Generating Custom Code for Efficient Query Execution on Heterogeneous Processors

Sebastian Breß · Bastian Köcher · Henning Funke · Tilmann Rabl · Volker Markl

Abstract Processor manufacturers build increasingly specialized processors to mitigate the effects of the power wall to deliver improved performance. Currently, database engines are manually optimized for each processor: A costly and error prone process.

In this paper, we propose concepts to enable the database engine to perform per-processor optimization automatically. Our core idea is to create variants of generated code and to learn a fast variant for each processor. We create variants by modifying parallelization strategies, specializing data structures, and applying different code transformations.

Our experimental results show that the performance of variants may diverge up to two orders of magnitude. Therefore, we need to generate custom code for each processor to achieve peak performance. We show that our approach finds a fast custom variant for multi-core CPUs, GPUs, and MICs.

Sebastian Breß
DFKI GmbH and TU Berlin E-mail: sebastian.bress@dfki.de

Bastian Köcher
TU Berlin E-mail: bastian.koecher@tu-berlin.de

Henning Funke
TU Dortmund E-mail: henning.funke@tu-dortmund.de

Tilmann Rabl
TU Berlin and DFKI GmbH E-mail: rabl@tu-berlin.de

Volker Markl
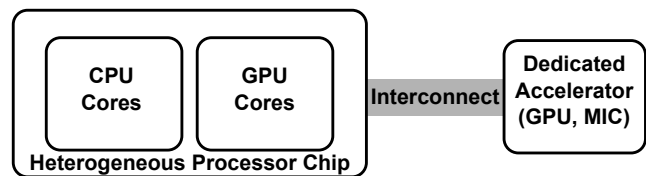TU Berlin and DFKI GmbH E-mail: volker.markl@tu-berlin.de

Fig. 1 Modern processors expose heterogeneity in the form of heterogeneous cores located on the same processor chip or specialized accelerator cards.

1 Introduction

The design of modern processors is primarily limited by a fixed energy budget per chip. This *power wall* forces vendors to explore new processor designs to stay in the energy budget [4,11]. One trend integrates heterogeneous processor cores on the same chip, e.g., combining CPU and GPU cores on the same chip as in Intel's processors with HD Graphics and AMD's Accelerated Processing Units (APUs). Another trend is *specialization*: processors are optimized for certain tasks, which already became commodity in the form of *Graphics Processing Units* (GPUs), *Multiple Integrated Cores* (MICs), or *Field-Programmable Gate Arrays* (FPGAs). These accelerators promise large performance gains because of their additional computational power and memory bandwidth. As a direct consequence of the power wall, current machines are built with a set of heterogeneous processors. Thus, from a processor design perspective, the *homogeneous many core age* ends [4,53]. The upcoming *heterogeneous many core age* forces database systems to embrace processor heterogeneity to achieve peak performance. We show such a heterogeneous processor system in Figure 1.

Parallel programming APIs such as OpenCL allow us to run single operators on a wide range of processors. However, we still need to customize the operator im-

plementation for peak performance [44]. In general, we need to select the most suitable parallelization strategies, data structures and code transformations for best performance.

Previous solutions either focused on generating highly efficient code for a single processor [33,51] or allowed database operators to run on multiple processors using the same operator code [20,56]. Code generation approaches suffered from a high compilation time, or were confined to a single processor by generating low-level machine code (e.g., LLVM [33]). Hardware-oblivious approaches suffer from limited performance portability [44]. As of now, we need to manually adapt the database system to every new processor (e.g., for Intels MIC architecture) to achieve peak performance.

Our long term goal is to enable database systems to automatically generate efficient code for any processor without *any a priori* hardware knowledge. To achieve this goal, we propose Hawk, a hardware-adaptive query compiler, which can generate variants of generated code. By executing different variants of a compiled query, Hawk can adapt to a wide range of different processors without any manual tuning. By compiling queries to OpenCL kernels, Hawk achieves low compilation times and can run queries on any OpenCL-capable processor.

In this paper, we make the following contributions:

1. We introduce *pipeline programs*, a new form of physical query plan. Pipeline programs store operations and implementation properties of a pipeline and are the basis for code generation. Furthermore, we demonstrate how we can systematically generate variants of pipeline programs (cf. Section 3).

2. We discuss the dimensions in which we can vary the generated code of *pipeline programs* (cf. Section 4, 5, 6) and show the impact of these variations on common processors and co-processors.

3. We compile all operators of a pipeline to kernel programs. These kernels can be compiled to a broad set of heterogeneous processors using parallel programming libraries such as OpenCL (cf. Section 7).

4. We present a learning strategy to automatically derive efficient variant configurations and incorporate them into a query optimizer (cf. Section 8).

5. We show the potential of a database system that rewrites its code until it runs efficiently on the underlying heterogeneous processor hardware (cf. Section 9).

2 Background

In this section, we discuss the background required for the remainder of the paper. Since we directly build
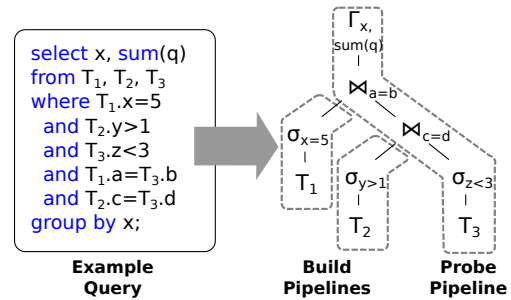

Fig. 2 Example for produce/consume style query compilation. The query plan is partitioned into three operator pipelines, two for building join hash tables, and one for probing both join hash tables.

on the produce/consume model of Neumann [33], we first provide an overview of query compilation using the produce/consume model. Second, we briefly discuss OpenCL, as it is the target language for our code generation.

2.1 Query Compilation

For our kernel compilation approach, we build on the produce/consume model for code generation by Neumann [33]. The goal of the produce/consume model is to partition a query into pipelines and to merge the operators belonging to the same pipeline into a single code fragment that implements the pipeline.

This code fragment iterates in a tight for-loop over all tuples of the input relation. Each tuple is pushed through all operators of the pipeline, before the next tuple is processed. This code generation achieves excellent data locality by keeping the tuple in the processor registers.

In the produce/consume model, each operator needs to provide a produce and a consume function. The *produce* function traverses the query plan top down from the root operator and creates a new pipeline for every pipeline breaking operator. When produce reaches a scan, the consume function of succeeding operators is called bottom up and generates the code for each operator in the current pipeline until a pipeline boundary is reached. Then, the code for the next pipeline is generated. Thus, the produce functions essentially partition the query plan into pipelines, whereas the consume functions fill the pipelines with operators and generate the code. In this paper, we refer to pipelines that were filled with operators by their consume functions as *operator pipelines*.

We illustrate the produce/consume style query compilation in Figure 2. We evaluate a query with two joins using hash joins. The hash tables are built on the left
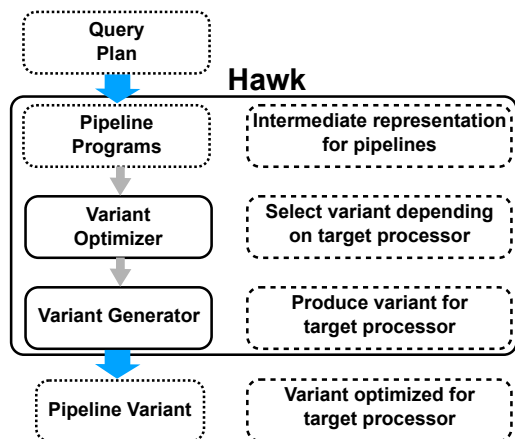
Fig. 3 Core concepts of Hawk and their role in the system.

sub plans. As building a hash table is a pipeline breaker, the produce/consume model creates two build pipelines on table T_1 and T_2 and one probe pipeline on table T_3.

2.2 Overview of OpenCL

The *Open Compute Language* (OpenCL) is a framework for massively parallel computing, which supports processors with different architectures to achieve functional portability [15]. Functional portability means that an OpenCL program developed for any OpenCL-capable processor will run on any other OpenCL-capable processor (e.g., a OpenCL program written for a GPU can also run on a CPU). OpenCL abstracts all processors as so-called devices. The CPU that executes the OpenCL API functions is called the host. All computations are expressed in special functions called *kernels*, which are then compiled just-in-time for all devices. The just-in-time compilation abilities make OpenCL especially interesting for our work, as it provides a native mechanism to compile generated code to a device. Furthermore, the functional portability allows us to run any variant of generated code on any device.

3 Hawk Compiler Architecture

In this section, we discuss how Hawk generates custom code for each processor by producing *variants* of operator pipelines (cf. Figure 3). As a core concept, Hawk represents operator pipelines of queries as *pipeline programs*. We use these programs as an intermediate representation that defines the semantics of a pipeline and serves as basis for target code generation.

Hawk produces variants along three major dimensions: execution strategy, data structure optimization,

and code transformations. Variations in these dimensions are implemented either by transforming or by reparameterizing pipeline programs. We can freely combine individual variations, this way spanning a large space of potential variants. The *variant optimizer* selects a variant configuration from this space. Then, the *variant generator* produces the code variant selected by the optimizer for the specified processor.

This architecture allows Hawk to create custom per-processor code for each pipeline in a query plan.

3.1 Pipeline Programs

Our goal is to create variants of a pipeline that can take into account various aspects of the underlying hardware. For this, we provide a precise definition of the semantics of an operator pipeline.

The produce/consume model fills a pipeline with operations by calling their consume functions during code generation. As such, the ordered sequence of operations added by the consume function defines the semantics of the pipeline. The calls to consume consist of *pipeline operations*, such as filtering, inserting a tuple into a hash table, probing a hash table, or aggregating tuples. Pipeline operations often have no direct equivalent in relational algebra, but are used to implement them. We briefly introduce the pipeline operations used in this work in Table 1.

Operations accept two categories of parameters:

1. **Regular parameters:** These parameters encode the semantics of the operation, such as the table scanned in LOOP or the filter predicate in FILTER. We format these parameters *italic*.

2. **Code generation modes:** These parameters define which code variant is generated by the operation, such as the hash table implementation used in HASH_PUT. We format these parameters **bold**.

We store the sequence of pipeline operations and their parameters as a *pipeline program*. We illustrate the process for the example query in Figure 2 and show the pipeline programs produced in Table 2. The query contains two joins, which forces the produce/consume model to create a new pipeline for each hash table build. The build pipelines iterate over their input tables (T_1 and T_2), apply their filters, insert the matching key into a hash table and materialize the result on the required attributes. The probe pipeline iterates over table T_3, applies its filter, probes the previously built hash tables, and performs the aggregation.

Pipeline Operation	Description
LOOP(T; **step**, **s**, **e**)	Iterate over input tuples of table T, using a loop increment of **step**, and a loop start index **s** and end index **e**
FILTER(F_σ; **m**, **o**)	Apply a filter condition F_σ using a branching mode **m** and an element access offset **o**
HASH_PUT(A; **h**, **p**)	Insert tuple in a hash table for attribute set A using hash table **h** with parameters **p**
HASH_PROBE(A; **h**, **p**, **m**, **o**)	Probe hash table for attribute set A using hash table **h** with parameters **p**, branching mode **m**, and element access offset **o**
CROSS_JOIN(T)	Loop over additional input table T and compute cross product of current tuple and T
ARITHMETIC(f; **o**)	Apply a computation $f : A \times B \to C$ of attributes A, B, C using element access offset **o**)
AGGREGATE(F; **m**, **o**)	Apply a non-grouping aggregation with aggregation expression $F = (f_1, f_2, \cdots, f_n)$, branching mode **m**, and element access offset **o**
HASH_AGGREGATE(G,F; **h**, **p**, **m**, **o**)	Apply an aggregation with grouping attributes G, aggregation expression $F = (f_1, f_2, \cdots, f_n)$, using hash table **h** with parameters **p**, branching mode **m**, and element access offset **o**
PROJECT(A; **m**, **o**)	Materialize tuples to output relation projecting attributes for attribute set A, branching mode **m**, and element access offset **o**

Table 1 Overview of pipeline operations in a pipeline program.

