

How to Write a Security Paper*

Patrick McDaniel

Systems and Internet Infrastructure Security Laboratory (SIIS)
Department of Computer Science and Engineering
Pennsylvania State University

mcdaniel@cse.psu.edu

Abstract -

1 Introduction

Whether graduate student, faculty, or industrial researcher all members of the research community must publish papers. If you are one of these happy few, publications are simply the “coin of the realm” in academia that determines when and if you graduate, what kind of job you receive after graduation, when and if you get tenure, what your raise is, etc. For this reason, it is essential that everyone learns the craft of writing papers. Unfortunately this is not something that is taught well or often. Thus, many of us are left fumbling in the dark trying to figure out how to be successfully at this craft.

This paper attempts in some small way to begin document this craft. It covers topics from why one should write a paper to how are its elements selected and organized. Gleaned from years of practice, we outline tricks, patterns, and means of communicating your ideas into a well constructed paper.

Let us begin with a central question that one must understand to be effective at publishing, “*why should you write paper?*” Apart from the obvious self-serving reasons defined above, the central reason is a dissemination of ideas. That is, at least with respect to your chosen community, the objective of the authorship exercise is to relating a new understanding of some scientific phenomena.

A more difficult question is “*when should you publish?*” There are several schools of thought on this, each leading to a different professional style. Some people publish each new observation idea as it occurs, leading to many incremental works, and others publish highly novel papers infrequently. Where you choose to publish along this spectrum is a matter of taste and necessity.¹ Note that

*This is a work in progress—updates will be made as possible.

¹Certain programs and academic positions have documented or undocumented requirements for the number of publications needed to be progress toward graduation, tenure, promotion, or other advancement. It is essential that an author be aware of those requirements and adjust their publication style as needed to achieve their professional goals. Ask

some fields have a history of publishing at lesser (e.g., operating system research) or greater (e.g., bioinformatics) frequency, and thus your chosen field may have built in publication expectations.

The words **incremental** and **novelty** have special meaning in the academic community. Often used in a derogatory way, incremental papers demonstrate a small scientific improvement over prior works. For example, if one were to publish an algorithm for routing packets in ad hoc networks, then publish a second paper that altered the protocol in a small way to make it more efficient, the second paper would almost certainly be deemed to be incremental over the first. Novelty is simply the measure of the amount of new scientific observation over previous works. Novelty is unquestionably the one of the two most important factors in determining a paper’s value (the other being **impact**, as discussed in the next section).

Ultimately, the goal of publishing is to make the reader understand. This goal is often in conflict with issues such as completeness, brevity, etc. However, you must always be looking to find the best way to speak to the readers and to help them appreciate the substance and implications of your work. Finding the balance between what you need to say to document your observations and saying it in a way that is easy to understand is key to becoming a good scientific writer. The latter sections of this paper outline some guidelines that are used to help strike this balance.

2 The Process of Publishing

An important precursor to understanding how to write a paper is an appreciation for how it will be evaluated and used. There are many different kinds of publications, each with a unique way of being created, evaluated, and eventually disseminated. This section covers the central types of publications and how they are evaluated.

Note that not all publications are valued equally: a single conference paper published at the right conference

multiple senior people in your organization for their thoughts.

containing an important idea can make a career. Conversely, a hundred papers published in lesser conferences may not be sufficient to land a good job or achieve tenure.

2.1 Technical Reports

A technical report is a paper published by an organization or department. These documents are not a *peer reviewed*, which means that they are not evaluated for correctness or presentation beyond a cursory inspection. Because these are seen as “self-publications”, these documents are generally not given much weight as scientific output.² Technical reports are generally disseminated freely from the parent institutions. However, it is sometimes the case technical reports are not distributed, such as the case where there are intellectual property considerations (e.g., the ideas are being evaluated for patenting or commercialization), or where the work will be submitted to a blind submission process (see below).

