

On improving the state of the art.

(A somewhat open letter to dr.Martin Rem.)

25th of October 1978

Dear Martin,

yesterday evening you raised the very broad question what directions computing science ought to take, suggesting that a well-considered answer to that question would be very useful. (You strengthened your point by the observation that of all the research papers we see, the overwhelming majority is very disappointing and that many research efforts seem misguided. And I agreed.) Since you raised the question (which I left unanswered) I have been thinking about it. While doing so I learned that, even if I were able to answer it more or less satisfactorily --one of these days I might-- , I was absolutely uncertain about my answer's "format". To open the discussion I have chosen the form of "a somewhat open letter", a form that seems to give me the latitude I need; please, don't regard this letter as my final answer to your question, as it is no more than a first exploration.

The first rule of the game is that in this discussion we must regard science as if it were an end in itself. We both know that it never is, yet we must do "as if", because, whatever the ends may be, you will never get any agreement on them. (At an American computer conference a few years ago, the keynote speaker explained that computing science was very important because only computing science would enable the USA to maintain its technological supremacy in the face of the Japanese threat. Some years later, at a computer conference in Tokyo, the keynote speaker explained that computing science was very important, because only computing science could free Japan from the American technological supremacy. So there you are!)

Your direct motivation was a profound dissatisfaction with the output of the computing science community, as it appears on the pages of our journals, and I agreed. Can we analyze the nature of that dissatisfaction? In a final answer to the question what directions computing science should take, such an analysis would probably be out of place; among ourselves a candid analysis could be illuminating.

1) There are papers of such peripheral interest that we just cannot read them for lack of knowledge of the subject matter. It may cover a very special practical field --say: an aspect peculiar to the processing of seismic data-- or a very special theoretical field --say: an obscure corner of automata theory-- . Usually we give such subcultures the benefit of the doubt. The problem is not serious in the sense that it takes very little time to come to the decision to skip such an article; my natural reaction is that, probably, the paper has been published in the wrong journal- It becomes a serious problem, however, when this skipping device is so frequently activated that whole issues of a journal remain unread: then either the editorial policy has failed, or the field is so incoherent that all topics are "peripheral ones". (Somewhere in the course of this analysis we should understand why I skip almost all articles in the Journal of the ACM. I still hesitate to accept my laziness or ignorance as full explanation of this disturbing fact. I should return to this later.)

2) There are papers that, even if they were well-written, are not worth studying because you could hardly learn anything worthwhile from them --say

data-flow analysis for optimization, as complicated by all the peculiarities of PL/I-- . Competent development often leads to such papers: vital for the man who has to solve exactly the same problem, but very little helpful everywhere else. Such papers, although not worthwhile to study, we should accept --up to a point, at least-- : they contribute to what is euphemistically described as "the accumulation of knowledge". They also contribute to its fragmentation! To counteract that fragmentation is one of the major calls of scientific research --a call that is easy to hear, but sometimes hard to obey: look at what happened to Nuclear Physics!-- . There is a fair chance that in computing science we should be more alert than in some other disciplines --such as, for instance, mathematics-- because confusing science and technology, and mistaking development for research seems in our field rather wide-spread, a fact for which "the physical machine" is undoubtedly responsible. You remember the colleague who, in a sad letter, confessed to me that, at last, his department too had bought a research computer, thus almost ensuring that its output at large would fall down to the general level of mediocrity? (If you don't I shall tell you.) God save us from the "research vehicle" --SRC terminology (UK)-- far, if you are not very careful, you will find yourself fully absorbed by "driving" it. (This, by the way, also explains the almost certain scientific sterility of computing science research establishments primarily manned by electronic engineers, when building something is the best they can think of to do.)

3) There are papers that are no more than an artfully disguised report on a not surprising failure. The paper starts very promising --as a matter of fact: it promises more than it can give-- and you are intrigued, because the claims have surprised you. As the paper continues the author's tools, techniques, and notations become more and more baroque, gradually you lose faith, and eventually you are given the meagre results with the difficulties encountered as "areas suggested for further research". Peter Naur has taught us a technique for identifying such papers quickly: start reading backwards. So we know how to deal with them. In a case like that I always blame the editor for having accepted the paper. Much more serious than the poor publication is, of course, that the author went into such an obviously blind alley. It probably means that, when faced with a problem, many people just attack it with the tools they happen to have, without giving any consideration to their adequacy. To teach people such an evaluation is one of the hardest tasks of education, it is good to be reminded of its necessity, but I don't think that the situation is peculiar to computing science.

