How to Write Your First Paper

Steven G. Krantz

This article is the third in an occasional series intended for graduate students. The series is coordinated by Associate Editor Lisa Traynor.

In today's world, most any math department wants each of its faculty to have a scholarly profile. If you wish to establish yourself in the profession, if you want to make your reputation, if you want to achieve tenure status, then you must publish. While it is okay to publish a "Letter to the Editor" in the *Notices*, or a recreational problem in the *American Mathematical Monthly*, the hard fact of the matter is that the publishing that really counts is that of a research article in a peer-reviewed mathematics journal. The purpose of this article is to tell you how to perform that task.

I have published more than 150 articles myself. So I guess that I know how to do it. I have never written an article and then been unable to publish it. The notion—that one reads about in stories or sees in popular movies—of a forlorn scholar languishing away because he/she cannot get his/her ideas in print is mostly nonsense. You cannot succeed at anything in life unless you understand what it is that you are trying to achieve. Once you understand what mathematical research is about, and how the publication process works, then you should be able to get your work into print. We shall lay out all the essential moves here.

Blood and Guts

First, you need to become involved in an ongoing research area of current interest. If you are lucky, you will have had a good and effective thesis advisor who will have given you a problem that is *not* a dead end. Then whatever you achieved in your thesis will have opened new doors, and suggested new questions, and you will certainly have interesting and productive things to think about. If, sadly, this is not the case, then you will have to do the job yourself. Go to one or more conferences, listen carefully to the best talks, and find out what people are thinking about. Pick two or three good papers and work through them in detail. *Talk to people.* Go to seminars in your own department. Get involved in some Internet chat groups. Immerse yourself in a field. In the best of all possible worlds, this should be a field that fascinates you, that gives you the proverbial "fire in the guts". Eventually you will find a problem that you cannot let go, that you must solve or else.

So solve it. Make sure it is right. Give a seminar on your result. Discuss it with some friends. When you are confident that you have a winner, then it is time to write it up. Again, it is essential that you understand what it is that you are trying to achieve.

You cannot write a good math paper by just picking up a pencil and starting to write. Some planning is definitely in order. Will this be a 10-page paper that simply states a theorem and proves it, or will this be a 50-page *magnum opus* that redefines a field and sets it in a new direction? If you are a beginner in the field, and if you are an ordinary mortal like most of us, then most likely your first paper will be of the first type. But we shall give here some advice for both types of work.

Begin by writing an outline of the paper. This could be as simple as

- Introduction
- Background
- Thanks

Steven G. Krantz is deputy director of the American Institute of Mathematics. His email address is skrantz@aimath.org.

- Definitions
- Statement of Main Results
- Indication of Methodology
- Details of Proof
- Concluding Remarks

At least now you will have an idea of what are the main ingredients of this new work. You can probably make an estimate of the length of the paper. And you can begin to write.

Some of us write directly on the computer (in TEX), without working from a paper draft written by hand. If you are doing serious, deep mathematics then you will certainly have to do some of your calculations by hand. You may also have to draft some of your theorems and your arguments by hand. That is just the way the human intellect works. You can hardly perform a delicate estimation of singular integrals on the fly at the keyboard. Some people will write out every word by hand before going to the computer. Others will combine the two media.

Think about what attributes make a paper readable. I have certainly seen papers which begin

Notation is as in my last paper.

Theorem: Let $\epsilon > 0$. …

This is okay if all you want to do is plant your flag. Back in the 1960s, there were many journals that Dutch theoretical computer scientist Edsger Dijkstra would have called "write only", and they would have published something like this. Today journals are more demanding, and in any event you should set a higher standard for yourself. Write a paper that you yourself would want to read. Make it accessible. Bear in mind that the referee for your paper will be a busy person who has no patience for a tract that he cannot fathom. Lay out the material so that it is rapidly apparent what your main result is, what the background for that result is, and how you are going to go about proving it. If the proof is long and complicated, then break it up into digestible pieces. Tell the reader what is going to happen *before* it happens. Tell the reader what has just happened before you go on to the next step. At the end of a long argument, summarize it.

Write a nice conclusion for your paper. A mathematical article that ends

and so $S \subseteq T$. □

has a certain *joie de vivre* to it, but leaves the reader hanging. Why not have a nice section of Concluding Remarks, telling the reader what you have accomplished and where things might go from here? Leave the reader with a forward-looking view of things; make him/her feel as though this is a field that he/she might want to get involved in.

