

# LoopBench: An Evaluation of Loop Acceleration in Heterogeneous Systems

Saman Biokaghazadeh  
School of Computing  
Arizona State University  
Tempe, US  
sbiokag@asu.edu

Fengbo Ren  
School of Computing  
Arizona State University  
Tempe, US  
renfengbo@asu.edu

Ming Zhao  
School of Computing  
Arizona State University  
Tempe, US  
mingzhao@asu.edu

**Abstract**—Computational intensive applications usually consist of multiple nested or flattened loops. These loops are the main building blocks of the applications and embody a specific type of execution pattern. In order to reduce the running time of the loops, developers are analyzing the loops in the code and try to parallelize them on the target hardware accelerators in a heterogeneous system, either spatially or temporally. Unfortunately, the lack of understanding of loop characteristics and the ability of hardware accelerators in handling these types of loops prevents application developers from choosing the right platform. In addition, developing an accelerator specific code is a time-consuming effort. To address this issue, we have developed LoopBench, which is a benchmarking tool to assess the effectiveness of available processors, in accelerating different common patterns of loops. LoopBench includes six important types of loops that commonly exist in real-world applications. Further, it evaluates different processors in accelerating these loop patterns. The result from LoopBench explains architectural differences between different accelerators with regard to different loop patterns. In addition, it provides insights for the developers to choose the right accelerators for their applications, before any coding. The current version of our benchmark supports both Field-Programmable Gate Arrays (FPGAs) and Graphics Processing Units (GPUs), which are the most versatile and available accelerators.

**Index Terms**—Heterogeneous, OpenCL, FPGA, GPU, Loop, Workload Characterization

## I. INTRODUCTION

Many applications can benefit from computing on hardware accelerators, ranging from high-performance computing (HPC) and cloud computing to big data and edge computing. Examples of these applications include (1) analysis of large quantity of data on big-data platforms, (2) training and running artificial intelligence (AI) and machine learning models in the cloud, (3) processing streams of requests and data from IoT devices, and (4) modeling and simulating the behaviors of scientific applications on HPC systems. By using accelerators, applications can achieve higher throughput [1], lower response time [2], and/or lower energy consumption [3].

A variety of accelerators are readily available for applications to choose for their computation needs. Graphics Processing Units (GPUs) are the most widely used and can be easily found in many HPC and cloud systems. Other types of accelerators are also becoming increasingly available, e.g., Tensor Processing Units (TPUs) [4] on the Google cloud and

Field-Programmable Gate Arrays (FPGAs) on the Amazon cloud (F1 nodes). These accelerators come with different capabilities and limitations. For example, FPGAs can be reconfigured to run any kinds of applications but can provide only low clock frequency; GPUs can be programmed using high-level languages to accelerate highly parallel applications; and TPUs are specifically designed for deep learning workloads. Although a general understanding of different accelerators is available, choosing the right accelerators for sophisticated applications is still a difficult problem.

Several related works have studied the performance of common algorithms on accelerators. For example, Rodinia benchmark [5] and its follow-up work [6] are designed to benchmark heterogeneous platforms including CPUs, GPU, and FPGAs. These benchmarks usually provide insights on a macro level, for a complete algorithm on a hardware platform, but they lack a thorough analysis of micro-level execution patterns that exist in different applications and the effectiveness of different hardware architectures in handling these patterns.

In order to address the above challenges, we study how the accelerators with different hardware architectures can accelerate different types of loops which are the basic building blocks of almost every computationally intensive application. These applications typically consist of one or many nested and flattened loops. These loops can embody different patterns in terms of types and degrees of dependency and concurrency, and some of these patterns can be found in many applications. For example, dynamic programming algorithms consist of one or more nested loops, where every iteration depends on another iteration that points diagonally in the iteration space. Therefore, abstracting the common loop patterns from applications and understanding how they perform on various hardware accelerators are essential steps towards optimally utilizing the accelerators for executing different applications.

Following this approach, we present a new benchmarking tool, named *LoopBench*, which includes five fine-grained loop patterns that commonly exist in real-world applications such as linear algebra, optimization, and data analytics algorithms. LoopBench parameterizes the key aspects of these loop patterns, including the type and degree of dependencies, operational intensity, and size of the iteration spaces. It allows them to be flexibly tuned to model diverse loop characteristics.

LoopBench provides optimized OpenCL implementations of these loop patterns for both GPU and FPGA. We focus on OpenCL because it is an important framework for the emerging heterogeneous computing paradigm.

In summary, LoopBench has demonstrated several key observations. First, for three out of five loop dependency patterns (intra-dimension dependency, conditional dependency, and half-parallelism half-dependency), the FPGA has the potential to outperform the GPU. For example, for the intra-dimension dependency pattern, FPGA outperforms GPU by 17.5x. Second, it shows that for various computational intensities, the FPGA can maintain an identical performance, whereas the GPU performance is highly variable. For example, having eight conditional statements can degrade the GPU performance by up to 45%. Third, it shows that increasing the input data size can increase the performance difference between these two accelerators. For example, for the diagonal dependency loop pattern, the performance gap increases by 51%, while changing the input data size from 4MB to 256MB.

In summary, here is the list of our contributions in LoopBench:

- Identification and classification of common loop patterns in computational intensive applications.
- Optimization of these loop patterns on the OpenCL-enabled FPGAs.
- Evaluation of the acceleration potentials of these loop patterns on two different accelerators, with regard to key configuration parameters, such as computational intensity, dependency and concurrency degrees, and input data size.

## II. BACKGROUND

### A. Accelerators

Acceleration is becoming a critical de-facto for many computationally intensive workloads on various computing systems. While there are different accelerators available, such as GPUs, FPGAs, TPUs, and DSPs, we focus on GPUs and FPGAs, since they are more general purpose than the others and can support a wide variety of applications.

