

This is Google's cache of <http://tagide.com/academia/the-story-of-a-journal-proposal/>. It is a snapshot of the page as it appeared on 21 Mar 2023 22:00:17 GMT. The [current page](#) could have changed in the meantime. [Learn more](#).

[Full version](#) [Text-only version](#) [View source](#)

Tip: To quickly find your search term on this page, press **Ctrl+F** or **⌘-F** (Mac) and use the find bar.

[Home](#)[Expertise](#)[Services](#)[Blog](#)[About](#)[Contact Us](#)

ACADEMIA CONFERENCES RESEARCH

November 9, 2013 / By / 20 Comments

The Story of a Journal Proposal

The Story of a Journal Proposal

This is the story of [Jan Vitek](#) and I approaching the [ACM Publications Board](#) with a proposal for publishing the OOPSLA papers in a journal. Even though we did this together, the analysis and opinions in this post are mine alone.

TL;DR: if there's any lesson to be learned from the ACM Publications Board's policies on scientific publication is this: pay attention to their policies, then do the exact opposite.

[WARNING: this post is irrelevant for all but academics]

Over the past year or so, I have been involved in making the [OOPSLA conference](#) become a venue for the publication of papers that are as good as they can be. This was achieved with the introduction of a [two-phase review process](#), where the set of papers selected in the first phase were given concrete editorial feedback for improvement, which was acted upon by the authors and was checked 2 months later by the reviewers. You know, journal like. I am very proud of the results this year: the papers presented people's work in their full potential. And this added no more time between submission and final acceptance! – 4 months.

Love or hate the work that has been published at OOPSLA this year, it is there in its best Sunday dress, as opposed to covered in rags with missing pieces of cloth. Only time will tell if the community values these papers by reading them and citing them

(or not). I'll keep an eye on the citation metrics! But that's not the purpose of this post.

In a [previous post](#), I talked about this conundrum of conferences vs. journals in Computer Science. For the longest time in Academic promotion cases, Computer Science department chairs have made the case that, unlike everyone else, Computer Science researchers publish their best work in conferences. It's a fact. Well, sort of. 10 years ago, it used to be a pervasive practice. Now, certain parts of Computer Science have finally given up the uphill battle, and have started to publish in journals [too]. They don't do it all in the same manner: machine learning people do truly publish in journals, in the sense that they send their original work to journals; but others (e.g. software engineering and programming languages people, for example) are unwilling to let go of their conferences, so they publish their original work at their best conferences and then patch that work up with 5 or 6 more pages of shinning cloth and bug fixes, and send that to a journal. While a lot of people do this last coat of veneer, [almost] everybody despises doing it. It's a waste of time and it dilutes citations. If we are going to publish our work, why not publish it in a form that's as best as it can be in the first place? I tell you why: internally within the community, what matters are the conference publications that describe original contributions and are selected by peers; externally, the journal articles give the necessary veneer for promotions. That's why people do it this way.

It's a perverse practice. It wastes people time. It's a sham. It should stop.

The truth of the matter is that in many parts of Computer Science, the publications selected for conference presentation are, indeed, the original research work. So why not just acknowledge that by (1) making it the best that it can be right there and then, and (2) tagging these collections of papers with the word "journal"? – because that's exactly what they are, especially when going to 2 phases of review, like OOPSLA and CSCW.

While giving this extra kick of quality to the OOPSLA papers, the OOPSLA Steering Committee gave some of us the blessing to go ahead and try to publish the OOPSLA papers directly in a journal. And hence the real story of this post begins...

From the beginning, I suspected this was not going to end well. I had read the information regarding the [creation of new journals within the ACM](#). Early on, I also came across very strong wording put out by the Pubs Board regarding the [publication of conference proceedings in ACM journals](#). This policy basically says: **no way in hell we are going to trust conference program committees to select the right papers for our journals** (*SIGGRAPH not included). OK!, I thought. We'll give it a try anyway, and hope to learn something along the way. And learn we did.