I promised to say something about longer papers, and I shall do so now. It is difficult to publish a long paper. If you write a 50-page paper, even if it is extremely good, you are going to have trouble getting it into print. Many journals have strict page limits, and the limit is usually about 15 or 20 pages. Many journals make it clear that, if they are going to publish your 50-page polemic, then it had better glow in the dark.[1] As a general strategy, it is best to break your ideas up into smaller pieces. Publish three 20-page papers rather than one 50-page paper. If you are into self-abuse and seek the defeatist situation of having a paper that you have labored over for two years and cannot get published, then writing a 50+ page tract is the way to go. You will get angry and frustrated and, in the end, be done in if you write a paper that nobody is willing to even consider.

But, if you have proved the Goldbach Conjecture, or found a zero of the zeta function that does not lie on the critical line, then you can probably justify writing a long tract. In this case, organization is particularly important. It is extremely useful for such a paper to have a thoughtful and detailed *Table of Contents*. You should be careful to isolate all your notation and definitions. Give an informal statement of your results before you give the detailed formulation. Give an outline of your proof before you trot out all the dirty details. Formulate a thoughtful and enlightening closing section. The book [KRA1] gives copious advice in these matters.

After your paper is completed, checked, and ready-to-go, a very natural thing to do is to post it on an electronic preprint server. Many specialty areas—such as K-theory and linear algebraic groups—have their own dedicated preprint servers. Also a number of mathematics institutes (such as the American Institute of Mathematics) and most math departments have their own preprint servers. If nothing else you will probably want to put the paper on your own webpage. But, as of this writing, the canonical place to put a new paper is on `arXiv`. Created by Alan Ginsparg and now based at Cornell University, `arXiv` is the standard repository for new papers in mathematics, physics, statistics, computer science, biology, and other disciplines. Posting a paper there is straightforward (and particularly easy if you use Greg Kuperberg's front end called `Front`).[2] And

[1] *Theodore Streleski garnered some notoriety in the early 1970s for murdering his graduate advisor Karel de Leeuw at Stanford University. He had always been rather unstable, but the straw that broke the camel's back was that the* Bulletin of the American Mathematical Society *had rejected his Ph.D. thesis. As you may know, the* Bulletin *specializes in research announcements, research expository articles, and book reviews. It simply does not publish entire theses. It seems that Streleski was the victim of bad advice or bad judgment or both.*

[2] *The preprint server* `arXiv` *can be accessed at* `http://arxiv.org` *and* `Front` *can be accessed at* `http://front.math.ucdavis.edu`.

then your paper will be freely accessible to all the world. Many journals allow you to submit a paper simply by pointing to the Web address of your paper on arXiv!

There are copyright issues to consider here. The moment you write something it is copyrighted to you. And you certainly then have the right to put it on a preprint server. But when your paper is accepted by a journal then you will probably be asked to sign a Transfer of Copyright Agreement. Then the paper is copyrighted to the journal. In principle the journal could ask you to take the paper *off* the preprint server. These days most journals have made peace with how the world works and they will not ask you to do so. You can leave your paper on arXiv and go ahead and have it published in a journal. An alternative approach is to decline to sign the Transfer of Copyright agreement and tell the journal that you wish to hold the copyright. Many journals will go along with that request (although many will not!). It may actually happen that a journal will ask you to take your paper off the Web, but it has never happened to me.

Practical Matters

My detailed thoughts about the chapter and verse of writing a paper are already recorded in [KRA1]. I shall not repeat them here. The main point of the present tract is to discuss how to submit your new paper and how to deal with the journal and its editors.

The choice of journal to which to submit your work is not a trivial matter. It is well known that the best journals are the *Annals of Mathematics*, *Acta Mathematica*, *Inventiones mathematicae*, *The Journal of the American Mathematical Society*, and a few others. It is quite an accomplishment to get a paper into one of these journals. But sending all your work into these recondite forums is not the way to go. If you have shown your work to colleagues, given some talks on it, and received copious praise and adulation, then perhaps it is appropriate to consider sending the paper to a top journal. Usually it is not, and you should set your sights a bit lower. As a beginning mathematician, you should be spending a good deal of time browsing journals, acquainting yourself with the literature, understanding what is published where. You should get a sense of where papers in your subject area are published. Some observations are obvious. *The Journal of Symbolic Logic* will not publish papers on pseudodifferential operators. *The Journal of Differential Geometry* will not publish papers on Moufang loops. With *The Transactions of the American Mathematical Society*, matters are less clear. Most journals have an *Instructions to Authors* page that will tell you this journal's conception of itself, and what types of papers it seeks. It will also acquaint you with the specific mechanism for actually submitting a paper to that particular journal.

