



Three Sins of Authors in Computer

Science and Math

Over the last seven years, I've read perhaps four hundred papers in computer science and math. Thirty or so were well-written. These anomalies aside, extracting meaning from most of the papers was like sucking a camel through the eye of the proverbial needle upon which a thousand angels were dancing on my head, if I mix my metaphors right.

Most of us learn how to write technical articles by osmosis. After reading a hundred or so papers, we unconsciously pick up the common patterns of the field and begin to imitate them. This phenomenon is largely beneficial, but there are several truly odious habits that are almost universally and inadvertently adopted without critical evaluation. Hence, I've distilled below a few of my opinions on what I see in the math and computing literature, and on how you can avoid annoying me in the future.

I realize that the first priority of computer scientists and mathematicians is not literary artistry. I am not complaining about mere bad writing; any scientific field will have its fair share, and I'm quite willing to accept the occasional incoherent paragraph and unlabeled axis as the price I have to pay for being more employable than a Literature major. Rather, my complaints here are about bad writing that's considered good. The crimes I enumerate below are practices that are accepted or encouraged and occasionally even enforced. Some are sins that even good authors feel obliged to commit.

All of the specific examples I give are taken from papers published in refereed conferences or journals; most of these papers have, in my opinion, high technical merit.

Grandmothering

Every paper needs an introduction. In fact, the introduction is the most important part of your paper, because few of your readers will ever read beyond it. And there's not much hope that any of them will if you don't grab their attention from the start. So it's a mystery why so many papers begin with twaddle like this excerpt from a conference on high performance computing.

Massively parallel computers (MPCs), characterized by their scalable architectures, are a viable platform on which to solve the so-called grand-challenge problems. These distributed-memory systems are expandable and can achieve a proportional performance increase without changing the basic architecture. In order to take full advantage of scalable hardware, the application software must also be scalable to exploit the increased computing capacity.

If you find your thoughts drifting away, don't feel bad; we have evidence that the authors felt the same way - consider the near-meaninglessness of the second sentence. The real weakness, though, is that this extreme form of "grandmothering" has no other function than to tell you something you already know, and wouldn't be comprehensible if you didn't already know. The paragraph is entirely superfluous to any reader who knows the meanings of "scalable" and "grand-challenge problems." Other readers, who don't know the buzzwords, are discouraged from continuing. Everybody loses.

I'm not going to give you the usual advice that you fuss and fret over your introduction until it's perfectly attuned to the psychological motivations of every potential reader. If everybody had time to do that, bad writing wouldn't be a problem. Instead, I'm going to offer advice that will *save* you time: **get to the point!**

The authors of the excerpt above should recognize that the primary objective of the first paragraph is to explain the purpose of their paper and thereby interest you in reading the second paragraph. But they don't. They understand that an introduction is obligatory, but they don't really know what to do with one, and they're too timid to jump directly to the central idea of their paper - perhaps they've never seen anyone else do that. But, hey, all the other supercomputing papers they've ever read start with the same paragraph, so it can't be too bad, right?

A more subtle form of grandmothering appears in this excerpt from a linear algebra conference.

In recent years, the study of preconditioners for iterative methods for solving large linear systems of equations, arising from discretizations of stationary boundary value problems of mathematical physics, has become a major focus of numerical analysts and engineers. ...

In a paper directed at newcomers to the field, this introductory sentence might be appropriate. However, the bulk of the paper is accessible only to those sufficiently expert in the field to know everything in the first two paragraphs of the introduction cold. So why bother?

Rather than telling the reader what the paper is about, this author begins by explaining how important and interesting his field of study is. This is an awful (and awfully common) habit. To justify your work by pointing out that it's "a major focus of numerical analysts and engineers" betrays a little insecurity, and isn't a good justification anyway.

In conclusion, every column inch devoted to convincing parallel computing experts of the importance of scalability, or introducing preconditioners to multigrid gurus, is fifteen seconds brutally stolen from their lifespans.

A table of contents in a paragraph

At the end of many introductions, I find an oozing cyst like this moribund specimen:

This paper is organized as follows. In Section 2, we describe local transformations in k dimensions. In Section 3, we describe an incremental approach for constructing k -D Delaunay triangulations using local transformations. In Section 4, we prove that this approach always constructs a Delaunay triangulation. In Section 5, we describe three algorithms and a data structure based on this approach. In Section 6, we discuss the time complexities of the algorithms and present experimental results from our implementation of these algorithms.