We wrote a [short proposal](#) as a means to start the conversation. Read it, if you have time! That proposal is so bluntly at odds with the Pubs Board's policies that they must have thought we came from Mars! But we didn't. That was an honest-to-good proposal for correcting the double standards that academics in the ACM community regularly have to resort to: internally valuing the peer-selected publications presented at conferences while externally presenting journal-versions of those articles that add very little or no value to the original conference publication; or getting away with publishing a semi-decent version of their work in a conference and moving on to the next project without bothering to apply veneer.

The conversation with the Pubs Board dragged for several months. First, we were encouraged to talk to the [ACM TACO](#) editors, which we did. The ACM Pubs Board

likes that model very much. The model is a bit contrived, as it is trying to achieve two conflicting goals at the same time: the speed of conference reviewing and the cover story of ACM journals. I won't try to explain it here. But the bottom line is this: an existing, conference (HiPEAC) was devoid of editorial decisions. Instead, that responsibility was moved to a journal, TACO, which started in 2004 and was having a hard time getting its feet off the ground (meaning: no submissions). TACO, like any ACM journal, has a fixed, 3-year term, [editorial board](#), headed by the 2 founders of the journal, who, as far as I can tell, are still the heart and soul of the resulting publication 10 years down the line (they are the ones we talked to, not the Editor-in-Chief).

So, here's my interpretation of the TACO/HiPEAC hybrid: an existing, perfectly healthy conference was devoid of all editorial decisions – basically the Program Chairs have no role in selecting papers and supervising the review process anymore, and there is no Program Committee. Those tasks were moved to a non-starter ACM journal with a fixed editorial board – essentially, a 3-year-term Program Chair called Editor-in-Chief (EiC), a 3-year-term Program Committee called Associate Editors (ACs), and the critically indefinite-term 2 Senior Editors, the 2 founders of the journal, who keep the ball rolling.

Cynical as I may sound, I believe the community around the HiPEAC conference likes this model! I was just left wondering what will happen to TACO when the original founders don't want to do the pushing anymore... That is, indeed, the real risk of letting outstanding individuals do exceptional things, instead of coming up with a model that scales beyond exceptions!

Our proposal to create a new journal for the OOPSLA papers was *almost* identical to the TACO model, except for a few details... and a whole different governance attitude: (1) editors would rotate on a yearly basis; (2) editors would be selected by the conferences' Steering Committees, which would also be part of the editorial board; (3) the journal would not accept ad-hoc submissions, only submissions intended to be presented at the conferences; (4) there would be no dragging of papers for the next year: they would either be accepted in 4 months, or they would be rejected.

These details, however, make the whole difference between accepting the top-down governance imposed by the Pubs Board, where the editorial board is under the control of a very small group of people who don't necessarily know anything about the topic, and maintaining the bottom-up governance that we have in conferences, where members of the research community who stand out over the years are asked to serve increasingly more important roles within the organization of the conferences and the selection of papers. This model has proven to work very well for the past 40 years – and it works without exceptional individuals!

The conversation with Pubs Board also took us to [TOPLAS](#), but that was a dead-end on arrival, since the current EiC's vision for TOPLAS is that it publishes "[the best work of the year in programming languages](#)", resulting in... less than 20 papers per year (!). OOPSLA alone publishes 50+ papers per year! Unlike TACO, TOPLAS has been around for a long time, so it has its own history and goals; that was also a strong factor in making a possible merge a non-starter.

OK, so our proposal was eventually shot down, as I suspected it would. But here's what I found along the way.

The governance surrounding scientific publications is really important, and the top-

down governance model used by the Pubs Board and so many other publishers is not only outdated, but it may very well explain the dismal performance of ACM journals in citation-based quality metrics. At least the commercial publishers have a strong incentive to make their journals perform well on those metrics, so chances are they choose editorial boards that can deliver. Not-for-profit publishers like the ACM have no such incentives, so if a journal is poorly cited over many years, no one really cares, and the whole issue is swept under the rug of the tagline/wishful thinking/self delusion "highest quality research in X".

