

"From my Life"

(Written because people ask me for these data.)

I was the third of four children. My father was an excellent chemist, who worked in the town of Rotterdam, first as a teacher of chemistry at a secondary school, later as its superintendent. His knowledge was vast and he was perfectly able (and willing) to teach cosmography or physics when needed. He made a few inventions, one of them really sophisticated. Among chemists all over the country he was well-known --he had, for instance, been President of the Dutch Chemical Society-- but steadfastly refused all forms of scientific promotion, explaining that he would rather be the first among the counts than the last among the dukes, but the net effect of adhering to this lofty principle was that he lost the colleagues that were intellectually his equals to the university campus. In the vague idealism after WWII Dutch secondary education deteriorated to such an extent that he forbade me to become a schoolteacher. My mother was a brilliant mathematician who had no job --as was customary in those days, even for most women with a college degree-- . Her mind was much less encyclopaedic than that of my father but very quick; she had a great agility in manipulating formulae and a wonderful gift for finding very elegant solutions, and I learned a lot from her though she had not the patience to be a good teacher.

During the last class of the Gymnasium --that was 1947/48-- I considered going to University to study law, as I hoped to represent my country in the United Nations. Then we did our final exams. Before we got our grades and were told whether we had passed, I was very depressed, for I reviewed all the mistakes and stupidities I had committed and I got convinced that I would fail. (My consolation was the conviction that I would get a second chance, always having been the best pupil in the class.) My grades turned out to be better than most teachers had ever seen at a final examination, and older and wiser people convinced me that it was safer to base my choice on ability than on idealism, and that it would be a pity if I did not devote myself to science. So I went to Leyden University to study mathematics and physics during the first years, and theoretical physics during the latter part.

On the whole I enjoyed the years at the university very much. We were very poor, worked very hard, never slept enough and often did not eat enough, but life was incredibly exciting. (Now, more than 40 years later, I still remember

the first "international" piano recital I attended: Clara Haskil playing Mozart!) Another thing I remember --because it surprised me-- was that a number of my professors --Kloosterman, Haantjes, Gorter, Kramers-- took me seriously.

My father, who was a subscriber to "Nature", had seen the announcement of a three-week course in Cambridge on programming for an electronic computer (for the EDSAC, to be precise), and when I had passed the midway exam before the end of my third year --which was relatively early-- he offered me by way of reward the opportunity to attend that course. I thought it a good idea, because I intended to become an excellent theoretical physicist and I felt that in my effort to reach that goal, the ability to use an electronic computer might come in handy. The course was in September 1951; by accident, A.van Wijngaarden, then Director of the Computation Department of the Mathematical Centre in Amsterdam, heard of my plans, invited me for a visit, and offered me a job. Due to obligations in Leyden I could accept that invitation only in March 1952, when I became the first Dutchman with the qualification "programmer" on the payroll. Professor H.A.Kramers, who had accepted me as student of theoretical physics, died, I did not like his successor, and for a while I considered continuing my studies in Utrecht, but it turned out that the laboratory work I would have to do in Utrecht was incompatible with my halftime job at the Mathematical Centre. That was the moment I had to choose between theoretical physics and computing; I chose the latter and my study of physics in Leyden became a formality. I got my degree in the spring of 1956, moved immediately from Leyden to Amsterdam, and from that moment on I worked full time at the Mathematical Centre.

My closest collaborators there were Bram J.Loopstra and Carel S.Scholten, who had been hired (while they were still students of experimental physics at the Municipal University of Amsterdam) to build a computer for the Mathematical Centre. They were about five years older than I, had been close friends since the age of 12, and had already worked on their machinery for quite a few years when I joined them. When I came their machine did not work, but they had learned a lot; among the things they had learned was to recognize that their first effort contained too many technological mistakes to be salvageable, and that they should start afresh. So, early in the game I became involved in the design of "the next machine", something that would be repeated a number of times. In retrospect, the most noteworthy aspect of that cooperation is that we accepted without any

discussion a discipline that seems rare today: the very first thing we did was to write a programmer's manual for the next machine, including its complete functional description. (For some reason, I was always the one that wrote that document.) This document then served as a contract between them and me: they knew what they had to build and I knew what I could build upon. And about a year later, we would be both ready: they had built their machine and I had all the basic software finished. (Not tested, of course, but that turned always out to be superfluous. At the time I did not realize how valuable this training was of programming for a machine that did not exist yet; now I know.)

