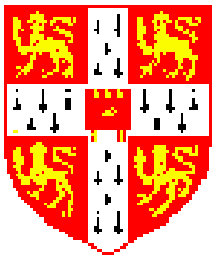

UNIVERSITY OF CAMBRIDGE

CENTRE FOR MATHEMATICAL SCIENCES

DEPARTMENT OF APPLIED MATHEMATICS
& THEORETICAL PHYSICS

HOW TO WRITE A PAPER?

Arieh Iserles

Cambridge

19th February 2007

Why write a paper?

Bad reasons:

- Everybody else does it: this is what academics do for living;
- I have been busy doing research, have been getting results and now I want everybody else to know about it;
- This is competitive world and unless I rush to record my achievements, they will be recognised as somebody else's achievements;
- I want dissertation, postdoc, career, tenure, fame! I want the world to be at my feet.

A good reason: (embarrassing in its simplicity)

- I want to communicate ideas to other people.

Why is this important?

Because every good text *must* be considered first and foremost from the point of view of its target audience.

The BGML holds about 75000 volumes of more than 2100 mathematics and theoretical physics journals. This is plenty of dead trees! Your paper needs to compete with thousands of other papers in a very crowded marketplace of ideas for the (very) limited attention of busy academics. It is your duty to persuade them that they should expand time and attention on your paper.

A good rule of a thumb is: *Ask yourself which sort of paper would you like to read – and then write it.*

Good reason 1 to read a paper:

Good title and abstract

A title and an abstract are your shop window. Most potential readers browse journals, whether in a library or on Internet, like shoppers on High Street. By this stage the purpose of the exercise is to grab their attention sufficiently to motivate them to read on:

1. Good title uses few well-chosen words to place your paper at the right corner of mathematical universe. Good abstract states clearly what is the subject area of the paper and its main results.

2. Use non-technical language. Your abstract should be understandable not just to few experts but to everybody in its wider constituency. This is particularly important if the results are relevant beyond your narrow specialism.

3. Never use in the abstract (or the title!) concepts that you define later in the paper.

4. Avoid complicated mathematical formulæ and references to bibliography. Titles and abstracts should stand on their own (hence no references) and be easily translatable to HTML (which, except in its XML manifestation, is hopeless with mathematical symbols).

Good reason 2 to read a paper:

Good introduction

OK, so you have tempted unwary punters to enter the shop. They browse around, finger the goods, enter the measuring booth: now you want them to surrender the plastic.

1. Explain very clearly what exactly is the problem that you are addressing in the paper – and why. Remember: you have worked on the wretched thing for months and months, for you the motivation is clear. But you are trying to persuade people that perhaps never thought about your problem. Mathematical results are worthwhile not just because they interest you. It is a good idea to commence with broad-brush context, explaining why the theme of paper is relevant and important.

2. Sketch the main results of the paper. No need to seek the greatest generality, which often depends on mathematical machinery that you'll introduce later. It is better to explain clearly and accessibly your main ideas and methodology.

3. Use profusely examples, toy problems and figures.

4. Place everything into context. You aren't the first to consider this problem, neither the first to say something worthwhile about it.

5. Provide continuity with relevant previous works. Not everybody remembers everything in papers underlying your argument. And sending the hapless reader to consult five different references in each paragraph is anti-social.

Good reason 3 to read a paper:
Good presentation

There are a number of issues that improve presentation:

Structure. Plan in advance the structure of the paper: sections, subsections and their content. Structure should be logical. It should not follow specific personal meanderings of your research (i.e., your personal logic) but the logic of a person trying to assimilate this information. Remember: Research is nonlinear, presentation should be linear.

Sections should not be too long: information is absorbed better in digestible chunks. Your argument should be illustrated by examples, tables and figures. It is sometimes a good idea to use a single example to illustrate different stages of your argument.

New concepts should be introduced at the point where they make sense and where their relevance is clear from the context.

