

Accelerator Utilization Technology That Cuts Costs, Reduces Power Consumption, and Shrinks Hardware Footprint

ISHIZAKA Kazuhisa, TAKENAKA Takashi, MORIYOSHI Tatsuji

Abstract

As cloud platforms are increasingly required to support social infrastructure services and data crunching, they need to be able to analyze large amounts of data collected from the real world, such as movies, images, audio, and sensor data at high speed and with low latency, and then feed the results back to the real world in a timely manner. Conventionally, distributed processing is performed using a large number of general-purpose servers, but there is a problem with this solution - equipment costs, power consumption, and space requirements simply become too great. In this paper, we discuss accelerator utilization technology that makes it possible to process and analyze massive volumes of data in the cloud computing, while actually lowering costs, reducing power consumption, and minimizing the platform footprint. We will also be introducing proposals for future developments in this area.



accelerator, many-core, FPGA, scheduling, CyberWorkBench, SQL

1. Introduction

In recent years, development of social infrastructure through information and communications technology (ICT) has been accelerating as more and more websites and mobile apps integrate social functionality into their services - all in an effort to make the user's life safer, more convenient, and more fascinating. For example, ICT is expected to contribute to various social solutions in areas such as public safety, infrastructure condition monitoring, fault sign detection, and accident prevention.

To support the ongoing development of social infrastructure, cloud platforms need to be able to process and analyze large amounts of data collected from the real world, such as movies, images, audio, and sensor data. And they need to be able to do it at high speed and with low latency so that the analysis results can be fed back to the real world in a timely manner.

This paper discusses accelerator utilization technology that achieves the high-speed, low-latency analytical processing of large-volume data at low cost with low power consumption and a small footprint.

2. Accelerator Utilization in Data Analysis

To handle large-volume data analysis with the general-purpose CPUs found in ordinary servers requires a lot of servers, resulting in increased equipment costs, increased power consumption, and increased installation space (footprint). Finding a way to lower the costs and power consumption of ICT systems is a critical issue for IT companies that wish to strengthen their competitiveness, not to mention an issue of increasing social and environmental concern. In addition to the problems of cost, power consumption, and space requirements, when distributed processing is performed with numerous servers, the time required for communication between servers and data input and output increases, resulting in extended delay time.

Today, there is growing interest in utilizing accelerators that can perform some processing with much better cost performance and power efficiency than general-purpose CPUs. These accelerators include a many-core coprocessor/GPU that conglomerates processor cores - ranging from a few dozen to a few hundred - into a single chip in order to achieve high performance by taking advantage of parallel processing, as well as a field-programmable gate array (FPGA) that combines high-speed processing equivalent to that of dedicated hardware

with the flexibility of software. The incorporation of the accelerators enhances the analytical processing capability of each server and makes it possible to perform processing with fewer processors, thereby improving the cost, power, footprint, and delay.

Nevertheless, since accelerators have significantly different structures and characteristics from those of general-purpose CPUs, the following problems - which normally are non-existent - need to be solved before employing accelerators for cloud platforms.

- Maximization of the performance of overall systems in which processors with different properties, such as general-purpose CPUs and many-core coprocessors, are both used.
- Development of software will be complicated by the addition of accelerators and will require improvement in design and productivity.

Below, we discuss NEC's efforts to solve these problems.

3. Many-Core Coprocessor Utilization Technology

The Intel Xeon Phi Coprocessor¹⁾ (hereinafter referred to as Xeon Phi) is a many-core coprocessor (hereinafter referred to as many-core) that is increasingly being used as an accelerator to meet the demands of high-performance computing (HPC). Featuring more than 60 processor cores based on the x86 architecture that forms the basis for the general-purpose CPUs powering conventional servers and boasting excellent power efficiency, the Xeon Phi is capable of delivering processing capability several times that of a general-purpose CPU with equivalent power consumption. Easily installed in a server extension slot, the Xeon Phi is optimized for high-density computing, delivering higher aggregate performance than can be obtained merely by increasing the number of servers, while the smaller footprint makes it suitable for applications where sever installation areas are limited.

