

Writing Tips

Andrzej Duda
Grenoble Institute of Technology
CNRS Grenoble Informatics Laboratory UMR 5217
Grenoble, France
Email: Andrzej.Duda@imag.fr

*“La forme, c’est le fond qui revient à la surface.”
“Style is the substance of the subject called unceasingly to
the surface.”*

— Hugo

Abstract—When working on many papers, I realized that I correct almost always the same mistakes. Here is a compilation of examples illustrating various types of mistakes. I also gathered a lot of useful tips on how to correct them. The report includes the text by Mark Allman on the view of a referee on the process of writing papers.

I. COMMON MISTAKES

Here are the most common mistakes that you can easily avoid.

- 1) Nouns in a nominal group are singular except the last one.

WRONG: packets delivery ratio

GOOD: packet delivery ratio

WRONG: The simulations results show...

GOOD: The simulation results show...

- 2) Resist the temptation to use long strings of nouns as adjectives.

WRONG: Consider the packet switched data communication network protocol problem.

- 3) “Which” vs. “That”.

WRONG: We present the structure of packets which contain the neighbor addresses.

GOOD: We present the structure of packets that contain the neighbor addresses.

More on this topic.

“That” is the defining or restrictive pronoun, “which” the non-defining or nonrestrictive.

The lawn mower that is broken is in the garage. (Tells which one.)

The lawn mower, which is broken, is in the garage. (Adds a fact about the only mower in question.)

If you can substitute ‘that’ for ‘which’, do it.

- 4) Punctuation in phrases with “which” and “that”.

GOOD: All the students that know when to use “which” and “that” will pass the quiz.

GOOD: The exam, which took place at the beginning of class, was not difficult.

- 5) Punctuation in enumerations.

GOOD: It had bells and lights.

GOOD: It had bells, whistles, and lights.

- 6) Do not omit “that” when it helps the reader to parse the sentence.

WRONG: Assume A is a group.

GOOD: Assume that A is a group.

- 7) Simplify “which is”.

WRONG: The figure presents the time interval which is available for other nodes.

GOOD: The figure presents the time interval available for other nodes.

- 8) Avoid passive voice.

WRONG: The simulation results are presented in Figure 1.

GOOD: Figure 1 presents the simulation results.

- 9) Check for missing articles, particularly if your native tongue does not have them. Roughly, concepts and classes of things do not, most everything else more specific does. Do not use articles in front of proper nouns and names.

GOOD: Routers route packets.

GOOD: The router architecture we consider uses small rodents.

GOOD: Internet Explorer is a popular web browser.

GOOD: The current version number is 5.0.

GOOD: Bill Gates did not write Internet Explorer.

- 10) Do not use “etc.” unless the remaining items are completely obvious.

ACCEPTABLE: We shall number the phases 1, 3, 5, 7, etc.

UNACCEPTABLE: We measure performance factors such as volatility, scalability, etc.

- 11) Avoid “etc.”; use “for example”, “such as”, “among others” or, better yet, try to give a complete list (unless citing, for example, a list of products known to be incomplete), even if abstract.

- 12) Uncountable nouns are substances or concepts that we cannot divide into separate elements (e.g., advice, information, news, work)

