Bachelors Thesis

Data Movement in Heterogeneous Memories with Intel Data Streaming Accelerator

Anatol Constantin Fürst

23rd January 2024

Technische Universität Dresden
Faculty of Computer Science
Institute of Systems Architecture
Chair of Operating Systems

Academic Supervisors:
Prof. Dr.-Ing. Horst Schirmeier
Prof. Dr.-Ing. habil. Dirk Habich
M.Sc. André Berthold

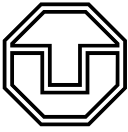

TECHNISCHE UNIVERSITÄT DRESDEN

Fakultät Informatik Institut für Systemarchitektur, Professur für Betriebssysteme

Aufgabenstellung für die Anfertigung einer Bachelor-Arbeit

Studiengang:	Bachelor
Studienrichtung:	Informatik (2009)
Name:	Constantin Fürst
Matrikelnummer:	4929314
Titel:	Data Movement in Heterogeneous Memories with Intel Data Streaming Accelerator

Developments in main memory technologies like Non-Volatile RAM (NVRAM), High Bandwidth Memory (HBM), NUMA, or Remote Memory, lead to heterogeneous memory systems that, instead of providing one monolithic main memory, deploy multiple memory devices with different non-functional memory properties. To reach optimal performance on such systems, it becomes increasingly important to move data, ahead of time, to the memory device with non-functional properties tailored for the intended workload, making data movement operations increasingly important for data intensive applications. Unfortunately, while copying, the CPU is mostly busy with waiting for the main memory, and cannot work on other computations. To tackle this problem Intel implements the Intel Data Streaming Accelerator (Intel DSA), an engine to explicitly offload data movement operations from the CPU, in their newly released Intel Xeon CPU Max processors.

The goal of this bachelor thesis is to analyze and characterize the architecture of the Intel DSA and the vendor-provided APIs. The student should benchmark the performance of Intel DSA and compare it to the CPU's performance, concentrating on data transfers between DDR5-DRAM and HBM and between different NUMA nodes. Additionally, the student should find out in what way and to what extent parallel processes copying data interfere with each other. Analyzing the performance information, the thesis should outline a gainful utilization of the Intel DSA and demonstrate its potential by extending the Query-driven Prefetching concept, which aims to speed up database query execution in heterogeneous memory systems.

Gutachter:	Prof. Dr.-Ing. Dirk Habich
Betreuer:	André Berthold, M.Sc.
Ausgehändigt am:	4. Dezember 2023
Einzureichen am:	19. Februar 2024

Prof. Dr.-Ing. Horst Schirmeier
Betreuender Hochschullehrer

Statement of Authorship

I hereby declare that I am the sole author of this master thesis and that I have not used any sources other than those listed in the bibliography and identified as references. I further declare that I have not submitted this thesis at any other institution in order to obtain a degree.

Dresden, 23rd January 2024

Anatol Constantin Fürst

Abstract

...abstract ...

write the abstract

Contents

List of Figures **XI**

1 Introduction **1**

2 Technical Background **3**
2.1 High Bandwidth Memory . . . 3
2.2 Query Driven Prefetching . . . 3
2.3 Intel Data Streaming Accelerator . . . 3
2.4 Programming Interface . . . 6
2.5 System Setup and Configuration . . . 8

3 Performance Microbenchmarks **9**
3.1 Benchmarking Methodology . . . 9
3.2 Benchmarks . . . 10
3.3 Analysis . . . 14

4 Design **17**
4.1 Detailed Task Description . . . 17
4.2 Cache Design . . . 17

5 Implementation **21**
5.1 Locking and Usage of Atomics . . . 21
5.2 Accelerator Usage . . . 25

6 Evaluation **27**

7 Conclusion And Outlook **29**
7.1 Conclusions . . . 29
7.2 Future Work . . . 29

Glossary **31**

Bibliography **33**

List of Figures

2.1 DSA Internal Architecture [5, Fig. 1 (a)]. Shows the components that the chip is made up of, how they are connected and which outside components the DSA communicates with. 4
2.2 DSA Software View [5, Fig. 1 (b)]. Illustrating the software stack and internal interactions from user applications, through the driver to the portal for work submission. 6
2.3 DML Memcpy Implementation Pseudocode. Performs copy operation of a block of memory from source to destination. The DSA executing this copy can be selected with the parameter `node`, and the template parameter `path` elects whether to run on hardware (DSA) or software (CPU). . . . 7

3.1 Benchmark Procedure Pseudocode. Timing marked with yellow background. Showing data allocation and the benchmarking loop for a single thread. 9
3.2 Throughput for different Submission Methods and Sizes. Performing a copy with source and destination being node 0, executed by the DSA on node 0. Observable is the submission cost which affects small transfer sizes differently, as there the completion time is lower. 11
3.3 Throughput for different Thread Counts and Sizes. Multiple threads submit to the same Shared Work Queue. Performing a copy with source and destination being node 0, executed by the DSA on node 0. 12
3.4 Xeon Max NUMA-Node Layout [13, Fig. 14] for a 2-Socket System when configured with HBM-Flat. Showing separate Node IDs for manual HBM access and for Cores and DDR memory. 13
3.5 Throughput for brute force copy from DDR to HBM. Using all available DSA. Copying 1 GiB from Node 0 to the Destination Node specified on the x-axis. Shows peak achievable with DSA. 14
3.6 Throughput for smart copy from DDR to HBM. Using four on-socket DSA for intra-socket operation and the DSA on source and destination node for inter-socket. Copying 1 GiB from Node 0 to the Destination Node specified on the x-axis. Shows conservative performance. 15
3.7 Throughput from DDR to HBM on CPU. Using 12 Threads spawned on Node 0 for the task. Copying 1 GiB from Node 0 to the Destination Node specified on the x-axis. 16
3.8 Throughput from DDR to HBM on CPU. Using the exact same code as smart copy, therefore spawning one thread on Node 0, 1, 2 and 3. Copying 1 GiB from Node 0 to the Destination Node specified on the x-axis. This shows the low performance of software path, when not adapting the code. 16

4.1 Public Interface of CacheData and Cache Classes. Colour coding for thread safety. Grey denotes impossibility for threaded access. Green indicates full safety guarantees only relying on atomics to achieve this. Yellow may use locking but is still safe for use. Red must be called from a single threaded context. 18

5.1 Sequence for Blocking Scenario. Observable in first draft implementation. Scenario where T_1 performed first access to a datum followed T_2 and T_3. Then T_1 holds the handlers exclusively, leading to the other threads having to wait for T_1 to perform the work submission and waiting before they can access the datum through the cache. 23

5.2 `CacheData::WaitOnCompletion` Pseudocode. Final rendition of the implementation for a fair wait function. 24

1 Introduction

write this chapter

2 Technical Background

write introductory paragraph

2.1 High Bandwidth Memory

High Bandwidth Memory is a novel memory technology promising an increase in peak bandwidth. It is composed of stacked DRAM dies [1, p. 1] and is slowly being integrated into server processors, the Intel® Xeon® Max Series [2] being one recent example. High Bandwidth Memory (HBM) on these systems may be configured in different memory modes, most notably, HBM Flat Mode and HBM Cache Mode [2]. The former gives applications direct control, requiring code changes while the latter utilizes the HBM as cache for the systems DDR based main memory [2].

2.2 Query Driven Prefetching

write this section

2.3 Intel Data Streaming Accelerator

> Intel DSA is a high-performance data copy and transformation accelerator that will be integrated in future Intel® processors, targeted for optimizing streaming data movement and transformation operations common with applications for high-performance storage, networking, persistent memory, and various data processing applications. [3, Ch. 1]

Introduced with the 4^{th} generation of Intel Xeon Scalable Processors, the DSA promises to alleviate the CPU from 'common storage functions and operations such as data integrity checks and deduplication' [4, p. 4]. To utilize the hardware optimally, knowledge of its workings is required. Therefore, we present an overview of the architecture, software, and the interaction of these two components, going into detail on the workings of the DSA engine itself. All statements are based on Chapter 3 of the Architecture Specification by Intel.

