How to write a frequently cited article

Henk van der Vorst

In August 2000 Henk van der Vorst received an email from the Institute for Scientific Information (ISI) about an article he published in 1992 in a SIAM journal: "Our analysis of high-impact papers in mathematics indicates that this paper has been cited 379 times to date, making it the most-cited mathematics paper of the last decade". The article introduced the Bi-CGSTAB algorithm. The ISI report gave rise to various newspaper articles. What made Van der Vorst's article the most-cited mathematics paper? Here the author reveals the tricks of the trade. (This article was previously published in the Dutch journal Nieuw Archief voor Wiskunde in March 2003).

In science it is essential to cite the sources upon which one's own contribution is based, so that the relationship between new and existing work and the authorship of presented results are clear. Usually these sources are provided in a separate bibliography at the end of the article. The listing of a publication in the bibliography is counted as one citation, no matter how often the publication in question is actually mentioned in the paper. If a citation refers to one's own work, this is called a self-citation.

The number of citations of a particular publication, counted over a certain period of time, supposedly measures the relative importance of the cited publication. After all, if a publication is cited it indicates that someone has not just thought it worth the trouble to read it: the cited publication has had value in placing the new work in the right context, or has even helped make the new work possible. In this modern age of mass scientific production, citation counts are a popular tool for measuring the effectiveness of a scientific group, or even an individual scientist. Citation behaviour itself is even the subject of scientific research.

The American Institute for Science Information (ISI) is world-leader and supplies its data to other institutions as a base for further research. The ISI also publishes the renowned Science Citation Index (SCI), which incidentally was created to find out where certain works are cited, and not to count citations as such. All the same, in December 2003 I was able to find out via the SCI that six of my publications were cited frequently (my standard for frequent citations being more than 100).

1 Measuring success and rankings

Citation counts are an obvious way of measuring the success of a researcher or even a whole institute. Managers and journalists like to use them because it saves them the trouble of making a more in-depth and subtle evaluation. Only recently, Dutch research groups were ranked by citation scores in Elsevier Magazine (March 16, 2002). The magazine examined what percentage of the articles of each institute belonged to the top 10% of the world's most cited articles in the relevant scientific field. As it turned out, mathematicians from Utrecht scored higher than their colleagues elsewhere in the country. This is of course great if you work in Utrecht, as I do, but what does it actually mean?

The larger the institute, the more it contributes – relatively speaking – to the overall production of articles and therefore the more that 10% is actually determined by that institute. If there were only one institute in the world, that institute would automatically have 10% of its articles in the top 10% of most cited articles. The larger the institute, the closer it should be to that 10%. If a large institute scored significantly lower than this standard it would mainly exist of nincompoops. Utrecht is above the standard, and moreover is the largest institute in The Netherlands, so on the whole it is a little more brilliant than the rest.

In mathematics, citation cultures differ enormously from subdiscipline to subdiscipline. Within statistics, operations research and numerical mathematics, citation scores are much higher than in pure algebra or symplectic geometry, for example. Below, I will elaborate on this. As it happens, it makes an enormous difference if results in a scientific field are of direct use outside mathematics, or mainly impact on colleagues. Statisticians and numerical analysts generally publish more than pure mathematicians, so I wouldn't be surprised if the 10% of most-cited articles are crowded with highly applicable mathematics. Utrecht has a relatively large production in the field of applied mathematics and can be expected to score higher than institutes that depend mostly on output in pure mathematics.

