

Some AMPL History

David M. Gay

ABSTRACT

This is a brief history of AMPL from my perspective.

My first brief encounter with optimization was in a high school, when I learned that smooth minima and maxima are critical points. In college I worked for a couple of summers in the office of the local refinery, where I heard — without details — that linear programming was used to choose the mix of feedstocks bought and products produced. In graduate school (I studied in the then Computer Science Department at Cornell University), I learned about nonlinear optimization algorithms and about the simplex method for linear programming. There was of course much interest in complexity. There was a Mathematical Programming Symposium (a meeting still held every three years) at Stanford University. I did not go, but the numerical analysis professors did and brought back a paper by Hugo Scolnik that claimed to describe a polynomial-time linear programming algorithm. I found faults with Scolnik's method and wrote a couple of technical reports thereon. Later when Leonid Khachiyan published a polynomial-time algorithm, I wanted nothing to do with the resulting excitement.

After teaching for a year, I went to work at the NBER Computer Research Center in Cambridge, Massachusetts (which became MIT/CCREMS after NBER got a new president who did not like computers). There I worked mostly on nonlinear data fitting (e.g., the NL2SOL algorithm), but also met Bob Fourer, who had recently graduated from MIT and at NBER was documenting an LP solver, SESAME, that William Orchard-Hays had written for NBER. Bob then went off to graduate school to get a PhD in Operations Research from Stanford University under George Dantzig. Bob subsequently became a professor at Northwestern University. At some point, he published a paper in *ACM Transactions on Mathematical Software*, in which he argued for use of a modeling language. I had meanwhile gotten a job at Bell Labs (Murray Hill) and had moved to New Jersey. A few years after I started at Bell Labs, Narendra Karmarkar, who was then also at the Murray Hill location of Bell Labs, announced a polynomial-time linear programming algorithm that he said could perform much better than Khachiyan's algorithm. I was in the Computer Science Research Center, where there was much interest in "little languages" to make various tasks easier to do on a computer. I saw Bob at a meeting and learned he was coming up for a sabbatical. I talked with my management, and they arranged for Bob to spend his sabbatical at Murray Hill. While he was there, he, Brian Kernighan, and I had many discussions and invented AMPL. Brian wrote the first

AMPL translator. To make giving data to it easier, I wrote a preprocessor that handled data sections. Later, to facilitate various extensions, I rewrote much of the AMPL translator and incorporated the data preprocessor.

Among other things, I was interested in using AMPL for nonlinear problems. I was aware of the AUGMENT preprocessor that could differentiate programs suitably written in Fortran. I wanted to do something similar with AMPL, which today we would call forward automatic differentiation (AD). I ran across Andreas Griewank at another Mathematical Programming Symposium and told him about my ideas for AD in AMPL. He told me about backwards AD and convinced me that on complexity grounds backwards AD would be much better than forward AD for computing gradients. I subsequently wrote a couple of papers on AD, added it to the AMPL/Solver Interface Library (ASL), and wrote several nonlinear solver interfaces that computed first and sometimes second derivatives by AD.

I would have loved to spend my whole career at Bell Labs — a fantastic place — but that was not to be. In 2003 the “dot com bubble” burst. Lucent Technologies, which after the breakup of the Bell System had become the parent company of Bell Labs, fell on hard times. Lucent had made loans to various customers so they could buy telephone switches and optical fiber from Lucent. After the “dot com bubble” burst, these customers could not make payments on their loans. Lucent had a lot of money in its retirement account and made generous retirement offers to more senior people, including me. I accepted the offer presented to me. At that point, Bell Labs had some paying customers for AMPL. While we worked on paperwork that would let the three co-authors set up a company to support AMPL, I was allowed to continue using my Bell Labs office. After a couple of years the Lucent lawyer working with us got very sick, and I feared nothing would happen with AMPL. I then pursued a job at Sandia National Labs in Albuquerque. After I got the job, Lucent woke up and the AMPL company got permission to deal with AMPL. At Sandia, I could not be paid for any work on AMPL, so I worked at Sandia by day and did things for AMPL without pay at night.

Eventually Bob Fourer got an NSF SBIR grant to make enhancements to AMPL (the constraint-programming extensions). I wanted to be involved and to work half-time for Sandia and half-time for AMPL, but Sandia would not allow this. I left Sandia in 2010 and have been working on AMPL from home ever since.

Sanjay Saigal introduced us to Bill Wells. With Bob Fourer, we all met at my house in Albuquerque in early 2010. We convinced Bill to become our CEO. The AMPL company has grown slowly ever since.

Much of my work over the years was facilitated by my wonderful wife, Tanner, who unfortunately died in January 2021. (See <https://daveandtanner.net> for much more about her.) Tanner gave me much time for technical work, and we also enjoyed doing many nontechnical things together. Tanner was an excellent writer and she often made useful comments on writings I had drafted. (Until we moved to New Jersey a little more than a year after we married, Tanner worked at publisher Houghton Mifflin Company in Boston as a senior editor in the School Reading Department. For many years

after we moved, she continued as a free-lance editor.) Early AMPL company meetings were held at our house, and Tanner helped make them a success. Though Tanner did not contribute technically to AMPL, she very much helped my contributions by helping make a wonderful environment in which I could work.

Many people remember 2020 because of Covid19. I barely noticed Covid because in early February, after I fell and broke my collarbone, I was found to have a slowly growing brain tumor, and I ended up spending more than 100 days in hospitals. Until then I was doing most of the technical work on AMPL and associated solver interfaces. Now, as I slowly recover, several other people are also involved.

Losing Tanner in early 2021 was very hard. To cope, I have talked and corresponded with many people; I have also gone to various events with friends and attended some grief counseling, all of which has been quite helpful. Working on the material in <https://daveandtanner.net> was also therapeutic.