2.3.1 Hardware Architecture

The DSA chip is directly integrated into the processor and attaches via the I/O fabric interface over which all communication is conducted. Through this interface, it is accessible as a PCIe device. Therefore, configuration utilizes memory-mapped registers set in the devices Base Address Register (BAR). Through these, the devices' layout is defined and memory pages for work submission set. In a system with multiple processing

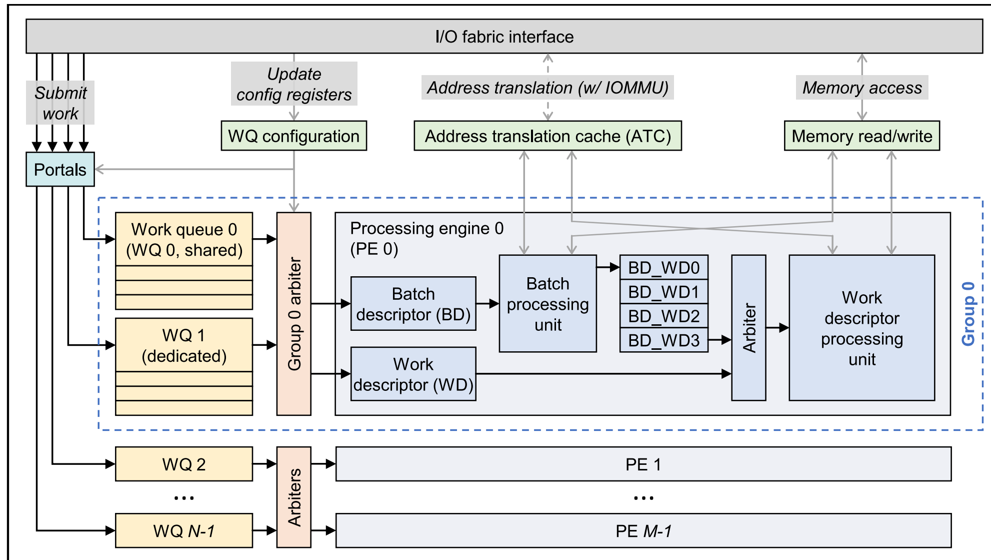

Figure 2.1: DSA Internal Architecture [5, Fig. 1 (a)]. Shows the components that the chip is made up of, how they are connected and which outside components the DSA communicates with.

nodes, there may also be one DSA per node, resulting in 4 being present on the previously mentioned Xeon Max CPU.

add citations to this section

To satisfy different use cases, the layout of the DSA may be software-defined. The structure is made up of three components, namely Work Queue (WQ), Engine and Group. WQs provide the means to submit tasks to the device and will be described in more detail shortly. They are marked yellow in Figure 2.1. An Engine is the processing-block that connects to memory and performs the described task. The grey block of Figure 2.1 shows the subcomponents that make up an engine and the different internal paths for a batch or task descriptor. Using Groups, Engines and WQs are tied together, indicated by the dotted blue line around the components of Group 0 in Figure 2.1. This means, that tasks from one WQ may be processed from multiple Engines and vice-versa, depending on the configuration. This flexibility is achieved through the Group Arbiter, represented by the orange block in Figure 2.1, which connects the two components according to the user-defined configuration.

too much detail for this being the first overview paragraph

A WQ is accessible through so-called portals, light blue in Figure 2.1, which are mapped memory regions. Submission of work is done by writing a descriptor to one of these. A descriptor is 64 bytes in size and may contain one specific task (task descriptor) or the location of a task array in memory (batch descriptor). Through these portals, the submitted descriptor reaches a queue. There are two possible queue types with different submission methods and use cases. The Shared Work Queue (SWQ) is intended to provide synchronized access to multiple processes and each group may only have one attached. A PCIe Deferrable Memory Write Request (DMR), which guarantees implicit synchronization, is generated via x86 Instruction ENQCMD and communicates with the device before writing [3, Sec. 3.3.1]. This may result in higher submission cost,

compared to the Dedicated Work Queue (DWQ) to which a descriptor is submitted via x86 Instruction MOVDIR64B [3, Sec. 3.3.2].

To handle the different descriptors, each Engine has two internal execution paths. One for a task and the other for a batch descriptor. Processing a task descriptor is straightforward, as all information required to complete the operation are contained within . For a batch, the DSA reads the batch descriptor, then fetches all task descriptors from memory and processes them [3, Sec. 3.8]. An Engine can coordinate with the operating system in case it encounters a page fault, waiting on its resolution, if configured to do so, while otherwise, an error will be generated in this scenario [3, Sec. 2.2, Block on Fault].

cite this

Ordering of operations is only guaranteed for a configuration with one WQ and one Engine in a Group when submitting exclusively batch or task descriptors but no mixture. Even then, only write-ordering is guaranteed, meaning that 'reads by a subsequent descriptor can pass writes from a previous descriptor'. A different issue arises, when an operation fails, as the DSA will continue to process the following descriptors from the queue. Care must therefore be taken with read-after-write scenarios, either by waiting for a successful completion before submitting the dependant, inserting a drain descriptor for tasks or setting the fence flag for a batch. The latter two methods tell the processing engine that all writes must be committed and, in case of the fence in a batch, abort on previous error. [3, Sec. 3.9]

An important aspect of modern computer systems is the separation of address spaces through virtual memory. Therefore, the DSA must handle address translation, as a process submitting a task will not know the physical location in memory which causes the descriptor to contain virtual values. For this, the Engine communicates with the Input/Output Memory Management Unit (IOMMU) and Address Translation Cache (ATC) to perform this operation, as visible in the outward connections at the top of Figure 2.1. For this, knowledge about the submitting processes is required, and therefore each task descriptor has a field for the Process Address Space ID (PASID) which is filled by the ENQCMD instruction for a SWQ [3, Sec. 3.3.1] or set statically after a process is attached to a DWQ [3, Sec. 3.3.2].

The completion of a descriptor may be signalled through a completion record and interrupt, if configured so. For this, the DSA 'provides two types of interrupt message storage: (1) an MSI-X table, enumerated through the MSI-X capability; and (2) a device-specific Interrupt Message Storage (IMS) table' [3, Sec. 3.7].

2.3.2 Software View

Since Linux Kernel 5.10, there exists a driver for the DSA which has no counterpart in the Windows OS-Family [6, Sec. Installation], meaning that accessing the DSA is only possible under Linux. To interface with the driver and perform configuration operations, Intel's accel-config [7] user space toolset may be used which provides a command-line interface and can read configuration files to set up the device as described previously. This can be seen in the upper block titled 'User space' in Figure 2.2. It interacts with the kernel driver, light green and labled 'IDXD' in Figure 2.2, to achieve this. After successful

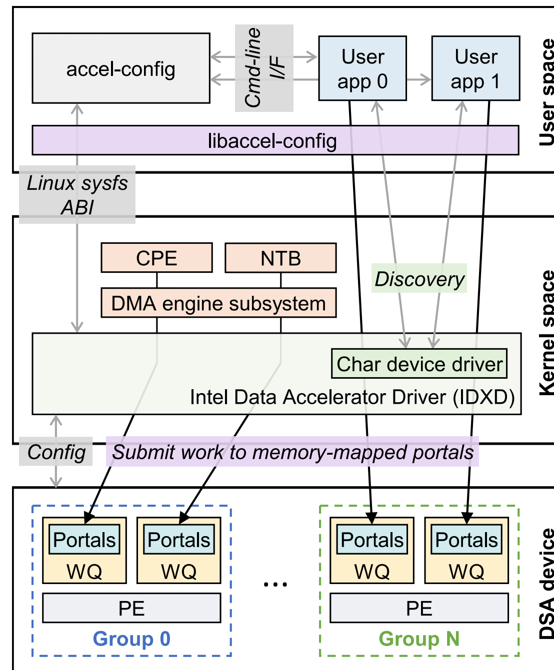

Figure 2.2: DSA Software View [5, Fig. 1 (b)]. Illustrating the software stack and internal interactions from user applications, through the driver to the portal for work submission.

configuration, each WQ is exposed as a character device by `mmap` of the associated portal [5, Sec. 3.3].

Given the file permissions, it would now be possible for a process to submit work to the DSA via either MOVDIR64B or ENQCMD instructions, providing the descriptors by manually configuring them. This, however, is quite cumbersome, which is why Intel Data Mover Library (Intel DML) exists.

With some limitations, like lacking support for DWQ submission, this library presents an interface that takes care of creation and submission of descriptors, and error handling and reporting. Thanks to the high-level-view the code may choose a different execution path at runtime which allows the memory operations to either be executed in hardware or software. The former on an accelerator or the latter using equivalent instructions provided by the library. This makes code using this library automatically compatible with systems that do not provide hardware support. [6, Sec. Introduction]

2.4 Programming Interface

As mentioned in Subsection 2.3.2, Intel DML provides a high level interface to interact with the hardware accelerator, namely Intel DSA. We choose to use the C++ interface and will now demonstrate its usage by example of a simple memcopy-implementation for the DSA.

```
template <path>
bool DsaMemcpy(char* dst, const char* src, size_t size, int node)

numa_run_on_node(node)

src_view ← dml::make_view(src, size)
dst_view ← dml::make_view(dst, size)

handler ← dml::submit<path>(dml::mem_copy.block_on_fault(), src_view, dst_view)

result ← handler.get()

return result.status == dml::status_code::ok
```

Figure 2.3: DML Memcpy Implementation Pseudocode. Performs copy operation of a block of memory from source to destination. The DSA executing this copy can be selected with the parameter node, and the template parameter path elects whether to run on hardware (DSA) or software (CPU).

