

THOUGHTS ON PAPER PUBLISHING IN THE DIGITAL AGE*

Sanjeev Arora
Princeton University, USA

<http://www.cs.princeton.edu/~arora>

What role should journals and conferences play in the age of arxiv, twitter and other yet-to-be-invented digital wonders? Detecting among many colleagues a general impatience with the status quo, I wrote this article to generate more public discussion. It is addressed solely to colleagues in theoretical computer science, not other fields. I also don't address legal issues such as copyright, costs, and free access because these have been extensively discussed elsewhere.

Despite its small size, theoretical CS has been remarkably successful. An incredible edifice of ideas was created together with an open culture that values the need to address papers and talks to nonspecialists. This allows ideas and techniques to jump rapidly across subspecialities. We should think hard about how to best continue that culture in the digital era. (Impact on promotion/tenure systems also must be carefully weighed.)

Below I will survey the major proposed approaches and their pros and cons, and my own thoughts on them. Anybody who has investigated this topic quickly discovers that it suffers from the *boolean algebra obstacle*: Given n binary options, you will find supporters—and good arguments—for all possible 2^n combinations. There is likely no universally accepted solution. I therefore propose starting a specific public process for continuing this discussion.

1 The three approaches

Let me list the three broad approaches I have encountered.

Arxiv-only. Conceivably, *both* journals and conferences could today be replaced by arxiv and other repositories. Some proponents of this view point to instant “impact metrics” (page views, number of tweets, etc.) this solution comes with, seemingly tailor-made for hiring and promotion cases. Machine learning researcher

*This article first appeared as a blog post on windowsontheory.org in October 2013.

Yann LeCun has taken this viewpoint to its logical end¹, arguing for a “free-market” system whereby papers appear only in digital repositories and are subject to a distributed market model for refereeing/commenting. Independent consortia of reviewers (basically a new name for journals and conferences of the future??) would decide to publish reviews of arxiv-ed papers, with or without the permission of the papers’ authors. Recent articles in *Nature* also explored such ideas (see the bibliography).

Journals-only. At the other end of the spectrum, Lance Fortnow is troubled by the proliferation of conferences and steady decline of journals. He thinks this has reduced research quality. In fact, since digital repositories allow very rapid dissemination, he thinks conferences are useful only to meet others and catch up on the latest research. He seeks to restore the importance of journals, lower the prestige of conferences by greatly raising acceptance rates, and have one flagship conference that everybody feels compelled to attend.

Fortnow’s views resonate with some colleagues I have talked to. Another point they raise is the costs (money and time) associated with conference presentations, which perhaps penalize less established researchers. This may be true, but we need more data on this point.

However, I cannot help noticing though that the arguments in favor of journals sometimes overlook (perhaps unconsciously) the possibility that the journal-only vision can be pursued with many radically different approaches: each of Math, Physics, Biology, Economics etc. has a different one. It is unclear *a priori* if all of these approaches are better than our conference-based approaches (anecdotally, colleagues from some these other fields do like aspects of our top conferences).

The current hybrid approach, adapted as needed to the digital era. This is the approach I have ended up favoring. Its advantage is that it can borrow good ideas from all other proposals while causing the least upheaval.

2 Is Arxiv enough?

Physics is one of several fields that have taken enthusiastically to e-publishing. A physics paper on arxiv may have follow-up papers within weeks or even days. Not all papers may appear in journals, and many that do, appear in truncated form.

While such arxiv-based models have plus points, one also sees dangers:(a) incentive to write shallow and incremental papers; (b) more priority disputes and the temptation to publish sketchy ideas in order to later claim full or at least partial credit; (c) lack of incentive for good reviewers to volunteer time reviewing enough papers (despite the attempted analogy to a free market in LeCun’s proposal).

¹<http://yann.lecun.com/ex/pamphlets/publishing-models.html>

Let me elaborate on (b). The following already seems to be an axiom among my younger colleagues: “If a result appears on arxiv, you have a few days to put up your independent manuscript. After that you can’t claim independent discovery.” Some other disciplines have already moved on to a more cut-throat model: “Whoever gets to arxiv first wins.”

