Publish and perish: why the current publication and review model is killing research and wasting your money

Fabio Casati, Fausto Giunchiglia, Maurizio Marchese
University of Trento
*firstname.lastname@*dit.unitn.it

Note: this is preliminary work (version 1.0, or rather 0.9). We release it anyway according to the concepts proposed in this document.

The research world, and specifically the academic world, is centered around the notion of *publication* as the basic mean to disseminate results, foster interaction among communities, and achieve international recognition (and career advancement).
Publications are done in conferences or journals, and are usually reviewed by a committee of experts, also referred as "peers". Typically, each paper is reviewed by 3 or 4 reviewers. The "best" papers among all the submitted ones are then accepted for publication in the journal or in the conference proceedings. In the computer science area, people typically publishes a dozen paper per year, and submit a little more than that (not all papers are accepted the first time around). Acceptance rates for conferences are often around 20% or lower[1].
There are three drivers behind this model:

1. **Disseminate ideas and make them visible**. Through publication and review, papers are made known to colleagues, and the review process is supposed to ensure that the best papers are more visible, so that researchers know where to go (good journals and conferences) if they want to read literature on certain topics. Publications also have legal implications as they "timestamp" work and ideas.
2. **Get credit, recognition.** Having papers accepted at prestigious conferences and journals is a way to prove (in theory) that the work is valuable. This in turn is a major criterion to determine career advancement.
3. **Meeting and networking.** Publications and conference participation leads to exchange of ideas with colleagues, and to networking. Conferences are also very useful for students to come and learn how the research community operates.

Highly Inefficient Publishing Process. This model is incredibly inefficient under every perspective, and results in a colossal waste of public funding, and forces researchers worldwide to waste countless hours that could be devoted to better research (or to have fun with family and friends). It is a system deeply rooted in the past, oblivious to the advent of the Web and related new forms of communication, information sharing, social networking and reputation. Here are some problems with the current state of affairs:

– **Too much time is spent writing papers rather than developing research**. Dissemination of results is important, and writing problem statements and results in a clear manner is also important. It is in integral part of the research work. This being said, one thing is to write papers with the purpose of making results available, and another is struggle to package and "sell" the work to try to get the highest number of papers published in the best conferences (or, in those conferences that guarantee career advancement in a certain institution). The latter is a huge effort and

[1] http://www.adaptivebox.net/research/bookmark/CICON_stat.html

often results in papers that are incremental work with respect to previous research by the same authors.

- **The reviewing process kills good papers and is inherently flawed**. In general, reviewing a paper is not easy, and it is rarely done properly. There are many problems with the peer review process today:
  1. Judging the impact of a paper is very hard, in general. Even smart people and great researcher have a hard time assessing whether a topic is interesting and relevant and likely to have an impact. See the reviews of the famous papers by Dijkstra on Goto statements, of the paper by Codd on the relational model, and many others [Santini, 2005].
  2. Sometimes good papers are cut because of bad reviews. It is not unheard of to have a paper rejected by a conference and win the best paper award at the next one. The main reason is that only one bad review is often enough to kill a paper. Reviews are often inconsistent, sometimes an author gets reviews criticizing the paper and saying opposite things.
  3. There are reviewers who are generally more negative and some that are more positive. So it is often a matter of luck to a certain extent whether your paper gets accepted. Clearly good papers eventually go through, but sometimes late and after a lot of reworks.
  4. Reviewing takes time, and is not necessarily time that results in better papers. Reviewers, especially scrupulous ones, spend a lot of time in doing reviews, and authors spend a lot of time adapting and tuning the paper not so much for the sake of making the best possible explanation, but to please reviewers and the conference style. While improving papers following comments is a good thing, very often one has to fight with meaningless or contrasting comments as well as space limitations that make the whole work cumbersome. Furthermore, sometimes there are certain styles of writing papers that is better accepted by reviewers, or that reviewers feel particularly bad in rejecting.
  5. A common effect of this review process is that many conferences tend to accept very detailed papers resulting from very detailed studies, rather than more innovative and creative papers.

- **Limited dissemination.** The entire review process itself limits dissemination (unless people post the papers on the web, which is a different kind of "publication", and likely a more appropriate one): reviewing introduces delays and if the paper is rejected then 6 more months will pass till the work has the chance to be published. Moreover, and very curiously indeed, research sponsored with public money is given to private publishing companies that profit from it and that sell papers. Furthermore, although it is nice to have papers in front when hearing presentations, printed proceedings by institutions tend to increase the cost of conferences.