Build Pipeline 1	Build Pipeline 2	Probe Pipeline
LOOP(T_1, ..)	LOOP(T_2, ..)	LOOP(T_3, ..)
FILTER(x=5, ..)	FILTER(y>1, ..)	FILTER(z<3, ..)
HASH_PUT(a, ..)	HASH_PUT(b, ..)	HASH_PROBE(a=c, ..)
PROJECT(a, x, ..)	PROJECT(b, ..)	HASH_PROBE(b=d, ..)
		HASH_AGGREGATE(x, sum(q), ..)

Table 2 Pipeline programs created for example query plan from Figure 2. Each pipeline program belongs to one pipeline.

3.2 Pipeline Variants

We define a *variation* as a systematic modification of a pipeline program that changes the generated target code or run-time parameters. One example is the predication mode, where we can either use conditional expressions or software predication to evaluate a selection. We define a pipeline variant (in short *variant*) by the set of variations applied to a pipeline program. From an implementation perspective, the set of program parameters and the sequence of transformation steps determine the variant. Program parameters are global properties of the pipeline program, such as the number of threads used or the execution strategy. Many variations require a fine-grained adaption of pipeline operations, such as software predication, loop unrolling, and vectorization. These variants are created by one transformation pass per variation. Each pass modifies the pipeline program and, thus, changes the produced result code. Therefore, a sequence of transformation steps defines the code that Hawk generates.

3.3 Dimensions of Variant Generation

We now discuss how we capture hardware properties in a generic way. We differentiate between execution strategies, data structure optimization, and code transformations.

Execution Strategy. The execution strategy defines how a pipeline is executed. As we will discuss in

Section 4, different strategies are optimal for various processors and have a strong impact on performance. Thus, a hardware-adaptive query compiler needs to cope with different execution strategies. For example, on CPUs, synchronization overhead is still cheap compared to co-processors: we have usually only tens of threads, so a single pass over the data is most efficient. However, on processors similar to GPUs or MICs (co-processors), it is much faster to make multiple passes over the data to avoid synchronization cost.

Data Structure Optimization. Depending on data and query characteristics, data structures with a certain parametrization are optimal. One critical case are hash tables, where we can choose between different hashing techniques [43] and specialized hash tables for a certain query [47]. Thus, an efficient query compiler needs to be able to exchange the data structures used in a query to generate efficient code.

Code Transformations. Optimizing code for certain processors usually involves many low-level code transformations. For example, we need to decide on the optimal memory access pattern or the predication mode. In general, we do not achieve the best performance by just applying all available optimizations. Therefore, a hardware-adaptive query compiler must be flexible enough to apply a certain subset of code transformations to the generated code.

In the following sections, we discuss how we can apply these variations to pipeline programs. Furthermore, describe how we keep each variation orthogonal to other variations. For example, the execution strategy should not depend on the hash tables or memory access pattern used.

3.4 Pipeline Variant Generation

We generate variants of a pipeline program in two steps: transformation and code generation. At first, we define which variant of the pipeline is compiled, because the
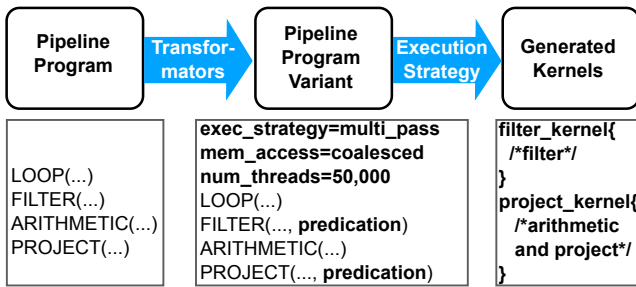
Fig. 4 Pipeline Variant Generation.

Fig. 5 Coarse-grained parallelism. One serial kernel is generated per pipeline. We parallelize by executing one kernel per physical core on different blocks of the input.

Listing 1 Projection Query 1.

```sql
select lo_linenumber, lo_quantity,
    lo_revenue
from lineorder where lo_quantity<25;
```

variant determines which transformation passes need to be executed.

In the *transformation step*, we execute a sequence of transformation passes. These passes can modify the pipeline program in two ways. First, they can set the global properties of the pipeline program (e.g., the memory access pattern). Second, they can (re-)configure a single pipeline operation (e.g., set the hash table in HASH_PUT).

In the *code generation step*, we create an interpreter for the selected execution-strategy. The interpreter traverses the pipeline program and calls the code generator for each pipeline operation. Depending on the number of kernels used by an execution strategy, the generated code is injected in one or more kernels (cf. Figure 16). We illustrate the process in Figure 4. We discuss target code generation in detail in Section 7.

Note that this code generation algorithm allows us to freely combine variations in the same pipeline program. For example, it is possible to generate a variant that uses coalesced memory access, software predication, with a fine-grained multi-pass execution strategy. Thus, it is a very flexible and powerful mechanism, which keeps the possible variations orthogonal to each other.

4 Execution Strategies

Efficient code generation for heterogeneous processors needs to trade-off two basic design dimensions: the degree of parallelism and synchronization overhead. Coarse-grained parallelism is usually sufficient on CPUs (e.g., spawning one thread per processor core). In contrast, co-processors require fine-grained parallelism (e.g., using all available SIMD lanes in a GPUs streaming multi processor and having enough thread blocks to hide memory latencies). Fined-grained parallelism requires ten thousand and more threads. Thus, synchronization among threads is very expensive compared to thread synchronization among tens of threads for coarse-grained parallelism.
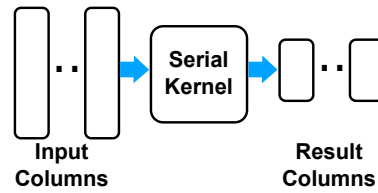
Algorithms optimized for co-processors avoid synchronization by not writing the result directly. Instead, the algorithms first compute unique write positions for each result tuple. Then, they repeat the computation and write the result tuples in parallel without any synchronization. Thus, these algorithms need to perform multiple passes over the data [16]. Due to the high memory bandwidth of co-processors, these multi-pass algorithms still achieve excellent performance. However, on CPUs, a multi-pass strategy introduces overhead. This is because a single pass over the data is typically more efficient, as the synchronization costs are moderate.

Depending on the type of operator pipeline, we generate different kernels. We differentiate between projection pipelines and aggregation pipelines.

4.1 Projection Pipelines

A projection pipeline is an operator pipeline in a query plan that does not perform aggregations. Thus, it projects matching tuples (filters and hash probes) in an output buffer. We show a simple query that creates a single projection pipeline in Listing 1. It consists of one filter predicate and projects three attributes.

Coarse-Grained Parallelism. On CPUs, it is common to generate a single for-loop per pipeline. This loop processes all input tuples and writes result tuples to the output buffers (cf. Figure 5). The coarse-grained strategy parallelizes query processing by concurrently executing the same pipeline on different chunks of the input relation [28].

Fine-Grained Parallelism. On processors with many light-weight cores (e.g., GPUs or MICs), the coarse-grained parallelization cannot utilize all cores. In this
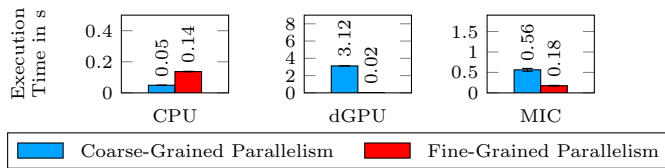
Fig. 6 Impact of different execution strategies on Projection Query 2 (cf. Listing 2).
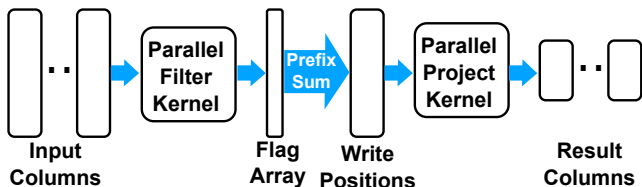
```
select lo_linenumber, lo_quantity,
    lo_revenue
from lineorder where lo_quantity<25 and
    lo_discount<=3 and lo_discount>=1 and
    lo_revenue>4900000;
```


Fig. 7 Fine-grained parallelism. We generate two kernels that are executed massively parallel.
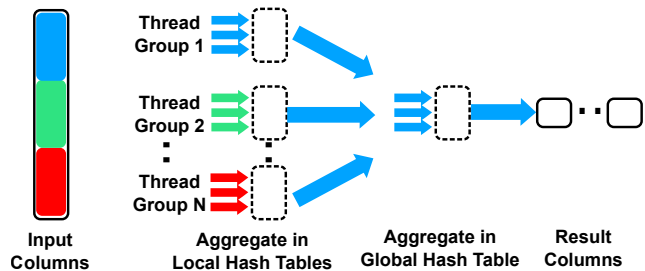

Fig. 8 Local hash table execution strategy for aggregation pipelines. For each of the N hash tables, M threads perform the aggregation.

4.2 Aggregation Pipelines

An aggregation pipeline is a pipeline where the last operator is an aggregation operator. Here, we materialize the result in a hash table and, therefore, we do not need to compute write positions in an output buffer.

4.2.1 Execution Strategies

Depending on the number of result groups, we use different aggregation strategies.

Local Hash Table Aggregation. If we expect few result groups, we perform the aggregation in two steps. First, we pre-aggregate the result in parallel in multiple local hash tables. Second, we merge the local hash tables into a global result hash table. We call this *local aggregation* and illustrate the principle in Figure 8. For each of the N hash tables, M threads perform the aggregation. The number of hash tables and threads per hash table are thus important tuning parameters (cf. Section 9). We synchronize concurrent operations on the aggregates using OpenCL's atomics.

Global Hash Table Aggregation. If we expect many result groups, we aggregate into a single global hash table. In this case, synchronization overhead is small and cost for merging large partial results high. We refer to this as *global aggregation*, which is a special case of local aggregation with a single local hash table. Thus, we only need to tune the number of threads per hash table.

case, we need fine-grained parallelism. We illustrate this trade-off in Figure 6, where we execute Projection Query 2 (cf. Listing 2) with the coarse-grained and the fine-grained strategy on different processors. We describe our detailed experimental setup in Section 9.1. The coarse-grained strategy outperforms the fine-grained strategy on CPUs by a factor of 2.8. The fine-grained strategy outperforms the coarse-grained strategy on a GPU by a factor of 148 and a MIC by a factor of 3.19.

Algorithms that use fine-grained parallelism avoid latching at all cost and are typically multi-pass strategies, consisting of three phases [16]. In the first phase, the operator is executed and all matching tuples are marked in a flag array. In the second phase, per-thread write positions are computed using an exclusive prefix sum. Finally, the operator is executed again, but this time, the threads can lookup globally unique write positions and can write their result.

We generalize this three-step processing technique to operator pipelines as follows. We generate two kernels, a filter and a projection kernel. In the first step, the *filter kernel* performs all operations that reduce the number of result tuples. These are essentially filter and hash probes (e.g., to conduct joins). All matching tuples are marked in a flag array. The second step computes the write positions for each thread by performing a prefix sum on the flag array. In the third step, the *projection kernel* repeats the hash probes to obtain the payload of matching join tuples. The projection kernel also performs arithmetic instructions and writes the result to the computed write positions. We illustrate the algorithm in Figure 7.
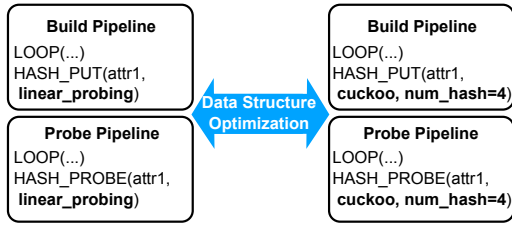
Fig. 9 Data Structure Optimization on the example of exchanging hash table implementations.