Surely, this must incentivize hasty writing and incomprehensible papers. Or maybe papers that even have errors, to be fixed in subsequent revisions. In the past we relied on conference committees and journals to adjudicate such disputes. How would that happen in a distributed market? One can imagine systems with feedback buttons, reliability ratings etc., but it seems a dubious method of doing science.

Several arguments have been cited in favor of the “free market reviewing” approach: unknown researchers publishing on equal footing with established ones; the superiority of “wisdom of the crowds” over the PC’s “groupthink.” But I suspect in reality things may turn out less fair than the current system. Prominent researchers with bigger megaphones (e.g., with more blog readers or twitter followers or friends willing to review their papers) will tend to benefit. Power centers will inevitably form.

By contrast, conferences in theoretical CS—perhaps because each PC is a fresh set of 20-25 individuals—have a good track record of showcasing great work by grad students and postdocs. Papers by unknown authors may get awards while those by Turing award winners may get rejected. (Aside: Remember, my comments pertain merely to theoretical CS. Possibly the dynamic was different in other fields, necessitating a switch to double-blind refereeing systems²)

To sum up, while very useful, digital repositories do not seem to me adequate replacements for conferences and journals.

3 Conference or Journals?

Historically, conferences came to dominate computer science because they allowed fast dissemination. Today, this goal can increasingly be met by other means, so how can we still justify conferences? Fortnow’s question is extremely pertinent.

Below, I list several reasons why conferences still make a lot of sense to me. My focus here is on promoting better, more creative science—I worry less about promotion/tenure policies since they will quickly adjust to accommodate any new dissemination method we choose (including arxiv and twitter). Also, I apologize in advance for occasional forays into pop psychology.

²For a quick introduction to arguments for and against double-blind reviewing systems see John Langford’s blog post at <http://hunch.net/?p=2656>.

A recommendation about which papers to read, plus 20-min introductory talks

This is arguably the biggest plus of our top conferences: providing us a map to help us navigate through the sea of papers that get written (especially outside our subspecialties). As a bonus, they force authors to write a 8-10-page description of this work, and record a 20-min introductory talk on it.

The 20-min talk has been good for our field both by propagating ideas across subspecialties, and training young researchers to interact with nonspecialists. In fact, the talks are the main reason I attend STOC/FOCS (though the social networking is a nice bonus). Of course, I could stay at home and watch videotaped talks but, really, who does that?

A caveat: Inevitably, any recommendation system—conference or journals—will end up shaping the research directions, and will be used for tenure and promotions. This vests some power in the recommendation system. But at least in the conference system the power rotates to a different groups of 25 people each time, instead of staying centralised with a (relatively fixed) journal board.

Equally obviously, the quest for such recommendations engenders competition. Others see this as problematic (notably, Oded Goldreich³), but I am more sanguine—humans (even toddlers) are naturals at competition, and also at cooperation.

Incentive system for researchers to produce a substantial piece of work, and then write it up, sort of comprehensibly, in 10 pages

The incentives in the arxiv model are quite the opposite—more frequent, insubstantial, and hastily written works.

Incentivizing substantial works is also the frequent justification for the move to journal-only model (as in Fortnow's article) but we should realize that while this may be true for some fields—maybe Math—it is not uniformly true across other journal-only fields.

The 10-page limit—archaic relic of the papyrus era—and the PC review model has led to our tradition of writing papers that are sort-of comprehensible to nonspecialists. Journals possibly reward a writing style geared to specialist reviewers.

³See <http://www.wisdom.weizmann.ac.il/~oded/PDF/struggle.pdf>

Incentive system for good researchers to (sort-of) review lots of papers

Every field has to find its own balance between refereeing capacity and number of papers that get written⁴. "Conferences vs Journals" is too simplistic a way of thinking about this issue, since journals in different fields differ a lot.

Journals in the life sciences (and many fields of physics) have more stringent page limits and faster refereeing time than our conferences. Volume of papers is high, and quality of reviews is variable. Economics journals have less stringent page limits (but still more stringent than our journals), require many rounds of rewrites, and consequently, have huge backlogs. This can make it difficult for young researchers to get published—many have no publications when they finish their PhD. Also, accept rates of < 10% in top journals force editors and reviewers to become risk-averse (according to my friends). Math journals seem to work well and have reasonable turnaround times, but then the system vests a lot more power and prestige with the editorial boards of top journals. Also, mathematicians seem to write fewer papers and are willing to spend a lot of time on refereeing.

