International Spring School on the Digital Library and E-publishing for Science and Technology
CERN, Geneva, Switzerland, 3 – 8 March 2002

Electronic Journals and The American Physical Society

Mark D. Doyle
Manager, Product Development, The American Physical Society, United States

INTRODUCTION

Founded in 1899, The American Physical Society is a scholarly society of physicists with more than 40,000 members. We are the publisher of the physics journals *Physical Review*, *Physical Review Letters*, and *Reviews of Modern Physics*. *Physical Review* was started in 1893 at Cornell University and APS took over its publication in 1913. *Physical Review Letters*, perhaps the premier journal in physics, was first published in 1958 and *Reviews of Modern Physics* began in 1929. Currently the APS editorial office receives more than 26,000 manuscripts annually of which about 14,000 are accepted for publication after peer review.

APS journal are international in scope. Submissions from the U.S. only account for about 30% of the submissions we receive and Western Europe submits more papers than the U.S. The publication percentages track very closely the submission percentages for most countries and regions. Furthermore, two-thirds of our subscribers are also outside of the U.S. We value the international character of the journals and this aspect influences many of the decisions we make.

All of our journals have been available online since 1997, with some being available electronically as early as 1995. During this time, we have changed how and where our journals are available several times and we continually evaluate how best to serve our end users. But even before 1995, APS has had a long history of using electronic publishing to produce the journals. In 1985, the typesetting of the journals was moved to using TROFF and then to TeX. And then in 1993 an effort was begun to integrate SGML into the production process. This latter effort is ongoing today – we are planning this year to initiate a new XML-based production process that will greatly enhance the quality of our electronic archive.

Another major electronic effort has been to make all APS journal content available online, in an effort called the *Physical Review Online Archive* (PROLA). Originally this project was a modest effort started in 1993 to make a searchable index out of the legacy electronic typesetting files mentioned above. Today, PROLA contains over 1.6 million scanned images and tens of thousands of PDF files covering 1893 to 1998. This archive has all the features of the platform used to make our current contents (1999-present) and it forms a core component of our long-term strategy for electronically archiving all of our content.

The advent of electronic publishing doesn't just mean a changing how publishers produce journals. Rather, it is having a much broader impact by enabling scholars to communicate more efficiently and directly with each other outside of the traditional journal system. Of course, the most well known example of this is Paul Ginsparg's development of the e-print

archives hosted at arXiv.org. First started in 1991, this new mode of communication continues to grow and to influence new efforts in fields outside of arXiv's focus on physics and math.

The rest of this paper will detail how the transformation from a print publisher to a curator of an electronic archive has changed how the APS carries out its primary mission ("the advancement and diffusion of the knowledge of physics"), the foundation we are building to maximize the benefits of electronic publishing, and the lessons we have learned in putting it all together.

APS EDITORIAL OFFICE AND PEER REVIEW

The APS mission is to "advance and diffuse the knowledge of physics." This led naturally to taking on a role in peer-reviewing articles and printing journals so that scholarly communication in a print environment would be of high quality and efficient.

Peer review for APS journals is carried out in the traditional manner. Manuscripts are screened by an editor who then selects between one and three referees (depending on the editor's judgment and the individual journal's policies). Referees return reports on the manuscripts and, after consideration by the editor, either the paper is accepted as is, is rejected outright, or sent back to the author for modifications. Accepted papers are then sent out to production vendors to copyedit and typeset the articles for print and online distribution. Authors of rejected papers have access to an appeals process in which external editorial boards can make further recommendations. The APS Editor-in-Chief is the final arbiter of all appeals.

Until the 1970's, APS journal operations were relatively small with only a handful of editors coordinating all of the peer-review work. The journal headquarters would follow the editor. However, by the late 1960's *Physical Review* had grown so large that the journal was split into four (and now five) subject-focused journals. The accompanying growth in editorial resources also led to the establishment of a central office with full-time editorial and support staffs. (Only Reviews of Modern Physics is still run by the editor from his location – the journal is a quarterly journal with less than 50 articles per year.)