4.2.2 Supporting different degrees of Parallelism

We implement *coarse-grained parallelism* by using the local hash table aggregation with one thread per hash table. Furthermore, we set the number of hash tables to the number of OpenCL compute units (e.g., the number of CPU cores).

Each thread group is responsible for one hash table. Thus, if we increase the number of threads per thread group, we achieve *fine-grained parallelism*.

5 Data Structure Optimization

Besides a well-selected execution strategy, high performance implementations require optimized data structures. A prominent example in a database context are hash tables, where different implementations are optimal depending on data and query characteristics [43]. We can also improve cache efficiency by removing unnecessary payloads, when we specialize hash tables for a certain query [47] (e.g., hash tables used for aggregations). A query compiler allows us to directly include these optimizations on a per-query basis.

5.1 Support of Different Hash Tables

We configure each HASH_PUT and HASH_PROBE operation with a hash table and its parameters, as we illustrate in Figure 9. Here, we exchange the linear probing hash table with a Cuckoo hash table. Note that we can also change the parametrization of a hash table (e.g., we can set the number of hash functions of Cuckoo hashing).

Build and probe pipeline operations need to work with the same hash table (and same parametrization). This introduces a dependency between pipeline programs. Thus, the query processor needs to ensure that corresponding HASH_PUT and HASH_PROBE operations use the same data structure.
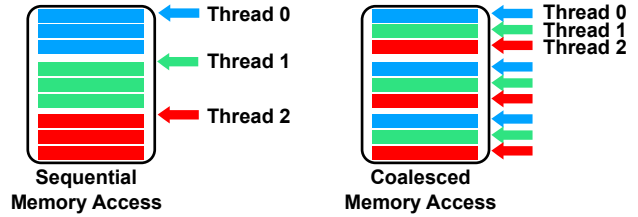

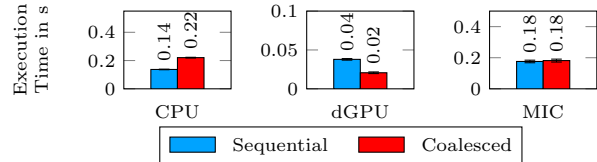**Fig. 10** Visualizing different memory access strategies.

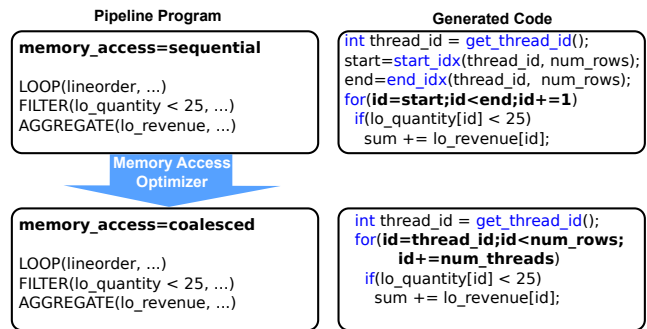**Fig. 11** Impact of memory access pattern on Projection Query 2 (cf. Listing 2).

Fig. 12 Effect of memory access pattern on generated code.

6 Code Transformations

We now discuss how a query compiler can capture different ways to exploit the hardware at the level of traditional code transformations (e.g., the memory access pattern).

6.1 Transformations

Adjusting the memory access pattern. Different processors prefer different ways of accessing the memory. In sequential access, each thread processes a continuous chunk of tuples, very similar to horizontal range partitioning. In coalesced memory access, every thread reads a neighbored location relative to other threads. We illustrate the principle of sequential and coalesced memory access in Figure 10.

We show the performance impact of the memory access pattern in Figure 11. On a CPU, sequential access outperforms coalesced memory access by a factor of 1.6. On a GPU, coalesced memory access outperforms sequential memory access by a factor of 1.8. In
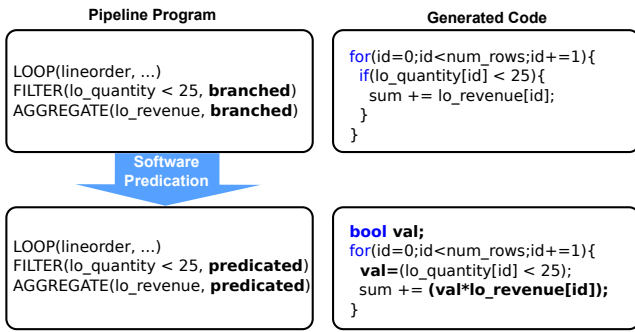
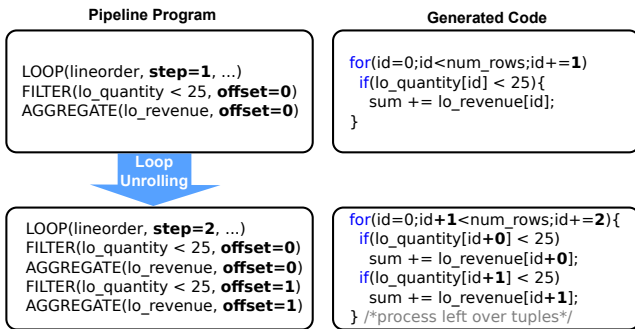**Fig. 13** Applying software predication transformation to a pipeline program.

Fig. 14 Applying loop unrolling transformation to a pipeline program.

Fig. 15 Query compilation times of a simple Projection Query (cf. Listing 2). OpenCL kernel compilation times for CPUs are in the same order of magnitude as HyPer's LLVM IR code generation.

this measurement, we see no significant difference for the MIC processor.

We rewrite the memory access pattern in a pipeline program by setting the memory access property. We show the impact on the generated code in Figure 12.

Applying software predication. Software predication is a common technique to avoid branch misprediction penalties. To support predication, each pipeline operation has a flag that determines whether code with branching (if statements) or with predication should be generated. In the predicated mode, the result of predicate evaluations is stored in a result value. This value is either added to the variable storing the result size (projection pipeline) or multiplied to the input values before an aggregation (aggregation pipeline). We illustrate the principle in Figure 13, where we switch a simple aggregation pipeline from branched to predication mode. Note that all pipeline operations have to be in the same mode. Otherwise, the result becomes incorrect. Thus, either the complete pipeline program uses predicated execution or not.

Other optimizations. We can also apply more complex code transformations, such as loop unrolling or vectorization. We exemplary show how loop unrolling can be supported by pipeline programs in Figure 14. Loop unrolling affects the original pipeline program be-
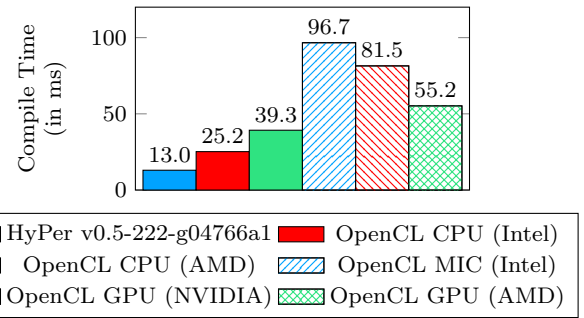
yond the choice of code generation parameters. However, loop unrolling does not limit the combinations with other variations.

The key point is that a pipeline program, which physically represents a pipeline, is a highly flexible representation. It stores low-level code transformations that are hard to represent in a physical query plan.

7 Target Code Generation

In this section, we discuss the target code generation of Hawk. We reason why we use OpenCL as target language and discuss how we generate kernels by fragment generation and assembly. Furthermore, we discuss how one can extend the code generator by new data structures and algorithms. Finally, we discuss implementation details of Hawk.

7.1 Target Code: OpenCL Kernel

The drawback of generating high-level code is usually high compilation time [26,33]. By compiling pipeline programs to OpenCL kernels, Hawk benefits from the JIT compilation capabilities and the performance portability of OpenCL. The latter allows Hawk to run any variant on any OpenCL-capable processor.

In Figure 15, we show query compilation times for a simple query (cf. Listing 2) for compiling OpenCL kernels for an Intel CPU, an AMD CPU, an Intel MIC, a NVIDIA GPU, and an AMD GPU. As reference, we also show the compilation time of HyPer [23] (v0.5-222-g04766a1), a state-of-the-art system for query compilation. Compiling OpenCL kernels for CPUs is in the same order of magnitude (slower by a factor of 2 to 3 for Intel and AMD OpenCL SDKs) as the LLVM IR query compilation used by HyPer [33]. Furthermore, we observe that compilation for GPUs and MICs is up
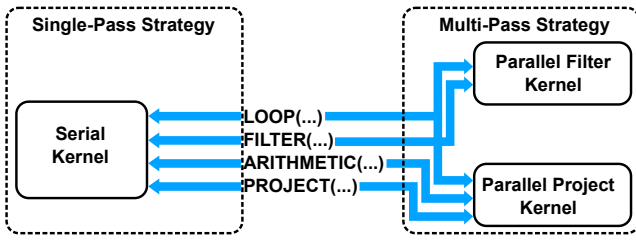
Fig. 16 Supporting multiple execution strategies. Each strategy is designed to act as an interpreter for a pipeline program, which routes generated code to the appropriate kernels.

to a factor of 4.3 to 7.4 slower compared to LLVM IR query compilation. The compilation times are consistently below 100ms. Thus, we conclude that query compilation using OpenCL is sufficiently efficient for compiling database queries to support interactive querying on GPUs and MICs. Note that the OpenCL compilation times can be further reduced. We can disable certain optimization passes, trading off runtime and optimization time, similar to optimization levels in some commercial database engines.

7.2 Fragment Generation and Assembly

We now discuss how we can generate code for projection and aggregation pipelines from pipeline programs. The code generation follows a two step scheme: *fragment generation*, followed by *fragment assembly*.

7.2.1 Fragment Generation

A code fragment (in short *fragment*) consists of six segments: host variable declarations, host initialization code, host cleanup code, kernel variable declarations, kernel code top, and kernel code bottom. These fine-grained separations allow us to route fragments into different kernels.

Each pipeline operation produces a fragment that implements its semantic. We retrieve the fragment for each pipeline operation to create all fragments.

Each operation can generate code for any part in the target source code (e.g., body of the for-loop, declarations, or cleanup operations). Furthermore, the fragment produced by a pipeline operation depends on the code generation modes: these modes allow us to adapt the fragment by re-parameterizing the pipeline operations or global properties of the pipeline program. Using this code generation scheme, it is straightforward to create systematic variants of a pipeline program to adapt to the underlying hardware.

7.2.2 Fragment Assembly

We combine fragments by assembling them into a single fragment. Note that this *fragment assembly* is essentially a string concatenation of code segments.

Our guiding idea is as follows. We provide an *interpreter* for pipeline-programs for each execution strategy. Each interpreter knows how many kernels are required for the strategy. The interpreter assigns the fragments, depending on the pipeline operation, to one or more kernels.

We illustrate this process in Figure 16. For the single-pass strategy, all fragments belong to the same kernel. The multi-pass strategy routes fragments from different pipeline operations to different kernels. Thus, a fragment can be part of multiple kernels, such as LOOP or HASH_PROBE.

For each kernel used by the execution strategy, the interpreter assembles all fragments assigned to the kernel to a result fragment. We create the final kernel from this result fragment. Note that Hawk's code generator is conceptionally not limited to OpenCL kernels, but could also produce code for frameworks such as CUDA.

Since we implement execution strategies as interpreters, we can apply different strategies to pipeline programs. Our design keeps the execution strategies orthogonal to any other variation on the pipeline program.

7.3 Example: Fragment Generation and Assembly

We now present an example that illustrates the code generation process. Consider the query *select b from t where a<5*, which will result in a pipeline program with three operations: LOOP, FILTER, and PROJECT. We show the generated fragments of the pipeline operations in Figure 17. The generated fragments can add code to two parts of the kernel: the variable declaration and initialization code block, and the for-loop. Code can be inserted into a for-loop at two positions: at the top position we insert the actual code; at the bottom position we insert closing brackets and perform operations (e.g., increasing counters) after an iteration. Generated code of succeeding operations is nested inside the brackets of previous operations. For example, the final projection is nested in the generated code of the filter operation.