In the function header of Figure 2.3 we notice two differences, when comparing with standard memcpy. The first is the template parameter named path and the second being the additional parameter int node which we will discuss later. With path the executing device, which can be the CPU or DSA, is selected, giving the option between dml::software (CPU), dml::hardware (DSA) and dml::automatic where the latter dynamically selects the device at runtime, preferring DSA over CPU execution. [6, Sec. Quick Start]

Choosing the engine which carries out the copy might be advantageous for performance, as we can see in Subsection 3.2.3. With the engine directly tied to the CPU node, as observed in Subsection 2.3.1, the CPU Node ID is equivalent to the ID of the DSA. As the library has limited NUMA support and therefore only utilizes the DSA device on the node which the current thread is assigned to, we must assign the currently running thread to the node in which the desired DSA resides. This is the reason for adding the parameter int node, which is used in the first step of Figure 2.3, where we manually set the node assignment according to it, using numa_run_on_node(node) for which more information may be obtained in the respective manpage of libnuma [8].

Intel DML operates on so-called data views which we must create from the given pointers and size in order to provide locations to the library. This is done using dml::make_view(uint8_t* ptr, size_t size) with which we create views for both source and destination, labled src_view and dst_view in Figure 2.3. [6, Sec. High-level C++ API, Make view]

We submit a single descriptor using the asynchronous operation from Intel DML in Figure 2.3. This uses the function dml::submit<path>, which takes an operation type and parameters specific to the selected type and returns a handler to the submitted

task which can later be queried for completion of the operation. Passing the source and destination views, together with the operation `dml::mem_copy`, we again notice one element sticking out of the call. This is the addition of `.block_on_fault()` which lets the DSA handle a page fault by coordinating with the operating system. This only works if the device is configured to accept this flag. [6, Sec. High-level C++ API, How to Use the Library] [6, Sec. High-level C++ API, Page Fault handling]

After submission, we poll for the task completion with `handler.get()` and check whether the operation completed successfully.

2.5 System Setup and Configuration

In this section we will give a brief step-by-step list of setup instructions to replicate the configuration being used for benchmarks and testing purposes in the following chapters. We found Intel's guide on DSA usage useful but consulted articles for setup on Lenovo ThinkSystem Servers for crucial information not present in the former. Instructions for configuring the HBM access mode, as mentioned in Section 2.1, may vary from system to system and can require extra steps not found in the list below.

1. Set 'Memory Hierarchy' to Flat [9, Sec. Configuring HBM, Configuring Flat Mode], 'VT-d' to Enabled in BIOS [10, Sec. 2.1] and, if available, 'Limit CPU PA to 46 bits' to Disabled in BIOS [11, p. 5]

2. Use a kernel with IDXD driver support, available from Linux 5.10 or later [6, Sec. Installation] and append the following to the kernel boot parameters in grub config: `intel_iommu=on,sm_on` [11, p. 5]

3. Evaluate correct detection of DSA devices using `dmesg | grep idxd` which should list as many devices as NUMA nodes on the system [11, p. 5]

4. Install `accel-config` from repository [7] or system package manager and inspect the detection of DSA devices through the driver using `accel-config list -i` [11, p. 6]

5. Create DSA configuration file for which we provide an example under `benchmarks/configuration-files/8n1d1e1w.conf` in the accompanying repository [12] that is used for most benchmarks available. Then apply the configuration using `accel-config load-config -c [filename] -e` [10, Fig. 3-9]

6. Inspect the now configured DSA devices using `accel-config list` [11, p. 7], output should match the desired configuration set in the file used

3 Performance Microbenchmarks

The performance of DSA has been evaluated in great detail by Reese Kuper et al. in [5]. Therefore, we will perform only a limited amount of benchmarks with the purpose of verifying the figures from [5] and analysing best practices and restrictions for applying DSA to Query-driven Prefetching (QdP).

3.1 Benchmarking Methodology

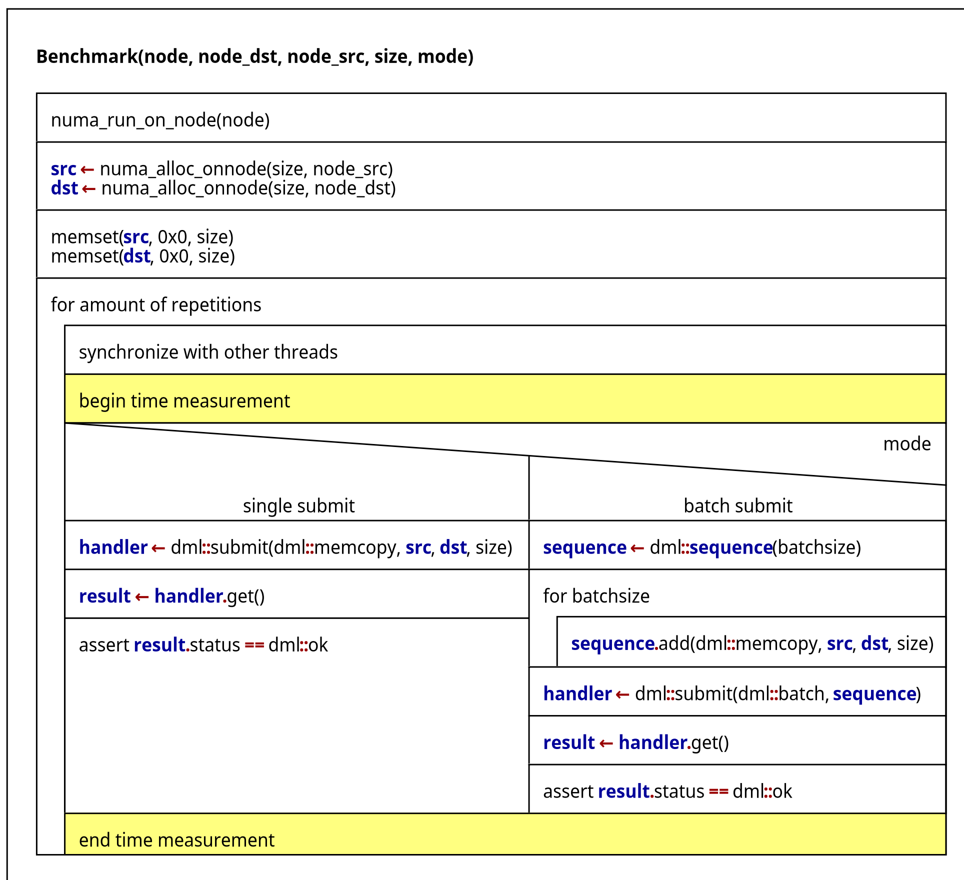

Figure 3.1: Benchmark Procedure Pseudocode. Timing marked with yellow background. Showing data allocation and the benchmarking loop for a single thread.

Benchmarks were conducted on an Intel Xeon Max CPU, system configuration following Section 2.5 with exclusive access to the system. As Intel DML does not have support for DWQ, we ran benchmarks exclusively with access through SWQ. The application written for the benchmarks can be obtained in source form under `benchmarks` in the thesis repository [12]. With the full source code available we only briefly describe a section of pseudocode, as seen in Figure 3.1, in the following paragraph.

The benchmark performs node setup as described in Section 2.4 and allocates source and destination memory on the nodes passed in as parameters. To avoid page faults affecting the results, the entire memory regions are written to before the timed part of the benchmark starts. These timings are marked with yellow background in Figure 3.1. To get accurate results, the benchmark is repeated multiple times. At the beginning of each repetition, all threads running will synchronize by use of a barrier. The behaviour then differs depending on the submission method selected which can be a single submission or a batch of given size. This can be seen in Figure 3.1 at the switch statement for 'mode'. Single submission follows the example given in Section 2.4, and we therefore do not go into detail explaining it here. Batch submission works unlike the former. A sequence with specified size is created which tasks are then added to. This sequence is then submitted to the engine similar to the submission of a single descriptor. Further steps then follow the example from Section 2.4 again.

3.2 Benchmarks

In this section we will describe three benchmarks. Each complete with setup information and preview, followed by plots showing the results and detailed analysis. We formulate expectations and contrast them with the observations from our measurements. Where displayed, the slim grey bars represent the standard deviation across iterations.

3.2.1 Submission Method

With each submission, descriptors must be prepared and sent off to the underlying hardware. This is expected to come with a cost, affecting throughput sizes and submission methods differently. By submitting different sizes and comparing batching and single submission, we will evaluate at which submission method performs best for each tested data size. We expect single submission to perform worse consistently, with a pronounced effect on smaller transfer sizes. This is assumed, as the overhead of a single submission with the SWQ is incurred for every iteration, while the batch only sees this overhead once for multiple copies.