I designed my first nontrivial algorithms. The algorithm for The Shortest Path was designed for the purpose of demonstrating the power of the ARMAC at its official inauguration in 1956, the one for The Shortest Spanning Tree was designed to minimize the amount of copper in the backpanel wiring of the X1. In retrospect, it is revealing that I did not rush to publish these two algorithms: at that time, discrete algorithms had not yet acquired mathematical respectability, and there were no suitable journals. Eventually they were offered in 1959 to "Numerische Mathematik" in an effort of helping that new journal to establish itself. For many years, and in wide circles, The Shortest Path has been the main pillar for my name and fame, and then it is a strange thought that it was designed without pencil and paper, while I had a cup of coffee with my wife on a sunny cafe terrace in Amsterdam, only designed for a demo....

Life became serious when Bram and Carel embarked on the design of the X1, which after a while they conceived with a real-time interrupt. I was frightened out of my wits, only being too much aware that this would confront me with a machine which was to all intents and purposes nondeterministic. For about three months I was able to delay the decision to include the interrupt, until eventually, they flattered me out of my resistance. (Their plea was along the lines "Yes, we understand that this makes your task much harder, but surely someone like you must be up to it, etc.") The real-time interrupt handler became my thesis topic; I earned my Doctorate in late 1959.

In the mean time, the ALGOL project was underway. I had attended a few conferences outside the Netherlands, but it was ALGOL that introduced me to frequent foreign travel. Those trips were very tiring, for my English was untrained.

(I am very grateful to Mike Woodger from the National Physical Laboratory, who quickly started correcting my English, but in the beginning only the errors he had heard me make more than once.) In the summer of 1959 I had discovered how to implement subroutines that could call themselves, by the end of the year I saw how to use a stack for the evaluation of expression and how to translate expressions from the usual infix notation into "reverse Polish" (only we did not call it that). For several months, Harry D. Huskey had been working at the Mathematical Centre, having been invited by van Wijngaarden in the hope that he would teach us how to write what was then called "an algebraic translator". Van Wijngaarden wanted us to base our ALGOL 60 compiler on Huskey's work, which was a mess (Huskey being an experimental programmer); I had seen that work and knew that we had to ignore it, but in order to convince van Wijngaarden I had to bang my fist on his table. (I remember this, because this has been the only time that I did so.) Late December I started with J.A. Zonneveld on an implementation of ALGOL 60 for the X1. We coded the translator in duplo, fully aware that we were making history: each evening each of us took his own copy of the manuscript home so that in case of a fire at the Mathematical Centre our translator would not be lost. (There was no fire.) In August 1960 our implementation was working, more than a year before the one of our nearest competitor, the group of Peter Naur in Copenhagen. In May 1960 I submitted the short article "Recursive Programming" to "Numerische Mathematik", which published it quickly; here I introduced the term "stack" ; while describing the quintessence of our implementation. For many years I thought that that paper had been ignored in the USA, but later I learned that it had established my name among America's scholarly computing scientists.

The ALGOL compiler had been a major effort on which I had worked day and night for 8 months, and I needed to recover; more than anything I needed to be free from the responsibility for a large project. For a few years I did not build anything sophisticated and ambitious, but thought, mainly about the logical intricacies of multiprogramming. This was when I formulated the Mutual Exclusion Problem, thought about synchronizing primitives and designed "semaphores" for that purpose. I did not publish that last invention because the company "Electrologica", for which I was a consultant, was going to incorporate them in their next machine, but I freely lectured about the topic. In 1961, at Brighton, Peter Naur and I lectured on ALGOL 60 and its implementation; at that occasion, Tony Hoare and I met for the first time.