Today, the APS editorial office employees about 145 people. Having such a large central office is somewhat rare. However, with 26,000 submissions annually, we find that having a central office increases efficiency and allows the journals' policies to be applied in a cohesive and uniform manner. Our journals also make use of "remote" editors who work part-time to help coordinate peer-review. For example, *Physical Review Letters* has about eleven full-time editors employed at the central office and six part-time remote editors. If PRL were to convert completely to a staff of remote editors, which typically devote at most 20% of their time to editorial work, PRL would require over 50 editors to handle the workload. Clearly coordinating 50 editors and maintaining cohesiveness among them would be a challenge.

There are some disadvantages of the APS approach. For one, it makes it fairly difficult to streamline and change the workflow through the office. In particular, APS has been undergoing a re-engineering of its editorial workflow (optimised for paper-based processing of manuscripts) for about four years now and we still have at least another year's worth of

work. However, once implemented, we will have achieved new efficiencies and, just as importantly, we will have made our work environment more flexible and easier to change.

The success of the APS centralization strategy is borne out by the extremely high quality of the journals despite their large size. There is an economic cost of course for maintaining such an office and the support staff. This cost is non-trivial and we will examine this point in more detail later. For now, suffice it to say that scale matters: One cannot simply scale up a small journal operation to handle the volume that APS does without incurring some diseconomies of scale. It should also be equally clear that taking a large journal and replacing it with many smaller journals also has some disadvantages as well.

ECONOMICS OF APS JOURNALS

Let us now focus our attention a bit more on the economics of peer review and how electronic publishing is affecting this. APS gets about $1750 per published article ($1000 per submitted article) in revenue. Since we are a non-profit society and our budgets are very close to break even, this is also our cost per article. Breaking down the costs further, peer review costs about $500 per article. This includes the entire overhead for the editorial office and rejected manuscripts incur this cost as well. Publishing an article ads another $1200 per article about half of which is for composition (copyediting and typesetting). The remaining half is split almost equally between online distribution and printing and print distribution. This is in our current production process, which is still optimised for print production. We are working with our vendors to develop new processes which should substantially reduce the composition costs and we are looking at ways to reduce our online costs as well. It is hoped that the extra costs associated with printing will eventually be shifted entirely to those who want print.

Not all of the composition costs are attributable to print. One of the outputs of our production process is a heavily marked up SGML file for each article. These files will eventually form the core of our long-term digital archive, but we are still developing this into a true archival deliverable – the new processes alluded to above are targeted at making the marked up file the primary output of our production (rather than a PostScript or PDF file) and ensuring that the files really are archival in nature. What we are finding is that we can make this major change and benefit from cost savings while doing it. This is a concrete example of how true electronic publishing, once optimised, can decrease costs. The projected 20-25% savings in composition costs is significant.

There are several observations that can be made at this point. First, we see that the cost associated with peer review is the same order of magnitude as the cost for print distribution, and that is dominated by labour costs. Thus, it makes very good sense in a print world to couple the two together. This leads naturally to the ubiquitous subscription model in which peer review costs are paid by the reader. Second, it does cost a good deal of money to build and maintain the infrastructure needed for online delivery. This includes all services associated with online delivery such as searching, linking, and mirroring.

Is the APS typical? Paul Ginsparg has done some analysis summarized in the chart shown in Figure 1. See: http://arXiv.org/blurb/pg01unesco.html for his explanation of the graph. One can debate the exact numbers in the graph, but I have found the numbers presented to be quite defensible. The aim of this chart is to highlight the differences in orders of magnitude for

various items. First, it is clear that the cost per article for the underlying research is at least an order of magnitude larger than the publishing costs. Second, it is clear that the more expensive commercial publishers seem to generate a lot more revenue per article than APS journals. One may ask whether their costs really are substantially greater than ours, and if so, why? Or, if their costs are similar to ours, what is the justification for such large margins?

