

Copyright Notice

The following manuscript

EWD 611: On the fact that the Atlantic Ocean has two sides
is held in copyright by Springer-Verlag New York.

The manuscript was published as pages 268–276 of

Edsger W. Dijkstra, *Selected Writings on Computing: A Personal Perspective*,
Springer-Verlag, 1982. ISBN 0-387-90652-5.

**Reproduced with permission from Springer-Verlag New York.
Any further reproduction is strictly prohibited.**

On the fact that the Atlantic Ocean has two sides.

Introduction and apology.

This is an open letter to my co-members of the IFIP Working Group 2.3 on "Programming Methodology". Among my writings thus far it will be an exception, because, up till now, it has been for me very rare to undertake a task of which I knew beforehand that I would not be able to do it well enough. The reason that, nevertheless, I have decided to undertake it is quite simple: it has to be done, and offhand I can think of no one else less unqualified to try to do so.

My subject should be very simple, for it is only the difference between the orientations of computing science at the two sides of the Atlantic Ocean. That there is a difference should not amaze us at all, for the Atlantic Ocean is very big. For a variety of reasons, however, this difference is a bit hard to discuss: the difference itself is no problem, but it becomes a problem when ignored or denied. It is a bit hard to discuss for about three reasons.

Firstly, we are comparing prevailing attitudes between continents. Everyone familiar with them is aware of the great diversity within each of them, and he knows that writing about "a European attitude" is as much writing about a literary fiction as writing about "an American attitude": in the kind of global comparison I feel forced to make, I simply have to do injustice to differences of continental significance only. I can only ask you to forgive me my gross oversimplifications, of which I am only too aware myself. Abstracting from the inhomogeneity of both of the continents, we can still observe considerable differences between the two continents, and those differences are the subject matter of this open letter.

Secondly, the difference between the Old and the New World has already been discussed so extensively, and by so many, that it is practically impossible to raise the subject without evoking all the cliché prejudices. And in this discussion we have to pay attention to the general cultural difference, to the different images of man and society, for they have a profound influence on computing science (much profounder on computing science than on a mere technical subject such as geology or medicine).

Thirdly, it is a subject that many people are a bit touchy about. Both continents have their inferiority complexes --overcompensated or not!-- and we are all "party" in the sense that we have been born at one side only! Fully aware of how firmly my roots are planted in Europe, I can only undertake this task with considerable trepidation, afraid as I am of failing to be fair and to do justice. (This fear of being unjust and thereby offensive has been so great that during the first years of my association with Burroughs I have subconsciously avoided comparing the two continents! Having just gathered my courage, I nearly lost it again when I received a letter from Jim Horning, to whom I had mailed a copy of my tripreport covering the last W.G.2.3 meeting. Jim wrote me "The analysis of the meeting in your tripreport is in substantial agreement with my own, although my report to the members wasn't quite as blunt." I was surprised: evidently my pen is sometimes sharper than intended or suspected.)

Whether we like it or not: it is a touchy subject. And that is exactly the reason why it is avoided, and why someone should bring it up. I became aware of this by a curious incident at our last meeting in St.Pierre-de-Chartreuse. After Mary Shaw's presentation a lengthy and, in its way, lively discussion ensued, but it was a very curious one. With the exception of two short questions for clarification posed by European participants, the discussion was entirely an American affair, and it was noteworthy for the inadequacy with which it was carried out. Among the European participants witnessing this discussion, the overwhelming feeling was one of embarrassment. (Some younger ones could hardly believe their ears and voiced their amazement/indignation later in private by comments "Some have to learn it the hard way..." or "Is this 1976 or 1966?" and cruder ones.) The bitter point of the whole incident, however, was that none of us did what should have been done, that none of us interrupted by remarking that this did not seem an adequate way of discussing this topic. That is what would have happened in an unhampered scientific discussion!

In retrospect I have wondered about our silence, and I have blamed myself for it. My conclusion is that by the time that certain topics are becoming so painful to discuss as to paralyze scientific meetings, something has to be done about it. This is my effort.