Here is some interesting data related to citation-based metrics of ACM publications.

From Thomson Web of Science's Journal Citation Reports (Impact Factor, JCR 2012):

Highest IF	50+
Median for all ACM publications	1.0
Highest of any ACM publication	3.5 (Computing Surveys)
ACM TOG, which largely publishes SIGGRAPH conference papers	3.36
CACM, a lowly "magazine"	2.7
JACM, the flagship publication	2.3
TOPLAS, "the best" in PLs, < 20 papers/year	1.03
SIGPLAN Notices, a lowly SIG "newsletter" that publishes all SIGPLAN conference proceedings, from POPL to PPOPP (over 300 papers/year)	0.7
TACO/HIPEAC	0.68

The high IF of TOG should send a strong signal to the Pubs Board about the effects of publishing the best conferences in ACM journals using the conferences' governance instead of the top down governance that it favors.

CACM vs. JACM is an eye opener. Finally, after soooo many years!, the ACM [acknowledged that CACM, not JACM, is, indeed, its "flagship" publication.](#)

TOPLAS (an elite publication) vs. SIGPLAN Notices (a huge patchwork of PL papers) is another interesting comparison: best as they are supposed to be, TOPLAS papers aren't that much more cited than the entire collection of 300+ papers carried by SIGPLAN Notices.

Not many people like Thomson's IF. Indeed, the representativeness of Computer Science material on that index is questionable, because, apart from SIGPLAN Notices, JCR doesn't contain any conference papers. So, let's look at other bibliometrics.

Google Scholar has [this rank](#) for CS/Software publications. Here is a snapshot of what it looks like today:

Top publications - Software Systems [Learn more](#)

Publication	h5-index	h5-median
1. International Conference on Software Engineering	58	99
2. IEEE Transactions on Software Engineering	48	72
3. SIGPLAN Conference on Programming Language Design and Implementation (PLDI)	47	72
4. ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)	44	66
5. Journal of Systems and Software	42	60
6. Information and Software Technology	41	61
7. ACM SIGPLAN Symposium on Principles & Practice of Parallel Programming (PPOPP)	37	61
8. IEEE Software	37	51
9. SIGSOFT International Symposium on the Foundations of Software Engineering	34	52
10. Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)	33	48
11. European Conference on Object-oriented Programming (ECOOP)	32	45
12. International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)	32	41
13. International Symposium on Software Testing and Analysis	31	56
14. Empirical Software Engineering	30	44
15. International Software Product Line Conference	30	43
16. IEEE/ACM International Conference on Automated Software Engineering (ASE)	30	40
17. International Conference on Functional Programming (ICFP)	29	41
18. IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)	28	41
19. Model Driven Engineering Languages and Systems	28	39
20. Science of Computer Programming	27	48

Dates and citation counts are estimated and are determined automatically by a computer program.

This ranking has it even worse for ACM journals: they are nowhere to be seen. The top 20 publications are dominated by ACM and IEEE conferences, along with some non-ACM journals.

Here is Microsoft Academic's ranking of all CS journals. Snapshot below:

Top journals in computer science

1–100 of 1,361 results

Journals	Publications	Field Rating
CACM - Communications of The ACM	13104	249
PAMI - IEEE Transactions on Pattern Analysis and Machine Intelligence	5389	246
TIT - IEEE Transactions on Information Theory	14743	241
PIEEE - Proceedings of The IEEE	17074	210
TCOM - IEEE Transactions on Communications	12787	190
AI - Artificial Intelligence	3411	189
JSAC - IEEE Journal on Selected Areas in Communications	5898	183
TSE - IEEE Transactions on Software Engineering	3925	183
JACM - Journal of The ACM	3035	178
SIGPLAN - Sigplan Notices	9704	176
BIOINFORMATICS - Bioinformatics/computer Applications in The Biosciences	7371	174
CCR - Computer Communication Review	2867	167
COMPUTER - IEEE Computer	8912	163
TC - IEEE Transactions on Computers	8633	162
Sigmod Record	3889	162
IEEE Transactions on Image Processing	5579	154
IJCV - International Journal of Computer Vision	2145	154
TSP - IEEE Transactions on Signal Processing	12427	150
CSUR - ACM Computing Surveys	1513	149
SIAMCOMP - Siam Journal on Computing	3326	147
ACM Siggraph Computer Graphics	2552	141