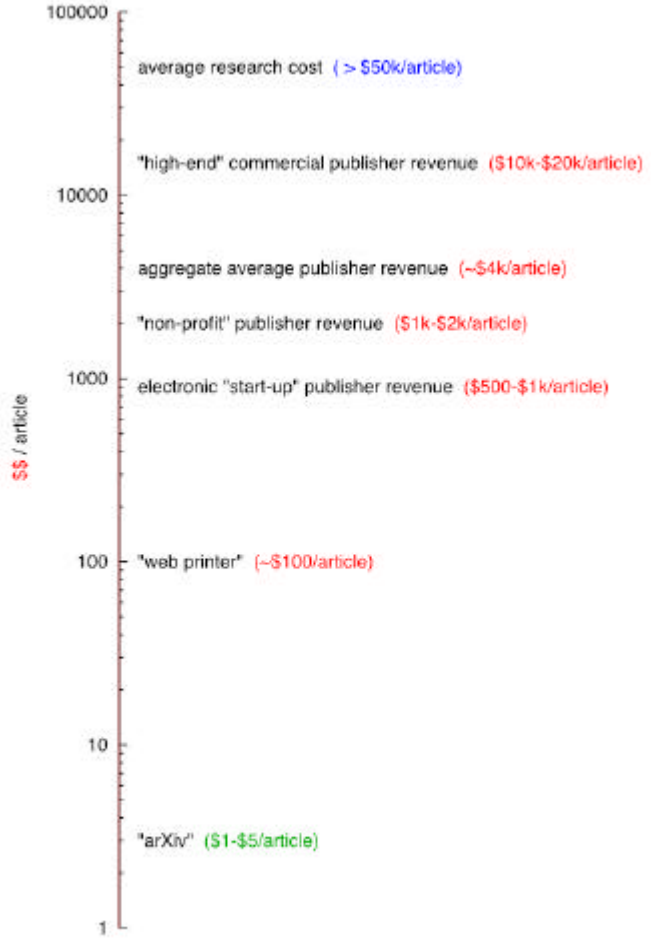

- Figure 1: A logarithmic plot of revenue per article (taken from http://arXiv.org/blurb/pg01unesco.html by P. Ginsparg with permission).

Third, small electronic journals do seem to be cheaper than a large operation like ours. However, it should be noted that the model journal for this analysis was the Journal of High Energy Physics, which has been freely available. Growing pains are forcing them to move to subscription model and a partnership with an established physics publisher (IOPP). I take this as an indication that they are closer to the upper end of the range in their costs. As I said above, APS is not as efficient as we could be because we still have some legacy issues related to print distribution and paper-based manuscript processing. I suspect that in the end APS's and JHEP's costs will converge to something close $1000/published article.

Finally, the graph shows that a pure electronic distribution via a service like arXiv.org is at least two orders of magnitude cheaper than APS costs. And this final point is the crux of the matter. With dissemination costs falling so low compared to composition and peer review costs, one really must re-examine whether there is enough added value to justify the expenses. APS believes that there is enough added value in both peer review and in preparing archival, richly tagged marked up files for what we peer review. On the other hand, it is clear that if the basic articles can be obtained so cheaply (over 90% of *Phys. Rev. D* our high energy physics journals is available from arXiv.org), then the subscription model is quite unstable as a way to recover these costs.

The subscription model is further unstable because of the well-known serials crisis. APS subscriptions have been decreasing at about 3% per year since the late 1960's. There have been many different reasons for this over the years, but the trend has been quite stable over 30 years. Most recently, this has been attributed to large institutions consolidating multiple subscriptions into a single subscription because online access makes the extra subscriptions redundant. I'll have more to say about this later when I talk about new pricing models. In any case, a decrease in the number of subscriptions means that subscriptions prices must go up in order to balance our journals budget. This coupled with a few percent per year growth in the number of articles we publish and the usual inflation in labour and other costs means that we had to increase prices on average about 8% per year. Clearly this is unsustainable in the long run.