Scales for comparison; general differences.

A very useful measure is --called after its inventor-- the "Buxton Index". John N. Buxton discovered that the most important one-dimensional scale along which persons are institutions to be compared, can be placed is the length of the period of time in the future for which a person or institution plans. This period, measured in years, gives the Buxton Index. For the little shopkeeper around the corner the Buxton Index is three-quarter, for a true Christian it is infinite, we marry with one near fifty, most larger companies have one of about five, most scientists have one between two and ten. (For a scientist it is hard to have a larger one: the future then becomes so hazy, that effective planning becomes an illusion.)

The great significance of the Buxton Index is not its depth, but its objectivity. The point is that when people with drastically different Buxton Indices have to cooperate while unaware of the concept of the Buxton Index, they tend to make moral accusations against each other. The man with the shorter Buxton Index accuses the other of neglect of duty, the man with the larger one accuses the other of shortsightedness. The notion of the Buxton Index takes the moral flavour away and enables people to discuss such differences among themselves dispassionately. There is nothing wrong with having different Buxton Indices! It takes many people to make a world. There is clearly no moral value attached to either a long or a short Buxton Index. It is a useful concept for dispassionate discussion.

In my own environment I have suffered from a relatively long Buxton Index --complete with accusations to and fro-- until the concept of the Buxton Index was brought to our attention. If, in the course of this discussion, I emerge as "very European", I think that among other things I do so on account of my large personal Buxton Index, because, on the average, the European Buxton Index seems to be larger than the American one. As an example I just mention the funding policy of the NSF and similar organizations --and it does not matter now whether we should regard this as cause or as symptom-- . The NSF policy states explicitly --and the need for the statement is significant-- that short-term goals at the expense of long-term concerns are not to be sponsored. Fine, but the majority of the research proposals aim at a tangible result within two or three years only. Personally I don't remember ever having seen a proposal for a grant beyond three years. The (to my taste) shortness of these periods has in the past been one of my main considerations for not joining the faculty of an American University, and as some of

them have tried hard enough to seduce me, I feel entitled to call the difference significant.

* * *

My first visit to the USA --in 1963-- was a shattering experience. (It was also frightening: I started with a few days all by myself in New York.) Of all memories from that visit, one is absolutely overpowering: for the first time in my life I was confronted with a civilization that did not give its scientists the automatic benefit of the doubt or the respect that I was used to. On that trip I learned the word "egg-head" as a truly untranslatable Americanism. (Untranslatability is always significant!) I was shocked to see how intellectuals could be --as it were-- by definition suspect, and I remember that the feeling of uncertainty from which I saw my colleagues suffer, worried me very much. It was the first time in my life that I realized what difference it makes to be a citizen of a very small monarchy in which each professorial appointment is confirmed by Her Majesty our Queen. (Again we need not argue here, whether Her Majesty's involvement is symptom or cause of our scientists's spiritual independence and feelings of social security.)

The above captures the overwhelming impression of my first visit to the USA; the assumption that it refers to a significant difference seems, therefore, safe. My many subsequent visits to the USA gave me some opportunity to figure out what I had seen that first time. The questions are: how does science justify itself, why does a society tolerate scientists? The way in which these questions are answered has a deep influence on the scientist's behaviour, not only on the way in which he presents his results, but also on his way of working and his choice of topics. Traditionally there are ~~two~~ ways in which science can be justified, the Platonic and the pragmatic one. In the Platonic way --"l'art pour l'art"-- science justifies itself by its beauty and internal consistency, in the pragmatic way science is justified by the usefulness of its products. My overall impression is that along this scale --which is not entirely independent of the Buxton Index-- Europe, for better or for worse, is more Platonic, whereas the USA, and Canada to a lesser extent, are more pragmatic. (Most of you must have been confronted with my Pan-Academic prejudices, which are most definitely Platonic, and by now you may wonder how in the world I could join not only an industrial organization --industrial organizations by their charter being more pragmatic-- but even an American one. But the answer is quite simple: in computing science the conflict need not

exists --and that is what makes the subject so fascinating!-- . To quote C.A.R. Hoare --from memory--: "In no engineering discipline the successful pursuit of academic ideals pays more material dividends than in software engineering." I could not agree more.)

