on the Net, I expect that separate journals with different prestige levels will be replaced by specialized electronic journals that will formally grade papers on their significance. However, that is not the only way to run things, and I can conceive of other procedures.) Later, as new information comes in, it would be attached to the paper. For example, if a mistake is found in the McKay-Radziszowski analysis a few years after the formal peer review process is completed, this would be noted, as would all future references to this paper, reviews, and so on. Readers could easily set up their search programs to pick out for them only those papers that had been formally reviewed, or had attained a high significance rating. The level at which readers would be notified would vary from field to field, depending on the reader's interests. I would set up my filters to alert me only to the real breakthroughs in low-dimensional topology, for example, but to tell me even of sci.math.research notes on public key cryptography.

The proposal above is designed to work within the confines of what we can expect both technology and ordinary fallible people to accomplish. It would integrate the roles of authors, casual readers, and official referees. When I reflect on my own papers, and mentally compare the preprints with the final journal versions, all the differences were due either to my own work, or to comments from colleagues, or to the referees. While I can't precisely quantify these three factors, the contribution of the referees was the smallest. The main advantage of my proposal is that it would provide a continuum of peer review that more closely matches the publication continuum that is likely to evolve on the Net.

9. The future of publishers, journals, and libraries

It is impossible to predict the date or speed of transition to systems like the one outlined in the previous section, but only because they will be determined primarily by sociological factors. The technology that is necessary for future systems is either already available or will be in a few years. The speed with which this technology will be adopted by scholars will depend on how quickly we are prepared to break with traditional methods in favor of a superior but novel system. For example, how quickly will tenure and promotion committees start accepting electronic publications as comparable to those in traditional journals?

The conversion to electronic publication would be easiest if we started with existing journals, made them electronic, possibly initially only as an alternative, keeping paper copies going for a while, and then proceeded to modify them. So far little has been done in this direction. Ann Okerson, in a Sept. 5, 1994 message to the Australian electronic publishing discussion list, aus-epub, described the reasons that print journals have difficulty in converting to electronic formats:

> "... Existing journals have a certain kind of process, have workers used to certain types of editing and production, have readers and writers who are accustomed to a certain way of doing things. They have a certain kind of management and mindset, perhaps. The new medium asks for a different skill set in production, a willingness to risk in a new format with few rules and standards, and readers who are agile on the Internet or can be persuaded to be. Moving existing titles into an electronic base and distributing them requires re-tooling and re-thinking of the procedures and outcomes of publishing and can be financially risky. ..."

Therefore it appears more likely that new electronic journals will pave the way. As a result of the extensive work by Neil Calkin and Herb Wilf, the founding editors, the *Electronic Journal of Combinatorics* (*ElJC*) currently has the most advanced features among the mathematical electronic journals, and more developments in the directions sketched here are likely to start showing up soon in *ElJC* and other journals. However, once print journal subscriptions start declining rapidly, one can expect that many traditional journals will convert to electronic form.

Although sociological factors will be most important, there are also various other technological and legal problems that will have to be faced. Some of the familiar institutions are going to be changed drastically. In this section I will discuss some of these issues.

9.1. Spending on new technologies

One point that deserves mention is that total spending by universities on computing and telecommunications is likely to increase. The requirements of systems that would replace our present scholarly journals can be met at a cost that will be rapidly decreasing. However, scholars will surely behave just like everyone else, and as new services become affordable, will use them. The recent dramatic price cuts in PCs have led to sales of more computers, and sales of more powerful computers, and thus an increase in total spending on PCs. Similarly, as cheap high speed networks proliferate, scholars will start transmitting videos of lectures, massive amounts of data, and so on. Therefore the savings in library budgets will have to be used partially to pay for an improved electronic infrastructure.

As one small example of the demands that are already arising, let us recall that a typical mathematical paper requires only 20 KB of storage in compressed TeX format. However, my colleague David Applegate has made available (using the system described in Section 3) proofs of optimality of certain traveling salesman tours that require about 20 MB of storage in compressed form. These proofs are generated with the assistance of a computer, and are only made to be verified with the use of a computer. Still, they are a novel type of publication and presage the demands that will soon arise.

9.2. Standards

Standards, or lack of them, can be a significant impediment to the adaption of new technologies. Several readers of earlier drafts of this essay commented that it was much easier to locate printed references than electronic ones. This deficiency will undoubtedly be overcome, possibly through wide use of URLs, but it is likely to be bothersome for a long time.

Mathematics, computer science, and physics all seem to have settled on TeX and its various dialects as the de facto typesetting standards. This makes it easier for these disciplines to move into electronic publishing than it is for others that have not converged on a solution. However, while TeX is adequate for almost all current papers, it may not suffice in the future as we move into a multimedia world. It is also possible that commercial packages such as Microsoft Word will be enhanced with addition of modules to handle scientific material, and may become the prevalent tools. (The era when scientists dominated electronic communications is coming to an end, and systems such as TeX, developed by scholars for scholars, might soon be eclipsed by general purpose packages.) We should not become too committed to any particular standard, as it may be transitional. My current preference is to make available both PostScript versions of my papers, to assure portability, and the TeX source files, to make text searches easier. (Provision of just PostScript is not much defense against plagiarism, since it is not that difficult to derive TeX code from PostScript, and in any event, wholesale conversion to electronics will make detection of plagiarism much easier through software tools.)

9.3. Books and the role of paper

While I do predict that most paper journals will disappear in the next two decades, paper is likely to persist far longer. Display technology is not moving anywhere near as fast as microprocessors or data transmission, and it will be a long time before you curl up in bed with a screen rather than a book. Eventually we are likely to have high-resolution displays that are light. At that stage paper will be truly obsolete. In the meantime, the dependence on paper will vary from person to person and from field to field. Most of the work on this essay was done at the keyboard, and it was much easier to handle the references that were available electronically than paper ones. The ability to search quickly, move large segments of text, and annotate were decisive. On the other hand, whenever I need to understand a hard mathematical proof, the kind that it takes a day or two of hard work to absorb, I need to work on it in the optimal setting, which in my case is sitting in my armchair, with a paper copy of the manuscript and a large pad of paper for the notes and calculations that are needed. (Electronic versions of papers are useful even then, though, since I can print a copy that has extra spaces in the crucial sections for my comments.)

Paper is likely to persist for a few decades. However, I expect its main application in scholarly communication to be for temporary uses, when scholars print out copies of papers they wish to study. Archival storage ought to be electronic, for compactness and ease of access.