7.4 Fragment Generation and Assembly Algorithms

We now discuss algorithms for fragment generation. We show pseudo code for each algorithm and highlight gen-
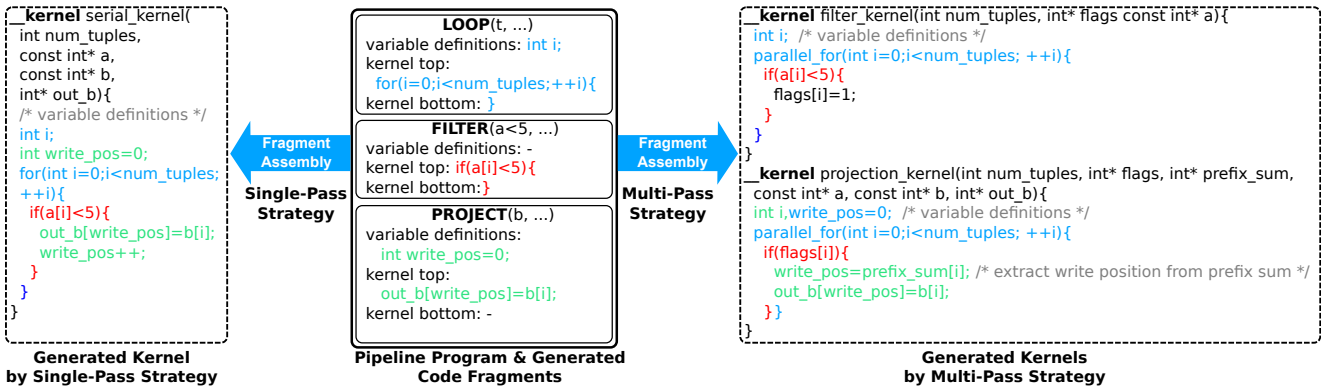
Fig. 17 Example for fragment generation and fragment assembly: Each pipeline operation generates fragments, which are then assembled to kernels. The single-pass strategy generates one kernel that includes all operations from the fragments. The multi-pass strategy generates a filter and a projection kernel which include different fragments.

Listing 3 Loop Fragment Generation:
LOOP(table, memory_access_pattern).

```
<thr_id = get_thread_id()>
if(memory_access_pattern==SEQUENTIAL){
  <start=start_idx(thr_id,num_rows)>
  <end=end_idx(thr_id,num_rows)>
  <for(id=start;id<end;id+=1)>
}else if(memory_access_pattern==COALESCED)
  {
  <for(id=thr_id;id<num_rows;id
  +=num_threads)>
}
```

Listing 4 Filter Fragment Generation:
FILTER(condition, predication_mode).

```
if(predication_mode==BRANCHED){
  <if(condition)>
}else if(predication_mode==PREDICATED){
  <result_increment=(condition)>
}
```

Generated Code: `<code>`

erated code by surrounding it with angle brackets and by coloring the background (`<generated code>`).

The LOOP operation generates code that iterates over every input tuple of a table in parallel. We can either iterate sequentially or interleaved over the tuples, which leads to sequential or coalesced memory access (cf. Listing 3). In case of sequential access, we compute the start and end offset of the partition that each thread processes. In case of coalesced access, each thread starts the iteration on its unique thread identifier and advances by adding the number of threads to the loop variable *id*.

The FILTER operation generates code that evaluates a selection predicate. It either generates an if-statement (no predication) or stores the result of the predicate evaluation in the variable *result_increment* (predication), as we illustrate in Listing 4.

The PROJECT operation generates code that copies the values of each projected attribute and writes them to the write position *write_pos* in the projection attribute's output array (cf. Listing 5). The generated code depends on the predication mode. If predication is disabled, we know the tuple passed all previous filters. Thus, we increment the write position after writing the tuple into the output buffer. If predication is enabled, we always write the result tuple but add the variable *result_increment* to *write_pos*. If the tuple passed all previous filters, *result_increment* is one and the write position is advanced by one row. In case the tuple did not match all filters, *result_increment* is zero and the write position is not changed, which discards the current tuple.

The HASH_PUT and HASH_PROBE operations generate code that insert/lookup tuples into/from a certain hash table (cf. Listing 6). For a linear probing hash table, we generate code that uses a single hash function. A Cuckoo hash table has a variable number of hash functions, which directly affects the generated code. We omit the detailed code for the sake of brevity.

The AGGREGATE operation generates code that computes the aggregates. The generated code depends on the predication mode. If predication is disabled, we evaluate the aggregate expression without any further modifications. In case of enabled predication, we need to take special care to not include a filtered out tuple in the aggregation. Therefore, we need to ensure that the aggregate is not changed in case the variable *result_increment* is zero. For example, for the count or sum aggregation functions, we multiply the tuple value with the *result_increment* before applying the aggregation function (cf. Listing 7). This way, the aggregate stays unchanged if and only if *result_increment* is zero.

Listing 5 Project Fragment Generation: PROJECT(proj_attributes, predication_mode).

```
<declare variable write_pos=0>
for(attribute in proj_attributes){
  <copy value of attribute to result
  column at position write_pos>
}
if(predication_mode==BRANCHED){
  <write_pos++>
}else if(predication_mode==PREDICATED){
  <write_pos+=result_increment>
}
```

Listing 6 Join Fragment Generation: HASH_BUILD(attr, hash_table) and HASH_PROBE(attr, hash_table).

```
if(hash_table==CUCKOO){
  <insert/lookup attr in cuckoo HT>
}else if(hash_table==LINEAR_PROBING){
  <insert/lookup attr in lin. probing HT>
}
```

Listing 7 Aggregate Fragment Generation: AGGREGATE(attr, SUM, predication_mode).

```
<declare variable aggregate=0>
if(predication_mode==BRANCHED){
  <aggregate+=attr[id]>
}else if(predication_mode==PREDICATED){
  <aggregate+=(attr[id]*result_increment)>
}
```

Generated Code: `<code>`

7.5 Extending the Code Generator

We now discuss how we can extend the code generator to support new data structures and algorithms.

7.5.1 Support for Different Data Structures

For each hash table supported, the query compiler requires the code templates for initializing, accessing, and modifying a data structure. These code templates may also be generated at run-time (e.g., to adapt the hash function depending on the data properties). The pipeline operations encapsulate the code generation for different data structures. For example, the HASH_PUT operation implicitly generates code for the hash table specified in its parameter. This variability allows us to select different hash table implementations depending on certain data characteristics [43]. For example, to add a robin hood hash table, we need to extend the existing HASH_PUT and HASH_PROBE operations by the respective code templates. Finally, we need to introduce a pipeline operation (e.g., an index scan).

7.5.2 Support for Different Algorithms

For each new algorithm, we extend the code generator by the data structures the algorithm uses. Then, the algorithm needs to be registered to the variant generator. We either add a new pipeline operation or include the algorithm in an existing pipeline operation. Finally, we provide the respective code templates.

For example, we can extend Hawk with the scan of Zhou and Ross [57], which uses SIMD instructions to check the predicate of multiple tuples at once. To support this SIMD scan in Hawk, we need to add a new code generation mode to the FILTER operation. The mode parameter allows us to select either the scalar or the SIMD code template. Furthermore, the code template for the SIMD scan has to be added to the FILTER operation. The same procedure applies for SIMD support for other pipeline operations supported by Hawk.

7.6 Hawk Implementation Details

We implemented Hawk as a prototype that targets main-memory database engines that store data in a column-major format. We show the viability of our approaches on the example of CoGaDB [5,6], as it fulfills our requirements and resulted in the smallest integration effort for us. Note that we can apply our concepts to any other system having an in-memory column store, including commercial systems such as SAP HANA [12], DB2 BLU [40], or the Apollo engine of SQL Server [27]. The main changes to the database engine consists of the extension of the query plan interface by the produce/consume code generation along with our proposed approach for variant generation. Furthermore, the execution engine has to be replaced by a run-time for the compiled queries.

Hawk supports all pipeline operations discussed in Table 1, which allows for producing code for selections, projections, joins, and aggregations. Aggregations are currently limited to distributive and algebraic aggregation functions (e.g., holistic aggregation functions such as the median are currently not supported).

8 Optimizing Pipeline Variants

Hawk can generate a large number of variants to adapt code to various processors. Consequently, we face a large optimization space that is the cross product of all values of all variation dimensions. Exploring the search space is very expensive for two reasons. First, we pay query compilation cost for each generated variant. Second, the execution time of variants may be significantly

slower than the optimal variant. It is especially problematic when we explore variants that are very slow on a certain processor (e.g., a serial implementation on a GPU).

In this section, we discuss how we can automatically find a fast-performing variant configuration for each processor for a given query workload.

Algorithm 1 Learning an efficient variant configuration for a processor.

Input: dimensions of variations: $D = \{D_1, \cdots, D_n\}$
Input: workload of k queries: $W = \{Q_1, \cdots, Q_k\}$
Output: variant configuration v
1: $v = (v_1, \cdots, v_n) \in D_1 \times \cdots \times D_n$
2: **for** $(iter = 0; iter < q; iter + +)$ **do**
3:     last_variant=v
4:     **for** $D_i \in D$ **do**
5:       execution_time=∞
6:       best_dimension_value=$\emptyset$
7:       **for** $d \in D_i$ **do**
8:         $v' = v$;
9:         $v'_i = d$;
10:        execution_time$'$ = executeQueries(W, v');
11:        **if** execution_time$'$¡execution_time **then**
12:          execution_time=execution_time$'$;
13:          best_dimension_value=d;
14:        **end if**
15:       **end for**
16:       /* Update configuration */
17:       v_i =best_dimension_value;
18:     **end for**
19:     **if** v ==last_variant **then**
20:       **return** v;
21:     **end if**
22: **end for**
23: **return** v

8.1 Navigating the Optimization Space

The key idea is that we explore the search space for a processor offline by executing a workload of representative test queries. We compile different variants of each query in the workload, and systematically explore which variations work best on a particular processor. We use the same strategy as we would in a structured experiment. We present the strategy in Algorithm 1. Initially, we have no knowledge about the performance behavior of the processor. We start from a base configuration (Line 1), which we initialize with the first parameter value in each variant dimension. We change one parameter at a time (Line 4–10), and select the parameter value with the best performance (Line 11–14). We perform this step for every variant dimension (e.g., execution strategy or memory access pattern). The best parameter values are stored in the variant configuration

(Line 16–17). A variant configuration is the abstract representation of a variant in the search space.

Note that different variations may influence each other. This means that a previously optimal parameter value of a variation may be sub-optimal in the new configuration. To make sure that our algorithm finds a fast performing variant, we repeat the core of the algorithm (Line 4–18) iteratively. The algorithm terminates in case we have not found any faster variant configuration (Line 3, 19–21) or reach a maximum number of iterations q (Line 2).

8.2 Reducing Variant Optimization Time

As the variant exploration requires us to execute the variants, very slow variants increase exploration time significantly. We can reduce the exploration time by applying *early termination* and *feature ordering.*

Early Termination: Our learning strategy allows us to systematically gain knowledge over the complete variant space. We can terminate the search early, when we reach a local optimum during an iteration. This *early termination* saves additional exploration time, but we may not reach the global optimum.

Feature Ordering: We can further optimize the search of the variant space when we take into account which variations typically have the most impact on performance. In this case, we explore the parameter values of the most critical variations first to find an efficiently performing variant faster. These variations are: the execution strategy, the number of threads, and the memory access pattern.

8.3 Building a Heuristic Query Optimizer

We now discuss how we can build a heuristic optimizer using the variant configurations learned. We learn variant configurations for a representative query workload. Thus, the resulting variant configuration is a *heuristic* that performs well for a workload. While the heuristic delivers good performance for the given queries, it may not be optimal, as query-dependent parameters can influence the optimal variant.

