

SEVENTH FRAMEWORK PROGRAMME
THEME ICT-2013.3.4
Advanced Computing, Embedded and Control Systems



Execution Models for Energy-Efficient Computing Systems
Project ID: 611183

D2.1

Models for energy consumption of data structures and algorithms

Phuong Ha, Vi Tran, Ibrahim Umar,
Philippas Tsigas, Anders Gidenstam, Paul Renaud-Goud,
Ivan Walulya, Aras Atalar



Date of preparation (latest version): 31.08.2014
Copyright© 2013 – 2016 The EXCESS Consortium

The opinions of the authors expressed in this document do not necessarily reflect the official opinion of EXCESS partners or of the European Commission.

DOCUMENT INFORMATION

Deliverable Number	D2.1
Deliverable Name	Models for energy consumption of data structures and algorithms
Authors	Phuong Ha Vi Tran Ibrahim Umar Philippas Tsigas Anders Gidenstam Paul Renaud-Goud Ivan Walulya Aras Atalar
Responsible Author	Phuong Ha e-mail: phuong.hoi.ha@uit.no Phone: +47 776 44032
Keywords	High Performance Computing; Energy Efficiency
WP/Task	WP2/Task 2.1
Nature	R
Dissemination Level	PU
Planned Date	01.09.2014
Final Version Date	31.08.2014
Reviewed by	Christoph Kessler (LIU) David Moloney (Movidius)
MGT Board Approval	YES

DOCUMENT HISTORY

Partner	Date	Comment	Version
UiT (P. Ha)	11.06.2014	Deliverable skeleton	0.1
CTH (A. Gidenstam)	22.07.2014	Added outline for concurrent data structures	0.2
UiT (V. Tran)	02.08.2014	Added sections 1.1, 1.3, 2.2, 3.1, 4.2	0.3
UiT (I. Umar, P. Ha)	05.08.2014	Revised D2.1 and added sections 1.2.1, 3.4, 4.1.1	0.4
CTH(A. Gidenstam)	05.08.2014	Added sections 1.2, 1.5, 2,1, 2.3, 3.2, 3.3, 3.5	0.5
UiT (P. Ha)	11.08.2014	Added Executive summary and sections 1.6 Incorporated Chalmers and UiT inputs	0.6
CTH, UiT	29.08.2014	Integration of internal review comments	0.7
CTH, UiT	30.08.2014	Final version	1.0

Abstract

This deliverable reports our early energy models for data structures and algorithms based on both micro-benchmarks and concurrent algorithms. It reports the early results of Task 2.1 on investigating and modeling the trade-off between energy and performance in concurrent data structures and algorithms, which forms the basis for the whole work package 2 (WP2). The work has been conducted on the two main EXCESS platforms: (1) Intel platform with recent Intel multi-core CPUs and (2) Movidius embedded platform.

Executive Summary

Computing technology is currently at the beginning of the disruptive transition from petascale to exascale computing (2010 - 2020), posing a great challenge on energy efficiency. High performance computing (HPC) in 2020 will be characterized by data-centric workloads that, unlike those in traditional sequential/parallel computing, are comprised of big, divergent, fast and complex data. In order to address energy challenges in HPC, the new data must be organized and accessed in an energy-efficient manner through novel fundamental data structures and algorithms that strive for the energy limit. Moreover, the general application- and technology-trend indicates finer-grained execution (i.e. smaller chunks of work per compute core) and more frequent communication and synchronization between cores and uncore components (e.g. memory) in HPC applications. Therefore, not only concurrent data structures and memory access algorithms but also synchronization is essential to optimize the energy consumption of HPC applications. However, previous concurrent data structures, memory access algorithms and synchronization algorithms were designed without energy consumption in mind. The design of energy-efficient fundamental concurrent data structures and algorithms for inter-process communication in HPC remains a largely unexplored area and requires significant efforts to be successful.

Work package 2 (WP2) aims to develop interfaces and libraries for energy-efficient inter-process communication and data sharing on the new EXCESS platforms integrating Movidius embedded processors. In order to set the stage for these tasks, WP2 needs to investigate and model the trade-offs between energy consumption and performance of data structures and algorithms for inter-process communication, which is Task 2.1. The energy models are developed in close cooperation with WP1 to ensure that they will be compatible with the energy modeling method of WP1.

The early result of Task 2.1 (PM1 - PM36) on investigating and modeling the trade-off between energy and performance in concurrent data structures and algorithms, as available by project month 12, are summarized in this report. The main contributions are the following:

- An improved and extended energy model for the CPU-based platform based on the model presented in EXCESS D1.1 [49]. This model decomposes the power into static, active and dynamic power, while classifying CPU-based platform components into three groups: CPU, main memory and uncore (e.g. shared cache, IMC, PCU, HA, etc.) (cf. Sec. 3.1). The experiment results confirm that static power is constant while active power depends on the frequency, the number of socket and not on the operations. Dynamic power is decomposed into dynamic CPU, dynamic memory and dynamic uncore power. Dynamic CPU power depends on the frequency and the operation type. It also shows almost linear behaviors to the number of threads. Dynamic power of memory and uncore components relate to the locality and bandwidth requirement of the implementations.
- A new power model for the Movidius Myriad platform that is able to predict power consumption of our micro-benchmarks with $\pm 4\%$ margin of measured power consumption

on the real platform (cf. Sec. 3.2). The new power model confirms the experimental power analysis of concurrent data structures such as concurrent queues: the dynamic power consumption is proportional to the number of SHAVE (Streaming Hybrid Architecture Vector Engine) processors used.

- A case study on how to choose the most suitable implementations for a multi-variant shared data structure in a certain application and context; and the prediction of the energy efficiency of different queue implementations through two metrics, namely throughput and power (cf. Sec. 4.1). The case study shows that the energy-efficiency is mainly ruled by the contention on the queue, which impacts both throughput and memory power dissipation.
- Implementation and evaluation of several different concurrent queue designs for the Myriad1 platform using three synchronizations primitives: mutex, message passing over shared variables and SHAVE FIFOs (a set of registers accessed in a FIFO pattern)(cf. Sec. 4.2). The valuations are performed on three metrics: execution time, power consumption and energy per operation. In terms of execution time, the implementation using mutex with two locks is the fastest and most scalable since it provides maximum concurrency. In terms of power, SHAVE FIFOs communication method is the most energy efficient. In terms of energy per operation, SHAVE FIFO implementation also consume the least energy.
- Investigation of the energy consumption and performance of concurrent data structures such as concurrent search trees (cf. Sec. 5). Based on our investigation, we have developed new locality-aware concurrent search trees called Δ Trees that are up to 140% faster and 80% more energy efficient than the state-of-the-art (cf. Sec. 5.6 and Sec. 5.7).

Contents

1	Introduction	9
1.1	Purpose	9
1.2	Concurrent Data Structures and Algorithms for Inter-process Communication	10
1.3	Micro-benchmarking	13
1.4	Metrics	14
1.5	Overview and Contributions	14
2	EXCESS Platforms and Energy Measurement Methodology	15
2.1	System A: CPU-based platform	16
2.1.1	System description	16
2.1.2	Measurement methodology for energy consumption	16
2.2	System B: Movidius Embedded Platform (Myriad1)	17
2.2.1	Myriad1 Platform Description	17
2.2.2	Measurement methodology	19
2.3	Significant Synchronization Hardware Differences Between the Two Systems	21
2.3.1	Fundamental synchronization primitives	22
3	Energy Models for EXCESS Platforms	23
3.1	Energy Models for CPU-based Platforms	23
3.1.1	General Power Model	23
3.1.2	Power Components Derivation	25
3.1.3	Dynamic CPU Power	26
3.1.4	Dynamic Memory and Uncore Power	27
3.1.5	Summary of Micro-Benchmarking for Power Modeling on CPU	32
3.2	Energy Models for Movidius Embedded Platforms	32
3.2.1	Description of Microbenchmarks	32
3.2.2	Movidius Power Model and Its Sanity Check	33
4	Modeling Energy Consumption of Concurrent Queue Implementations	35
4.1	Concurrent Queues on CPU-based Platform	35
4.1.1	Objective and Process	35
4.1.2	Throughput	38
4.1.3	Power Given Throughput Measurement	43
4.1.4	Complete Prediction	49
4.1.5	Description of the Implementations	51
4.1.6	Towards Realistic Applications: Mandelbrot Set Computation	53
4.2	Concurrent Queues on Movidius Embedded Platform	61
4.2.1	Concurrent Queue Implementations	61
4.2.2	Experimental Evaluation	64

5	Analysis of Energy Consumption of Concurrent Search Trees: ΔTrees	71
5.1	Introduction	71
5.1.1	I/O model	71
5.1.2	Ideal-cache model	72
5.2	Overview on the van Emde Boas layout	74
5.2.1	Static van Emde Boas (vEB) layout	75
5.2.2	Relaxed cache-oblivious model and dynamic vEB layout	76
5.3	Δ Tree implementation	78
5.3.1	Data structures	78
5.3.2	Function specifications	80
5.4	Balanced Δ Tree	86
5.4.1	<i>Map</i> instead of pointers	87
5.4.2	Concurrent and balanced tree	87
5.5	Heterogenous balanced Δ Tree	89
5.6	Performance evaluation	89
5.6.1	Experimental results	91
5.6.2	On the worst-case insertions	93
5.6.3	Performance on different upper-bounds <i>UB</i>	94
5.7	Energy consumption evaluation	94
5.8	Conclusion	97
6	Conclusion	98

1 Introduction

1.1 Purpose

In order to address energy challenges in HPC and embedded computing, data must be organized and accessed in an energy-efficient manner through novel fundamental data structures and algorithms that strive for the energy limit. Due to more frequent communication and synchronization between cores and memory components in HPC and embedded computing, not only concurrent data structures and memory access algorithms but also synchronization is essential to optimize the energy consumption. However, previous concurrent data structures, memory access algorithms and synchronization algorithms were designed without considering energy consumption. Although there are existing studies on the energy utilization of concurrent data structures demonstrating non-intuitive results on energy consumption, the design of energy-efficient fundamental concurrent data structures and algorithms for inter-process communication in HPC and embedded computing is not yet widely explored and becomes an challenging and interesting research direction.

EXCESS aims to investigate the trade-offs between energy consumption and performance of concurrent data structures and algorithms as well as inter-process communication in HPC and embedded computing. By analyzing the non-intuitive results, EXCESS devises a comprehensive model for energy consumption of concurrent data structures and algorithms for inter-process communication, especially in the presence of component composition. The new energy-efficient technology will be delivered through novel execution models for the energy-efficient computing paradigm, which consist of complete energy-aware software stacks (including energy-aware component models, programming models, libraries/algorithms and runtimes) and configurable energy-aware simulation systems for future energy-efficient architectures.

The goal of Work package 2 (WP2) is to develop interfaces and libraries for inter-process communication and data sharing on EXCESS new platforms integrating Movidius embedded processors, along with investigating and modeling the trade-offs between energy consumption and performance of data structures and algorithms for inter-process communication. WP2 also concerns supporting energy-efficient massive parallelism through scalable concurrent data structures and algorithms that strive for the energy limit, and minimizing inter-component communication through locality- and heterogeneity-aware data structures and algorithms.

The first objective of WP2 (Task 2.1) is to investigate and model the trade-off between energy and performance in concurrent data structures and algorithms. In order to model energy and performance, the analysis is conducted for non-intuitive results and their trade-offs to devise comprehensive models for energy consumption of concurrent data structures and algorithms of inter-process communication. The energy models are developed in close cooperation with WP1, ensuring that they will be compatible with the modeling method of WP1.

This report summarizes the early results of Task 2.1 on investigating and modeling the consumed energy of concurrent data structures and algorithms. The work of Task 2.1 forms

the theoretical basis for the whole work package.

1.2 Concurrent Data Structures and Algorithms for Inter-process Communication

Concurrent data structures are the data sharing side of parallel programming. Data structures give the means to the program to store data but also provide operations to the program to access and manipulate these data. These operations are implemented through algorithms that have to be efficient. In the sequential setting, data structures are crucially important for the performance of the respective computation. In the parallel programming setting, their importance becomes even more crucial because of the increased use of data and resource sharing for utilizing parallelism. In parallel programming, computations are split into subtasks in order to introduce parallelization at the control/computation level. To utilize this opportunity of concurrency, subtasks share data and various resources (dictionaries, buffers, and so forth). This makes it possible for logically independent programs to share various resources and data structures. A subtask that wants to update a data structure, say add an element into a dictionary, that operation may be logically independent of other subtasks that use the same dictionary.

Concurrent data structure designers are striving to maintain consistency of data structures while keeping the use of mutual exclusion and expensive synchronization to a minimum, in order to prevent the data structure from becoming a sequential bottleneck. Maintaining consistency in the presence of many simultaneous updates is a complex task. Standard implementations of data structures are based on locks in order to avoid inconsistency of the shared data due to concurrent modifications. In simple terms, a single lock around the whole data structure may create a bottleneck in the program where all of the tasks serialize, resulting in a loss of parallelism because too few data locations are concurrently in use. Deadlocks, priority inversion, and convoying are also side-effects of locking. The risk for deadlocks makes it hard to compose different blocking data structures since it is not always possible to know how closed source libraries do their locking. It is worth noting that in graphics processors (GPUs) locks are not recommended for designing concurrent data structures. GPUs prior to the NVIDIA Fermi architecture do not have writable caches, so for those GPUs, repeated checks to see if a lock is available or not require expensive repeated accesses to the GPU's main memory. While Fermi GPUs do support writable caches, there is no guarantee that the thread scheduler will be fair, which can make it difficult to write deadlock-free locking code. OpenCL explicitly disallows locks for these and other reasons.

Lock-free implementations of data structures support concurrent access. They do not involve mutual exclusion and make sure that all steps of the supported operations can be executed concurrently. Lock-free implementations employ an optimistic conflict control approach, allowing several processes to access the shared data object at the same time. They suffer delays only when there is an actual conflict between operations that causes some operations to retry. This feature allows lock-free algorithms to scale much better when the number of processes increases.

An implementation of a data structure is called *lock-free* if it allows multiple processes/threads to access the data structure concurrently and also guarantees that at least one operation among those finishes in a finite number of its own steps regardless of the state of the other operations. A consistency (safety) requirement for lock-free data structures is *linearizability* [44], which ensures that each operation on the data appears to take effect instantaneously during its actual duration and the effect of all operations are consistent with the object's sequential specification. Lock-free data structures offer several advantages over their blocking counterparts, such as being immune to deadlocks, priority inversion, and convoying, and have been shown to work well in practice in many different settings [73, 69]. They have been included in Intel's Threading Building Blocks Framework [46], the NOBLE library [69, 70], the Java concurrency package [53] and the Microsoft .NET Framework [61]. They have also been of interest to designers of languages such as C++ [21].

The focus here is on concurrent implementations of common abstract data types, such as queues, stacks, other producer-consumer collections, dictionaries and priority queues, which can act as a communication "glue" in parallel applications. Moreover, practical lock-free protocols/algorithms are preferred due to their desirable qualities in terms of performance and fault-tolerance [73, 69]. For each of these abstract data types there exist a considerable number of proposed protocols/algorithms in the literature, see, e.g., the surveys in [43, 15]. As each implementation of an abstract data type has different qualities depending on how it is actually used, e.g. its contents; the level of concurrent accesses to it; the mix of read or update operations issued on it etc., it makes good sense to view them as multi-variant "components". However, a concurrent shared data structure does not match the notion of a component in the EXCESS component model defined in EXCESS D1.2 [48] since an EXCESS component is a computational entity while a concurrent shared data structure is a data storage and communication entity. Hence, multi-variant concurrent shared data structures need to enter the framework as something different from components. There are a number of ways concurrent shared data structures can be used in the EXCESS programming model and runtime system:

- As internal communication medium, "glue", inside component implementations. The variant selection for the concurrent shared data structure can then either be performed inside the component implementation or the component implementation itself treated as a template and expanded into one actual component implementation for each variant of the data structure available. In the latter case variant selection for the data structure would reduce to component variant selection.
- As parameters to components. In the EXCESS component model parameters to components are used to pass data in and out of components. Concurrent shared data structures can be used in this role and would support concurrent updates of the data structure from the inside and/or outside of the component during its execution. This use could be integrated in a similar way to the smart containers discussed in EXCESS D1.2 [48].
- As part of the implementation of the runtime system itself.

To select the most suitable data structure implementation for a given situation is not an easy problem as many aspects of the subsequent use of it impacts the time and energy costs for operations on the data structure. For example, superior operation throughput at high contention (a common selling point of new algorithms for concurrent shared data structures) from which usually (due to system static power) also follows superior energy efficiency in that state does not necessarily translate to superior energy efficiency at lower levels of contention as the empirical case study in Section 4.1.6 below indicates.

In EXCESS D1.1 [49] we determined a number of energy-affecting factors for System A, an Intel server system. Below follows a discussion of each of these factors in the context of a concurrent shared data structure used as communication “glue” in a parallel computation.

General

- *Execution time.* The time spent executing operations on the data structure depends on the number of calls to execute and the duration of each operation/call. The latter is often difficult to predict as it may depend on many aspects of the state of the system and the data structure, such as the algorithm for the operation combined with the current state(/contents) of the data structure (cf. time complexity of sequential data structure operations) and the interference from other concurrent operations on the data structure.
- *Number of sockets used.* Determined by the scheduling of the tasks using the data structure.
- *Number of active cores.* Determined by the scheduling of the tasks using the data structure.

Functional units

- *Instruction types.* Depends among other things on the algorithm for the operation combined with the current state(/contents) of the data structure and the interference from other concurrent operations on the data structure as these may activate different code-paths.
- *Dependency between operations.* Depends among other things on the algorithm for the operation combined with the current state(/contents) of the data structure and the interference from other concurrent operations on the data structure as these may activate different code-paths.
- *Branch prediction.* Failed predictions depend among other things on the algorithm for the operation combined with the current state(/contents) of the data structure and the interference from other concurrent operations on the data structure. The last may be particularly difficult since values are changed outside the current instruction sequence.
- *Clock frequency.* Determined by the system configuration.

Memory

- *Resource contention.* Resource contention in the memory hierarchy can be either accidental, such as cache line eviction due to the limited size of the cache or due to placement policy restrictions forcing otherwise independent cache lines to contend for a particular slot, or deliberate as is often the case in shared data structure implementations where code running on different cores tries to touch the same cache lines at the same time.
- *Number of memory requests.* Depends among other things on the algorithm for the operation combined with the current state(/contents) of the data structure and the interference from other concurrent operations on the data structure as these may activate different code-paths and force retries.
- *Level of memory request completion.* Interference from concurrent operations introduce additional cache line invalidations and, hence, cache misses.
- *Locality of memory references.* Interference from concurrent operations introduce additional cache line invalidations and coherence traffic.

As can be seen above, for concurrent shared data structures several of these factors are affected by the concurrent operations on/dynamic state of the data structure. This means that to estimate the time and energy cost for one operation this state must be known or estimated. As a first approach we will consider concurrent shared data structures in a “steady state”, that is, exposed to an unchanging mix of operations issued by an unchanging set of tasks at an unchanging rate. In this case we can then assign average time and energy costs to operations based on the total throughput and power use. Empirical data for a selection of such “steady states” on a particular system can be collected with micro-benchmarks.

For the initial work on energy efficiency prediction for concurrent shared data structures we have picked some commonly used collection data types as case studies. These are concurrent producer/consumer collections, such as queues, and concurrent search trees. Most of the data structure implementations that we use are part of the NOBLE library [69] which is described in Section 4.1.5 below.

1.3 Micro-benchmarking

In general, micro-benchmarking is conducted to discover targeted properties of a specific system. A *micro-benchmark* in this work context is a small piece of code designed to measure the performance or energy of basic operations of hardware or software systems.

The micro-benchmarks are developed by a loop over a low-level operation (e.g. xor, mul) with N iterations. This description of micro-benchmarks can be represented as EBNF-inspired formal notation $(op)^N$ which means the body op of a loop is executed with N iterations. In order to analyze experimental results, micro-benchmarks must work with a fixed size data-set, perform a constant amount of work per iteration and run for a reasonable

amount of time. In this work package, micro-benchmarks are used to find out the key features and properties of the components that affect energy efficiency of the systems. From the measurement results of micro-benchmarks, the parameters in a proposed analytic energy model are derived.

Micro-benchmarking is a common method used for performance and energy modeling of a given computer system described at an abstract level. It can work directly on the a specific component and does not require detailed simulation. Micro-benchmarking can be applied to different systems since its code is portable. By using the micro-benchmarking method, offline predictors of the energy model are built to predict energy and performance and to support energy optimization.

1.4 Metrics

We rely on two original metrics:

- Throughput: it is a natural metric, used extensively in the performance analysis of data structures, and which measures the number of operations that has been done on the data structure per second.
- Power.

The experiments consist of running a benchmark on the data structure during a given time. Then we count the number of successful operations, which gives the throughput. As we work under a constant execution time, power and energy are equal within a multiplicative factor that is the execution time.

However, another metric is studied in this deliverable, so that we are able to evaluate the energy efficiency of different implementations of the same data structure: the energy per operation. This metric can be useful in the following case: we are given a workload that needs to be executed on a given platform, and there is no time requirement. Then, in terms of energy savings, the implementation that uses the minimum energy per operation is the best implementation.

In our case, the energy per operation is obtained by simply dividing the power by the throughput.

Finally, if we are interested in the bi-criteria problem that mixes energy and performance (*i.e.* where we aim at optimizing both energy and performance), we can plot the energy per operation according to the throughput. By doing this we can also trace the Pareto-optimal frontier for this bi-criteria problem (we eliminate every point such that there exists another point that is better both in terms of energy per operation and throughput).

In this deliverable, we model the two original metrics, namely throughput and power, and derive the other ones from those two.

1.5 Overview and Contributions

As a first step to investigate the energy and performance trade-offs in concurrent data structures, we have developed a new power model for the Movidius Myriad platform that is able

to predict power consumption of our micro-benchmarks with $\pm 4\%$ margin of measured power consumption on the real platform (cf. Sec. 3.2). The new power model confirms the experimental power analysis of concurrent data structures such as concurrent queues: the dynamic power consumption is proportional to the number of SHAVE processors used.

We have performed micro-benchmarks on CPU as well, where we have decomposed power (CPU power, memory power and uncore power) into several parts, namely static, active and dynamic part. In this decomposition, we are able to split the dependencies of power dissipation according to the kind of operation, locality of operands, number of active cores and sockets.

This micro-benchmark study is a preliminary step towards the modeling of performance and power dissipation of several concurrent queue implementations. We define parameters that rule the behavior of the queues, and show how to extrapolate both throughput and power values, by relying on only a few measurements.

Moreover, we have analyzed the feasibility of porting concurrent data structures onto low energy embedded platforms (cf. Sec. 4). We have selected a few synchronization mechanisms from the HPC domain that could easily be replicated on a Movidius Myriad MPSoC and analyzed the performance when used to implement concurrent FIFO queues. This work has been continued to investigate the energy consumption of these synchronization mechanisms and we have then proceeded to determine energy-performance trade-offs.

On Intel platform, we have investigated the energy consumption and performance of concurrent data structures such as concurrent search trees (cf. Sec. 5). Based on our investigation, we have developed new concurrent search trees called Δ Trees that are up to 140% faster and 80% more energy efficient than the state-of-the-art concurrent search trees.

The remaining of the report is organized as follows. Section 2 describes the methodology to measure the consumed energy for two EXCESS platforms, namely CPU-based and Movidius. Section 3 presents energy models that can be used to predict the power consumed on each platform. In section 4, the performance and energy-efficiency of concurrent queue data structures are investigated on CPU-based and Movidius platforms. The investigation results are supported by experimental evaluations. Section 5 analyzes the performance and energy-efficiency of concurrent tree data structures and introduces a novel locality-aware concurrent search tree. The conclusions and future work are provided in section 6.

2 EXCESS Platforms and Energy Measurement Methodology

In this section we summarize the EXCESS systems and the energy measurement methodology used for the experiments in this report. A summary is given here as these systems have previously been described in EXCESS D1.1 [49] and D5.1 [71].

- **System A:** An Intel multicore CPU server (located at Chalmers);
- **System B:** Movidius Myriad1 MV153 development board and simulator (evaluated at Movidius and UiT).

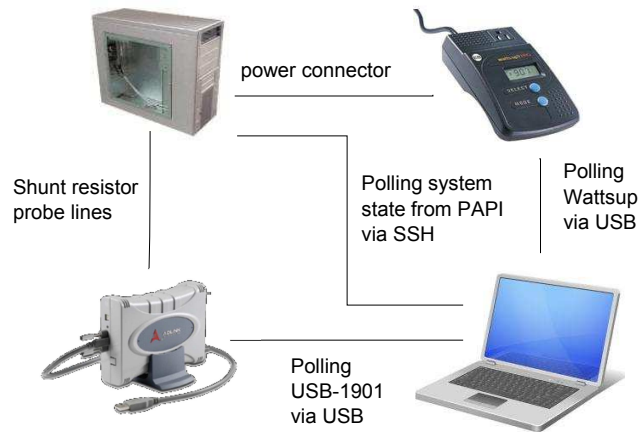


Figure 1: Deployment of energy measurement devices for System A.

2.1 System A: CPU-based platform

2.1.1 System description

- CPU: Intel(R) Xeon(R) CPU E5-2687W v2
 - 2 sockets, 8 cores each
 - Max frequency: 3.4GHz, Min frequency: 1.2GHz, frequency speedstep by DVFS: 0.1-0.2GHz. Turbo mode: 4.0GHz.
 - Hyperthreading (disabled)
 - L3 cache: 25M, internal write-back unified, L2 cache: 256K, internal write-back unified. L1 cache (data): 32K internal write-back
- DRAM: 16GB in 4 4GB DDR3 REG ECC PC3-12800 modules run at 1600MTransfers/s. Each socket has 4 DDR3 channels, each supporting 2 modules. In this case 1 channel per socket is used.
- Motherboard: Intel Workstation W2600CR, BIOS version: 2.000.1201 08/22/2013
- Hard drive: Seagate ST10000DM003-9YN162 1TB SATA

2.1.2 Measurement methodology for energy consumption

The energy measurement equipment for System A at CTH, described in Section 2.1.1, is shown in Figure 1 and outlined below. It has previously been described in detail in EXCESS D1.1 [49] and D5.1 [71].

The system is equipped with external hardware sensors for two levels of energy monitoring as well as built in energy sensors:

- At the system level using an external Watts Up .Net [23] power meter, which is connected between the wall socket and the system.
- At the component level using shunt resistors inserted between the power supply unit and the various components, such as CPU, DRAM and motherboard. The signals from the shunt resistors are captured with an Adlink USB-1901 [1] data acquisition unit (DAQ) using a custom utility.
- Intel’s RAPL energy counters are also available for the CPU and DRAM components. A custom utility based on the PAPI library [14, 79] is used to record these counters and other system state parameters of interest.

For the work presented in this report the component level hardware sensors and the RAPL energy counters have mainly been used.

2.2 System B: Movidius Embedded Platform (Myriad1)

2.2.1 Myriad1 Platform Description

The Myriad1 platform developed by Movidius contains a total of 8 separate SHAVE (Streaming Hybrid Architecture Vector Engine) processors (see Figure 2), each existing on solitary power islands.

The SHAVE processor contains a mix of RISC, DSP, VLIW and GPU features and supports the following data types: (float) f16/32, (unsigned) u8/16/32, and (int) i8/16/32. The SHAVE architecture uses Very Long Instruction Words (VLIWs) as input. The processor is designed to provide a platform that excels in multimedia and video processing. Each SHAVE has its own Texture Management Unit (TMU).

SHAVE also contains wide and deep register files coupled with a Variable-Length Long Instruction-Word (VLLIW) for code-size efficiency. As shown in Figure 2 VLLIW packets control multiple functional units which have SIMD capability for high parallelism and throughput at a functional unit and processor level.

Functional Units of SHAVE

- Integer Arithmetic Unit (IAU) Performs all arithmetic instructions that operate on integer numbers, accesses the IRF.
- Integer Arithmetic Unit (IAU) Performs all arithmetic instructions that operate on integer numbers, accesses the IRF.
- Scalar Arithmetic Unit (SRF) Performs all Scalar integer/floating point arithmetic and interacts with the SRF or IRF depending on what values are used.
- Vector Arithmetic Unit (VAU) Performs all Vector integer/floating point arithmetic and interacts with the VRF.

- Load Store Unit (LSU) There are 2 of these (LSU0 & LSU1) and they perform any memory IO instructions. This means that it interacts with the 128kB CMX memory tile located in the SHAVE.
- Control Move Unit (CMU) This unit interacts with all register files, and allows for comparing and moving between the register files.
- Predicated Execution Unit (PEU) Performs operations based on condition code registers.
- Branch Repeat Unit (BRU) Manages instructions involving any loops, as well as branches.
- Instruction Decoding Unit (IDC) This unit takes a SHAVE variable-length instruction as input and decodes it to determine which functional units are being utilised by the inputted instruction.
- Debug Control Unit (DCU) Used for monitoring the execution of the program, takes note of interrupts and exceptions.

Register Files

- Integer Register File (IRF) Register file for storing integers from either the IAU or the SAU. Can hold up to 32 words which are each 32-bits wide.
- Scalar Register File (SRF) Register file for storing integers from either the SAU. Can hold up to 32 words which are each 32-bits wide.
- Vector Register File (VRF) Register file for storing integers from either the VAU. Can hold up to 32 words which are each 128-bits wide.

The additional blocks in the diagram are the Instruction DeCode (IDC) and Debug Control Unit (DCU). An instruction fetch width of 128-bits and 5-entry instruction pre-fetch buffer guarantee that at least one instruction is ready taking account of branches. Data and instructions reside in a shared Connection MatriX (CMX) memory block which can be configured in 8kB increments to accommodate different instruction/data mixes depending on the workload. The CMX also includes address-translation logic to allow VLLIW code to be easily relocated to any core in the system.

In the 65nm System-on-Chip (SoC), eight SHAVE processors are combined with a software-controlled memory subsystem and caches which can be configured to allow a large range of workloads to be handled, providing exceptionally high sustainable on-chip bandwidth to support data and instruction supply to the 8 processors. Data is moved between peripherals, processors and memory via a bank of software-controlled DMA engines. The device supports 8, 16, 32 and some 64-bit integer operations as well as fp16 and fp32 arithmetic and is capable of aggregate 1 TOPS/W maximum 8-bit equivalent operations in a low-cost plastic BGA package with integrated 128Mbit or 512Mbit Mobile DDR2 SDRAM.

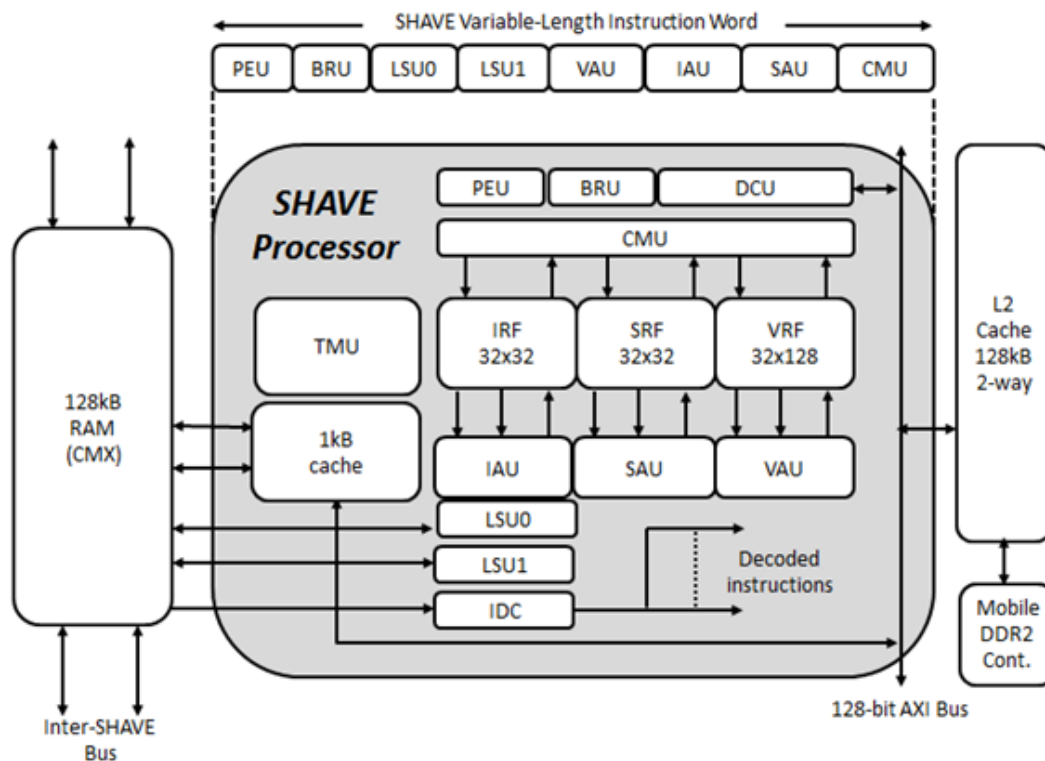


Figure 2: SHAVE Instruction Units

As power efficiency is paramount in mobile applications, in addition to extensive clock and functional unit gating and support for dynamic clock and voltage scaling for dynamic power reduction, the device contains a total of 11 power-islands: one for each SHAVE, one for the CMX RAM, one for the RISC and peripherals and one always-on domain. This allows very fine-grained power control in software with minimal latency to return to normal operating mode, including maintenance of SRAM-state eliminating the need to reboot from external storage.

2.2.2 Measurement methodology

An extensive effort has been made to measure Myriad1 performance in a semi-automated way in order to produce better power estimates.

The Power Measurements tests were introduced in order to have an insight into the power consumed by Myriad1 in several basic cases in order to be able both to decompose the Myriad1 power consumption into power components and characterize power consumption in such basic operations of Myriad1. The tests set were devised to characterize the power consumed when running SHAVE code without DDR data accesses.

The modifications were made to the MV153 to bypass the on-board voltage regulator

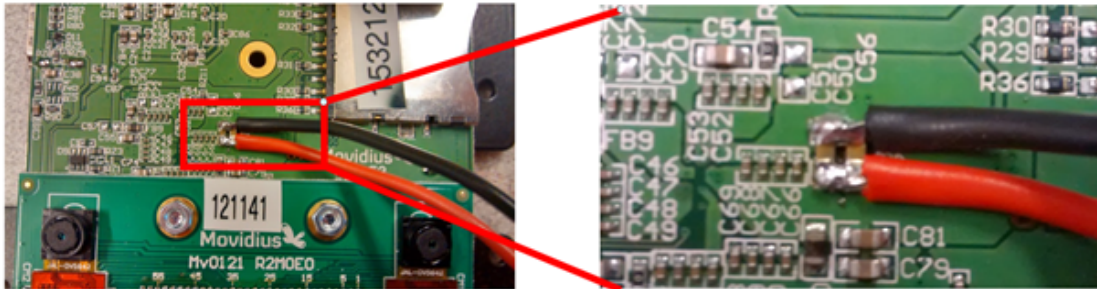


Figure 3: Power Supply Modification

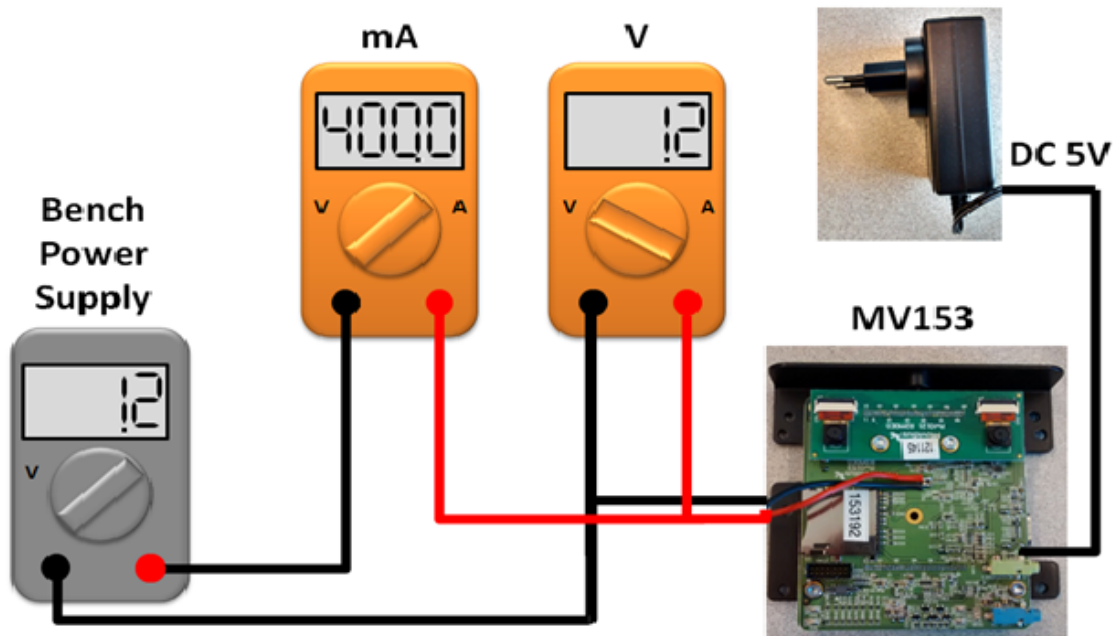


Figure 4: Bench setup for MV153 Power Measurement Schematic

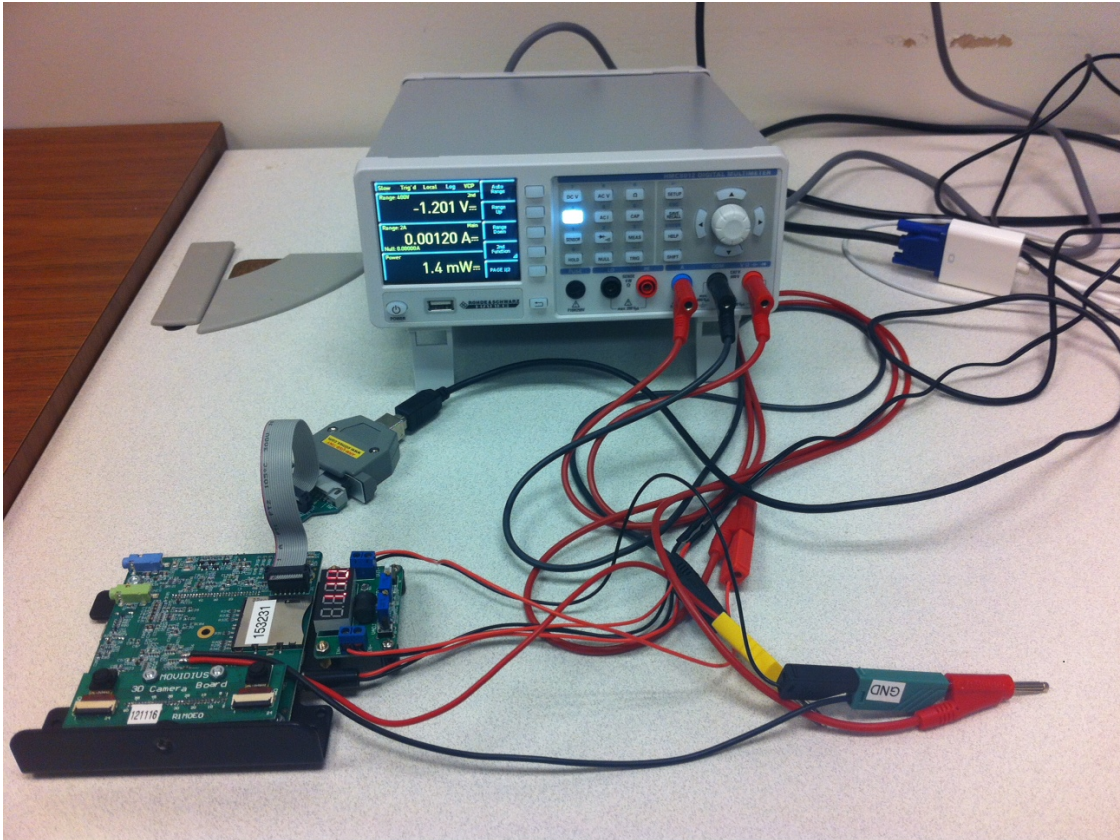


Figure 5: Bench setup for MV153 Power Measurement

which down-regulates the 5V wall PSU to the 1.2V core voltage required by Myriad1 allowing an external bench power-supply to be used in its place as shown in Figure 3.