So far this essay has dealt exclusively with scholarly journals. However, the same trends are going to influence book publishing. Books come in all forms, from specialized monographs to advanced textbooks, to college textbooks, to the popular press. They will be affected by technology in different ways. I expect that they will all be displaced by electronics eventually, but the displacement will start with journals (soon), then move on to monographs, and only much later will it reach college textbooks. Further, the change to electronic publishing of college textbooks and the popular press will be affected by their different economic situation, since there the authors do expect to earn substantial financial remuneration. (See [Rawlins] for a discussion of a possible future of book publishing.)

Although books will change less slowly than journals, some changes are likely to show up soon. For example, we are likely to see soon textbooks customized for particular courses. An instructor will choose from a publisher's catalog what sections are to go into the students' textbook. The students will then go to their school's bookstore and have copies printed for them on the new machines that are becoming available.

9.4. Costs of present system

To understand the present and future of publishing, we need to have a clear picture of the costs involved. This involves both the explicit costs, such as those for journal subscriptions, where money that is allocated for publications changes hands, and the implicit costs, such as those for the time of authors and referees.

What is the cost of producing a typical mathematical paper? It appears that the average researcher publishes two or three papers per year. The total cost of employing such a person is at least $150 K per year (this is a conservative estimate, over twice the average salary of the average mathematician in the US, as it is meant to include standard salary, grant support, benefits, as well as all the office space, libraries, and university administration costs). Let us assign one third of this cost to research activities. If we do that, we conclude that each paper costs at least $20,000, and this cost is born by taxpayers, students' parents, or donors to universities.

The reviewers of scholarly papers are almost uniformly unpaid, and so are most editors. It is hard to estimate how much effort they devote to a paper. It appears that about half of the papers in mathematics that are submitted are accepted at a typical journal. Since many submissions require several revisions and extensive correspondence, and increasingly papers are sent to two referees, it seems reasonable to estimate that between one and two weeks' time by the editors and referees is devoted to their jobs for each accepted paper. Thus the value of their time is around \$4,000.

In mathematics, there are two main review journals, *Mathematical Reviews* (*Math. Rev.*) and *Zentralblatt für Mathematik* (*Zbl.*). Both rely primarily on unpaid outside reviewers. If a reviewer spends a day preparing a review, reading the paper, locating additional references, and so on, then the implicit cost is around \$500 for each review journal, or \$1,000 for each published paper.

We next turn to the explicit costs of scholarly publishing. How much does it cost society to publish a paper in a research journal? For an estimate, I will use the figures in [AMSS]. (The only comprehensive studies that I was able to find, in [KingMR] and [Machlup], are over 15 years old, pre-dating the arrival of electronic typesetting, and so are of interest, but of little current relevance.) If we assume that a paper is typeset with 50,000 characters, and multiply that by the cost per character given in [AMSS] for any given journal, and then multiply by the circulation for that journal, we obtain an estimate for how much it costs to publish at article in that journal. For example, according to [AMSS], *Amer. J. Math.* has a circulation of 1458, and an annual subscription costs \$0.048 per 1,000 characters. This produces an estimate of $\$1458 * 50 * 0.048$, or about \$3,500 for the cost of publishing a single article there. This figure includes all the editorial, printing, and mailing expenses. Doing this for the other journals listed in [AMSS] for which both costs and circulation are given (but excluding *Bull. Amer. Math. Soc.*, which differs substantially in scope and especially circulation from the standard research journals) produces estimates for single article costs between \$900 (for the *Notre Dame J. Formal Logic*) and \$8,700. The median cost figure is about \$4,000, and is the one I will use.

For comparison, *Physical Review Letters* (*PRL*) has about 2,500 institutional subscriptions at about \$1,500 each, plus 4,500 individual subscriptions at \$140 each, which together with page charges produce annual revenue of about \$5 M. Since *PRL* publishes about 2,400 out of the 6,000 submissions it receives each year, the cost per published paper is about \$2,000. Since *PRL* publishes only short papers (extended abstracts), this cost is surprisingly high. It is presumably caused by *PRL* having several paid full-time professional editors.

Review journals are comparatively cheap, if one considers the cost per paper reviewed. *Math. Rev.* has annual revenue of around \$5 M (about 1,000 subscriptions at \$5,000 each), so the explicit cost for each of the 50,000 reviews published each year is \$100. The figure for *Zbl.* is presumably similar. For comparison, *Chemical Abstracts* publishes about 500,000 reviews per year at a total cost of \$150 M, for a cost per review of \$300. The much higher monetary cost of these reviews than those in *Math. Rev.* is probably caused by having a paid in-house staff prepare them instead of relying on the unpaid help of outside scholars.

Most of the cost of producing traditional scholarly journals is in the processing of the manuscripts (the "first-copy" cost [Grycz]), and little in printing and distribution. The *PRL* individual subscription price of \$140 per year is set to cover the marginal cost of printing and mailing one additional copy of each issue, and the bulk of the money collected from libraries goes for all the editing and overhead costs. Similar figures seem to apply to other publishers.

In general, even approximate cost estimates are seldom released, as they are regarded as proprietary. One audited accounting that is available is that for *Physical Review B* (*PRB*) [APSPRB]. In 1993, the revenue for this journal was \$7.3 M (which is atypical for scholarly journals, which are usually far smaller). Of this amount, \$1.4 M was profit, and expenses were as follows:

editorial	\$1.6 M
composition and production	\$2.3 M
printing and distribution	\$1.6 M
other	\$0.3 M
total	\$5.9 M

(The total figure is not exactly the sum of the individual items because of rounding.) Thus of the total cost of this journal of \$7.3 M (this is the cost to society in general, and is the relevant one to use for our purposes) only \$1.6 M, or 22%, is the cost of producing the physical copy and distributing it. The editorial costs at *PRB* cover the expenses of over half a dozen professional physicists who work as editors, handling the 6,000 manuscripts that are submitted each year.

Profits in scholarly publishing are common. Even non-profit organizations such as the APS make substantial profits from journals such as *PRB* , and these profits are then used to subsidize other activities. Such organizations therefore have a vested interest in the preservation of the present system. Even some of the small scholarly journals published by universities are often not truly non-profit. They may not produce monetary revenue, but they are often exchanged for copies of other journals, thereby saving the costs of subscribing to those journals.