In Figure 3.2 we conclude that with transfers of 1 MiB and upwards, the submission method makes no noticeable difference. For smaller transfers the performance varies greatly, with batch operations leading in throughput. This is aligned with the finding that 'SWQ observes lower throughput between 1-8 KB [transfer size]' [5, p. 6 and 7] for normal submission method.

Another limitation may be observed in this result, namely the inherent throughput limit per DSA chip of close to 30 GiB/s. This is apparently caused by I/O fabric limitations [5, p. 5]. We therefore conclude, that the use of multiple DSA is required to

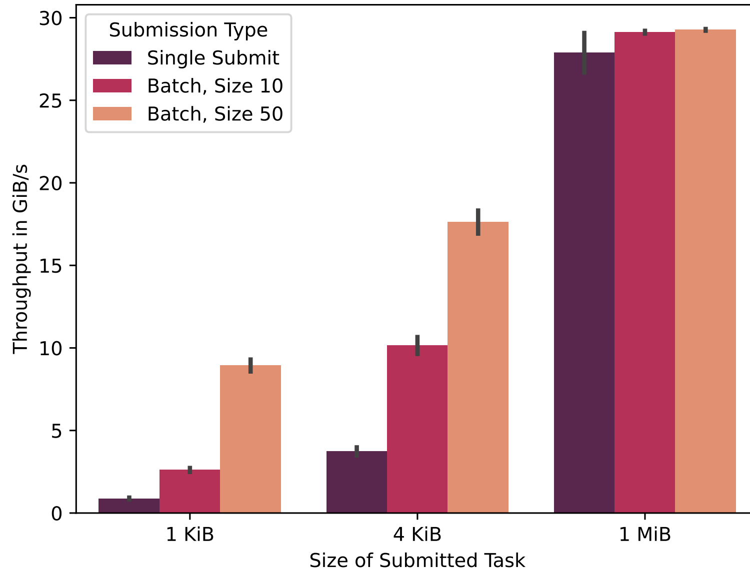

Figure 3.2: Throughput for different Submission Methods and Sizes. Performing a copy with source and destination being node 0, executed by the DSA on node 0. Observable is the submission cost which affects small transfer sizes differently, as there the completion time is lower.

fully utilize the available bandwidth of HBM which theoretically lies at 256 GB/s [1, Table I].

3.2.2 Multithreaded Submission

As we might encounter access to one DSA from multiple threads through the Shared Work Queue associated with it, determining the effect this type of access has is important. We benchmark multithreaded submission for one, two and twelve threads. The number twelve representing the core count of one node on the test system. Each configuration gets the same 120 copy tasks split across the available threads, all submitting to one DSA. We perform this with sizes of 1 MiB and 1 GiB to see, if the behaviour changes with submission size. As for smaller sizes, the completion time may be faster than submission time. Therefore, smaller task sizes may see different effects of threading due to the fact that multiple threads can work to fill the queue and ensure there is never a stall due it being empty. On the other hand, we might experience lower-than-peak throughput with rising thread count, caused by the synchronization inherent with SWQ.

In Figure 3.3 we see that threading has no negative impact. The synchronization therefore affects single threaded access in the same way it does for multiple. For the smaller size of 1 MiB our assumption came true and performance actually increased with threads which we attribute to queue usage. We did not find an explanation for the speed difference between sizes, which goes against the rationale that higher transfer size would lead to less impact of submission time and therefore higher throughput.

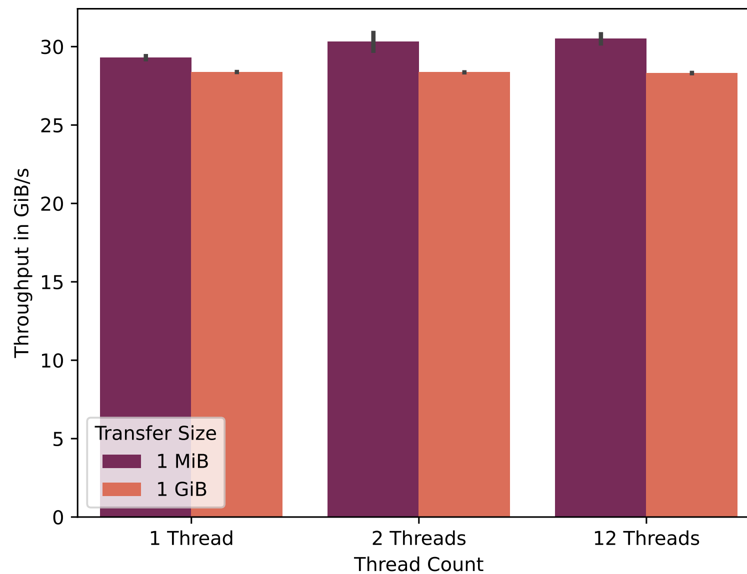

Figure 3.3: Throughput for different Thread Counts and Sizes. Multiple threads submit to the same Shared Work Queue. Performing a copy with source and destination being node 0, executed by the DSA on node 0.

3.2.3 Data Movement from DDR to HBM

Moving data from DDR to HBM is most relevant to the rest of this work, as it is the target application. As we discovered in Section 3.2.1, one DSA has a peak bandwidth limit of 30 GiB/s. We write to HBM with its theoretical peak of 256 GB/s [1, Table I]. Our top speed is therefore limited by the slower main memory. For each node, the test system is configured with two DIMMs of DDR5-4800. We calculate the theoretical throughput as follows: $2\ DIMMs * \frac{4800\ Megatransfers}{Second\ and\ DIMM} * \frac{64-bitwidth}{8bits/byte} = 76800\ MT/s = 75\ GiB/s$. We conclude that to achieve closer-to-peak speeds, a copy task has to be split across multiple DSA.

how to verify this calculation with a citation? is it even correct because we see close to 100 GiB/s throughput?

Two methods of splitting will be evaluated. The first being a brute force approach, utilizing all available for any transfer direction. The seconds' behaviour depends on the data source and destination locations. As our system has multiple sockets, communication crossing the socket could experience a latency and bandwidth disadvantage . We argue that for intra-socket transfers, use of the DSA from the second socket will have only marginal effect. For transfers crossing the socket, every DSA chip will perform badly, and we choose to only use the ones with the fastest possible access, namely the ones on the destination and source node. This might result in lower performance but also uses one fourth of the engines of the brute force approach for inter-socket and half for intra-socket. This gives other threads the chance to utilize these now-free chips.

cite this

For this benchmark, we copy 1 Gibibyte of data from node 0 to the destination node, using the previously described submission method. For each node utilized, we spawn one thread pinned to it, in charge of submission. We display data for nodes 8, 11, 12 and 15.

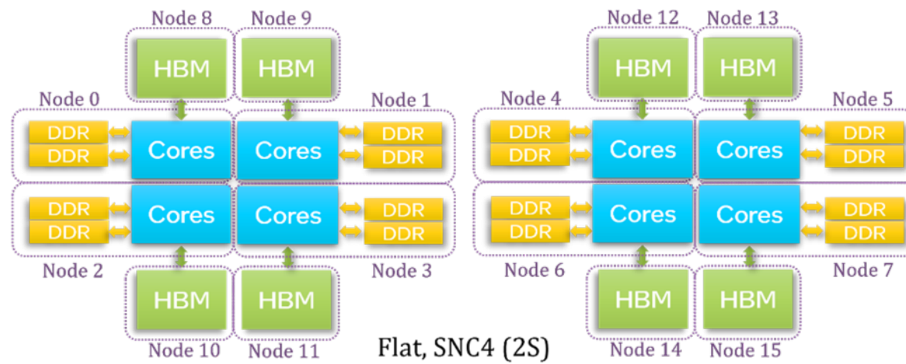

Figure 3.4: Xeon Max NUMA-Node Layout [13, Fig. 14] for a 2-Socket System when configured with HBM-Flat. Showing separate Node IDs for manual HBM access and for Cores and DDR memory.

To understand the selection, we refer to Figure 3.4 which shows the configured systems' node ids and the storage technology they represent. Node 8 accesses the HBM on node 0 and is therefore the physically closest destination possible. Node 11 is located diagonally on the chip and therefore the furthest intra-socket operation benchmarked. Node 12 and 15 lie diagonally on the second sockets CPU and should therefore be representative of inter-socket transfer operations.

From the brute force approach in Figure 3.5, we observe peak speeds of 96 GiB/s when copying across the socket from Node 0 to Node 15. This invalidates our assumption, that peak bandwidth would be achieved in the intra-socket scenario and goes against the calculated peak throughput of the memory on our system . The results however, align with findings in [5, Fig. 10].

find out why maybe?

calculation wrong? other factors?