WRONG: informations, interferences

GOOD: information, interference

WRONG: an information

GOOD: a piece of information

GOOD: the information

- 13) Possessive 's. Concerns persons not things.
 WRONG: packet's size
 GOOD: packet size
- 14) Allow. Permit. Permit is more formal than allow. Enable.
 GOOD: They do not allow smoking here.
 GOOD: They do not allow us to smoke here.
 GOOD: They do not permit smoking here.
 GOOD: They do not permit people to smoke here.
 GOOD: It enables nodes to detect collisions.
- 15) Bibliographic references are not nouns. The phrase needs to make sense if you delete them.
 WRONG: The authors in [7] proposed the best protocol for flooding.
 GOOD: The authors proposed the best protocol for flooding [7].
 WRONG: In [7], Jain et al. proposed the best protocol for flooding.
 GOOD: Jain *et al.* proposed the best protocol for flooding [7].
- 16) No articles before symbols.
 WRONG: Based on the distance d , we can compute...
 GOOD: Based on distance d , we can compute...
- 17) Symbols in different formulae must be separated by words.
 WRONG: Consider S_q , $q < p$.
 GOOD: Consider S_q , where $q < p$.
- 18) Protocol abbreviations typically do not take an article, even if the expanded version does.
 GOOD: The Transmission Control Protocol delivers a byte stream.
 GOOD: TCP delivers a byte stream.
 GOOD: The TCP design has been successful.
- 19) Check that abbreviations are always explained before use.
 WRONG: IEEE has developed a standard for LR-WPANS.
 GOOD: IEEE has developed a standard for Low-Rate Wireless Personal Area Networks (LR-WPANS).
- 20) Common error.
 WRONG: comprised of
 GOOD: composed of
- 21) Sections and figures.
 WRONG: In section 5, we show...
 GOOD: In Section 5, we show...
 WRONG: We can see in figure 5...
 GOOD: We can see in Figure 5...
 but
 GOOD: We will explain in the section that...
 GOOD: We can see in the figures that...
- 22) The text should make sense if we read through it omitting the titles of subsections.

WRONG: 2. Contour Integration. This technique, invented by Cauchy, defines...

GOOD: 2. Contour Integration. The technique of integrating along curves in the complex plane, invented by Cauchy, defines...

- 23) Capitalize titles (also in bibtex references). To keep titles capitalized use double `{ { }` in a bibtex item:

```
@article{Jain2001,
  Author = {R. Jain and A. Puri and R. Sengupta},
  Title = {{Geographical Routing Using Partial
  Information for Wireless Ad Hoc Networks}},
  Journal = {IEEE Personal Comm.},
  Month = {February},
  Year = {2001}
}
```

WRONG: Classifying service flows in the encrypted skype traffic.

GOOD: Classifying Service Flows in the Encrypted Skype Traffic.

- 24) Adverbs before verbs.

WRONG: The protocol operates usually in indoor environments.

GOOD: The protocol usually operates in indoor environments.

- 25) Frequent repetition of a word like “this”, “they”, “just”, or “then”.

Avoid nonreferential use of “this”, “that”, “these”, “it”, and so on. Requiring explicit identification of what “this” refers to enforces clarity of writing. Here is a typical example of nonreferential “this”:

WRONG: Our experiments test several different environments and the algorithm does well in some but not all of them. This is important because...

- 26) Use strong verbs instead of lots of nouns and simple terms rather than fancy-sounding ones.

WRONG: make assumption

GOOD: assume

WRONG: is a function of

GOOD: depends on

WRONG: is an illustration

GOOD: illustrates, shows

WRONG: is a requirement

GOOD: requires, needs to

WRONG: utilizes

GOOD: uses

WRONG: had difference

GOOD: differed

II. VARIOUS TIPS

Here is a compilation of several tips on writing style [1], [2]:

- 1) Avoid words like “this” or “also” in consecutive sentences.
- 2) Use “Eq. 7”, not “Equation (7)”, unless you need to fill empty pages.
- 3) Don't start sentences with “That's because”.