The schematic for the connection of the Power Supply Unit (PSU), multimeters and MV153 for power measurement are shown in Figure 4. Note the standard DC wall supply is required in addition to the bench PSU in order to supply the other elements of the system.

The bench setup consists of a modified MV153 board, a DC step down converter down-regulating the 5V wall PSU to the 1.2V core voltage and one HAMEG multimeter measuring all the voltage, current and consumed power values as shown in Figure 5.

2.3 Significant Synchronization Hardware Differences Between the Two Systems

To synchronize processes efficiently, multi-/many-core systems usually support certain synchronization primitives. This section discusses the fundamental synchronization primitives, which typically read the value of a *single* memory word, modify the value and write the new value back to the word *atomically*. Different architectures support different synchronization

primitives in hardware.

2.3.1 Fundamental synchronization primitives

The definitions of the primitives are described in Figure 6, where x is a memory word, v, old, new are values and op can be operators *add*, *sub*, *or*, *and* and *xor*. Operations between angle brackets $\langle \rangle$ are executed atomically.

<p>TAS(x) /* <i>test-and-set, init: $x \leftarrow 0$</i> */ $\langle oldx \leftarrow x; x \leftarrow 1; \mathbf{return} \ oldx; \rangle$</p>	<p>LL(x) /* <i>load-linked</i> */ $\langle \mathbf{return} \ \text{the value of } x \ \text{so that}$ <i>it may be subsequently used</i> with SC \rangle</p>
<p>FAO(x, v) /* <i>fetch-and-op</i> */ $\langle oldx \leftarrow x; x \leftarrow op(x, v); \mathbf{return} \ oldx; \rangle$</p>	<p>SC(x, v) /* <i>store-conditional</i> */ $\langle \mathbf{if} \ (\text{no process has written to } x$ <i>since the last LL(x))</i> $\{x \leftarrow v;$ $\mathbf{return}(true)\};$ $\mathbf{else} \ \mathbf{return}(false); \rangle$</p>
<p>CAS(x, old, new) /* <i>compare-and-swap</i> */ $\langle \mathbf{if}(x = old) \{x \leftarrow new; \mathbf{return}(true)\};$ $\mathbf{else} \ \mathbf{return}(false); \rangle$</p>	

Figure 6: Synchronization primitives

Synchronization power The primitives are classified according to their synchronization power or *consensus number* [57], which is, roughly speaking, the maximum number of processes for which the primitives can be used to solve a *consensus problem* in a fault tolerant manner. In the consensus problem, a set of n asynchronous processes, each with a given input, communicate to achieve an agreement on one of the inputs. A primitive with a consensus number n can achieve consensus among n processes even if up to $n - 1$ processes stop [74].

According to the consensus classification, read/write registers have consensus number 1, i.e. they cannot tolerate any faulty processes in the consensus setting. There are some primitives with consensus number 2 (e.g. *test-and-set* (*TAS*) and *fetch-and-op* (*FAO*)) and some with infinite consensus number (e.g. *compare-and-swap* (*CAS*) and *load-linked/store-conditional* (*LL/SC*)). It has been proven that a primitive with consensus number n cannot implement a primitive with a higher consensus number in a system of more than n processes [57]. For example, the *test-and-set* primitive, whose consensus number is two, cannot implement the *compare-and-swap* primitive, whose consensus number is unbounded, in a system of more than two processes. Most modern general purpose multiprocessor architectures support *compare-and-swap* (*CAS*) in hardware. *compare-and-swap* (*CAS*) is also the most popular synchronization primitive for implementing both lock-based and nonblocking concurrent data structures. For many non-blocking data structures a primitive with a consensus number n is needed.

The Myriad platform, as many other embedded platforms, avails *test-and-set (TAS)* registers which have consensus number 2 and not *compare-and-swap (CAS)*. These Test-and-Set registers can be used to create spin locks, which are commonly referred as "mutexes". Spin-locks are used to create busy-waiting synchronization techniques: a thread spins to acquire the lock so as to have access to a shared resource.

The Myriad platform also avails a set of registers that can be used for fast SHAVE arbitration. Each SHAVE has its own copy of these registers and its size is 4x64 bit words. An important characteristic is that they are accessed in a FIFO pattern, so each one of them is called a "SHAVE's FIFO". Each SHAVE can push data to the FIFO of any other SHAVE, but can read data only from its own FIFO. A SHAVE writes to the tail of another FIFO and the owner of the FIFO can read from any entry. If a SHAVE attempts to write to a full FIFO, it stalls. Finally, the LEON processor cannot access the FIFOs. SHAVE FIFOs can be utilized to achieve efficient synchronization between the SHAVEs. Also, they provide an easy and fast way for exchanging data directly between the SHAVEs (up to 64 bits per message), without the need to use shared memory buffers.

Analysis of experiments on the Myriad platform shows that the mutex implementation is a fair lock with round-robin arbitration. But most scalable designs for concurrent data structures require , a hardware primitive that has unbounded consensus number. Because of the lack of support of strong synchronization primitives, from the embedded hardware side, we had to come with new algorithmic designs for the data structures under consideration fitting the capabilities of the embedded systems area.

3 Energy Models for EXCESS Platforms

3.1 Energy Models for CPU-based Platforms

3.1.1 General Power Model

The power model that is presented in EXCESS D1.1 [49] decomposes the total power into static, socket activation and dynamic power, as recalled in Equation 1. In this equation, f is the clock frequency, soc the number of activated sockets on the chip, op is the considered operation and thr is the number of active cores; the active power is proportional to the number of active sockets, while the dynamic power is proportional to the number of active cores.

For modeling power consumption of data structures, we need to estimate the dynamic component which depends on the frequency, number of active cores, locality and amount of memory requests together with the instruction type as is also mentioned in D1.1 [49]. Therefore, we improve the model with an additional parameter loc to represent the locality of operands for instructions that can transfer data between memory and registers, such a move from L1, L2, last level cache, main memory or remote memory. This parameter was not included in D1.1 because we consider only the total power in which loc parameter plays a negligible role. For fine-grained analysis, we include this parameter in our model.

$$\begin{cases} P(f, op, soc, loc, thr) = P_{stat} + P_{active}(f, soc) + P_{dyn}(f, op, loc, thr) \\ P_{active}(f, soc) = soc \times P_{active}(f) \end{cases} \quad (1)$$

	Static	Active	Dynamic
CPU	$P_{stat}^{(C)}$	$P_{active}^{(C)}$	$P_{dyn}^{(C)}$
Memory	$P_{stat}^{(M)}$	$P_{active}^{(M)}$	$P_{dyn}^{(M)}$
Uncore	$P_{stat}^{(U)}$	$P_{active}^{(U)}$	$P_{dyn}^{(U)}$

Table 1: Power views

As another improvement to the previous power model, we decompose the power into two orthogonal bases, each base having three dimensions. On the one hand, we define the model basis by separating the power into static, active and dynamic power, such that the total power is computed by:

$$P = P_{stat} + P_{active} + P_{dyn}.$$

On the other hand, the measurement basis corresponds to the components that actually dissipates the power, *i.e.* CPU, memory and uncore. The power dissipation measurement is done through Intel’s RAPL energy counters read via the PAPI library [14, 79]. These counters reflect this discrimination by outputting the power consumption along three dimensions:

- power consumed by CPU, which includes the consumption of the computational cores, and the consumption of the first two level of caches;
- power consumed by the main memory;
- remaining power, called “uncore”, which includes the ring interconnect, shared cache, integrated memory controller, home agent, power control unit, integrated I/O module, config agent, caching agent and Intel QPI link interface.

Also, total power is obtained by the sum:

$$P = P^{(C)} + P^{(M)} + P^{(U)}.$$

This latter additional orthogonal dimension will provide a better perspective for modeling power consumption of data structures, especially for the dynamic component. Table 1 sums up both dimensions.

In this section, we study each dimension, in each base, so that we are able to express the power dissipation from any perspective:

$$P = \sum_{X \in \{C, M, U\}} \left(P_{stat}^{(X)} + P_{active}^{(X)} + P_{dyn}^{(X)} \right).$$

3.1.2 Power Components Derivation

By definition, only the dynamic component of power is dependent on the type of instruction or more generally the executing program. In order to obtain dynamic component P_{dyn} , we first have to determine static P_{stat} and socket activation P_{active} costs. This was done in D1.1 but the derivation process was not described in detail. Therefore, we will explain it briefly in this subsection.

In D1.1, a large variety of instructions were examined with respect to their power and energy consumption. We have observed a linear relation between the number of threads and power for instructions that do not lead to data transfer between the memory hierarchy and registers. For instance, addition operates on two registers and do not lead to data transfer. Data transfer is done via move instructions before or after addition instruction if required. So, the locality parameter loc is only valid for instructions that is dependent on the locality of data, like variants of the move instruction. These operations are also prone to variability due to cache and memory states which can also change with the interaction between threads. Briefly, $P_{dyn}^{(M)}$ and $P_{dyn}^{(U)}$ is significant only for the instructions that lead to data transfer in the memory hierarchy. Also the loc parameter is only meaningful for $P_{dyn}^{(M)}$ and $P_{dyn}^{(U)}$. For derivation of P_{stat} and P_{active} , we just use the instructions that operate on the registers because the $P_{active}^{(M)}$ and $P_{active}^{(U)}$ parts can be neglected for these instructions. We refer to these instructions as op_{reg} and utilize them to obtain static and socket activation costs for each component (CPU, memory, uncore) of the orthogonal decomposition. A bunch of instructions belonging to op_{reg} is executed repeatedly for some time interval with varying number of threads for each frequency. We formulate the derivation process as, for all $X \in \{C, M, U\}$:

$$\begin{aligned}
 P_{dyn}(f, op, loc, thr) &= P_{dyn}^{(M,U)}(f, op, loc, thr) + P_{dyn}^{(C)}(f, op, thr) \\
 P_{dyn}^{(M,U)}(f, op_{reg}, loc, thr) &= 0 \\
 P_{dyn}^{(C)}(f, op_{reg}, thr) &= thr \times P_{dyn}^{(C)}(f, op_{reg}) \\
 P_{dyn}^{(X)}(f, op_{reg}) &= \frac{1}{2} (P^{(X)}(f, op_{reg}, soc, loc, 16) - P^{(X)}(f, op_{reg}, soc, loc, 14)) \\
 P_{active}^{(X)}(f) &= P^{(X)}(f, op_{reg}, 2, loc, 10) - P^{(X)}(f, op_{reg}, 1, loc, 8) - P_{dyn}^{(X)}(f, op_{reg}) \times 2 \\
 P_{stat}^{(X)}() &= P^{(X)}(f, op_{reg}, soc, loc, thr) - soc \times P_{active}^{(X)}(f) - thr \times P_{dyn}^{(X)}(f, op_{reg})
 \end{aligned}$$

Using above equations, we verified that $P_{stat}^{(X \in \{C, M, U\})}$ is approximately constant according to instruction type, pinning, number of threads and frequency thus we take the mean of the values of P_{stat} over the whole space to find P_{stat} . We apply the same approach to find $P_{active}^{(X \in \{C, M, U\})}$ which only depends on frequency, and not on the operation. Having obtained P_{stat} and $P_{active}^{(X)}$, we extract $P_{dyn}^{(X \in \{C, M, U\})}$ for “all” types of instructions, thread, pinning and frequency setting, by removing the static and active part from the total power.

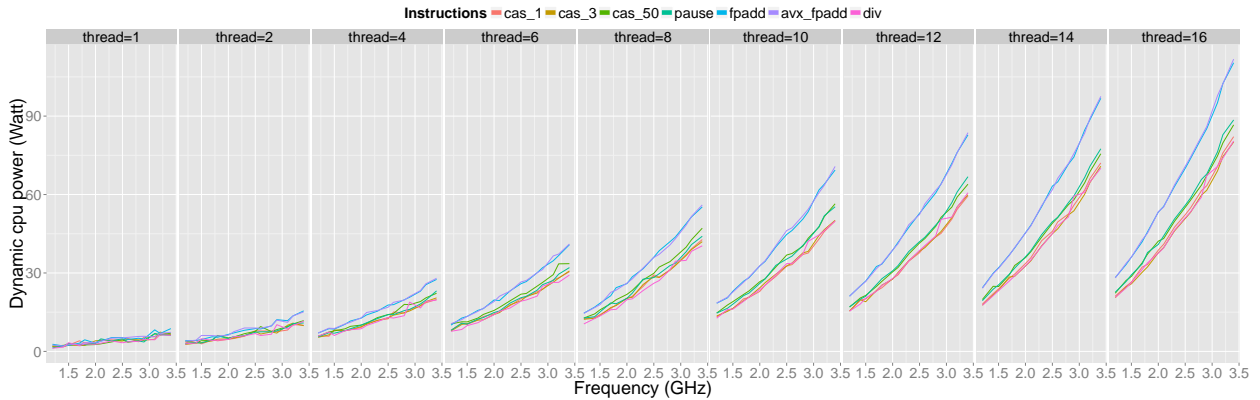


Figure 7: Dynamic CPU power for micro-benchmarks

3.1.3 Dynamic CPU Power

Having determined and excluded static and socket components, we obtain the dynamic power component for each instruction, thread count, pinning and frequency setting. In the micro-benchmarks of D1.1, a large variety of instructions are surveyed. Among them, we pick a small set of instructions that can be representative for data structure implementations, namely *Compare-and-Swap*, pause, floating point division, addition together with vector addition. *Compare-and-Swap* can be representative for the retry loops and divisions/additions can be used to represent the parallel work which determines the contention on the data structures. The decomposition of dynamic power in terms of CPU, memory and uncore components for these instructions are illustrated in Figures 7, 8 and 9.

Based on the observation that $P_{dyn}^{(C)}$ shows almost linear behavior with respect to number of threads, we model the convex $P_{dyn}^{(C)}$ as:

$$P_{dyn}^{(C)}(f, op) = (A \times f^\alpha + B)$$

Each instruction might provide different power behavior as illustrated in Figure 7, therefore we find A , B , α for each instruction separately. B could be different for each instruction because of the activation of different functional units, this is also why we included this constant in $P_{dyn}^{(C)}$.

To obtain A , B and α , we proceed in the following way. We are given an operation op , and we consider the executions of this operation with 16 threads on 2 sockets. Let $v^{(freq)}$ be the vector of frequencies where we want to estimate the dynamic power (we dispose F different frequencies, expressed in 10^{-1} GHz, such that $v_1^{(freq)} = 12$ and $v_F^{(freq)} = 34$). We note $v^{(meas)}$ the vector of dynamic powers that have been computed from the measurements through the process described above, and $v^{(est)}(A, B, \alpha)$ the vector of estimated dynamic

powers. More especially, for all $i \in \{1, \dots, F\}$:

$$v_i^{(meas)} = P_{dyn}^{(C)}(v_i^{(freq)}, op)$$

$$v_i^{(est)}(A, B, \alpha) = \left(A \times \left(v_i^{(freq)} \right)^\alpha + B \right)$$

The Euclidean norm of a vector v is denoted $\|v\|$.

We solve the following minimization problem, with the help of the Matlab “fminsearch” function:

$$\min_{A, B, \alpha} \|v^{(meas)} - v^{(est)}(A, B, \alpha)\|$$

Table 2 provides the values for power constants and exponent for selected instructions.

3.1.4 Dynamic Memory and Uncore Power

In the micro-benchmarks, we observe that many instructions do not lead to an increase in dynamic memory and uncore power because the operands of the instructions, except *Compare-and-Swap*, presumably reside in the core. On the other hand, *Compare-and-Swap* lead to an increase in memory and uncore power only when the threads are pinned to different

	A	α	B
cas	0.001392	1.6415	0.0510
fpdiv	0.001038	1.7226	0.0585
add	0.001004	1.8148	0.0912
avx-add	0.001130	1.7828	0.0894
pause	0.000854	1.7920	0.0736

Table 2: Instruction power coefficients

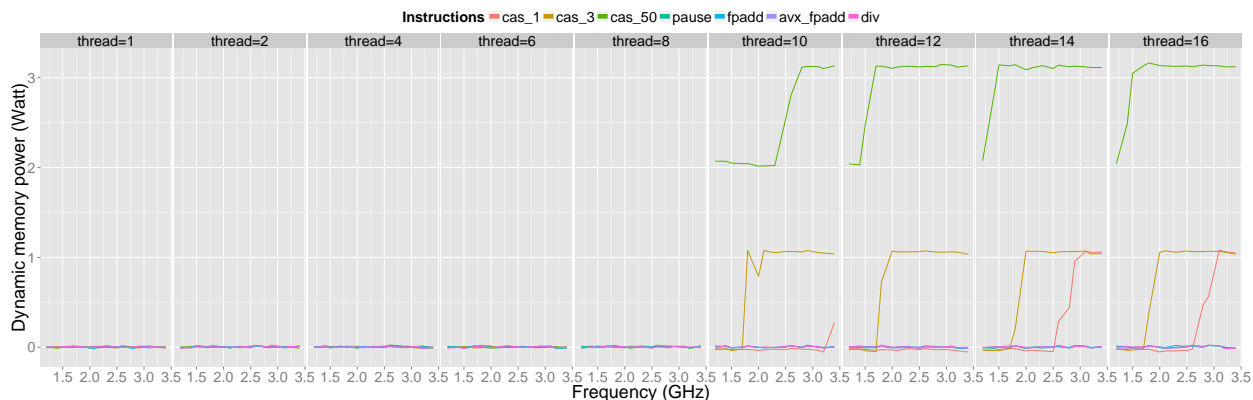


Figure 8: Dynamic memory power for micro-benchmarks

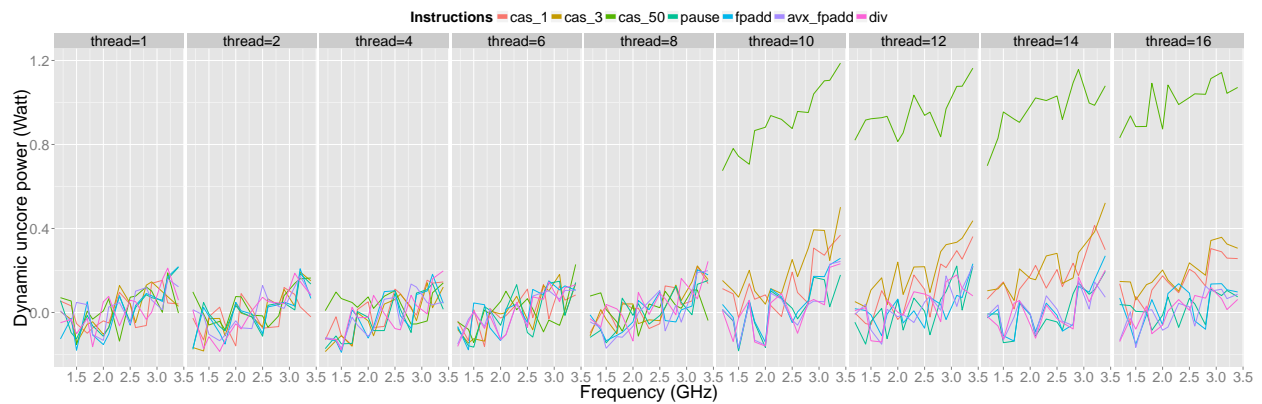


Figure 9: Dynamic uncore power for micro-benchmarks

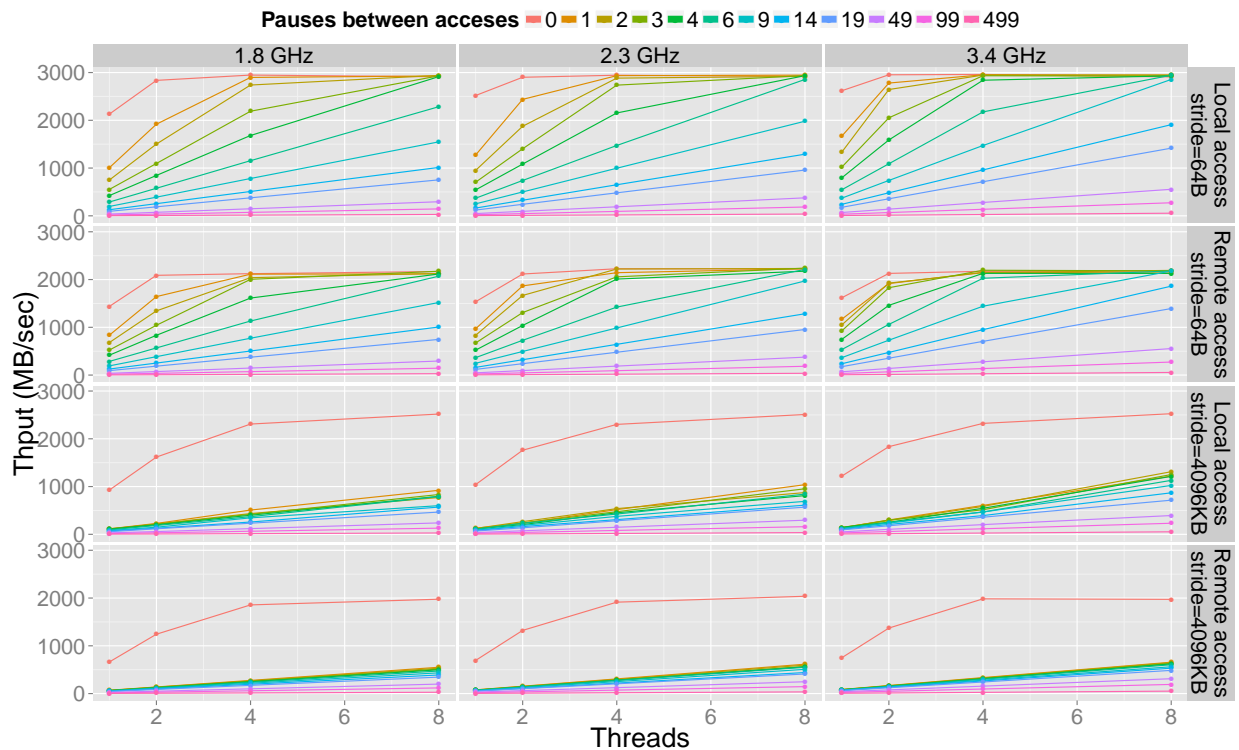


Figure 10: Throughput for array traversal benchmark

sockets. The resulting *ping-pong* of the updated cache line between sockets is responsible for this effect. As mentioned in D1.1, *Compare-and-Swap* micro-benchmarks are indeed prone to unfairness among threads. When *Compare-and-Swap* is executed repeatedly by all threads on the same cache line without any work in between *Compare-and-Swap* attempts, the thread which gets the ownership of cache line succeeds repeatedly while others starve. This fact decreases the transfer rate of the cache line between local caches. Due to this, we introduce 3

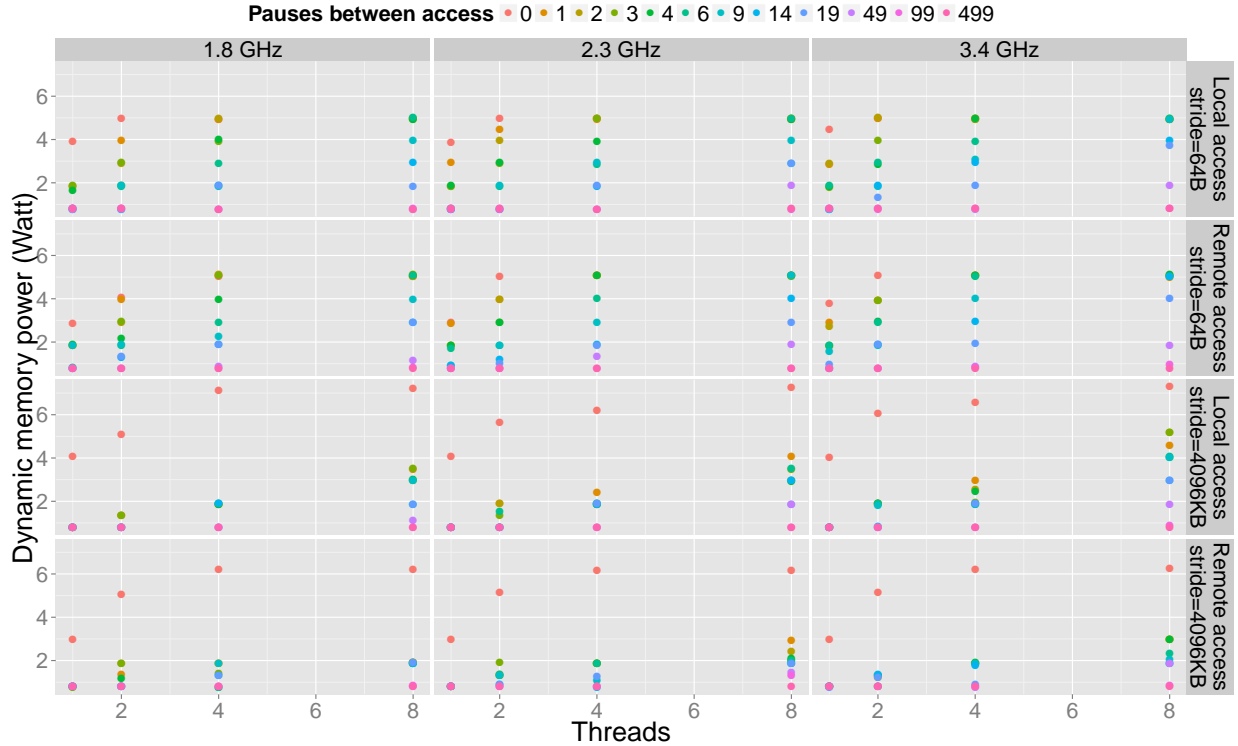


Figure 11: Step-like power for array traversal benchmark

different *Compare-and-Swap* micro-benchmarks looping on 1, 3 and 50 shared variables that are aligned to different cache lines. By doing so, we aim at increasing the traffic between cores and sockets together with the amount of memory accesses. Figures 8 and 9 provide the $P_{dyn}^{(U)}$ and $P_{dyn}^{(M)}$ values. It can be observed that all parameters including number of threads, frequency, pinning play a role for *Compare-and-Swap*. As a remark for the provided figures, threads are pinned using a dense mapping strategy that leads to inter-socket communication only after 8 threads.

$P_{dyn}^{(U)}$ and $P_{dyn}^{(M)}$ do not increase when threads are pinned to same socket. In this case, the intra-socket communication between threads takes place via the ring interconnect without introducing a memory access. Thus, absence of increase in memory power is reasonable. However, one might expect an increase in uncore power for these cases because RAPL uncore power presumably includes LLC and ring interconnect power. We do not observe this probably because the main components that attached to the ring are not used. The increase of uncore and memory power can be observed when threads are pinned to different sockets, due to remote memory accesses which uses important uncore components such as the QPI link interface and home agent. An interesting observation regarding memory power is that it shows a step function behavior. We think that this is because of the RAPL power capping algorithm which determines a power budget based on memory bandwidth, as presented in the work of David et al. [20]. The RAPL algorithm specifies a power cap for a time win-

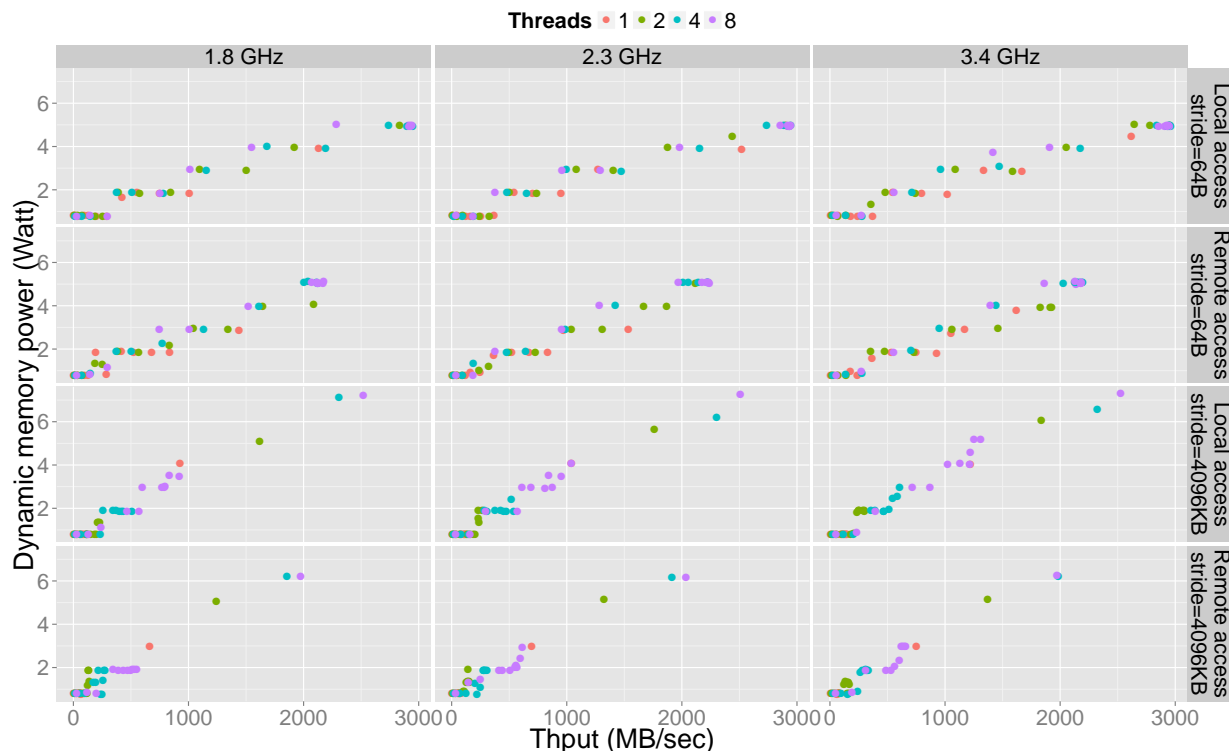


Figure 12: Memory power for array traversal benchmark

down depending on the memory bandwidth requirements of previous time intervals and sets the memory in a power state that is expected to maximize energy efficiency. Based on the amount of memory accesses, it jumps between states finding a trade-off between bandwidth and power. The finite number of states leads to the step-like power curves in Figure 8. Thus, the memory power seems to be determined by the amount of memory accesses per unit of time which is dependent on frequency, number of threads and the amount of shared variables for our *Compare-and-Swap* experiments.

To justify this observation, we use a benchmark which stresses the main memory. We allocate a huge contiguous array and align each element of the array to a separate cache line. In addition, we force the array to be allocated in the memory module residing in the first socket. Thus, we can regulate remote and local memory accesses by pinning strategies. We pin all threads either to first or second socket. We also change the number of threads, frequency and interleave varying amount of pause operations between array accesses to change the bandwidth requirements of the benchmark. Threads access independent portions of the array with a stride. The hardware prefetcher increases the performance remarkably when adjacent cache lines are accessed while traversing the array and a stride of page size can be used to disable the hardware prefetcher. We run the same experiment both with a stride of 64 Bytes, which is the size of a cache line, and 4096 kBytes which is the page size, to reveal effect of prefetching. As provided in Figure 10, the system reaches its peak bandwidth

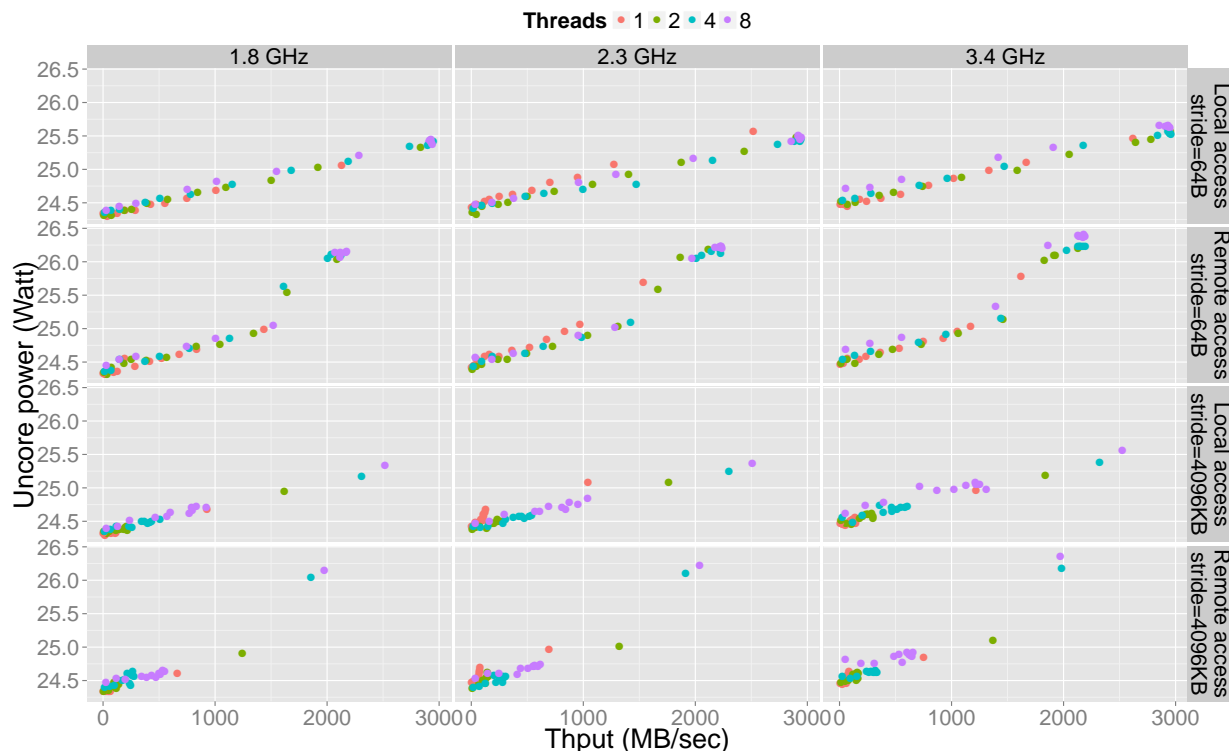


Figure 13: Uncore power for array traversal benchmark

more rapidly when the prefetcher is activated and attains better bandwidth. Moreover, the bandwidth difference between completely remote and local accesses is noticeable. Another point is that frequency does not influence the maximum achievable bandwidth. This fact means that there is opportunity for energy savings with DVFS for memory-bounded regions of applications. We also observe the step-like power behavior, due to the RAPL algorithm, with this benchmark in Figure 11.

