Chapter 40 Database Techniques for New Hardware

Xiongpai Qin Renmin University of China, China

Yueguo Chen Renmin University of China, China

ABSTRACT

In the last decade, computer hardware progressed by leaps and bounds. The advancements of hardware include the application of multi-core CPUs, use of GPUs in data intensive tasks, bigger and bigger main memory capacity, maturity and production use of non-volatile memory, etc. Database systems immediately benefit from faster CPU/GPU and bigger memory and run faster. However, there are some pitfalls. For example, database systems running on multi-core processors may suffer from cache conflicts when the number of concurrently executing DB processes increases. To fully exploit advantages of new hardware to improve the performance of database systems, database software should be more or less revised. This chapter introduces some efforts of database research community in this aspect.

INTRODUCTION

In the last decade, computer hardware progresses with leaps and bounds. The advancements of hardware include: widely application of multi-core CPUs, using of GPUs in data intensive tasks, bigger and bigger main memory capacity, maturity and production use of non-volatile memory etc.

Database systems immediately benefit from faster CPU/GPU and bigger memory, and run faster. However, there are some pitfalls. For example, database systems running on multi-core processors may suffer from cache conflicts when the number of concurrently executing DB processes increases. To fully exploit advantages of new hardware to improve the performance of database systems, database software should be more or less revised (Ailamaki, 2004).

This chapter introduces some efforts of database research community in this aspect. And the following section gives a brief introduction to some new hardware technologies that database systems could utilized.

DOI: 10.4018/978-1-5225-7598-6.ch040

BACKGROUND

Multi-Core CPU

Improving the performance of CPU through increasing its clock frequency becomes more and more difficult. Researchers and engineers bring forth multi-core technology. In a typical multi-core CPU, 2, 4, 8 or more cores are integrated on one chip. Putting several cores on one die allows for higher communication speeds between the cores, which will benefit many computing tasks.

The cores have their own private caches (Level1 Cache, or L1 Cache), and share some larger but slower caches (L2 Cache). They access the shared main memory for parallel data processing. Database systems are basically multi-threaded, and they can benefit from multi-core CPUs without any modification. However, to fully utilize the cores to boost database performance, there is much work to do.

GPU for General Tasks

GPU is traditionally used to accelerate the specific task of graphic rendering. GPU venders have integrated many computing units in a single die, and optimized the bandwidth to process large volume of graphic data. The highly parallelism of GPUs is exploited to speed up data intensive tasks as well, and GPU becomes GPGPU (General Purpose GPU).

Since GPU is designed primarily for graphic processing tasks instead of general tasks, the architecture of a GPU is rather different from CPU. It has its own unique thread hierarchy and memory hierarchy, which should be taken into account when using GPU for data processing tasks.

Bigger Memory

The price of memory is going down, now people can install more memory in a single server. It is rather common for a single server to possess a memory capacity as large as hundreds of gigabytes, or even up to terabytes. For moderate-size applications, it is possible to load the whole dataset into memory for fast access.

Non-Uniform Memory Access

Non-Uniform Memory Access (NUMA) machine is becoming more and more common. The NUMA architecture consists of a small number of processors, each having its own memory and I/O channels. Each group of CPUs is called a 'node'. Memory that is local to a node is called local or near memory, while memory outside of a node is called foreign or far memory (Golding, 2010). Accessing foreign memory is much slower than accessing local memory. NUMA architecture requires changes of memory management.

Non-Volatile Memory, Solid State Disks, and Changes of the Memory Hierarchy

Non-volatile memory (NVRAM) is a type of computer memory that can keep the information even after the power if turned off. There are several flavors of NVRAM, including flash memory, and phase change memory (PCM) etc.

15 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/database-techniques-for-new-hardware/214642

Related Content

A Model for Mobile Learning Service Quality in University Environment

Nabeel Farouq Al-Mushashaand Shahizan Hassan (2009). *International Journal of Mobile Computing and Multimedia Communications (pp. 70-91).* www.irma-international.org/article/model-mobile-learning-service-quality/4064

The Economics of Internetization

Constantine E. Passaris (2019). Advanced Methodologies and Technologies in Network Architecture, Mobile Computing, and Data Analytics (pp. 1714-1729). www.irma-international.org/chapter/the-economics-of-internetization/214734

Real-Time 3D Design Modelling of Outdoor Structures Using Mobile Augmented Reality Systems

Wayne Piekarski (2009). Mobile Computing: Concepts, Methodologies, Tools, and Applications (pp. 937-951).

www.irma-international.org/chapter/real-time-design-modelling-outdoor/26558

Social Networks on the Internet: Twitter Coverage of the Exile of the Peruvian Indigenous Leader Alberto Pizango to Nicaragua

Débora Pereira (2011). ICTs for Mobile and Ubiquitous Urban Infrastructures: Surveillance, Locative Media and Global Networks (pp. 108-127).

www.irma-international.org/chapter/social-networks-internet/48347

A Local Statistical Information Active Contour Model for Image Segmentation

Shigang Liu, Yali Peng, Guoyong Qiuand Xuanwen Hao (2014). *International Journal of Mobile Computing and Multimedia Communications (pp. 33-49).*

www.irma-international.org/article/a-local-statistical-information-active-contour-model-for-image-segmentation/128999