At the end of the summer of 1962 I attended the IFIP Congress in Munich, where --thanks to van Wijngaarden-- I was given the opportunity to deliver an invited speech. My speech was well-received, and the congress was the occasion at which Niklaus Wirth and I met each other for the first time. When I came back from Munich, it was September, and I was Professor of Mathematics at the Eindhoven University of Technology.

Later I learned that I had been the Department's third choice, after two numerical analysts had turned the invitation down; the decision to invite me had not been an easy one, on the one hand because I had not really studied mathematics, and on the other hand because of my sandals, my beard and my "arrogance" (whatever that may be). I have also been told that I should be very grateful to the Department because it had taken a great risk by appointing me. For two years I gave the courses in Numerical Analysis because they had no one else to do so, and at the same time I began building a little group of computing scientists. In the years that followed, I built with Cor Ligtmans, Piet Voorhoeve, Nico Habermann and Frits Hendriks what became known as the "THE Multiprogramming System" for the X8 (the next machine of Electrologica, of which the university had ordered a copy); machine and operating system served the computation centre of the university very well for about half a decade, until the X8 had to be replaced by a more powerful engine.

With the THE Multiprogramming System we also knew we were making history, because it was the world's first operating system conceived (or partitioned) as a number of loosely coupled, explicitly synchronized, co-operating sequential processes, a structure that made proofs of absence of the danger of deadlock and proofs of other correctness properties feasible. I presented the quintessence of this design in Teddington, in Edinburgh, and in Gatlinburg, Tennessee, (at an ACM Conference on Operating Systems Principles), and it got all the attention and recognition it deserved. The Department saw it differently and disbanded the little group, in reaction to which most of them left the University. It was obviously too difficult for the Department to honour what had been achieved: as classical mathematicians they were not used to group efforts at all, and furthermore it had no connection to the kind of mathematics they were used to. In addition, some began to fear that, instead of a strong component of the Department, Computing Science could become too powerful a competitor for the true mathematics that were the Department's real concern. For me, life became a little bit difficult.

For more than half a year I suffered from a rather severe depression, but things became better when I realized that I was frightened by the next challenge, viz. the study of difficulty as such, and of how to overcome it. How do we do difficult things? My previous more ambitious projects acquired the status of finger exercises, and for therapeutic reasons I wrote "Notes on Structured Programming"; as far as my depression was concerned, the therapy worked.

Of the "Notes on Structured Programming" I sent only about 20 copies to friends abroad, but by that time the copier had become ubiquitous, and (without my knowledge) the text spread like wildfire: later I met people in far-away places who treasured a fifth- or sixth-generation copy! It was not a great text and its spontaneous wide distribution surprised me when I heard of it; the only explanation I have been able to think of is that it was the first text that openly took the programming challenge very seriously. It has certainly played a role in the founding of the IFIP Working Group 2.3 on "Programming Methodology", a topic that would only in the 70s become a subject of academic concern.

Late one afternoon in 1972, the telephone in my office at the University rang; I picked up the receiver, and this was Franz L. Alt of the ACM, calling from the USA to tell me that I had won the ACM Turing Award. It took me by complete surprise for I had never considered myself as a potential recipient of that most prestigious award, and I could hardly believe it, but Alt convinced me that it was really true, and then we disconnected. When I realized that my six predecessors had all been native English speakers working at famous institutes in the USA or the UK, I got even a bit overwhelmed. When I had calmed down I went to the office of the Department's Chairman to inform him if he was still in. He was, and I told him. His first reaction was "In the world of computing, they are rather lavish with prizes, aren't they?". When in the fall of that year, another colleague, N.G. de Bruijn, launched a plan in which a massive undergraduate course in programming would earn the money for the graduate program in "genuine mathematics", I knew it was time to go. After having given Philips the first right of refusal, I became Burroughs Research Fellow on the 1st of August 1973.