To avoid high overhead during query processing, we execute the learning algorithm before query processing. We use the best found variant (heuristic) of a processor to produce a custom variant of the generated code as discussed in Section 7.

Processor	Short	Architecture	Vendor
A10-7850K (CPU)	CPU	Kaveri	AMD
A10-7850K (GPU)	iGPU	CGN 1.1	AMD
Tesla K40M	dGPU	Kepler	Nvidia
Xeon Phi 7120	MIC	Knights Corner	Intel

Table 3 Processors used in evaluation.

9 Evaluation

In this section, we discuss our experimental setup and design, present our results on hardware adaption, and discuss the implications of these results.

9.1 Experimental Setup

As evaluation platform, we use two machines that have several heterogeneous processors installed. In total, we consider four different processor types with varying architectures: a CPU, an integrated GPU (iGPU), a dedicated GPU (dGPU), and a MIC, as shown in Table 3. All machines run Ubuntu 16.04 LTS (64bit). Depending on the processor's vendor, we have to use a certain OpenCL SDK and driver to compile and run our kernels. For CPU and iGPU from AMD, we use the AMD APP SDK version 3.0. For the MIC, we use the Intel OpenCL SDK Version 4.5.0.8. For the dGPU from Nvidia, we use the CUDA SDK version 8.

For processors with dedicated main memory, we cache all input data before running the variants to avoid biased observations because of PCIe transfers. Our goal is to evaluate the performance of queries on heterogeneous processors, rather than bottlenecks in current interconnects. We run each variant of a pipeline program 5 times and report the mean and the standard deviation. We prune the variant space if we detect a very slow variant (execution time greater than one second) to keep the run-time of the benchmark in a reasonable time frame.

As evaluation datasets, we use the Star Schema Benchmark [34] and the TPC-H Benchmark [49]. We use Scale Factor 1 for the experiments including a full exploration of all variants. With larger scale factors, poor performing variants would not finish in reasonable time. As OpenCL does not provide any mechanism to abort a kernel, we have to wait until the kernel finishes. For all other experiments, we use a scale factor of 10. The main memory of the iGPU usable by OpenCL is limited to 2.2GB and thus, we can not use a larger scale factor.

Listing 8 Grouped Aggregation Query 1

```sql
select lo_shipmode, sum(lo_quantity) from
    lineorder group by lo_shipmode;
```

Listing 9 Grouped Aggregation Query 2

```sql
select lo_partkey, sum(lo_quantity) from
    lineorder group by lo_partkey;
```

9.2 Experimental Design

We now discuss our evaluation workload and the variants we generate for our test queries.

9.2.1 Queries

All SQL queries can be split in a series of projection and aggregation pipelines. Thus, we evaluate our approaches for pipeline variant generation and optimization on simple queries representing a single pipeline. These single-pipeline queries allow for unbiased observation of hardware adaption using pipeline variations. Additionally, we validate the results on complex benchmark queries.

Projection Pipelines. We take as representatives for projection pipelines one query with 50 % selectivity (Projection Query 1, cf. Listing 1) and one filter query with very high selectivity (<0.01 %, Projection Query 2, cf. Listing 2). While Projection Query 1 is read and write intensive, Projection Query 2 is read intensive.

Aggregation Pipelines. As representatives for aggregation pipelines, we use one query with few result groups (Aggregation Query 1, cf. Listing 8) and one query with many result groups, e.g., several hundred thousand (Aggregation Query 2, cf. Listing 9). The first query is common for the final group by in an OLAP query. The second query is common in sub-queries using a group by (e.g., TPC-H Query 15). Having so many groups, Aggregation Query 2 is latency bound.

TPC-H Q1 and SSB Q4.1. We perform a full variant exploration for TPC-H Q1 and SSB Q4.1. The TPC-H query is a compute intensive aggregation query. It consists of a single pipeline with a FILTER, several ALGEBRA operations and a grouped aggregation with multiple aggregation functions. The SSB query is a join dominated query (four joins), consisting of four projection pipelines and one aggregation pipeline. The projection pipelines build the hash tables, whereas the aggregation pipeline probes each hash table.

Other Queries. We also evaluate the performance of other Star Schema Benchmark and TPC-H queries.

Due to current implementation restrictions of our prototype system (e.g., a missing LIKE operator), we limit the evaluation queries to a representative subset. For the Star Schema Benchmark, we use the queries Q1.1-Q1.3, Q3.2-Q3.4 and Q4.1-Q4.3. For the TPC-H benchmark, we use the queries Q5, Q6 and Q7.

9.2.2 Variant Space of Generated Variants

For all pipeline types, we vary the memory access pattern (sequential and coalesced) and the branch evaluation mode (branched predicate evaluation and software predication). The total number of variants multiply with each new variant dimension. We encode the number of variants in brackets [x variants].

Projection Pipelines. For projection pipelines, we additionally vary the execution strategy (single pass for coarse-grained parallelism and multi pass for fine-grained parallelism) [2 variants]. For the single-pass strategy, we set the number of parallel running pipelines to the number of *maximal compute units* of the OpenCL device [1 variant]. Thus, we generate 4 variants that use the single-pass strategy. The multi-pass strategy uses a multiplier (1, 8, 64, 256, 1024, 16384, 65536) that is multiplied with the number of *maximal compute units* of the OpenCL device to calculate the number of threads [7 variants]. We generate 28 variants that use the multi-pass strategy. In total, we generate 32 variants for a projection pipeline.

Aggregation Pipelines. For aggregation pipelines, we additionally vary the aggregation execution strategy, the hash table implementation, and the hash function. For the hash table implementation, we vary between linear probing and Cuckoo hashing [2 variants]. The hash function is either Murmur hashing or Multiply-Shift hashing [2 variants]. For the execution strategy we vary between local and global aggregation. In case of a local aggregation, we optimize the number of hash tables (1, 8, 64, 256, 1024, 16384, 65536) as a multiplier of the number of *maximal compute units* of the OpenCL device to test different levels of thread over-subscription. We also optimize the number of threads per hash table (16, 32, 64, 128, 256, 512, 1024) to find the best configuration between high parallelism and synchronization overhead [7x7 variants]. In case of global aggregation, we optimize the number of threads per hash table, which configurations are identical to local aggregation [7 variants]. We generate [2x2x2x2x7x7 variants] for the local aggregation and [2x2x2x2x7 variants] variants for the global aggregation. In total, we generate 896 variants for an aggregation pipeline.

9.3 Results

We validate our concepts as follows. First, we evaluate kernel compilation times for all generated kernels. Second, we evaluate all variants on representative queries: two projection queries, two aggregation queries, TPC-H Query 1, and SSB Query 4.1. We determine the optimal variant of a pipeline program by performing a full search. This means that we generate all possible variants for a pipeline and execute them multiple times. The in average fastest variant is reported in the plots as CPU, iGPU, dGPU, and MIC optimized. Furthermore, we evaluate our learning strategy for automatic hardware adaption. We report the learned variant configurations and the query execution times on different processors.

9.3.1 Compilation Times

In Figure 18, we show for each of our evaluation queries and processors the compilation time of all variants in a box plot. The boxes include 50 % of the observations, whereas the upper and lower whiskers mark a 99 % confidence interval. It becomes clearly visible that all compilation times for kernels for the projection query are below 70ms for CPU, iGPU, and dGPU, and below 100 ms for the MIC processor. For the aggregation query, we observe that except for the MIC processor, kernel compilation times are either 100ms, or below. As we need to generate more code for aggregation pipelines compared to projection pipelines, the compilation time increases.

Compiling TPC-H Query 1 takes longer compared to the aggregation queries. This is because the TPC-H query results in a larger kernel due to many additional computations. We observe 66ms on the CPU, 113ms on the iGPU, 216ms on the dGPU and 4.9s on the MIC. Compiling SSB Query 4.3 is even more time intensive, as we have to compile four projection and one aggregation pipelines. We observe 245ms on the CPU, 380ms on the iGPU, 818ms on the dGPU and 1.8s on the MIC. Note that we can compile multiple pipelines in parallel to reduce the compilation time.

Compiling for the MIC is very expensive, and may take longer than a second, even for a single pipeline. However, this is the only processor where we observed this behavior. We repeated our experiments on other machines using NVIDIA GPUs and Intel CPUs, and measured similar kernel compilation times reported here for CPU, iGPU, and dGPU. Thus, we assume that the high compilation time for the MIC is an implementation artifact, which we expect will be resolved in future versions of the Intel OpenCL SDK.
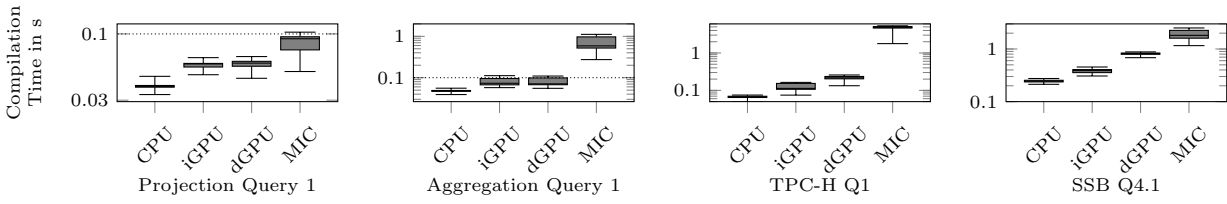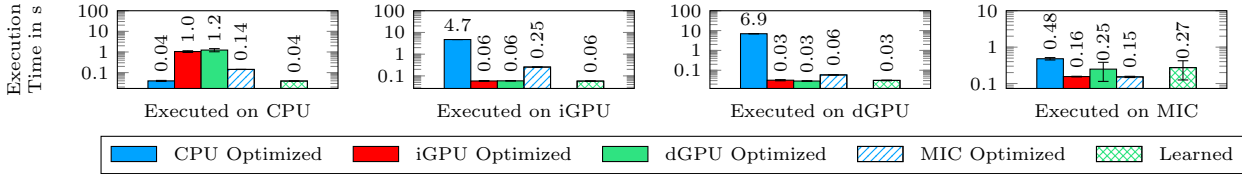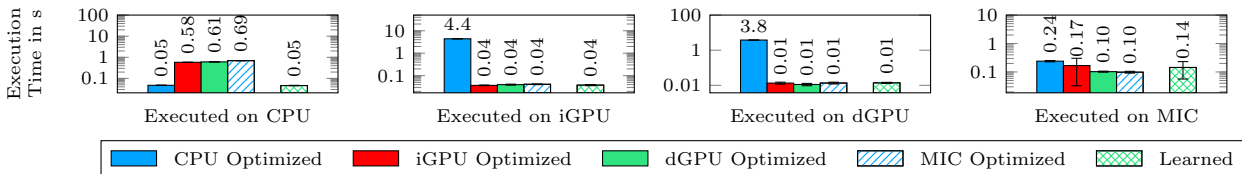
Fig. 18 Compilation times for all generated kernel variants for each processor and query pipeline. Most kernels can be compiled in less than 100ms, which allows for fast query compilation.

(a) Execution times of variants optimized for CPU, iGPU, dGPU, and MIC for Projection Query 1.

(b) Execution times of variants optimized for CPU, iGPU, dGPU, and MIC for Projection Query 2.
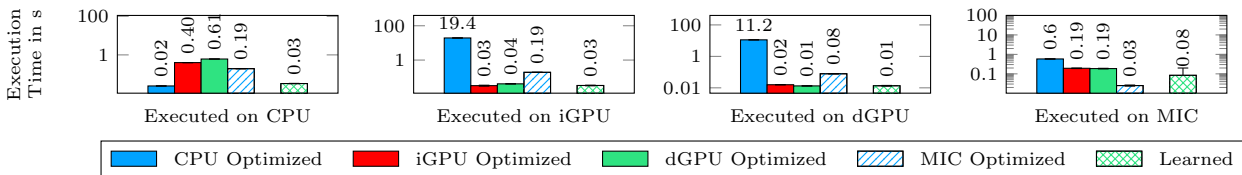
(c) Execution times of variants optimized for CPU, iGPU, dGPU, and MIC for Aggregation Query 1.