In Figure 12, dynamic memory power consumption is shown for the array traversal benchmarks. From the analysis of the results, it can be deduced that the memory power is strongly correlated with the number of bytes accessed per second. There is no clear impact of the number of threads and frequency to the memory power except their indirect effect on bandwidth. In contrast, access stride has a direct, though limited, impact on the memory power together with its indirect impact as it increases the bandwidth. By accessing data with a stride of a cache line, we possibly make use of the open page mode of DRAM which could be influential in terms of energy efficiency due to avoidance of bit-line precharge and row access cost. But, we still observe a linear relation between throughput and memory power for both strides. In addition, remote or local accesses do not provide a noticeable difference for memory power. On the other hand, it is observed that uncore power depends on the frequency, presumably due to the traffic on the ring interconnect and components attached to it such as the Home Agent and Integrated Memory Controller. Furthermore, remote memory accesses

increase the uncore power consumption because they use the QPI link interface which adds an additional cost compared to local memory accesses as shown in Figure 13. All these observations regarding memory and uncore power will shed light to the analysis and modeling of data structures in Section 4. One major source of differences in power consumption between different implementations is the memory and uncore consumption, which is related to locality and bandwidth requirements of the implementations.

3.1.5 Summary of Micro-Benchmarking for Power Modeling on CPU

Figure 14 recalls the main achievements of the micro-benchmark study on CPU, where d is the amount of memory accessed per unit of time in the main memory or through QPI link.

	Static	Active	Dynamic
CPU	$P_{stat}^{(C)}(f, op, soc, loc, n)$	$P_{active}^{(C)}(f, op, soc, loc, n)$	$P_{dyn}^{(C)}(f, op, soc, loc, n)$
Memory	$P_{stat}^{(M)}(f, op, soc, loc, n)$	$P_{active}^{(M)}(f, op, soc, loc, n)$	$P_{dyn}^{(M)}(f, op, soc, loc, n)$
Uncore	$P_{stat}^{(U)}(f, op, soc, loc, n)$	$P_{active}^{(U)}(f, op, soc, loc, n)$	$P_{dyn}^{(U)}(f, op, soc, loc, n)$

Dependency removal

	Static	Active	Dynamic
CPU	$P_{stat}^{(C)}()$	$soc \times P_{active}^{(C)}(f)$	$n \times (A(op) \times f^{\alpha(op)} + B(op))$
Memory	$P_{stat}^{(M)}()$		$d \times P_{dyn}^{(M)}(op, loc)$
Uncore	$P_{stat}^{(U)}()$		$d \times P_{dyn}^{(U)}(op, loc)$

Figure 14: Dependency shrinking

3.2 Energy Models for Movidius Embedded Platforms

3.2.1 Description of Microbenchmarks

Regarding the assembly files used in the test execution process, a fixed number of instructions in the loop was established for all the tests, meaning that each assembly file contains six instructions in the loop that is infinitely repeated. This convention was made in order to keep a continuity and a consistency of tests, by giving an insight in measuring the consumption on different SHAVE units, enabling to make comparisons between SHAVE units.

The assembly files used in the testing process contain code that test the instruction power decode and the instruction fetch. The majority of tests use pseudo-realistic data, by pseudo-realistic data we understand having as many non-zero values as possible and avoiding data value repetition at different offsets.

Below are the used test cases as micro-benchmarks for Movidius platform.

- SauMul, SauXor: SAU operations - with instruction fetch and different data values (XOR-MUL)
- IauMul, IauXor: IAU operations - with instruction fetch and different data values (XOR-MUL)
- VauMul, VauXor: VAU operations - with instruction fetch and different data values (XOR-MUL)
- CmuCpss, CmuCpivr: CMU operations — with instruction fetch and different data values (CPSS - CPIVR)
- SauXorCmuCpss: SAU & CMU — with instruction fetch and different data values (SAU.XOR || CMU.CPSS)
- SauXorCmuCpivr: SAU & CMU — with instruction fetch and different data values (SAU.XOR || CMU.CPIVR)
- SauXorIauXor, IauXorCmuCpss: SAU & IAU — with instruction fetch and different data values (SAU.XOR || IAU.XOR & IAU.XOR || CMU.CPSS)
- SauXorVauXor, SauXorVauMul: SAU & VAU — with instruction fetch and different data values (SAU.XOR || VAU.XOR & SAU.XOR || VAU.MUL)
- SauXorCmuIauXor: SAU & IAU & CMU — with instruction fetch and different data values (SAU.XOR || CMU.CPIS || IAU.XOR)

3.2.2 Movidius Power Model and Its Sanity Check

The experiments are conducted for benchmarks with single unit and multiple units. Each benchmark is tested by running Myriad1 with 1, 2, 4, 6 and 8 SHAVE cores.

From the experiment results, we observe that the power consumption of Movidius Myriad1 platform is ruled by the following model:

$$P = P^{stat} + \#\{\text{active SHAVE}\} \times (P^{act} + P_{SHAVE}^{dyn}) \quad (2)$$

The operands in the formula are explained as below. The static power P^{stat} is the needed power when the Myriad1 processor is on. The P^{act} is the power consumed when a SHAVE core is on. Therefore, this active power is multiplied with the number of used SHAVE core when the benchmark is run with several cores.

The dynamic power P_{SHAVE}^{dyn} of each SHAVE is the power consumed by all working operation units working on SHAVE. As described in the previous section, each SHAVE core has several components, including LSU, PEU, BRU, IAU, SAU, VAU, CMU, etc. Different arithmetic operation units have different P^{dyn} values. In this power model, we focus on the

experiments with the benchmarks performing different arithmetic operations such as IAU, VAU, SAU and CMU.

The dynamic power P^{dyn} of each SHAVE is the power consumed by all working operation units. As described in the previous section, each SHAVE core has several components, including LSU, PEU, BRU, IAU, SAU, VAU, CMU, etc. Different arithmetic operation units have different P^{dyn} values. In this power model, we focus on the experiments with the benchmarks performing different arithmetic operations such as IAU, VAU, SAU and CMU.

When adding one more SHAVE, we can identify the sum of SHAVE P^{act} and P^{dyn} which is the power level difference of the two runs (with one SHAVE core and with two SHAVE corer). Given the sum of P^{act} and P^{dyn} , P^{stat} is derived from the formula. Then, the average value of P^{stat} from all micro-benchmarks experimental results is computed:

$$P^{stat} = 62.63 \text{ mW} \quad (3)$$

Then for each operation unit, we obtain the two parameters P_{op}^{dyn} and P^{act} by using the actual power consumption of the benchmark for individual units and multiple units.

$$\begin{cases} P_{(IauXor)}^{dyn} = P^{stat} + \#\{\text{active SHAVE}\} \times (P^{act} + P_{(IauXor)}^{dyn}) \\ P_{(SauXorIauXor)}^{dyn} = P^{stat} + \#\{\text{active SHAVE}\} \times (P^{act} + P_{(SauXorIauXor)}^{dyn}) \\ P_{(SauXor)}^{dyn} = P^{stat} + \#\{\text{active SHAVE}\} \times (P^{act} + P_{(SauXor)}^{dyn}) \end{cases} \quad (4)$$

Then, the average value of P^{act} among all operation units is calculated.

$$P^{act} = 51.4 \text{ mW} \quad (5)$$

At this point, we also have the P_{op}^{dyn} of every arithmetic unit op . Applying the model to calculate the consumed energy, the results showed the deviation from the measured data, especially the benchmarks running a single arithmetic unit. We attribute this difference to the inter-operation cost when more than one unit of the SHAVE core work in a combination. This inter-operation cost is also mentioned in the model suggested by Movidius in EXCESS deliverable D4.1. The P^{dyn} of SHAVE running multiple units is then computed by the formula below:

$$P_{SHAVE}^{dyn} = \sum_i P_i^{dyn}(op) + \max_i \{O_i(op)\} \quad (6)$$

The combined P^{dyn} is the sum of P^{dyn} of each unit, plus the highest inter-operational cost O_i among the units in the combination. By using the highest inter-operation cost, P^{dyn} after computed is more accurate than using the sum of inter-operation cost from all units in the combination. E.g. $P_{(SauXorIauXor)}^{dyn} = P_{(SauXor)}^{dyn} + P_{(IauXor)}^{dyn} + \max(O_{SauXor}, O_{IauXor})$.

Given the P^{stat} and P^{act} , the P_{op}^{dyn} of an operation unit is computed based on the actual power consumption of the benchmark using a single unit (e.g. SauXor, IauXor, etc.). Then, its inter-operational cost is computed based on the actual power consumption when this unit works in a combination with other units.

The Table 3 lists the inter-operational cost of each unit when it works in a combination with other units.

Operation Unit	P_{op}^{dyn} (mW)	O_{op} (mW)
SauXor	3.05	1.15
SauMul	6.97	1.83
VauXor	17.57	13.12
VauMul	32.78	11.62
IauXor	4.53	1.07
IauMul	3.98	4.42
CmuCpss	1.00	4.60
CmuCpivr	6.41	5.69

Table 3: Inter-operational cost (in mW) of SHAVE operation units

From Equations 3, 5 and 6, the complete model for Movidius Myriad1 is derived as follows:

$$P = P^{stat} + \#\{\text{active SHAVE}\} \times (P^{act}) + \#\{\text{active SHAVE}\} \times \left(\sum_i P_i^{dyn}(op) + \max_i \{O_i(op)\} \right) \quad (7)$$

Applying this formula to different combinations of operation units in the SHAVE core, we plot the relative error of this model in the Figure 15. The relative error is the difference between the actual power consumption measure by device and the predicted power consumption computed through Equation 7, then divided by the actual power consumption. Under this model, the relative error varies within $\pm 4\%$. This model is not only applicable for a single unit but also the combination of two or three units.

4 Modeling Energy Consumption of Concurrent Queue Implementations

4.1 Concurrent Queues on CPU-based Platform

Concurrent FIFO queues and other producer/consumer collections are fundamental data structures that are key components in applications, algorithms, run-time and operating systems. The Queue abstract data type is a collection of items in which only the earliest added item may be accessed. Basic operations are *Enqueue* (add to the tail) and *Dequeue* (remove from the head). Dequeue returns the item removed. The data structure is also known as a “first-in, first-out” or FIFO buffer.

4.1.1 Objective and Process

As explained in Section 1.2, we aim at predicting the energy efficiency of different queue implementations through several metrics. We have seen in D1.1 [49] that the energy efficiency

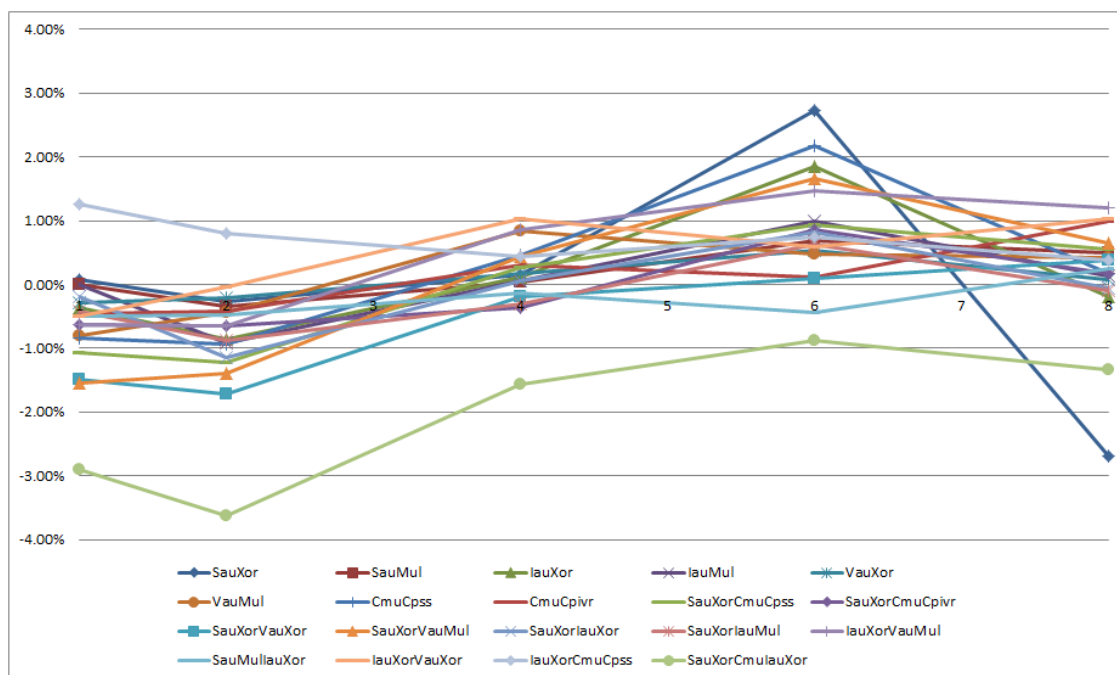


Figure 15: Relative error of Myriad1 power prediction

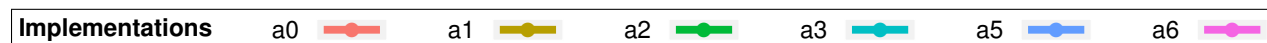


Figure 16: Key legend of the graphs

is strongly related to both performance of the considered algorithm and power dissipation of the architecture. Hence we naturally decompose this blurry notion of energy-awareness, or energy efficiency, into these two dimensions.

We study the problem by modeling the behavior of the implementations, from both the performance and the power point of view. In both cases, we run the implementations on some problem inputs, and measure performance and power dissipation so that we can instantiate the parameters of the model. Then, once the model is instantiated we are able to predict the queue implementation’s energy efficiency on any problem instance.

On the one hand, the performance-related metric that we use in this deliverable is throughput, which represents the number of successful operations per unit of time (here, per second). We consider the dequeuing of an element from the queue as the successful operation, and the measurement is a simple counter of successful operations. On the other hand, the power dissipation measurement is done through Intel’s RAPL energy counters read via the PAPI library [14, 79], as explained in Section 3.1.1.

Finally, we combine those two metrics into an energy-efficiency-related one: energy consumed per successful operation, which is the ratio between power dissipation and throughput.

We dispose of a framework in which we have implemented the most well-known queue

algorithm of the literature. We consider the following implementations of queues, which are described in some detail in Section 4.1.5 below:¹

- **a0**. Lock-free and linearizable queue by Michael and Scott [60],
- **a1**. Lock-free and linearizable queue by Valois [76],
- **a2**. Lock-free and linearizable queue by Tsigas and Zhang [72],
- **a3**. Lock-free and linearizable queue by Gidenstam et al. [29],
- **a5**. Lock-free and linearizable queue by Hoffman et al. [45],
- **a6**. Lock-free and linearizable queue by Moir et al. [62].

We use the same legend for all the graphs in this section, except from Figure 18. It is depicted in Figure 16. The idea here is to use as little knowledge as possible about the algorithms to predict their behaviors, so that if a new algorithm is implemented its behavior can be predicted as well, without changing the model that we present in the following sections.

We run a simple benchmark composed of the two functions described in Figure 17. Half of the threads are assigned to be enqueueers while the remaining ones are dequeuers. We disable multi-threading and map separate threads into separate cores, also the number of threads never exceeds the number of cores. In addition, the mapping is done in the following way: when adding an enqueueer/dequeuer pair, they are both mapped on the most filled but non-full socket.

The parallel work shall be seen as a processing activity, pre-processing for the enqueueers before it enqueues an element, and post-processing on an element of the queue for the dequeuers. This activity is presumed to be computation-intensive, that is why we coded it in the benchmark as a sequence of floating point divisions.

As explained previously in Section 1.2, the benchmark represents an application that uses the queue in a steady-state manner; however, the behavior of the queue is likely to differ from one application to the other, according to the amount of work in the parallel section; also, in the experiments, this amount of work will be part of the variables.

Two more clarifications are necessary. On the one hand, when we speak about implementations of the queues, we actually refer to the different implementations of enqueueing and dequeuing functions, along with their memory management schemes. On the other hand, the slowest of the two function calls (enqueue and dequeue) is the bottleneck of performance and hence determines the throughput of the queue. Also when we reason about the “retry-loop” in the following, we imply “retry-loop of the slowest function call”. Notice that only half of the threads are competing for this operation.

We conclude this introduction by defining some parameters that we use extensively in the next subsections. We denote by n the number of running threads that call the same

¹The **a0**, **a1**, etc., designations refer to the micro-benchmark’s naming of the algorithms. Algorithm **a4** is a single-producer/single-consumer queue that is unsuitable for our purposes and has therefore been left out.

operation, and by f the clock frequency of the cores (we only consider the case where all cores share the same clock frequency, even if an asymmetric setting of the frequencies between the sockets could be of interest). We note pw the amount of work in the parallel section, and cw the amount of work in one try of the retry-loop in the considered implementation.

More generally, an exponent “(off)” refers to an inter-socket execution, while “(on)” refers to an intra-socket one; a subscript “lc” denotes an execution in low-contention state.

In the following subsections all experiments and their underlying predictions are done on Chalmers’ platform, which is described in Section 2.1.1. We run the implementations within a set of three frequencies {1.2 GHz, 2.3 GHz, 3.4 GHz}, for all possible even total numbers of threads, from 2 to 16, *i.e.* for $n \in \{1, \dots, 8\}$.

<pre> 1: procedure ENQUEUEER 2: Initialization() 3: while execution time < t do 4: Parallel_Work() 5: Enqueue() </pre>	<pre> 1: procedure DEQUEUEER 2: Initialization() 3: while execution time < t do 4: $res \leftarrow$ Dequeue() 5: if $res \neq$ NULL then 6: Parallel_Work() </pre>
--	--

Figure 17: Queue benchmark

4.1.2 Throughput

We start this section by underlining some interactions between what we called “work” previously, and the actual execution time of those pieces of code according to different parameters. Then we describe the throughput model under two distinct states that the queue can experience, and finally, we exhibit our results.

Preliminaries First of all, we have seen that the parallel section is full of computations, thus the amount of work in it is actually the number of bunches of 10 floating point divisions that we operate; those operations are perfectly scalable, meaning that the time t_{PS} spent in a given parallel section is proportional to $\frac{pw}{f}$.

In order to obtain a stable execution time, which is linear with the number of bunches of divisions despite compiler and runtime optimizations, we have used CPUID and RDTSCP assembly primitives. As a consequence, the multiplicative factor that correlates t_{PS} and pw/f depends on the number of threads that are running on the platform. Finally, the time spent in a parallel section can be computed thanks to:

$$t_{PS} = \frac{pw}{\lambda \times f}, \quad (8)$$

where λ depends on the number of threads.

The execution time of dequeue and enqueue operations is more problematic, for three main reasons. Primo, because of the lock-free nature of the implementations. From a high-level perspective, those two functions calls are both retry-loops: the thread reads a data,

then works with this version of the data, modifies it, and finally tries to operate a *Compare-and-Swap* on it. If the *Compare-and-Swap* fails, then it reads it again, and re-iterates the process. It exits the function when the *Compare-and-Swap* is a success. As the number of retries is unknown, the time spent in the function call is not straightforwardly computable. This behavior leads us to distinguish two cases: low-contention case, where we are ensured that no retry-loop will fail, and high-contention case, where the threads will generally fail one or several times before succeeding. Secundo, in the high-contention case, the threads compete for accessing a shared data, and they wait for some time before actually being able to access data. We name this as the *expansion*, as it leads to an increase in the execution time of one try of the retry-loop. Tertio, the time before obtaining the data changes, depending whether the data is located in the same socket. This last pathology is however benign, if we look at the following experiment.

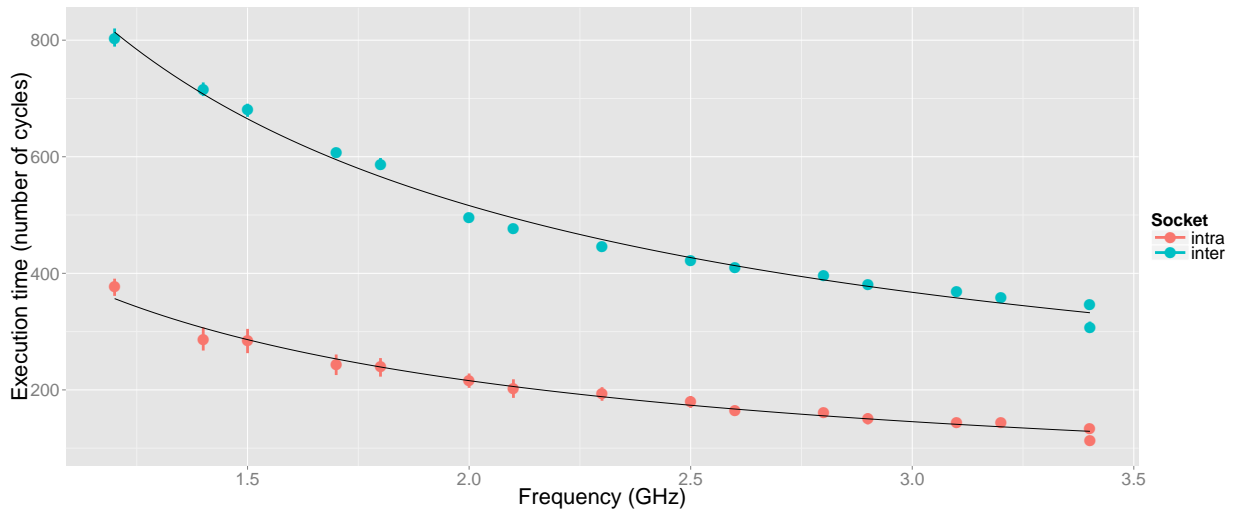


Figure 18: Execution time of *Compare-and-Swap*

We envision the approximation where the retry-loop is a mixed sequence of *Compare-and-Swap* and other shared memory accesses, which is supposedly not too far from the reality. We have measured the execution time of a *Compare-and-Swap* operation, on the one hand when the data is initially located in the same socket as the requester, and on the other hand when the data is in the other socket. In Figure 18, we plot the execution time according to the clock frequency. On-socket, the cost can be fitted with a function $f \mapsto a/f$, while the cost of an off-socket access is fitted by $f \mapsto a'/f + b'$. In other words, the off-socket access includes a non-scalable component that the QPI link is responsible for.

As a consequence, in the low-contention case, *i.e.* when we know that the function call contains only one single try of the retry-loop and that there is no expansion, if we assume that cw is the equivalent of the number of *Compare-and-Swap* inside the retry-loop, we have:

$$\begin{cases} t_{\text{RL}} = cw \times \frac{a}{f}, & \text{if there are not more than 8 cores, and} \\ t_{\text{RL}} = cw \times \left(b' + \frac{a'}{f} \right), & \text{otherwise.} \end{cases} \quad (9)$$

Low Contention We study in this section the low-contention case, *i.e.* when (i) the threads does not suffer from expansion and (ii) a success is obtained with a single try of the retry-loop. As it appears on the scheme in Figure 19, we have a cyclic execution, and the length of the shortest cycle is $t_{PS} + t_{RL}$. Within each cycle, every thread performs exactly one successful operation, thus the throughput is easy to compute thanks to:

$$\mathcal{T} = \frac{n}{t_{PS} + t_{RL}}. \quad (10)$$

This model includes two parameters: λ , which depends only on the number of threads n , and cw (through Equation 9), which depends only on the implementation. The constants of Equation 9 can indeed be determined beforehand.

We determine λ by running anyone of the implementations with a very large parallel work pw_∞ , at a given frequency f_0 , for every number of threads. We note $t_{PS,\infty}$ the execution time of the parallel section of size pw_∞ at the frequency f_0 , we measure the throughput \mathcal{T}_∞ and approximate the Equation 10 with $\mathcal{T}_\infty = \frac{n}{t_{PS,\infty}}$, since the execution time of the parallel section $t_{PS,\infty}$ is such that $t_{PS,\infty} \gg t_{RL}$. Therefore, for each *thr*, we can obtain λ from Equation 8, with the following equation:

$$\lambda = \frac{\mathcal{T}_\infty}{n} \times \frac{pw_\infty}{f_0}.$$

Concerning the amount of work cw in the retry-loop, we have observed that the model is very sensitive to cw , which is why we consider that the amount of work in the retry-loop differs from an intra-socket to an inter-socket execution. We note $cw_{lc}^{(on)}$ the former one and $cw_{lc}^{(off)}$ the latter one. Those two values are obtained by running each implementation in low-contention state. In other words, we pick an amount of parallel work pw_{lc} , which is big enough so that the queue is lowly congested. At frequency f_0 , we run the implementation once with n_0 threads such that $2n_0 > 8$ (leading to the throughput measurement $\mathcal{T}_{lc}^{(off)}$), and once with n_1 threads such that $2n_1 \leq 8$ (leading to the throughput measurement $\mathcal{T}_{lc}^{(on)}$). Then the system in Equation 9 implies that:

$$\left\{ \begin{array}{l} \mathcal{T}_{lc}^{(on)} = \frac{n_0}{\frac{pw_{lc}}{\lambda \times f_0} + cw_{lc}^{(on)} \times \frac{a}{f_0}} \\ \mathcal{T}_{lc}^{(off)} = \frac{n_1}{\frac{pw_{lc}}{\lambda \times f_0} + cw_{lc}^{(off)} \times \left(b' + \frac{a'}{f_0}\right)} \end{array} \right. , \quad \text{hence}$$

$$\left\{ \begin{array}{l} cw_{lc}^{(on)} = \frac{n_0 \times f_0}{a \times \mathcal{T}_{lc}^{(on)}} - \frac{pw_{lc}}{a \times \lambda} \\ cw_{lc}^{(off)} = \frac{1}{b' + a'/f_0} \times \left(\frac{n_1}{\mathcal{T}_{lc}^{(off)}} - \frac{pw_{lc}}{\lambda \times f_0} \right) \end{array} \right.$$

Finally, given a frequency f , a parallel work pw , and number of threads n , the evaluation of the throughput in low-contention state is done thanks to:

$$\left\{ \begin{array}{l} \mathcal{T} = \frac{n}{\frac{pw}{\lambda f} + \frac{a \times cw_{lc}^{(on)}}{f}} \quad \text{if } 2n < 9 \\ \mathcal{T} = \frac{n}{\frac{pw}{\lambda f} + cw_{lc}^{(off)} \times (b' + a'/f)} \quad \text{if } 2n > 8 \end{array} \right. \quad (11)$$

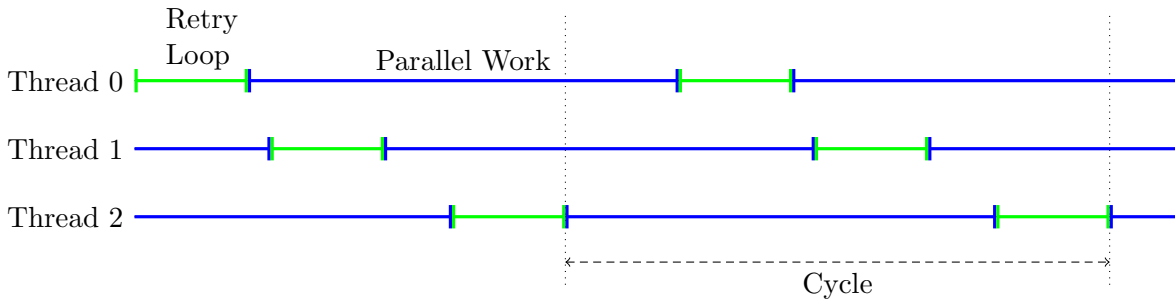


Figure 19: Cyclic execution under low contention

High Contention In the preliminaries, we have explained why the evaluation of the throughput is complex when contention is high: because of the expansion that changes the execution time of one try of the retry-loop, and because the number of tries before a success in the retry-loop is variable.

However, in previous studies, we have seen that the throughput is approximately linear with the expected number of threads that are in the retry-loop at a given time. In addition, this expected number is almost proportional to the amount of work in the parallel section. As a result, a good approximation of the throughput, in high-congestion cases, is a function that is linear with the amount of work in the parallel section.

There remains that the way the threads interfere in the chip, hence the relation between the slope of this straight line and the different parameters, is very dependent on the architecture. That is why, for each frequency, each number of threads and each implementation, we interpolate this line by measuring the throughput for two small amounts of work in the parallel section.

Frontier We now have to estimate when the queue is highly congested and when it is not. We recall that, generally speaking, long parallel sections lead to a low-congested queue since threads are most of the time processing some computations and do not try to access to

the shared data. Reversely, when the parallel section is short, the ratio of time that threads spend in the retry-loop is higher, and gets even higher because of both expansion and retries.

That being said, there exists a simple lower bound on the amount of work in the parallel section, such that there exists an execution where the threads are never failing in their retry-loop. Let us note $t_{\text{RL,LC}}$ the execution time of the retry-loop in low-contention case (we recall that we are able to compute this value as we know the amount of work in the retry-loop), and its relation with the clock frequency. We plot in Figure 20 an ideal execution with $n = 3$ threads and $t_{\text{PS}} = (n - 1) \times t_{\text{RL,LC}}$. In this execution, all threads always succeed at their first try in the retry-loop. Nevertheless, if we make the parallel section shorter, then there is not enough parallel potential any more, and the threads will start to fail: the queue enters the high-congested state.

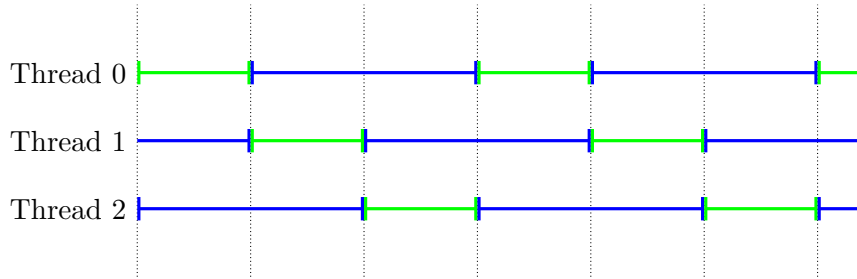


Figure 20: Critical contention

This lower bound ($t_{\text{PS}} = (n - 1) \times t_{\text{RL,LC}}$) is actually a good approximation for the critical point where the queue switches its state. Altogether, we evaluate the throughput in the following way:

- If the execution is intra-socket, *i.e.* if $2n \leq 8$, then
 - if $t_{\text{PS}} \geq (n - 1) \times t_{\text{RL,LC}}^{(\text{on})}$, use Equation 10
 - if $t_{\text{PS}} < (n - 1) \times t_{\text{RL,LC}}^{(\text{on})}$, use Equation 11;
- If the execution is inter-socket, *i.e.* if $2n > 8$, then
 - if $t_{\text{PS}} \geq (n - 1) \times t_{\text{RL,LC}}^{(\text{off})}$, use Equation 10
 - if $t_{\text{PS}} < (n - 1) \times t_{\text{RL,LC}}^{(\text{off})}$, use Equation 11.

Results The throughput prediction is plotted in Figure 21 (we recall that the key can be found in Figure 16, page 36). Points are measurements, while lines are predictions. We will follow this rule for all comparisons between prediction and measurement. In the actual execution, the queue goes through a transient state when the amount of work in parallel section is near the critical point, but the prediction is not so far from it.

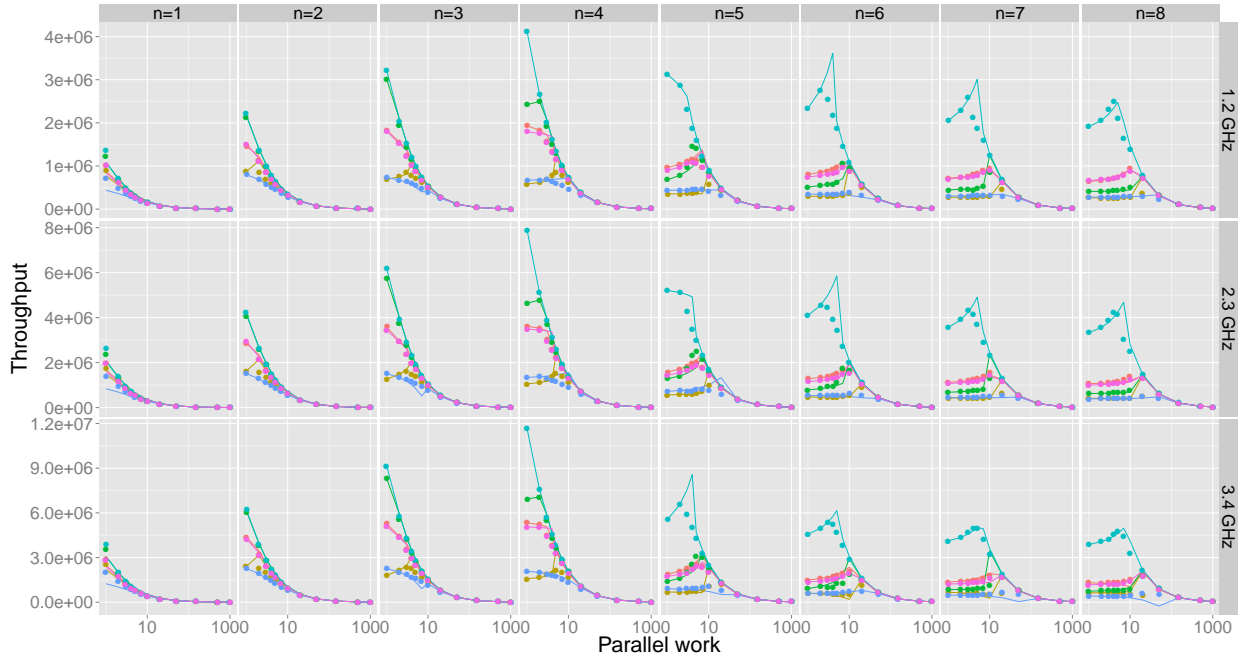


Figure 21: Throughput

Let us summarize the measurements that are needed to instantiate the model, and note M the number of measurements, decomposed into groups M_i . We denote by N , F and A the cardinality of the sets of all possible numbers of respectively threads, frequencies and implementations where we want to predict the throughput.

- λ is found by running only one implementation, for all number of threads, hence $M_1 = N$ measurements.
- For low-contention case, for every implementation, for two different number of threads we run the benchmark to obtain $cw_{lc}^{(on)}$ and $cw_{lc}^{(off)}$, thus $M_2 = 2A$ measurements.
- For high-contention case, for every implementation, every frequency and every number of threads, we launch two runs, leading to $M_3 = 2AFN$ measurements.

In total, we need $M = \sum_i M_i = N + 2A + 2AFN$ runs. In our case, with three frequencies, six implementations, 8×2 threads, and one second per run, it represents around 300 seconds.

4.1.3 Power Given Throughput Measurement

Power Model We use here the same power model that we have exposed in Section 3.1. We recall that we decompose the power into two orthogonal basis, each base having three dimensions. On the one hand, we define the model base by separating the power into static, active and dynamic power, such that the total power is computed by:

$$P = P_{stat} + P_{active} + P_{dyn}.$$

On the other hand, the measurement base corresponds to the hardware that actually dissipates the power, *i.e.* CPU, memory and uncore. Also, power is obtained by the sum:

$$P = P^{(C)} + P^{(M)} + P^{(U)}.$$

In this section, we study each dimension, in each base, so that we are able to express the power dissipation from any perspective:

$$P = \sum_{X \in \{C, M, U\}} \left(P_{stat}^{(X)} + P_{active}^{(X)} + P_{dyn}^{(X)} \right).$$

Thanks to the micro-benchmarking of Section 3.1, we know already all static and active powers, therefore the whole point of this section will be to determine the dynamic power of CPU, memory, and uncore.

On the other hand, the power model was only tailored for micro-benchmarking. However, in this more involved case of power modeling of data structures, we take a single step further towards a more realistic application: we can see both pairs parallel section- enqueue, and dequeue - parallel section as two operations at a higher level, and we keep the steady-state property, which is important in micro-benchmark philosophy.

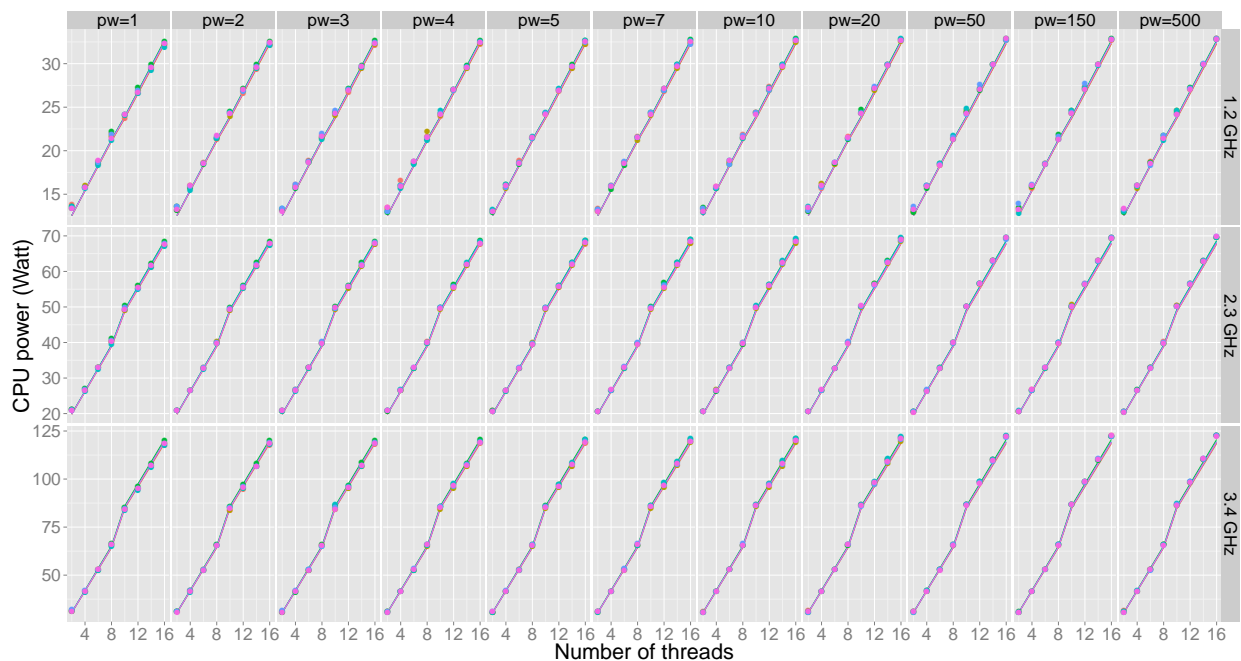


Figure 22: CPU power

CPU Power We have seen in D1.1 [49] that in Chalmers' platform, most of the power is dissipated in the CPU. Additionally, in Section 3.1, we have successfully modeled the

dynamic power of CPU for several operations thanks to the generic formula:

$$P_{dyn}^{(C)} = n \times (A \times f^\alpha + B),$$

where A , B and α are three numbers that depend only on the operation that is executed on the CPU.