That same summer I designed the first so-called "self-stabilizing systems". It was a new problem of which no one knew whether it could be solved at all, and the solutions were at the time a tour de force. I submitted a short article

to the Communications of the ACM, which published it in 1974; it was ignored until Leslie Lamport drew attention to it in 1984, and since then it has added a new dimension to ways of dealing with error recovery.

In the fall of 1973 I designed "predicate transformers" as a tool for defining program semantics in a way that would provide a suitable basis for a calculus for the derivation of programs. In the late 60s, Tony Hoare had provided a basis for this work: he taught us the relevance of the predicate calculus for reasoning about programs and showed us how inference rules could be tied to the formal syntax of the programming language, but it was at the time not clear to which extent his "axiomatic basis" defined the programming language semantics. I tightened that up and introduced in passing nondeterminacy. This formed the starting point for the book "A Discipline of Programming", which kept me busy for the next few years. Prentice-Hall published it in 1976, it is still for sale, and it has been quite influential; the organization that publishes the Science Citation Index has identified it as a "Citation Classic".

I did the "On-the-fly garbage collection", which at the time was quite a feat of fine-grained interleaving, and the distributed termination detection of diffusing computations, which emerged as a fundamental problem. In 1981 I designed a surprising sorting algorithm.

On the whole, my years as Burroughs Research Fellow were a wonderful time. I would visit the company in the USA a few times a years and my further assignment was to do my own thing, which I did in the smallest Burroughs research outfit in the world, viz. my study on the second floor of our house in Nuenen. I maintained my contacts by mail, by going to the University once a week, and by travelling all over the world. I lectured all over the world, but foreign travel lost what little attraction it had (I never was a very good tourist: I managed to visit Moscow without being dragged to the Kremlin). Whenever I was in the country, I would work on Friday with Carel S.Scholten; others I saw as regularly were W.H.J.Feijen, A.J.M.van Gasteren and Alain J.Martin.

It was too good to last, and it didn't. When Michael W.Blumenthal took over, Burroughs Corporation changed from an enlightened computer manufacturer

into a financial organization, and precisely the kind of people I had established relations with were the first to leave the company: I lost my audience there. Whereas Burroughs's intellectual horizon was shrinking, my own was widening, as my interests shifted from computing and programming methodology to mathematical methodology in general. A return to the University Campus seemed appropriate, and when The University of Texas at Austin offered me the Schlumberger Centennial Chair in Computer Sciences, the offer was actually welcome. The offer was not so surprising, nor my acceptance, for I knew Austin, I knew UT and they knew me. All through my years at Burroughs I had made it a rule to visit, wherever I went, the local university and since the late 70s Burroughs had had an Austin Research Center, a visit to which was included in most of my trips to the USA.

In my interest in mathematical methodology I had been encouraged by a grant that I received thanks to the efforts of Dr. Donald W. Braben of the Venture Research Unit of BP, a grant that enabled me to offer a position to Netty van Gasteren (who thus became the first BP Venture Research Fellow). Another strong stimulus came from my efforts --in which I would soon be joined by Carel S. Scholten-- to give my "predicate transformer semantics" of the 70s the necessary mathematical underpinning; these efforts became an exciting experience when we discovered (after a few notational innovations) the surprising utility of calculational proofs. In 1988, Netty van Gasteren successfully defended her Ph.D. thesis "On the shape of mathematical arguments".

In the summer of 1984 we moved to Austin, Texas, where we have been happy ever since, partly by the simple device of having warned ourselves beforehand that The University of Texas at Austin would not be Heaven-on-Earth. (It is not.) With the prolonged recession, the pressures to do the wrong things, always already strong, are mounting, and part of the university gives in. And in computing, the excessive preoccupation with speed is, of course, a little bit vulgar.

Austin, 29 November 1993

prof.dr.Edsger W.Dijkstra
Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712-1188
USA