- 4) In formal writing, contractions like “don’t”, “doesn’t”, “won’t” or “it’s” are generally avoided.
- 5) Be careful not to confuse “its” with “it’s” (it is).
- 6) Avoid clichés like “Recent advances in ...”.
- 7) Units are always in roman font, never italics or LaTeX math mode. Units are set off by one (thin) space from the number. In LaTeX, use `~` to avoid splitting number and units across two lines. `\;` or `\,` produces a thin space.
- 8) Use “kb/s” or “Mb/s”, not “kbps” or “Mbps”—the latter are not scientific units. Be careful to distinguish “Mb” (Megabit) and “MB” (Megabytes), in particular “kb” (1,000 bits) and “KB” (1,024 bytes).
- 9) Avoid excessive use of “i.e.”. Vary your expression: “such as”, “this means that”, “because”, “I.e.” is not the universal conjunction!
- 10) Remember that “i.e.” and “e.g.” are always followed by a comma.
- 11) Do not use ampersands (&) or slash-abbreviations (such as s/w or h/w) in formal writing.
- 12) “Respectively” is preceded by a comma, as in “The light bulbs lasted 10 and 100 days, respectively.”
- 13) “Therefore”, “however”, “hence”, and “thus” are usually followed by a comma, as in “Therefore, our idea should not be implemented.”
- 14) Never use “related works” unless you are talking about works of art. It’s “related work”.
- 15) Do not refer to colors in graphs. Most people will print the paper on a monochrome (black and white) printer and will have no idea what you are talking about. Make sure that graph lines are easily distinguishable when printing on a monochrome printer.
- 16) Use hyphens for concatenated words: “end-to-end architecture”, “real-time operating system” (but “the computer may analyze the results in real time”), “per-flow queuing”, “flow-enabled”, “back-to-back”, ...
- 17) Don’t overuse dashes for separation, as they interrupt the flow of words. Dashes may be appropriate where you want to contrast thoughts very strongly or the dash part is a surprise of some sort. Think of it as a very long pause when speaking. In many cases, a comma-separated phrase works better. If you do use a dash, make sure it’s not a hyphen (- in LaTeX), but an em-dash (--- in LaTeX).

GOOD: Learn to use the em-dash—it is a good friend.
- 18) “While” instead of “and”, “but”, “although”. In general “while” should be used only in the strict sense of “during the time time”; so search for “while” in your text to see if the sentence is about time or could be replaced with “although”.
- 19) Provide sufficient information in the caption of a figure so that the reader does not need to look for the description in the text.

GOOD: Figure 5. Outstanding data (*cwnd*) and queue size at R1 and R2. Equal buffer size

(30,000 B) on both sides. The downlink buffer remains empty for significant periods of time, whereas the uplink buffer never empties.

- 20) Motivate the reader for what follows. Perhaps the most important principle of good writing is to keep the reader uppermost in mind: What does the reader know so far? What does the reader expect next and why?
- 21) Get the attention of your readers immediately. Snappy titles, arresting first sentences, and lucid initial paragraphs are all methods of doing this.
- 22) Don’t use the same notation for two different things. Conversely, use consistent notation for the same thing when it appears in several places.

III. TERMS AND PHRASES TO AVOID IN A FORMAL DOCUMENT

Douglas Comer gives a list of terms and phrases to avoid in a PhD dissertation [3]. It contains many examples that help to formalize scientific discourse, however, the current usage may evolve to something more relaxed, so use your judgement.

- adverbs

Mostly, they are very often overly used. Use strong words instead. For example, one could say, “Writers abuse adverbs.”
- jokes or puns

They have no place in a formal document.
- “bad”, “good”, “nice”, “terrible”, “stupid”

A scientific dissertation does not make moral judgements. Use “incorrect/correct” to refer to factual correctness or errors. Use precise words or phrases to assess quality (e.g., “method A requires less computation than method B”). In general, one should avoid all qualitative judgements.
- “true”, “pure”,

In the sense of “good” (it is judgemental).
- “perfect”

Nothing is.
- “an ideal solution”

You’re judging again.
- “today”, “modern times”

Today is tomorrow’s yesterday.
- “soon”

How soon? Later tonight? Next decade?
- “we were surprised to learn...”