**GPUs** have been well studied and widely used as accelerators. While GPUs are highly effective in handling applications with high level of concurrency and regular memory access patterns, they come short for applications with low computational intensity, a high degree of dependency, and/or a high number of conditional branches. Examples of these applications include graph processing [7], sorting [8], small signal processing problems [9], and sparse linear algebra [10].

**FPGAs** are each a farm of logic, computation, and storage resources that can be configured dynamically. Different from widely-adopted GPUs, FPGAs can accelerate almost all types of algorithms (irrespective to their computational pattern), due to their reconfigurability. Several related works have demonstrated the usability of FPGAs to accelerate HPC applications, using hardware description languages (HDL), such as VHDL and Verilog [11]. Despite their impressive acceleration power, programming and optimization difficulties have been a serious

obstacle to their wider adoption. Recent FPGA advancements in supporting high-level synthesis (HLS) have made it possible to program FPGAs using high-level languages, especially OpenCL [12], which has made FPGAs much easier to use and much more accessible to applications. Even though an HLS-based program may not perform as well as a carefully hand-crafted HDL program, the productivity enabled by HLS is often far more important.

### B. OpenCL

OpenCL is a versatile C-based programming model that can execute across heterogeneous platforms, including CPUs, GPUs, and Digital Signal Processors (DSPs). CPU and GPU vendors, such as Intel, AMD, and NVIDIA have been supporting OpenCL on their platforms for over a decade. The recently-extended support of OpenCL to FPGAs has opened the gate for conveniently integrating FPGAs into the heterogeneous computing paradigm. Using OpenCL, programmers do not need to make any major changes to their code, when porting it across different platforms. Moreover, developers can split their applications and deploy the parts on different accelerators to make optimal use of the accelerators' different capabilities.

OpenCL's ease of programming and portability across platforms unlock a whole new level of productivity, even though it might lose some performance compared to the traditional frameworks for accelerator programming. Compared to CUDA-based GPU programming, related works have shown OpenCL has only a slightly worse performance on GPUs [13]–[16]. For example, the optimized SGEMM routine on OpenCL [13] performs only 12.28% slower than the same routine on CUDA. Compared to C-based HLS on FPGAs, OpenCL-based FPGA programming has about the same level of performance. For example, consider the disparity map calculation algorithm [17]. For the window size of  $7 \times 7$  the OpenCL implementation is faster than the C one by 6.14%; For the window size of  $9 \times 9$ , the OpenCL implementation can be slower by up to 15.3%.

Therefore, in this study, we focus on the OpenCL-based GPU and FPGA computing and study their effectiveness for accelerating common algorithmic patterns.

### C. Loop Parallelism

Algorithms are composed of one or many loops, either nested or flattened. The acceleration of algorithms is the process of acceleration of the loops, using parallelization and pipelining methods. Algorithms can be parallelized either *temporally* or *spatially*.

**Spatial Parallelism.** In *spatial* parallelism [18], processing elements (PEs) are executing the same task (SIMD) or multiple different tasks (MIMD), simultaneously. Both GPU and FPGA are able to exploit spatial parallelism in algorithms. The amount of data dependency between the iterations of the loops in the algorithm can decide the level of achievable spatial parallelism on the target architecture. In another word, having less data dependency increases the opportunity of speedup on parallel architectures, such as GPUs and FPGAs. In general,

GPUs are better at exploiting spatial parallelism, because FPGAs cannot adopt as many compute cores as GPUs, and FPGAs also tend to operate at a lower clock frequency, up to 2-5 times slower than GPUs.

**Temporal Parallelism.** In *temporal* parallelism [18], processing tasks that have a dependency on each other are mapped onto different PEs and execute in parallel in a pipeline fashion. Data processing has multiple stages, and each stage is being handled by one PE. In this multi-stage pipeline, as data is processed by the element  $PE_i$ , it is sent to the next element  $PE_{i+1}$  and element  $PE_i$  moves on to handle new data coming from the previous stage. In the cases where a single task cannot fully occupy the available PEs, multiple tasks can be interleaved and mapped onto the PEs to increase the *temporal* parallelism. Among general purpose accelerators, FPGAs are exclusively able to exploit coarse-grained temporal parallelism in the algorithms, due to their reconfigurability. SIMD platforms like GPU can perform at most one instruction at a time on each available core, whereas FPGA can execute hundreds of operations on all available stages in the pipeline.

Figure 1 depicts both parallelism dimensions. Each circle represents an individual iteration in a set of nested loop blocks. The  $(i, j)$  pair in each circle represents the  $i$ th iteration in the first dimension and the  $j$ th iteration in the second dimension. The arrow represents the dependency of one iteration on another, e.g.,  $(1,2)$  depends on  $(1,1)$ . Each iteration usually involves separate calculation for a specific indexed item or accumulation on a shared value among iterations of a loop block. The dashed box contains iterations with zero dependency, which can be easily parallelized spatially. On the other hand, the dotted box contains iterations with data dependency, which cannot be parallelized spatially but may have the potential to be parallelized temporally. We use the above format throughout the paper to represent the dependency flow in the loops iteration space.

In summary, GPUs excel at exploiting spatial parallelism but cannot utilize temporal parallelism, whereas FPGAs can take good advantage of both types of parallelism. However, despite this general understanding of GPU and FPGA’s different strengths, it is still difficult to understand which accelerator works the best for which algorithm. Every single application consists of different types and degrees of conditional and data dependencies. Developers usually need to implement the code for different accelerators and then apply several different transformations on the algorithm to assess the acceleration potentials on different devices. Understanding the relationship between common micro-level patterns such as loop patterns and their potential acceleration can reduce the effort of choosing the right device. These are the motivations for our study on loop acceleration using GPUs and FPGAs, which, to the best of our knowledge, is the first.