We rely on these observations to model the dynamic power of CPU for more complex applications, especially in this deliverable for the queue implementations. We recall that the parallel section is filled with floating point divisions and our assumption such that the retry-loop can be viewed approximately as a sequence of *Compare-and-Swap* has not been checked yet. On the other side, we remark that both *Compare-and-Swap* and floating point divisions are modeled with a similar α , which is around $1.7 = \alpha_0$. As a consequence, we consider now the queue at a higher level and view it as a single complex operation that we can model through:

$$P_{dyn}^{(C)} = n \times (A \times f^{\alpha_0} + B),$$

where A and B have to be determined. One can notice that we have kept the linearity according to the number of threads; this is because all threads in the queue implementation have the same behavior, exactly in the same way as in the micro-benchmark case.

In order to instantiate these parameters, at two frequencies f_0 and f_1 , for a given work in the parallel section and a given number of threads, and for every implementation, we measure the CPU power and extract the dynamic parts p_0 and p_1 . Then we solve the system:

$$\begin{cases} p_0 = n_0 (B + A \times f_0^{\alpha_0}) \\ p_1 = n_0 (B + A \times f_1^{\alpha_0}) \end{cases}, \quad \text{which leads to}$$

$$\begin{cases} A = \frac{p_1 - p_0}{f_1^{\alpha_0} - f_0^{\alpha_0}} \\ B = \frac{p_0 \times f_1^{\alpha_0} - p_1 \times f_0^{\alpha_0}}{f_1^{\alpha_0} - f_0^{\alpha_0}} \end{cases}.$$

The prediction and measurements are plotted in Figure 22.

Memory Power As the retry-loop, which is particular to each implementation, is mainly composed of memory operations, the main difference between the various implementations in terms of power happens in the dynamic power of memory.

Generally speaking, we have shown show in Section 3.1 that the power dissipation of the memory is due to both accesses to main memory and remote accesses to memory. Those accesses are characterized by the amount of data d that is accessed remotely per second, and dynamic power dissipation is considered as proportional to this amount. As in Section 3.1.4, we notice that the data structure does not imply accesses to main memory, hence the power dissipation in memory is only due to remote accesses, which only appear when the threads are spread across sockets.

As the parallel section is full of pure computations, communications can only take place in the retry-loop. We make one last assumption: the amount of data that are accessed per

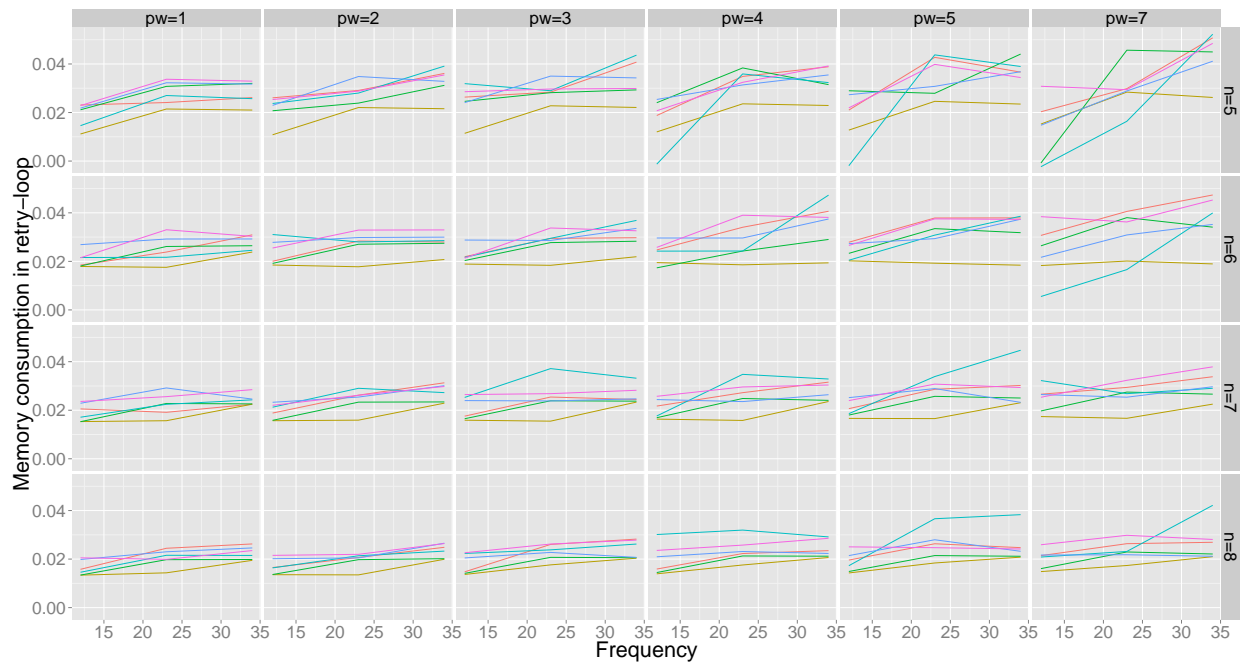


Figure 23: Memory consumption

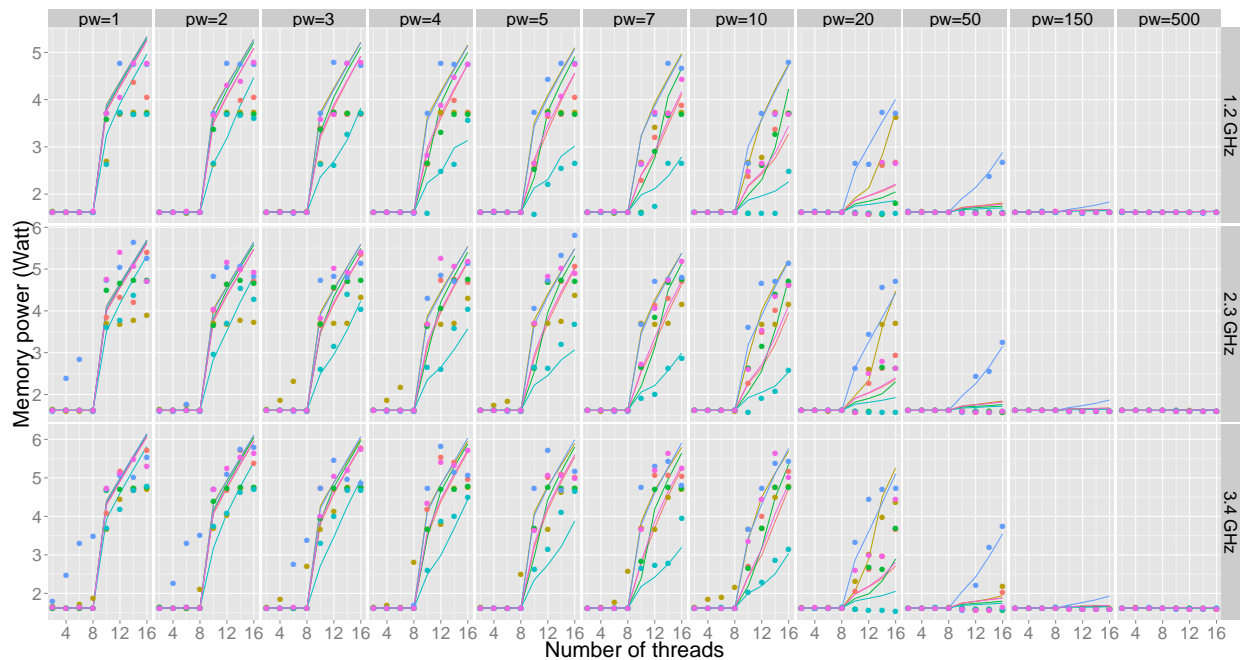


Figure 24: Memory power based on measured throughput

second in a retry-loop depends on the implementation, but given an implementation, once a thread is in the retry-loop, it will always try to access to the same amount of data per

second. When the queue is highly congested, if a thread fails then it will retry and will access the data in the same way as the previous try; and if there is expansion, then the thread will still try to access the data for the whole time it is in the retry-loop. As a consequence, the amount of data that are accessed remotely, hence the dynamic power of memory, is strongly related with the ratio r of the time spent in the retry-loop over the time spent in the parallel section. The dynamic power of memory can be computed by:

$$P_{dyn}^{(M)} = \rho \times n \times r,$$

where ρ is a number depending only on the implementation, and represents the memory access intensity in the retry-loop. We have again n as a multiplicative factor since all the threads have the same behavior.

Now we rely on the structure of the benchmark and on the throughput to compute this ratio. We do not know how many retries are necessary for a thread to successfully exit the retry-loop; however, we know that the thread performs exactly one parallel section per successful operation. As the throughput \mathcal{T} is the number of successful operations per second, the ratio is found by:

$$1 - r = \frac{\mathcal{T} \times pw}{n \times \lambda \times f}. \quad (12)$$

We still have indeed that $\mathcal{T} = n/(t_{PS} + t_{RL})$. This leads to:

$$1 - r = \frac{t_{PS}}{t_{PS} + t_{RL}} = \frac{t_{PS} \times \mathcal{T}}{n},$$

and Equation 12 is derived from the expression of the time spent in the parallel section of Equation 8.

In Figure 23, we plot:

$$\frac{P^{(M)} - P_{stat}^{(M)}}{r \times n},$$

where $P^{(M)}$ is the measured power dissipated by the memory and r is computed through Equation 12. Firstly we remark that the ideas in the model are not contradicted by the graph: everything is almost constant, and the power dissipated by the memory seems to be ruled indeed by the considered ratio. A priori, this ratio should depend on the implementation, but we observe that there is no clear trend, and implementations are very close to each other. This means that all implementations behave in a similar way concerning the amount of data accessed remotely per second.

That is why we only need, if the throughput is known, to run the benchmark for a given implementation, a given size of parallel section, with a given n , at a given frequency f , in order to find the unique ρ , common to all implementations.

The comparison between the measured and the estimated power is plotted in Figure 24. Two noticeable observations should be added: first, as in the micro-benchmark experiments, we remark some steps in the measured power, but we prefer to keep a continuous estimate. Second, we see that implementations **a1** and **a5** sometimes consume memory power even for

intra-socket execution. This could be due to the fact that these versions implement reference counting in their memory management, which could lead to the use of main memory due to overly long chains of unreclaimed nodes.

Uncore Power We predict the uncore power in the same way as the memory power, except that we have an additive component which is linear with the number of threads. This linear component is due to the RDTSCP utilization in the parallel section, and the remaining part may be related to the ring utilization when the threads access the shared data, both inter- and intra-socket.

Briefly, we take the uncore measurement, from which we subtract the static uncore power and the linear component, then operate in the same way and find a new constant ρ' .

Results are pictured in Figure 25, where we notice that, even if the behavior is similar, the amplitude of uncore power variations is relatively smaller than the memory power, and almost negligible in front of CPU power.

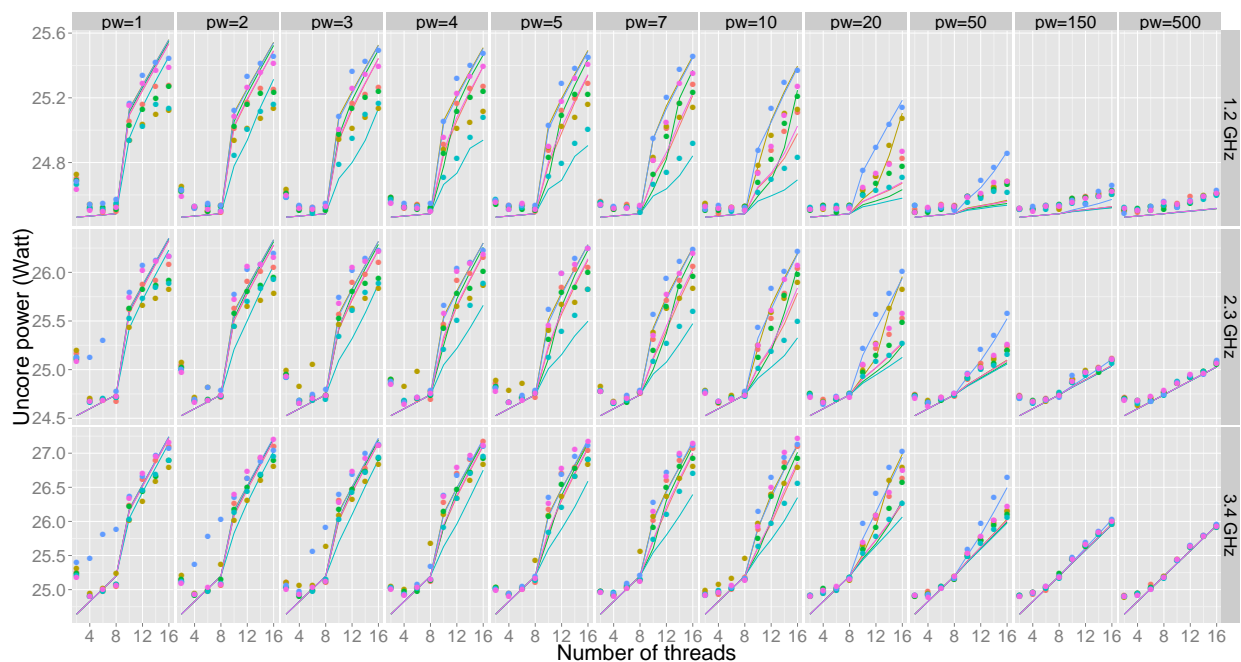


Figure 25: Uncore power

Total Power In terms of number of measurements, estimating the power dissipation is not costly: we can choose a parallel section and n not less than 4, then run each implementation at two different frequencies. This enables the prediction of CPU power, and we can use one of those measurements to predict the memory and uncore powers. Altogether, we only need $M = 2A$ measurements. We plot the comparison of total power in Figure 26 to appreciate the quality of the estimation.

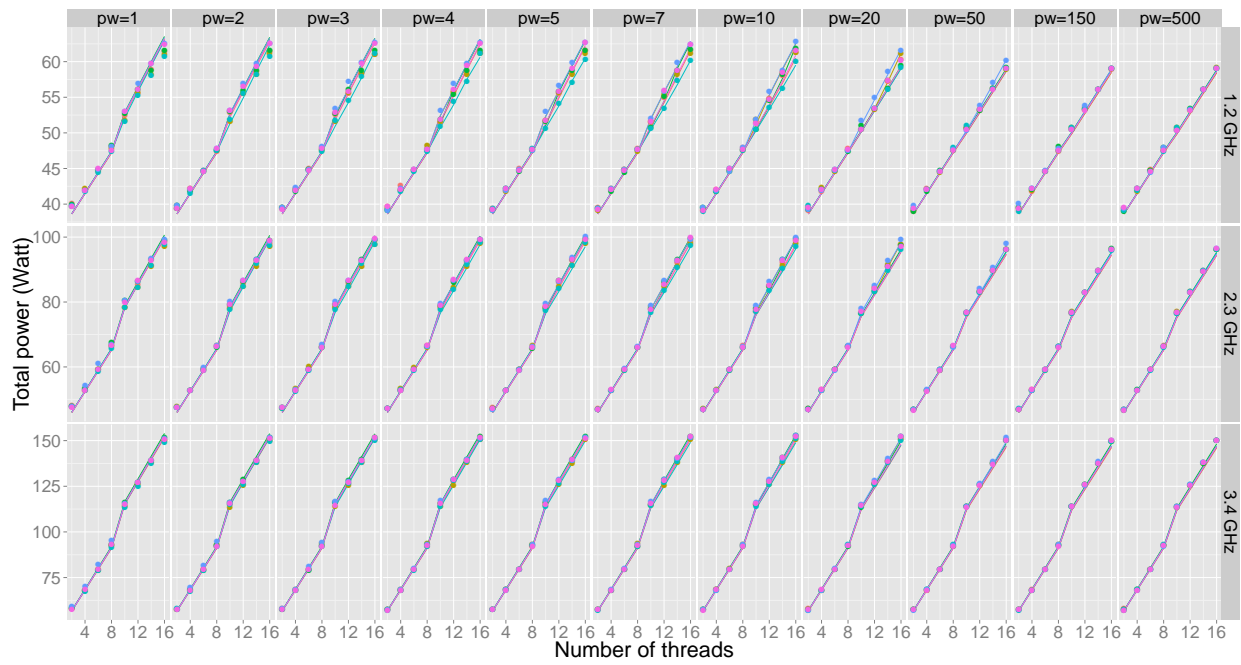


Figure 26: Total power

4.1.4 Complete Prediction

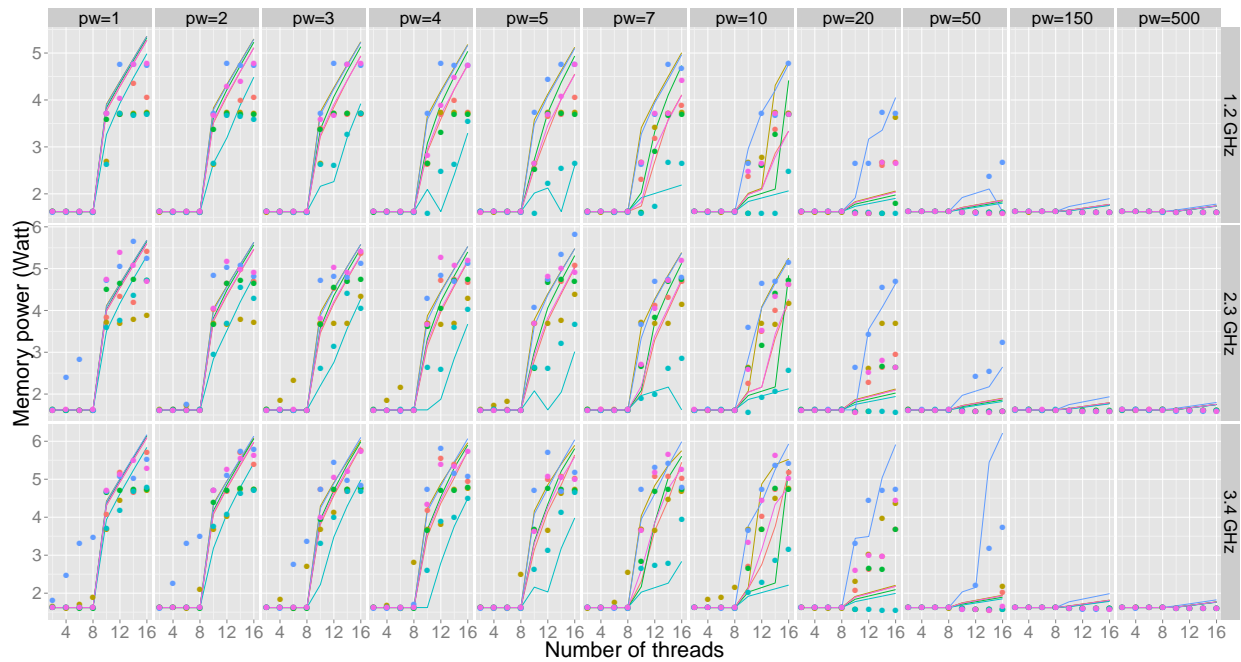


Figure 27: Memory power based on estimated throughput

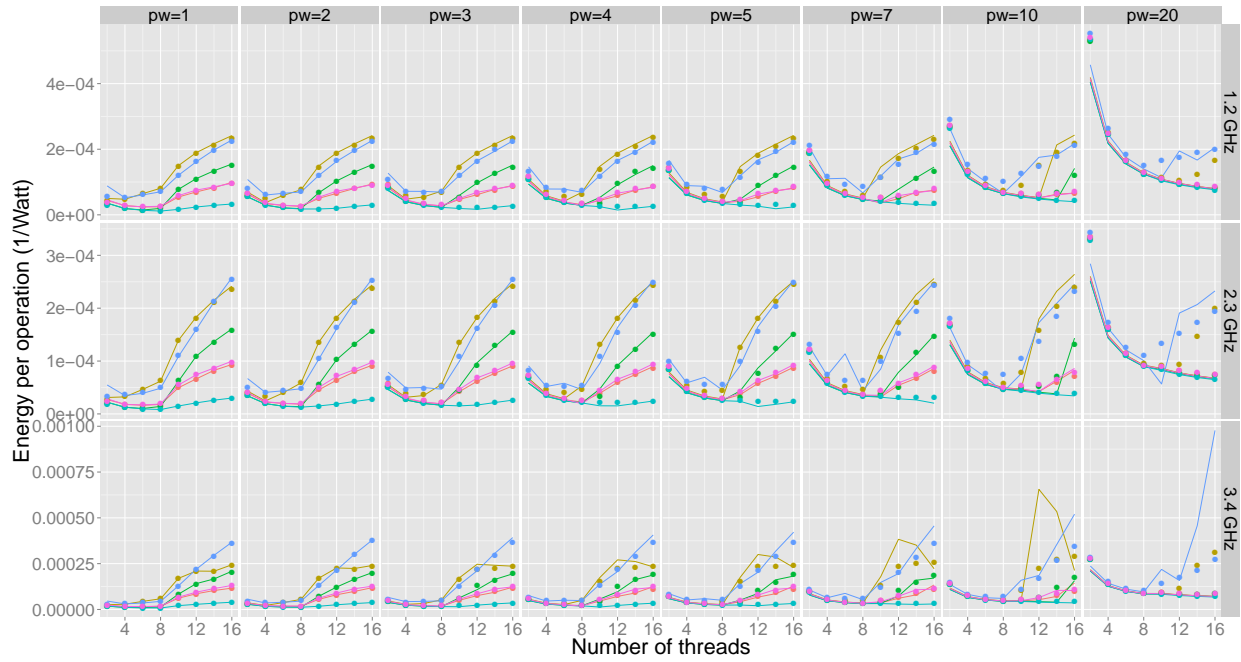


Figure 28: Energy per operation

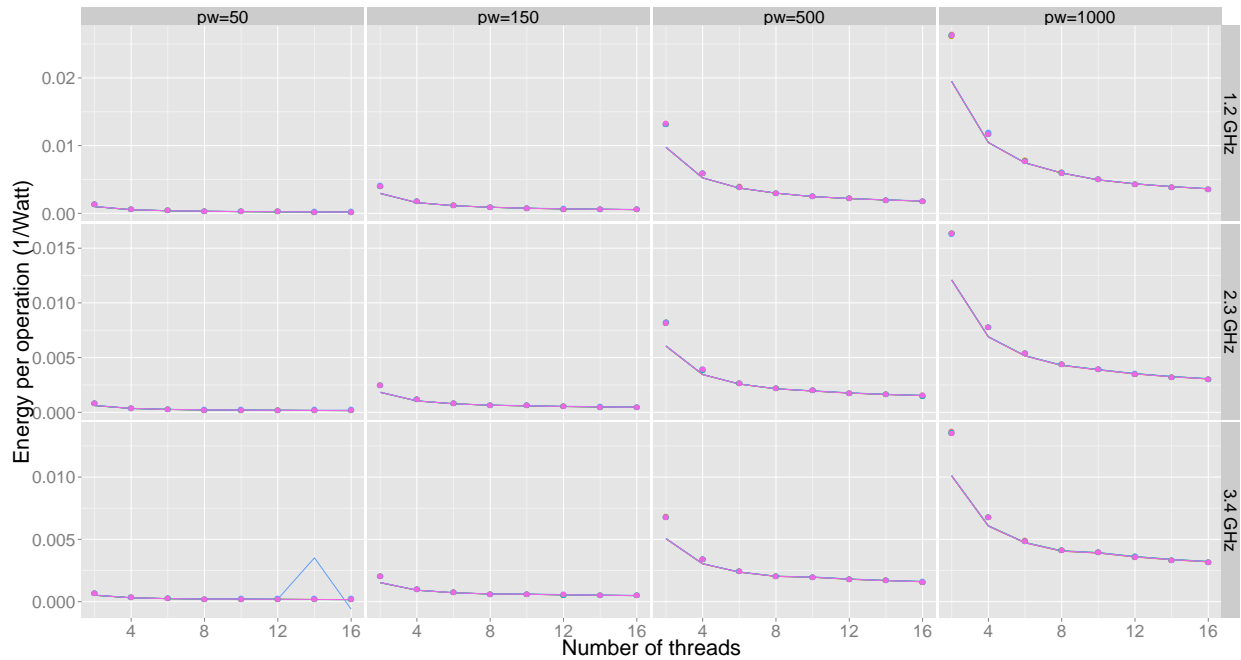


Figure 29: Energy per operation

We plot in Figure 27 the estimate of the memory power dissipation when we use the predicted throughput while computing the ratio in Equation 12, instead of the measured

throughput. It shows that the throughput prediction is good enough, since there is no clear difference in the memory power, and we do not need to rely on the throughput measurement. This is an important property since otherwise, we should have run the benchmark for each value of the variables and measure the throughput, in order to be able to compute the ratio, and then the memory power.

In Figures 28 and 29 is represented the energy per operation. Except from the pathological case with $pw = 20$, the prediction is accurate. This mistake in estimation occurs because the implementations **a1** and **a5** have a long transient behavior between high and low contention cases, and the throughput is harder to estimate in this range.

4.1.5 Description of the Implementations

NOBLE [69, 70] Most of the implementations that we use are part of the NOBLE library. The NOBLE library offers support for non-blocking multi-process synchronization in shared memory systems. NOBLE has been designed in order to: i) provide a collection of shared data objects in a form which allows them to be used by non-experts, ii) offer an orthogonal support for synchronization where the developer can change synchronization implementations with minimal changes, iii) be easy to port to different multi-processor systems, iv) be adaptable for different programming languages and v) contain efficient known implementations of its shared data objects. The library provides a collection of the most commonly used data types. The semantics of the components, which have been designed to be the very same for all implementations of a particular abstract data type, are based on the sequential semantics of common abstract data types and adopted for concurrent use. The set of operations has been limited to those which can be practically implemented using both non-blocking and lock-based techniques. Due to the concurrent nature, also new operations have been added, e.g. Update which cannot be replaced by Delete followed by Insert. Some operations also have stronger semantics than the corresponding sequential ones, e.g. traversal in a List is not invalidated due to concurrent deletes, compared to the iterator invalidation in STL. As the published algorithms for concurrent data structures often diverge from the chosen semantics, a large part of the implementation work in NOBLE, besides from adoption to the framework, also consists of considerable changes and extensions to meet the expected semantics.

The various lock-free concurrent queue algorithms that we include in this study are briefly described below.

Tsigas-Zhang [72] Tsigas and Zhang [72] presented a lock-free extension of [51] for any number of threads where synchronization is done both on the array elements and the shared head and tail indices using *CAS*, and the ABA problem is avoided by exploiting two (or more) null values. In [72] synchronization is done both directly on the array elements and the shared head and tail indices using *CAS*², thus supporting multiple producers and consumers. In

²The Compare-And-Swap (*CAS*) atomic primitive will update a given memory word, if and only if the word still matches a given value (e.g. the one previously read). *CAS* is generally available in contemporary

order to avoid the ABA problem when updating the array elements, the algorithm exploits using two (or more) null values; the ABA problem is due to the inability of *CAS* to detect concurrent changes of a memory word from a value (A) to something else (B) and then again back to the first value (A). A *CAS* operation can not detect if a variable was read to be A and then later changed to B and then back to A by some concurrent processes. The *CAS* primitive will perform the update even though this might not be intended by the algorithm's designer. Moreover, for lowering the memory contention the algorithm alternates every other operation between scanning and updating the shared head and tail indices.

Valois [76] Valois [76, 77] makes use of linked list in his lock-free implementation which is based on the *CAS* primitive. He was the first to present a lock-free implementation of a linked-list. The list uses auxiliary memory cells between adjacent pairs of ordinary memory cells. The auxiliary memory cells were introduced to provide an extra level of indirection so that normal memory cells can be removed by joining the auxiliary ones that are adjacent to them. His design also provides explicit cursors to access memory cells in the list directly and insert or delete nodes on the places the the cursors point to.

Michael-Scott [60] Michael and Scott [60] presented a lock-free queue that is more efficient, synchronizing via the shared head and tail pointers as well as via the next pointer of the last node. Synchronization is done via shared pointers indicating the current head and tail node as well via the next pointer of the last node, all updated using *CAS*. The tail pointer is then moved to point to the new element, with the use of a *CAS* operation. This second step can be performed by the thread invoking the operation, or by another thread that needs to help the original thread to finish before it can continue. This helping behavior is an important part of what makes the queue lock-free, as a thread never has to wait for another thread to finish. The queue is fully dynamic as more nodes are allocated as needed when new items are added. The original presentation used unbounded version counters, and therefore required double-width *CAS* which is not supported on all contemporary platforms. The problem with the version counters can easily be avoided by using some memory management scheme as e.g. [59].

Moir-et-al. [62] Moir et al. [62] presented an extension of the Michael and Scott [60] lock-free queue algorithm where elimination is used as a back-off strategy to increase scalability when contention on the queue's head or tail is noticed via failed *CAS* attempts. However, elimination is only possible when the queue is close to empty during the operation's invocation.

Hoffman-Shalev-Shavit [45] Hoffman et al. [45] takes another approach in their design in order to increase scalability by allowing concurrent *Enqueue* operations to insert the new node at adjacent positions in the linked list if contention is noticed during the attempted

systems with shared memory, supported mostly directly by hardware and in other cases in combination with system software.

insert at the very end of the linked list. To enable these "baskets" of concurrently inserted nodes, removed nodes are logically deleted before the actual removal from the linked list, and as the algorithm traverses through the linked list it requires stronger memory management than [59], such as [28] or [41] and a strategy to avoid long chains of logically deleted nodes.

Gidenstam-Sundell-Tsigas [29] Gidenstam et al. [29] combines the efficiency of using arrays and the dynamic capacity of using linked lists, by providing a lock-free queue based on linked lists of arrays, all updated using *CAS* in a cache-aware manner. In resemblance to [51][27][72] this algorithm uses arrays to store (pointers to) the items, and in resemblance to [72] it uses *CAS* and two null values. Moreover, shared indices [27] are avoided and scanning [72] is preferred as much as possible. In contrast to [51][27][72] the array is not static or cyclic, but instead more arrays are dynamically allocated as needed when new items are added, making the queue fully dynamic.

4.1.6 Towards Realistic Applications: Mandelbrot Set Computation

Mandelbrot Set Description As a simple case-study of parallel applications that use concurrent data structures for communication we have used an application that computes and renders an image of the Mandelbrot set [58] in parallel using the producer/consumer pattern. The program, also used as part of the evaluation in [68], uses a shared collection data structure for communication between the program's two major phases:

- Phase 1 consists of computing the number (with a maximum of 255) of iterations for a given set of points within a chosen region of the image. The results for each region together with its coordinates are then put in the collection data structure.
- Phase 2 consists of, for each computed region stored in the collection, computing the RGB values for each contained point and draw these pixels to the resulting image. The colors for the corresponding number of iterations are chosen according to a rainbow scheme, where low numbers are rendered within the red and high numbers are rendered within the violet spectrum.

Phase 1 is performed in parallel with phase 2, i.e., like a pipeline. Half of the threads perform phase 1 and the rest perform phase 2. The application is implemented in C, and renders a 32-bit color image of 8192 times 8192 pixels of the Mandelbrot set. The size of each square region is chosen to be one of 16×16 (i.e., 16 by 16), 8×8 , 4×4 , or 2×2 pixels which also determines the number of work units and, hence, the level of contention on the shared collection. The whole image is divided into a number (equal to half the number of threads) of larger row-oriented parts³, where each producer thread (i.e., phase 1) work sequentially on the regions contained within its own part. The consumer threads (i.e., phase

³Due to the nature of the Mandelbrot set, this way of deciding each part might not be fair in respect of workload per thread. As can be seen in the experimental results, this partition pattern causes 3 parts to take longer time than 2 parts in parallel, because the total execution time depends on the slowest part.

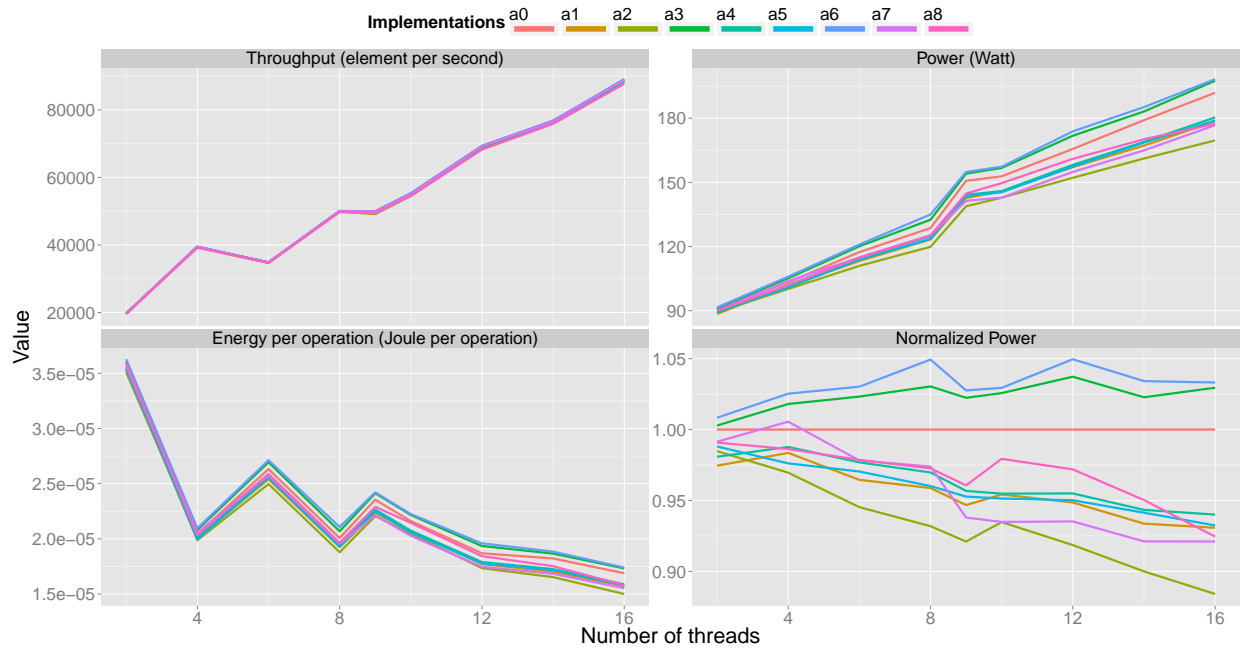


Figure 30: Mandelbrot using 16×16 pixel regions.

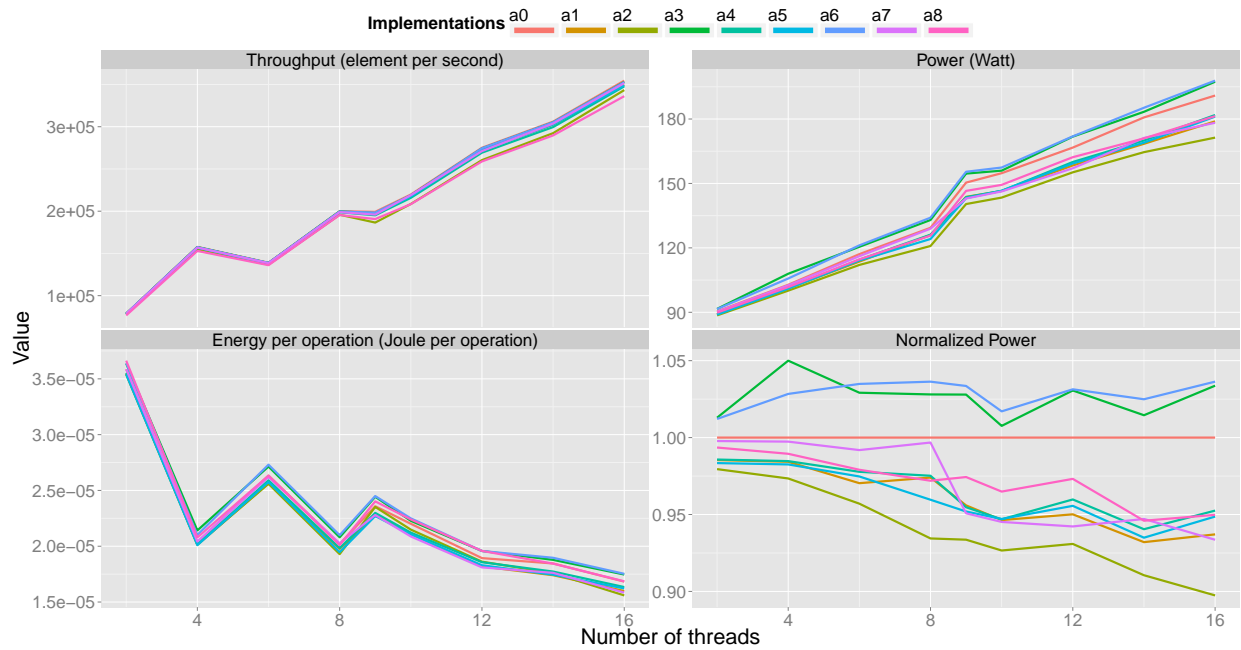


Figure 31: Mandelbrot using 8×8 pixel regions.