The main point is that you are trying to establish yourself in the profession, your tenure clock is running, and you cannot afford to fritter away five years getting your first paper published. You want to handle the matter expeditiously so that you can move on to the next project. Therefore choose a journal that is (a) well-suited to the subject matter of your paper and (b) at the right level. It helps if you know one or more of the editors. That will make you feel more comfortable with the process, and also will perhaps suggest that this is a periodical that will appreciate your work.

Another consideration if your tenure clock is ticking away is how long it will take any given journal to get your paper into print. Some deans are extremely punctilious and only believe that a paper exists when they hold the reprint in their hands. So you do not simply want to have your paper accepted, you want it to be in print. There *is* information available about journal backlogs. The *Notices* regularly collects and publishes such data. And many journals put backlog information on their webpages.

I must stress here that it is a hard and fast rule in academics—and most journals will state this explicitly on their Instructions to Authors page—that you may submit a paper to only one journal at a time. This dictum is in place partly because of tradition, but primarily because the journal does not want to waste referees' time nor its own time. And journals certainly want to prevent various forms of academic dishonesty that could propagate from multiple submissions.[3] The books [KRA1] and [KRA2] discuss these matters in some detail.

The traditional way to submit a paper is in hard copy. *In a single envelope*, you send in two—or perhaps more!—copies of the paper printed one side only and a cover letter telling the editor or secretary what he/she is receiving. The cover letter should give your name (and those of your co-authors), the paper's name, your affiliation, and all your contact information (mailing address, email address, phone number, fax number, and so forth). If you are going to be traveling, or going on sabbatical, that should be mentioned in the letter. At various times the journal will need to contact you (to read proofs, sign copyright transfer forms,

[3] *Matters are different when you are attempting to get a book manuscript published by a commercial publisher. Then it is allowed, and indeed expected, that you will submit your project to more than one publisher at once. The reason for this difference in the rules is partly custom, but also that referees for book manuscripts are paid for their work.*

and so forth). Most likely email will expedite communication, but it is always a good idea to have all the key information about yourself on a single sheet of paper that the journal has on file.

These days many journals will accept a paper electronically. That means that you send an email to a designated address (the Instructions to Authors page will provide that information), and include the paper as an attachment. Well, electronics are confusing. Should you send your TEX source file and all the `*.eps` files for your figures plus all your style files and your font files? Decidedly not. This would give the recipient myocardial infarction, and it is highly inappropriate. The right thing to send in at first[4] is an `Acrobat` or `*.pdf` file. If properly prepared, this will have all the graphics and fonts embedded in it, so that anyone with an `Acrobat` reader (freely downloadable from the Web) can read it or print it out just as it was meant to appear. The referees will have no trouble reading the file, and neither will the editors or the clerical staff.[5] Be sure that the cover email contains all the contact information that was described above.

Patience

Dealing with an academic journal requires a good dose of patience. Generally speaking such operations are understaffed, or perhaps only staffed part-time. You may have to wait a month just to receive an acknowledgement that your paper has been received. That communication will often contain some generalized platitudes about when you can expect a referee's report, but they will usually not be very specific. And that is because they do not know. The journal will be well acquainted with its associate editors—those who handle the papers and do the legwork of getting the papers refereed—but they have no control whatever over the referees themselves. Even a well-meaning referee will have many distractions—a cat that is about to have kittens, a child graduating from high school, a house renovation, an upcoming surgery, or any number of other vicissitudes of life. Worse, the referee may be recovering from a drug dependency or getting a divorce. Or he/she may be terminally disorganized or hopelessly irresponsible. Who knows? So, if you are lucky, you will get a referee's report in three to four months. If you are not, it could take a year or more.

Most math journals use just one referee for a paper (in other disciplines this is *not* at all the norm; biology journals typically use two or three). But the *Annals* has an extremely high standard, and often uses at least two. The *Monthly* also has a high standard (of a somewhat different sort) and typically uses at least two referees. Thomas Hales's solution of the Kepler sphere packing problem, which appeared recently in the *Annals*, posed a particularly thorny refereeing problem (because the work involves massive computer calculations). The *Annals* enlisted a team of about a dozen Hungarian mathematicians who spent several years at the refereeing task.

I am an old dead white guy, and I usually do not care how long it takes for my papers to be refereed. I am busy writing other papers or doing other interesting tasks. I can wait. If your career is hanging in the balance, however, your view may be somewhat different. As a journal editor, I have certainly received very sincere and fervent letters from authors that said, "My tenure case is coming up in two months" or "my promotion is imminent" or "my grant is in the offing" and "I really need a decision." I do what I can to help—in some cases refereeing the paper myself—but in most instances I am at the mercy of the referee. I have had to abandon some referees—because they were so unresponsive—in effect forgetting that I had sent the paper to referee *A* and just starting again from scratch and sending the paper to referee *B*. Because of my experience and my contacts, I can usually get a paper refereed fairly expeditiously. In a crisis situation, I can usually help out. But you may be dealing with an editor who is less in tune, or less effectual, and your choices may be limited.