There's CACM

topping the list. SIGPLAN Notices (listed as a journal) comes in #10 (!), almost tied with JACM. No other ACM journal makes the top 20. ACM Computing Surveys comes at #19 [Thanks for pointing that out, Jonathan!] No other ACM journal makes the top 20.

Finally, let's zoom in on [Programming Languages](#) in Microsoft Academic:

Top journals in programming languages

1-20 of 20 results

All Years

Journals	Publications	Field Rating
SIGPLAN - Sigplan Notices	9704	176
TOPLAS - ACM Transactions on Programming Languages and Systems	1135	110
Mathematical Programming	4089	109
SCP - Science of Computer Programming	1562	67
JLP - The Journal of Logic and Algebraic Programming	696	64
JFP - Journal of Functional Programming	882	54
IJPP - International Journal of Parallel Programming	1055	42
JOOP - Journal of Object-oriented Programming	649	35
ACM Sigplan Lisp Pointers	296	25
CL - Computer Languages, Systems & Structures	464	24
TPLP - Theory and Practice of Logic Programming	333	24
LOPLAS - ACM Letters on Programming Languages and Systems	49	22
JALC - Journal of Automata, Languages and Combinatorics	304	19
JPL - Journal of Programming Languages	60	18
ACM Sigapl Apl Quote Quad	2077	17
JFLP - Journal of Functional and Logic Programming	78	15
OOPSM - Oops Messenger	377	14
Object Oriented Systems	34	12
ACM Sigplan Fortran Forum	287	9
ACM Lisp Bulletin	7	6

SIGPLAN Notices comes on top.

In others words, the data from both Google and Microsoft Academic is showing, in different ways, that the large patchwork of papers that are presented in ACM software conferences has higher citation-based rankings than the carefully selected collections published in the ACM journals focusing on software.

And here is where the Guardians of Quality come in and say: "But these rankings are meaningless! That is a popularity contest! Everyone knows that IF and the like is a horrible way to measure research quality! The ACM journals are carefully refereed and the conferences aren't!" Perhaps, but many people disagree on that refereeing statement (and clearly SIGGRAPH's papers are refereed as poorly or as well as any other top conference's papers). In my book, after data like this, that kind of defense sounds like elitism for the sake of elitism, especially given that when impact factors come out, [the ACM checks how it's doing](#). In my book, data like this gives a strong hint about how ACM journals are simply doing it wrong – and how the Pubs Board is missing its opportunity to improve the situation by taking a more community-oriented approach to running its publications: taking the best conferences, **with their governance unchanged**, their hits and misses, and publishing their papers as journals, which is exactly what we proposed to do for OOPSLA. Perhaps, just perhaps, if the ACM does that, we will start seeing ACM journals routinely getting IFs above 3, like TOG.

Community-oriented, bottom-up, vote-with-your-citations-and-your-downloads scientific publications [is in the spirit of how other fields are now starting to do it](#). We've been doing it for a while... in conferences.

Share this post with your friends

 Facebook

 Twitter

 LinkedIn

[Prev Post](#)[Next Post](#)

The Future of Conferences | Immersive VR is here and it's awesome

Comments



James Noble, 10 Nov 2013

This looks like a pretty good journal to me:

<http://www.informatik.uni-trier.de/~ley/db/journals/sigplan/sigplan47.html>

and has the interesting consequence that it needs no further effort from us. The more I look at reviewing practices in other disciplines' "prestigious" journals – the more I'm happy to stand behind CS & SE's practices in places like ECOOP, OOPSLA & ICSE.



crista, 10 Nov 2013

yep!