4) There are ugly papers to which I have methodological objections.

4a) Lousy use of English, that makes the paper ranging from empty to ambiguous. (In the second case you get the feeling that by rewriting it you could perhaps salvage parts of it.) Such poor use of language puts me off very rapidly. It makes the reading very much harder and the effort/gain ratio goes up very quickly, also because the gain goes down: it is impossible to carve a gem with a blunt instrument. Careful use of language is not a dispensable luxury but a vital matter, often deciding between success and failure. It is in this connection worth noticing that in the Comm.ACM the papers on data bases, AI and information retrieval, and social implications are of a markedly lower linguistic quality than the others. A subcategory are the papers with a shoddy motivation which a more careful linguistic analysis would have revealed as a slogan (readability, understandability, natural and intuitive, programming for the non-programmers, etc.).

- 4b) Pompous notations. This makes the paper boring to read and shows a lack of good taste. Sometimes the paper can be saved by translating it formula by formula. Very irritating are such papers when they are essentially an exercise in sterility, such as "the precise formal definition of XXXX" with which subsequently nothing is done because the formalism is absolutely unmanageable.
- 4c) Clumsy mathematics.
- 4d) Intertwining of concerns that you have learned to separate. This leads to irritation, for the paper is always confusing, unnecessarily long and clumsy (and sometimes nonsensical).
- 4e) Use of misleading metaphors, such as anthropomorphic terminology. (Most diagrams and pictures, needless to say, fall in the same category.)

I think that the above --although certainly not complete-- is enough of it. Part of the problem is certainly not a scientific one. Many papers are offered by authors who will never become scientists and should not have been manoeuvred into positions in which they are invited to behave as if; this is not a scientific problem. Also, publishing is a business in its own right, not obliged to care if the material published is junk, as long as it sells --some publishers, by the way, do care!-- and this explains (I refuse to say: justifies!) the editorial fear of frightening the readership by publishing material the appreciation of which would require thinking or, even, that could really educate its readers. Educating people is a very risky business, as most people resist it and object to it. This is a dilemma, and many politicians believe that the safest way out of it is fraud: you talk about your activities as if they are very educational because that sounds very respectable, while in fact actually educating people is what you most carefully avoid. (Observe, for instance, our dear current Minister of Education and Sciences.) So much for the business of publishing. That is not a scientific problem either.

* * *

(Continued on 9th of November 1978.)

In the meantime two requests to give an Invited Speech (Munich and London, both in September 1979) have gathered on my desk, and while thinking about them --why again I had been invited, whether I should accept or not, what kind of speech I would deliver when I accepted, etc.-- I suddenly remembered my uncompleted letter to you, and decided that this was a most appropriate occasion for its completion.

The two invitations brought back the memory of my first Invited Speech when I had about your age and was hardly known (IFIP 62, in Munich), and I decided that times indeed have changed. Very much to my surprise that talk created a sensation, and it did so because it was unusual. (And that was largely due to my lack of experience and my unfamiliarity with the customs of the international scientific scene: I was just naive and honest and tried to explain, because it seemed to be glossed over, why programming was difficult.) It was a great success, it was immediately accepted as refreshing, as a new and promising way of talking about an important aspect of our field of which it was almost unanimously accepted that, indeed, it had been neglected. I can only conclude that at that time I had a very flexible audience. Six years later the scene had already changed: my --only explanatory!-- 1968 Letter to the Editor of the Comm.ACM evoked primarily abusive letters. And now, another ten years later, we seem to have reached the stage in which some referees of the Comm.ACM only deem acceptable things they have already seen before. Clearly the great difference is that in the meantime a computing science establishment has emerged, which, like all establishments, is very

conservative and perhaps even frightened by the thought of change. (This very morning I read a long paper by A. Endres (IBM, Böblingen) who seems to suggest that our profession has now evolved to the stage in which the design of new programs has become less important than learning to live with the existing ones. Mills's Software Panel had a section in the same vein. The usual, economic justification --referring to our "investments"-- has never been able to convince me; pandering to the fears of a conservative establishment is a much more likely explanation. The latter is at least one that I can understand.) The existence of the computing science establishment could also explain why all invited speakers --in Munich I was envisaged as a member of a quartet, in London as a member of a trio-- were on the average almost of my generation instead of at least fifteen years younger. If today a youngster like yourself would try to present something refreshing, he would be in danger of being shouted away from the platform, and in order to get things moving you now need people with a firmly established international reputation. This, of course, is very sad. But it gives an indication of what should be done.