Several related works [2], [19] have studied the acceleration of specific applications on FPGAs and compared it to GPUs. Other related works [20], [21] have proposed frameworks for the application runtime prediction. However, these works do

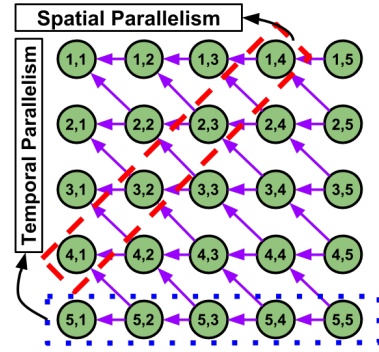


Fig. 1: Spatial and temporal parallelism in multiple iteration dimensions.

not provide a thorough insight into the correlation between loop patterns and modern accelerators, and the extent of potential acceleration of different loops on a target accelerator. The lack of such insights has motivated the development of LoopBench. (See Section IV for a detailed examination of the related works.)

### III. LOOP ANALYSIS

#### A. Methodology

Our approach to understanding how to choose the optimal accelerator for a given algorithm is by studying the performance characteristics of common loop patterns on GPUs and FPGAs. Following this approach, we designed *LoopBench*, a new benchmark suite that captures the key loop patterns extracted from real-world algorithms, and allows flexible testing of each type of loops by varying the following key parameters:

- 1) *Computational intensity*, which is the total number of computational operations that each iteration of the algorithm performs. In our benchmark, it is defined as the number of cosine functions. The computational intensity can affect the size of the pipeline and the number of instructions on both FPGA and GPU. Changing this parameter can show how both platforms performances are susceptible to the amount of computation;
- 2) *Dependency and concurrency degrees*, which defines how many iterations depends on each other and how many other iterations can be executed separately.
- 3) *Input data size*, which specifies the total number of floating-point variables that the algorithm processes. The size of the input data can affect the load of computation on a target platform, which can decide the suitability of one device over another.

LoopBench includes optimized implementations of each loop type for GPU and FPGA. The rest of this section details each loop type and its GPU and FPGA implementations, and presents experiments from running them on real devices. While optimizing GPU programming has been well studied, OpenCL-based FPGA optimization is not well explored and not trivial. In our discussions, we will also detail how we performed the optimizations for each key loop type.

All GPU-related experiments were conducted on two server nodes with two type of GPUs. One server is equipped with an

Nvidia Tesla K40m GPU, dual Intel Xeon E5-2637 v4 CPU, and 64GB of DDR4 main memory (2133MHz). The other server is equipped with an Nvidia Geforce Titan X, Intel Xeon E5-2650 v3 CPU, and 198GB of main memory. All the FPGA-related experiments were conducted on an Intel Fog Reference Design unit [22], equipped with two Nallatech 385A FPGA Acceleration Cards (Intel Arria 10 GX1150 FPGA), and Intel Xeon E5-1275 v5 CPU, and 32GB of DDR4 main memory (2133 MHz). Need to mention that the difference between these two hosts does not affect the results since we measured only the benchmark’s runtime on the devices. The OpenCL kernels for FPGAs were compiled using Intel FPGA SDK for OpenCL (version 16.0) with Nallatech p385a\_sch\_ax115 board support packages (BSP). The GPU OpenCL kernels were compiled just-in-time at runtime using available OpenCL library in CUDA Toolkit 9.0. For the FPGA, we implemented all the kernels in the single-thread mode. Single-thread kernels on the FPGA typically have much less overhead and can achieve much higher clock frequency rate, compared to multi-threaded kernels. On the GPU, we implemented the kernels in the NDRange mode which in the OpenCL will deploy concurrent threads on the available compute units.

### B. Intra-Dimension Dependency

**Definition.** This type of loops is usually composed of two or more nested iterative blocks, where each level of iterative blocks is considered a *dimension*. In this pattern there exist a *loop-carried data dependency*, which is a dependency of one iteration on the output of the previous iterations (read-after-write), in one or more dimensions, while at the same time one or more dimensions have no dependency between their iterations. In another word, we can observe both dependency and concurrency in the overall iteration space.

For example, in Algorithm 1, the dependency exists between iterations with the index of  $i$ . In this algorithm, updating every element of the array  $A$  with the index of  $i$  on the first dimension depends on the value of the element with the index of  $i - 1$ . Elements in the second dimension with the index of  $j$  do not carry any dependency. In this case, the dependency exists on the dimension with the index of  $i$  and the concurrency exists on the dimension with the index of  $j$ . Figure 2 illustrates the iteration space and the dependency graph of intra-dimension dependent loops. Although in this example, the nested loops have only two dimensions, indexed by  $i$  and  $j$ , in reality, the algorithm can have multiple dimensions and dependency within any one of the dimensions.

Simple linear algebraic algorithms [23], such as matrix-matrix or matrix-vector multiplications are following this type of loop pattern. For example, in matrix-matrix multiplication, each cell of the output matrix can be computed separately (concurrency), while the dot multiplication of one row and one column can only be performed sequentially in a single thread (dependency).

The degree of spatial and temporal parallelism, combined with the arithmetic intensity, can determine the choice of deployment on either FPGA or GPU. Algorithms with a high

---

### Algorithm 1 Intra-dimension dependency algorithm

---

```

 $i \leftarrow 1$ 
 $j \leftarrow 1$ 
for  $i \leq n$  do
  for  $j \leq m$  do
     $A[i][j] = \text{func}(A[i-1][j], B[i][j], \dots)$ 
  end for
end for

```

---




Fig. 2: Intra-dimension dependent loop pattern

degree of dependency can usually finish faster on FPGAs, while algorithms with a high degree of concurrency can utilize the available farm of SIMD compute units on the GPUs and accelerate their execution.