Eric Eide, 10 Nov 2013

Thank you for an interesting article!

I've been thinking about CS publishing models for the past little while, because I was part of the [TRIOS](#) experiment at SOSP '13. My opinions are in flux, but definitely it seems that things need to change.



Jonathan Grudin, 10 Nov 2013

Nice essay and analysis, thanks. Very minor correction, the Microsoft Academic top 20 did also include Computing Surveys.

For those interested, related essays that appeared in the January 2013 issues of CACM and Interactions can be accessed (free) from my web site.



Shriram Krishnamurthi, 10 Nov 2013

Crista – I respect the effort Jan and you have put in over the past year, and I am full of nothing but sympathy for people who have to do battle with the ACM.

However, I have to say, I don't really agree with your thrust. The solution to this problem (or this very large set of interlinked problems) is definitively not, imo, to conflate conferences and journals, even with an extra round of reviewing.

I have written several journal papers, and I actually really enjoy writing them. But I write them 3-4 years after the work, not in six months. And by the time I get around to writing them, they look almost nothing like the original conference papers. In fact, every time I've tried to lazily reuse the original paper text, I've found the process much more painful than if I'd just started from scratch.

Essentially, my recipe is, do the work, present it a few times, write and publish the conference paper, present it many more times, figure out what it is we actually did, and when I have both the presentation and figuring out done, then I write the journal paper. Your model would make this virtually impossible.

I am perfectly fine with Jan and you saying, "Lookit, Shriram, you're such a tiny minority, and even then you produce so few journal papers, that you're irrelevant to the grand scheme of things". I would even agree with you. But I do object to thinking that a journal paper is really just a glorified conference paper, "almost there" but for one more round of reviewing.

This process sends exactly the wrong message, and as tilter against windmills, I say no, that's what computer science has perverted the journal paper into being, and we should stand athwart yelling stop.

Separately, I recognize that part of the problem here is the administrators of research who mandate citation indices and the like, and that this causes problems for people in certain countries. I feel their pain...but the solution is surely not to pervert the whole system because of some wacko administrators, because where does that end? One does not negotiate with terrorists. Yes, it takes time to

educate these administrators, but I'd say it's time well-spent, because I believe it will result in better research overall (because this won't be the only issue on which they need education).

So, again: my sympathies, and I hate that you lost because you ran up against the immovable force of the ACM, but I'm not sorry about the outcome.



crista, 10 Nov 2013

Shriram, the reply to your comment requires another blog post 😊 Some other time!



Sophia Drossopoulou, 11 Nov 2013

Hi Shriram,

I think that the model of "write a journal paper 3-4 years after the work", is certainly appealing. But does it reflect what happens in other science disciplines?

We can still have the mature, 3-4 years after the work papers, in some other kind of journals, esp. created for this purpose.



Andrew Myers, 10 Nov 2013

Your basic point may be right, but the way that Microsoft Academic Search and Google Scholar rank venues appears to be misleading, biasing in favor of large venues with lots of papers. That is, if you publish enough papers, you will get a lot of citations even if the papers are mostly poor quality. I scraped the MAS data and got perhaps a more meaningful ordering on CS publication venues:

<http://www.cs.cornell.edu/andru/csconf.html>. After Computing Surveys, the top journal I recognize is TOCS, with an average of 33 citations per article. It is only #14 on the list.



Jonathan Aldrich, 10 Nov 2013

Hi Andrew,

Your list is an interesting point of comparison, but in my opinion the metric you chose is quite problematic, and it's not one I'd like to see widely used. A "citations per article" metric creates a perverse incentive: to keep the proceedings of a conference or journal as small as possible. I'd rather see a metric that considers both citations per article and number of articles. For example, the HCI community made a conscious decision that the CHI conference would accept as many quality papers as were submitted, rather than fragmenting into many conferences (as SIGPLAN has arguably done). I admire this decision, but it makes them (incorrectly, in my view) look bad on your list. Of course, no metric will ever be perfect, and many are worse than yours, but fundamentally we should be judging venues on their total impact, not on the impact per paper.