Typically, accelerators are used to execute processing tasks at high speed. However, to maximally exploit a system's performance, it is essential to use both the accelerator and host processor (CPU). With source code fully compatible with CPUs, Xeon Phi coprocessors can be used together with standard CPUs to optimize any workload. By taking advantage of this capability, we have developed technology that significantly reduces costs while offering dramatic overall performance gained by flexibly distributing processing tasks both the CPU and the many-core.

3.1 Offload Scheduling

When multiple programs offload processing to the many-core, the load on the many-core creates a bottleneck and idle time is generated in the CPU, preventing maximization of

system performance (Fig. 1 a).

To solve this problem, we have developed offload scheduling technology.²⁾ Featuring dynamic decision-making, this technology balances offloading according to load conditions, determining whether processing will be executed in the many-core or on the CPU. When the many-core is busy, total processor usage efficiency can be improved by executing the offload processing on the CPU as shown in Fig. 1 b.

When technology like this is used, total execution time can be reduced by balancing the loads on the CPU and many-core as shown in Fig. 2.

3.2 Virtual Pipeline Execution Model

In offload execution, specific processing is executed with the many-core as shown in Fig. 3 a. The efficiency of this can vary depending on the application.

For example, when performing super-resolution image processing (to increase the resolution) that consists of multiple image processing tasks, processing throughput for each task must be considered when allocating task execution in order to balance the load between the CPU and the many-core. This is inefficient because redeveloping relevant application will be

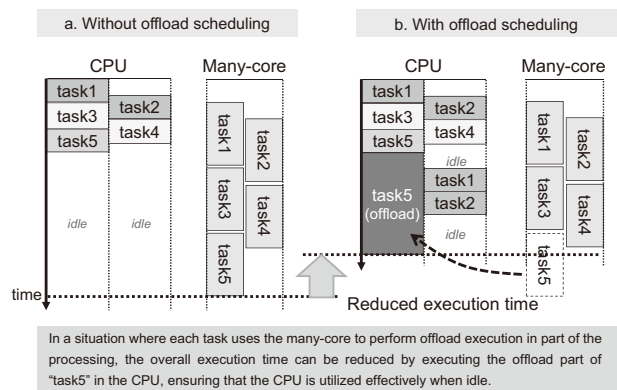


Fig. 1 Offload scheduling technology.

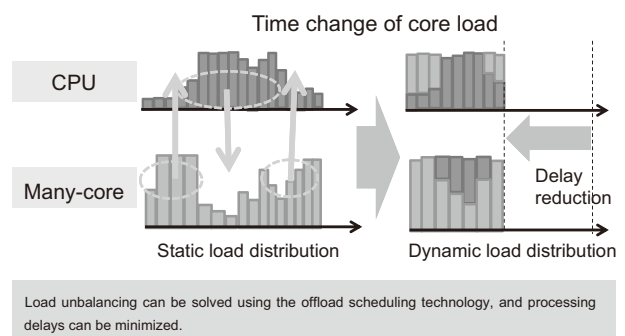


Fig. 2 Dynamic load distribution between CPU and many-core.

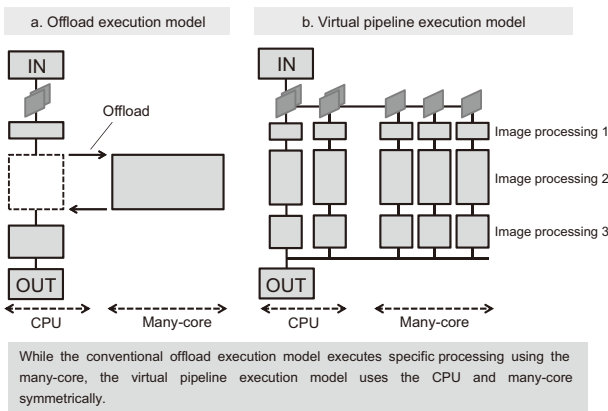


Fig. 3 Virtual pipeline execution model.