Notation. Much of the clarity of your argument depends on notation. It is a good policy to decide on notation in advance (think how messy it is to change it mid-stream) according to a number of general rules of the thumb:

- Consider the use of notation as an opportunity to clarify your argument and inform the reader, not just as a technical necessity.
- Don't invent your personal notation if agreed notation already exists in your subject.
- Keep to convention: thus, "for every $\varepsilon > 0$ there is a $\delta > 0$" (not the other way around), $\varepsilon > 0$ is small and $R > r$. Clever dickery is a silly distraction.
- Consistency in notation helps the reader to organise information. For example, **boldface** for vectors, UPPER CASE for matrices, $\mathbb{SHELL\ CAPITAL}$s for sets.
- Avoid a profusion of $_{\text{subscripts}}$, $^{\text{superscripts}}$, multiple indices, hats, inverted hats, tildes, double tildes, ... The fact that it can be done with LaTeX is not a sufficient reason to do it!
- You might have defined something-or-other on page 5. This doesn't mean that the reader will remember it on page 27. Sometimes it is a good policy to recall definitions and notation.
- Never, but never, overload notation!

Language. Papers are not written in Mathematese, they are written in English. Remember: language is the main tool to convey information. Endless bright ideas have been rendered obscure and impenetrable by poor language.

- First and foremost, even if your native language isn't English, avoid poor or careless linguistic presentation. Be sensitive to the language, its idiom and cadences.

- Presentation shouldn't be overly flowery or informal: this is not a paper in literary criticism and you are judged on your ideas and their clear presentation, not on linguistic virtuosity. The language should be clear, unambiguous and informative.

- Avoid like the plague the sort of lifeless formalism and dry linguistic economy that made Bourbaki books declared as torture under the European Human Rights Act.

- Occasional flash of lighthearted humour or informal lingo is fine. Mathematics stand-up style is not.

- Be verbose enough to be clear – yet concise enough to privilege your core mathematical argument over its presentation.

- Not using a spell-checker is major folly. Relying totally on a spell-checker is carelessness: no spell-checker will distinguish between "some" and "same".

Proofs & co. Your argument – in particular your own contribution to the subject – should be self-contained. Only trivial issues may be "left to the reader". If you present a statement, say, and promise a proof elsewhere, there is a good reason to believe that your paper should not have been written at the first place. Having said so, there is often the case to relegate gory technicalities to an appendix, while providing enough information in the body of the paper for the reader to understand the gist of your argument.

Once presenting material of others – as is sometimes necessary, whether to explain your narrative or place it in context – you may skip technicalities and unnecessary details or proofs (with proper reference).

Presenting your own results, it is a good strategy in long proofs or constructions to explain in advance the main chain of argument. You are not writing a detective story but a mathematical paper! Also, it is a good idea to explain exactly how your statements, definitions, theorems, proofs and numerical results differ (for better and for worse) from other work.

Your claims. Be careful to give credit where credit is due, refer generously to the work of others* and graciously acknowledge their help. This is not just a matter of basic decency or of long-term calculation to generate goodwill and avoid tit-for-tat spats. Clarity in referring to the work of others helps the reader to identify your own contribution to the subject.

Nobody expects you to change human knowledge as we know it in your paper. You are standing on the shoulders of giants: what was good for Isaac Newton is good enough for you. Never overstate your results: you may occasionally mislead the ignorant but experts will laugh with scorn.

Remember: Every mathematical innovation has limited scope. A new theory or result is typically good for something, inferior for something else. Being honest about the scope of your new results is basic academic duty, but it is also the right long-term strategy.

*Incidentally, referring to detail in a book, it is useful to provide a page.

Multimedia. A good, well-crafted picture is not worth a thousand words but it can admirably illustrate and clarify your argument.

- The correct way of displaying graphic information often calls for a great deal of imagination. It is not just displaying a number of curves on a graph in an obscure fashion. A picture should make a specific point and its presentation should be geared towards this.

- Make the picture easy to understand. Thus, avoid information overload. Thus, label precisely and concisely. Thus, avoid a plethora of solid, dashed, dotted, dot-dashed, dot-dash-dot-dotted etc. lines – you are not sending a Morse code message!. Thus, remember that your figure might be in the most vivid technicolour but the journal will be printed in monochrome and information might be lost. Thus, explicitly link figures to text and explain in words their significance.

