Web performance metrics for online journals: monitoring and improving accessiblity

Arthur P. Smith

The American Physical Society

The American Physical Society publishes the Physical Review series of journals (Letters, A,B,C,D,E, and a new special topics series) and Reviews of Modern Physics. Physical Review Letters first appeared electronically, over a proprietary network, in 1995. By the end of 1997 all the journals were online on the World Wide Web with freely available tables of contents and HTML abstract pages, and PDF files for the full text of articles. The full text is actually produced in SGML now, but current web browsers are unable to handle the mathematical notations and other important "look-and-feel" characteristics of the print journals, so we have had to rely on PDF for actual display of the articles.

We provide the journals over the web through three different servers, two of them (AIP and Highwire) subcontracted for this service, and one of them a server under our direct control. This was done because of perceived differences in requirements for the different journals, and in particular the journal that we are running directly (Physical Review D) addresses the high energy physics community where specific functionalities and time constraints were perceived to be critically important. This journal is now the first of our journals to provide "e-first" publication: articles are considered published as soon as they are available electronically, which averages at least two weeks earlier than print and can be a month or more in advance of the print issues.

Requests for the online journals have been growing quite rapidly - as the graph here indicates, monthly requests via the web for articles from Phys Rev D have gone up roughly a factor of 10 in the last year. Articles are available in PostScript and PDF formats - until recently the average size of a downloaded article was between 1/2 and 1 Megabyte. The recent drop in megabyte count is because we now use level-2 postscript, which can have much smaller files due to image compression.

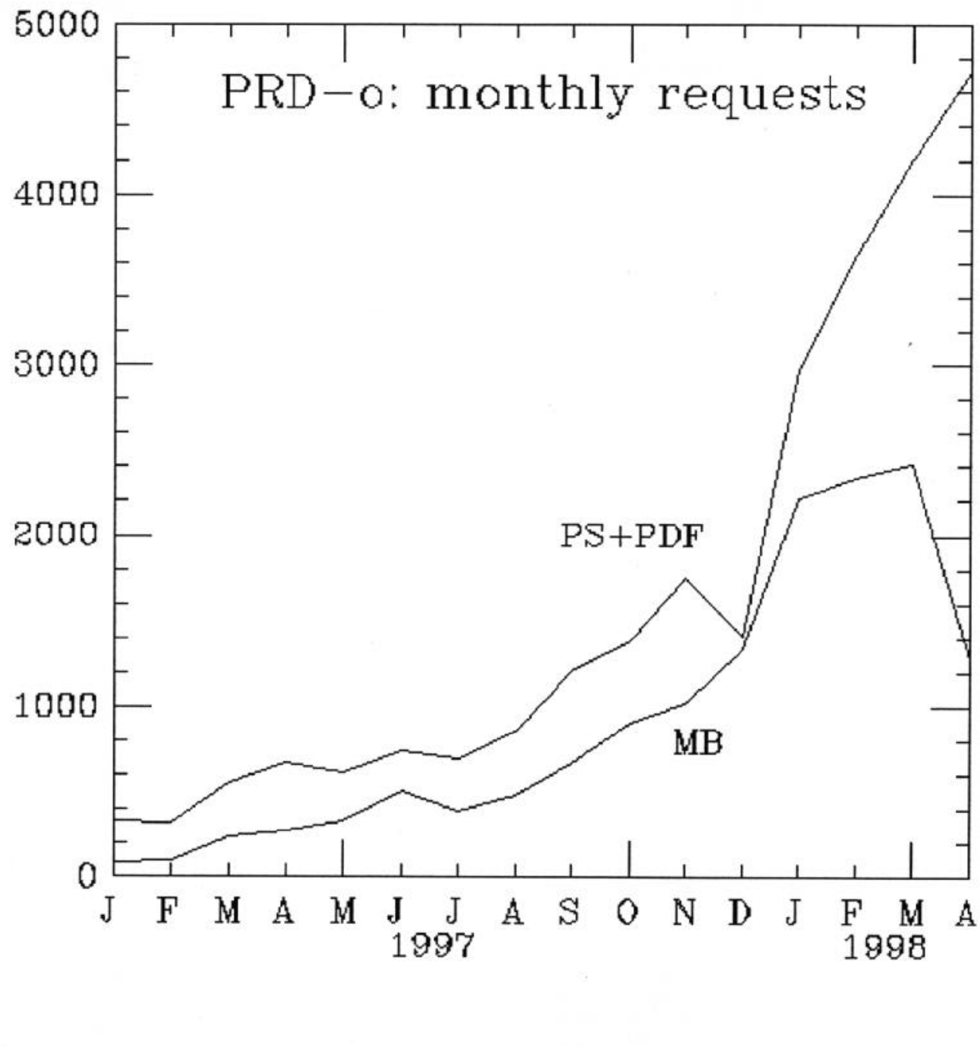

Now that the basic online functionality is there for our journals, we have turned to issues of improving performance. In particular, the most glaring performance problem is the length of time it takes a typical user to download a paper from one of our journals. Many of our users see persistently slow connections (1 kB/sec or lower) so downloading a file approaching 1 MB in size can take from minutes to hours.