Fig. 19 Execution times of variants optimized for CPU, iGPU, dGPU, and MIC for different queries, executed on all processors.

9.3.2 Full Variant Exploration

We show the run-time of all variants optimized for a particular processor and query. We show these variants for the following queries: the projection queries (Listing 1 and 2), the aggregation queries (Listing 8 and 9), TPC-H Query 1, and SSB Query 4.3.

Observations Projection Query 1. We show in Figure 19(a) that the CPU-optimized variant outperforms the variants optimized for the iGPU, dGPU, and MIC by a factor of 25, 30, and 3.5, respectively. However, we see that the same implementation performs more slowly compared to optimized variants on the iGPU, dGPU, and MIC by a factor of up to 81, 237, and 3.1, respectively. The large performance difference between CPU and the other processors is mainly due to the execution strategy: CPUs prefer the single-pass strategy using coarse-grained parallelism, whereas GPUs, and MICs prefer the multi-pass strategy using fined-

grained parallelism. On the iGPU, we observe that the variant optimized for iGPU outperforms the variant optimized for MIC by a factor of 4.3. For the dGPU optimized variant on the iGPU, we observe that the performance is equal to the iGPU optimized variant. The main difference among the variants optimized for iGPU, dGPU, and MIC is in the optimal number of threads. Furthermore, the MIC prefers sequential memory access, similar to CPUs, whereas the GPU variants prefer coalesced memory access. We do not observe further performance gaps on the dGPU and MIC processor. Our learning strategy found a configuration that performs closely to the optimal variant on CPU, iGPU and dGPU. On the MIC, the found variant is by a factor of 1.8 slower than the optimum.

Observations Projection Query 2. In Figure 19(b), we make the same basic observation for Projection Query 2 (very high selectivity, <0.001 %) as for Projection Query 1 (50 % selectivity). The CPU-optimized variant
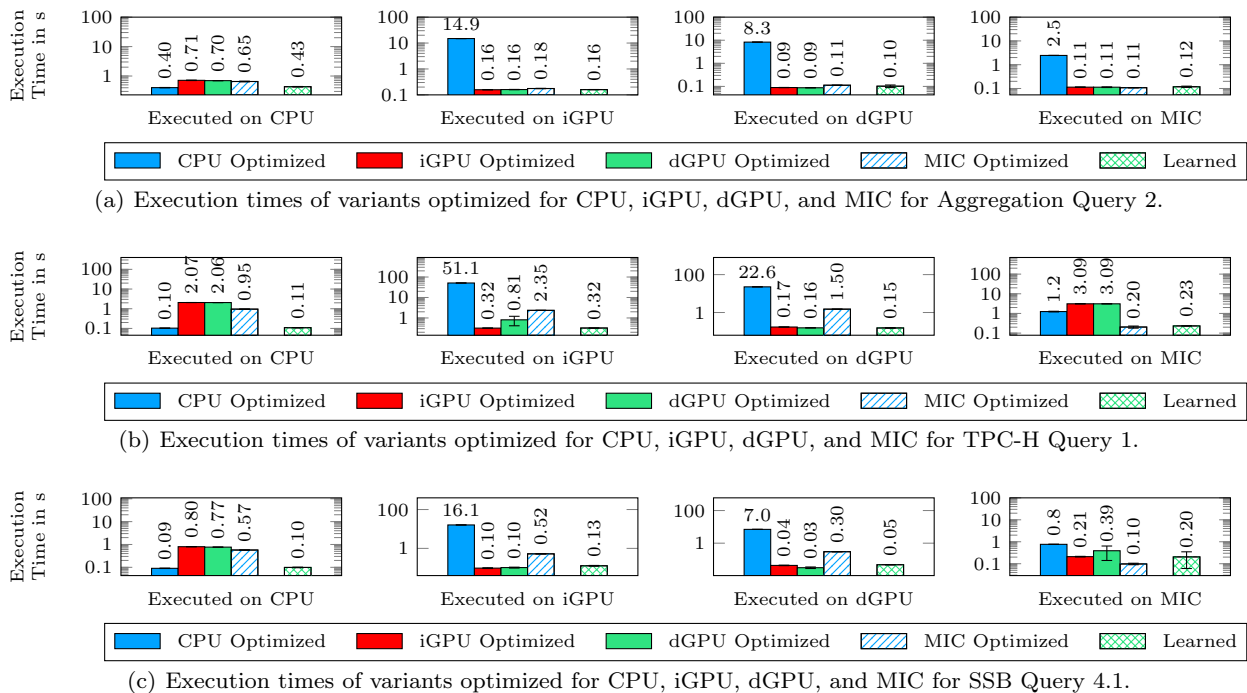
(a) Execution times of variants optimized for CPU, iGPU, dGPU, and MIC for Aggregation Query 2.

(b) Execution times of variants optimized for CPU, iGPU, dGPU, and MIC for TPC-H Query 1.

(c) Execution times of variants optimized for CPU, iGPU, dGPU, and MIC for SSB Query 4.1.

Fig. 20 Execution times of variants optimized for CPU, iGPU, dGPU, and MIC for different queries, executed on all processors.

Processor	Backtracking (in seconds)	Feature-Wise (in seconds)	Factor Improved
CPU	197,139	479	411
iGPU	78,388	1,219	64.3
dGPU	52,897	1,036	51
MIC	177,914	3,390	52.5

Table 4 Variant exploration times for SSB Q4.1 on SF1. Our learning strategy outperforms the backtracking search by up to two orders of magnitude.

outperforms variants optimized for iGPU, dGPU, and MIC by a factor of 12, 12, and 14, respectively. On the other processors, the CPU-optimized variant is slower by a factor of 118, 345, 2.4 on the iGPU, dGPU, and MIC, respectively. Our learning strategy found a configuration that performs closely to the optimal variant on CPU, iGPU, dGPU, and MIC.

Observations Aggregation Query 1. We show in Figure 19(c) that on the CPU the CPU-optimized variant outperforms variants optimized for iGPU, dGPU, and MIC by a factor of 16, 24.4, and 7.6, respectively. However, we see that the same implementation performs significantly slower compared to optimized variants on the iGPU, dGPU, and MIC by a factor of up to 606, 861, and 24, respectively. We also see significant differences between the variants optimized for iGPU, dGPU, and MIC: On the iGPU, the iGPU-optimized variant outperforms the variants optimized for dGPU

and MIC by a factor of 1.2, and 6, respectively. On the dGPU, the dGPU-optimized variant outperforms variants optimized for iGPU and MIC by a factor of 1.2, and 6. On the MIC, the MIC-optimized variant outperforms variants optimized for iGPU and dGPU by a factor of 7.6 and 7.6, respectively. Our learning strategy found a configuration that performs closely to the optimal variant on CPU, iGPU and dGPU. On the MIC, the found variant is by a factor of 3.4 slower than the optimal variant.

Observations Aggregation Query 2. We make the same basic observation as for Aggregation Query 1 (cf. Figure 20(a)): On the CPU, the CPU-optimized variant outperforms variants optimized for iGPU, dGPU, and MIC by a factor of 1.7, 1.7, and 1.6, respectively. The same CPU-optimized variant is significantly slower compared to variants optimized for iGPU, dGPU, and MIC by a factor of up to 94, 94, and 23. Note that for this query, the optimal variant of iGPU and dGPU is the same, thus we will report numbers only once (GPU). On the GPUs, the GPU-optimized variant outperforms the MIC by a factor of 1.1 (iGPU) and 1.2 (dGPU). On the MIC, the MIC-optimized variant achieves the same performance as the GPU variant. Our learning strategy found a configuration that performs closely to the optimal variant on CPU, iGPU, dGPU and MIC.

Observations on Complex Queries. We show that the variant exploration has the same impact on

Variation Dimension	Learned Optimizers		
	cpu-o	dgpu-o	mic-o
Execution Strategy (Projection)	single-pass	multi-pass	multi-pass
Execution Strategy (Aggregation)	local hash table	local hash table	local hash table
Memory Access Pattern	sequential	coalesced	coalesced
Hash Table Implementation	linear probing	Cuckoo hashing	Cuckoo hashing
Predication Mode (Query Dependent)	branched	branched	branched
Thread Multiplier (Projection)	1	16384	65536
Thread Multiplier (Aggregation)	1	1	1
Work Group Size (Aggregation)	256	1024	64

Table 5 Per processor variant configurations identified by the variant learning strategy for the SSB and TPC-H workload.

more complex queries. Thus, we present the result of the variant exploration for two OLAP queries: TPC-H Q1 (cf. Figure 20(b)) and SSB Q4.1 (cf. Figure 20(c)). For every processor, we observe similar factors between the optimal variant and the other variants. We also see that on the MIC, the variance of execution time for some variants is very high. We do not observe this issue on any other processor.

9.3.3 Optimization Time

We now investigate how long the variant exploration itself takes. We show in Table 4 the time to explore the best variant for SSB Query 4.1. We compare backtracking (executing every possible variant and choosing the fastest) with our learning strategy. We observe that our strategy improves the search time by up to a factor of 411. While the longest exploration took more than two days, our strategy finished within an hour. Thus, we can run our calibration benchmarks offline (e.g., as part of the database installation process).

9.3.4 Hardware Adaption on Full Queries

Learned Variant Configurations. We derive processor-specific optimizers using our learning strategy from Section 8. We explore the same variant space as the full exploration, which we discuss in Section 9.2.2. As training workload, we use the Query Groups 1, 3, and 4 of the Star Schema Benchmark and Queries 5, 6, 7 from the TPC-H benchmark. In this experiment, we use a scale factor of 10 for both benchmarks. We show the learned variant configurations optimized for the CPU (cpu-o), for the dGPU (dgpu-o) and for the MIC (mic-o) in Table 5.

The learning strategy correctly identifies that for projection pipelines, CPUs prefer single-pass strategies with coarse-grained parallelism, whereas the GPUs and the MIC prefer multi-pass strategies with fine-grained parallelism. For aggregation pipelines, the CPU, GPU, and MIC prefer local hash table aggregation. The learning strategy also found that the CPU prefers sequential

memory access, whereas the GPUs and MIC prefer coalesced memory access. The preferred aggregation hash table for CPUs is linear probing, whereas the GPUs and the MIC are more efficient when using Cuckoo hashing. For the query workload, the learning strategy found that the branched predication mode (evaluation using if-statements) outperforms variants that use software predication.

We implement the number of threads as a multiplier of the number of OpenCL compute units ("cores"), as the multiplier quantifies the degree of over-subscription required for a processor. CPUs prefer no over-subscription (one thread per core), whereas the GPUs and the MIC need a large multiplier (over-subscription) to have enough thread blocks ready to hide memory access latencies. Additionally, we need to specify the work group size for aggregation pipelines, which also strongly differs between the different processors.

Performance. We execute for each query a variant optimized for CPU, dGPU, and MIC and measure the execution times on CPU, dGPU, and MIC without compilation times. Note that each variant is optimized for a complete workload (cpu-o, dgpu-o, and mic-o). We call these variants per-workload variants. We include measurements of a per-query optimized variant for each query (q-o) to show additional optimization potential compared to the per-workload variants.

We illustrate the results in Table 6 and include measurements of HyPer (v0.5-222-g04766a1) with the same queries on the same dataset on the CPU.[1] We observe that the code generated by Hawk on a CPU is in the same order of magnitude as an optimized state-of-the-art query compiler.

Most queries are executed faster when we use the per-workload variant of the target processor. On the CPU, the performance of a CPU-optimized variants outperforms GPU and MIC-optimized variants by up to a factor of 5.5 (SSB Query 3.4). On the GPU, the GPU-optimized variant outperforms the other per-workload

---

[1] Note that this comparison is not intended to be an end-to-end measurement of system performance.