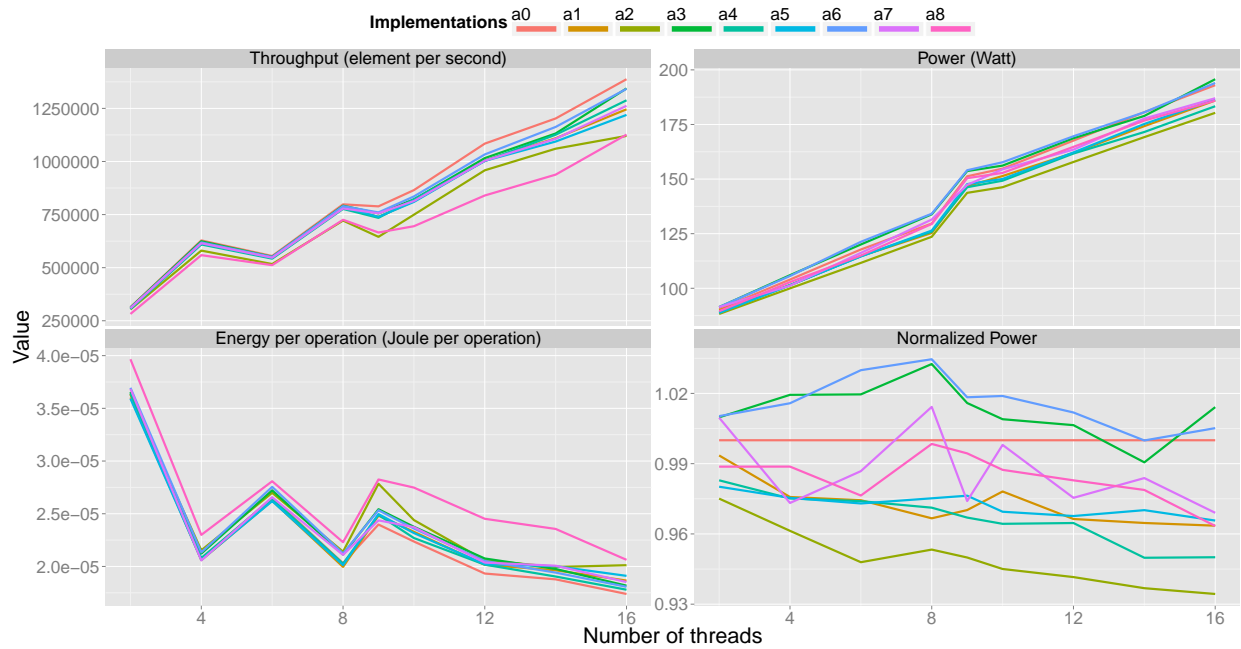


Figure 32: Mandelbrot using 4×4 pixel regions.

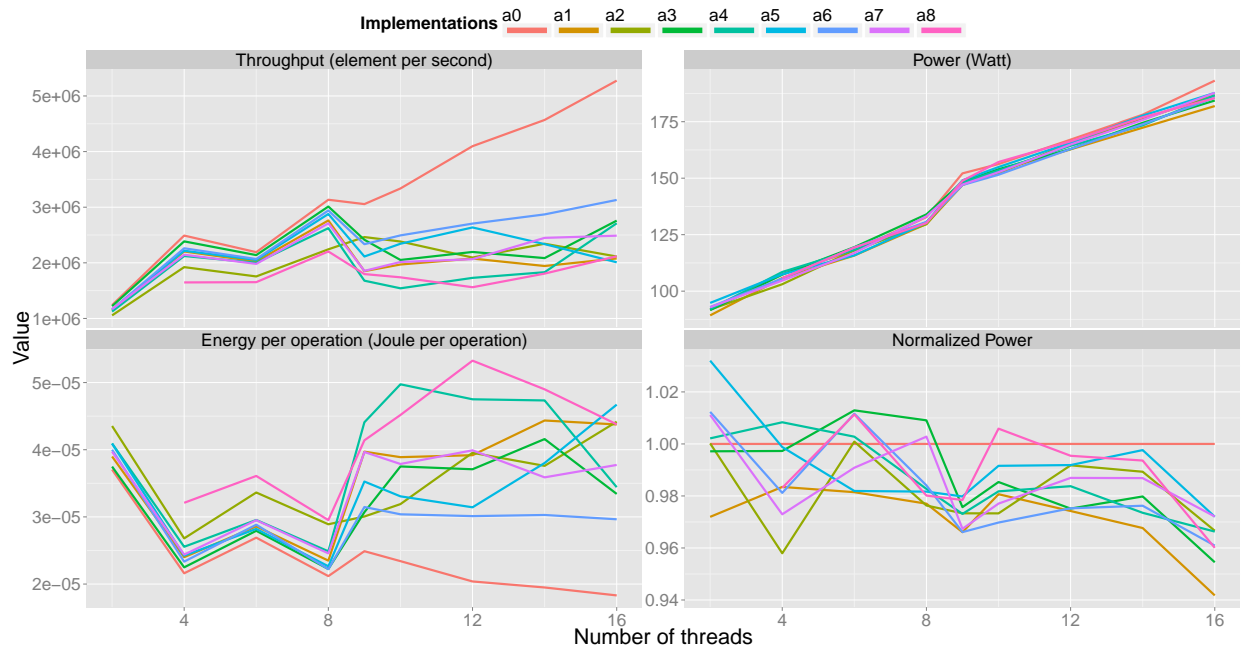


Figure 33: Mandelbrot using 2×2 pixel regions.

2) render the regions got from the collection in the order that they were obtained, until the producer threads have finished and the collection is empty. The application uses a dense pinning strategy, pinning the producers and then the consumers to consecutive cores, e.g. when 16 threads are used the producers are pinned to cores on the first socket while the consumers are pinned to cores on the second socket. This is just one of many possible ways to divide the work and pin threads, it remains as future work to explore other ways.

In this application the shared data structure used for communication only need to offer one *Add* operation that adds an element to the collection and one *TryRemove* operation that removes and returns one element from the collection. The minimal semantic requirements are that at most one *TryRemove* returns each *Added* element and that a surplus of *TryRemove* operations eventually (e.g. after all *Adds* have been issued) returns all *Added* elements.

Any linearizable concurrent queue, stack or bag data structure meets these requirements and could be used as the shared collection. There are even some non-linearizable data structures that could meet them.

The following concurrent shared collection data structures, most of which are described in Section 4.1.5, have been considered:

- **a0**. Lock-free and linearizable bag by Sundell et al. [68].
- **a1**. Lock-free and linearizable queue by Michael and Scott [60].
- **a2**. Lock-free and linearizable queue by Valois [76].
- **a3**. Lock-free and linearizable queue by Tsigas and Zhang [72].
- **a4**. Lock-free and linearizable queue by Hoffman et al. [45].
- **a5**. Lock-free and linearizable queue by Moir et al. [62].
- **a6**. Lock-free and linearizable stack by Michael [59].
- **a7**. Lock-free and linearizable stack by Hendler et al. [40].
- **a8**. Lock-free EDTree (a.k.a. pool or bag) by Afek et al. [2].

Each implementation has been run at each of the 4 work unit sizes (2×2 , 4×4 , 8×8 and 16×16 pixels) and with 2, 6, 8, 9, 10, 12, 14 and 16 threads on the EXCESS server at Chalmers. The results are presented in Figures 30 to 33 in order of decreasing work unit size, i.e. increasing contention. For each case the following metrics are shown (clockwise starting from the top left): i) throughput in pixels per second; ii) total system power in Watts; iii) total system power normalized by **a1** power; and iv) total energy in Joules consumed per pixel.

As mentioned above the method used to divide the Mandelbrot set into regions does not share the work equally among the producer threads which results in the decreases in throughput for 6 and 9 threads.

When the work units are large, such as in Figure 30, the difference in throughput between the different collection implementations is very small indeed for any number of threads. The work load is dominated by independent parallel computation and consequently the level of contention on the collection is low. There is, however, a somewhat larger difference in energy per pixel. This difference is interesting as it ought to be directly related to properties of the collection implementation as all implementations carry out the same total amount of parallel work and a very similar number of successful collection operations per second. Moreover, the lowest energy per pixel costs are achieved by the implementation, **a2**, which is among the worst at high contention (compare with Figure 33). In this particular application the producers do a larger part of the total work than the consumers which can lead to the shared collection becoming empty at times. However, the cost is not distributed equally across all work units – some are cheaper for the producers than others. Consumers finding the shared collection empty will retry the *TryRemove* operation in a tight loop. This could be one reason for the difference in power as the effort needed to determine that the collection is empty varies among the different algorithms. E.g. for **a1** and **a2** this just requires reading a small number of pointers (2 to 3), which however invokes memory barriers, while for **a0** it entails scanning through (while invoking few memory barriers) at least one block of pointers per thread using the data structure.

When the work units are small, such as in Figure 33, there are large differences in throughput from 4 threads and up. This together with the fact that the total system power for the different implementations (at the same number of threads) is even more close together than when using larger work units the differences in energy per pixel varies considerably. Here the contention level on the collection is higher, above 8 threads where the throughput of the less scalable implementations flatten or decrease it can be considered high. In this case all but one of the implementations have their energy per pixel sweetspot at less than or equal to 8 threads (i.e. when using cores in only one socket). Implementation **a0** (the bag) is the only one that delivers the lowest energy per pixel when using all cores of the machine. It is worth noting that the bag data type has a potential to use less synchronization than a queue or stack data type that must enforce an (illusion of) total order among all their elements.

From this case-study some observations can be made about the problem of making an informed selection of implementation for a multi-variant shared data structure in a certain application and context:

- the semantic requirements of the application must be known (naturally) but should also not be overstated as that would limit the choices of implementation;
- the required throughput of data structure operations (and their mix) needs to be predicted (bounded) from the parallel work-load to estimate the level of contention (which if too high would further bound the achievable throughput of data structure operations); and that, consequently,
- a good prediction of achievable data structure operation throughput for each implementation and for a certain state will be needed to do that.

Simplified implementation As mentioned above, realistic applications introduce variety of additional parameters that hardens the estimation of throughput and power. The Mandelbrot application has two main differences from synthetic tests. In the first place, the parallel sections are composed of a mixture of computations and memory accesses. It is hard to estimate the memory access delays and intensity which are important for our model. These metrics are used to determine the parallel section size and the bandwidth requirements which are used to obtain the memory power consumption.

Another complexity regarding the Mandelbrot application is the unequal load balance among producer threads. Even though the problem domain is decomposed into equally sized chunks, some radians require less work than others because they diverge rapidly and require less iterations before determining that they do not belong to the Mandelbrot set. This fact creates variability in the parallel section size which does not occur in the synthetic tests. There are some ways to eliminate this load balance problem. One very simple way is to force each thread to iterate until maximum count even after determining that the point does not belong to the Mandelbrot set. However, this is not an ideal approach since it leads to waste of resources. Instead, one can decompose the domain in an interleaved manner to obtain a better load balance.

For now, we apply the simple approach and leave the interleaved decomposition as future work. Having obtained load balance with simple modifications, we make use of synthetic applications to predict power and throughput values for the Mandelbrot application. We determine the parallel section size and extrapolate the Mandelbrot power and throughput metrics from the corresponding synthetic application.

Mandelbrot Prediction There are slight differences between the simplified Mandelbrot implementation that we consider in this paragraph and the synthetic benchmarks that we have analyzed in the previous subsection. Those differences have an impact on power dissipation through two main components:

- CPU power. What is considered as parallel section in Mandelbrot differs from the synthetic test since the operations that reside in this parallel section are of a different nature: only floating point divisions for synthetic benchmark, and a complete mixture of arithmetic and memory operations for Mandelbrot. The dynamic CPU power that we have measured and extrapolated in synthetic test is then no longer valid for the new application.
- Memory power. Again, as some memory operations take place inside the parallel section, the amount of remote accesses per unit of time in the whole program changes; and we have seen that this metric impacts directly the memory power dissipation.

Because of those variations between the synthetic test and the new application, we need to measure the power dissipation of memory and CPU for some more values of the parameters. This requirement of new power measurements comes however naturally in the process; we cannot expect to be able to predict the power dissipation of any application that uses a queue without having any knowledge about the characteristics of the application according

to power. We are then able to extract from those power measurements both power dissipation of the retry-loop (which is correlated to the queue implementation) and power dissipation of the parallel section (which depends on the application that actually uses one of the queue implementations).

Concerning the CPU power dissipation, we do not reconsider the assumption that it mostly depends on the CPU power dissipated in the parallel section, *i.e.* there is no clear difference of CPU power in the different queue implementations that we have studied in this deliverable. However, contrary to the floating point division case in synthetic benchmark, we do not know what is the relation between frequency and CPU power. We then still rely on the following equation:

$$P^{(C)} = P_{stat}^{(C)} + soc \times P_{active}^{(C)} + n \times P_{dyn}^{(C)}(f, prog), \quad \text{which is equivalent to}$$

$$P_{dyn}^{(C)}(f, prog) = \frac{1}{n} \times \left(P^{(C)} - P_{stat}^{(C)} - soc \times P_{active}^{(C)} \right)$$

Hence, we choose a pattern and a queue implementation, and for each value of the frequency, we run the application and obtain the corresponding value of dynamic CPU power.

Regarding the memory consumption, we rely on a simple memory consumption model, in order to estimate the intensity of remote accesses in the parallel section. In the Mandelbrot application, the number of cache misses in a given parallel section is expected to depend on the pattern that is used; more precisely, it should be roughly proportional to the size of the subregion (2×2 , 4×4 pixels, ...) that is used in the application. We recall that (i) the memory power is computed through $P^{(M)} = P_{stat}^{(M)} + P_{dyn}^{(M)}$, where $P_{dyn}^{(M)}$ is proportional to the number of accesses to remote or main memory per unit of time, (ii) the ratio of time spent in the retry-loop is $r = 1 - (Tpw)/(n\lambda f)$, (iii) in the synthetic benchmark, *i.e.* without memory accesses in the parallel section, thanks to memory power measurements, we have instantiated the value of ρ , such that the dynamic memory power is $P_{dyn}^{(M)} = \rho nr$.

Now, for the parallel section of the Mandelbrot application, which contains memory accesses, we define ρ' , such that $\rho'n$ would be the dynamic memory power if there were only parallel sections and no retry-loop. As the application spends $r\%$ of the time in the parallel section, and $(1-r)\%$ in the retry-loop, the dynamic memory power can be computed through:

$$P_{dyn}^{(M)} = \rho \times n \times (1 - r) + \rho' \times n \times r.$$

Then, we run the application only once to obtain the value of ρ' .

The results are presented in Figure 34, where dashed lines and points are the actual measurements, and plain lines are predictions. The key is again the one represented in Figure 16, on page 36.

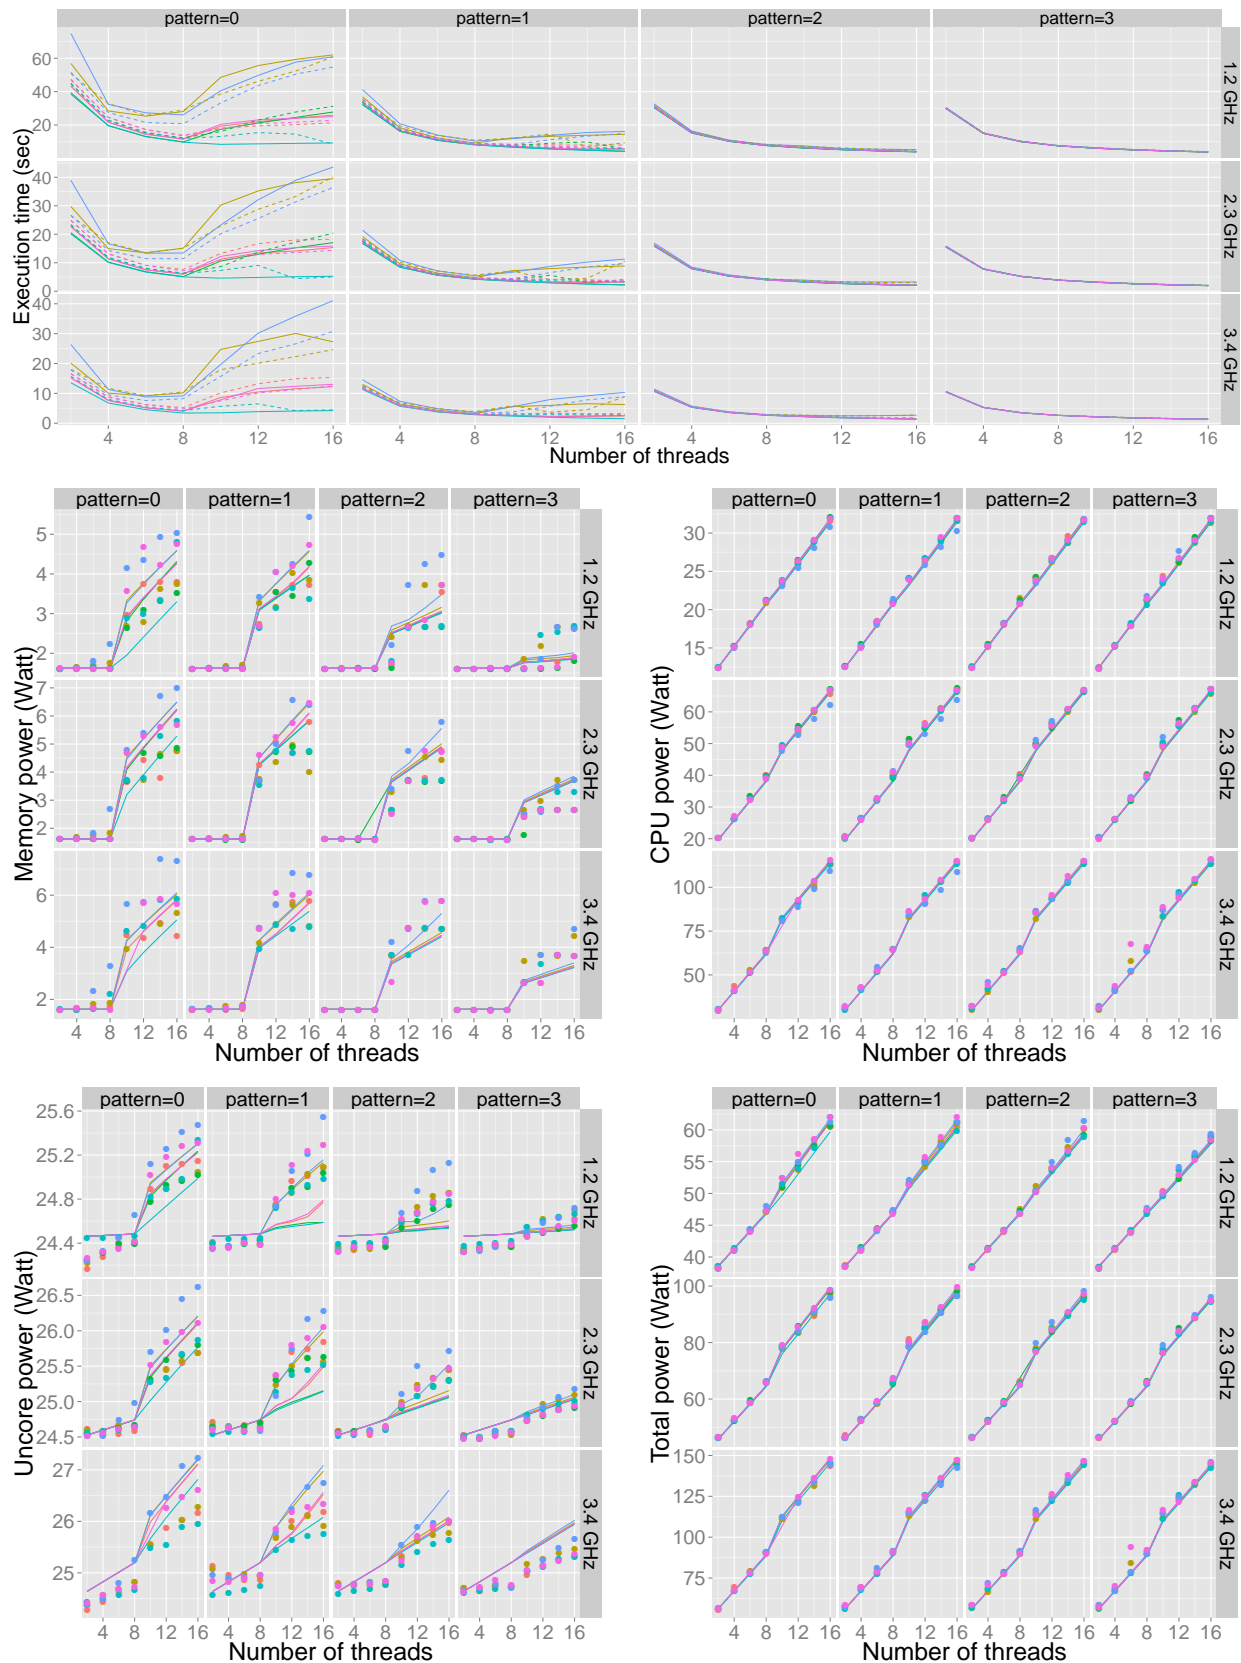


Figure 34: Prediction on Mandelbrot simple implementation

4.2 Concurrent Queues on Movidius Embedded Platform

As described in Section 2.3 the Myriad platform avails a number of different options for synchronizing SHAVEs and the LEON processor. Based on these we propose a number of different concurrent queue implementations and evaluate them experimentally with respect to performance and power consumption.

4.2.1 Concurrent Queue Implementations

In this section we describe the concurrent queue implementations we evaluated on Myriad platform in the context of this work. Concurrent queues are used in a wide range of application domains, especially in the implementation of path-finding and work-stealing algorithms. The queue is implemented as a bounded cyclic array, accessed by all SHAVE cores. SHAVEs request concurrently access to the shared queue for inserting and removing elements.

Table 4 summarizes all different queue implementations we developed and evaluated on the platform. We used three different kinds of synchronization primitives: mutexes, message passing over shared variables and SHAVEs' FIFOs. Mutexes and SHAVEs' FIFOs were described in Section 2.3. With respect to the shared variables, we implemented communication buffers between the processors, used to exchange information for achieving synchronization. To reduce the cost of spinning on shared variables, we allocated these buffers in processor local memories, to avoid the congestion of the Myriad main buses.

The queue implementations can be divided in two basic categories: Lock based and Client-Server.

Lock-based Implementations The lock-based implementations of the concurrent queue utilize the Mutex registers provided by the Myriad architecture. We implemented two different lock-based algorithms: In the first one, a single lock is used to protect the critical section of the *enqueue()* function and a second one to protect the critical section of the *dequeue()*. Therefore, simultaneous access to both ends of the queue can be achieved. The second implementation utilizes only one lock to protect the whole data structure.

Client-Server Implementations The Client-Server Implementations are based on the idea that a server arbitrates the access requests to the critical sections of the application and executes them. Therefore, the clients do not have direct access to the critical section. Instead, they provide the server with the required information for executing the critical section. This set of implementations is an adaptation of the Remote Core Locking algorithm (RCL) [63][56]. In Myriad platform the role of the server can be played either by LEON or a SHAVE core, as shown in Table 4. The SHAVEs are the clients, requesting access to the shared data from the server.

To maximize the efficiency of the Client-Server implementations, each SHAVE allocates the elements to be enqueued in its local CMX slice. Although the CMX is much smaller in comparison with the DDR memory, it provides much smaller access time for the SHAVEs

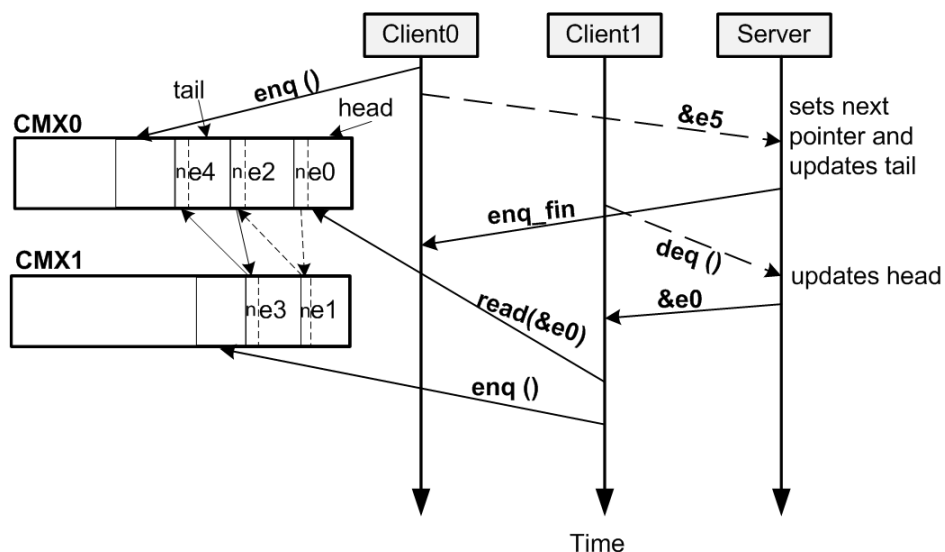


Figure 36: Client-Server Implementation: The server maintains FIFO order of the allocated objects by pointing to the first and the last enqueued elements.

primitive used. We expect SHAVEs' FIFOs primitive to be efficient both in terms of performance and power consumption, since it avoids memory accesses during the exchange of information between the clients and the server. However, the main disadvantage in this case is that it can be only implemented using a SHAVE as a server.

In the second Client-Server implementation we altered the queue structure as follows: The server, instead of managing a queue to store object addresses, utilizes two pointers: *head* and *tail* that point at the first and the last element allocated respectively. Additionally, each element has a *next* pointer which points to the next element, keeping in this way the FIFO order. All these pointers are managed by the server, in order to improve the application parallelism by allowing the clients to perform tasks only outside the critical section.

When a client allocates an object in its local queue stored in CMX memory, it sends the address to the server, as in the first implementation. When the server receives the address, first it updates the *next* pointer of the last element to point to the newly allocated element. Then, it updates the *tail* pointer to point to the new element. (This is the same that happens in a singly linked list FIFO data structure). In the dequeue, as soon as the server receives the request, it sends to the client the address of the element pointed by the *head* and then updates the *head*, using the *next* pointer of the dequeued element.

To illustrate this algorithm, Figure 36 shows an example. Initially, five elements exist in the queue. e_0 is the first allocated and e_4 is the last one. Therefore, head points to e_0 and tail to e_4 . Client0 allocates element e_5 an element in CMX0 and sends the address to the server. The server sets the *next* pointer of e_4 to point to e_5 and updates the tail pointer to e_5 as well. Then, it sends an *enq_fin* message to Client0. Client1 requests a dequeue. The server receives the message, updates the head pointer and sends the address of e_0 to the client.

	Synchronization Primitive		
	Mutex	Shared Var	SHAVE FIFO
no server	Y	-	-
Leon-srv-addr	-	Y	-
SHAVE-srv-addr	-	Y	Y
Leon-srv-h/t	-	-	-
SHAVE-srv-h/t	-	Y	Y

Table 4: Concurrent Queue Implementations: (“Y” indicates the ones that are evaluated in this work.)

In comparison with the *Leon-Srv-addr* or *SHAVE-Srv-addr*, this implementation consumes less memory space. Therefore, the space available in each local CMX slice is affected only by the number of allocated elements of the corresponding client, unlike the previous implementation where the queue of stored addresses reduced the available memory of the slice where it was allocated. It is important to mention that each client accesses only its local CMX slice during the enqueue and the dequeue operations. Only the server makes inter-slice accesses. The disadvantage of this implementation is that it cannot be efficiently implemented using LEON as a server, since it accesses the *next* pointers of each element with high cost. The implementation was designed using both shared variables and SHAVEs’ FIFOs for communication between the server and the clients. In Table 4 is displayed as *SHAVE-srv.ht*.

4.2.2 Experimental Evaluation

The concurrent queue implementations were evaluated using a synthetic benchmark, which is composed by a fixed workload of 20,000 operations and it is equally divided between the running SHAVEs. In other words, in an experiment with 4 SHAVEs each one completes 5,000 operations, while in an experiment with 8 SHAVEs, each one completes 2,500 operations. In the implementations where a SHAVE is utilized as a server, we run the experiments using up to 6 clients (to have the same number of enqueueers and dequeueers).

All algorithms were evaluated in terms of time performance, for the given fixed workload, which is expressed in number of execution cycles. More specifically, in Myriad platform the data flow is controlled by LEON. SHAVEs start their execution when instructed to do so by LEON and then LEON waits for them to finish. The number of cycles measured is actually LEON cycles from the point that SHAVEs start their execution until they all finish. This number represents accurately the execution time. Power consumption was measured using a shunt resistor connected at the 5V power supply cable. Using a voltmeter attached to the resistor’s terminals we calculated the current feeding the board and therefore the power consumed by the Myriad platform.

We performed two sets of experiments for evaluating the behavior of the designs: dedicated SHAVEs and random operations. In the “dedicated SHAVEs” experiment each SHAVE

performs only one kind of operations. In other words, half of the SHAVEs enqueue and half dequeue elements to / from the data structure. “Random operations” means that each SHAVE has equal probability to perform either an enqueue or a dequeue each time it prepares its next operation.

Execution Time Evaluation In this subsection we present the experimental results for performance. *mtx-2-locks* is the lock-based queue implementation with 2 locks, while *mtx-1-lock* is the same implementation with a single lock. *leon-srv-addr* refers to the Client-Server implementation, where the server is LEON and stores the addresses of the objects in a queue, while *SHAVE-srv-addr* is the same implementation where a SHAVE is the server. *leon-srv-ht* and *shave-srv-ht* refers to the Client-Server implementation where the server (LEON and a SHAVE respectively) manages a head and a tail pointer to maintain the FIFO order. Finally, “shared-var” means that the communication is achieved using a shared buffer (i.e. shared variables) and “sf” means that the communication is made through the SHAVEs’ FIFOs.

Figure 37 and Figure 38 show the execution time of dedicated SHAVEs and random operations respectively. We notice that the *mtx-2-locks* implementation is the fastest one in the case of 8 SHAVEs and seems to scale well. This is expected, since it provides the maximum possible concurrency. It requires about half the number of execution cycles in comparison with the *mtx-1-lock*.

The SHAVEs’ FIFOs implementations perform well, especially in the case of 4 SHAVEs in the random operations experiment (28.3% in comparison to the *mtx-2-locks*). The utilization

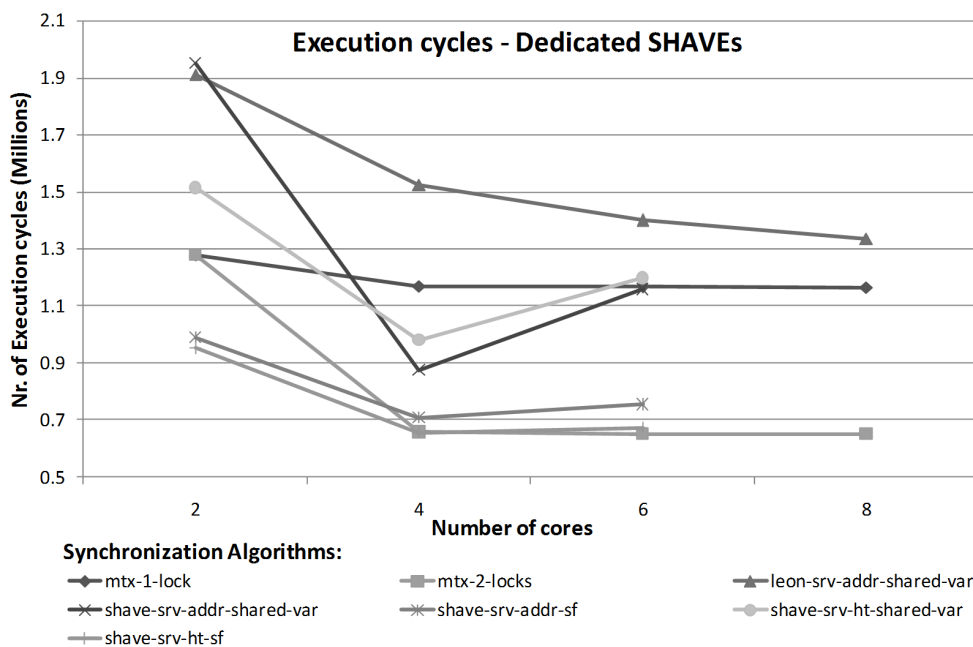


Figure 37: Execution cycles for different synchronization algorithms, when half of the SHAVEs perform enqueue and half dequeue operations.

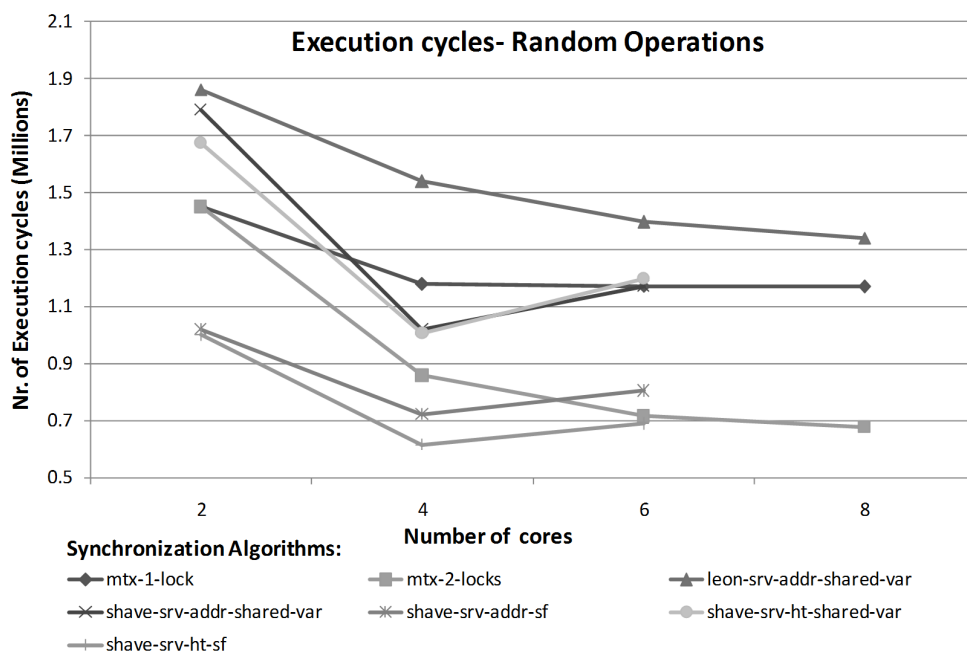


Figure 38: Execution cycles for different synchronization algorithms, when the SHAVES perform randomly enqueue and dequeue operations.

of SHAVES' FIFOs for communication seems to be very efficient in terms of execution time. On the other hand, shared variables provide much lower execution time: For example, *shave-srv-addr-shared-var* leads to 53.3% more execution cycles in comparison with the *shave-srv-addr-sf* in the dedicated SHAVES experiment with 6 SHAVES. Also, the implementations where the server maintains a head and tail pointer performs slightly better in most cases in comparison with the one where the server stores the addresses of the elements (up to 16.7% in random operations for the 6 SHAVES experiment).

In most implementations, we notice a very large drop in the execution time from 2 to 4 SHAVES, due to the increase of concurrency. In other words, in the 2 SHAVES experiments there are time intervals where no client requests access to the shared data. However, in case of 4 clients or more, there is always a SHAVE accessing the critical section. Since the workload is fixed, there is a large drop in the execution time compared to the 2 client experiment. However, since access to the critical section is serialized, the execution time drop for more than 4 SHAVES is much smaller (e.g. in case of *mtx-2-locks*) or even non-existent (e.g. in *mtx-1-lock*).

Finally, in all experiments where a SHAVE is utilized as a server there is an increase in execution time from 4 to 6 SHAVES. The reason for that is the overhead added by inter-slice accesses, which is larger than the decrease in execution time due to increased concurrency. The utilization of LEON as a server using shared variables for communication (i.e. *leon-srv-addr-shared-var*) is inefficient since two overheads are accumulated: LEON is accessing variables in the CMX memory (which is more costly in comparison with SHAVES) and the

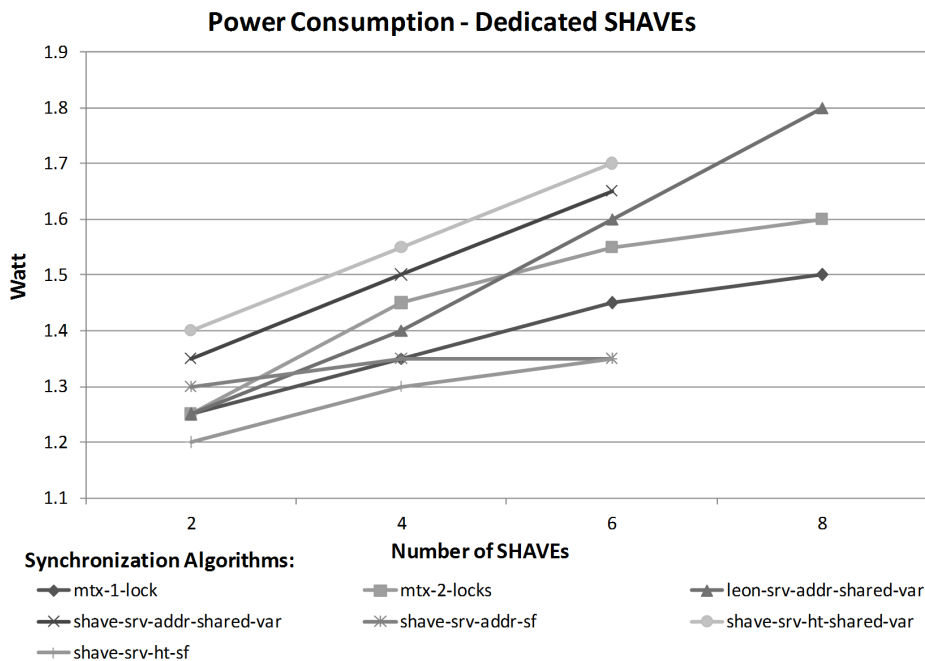


Figure 39: Power consumption for different synchronization algorithms, when half of the SHAVES perform enqueue and half dequeue operations.

spinning on shared variables for achieving communication.