The accompanying graph shows the number of requests (normalized per factor of 10 in time) for PRD-o papers in March 1998 that took roughly the time shown to download. These times were measured with the "%T" configurable logging directive in our Apache server. About two thirds of the downloads take less than 30 seconds, but there is a long tail to very long download times, which afflict many users on a regular basis.

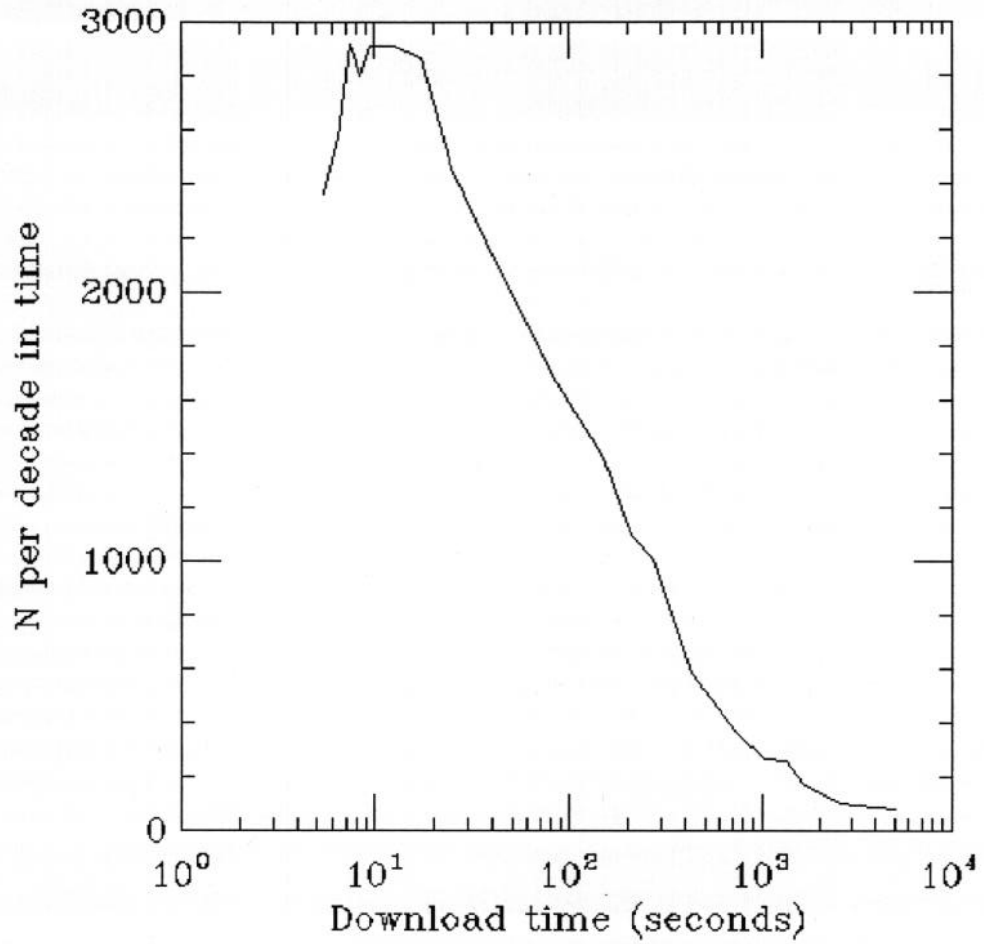

Our goal is to get times down to a few seconds, which should remove the perception of slowness. This means cancelling at least one, and in some cases two orders of magnitude in the download time. Perhaps a factor of 2-4 could be gained by reducing file sizes - this may be possible through separate distribution of font information, for example with a compressed XML. Byte-serving PDF files can bring the initial chunk needed for display down to the 30-50 kB range, so that also can give a perceived factor of 2-3 (not "real" since you do not receive the entire file, only one page at a time). But most of the improvement is going to have to come from getting data transfer rates for everyone up to the 20 kB/sec level - and this clearly depends on details of network connectivity. How much control can we have over this?