It is here that I must mention three general phenomena that go hand in hand with greater pragmatism. I must mention them, because they seem all relevant for computing science.

The first phenomenon is a greater tolerance for the soft sciences which purport to contribute to the solutions of "real" problems, but whose "intellectual contents" are singularly lacking. (When I was a student at Leyden, a quarter of a century ago, economy and psychology had been admitted to the campus, but only with great reservations and absolutely no one considered them as respectable; we had not dreamt of "management science" --I think we would have regarded it as a contradiction in terms-- and "business administration" as an academic discipline is still utterly preposterous.)

The second phenomenon is the one for which I had to coin the term "integralism". Scientific thought, as I understand it, derives its effectiveness from our willingness to acknowledge the smallness of our heads: instead of trying to cope with a complex, inarticulate problem in a single sweep, scientific thought tries to extract all the relevant aspects of the problem, and then to deal with them in turn in depth and in isolation. (And everytime a significant aspect of a complex problem has been isolated successfully, this is ranked as an important scientific discovery. As an example I mention John Backus's introduction of BNF, capturing the context-free aspects of programming language syntax.) Dealing with some aspect of a complex problem "in depth and in isolation" implies two things. "In isolation" means that you are (temporarily) ignoring most other aspects of the original total problem, "in depth" means that you are willing to generalize the aspect under consideration, are willing to investigate variations that are needed for a proper understanding, but are in themselves of no significance within the original problem statement. The true integralist becomes impatient and annoyed at what he feels to be "games"; by his mental make-up he is compelled to remain constantly aware of the whole chain, when asked to focus his attention upon a single link. (When being shown the derivation of a correct program he will interrupt: "But how do you know that the compiler is correct?") The rigorous separation

of concerns evokes his resistance because all the time he feels that you are not solving "the real problem".

The third phenomenon that goes hand in hand with a greater pragmatism is that universities are seen less as seats of learning and centres of intellectual innovation and more as schools preparing students for well-paid jobs. If industry and government ask for the wrong type of people --students, brain-washed by COBOL and FORTRAN-- that is then what they get. I know that the perpetuation of obsolete programming habits in the U.S.A. is beginning to be considered as a matter of serious concern, because in the triangle computer users/computer manufacturers/universities, no single party seems able any longer to interrupt the vicious circle. (The moral of the text I read was that, therefore, here was a federal responsibility, because otherwise the USA could be overtaken by in this respect still more flexible nations. An outsider's corollary of this deadlock situation is that --in no field!-- Universities should forsake their role of intellectual innovators.)

* * *

A third difference between the USA and Europe must be mentioned because it has such profound consequences. The USA are very large, and, compared to Europe, much more homogeneous. Please don't accuse me of the gross oversimplification "When you have seen one American, you have seen them all". I have now been in so many states of the US and seen so many differences between them that I have concluded that, with my values of the terms, it is better for me to consider the USA not as "a country" but as "a continent". It is more that, besides all the local diversity, there are homogenizing forces in the USA that are absent in Europe. All American computing scientists write, speak and publish in the same language, they all see the publications from the same ACM and IEEE, the manuals from the same computer manufacturers, their academic research is supported by the same central funding organizations, etc. This large and relatively homogeneous continent tends to become a law unto itself; the American computing community is, therefore, in a greater danger of regarding its mode of behaviour as the mode of behaviour, it is in a greater danger of becoming provincial and parochial. (Deviation from The Standard then becomes to be considered as wrong: in the Computing Reviews of the ACM British authors of British publications are regularly being blamed for their Britishisms! See for a recent instance for example CR 30214.)