- Too many figures spoil the argument: you are writing for mathematicians, not for the *Hello* magazine.

- There is no need for trivial figures: do you really think that plotting a linear function will add to your presentation?

- The practice of following the main body of the paper by endless figures and drawings is singularly unhelpful, often seen (often rightly) as bulking lightweight material with graphics and often disregarded by weary readers. Consider graphic information as an integral part of your information flow, rather than as an add-on.

TEXpertise. Write yours papers in LATEX, $\mathcal{AMS}$TEX or plain TEX. Not in MS Word unless you wish your paper to look clumsy and unprofessional.

- It is a good idea to use from the outset the class file of the relevant journal. Once the paper is accepted, this will make copy-editing easy but (more importantly from your point of view) prevent massive changes over which you might have little control and which can introduce unexpected errors.
- It is sound policy to use macros since this practice minimises typos and makes late changes safer. (Note that some journals frown upon this.)
- Unless you know exactly what you are doing, don't tinker with build-in parameters, like `\parindet`, `\parskip` or `\topfraction`: they are there for a reason.
- Avoid **too** *many* `font` CHANGES, funny SIZE letters and other kindergarten tricks.
- Clumsy, careless TEXing is obvious to the experts (i.e., most of us) and displays lack of respect to readers (and to referees!). Thus, avoid like a minor plague overfull or underfull `\hboxes` and `\vboxes`. Break equations sensitively in regard to both mathematical content and esthetics. Always typeset maths in maths style.
- Using BibTEX is good policy, both in minimising eventual effort and in making stylistic changes easier.
- Don't let a preoccupation with endless minutiae of TEXing become an obsession. With all its importance, typesetting is just a tool.

That extra professional touch. A carefully written paper is like roses for Valentine: it demonstrates to your readers (in particular, to your referees) that you care and that you respect them.

- You may write in English or in American, but never mix the two in a single paper. This is not just "behaviour" *vs* "behavior" or "maths" *vs* "math" but more subtle issues. For example, in OED English an abbreviation that ends with the same letter as the original word (Mr, Ms, Dr, Revd – but not Prof.) comes without full stop.

- There is also difference between UK and US maths presentation. Thus, in UK maths, constants are typeset in Roman (`\mathrm{}`):

  in Cambridge, Mass $e^{i\pi} = -1$, in real Cambridge $\mathrm{e}^{\mathrm{i}\pi} = -1$.

  After a while, this becomes second nature.

- Typically, figures are produced with MATLAB, MAPLE or MATHEMATICA. None of these uses the same font as TEX (`Computer Modern`) and some make the use of Greek or nonstandard letters difficult in captions. It is possible, though (with moderate difficulty) to tweak the `.eps` file, changing the font from `Helvetica` to `Times-Roman` and even generating Greek letters and nonstandard symbols with the `Symbol` font.

- In multi-author papers it is vital that the final version is written (or at least thoroughly copy edited) by a single person. You know, clumsy stitches show and they are ugly!

Good reason 4 to read a paper:

Relevance

To write or not to write? There exists natural temptation to take any half-decent piece of research and use it like soup concentrate: add huge amounts of water and publish it as several incremental papers. By the end of your PhD you'll have a longer publication list: enough to impress the careless or the unwary. But no professional will fall for it. Unless you have something both new and substantive to say, don't say it! Good reputations are built on good publications, not on numerous publications.

If more than a third of your paper is devoted to reviewing your former work, rather than to new results, you probably don't have enough 'meat' to justify a paper. And referees will notice it.

On the other hand, don't be intimidated by all the wonderful and hard-to-understand papers that you've read. Yes, you can understand your paper easily – but that's because you've been working on it for a long while. As long as it has substantive body of significant results, it will probably be an important (and difficult) contribution to scholarship and a good publication!

The paper mill

Every good paper should mature in oak. Have you finished a paper? Read it again. And again. And again. Correct it. Give it to colleagues for their reaction. Correct it again. Then let it mature in oak (or chipboard): sit in a filing cabinet for a week or two. Then read it again – you'll be amazed by the number of corrections! And then stop. There comes a point where you can't go on and on with corrections.