Even if you were, so what?
- “seems”, “seemingly”,

It doesn’t matter how something appears.
- “would seem to show”

All that matters are the facts.
- “in terms of”

Usually vague.
- “based on”, “X-based”, “as the basis of”

- Careful; can be vague.
- “different”
 - Does not mean “various”; different than what?
- “in light of”
 - Colloquial.
- “lots of”
 - Vague and colloquial.
- “kind of”
 - Vague and colloquial.
- “type of”
 - Vague and colloquial.
- “something like”
 - Vague and colloquial.
- “just about”
 - Vague and colloquial.
- “number of”
 - Vague; do you mean “some”, “many”, or “most”? A quantitative statement is preferable.
- “due to”
 - Colloquial.
- “probably”
 - Only if you know the statistical probability (if you do, state it quantitatively).
- “obviously, clearly”
 - Be careful: obvious/clear to everyone?
- “simple”
 - Can have a negative connotation, as in “simpleton”.
- “along with”
 - Just use “with”.
- “actually, really”
 - Define terms precisely to eliminate the need to clarify.
- “the fact that”
 - Makes it a meta-sentence; rephrase.
- “this”, “that”
 - As in “This causes concern.” Reason: “this” can refer to the subject of the previous sentence, the entire previous sentence, the entire previous paragraph, the entire previous section, etc. More important, it can be interpreted in the concrete sense or in the meta-sense. For example, in: “*X does Y. This means ...*” the reader can assume “this” refers to Y or to the fact that X does it. Even when restricted (e.g., “this computation...”), the phrase is weak and often ambiguous.
- “You will read about...”
 - The second person has no place in a formal dissertation.
- “I will describe...”
 - The first person has no place in a formal dissertation.

- If self-reference is essential, phrase it as “Section 10 describes...”
- “Hopefully, the program...”
 - Computer programs do not hope, not unless they implement AI systems. By the way, if you are writing an AI thesis, talk to someone else: AI people have their own system of rules.
- “...a famous researcher...”
 - It does not matter who said it or who did it. In fact, such statements prejudice the reader.
- Be careful when using “few, most, all, any, every”.
 - A dissertation is precise. If a sentence says “Most computer systems contain X”, you must be able to defend it. Are you sure you really know the facts? How many computers were built and sold yesterday?
- “must”, “always”
 - Absolutely?
- “should”
 - Who says so?
- “proof”, “prove”
 - Would a mathematician agree that it is a proof?
- “show”
 - Used in the sense of “prove”. To “show” something, you need to provide a formal proof.
- “can/may”
 - Your mother probably told you the difference.

IV. WRITING THE INTRODUCTION TO A PAPER

There are a lot of good tips for writing papers [4], [5], [6], [7], [8], [9], [10], [11], [12], [13]. As “*Abstract and Introduction are the most important sections*” [14], here are just some clues about writing introductions.

Unless there is a good argument against it, the Introduction should consist of five paragraphs answering the following five questions [4]:

- 1) What is the problem?
- 2) Why is it interesting and important?
- 3) Why is it hard? (e.g., why do naive approaches fail?)
- 4) Why has not it been solved before? (Or, what’s wrong with previous proposed solutions? How does mine differ?)
- 5) What are the key components of my approach and results? Also include any specific limitations.

Then, have a final paragraph or subsection: “Summary of Contributions”. It should list the major contributions in bullet form, mentioning in which sections they can be found. This material doubles as an outline of the rest of the paper, saving space and eliminating redundancy.

V. FINAL TRUTHS

Here are some handy maxims (notice that each sentence below is “wrong” in the sense that it illustrates bad usage).

However, remember that you can brake rules if it improves clarity.

- Watch out for prepositions that sentences end with.
- When dangling, consider your participles.
- About them sentence fragments.
- Make each pronoun agree with their antecedent.
- Don't use commas, which aren't necessary.
- Try to never split infinitives.

Put yourself in the reader's place! Ask yourself what the reader knows and expects to see next at some point in the text.

Do organize and do not distract.

Bad writing comes from bad thinking, and bad thinking never produces good writing.

Easy writing makes damned hard reading.

