Processor Research Team

K computer
Computer simulations create the future

The performance of parallel processing supercomputers can be raised by increasing the number of compute nodes employed. There are two broad approaches used to raise performance: the “weak scaling” approach employs more nodes to deal with a larger portion of the computation task, while maintaining the same computation time; the “strong scaling” approach aims to attain faster computation for the same sized task by increasing the number of nodes used. The latter is more difficult to accomplish. In order to realize strong scaling performance, we seek to fine-tune application programs and improve the efficiency of a processor, a major component of the compute node, with the aim of achieving full utilization of the K computer and to set the course for a next-generation supercomputer.

Team Leader
Makoto TAIJI

What Kinds of Computation Require Strong Scaling Performance?

The method that enables a task to be solved faster by increasing the number of compute nodes is called strong scaling performance. A typical example of calculations that require strong scaling performance is the molecular dynamics calculation of proteins. The calculation is implemented by moving atoms of molecules step by step according to Newton's equation of motion. Since each successive computation must await the result of the previous computation, it is difficult to attain fast computation by increasing the number of compute nodes. Furthermore, although a protein molecule can be divided into small units of data, each of which is assigned to one of 80,000 compute nodes of the K computer so as to attain fast computation, dividing a molecule into such small units is not effective.

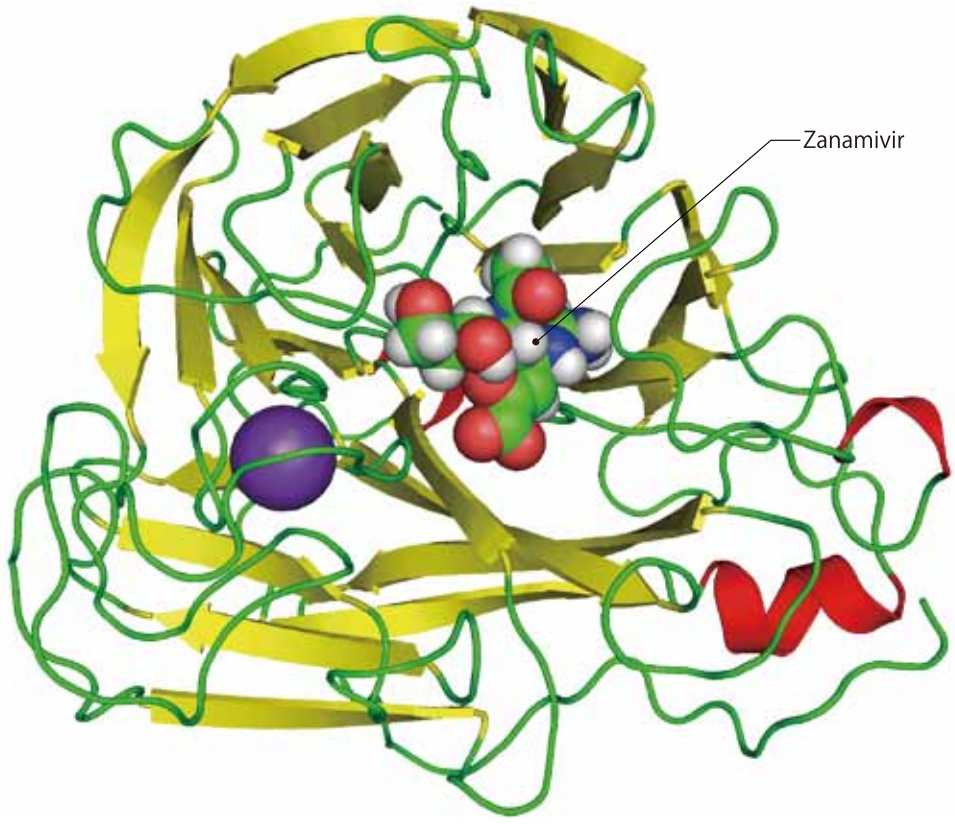

Shown opposite is the structure of neuraminidase (a protein that plays a role in influenza virus infection) in a complex interaction with Zanamivir (an inhibitor of the neuraminidase). The image is a snapshot of on-going change over the course of time, as simulated by molecular dynamics calculation.

Resolving problems associated with parallel processing supercomputers by improving compute node architecture

Program Optimization

Pursuing strong scaling performance can increase the proportion of compute nodes in a non-parallel state, which remain idle until receiving results from other compute nodes. In order to minimize this, program optimization is essential.

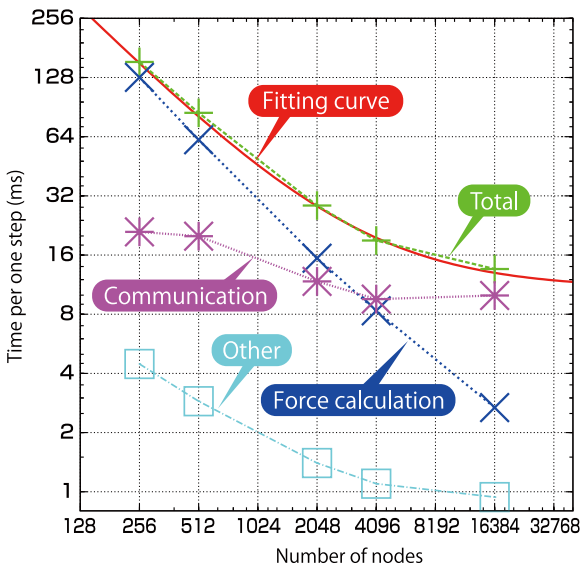

A molecular dynamics calculation (using an existing program) of a molecular system consisting of more than 1.6 million atoms resulted in 0.03% of the compute nodes falling into a non-parallel state. At this rate of ineffectiveness, overall computation speed will plateau at around 4,000 nodes and not increase further, no matter how many more nodes are added.

Developing a New Processor

We are using system-on-chip technology to develop a new processor in which all of the general-purpose cores, application-specific cores (accelerators), memory units, and networks are mounted on a single semiconductor chip. This processor will minimize delay during computation. The accelerators can also be adjusted, depending on the kind of computation being performed.