While using significantly more resources, the brute force copy shown in Figure 3.5 outperforms the smart approach from Figure 3.6. We observe an increase in transfer speed by utilizing all available DSA of 2 GiB/s for copy to Node 8, 18 GiB/s for Node 11 and 12 and 30 GiB/s for Node 15. The smart approach could accommodate another intra-socket copy on the second socket, we assume without seeing negative impact. From this we conclude that the smart copy assignment is worth to use, as it provides better scalability.

3.2.4 Data Movement using CPU

For evaluating CPU copy performance two approaches were selected. For Figure 3.8 we used the smart copy procedure as in 3.2.3, just selecting the software path, with no other changes to the test. In Figure 3.7 the test was adapted to spawn 12 threads on Node 0.

The benchmark resulting in Figure 3.7 fully utilizes node 0 of the test system through spawning 12 threads. We attribute the low performance of the inter-socket copy operation to overhead of the interconnect. The slight difference between node 12 and 15 may be explained using Figure 3.4, where we observe that physically node 12 is closer to node 0 than node 15.

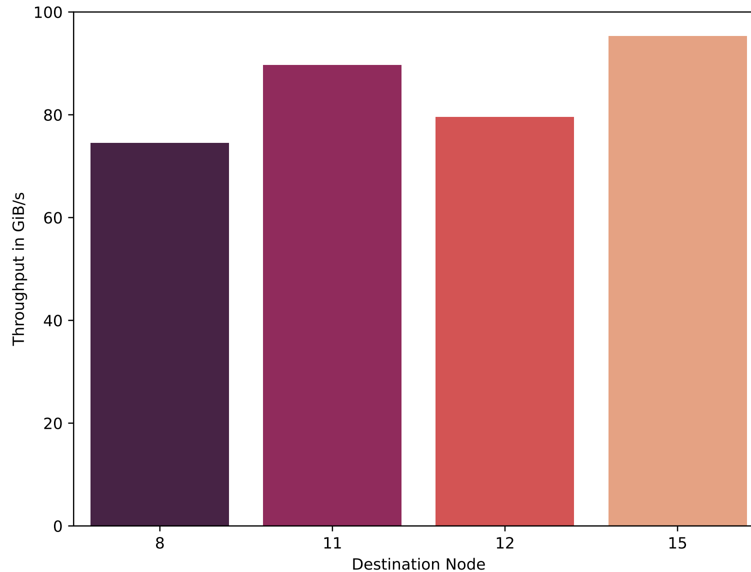

Figure 3.5: Throughput for brute force copy from DDR to HBM. Using all available DSA. Copying 1 GiB from Node 0 to the Destination Node specified on the x-axis. Shows peak achievable with DSA.

Comparing the results in Figure 3.8 with the data from the adapted test displayed in Figure 3.7 we conclude that the software path just serves compatibility. The throughput is comparatively low, requiring code changes to increase performance. This leads us not to recommend using it.

3.3 Analysis

In this section we summarize the conclusions drawn from the three benchmarks performed in the sections above and outline a utilization guideline. We also compare CPU and DSA for the task of copying data from DDR to HBM.

1. From 3.2.1 we conclude that small copies under 1 MiB in size require batching and still do not reach peak performance. Datum size should therefore be at or above 1 MiB if possible.

2. Subsection 3.2.2 assures that access from multiple threads does not negatively affect the performance when using Shared Work Queue for work submission.

3. In 3.2.3, we chose to use the presented smart copy methodology to split copy tasks across multiple DSA chips to achieve low utilization with acceptable performance.

We again refer to Figures 3.5 and 3.7 which both represent the maximum throughput achieved with utilization of either DSA for the former and CPU for the latter. Noticeably, the DSA does not seem to suffer from the inter-socket overhead like the CPU. To the

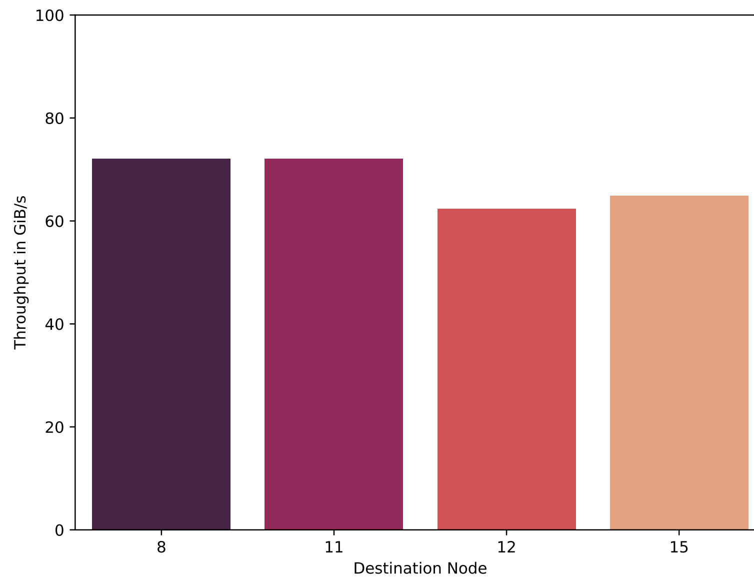

Figure 3.6: Throughput for smart copy from DDR to HBM. Using four on-socket DSA for intra-socket operation and the DSA on source and destination node for inter-socket. Copying 1 GiB from Node 0 to the Destination Node specified on the x-axis. Shows conservative performance.

explanation?

contrary, we do see the highest throughput when copying across sockets . In any case, DSA outperforms the CPU for transfer sizes over 1 MiB, demonstrating its potential to increase throughput while at the same time freeing CPU cycles.

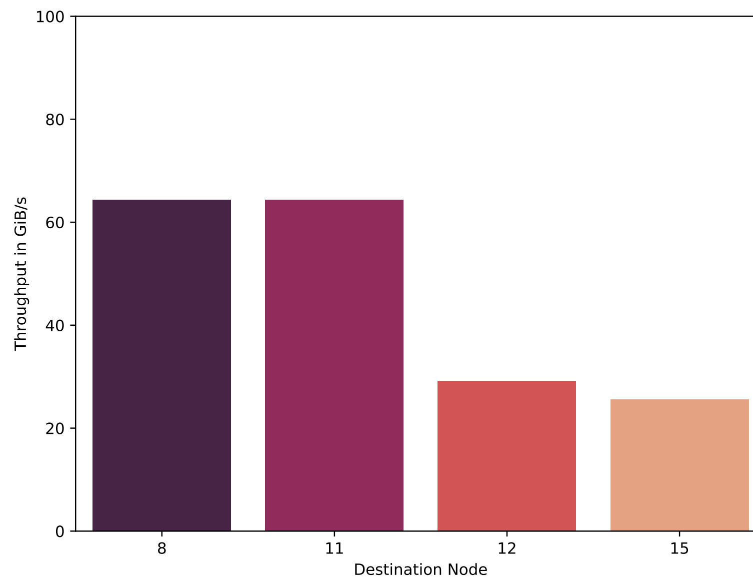

Figure 3.7: Throughput from DDR to HBM on CPU. Using 12 Threads spawned on Node 0 for the task. Copying 1 GiB from Node 0 to the Destination Node specified on the x-axis.

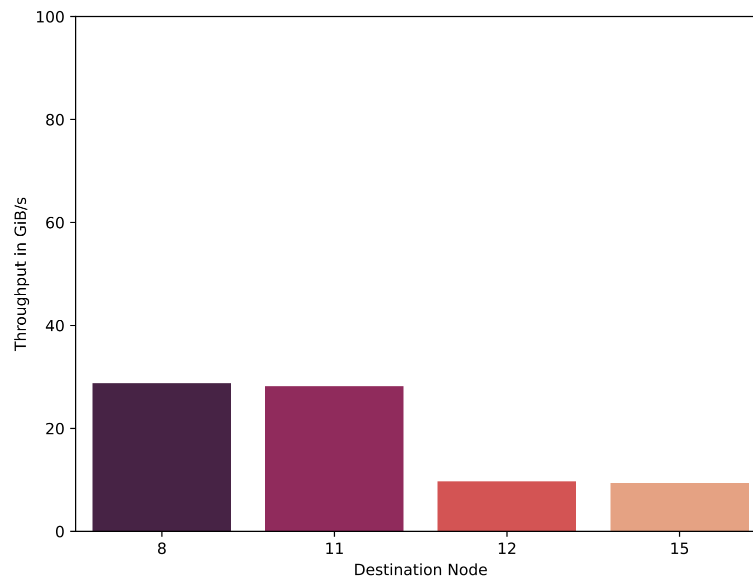

Figure 3.8: Throughput from DDR to HBM on CPU. Using the exact same code as smart copy, therefore spawning one thread on Node 0, 1, 2 and 3. Copying 1 GiB from Node 0 to the Destination Node specified on the x-axis. This shows the low performance of software path, when not adapting the code.