Off to Internet. Even before you send the paper to a journal, it is a good idea to display it on your (or your group's) website. You may also deposit it in one of subject-specific Internet archives, e.g. `ArXiv`.

You can let interested parties (or appropriate websites) know about your paper and send them the abstract – but unless you are completely confident that they will welcome it, don't jam their in-boxes with your source file. They will not thank you for this.

How to choose a journal? Not all journals are equal. In each and every discipline there exists an implicit pecking order. (And an explicit “impact factor”, for whatever it is worth.) It is surprisingly easy to publish in a mediocre journal but becomes increasingly more difficult the higher you go in respectability stakes.

It is a good idea to ask around and to form your own impressions: Where have you seen the best papers? Which journals are referenced more frequently? Have a look at editorial boards, both to discern quality and to identify editors with an expertise to handle your papers.

So, should you submit to the best journal? Not necessarily, unless your paper is of the highest quality – and it is good to consult experts on that, not just following your intuition. However, unless you feel confident to submit to a good journal, you shouldn’t have probably written the paper at the first place.

The "process" Once you've submitted a paper to a journal, it follows a set pattern:

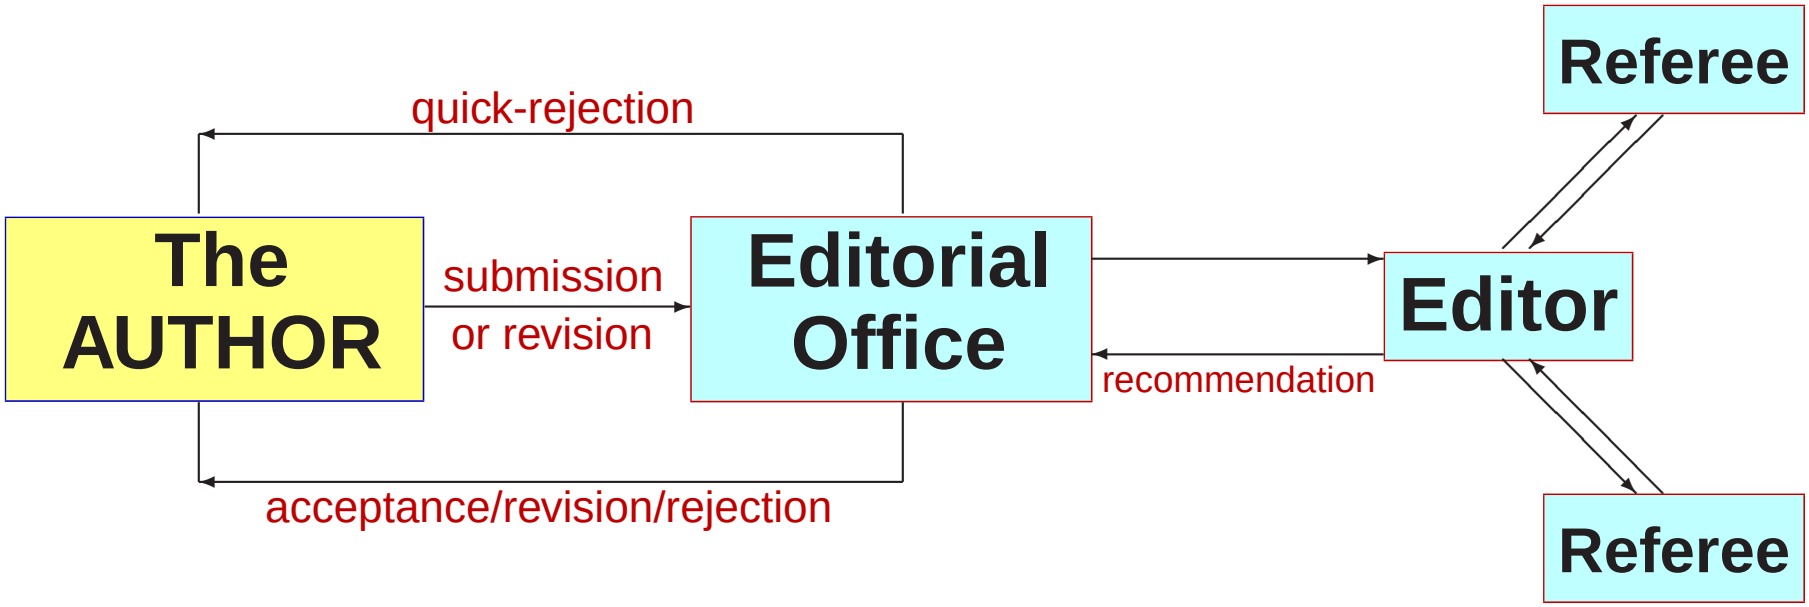

This procedure might have mild variations (e.g., the paper might be rejected with resubmission being encouraged). These days it is mostly conducted via the Internet, often through dedicated websites.

Rejection: Whatever I might tell you,

REJECTION = DEJECTION.

Your natural defence mechanisms will kick in: the referees were ignorant, the editor a prat, the editor in chief incompetent. Well, perhaps. . . . Yet, you should not jump into two extreme alternatives: neither send the paper at once to a different journal nor give up research altogether.

Read carefully the reports. Not everything in a referee report is gospel truth but usually referees make important points. Also, take it as an axiom: If the referee misunderstood you, this is your fault – next time explain better and more clearly. Write a new, revised version of the paper, taking on board all valid criticisms. This might require more research or computation – do it.

Never resubmit to the same journal or argue with editors. Submit to a different journal, but only once you are confident that your paper is better: the chances of it reaching the same referee again are non-negligible.