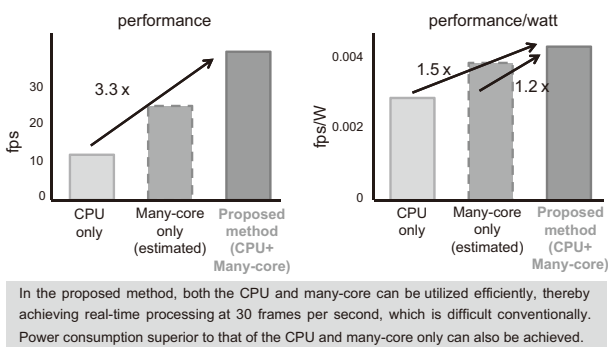


Fig. 4 Performance of super-resolution processing using virtual pipeline execution model.

required each time the image processing parameters are altered or a new image processing task is going to be added.

Our solution is a virtual pipeline execution model³⁾ that allocates multiple pipelines to each processor while putting a series of image processing tasks together as a bundle of pipelines so that those pipelines are virtually shown as one pipeline (Fig. 3 b). This system is flexible because the load balance can be adjusted by changing the number of frames of images to be transmitted to each pipeline.

Fig. 4 shows the performance when super-resolution processing is executed using this execution model and the power consumption. From this, it can be seen that the execution model makes it possible to utilize both the CPU and many-core at almost 100 percent, delivering higher performance and greater power consumption efficiency than when the CPU or many-core is used on its own.

4. Utilization Technology of FPGA

A field-programmable gate array (FPGA) is a programmable IC which has the characteristics of both hardware and software,

in that it is capable of massively parallel processing (MPP) and can be reprogrammed to perform different tasks. It features low latency, predictable processing time, dedicated arithmetic processing performance, and low power consumption.

FPGAs have conventionally been used as a substitute for customized large-scale integrated circuit chips (LSIs). For example, they have been used for prototyping of customized LSIs, as well as being incorporated in systems that require low power consumption such as digital home electric appliances, and in mobile phone base stations that require high-speed signal processing.

Thanks to their low latency and low power consumption (performance/power ratio), FPGAs are also being increasingly applied in data centers and cloud platforms. For example, they have reportedly been utilized to speed up stock exchange operation where low latency in the order of milliseconds is required. They are also used to improve the efficiency of search engines at large-scale data centers.

However, in order to utilize FPGAs in cloud platforms, the problem of low productivity in the basic FPGA design must be overcome. Logical designing (programming) of FPGAs needs to be performed with a dedicated language (hardware description language, the typical examples of which are VHDL and Verilog-HDL), while being conscious of the hardware at a low abstraction level (as if writing an assembler), so it can take specialized hardware designers a few weeks to a few months to carry out the task.

To solve this problem, we are researching and developing design technology that will make the high-speed processing/low latency of the FPGA compatible with its design productivity. For example, a design tool called CyberWorkBench allows software designers to perform logical designing (programming) of FPGAs using the C programming language, which is one of the most common languages used by software designers to write an algorithm. Also underway is R&D into the technology to make it possible to design high-speed circuits on FPGAs using SQL, which is widely used for complex event processing that extracts meaningful information from sensor data - demand for which is expected to increase explosively with the advent of the Internet of Things (IoT) era.⁴⁾ This technology even makes it possible for data analysis personnel who have no knowledge of hardware design to code FPGAs using SQL and speed up processing with just a few hours of programming work.

5. Conclusion

We have introduced accelerator utilization technology that enables the construction of cloud platforms that cost less, use less power, and have a smaller footprint, while being capable of advanced analysis processing of large-volume data. As demand for increased processing data volume, increased complexity, and advanced analysis processing will continue to

grow, we are committed to continuing our efforts to develop versatile optimal accelerator application technology, as well as collaboration technology to facilitate the combined use of different types of accelerators.

* Intel and Xeon Phi are trademarks of Intel Corporation in the U.S. and other countries.