I have already mentioned that we are beginning to see significant cost savings in some of our journal production. Are there other places in which there could be substantial cost savings? The peer review component is quite labour intensive. Increasing submission numbers means more labour. Also, the computer resources and other infrastructure are also rather substantial. Both the labour and support costs are rather inelastic. We may be able to reduce these costs by a few hundred dollars per article, but it would take a major restructuring (which we are working on) and quite some time before we would size savings of this size, and the cost would still be two orders of magnitude above the pure electronic dissemination of arXiv.org.

Getting rid of print distribution would save us about 15%. There isn't much more one could do here. In fact, we now offer online-only subscriptions at a 15% discount to pass the savings along directly to the subscribers. However, only about 5% of our subscribers have decided to take online-only suggesting that it will be quite some time before we can do away with print. Some our electronic efforts, especially those relating to long-term archiving, are aimed on making this transition more comfortable for libraries and researchers.

The only other place to see large savings would be to further reduce our production costs. A major change could be accomplished if authoring tools were developed to the point where authors could deliver to APS a highly marked up SGML (or XML) file that could be used directly in production. We have done some development work in this direction, but we are still years away from this ideal system. Effectively, we would be transferring our production costs back to the author (although these costs would be greatly reduced because the tool would make adding the markup transparent to the author – the earlier mark-up is added in a document's life cycle the cheaper it is).

To sum up, then, it is clear that there are real costs associated producing high-quality peer review on the scale of our journals. Furthermore, there are nontrivial costs associated with creating an archival document. These costs are the primary reason APS costs are two or three orders of magnitude about those of arXiv.org. It should be noted that if authoring tools were developed as mentioned above, a server such as arXiv.org would be able to make substantial use of the mark-up to offer many of the services that differentiate high-end online journals from what arXiv.org offers.

The final point I would like to make in this section is that the APS's mission is not to produce journals. We still believe that our journals are of sufficient use and importance in "advancing and diffusing knowledge of physics" that we will continue to produce them. But if the physics community were to decide that they no longer serve this purpose, that traditional peer review is no longer warranted, or if the journals were to become economically unviable, there is little doubt that we would cease publishing them. Of course, our fervent hope is that we will be able to transform ourselves to use better processes and better economic models resulting in journals that are sustainable into the indefinite future.

ALTERNATIVE ECONOMIC MODELS

Now that we have some idea what it costs to do peer review and create archival versions of manuscripts, as well as the shortcomings of the papyrocentric subscription model, we can look at some alternatives made possible by the fact that online versions are becoming the dominant way of accessing journal content.

TWEAKING THE SUBSCRIPTION MODEL

To begin with, the most straightforward thing to try is to adjust the subscription model to better match how the material is being used. Historically, larger research institutions had a larger role in supporting the journals through multiple subscriptions. Now, however, these institutions have cancelled their multiple subscriptions. This has the effect of transferring the burden to the many smaller institutions that subscribe to our journals. These smaller institutions may have limited research programs, but they still need to give students and faculty access to journal material. To help reverse this shift, APS introduced several years ago a tiered pricing model in which larger institutions pay more than smaller institutions. Furthermore, larger prices increases are applied to the higher tiers to lessen the pressure on smaller institutions. Thus, while a large institution may see double digit percentage increases, smaller institutions might see only a 2% increase. Overall, we still have to maintain the 7 or 8% annual growth in income pointed out above. Currently tier assignments for U.S. institutions are based on the Carnegie Classification (http://librarians.aps.org/institutional.html, http://www.carnegiefoundation.org/). The tire for non-U.S. institutions is based upon a comparison of an institution's downloads and the median for each U.S. tier.