Although profits from scholarly journals do exist, the primary reason for the high total cost of these journals is that the first copy costs are high. Those costs are high both because of the overhead costs of running a publishing organization, and because of the variety of specialists involved in typesetting a paper, editing it before and after typesetting, proofreading, and so on. The outcries against profit-hungry commercial publishers have only some basis in fact. Costs are high even at nonprofit publishers, and they are high primarily because of the skilled labor that goes into producing today's journals.

It is helpful to estimate other costs of publishing scholarly papers. Typesetting mathematical papers costs between \$10 and \$20 per page in the US, depending on whether one counts all the overheads of a fully loaded salary or considers just the cost of employing a part-timer on an occasional basis. Therefore the cost of preparing a typical 20-page paper is \$200–400. (When this same paper is typeset by the author, the implicit cost is likely to be \$1,000–3,000, both because of higher wages and of lower speed, but the comparison is not relevant, since scholars who do their own typesetting mix it with the basic composition of the manuscript.)

While few editors of mathematical journals are paid, most have secretarial support, supplied either by their home institution or paid by the publisher. Based on data from Walter Gautschi, the Managing Editor of *Math. Comp.*, and Andrew Appel, the Editor-in-Chief of *ACM Trans. Progr. Languages*, it appears that it is possible to have an editorial assistant that handles all the correspondence involved in editing a journal for \$100–200 per published paper. Moreover, these costs are likely to decline as more of the correspondence moves to electronic media and better software is developed.

As a final remark, costs of scholarly publications also include costs to readers, namely the time they spend scanning or studying a paper. One estimate of total costs for 1977 (Table 8.3 in [KingMR]) was as follows:

authors	12%
publishers	14%
libraries	10%
users	64%

(The author costs in [KingMR] include only the costs of preparing the manuscript, and do not include the costs of the research that was required, and are thus not comparable to the $20,000 per paper figure that I use.) Because user costs are so much higher than any others, there is real value in publishing high quality papers that are easy to absorb.

9.5. How much should journals cost?

Is $4,000 per article too much to pay for a scholarly journal publication? This sum is only a small fraction of the $25,000 that doing the research, writing the paper, and having it reviewed cost. However, it is an extraordinarily high sum if looked at from another point of view. If indeed only 20 scholars read the typical paper, then this means the cost of each reading is $200. How many scholars would not flinch if, on approaching a library shelf, they had to insert $200 into a meter for each article they wanted to read, even if the money were coming from their research grant or their department's budget? Line-by-line readers of papers are not the only constituency for journals, of course. Tremendous value is derived by scholars from skimming articles to learn what has been done and approximately how. Many more papers are skimmed than read. However, even if the typical article is skimmed by 200 scholars, then the cost per article is still $20. (Estimates in [KingMR] and figures on paper retrievals from Ginsparg's preprint server do show that papers on average are looked at by several hundred scholars.) How many scholars would be willing to pay that, if the cost were stated explicitly this way? Of course, they or their institutions are paying this sum, but the cost is concealed in separate budgets. This high cost, and scholars' general unwillingness to pay such astronomical prices, is likely to doom any efforts to have pay-per-view in scholarly publications, at least with present prices. Charges sufficient to recover current high costs would deter readers.

Another way to look at the cost figures is to consider the total cost of journals. Since about 50,000 mathematical papers are published each year, the total cost of traditional mathematical journals is about $200 M per year. If we assume that 35% of this cost is paid by subscribers in the US (which is probably a low estimate), then we find that US universities, laboratories, and individuals spend $70 M per year for mathematical journals. That is almost exactly the same as the NSF budget for mathematical research. If we could eliminate this cost, we could potentially double the NSF budget for mathematics at no extra cost to society at large. (If we eliminated all the other library costs by converting even old publications to electronic formats, we could save another $70-150 M per year in the US, which could be used to increase research funding even further.) Of course, the tradeoff is not that simple, since journals are paid for from different sources than research grants, and we do not have much freedom in moving public and private funding around. However, the \$70 M figure for just journal subscriptions is sufficiently large that it cannot easily grow, and there will be pressure to lower it if some method for doing so can be found. A significant influence of new electronic technologies is likely to be to focus attention on the high costs of scholarly publishing.

To what extent can publishing costs be lowered? Some publishers have predicted that electronic publication of journals would lower the costs only by around 30%, which is only slightly more than the cost of printing and mailing. However, that assumes that the current editing procedures are followed. It also ignores the enormous disparity in costs of different journals. Some journals are produced much more efficiently than others, which shows that one can lower costs even with print journals. If one lowers the standards, then even traditional paper journals could be produced at a fraction (between one tenth and one quarter) of the present \$4,000 per article. After all, one can xerox author-prepared manuscripts, staple them together, and call that a journal issue. The present scholarly publishing system has evolved into its present high-cost state only because of its strange nature. The ultimate consumers, the scholars, do not pay directly for the journals, and are seldom even aware of the costs of the publications they use. They have never been presented with a range of cost and quality options and asked to choose among them. The publishers do not compete on price, but engage in "monopolistic competition," [Grycz], in which different journals present incomparable material, and strive to be the best in a specialized area. This system has many features in common with the US medical establishment, another producer of high-quality but also extremely high-priced services. If there were a single US federal agency responsible for scholarly publications, there would undoubtedly be Congressional hearings featuring outcries against the \$200 cost of each article read by a scholar instead of against the \$500 flashlight for the military. There would be questions whether scholars were not using Cadillacs where Chevys (or even bicycles) would do. Personally, I do like the high quality of present journals. I have published many papers in them, and expect to publish many more before they disappear. However, I suspect that society will not be willing to continue paying the price for them.

A few years ago, drastic decreases in the costs of journals would have meant going from Cadillacs to bicycles, with journals consisting of stapled collections of mimeographed copies. However, with the advances in technology described in sections 3–5, we can now easily move to something that is at least at the level of a Chevy in luxury, and in addition has the cross-country capabilities of a helicopter. One solution is to transform current journals into much cheaper electronic ones. Eliminating printing and distribution would by itself save 15–30% of the present costs. However, much larger savings should be possible. This would require reengineering the entire publishing enterprise to eliminate whole layers of specialists. This has been done in many other businesses, Recall that keyboarding a paper costs only \$200–400, and is currently mostly provided by the author's institution or is done by the author. All the correspondence about the paper can be handled by an assistant for \$100–200 per paper, and this cost is likely to decrease as communications moves further towards email. It should be possible to provide editing assistance for \$200–600 per (already typeset) paper that would achieve reasonable quality. What would be lost? Many of the features of the existing system would be gone, as a typical paper might be processed by just a single editing generalist who would combine many of the roles of today's editors, copy editors, and proofreaders. Some of the uniformity of appearance of papers in a journal might be gone. Would that be a great loss, though? Should not the unit of scholarly publication be the individual paper, and not the journal issue? For paper publications, it was natural to bundle them into larger packages to achieve economies of scale in distribution and storage. Most of the time, though, a scholar reads or even skims only a couple of articles per issue. Since most of the literature searching involves moving between different journals with different formats, why bother to keep uniform style in each journal? A uniform style of journal references also contributes to the quality of present publications. However, just how valuable is it, and how valuable will it be in the future, when each reference might have a hypertext link to the paper being referenced, or at least something like the URL address?