It is perfectly acceptable, after four or five or six months have elapsed, to write a note (by email or snail mail) to the editor or secretary who received your paper and inquire about its status. For this reason it is essential that you keep good records. Save all your correspondence concerning the paper. That way you will have, for example, a printout of the email acknowledging receipt (which will usually include the all-important manuscript number), and you will know just whom to write to and just what to say.

Of course always be polite in your correspondence. A two- or three-sentence note saying, "I submitted MS #xyzw on this date. Can you bring me up to date on the refereeing process? When might I expect a report?" will certainly do the job.

I once waited four years for a referee's report on a pretty good paper (from a journal that I am now too polite to name). I finally sent them a letter saying that if they could not come up with a report in six months then I would withdraw my paper and submit my work elsewhere. They got a report back to me in three months. I am not sure that this

[4] *Likely as not, as a result of the referee's reports, your paper will be revised. So it makes no sense to send in your TEX and other source files at the first submission. When the situation is finalized, and your paper is accepted, then the journal will certainly want your source files.*

[5] *Some journals have a webpage and accept submissions by way of* `ftp`. *The process is usually self-explanatory, and the Instructions to Authors page will tell you all the steps.*

is a good role model for you. Threatening people is no way to do business. But I was desperate.

The Denouement

Eventually your paper will be refereed and you will hold in your hands a referee's report (assuming that you know how to print out email text). You must learn to read referee's reports dispassionately. I have very occasionally received a report that said, "Krantz is a hail fellow well met. We are so lucky to have this paper. Publish it with all dispatch." But such is not the norm. Most referee's reports will contain a mixture of praise and constructive criticism. More often than not (assuming that the paper was accepted) you will have to revise the paper as a result of the report. Do so earnestly. Take all of the recommendations very seriously. Re-submit the paper with a careful log of how you handled each suggestion. If you disagree with a suggestion, say why (politely). Usually the editor will respect your judgment. Frequently the editor will send the paper back to a referee (it really ought to be the *same* referee, and that is usually who it will be, but not always). If you have played the game sincerely and carefully, the referee should then give your work his seal of approval and the paper will be accepted.

If your paper is rejected, do not lose heart. Most everyone has had papers rejected. Some of my most important and most influential papers have not only been rejected, but were treated rather shabbily. Celebrated authors from Jane Austen to Agatha Christie to Henrik Ibsen had their best work rejected. Be of stout heart. Learn what you can from the referee's report, and from the editor's comments, and then make the paper stronger. Submit to another journal. Do not waste time bemoaning your fate. Get the paper re-submitted and move on to more fertile territory.

Now let us return to the happy situation in which your paper (at least in principle) has been accepted. The rest is a formality. You will eventually receive page proofs of your paper (usually in electronic form, as a `*.pdf` file). You must respond to the proofs in a timely manner, offering any corrections that you may have. Then that is it. Eventually your paper will appear—either in the printed journal or the electronic journal or both. In the old days you would also get about 50 reprints of your paper, though this is a tradition that is fast vanishing. Many journals now will send you a `*.pdf` file of the paper in final form, and you can then print out reprints if you wish. The fact of the matter is that, if you want to distribute the paper among friends and colleagues, it is probably most expeditious to just send them the `*.pdf` file.[6] Or post the paper on a website like `arXiv`, or both. In fact you can post the paper on `arXiv` as soon as it is completed—even before you submit it to a journal. And you should. For the most important feature of the Web is that it can make your work available to a broad audience (essentially the entire world) rapidly and at no charge. Since you are trying to establish your reputation, you should do so.

[6] *On the other hand, if you want to give a copy of your paper to your grandmother, then hard copy is probably the way to go.*

Closing Thoughts

Developing your own ideas and publishing them is one of the most important and rewarding parts of academic life. For me, the best part is receiving feedback from students and colleagues and then engaging in useful repartee. This often leads to new insights and new collaborations, and makes the whole exercise productive and worthwhile. I hope that this article will have made it easier for you to become a part of this happy process.

References

[KRA1] S. G. Krantz, *A Primer of Mathematical Writing*, American Mathematical Society, Providence, RI, 1996.

[KRA2] ———, *Mathematical Publishing, A Guidebook*, American Mathematical Society, Providence, RI, 2005.