**Implementation.** Our benchmark contains the GPU and FPGA versions of the intra-dimension dependent loop. For the GPU version, the loop is unrolled spatially over the non-dependent dimension. Each independent iteration is deployed as a work-item (unit of a task in the OpenCL), and the total number of work-items are grouped into several work-groups (unit of execution on a single compute unit). Also, we specifically order the memory access indexes to enable memory access coalescing among work-items in a work-group for better performance. For the FPGA version, we first apply statement re-ordering to place the dependent loop as the inner-most loop, which enables interleaving of the outer-loop iterations (non-dependent) inside the inner-loop cycle. It also helps achieve the initiation interval of one in the inner-most loop. In loop pipelining, the initiation interval is the number of clock cycles between the start times of consecutive loop iterations. Having an initiation interval of one enables the FPGA to push one iteration into the pipeline at every clock cycle and achieve the highest performance, which is the ultimate goal for every design. Further, we apply *loop blocking* (also known as loop tiling) on the outer for loop. Doing so enables allocating non-chip registers of size block and copy the required data for all the iterations into that block, as a whole into the on-chip register (of the size of the block) and reduce the contention on the DRAM.

**Experiment.** We deployed FPGA and GPU kernels, resembling Algorithm 1. Input data is an array of floating-point variables of a specific size (4, 32, 256 MB). Every single iteration in the algorithm is responsible for a single element in the array. As a result, the total number of iterations is equal

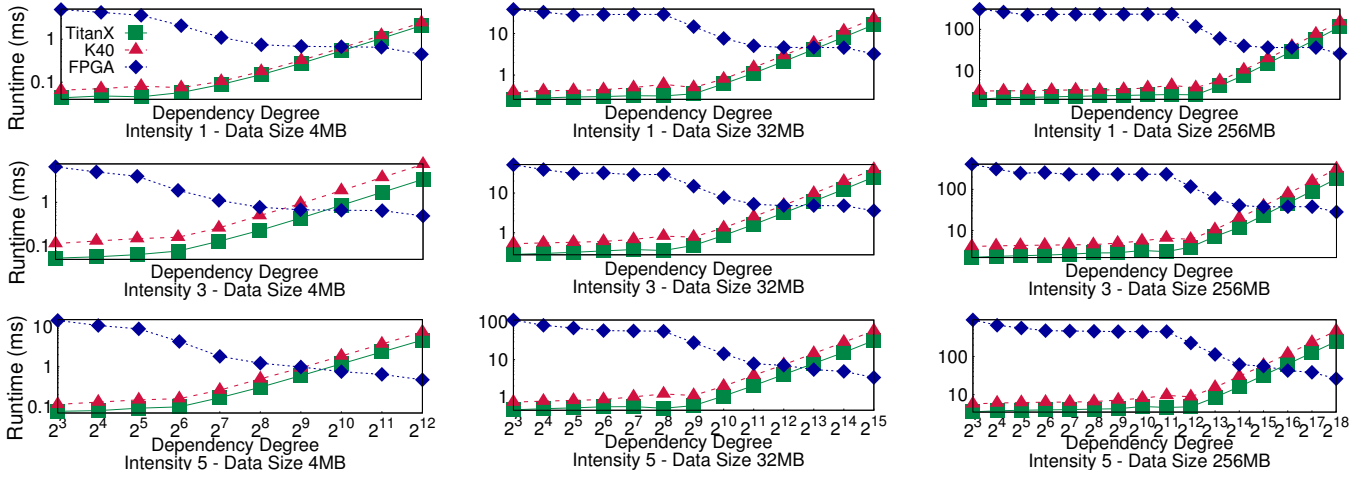


Fig. 3: Intra-dimension dependency performance on the GPU and the FPGA

to the number of input values. As shown in the algorithm, the dependency and concurrency degrees are configured by changing the number of iterations,  $n$  and  $m$ , respectively. Figure 3 shows the runtime of this intra-dimension dependent loop on both FPGA and GPU.

We can make several key observations from the results. First, GPU does excel at accelerating the loop with a high degree of concurrency. More concurrency can lead to better spatial parallelization, which makes the GPU a great candidate for deployment. In contrast, with the increase in the dependency degree, the FPGA can take advantage of the configured long pipeline and parallelizing the dependent iterations. In this case, with a high degree of dependency, the FPGA can outperform both Tesla K40 and Titan X up to 17.5x and 9.6x. With a high degree of concurrency, both Tesla K40 and Titan X perform better than the FPGA, by up to 154x and 247x, respectively.

The second observation is about the effect of computational intensity (the total number of computational operations in each iteration) on the final performance. Higher intensity means more computations, which leads to more pipeline stages. With more pipeline stages, FPGA can handle more dependent iterations and achieve higher performance. Need to mention, the available hardware resources on the FPGA are limited and may block developers from configuring a large number of pipeline stages. As a result, developers may need to adopt a smaller loop block size, which leads to the reduction of the performance. Compared to FPGA, the GPU has to spend more time executing each loop iteration, with no opportunity for pipelining the iteration pipelining. For example, Figure 3 shows that going from the intensity of 1 to 5, the performance drops by up to 3x and 2.1x, on Tesla K40 and Titan X, respectively.

The third observation is the performance reduction of FPGA for kernels with low dependency, because there are not enough dependent iterations to fully saturate the configured pipeline. In this situation, developers may want to switch into the NDRange mode kernels, which can interleave the parallel iterations into the pipeline and keep it saturated. In comparison,

GPU can utilize the massive farm of cores to exploit a large degree of parallelism when the dependency is low. Therefore, as shown in Figure 3, FPGA’s performance is worse with lower dependency degree whereas GPU’s performance is not affected.