A large part of our Treasurer/Publisher's time is spent looking at how we can make this tiered pricing even more equitable. For instance, we can track submissions from institutions as well as downloads. These data could then be factored into deciding which tier a particular institution should be assigned to. We are very careful to make such criteria as objective as

possible to avoid arbitrariness. Furthermore, we want a stable model in which an institution's assigned tier is stable over time. The goal is always the same. We would like an institution's price to reflect how it uses the journals and the size of its scientific output. Tiered pricing is certainly an improvement over the one-size-fits-all model, but it has its limitations. It is still rather coarse (we are moving to increasing from three to five tiers) and it is difficult to really match what an institution pays to what it costs the APS for publish its research and to deliver the journals to the researchers.

Other variations to the subscription model are focused on moving beyond single institutions as subscribers. Instead, we sometimes negotiate with library consortia or even provide countrywide site licensing. Library consortia have their drawbacks from our point of view though. Each deal requires extensive negotiation and most consortia are more geared towards getting more access to material for a modest increase. However, APS does not have a large inventory of journals to sweeten the pot as some large publishers do. What we do see are large chunks of library acquisition budgets being soaked up by these large consortia deals which then increases the pressure on our subscriptions. OhioLink is one notable exception. We have had very good success in working with them in order to broaden access to our journals.

Broad site licenses such as OhioLink or countrywide agreements such as we have with Russia and Taiwan is actually quite appealing in my view. The key problem inherent in a pure subscription model are that the price an institution pays doesn't scale properly with its research output or with its use of the journals. The tiered pricing outlined above aims to ameliorate some of these problems. But the advantage of a large site license is that one can endeavour to make the negotiated fee match closely what it actually costs to serve that particular region. The main disadvantages to this approach are finding sustainable government or other alternative funding sources with which fees can be negotiated to balance the costs appropriately. I think with time we will be able to solve some these problems. It is good that such options exist. For when we look at replacing the subscription model with something else, one quickly realizes that it will be very difficult to come up with a single model that can be applied in all countries. Thus, we can have a variety of models, each of which endeavours to balance the costs for a particular region. With this in mind, we can now focus on alternatives to the subscription model.

REPLACING THE SUBSCRIPTION MODEL

There are three main strategies the APS is working on to replace or reduce the need for subscription revenue:

- sponsorship in which all costs associated to a journal are funded by grants;
- submission and/or publication charges for each article;
- development of new derived deliverables which can generate income.

An example of the first approach is our new journal *Physical Review Special Topics: Accelerators and Beams* started in 1998. This journal is an online-only journal that is sponsored by eight large accelerator laboratories. This works well for the journal because of its small size – it has only published about 225 articles. But it is serving an important role in a field that didn't have any low cost journals dedicated to the subject area.

The second approach has the most promise for grounding the journals in a financially stable model. A rather obvious fact is that if we were able to get authors to pay all the costs associated with publishing their work, then we wouldn't need to have subscriptions at all. There are many obstacles to overcome though. Historically, there were page charges authors had to pay in order to get their paper published. This defrayed some of the cost for publishing the work. At the present time, only *Phys. Rev. Letters* has mandatory page charges. The other journals have only voluntary page charges and the honouring rate for these has been steadily decreasing.

Why are page charges voluntary? This is one item that the high-energy physics community got wrong (as opposed to being the main driving force behind the growth of arXiv.org and the movement it represents). Commercial publishers didn't have page charges and physicists faced with a choice of paying or not paying elected not to pay. This problem was particularly acute for theorists who publish frequently and have small budgets (as opposed to experimentalists). Money for page charges had to come out of the research groups' budgets. The result was a rise in prominence of some commercial journals at great expense to the libraries and page charges became voluntary. When APS governing bodies tried to reinstate page charges in the mid-1990's for the high-energy physics journal *Phys. Rev. D*, the editorial board resigned in protest.

So this is what doesn't work, at least in physics. (It should be noted that some fields such as astronomy and astrophysics, substantial page charges are routinely paid by authors.) There are some clear lessons. First, if the authors' grants have to directly cover the charges, then authors will seek other venues for publishing their work. Second, the authors didn't get anything in return for their page charges. Both of these can be remedied by trying a new approach. Ideally, it would be the authors' institutions or a group of institutions that would pay the charges. Furthermore, in exchange for these payments, a publisher could give something back to the institutions such as rebates on subscriptions, free access to the articles in the online journal, or even the archival marked up file and associated PDF for local mounting at the institution. APS is pursuing just such a model with some U.S. institutions, but it is too early to report on this effort.