Power Consumption Evaluation As previously stated, power consumption was measured using a shunt resistor connected to the power supply of the platform. Figure 39 and Figure 40 show the power consumption for dedicated SHAVES and random operations. For the 8 SHAVES experiment the most power efficient implementation is the lock-based with a single lock (6.25% in dedicated SHAVES in comparison with the *mtx-2-locks*). However, for a smaller number of clients, the SHAVES' FIFOs implementations are the most power efficient. Indeed, power consumption drops up to 6.8% for 6 SHAVES in the dedicated SHAVES experiment.

We notice that the lock-based implementation with a single lock consumes less power than the 2-lock implementation. This is due to the fact that the power consumption is affected by the number of SHAVES accessing the memory concurrently. In the single lock implementation only one SHAVE accesses the memory for performing operations. However, in the 2-locks implementation there are 2 SHAVES which perform operations concurrently, while the rest are spinning on the locks. Therefore, the 2-lock algorithm consumes more power.

Spinning on a lock consumes very low power, because mutexes are hardware implemented. Microbenchmarking experiments show that 8 SHAVES spinning on a lock concurrently, consume about 20% less power in comparison with the case where 8 SHAVES access the memory concurrently. In fact, this is the case with the shared variables synchronization primitive.

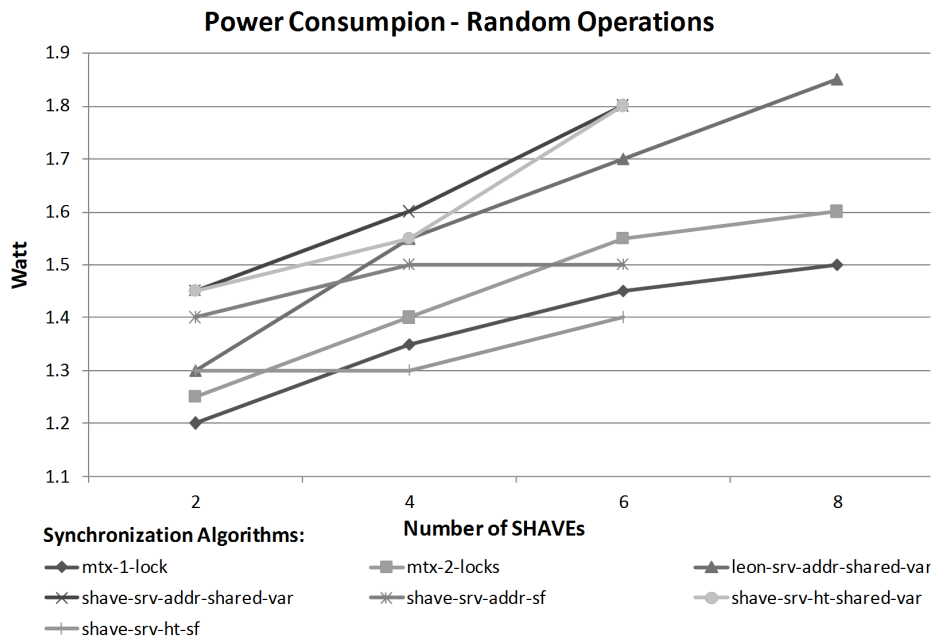


Figure 40: Power consumption for different synchronization algorithms, when the SHAVEs perform randomly enqueue and dequeue operations.

All shared variable implementations consume a lot of power, because spinning on a memory location is energy inefficient, even if the spinning takes place in a local CMX slice.

SHAVEs' FIFOs communication method is the most energy efficient. When a SHAVE tries to write in a full FIFO or read from an empty FIFO stalls, until the FIFO gets non-

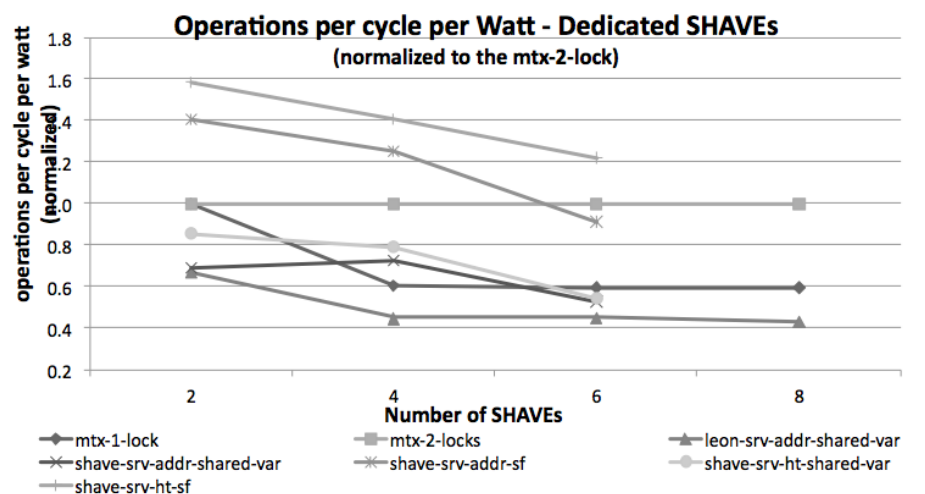


Figure 41: Normalized energy per operation of the synchronization algorithms, when half of the SHAVEs perform enqueue and half dequeue operations.

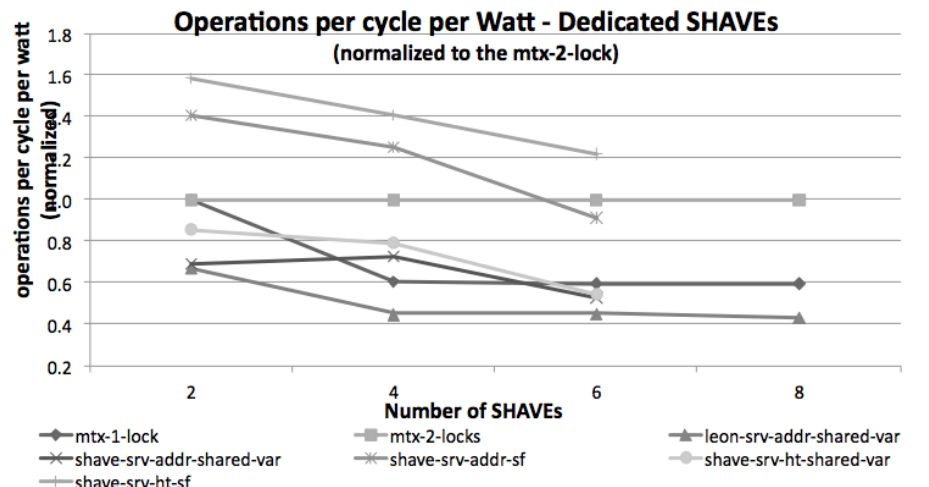


Figure 42: Normalized energy per operation of the synchronization algorithms, when the SHAVES perform randomly enqueue and dequeue operations.

full or non-empty respectively. Microbenchmarking experiments we performed show that 8 SHAVES stalling in a FIFO consume about 28% less power than spinning on a mutex. Indeed, stalling in FIFOs is common in our experiments, where the contention is high. The fact that this synchronization algorithm avoids spinning on memory locations and set SHAVES to stall mode makes it very power efficient.

Energy per operation Evaluation To evaluate in more depth the synchronization algorithms, we present the energy per operation results in Figure 41 and Figure 42 for the dedicated and the random operations experiments respectively. The results are normalized to the *mtx-2-locks* calculated values. We notice that the RCL implementations that utilize the SHAVE’s FIFOs for communication between the clients and the server consume in almost all cases less energy per operation than the *mtx-2-locks*. In the random operation experiment, *shave-srv-ht-sf* consumes 33% less energy per operation in comparison with the *mtx-2-locks*. Indeed, when utilizing SHAVES FIFOs instead of memory buffers for arbitration between the SHAVES, the energy consumption is low. The shared buffer communication scheme is proven to be inefficient in terms of energy consumption. For instance, *leon-srv-addr-shared-var* consumes more than two times energy per operation in comparison with the *mtx-2-locks*.

Discussion The mutex synchronization primitive is indeed efficient for the concurrent queue implementation in terms of both performance and power consumption (especially in the Myriad platform, where mutexes are hardware implemented and very optimized). However, our results show that RCL implementations provide very promising results for the queue implementations and in most cases perform similar to the lock-based ones. We expect that in future MPSoCs, where the number of cores will increase even further, client-server

implementations will become even more efficient.

The reason that the RCL implementations seem an attractive alternative to the lock-based ones is the fact that they transfer computational overhead of the critical sections from the cores, which are the queue workers (clients) to another dedicated core that plays the role of the server. The computational overhead of inserting and removing elements to / from the queue is transferred from the clients to the server. Therefore, while the server executes the critical section, the clients can proceed with other computations, thus increasing the parallelism and reducing the application execution time. Furthermore, by minimizing the communication overhead between the clients and the server (e.g. by utilizing the SHAVE's FIFOs), the results are very satisfactory. On the other hand, in the lock-based implementations, the computational overhead of accessing the queue elements is handled by the workers. However, in this case, simultaneous accesses to the data structure can be achieved, which is obviously not possible in the RCL algorithm. However, with these experiments we show that the RCL algorithm should be evaluated in embedded systems along with the lock-based solutions, especially in applications that use data structures which allow relatively low level of parallelism.

5 Analysis of Energy Consumption of Concurrent Search Trees: Δ Trees

5.1 Introduction

Concurrent trees are fundamental data structures that are widely used in different contexts such as load-balancing [22, 36, 66] and searching [3, 12, 13, 18, 24, 25]. Most of the existing highly-concurrent search trees are not considering the fine-grained data locality. The non-blocking concurrent search trees [13, 25] and Software Transactional Memory (STM) search trees [3, 12, 18, 24] have been regarded as the state-of-the-art concurrent search trees. They have been proven to be scalable and highly-concurrent. However these trees are not designed for fine-grained data locality. Prominent concurrent search trees which are often included in several benchmark distributions such as the concurrent red-black tree [24] by Oracle Labs and the concurrent AVL tree developed by Stanford [12] are not designed for data locality either. It is challenging to devise search trees that are portable, highly concurrent and fine-grained locality-aware. A platform-customized locality-aware search trees [50, 65] are not portable while there are big interests of concurrent data structures for unconventional platforms [37, 34]. Concurrency control techniques such as transactional memory [42, 38] and multi-word synchronization [39, 33, 52] do not take into account fine-grained locality while fine-grained locality-aware techniques such as van Emde Boas layout [64, 78] poorly support concurrency.

5.1.1 I/O model

One of the most cited memory models is the two-level I/O model of Aggarwal and Vitter [4]. In their seminal paper, Aggarwal and Vitter postulated that the memory hierarchy consists of two levels, a main memory with size M (e.g. DRAM) and a secondary memory of infinite size (e.g. disks). Data is transferred in B -sized blocks between those two levels of memory and CPUs can only access data which are available in the main memory. In the I/O model, an algorithm time complexity is assumed to be dominated by how many block transfers are required.

The simplicity and feasibility of this model has made this model popular. However to use this model, an algorithm has to know the B and M parameters in advance. The problem is that these parameters are sometimes unknown and most importantly not portable between platforms. For this I/O model, B-tree [5] is an optimal search tree [17].

Concurrent B-trees [9, 16, 30, 31] are optimised for a known memory block size B (e.g. page size) to minimise the number of memory blocks accessed during a search, thereby improving data locality. In reality there are different block sizes at different levels of the memory hierarchy that can be used in the design of locality-aware data layout for search trees. For example in [50, 65], Intel engineers have come out with very fast search trees by crafting a platform-dependent data layout based on the register size, SIMD width, cache line size, and page size. Namely, existing concurrent B-trees limits its spatial locality optimisation to the memory level with block size B , leaving access to other memory levels with a different

block size unoptimised.

For example in this I/O model, a traditional B-tree that is optimised for searching data in disks (i.e. B is page size), where each node is an array of sorted keys, is optimal for transfers between a disk and RAM (cf. Figure 44c). However data transfers between RAM and last level cache (LLC) is no longer optimal. For searching a key inside each B -sized block in RAM, the transfer complexity is $\Theta(\log(B/L))$ transfers between RAM and LLC, where L is the cache line size. Note that a search with optimal cache line transfers of $O(\log_L B)$ is achievable by using the van Emde Boas layout [11].

5.1.2 Ideal-cache model

The *ideal-cache* model was introduced by Frigo et. al. in [26], which is similar to the I/O model except that the block size B and memory size M are unknown. This paper has coined the term *cache-oblivious algorithms*. Using same analysis of the Aggarwal and Vitter's two-level I/O model, an algorithm is categorised as *cache-oblivious* if it has no variables that need to be tuned with respect to hardware parameters, such as cache size and cache-line length in order to achieve optimality, assuming that I/Os are performed by an optimal off-line cache replacement strategy.

Cache-oblivious algorithms by default have the optimal temporal locality, mainly because of the unknown M . Thus, cache-oblivious algorithms mainly concentrate on optimising spatial locality. Because it is optimal for an arbitrary size of the two levels of memory, a cache-oblivious algorithm is also optimal for any adjacent pair of available levels of the memory hierarchy. Therefore without knowing anything about memory level hierarchy and the size of each level, a cache-oblivious algorithm can automatically adapt to multiple levels of the memory hierarchy.

Empirical results have showed that a cache-oblivious algorithms are often outperform the basic RAM algorithms but not always as good as the carefully tuned (cache-aware) algorithms. However cache-oblivious algorithms perform better on multiple levels of memory hierarchy and are more robust despite changes in memory size parameters compared to the cache-aware algorithms [11].

It is important to note that in the ideal-cache model, B-tree is no longer optimal because of the unknown B . Instead, the trees that are described in the seminal paper [78] by Peter van Emde Boas, are optimal. The van Emde Boas (vEB) tree is an ordered dictionary data type which implements the idea of a recursive structure of priority queues. The efficient structure of the vEB tree arranges data in a recursive manner so that related values are placed in contiguous memory locations (cf. Figure 43). This work has inspired many cache-oblivious data structures such as cache-oblivious B-trees [6, 7, 8] and cache-oblivious binary trees [10]. These researches have demonstrated that vEB layout is suitable for cache oblivious algorithms as it lowers the upper bound on memory transfer complexity.

For example in a system where block size $B = 3$, a search tree with Breadth First Search layout (or BFS tree for short) (cf. Figure 44a) with height 4 will need to do three memory transfers to locate the key in leaf-node 13 in top-down traversing. The first two levels with three nodes (1, 2, 3) will fit within a single block transfer, while the other

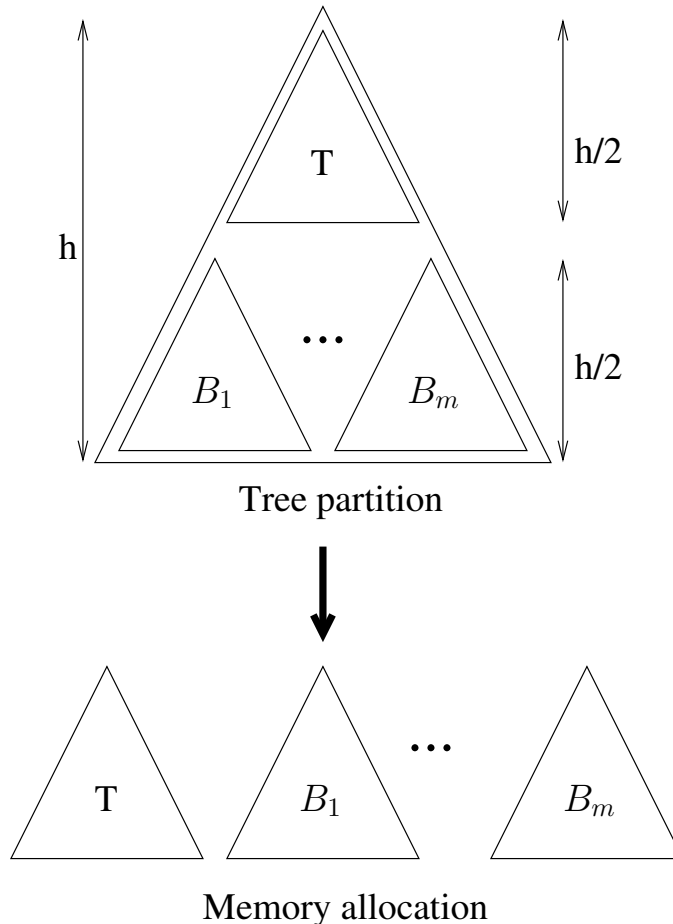


Figure 43: Static van Emde Boas (vEB) layout: a tree of height h is recursively split at height $h/2$. The top subtree T of height $h/2$ and $m = 2^{h/2}$ bottom subtrees $B_1; B_2; \dots; B_m$ of height $h/2$ are located in contiguous memory locations in the order of $T; B_1; B_2; \dots; B_m$.

two levels need to be loaded in two separate memory transfers, where each of the transfer contains (6, 7, 8) and (13, 14, 15) nodes, respectively. Therefore, required memory transfers is $(\log_2 N - \log_2 B) = \log_2 N/B \sim \log_2 N$ for $N \gg B$.

However, for a vEB tree with the same height, the required memory transfers is only two. As seen in Figure 44b, locating the key in leaf-node 12 requires only (1, 2, 3) nodes transfer followed by (10, 11, 12) nodes transfer. This means the transfer complexity is now reduced to $\frac{\log_2 N}{\log_2 B} = \log_B N$, simply by crafting an efficient data structure so that nearby nodes are located in adjacent memory locations. If $B = 1024$, traversing a BFS tree requires 10x more I/Os than a vEB tree.

So far the vEB layout is shown to have $\log_2 B$ less I/Os for two-level memory. On commodity machines where exists multiple-level memory, the vEB layout is outright efficient. In a typical machine having three levels of cache (with 64B cache line size), a RAM (with 4KB page size), and a disk; vEB tree can deliver up to 640x less I/Os than BFS tree,

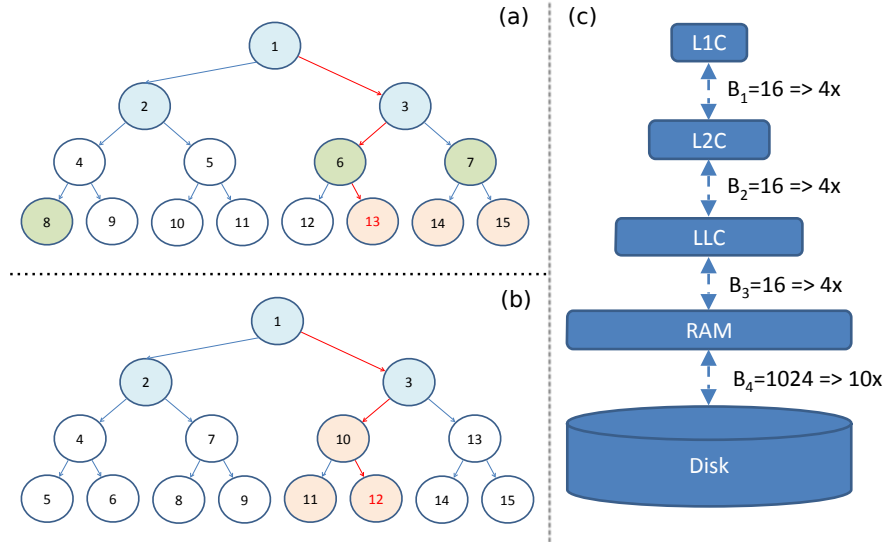


Figure 44: Illustration of required memory transfers in searching for key 13 in (a) BFS tree layout and (b) vEB tree layout. An example of multiple levels of memory is shown in (c). B_x is the block size B between levels of memory.

assuming node size is 4 bytes (Figure 44c).

However, while proven to perform well in searching, vEB trees poorly support concurrent update operations. Inserting or deleting a node in a tree may result in relocating a large part of the tree in order to maintain the vEB layout. The work in [8] has discussed this problem but a feasible implementation hasn't been reported yet [9]. We would like to refer the readers to [11, 26] for a more comprehensive overview of the I/O model and the ideal-cache model.

We introduce Δ Tree family, novel locality-aware concurrent search trees. Δ Tree is an unbalanced locality-aware concurrent search tree of Δ Nodes whose Search, Insert, and Delete are non-blocking to each other, but Insert and Delete may be occasionally blocked by maintenance operations within a Δ Node. Δ Node is a fixed size tree-container with the van Emde Boas layout (cf. Figure 45(left)).

BalancedDT is a balanced Δ Tree with pointer-less Δ Nodes, enabling $\sim 200\%$ more keys to fit in a Δ Node, resulting in 90% improvement in performance and 2x improvement in energy efficiency compared to Δ Tree.

HeterogeneousBDT is a "heterogeneous" BalancedDT where its inner Δ Nodes are leaf-oriented but its leaf Δ Nodes are not, enabling 100% more keys fitting in the leaf Δ Nodes, resulting in 20% improvement in performance and 50% improvement in energy efficiency compared to BalancedDT.

5.2 Overview on the van Emde Boas layout

We propose a modification to the original (static) van Emde Boas (vEB) layout to support high concurrency and fast update operations. This effort results in cache-oblivious concurrent

search trees in the form of a *dynamic* vEB layout. We first define the following notations that will be used to elaborate more on the idea:

- b_i (unknown): block size in terms of the number of nodes at level i of the memory hierarchy (like B in the I/O model [4]), which is unknown as in the cache-oblivious model [26]. When the specific level i of the memory hierarchy is irrelevant, we use notation B instead of b_i in order to be consistent with the I/O model.
- UB (known): the upper bound (in terms of the number of nodes) on the block size b_i of all levels i of the memory hierarchy.
- $\Delta Node$: the largest recursive subtree of a van Emde Boas-based search tree that contains at most UB nodes (cf. dashed triangles of height 2^L in Figure 45). $\Delta Node$ is a fixed-size tree-container with the vEB layout.
- "level of detail" k is a partition of the tree into recursive subtrees of height at most 2^k .
- Let L be the level of detail of $\Delta Node$. Let H be the height of a $\Delta Node$, we have $H = 2^L$. For simplicity, we assume $H = \log_2(UB + 1)$.
- N, T : size and height of the whole tree in terms of basic nodes (not in terms of $\Delta Nodes$).

5.2.1 Static van Emde Boas (vEB) layout

The conventional *static* van Emde Boas (vEB) layout has been introduced in cache-oblivious data structures [6, 7, 8, 10, 26]. Figure 43 illustrates the vEB layout. Suppose we have a complete binary tree with height h . For simplicity, we assume h is a power of 2, i.e. $h = 2^k$. The tree is recursively laid out in the memory as follows. The tree is conceptually split between nodes of height $h/2$ and $h/2 + 1$, resulting in a top subtree T and $m_1 = 2^{h/2}$ bottom subtrees B_1, B_2, \dots, B_m of height $h/2$. The $(m_1 + 1)$ top and bottom subtrees are then located in consecutive memory locations in the order of subtrees T, B_1, B_2, \dots, B_m . Each of the subtrees of height $h/2$ is then laid out similarly to $(m_2 + 1)$ subtrees of height $h/4$, where $m_2 = 2^{h/4}$. The process continues until each subtree contains only one node, i.e. the finest *level of detail*, 0.

The main feature of the vEB layout is that the cost of any search in this layout is $O(\log_B N)$ memory transfers, where N is the tree size and B is the *unknown* memory block size in the I/O model [4] or ideal-cache [26] model. Namely, its search is cache-oblivious. The search cost is the optimal and matches the search bound of B-trees that requires the memory block size B to be known in advance. Moreover, at any level of detail, each subtree in the vEB layout is stored in a contiguous block of memory.

Although the static vEB layout is helpful for utilising data locality, it poorly supports concurrent update operations. Inserting (or deleting) a node at position i in the contiguous block storing the tree may restructure a large part of the tree. For example, inserting new nodes in the full subtree B_1 (a leaf subtree) in Figure 43 will affect the other subtrees B_2, B_3, \dots, B_m by rebalancing existing nodes between B_1 and the subtrees in order to

have space for new nodes. Even worse, we will need to allocate a new contiguous block of memory for the whole tree if the previously allocated block of memory for the tree runs out of space [10]. Note that we cannot use dynamic node allocation via pointers since at *any* level of detail, each subtree in the vEB layout must be stored in a *contiguous* block of memory.

5.2.2 Relaxed cache-oblivious model and dynamic vEB layout

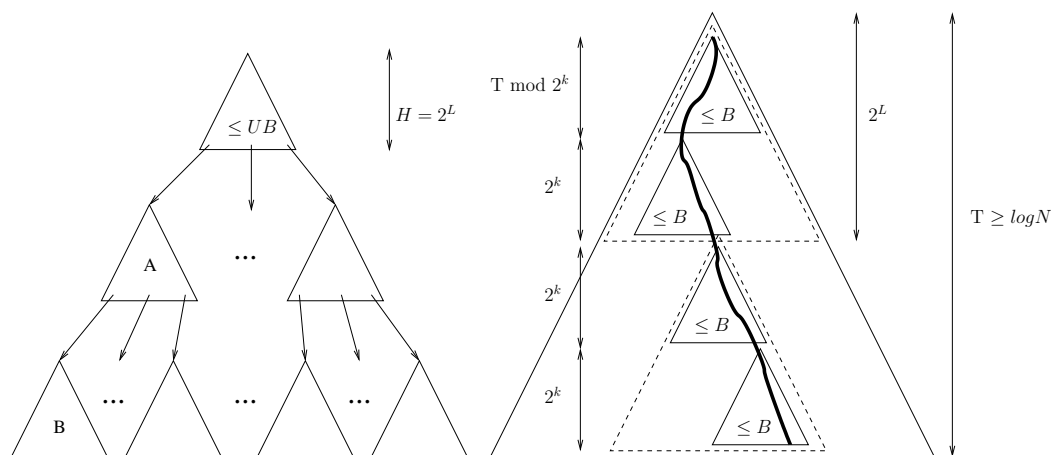


Figure 45: (left): New dynamic vEB layout. (right): Search using dynamic vEB layout.

In order to make the vEB layout suitable for highly concurrent data structures with update operations, we introduce a novel *dynamic* vEB layout. Our key idea is that if we know an upper bound UB on the unknown memory block size B , we can support dynamic node allocation via pointers while maintaining the optimal search cost of $O(\log_B N)$ memory transfers without knowing B (cf. Lemma 5.1).

We define *relaxed cache oblivious* algorithms to be cache-oblivious (CO) algorithms with the restriction that an upper bound UB on the unknown memory block size B is known in advance. As long as an upper bound on all the block sizes of multilevel memory is known, the new relaxed CO model maintains the key feature of the original CO model [26]: First, temporal locality is exploited perfectly as there is no constraints on cache size M in the model. With this an optimal offline cache replacement policy can be assumed. In practice, the Least Recently Used (LRU) policy with memory of size $(1 + \epsilon)M$, where $\epsilon > 0$, is nearly as good as the optimal replacement policy with memory of size M [67]; Second, analysis for a simple two-level memory are applicable for an unknown multilevel memory (e.g. registers, L1/L2/L3 caches and memory). Namely, an algorithm that is optimal in terms of data movement for a simple two-level memory is asymptotically optimal for an unknown multilevel memory. This feature enable designing algorithms that can utilise fine-grained data locality in deep memory hierarchy of modern architectures. In practice, although the exact block size at each level of the memory hierarchy is architecture-dependent (e.g. register size, cache line size), obtaining a single upper bound for all the block sizes (e.g. register size,

cache line size and page size) is easy. For example, a page size obtained from the operating system is such an upper bound.

Figure 45 illustrates the new dynamic vEB layout based on the relaxed cache oblivious model. Let L be the coarsest level of detail such that every recursive subtree contains at most UB nodes. Namely, let H and S be the height and size of such a balanced subtree then $H = 2^L$ and $S = 2^H < UB$. The tree is recursively partitioned into level of detail L where each subtree represented by a triangle in Figure 45, is stored in a contiguous memory block of size UB . Unlike the conventional vEB, the subtrees at level of detail L are linked to each other using pointers, namely each subtree at level of detail $k > L$ is not stored in a contiguous block of memory. Intuitively, since UB is an upper bound on the unknown memory block size B , storing a subtree at level of detail $k > L$ in a contiguous memory block of size greater than UB , does not reduce the number of memory transfers, provided there is perfect alignment. For example, in Figure 45, a travel from a subtree A at level of detail L , which is stored in a contiguous memory block of size UB , to its child subtree B at the same level of detail will result in at least two memory transfers: one for A and one for B . Therefore, it is unnecessary to store both A and B in a contiguous memory block of size $2UB$. As a result, the memory transfer cost of any search in the new dynamic vEB layout is intuitively the same as that of the conventional static vEB layout (cf. Lemma 5.1) while the dynamic vEB supports highly concurrent update operations.

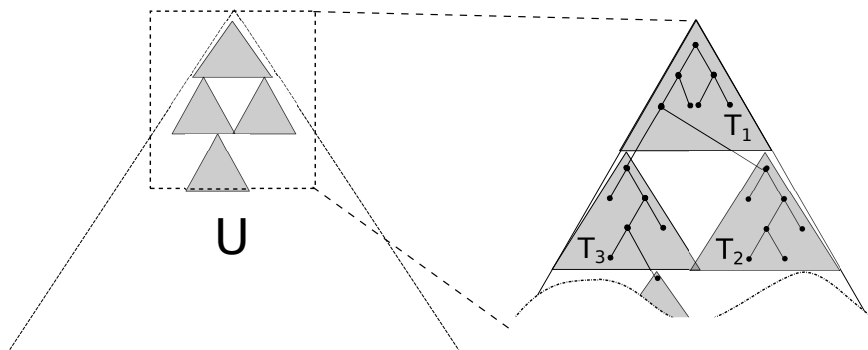


Figure 46: Depiction of a DeltaNode U . Triangles T_x represent the Δ Nodes.

Let Δ Node be a subtree at level of detail L , which is stored in a contiguous memory block of size UB and is represented by a triangle in Figure 45.

Lemma 5.1. *A search in a complete binary tree with the new dynamic vEB layout needs $O(\log_B N)$ memory transfers, where N and B is the tree size and the unknown memory block size in the ideal cache model [26], respectively.*

Proof. (Sketch) Figure 45 illustrates the proof. Let k be the coarsest level of detail such that every recursive subtree contains at most B nodes. Since $B \leq UB$, $k \leq L$, where L is the coarsest level of detail at which every recursive subtree contains at most UB nodes. That means there are at most 2^{L-k} subtrees along the search path in a Δ Node and no subtree of

depth 2^k is split due to the boundary of Δ Nodes. Namely, triangles of height 2^k fit within a dashed triangle of height 2^L in Figure 45.

Due to the property of the new dynamic vEB layout that at any level of detail $i \leq L$, a recursive subtree of depth 2^i is stored in a contiguous block of memory, each subtree of depth 2^k within a Δ Node is stored in at most 2 memory blocks of size B (depending on the starting location of the subtree in memory). Since every subtree of depth 2^k fits in a Δ Node (i.e. no subtree is stored across two Δ Nodes), every subtree of depth 2^k is stored in at most 2 memory blocks of size B .

Since the tree has height T , $\lceil T/2^k \rceil$ subtrees of depth 2^k are traversed in a search and thereby at most $2\lceil T/2^k \rceil$ memory blocks are transferred.

Since a subtree of height 2^{k+1} contains more than B nodes, $2^{k+1} \geq \log_2(B+1)$, or $2^k \geq \frac{1}{2} \log_2(B+1)$.

We have $2^{T-1} \leq N \leq 2^T$ since the tree is a *complete* binary tree. This implies $\log_2 N \leq T \leq \log_2 N + 1$.

Therefore, the number of memory blocks transferred in a search is $2\lceil T/2^k \rceil \leq 4\lceil \frac{\log_2 N+1}{\log_2(B+1)} \rceil = 4\lceil \log_{B+1} N + \log_{B+1} 2 \rceil = O(\log_B N)$, where $N \geq 2$. \square

5.3 Δ Tree implementation

Figure 46 illustrates a Δ Tree U . U uses the dynamic vEB layout presented in Section 5.2.2. The Δ Tree consists of $|U|$ Δ Nodes of fixed size UB each of which contains a *leaf-oriented* binary search tree (BST) $T_i, i = 1, \dots, |U|$. Δ Node's internal nodes are put together in cache-oblivious fashion using the vEB layout.

The Δ Tree U acts as a dictionary of abstract data types. It maintains a set of values which are members of an ordered universe [25]. The Δ Tree U provides the following operations: `INSERTNODE(v, U)`, which adds value v to the set U , `DELETENODE(v, U)` for removing a value v from the set, and `SEARCHNODE(v, U)`, which determines whether value v exists in the set. We use the term *update* operation for either insert or delete operation. We assume that duplicate values are not allowed inside the set and a special value, for example 0, is reserved as an indicator of an `EMPTY` value.

5.3.1 Data structures

The implementation of Δ Tree utilises the data structures defined in Figure 47. The topmost level of Δ Tree is represented by a struct `UNIVERSE` (line 20) that only contains a pointer to the root of the topmost Δ Node.

Each Δ Node that forms the Δ Tree is represented by the struct `Δ NODE` (line 12). Each Δ Node has an associated mirror. This structure consists of a field `opcount`, which is a counter that indicates how many insert/delete threads that are currently operating within that Δ Node; field `locked` that indicates whether a Δ Node is currently locked by maintenance operations: when it is set to *true*, no insert/delete threads are allowed to get in. The `root` pointer serves as the root of a Δ Node, while the pointer `mirror` references root of the

- 1: **Struct node n :**
- 2: member fields:
- 3: $tid \in \mathbb{N}$, if > 0 indicates the node is *root* of a Δ Node with an id of tid (T_{tid})
- 4: $value \in \mathbb{N}$, the node value, default is **empty**
- 5: $mark \in \{true, false\}$, a value of **true** indicates a logically deleted node
- 6: $left, right \in \mathbb{N}$, left / right child pointers
- 7: $isleaf \in true, false$, indicates whether the node is a leaf of a Δ Node, default is **true**

- 8: **Struct triangle S :**
- 9: member fields:
- 10: $nodes$, a group of pre-allocated node n $\{n_1, n_2, \dots, n_{UB}\}$
- 11: $buffer$, an array of value with a length of the current number of threads

- 12: **Struct Δ Node T :**
- 13: member fields:
- 14: $locked$, indicates whether a Δ Node is locked
- 15: $opcount$, a counter for the active update operations
- 16: $root$, pointer the root node of the Δ Node ($S_x.n_1$)
- 17: $rootbuffer$, pointer the buffer of the Δ Node ($S_x.buffer$)
- 18: $mirror$, pointer to root node of the Δ Node's mirror ($S_{x'}.n_1$)
- 19: $mirrorbuffer$, pointer to buffer of the Δ Node's mirror ($S_{x'}.buffer$)

- 20: **Struct universe U :**
- 21: member fields:
- 22: $root$, pointer to the *root* of the topmost Δ Node ($T_1.root$)

Figure 47: Δ Tree's data structures.

Δ Node's mirror. Also there is **rootbuffer** and **mirrorbuffer** pointers that reference the Δ Node's buffer and the mirror's buffer, respectively.

Each NODE structure (line 1) contains field **value**, which holds a value that will guide the search or a data value at a leaf-node. Field **mark** is used to indicate whether a value is logically deleted. A *true* value of **isleaf** indicates a leaf node (as in the leaf-oriented tree), and *false* otherwise. Field **tid** is a unique identifier of a corresponding Δ Node and it is used to let a thread know whether it has moved between Δ Nodes.