2 Citation scores

How about the citation scores in other disciplines? Both the second and third most-cited mathematics articles in the nineties were articles in the field of statistics. In contrast, in August 2000 Wiles' famous article (Annals of Mathematics, 1995) boasted fewer than 100 citations [1]. Within my own discipline my citation score was less striking: the total citation scores (measured in early 2001) were about 1100 for an article by Saad and Schultz (1986, GMRES-method), Peter Sonneveld's article on CGS (1989) scored 338 times and Freud and Nachtigal's work (1991, QMR method) received 215 citations. Note that the first two articles did not appear in the 1990's, which was a condition for ISI's hitlist. To get an impression of citation scores inspect the SCI or surf to ISIHighlyCited.com. On those pages the ISI gives all kinds of information about 'Highly Cited Researchers'. Via the NEC CiteSeer: http://citeseer.nj.nec.com you can find out how often articles in electronically available documents in institutional databases of Computer Science are cited. Because these institutes are often linked to mathematical institutes they include many mathematical articles. The system is far from perfect, but gives a fairly good idea of where and how work is cited: it is possible to view the five lines or so of the article in which the citation takes place.

On top of this, CiteSeer shows a citation histogram for every scientist that is cited. In it, all citations of an article are set against the year of the publication − very instructive for getting an impression of the most fruitful years of the researcher.

Back to the official SCI. In mid June 2002, Wiles' article totalled 147 citations, and for the possibly even more important article of Taylor and Wiles on ring-theoretic properties of Hecke algebras I found 75 citations. To put it bluntly, the top of pure mathematics didn't even make it to the ground level of numerical linear algebra. Outside mathematics, on the other hand, it is possible to generate much higher scores. Top articles in chemistry or medical sciences can easily obtain scores of 30 a month, albeit during a relatively short period (1 to 2 years). For an outsider it would be very easy to conclude that mathematics is a relatively unimportant science.

3 Citation examples

Before I continue to explain how to deal with citations and citation behaviour, it is useful to explicitly point out that a citation, i.e., inclusion in a list of citations, just means that the cited publication is mentioned somewhere in the text. This can be done in many different ways, and below are some real-life examples (the citation numbers are the original numbers):

- *In an effort to overcome this, in many areas a suitable stochastic model is employed to describe the uncertainty of the data, see e.g., [4,8,10,11,20-23,27].*
- *The patch test is conducted for the plate problems described in [16,28].*
- *It turns out that GMRES [13] took roughly twice as much computing time than BiCGSTAB [16] in all our test-cases.*
- *Although the perturbation technique is often used in engineering, it is usually applied for a finite number of random*

The subject of the much cited Bi-CGSTAB article

The cooling of baby brains, ocean currents, the force distribution on a launch missile, the contamination spread in groundwater and the surgical cutting in a patient's body are all examples of problems which after modelling lead to very large systems of linear equations. These days 'very large' means in the order of one million to one billion unknowns. Many such systems can be solved quickly on a computer using iterative algorithms. Amongst the most efficient algorithms are the so-called Krylov-methods, which, among others, comprise Conjugate Gradients, GMRES and Bi-CGSTAB. These methods roughly work as follows.
Write the system to be solved as $Ax = b$, with A a real, $n \times n$ non-singular matrix and b given. Starting with an initial guess x_0 for the solution x, a search space is constructed, in which a new x_{i+1} is determined, by extending the available search space with the residue $r_i = b - Ax_i$. After this extension the search space, denoted by K^{i+1}, is spanned by the $i + 1$ vectors $r_0, r_1, \ldots, r_i$. This space is a so-called Krylov space. The new iterative is determined, for example, by demanding that $r_{i+1} \perp K^{i+1}$. If A is symmetric and positive, this in principle leads to the Conjugate Gradients method. Another suitable criterion is the requirement that the Euclidian length of r_{i+1} is minimized over all possible solutions in K^{i+1}. This leads to the GMRES method.
For non-symmetric A, GMRES is an expensive method, but it can be shown that useful iterates can be constructed more cheaply by keeping track of two spaces: the Krylov space for $Ax = b$, and another Krylov space for the otherwise uninteresting system $A^T y = b$. This leads to the Bi-CG method. Sonneveld (Technical University Delft) has shown how the computations with A^T can be replaced by computations with A. This makes it possible to determine approximate solutions in search spaces twice as large, at almost the same computational cost.
Following this principle, the Bi-CGSTAB method cheaply constructs very good approximate solutions in that double dimensional search space. Although there are now variants which are even better, Bi-CGSTAB has remained very popular because of its great simplicity. A computer code for the method itself does not take up more than about twenty lines.

parameters, see e.g., [17], and often without rigorous theory.