Revision: Resist two natural temptations, either to revise and reply by return of post or to put the job to the side (after all, you are busy with something new and exciting!), where it will languish until further notice.

Read very carefully the referee reports and the editorial letter. Address each issue therein and carefully prepare a revision.

Write a detailed record of how you've addressed the referees' comments, point-by-point. There is no need to agree to everything a referee said but there is an absolute need to explain (firmly yet respectfully) why you disagree. Enclose this record with your revision: it is immensely useful for editors and referees and will speed up the second refereeing round: the referees will need to read only the relevant bits, rather than to read the entire paper afresh.*

*And you don't want them to read it afresh anyway, because they will find more to disagree with. . . .

Acceptance: It happens very rarely that a paper is accepted at once in a quality journal but eventually, hopefully, after a round or two of revision, you'll receive "I am glad to let you know" letter.

Your labour is not over yet.

- Update your CV and publications' list.
- You'll need to provide a 'clean" source file and all other files that you've used: class files, macros, graphics,.... Now, if you were clever and used the journal's house style and class files, this should be a painless exercise. If you need to convert to the house style you deserve all the extra grief.
- You'll also need to sign and send a copyright form: this is the norm, don't argue.
- Some top journals (*Nature, Proc. Royal Soc.,* ...) insist on a media embargo until the publication day. Thus, you'll be unable to sell the story to *Daily Star.* It is tough being an academic....
- Once your paper has been processed, possibly reset, possibly copy-edited, you'll receive galley proofs, together with a stern letter telling you to correct them and return within a couple of days. If you've sent a clean source file in the house style, life is easy: just read quickly, answer few queries and return the lot. However, if the paper has been reset (probably introducing fresh errors) you'll have your work cut out.

And then, few months later, you'll open a fresh, nice-smelling volume and here is your paper smiling at you, and you'll realise that

Happiness is a warm paper!

References:

P. Halmos, "How to write mathematics", *L'Enseignement Mathematique* **16** (1970), 123–152.

N.J. Higham, *Handbook of Writing for the Mathematical Sciences,* SIAM, Philadelphia PA (1998).

S.B. Stechkin, " Как писать работы ", *Fundamental & Appld Maths* **4** (1997), 1261–1265.

N.E. Steenrod, *How to Write Mathematics,* Amer. Math. Soc., Providence, RI (1983).