As a companion point, in reviewing we should reject papers because they are not to the standard where they can make an impact on the field, not just because there is some other paper that we think is better.



Sean McDirmid, 11 Nov 2013

Hi Jonathan,

I find that the CHI mega-conference format is very bad: it encourages LPUs (least publishable units), little follow through on most ideas, and a loosely filtered proceedings (lots of low quality 4 page notes). Contrast to OSDI/SOSP 30 papers/year and see how the attention deficit disorders are correspondingly reduced.



David Evans, 10 Nov 2013

I appreciate your efforts on this, but it seems to me that the the real solution to this conundrum - "That was an honest-

to-good proposal for correcting the double standards that academics in the ACM community regularly have to resort to: internally valuing the peer-selected publications presented at conferences while externally presenting journal-versions of those articles that add very little or no value to the original conference publication; or getting away with publishing a semi-decent version of their work in a conference and moving on to the next project without bothering to apply veneer.” – is for our community to make it the case that publishing worthless journal papers is viewed as a strongly negative thing to have on one’s tenure case, at least negative enough to counteract any positive value they are perceived to have at higher levels in that process. Now, untenured faculty who feel pressured to publish in journals to fill up their tenure CV only have to weigh the tradeoff as one between wasting a bit of their own time, a lot of effort from reviewers and editors, and a little bit of time for everyone else who might read the paper, against the selfish career benefits it provides. If enough potential tenure-letter writers adopt the view that a candidate wasting their own and their communities time by publishing worthless journal papers should be viewed negatively, maybe this will change the tradeoffs enough to dis-incentivize the practice.



Laurence Tratt, 10 Nov 2013

This is a very timely post, and touches on a lot of issues that are causing us collective grief.

From my selfish perspective, I think the 2 stage OOPSLA process this year worked very well. I think most conference papers have needed a little more “comprehensibility” polishing than they have traditionally received. The two stage process certainly forced the paper I was involved with to do a better job in that regard and I hope more conferences adopt a similar process in the future.

That said, I think Shriram indirectly makes an interesting point. Our local circumstances can effect what we perceive to be problems in the current publishing process. For example, I’ve increasingly noticed the different effects the hiring / tenure / promotion has on my US-based colleagues than those in the UK or Europe. To give one simple example: the UK no longer has formal tenure; being employed as even a junior lecturer is akin to being tenured. There is, therefore, a very low risk of being shunted out of the system after a few years (I’ve never heard of it happening). Perhaps because of that, I feel less direct

pressure to churn out journal publications than some of my US-based colleagues and was able to take a somewhat relaxed approach to publishing in my first few years as an academic. That said, in the UK we're measured every 5 years on "research quality", via our 4 best publications. It is often thought that journal publications are the best signifier of such quality. In the future, it seems likely that citation counts of these publications will also play an important part, quite possibly overriding the perceived quality of the venue itself. My incentive is thus edging closer to the process Shriram outlined: synthesising, over time, several papers into a single journal publication (and then hoping that the citation counts concentrate on it).

Thus different countries are gradually providing different incentives to researchers. I have to produce a small number of somewhat respected / well cited publications; some countries want to see a large number of publications in highly respected venues. Doubtless there are other schemes of which I am ignorant. I think it is possible that we might unintentionally optimise for one scheme over another. Of course, I'm doing the easy thing which is pointing out some problems, without providing any hint of a solution, as I have none. I'm glad to see that people such as Crista and Jan are trying new approaches, even if it turned out to be a lot of work for not much gain!

[My internal lawyer asked me to add this: my detailing the incentives of a scheme does not mean that I endorse it.]