Less is more.

Mais malheur à l'auteur qui veut toujours instruire ! Le secret d'ennuyer est celui de tout dire.

Voltaire, *De la Nature de l'Homme*.

VI. MARK ALLMAN'S PLEA

Here are the comments of Mark Allman quoted inline [15]. They are specific to the Networking community, but most of the remarks also concern other domains.

I have just completed my first (maybe only!) stint on the program committee for the ACM SIGCOMM annual symposium. While I have refereed many papers over the years I have never before read so many submissions in such a short span of time (I filed 20 reviews within about a month and read/skimmed many more papers). This experience was eye-opening in that I now realize that the community generates a large amount of *junk*. And, I don't mean "junk" in the sense of bad ideas. In general, each paper I reviewed for SIGCOMM this year had at least a nugget of an interesting idea. I mean "junk" in the "it took me twice as long to read this as it should have" sense. In other words, the paper was sloppy. I had a very difficult time trying to figure out what the paper was proposing or what the experiments were really showing. The following is a short list of *easy* ways to improve your research papers.

Most of the following comments are general and I presume would apply to any situation in which one is attempting to publish scholarly work. However, the last comment is directed specifically at the internetworking research community.

A. Fire Up the Spell Checker, Please!

I cannot spell either. I have, however, mastered `ispell`. It seems to me that everyone has access to a good spell checker these days. Not using the spell checker just indicates that either you are sloppy or you do not care. If you're sloppy with your writing why should I believe you are not sloppy with your

experiments/analysis as well? If you don't care about your paper, why should I?

B. Sloppy Papers Are Long Shots

Proofread your paper. I see lots of mistakes that indicate the paper was not looked over before it was submitted (e.g., "The foo algorithm works better than the the bar algorithm."). Again, if you do not take enough care in presenting your work, why should I be compelled to recommend accepting it? Sloppy papers give the indication of sloppy research. As well as reflecting poorly on your work, such papers are often very difficult to read and understand. Presumably, the point of writing a paper is to convey new information to the world. Therefore, there is major value in clear writing. If a referee can't get the point of the paper, or has a very difficult time getting the main idea, there is a high likelihood that the paper will be rejected.

C. I Read Submissions in My Lazy Boy

I print out all papers I review. I read them in my lazy boy¹. I scribble notes on them. Do not assume that I am looking at color plots! Color is certainly nice and helps out immensely in conference presentations. But, color plots come out of my black and white printer very badly in most cases. Besides, most conferences/journals publish black and white papers. Make it easy for my weary eyes to tell the difference between the lines on your plots. Do not plot on a grey background. Make the font on the plots readable without a magnifying glass. Help me out! Please!

Often the plots are key to understanding the paper. And, at the risk of being redundant: If I cannot understand the paper I will recommend rejecting it.

D. Don't Waste My Time I

You do not need to give an in-depth tutorial of all background material. I cannot give you a good rule of thumb to help you know how much background material to include. It depends on the conference to which you are submitting, the maturity of the background work, etc. But, I can tell you that there are some topics (e.g., TCP congestion control) that are well understood by people attending any networking conference. You do not need to describe the algorithms in painful detail. A short paragraph re-capping the main ideas and providing references to the original works should suffice.

E. Don't Waste My Time II

You do not need to provide me with an example of every kind of plotting strategy you use. When you first show a plot with notation that I might not

¹a comfortable coach

understand, explain the notation. But, using half a page to show me a plot that adds nothing to my understanding of the topic is a less than compelling use of real estate.

F. Don't Waste My Time III

Roughly stick to the page limit and the requested format. Don't make me try to guess how long your paper really is because you 1.75-spaced the paper rather than double-spacing it as requested in the call for papers.

(Note to program committee chairs and journal editors: This reviewer's opinion is that if a submission is not obviously in the ball-park of the page limit it should be immediately returned to the authors until such time that it is in the ball-park.)