Furthermore, the current publication model, and even the notion of "publication", are rooted in the past. If academic research was born after the Web, we would not even be talking about publications as they are today. With a printed paper model, typical of journals, one needs to have the notion of publication, which happens periodically.
If the authors do some extra work or have new findings, they need to write another paper, they cannot update or extend the current one. If people want to comment or discuss on the paper, they need to do this via email and via private discussions with the authors. Of course there is the issue of how to evaluate and give credits to people, but that is a separate matter. With the Web, this is not the case, and there is no reason for the "publication" model to go on unchanged.

Failures of the past. Despite these very significant shortcomings, the research community has been unable to come up with a better model. This is certainly also

because the problem is hard in itself, but we suspect a significant reason is that people respected in the community are successful in the current system, and hence are not very interested in changing it. Besides, people are always so busy writing papers that it is hard to take a break and think about creating and pushing for a better system.
This does not mean to say that no attempts have been made or that the problem has not been studied. Over the last decades, there have been a few attempts to experiment with different models as well as to study in a scientific way the effectiveness of the current approach to paper evaluation and publication.

In terms of conference models, variations include:

- **Peer-review with rebuttal** (e.g., ICSOC'05) **or double blind review** (e.g., Sigmod): unlike traditional conference review models where authors cannot reply, some conferences are experimenting today with rebuttal,, where authors have a few days to reply, in a few lines, to the reviewers to correct errors in the review. In theory, this is used as input in the discussion among PC members. In practice, rebuttal rarely leads to reviewers changing their minds, but it affects PC chairs when making decisions and, most importantly, leads to better reviews in the first place. Double blind reviews occurs when reviewers do not know the name of authors. There is contradicting research on whether double blind improves the fairness of the selection process.
- **Community review** (e.g., eclipseCon 2006): the community can vote on papers or on abstracts. There is no restricted program committee, the community decides what they want to be presented. This approach had very little success, for reasons yet to be fully studied and understood.
- **Open** (e.g., INFORMS): There is little to no selection, everybody can go to present. Participants can read abstract and exercise their own judgment with respect to what presentation they will listen to. Open conferences do not assign credit to the papers, though they are great for dissemination and networking.
- **By invitation** (e.g. in physics): the conference organizers invite people to come and give presentations. This appears to be good as it is a freeform way for the community to select top researchers to come to conferences. However it is not clear how to distinguish good conferences/meeting from average ones and at times, if people are not serious, it may be more based on friendships rather than scientific merit.

Journals also experimented with alternative models. The most significant one is ETAI, where papers are first put online and then reviewed, with comments openly posted on the pages (open reviewing) before a review process begin. For reasons that are still unclear, but probably related to the fact that people were posting comments in the open, this approach did not succeed and ETAI stopped publishing in 2002.

In terms of research on this topic, a few papers have been published on various aspects of the reviewing process, sometimes with contradicting results (see e.g., papers on double blind reviewing or repeatability of the review process [Tung, 2006; Madden, 2006; Fisher, 1994; Rothwell, 2000]. The conclusions are sometimes contradictory. There are no indications on which review process and model works best and no clear evaluation of benefits and shortcomings of each, so that program chairs and journal editors are still left in the dark and, in the absence of a clearly stated "better way", proceed with the status quo. This is often the approach that generates the least discussions: even if most people want a different model, they disagree on which one, so in the end it is sometimes just "easier" to keep going with the same old approach. However a large-scale study is still missing, and contributions mostly focus on small samples of reviews.

Thoughts towards new models *extreme writing and paper as software.* We are in the initial stages of an investigation on innovative publication and review model. Our exploratory direction will be initially based on two main ideas:

Separate the dissemination, evaluation/recognition, and retrieval aspects: today, with a publication, researchers achieve all of them. A publication disseminates the work, causes recognition for the authors (the peer evaluation recognizes it as quality work), and makes the paper "visible" in that people can look on papers published in "good" conferences or journals if they want to find "good" work in a certain area. However, there is no reason for these three aspects to be tied now that dissemination is not necessarily related to the physical, paper printing of the scientific contribution in a journal.

