Boris Chirikov and the Cray I

by Franco Vivaldi

In the summer of 1979, Boris Chirikov visited the United States, his first and only visit to North America. He arrived in Atlanta, where he was greeted by the late Joseph Ford, his host and good friend. At the time, I was a postdoctoral fellow with Joe, at Georgia Tech.

Boris went to America to use the Cray I supercomputers, then the fastest on Earth. They were mind-bogglingly faster than anything available in Siberia at the time. For years, Boris had performed miracles with the spartan computing facilities of the Institute of Nuclear Physics in Novosibirsk. He told me that sometimes they wrote the inner loop of programs in assembler, to improve speed. Some of Boris' 'tricks' —such as considering Poincaré recurrence instead of correlations decay— were born out of computational necessity, but ended up affecting subsequent developments of the theory.

I followed Boris and Joe during their trip West, visiting the Oak Ridge National Laboratory in Tennessee, the National Center for Atmospheric Research at Boulder, Colorado, and Berkeley University. Oak Ridge was linked to the Cray I computer of the Los Alamos National Laboratory. It was the middle of the Cold War, and it was an unusual situation to have a prominent Soviet scientist using state-of-the-art facilities in a US national lab. It's indeed a tribute to Joe's skills and perseverance if this visit took place at all.

We arrived at Oak Ridge with a Fortran program that performed long-time dynamical evolution of a driven two-dimensional map. A fast computer was needed to measure the subtle phenomenon of Arnold diffusion. A local Cray expert briefed us on how to approach the awesome machine. He showed us graphs and statistics, detailing the machine performance; he then stated the maximal speed attainable in optimal situations, when the so-called full vectorization of a code was implemented.

I sat in front of a terminal; Boris behind me, looking over my shoulders. Optimization of computer programs had never been a major concern of mine. Boris suggested several of small changes in my code: so divisions turned into multiplications, scalar variables became arrays, statistical tools were fine-tuned. Finally, there were two independent nested loops —the core of the program. Boris suggested that I swap the order of the loops. Mathematically, it made no difference, but he explained that, given how matrix elements are accessed, reversing the order would speed up the program.

In an hour or so, we were ready for a trial run on the Cray I. At the first go, the program ran faster that the declared theoretical limit. The Cray experts were stunned: they ended up downloading and analyzing the assembler code, to figure out how this could have happened.

This was vintage Boris Chirikov.

School of Mathematical Sciences
Queen Mary, University of London
London, November 2008