### C. Diagonal Dependency

**Definition.** Diagonal dependent loops are following almost the same pattern as intra-dimensions dependent loops, except that the dependency is diagonal instead of horizontal or vertical in the iteration space. As illustrated in Figure 4, horizontal (vertical) dependency refer to the dependency of an iteration on the left (top) neighbor iterations with the same  $i$  ( $j$ ), respectively. For example, in the aforementioned intra-dimension dependency, there is horizontal dependency among the iterations as shown in Figure 2. Diagonal dependency means that an iteration depends on its relative top-left iteration which has both different  $i$  and  $j$  indexes. For example, in Figure 4, iteration (2, 2) depends on its diagonal neighbor iteration (1, 1). Algorithm 2 shows an example of this kind of loops, where the computation requires data from its diagonal neighbor in the iteration space. In specific cases, the dependency can be extended and include either horizontal or vertical, as well.

Parallelization of these types of loops on SIMD architectures, such as GPU, is not straightforward. Depending on the type of diagonal dependency, developers can either parallelize the diagonals or use the wavefront technique [24] for parallelization. In the wavefront parallelism mode, kernels are enqueued back to back to the GPU, each computing one set of independent iterations. The number of the kernels is equal to the length of the diagonal.

Dynamic programming algorithms [25] are usually composed of diagonal dependent iterations. A specific example of such algorithm is Needleman-Wunsch [26], which performs matching between two input strings while minimizing the penalty.

**Implementation.** For the GPU implementation, the parallelization method depends on the existence of vertical or



**Algorithm 2** Diagonal dependency algorithm

---

```

i ← 1
j ← 1
for i ≤ n do
  for j ≤ m do
     $A[i][j] = \text{func}(A[i-1][j-1], B[i][j], \dots)$ 
  end for
end for

```

---




Fig. 4: Diagonal dependency loop pattern

horizontal dependency. In the absence of both of these dependencies, each thread can take care of one diagonal, in parallel. The existence of any of the mentioned dependencies (in addition to diagonal dependency) would force the GPU to perform *anti-diagonal parallelism*. As shown in Figure 4, the independent iterations that can be parallelized form a line that is perpendicular to the diagonal dependent iterations.

For the FPGA implementation, we first performed loop blocking on the first dimension, which enables caching of the input data for each iteration of the second dimension’s iterations. Later, we copy the required data for the second dimension’s computation into the allocated on-chip registers of the size block. Every iteration of the second dimension first reads the data from the registers, performs the calculation, and writes back the data to the registers and the DRAM. To handle all elements in the block, each iteration of the second dimension contains a nested loop of size block, which is fully unrolled. In this implementation, the iterations of the second dimension have a loop-carried data dependency. Unfortunately, the compiler cannot infer an initiation interval of one for this loop body, due to the existence of large latency. To overcome this issue, we interleaved the execution of the block iterations inside the second dimension loop, which enables full exploitation of the available pipeline stages. Doing so reduces memory accesses and leads to higher operating frequency and fewer stalls in the pipeline.

**Experiment.** Figure 5 shows the performance of the diagonal dependent loops on the FPGA and the GPUs, where the dependency only exists diagonally. We did measurements for three different computational intensities (1, 3, and 5) and three different input sizes (4, 64, and 512 MBs). Based on the results, the GPU outperforms the FPGA in almost all cases, except for the experiment with high computational intensity and small data size. In this type of dependency, GPU can

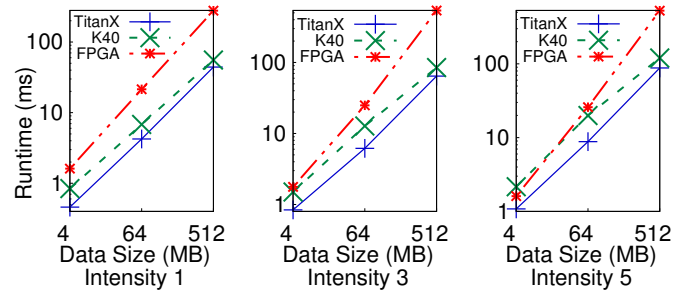


Fig. 5: Diagonal dependency runtime on both FPGA and GPU. The dependency is only diagonal.

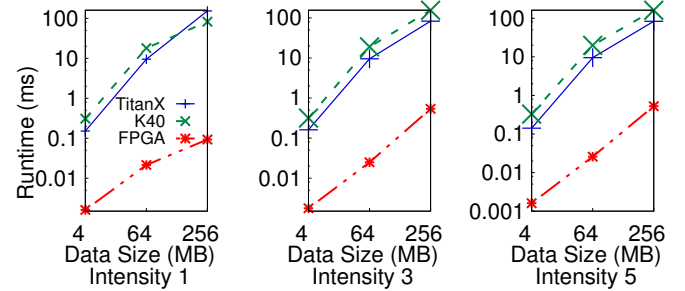


Fig. 6: Diagonal dependency runtime on both FPGA and GPU. The dependency also includes horizontal and vertical.

assign one diagonal set of iterations to one work-item and exploit high degree of parallelism on all the available cores. In this case, Tesla K40 and Titan X are outperforming the FPGA by up to 4.3x and 6x, respectively.

Figure 6 shows the performance of the same loop pattern but with additional horizontal and vertical dependencies between the iterations. We modified the function  $f$  in Algorithm 2 to include both  $A[i-1][j]$  and  $A[i][j-1]$ , in addition to  $A[i-1][j-1]$ , as its parameters to introduce these dependencies between  $A[i][j]$  and its horizontal, vertical, and diagonal neighbor iterations. In this case, the FPGA can utilize the same pipelining method to accelerate the execution, while both GPUs need to use wavefront parallelism and parallelize computation for each anti-diagonal. Unlike the case with diagonal dependency, the wavefront parallelism model cannot exploit a large number of parallel threads. In addition, it needs to repetitively deploy the same kernel to calculate a new set of anti-diagonal iterations. As a result, the FPGA outperforms both Tesla K40 and Titan X by up to 322x and 165x, respectively.