- *The classical solution given by Chrisfield [9] is used as the reference solution.*
- *In 1992 Bettess published the first book on infinite elements [21], which set forth the state of the art at that time.*
- *The fundamental paper [4] of Dörfler for the Poisson equation shows...*
- *The above equation is very similar to the form proposed by McMeeking and Rice [11].*

It should be clear that not all citations are the same and that citation counts should be approached with caution. Furthermore, the value of a citation is partly determined by the quality of the medium in which the new work is published. A citation in an article in the Northern Territory Mathematical Journal, that publishes everything without peer review, certainly represents less value than a citation in an article published in the prestigious and extensively refereed Annals of Mathematics. Therefore, the ISI only counts citations in journals that fulfil a number of quality conditions. However, even within this list lower-quality journals exist next to better ones.

Despite these objections, on the whole it does of course mean something if an article is cited frequently, just as it means

something if an article is never cited, self-citations always excluded. Frequent citations can contribute to a comfortable feeling: the article in question has been (at least of some) use to many researchers, which makes one's own contribution to science a little more visible. Neither more nor less than that. Wiles' publication, which definitively confirmed Fermat's last theorem, has led to (far) fewer citations than my own article, in which I proposed a slightly faster method for solving linear systems. Wiles' publication solved an unbelievably difficult problem, and I don't expect to be accused of false modesty when I say his work is of a higher order. His work deservedly received enormous attention worldwide, even in the regular press; my work enabled a number of other scientists to carry out computations more efficiently, which enhanced their possibilities (somewhat). Again, citation counts can't be put into enough perspective.

4 Citation behaviour

I now have set an elaborate background for some pointers regarding the ambitious title of this article. I will start by saying that you must hope that your brainchildren will be cited a lot, but that it is also possible to do something to have them receive more attention (=citations). You should do this appropriately though, and avoid forcing citations to your own work upon others. I am thinking about the following, not in the least hypothetical, situations.

You can encourage PhD students to cite the work of their Great Helmsman more than is strictly necessary. This is reprehensible because it contributes negatively to the scientific education of young scientists. It may lead to a few extra citations, but also in the long term to a dubious reputation.

It is also tempting to suggest citations to one's own work, where it is not strictly necessary, under the cover of anonymity while acting as a referee. It happens and often it has the desired effect, since many authors are inclined to satisfy their referees. I find this behaviour aggressive and editors should be sure to prevent it. You could also consider having your work generously cited by a contracted scientist from a poor country (this also happens) but that is a dead end street as well, and further comments are not really necessary. It is very suspicious if an author is cited noticeably often by an unknown source.

In brief, every self-influence of the citation profile is unethical and, in the shorter or longer term, will backfire against the perpetrator.

5 Attracting attention

I now come to the more positive steps authors can take to give their work the attention it deserves. Nowadays there is an abundance of publications, and certain articles easily go unnoticed unnecessarily. Of course a real breakthrough gets known quickly; it is usually not necessary to give it extra publicity in scientific circles. With less spectacular steps forward however this certainly can do no harm, but it has to be done with due caution.

Let the following anecdote serve as an example of how not to proceed. In an invited talk at a big conference I once explained how lattice points should be numbered so that certain problems can be solved faster numerically, and moreover in parallel (the accompanying article appeared later in a well-known journal). My description was recursive so that the numbering was defined directly for all dimensions. In the evening after the talk I explained to a small group how the trick worked for two-dimensional lattices. Two listeners started working with this almost immediately and published on it; since then this numbering is often referred to as the 'vdv-ordering' in the literature, usually with a citation to their article. In this article the discussion with me is acknowledged, but my own publication is not cited. This was not intentional; they simply had not understood that my talk related to the same numbering as in the example during the evening discussion. This of course illustrates the (in)accessibility of my formal recursive description.

