

Historic documents

Tony Hoare

October 2006

held in
Computer Museum
Mountain View.

[1] Russian article.

The article was written for a Russian Journal on Machine Translation and Applied Linguistics. The displayed typescript was typed by me in 1960, when I was a graduate student at Moscow State University. I used a Russian manual typewriter lent by my friend Volodya Demin. At the end, I acknowledge the help of Igor Melchuk, now a distinguished figure in linguistics.

The article describes a rather trivial algorithm for translating a tree structured syntactic analysis to a linear sentence, with the words in the correct order for the target language. The algorithm is much less significant than that described in the next exhibit, which was also motivated by application to machine translation.

[2] Quicksort.

In 1960, the standard method of machine translation was to sort the words of a source language sentence into order, so that they could be looked up on a single pass of the dictionary stored on magnetic tape. This algorithm for sorting was simply the second idea that occurred to me in my room at Moscow State University in 1960. I had rejected the first idea as too slow.

But I did not know exactly how fast it was until a year later in England, when one Sunday afternoon I was playing around with some formulae and discovered that I had (again) proved that $0 = 0!$ Fortunately, a careful analysis of the proof showed that it could be played backward as a valid proof of the formula on page 12. This discovery encouraged me to write up the article for publication in the British Computer Journal. More complete analysis of the algorithm by Don Knuth later confirmed its efficiency.

[3] Report on the Elliott ALGOL translator.

In 1960 I was employed as a programmer by Elliott Brothers of London Ltd., a small British scientific computer manufacturer. After a few months experience programming, I was put in charge of the implementation of the programming language ALGOL 60, which had just been designed by an international panel of experts: the language manual was only 23 pages long, written by Peter Naur.

This paper was published by the British Computer Journal in July 1962. It describes the background for the design decisions of the Elliott ALGOL compiler, and reports on experience of the implementation project. Things were so much simpler and smaller in those days -- but even so, we made the same mistakes (see [5] below).

[4] A programming language for processor construction.

In 1965, I wanted to attend the IFIP Congress in New York; and to obtain support from my employers (Elliott Bros., Ltd.) I had to write a contribution: here it is.

The aim of the contribution was to extend the applicability of single-pass syntax-directed compilers, by writing them in a single-assignment language, where the values of variables could be used before they were assigned. Peter Naur was in the audience, and at the end he stood up to give some friendly but devastating advice: that writing a nine-pass compiler was in fact easy -- he had now just done it. As a result, I never wrote up the article for publication, and this is the only copy of the text. But I took my wife and infant son to New York, and we both enjoyed our trip.

[5] Problems of Software Development.

The main part of this exhibit is a cyclostyled memorandum distributed among the senior programmers and divisional management of Elliott Brothers' Computer Division. It reports the findings of a post mortem meeting held on October 22 1965, to investigate the collapse of the entire project to produce 'second generation' software for the Company's main computer product.

Some years later, I submitted this article for publication in 'Software Practice and Experience', because I thought one could learn more from failure than success -- as I myself certainly did. But the referees of the journal thought otherwise, and the story had to wait for public airing in my Turing Award address in 1980.

[6] Programming Languages: Predictions and Prospects.

In 1968, I attended the IFIP Congress in Edinburgh. I was sitting in the audience in the interval between presentations, when my neighbour drew my attention to an article in the Computer Weekly. 'Good stuff, this,' he said. I'm afraid my unconsidered reply was 'Yes, I know, I wrote it'. Here it is.

Some of the predictions made in the article have come true quite quickly, but many have not. Unfortunately, it is the optimistic predictions that have not been realised.