My general conclusion is that it should be possible to publish scholarly journals electronically for well under $1,000 per article, and probably under $500, without losing much quality. This agrees with Harnad's contention [Harnad2] that electronic scholarly publication should not cost more than one quarter of what it does now. The cost of reading a single article is still going to be so high that pay-per-view schemes will be impractical. However, various of the site-license models discussed in [Grycz] should allow for recovery of these costs.

Can explicit costs be lowered even further? One approach that is already widely used among existing electronic journals is to provide free access, with all the labor involved in running them performed by scholars. My feeling is that this model is likely to predominate. This will mean creating some additional costs for the authors, editors, and their institutions, but those are not likely to be large.

Publishers and librarians often scoff at the idea of scholars being their own publishers. However, they appear to be underestimating how easy that is with the recent advances in technology. Editors and referees already put about as much effort into running scholarly journals as do the publishers. The additional work needed to publish an electronic journal is slight, especially if we relax standards about appearance. (It's worth emphasizing that we are not talking of changing the peer review standards. Those are maintained by the unpaid editors and referees anyway, and can be maintained at the same level in free electronic journals.) The example in Section 3 of the system that my colleagues and I use for preprint distribution shows just how easy it is to operate a rudimentary electronic journal.

Will scholars accept the quality of papers produced by other scholars without the help of the skilled professionals that publishers provide? My feeling is that they will, especially since the gap between what an author can produce and what publishers provide is steadily narrowing with advances in hardware and software. Most of the communication among researchers in many fields already takes place through author-prepared preprints. As further evidence, consider the increasing numbers of books that are produced by publishers taking the authors' PostScript files and printing them without any editing. Also, the proceedings of theoretical computer science conferences, such as FOCS and STOC, are produced from author-prepared manuscripts, and their quality has been improving over the years, and is not that much inferior to that of journal papers. Authors have to write their manuscripts anyway, and with modern tools this is becoming as easy to do at the terminal as on paper, even for technical material. Most of the high cost of traditional publishing is caused by the need for communication and cooperation among the many experts involved in the process. With modern technology, doing something is becoming easier than explaining to another person what to do.

Some of the electronic journals in mathematics operate with some subsidies, and warn that they might eventually start charging subscription fees. However, it is not clear whether such fees will be necessary. The hardware is already cheap and getting cheaper so that no special computers for electronic journals should be required. Ginsparg's electronic preprint service, which was described earlier, uses only a small fraction of the capacity of a workstation that is devoted to other purposes. The netlib system, developed by Jack Dongarra at Argonne and Eric Grosse at Bell Labs, which has been used widely by the applied mathematics specialists for several years, ran for a while on a Sequent multiprocessor that was donated just for that purpose. Today, although the traffic through the system has increased tremendously, a fraction of a workstation handles it comfortably. Thus equipment costs in the future should be absorbed in the general overhead of providing a computing environment for editors in their normal professional jobs. Networking is currently supported by federal subsidies and general university overhead expenses. With telecommunications going down, this should remain true.

While *ETNA* has received financial support (most of which has been spent on equipment), *ElJC* and UQ have from the beginning operated without any financial subsidies. They are operated from their editors' computers. Moreover, *ElJC* was set up in the space of a few weeks, primarily through the efforts of the Editor-in-Chief, Herb Wilf, and the Managing Editor, Neil Calkin. While substantial software work was done, the main reason for the rapid startup was that so much of the software necessary to run the journal was already freely available on the Net.

All the present electronic journals in mathematics operate the conventional peer-review system, and maintain standards comparable to those of conventional journals. Where they differ from most paper journals is in the preparation of manuscripts. Arden Ruttan (private communication) reports that the main negative comments about *ETNA* have been about the lack of copy editing, and the requirement that all manuscripts be in LaTeX or TeX. The restriction on input format is something that is likely to become less of a problem, in that I expect that mathematicians will have almost all their papers prepared in one of a few typesetting systems (which may include something other than just LaTeX or TeX). The big question about the form of scholarly publishing in the future, it appears, is what will happen to manuscript editing (by which I mean all the different copy editing, proofreading, and related functions). At *ETNA*, Ruttan does extensive editing himself. To what extent will scholars be willing to do this? There are editors of paper journals who also devote extraordinary amounts of time to their jobs. For example, Walter Gautschi, the Managing Editor of *Math. Computation*, checks meticulously all accepted manuscripts. He has occasionally picked up mistakes that I had missed in my own papers, which are far outside his specialty. He devotes about half of his time to the journal, far more than most editors are willing to spend. Can we find enough scholars willing to do for free the editorial work that Gautschi and Ruttan do? Even Gautschi has a part-time editorial assistant (paid for by the AMS, the publisher) to handle all the correspondence of the journal.

Given the pressure to reduce costs, there seem to be two main possibilities for future scholarly journals.

(A) There may be no manuscript editing beyond what the authors and editors provide, which usually will probably not be much. Then the journals can be free and available to anyone.

(B) A modest amount of manuscript editing, costing \$300–1,000 per paper, is provided. There are two ways this can be arranged.

(B1) The editing might be done by publishers, either commercial companies or nonprofit professional societies, who will then recover the costs from users. This means that some subscription fee will have to be collected. This will force restrictions on access to the information, most likely through some form of site licenses.

(B2) The editing might be done either at the author's institution, or else by publishers but with costs covered by the author or author's institution through page charges. Then the journals could also be distributed without restrictions.

It is hard to predict which model will be dominant, and it is possible that all may coexist. I share the opinion of Harnad [Harnad2] and many other scholars that (B2) is likely to be the most popular. Notice that the estimated editing costs per manuscript are only about double the costs of having the manuscript typeset, a service that most institutions already provide for their scholars. The main advantage of models (A) and (B2) is that they provide for free unlimited distribution. Model (B1) requires all the cumbersome machinery of copyright transfers, access restrictions, and fees, which is costly to operate and prevents the widest possible dissemination of publications, which is the typical scholar's goal.