4 Design

write introductory paragraph

4.1 Detailed Task Description

write this section

- give slightly more detailed task Description
- perspective of "what problems have to be solved"
- not "what is querry driven prefetching"

4.2 Cache Design

The task of prefetching is somewhat aligned with that of a cache. As a cache is more generic and allows use beyond Query Driven Prefetching, the choice was made to solve the prefetching offload by implementing an offloading `Cache`. When referring to the provided implementation, `Cache` will be used from now on. The interface with `Cache` must provide three basic functions: requesting a memory block to be cached, accessing a cached memory block and synchronizing cache with the source memory. The latter operation comes in to play when the data that is cached may also be modified, requiring the entry to be updated with the source or the other way around. Due to the many possible setups and use cases, the user should also be responsible for choosing cache placement and the copy method. As re-caching is resource intensive, data should remain in the cache for as long as possible while being removed when memory pressure due to restrictive memory size drives the `Cache` to flush unused entries.

4.2.1 Interface

To allow rapid integration and ease developer workload, a simple interface was chosen. As this work primarily focuses on caching static data, the choice was made only to provide cache invalidation and not synchronization. Given a memory address, `Cache::Invalidate` will remove all entries for it. The other two operations are provided in one single function, which we shall call `Cache::Access` henceforth, receiving a data pointer and size it takes care of either submitting a caching operation if the pointer received is not yet cached or returning the cache entry if it is. The cache placement and assignment of the task to accelerators are controlled by the user. In addition to the two basic operations outlined before, the user also is given the option to flush the cache using `Cache::Flush` of unused elements manually or to clear it completely with `Cache::Clear`. This interface is represented on the right block of Figure 4.1 labelled 'Cache'.

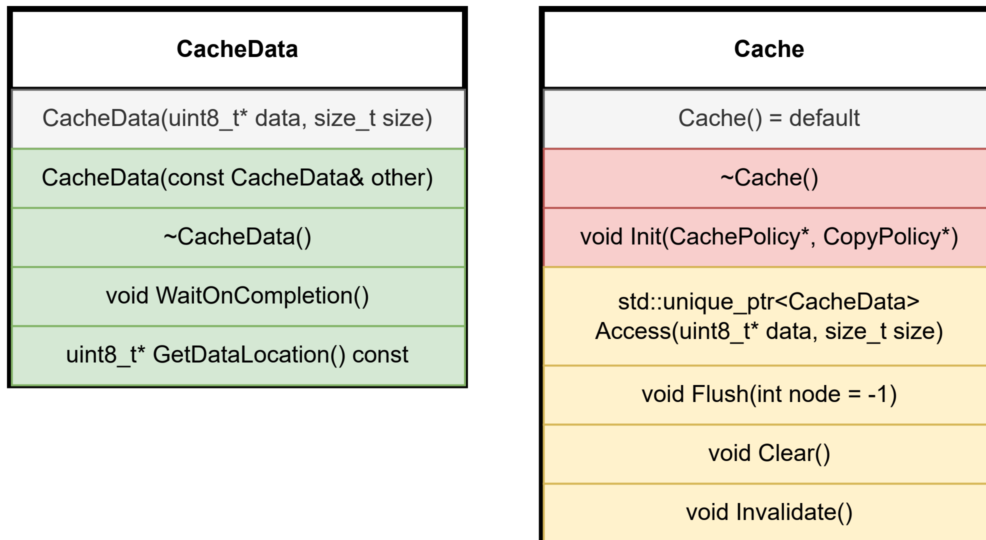

Figure 4.1: Public Interface of CacheData and Cache Classes. Colour coding for thread safety. Grey denotes impossibility for threaded access. Green indicates full safety guarantees only relying on atomics to achieve this. Yellow may use locking but is still safe for use. Red must be called from a single threaded context.

As caching is performed asynchronously, the user may wish to wait on the operation. This would be beneficial if there are other threads making progress in parallel while the current thread waits on its data becoming available in the faster cache, speeding up local computation. To achieve this, the `Cache::Access` will return an instance of an object which from hereinafter will be referred to as `CacheData`. Through `CacheData::GetDataLocation` a pointer to the cached data will be retrieved, while also providing `CacheData::WaitOnCompletion` which must only return when the caching operation has completed and during which the current thread is put to sleep, allowing other threads to progress. Figure 4.1 also documents the public interface for `CacheData` on the left block labelled as such.

4.2.2 Cache Entry Reuse

When multiple consumers wish to access the same memory block through the `Cache`, we could either provide each with their own entry, or share one entry for all consumers. The first option may cause high load on the accelerator due to multiple copy operations being submitted and also increases the memory footprint of the system. The latter option requires synchronization and more complex design. As the cache size is restrictive, the latter was chosen. The already existing `CacheData` will be extended in scope to handle this by allowing copies of it to be created which must synchronize with each other for `CacheData::WaitOnCompletion` and `CacheData::GetDataLocation`. This is shown by the green markings, signalling thread safety guarantees for access in Figure 4.1.

4.2.3 Cache Entry Lifetime

By allowing multiple references to the same entry, memory management becomes a concern. Freeing the allocated block must only take place when all copies of a `CacheData` instance are destroyed, therefore tying cache entry lifetime to the lifetime of the longest living copy of the original instance. This makes access to the entry legal during the lifetime of any `CacheData` instance, while also guaranteeing that `Cache::Clear` will not have any unforeseen side effects, as deallocation only takes place when the last consumer has `CacheData` go out of scope or manually deletes it.

4.2.4 Usage Restrictions

As cache invalidation applies mainly to non-static data which this work does not focus on, two restrictions are placed on the invalidation operation. This permits drastically simpler cache design, as a fully coherent cache would require developing a thread safe coherence scheme which is outside our scope.

Firstly, overlapping areas in the cache will cause undefined behaviour during invalidation of any one of them. Only the entries with the equivalent source pointer will be invalidated, while other entries with differing source pointers which, due to their size, still cover the now invalidated region, will not be invalidated. At this point, the cache may and may continue to contain invalid elements.

Secondly, invalidation is to be performed manually, requiring the programmer to remember which points of data are at any given point in time cached and invalidating them upon modification. No ordering guarantees will be given for this situation, possibly leading to threads still having a pointer to now-outdated entries and continuing their progress with this.

Due to its reliance on libnuma for memory allocation and thread pinning, `Cache` will only work on systems where this library is present, excluding, most notably, Windows from the compatibility list.

4.2.5 Accelerator Usage

Compared with the challenges of ensuring correct entry lifetime and thread safety, the application of DSA for the task of duplicating data is simple, thanks partly to Intel DML [6]. Upon a call to `Cache::Access` and determining that the given memory pointer is not present in cache, work will be submitted to the Accelerator. Before, however, the desired location must be determined which the user-defined cache placement policy function handles. With the desired placement obtained, the copy policy then determines, which nodes should take part in the copy operation which is equivalent to selecting the Accelerators following 2.3.1. This causes the work to be split upon the available accelerators to which the work descriptors are submitted at this time. The handlers that Intel DML [6] provides will then be moved to the `CacheData` instance to permit the callee to wait upon caching completion. As the choice of cache placement and copy policy is user-defined, one possibility will be discussed in 5.

5 Implementation

write introductory paragraph

5.1 Locking and Usage of Atomics

The usage of locking and atomics proved to be the challenging. Their use is performance critical and mistakes may lead to deadlock. Therefore, they also constitute the most interesting part of the implementation which is why this chapter will focus extensively on the details of the implementation in regard to these.

5.1.1 Cache State Lock

To keep track of the current cache state the `Cache` will hold a reference to each currently existing `CacheData` instance. The reason for this is twofold: In 4.2 we decided to keep elements in the cache until forced by memory pressure to remove them. Secondly in 4.2.2 we decided to reuse one cache entry for multiple consumers. The second part requires access to the structure holding this reference to be thread safe when accessing and extending the cache state in `Cache::Access`, `Cache::Flush` and `Cache::Clear`. The latter two both require a unique lock, preventing other calls to `Cache` from making progress while the operation is being processed. For `Cache::Access` the use of locking depends upon the caches state. At first only a shared lock is acquired for checking whether the given address already resides in cache, allowing other `Cache::Access`-operations to also perform this check. If no entry for the region is present, a unique lock is required as well when adding the newly created entry to cache.

A map was chosen to represent the current cache state with the key being the memory address of the entry and as value the `CacheData` instance. As the caching policy is controlled by the user, one datum may be requested for caching in multiple locations. To accommodate this, one map is allocated for each available node of the system. This can be exploited to reduce lock contention by separately locking each nodes state instead of utilizing a global lock. This ensures that `Cache::Access` and the implicit `Cache::Flush` it may cause can not hinder progress of caching operations on other nodes. Both `Cache::Clear` and a complete `Cache::Flush` as callable by the user will now iteratively perform their respective task per nodes state, also allowing other nodes to progress.