So why publish technical reports? First, technical reports are a way of disseminating results before they are mature enough for publication. Technical reports can be updated many times as the work matures toward eventual publication. The second reason to publish technical reports to develop a paper trail for your progress as a researcher. This is particularly important for graduate students and junior faculty, who will be making the case that they have an established *research agenda*.³

Citations for technical reports generally look like this [1]:

Kevin Butler, Stephen McLaughlin, and Patrick McDaniel. Rootkit-Resistant Disks. Technical Report NAS-TR-0089-2008, Network and Security Research Center, Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA, USA, April 2008.

where the *technical report number* is NAS-TR-0089-2008, the authors are Steve, Kevin, and Patrick, and the institution and date of initial creation are identified.

One school of thought in treating technical reports is to create a technical report when some preliminary version of the paper is ready and list those on your professional vita. As the research (and paper) matures, the report is updated. When the paper is accepted, the author generally replaces the technical report with conference version. Listing both a published and TR version of a paper on

²There are many exceptions to this rule. Particularly in industrial research laboratories where ideas and systems often carry more weight than academic publications.

³A research agenda is the body of work you pursue as a researcher. A research agenda that shows a coherent line of inquiry on a single or set of problems. Having many disconnected papers on different topics is generally viewed as less effective, but some have built substantial careers this way. The latter approach is much harder.

your vita is generally not viewed positively (some view this behavior as an attempt to “pad” a vita).

2.2 Conference Papers

The main vehicle for publication in computer science is conferences. Conferences are meetings generally held once a year and attended by members of the specific community. For example, the USENIX Security Symposium⁴ is held every year and publishes systems and applied papers in computer and network security. Most of the security community attend the conference and exchange ideas, listens to talks, socialize, etc. Attending conferences is absolutely essential to professional development (you need to be seen to become known).

Most of the time a conference’s quality determines the professional value of the paper. In general, conferences in a particular field are considered **first tier** (the top conferences), **second tier** (solid conferences, but not as selective), and **third tier** (conference that are not very discerning). Again deciding where to submit your papers is a matter of taste and necessity. Note that some institutions will ignore some classes of papers in evaluating, for example, tenure cases.

Fourth tier conferences fulfill a special class. These conferences are set up to make money for the organizers, and are often held in desirable places such as Hawaii. The conferences generally accept papers without serious review (although they often pretend to make an effort). Examples of papers that repeat the same sentence over and over for 10 pages [] or are based on random language generators [] represent some of the most entertaining lore in academia. The overwhelming number of these papers are simply garbage created by people gaming the system to create a bogus publication record or to use research money to pay for a vacation. *Having these papers on your resume is very, very dangerous professionally.* Submit and attend at your own peril.

One of measure of a conference is its selectivity. This is grossly measured in an **acceptance rate**, which is the percentage of submitted papers that are accepted. Generally the lower acceptance rate the better the conference. Be careful in looking only at this metric however, as some of the worst conferences receive many, many bad papers resulting a rather artificially low acceptance rate. First tier conferences have acceptance rates generally below 20%, but many of the best are below 10%.

While each conference has its own formula, most papers are judged based on the following criteria:

- **Novelty** - As mentioned above, novelty is a measure of how “new” the paper content is. A novel paper can identify a new problem, or address an old paper in

⁴A symposium is just another name for a conference.

a new way. Incremental (papers with little novelty) are difficult to publish.

- **Importance** - Also known as impact, this measures the scale of the consequence of the result. For example, a paper outlining an attack on a widely used operating system may have tremendous importance (presumably because it affects many people). Similarly, a theoretical paper that identifies a faster algorithm for routing may impact the way that network devices are constructed.
- **Correctness** - Correctness is the degree to which the science is correctly executed and the right conclusions drawn. Papers that use the wrong methodology, or incorrectly or inappropriately apply methodology will be deemed incorrect, and thus scientifically deficit, e.g., errors on proofs or misread data. Incorrect papers are almost universally rejected. Be safe, get it right.
- **Presentation** -
- **Relevance** - Every venue, be it conference or journal or otherwise has a given "scope". The scope of the venue dictates the kinds or problems considered and the kinds of methodologies appropriate for submission. While there are sometimes exceptions, topics outside the scope of the venue will be discarded by the reviewers. Note that general conferences in topics like security are pretty flexible about scope, whereas smaller, more focused venues like workshops are less so.⁵
- **Excitement** - Sometimes related to novelty, this is the measure of how timely the paper is. As I write this, topics in cellular phones, cloud computing and social networks are very timely—people get excited reading about these areas because they are popular and changing the way we interact with computational infrastructure. Thus, they make good paper because they naturally draw an audience. As these technologies mature (and more is known about them and the security they refer to), papers become less exciting.
- **Overall** - This is cumulative score of the paper after taking into consideration all of the above elements. This is the score that is used in almost all deliberations in a program committee meeting.