While model (B2) does have its attractions, it might not dominate. Scholars' institutions might not be willing to provide the necessary support, and model (A) might not yield satisfactory quality, as scholars might not be willing to put in the necessary effort themselves, either as authors or as editors themselves, either as authors or as editors. An instructive example is that of computer algebra systems. There are many outstanding public domain systems available that were written by scholars, such as UBASIC, PARI, GAP, and Macauley. However, the most widely used packages are Mathematica and Maple, which are sold by commercial organizations. The license fees that are collected are used to provide support in the form of steady upgrades, bug fixes, and advice to users, services that scholars are neither willing nor (because of the scale of these operations) able to offer.

Software support is important to scholars. Note that even though the basic TEX and several versions of it are in the public domain, there are companies that successfully sell their support services or add-ons.

The examples of computer algebra and typesetting systems indicates that scholars are willing to pay for services that are truly helpful to them. However, it also shows that they are not willing to pay too much. For a single-user license for a PC, say, a scholar might pay around \$600 initially, say, and then another \$100–400 per year for upgrades and additional packages. Similar prices apply for other software apckages. This is far from the \$5,000–10,000 per scholar that the present publishing system costs.

9.6. Publishers

> "Well, in *our* country," said Alice, still panting a little, "you'll generally get to somewhere else - if you ran very fast for a long time as we've been doing."
>
> "A slow sort of country!" said the [Red] Queen. "Now, *here*, you see, it takes all the running *you* can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!"
>
> Lewis Carroll, *Through the Looking Glass.*

The publishers' future in scholarly journal publishing looks bleak. Even if they run much faster than they do now, I don't think they can avoid shrinking. Revenues and profits will shrink, and many skilled professionals who have dedicated their careers to producing high quality publications will lose their jobs. This will be a tragedy, but I do not see any way to stop it. The main role of electronic technology has been to eliminate the middlemen, such as the bank tellers who have been replaced by automatic teller machines that enable customers to deal directly with the banks' computers. The publishers and librarians have been the middlemen between the scholars as producers of information and the scholars as consumers, and are likely to be largely squeezed out of this business.

The costs of publishing will have to go down. In general, when costs go down, there are many possible outcomes. In the PC industry, prices have been falling at fantastic rates. However, the lower costs have stimulated enough new sales, and sales of more powerful computers, so that total revenues of the industry have continued to climb. On the other hand, when tire manufacturers doubled the life of the average car tire, their industry had to shrink by half, since the slight lowering of the costs of car ownership that resulted from longer-lasting tires did not translate into twice as many miles being driven. The question is, where in this spectrum of responses will publishing fit in?

My feeling is that publishers will have to shrink. Scholars are not going to increase substantially the number of papers they publish in response to lower subscription prices and more space available in journals. At least in the scientific and technical areas, it appears that just about everything that is written that is novel and correct already gets published, and most of the complaints one hears are about standards being too low, not too high. Therefore the number of the scholarly journal papers is likely to grow only with the growth in the number of scholars.

My projection could be wrong, and publishers might manage to hold onto their present revenues. To do that, though, they would need to defeat the subversive influence of preprints. Widespread electronic distribution of preprints seems bound to destroy all high-cost journals. In mathematics, computer science, electrical engineering, and those areas of physics that I know, it seems that preprints are so common that they cannot be uprooted. Their existence has typically been ignored in all the copyright agreements. However, they are there, and likely to keep spreading. The AMS has in effect conceded that preprint distribution cannot be controlled, and its latest copyright transfer form explicitly reserves to authors the right to distribute preprints on paper or electronically either before or after publication in AMS journals. Some other publishers still seem to think that the clock can be rolled back. This might be true, at least for a while, and especially in those fields where not much rolling back is required. I am told that in chemistry preprints are rare, and that publication usually takes place within three to four months of submission. Publishers could explicitly require that authors submitting their manuscripts to their journals not distribute preprints. By moving to electronic distribution, they could also shorten the already rapid publication cycle. They could offer many of the novel features that electronic publication makes possible. By thus rigidly controlling the distribution of information, they could maintain their position as indispensable intermediaries.

For the strategy of stopping preprint distribution to work, the time between submission of a manuscript and publication has to be short. Delays caused by backlogs at publishers can be eliminated with electronic publications. The main problem would be to shorten the time peer review takes. In chemistry that time is already short, but the quality of the review process leaves much to be desired [Quinn], since the referees usually only check the methodology and conclusions. Experiments are not expected to be duplicated. In mathematics the expectation is that the referees will check the details of the proof. The review typically takes half a year to a year, and there are often additional iterations of the process. Would a shorter review cycle be acceptable? Andrew Appel, the Editor-in-Chief of ACM Trans. Progr. Lang., has managed (personal communication) to bring the time for peer review at his journal down to three to four months by exclusive use of electronic communication, and asking referees for commitments to review a paper within a month. (When they refuse, he and his associate editors search out new ones. This means that for any single paper, they often have to approach half a dozen referees before finding the two willing to do the job in the required time.) Would this process be acceptable to other scholarly communities that are accustomed to leisurely refereeing? Would it be felt to produce acceptable quality of review? This seems doubtful.

I suspect that the strategy of attempting to prohibit free preprint distribution will fail. Scholars in those areas that already use preprints widely are not likely to tolerate a significant restriction on their accustomed freedom of distributing their work. There are just too many journals for them to publish in, and they would likely redirect their papers to those with least restrictive policies. Even in areas where preprints are not common, the dam will likely burst in the near future. The problem is that scholarly publishing differs greatly from the trade press. Section 2 already discussed one significant difference, namely the constant size of the audience for a typical scholarly publication, no matter how much the total number of scholars grows. Another difference is that scholars do not get direct financial benefits from their papers. The scholarly community is in the business of writing papers, giving them to publishers, and then buying them back as journals and books. Scholars and publishers have developed a relationship that used to be mutually beneficial. However, given the crisis in the present system, this relationship is under strain. The publishers are interested in maintaining their revenues. On the other hand, scholars (other than those writing textbooks or other popular books that sell in large quantities) are interested in obtaining recognition for their work from other scholars. If authors were getting as royalties \$2,000 of the \$4,000 that publishers charge for each article, scholars might possibly have a sufficiently strong incentive to cooperate with publishers in preserving the current high cost system. That is far from the present situation, though. The most cynical view of scholars says that they are only interested in getting their papers into the most prestigious journals they can, and do not care about delays in publication or limited availability of their work. Even if we do accept this view, this does not provide much comfort for publishers. The problem is that the prestige attained by a journal is the work primarily of scholars, and is not difficult to achieve in a short period of time. In mathematics, we have the examples of *Inventiones math.* (published by Springer) and *J. AMS*, both of which attained high status right after they were established. This was achieved not so through the prestige or economic investment of the publishers (although that helped), but through the prestige and policies of the scholars who were the editors. Thus it is not necessary to have a long history to achieve high status, and free electronic journals run by scholars will be able to do it as well.