5.3.2 Function specifications

```

1: function SEARCHNODE( $v, U$ )
2:    $lastnode, p \leftarrow U.root$ 
3:   while  $p \neq \text{end of tree} \ \& \ p.isleaf \neq \text{TRUE}$  do
4:      $lastnode \leftarrow p$ 
5:     if  $p.value < v$  then
6:        $p \leftarrow p.left$ 
7:     else
8:        $p \leftarrow p.right$ 
9:     if  $lastnode.value = v$  then
10:      if  $lastnode.mark = \text{FALSE}$  then                                 $\triangleright$  lastnode is not deleted
11:        return TRUE
12:      else
13:        return FALSE
14:     else
15:       Search (last visited  $\Delta$ Node's rootbuffer) for  $v$ 
16:       if  $v$  is found then
17:         return TRUE
18:       else
19:         return FALSE

```

Figure 48: Δ Tree's wait-free search algorithm.


```

1: function INSERTNODE( $v, U$ )
2:    $t \leftarrow U.root$ 
3:   return INSERTHELPER( $v, t$ )
4: function DELETENODE( $v, U$ )
5:    $t \leftarrow U.root$ 
6:   return DELETEHELPER( $v, t$ )
7: function DELETEHELPER( $v, node$ )
8:    $success \leftarrow \text{TRUE}$ 
9:   if Entering new  $\Delta$ Node  $T_x$  then
10:     $T'_x \leftarrow \text{GETPARENT}\Delta\text{NODE}(T_x)$ 
11:    DECREMENT( $T'_x.opcount$ )
12:    WAITANDCHECK( $T_x.locked, T_x.opcount$ )
13:    if ( $node.isleaf = \text{TRUE}$ ) OR ( $!node.left$  &
14:     $!node.right$ ) then
15:      if  $node.value = v$  then
16:        if CAS( $node.mark, \text{FALSE}, \text{TRUE}$ ) !=
17:        FALSE) then
18:           $success \leftarrow \text{FALSE}$   $\triangleright$  already deleted!
19:          DECREMENT( $T_x.opcount$ )
20:        else
21:          if ( $node.left.value \& node.right.value = \text{empty}$ )
22:            then
23:              DECREMENT( $T_x.opcount$ )
24:              MERGENODE(PARENTOF( $T_x$ ))  $\leftarrow$ 
25:              TRUE
26:            else
27:               $success \leftarrow \text{DELETEHELPER}(v, node)$   $\triangleright$ 
28:              re-try
29:            else
30:              Search ( $T_x.rootbuffer$ ) for  $v$ 
31:              if  $v$  is found in  $T_x.rootbuffer.idx$  then
32:                 $T_x.rootbuffer.idx \leftarrow \text{empty}$   $\triangleright$  buffered
33:              delete
34:            else
35:               $success \leftarrow \text{FALSE}$ 
36:              DECREMENT( $T_x.opcount$ )
37:            else
38:              if  $v < node.value$  then
39:                 $success \leftarrow \text{DELETEHELPER}(v, node.left)$   $\triangleright$  go
40:                left
41:              else
42:                 $success \leftarrow \text{DELETEHELPER}(v, node.right)$   $\triangleright$ 
43:                go right
44:            return  $success$ 
45: function INSERTHELPER( $v, node$ )
46:    $success \leftarrow \text{TRUE}$ 
47:   if Entering new  $\Delta$ Node  $T_x$  then
48:     $T'_x \leftarrow \text{GETPARENT}\Delta\text{NODE}(T_x)$ 
49:    DECREMENT( $T'_x.opcount$ )
50:    WAITANDCHECK( $T_x.locked, T_x.opcount$ )
51:    if  $node.left$  &  $node.right$  then
52:      if  $v < node.value$  then
53:        if ( $node.isleaf = \text{TRUE}$ ) then  $\triangleright$  insert to
54:        the left:
55:          if CAS( $node.left.value, \text{empty}, v$ ) =
56:          empty then
57:             $node.right.value \leftarrow node.value$ 
58:             $node.right.mark \leftarrow node.mark$ 
59:             $node.isleaf \leftarrow \text{FALSE}$ 
60:            DECREMENT( $T_x.opcount$ )
61:          else
62:             $success \leftarrow \text{INSERTHELPER}(v, node)$   $\triangleright$ 
63:            re-try
64:          else
65:             $success \leftarrow \text{INSERTHELPER}(v, node.left)$   $\triangleright$ 
66:            go left
67:          else if  $v > node.value$  then
68:            if ( $node.isleaf = \text{TRUE}$ ) then  $\triangleright$  insert to
69:            the right:
70:              if CAS( $node.left.value, \text{empty},$ 
71:               $node.value$ ) =
72:              empty then
73:                 $node.right.value \leftarrow v$ 
74:                 $node.left.mark \leftarrow node.mark$ 
75:                ATOMIC {  $node.value \leftarrow v$ 
76:                 $node.isleaf \leftarrow \text{FALSE}$  }
77:                DECREMENT( $T_x.opcount$ )
78:              else
79:                 $success \leftarrow \text{INSERTHELPER}(v, node)$   $\triangleright$ 
80:                re-try
81:              else
82:                 $success \leftarrow \text{INSERTHELPER}(v, node.right)$ 
83:                 $\triangleright$  go right
84:              else if  $v = node.value$  then
85:                if ( $node.isleaf = \text{TRUE}$ ) then
86:                  if  $node.mark = \text{FALSE}$  then  $\triangleright$  is
87:                  deleted?
88:                     $success \leftarrow \text{FALSE}$   $\triangleright$  value exist!
89:                    DECREMENT( $T_x.opcount$ )
90:                  else
91:                    Goto 57  $\triangleright$  goto insert right
92:                  else
93:                     $success \leftarrow \text{INSERTHELPER}(v, node.right)$   $\triangleright$ 
94:                    go right
95:                else
96:                  if  $v$  already in  $T_x.rootbuffer$  then
97:                     $success \leftarrow \text{FALSE}$ 
98:                    DECREMENT( $T_x.opcount$ )
99:                  else
100:                    $put\ v\ inside\ T_x.rootbuffer$   $\triangleright$  buffered
101:                   insert
102:                   if TAS( $T_x.locked$ ) then  $\triangleright$  Acquired
103:                   maintenance lock
104:                     DECREMENT( $T_x.opcount$ )
105:                     SPINWAIT( $T_x.opcount$ )  $\triangleright$  Waits updates
106:                   to finish
107:                     ... do REBALANCE( $T_x$ ) or EXPAND( $node$ )
108:                   ...
109:                   return  $success$ 

```

Figure 49: Δ Tree's update algorithms and their helper functions.

Operation $\text{SEARCHNODE}(v, U)$ (cf. Figure 48) is going to walk over the ΔTree (Figure 48, lines 3–8) to find whether the value v exists in U . This particular operation is guaranteed to be wait-free (cf. Lemma 5.2), and it returns **true** whenever v has been found, or **false** otherwise (Figure 48, line 10). Operation $\text{INSERTNODE}(v, U)$ (cf. Figure 49, line 37) inserts value v at a leaf of ΔTree , provided v does not yet exist in the tree (Figure 49, line 70). Following the nature of a leaf-oriented tree, a successful insert operation will replace a leaf with a subtree of three nodes [25] (cf. Figure 50a and pseudocode in Figure 49, line 46 & 57). $\text{DELETENODE}(v, U)$ (cf. Figure 49, line 7) simply *mark* the leaf that contains value v as deleted (Figure 49, line 15), instead of physically removing the leaf or changing its parent pointer (as in [25]).

Lemma 5.2. *ΔTree search operation is wait-free.*

Proof. (Sketch) The proof can be served based on these observations on Figure 48:

1. SEARCHNODE and invoked SEARCHBUFFER (line 15) don't wait for any locks.
2. The number of iterations in the *while* loop (line 3) is bounded by the *height* of the tree, $\mathcal{O}(N)$.
3. SEARCHBUFFER time complexity is bounded by the buffer size, which is a constant.

Therefore the SEARCHNODE time is bounded by $\mathcal{O}(N)$. □

There is a difference between inserting to the left (Figure 49, lines 46–49) and inserting to the right (Figure 49, lines 57–61) because an insert to the right will need to change the value of the root of the new subtree in order to guide the tree search. And it's not necessary to do that when inserting a value to the left.

Maintenance functions Apart from the basic operations, three maintenance ΔTree operations are invoked in certain cases of inserting and deleting a node from the tree. Operation $\text{REBALANCE}(T_v.\text{root})$ (cf. Figure 49 line 84) is responsible for rebalancing a ΔNode after an insertion. Figure 50a illustrates the rebalance operation. Whenever a new node v is to be inserted at the last level H of ΔNode T_1 , the ΔNode is rebalanced to a complete BST by setting a new depth for all leaves (e.g. a, v, x, z in Figure 50a) to $\log N + 1$, where N is the number of leaves. In Figure 50a, right after the rebalance operation, tree T_1 becomes balanced and its height reduces from 4 to 3.

We also define $\text{EXPAND}(l)$ operation (cf. Figure 49 line 84), that will be responsible for creating a new ΔNode and connecting it to the parent of a leaf node l (cf. Figure 50b). Expanding will be triggered only if after $\text{INSERTNODE}(v, U)$, leaf l will be at the last level of a ΔNode and rebalancing will no longer reduce the current height of the subtree T_i stored in the ΔNode . For example if expanding is taking place at a node l located at the bottom level of the ΔNode (Figure 50b, node l contains value v), or $\text{depth}(l) = H$, then a new ΔNode (T_2 for example) will be referred by the parent of node l . Namely, the parent of l swaps one of its pointer that previously points to l , into the root of the newly created ΔNode , $T_2.\text{root}$.

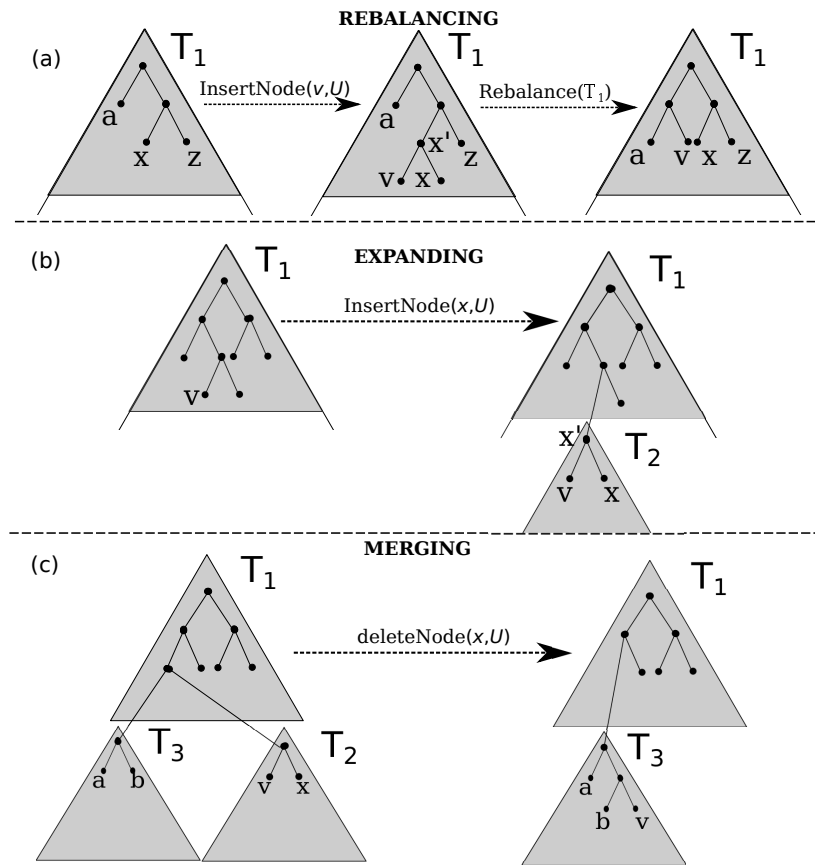


Figure 50: (a) *Rebalancing*, (b) *Expanding*, and (c) *Merging* operations on Δ Tree.

Operation $\text{MERGE}(T_x.\text{root})$ (cf. Figure 49 line 21) is for merging T_x with its sibling after a node deletion. For example in Figure 50c, T_2 is merged into T_3 . Then the pointer of T_3 's grandparent that previously points to the parent of both T_3 and T_2 is replaced by a pointer to T_3 . Merge operation is invoked provided that a particular ΔNode , where a deletion has taken place, is filled less than 2^t of its capacity (where $t = 1/2H$) and all values of that ΔNode and its siblings could be fitted into a ΔNode .

To minimise block transfers required during tree traversal, the height of the tree should be kept minimal. These auxiliary operations are the unique feature of ΔTree in the effort of maintaining a small height.

These INSERTNODE and DELETENODE operations are linearisable to other SEARCHNODE , INSERTNODE and DELETENODE operations (cf. Lemma 5.3 and 5.4). Both of the operations are using single word CAS (Compare and Swap) and "leaf-checking" (cf. Figure 49, line 15 & 19 for delete and 46 & 57 for insert) to achieve that.

Lemma 5.3. *For a value that resides on the leaf node of a ΔNode , SEARCHNODE operation (Figure 48) has the linearisation point to DELETENODE at line 10 and the linearisation point to INSERTNODE at line 9. For a value that stays in the buffer of a ΔNode , SEARCHNODE operation has the linearisation point at line 16.*

Proof. (Sketch) It is trivial to demonstrate this in relation to deletion algorithm in Figure 49 since only an atomic operation is responsible for altering the *mark* property of a node (line 15). Therefore DELETENODE has the linearisation point to SEARCHNODE at line 15.

For SEARCHNODE interaction with an insertion that grows new subtree, we rely on the facts that:

1. a snapshot of the current node p is recorded on *lastnode* as a first step of searching iteration (Figure 48, line 4);
2. *v.value* change, if needed, is not done until the last step of the insertion routine for insertion of $v > \text{node.value}$ and will be done in one atomic step with *node.isleaf* change (Figure 49, line 60); and
3. *isleaf* property of all internal nodes are by default **true** (Figure 47, line 7) to guarantee that values that are inserted are always found, even when the leaf-growing (both left-and-right) are happening concurrently.

Therefore INSERTNODE has the linearisation point to SEARCHNODE at line 46 when inserting a value v smaller than the leaf node's value, or at line 57 otherwise.

A search procedure is also able to cope well with a "buffered" insert, that is if an insert thread loses a competition in locking a ΔNode for expanding or rebalancing and had to dump its carried value inside a buffer (Figure 49, line 80). Any value inserted to the buffer is guaranteed to be found. This is because after a leaf *lastnode* has been located, the search is going to evaluate whether the *lastnode.value* is equal to v . Failed comparison will cause the search to look further inside a buffer ($T_x.\text{rootbuffer}$) located in a ΔNode where the

leaf resides (Figure 48, line 15). By assuring that the switching of a root Δ Node with its mirror includes switching $T_x.rootbuffer$ with $T_x.mirrorbuffer$, we can show that any new values that might be placed inside a buffer are guaranteed to be found immediately after the completion of their respective insert procedures. The "buffered" insert has the linearisation point to SEARCHNODE at line 80.

Similarly, deleting a value from a buffer is as trivial as exchanging that value inside a buffer with an **empty** value. The search operation will failed to find that value when doing searching inside a buffer of Δ Node. This type of delete has the linearisation point to SEARCHNODE at the same line it's emptying a value inside the buffer (line 27). \square

Lemma 5.4. *In the absence of maintenance operations, the linearisation point of Δ Tree's INSERTNODE to DELETENODE is in line 46 and 57 of Figure 49. Linearisation points of DELETENODE operations to INSERTNODE are in line 15 of Figure 49.*

Proof. (Sketch) In a case of concurrent insert operations (Figure 49) at the same leaf node x , insert threads that need to "grow" the node (for illustration, cf. Figure 50) are going to do $CAS(x.left, \mathbf{empty}, \dots)$ (line 46 and 57) as their first step. This CAS is the only thing needed since the whole Δ Node structure is pre-allocated and the CAS is an atomic operation. Therefore, only one thread will succeed in changing $x.left$ and proceed populating the $x.right$ node. Other threads will fail the CAS operation and they are going to try restart the insert procedure all over again, starting from the node x .

To assure that the marking delete (line 15) behaves nicely with the "grow" insert operations, $DELETENODE(v, U)$ that has found the leaf node x with a value equal to v , will need to check again whether the node is still a leaf (line 19) after completing $CAS(x.mark, FALSE, TRUE)$. The thread needs to restart the delete process from x if it has found that x is no longer a leaf node. \square

Mirroring Whenever a Δ Node is undergoing a maintenance operation (balancing, expanding, or merging), a mirroring operation also takes place. Mirroring works by maintaining the original nodes but write the results into the mirror nodes. After this is done, the pointer will be switched and now the mirror nodes become the original, vice-versa. The Δ Node's lock will be released, and all the waiting update threads can continue with their respective operation. The original nodes and helping buffer served as the latest snapshot, which enables wait-free search on that Δ Node.

Despite INSERTNODE and DELETENODE are non-blocking, they'll still need to wait at a tip of a Δ Node whenever a maintenance operation is currently operating within that Δ Node. We employ TAS (Test and Set) on Δ Node's *locked* field (cf. Figure 49, line 81) before any maintenance operation starts. Advanced locking techniques [35, 47, 55] can also be used. This is to ensure that no basic update operations will interfere with the maintenance operation. This is also necessary to prevent the buffer overflow.

Performance concerns Note that the previous description has shown that REBALANCE and MERGE execution are actually sequential within a Δ Node. Rebalancing and merging

only involve a maximum of two Δ Node with size UB . Their operation consist of traversing and re-inserting all members of one or two Δ Nodes. Because $UB \ll N$, REBALANCE and MERGE operations are not affecting much on the Δ Tree performance.

5.4 Balanced Δ Tree

Δ Tree implementation served as an initial proof of concept of a dynamic VEB-based search trees. This tree however has major weaknesses that would affect its performance. The first is the fact that the *left* and *right* pointers occupies too much space. In a Δ Node with 127 nodes, the sets of pointers will occupy 2032 bytes of memory in a 64-bit operating system, twice of the *key* nodes (in integer) that only require 1016 bytes. A cache oblivious data structures will gain most of its benefits if only more data could occupy a small amount of memory space. Thus, a single cache-line transfer will relocate more data between any levels of memory. Secondly, inserting a consecutive increasing or decreasing numbers into Δ Tree will results a linked-list of Δ Node. It will be hard to guarantee the optimal search performance of Lemma 5.1 in this particular case.

```

1: Struct Map:
2:   member fields:
3:     left  $\in \mathbb{N}$ , interval of the left child pointer address
4:     right  $\in \mathbb{N}$ , interval of the right child pointer address

5: function RIGHT(p, base)
6:   nodesize  $\leftarrow$  SIZEOF(single node)
7:   idx  $\leftarrow$  (p - base)/nodesize
8:   if (map[idx].right  $\neq$  0) then
9:     return base + map[idx].right
10:  else
11:    return 0

12: function LEFT(p, base)
13:   nodesize  $\leftarrow$  SIZEOF(single node)
14:   idx  $\leftarrow$  (p - base)/nodesize
15:   if (map[idx].left  $\neq$  0) then
16:     return base + map[idx].left
17:   else
18:     return 0

```

Figure 51: *Mapping* functions.

5.4.1 *Map* instead of pointers

We have developed an improved Δ Tree, namely the balanced DT by completely eliminating (*left* and *right*) pointers within a Δ Node. We replaced them with `LEFT` and `RIGHT` functions instead (Figure 51, lines 12 & 5). These two functions, given an arbitrary node and its container Δ Node *root* memory address, will return the left and right child node address of that arbitrary node, respectively. A Δ Node is now slim-lined into just an array of keys. Each Δ Node is also coupled with a metadata that contains an array of pointers for the inter- Δ Node connection, and a structure that holds lock and counters.

With this mapping, we need only a single Δ Tree's pointer-based Δ Node to be created in the initialisation phase. This Δ Node is used to populate the *map*, by calculating the memory address differences between a node and its left and right children, respectively. This *map* is then used for every Balanced DT Δ Node's `LEFT` and `RIGHT` operations. The memory for the pointer-based Δ Node can be freed after a *map* is created. And since we are re-using one *map* array of size UB for traversing, memory footprint for the Balanced DT's Δ Node operations can be kept minimum. We ended up having 200% more node counts in a Δ Node given the same UB , compared to the Δ Tree.

The inter- Δ Node connection works by using the tree encoding. Here we gave colour to each nodes using either 0 or 1 with a condition that adjacent nodes at the same level have different colours. With this the path traversed from the *root* of a Δ Node to reach any internal node will produce a bit-sequence of colours. This bit representation will be translated into an array index that contains a pointer of another Δ Node. We are using leaf-oriented tree and allocate a pointer array with the length equal to the number of nodes in that Δ Node. Figure 52 illustrate how the inter- Δ Node connection works in a pointer-less search function.

5.4.2 Concurrent and balanced tree

To solve the the worst case of consecutive numbers insertion, we adopt the structure and the algorithm of B-Link trees [54] coined by Lehman and Yao. This tree is a highly concurrent variation of B-Tree which sometimes referred as Blink tree. We maintain the concept of a dynamic-vEB Δ Node and used these in place of the array nodes of B-Link trees.

To implement this, two new variables were added into the Δ Nodes' metadata, namely *nextRight* or pointer to the right sibling Δ Node and a *highKey* value that contains the upper-bound value of that specific Δ Node. The insertion were done bottom-up and searches were in top-down, left-to-right direction. With these additional variables and restrictions, `SEARCH` operations are guaranteed as *wait-free* [54]. Bottom-up insertion also ensures that the tree is always in a balanced condition as mandated by Lemma 5.1. The same rebalancing procedures (Figure 50a) were also employed to ensure a Δ Node is full before it splits. The rebalancing also help to clean-up the nodes marked for deletion, keeping Δ Nodes always in good shape.

```

1: function POINTERLESSSEARCH(key,  $\Delta$ Node, maxDepth)
2:   while  $\Delta$ Node is not leaf do
3:     bits  $\leftarrow$  0
4:     depth  $\leftarrow$  0
5:     p  $\leftarrow$   $\Delta$ Node.root;
6:     base  $\leftarrow$  p
7:     link  $\leftarrow$   $\Delta$ Node.link
8:     while p & p.value  $\neq$  EMPTY do ▷ continue until leaf node
9:       depth  $\leftarrow$  depth + 1 ▷ increment depth
10:      bits  $\leftarrow$  bits << 1 ▷ either left or right, shift one bit to the left
11:      if key < p.value then
12:        p  $\leftarrow$  LEFT(p, base)
13:      else
14:        p  $\leftarrow$  RIGHT(p, base)
15:        bits  $\leftarrow$  bits + 1 ▷ right child colour is 1
16:      bits  $\leftarrow$  bits >> 1
▷ pad the bits to get the index of the child  $\Delta$ Node:
17:      bits  $\leftarrow$  bits << (maxDepth - depth)
▷ follow nextRight if highKey is less than searched value:
18:      if  $\Delta$ Node.highKey <= key then
19:         $\Delta$ Node  $\leftarrow$   $\Delta$ Node.nextRight
20:      else
21:         $\Delta$ Node  $\leftarrow$  link[bits] ▷ jump to child  $\Delta$ Node
22:      return  $\Delta$ Node

```

Figure 52: Search within pointer-less Δ Node. This function will return the *leaf* Δ Node containing the searched key. From there, a simple binary search using LEFT and RIGHT functions is adequate to pinpoint the key location

5.5 Heterogenous balanced Δ Tree

The reason why we maintained the leaf-oriented (or external tree) layout for Δ Node is to make sure the inter- Δ Node mechanism works. Thus, it is not necessary for leaf Δ Nodes or the last level Δ Nodes to have leaf-oriented layout since they don't have any child Δ Nodes.

Based on this observation, we implement a special layout for the leaf Δ Nodes, making the balanced Δ Tree is having heterogenous Δ Nodes. This special layout is using internal tree for the key nodes, therefore 100% more key nodes can fit into leaf Δ Nodes compare to the non-leaf Δ Nodes given the same UB limit. To save space even more, we omit the array of pointers for intra- Δ Node connection in the leaf Δ Nodes' metadata.

Stepping up to this improved version of Δ Tree, we found that the efficiency of searches were greatly improved. Compared to original pointer-based Δ Tree and balanced DT, this heterogenous BDT delivered lower cache misses and more efficient branching.

5.6 Performance evaluation

To evaluate our conceptual ideas of dynamic-vEB implemented in the Δ Trees (section 5.3), balanced Δ Tree (plDTv1) (section 5.4), and heterogenous BDT (plDTv2) (section 5.5), we compare their performance with other prominent concurrent trees. The benchmark include the non-blocking binary search tree (NBBST) [25], concurrent AVL tree (AVLtree) [12], concurrent red-black tree (RBtree) [24], and speculation friendly tree (SFtree) [18] from the Synchrobench benchmark [32]. We also develop a concurrent version of the static vEB binary search tree in [10] using software transactional memory (STM). We utilise the GNU C Compiler's STM implementation from the version 4.9.1 for this tree and named it VTMtree. An optimised Lehman and Yao concurrent B-tree [54] implementation (CBTree) is also included in the benchmark. The tree (sometimes known as B-link tree) is a highly-concurrent B-tree implementation and it is being used as the back-end in popular database systems such as PostgreSQL⁴. We use Pthread for concurrent threads and pin the threads to distinct available physical cores. We use GCC 4.9.1 with -O2 for all compilations.

The base of the conducted experiment consists of running a series of ($rep = 5,000,000$) operations. Assuming we have nr as the number of threads, the maximum time for any of the threads to finish a sequence of $\frac{rep}{nr}$ operations is recorded. We also define an update rate u that translates to $upd = (u\% \times rep)$ number of insert and delete operations and $src = (rep - upd)$ number of search operations. We conduct experiments based on the combinations of update rate $u = \{0, 20, 50\}$ and the number of thread $nr = \{1, 2, \dots, 16\}$. Update rate of 0 means that only searching operations are conducted (100% search), while update rate 50 indicates that 50% insert and delete operations are being done out of rep operations. For each of the combination above, we pre-fill the tree with $(2^{22} - 1)$ (or 4,194,303) random values before starting the benchmarks.

The initial size ($init$) of $(2^{22} - 1)$ was chosen to simulate initial trees that partially fit into the last level cache (LLC). All involved operations, namely search, insert, and delete invoked during the tests, use random values $v \in (0, init \times 2], v \in \mathbb{N}$, as their parameter. Note that

⁴<https://github.com/postgres/postgres/blob/master/src/backend/access/nbtree/README>

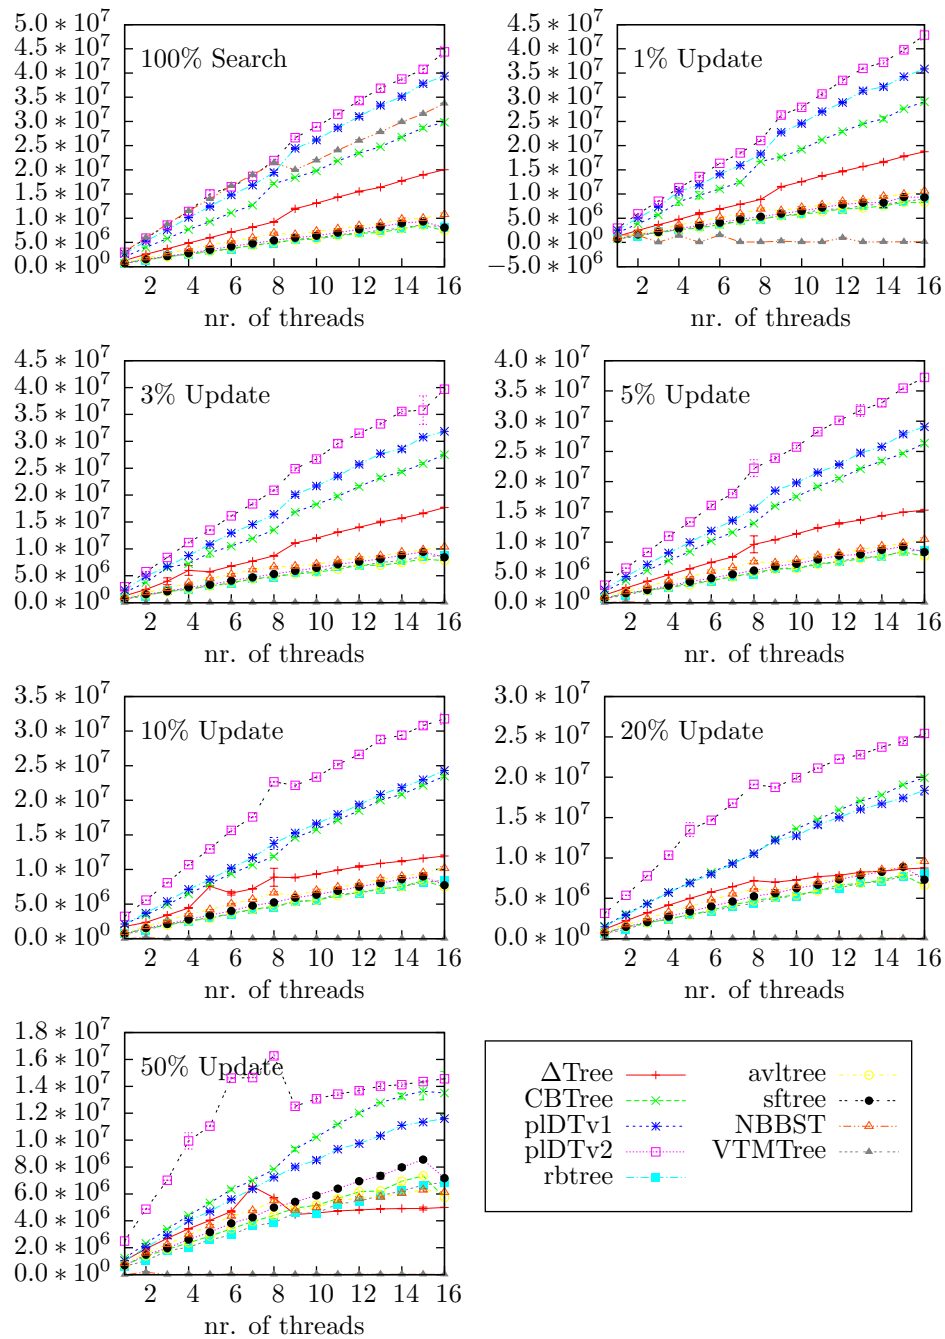


Figure 53: Performance rate (operations/second) of the tested trees with 2,500,000 initial members. The y-axis indicates the rate of operations/second.

since VTMtree’s static vEB layout is fixed, we set its layout size to $(2^{23} - 1)$ for running the experiments. Namely, this setting is the best case for VTMtree since the memory allocated for its static vEB layout is large enough to accommodate all the values $v \in (0, \text{init} \times 2]$ and therefore its layout never needs to expand and rebuild during the experiments. To make a fair comparison, we set the UB values of the Δ Nodes and the CBTree’s node-size to respective values so that each Δ Node and each CBTree node will fit into the system page size of 4KB.

The conducted experiments run on a dual Intel Xeon CPU E5-2670 machine, with total of 16 available cores. The machine has 32GB of memory, with a 2MB (8×256 KB) L2 cache and a shared 20MB L3 cache for each processor. The Hyperthread feature of the processor is turned off to make sure that the experiments only run on physical cores. Linux OS with Red Hat’s kernel version 2.6.32-358 are installed in this system. All performance result (in operations/second) for concurrent operations are calculated by dividing the number of iterations by the maximum time to finish the whole operations.

5.6.1 Experimental results

Before comparing the trees performance, there is one interesting thing to note from the benchmark results. The x86 test system consists of 2 CPUs with 8-cores each. Therefore we could see some ”spikes” in performance for couple of trees when it goes from 8-thread to 9-thread. This is expected though. For trees exploiting data-locality in cache (such as the VTMTree’s 100% search and plDTv2’s 50% update in Figure 54), maintaining the cache-coherence between the two CPUs (when more than 8 threads are used) via shared memory reduces the benefit gained from cache-locality. Among the proposed trees, plDTv1 is up to 100% faster than Δ Tree for 100% search and is up to 2.5x and 3.5x faster in 20% and 50% updates, respectively. Using map instead pointers and keeping the tree balanced manage to lower the cache-misses of plDTv1 by 40%. The plDTv2 is up to 5% faster than plDTv1 in 100% searching. However in 20% and 50% updates, plDTv2 is up to 40% faster than the plDTv1 (cf. Figure 54). It is because heterogeneous leaf Δ Nodes that can hold twice many keys, manage to lower the search time and Δ Nodes’ re-balancing overheads. Unix *perf* utility shows that plDTv2 has up to 30% less cache-misses and 20% more efficient branching than plDTv1.

Δ Trees versus VTMtree It is expected that VTMtree is among the fastest in 100% search along with the plDTv2 (Figure 54). As the cache-oblivious tree implementation, VTMtree is able to exploit perfectly data-locality in all levels of memory. Our dynamic-vEB plDTv2 is the only contender and even beats VTMtree by up to 20% past the 8-thread mark. In $\text{init} = 4, 194, 303$, the performance gap is even higher with plDTv2 leading by up to 30%, after 9-thread. The plDTv1 and Δ Tree are both beaten by VTMtree since it can only pack less data inside a memory page. CBTree and other trees are not exploiting data-locality, which makes them slower in 100% search benchmark.

All other trees are demonstrating better performance compared to VTMtree whenever update operations are involved. The bad performance of VTMtree’s concurrent update is inevitable because of its static vEB tree layout. With this, the VTMtree needs to always

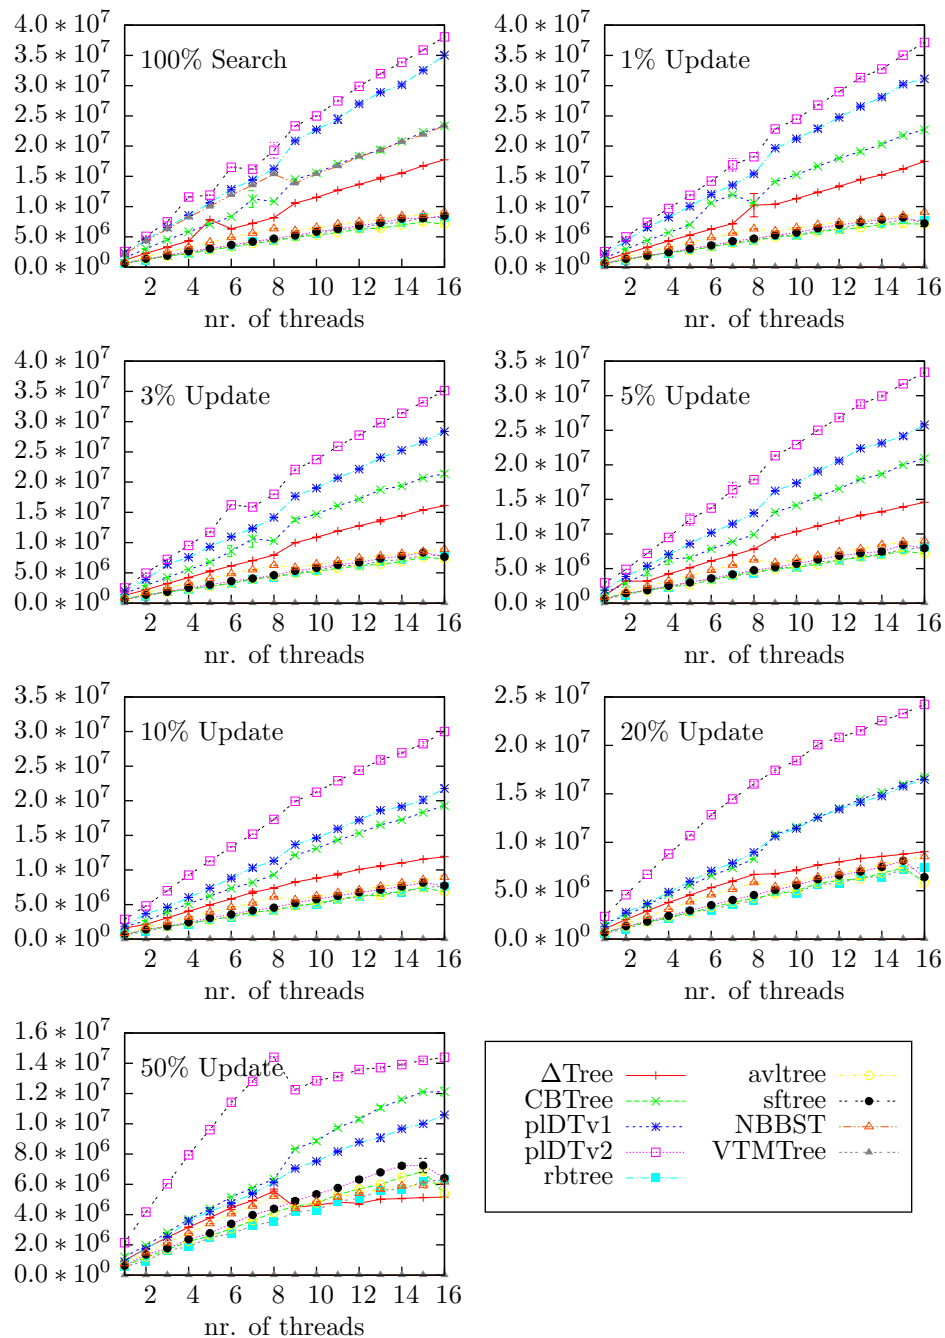


Figure 54: Performance comparison of the tested trees with 4,194,303 initial members using X86 based Intel Xeon system. There were 2 CPUs with 8 cores each and threads were pinned to the cores, therefore several "spikes" in performance could be observed for 9-thread tests. The y-axis indicates operations/second.

maintain a small height, which is done by *incrementally* re-balancing different portions of its structure [10]. However in the worst case the whole tree must be blocked whenever a rebalance operation is being executed, blocking other operations as a result. While [10] explained that amortised cost for this is small, it will hold true only when implemented in the sequential fashion. In all variants of Δ Tree, maintenance operations only block a Δ Node, which is discernible in size compared to the whole tree.

Δ Trees versus other trees In comparison with the other trees, the benchmark result in Figure 53 and 54 shows that the heterogenous BDT (plDTv2) is the fastest among other trees. In 50% update using a single CPU socket (8 threads), plDTv2 is up to 140% faster than CBTree. plDTv1 and CBTree are trailing closely behind plDTv2. CBTree manages to outperform plDTv1 in higher update ratios because plDTv1 sometimes need to do rebalancing, which is more expensive compared to array shifting in CBTree. However rebalancing doesn't affect much the plDTv2 because its leaf Δ Nodes are not leaf oriented. Thus, the amortised cost of rebalancing is much lower compared to plDTv1.

NBBST performs similarly with the Δ Tree, mainly because the latter is modelled after NBBST for achieving concurrency. However being not a locality-aware structures makes NBBST unable to lead in the search-intensive benchmarks. The balanced Δ Tree (plDTv1) performs well only in the low-contention situations, as it is able to deliver good performance only up to 20% update.