	Hyper (CPU)	Executed on CPU				Executed on dGPU				Executed on MIC			
		cpu-o	dgpu-o	mic-o	per-q	cpu-o	dgpu-o	mic-o	per-q	cpu-o	dgpu-o	mic-o	per-q
Q1.1	0.149	0.186	0.441	0.342	0.189	0.067	0.015	0.015	0.015	0.055	0.062	0.057	0.057
Q1.2	0.099	0.113	0.271	0.272	0.114	0.052	0.047	0.047	0.013	0.046	0.061	0.06	0.064
Q1.3	0.092	0.111	0.25	0.248	0.109	0.052	0.022	0.023	0.013	0.047	0.056	0.049	0.051
Q3.2	0.2	0.21	2.258	0.885	0.206	56.697	0.191	1.724	0.138	5.021	0.247	0.221	0.155
Q3.3	0.146	0.115	1.467	0.61	0.114	53.543	0.073	0.472	0.052	3.781	0.14	0.14	0.114
Q3.4	0.146	0.111	1.468	0.615	0.114	53.646	0.073	0.471	0.053	3.795	0.132	0.128	0.109
Q4.1	0.654	0.567	2.186	1.559	0.567	77.701	0.743	5.188	0.25	11.146	0.423	0.397	0.417
Q4.2	0.588	0.444	1.758	1.272	0.45	78.042	0.523	1.704	0.111	8.552	0.341	0.322	0.351
Q4.3	0.316	0.195	2.421	1.073	0.212	58.718	0.764	4.816	0.286	4.709	0.435	0.415	0.343
Q5	0.857	0.934	5.105	4.095	1.033	73.572	7.091	10.605	0.261	10.874	0.905	0.907	0.838
Q6	0.147	0.185	0.257	0.258	0.195	0.063	0.009	0.009	0.011	0.033	0.037	0.036	0.036

Table 6 Execution times in seconds of variants optimized for CPU (cpu-o), dGPU (dgpu-o), and MIC (mic-o) for selected queries of the star schema and TPC-H benchmark (Scale Factor 10), executed on a CPU, a dedicated GPU, and a MIC processor.

variants by up to a factor of 9 (SSB Query 3.2). On the MIC, the MIC-optimized variant outperforms the other per-workload variants by up to a factor of 1.12 (SSB Query 3.2). The reason for this low factor is that GPU variants are typically also fast on a MIC (but not the other way around). However, we can still improve the performance with a custom variant for the MIC. We occasionally observe a better performance of another variant for some queries, such as TPC-H Q5 and Q6 for the CPU. The reason for this is that a variant optimized for several queries may be sub-optimal for a particular query. We conclude that we achieve the best performance when we use a variant optimized for a target processor.

If we additionally tune the variant to a particular query, we observe for our workload speedups on the CPU by up to a factor of 1.02, on the GPU by up to a factor of 27 (TPC-H Query 5), and on the MIC by up to a factor of 1.43 (SSB Query 3.2). The per-query variants differ mainly in the thread multipliers, as different degrees of parallelism are optimal for different queries on GPU and MIC. Additionally, some queries are faster with enabled predication or prefer global instead of local hash table aggregation specific queries, such as TPC-H Q5 on GPUs and MICs. We conclude that the optimal variant is query-dependent on GPU and MIC. On the CPU, a per-query variant provides small benefit over a generic per-workload variant.

Summary. In general, we conclude that we need to create custom variants per processor in order to reach peak performance. Furthermore, we observe that the algorithm derives an efficient configuration, but the optimal variant is to some extend query dependent, e.g., number of threads, branching mode [46] and hash table implementation [43]. This limitation can be lifted by adding a run-time optimizer that performs per-query

variant optimization, similar to the work of Raducanu [46] and Zeuch [54].

9.4 Discussion

In our experiments, we observed that most compilation times for single pipelines are very fast (< 100 ms). OpenCL could compile even complex queries in several hundred milliseconds, if we disregard vendor-specific artifacts. We conclude that efficient query compilation is possible using OpenCL. This ensures that the database engine still allows for interactive querying despite using query compilation.

Furthermore, we observe large performance differences among variants optimized for a CPU, a GPU, and a MIC by up to two orders of magnitude. Thus, we conclude that a hardware-adaptive query compiler can achieve high performance gains. This is because it can optimize for various processors of different architectures with previously unknown performance behavior *without any manual tuning.*

The diversity of the optimized variants shows that we need to support the discussed dimensions of variant generation and their individual variations.

Finally, we find that our learning strategy detected all major preferences of all processors. The strategy derived efficient per-processor variants without having to explore all variants. In future work, we will research a run-time optimizer that also captures query-dependent parameters.

We observed on the MIC that the execution time of some variants have a high variance (up to a factor of four). We could not pinpoint the exact cause. We run the same OpenCL code on different CPUs and GPUs with different OpenCL vendors and observe this behavior with the MIC only. Thus, the cause is likely to be

an implementation artifact of the Intel OpenCL SDK. The unreliable performance behavior of the MIC processor makes it difficult to find a fast variant on the MIC for any search strategy. We mitigate the problem using outlier detection on the execution times. However, with growing randomness of the processors performance, optimization gets increasingly difficult.

10 Related Work

In this section, we discuss related work on query compilation, compiling programs to heterogeneous processors, data processing on heterogeneous processors, and automatic optimization of variants.

10.1 Query Compilation

Query compilation goes back to System R [8], and was re-investigated in the 80s [14]. With the upcoming of main-memory databases, query compilation received new attention as reducing main-memory traffic and executed CPU instructions became increasingly important. Rao and others generated query-specific code using the just-in-time compilation capabilities of Java [41]. Krikellas and others used a template-based code generation approach to compile queries to C code, which was then compiled by a C compiler to machine code [26]. Neumann introduced the produce/consume model, which provides a systematic way to generate code that allows for data-centric query processing by fusing all operators in an operator pipeline. Additionally, Neumann proposed to generate LLVM IR code instead of C code to achieve low compilation times [33]. Leis and others proposed the morsel framework, which introduces NUMA-aware parallelization of compiled operator pipelines [28].

Sompolski and others carefully studied vectorized and compiled execution [48]. They observe that compilation is not always superior to vectorization and conclude that compilation should always be combined with block-wise query processing. Dees and Sanders compiled the 22 TPC-H queries by hand to C code and showed large performance potentials for query compilation [10]. Nagel and others investigated query compilation in the context of language-integrated queries in managed run-times [32]. Amad and others developed DBToaster, which uses code generation to compile view maintenance queries to efficient machine code [1]. Query compilation also found its way into commercial products such as Hekaton [13] and Impala [50].

Weld is a run-time that efficiently executes data-intensive applications [35]. The key idea is to compile code to a common intermediate representation. Weld removes data movement between functions in a work-flow and generates efficient parallel code for CPUs. In contrast to Hawk, Weld cannot generate custom code for different heterogeneous processors. However, Welds code-generation backend can be enriched by the variant generation concepts introduced in this paper to efficiently support GPUs and MICs.

10.2 Query Compilation for CPUs and GPUs

Wu and others proposed Kernel Weaver, a compiler framework that can automatically fuse the kernels of relational operators and kernels of other domains [51]. In contrast to kernel weaver, Hawk uses our concept of execution strategies to generate a minimal number of kernels. We see Kernel Fusion as a complementary building block. Another key difference is that Kernel Weaver targets GPUs only, whereas Hawk executes efficiently on CPUs, GPUs, and MICs.

Rauhe and others proposed the compute/accumulate model to compile queries to GPU code [42]. Here, query operations are put in three kernels. The compute phase generates a local compute and a local accumulate kernel. The accumulate phase generates a global accumulate kernel. In contrast to Hawk, they cannot generate highly tailored code for CPUs and GPUs.

A new line of research focuses on writing database systems in a high-level language [25]. The LegoBase system uses generative programming to generate efficient low-level C code for a database implementation in a high-level language [24]. Shaikhha and others further refine this principle in DBLAB [47] by introducing a stack of multiple Domain Specific Languages (DSLs) that differ in the levels of abstraction. Here, high-level code is progressively lowered to low-level code, by compiling code in multiple stages, where each stage compiles to a DSL of lower abstraction level, until the final code is generated.

Pirk and others propose the Voodoo framework, which consists of an intermediate algebra representation based on vectors and a code generator for OpenCL [39]. Based on the algebra, Voodoo is capable of generating code for different processors, including CPUs and GPUs. The voodoo algebra and our pipeline programs are conceptually similar, albeit on different abstraction levels. Note that pipeline programs and voodoo algebra are complementary, as we could generate voodoo algebra from pipeline programs. The key difference between Hawk and Voodoo is that Hawk creates a large space of potential variants and can systematically fine-tune the generated code to the underlying processors.

In summary, existing query compilation approaches generate efficient code for a single processor. Hawk is

the first hardware-adaptive query compiler that can produce variants of code to run efficiently on different processors.

10.3 Compilers

Brown and others developed Delight, a framework that allows to build, compile and execute DSLs which enable users to program at a high-abstraction level [7]. The key idea is to compile domain-specific languages to a common intermediate representation. From the intermediate representation, Delight generates code for CPUs and GPUs. However, Delight does not optimize for heterogeneous processors to the degree Hawk does, such as changing execution strategies. The concepts of Delight and Hawk complement each other.

Dandelion is a general purpose compiler based on .NET LINQ that compiles data-parallel programs to multiple heterogeneous processors, such as CPUs, GPUs, and FPGAs and automatically distributes data processing on different processors, be it in a single machine or a cluster [45]. While Dandelion uses cross compilation to support GPUs, Hawk profits from the functional portability of OpenCL, which allows Hawk to run code on any OpenCL-capable processor. At the same time, Hawk can generate custom code for different processors using the same code generator.

Jacc is a compiler framework that can compile and run Java code on GPUs. Jacc compiles Java byte code directly to NVIDIA's PTX code and manages computation and data transfers transparantly to the user [9]. In contrast to Jacc, Hawk generates variants of code to find a high-performance implementation.

10.4 Databases on Heterogeneous Hardware

Balkesen and others studied efficient hash joins [3] and sort-merge joins on multi-core CPUs [2]. He and others developed efficient algorithms for joins [17,18] and other relational operators [16] on GPUs. He and others also studied efficient co-processing on APUs [19]. Pirk and others studied common database operations on the Intel Xeon Phi (MIC) and compared them to GPUs [37]. Jha and others investigated hash joins on the Intel Xeon Phi [21].

Paul and others investigated the effect of pipelining between multiple GPU kernels using the channel mechanism provided by OpenCL 2.0 pipes [36].

Meraji and others implemented support for GPU acceleration into DB2 with BLU acceleration and observed significant performance gains using GPUs for query processing [29].

Karnagel and others analyzed hash-based grouping and aggregation on GPUs [22]. This work was the basis for Hawk's execution strategies for grouped aggregation. Müller and others studied database query processing on FPGAs [30] and developed Glacier, a query compiler that generates logic circuits for queries to accelerate stream processing [31].

Many database prototypes were developed to study different aspects of query processing on CPUs and GPUs, such as GDB [16], GPUDB [52], OmniDB [56], Ocelot [20], CoGaDB [6], and HeteroDB [55].

Pirk and others introduced the approximate and refine technique, which only keeps the higher bits of a value in GPU memory, which allows to keep more data on GPU. The lossily compressed data is stored on the GPU [38]. Queries on the GPU return an approximate result, which has to be refined on the CPU using the uncompressed data.

Heimel and others showed the feasibility of building a database engine in OpenCL, which allows to run a database engine with the same operator code base on any OpenCL-capable processor [20]. To achieve high efficiency on CPUs and GPUs, they left the memory access pattern configurable so they can adapt it to the processor type. The core difference between Ocelot and Hawk is that Ocelot provides the same operator implementations for each processor, while Hawk can generate custom per-processor variants for each query.