The fact that the majority of the American computing scientists are essentially monolingual, is in this discussion about computing science of special significance. A thorough study of one or more foreign languages makes one much more conscious about one's own; because an excellent mastery of his native tongue is one of the computing scientist's most vital assets, I often feel that the American programmer would profit more from learning, say, Latin than from learning yet another programming language.

* * *

Finally a difference that is very specific to academic computing science: in Europe, Artificial Intelligence never really caught on. All sorts of explanation are possible: Europe's economic situation in the early fifties when the subject emerged, lack of vision of the European academic or military world, European reluctance to admit soft sciences to the university campus, cultural resistance to the subject being more deeply rooted in Europe, etc. I don't know the true explanation, it is probably a mixture of the above and a few more. We should be aware of this difference, whether we can explain it or not, because the difference is definitely there and it has its influence on the outlook of the computing scientist.

How difficult is programming?

When, in the late sixties, it became abundantly clear that we did not know how to program well enough, people concerned with Programming Methodology tried to figure out what a competent programmer's education should encompass. As a result of that effort programming emerged as a tough engineering discipline with a strong mathematical flavour. This conclusion has never been refuted, many, however, have refused to draw it for the unattractiveness of its implications, such as

- 1) good programming is probably beyond the intellectual abilities of today's "average programmer"
- 2) to do, hic et nunc, the job well with today's army of practitioners, many of whom have been lured into a profession beyond their intellectual abilities, is an insoluble problem
- 3) our only hope is that, by revealing the intellectual contents of programming, we will make the subject attracting the type of students it deserves, so that a next generation of better qualified programmers may gradually replace the current one.

The above implications are certainly unattractive: their social implications are severe, and the absence of a quick solution is disappointing to the impatient. Opposition to and rejection of the findings of programming methodology are therefore only too understandable. We should remember that the conclusion about the intrinsically mathematical nature of the programming task has been made on technical grounds, and that its rejection is always for political or emotional reasons.

The rejection takes place at both sides of the Atlantic. It was a British programmer that commented on my book that "it would be of no meaningful benefit to the programming profession as a whole" because "its techniques are mathematical, whereas the majority of today's programmers are not." (I regard this less as a comment on my work than as a statement from an English programmer that, in his view, his current colleagues are fairly education-resistant.) It was my own Department of Mathematics in Eindhoven that needs in 1972 an easier subject than "true mathematics" in order to enlarge its undergraduate enrollment drastically and chose..... programming! (This was a very extreme case.)

On the whole, the underestimation of the mathematical maturity that is required for the programming task, seems somewhat stronger in the ASU than in Europe. In view of earlier remarks about the differences between those two continents this is understandable. Our "solution" 3 -see above-- is a long-range one and it requires a large Buxton Index to appreciate it as such. It is more Platonic than pragmatic, it is the result of a rigorous separation of concerns --abstracting from today's average programmers and also from today's average machines-- . It does an open appeal to the innovating role of Universities. It favours the careful development of "natural intelligence" based on the conviction that "artificial intelligence" will never be able to do the job.

* * *

The first series of machines --that of the singletons-- was mainly developed in the USA shortly after World War II, while a ruined continental Europe had neither the technology, nor the money, to start building computers: the only thing we could do was thinking about them. Therefore it is not surprising that many US Departments of Computer Science are offsprings of Departments of Electrical Engineering, whereas those in Europe started (later) from Departments of Mathematics (of which they are often still a part). This different heritage still colours the departments, and could provide an acceptable explanation that in the USA Computing

Science is viewed more operationally than in Europe.

We may add to this that John von Neumann's habit to describe computing systems and their parts in an anthropomorphic terminology has been adopted more generally in the USA than in Europe. (I was first exposed to the American's use of anthropomorphic terminology in the late fifties --when the CommACM started to appear-- and I remember that I was shocked by it. In the meantime, a less anthropomorphic terminology had already been established in my environment.) The problem caused by this metaphor is, that it invites us to identify ourselves with programs, with processes, etc, because "existing" is one of our most intrinsic "activities". (That is, why death is so hard to grasp.) The prevailing anthropomorphism erects another barrier to abstraction from program execution and computational histories.