G. Answer All the Easy Questions

If you need a constant for the foo algorithm and you define $c = 17$, I want to know why you chose 17. Also, some items in papers jump out at reviewers as odd. For instance, I was recently reviewing a paper that simulated a network with a bottleneck bandwidth of "about T1 rate". The actual rate of the link turned out to be something like 1.3 Mbps. This begs the obvious question of why the authors used 1.3 Mbps rather than using T1 rate (~1.5 Mbps). Even if the explanation you end up with is rough, it is better than no explanation at all in the vast majority of the cases.

H. Make My Life Easy

It is not a super-big deal, but it is always nice to get a paper that is easy to read. This includes things like using a very clear font that is of readable size and numbering the pages.

I. A Little Restraint, Please

Your paper does not solve the entire world networking problems. Really. Many papers rub reviewers the wrong way by making bold claims about solving all sorts of problems without any sort of evidence (or only weak evidence). It is much better to show some restraint and perspective when discussing your results. It is alright to solve small problems, or to only show preliminary results that *may indicate* a solution. Just don't over claim what you have shown.

J. Use Units

"The length of our data transfers is 1" means nothing. 1 what? 1 packet? 1 KB? 1 minute? Sometimes it is obvious from the context—sometimes it isn't. If you use units you don't have to worry about it.

K. We're All Not Mathematicians

Theorems are quite a necessary part of many papers. However, it is not necessary to give *only* a theorem with no summary or context. Summaries aid me as a reviewer to gain an intuitive understanding before trying to dig into the math to make sure what you have it correct—thus making my task easier and less error prone, I believe. As a reader of a paper, I often just want the high-order bit without slogging into the theorems and proofs to get said bit.

L. On References

One common mistake that authors make is failing to include relevant references. It is better to err on the side of citing too much related work.

Next, when previous work must be criticized, do so in a constructive manner. In other words, do not say something like "Zevon's [1] work is easily shown to be less effective than the work presented in this paper". Mr. Zevon might be reviewing your paper! Something more along the lines of: "Zevon [1] first introduced the foo algorithm for decreasing the size of routing tables. We show a more robust algorithm, bar, that goes a step further and reduces the size of routing tables by 50% when compared to the foo algorithm."

Finally, do not overstate what a paper you are citing says. It is often better to stick with the author's conclusions and not come up with your own. For instance, many times I see simple simulation studies cited as "proving" the algorithm in question is safe and works well for the entire Internet. If one goes back to the original paper the conclusion is not nearly as bold, often claiming that the simple simulations are an indication that the algorithm is workable and that future work on testing the algorithm in the Internet is needed. Overstating previous work may make your argument stronger if you are not caught. But, when I notice such references they cast your paper in a not-so-wonderful light in my eyes.

M. Read These Papers

I have noticed that a number of other reviewers are including pointers to the papers on TCP [16] and on simulating the Internet [17] in their reviews. They do not necessarily provide hard-and-fast rules for how to conduct research. However, they do provide some collected wisdom that should at least be considered in your research efforts.

In addition, the note by Partridge provides specific guidance on getting a paper accepted for presentation at SIGCOMM [18].

I don't want to make it sound like I think you need to write *perfect* papers when submitting to a conference or journal. A mis-spelling here or there, an extra word, an extra page, etc. are all OK. They

will not be cause for your paper to be rejected. But, it is surprising how often authors make many of the mistakes listed above in a single submission. The bottom line is that if it is difficult for me to make it through your paper (because the writing is poor or I can't read the plots or whatever) there is a higher chance that I will err on the side of recommending that the paper be rejected. Furthermore, the above items are *easy* things to fix when compared with the *hard* task of conducting your investigation. As I said in the opening, I believe most of the papers I review have a decent idea – yet I recommend rejecting the vast majority of the papers I review. I believe many papers would have a much better chance at being accepted if the authors simply take a little more time in preparing their manuscript.