Extreme writing and papers as software: we can make a parallel between paper writing and software development. In software, the code is developed and then improved. New functionality is added with time, and the artifact is released and then improved. In extreme programming approaches [Beck, 1999], the code is also "evaluated" quickly in the process, rather than waiting till development is complete. Taking into account differences that do exist, one can borrow ideas from software development and try to apply them to writing. In software development, we do not change the name of a class each time we make a change to a function. We just release a new version of the class. Once a certain amount of functionality is developed, then the code is released for "testing". Similarly, with scientific papers, an approach that seems sensible is to publish versions of the paper when the work is sufficiently mature and clear so that somebody can read and gain insights from it, and then improve it. More importantly, minor changes (delta contributions) should not result in yet another paper (class) and yet another set of peer reviews as it is always the case today, but in variations or extensions to (versioning of) an existing work.
Of course the development of a large program is a cooperative effort, while researchers compete more than cooperate, so this has to be taken into account. One sometimes does not want to release initial ideas for fear that they are copied, but usually this does not happen and whoever posts a version of a work has a significant lead on others. Besides, early posting, coupled with a secure and community trusted timestamp mechanism, gives people the right to claim that they have been the "first" to a certain discovery. Furthermore, the researchers keep the control on when they want to release the new version of a paper. Needless to say, early releases contribute to science more than late releases.

Other interesting analogies are with web search and open source software development.
Open source development can provide interesting insights for the way people cooperate to provide feedback and improve the development. Again this is challenged by the fact that researchers are not very cooperative while open source development is often led by enthusiast that really use the results of what they develop. Still, it is a very effective way to improve and extend an artifact and it would be interesting to see what can be "reused" for paper evaluation and even improvement.
Web search gives an almost instantaneous way to identify significant documents. One wonders how much of this can be applied to evaluate posted versions of papers. Today's approaches use page rank to rate documents [Brin, 1998] and citation/impact factors to evaluate papers (research document). The problem here is how much of these can be leveraged to either "automatically" evaluate papers, or at least to assist reviewers or perform a preliminary screening.

Preliminary work on this topic is starting to appear. Chen et al [Chen, 2007] studied alternative metrics of paper quality and impact. They have applied a variant of the

PageRank algorithm [Rodriguez, 2006; Ball, 2006] to assess the relative importance of all publications in the Physical Review family of journals from 1893–2003. PageRank number and the number of citations for each publication are in fact positively correlated. Furthermore, outliers from this linear relation identify other exceptional papers or "gems" that are not easily found with traditional citation/impact factors. The reasoning behind this approach is that the situation in citation networks is not that dissimilar from that in WWW links: scientists commonly discover relevant publications by simply following chains of citation links from other papers. Thus it is reasonable to assume that the popularity or "citability" of papers may be well approximated by the random surfer model that underlies the PageRank algorithm.
One meaningful difference between the WWW and citation networks is that citation links cannot be updated after publication, while WWW hyperlinks keep evolving together with the webpage containing them. Another limitation of citations is that in the current publication models they cannot be used directly for evaluation in the extreme writing model as they assume that a paper is published, visible, and with an "identifier" (published in a journal/conference or at least as a technical report), because before the paper has a high citation count it has to be above the noise level among all documents, and because this is a slow process (you need for many referring papers to be released before you can assess the quality of a paper).

Pre-print repositories, such as e-Prints[2] and academic digital libraries and academic web search services, like CiteSeer.IST[3], Google Scholar[4] and Windows Academic Live[5], have also seen a significant increase in use over the past years across multiple research domains. Furthermore, emerging standard, like the DOI®[6] (Digital Object Identifier) are appearing and acquiring momentum to provide a system for persistent and actionable identification and interoperable exchange of managed information on digital networks. On this basis, researchers are beginning to develop applications capable of using these repositories to assist the scientific community above and beyond the pure dissemination of information. In [Rodriguez, 2006] a deconstructed publication model is presented in which the peer-review process is mediated by an Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) peer-review service. This peer-review service uses a social-network algorithm to determine potential reviewers for a submitted manuscript and for weighting the influence of each participating reviewer's evaluations.

In summary, it seems that the road towards an alternative review and publication model has received so far too little attention. There are spot studies on small numbers of cases and a few proposals that quickly lost appeal or that for reasons not entirely clear failed to stick. There is also evidence that the traditional process has flaws and that the famous "publish or perish" approach is a waste of time and money. Online communities have found many alternative ways to solve analogous problems, but these solutions have failed to reach the world of academia, or at least to be transformed in a way that could be applicable with success. With this paper we hope to raise awareness and stimulate researchers to join our currently ongoing search for a better approach to publication and review. We also hope to post soon, on this same forum, a contribution that presents the results of our efforts.