Later this year we should also have available an archive back to at least 1985 with scanned images of our journals - these are 4-10 times larger than the current PDF files being downloaded, so users will need 100 to 200 kB/sec data rates to view these files without feeling that the connection is detrimentally slow.

Even 20 kB/sec is large compared to typical modem speeds - a 28.8k modem sees at most 3.6 kB/sec - so we have to face the fact that until home users typically have ISDN, DSL or cable modems, improving our data capacity will have at most a limited effect on them. However, most of our current users are primarily at universities, laboratories, or institutions with, in principle, much faster connectivity. Since we (so far) never have more than a handful of users simultaneously downloading files, the bottlenecks are somewhere after data leaves our T-1 (~ 200 kB/sec) line and before it reaches the university or laboratory's local internet connection.

Before trying to improve things, we needed a way to test or monitor to find out where the problems were. We explored four different approaches to the testing/monitoring issue:

1. Requesting that some of our users (particularly from outside the US) time downloads of large files.
2. Using the Apache web server's '%T' logging directive to record the time taken by every transaction.
3. Using a service that regularly requests one of our web pages from tens of different locations around the internet.
4. Sending relatively large files by FTP to several test sites around the world.

1. User testing

We conducted two user tests, one in March/April 1997, and the second in December 1997. Both tests had roughly 30 participants, a few percent of those solicited. In both cases we made data available via two or more "mirror sites" (more on this below), and requested that users download files, sending us the timing information via a web form.

Even though a few users performed the tests multiple times, we soon realized even 30 volunteers was not enough to get more than a few snapshots of worldwide internet performance. More on the test results below. At the least, it was clear that more systematic performance measurements were in order.

2. Web logs

Recording the timing for all downloads produces at least a large quantity of data. We encountered several problems however that required some analysis to identify and resolve.

1. Accuracy of the logs. Times are recorded only to the nearest second, resulting in errors for small files and fast connections. Sizes of files sent are recorded accurately (in bytes) except when a file was only partially downloaded - in that case the TCP send buffer caused the recorded size to exceed (by up to the buffer size) the quantity of data actually transferred. We attempted to avoid both problems by only measuring data rates for "larger" files (> 10 kB).
2. Request method. Some "POST" requests sent large quantities of data from the user to our server, but this size was not included in the bytes sent recorded in the log. Avoiding this problem meant restricting the analysis to the "GET" requests.
3. Self-selection in quantity and timing of usage. People come in more frequently at "off-hours", when they know the connection is faster. Also, the faster the link, the more they tend to download. This was evident in comparisons of our regular and mirror server logs, where when one was significantly better it tended to have at least double the traffic from the same group of users. This means that simply averaging over all usage will understate any connectivity problem. We eventually resolved this problem to some extent by averaging hour-by-hour on a weekly basis - but this only works if the average is restricted to a group of users with similar connectivity.

In the end the web logs provided (and continue to provide) valuable data, but the above problems complicate the analysis enough that we regard the results as indicators, rather than as accurate performance measurements.

3. Monitoring service

We evaluated "Keynote", the producer of a weekly "Internet business 40" index, as a monitoring service for our web site. They would download the page we requested every 15 minutes from several dozen sites around the country. Again it provided some useful information, but we decided not to use their service for several relatively minor reasons and also because their international coverage was very limited (2-3 sites in Europe at the time). Their averaging method, looking at the average of the time taken to download a page, is probably a pretty good measure of user aggravation, but can be too strongly weighted by a few bad connections.

4. FTP testing

By contacting a handful of willing participants in Europe and Japan we were able to run our own version of the Keynote tests, using FTP rather than the Web for convenience (and also to see if HTTP and FTP performance differed significantly). The test consisted of transferring a 100 kB file to the test sites at regular periods throughout the day, and ran through most of the month of February. Being a straightforward and well-controlled test, we found it much easier to analyze the results.

Note that with the several orders of magnitude range in performance, different methods of averaging can give results different enough to affect the conclusions drawn. For example for our French ftp test, an average of the data rate values over the whole week gives (0.86 kB/sec - regular web site)/(1.02 kB/sec - mirror site), suggesting that the mirror is not much better than the regular connection. On the other hand, using the inverse of the Keynote-style average of the time (user aggravation) gives data rates of (0.28 kB/sec)/(0.90 kB/sec), indicating a tremendous improvement with the mirror. Rather than taking averages, the results tabulated below try to give a picture of the distribution of download times (for example, the fraction of the time that a 100 kB download took more than 5 minutes).