6 Theorem-proof-remarks

The lesson to be learned from the above is: first of all, try to write clearly and accessibly. Do not hide your ideas behind notation, but accompany them with sufficient explanations and examples. While writing, it can be helpful to envisage your audience as if they are reading over your shoulder, and to imagine their possible reactions.

I was lucky enough to have had a critical supervisor and with many an article it has helped me to ask myself: "What would he say about this sentence were he to read it?". Do not let the reader guess your intentions, and let what is important clearly come to the fore. I regularly publish on algorithms and I try to describe them in such a way that it is simple for the reader to evaluate them on a computer. Preferably I would also put such an algorithm in a separate box. If it is necessary to give a proof, it would get a less prominent place unless the proof itself were very instructive. In fact, I try to keep the person who has started reading my article captivated for as long as possible. If I can reach that goal by shifting the mathematical depth backwards, I will happily do so.

Of course, whatever you write has to be sound and verifiable, but it does not have to be presented via the classical theorem-proof-remark approach. Personally I prefer an accessible formulation of the problem first, and after that a motivating example, only then followed by a theorem (or algorithm), which does not seem to appear out of nowhere after the introduction and example. I then prefer some more comments and remarks before providing the proof. Usually I finish with a few examples with applications of the theorem (the algorithm).

There exists literature on writing scientific publications, and personally I enjoyed using Higham's book [2].

7 Promoting personal work

Of course I distinguish between articles that I am really proud of, in the sense that they contain truly new ideas, and those in which I elaborate on previous work. The latter are mostly structured around highlights, serve to demonstrate the relevance of the main articles, and to extend them. Often these are papers presented at conferences which are included in proceedings.

The other things I use to further draw attention to my personal work are follow-up publications, lectures, the internet and email. In the following I will comment on each of these.

Publications

Follow-up publications can be used to enhance the profile of previously presented results, provided of course that this is useful. Personally I always try to use self-citations sparingly. There is a penalty on exaggeration. Citation evaluations often divulge the percentage of self-citations. Keep in mind that self-citations are not included in your citation score. Realize that a high self-citation score may well make you look as if you have got blinkers on (the closed research culture), and be aware of this whenever you refer to prior work of your own.

Writing review articles is especially effective, but this usually happens by invitation only. Also very effective might be the writing of a book. In a book, the author has the opportunity to pursue a subject in greater depth than is allowed in an article. The author can choose to target a broader audience than that of a journal. I have used this possibility to (co-)write a book in which a number of numerical methods were described at such a level that they could be used by mathematicians without prior knowledge of numerical mathematics. In my opinion this has contributed enormously to the popularity of the algorithms, of which "my" often cited Bi-CGSTAB method is one.

This brings me to a "negative" aspect of popularity. It sometimes happens that a product of a manufacturer becomes so popular that the brand name becomes a generic name; an example of this is Aspirin. The brand name is no longer protected and other producers are allowed to use it. The same can happen to the author of an algorithm.

When solving sets of equations by elimination hardly anyone still cites Gauss. The Bi-CGSTAB algorithm has now been included in Matlab and not everyone still refers to its origin. This leads to a natural limitation on the citation scores of the source.

Talks

I see talks as an excellent means of me selling my work as a kind of salesman. I always hope that some of my audience decide to examine my work more closely. Please read Klaas Landsman's excellent article in NAW on how to give an effective talk [3]: I do not think there is anything to add to that. Recently I tried out a talk on some student volunteers, asking them to "shoot at will". This has given me plenty of useful information.

I always finish my talks with a referral to my web pages, where the audience can find more information. This brings me to the third item.