Computing practice on the average has dug itself into a hole of which it does not see how to come out again; by way of consolation it deceives itself by praising the hole for the protection it gives. Computer science --much more strongly influenced by the prevailing prejudices than most of its practitioners are able to perceive or willing to admit-- has got stuck in its wake. Clearly some refreshment is urgently needed. (Remember the rule we have already agreed upon that in this letter science is to be regarded as if it were an end in itself.)

How do we refresh? Well, by the choice of unpopular targets and the choice of unpopular means. After that the combination of talent, good taste, and faith in the worthiness of your cause is almost certain to do the rest. Shall we try --in an effort to be constructive!-- a few unpopular suggestions?

- 1) The use of computers should be avoided wherever possible; we don't object to computer usage by scientists being a component in the determination of their scientific status, but the sign of the component should be inverted. (This is guaranteed to make almost all high-energy physicists your declared enemies and, therefore, has tremendous advantages. Another advantage is that some computer manufacturers may not like you either.)
- 2) The purpose of scientific research is to gain understanding, an obligation no person can delegate without giving up. (That, together with (1) does away with the "real verifiers", the "automatic theorem provers/checkers" and most of the other AI activities. This clears the sky tremendously.)
- 3) As long as the problem of quality control for LSI chips is unsolved, technology's current promises should be regarded as lies; the concurrency bandwagon should only be joined with great hesitation --the ability to do something no one else has been able to do might be used as excuse-- ; all large scale applications of LSI chips are by definition highly suspect. (That does away with "personal computing", "home computers", "the information society" and all that jazz. Very effective.)
- 4) Argue that complexity is evil, and any system needing yards of documentation is no longer fit for use. (This will make you very unpopular with all those who have just followed the 20-26 hours Advanced Systems Inc. course on how to interpret MVS/JES 3 ABEND core dumps.)

- 5) Argue that most of the problems in industrial system development are due to an unfortunate interface between managers and technicians, an interface that has not been designed but grew historically at a time that the nature of system development was not well understood. (This is nicely alarming; it should be equally alarming for the managers and the technicians, by virtue of the symmetric way in which it has been phrased. I promise that the managers will dislike you for this much more strongly than the technicians, for the managers feel themselves much more threatened. On the University Campus this attitude is particularly effective, as it is guaranteed to isolate you fully and effectively from the Department of Management Science, a gain that should not be ignored.)
- 6) Stress the importance of perfect mastery of one's native tongue, and correct your countrymen when they speak a lousy Dutch. This is guaranteed to make you a number of enemies. (If, upon correction, someone reacts with "Why care, you understand me, don't you?" ridicule him ruthlessly. If necessary, quote Wittgenstein.)
- 7) Stress the importance of mathematics and agility in formal manipulations in general. (This will, with the exception of the mathematicians, estrange you from the rest.)
- 8) Demonstrate how clumsily the average mathematician writes, thinks and manipulates; to add insult to injury you might explain that up till now he could get away with those techniques because he has ignored the intellectual challenge embodied by powerful computing equipment. (Then you have estranged yourself from the mathematicians as well, and at last you are free!)

Well, how was that for a first effort to be abusive? Not bad: when you think about these eight points a little bit further, I guess that each, or almost each of them could inspire a refreshing and potentially very valid piece of research. (I have most hesitations about (5): it seems to me so obvious that surprising insight from further investigations is probably too much to hope for. Its political usefulness seems unquestionable, its scientific potential seems meagre.)

Dear Martin, I think that by now my letter is long enough: after all I did not intend more than a first exploration. Today Ria visits her mother in Amsterdam, which implies that I am responsible for feeding the dogs. It is close to 5 p.m. and particularly Bamse --who, as is easily seen, is always willing to eat-- is definitely getting restless.

It was a pleasure to write to you. Thank you! With my greetings and warmest regards,

yours ever,

Edsger.

Plataanstraat 5
5671 AL Nuenen
The Netherlands

prof.dr.Edsger W.Dijkstra
Burroughs Research Fellow

PS. I now realize that in my eight-point program I have forgotten to insult Electrical Engineering. Sorry!

EWD