A closely related model is the BioMedCentral model in which there are institutional members who pay annual fee. Authors from a participating institution then don't have to pay the article-processing fees. Institutional members also get certain rights for reusing the content. The main issue with this model (quite simple) is that it doesn't accurately reflect what an institution actually submits to the BioMedCentral journals. More information can be found at http://www.biomedcentral.com/info/instmembership.asp.

The third alternative is to generate more revenue from non-traditional sources. For instance, one derived product might be richly marked up bibliographic records that could be marketed to abstracting and indexing services. Different levels of mark-up would have different price points: Table of contents level data might be free, abstracts might be added for a modest fee, marked up references for linking might be more expensive, and the full article content for use in local platforms (searching for instance) might be sold at some multiple of a standard subscription price. APS has started to generate some income in this way, but we are long way from covering all of our costs. However, whatever we can generate this way will reduce the pressure to increase our subscription prices. It isn't clear to me how this market will develop

or whether it will always be there, so I would rather focus on the second alternative. Still, for this transition period, any additional income is quite important.

INTELLECTUAL PROPERTY ISSUES

Leaving behind economic issues, let us turn our attention to other issues of great importance in the digital age. One of the most important issues is that of intellectual property. It is often argued that authors should retain their copyright and not transfer it to publishers. However, the issue isn't really about copyright, but what a publishers and authors can do with the work. APS still requires transfer of copyright to us for articles that we publish (except for articles covered by statutory exceptions for research conducted by government workers). Our copyright transfer agreement immediately transfers back to the author many important rights. For instance, authors may keep their papers on the e-print archives and even update them to reflect the final peer reviewed content. They may not post the APS formatted version of the article to the e-print archive; however, they can make it available from their own personal web site or use it in a personal compilation of their work. Other rights including authors being able to grant third parties the right to reuse figures or tables.

Why do we require the copyright? There are several reasons. To begin with, we feel we are in a better position to defend the copyright than the author. Also, a license to publish may not give us the rights to do what we like with the articles. For instance, without copyright, it would have been nearly impossible for us to create PROLA. Some countries do not allow authors to sign unlimited licenses and so a license to publish would be restricted to only currently known uses of the manuscript and its contents. But as technology changes, we would like to have complete freedom in transforming the articles to new formats and repackaging the articles to best serve the physics community.

By granting back to the authors all of the rights which effectively allows them to do as they please with the contents of the article, but with some restrictions on what can be done with our value-added formatting, we feel that we have struck the proper balance between our needs and authors' needs.

The APS copyright transfer form is available from http://forms.aps.org/author.html.

APS ONLINE JOURNALS

I would now like to turn to a brief summary of APS strategies and decisions for making content available online. The key underlying strategy in all that we have done is to maintain flexibility and simplicity in choosing how we make our journals available online. This has served us quite well as should be apparent in the following.

JOURNAL DEPLOYMENT

Since making our first journal available electronically in 1995, APS has hosted our journals on different online publishing platforms maintained by external vendors as well as on our own servers. Most of our "current contents" journals (1999 –present) are now hosted on AIP's

Online Journal Publishing System (OJPS) platform, but this has not always been the case. Early on we experimented with hosting several journals ourselves and we also hosted Reviews of Modern Physics on Highwire Press's platform for several years. This gave us increased leverage in negotiating with the vendors and helped us gain experience in putting content online.

We have also learned a lot by hosting journal content ourselves. *Phys. Rev. D* and the *Rapid Communications* section of *Phys. Rev. B* were first served directly from our editorial office. Today we still PROLA and *Phys. Rev. ST Accel. Beams* ourselves. PROLA also gives us extra flexibility. Currently, our policy is only to serve the 3 past years plus the current partial year in our current contents offerings. Each January we shift one year's worth of content into PROLA. Furthermore, PROLA actually contains all of the current content as well so that it can function as an emergency fail over for the content on OJPS. This also means that we can periodically re-evaluate whether to continue to host our content externally.