The Web

The power and might of the web are still largely underestimated in our circles. Many scientists search for information on the web and hardly frequent libraries anymore. In 2001 Steve Lawrence published a nice article in Nature called 'Online or Invisible?'[4]. As well as all kinds of nice statistics, his research into citations of 119,924 publications in total (included in Conference proceedings), was striking. The average number of citations of electronically available publications (i.e., via the web) was 7.03 versus only 2.74 for publications that were exclusively available in print. Important lesson: make sure that your publications are readily available on the web, preferably in a directly printable or readable format. Nowadays most publishers allow this.

Personally I also use the web to make my research fields more accessible via so-called *Lecture Notes*. In these Lecture Notes I can take explanations and illustrations as far as I want to; at least further than in regular publications. Of course I also use them to clarify my work more comprehensively for a broader audience. The Lecture Notes are fairly popular; they are downloaded from my web pages about 45 times a month. I have a strong feeling that this has helped make my work better known and more accessible, which certainly has pushed my citation scores upwards. In writing the Lecture Notes I always keep a good eye on quality. Lack of it is punished ruthlessly in the scientific world. Be extremely careful with what you make available on the internet. It cannot be withdrawn: after a (short) while electronic copies of your work seem to surface everywhere. For this reason, do not put rough drafts on the web, but restrict yourself to contributions you fully endorse.

I keep a regular eye on which articles are downloaded from my web pages and how often this happens (in our system group we receive up-to-date daily reports). This way I see where the demand is and which subjects might warrant further attention. Initially it seemed to be a bit of a waste of time to write review articles ('State-of-the-art papers'), but when I noticed that they were relatively popular this increased my motivation to write some more (although at the request of editors). It can be of tremendous help to others in opening up a large part of a field and that is why review articles are usually cited frequently.

Because of their popularity, I regularly expand the Lecture Notes with new material, or make corrections that readers suggest to me. After learning that the notes are sometimes used as educational materials I have started to include exercises. Recently this has resulted in a book [5], but a limited preliminary version of it still remains on the web page as a kind of tempter. Tip: make sure that your publications on the web do not become too long. I have divided my Lecture Notes into three parts to prevent irritation while downloading.

Email

I have already mentioned reader feedback; this happens via verbal contact at workshops and conferences, but primarily via email. Another possibility to make your

work more widely known is to circulate postscript or pdf files by email. Personally I would be reluctant to do this. I find it quite annoying to be sent large unsolicited files. It is much nicer to be notified by colleagues as to where you can download new work that might be of relevance to you. You might still circulate such succinct notices selectively; but personally I always find it pleasant if colleagues close their emails to me with their address and the URL of their web pages, and leave it at that.

8 Epilogue

"How to write a frequently cited article?" I would rather interpret the title as being *"How to make sure my article will receive attention?"* It is such a pity if you come up with new ideas and nobody finds out about them. In this article I have explained which line(s) of action I would personally follow. Of course I do not have a formula to write a frequently cited article and therefore I am curious to await this article's fate.

References

[1] Howard Elman, *What's Hot? Study Points to Iterative Methods for Nonsymmetric Linear Systems*, SIAM News **34** (2001), 2.
[2] Nicholas J. Higham, *Handbook of Writing for the Mathematical Sciences*, Second Ed., (SIAM Philadelphia 1998).
[3] Klaas Landsman, *How to give a mathematical talk* (in Dutch), Nieuw Archief voor Wiskunde, Vijfde Serie, **2** (2001), 351-355.
[4] Steve Lawrence, *Online or Invisible?*, Nature, **411** (2001), 521.
[5] Henk van der Vorst, *Iterative Krylov Methods for Large Linear Systems*, (Cambridge University Press, Cambridge 2003).

Mathematical Institute, Utrecht University, Budapestlaan 6, NL-3584 CD Utrecht, the Netherlands
E-mail: vorst@math.ruu.nl

Translation by Maaike Wienk