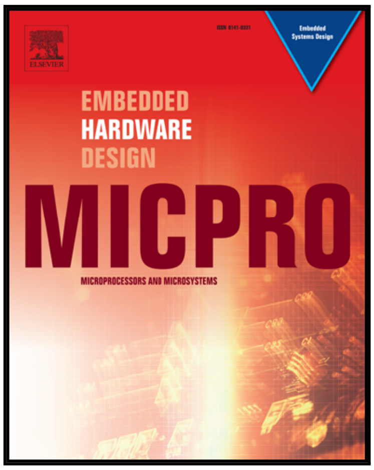

Application-Specific Architectures for Energy-Efficient Database Query Processing and Optimization

Sebastian Haas, Stefan Scholze, Sebastian Höppner, Annett Ungethüm, Christian Mayr, René Schüffny, Wolfgang Lehner, Gerhard Fettweis

PII: S0141-9331(17)30104-7
DOI: 10.1016/j.micpro.2017.10.002
Reference: MICPRO 2620

To appear in: *Microprocessors and Microsystems*

Received date: 20 February 2017
Revised date: 8 August 2017
Accepted date: 11 October 2017

Application-Specific Architectures for Energy-Efficient Database Query Processing and Optimization⋆

Sebastian Haas*, Stefan Scholze†, Sebastian Höppner†, Annett Ungethüm‡, Christian Mayr†, René Schüffny†, Wolfgang Lehner‡, Gerhard Fettweis*

*Vodafone Chair Mobile Communications Systems, Technische Universität Dresden, Germany
†Chair for Highly-Parallel VLSI-Systems and Neuro-Microelectronics, Technische Universität Dresden, Germany
‡Database Systems Group, Technische Universität Dresden, Germany
Center for Advancing Electronics Dresden (cfaed)
Email: {first.last}@tu-dresden.de

***Abstract*—Data processing on a continuously growing volume of data and the increasing power restrictions have become an ubiquitous challenge in our world today. Besides parallel computing, a promising approach to improve the energy efficiency of current systems is to integrate specialized hardware. This paper presents two application-specific architectures to accelerate basic database operators frequently used in modern database systems: an extended instruction set based on a given Cadence Tensilica processor (ASIP) and a comparable application-specific integrated circuit (ASIC). The ASIP is implemented in a system-on-chip and manufactured in a 28 nm CMOS technology to realize measurements of performance and power consumption. Furthermore, the comparison with the ASIC blocks allows to quantify the results with the ASIP approach in terms of throughput, area, and energy efficiency as well as to discuss the capabilities and limitations when accelerating selected database operators.**

***Keywords*—Application-Specific Instruction-Set Processor, Application-Specific Integrated Circuit, System-on-Chip, Database Systems, Query Processing**

I. Introduction

Today, we live in the big data era which demands database systems and their underlying solid-state circuits to deal with an ever growing amount of data. Modern database systems deploy general-purpose CPUs for data processing tasks such as query optimization and execution. The increase of the processor's clock frequencies by exploiting smaller CMOS technologies has stalled since about 2005 [2]. Although the transistor count per chip still increases (Moore's law [3]), higher static power losses hinder the increase of the clock frequency and thus to improve the processor's single-threaded performance. The approach to use parallelism within multiprocessor systems to provide the required performance is limited as well. The solution to not exceed the power budget is turning off parts of the chip which is referred to as dark silicon [4].

When investigating novel solutions, integrating customized hardware is one promising approach to address the aforementioned challenges. Application-specific integrated circuits (ASICs) can be highly adapted by exploiting dedicated hardware with tailored vectorization methods and parallel concepts. However, besides a complex and time-consuming development, ASICs are designed for only one particular use case. In contrast, general-purpose processors provide this needed flexibility. Hence, the promising combination of both leads to an application-specific instruction-set processor (ASIP). ASIPs contain a base instruction set and can be extended by additional hardware components accessed via new instructions.

This paper mainly pursues these two approaches: First, we use a customizable Cadence Tensilica processor and extend the core with additional instructions specialized for fundamental database operators used in query processing. Second, we design ASIC blocks with comparable features such as register width and memory bandwidth performing a subset of the algorithms. In comparison to the ASIP, we expect that the ASIC can achieve further energy enhancements. However, the ASIP is more flexible since it still includes the general-purpose instruction set to support non-accelerated algorithms.

A. Query Processing in DBMS

A database management system (DBMS) is a complex software program and is responsible for defining, storing, processing, and sharing data [5]. Query processing is one of the key tasks in a DBMS and includes the complete procedure from parsing and optimizing a given query towards the execution of the underlying database operators such as selection, sorting, and join. We refer to two essential requirements for an effective query processing. (1) Query throughput: The performance has to be adapted to the continuously increasing amount of data and complexity produced by applications such as the Internet of Things (IoT) [6], deep learning [7], and 5G mobile networks [8]. (2) Query latency: The response time from the database system to the user should be minimized. This

⋆This paper is an extended version of the conference paper "A Database Accelerator for Energy-Efficient Query Processing and Optimization" [1] published at NorCAS'16.