One might wonder how we manage to shift content around so much. Doesn't it break links? The answer to this question reveals another part of our strategy of simplicity and flexibility.

LINKING

While we were experimenting with hosting our journals in various places, it was clear that it would be very difficult to link to our journal articles. Each deployment platform had its own URL scheme and, more importantly, the URLs sometimes contained information that wasn't readily available in most citations to our articles. In particular, issue numbers were used in the URL, but in physics, issue numbers are almost never part of an article's citation in a bibliography.

The solution to this is rather simple and it occurs again and again on the web. Namely, one creates another simpler URL that points to a program that then determines where to the user should be redirected. In our case, this resulted in the APS Link Manager, which has a very simple URL structure based only upon journal, volume, and page number. For instance, to link to the article that starts on page 1 of volume 85 of Physical Review Letters, the URL would simply be http://link.aps.org/abstract/PRL/v85/p1.

There are several things to note. First, the URL is quite simple and any article published by the APS (including those all the way back to 1893) can be linked using the same URL. In addition, it points the user to the abstract page, but more generally this is a wrapper page in which there are many useful links and features for the user. In order for this to work as seamlessly as possible, we also developed very early on a policy of making these basic wrapper pages freely available - no subscription is needed. Of course, there are some features on the wrapper page (for example, displaying the bibliography of the article with links to reference works or displaying a list of citing articles) that are only available to subscribers.

More recently we have extended this simple notion to our scheme for Digital Object Identifiers (DOI). APS is a member of CrossRef (http://www.crossref.org/) which was formed by publishers to make linking easier. CrossRef is based on DOIs and the main functions of CrossRef are to register DOIs for publishers and let others discover the DOI for an article by querying with the basic metadata found in a citation. APS link manager links don't really

require this infrastructure, but CrossRef offers other advantages by making it easier for publishers to work together. In any case, an APS DOI is very similar to the link manager URL. The example above would have the DOI 10.1103/PhysRevLett.85.1 and would be accessible through the URL http://dx.doi.org/10.1103/PhysRevLett.85.1. The 10.1103 is a numeric prefix that is unique to each DOI registrant. The DOI URL then resolves to an APS URL, in this case another APS Link Manager URL (same software, but the URL syntax is different), which then takes the end-user to the final destination. So once again, redirection is used to ensure that URLs remain reliable.

Many other publishers do not use such a simple scheme. Rather, they construct DOIs, which have little to do with the bibliographic metadata of the article. It is often argued that DOIs should be opaque identifiers with no discernible intelligence. However, I take the opposite view for journal articles. Once we make an article citeable by assigning it a journal's volume and giving it a page number, this information is permanent. Even if we were to transfer the journal to another publisher, the article would still be cited by the same information. Thus another important concept in our strategy is that an article's citation (resulting from the peer-review process) is a persistent name and it is natural to use it as such.

The other argument that publishers make when they assign opaque DOIs is that they would like to make the articles available online before the print issue is complete. Thus the page number cannot be used as an ingredient for the DOI. We had the same problem at the APS, but rather than introducing multiple ways of citing a paper, we made a radical but simple change: We got rid of page numbers.

E-FIRST PUBLISHING

There are clear advantages to making articles available as soon as they are accepted for publication rather than waiting until an entire issue is complete. Articles can appear up to six weeks ahead of the print issue for some of our journals. However, as explained above, we were loath to introduce multiple citations for the same article. The solution was to do away with page numbers and instead use article identifiers. These identifiers look and function like six digit page numbers (except they may contain a leading zero). There were some constraints on how these numbers could be constructed. First, they must be unique for each volume of course. Second, they must allow articles on similar topics to appear adjacent to each other in the table of contents. Third, users of the printed journals should be able to quickly pick out the correct issue and locate the article.

