CONQUEST: A Quantum Leap in Atomic-Scale Simulation

Dr David Bowler, Prof. Mike Gillan
University College London

Simulations of matter on the atomic scale based on quantum mechanics have made a colossal impact on a wide range of basic sciences over the past 10 years, including physics, chemistry, biology and earth sciences, and the CSAR service and its predecessors have been crucial for this work in the U.K. All kinds of scientific problems have been tackled, ranging all the way from chemical reactions on surfaces to the physics of the Earth's core.

But up to now, the algorithms have suffered from a major bottle-neck, because the computer effort needed increases too fast with the number of atoms being simulated. This has meant that large systems containing thousands of atoms, which are essential for studying complex materials, nanostructures or biological systems, have been out of reach. This bottle-neck has now been well and truly broken. With the CONQUEST code running on the Cray-T3E turing, practical quantum calculations have been achieved on systems of over 16,000 atoms, and the results show that the scaling with numbers of atoms and processors is excellent, even for calculations using nearly the entire machine.

The CONQUEST code has been developed over the past five years, with EPSRC support via the U.K. Car-Parrinello consortium, one of the largest consortia working at CSAR. The methods are based on those used in the well-known CASTEP code, but the key difference is that the computer effort needed by CONQUEST is strictly proportional to the number of atoms, instead of CASTEP's increase as the cube of this number.

Both codes perform fully ab initio quantum calculations, using the density-functional theory of many-electron systems, and they can achieve similar precision. The CONQUEST code relies heavily on the parallel manipulation of sparse matrices, and CSAR support in optimising that part of the code was vital in making the code run efficiently on turing.

As an illustration of the excellent scaling with number of atoms, the graph in figure 1 below shows timings for calculations on silicon systems containing from 4096 to 12288 atoms run on 512 processors of turing. Remarkably, the scaling with number of atoms is actually slightly sub-linear. Other tests show that, in terms of Amdahl's law, the unparallelised part of the code is completely negligible when we run it on turing. With further code improvements now on the way, we are optimistic that ab initio calculations on systems of up to 50,000 atoms will be feasible in the near future - an increase of nearly two orders of magnitude over what was possible before!

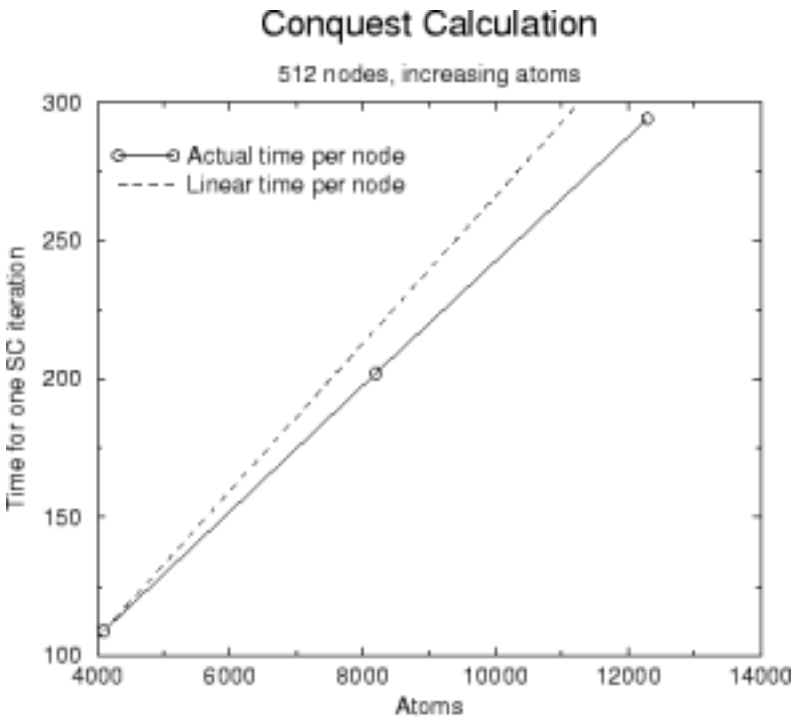

Figure 1: Time for CONQUEST calculation on systems of silicon atoms containing from 4096 to 12288 atoms run on 512 processors of turing, showing the slightly sub-linear increase of time with number of atoms.