UNIX

A History and a Memoir

Brian Kernighan

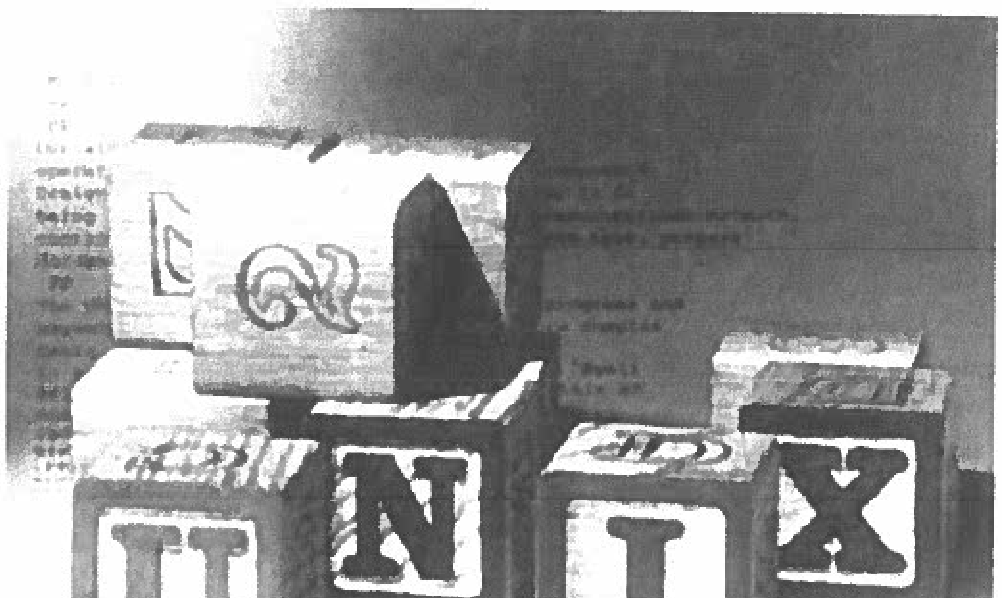

This was enough justification for a PhD thesis topic, but I wasn't making much progress. When I returned to the Labs for a second internship in the summer of 1968, I described the problem to Shen Lin (Figure 1.8), who had recently developed the most effective known algorithm for the classic Traveling Salesman problem: given a set of cities, find the shortest route that visits each city exactly once, and then returns home.

Shen came up with an approach for graph partitioning that looked promising, though there was no assurance that it would produce the best possible answers, and I figured out how to implement it efficiently. I did experiments on a large number of graphs to assess how well the algorithm worked in practice. It seemed highly effective, but we never discovered a way to guarantee an optimum solution. I also found a couple of interesting special-case graphs where I could devise algorithms that were both fast and guaranteed to produce optimal solutions. The combination of results was enough for a thesis, and by the end of the summer I had everything I needed. I wrote it up over the fall, and had my final oral exam late in January 1969. (Princeton's optimistic estimate of three years had turned into four and a half.)

A week later, I started work in the Computing Science Research Center at Bell Labs. I never had an interview; the Labs sent me an offer sometime in the fall, though with a caveat: my thesis had to be finished. Sam Morgan, the director of the Center and thus my boss two levels up, told me, "We don't hire PhD dropouts." Finishing the thesis was definitely a good thing; I got another letter in December saying that I had received a substantial raise, well before I reported for work!

As an aside, although Shen and I didn't know it at the time, there was a reason why we were unable to find an efficient graph-partitioning algorithm that would always find the best possible answer. Others had been puzzling over the inherent difficulty of combinatorial optimization problems like graph partitioning, and had discovered some interesting general relationships.

In a remarkable 1971 result, Stephen Cook, a mathematician and computer scientist at the University of Toronto, showed that many of these challenging problems, including graph partitioning, are equivalent, in the sense that if we could find an efficient algorithm (that is, something better than trying all possible solutions) for one of them, that would enable us to find efficient algorithms for all of them. It's still an open problem in computer science whether such problems are truly hard, but the betting is that they are. Cook received the 1982 Turing Award for this work.

When I got to Bell Labs as a permanent employee in 1969, no one told me what I should work on. This was standard practice: people were introduced to other people, encouraged to wander around, and left to find their own

Figure 1.8:

research topics and collab
been daunting, but I don't
going on that it wasn't har
with, and after two summ
projects.

This lack of explicit
Projects in 1127 were not
tom up, coalescing a grou
same was true for work w
some development group.
me, but they would be vol

In any case, for a whil
optimization. Shen was e
sense a promising line of a
had a new idea for the tra
improved on his previous
I implemented it in a Fort
was the state of the art.

This kind of work was
ideas into working code
parts. So I gradually drift
specialized programming l

I did come back to wo
complex tool for optimizi
tomers. It was good to fli
puter science and systems