Reference

- 1) Intel Corporation: Intel Xeon Phi Coprocessor Datasheet, November 2012, reference number 328209-001EN
- 2) T. Miyamoto, K. Ishizaka, T. Hosomi (NEC), "A Dynamic Offload Scheduler for spatial multitasking on Intel Xeon Phi Coprocessor", The 18th Workshop on Synthesis And System Integration of Mixed Information Technologies (SASIMI2013), Oct.2013
- 3) K. Ishizaka, et al.: Power Efficient Realtime Super Resolution by Virtual Pipeline Technique on a Server with Manycore Coprocessors, CoolChips XVI, April 2014
- 4) T. Takenaka, M. Takagi, and H. Inoue: A Scalable Complex Event Processing Framework For Combination of SQL-based Continuous Queries and C/C++ Functions, IEEE International Conf. on Field Programmable Logic and Applications, pp.237-242, August 2012

Authors' Profiles

ISHIZAKA Kazuhisa

Assistant Manager
Green Platform Research Laboratories

TAKENAKA Takashi

Principal Researcher
Green Platform Research Laboratories

MORIYOSHI Tatsuji

Principal Researcher
Green Platform Research Laboratories

The details about this paper can be seen at the following.

Related URL:

NEC develops big data software technology that enables real-time processing at world-class speed
http://www.nec.com/en/press/201312/global_20131203_02.html

NEC Technology Enables the Design of High-Speed Big Data Processing Hardware in 1/50 the Time
http://www.nec.com/en/press/201208/global_20120831_01.html

CyberWorkBench
<http://www.nec.com/en/global/prod/cwb/index.html>

Information about the NEC Technical Journal

Thank you for reading the paper.

If you are interested in the NEC Technical Journal, you can also read other papers on our website.

Link to NEC Technical Journal website

Japanese

English

Vol.9 No.2 Special Issue on Future Cloud Platforms for ICT Systems

Remarks for Special Issue on Future Cloud Platforms for ICT Systems

NEC's Approach to Orchestrating the Cloud Platform

NEC C&C cloud platforms ? NEC Cloud IaaS Services

Portal Services Integrate Multi-Cloud Environments

A Hybrid Server Hosting Which Have Broader Range of Applications

Network Service That Offers a Versatile Network Environment

Dependable Security Service That Takes Advantage of Internal Control Methodology

Data Center Service That Supports Cloud Infrastructure

Products and latest technologies supporting NEC C&C cloud platforms

MasterScope Virtual DataCenter Automation - Entire IT System Cost Optimization by Automating the System Administration

Integrated Operation and Management Platform for Efficient Administration by Automating Operations

Micro-modular Server and Phase Change Cooling Mechanism Contributing to Data Center TCO Reduction

iStorage M5000 Providing a High-Reliability Platform for the Cloud Environment

The iStorage HS Series Features the Superior Data Compression and High-Speed Transmission Capabilities that are Essential Functions of Big Data Storage

SDN Compatible UNIVERGE PF Series Supports Large-Scale Data Centers by Automating IT System Management

Phase Change Cooling and Heat Transport Technologies Contribute to Power Saving

Future technology for NEC's C&C cloud platforms

Accelerator Utilization Technology That Cuts Costs, Reduces Power Consumption, and Shrinks Hardware Footprint

Scalable Resource Disaggregated Platform That Achieves Diverse and Various Computing Services

Support Technology for Model-Based Design Targeted at a Cloud Environment

Cloud-based SI for Improving the Efficiency of SI in the Cloud Computing by Means of Model- Based Sizing and Configuration Management


Big Data Analytics in the Cloud - System Invariant Analysis Technology Pierces the Anomaly -

Case Studies

Using Cloud Computing to Achieve Stable Operation of a Remote Surveillance/Maintenance System Supporting More Than 1,100 Automated Vertical Parking Lots throughout Japan

Meiji Fresh Network's Core Business Systems are Transitioned to NEC Cloud IaaS NEC's Total Support Capability is Highly Evaluated.

Sumitomo Life Insurance Uses NEC's Cloud Infrastructure Service to Standardize IT Environments across the Entire Group and Strengthen IT Governance



Vol.9 No.2

June, 2015

Special Issue TOP

NEC Information

NEWS

2014 C&C Prize Ceremony
