My first supercomputer … 1982-4

… was a Cray 1S which our company bought in 1982 I think. Its on the left with an old colleague taking advantage of the most expensive seating in the world. This beastie cost around 3million pounds which was a very considerable sum in those days. Ours was only the 6th in the UK. Since the others were owned by companies like Shell and a certain code-cracking organisation in Cheltenham, we were actually vetted to see if we were suitable. In the end it was deemed we were and so my team and I got down to making the thing pay its way. The difference was that Shell etc. had magnificent machine room empires which probably cost more than the Cray, whereas ours was housed over the Woolwich building society in the middle of Woking, Surrey, UK. The damn thing was so heavy we had to put spreader plates on the floor as it exceeded the floor loading as it stood. Had these not worked, it would have been the biggest deposit the Woolwich ever had. It also consumed an inordinate amount of electricity (around 130kW) and we had to lay a special cable to a local substation. You could always tell where this was because when it snowed, the snow always melted over our cable.

I had always had a vision of software that just moves without incident between machines, after all, why shouldn't it? We had international standards for languages but when you are talking to various kinds of mass storage and manufacturers deliberately introduce features that trap you on their architecture and so on, its not always so simple, however by then I knew how to do it and developed tools for my team to use which enforced the necessary constraints. One of these eventually became *Flint*, a Fortran portability checker and was quite widely used by the aerospace industry including a number of NASA sites https://leshatton.org/Flint_0489.html, another became *Safer C* https://leshatton.org/index_SA.html which was used widely in the automotive industry.

Our struggles with design for portability eventually got published in *Software Practice and Experience*, https://leshatton.org/SPE_488.html. Today, this problem has been greatly ameliorated with the growth of portable open source scripting languages such as perl and Python.

Our main computer centre at Woking houses a Cray 1-S/1300 supercomputer with Sperry and Data General systems for front-end processing. Facilities are the basis of a service – but quality of staff is all important, and our 140 people can handle the most complex large seismic processing contracts from anywhere in the world.

Seismic processing at Merlin begins with a meeting to define client objectives and assemble all the background information needed to produce the best results.

Intermediate stages include: designing and testing the processing route; demultiplexing on a dedicated Perkin Elmer computer; pre-stack processing to enhance signal to noise ratio on the raw data; area adaptive analysis to estimate exact dynamic corrections; multiple attenuation; stacking; a second, post-stack procedure to improve signal content; migration to a true zero offset section; high resolution laser plotting onto reproducible film.

Every stage is accompanied by simultaneous quality control checks, and the final product is a fully documented report, including all trials testing, data examples, processing sequences with parameters, and our opinion of the results if requested.

We offer our clients only the best quality, and recognise inside Merlin that the greatest influences on that are innovative software design, capable and well motivated staff, and service-oriented management.

The success of our policy is apparent from the excellent work produced and the grow-

Seismic Kernel System – SKS – is a complete suite of seismic processing software designed by Merlin. Although this is surely one of the most complex applications in the computer world, we have been able to make SKS portable across most mainframes and minicomputers.

We are not the only company claiming to have portable programs, but we are the only company to prove it by porting our seismic package from a mixed configuration of conventional mainframes onto a vector-processing Cray supercomputer – in a matter of a few man-weeks.

All our software packages have this essential feature of portability.

We offer them to the industry at large, which gives clients the opportunity to update software capability without having to buy new hardware.

Information on Merlin Profilers software packages is published in a regular series of technical bulletins and is available on request to oil, seismic and computer managers and technical specialists.

We took delivery of the Cray 1-S/1300 supercomputer in mid July 1984. Largely because we had already designed portability into our software, a very small team of Merlin programmers had the Cray working by mid September. We were running major seismic contracts on it by mid October, only 3 months after delivery.

From a company brochure for Merlin Profilers SA. I am on the right chatting with my colleague. The machine is a Cray 1S with a dual Univac 1106 front end. It had twin 600MB discs and 8MB of memory. The whole lot cost around 3.5 million quid. To put it into perspective, this machine is approximately 10,000 times slower than the desktop I am typing this on and has 800 times less disc space and 4,000 times less main memory. How times have changed.

25-Jul-2024