**D. Conditional Dependency**

**Definition.** The existence of conditional statements in loop bodies can alter the extent of parallelization on certain accelerators. In loops with a conditional statement, every iteration diverges in the execution path, depending on the specific conditions. Algorithm 3 represents an example, where every iteration performs either the first or the second statement based on the content of an array in that specific iteration index.

Algorithms such as K-means and single-source shortest path (SSSP) consist of many conditional decisions. In the K-means, the clustering of the observations requires many

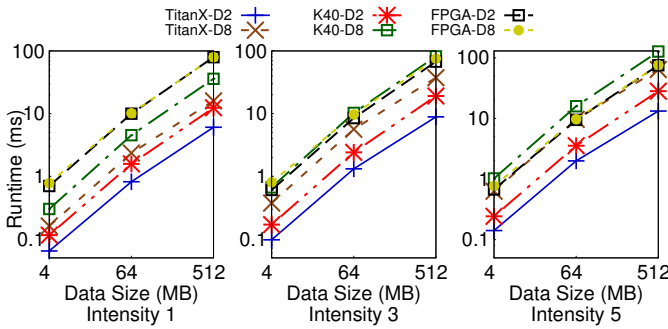


Fig. 7: Conditional dependency runtime on both FPGA and GPU, for different intensities.

comparisons, based on the distance; SSSP relies on the sparse matrix multiplication, where the number of iterations for each output calculation is non-deterministic.

---

#### Algorithm 3 Conditional dependency algorithm

---

```

i ← 1
for i ≤ n do
  if B[i] > 0.0f then
    A[i] = f(B[i], D[i], ...)
  else
    A[i] = f(C[i], D[i], ...)
  end if
end for

```

---

**Implementation.** The conditional dependency is introduced by an if-else statement in the kernel. On the GPU, the loop is simply parallelized on different cores, and each thread performs the if-else comparisons. But the SIMD architecture in the GPU cannot efficiently handle the conditional statements in the work-items, due to thread divergence issue. In the FPGA implementation, the kernel is developed in a single-thread mode and the loop is unrolled to the limit of the FPGA area and available DRAM bandwidth. In contrast to the GPU implementation, FPGAs can handle numerical conditional statements, using look-up tables and a simple multiplexer. More specifically, the FPGA can map all different paths of the execution in the design and enables different threads running simultaneously in different conditional blocks.

**Experiment.** Figure 7 shows the runtime of the conditional dependent loop on the GPU and FPGA, with various computational intensities (one, three, and five) number of conditional branches (two and eight) within each iteration as well as various total input data sizes (4, 64, and 512 MB). The number of conditional statements is represented as D2 and D8, for two and eight conditional decisions, respectively. The results show that the FPGA can sustain the same performance among kernels with different conditional branches, whereas the GPU suffers more performance degradation for kernels with more conditional branches (up to 45% slowdown). As a result, the FPGA outperforms the GPU with a higher number of conditional dependencies; e.g., 40% better for a dependency level of eight. This observation suggests the suitability of FPGAs for algorithms with a high degree of decision making

during the execution. These types of applications usually cannot exploit the massive parallelism in SIMD architectures, thus better handle them with reconfigurable processors such as FPGAs.

#### E. Anti-dependency

**Definition.** In this loop pattern, every iteration consists of more than one statements. Unlike the intra-dimension dependent loops, where the dependency is read-after-write, this pattern carries write-after-read dependency. In this pattern, one statement of an iteration reads a data item that is going to be updated by the other statement in the next iteration. It is named anti-dependency because the statements in different iterations are following the write-after-read pattern, as opposed to read-after-write in the typical dependency patterns. Algorithm 4 demonstrates a general example of such loops. The existence of read-after-write dependency creates an anti-dependent loop pattern.

Anti-dependent loops are carrying a unique feature. It is possible to face race condition in case of parallelization of all the iterations. More specifically, the first iteration reads the old value of an array element (e.g.,  $A[i]$  depends on  $B[i+1]$  in Algorithm 4), while the second iteration updates the same value, and so on and so forth. When these iterations are executed on different threads to achieve parallelism, the dependent read and write might be executed out of order, which damages the correctness.

---

#### Algorithm 4 Anti dependency algorithm

---

```

i ← 1
for i ≤ n do
  A[i] = B[i + 1] + C[i] * D[i]
  B[i] = B[i + 1] - E[i] * D[i]
end for

```

---

**Implementation.** These types of loops can be parallelized on vector processors with a global barrier mechanism among all SIMD threads. Unfortunately, both the FPGA and the GPU lack such global barrier mechanism between all threads. An approach for parallelizing inter-iteration dependent loops is loop-splitting. In this approach, the loop can be divided into multiple separate loops, where none of them carries any dependency. In this situation, loops should run sequentially on the target processor (to guarantee the correctness of the execution), but each loop can fully exploit the available spatial core units. Figure 8 represents the execution and the dependency of the original loop, along with the transformed version of it. The dotted blue box and the solid red box represent different statements in the loop body. The arrow shows the anti-dependency between different statements of consecutive iterations.

To implement this type of loop on GPU and FPGA, we applied statement re-ordering and loop splitting. The transformation creates multiple flattened loops, where each of them represents a stage of the execution. The lack of global barriers prevents both platforms from co-locating the execution




Fig. 8: Anti dependency loop pattern.

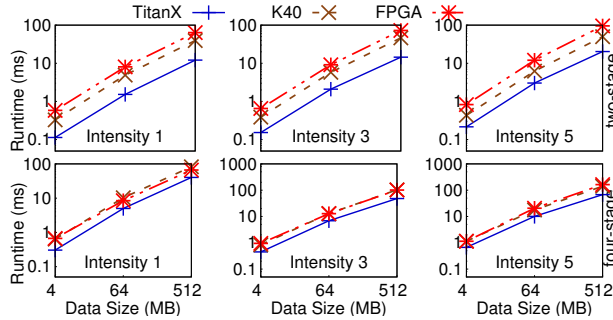


Fig. 9: Anti-dependency results for two and four stages.




