

Note: *This is a recently launched work in progress—updates will be made frequently. Please check URL for updates.*

<http://www.cse.psu.edu/~mcdaniel/cse543-f08/docs/howto-paper.pdf>

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 senior or junior, or graduate student, faculty, or 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, 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.

This paper attempts to 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 exercises 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 published along

this spectrum is a matter of taste and necessity.¹ Note that 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 will have some impact on the publication expectations.

The words **incremental** and **novelty** have special meaning in the 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. Note that 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 content and implication of your ideas. 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 ways that one can strike this balance.

¹Certain programs and academic positions have documented or undocumented requirements for the number of publications needed to be progressing adequately toward graduation, tenure, promotion, etc.. It is essential that an author be aware of those requirements and adjust their publication style as needed.

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

2.1 Technical Reports

2.2 Conference Papers

While each conference has its own ways of measuring a paper, the following are frequently seen as the key dimensions.

- **Novelty** -
- **Importance** - Also known as impact, this measures the scope 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** -
- **Presentation** -
- **Relevance** -
- **Excitement** -
- **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.

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.²

One often has to be aware of the makeup of the program committee to be successful. 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 Journal

2.4 Books

3 Structure

How much to write - the bell curve of writing

4 Elements of a Paper

4.1 Abstract

[1]

²Pick 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 Conclusions

References

- [1] 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.