The PC review system in theoretical CS is not perfect, but this semi-refereed model represents our own best effort to balance refereeing capacity with the number of papers. It is generally considered bad form to turn down requests from PC members to review papers. Anecdotal evidence suggests to me that the quality of reviews at our top conferences is not worse than in other fields with a similar volume of papers.

Also, many researchers seem happier to serve on a STOC/FOCS PC once every 3–4 years rather than on a journal board for 3–4 years. Maybe humans prefer shorter but more intense pain to a longer and less intense one. Or perhaps journal boards are less interesting because you end up handling papers in your own subspeciality, including those you already saw 2+ years ago.

Clearing point for deciding upon priority, novelty, correctness etc. of claimed results

Conferences can do it faster and better than journals in most cases (but the systems are different: a jury of 20–25 PC members + specialist subreviewers versus a jury of one editor and 3 specialist reviewers). The informal refereeing system at conferences at first glance seems to invite abuse but I can think of very few accepted papers at STOC/FOCS in the last 30 years that turned out to be seriously flawed (and often those were recognized as controversial when accepted).⁵

⁴Fortnow and others have suggested that conferences cause the field to produce too many papers. This thesis deserves wider discussion; I suspect tenure/promotion/grant policies play a bigger role.

⁵Again, my conclusions are not universally applicable. Colleagues in some other fields think that flawed papers slip too often into their top conferences. Even in theoretical CS, fields like

A synchronization mechanism for our field

Is it just my imagination, or do conference deadlines actually enhance collaborations and improve creativity? Half-imagined results get fleshed out as people get together in the months or weeks before the deadline (and I am not referring to caffeine-fueled late-night finishes, which I avoid). We need this synchronization to structure our busy lives, and neither arxiv nor journals provide it. If you don't care for the human weaknesses this argument stands on (or believe in a platonic ideal of scientists engaged in heroic quests unfazed by mundane complications) I should mention Boaz Barak's alternative explanation: sometimes correlated equilibria⁶ are superior to Nash equilibria.

4 Needed: Process to rethink the conference system

I believe that our conference system led to the current innovative culture in theoretical CS that allows radically new research directions to pop up every few years. It would be good to update these conferences for the digital era, while maintaining their best qualities.

Currently, the process for improving the conference system has been delegated to PC chairs and has resulted in some welcome experimentation: a day of workshops, poster session, recorded talks, no paper proceedings, better feedback for authors, etc.

But incremental experimentation has the drawback of sowing confusion and cynicism. For example: Is the 10-page conference version important or not? Should the number of accepted papers be increased or decreased? It is hard to know by experimenting for a year or two. Also, beware of the *boolean algebra obstacle* mentioned earlier. There is likely no universally accepted solution and so any structure arrived at via experimentation will appear ad hoc, and garner little respect.

The only feasible alternative seems to be a *centrally computed solution*. ACM SIGACT (maybe in partnership with EATCS) should create a group of, say, 7 – 9 people —suitably diverse in terms of seniority level, areas of expertise, gender, countries of residence etc.—that will suggest a future blueprint for theory conferences. This group should produce such a document after conducting an open discussion, while accepting public input —including fleshed-out proposals—via a moderated forum. Such a public process could have a rejuvenating effect on our field and its conferences.

cryptography seems to need more careful refereeing.

⁶http://en.wikipedia.org/wiki/Correlated_equilibrium

While no single solution will be perfect or universally loved, at least if it is found via the above open process, we will all come to appreciate the competing factors being balanced.

4.1 My own thoughts for further improvement

Here are some of my thoughts, but I remain open to other opinions.

- Keep the conference format (say 12 pages, 11pt) —independent of whether or not the conference has a published proceeding. Why: a) Promotes tighter, focused writing; (b) Improves our ability to keep abreast of research outside our specialities.

But to reduce work and give authors an incentive to produce a readable version, make the submission format identical to the published format (this is the norm at ML conferences).

Having done my share of grumbling about the conference format, on balance I think it is important for our field that 50-page arXiv papers should be accompanied by shorter, more readable, versions. If you think 12 pages is too few, try vying for the privilege of publishing your result in *Science* and *Nature*, in 2–4 pages!

