



Preface

Special Issue on the K computer

Yuji Oinaga
Head of Next Generation Technical Computing Unit
Fujitsu Ltd.

In 2006, Fujitsu joined “the Next-Generation Supercomputer Project” being carried out by RIKEN. Its objective was to design, build, and set up the world’s fastest and most advanced computer, with a speed of 10 petaflops, more than 35 times the performance of the top supercomputer in the world at that time. The project was considered a “key technology of national importance,” serving as a foundation of national scientific technology and industrial competitive strength. Its completion was specifically designated a long-term strategic national goal by Japan’s Council for Science and Technology Policy. This supercomputer came to be called the “K computer”.^{note)}

The K computer was not developed just to achieve a high level of performance using a specific benchmark program. The aim was threefold: to provide a next-generation supercomputer platform on which many applications could be developed, to enable many people to obtain beneficial research results by using those applications, and to achieve breakthroughs in supercomputer development and usage technologies.

With this in mind, RIKEN and Fujitsu set out to develop a practical and useful supercomputer with a high performance target of 10 PFLOPS while having to solve a number of problems related to power consumption and installation space restriction. Along the way, however, the development process was affected by a number of factors,

note) “K computer” is the English name that RIKEN has been using for the supercomputer of this project since July 2010. “K” comes from the Japanese word “Kei,” which means ten peta or 10 to the 16th power.

such as a change in system configuration from a hybrid system of scalar computers and vector computers to one with only scalar computers, debates on revising the project, and manufacturing centers that were adversely affected by the Great East Japan Earthquake. Despite these severe challenges, the project was nevertheless able to proceed due to the full support of many intellectuals in various fields and the concerted efforts of everyone involved, who shared a firm resolution to get the mission done and achieve the ultimate goal.

In June 2011, the K computer, despite being only partially completed, was awarded the No. 1 position on the TOP500 List that ranks supercomputer performance. This ranking marked the first time in seven years that a Japanese supercomputer had taken the top spot. Then, in November of the same year, the K computer was again ranked No. 1 as it became the first system in the world to exceed a measured performance of 10 PFLOPS—its original target value. Around the same time, the K computer also demonstrated high performance in practical fields by winning the Gordon Bell Peak Performance prize and receiving a No. 1 ranking in all four benchmarks of the HPC Challenge Award.

This special issue provides a broad overview of the K computer. It outlines the project, describes RIKEN and Fujitsu's approach to developing this supercomputer, and covers hardware, software, and applications, including key papers written by RIKEN personnel.

The K computer is scheduled to begin shared-use operations in the fall of 2012. I am convinced that many beneficial results will be forthcoming through the development of many applications and the achievement of high execution performance as originally planned. Looking to the future, Fujitsu plans to achieve a supercomputer with even higher performance and to develop even more beneficial applications. We will accomplish this by refining the technologies introduced in this special issue and leveraging our strengths in developing entire computer systems from CPU to software. On the occasion of this special issue, I would like to extend my deep appreciation to all concerned for their valuable cooperation in the development of the K computer and their deep interest in its core technologies. I look forward to their continued guidance and collaboration.