Even with this optimization, in scenarios where the `Cache` is frequently tasked with flushing and re-caching by multiple threads from the same node lock contention will negatively impact performance by delaying cache access. Due to passive waiting, this impact might be less noticeable when other threads on the system are able to make progress during the wait.

5.1.2 CacheData Atomicity

The choice made in 4.2.2 requires thread safe shared access to the same resource. `std::shared_ptr<T>` provides a reference counted pointer, which is thread safe for the required operations, making it a prime candidate for this task. An implementation using it was explored but proved to offer its own set of challenges. As we wish to reduce time spent in a locked region, the task is only added to the nodes cache state when locked. Submission takes place outside, which is sensible, as this submitting one task should not hinder accessing another. To achieve the safety for `CacheData::WaitOnCompletion` outlined in 4.2.2 this would require the threads to coordinate which thread performs the actual waiting, as we assume the handlers of Intel DML to be non-threadsafe. In order to avoid queuing multiple of the same copies, the task must be added before submission. This results in a `CacheData` instance with invalid cache pointer and no handlers to wait for being available, requiring additional usage of synchronization primitives. With using `std::shared_ptr<T>` also comes the uncertainty of relying on the implementation to be performant. The standard does not specify whether a lock-free algorithm is to be used and [14] suggests abysmal performance for some implementations, although the full article is in Korean. No further research was found on this topic.

It was therefore decided to implement atomic reference counting for `CacheData` which means providing a custom constructor and destructor wherein a shared (through a standard pointer however) atomic integer is either incremented or decremented using atomic fetch sub and add operations [15] to increase or deacrease the reference counter and, in case of decrease in the destructor signals that the destructor is called for the last reference, perform actual destruction. The invalid state of `CacheData` achievable is also avoided. To achieve this, the waiting algorithm requires the handlers to be contained in an atomic pointer and the pointer to the cache memory be atomic too. Through this we may use the atomic wait operation which is guaranteed by the standard to be more efficient than simply spinning on Compare-And-Swap [16]. Some standard implementations achieve this by yielding after a short spin cycle [17].

Designing the wait to work from any thread was complicated. In the first implementation, a thread would check if the handlers are available and if not atomically wait [16] on a value change from nullptr. As the handlers are only available after submission, a situation could arise where only one copy of `CacheData` is capable of actually waiting on them. Lets assume that three threads T_1, T_2 and T_3 wish to access the same resource. T_1 now is the first to call `CacheData::Access` and therefore adds it to the cache state and will perform the work submission. Before T_1 may submit the work, it is interrupted and T_2 and T_3 obtain access to the incomplete `CacheData` on which they wait, causing them to see a nullptr for the handlers but invalid cache pointer, leading to atomic wait on the cache pointer (marked blue lines in 5.1). Now T_1 submits the work and sets the handlers (marked red lines in 5.1), while T_2 and T_3 continue to wait. Now only T_1 can trigger the waiting and is therefore capable of keeping T_2 and T_3 from progressing. This is undesirable as it can lead to deadlocking if by some reason T_1 does not wait and at the very least may lead to unnecessary delay for T_2 and T_3 if T_1 does not wait immediately.

To solve this, a different and more complicated order of waiting operations is required. When waiting, the threads now immediately check whether the cache pointer contains a

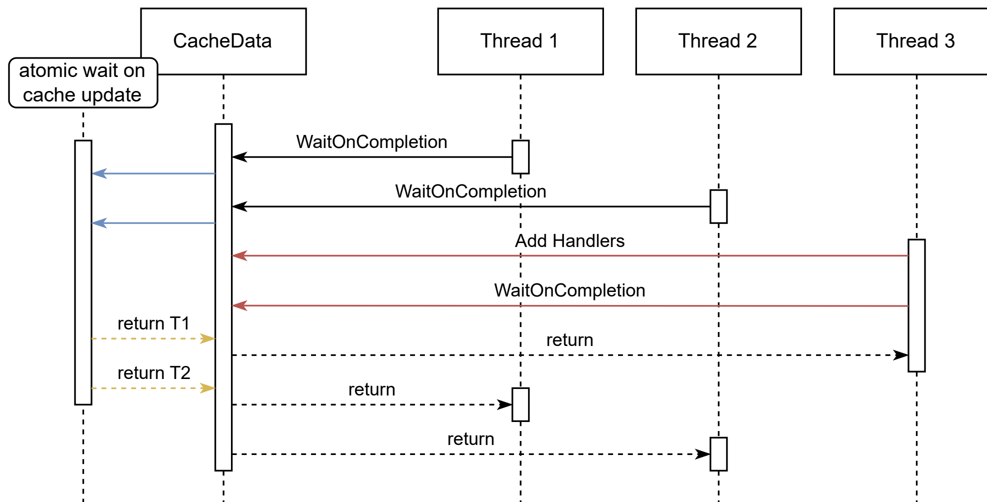

Figure 5.1: Sequence for Blocking Scenario. Observable in first draft implementation. Scenario where T_1 performed first access to a datum followed T_2 and T_3. Then T_1 holds the handlers exclusively, leading to the other threads having to wait for T_1 to perform the work submission and waiting before they can access the datum through the cache.

valid value and return if it does, as nothing has to be waited for in this case. Let's take the same example as before to illustrate the second part of the waiting procedure. T_2 and T_3 now both arrive in this latter section as the cache was invalid at the point in time when waiting was called for. They now atomically wait on the handlers pointer to change, instead of doing it the other way around as before. Now when T_1 supplies the handlers, it also uses `std::atomic<T>::notify_one` [18] to wake at least one thread waiting on value change of the handlers pointer, if there are any. Through this the exclusion that was observable in the first implementation is already avoided. If nobody is waiting, then the handlers will be set to a valid pointer and a thread may pass the atomic wait instruction later on. Following this wait, the handlers pointer is atomically exchanged [19] with nullptr, invalidating it. Now each thread again checks whether it has received a valid local pointer to the handlers from the exchange, if it has then the atomic operation guarantees that is now in sole possession of the pointer. The owning thread is tasked with actually waiting. All other threads will now regress and call `CacheData::WaitOnCompletion` again. The solo thread may proceed to wait on the handlers and should update the cache pointer.

Some two additional cases must be considered for the latter implementation to be safe. The wait operation first checks for a valid cache pointer and then waits on the handlers becoming valid. After processing the handlers, they are deleted and the pointer therefore invalidated. Should the cache pointer now be invalid as well, deadlocks would ensue. Therefore, the thread which exchanged the handlers pointer for a valid local copy must set the cache pointer to a valid value. Should one of the offloaded operations have failed, using the cache pointer is out of question as the datum it references might be

void CacheData::WaitOnCompletion()
is cache valid (not nullptr)
T
F
return
while handlers is not valid (equals nullptr)
handlers.wait()
local_handlers ← handlers.exchange(uint64::max)
handlers ← uint64::max
is local_handlers valid (not nullptr and not uint64::max)
T
F
foreach handler in local_handlers
status ← handler.wait()
status contains error
T
F
cache ← source pointer
free(destination pointer)
cache ← destination pointer
signal threads waiting on cache or handlers
while cache is not valid (equals nullptr)
cache.wait()

Figure 5.2: CacheData::WaitOnCompletion Pseudocode. Final rendition of the implementation for a fair wait function.

invalid. The cache is set to the source address in this case. Secondly, after one thread has exchanged the pointer locally, threads may collect waiting on the handlers to become available. This can happen when the wait on the handlers takes sufficient amount of time during which both handlers and cache pointer are invalid. After waiting, the responsible thread must therefore signal all [20] threads waiting on the handler to continue.

Two types of deadlocks were encountered during testing and have been accounted for. On one hand, it was found that the guarantee of std::atomic<T>::wait to only wake up when the value has changed [16] is stronger than the promise of waking up all waiting threads with std::atomic<T>::notify_all [20]. The value of the handler pointer may therefore not be exchanged with nullptr which is the value we wait on. As the highest envisioned address requires the lower 52-bits of current 64-bit wide systems [21, p. 120] [22, p. 4-2] setting all bits of a 64-bit-value yields an invalid pointer which is used as the second invalid state possible. The second type was encountered when after creating a CacheData instance it was determined this exists in the cache already and dropped. As destruction waits on completion in order to ensure that no further jobs require the memory held, a deadlock would arise from the cache and handler pointers both being null and no handlers ever being set due to the instance being deleted immediately. To circumvent this, the constructor of CacheData was modified to point to source memory

by default. Only after calling a separate initialization function will `CacheData` replace this with nullptr, therefore readying the instance for multithreaded usage.