VII. CONCLUSIONS

After reading all this stuff, you may think that some remarks are contradictory. Yes, maybe—use your common sense to choose the right word, or better check on the Internet.

Good writing requires practice. A good starting point is to read well written papers, analyze, and use them as inspiration.

ACKNOWLEDGEMENTS

The report contains many examples gathered from various sources on the Internet. Their authors are gratefully acknowledged: Mark Allman, Douglas Comer, Donald Knuth, Daniel Lemire, David Patterson, and Henning Schulzrinne.

REFERENCES

- [1] H. Schulzrinne, “Common Bugs in Writing,” <http://www.cs.purdue.edu/homes/dec/essay.dissertation.html>, 2011.
- [2] D. Knuth, T. Larrabee, and P. Roberts, “Mathematical Writing,” http://jmlr.csail.mit.edu/reviewing-papers/knuth_mathematical_writing.pdf, 1987.
- [3] D. Comer, “How to Write a Dissertation or Bedtime Reading for People Who Do Not Have Time to Sleep,” <http://www.cs.purdue.edu/homes/dec/essay.dissertation.html>, 2011.
- [4] J. Widom, “Tips for Writing Technical Papers,” <http://infolab.stanford.edu/~widom/paper-writing.html>, 2006.
- [5] D. Lemire, “Write Good Papers,” <http://lemire.me/blog/rules-to-write-a-good-research-paper>, 2011.
- [6] H. Schulzrinne, “Writing Technical Articles,” <http://www.cs.columbia.edu/~hgs/etc/writing-style.html>, 2011.
- [7] S. Bloomer and M. Haas, “How to Write a Scientific Paper,” <http://www.aocs.org/files/PressPDFs/howto.pdf>, 2004.
- [8] M. Morrison, “Tips on Scientific Writing,” <http://www.nhn.ou.edu/~morrison/Teaching/WritingTips.pdf>, 2004.
- [9] M. Faloutsos, “The Structure of Paper/Report in Systems,” <http://www.cs.ucr.edu/~michalis/TECHWRITING/structure.html>, 2011.
- [10] —, “Good Sites for Researchers and Students,” <http://www.cs.ucr.edu/~michalis/techwriting.html>.
- [11] A. Zale, “How to Write Good,” <http://www.montana.edu/mtcfru/Zalefrontpage/technicalwritingtips.pdf>, 2004.
- [12] S. Jones, “How to write a good research paper,” <http://research.microsoft.com/en-us/um/people/simonpj/papers/giving-a-talk/writing-a-paper-slides.pdf>, 2004.
- [13] D. Patterson, “Dave Patterson’s Writing Advice,” <http://www.eecs.berkeley.edu/~pattrsn/talks/writingtips.html>, 2011.
- [14] M. Faloutsos, “Writing a Paper,” <http://www.cs.ucr.edu/~michalis/TECHWRITING/doing-research/writingPapers-michalis.pdf>, 2011.
- [15] M. Allman, “A Referee’s Plea,” <http://www.icir.org/mallman/plea.txt>, 2001.
- [16] M. Allman and A. Falk, “On the Effective Evaluation of TCP,” *SIGCOMM Comput. Commun. Rev.*, vol. 29, no. 5, 1999, <http://www.icir.org/mallman/papers/tcp-evaluation.pdf>.
- [17] S. Floyd and V. Paxson, “Difficulties in Simulating the Internet,” *IEEE/ACM Trans. Netw.*, vol. 9, August 2001, <http://www.icir.org/vern/papers/sim-difficulty.TON.2001.pdf>.
- [18] C. Partridge, “How to Increase the Chances Your Paper Is Accepted at ACM SIGCOMM,” *SIGCOMM Comput. Commun. Rev.*, vol. 28, no. 3, 1998, <http://www.acm.org/sigcomm/ccr/archive/1998/jul98/ccr-9807-partridge.html>.