The final numbering scheme satisfied all three of these constraints. APS article identifiers consist of six digits. The first pair of digits corresponds to the issue number the article appears in (perhaps with a leading zero), the middle pair of digits corresponds to a subject classification scheme for the tables of contents, and the final pair of digits is a sequence number. Thus an identifier such as 112504 would mean the fourth article in the subject area corresponding to '25' in issue 11. It must be emphasized that it is not at all necessary for an end user to understand the internals of the scheme. Instead, these numbers will appear in tables of contents in numerically sorted order allowing one to navigate a print issue in the same way one does with page numbers.

All of the Physical Review journals now use this scheme. In the DOI and the link manager URLs, one simply uses the identifier in place of the page number. More information about the APS Link Manager can be found at http://publish.aps.org/linkfaq.html.

JOURNAL AVAILABILITY

An important aspect of electronic publishing is to make sure the journals are accessible on a 24 x 7 basis. APS treats our online journals and the associated infrastructure such as the link manager as mission critical applications. This naturally leads to additional expenses. There are two prongs to our approach to ensuring the availability of the APS electronic archive. For the short term, we are in the process of establishing a network of mirrors. The first mirror of PROLA was established in May, 2001 at Cornell University and we are evaluating possible European and Far East mirrors.

For the longer term, we have an agreement with the U.S. Library of Congress in which we will deposit a full copy of our electronic archive at the library. This will be available to on site patrons of the library as well as the library's statutory users (Congress for example). Both the short and long term strategies are aimed at making researchers and the library community more comfortable with the notion of a print-free future.

VERSION OF RECORD

Another important change that is a consequence of electronic publishing has to do with which version of an article is the "version of record." Should it continue to be the print version or should it be the online version? APS has declared that the online version is the historical version. There are several reasons for this, the most important of which is that articles now appear online with features that don't appear in print. A simple example is the use of colour figures. Since colour is expensive to print, we charge extra fees to authors for colour figures. This sometimes discourages the use of colour. So we would like to offer authors a choice of having colour in the online version only (authors still need to supply two versions of the figure and perhaps two different captions since naively rendering a colour figure in greyscale may give poor results).

Another example is the use of multimedia in articles. We are developing policies for articles in which video or other multimedia content is an integral (in contrast to a supplemental) part of an article. *Phys. Rev. ST Accel. Beams* has 3 articles that have video content. This journal of course doesn't have a print version, but we would like to extend it to the other journals that do have print. Once we do so, the online version will be richer and it will be necessary to read the article electronically.

LOOKING INTO THE FUTURE

As should be clear, I believe we are still undergoing a transition and that this transition is inevitable. Already we have seen how we have had to take on more post-publication responsibilities – no longer can we just print and mail books. Instead, we are becoming curators for an electronic archive that will have to be useable into the far future. Researchers

are changing the way they communicate, and we in turn are changing the way we work. After almost a decade of changing our processes, we are finally getting to the point where print will become secondary, and, more importantly, instead of just seeing the expenses associated with investments in new electronic infrastructure, we are beginning to see real cost savings.

Ultimately we will need a new economic model. The subscription model is incommensurate with the extremely low-cost arXiv-style dissemination of article content. We will have to find value-added services that overlay such dissemination, which the community will truly find useful, services that will have to recover their own costs. It is an open question whether the expense of peer review and the long-term archiving that we build into our processes will remain valuable to the physics community. If not, we are prepared to move on to other ways of achieving our mission. However, I believe that peer review in some form will remain desirable and that the society will continue to play a role in coordinating it. Furthermore, the issue of archiving is essential and until authors can create true archival documents, there will be a need for us to ensure that a proper archive is built. What better guardian for the physics literature is there than a self-organized society of physicists?

APS has prepared itself for whatever the future brings by building flexibility and simplicity into our electronic endeavours. In addition, we are closely aligned with our membership and the trends they initiate. It is for this reason that I believe we will survive the transition while at the same time transforming how physicists work and communicate.