Lesson 1 Simple averages can be misleading.

Results

One of the "mirror sites" we tested in early '97 was with the Digital Island service, also used by Highwire Press at Stanford University. Rather than actually locating the data in multiple countries, their concept was to use a single server with good connections to a wide variety of overseas networks. Indeed, at some times the user testing in early '97 indicated bandwidth improvements of 2-3 orders of magnitude, particularly for users in Japan.

Lesson 2 Poor internet performance can be immensely improved by being "close" to your users.

When we started regularly using this service in late '97, however, several things had changed, as was quickly evident in the web log "indicators" and the December user tests. Our Japanese and German users were seeing considerably improved connectivity to our regular web site because the research networks in those countries (SInet and DFN, respectively) had significantly upgraded their connections to the US in the interim. Both groups were also seeing very poor connections to the Digital Island network because those research networks were not well connected in their own countries to the commercial networks Digital Island had contracted with. As a result, our new "mirror" actually had worse performance than our old US site, for those users. The German situation soon corrected itself with the installation of a new major internet exchange in that country, but we still have problems in Japan. The February FTP tests give a reasonable picture of these issues in the table below.

FTP download timing distribution (old:mirror numbers)

Country	> 10 min	> 5 min	> 1 min
Japan	0:9	0:15	26:38
Germany	2:0	2:0	24:2
France	36:2	46:3	168:168
Russia	45:1	57:3	89:83

This shows the number of hours in a week (out of 168) that the time to download a 100 kB file was greater than the indicated number of minutes. We would like these times to get down to the range of just a few seconds, but for now minutes are the relevant time scale to measure user aggravation. These numbers show that while our "mirror" was effectively available from a commercial network in Tokyo, in fact our Japanese academic users had a much easier time getting to our regular US-based web site.

Lesson 3 "Close" doesn't necessarily mean geographically close. Being on (or well connected to) the same network as your users is critically important.

While tracking down these problems we realized that the long round-trip-times (1/2 second or more) could also have been limiting our bandwidth - interestingly also related to the TCP send buffer size. We increased that (a server configuration parameter) in mid-January and saw an immediate improvement in measured bandwidth to many locations.

Lesson 4 Server configuration really can have an impact on web performance.

Future

Our next step is to try to get ourselves on the "right" networks, both in the US and overseas. The following list gives the networks used by the top 100 institutional users of our web site in March 1998:

1. DFN: 17 institutions (German research network)
2. MCI: 13
3. ESnet: 9
4. Teleglobe/NORDU/Ten-34: 9 (Canadian net with ties to north Europe)
5. UUNet: 8
6. SInet: 7 (Japanese academic network)
7. BBN: 6
8. CERFnet: 3
9. AT&T: 3
10. JAnet: 3 (British academic network)
11. Renater: 3 (French academic network)
12. ANS: 2
13. NYSERNET: 2 (NY State network - close ties to Sprint)
14. Digex: 2
15. Sprintlink: 2
16. AGIS: 2
17. USPnet: 2 (Brazilian network)
18. Genuity: 1
19. KREN: 1 (Korean research network)
20. IMnet: 1 (Japanese government network)
21. Telintar: 1 (Argentine network)
22. IIJ: 1 (Japanese commercial network)
23. SURF: 1 (Netherlands network)
24. IIX: 1 (Israel)

Note that ESnet is quite prominent. In fact, lanl.gov was the top user site in March. Across the online Physical Review journals, the following are the rankings of DOE labs among the top 25 users of each journal (as of March 1998), in terms of downloaded articles:

Rankings of DOE labs among top 25 users of Physical Review online

Lab	PRA	PRB	PRC	PRD	PRE	PRL	Overall
lbl.gov	-	16	4	21	-	3	4
bnl.gov	-	12	3	10	-	12	7
anl.gov	-	14	-	-	-	14	13
nrel.gov	-	8	-	-	-	-	29
ornl.gov	4	-	13	-	-	-	41
cebaf.gov	-	-	21	-	11	-	48
lanl.gov	-	-	-	14	-	-	61
fnal.gov	-	-	-	23	-	-	73

There are also 2 NIST locations and one DOD site in the top-25 rankings. DOE labs generally account for roughly 10 percent of our web volume, or roughly a quarter of all US activity.

Of course, the future for academic and government research in the US also holds the promise of Internet-2 and the NGI project - we hope to be able to provide our journals over these new networks at the earliest possible date, and then we may even reach our goal of 100-200 kB/sec to the user.