The good performance of Δ Tree, plDTv1, and plDTv2 can be attributed to the dynamic vEB layout that permits fast search. Also the fact that several Δ Node can be concurrently updated and restructured is also one of the leading factor over the static vEB layout. The CBTree layout, although fast and highly-concurrent, still suffers from high branching operations, based on Unix perf profiling. CBTree branching is up to 90% more than plDTv1's and plDTv2's. Its cache references is also 150% higher than plDTv1's and plDTv2's, as expected in B-tree since block B is not optimal for transfers between memory levels.

The software transactional memory (STM) based trees have a significant overhead in maintaining transaction. Therefore their performances are the slowest.

5.6.2 On the worst-case insertions

One of the motivations in improving the Δ Tree into balanced DT and subsequently heterogenous BDT, is to solve the poor performance of Δ Tree in the worst-case insertions. Inserting a sequence of increasing numbers to the Δ Tree will result in a linked-list of Δ Nodes.

Therefore we compare CBTree and plDTv2 for the worst-case insertions. GCC standard library `std::set` is included as the baseline. Starting from a blank tree, we insert a sequence of increasing number within $(0, 5, 000, 000]$ range using single thread.

The result in Figure 55 shows that plDTv2 is 50% faster than CBTree and `std::set` from GCC standard library. This test is done using the same x86 experimental system as in Section 5.6.1.

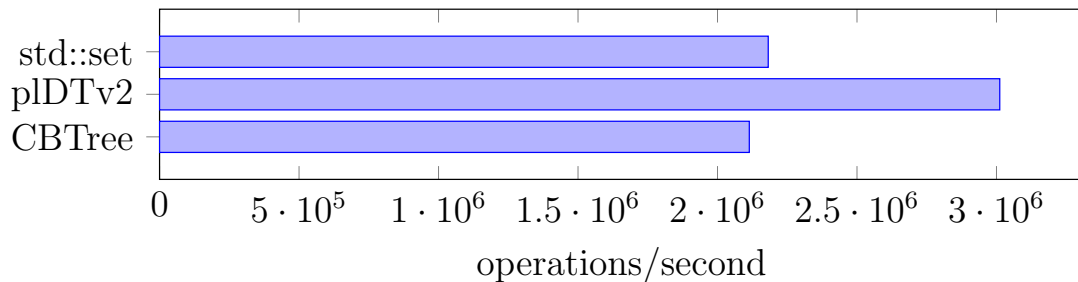
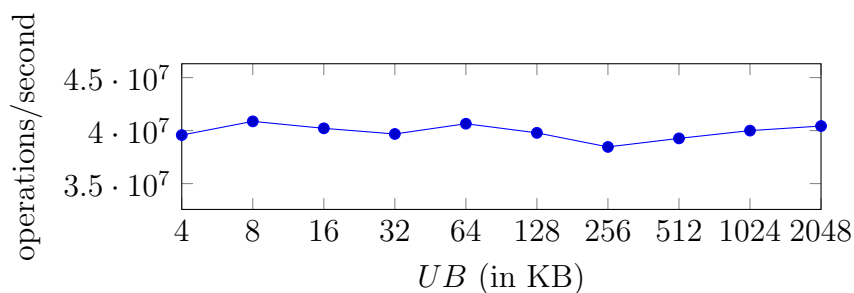


Figure 55: Worst case insertion of 5,000,000 increasing numbers

5.6.3 Performance on different upper-bounds UB

Figure 56: Benchmarks of 100% search using different sizes of UB . Tested on X86 platform.

Dynamic vEB requires that an upper-bound UB be specified or known in advance. One may argue that UB is a fine-tuned value that will determine the performance of dynamic vEB trees. To get the conclusion of whether this is the case, we test the heterogeneous BDT using different upper-bounds UB , starting from 4KB (normal page size) up to 2MB (huge page size). This tree is filled with $(2^{22} - 1)$ random values and time is recorded to conduct 5 million operations of 100% search.

The result shows that the heterogeneous BDT is resilient to different upper-bounds UB (cf. Figure 56). This is an expected result according to Lemma 5.1. In fact, UB can be as big as the whole tree and searching performance is still optimal, provided the meta-data (e.g. the tree map in Section 5.4.1) occupies only a small fraction of the last level cache LLC. As mentioned in Section 5.2.2, small UB benefits concurrent tree updates.

5.7 Energy consumption evaluation

To assess the energy consumption of the trees, the energy indicators are subsequently collected during specified benchmarks. For these tests we use a specialised server with 2x Intel Xeon E5-2690 v2 for 20 total cores. We use Intel PCM library that can measure the energy for each CPU and DRAMs using the built-in CPU counters. The collected energy measurements do not include the initialisations of trees.

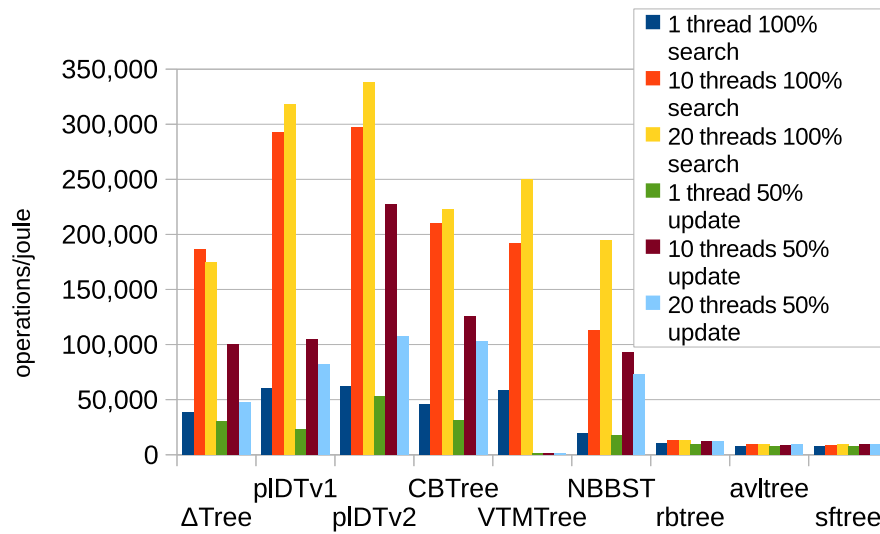


Figure 57: Energy profile on X86 processor. Note that the energy efficiency goes down for some trees in 50% update on 20 threads (with 2 CPUs) because of the same reason discussed in Section 5.6.1

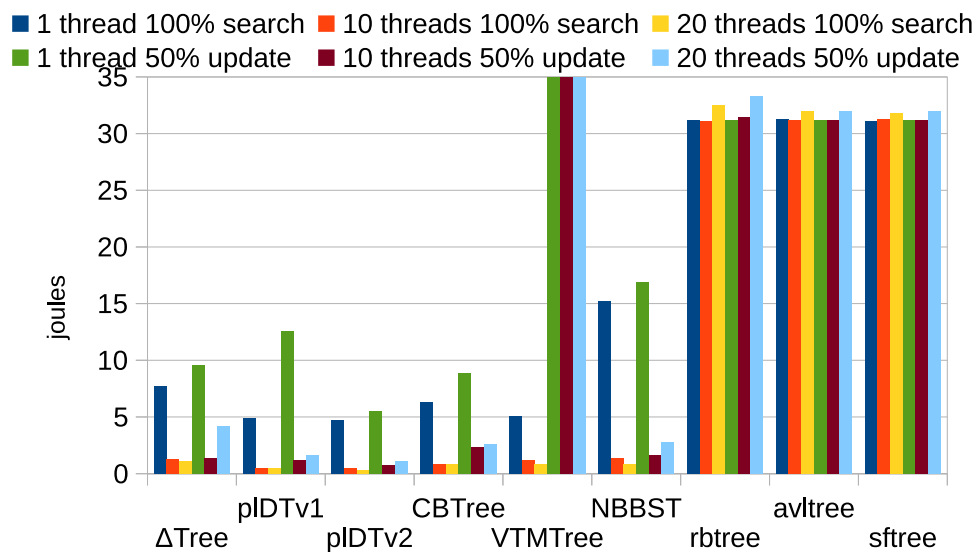


Figure 58: Memory (DRAM) energy profile on X86 platform. DRAM energy requirement goes up considerably for some trees in 50% update on 20 threads (with 2 CPUs) because of the cache-coherence mechanism (cf. Section 5.6.1). Also it is reflected in the memory write counters chart on Figure 60. Measured using Intel PCM.

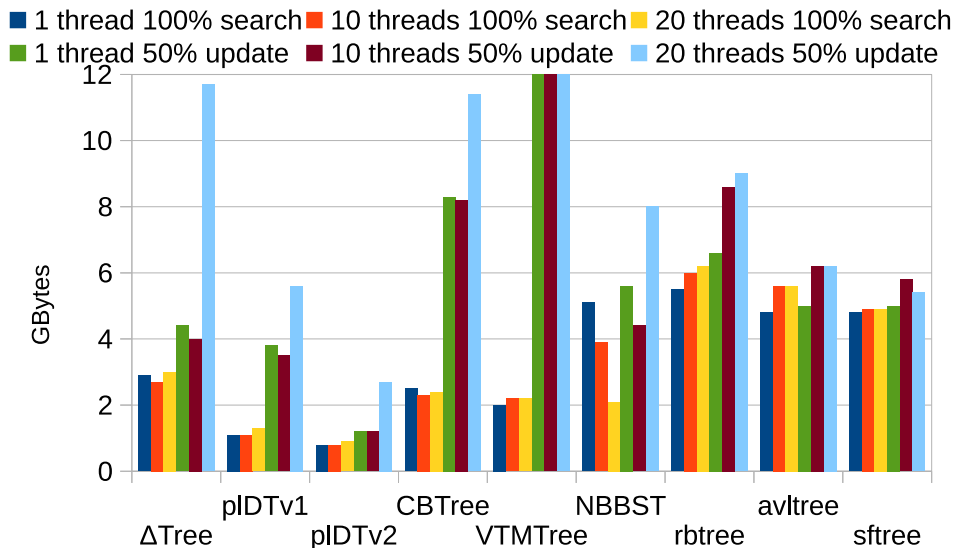


Figure 59: Amount data transferred between RAM and CPU for Read + Write operations. Measured using Intel PCM.

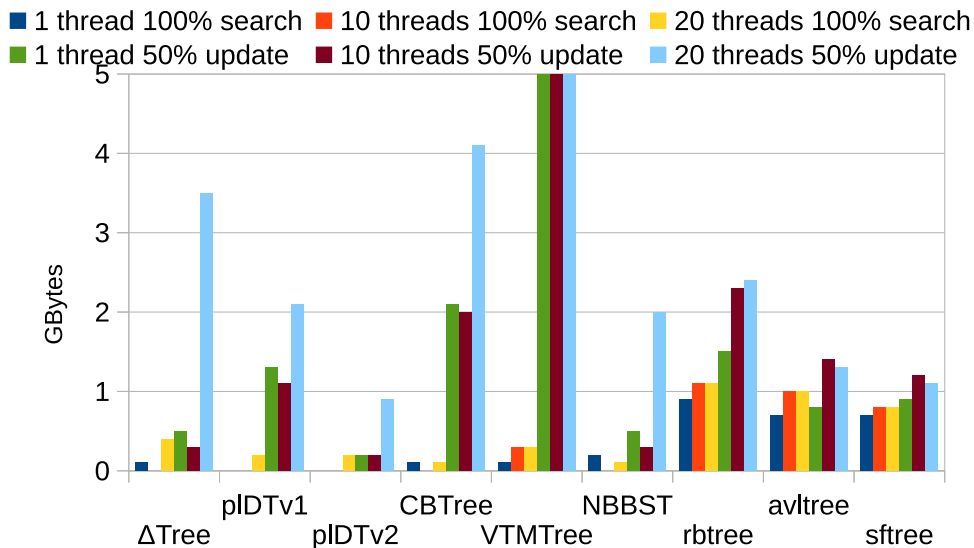


Figure 60: Amount data transferred between RAM and CPU for write-only operations as measured using Intel PCM. Data transfers goes up considerably for some trees in 50% update on 20 threads (with 2 CPUs) because of the cache-coherence mechanism.

In this experiment, we conduct 5 million operations of 100% search and 50% update on the trees. The trees are pre-filled with initial $(2^{22} - 1)$ random values. A combination of minimum and maximum available physical cores are used as one of the benchmarks parameters. The

total energy used for all CPUs and memory (in Joule) are divided by the number of operations to produce operations/joule results.

The experimental result over the test system (Figure 57) shows that using the dynamic-vEB layouts is able to reduce the CPU and memory energy consumption. In the search-only benchmark, plDTv1, plDTv2, and Δ Tree’s actual energy efficiency is comparable to that of the static vEB-based VTMtree and leads over the other trees by up to 33%. plDTv2 energy efficiency is 80% better than CBTree’s and other trees in 50% update using 10-threads. As expected, the VTMtree concurrent update results are very poor for the same reason discussed on section 5.6.1.

Δ Trees are locality-aware trees, and their memory access pattern are more efficient to the other trees. Profiling data showed that branching and cache-misses are kept low. Assessments on memory transfers (Figure 59 and 60) suggests that all Δ Tree versions are transferring less data between RAM and CPU compared to other trees. Thus the energy consumed by the Δ Trees on DRAM operation is also efficient and comparable to the static-vEB’s VTMTree in the searching only case (cf. Figure 58). The reader is referred to [75] for more details about Δ Trees.

5.8 Conclusion

We have introduced a new *relaxed cache oblivious* model that enables high parallelism while maintaining the key feature of the original cache oblivious (CO) model [26]: an algorithm that is optimal in terms of data movement for a simple two-level memory is asymptotically optimal for an unknown multilevel memory. This feature facilitates the development of fine-grained locality-aware algorithms for deep memory hierarchy in modern architectures as desired by energy efficient computing [19]. Unlike the original CO model, the relaxed CO model assumes a known upper bound on unknown memory block sizes B of multilevel memory systems.

Based on the relaxed CO model, we have developed a novel *dynamic* van Emde Boas (dynamic vEB) layout that makes the vEB layout suitable for highly-concurrent data structures with update operations. The dynamic vEB supports dynamic node allocation via pointers while maintaining the optimal search cost of $O(\log_B N)$ memory transfers for vEB-based trees of size N without knowing memory block size B .

Using the dynamic van Emde Boas layout, we have developed a pointer-based Δ Tree, a balanced Δ Tree, and a heterogenous version of the latter that support both high concurrency and fine-grained data locality. Both pointer-based Δ Tree as well as balanced DT and heterogenous BDT *Search* operations are wait-free. Only Δ Tree has non-blocking *Insert* and *Delete* operations. All 3 versions of Δ Tree are relaxed cache oblivious: the expected memory transfer costs of its *Search*, *Delete* and *Insert* operations are $O(\log_B N)$, where N is the tree size and B is unknown memory block size in the ideal cache model [26]. Our experimental evaluation comparing the Δ Trees with non-blocking binary search tree of [25], concurrent AVL tree [12], concurrent red-black tree [24], and speculation-friendly trees [18] from the the Synchrobench benchmark [32], and the highly-concurrent B-tree of [54] has shown that the best version of Δ Tree achieves the best performance and the highest energy efficiency.

6 Conclusion

In this work, we have presented our current results on the investigation and modeling of the trade-off between energy and performance as follows.

- A new power model for the Movidius Myriad platform has been proposed. The model can predict power consumption of our micro-benchmarks with $\pm 4\%$ accuracy compared to the measured values on the real platform. The new power model confirms that the dynamic power consumption is proportional to the number of SHAVE (Streaming Hybrid Architecture Vector Engine) processors used.
- Inspired by EXCESS D1.1 [49], a new version of the energy model for the CPU-based platform has been developed. This new model decomposes the power into static, active and dynamic power, and considers power as the sum of power from CPU, main memory and uncore. The model parameters have been derived and evaluated; dynamic powers roughly depend on nature of operation, amount of memory accessed per unit of time and remote accesses, respectively for CPU, memory and uncore.
- We have made a first step towards realistic application by modeling the performance and power dissipation of synthetic applications, whose design is based on concurrent queues. We have exhibited a small set of parameters that rules both performance and power consumption of the application and discriminates the different queue implementations. It leads to a good prediction of those two key metrics on the whole space of study, while needing only a few measurements from this space.
- Several concurrent queue designs have been transferred to Myriad1 platform. Mutex with two locks implementation is the fastest and most scalable since it provides maximum concurrency until 8 SHAVEs. With 4 SHAVEs, FIFO-based implementation performs well and is 28.3% faster than the mutex with two locks. Shared variable based implementation has the worst performance. In terms of power, SHAVE FIFO-based communication method is the most energy efficient. When contention is high and stall is common, it is power efficient to avoid spinning and set SHAVEs to stall. Communication via shared variables consumes more power because spinning on a memory location is energy inefficient, even when the spinning happens in a local CMX slice.
- Another data structure that has been investigated and analyzed in this work is concurrent search tree. Our new concurrent search trees show the improvement on performance and energy consumption. Based on experimental evaluations, our new concurrent search trees called Δ Trees that are up to 140% faster and 80% more energy efficient than the traditional search trees.

In the next steps of this work, WP2 aims to identify other essential concurrent data structures and algorithms for inter-process communication in HPC and embedded computing and focus on customizing them. We will exploit common data-flow patterns to create a generalized communication abstraction with which application designers can easily create

and exploit the customization for the data-flow patterns. The results will also constitute a white-box methodology for tuning energy efficiency and performance of concurrent data structures and algorithms, and programming abstractions with which application designers can easily create and exploit the customization for common data-flow patterns. The novel concurrent data structures and algorithms will constitute libraries for inter-process communication and data sharing on EXCESS platforms.

References

- [1] Adlink Technology Inc. *USB-1900 Series 16-bit 250kS/s USB 2.0-based High-performance DAQ Module USB-1901/1902/1903 User's Manual*, 2011.
- [2] Y. Afek, G. Korland, M. Natanzon, and N. Shavit. Scalable Producer-Consumer Pools Based on Elimination-Diffraction Trees. In *Euro-Par 2010*, volume 6272 of *Lecture Notes in Computer Science*, pages 151–162. Springer, 2010.
- [3] Yehuda Afek, Haim Kaplan, Boris Korenfeld, Adam Morrison, and Robert E. Tarjan. Cbtree: a practical concurrent self-adjusting search tree. In *Proceedings of the 26th international conference on Distributed Computing, DISC'12*, pages 1–15, Berlin, Heidelberg, 2012. Springer-Verlag.
- [4] Alok Aggarwal and S. Vitter, Jeffrey. The input/output complexity of sorting and related problems. *Commun. ACM*, 31(9):1116–1127, 1988.
- [5] R. Bayer and E.M. McCreight. Organization and maintenance of large ordered indexes. *Acta Informatica*, 1(3):173–189, 1972.
- [6] Michael Bender, Erik D. Demaine, and Martin Farach-Colton. Cache-oblivious b-trees. *SIAM Journal on Computing*, 35:341, 2005.
- [7] Michael A. Bender, Martin Farach-Colton, Jeremy T. Fineman, Yonatan R. Fogel, Bradley C. Kuszmaul, and Jelani Nelson. Cache-oblivious streaming b-trees. In *Proceedings of the 19th annual ACM symposium on Parallel algorithms and architectures, SPAA '07*, pages 81–92, 2007.
- [8] Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Bradley C. Kuszmaul. Concurrent cache-oblivious b-trees. In *Proceedings of the 17th annual ACM symposium on Parallelism in algorithms and architectures, SPAA '05*, pages 228–237, 2005.
- [9] Anastasia Braginsky and Erez Petrank. A lock-free b+tree. In *Proceedings of the 24th ACM symposium on Parallelism in algorithms and architectures, SPAA '12*, pages 58–67, 2012.
- [10] Gerth Stølting Brodal, Rolf Fagerberg, and Riko Jacob. Cache oblivious search trees via binary trees of small height. In *Proceedings of the 13th annual ACM-SIAM symposium on Discrete algorithms, SODA '02*, pages 39–48, 2002.
- [11] GerthStølting Brodal. Cache-oblivious algorithms and data structures. In Torben Hagerup and Jyrki Katajainen, editors, *Algorithm Theory - SWAT 2004*, volume 3111 of *Lecture Notes in Computer Science*, pages 3–13. Springer Berlin Heidelberg, 2004.
- [12] Nathan G. Bronson, Jared Casper, Hassan Chafi, and Kunle Olukotun. A practical concurrent binary search tree. In *Proceedings of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP '10*, pages 257–268, 2010.

- [13] Trevor Brown and Joanna Helga. Non-blocking k-ary search trees. In *Proceedings of the 15th international conference on Principles of Distributed Systems*, OPODIS'11, pages 207–221, Berlin, Heidelberg, 2011. Springer-Verlag.
- [14] S. Browne, J. Dongarra, N. Garner, G. Ho, and P. Mucci. A portable programming interface for performance evaluation on modern processors. *Int. J. High Perform. Comput. Appl.*, 14(3):189–204, August 2000.
- [15] Daniel Cederman, Anders Gidenstam, Phuong Ha, Hakan Sundell, Marina Papatriantafylou, and Philippas Tsigas. Lock-free concurrent data structures. In Sabri Pllana and Fatos Xhafa, editors, *Programming Multi-Core and Many-Core Computing Systems*. Wiley, 2013. To appear.
- [16] Douglas Comer. Ubiquitous b-tree. *ACM Comput. Surv.*, 11(2):121–137, 1979.
- [17] Thomas H. Cormen, Clifford Stein, Ronald L. Rivest, and Charles E. Leiserson. *Introduction to Algorithms*. McGraw-Hill Higher Education, 2nd edition, 2001.
- [18] Tyler Crain, Vincent Gramoli, and Michel Raynal. A speculation-friendly binary search tree. In *Proceedings of the 17th ACM SIGPLAN symposium on Principles and Practice of Parallel Programming*, PPOPP '12, pages 161–170, New York, NY, USA, 2012. ACM.
- [19] Bill Dally. Power and programmability: The challenges of exascale computing. In *DoE Arch-I presentation*, 2011.
- [20] Howard David, Eugene Gorbatov, Ulf R. Hanebutte, Rahul Khanaa, and Christian Le. Rapl: memory power estimation and capping. In *ISLPED*, pages 189–194, 2010.
- [21] Damian Dechev, Peter Pirkelbauer, and Bjarne Stroustrup. Lock-free dynamically resizable arrays. In *Proceedings of the 10th International Conference on Principles of Distributed Systems (OPODIS '06)*, Lecture Notes in Computer Science, pages 142–156. Springer Verlag, 2006.
- [22] Giovanni Della-Libera and Nir Shavit. Reactive diffracting trees. *Journal of Parallel and Distributed Computing*, 60(7):853 – 890, 2000.
- [23] Electronic Educational Devices. *Watts up Operators manual*, 2003.
- [24] Dave Dice, Ori Shalev, and Nir Shavit. Transactional locking ii. In *Proceedings of the 20th international conference on Distributed Computing*, DISC'06, pages 194–208, 2006.
- [25] Faith Ellen, Panagiota Fatourou, Eric Ruppert, and Franck van Breugel. Non-blocking binary search trees. In *Proceedings of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, PODC '10, pages 131–140, 2010.

- [26] Matteo Frigo, Charles E. Leiserson, Harald Prokop, and Sridhar Ramachandran. Cache-oblivious algorithms. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science, FOCS '99*, page 285, Washington, DC, USA, 1999. IEEE Computer Society.
- [27] John Giacomoni, Tipp Moseley, and Manish Vachharajani. Fastforward for efficient pipeline parallelism: a cache-optimized concurrent lock-free queue. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming (PPoPP '08)*, pages 43–52, New York, NY, USA, 2008. ACM.
- [28] Anders Gidenstam, Marina Papatriantafidou, Håkan Sundell, and Philippas Tsigas. Efficient and reliable lock-free memory reclamation based on reference counting. *IEEE Transactions on Parallel and Distributed Systems*, 20(8):1173–1187, 2009.
- [29] Anders Gidenstam, Håkan Sundell, and Philippas Tsigas. Cache-aware lock-free queues for multiple producers/consumers and weak memory consistency. In *Proceedings of the 14th International Conference on Principle of Distributed Systems (OPODIS 2010)*, volume 6490 of *Lecture Notes in Computer Science*, pages 302–317. Springer, 2010.
- [30] Goetz Graefe. A survey of b-tree locking techniques. *ACM Trans. Database Syst.*, 35(3):16:1–16:26, July 2010.
- [31] Goetz Graefe. Modern b-tree techniques. *Found. Trends databases*, 3(4):203–402, April 2011.
- [32] Vincent Gramoli. Synchronobench: A benchmark to compare synchronization techniques for multicore programming. <https://github.com/gramoli/synchrobench>.
- [33] P. H. Ha and P. Tsigas. Reactive multiword synchronization for multiprocessors. In *2003 12th International Conference on Parallel Architectures and Compilation Techniques*, pages 184–193, 2003.
- [34] P. H. Ha, P. Tsigas, and O. J. Anshus. The synchronization power of coalesced memory accesses. *IEEE Transactions on Parallel and Distributed Systems*, 21(7):939–953, 2010.
- [35] Phuong Hoai Ha, Marina Papatriantafidou, and Philippas Tsigas. Efficient self-tuning spin-locks using competitive analysis. *Journal of Systems and Software*, 80(7):1077 – 1090, 2007.
- [36] Phuong Hoai Ha, Marina Papatriantafidou, and Philippas Tsigas. Self-tuning reactive diffracting trees. *Journal of Parallel and Distributed Computing*, 67(6):674 – 694, 2007.
- [37] Phuong Hoai Ha, P. Tsigas, and O. J. Anshus. Wait-free programming for general purpose computations on graphics processors. In *2008 IEEE International Symposium on Parallel and Distributed Processing*, pages 1–12, 2008.

- [38] Phuong Hoai Ha, Philippas Tsigas, and Otto J. Anshus. Nb-feb: A universal scalable easy-to-use synchronization primitive for manycore architectures. In *Proceedings of the 13th International Conference on Principles of Distributed Systems*, OPODIS '09, pages 189–203, 2009.
- [39] Phuong Hoai Ha, Philippas Tsigas, Mirjam Wattenhofer, and Rogert Wattenhofer. Efficient multi-word locking using randomization. In *Proceedings of the Twenty-fourth Annual ACM Symposium on Principles of Distributed Computing*, PODC '05, pages 249–257, 2005.
- [40] D. Hendler, N. Shavit, and L. Yerushalmi. A Scalable Lock-Free Stack Algorithm. *Journal of Parallel and Distributed Computing*, 70(1):1–12, 2010.
- [41] M. Herlihy, V. Luchangco, P. Martin, and M. Moir. Nonblocking Memory Management Support for Dynamic-sized Data Structures. *ACM Transactions on Computer Systems*, 23:146–196, May 2005.
- [42] Maurice Herlihy and J. Eliot B. Moss. Transactional memory: Architectural support for lock-free data structures. In *Proceedings of the 20th Annual International Symposium on Computer Architecture*, ISCA '93, pages 289–300, 1993.
- [43] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming*. Morgan Kaufmann Publishers, 2008.
- [44] Maurice Herlihy and J. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, 1990.
- [45] Moshe Hoffman, Ori Shalev, and Nir Shavit. The baskets queue. In Eduardo Tovar, Philippas Tsigas, and Hacène Fouchal, editors, *Proceedings of the 11th International Conference on Principles of Distributed Systems (OPODIS 2007)*, volume 4878 of *Lecture Notes in Computer Science*, pages 401–414. Springer, 2007.
- [46] Intel. Threading Building Blocks, 2009.
- [47] Anna R. Karlin, Kai Li, Mark S. Manasse, and Susan Owicki. Empirical studies of competitive spinning for a shared-memory multiprocessor. In *Proceedings of the Thirteenth ACM Symposium on Operating Systems Principles*, SOSP '91, pages 41–55, 1991.
- [48] Christoph Kessler, Lu Li, Usman Dastgeer, Anders Gidenstam, and Aras Atalar. D1.2 Initial specification of energy, platform and component modelling framework. Technical Report FP7-611183 D1.2, EU FP7 Project EXCESS, August 2014.
- [49] Christoph Kessler, Lu Li, Usman Dastgeer, Philippas Tsigas, Anders Gidenstam, Paul Renaud-Goud, Ivan Walulya, Aras Atalar, David Moloney, Phuong Ha Hoai, and Vi Tran. D1.1 Early validation of system-wide energy compositionality and affecting

- factors on the EXCESS platforms. Project Deliverable, EU FP7 project Execution Models for Energy-Efficient Computing Systems (EXCESS), www.excess-project.eu, April 2014.
- [50] Changkyu Kim, Jatin Chhugani, Nadathur Satish, Eric Sedlar, Anthony D. Nguyen, Tim Kaldewey, Victor W. Lee, Scott A. Brandt, and Pradeep Dubey. Fast: fast architecture sensitive tree search on modern cpus and gpus. In *Proceedings of the 2010 ACM SIGMOD Intl. Conference on Management of data*, SIGMOD '10, pages 339–350, 2010.
- [51] Leslie Lamport. Specifying concurrent program modules. *ACM Trans. Program. Lang. Syst.*, 5(2):190–222, 1983.
- [52] Andreas Larsson, Anders Gidenstam, Phuong H. Ha, Marina Papatriantafidou, and Philippas Tsigas. Multiword atomic read/write registers on multiprocessor systems. *J. Exp. Algorithmics*, 13:7:1.7–7:1.30, 2009.
- [53] Doug Lea. The Java Concurrency Package (JSR-166), 2009.
- [54] Philip L. Lehman and s. Bing Yao. Efficient locking for concurrent operations on b-trees. *ACM Trans. Database Syst.*, 6(4):650–670, December 1981.
- [55] Beng-Hong Lim and Anant Agarwal. Reactive synchronization algorithms for multiprocessors. In *Proceedings of the Sixth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS VI, pages 25–35, 1994.
- [56] Jean-Pierre Lozi, Florian David, Gaël Thomas, Julia Lawall, and Gilles Muller. Remote core locking: Migrating critical-section execution to improve the performance of multithreaded applications. In *Proceedings of the 2012 USENIX Conference on Annual Technical Conference*, USENIX ATC'12, pages 6–6, Berkeley, CA, USA, 2012. USENIX Association.
- [57] M-Herlihy. Wait-Free Synchronization. *ACM Transactions on Programming Languages and Systems*, 11(1):124–149, January 1991.
- [58] Benoit B. Mandelbrot. Fractal aspects of the iteration of $z \rightarrow \lambda z(1 - z)$ for complex λ and z . *Annals of the New York Academy of Sciences*, 357:249–259, 1980.
- [59] Maged M. Michael. Hazard pointers: Safe memory reclamation for lock-free objects. *IEEE Transactions on Parallel and Distributed Systems*, 15(8), August 2004.
- [60] Maged M. Michael and Michael L. Scott. Simple, Fast, and Practical Non-Blocking and Blocking Concurrent Queue Algorithms. In *PODC '96: Proceedings of the fifteenth annual ACM symposium on Principles of distributed computing*, pages 267–275. ACM Press, 1996.
- [61] Microsoft. Parallel Computing Developer Center, 2009.

- [62] Mark Moir, Daniel Nussbaum, Ori Shalev, and Nir Shavit. Using elimination to implement scalable and lock-free fifo queues. In *Proceedings of the 17th annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 253–262, 2005.
- [63] Darko Petrović, Thomas Ropars, and André Schiper. Leveraging hardware message passing for efficient thread synchronization. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP '14, pages 143–154, New York, NY, USA, 2014. ACM.
- [64] Harald Prokop. Cache-oblivious algorithms. Master's thesis, MIT, 1999.
- [65] Jason Sewall, Jatin Chhugani, Changkyu Kim, Nadathur Rajagopalan Satish, and Pradeep Dubey. Palm: Parallel architecture-friendly latch-free modifications to b+ trees on many-core processors. *Proceedings of the VLDB Endowment*, 4(11):795–806, 2011. VLDB 2011.
- [66] Nir Shavit and Asaph Zemach. Diffracting trees. *ACM Trans. Comput. Syst.*, 14(4):385–428, 1996.
- [67] Daniel D. Sleator and Robert E. Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, February 1985.
- [68] Håkan Sundell, Anders Gidenstam, Marina Papatrifiantafilou, and Philippas Tsigas. A lock-free algorithm for concurrent bags. In *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. ACM, 2011.
- [69] Håkan Sundell and Philippas Tsigas. NOBLE: A Non-Blocking Inter-Process Communication Library. *Lecture Notes in Computer Science*, 2002.
- [70] Håkan Sundell and Philippas Tsigas. NOBLE: Non-blocking programming support via lock-free shared abstract data types. *SIGARCH Computer Architecture News*, 36(5), 2008.
- [71] Salam Traboulsi (ed.) and Yosandra Sandoval (ed.). D5.1 Report on specification of evaluation criteria. Technical Report FP7-611183 D5.1, EU FP7 Project EXCESS, February 2014.
- [72] Philippas Tsigas and Yi Zhang. A simple, fast and scalable non-blocking concurrent FIFO queue for shared memory multiprocessor systems. In *Proceedings of the 13th annual ACM Symposium on Parallel Algorithms and Architectures*, pages 134–143, 2001.
- [73] Philippas Tsigas and Yi Zhang. Integrating non-blocking synchronisation in parallel applications: performance advantages and methodologies. In *Proceedings of the third international workshop on Software and performance*, pages 55–67. ACM Press, 2002.
- [74] J. Turek and D. Shasha. The many faces of consensus in distributed systems. *IEEE Computer*, 25(2):8–17, 1992.

- [75] Ibrahim Umar, Otto Anshus, and Phuong Ha. Deltatree: A practical locality-aware concurrent search tree. techreport 2013-74, University of Tromso, 2013. arXiv:1312.2628.
- [76] J. D. Valois. Implementing Lock-Free Queues. In *Proceedings of the 7th International Conference on Parallel and Distributed Computing Systems*, pages 64–69, 1994.
- [77] J. D. Valois. *Lock-Free Data Structures*. PhD thesis, Rensselaer Polytechnic Institute, Troy, New York, 1995.
- [78] P. van Emde Boas. Preserving order in a forest in less than logarithmic time. In *Proceedings of the 16th Annual Symposium on Foundations of Computer Science, SFCS '75*, pages 75–84, Washington, DC, USA, 1975. IEEE Computer Society.
- [79] Vincent M. Weaver, Matt Johnson, Kiran Kasichayanula, James Ralph, Piotr Luszczek, Dan Terpstra, and Shirley Moore. Measuring energy and power with papi. In *Proceedings of the 2012 41st International Conference on Parallel Processing Workshops, ICPPW '12*, pages 262–268, Washington, DC, USA, 2012. IEEE Computer Society.

Glossary

BRU	Branch Repeat Unit (on SHAVE processor)
CAS	Compare-and-Swap instruction
CMX	Connection MatriX on-chip (shared) memory unit, 128KB (Movidius Myriad)
CMU	Compare-Move Unit (on SHAVE processor)
Component	1. [hardware component] part of a chip's or motherboard's circuitry; 2. [software component] encapsulated and annotated reusable software entity with contractually specified interface and explicit context dependences only, subject to third-party (software) composition.
Composition	1. [software composition] Binding a call to a specific callee (e.g., implementation variant of a component) and allocating resources for its execution; 2. [task composition] Defining a macrotask and its use of execution resources by internally scheduling its constituent tasks in serial, in parallel or a combination thereof.
CPU	Central (general-purpose) Processing Unit
uncore	including the ring interconnect, shared cache, integrated memory controller, home agent, power control unit, integrated I/O module, config Agent, caching agent and Intel QPI link interface
CTH	Chalmers University of Technology
DAQ	Data Acquisition Unit
DCU	Debug Control Unit (on SHAVE processor)
DDR	Double Data Rate Random Access Memory
DMA	Direct (remote) Memory Access
DRAM	Dynamic Random Access Memory
DSP	Digital Signal Processor
DVFS	Dynamic Voltage and Frequency Scaling
ECC	Error-Correcting Coding
EXCESS	Execution Models for Energy-Efficient Computing Systems
GPU	Graphics Processing Unit
HPC	High Performance Computing
IAU	Integer Arithmetic Unit (on SHAVE processor)
IDC	Instruction Decoding Unit (on SHAVE processor)
IRF	Integer Register File (on SHAVE processor)
LEON	SPARCV8 RISC processor in the Myriad1 chip
LIU	Linköping University
LLC	Last-level cache
LSU	Load-Store Unit (on SHAVE processor)
Microbenchmark	Simple loop or kernel developed to measure one or few properties of the underlying architecture or system software
PAPI	Performance Application Programming Interface

PEPPHER	Performance Portability and Programmability for Heterogeneous Many-core Architectures. FP7 ICT project, 2010-2012, www.peppher.eu
PEU	Predicated Execution Unit (on SHAVE processor)
Pinning	[thread pinning] Restricting the operating system's CPU scheduler in order to map a thread to a fixed CPU core
QPI	Quick Path Interconnect
RAPL	Running Average Power Limit energy consumption counters (Intel)
RCL	Remote Core Locking (synchronization algorithm)
SAU	Scalar Arithmetic Unit (on SHAVE processor)
SHAVE	Streaming Hybrid Architecture Vector Engine (Movidius)
SoC	System on Chip
SRF	Scalar Register File (on SHAVE processor)
SRAM	Static Random Access Memory
TAS	Test-and-Set instruction
TMU	Texture Management Unit (on SHAVE processor)
USB	Universal Serial Bus
VAU	Vector Arithmetic Unit (on SHAVE processor)
Vdram	DRAM Supply Voltage
V_{in}	Input voltage level
V_{io}	Input/Output voltage level
VLIW	Very Long Instruction Word (processor)
VLLIW	Variable Length VLIW (processor)
VRF	Vector Register File (on SHAVE processor)
Wattsup	Watts Up .NET power meter
WP1	Work Package 1 (here: of EXCESS)
WP2	Work Package 2 (here: of EXCESS)