Jonathan Aldrich, 10 Nov 2013

Crista mentioned TOG, so I looked it up to see the details- it's a more illustrative case than I could have imagined.

From the Journal home page <http://tog.acm.org/>:

"As illustrated below, the TOG journal has a strong synergy with ACM SIGGRAPH, the premiere conference organization in graphics. Of the six issues published by TOG each year, two are special issues containing the papers presented at the annual SIGGRAPH and SIGGRAPH Asia conferences. Conversely, authors of papers published in the regular issues of TOG can present their work at either of these two conferences. Also, several paths provide reviewer continuity between the conference and the journal."

This is why TOG blows away most every other CS journal in various metrics (of course the metrics are far from perfect, but they are not completely meaningless either).

What TOG is doing sounds like a great solution. I can hardly see any downside to it—even Sriram’s objection, which is reasonable, is less of an issue because authors who wish to do so can still use the usual journal process (and even get to present at SIGGRAPH). Why don’t we do this in SIGPLAN/SIGSOFT?



crista, 10 Nov 2013

Jonathan, because the Pubs Board, as is, doesn’t let us do it. SIGGRAPH was the first SIG to realize the bizarre situation coming from artificially separating conferences (community) from journals. They were smart enough to fix it, although I heard that some people don’t like it that there are 2 publications instead of just one (dilution of citations). I don’t know the back story, but I know that the Pubs Board suddenly realized that it was going to lose control over “quality” of the papers, and wrote that policy I refer to. By the time they realized it, it was too late to prevent SIGGRAPH from doing the right thing. So the SIGGRAPH/TOG situation comes as an exception on that policy.

The philosophy underlying that policy is so broken and ill-informed, it’s sort of hard to even know where to start. When I asked the Pubs Board members what were the reasons for that policy, they said that it comes from the “highest ranks of the ACM”. This answer left me wondering whether it really comes from higher up or whether the person we were talking to simply didn’t want to discuss it with us.



Ryan schmidt, 11 Nov 2013

You may want to talk to some graphics researchers before taking SIGGRAPH/TOG as your model. Yes it is great for tenure and grants, but many I have discussed with (including some involved in the decision) consider it a huge mistake.

Basically, it raises the stakes. People naturally become more conservative (ie meaner reviewers, followed by safer submissions). It has killed all our small conferences, everybody instead has a backlog of rejected siggraph

papers that they are hoping might make it next year.

Journals are for archival work. Conference papers are for active research. ML and Vision have it right.



crista, 11 Nov 2013

Only in CS is the word “journal” used to denote “non-active research”. In all other fields, the word “journal” is used to denote a large spectrum of publications, from active research to retrospectives to surveys... Our colleagues list all of their publications under “Journal”, and only their 2-page abstracts (if any) under “Conferences”. We are *all* doing it wrong.



James Noble, 12 Nov 2013

I don’t know what it means to be “doing it wrong”.

Publication practices aren’t handed down on tablets of stone. We may be doing things differently to the majority of the physical sciences – say – but they do things differently to the arts, who themselves differ from law, or literature, or architecture, or engineering. Classical physics came of age with the sailing ship. Computer Science came of age with the 747. That explains a lot.

Does this matter? I don’t see how it does within the discipline, so long as we’re all aware of how publication actually rate. For example, my thesis publication is well buried in a journal special issue no-one has ever read. The papers rejected from that issue went into a book that had rather more impact (still close to zero thought).

Comparing across disciplines, well I always think that’s political. Deans or other authority figures who want to support computer science & software engineering will take account of our publication policies. Those who don’t want to support us will always find

reasons, whatever our publication venues are called. And if we change things wholesale, someone who had always rejected conference publications will most likely come around and say “well you guys always said your journals were junk and the conferences are what counted. Now you’re saying its the other way round – why should I believe anything except that all CS is junk?” (presumably what they thought all along)

The good news for us in SIGPLAN and PL work is that we’ve got the best of both worlds. With SIGPLAN Notices, the problem really is moot.



crista, 12 Nov 2013

Yes, James, this is all about politics indeed. And I would like to start pointing the spotlight away from our own belly button and out to the rest of the Academic world. Our internal politics are making us all loose.