All the aforementioned techniques improve the absolute performance of a database system running on heterogeneous processors. These optimizations are orthogonal to the concepts presented in this paper.

10.5 Variant Optimization

Raducanu and others propose Micro Adaptivity, a framework that provides alternative function implementations called *flavors* (equivalent to our term variant) [46]. Micro Adaptivity exploits the vector-at-a-time processing model and can potentially exchange a flavor at each function call, which allows for finding the best implementation for a certain query and data distribution.

Rosenfeld and others showed for selection and aggregation operations that many operator variants can be generated and that different code transformations are optimal for a particular processor [44].

Zeuch and others propose to use performance counters of modern CPUs for progressive optimization. They introduce cost models for cache accesses and branch mispredictions and derive selectivities of predicates at query run-time to re-optimize predicate evaluation orders [54].

The techniques for variant optimization from Raducanu [46], Rosenfeld [44], and Zeuch [54] are orthogonal to the variant generation of this paper.

11 Summary

In this paper, we describe a hardware-adaptive query compiler that can generate code for a wide range of heterogeneous processors. Through hardware-tailored implementations, our query compiler produces fast code *without manual tuning* for a specific processor.

Our key findings are as follows. Our abstraction of pipeline programs allows us to flexibly produce variants of pipelines while keeping a clean interface and a maintainable code base. Pipeline variants optimized for a particular processor can result in performance differences of up to two orders of magnitude on the same processor. Therefore, it is crucial to optimize the query compiler to each processor. Consequently, we proposed a learning strategy that automatically derives an efficient variant configuration for a processor. Based on this algorithm, we derived efficient variant configurations for three common processors. Finally, we incorporated the variant configurations into a heuristic query optimizer.

References

1. Y. Ahmad and C. Koch. DBToaster: A SQL compiler for high-performance delta processing in main-memory databases. *PVLDB*, 2(2):1566–1569, 2009.
2. C. Balkesen, G. Alonso, J. Teubner, and T. Özsu. Multi-core, main-memory joins: Sort vs. hash revisited. *PVLDB*, 7(1):85–96, 2013.
3. C. Balkesen, J. Teubner, G. Alonso, and M. T. Özsu. Main-memory hash joins on multi-core CPUs: Tuning to the underlying hardware. In *ICDE*, pages 362–373, 2013.
4. S. Borkar and A. Chien. The future of microprocessors. *Communications of the ACM*, 54(5):67–77, 2011.
5. S. Breß. The design and implementation of CoGaDB: A column-oriented GPU-accelerated DBMS. *Datenbank-Spektrum*, 14(3):199–209, 2014.
6. S. Breß, H. Funke, and J. Teubner. Robust query processing in co-processor-accelerated databases. In *SIGMOD*, pages 1891–1906. ACM, 2016.
7. K. Brown, A. Sujeeth, H. Lee, T. Rompf, H. Chafi, M. Odersky, and K. Olukotun. A heterogeneous parallel framework for domain-specific languages. In *PACT*, pages 89–100. IEEE, 2011.
8. D. Chamberlin et al. A history and evaluation of System R. *Commun. ACM*, 24(10):632–646, 2013.
9. J. Clarkson, C. Kotselidis, G. Brown, and M. Luján. Boosting Java performance using GPGPUs. In *International Conference on Architecture of Computing Systems (ARCS)*, pages 59–70. Springer, 2017.
10. J. Dees and P. Sanders. Efficient many-core query execution in main memory column-stores. In *ICDE*, pages 350–361. IEEE, 2013.
11. Esmaeilzadeh et al. Dark silicon and the end of multicore scaling. In *ISCA*, pages 365–376. ACM, 2011.
12. F. Färber, N. May, W. Lehner, P. Große, I. Müller, H. Rauhe, and J. Dees. The SAP HANA database – an architecture overview. *IEEE Data Eng. Bull.*, 35(1):28–33, 2012.
13. C. Freedman, E. Ismert, and P. Larson. Compilation in the microsoft SQL server hekaton engine. *IEEE Data Eng. Bull.*, 37(1):22–30, 2014.
14. J. Freytag and N. Goodman. On the translation of relational queries into iterative programs. *ACM Trans. Database Syst.*, 14(1):1–27, 1989.
15. B. Gaster, L. Howes, D. Kaeli, P. Mistry, and D. Schaa. *Heterogeneous Computing with OpenCL: Revised OpenCL 1.2*. Newnes, 2012.
16. B. He, M. Lu, K. Yang, R. Fang, N. Govindaraju, Q. Luo, and P. Sander. Relational query co-processing on graphics processors. In *ACM Trans. Database Syst.*, volume 34. ACM, 2009.
17. B. He, K. Yang, R. Fang, M. Lu, N. Govindaraju, Q. Luo, and P. Sander. Relational joins on graphics processors. In *SIGMOD*, pages 511–524. ACM, 2008.
18. J. He, M. Lu, and B. He. Revisiting co-processing for hash joins on the coupled CPU-GPU architecture. *Proc. VLDB Endow.*, 6(10):889–900, 2013.
19. J. He, S. Zhang, and B. He. In-cache query co-processing on coupled CPU-GPU architectures. *PVLDB*, 8(4):329–340, 2014.
20. M. Heimel, M. Saecker, H. Pirk, S. Manegold, and V. Markl. Hardware-oblivious parallelism for in-memory column-stores. *PVLDB*, 6(9):709–720, 2013.
21. S. Jha, B. He, M. Lu, X. Cheng, and H. Huynh. Improving main memory hash joins on Intel Xeon Phi processors: An experimental approach. *PVLDB*, 8(6):642–653, 2015.
22. T. Karnagel, R. Müller, and G. Lohman. Optimizing GPU-accelerated group-by and aggregation. In *International Workshop on Accelerating Data Management Systems Using Modern Processor and Storage Architectures (ADMS)*, pages 13–24, 2015.
23. A. Kemper and T. Neumann. HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In *ICDE*, pages 195–206. IEEE, 2011.
24. Y. Klonatos, C. Koch, T. Rompf, and H. Chafi. Building efficient query engines in a high-level language. *PVLDB*, 7(10):853–864, 2014.
25. C. Koch. Abstraction without regret in database systems building: a manifesto. *IEEE Data Eng. Bull.*, 37(1):70–79, 2014.

26. K. Krikellas, S. Viglas, and M. Cintra. Generating code for holistic query evaluation. In *ICDE*, pages 613–624. IEEE, 2010.

27. P.-A. Larson, A. Birka, E. N. Hanson, W. Huang, M. Nowakiewicz, and V. Papadimos. Real-time analytical processing with sql server. *Proc. VLDB Endow.*, 8(12):1740–1751, 2015.

28. V. Leis, P. Boncz, A. Kemper, and T. Neumann. Morsel-driven parallelism: A NUMA-aware query evaluation framework for the many-core age. In *SIGMOD*, pages 743–754. ACM, 2014.

29. S. Meraji et al. Towards a hybrid design for fast query processing in DB2 with BLU acceleration using graphical processing units: A technology demonstration. In *SIGMOD*, pages 1951–1960. ACM, 2016.

30. R. Müller, J. Teubner, and G. Alonso. Data processing on FPGAs. *PVLDB*, 2(1):910–921, 2009.

31. R. Müller, J. Teubner, and G. Alonso. Streams on wires - A query compiler for FPGAs. *PVLDB*, 2(1):229–240, 2009.

32. F. Nagel, G. Bierman, and S. Viglas. Code generation for efficient query processing in managed runtimes. *PVLDB*, 7(12):1095–1106, 2014.

33. T. Neumann. Efficiently compiling efficient query plans for modern hardware. *PVLDB*, 4(9):539–550, 2011.

34. P. O'Neil, E. J. O'Neil, and X. Chen. The star schema benchmark (SSB), 2009. Revision 3, `http://www.cs.umb.edu/~poneil/StarSchemaB.PDF`.

35. S. Palkar, J. Thomas, A. Shanbhag, D. Narayanan, H. Pirk, M. Schwarzkopf, S. Amarasinghe, and M. Zaharia. Weld: A common runtime for high performance data analytics. In *CIDR*, 2017.

36. J. Paul, J. He, and B. He. GPL: A GPU-based pipelined query processing engine. In *SIGMOD*, pages 1935–1950. ACM, 2016.

37. H. Pirk, S. Madden, and M. Stonebraker. By their fruits shall ye know them: A data analyst's perspective on massively parallel system design. In *DaMoN*, pages 5:1–5:6. ACM, 2015.

38. H. Pirk, S. Manegold, and M. Kersten. Waste not... efficient co-processing of relational data. In *ICDE*. IEEE, 2014.

39. H. Pirk, O. Moll, M. Zaharia, and S. Madden. Voodoo - a vector algebra for portable database performance on modern hardware. *PVLDB*, 9(14):1707–1718, 2016.

40. V. Raman and others. DB2 with BLU acceleration: So much more than just a column store. *PVLDB*, 6(11):1080–1091, 2013.

41. J. Rao, H. Pirahesh, C. Mohan, and G. Lohman. Compiled query execution engine using JVM. In *ICDE*, pages 23–. IEEE, 2006.

42. H. Rauhe, J. Dees, K.-U. Sattler, and F. Faerber. Multi-level parallel query execution framework for CPU and GPU. In *ADBIS*, pages 330–343. Springer, 2013.

43. S. Richter, V. Alvarez, and J. Dittrich. A seven-dimensional analysis of hashing methods and its implications on query processing. *PVLDB*, 9(3):96–107, 2015.

44. V. Rosenfeld, M. Heimel, C. Viebig, and V. Markl. The operator variant selection problem on heterogeneous hardware. In *ADMS@VLDB*. VLDB Endowment, 2015.

45. C. Rossbach, Y. Yu, J. Currey, J.-P. Martin, and D. Fetterly. Dandelion: A compiler and runtime for heterogeneous systems. In *SOSP*, pages 49–68. ACM, 2013.

46. B. Răducanu, P. Boncz, and M. Zukowski. Micro adaptivity in Vectorwise. In *SIGMOD*, pages 1231–1242. ACM, 2013.

47. A. Shaikhha, Y. Klonatos, L. Parreaux, L. Brown, M. Dashti, and C. Koch. How to architect a query compiler. In *SIGMOD*, pages 1907–1922. ACM, 2016.

48. J. Sompolski, M. Zukowski, and P. Boncz. Vectorization vs. compilation in query execution. In *DaMoN*, pages 33–40. ACM, 2011.

49. Transaction Processing Performance Council. TPC Benchmark H, 2012.

50. S. Wanderman-Milne and N. Li. Runtime code generation in cloudera impala. *IEEE Data Eng. Bull.*, 37(1):31–37, 2014.

51. H. Wu, G. Diamos, S. Cadambi, and S. Yalamanchili. Kernel weaver: Automatically fusing database primitives for efficient GPU computation. In *MICRO*, pages 107–118. IEEE, 2012.

52. Y. Yuan, R. Lee, and X. Zhang. The yin and yang of processing data warehousing queries on GPU devices. *PVLDB*, 6(10):817–828, 2013.

53. M. Zahran. Heterogeneous computing: Here to stay. *Commun. ACM*, 60(3):42–45, 2017.

54. S. Zeuch, H. Pirk, and J. Freytag. Non-invasive progressive optimization for in-memory databases. *PVLDB*, 9(14):1659–1670, 2016.

55. K. Zhang, F. Chen, X. Ding, Y. Huai, R. Lee, T. Luo, K. Wang, Y. Yuan, and X. Zhang. Hetero-DB: Next generation high-performance database systems by best utilizing heterogeneous computing and storage resources. *J. Comput. Sci. Technol.*, 30(4):657–678, 2015.

56. S. Zhang et al. OmniDB: Towards portable and efficient query processing on parallel CPU/GPU architectures. *PVLDB*, 6(12):1374–1377, 2013.

57. J. Zhou and K. Ross. Implementing database operations using SIMD instructions. In *SIGMOD*, pages 145–156. ACM, 2002.