5.1.3 Performance Guideline

Atomic operations come with an added performance penalty. No recent studies were found on this, although we assume that the findings of Hermann Schweizer et al. in "Evaluating the Cost of Atomic Operations on Modern Architectures" [23] still hold true for today's processor architectures. Due to the inherent cache synchronization mechanisms in place [23, Subsection IV.A.3 Off-Die Access], they observed significant access latency increase depending on whether the atomic variable was located on the local core, on a different core on the same chip or on another socket [23, Fig. 4]. Reducing the cost of atomic accesses would require a less generic implementation, reducing some of the guarantees we give in 4.2. This would allow reducing the amount of atomics required but is outside the scope of this work.

With the distributed locking described in 5.1.1, lock contention should not have a significant impact, although this remains to be tested. In addition to that, passive waiting at the contended section will benefit other threads and might allow overall progress to continue, if the application utilizing the cache has a threading model supporting this. These two factors lead us to classify lock contention as only a minor performance problem.

5.2 Accelerator Usage

After 4.2.5 the implementation of `Cache` provided leaves it up to the user to choose a caching and copy method policy which is accomplished through submitting function pointers at initialization of the `Cache`. In 2.5 we configured our system to have separate NUMA-Node (Node)s for accessing HBM which are assigned a Node-ID by adding eight to the Nodes ID of the Node that physically contains the HBM. Therefore, given Node 3 accesses some datum, the most efficient placement for the copy would be on Node $3 + 8 == 11$. As the `Cache` is intended for multithreaded usage, conserving accelerator resources is important, so that concurrent cache requests complete quickly. To get high per-copy performance while maintaining low usage, the smart-copy method is selected as described in 3.2.3 for larger copies, while small copies will be handled exclusively by the current node. This distinction is made due to the overhead of assigning the current thread to the selected nodes, which is required as Intel DML assigns submissions only to the DSA engine present on the node of the calling thread [6, Section "NUMA support"]. No testing has taken place to evaluate this overhead and determine the most effective threshold.

6 Evaluation

...evaluation ...

write this chapter

7 Conclusion And Outlook

write introductory paragraph

7.1 Conclusions

write this section

7.2 Future Work

write this section

- analyse whether prefetching yields better performance than hbm caching mode [2]
- evaluate impact of lock contention and atomics on performance
- provide optimized use case specific versions with less locking
- extend the cache implementation use cases where data is not static

Glossary

A

ATC

... desc ...

B

BAR

... desc ...

D

DMR

... desc ...

DSA

... desc ...

DWQ

... desc ...

E

ENQCMD

... desc ...

H

HBM

... desc ...

I

Intel DML

... desc ...

IOMMU

... desc ...

M

MOVDIR64B

... desc ...

N

Node

... desc ...

P

PASID

... desc ...

Q

QdP

... desc ...

S

SWQ

... desc ...

W

WQ

... desc ...

Bibliography

[1] H. Jun, J. Cho, K. Lee, H.-Y. Son, K. Kim, H. Jin and K. Kim, 'Hbm (high bandwidth memory) dram technology and architecture', in *2017 IEEE International Memory Workshop (IMW)*, 2017, pp. 1–4. DOI: 10.1109/IMW.2017.7939084.

[2] Intel. 'Intel® Xeon® CPU Max Series Product Brief'. (6th Jan. 2023), [Online]. Available: https://www.intel.com/content/www/us/en/content-details/765259/intel-xeon-cpu-max-series-product-brief.html (visited on 18th Jan. 2024).

[3] Intel. 'Intel® Data Streaming Accelerator Architecture Specification'. (16th Sep. 2022), [Online]. Available: https://www.intel.com/content/www/us/en/content-details/671116/intel-data-streaming-accelerator-architecture-specification.html (visited on 15th Nov. 2023).

[4] Intel. 'New Intel® Xeon® Platform Includes Built-In Accelerators for Encryption, Compression, and Data Movement'. (Dec. 2022), [Online]. Available: https://www.intel.com/content/dam/www/central-libraries/us/en/documents/2022-12/storage-engines-4th-gen-xeon-brief.pdf (visited on 15th Nov. 2023).

[5] R. K. et al. 'A Quantitative Analysis and Guideline of Data Streaming Accelerator in Intel® 4th Gen Xeon® Scalable Processors'. (May 2023), [Online]. Available: https://arxiv.org/pdf/2305.02480.pdf (visited on 7th Jan. 2024).

[6] Intel, *Intel Data Mover Library Documentation*, https://intel.github.io/DML/documentation/api_docs/high_level_api.html. (visited on 7th Jan. 2024).

[7] Intel, *Intel IDXD User Space Application*, https://github.com/intel/idxd-config. (visited on 7th Jan. 2024).

[8] Debian, *Debian manpage 3 for libnuma-dev.* [Online]. Available: https://manpages.debian.org/bookworm/libnuma-dev/numa.3.en.html (visited on 21st Jan. 2024).

[9] J. C. Sam Kuo. 'Implementing High Bandwidth Memory and Intel Xeon Processors Max Series on Lenovo ThinkSystem Servers'. (26th Jun. 2023), [Online]. Available: https://lenovopress.lenovo.com/lp1738.pdf (visited on 21st Jan. 2024).

[10] Intel. 'Intel® Data Streaming Accelerator User Guide'. (11th Jan. 2023), [Online]. Available: https://www.intel.com/content/www/us/en/content-details/759709/intel-data-streaming-accelerator-user-guide.html (visited on 15th Nov. 2023).

[11] A. Huang. 'Enabling Intel Data Streaming Accelerator on Lenovo ThinkSystem Servers'. (), [Online]. Available: https://lenovopress.lenovo.com/lp1582.pdf (visited on 18th Apr. 2022).

[12] A. C. Fürst, *Accompanying Thesis Repository*. [Online]. Available: https://git.constantin-fuerst.com/constantin/bachelor-thesis.

[13] Intel. ‘Intel® Xeon® CPU Max Series Configuration and Tuning Guide’. (Aug. 2023), [Online]. Available: https://cdrdv2-public.intel.com/787743/354227-intel-xeon-cpu-max-series-configuration-and-tuning-guide-rev3.pdf (visited on 21st Jan. 2024).

[14] T. Ku and N. Jung, ‘Implementation of Lock-Free shared_ptr and weak_ptr for C++11 multi-thread programming’, in *Journal of Korea Game Society*, vol. 21, 28th Feb. 2021, pp. 55–65. DOI: 10.7583/jkgs.2021.21.1.55..

[15] Unknown. ‘CPP Reference List of Atomic Operations’. (), [Online]. Available: https://en.cppreference.com/w/cpp/thread#Atomic_operations (visited on 18th Jan. 2024).

[16] Unknown. ‘CPP Reference Entry on std::atomic<T>::wait’. (), [Online]. Available: https://en.cppreference.com/w/cpp/atomic/atomic/wait (visited on 18th Jan. 2024).

[17] T. Rodgers. ‘Implementing C++20 atomic waiting in libstdc++’. (6th Dec. 2022), [Online]. Available: https://developers.redhat.com/articles/2022/12/06/implementing-c20-atomic-waiting-libstdc#how_can_we_implement_atomic_waiting_ (visited on 18th Jan. 2024).

[18] Unknown. ‘CPP Reference Entry on std::atomic<T>::notify_one’. (), [Online]. Available: https://en.cppreference.com/w/cpp/atomic/atomic/notify_one (visited on 18th Jan. 2024).

[19] Unknown. ‘CPP Reference Entry on std::atomic<T>::exchange’. (), [Online]. Available: https://en.cppreference.com/w/cpp/atomic/atomic/exchange (visited on 18th Jan. 2024).

[20] Unknown. ‘CPP Reference Entry on std::atomic<T>::notify_all’. (), [Online]. Available: https://en.cppreference.com/w/cpp/atomic/atomic/notify_all (visited on 18th Jan. 2024).

[21] AMD. ‘AMD64 Programmer’s Manual Volume 2: System Programming’. (Dec. 2016), [Online]. Available: https://support.amd.com/TechDocs/24593.pdf (visited on 18th Jan. 2024).

[22] Intel. ‘Intel 64 and IA-32 Architectures Software Developer’s Manual Volume 3A: System Programming Guide, Part 1’. (Dec. 2016), [Online]. Available: https://support.amd.com/TechDocs/24593.pdf (visited on 18th Jan. 2024).

[23] H. Schweizer, M. Besta and T. Hoefler, ‘Evaluating the Cost of Atomic Operations on Modern Architectures’, in *2015 International Conference on Parallel Architecture and Compilation (PACT)*, 2015, pp. 445–456. DOI: 10.1109/PACT.2015.24.