Fig. 10: Half-parallelism half-dependency loop pattern.

of the generated sub-loops after the main loop distribution, except for using channels in FPGA which is a mechanism for passing data between kernels and synchronizing kernels with high efficiency and low latency. Usually, kernels need to communicate through DRAM, which increases the application runtime. By using channels, loops can start pipelining their partial results to the next loop, which enables co-location of the computation and communication and reduces the application runtime. Unlike the FPGA, GPU should execute the flattened loop sequentially, but each stage can be fully parallelized spatially.

**Experiment.** Figure 9 shows the runtime of the FPGA and GPU in accelerating these loops. We varied the degree of anti-dependency which is the number of statements involved in the anti-dependency. For example, in Algorithm 4 the dependency exists between two statements, which yields into the anti-dependency degree of two. As a result, the main loop in the benchmark can be split into several separate and parallelizable loops, depending on the number of anti-dependent statements in the loop body. We also varied the intensity level and input data size.

Comparing the runtimes for the case of four anti-dependencies, the FPGA can outperform the Tesla K40 GPU for kernels with low intensity (up to 20% speedup), whereas it performs close to the GPU for higher intensities (up to 15% speed degradation). Comparing to Titan X, the FPGA performs 1.6x slower. Kernels with higher intensities lead into larger area consumption and limit the parallelism level in each stage, which further results in the reduction of the channels widths. Need to mention, increasing the channel width can directly affect the total required resources on the FPGA. Overall, higher intensities in the inter-iteration dependent type of loops reduce the chance of outperforming GPUs.

Increasing the number of statements with anti-dependencies will result in more separate loops. Based on Figure 9, increasing the degree of anti-dependency reduces the gap between the FPGA and GPU. We can expect that by following this trend, the FPGA will eventually outperform the GPU.

#### F. Half-Parallelism Half-Dependency

**Definition.** Half-parallel half-dependent loops usually include the dependent and the parallel statements, simultaneously,

and are composed of only one loop, with no nested loop. Algorithm 5 lists an example of this type of loops. The existence of loop-carried dependent statements (read-after-write) prevents the spatial parallelization of the algorithm, as a whole. Transforming the loop into multiple flattened loops enables the execution of the loop in two different stages. Dissimilar to anti-dependent loops, the loop-splitting process does not enable spatial parallelization opportunity for all the loops, since part of the algorithm carries read-after-write dependency. After the splitting, the parallel portion of the loop can be deployed on processors with a high number of parallel compute units, e.g., GPUs, while the dependent portion can be handled by processors that are suitable for sequential execution, e.g., on CPUs and FPGAs.

Applications such as K-nearest neighbor (KNN) are composed of both parallel parts (distance computation) and dependent parts (sorting). These applications can utilize one or more hardware accelerators for an efficient acceleration.

**Implementation.** We applied loop splitting to separate the parallel section from the dependent section. For the GPU, we first compute the parallel part on the GPU and then transfer the data back to the main memory of the host and execute the dependent part on the CPU. Running the dependent block of code on the GPU is not efficient and will lead to poor performance. For the FPGA we have multiple options, (1) running the parallel and dependent blocks of the loop serially on the FPGA, (2) running the parallel block on the FPGA and the dependent block on the CPU, and (3) using channel to pipeline the intermediate result from the parallel part to the dependent part and decrease the running time overhead. Using the channels is the best available option to co-locate computation and communication, and achieve the highest possible performance.

---

#### Algorithm 5 Half-parallelism half-dependency algorithm

---

```

i ← 1
for i ≤ n do
  A[i] += C[i] * D[i]
  sum += B[i] + A[i] + D[i]
end for

```

---

**Experiment.** Figure 10 represents the half-parallelism half-



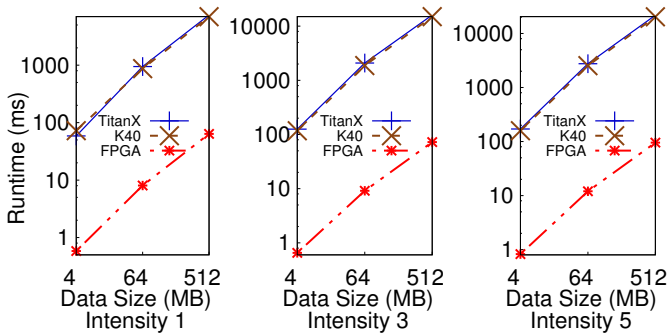


Fig. 11: Half-parallelism half-dependency runtime on both FPGA and GPU, for different intensities.

dependent loop pattern. For this pattern, each red box in an iteration depends on another red box from the previous iteration. Furthermore, Each red box depends on the value of the blue box in the same iteration. For this experiment, we provided input data with a size of 1 to 1024 MB. The FPGA can outperform both Titan X and Tesla K40 GPUs, by up to 118x and 110x, respectively. The overhead of the data transfer from the GPU to CPU reduces both GPUs performance significantly. It is worth noticing that Tesla K40 experiment is running faster than the Titan X because it is hosted on a host with a faster CPU. As a conclusion, co-locating the parallel and the dependent sections of the code on the FPGA can yield into much higher performance, compared to utilizing two different types of accelerators with a much slower communication channel.