Even if the total journal publishing revenue goes down, there might be room for publishers to provide new services. With advances in technology, various multimedia publications will be developed, even in the scholarly arena. This is an area where specialized skills are needed that cannot yet be bought in a \$500 software package, as they can in traditional publishing. Publishers could develop and market such products, since one of their strengths is in assembling teams of specialist professionals. There are other opportunities of this type as well, but I doubt they will be large enough to compensate for the shrinking of the traditional journal business.

Although scholarly publishing seems bound to shrink, so that many highly trained experts will lose their jobs, and many publishers are likely to leave the business, some might be able to prosper. The example of encyclopedias is instructive in several respects. In 1993 more CD-ROM encyclopedias were sold in the US than paper ones. The advantages of CD-ROM technology involve less space, better searching facilities, multimedia features, and lower price, as well as the sheer novelty of the concept, and they apparently outweigh the disadvantages of poorer display quality and breaking with tradition. The most scholarly of the English-language publications, *Encyclopaedia Britannica*, was slow to embrace the new medium, its sales have nosedived in the last two years, and it is losing money [Samuels]. On the other hand, one of its rivals, *World Book*, is selling more copies than before. Individual prices are lower so that total revenues are smaller than they used to be, but profits are higher. There is of course the question whether *World Book* profits can stay at their level, since in an equilibrium competition is likely to press them down. However, equilibrium will not hold in the next decade or two in any areas of the economy that are affected by the revolution in electronics, and only nimble action of the sort taken by the *World Book* offers any hope for survival.

9.7. Review journals

> "Doing good is noble. Telling others to do good is nobler and less trouble."
>
> attributed to Mark Twain

It is easy to say how scholarly publishing ought to be done. However, who will do all the necessary work to keep everything running, to filter information and make it available in a systematic way? Currently scholars do most of the work themselves, but do have the help of publishers, review journals, and librarians. Will scholars be able to do everything by themselves in the future? At the moment they clearly cannot. Even in areas where scholars have a strong self interest, they do not always act as they should. Most journal editors have stories of authors whose papers never got published because those authors never revised them to take into account referees' comments. There are also all the annoying problems that crop up of obscure papers, foreign language publications, and so on. While many of these problems are likely to disappear, probably there will always be some left, especially since electronic publishing and natural growth in the scholarly community will make much more information available on the Net. I feel there will always be a need for a central organization or organizations to provide a coherent framework to the technical information on the Net.

Once journals move to an electronic format, will we need review publications such as *Math. Rev.* and *Zbl.*? It is possible that they will disappear. These journals play an invaluable role in informing scholars of the latest developments in mathematics. However, just how valuable is this contribution going to be in the new era in which everything is available on-line? If I can do text searches over the whole mathematical literature, how important will the abstracts be? Computerized searches can take over many of the functions of review journals. As an example of what is already possible, Marc VanHeyningen, a graduate student at Indiana University, has in his spare time set up an automated search service. It contains about 10,000 abstracts of computer science technical reports that are available via anonymous ftp, together with locations of the full reports, and makes possible various keyword searches. (The URL is http://cs.indiana.edu/cstr/search.) VanHeyningen has to tell the program what sites to search and how to look for reports at any site, but additions of reports to the sites in the database are handled automatically. In the future we can expect programs of this type to be much better, and also to have much more data available to them in standardized ftp directories (such as that described in Section 3).

While it is possible that *Math. Rev.* and *Zbl.* will disappear, I suspect they won't, and they might even flourish, and become (in Frank Quinn's words) "gateways to published mathematics." There are two reasons for this. One has to do with the costs of these journals. Each one is priced at about $5,000 per year. With the shift towards electronic journals and electronic communications, their costs for mailings, physical storage, and the like are likely to drop, so their prices might not have to rise even with a substantial increase in the volume of publications they handle. Now $5,000 per year is high when compared with the costs of individual journals, but is sufficiently low that it can easily be absorbed in a department's budget. Thus the review journals need only to provide some slightly unique services to justify paying for them, especially if library budgets do drop significantly. The other reason that review journals are likely to survive is that they provide valuable services that might not be easily derivable from the information supplied by authors in their papers. The bulk of the costs of the review journals is in the maintenance of their data base. In addition to all the people who handle keyboarding, acquisitions, and the like, these journals employ teams of professional mathematicians whose job it is to classify articles, find appropriate reviewers for them, check references, translate titles and references in languages other than English, and so on. Little of this work can be automated right now. Eventually, as software improves, and the input data becomes more standardized, we might see more of this work done by computers, but the chances are there will always be corners of the electronic universe that will require human attention. There will be strange databases that need special methods to access, changes in locations of scholars, different versions of the same papers on different machines, and so on. With increasing specialization and growth in literature, there will be increased need for timely surveys. These will have to be written by scholars, but review journals might be the ones taking the lead in finding and persuading experts to undertake the task so as to provide even coverage. Even though free automated search systems might provide 80% of what scholars need, review journals might well justify their prices by providing the extra 20%.

To survive in the coming world of electronic publishing, review journals will have to change. They will need to be accessible electronically, and will most likely be paid for by a site license fee, giving unlimited access to the database to all scholars affiliated with the customer institution. They will provide much more current information than is true today, since there will be no publications delays. (*Math. Rev.* already provides a paper publication, *Current Math. Pub.*, which lists the latest mathematical papers to appear. In the future this should be replaced by an electronic version that will eliminate all publication delays other than those needed to verify the subject matter is of mathematical interest, and should include electronically available preprints.) The formats of reviews might vary from those used today. The main distinction from today is likely to be the presence of hypertext links from reviews to the papers and the commentaries associated to those papers. Combined with easy electronic access to the primary materials, review journals will then provide all the functions of a specialized library.