Note that you do not need to have a high score in all elements to be published—a very innovative paper that is exciting can overcome limited problems with, for example, relevance or presentation. However, be warned: low novelty or low correctness is almost certainly deadly in all but the least competitive venues.

⁵The scope of the venue you are attempting to published in will almost certainly be defined on the accompanying website. When in doubt, contact the organizers/editors of the target venue.

Almost all competitive conferences hold program committee meetings. The PC meetings almost follow a familiar pattern, with small variations. The program chair(s) will select a subset of papers to be evaluated that scored on average highly, were interesting, or where there was disagreement on the value of the paper. This set typically is about twice the size of the eventual program (e.g., if the conference typically accepts 25 papers, you can assume the PC will discuss 50 of them during the PC meeting). Over the course of the one or two day meeting, each of the papers will be discussed by the PC and a determination made. Note that this process is a social one, which often leads to conflict and compromise.⁶

One often has to be aware of the makeup of the program committee to be success. Know who is likely to receive your paper and make sure you understand their prior work in the area. Be sure to cite the relevant works of the PC members—almost every PC member will check to make sure of this even before they read the abstract.

Warning: Beware of conferences that accept far too many papers, have little or no quality control, and often are set to make money for the organizers. Such publications on a CV can often raise red-flags with employers or promotion committees. Always understand the quality of the conference before you submit. Remember, your publication acts as an implicit endorsement of the venue. For this reason, people who regularly publish at these venues are almost certainly going to be judged harshly.

2.3 Workshop Papers

2.4 Journal

2.5 Books

3 Structure

How much to write - the bell curve of writing

4 Elements of a Paper

4.1 Abstract

[2]

⁶Pick up any conference proceedings and there will be at least one paper which does not have the same quality of the remaining papers. That paper is almost certainly the result of a compromise or one of the late papers accepted when the PC was tired.

4.2 Introduction and Conclusions

4.3 Related Work

4.4 Problem Statement

4.5 Solution

4.6 Evaluation

5 Evaluation

6 Hints on Successful Writing

Becoming a successful paper writer requires mastering a number of skills and strategies that make your work more effective at communicating science and attractive for publication. Below are a list of hints and observations I have learned over many years that may aid you in improving these skills.

Writers, particularly new ones, will be inclined to try to write text that is "sophisticated". You must resist this urge.

One way to make your paper more understandable is to write short sentences with clearly defined scope. Long multi-topic sentences can be effective when used infrequently, but make reading very difficult.

Don't pad.

Know your terminology.

7 Conclusions

It is important to remember that even the best researchers in the world fail more often than they are successful. Sometimes even the best papers take two or more tries to be accepted. Learning to accept this reality is often one of the most difficult tasks in becoming a successful researcher. People that dwell on the failures too long waste effort and end of being miserable in their vocation. This is particularly true of new graduate students: I have seen numerous graduate students who simply could not come to terms with paper rejections and left the program before really learning how to write papers. *Failing to learn from failure* is an equally fatal problem—those who do not recognize why their papers are being rejected are likely to repeat these problems.

References

- [1] Kevin Butler, Stephen McLaughlin, and Patrick McDaniel. Rootkit-Resistant Disks. Technical Report NAS-TR-0089-2008, Network and Security Research

Center, Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA, USA, April 2008.

- [2] Patrick McDaniel, Atul Prakash, and Peter Honeyman. Antigone: A Flexible Framework for Secure Group Communication. In *Proceedings of the 8th USENIX Security Symposium*, pages 99–114, August 1999. Washington, DC.