---

[2] http://www.eprints.org/
[3] http://citeseer.ist.psu.edu/
[4] http://scholar.google.com/
[5] http://academic.live.com/
[6] http://www.doi.org/

References

[Blog, 2005] Academia's Conflicted Reaction To Blogging. 2005 http://acrlblog.org/2005/11/28/academias-conflicted-reaction-to-blogging/

[Beck, 1999] K. Beck. *Extreme Programming Explained: Embrace Change*. Addison-Wesley Professional, Reading. 1999.

[Chen, 2007] P. Chen, H. Xie, S. Maslov, and S. Redner. *Finding scientific gems with Google. Journal of Informetrics,* to appear (2007); (http://physics.bu.edu/~redner/pubs/ps/google.ps).

[Brin, 1998] S. Brin and L. Page. The Anatomy of a Large-Scale Hypertextual Web Search Engine. Computer Networks and ISDN Systems, 30, 107 (1998).

[Ball, 2006] Philip Ball. *Prestige is factored into journal ratings*, Nature 439, 770–771 (2006).

[Fisher, 1994] Martin Fisher, MD; Stanford B. Friedman, MD; Barbara Strauss. *The Effects of Blinding on Acceptance of Research Papers by Peer Review. JAMA.* 1994. http://www.ama-assn.org/public/peer/7_13_94/pv3058x.htm

[Madden, 2006] Madden and DeWitt. *Impact of Double-Blind Reviewing on SIGMOD Publication Rates*. Sigmod Record, Sept 2006.

[Rodriguez, 2006] Marko A. Rodriguez, Johan Bollen, Herbert Van de Sompel. *The convergence of digital libraries and the peer-review process*. Journal of Information Science, Vol. 32, No. 2, 149-159 (2006)

[Rothwell, 2000] P. Rothwell and C. Martyn. *Reproducibility of peer review in clinical neuroscience*. Brain, Vol. 123, No. 9, 1964-1969, Sept. 2000 http://brain.oxfordjournals.org/cgi/content/abstract/123/9/1964

[Santini, 2005] Simone Santini. *We are sorry to inform you…* IEEE Computer, 38(12). Dec 2005.

[Tung, 2006] A. Tung. *Impact of Double Blind Reviewing on SIGMOD Publication: A More Detail Analysis*. 2006. http://www.comp.nus.edu.sg/~atung/

AUTHOR BIOS

Fabio Casati.
Professor of Computer Science at the University of Trento. Fabio got his PhD from Politecnico di Milano in 1999. After that, he joined HP Labs, Palo Alto (1999-2006) and the University of Trento (2006-). He has been working in the fields of of workflows, web services, data warehousing, and business process intelligence. He is member of the editorial board of ACM TWEB and in the steering committee of the ICSOC and BPM conferences. Fabio has acted as program chair, industrial chair, or in other officer positions in many conferences including ICDE,

ICSOC, BPM, ICWE, and CEC/EEE. More details at the URL: http://www.dit.unitn.it/~casati

Fausto Giunchiglia.
Professor of Computer Science at the University of Trento, ECCAI Fellow. He has done research in various related areas including knowledge management, reasoning with context and formal methods. He has been program or conference chair various events, including: IJCAI 2005, Context 2003, AOSE 2002, Coopis 2001, KR&R 2000. He has been editor or editorial board member of around ten journals, including: Journal of Autonomous Agents and Multi-agent Systems, Journal of applied non Classical Logics, Journal of Software Tools for Technology Transfer, Journal of Artificial Intelligence Research. He has been Member of the ECCAI Fellows Selection Committee, of the IJCAI Board of Trustees member (01-11), President of IJCAI (05-07), President of KR, Inc. (02-04), Advisory Board member of KR, Inc., Steering Committee member of the CONTEXT conference. More details at the URL: http://www.dit.unitn.it/~fausto

Maurizio Marchese.
Maurizio Marchese graduated with full honor in Physics in 1984 at the University of Trento, Italy. He has been Visiting Researcher at the National Research Council of Canada, Ottawa, Canada; Post-Doctoral Research Associate at the Material Research Laboratory, University of Urbana-Champaign, USA; Visiting Researcher at the Institute for Computer Applications, University of Stuttgart, Germany. He is currently Assistant professor at the Department of Information and Communication technologies at the University of Trento, Italy. Current research interests are: architectures for web services, distributed architectures for digital libraries, service integration in Geographical Information Systems (GIS) environments. He has published more than 60 papers in international journals and conferences. Dr. Marchese is a member of IEEE the Computer Society and ACM.