9.8. Libraries

Advances in technology will also drastically affect libraries. Libraries are expensive. While mathematical journals cost the US mathematical community about \$70 M per year just in subscription costs, the other costs of maintaining library collections (both people and space, including overheads, but excluding subscription costs) are usually at least twice that. Right now librarians are the ones facing the crisis in publishing, since they have to deal with increasing prices and increasing scholarly output. They are squeezed between publishers on one hand, and scholars on the other. Technology will solve the librarians' problem, but will also eliminate most of their jobs. Just like publishers, libraries will have to shrink and change their role.

The transition to the new system is likely to be much less painful for libraries than for publishers. There is much more inertia in the library system, with old collections of printed material that will need to be preserved and converted to digital formats. Further, most of the jobs that will be lost initially will be those of low-skilled workers, typically students working part-time, who check out and reshelve books, process the paperwork for interlibrary loans, and so on. (According to the ARL, about half of the personnel budget in a typical research library is for professionals, and half for support staff.) However, there are bound to be changes, and they will be painful. As one sign of what is likely to happen, the New York Times on Oct. 7, 1993, reported that Time, Inc., the parent company of Time, Sports Illustrated, and several other magazines, was cutting its library staff from 120 to 80. The justification for the cut was that many of the functions performed by the staff, such as cutting out and filing stories from various sources were being taken over by electronic searches. This is just the beginning of the trend that will see the demise of the traditional library as well as the traditional journal.

University libraries have already lost some of their importance. Spending on libraries has been increasing rapidly, much faster than inflation. Still, Albert Henderson has pointed out that over the last 25 years, the fraction of budgets of research universities in the US that are devoted to libraries has declined from 6One could therefore argue that everything would be fine with scholarly publishing if only libraries regained their "rightful share" of university budgets. My opinion is that this is unrealistic, and that the decline in the relative share of resources devoted to libraries resulted from their decreasing importance. The increasing availability of phone, fax, email, interlibrary loan, and other methods of obtaining information, and the inability of any single library to satisfy scholars' needs, may mean that scholars do not need the library as much, and as a result do not fight for it. In the best of all possible worlds, there would be resources to acquire everything, but in practice, choices have to be made, and at some level in the university power structure, libraries compete for money with faculty salaries, student scholarships, and so on. That libraries have been losing this competition probably means that they have already lost some of their constituency, and will have to change.

The traditional library, a building full of paper, will lose its importance in scholarly work. Its place will be taken by the "virtual library," consisting of digital information distributed over the globe and connected by efficient networks. It is not certain just how many professional librarians will be needed in this environment. They will certainly be needed for interdisciplinary endeavors, which are likely to continue increasing in importance, as well as in commercial uses of the virtual library, which will certainly explode. In specific scholarly disciplines, though, there is likely to be much less demand for them. If the review journals evolve the way I project in the Section 9.7, they will provide directly to scholars all the services that libraries used to. With immediate electronic access to all the information in a field, with navigating tools, reviews, and other aids, a few dozen librarians and scholars at review journals might be able to substitute for a thousand reference librarians.

10. Conclusions

We are on the threshold of dramatic changes in scholarly publications. Current scholarly journals are a refined product of long evolution. However, the environment they operate in is changing radically, and they will need to change as well. The question is whether they will die out like the dinosaurs, or whether they will adopt to the new era.

Acknowledgements: This article provides my own personal view of the future of mathematical journals. Few of the observations and predictions are original, and I have freely drawn on the ideas in the papers listed below. I have benefited greatly from extensive email correspondence with Paul Ginsparg and Frank Quinn, and especially with Stevan Harnad. Helpful comments and useful information were provided

by Andrew Appel, Robert Badger, Henry Baird, Margaret Baker, Michael Barr, Judith Bausch, Rick Becker, Edwin Beschler, Ken Birman, Carol-Ann Blackwood, Piotr Blass, Enrico Bombieri, Dan Boneh, Fred Bookstein, Joe Buhler, Catriona Byrne, David Cantor, Donald Case, Steve Cavrak, Lorinda Cherry, Peter Cziffra, Keith Dennis, Howard Dess, Whit Diffie, Peter Doyle, Tom Duff, David Elliott, Lloyd Etheredge, Emanuel Farber, Joan Feigenbaum, Janet Fisher, Dominique Foata, Walter Gautschi, Sergei Gelfand, Mark Ginsburg, Andrew Granville, Dan Grayson, Albert Greenberg, Martin Groetschel, Eric Grosse, Jean-Claude Guedon, Leo Guibas, Richard Guy, Phil Hanlon, Tony Hearn, Joachim Heinze, Dennis Hejhal, Albert Henderson, Ross Hoffman, Andrew Hume, Carol Hutchins, Patrick Ion, Lee Jaffe, David Johnson, Jerry Kazdan, Arnoud de Kemp, Martin Kesselman, Jane Kister, Steve Krantz, Greg Kuperberg, Bob Kurshan, Jeff Lagarias, Brian LaMacchia, Leslie Lamport, Susan Landau, Jim Landwehr, Joshua Lederberg, Silvio Levy, Bernhard Lewerich, Jack van Lint, Dick Lipton, Daniel Loeb, Bruce Long, Barbara Longhurst, Joachim Luegger, Jeffrey MacKie-Mason, Dan Madden, Andy Magid, Clyde Martin, John Martindale, Yossi Matias, Brendan McKay, Norman Megill, Jill Mesirov, Jerry Morine, Kent Morrison, Jim Muckerheide, Gary Mullen, Bernard Naylor, Jakob Nielsen, Paul Odlyzko, Ann Okerson, Dick Palais, Tony Parsons, Paul Pedersen, Arno Penzias, Chris Phillips, Rob Pike, Barry Rackner, Stanislaw Radziszowski, Arden Ruttan, Tony Ryan, Chih-Han Sah, Ann Schaffner, Wilfried Schmid, Claus Schnorr, Bob Sedgewick, Mark Sheingorn, Alan Sherman, Larry Siebenmann, Daniel Sleator, Neil Sloane, Warren Smith, Wolfram Sperber, David Stodolsky, Daniel Szyld, Jean Taylor, Emre Telatar, Bill Thurston, Flemming Topsoe, Martha Tucker, Elspeth Tyler, Marc VanHeyningen, Hal Varian, Edward Vielmetti, Leo Waaijers, Bernd Wegner, Gene Wells, Joerg Winkler, Scott Wolpert, Dan Woodard, Ian Worthington, Dave Wuolu, and Jeff Young.