Upon close examination one finds that this is really a table of contents, in paragraph format and without page numbers (yechh). It's choppy and, even though there is some logical flow from one sentence to the next, it's difficult to read. It's barely useful as an index anyway, as it does little more than repeat the section titles, which are more easily absorbed by skimming the article. I've learned to skip these eyesores without so much as a direct glance, and I bet you have too. Unfortunately, they've become institutionalized, and some researchers even hold them up as a staple of proper technical writing. I bet these same people wear polyester to the lab each day.

Some readers, having had no better medium from which to form a sense of aesthetics, might not see why I find these shotgun summaries so atrocious. Allow me to put forth a better alternative. The odd thing about the paragraph quoted above is that it appears at the end of an otherwise well-written two-page introduction. I submit that the references to each section of the paper should have been folded into the introduction, each appearing in its logical place. The sections of the paper follow a clear logical progression; the introduction should echo that progression, and include references to sections of the paper as appropriate. If it is clumsy to have the introduction echo the paper precisely, it's okay to refer to a few sections out of order.

Furthermore, each section number should be listed after its description, and not before. The formula is "Get them interested, then tell them where the information is." Nobody on the face of the earth gives a damn what Section 5 is about. But once you've got them interested in the cool results, they'll wanna know where to find them. A sentence beginning, "In Section 5, we..." is off-putting because there's no compelling reason to discuss Section 5 before its contents. You should write every sentence as if they'll toss your paper if the first half of the sentence isn't interesting.

Don't mention your "Conclusions" section at all unless you want to point out something specific about it. Your readers can find it without your help, thank you.

I shan't torture you by rewriting the author's entire introduction, but here's a sample paragraph, lifted from the middle of the introduction, which I've modified (by adding three words) to give some idea of what I have in mind. It still doesn't sound as good as it would had the introduction been written from scratch with my suggestions in mind, but it'll do.

We are interested in algorithms that can construct k -D Delaunay triangulations...[five and a half sentences omitted]. In this paper, we prove that local transformations can be used to construct k -D Delaunay triangulations using an incremental approach, and present algorithms **in Section 5** that are k -D versions of those in [13].

(Actually, I wouldn't have written it like this. For a single author to refer to himself as "we" is a custom that continues to mystify me, and I don't agree that "[13]" is a noun. But enough, already.)

Conclusions that don't

What's wrong with this excerpt, which occurs at the beginning of a section entitled "Conclusions"?

We have proven tight bounds on the cardinality of a triangulation in terms of local feature size and the smallest angle, up to constant factors. We have also shown that two triangulations with similar local feature size must have similar cardinality, up to a $1/a$ factor.

The answer? Why, it's not a conclusion at all. It's an introduction. Indeed, all the information in these two sentences can be found in this paper's introduction, and in its abstract too.

I once read a paper in which the conclusions were an almost exact copy of the introduction, changing only the tense of the verbs. That's unforgivable.

A lot of people have picked up the misconception that they should conclude their paper with a summary. The dictum is, "Tell them what you're going to tell them, then tell them, then tell them what you told them." Well, yes, yes, and no.

Conclusions should synthesize the results of your paper and separate what is significant from what is not. Ideally, they should add *new* information and observations that put your results in perspective. Here's a simple test: if somebody reads your conclusions before reading the rest of your paper, will they fully understand them? If the answer is "yes," there's probably something wrong. A good conclusion says things that become significant after the paper has been read. A good conclusion gives perspective to sights that haven't yet been seen at the introduction. A conclusion is about the implications of what the reader has learned. Of course, a conclusion is also an excellent place for conjectures, wish lists, and open problems.

If you don't have any conclusions, be honest with yourself and don't write a "Conclusions" section. You've probably been indoctrinated with the notion that it's bad to end a paper without conclusions. I absolutely disagree. But whether it's bad or not, it's surely worse to end a paper without conclusions and yet include a section entitled "Conclusions" anyway.

A postscript on FBAs

Really Bad Acronyms, or FBAs, are spawned by FNPLs (Nerdy Project Leaders) when naming new systems. I cannot comprehend why projects like FTMPs and NUMachine were not given more charming monikers such as Infectoid or Puggsley or Vomitsauce. These names would stick in people's minds as FTMPs can never dream.

Jonathan Shewchuk, 1997
jrs@cs.berkeley.edu