Choosing hardware for Mauve

Occasionally a heavy Mauve addict asks what sort of compute hardware would be ideal for computing Mauve alignments. The answer depends in part on the length and number of sequences being aligned. When aligning a large number of sequences, Mauve is both compute intensive and memory intensive. When aligning very large genomes such as mammalian genomes, Mauve can utilize multiple independent hard drives to perform external sort operations.

PC hardware

The single most important hardware feature for Mauve is speed of memory access. CPUs that have an integrated memory controller will run Mauve faster. Examples of such CPUs include Intel Core i5, Core i7, and any CPU made by AMD since about 2003. Total memory capacity also affects the number and to a lesser extent the size of genomes that are practical to align. Aligning 20 or more 5Mbp-long genomes currently (2009) requires 5-10GB of memory, but two or three genomes of such size requires only a few hundred MB of RAM.

Mac hardware

Since its transition to Intel CPUs, Apple hardware is very similar to Dell hardware. Apple offers Mac Pro machines which are comparable to the eight-core machine described in the previous section.

Other hardware

Companies such as IBM, SGI, Sun, and others manufacture large memory model, multi-CPU machines. Such hardware is typically much more expensive than a PC, but usually has the advantage of high-throughput memory I/O from each CPU. Although pre-compiled binaries are not available for these architectures, the aligner source code can be compiled using gcc. Other compilers have not been tested.