References

[AMSS] Survey of American research journals, *Notices Amer. Math. Soc.* 40 (1993), 1339-1344.

[APS] Report of the APS task force on electronic information systems, *Bull. Amer. Phys. Soc. 36, no. 4* (April 1991), 1119-1151.

[APSPRB] Report of the APS task force to review Physical Review B, available at URL http://aps.org/PRBREPORT/prb.html.

[Babai] L. Babai, E-mail and the unexpected power of interaction, pp. 30-44 in *Proc. 5th IEEE Structures in Complexity Theory Conf.*, Barcelona 1990.

[DavoustS] E. Davoust and L. D. Schmadel, A study of the publishing activity of astronomers since 1969, *Scientometrics 22* (1991), 9-39.

[Franks] J. Franks, The impact of electronic publication on scholarly journals, *Notices Amer. Math. Soc. 40* (1993), 1200–1202.

[Franks2] J. Franks, What is an electronic journal?, unpublished report, available at gopher://gopher.cic.net/11/e-serials/related.

[Ginsparg] P. Ginsparg, First steps towards electronic research communication, *Computers in Physics, 8, no. 4*, July/August, (1994), 390-396. Also available at URL http://xxx.lanl.gov/blurb/.

[GroetschelLS] M. Groetschel, J. Luegger, and W. Sperber, Scientific publishing and electronic information at a turning point: A report on the situation from the point of view of mathematics, German original May 1993, English translation by Regine Fadiman of Math. Rev., Jan. 1994.

[Grycz] C. J. Grycz, ed., Economic Models for Networked Information, special issue of *Serials Review, 18*, no. 1-2 (1992).

[Hall] D. R. Hall, Rate of growth of literature in geoscience from computerized databases, *Scientometrics 17* (1989), 15–38.

[Harnad1] S. Harnad, Scholarly skywriting and the prepublication continuum of scientific inquiry, *Psychological Science 1* (1990), 342-343. Reprinted in *Current Contents 45* (November 11, 1991), 9-13.

[Harnad2] S. Harnad, Implementing peer review on the Net: Scientific quality control in scholarly electronic journals, *Proc. Intern. Conf. on Refereed Electronic Journals: Towards a Consortium for Networked Publications*, to appear. (Available via anonymous ftp, along with [Harnad1] and other related papers, from princeton.edu, in directory pub/harnad/Harnad.)

[Harnad3] S. Harnad, Publicly Retrievable FTP Archives for Esoteric Science and Scholarship: A Subversive Proposal. To be presented at: Network Services Conference, London, England, 28-30 November 1994. (Available via anonymous ftp from princeton.edu, in directory pub/harnad/Psycoloquy/Subversive.Proposal.)

[JaffeQ] A. Jaffe and F. Quinn, "Theoretical mathematics": toward a cultural synthesis of mathematics and theoretical physics, *Bull. Amer. Math. Soc. (NS) 29* (1993), 1-13.

[JaffeQ2] Responses to "Theoretical mathematics: toward a cultural synthesis of mathematics and theoretical physics," by M. Atiyah et al., and "Response to comments" by A. Jaffe and F. Quinn, *Bull. Amer. Math. Soc. (NS) 30* (1994), 178-211.

[KingMR] D. W. King, D. D. McDonald, and N. K. Roderer, *Scientific Journals in the United State. Their production, use and economics*, Hutchinson Ross, 1981.

[Lamport] L. Lamport, How to write a proof, *Amer. Math. Monthly*, to appear.

[Lancaster] F. W. Lancaster, *Toward Paperless Information Systems*, Academic Press, 1978.

[Lederberg] J. Lederberg, Digital communications and the conduct of science: The new literacy, *Proc. IEEE 66* (1978), 1314-1319.

[Licklider] J. C. R. Licklider, *Libraries of the Future*, MIT Press, 1965.

[Loeb] D. Loeb, An Electronic Journal of Mathematics: Feasibility Report 5, electronically circulated manuscript, 1991.

[Machlup] F. Machlup, K. Leeson, and Associates, *Information Through the Printed Word: The Dissemination of Scholarly, Scientific, and Intellectual Knowledge*, vol. 2: Journals, Praeger, 1978.

[MacKieV] J. K. MacKie-Mason and H. R. Varian, Some economics of the Internet, in *Networks, Infrastructure and the New Task for Regulation*, W. Sichel, ed., to appear. (Available via gopher or ftp together with other related papers from gopher.econ.lsa.umich.edu in /pub/Papers.)

[MR50] *Mathematical Reviews, 50-th Anniversary Celebration*, special issue, Jan. 1990.

[NyceK] *From Memex to Hypertext: Vannevar Bush and the Mind's Machine*, J. M. Nyce and P. Kahn, eds., Academic Press, 1991.

[Okerson] A. Okerson, The electronic journal: What, whence, and when?, *The Public-Access Computer Systems Review 2*, no. 1 (1991), 5–24.

[Price] D. J. Price, The exponential curve of science, *Discovery 17* (1956), 240-243.

[Quinn] F. Quinn, Roadkill on the electronic highway? The threat to the mathematical literature, *Notices Amer. Math. Soc.*, Jan. 1995, to appear.

[Rawlins] G. J. E. Rawlins, Publishing over the next decade, *J. Amer. Soc. Information, 44* (1993), 474-.

[Samuels] Gary Samuels, CD-ROM's first big victim, *Forbes*, Feb. 28, 1994, pp. 42-44.

[Schaf1] A. C. Schaffner, Electronic journals in the sciences, Brandeis internal report, available from the author, schaffner@logos.cc.brandeis.edu.

[Schaf2] A. C. Schaffner, The future of scientific journals: Lessons from the past, *Information Technology and Libraries*, vol. 13, Dec. 1994, to be published. Preprint available from the author, schaffner@logos.cc.brandeis.edu.

[Story] G. A. Story, L. O'Gorman, D. Fox, L. L. Schaper, and H. V. Jagadish, The RightPages image-based electronic library for alerting and browsing, *IEEE Computer 25* (no. 9, Sept. 1992), 17-26.

[Sci1] G. Taubes, Publication by electronic mail takes physics by storm, *Science 259* (Feb. 26, 1993), 1246-1248.

[Sci2] G. Taubes, E-mail withdrawal prompts spasm, *Science 262* (Oct. 8, 1993), 173-174.

[Thurston] W. P. Thurston, On proof and progress in mathematics, *Bull. Amer. Math. Soc. (NS) 30* (1994), 161-177.