To forget that program texts can also be interpreted as executable code, to define program semantics as a direct derivation from the program text and not via the detour of the class of possible computations, to define programming semantics independent of any underlying computational model, these are difficult abstractions to get used to. I have the impression that for an American computing scientist it is still harder than for a European one. Yet it is one of the most vital abstractions, if any significant progress is to be made at all.

It was the complete entanglement of language definition and language implementation that characterized the discussion after Mary Shaw's presentation, and it was this entanglement that left many of the Europeans flabbergasted. It was also this entanglement that made it impossible for me to read the LISP 1.5 Manual: after an incomplete language definition, that text tries to fill the gaps with an equally incomplete sketch of an --of the?-- implementation. Yet LISP 1.5 conquered in the decade after the publication of that manual a major portion of the American academic computing community. This, too, must have had a traceable influence. Why did LISP never get to that position in Europe? Perhaps because in the beginning its implementation made demands beyond our facilities, while later many had already learned to live without it. (I myself was completely put off by the Manual.)

* * *

My first visit to the USA, in 1963, was the result of an amazing invitation

from the ACM. Without the obligation to present a paper I was asked to attend --as "invited participant", so to speak-- a three-day conference in Princeton: for the opportunity of having me sitting in the audience and participating in the discussions, my hosts were willing to pay my expenses, travel included! As you can imagine, I felt quite elated, but shortly after the conference had started, I was totally miserable: the first speaker gave a most impressive talk with wall-to-wall formulae and displayed a mastery of elaborate syntax theory, of which I had not even suspected the existence! I could only understand the first five minutes of his talk, and realized that I was only a poor amateur, sitting in the audience on false pretences.

I skipped lunch, walking around all by myself, trying to make out what that first speaker had told us. I got vaguely funny feelings, but it was only during the cocktail party that evening, when I had recovered enough to dare to consider that it had all been humbug. Tentatively I transmitted my doubts to one of the other participants. He was amused by my innocence. Didn't I know that the first performer was a complete bogus speaker? Of course it was all humbug, everybody in the audience knew that! Puzzled I asked him why the man had been invited and why, at the end, some of the participants had even faked a discussion. "Oh, on occasions like that, we just go through the motions. IBM is one of the sponsors of this conference, so we had to accept an IBM speaker. He was given the first slot, because the sooner that is over, the better." I was flabbergasted.

Since then I have learned that this "going through the motions" is, indeed, a typical habit of the American scientific community. Whenever a large project is sponsored by a sufficiently prestigious or powerful body (MIT, ARPA, IBM, you name it), it is officially treated as sound and successful. The above story illustrates how utterly misleading that habit can be for an innocent European. By European standards, that habit is nearly fraudulent. But if Americans have a capacity for greater dishonesty, they have also a capacity for greater honesty! From American sources --both private and public-- I can quote many comments on the Americans so candid, that I cannot imagine a European discussing his own country in similar terms.

In other words, the rules that govern when to be explicit and when to be silent, when to exaggerate for the sake of emphasis and when to use euphemisms,

differ in the two continents. In international groups, this can cause endless confusion, and I see only one way out: to make it for such a group an explicitly stated rule that everybody be outspoken and as clear as possible.

I don't remember whether it is the result of a consciously taken decision or whether the tradition just grew, but in W.G.2.3 we certainly used to apply such a rule, knowing full well that we would often display what looked like inconsiderate behaviour. I now understand why in a group like W.G.2.3 such a rule is absolutely essential, and I would like you to share that understanding with me. I also suspect that its former application is largely responsible for W.G.2.3's former success, and I would like you to share that suspicion with me.

Plataanstraat 5
5671 AL NUENEN
The Netherlands

prof.dr.Edsger W.Dijkstra
Burroughs Research Fellow

P.S. I apologize for having been so often so apologetic.

EWD.