We are doing it wrong in the sense that all of our colleagues use the words “journal” and “conference” to mean things completely different from what we mean when we use them. To the rest of the word, a “journal” is any publication that includes peer reviewed articles; a “conference” is a meeting where you present by sending a 2 or 3 page abstract. For a “journal” in the broad sense, the process of producing peer reviewed articles – deadlines, no deadlines, 1 cycle of revisions, many cycles, etc. – doesn’t matter much; all that matters is that they are peer reviewed and published. PLOS One recently went all the way to establish a publishing platform where ALL papers that meet only a minimum criteria of quality are published, without editorial judgments. That’s one extreme; but the spectrum is very wide.

We, on the other hand, define “journal” as

“a place where inactive research gets buried.” Really? Who in their right mind likes to spend most of their time publishing dead work?

Other fields’ “journals” are alive and kicking, and our colleagues proudly list their articles there under “Journal Publications” on their CVs. As such, when a pile of CS CVs with dead work under “Journals” goes against a pile of CVs with active research under “Journals”, CS loses. Always. We see this at University level and we are starting to see it at national levels.

This is the real issue underlying this story of the journal proposal.



James Noble, 13 Nov 2013

Yes, I figured that the problem is political – but it’s good to have that out in the open.

In the short term, at least 2/3rds of what I do is covered. I have in the past (and may do so again) list OOPSLA & co as Journal Articles in SIGPLAN Notices, and ECOOP etc as Book Chapters in LNCS. The IEEE is I guess more interesting here, as they don’t have a nominal journal for their conference proceedings.

In the long term, people like you & Jan have to be willing to play the politics and get into positions of power (where you can dismiss other disciplines for not publishing in conferences :-). But I’m a little bitter, tomorrow I’m off to a meeting tomorrow on the ‘National Science Challenges’ where the briefing papers assume that NZ research in CS, SE, or PL simply does not exist. I doubt having published in more titular

journals would have changed
that one jot.



Jan Vitek, 17 Nov 2013

I like to separate the practices from the bean counting.

* Practices: Shriram's point is that scientific publications come in different shapes and kind. There are timely papers that report on fresh results and there are more reflective write-ups that summarize years of research. We agree there. Both kinds of paper have a role in science: the former is a way to get rapid feedback about our research ideas, the latter is intended to be a definitive account of a significant research effort. Nothing that we are proposing prevents that practice. TOPLAS will always be there for long, best-of-breed, papers, or you can go with Computing Surveys if the paper is a broader study.

The practices that I would like to discourage is the bean-counting-driven addition of 20% filler that sometime passes for a "journal" paper. It is waste all around: author's time, reviewers' time, paper, energy. (I think that we do publish too much and too early, but that's not a fight I am picking, 'cause I have no idea how to change that behavior)

When I came to Purdue, I was told that journal publications are important for tenure. So, over the years I have committed 25 journal papers. Out of those, there are two for which we did additional work that was worth it. Three papers were summaries of large projects. Two were papers that could not have been published in a conference. But, to be honest I don't think our field would have missed any of them. Because the ideas were already out and the delta added by the journal versions was more about details than about the core ideas. Crisper explanations, additional experiments, proofs... but nothing really crucial.

Traditional CS Journals are occasionally handy, and it is good to have them as an outlet. Did I mention that I don't read journals anymore. I used to as a student. Is it just me?

* Bean counting: We are shooting ourselves in the foot by trying to fight the accepted labeling system in all the other sciences were are in competition with for resources. For them: Journals are where science happens, Conference are where you go to socialize. We can claim that we have "trained" our administrations to 'understand' that Computer Science is different, but that's bull. Administrators let us do