Unfortunately, we seem to be sliding into a hodge-podge system which I attribute to poor communication and feedback by conference PCs (e.g., unwillingness to penalize bad writing).

- Respond to the arXiv challenge.

For example, why delay the conference until the proceedings are ready? Have the submission deadline 3 months before the conference — this gives plenty of time for review and revision. The proceeding is less and less important, and could even be turned into a fully reviewed "journal."

A related idea (tested in fields such as databases and ML) combines some of the speed of arXiv with reliable timestamping. The conference has two or more submission deadlines a year—papers appear on the website as soon as they are accepted in the first cycle, and papers rejected in the first cycle cannot be resubmitted for another year. This spreads out a PC's work over a longer period, whose pluses and minuses have to be studied.

- Increase number of acceptances moderately to accommodate the increase in the size of the field. (Beware though of Parkinson's law⁷: submissions

⁷Parkinson's law: "Work expands so as to fill the time available for its completion".

increase to fill all available refereeing capacity. So don't agonize if acceptance rates stay below 30% despite this increase.) Do not treat all accepted papers equally: have some presented in plenary sessions, and rest in 3 parallel sessions⁸.

Stretching FOCS beyond 3 days seems problematic since it is in the middle of the Fall term, but possibly STOC could be moved into June and lengthened to 4 or even 5 days? (Unfortunately this would bring it into even more direct conflict with ICALP than it currently is.)

- Avoid the temptation to delegate most reviewing to students/postdocs and conference committee work to junior people. In the past this was done as a welcome way to empower junior people, but possibly this correction may have now gone too far.
- Prepare a clear statement on how priority/correctness should be judged in the age of arxiv. Avoid a slide to the system "First to get on arxiv wins."
- It would be nice—perhaps independent of conferences—to have a forum for posting reviews/comments on theory papers. (Hints of LeCun's ideas here. Tim Roughgarden's phrasing is that he'd like to see "more radio channels.") To be useful this forum must avoid the vicious smallness of blog comments. Requiring users to use verifiable identities should preclude the worst abuses (the system only needs to scale to a couple thousand users).
- Poll senior researchers for changes that would attract them back to conferences. I know several who have stopped participating in STOC/FOCS (people 40 and older seem a distinct minority even on PCs in recent years, and certainly on the conference floor).
- Last but most important: keep the various points made in this article (or any other set of principles discussed and agreed upon collectively) in mind when proposing new changes.

5 Needed: less cynicism

I must admit that when I started this thought process and discussions with colleagues, I started out somewhat skeptical of conferences. I ended up strongly in favor. I decided to be more proactive in combating the cynicism or pessimism I often see in such discussions; hence this essay.

⁸Interestingly, Fortnow and Goldreich have made similar suggestions. Thus I agree with some of their "remedies" even though I don't agree with the "diagnoses."

People's views tend to be colored by their last conference rejection. Typically, senior people complain about inexperienced PCs valuing technical sophistication over conceptual contributions. Young researchers wish to publish more to establish themselves and feel anxious about being judged by a power structure that they don't fully understand or feel part of. Such anxieties have existed since prehistoric times—there is no way to do research and not have it be misjudged at times. Cynicism is not a good response. Stay involved!

6 Bibliography/Further reading

- Lance Fortnow. *Time for computer science to grow up*. Communications of the ACM 52(8):33–35, 2009.
- Yann LeCun. *A New Publishing Model in Computer Science*. Pamphlet, available from the author's web site⁹, that proposes a new publishing model based on an open repository and open (but anonymous) reviews.
- *The Future of Publishing*. Nature 495(7442), 28 March 2013. This special issue of Nature is devoted to the topic of this article¹⁰.
- W.S. Brown, J.R. Pierce and J.F. Traub. *The future of scientific journals. A computer-based system will enable a subscriber to receive a personalized stream of papers*. Science 158(3805):1153–1159, 1 December 1967.

Acknowledgements I have had useful initial discussions with colleagues, such as Moses, Bernard, Avi, Mark, Zeev, Boaz and Ankur. I have also benefited from online feedback from blog comments and personal communications from many other colleagues.

⁹<http://yann.lecun.com/ex/pamphlets/publishing-models.html>

¹⁰For instance, see <http://www.nature.com/nature/journal/v495/n7442/full/495437a.html>.