#### IV. RELATED WORK

To the best of our knowledge, LoopBench is the first to provide a comprehensive study of common loop patterns on important hardware accelerators, including both GPUs and FPGAs. There are a number of related works that are complementary to the focus of LoopBench. Roofline modeling [27], [28] was first designed to provide insights into the performance of multicore architectures, utilizing a parameter, operational intensity. It helps to understand the potential bottlenecks and improvement opportunities for an application on different families of CPUs. Other efforts [4], [29], [30] extended this model to accelerators, such as the GPU and TPU. The roofline model does not provide insights into the loop-level acceleration opportunity on different hardware accelerators. These limitations prevent the developer from choosing the right accelerator, prior to any development. In comparison, LoopBench provides optimization details in loop-level granularity (not the whole application) and does not rely on the real implementation of the algorithm.

Existing benchmarks adopted widely-used algorithms or computational patterns to draw comparison lines between different processors. Some of these works [5], [31]–[33] focused only on a particular type of processor, whereas others [14] were designed to compare different families of processors, e.g., CPUs vs. GPUs. These benchmarks help understand the performance differences between accelerators while executing certain types of applications, but their insights are limited to

specific applications. It is difficult for the developer to use these benchmarks to decide which accelerator has more potential to accelerate a new type of application. In comparison, LoopBench’s insights are not limited to a particular application and can be applied to any new algorithm.

Closely related to our approach, the TSVC benchmark [34] includes a suite of various types of loops, which has inspired some of the loop patterns considered by our LoopBench. TSVC was mainly designed to evaluate the efficiency of compilers on detecting and vectorizing such loops on SIMD architectures. In comparison, the goal of LoopBench is to evaluate the correlation between common loop patterns and the extent of accelerating such loop on different hardware platforms. As a result, LoopBench represents loop patterns that draw a significant difference between GPU and FPGA. In addition, it provides an in-depth analysis of how loop characteristics impact the accelerator performance, all of which are not possible by simply applying or porting the TSVC.

Related works studied the identification and performance analysis of certain memory access patterns, know as *idioms*, on a target platform [20], [35]. The idioms are an abstraction of memory accesses patterns that are common among applications. In comparison, to LoopBench is designed to explain the capability of accelerators in handling loops with common dependency patterns.

There are efforts in predicting the performance of a complete application on a target platform [21], [36]–[38]. These solutions require access to the real implementation of the application, and the prediction is specific to the application. In comparison, LoopBench is able to give insight on the acceleration opportunities of different loop patterns, on an abstract level. It also helps the developer better chose the right accelerator, prior to implementing the actual code.

Finally, several recent works demonstrated the benefits of FPGA-based accelerations for a variety of applications, compared to GPUs [2], [6], [19]. These works are mainly discussing the feasibility of using OpenCL-enabled FPGAs for certain workloads or a limited set of loop patterns. In comparison, LoopBench provides a comprehensive coverage of common loop patterns, which can help understand the effectiveness of using FPGA to accelerate diverse applications.

#### V. CONCLUSIONS AND FUTURE WORK

In this work, we designed LoopBench for studying common loop patterns on important GPU and FPGA accelerators. We identified and analyzed five common loop patterns, along with the key configuration parameters in these patterns. We then studied the acceleration opportunities for these loop patterns and how the loop configurations and accelerator platforms affect the effectiveness of acceleration. Using LoopBench, developers can gain a good understanding of the acceleration potential of their algorithms on different platforms, without having to implement them for any specific platform, based on the loop patterns that these algorithms embody. LoopBench is open source and publicly available [39].

## REFERENCES

- [1] J. D. Owens, M. Houston, D. Luebke, S. Green, J. E. Stone, and J. C. Phillips, "GPU computing," *Proceedings of the IEEE*, vol. 96, no. 5, pp. 879–899, 2008.
- [2] S. Biokaghazadeh, F. Ren, and M. Zhao, "Are FPGAs suitable for edge computing?" *arXiv preprint arXiv:1804.06404*, 2018.
- [3] J. Fowers, G. Brown, P. Cooke, and G. Stitt, "A performance and energy comparison of FPGAs, GPUs, and multicores for sliding-window applications," in *Proceedings of the ACM/SIGDA international symposium on Field Programmable Gate Arrays*. ACM, 2012, pp. 47–56.
- [4] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers *et al.*, "In-datacenter performance analysis of a tensor processing unit," in *Computer Architecture (ISCA), 2017 ACM/IEEE 44th Annual International Symposium on*. IEEE, 2017, pp. 1–12.
- [5] S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, S.-H. Lee, and K. Skadron, "Rodinia: A benchmark suite for heterogeneous computing," in *Workload Characterization, 2009. IISWC 2009. IEEE International Symposium on*. Ieee, 2009, pp. 44–54.
- [6] H. R. Zohouri, N. Maruyama, A. Smith, M. Matsuda, and S. Matsuoka, "Evaluating and optimizing opencl kernels for high performance computing with fpgas," in *High Performance Computing, Networking, Storage and Analysis, SC16: International Conference for*. IEEE, 2016, pp. 409–420.
- [7] J. Cong, Z. Fang, M. Lo, H. Wang, J. Xu, and S. Zhang, "Understanding performance differences of FPGAs and GPUs," in *2018 IEEE 26th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 2018, pp. 93–96.
- [8] D. Koch and J. Torresen, "FPGASort: A high performance sorting architecture exploiting run-time reconfiguration on FPGAs for large problem sorting," in *Proceedings of the 19th ACM/SIGDA international symposium on Field programmable gate arrays*. ACM, 2011, pp. 45–54.
- [9] B. Duan, W. Wang, X. Li, C. Zhang, P. Zhang, and N. Sun, "Floating-point mixed-radix FFT core generation for FPGA and comparison with GPU and CPU," in *Field-Programmable Technology (FPT), 2011 International Conference on*. IEEE, 2011, pp. 1–6.
- [10] Y. Zhang, Y. H. Shalabi, R. Jain, K. K. Nagar, and J. D. Bakos, "FPGA vs. GPU for sparse matrix vector multiply," in *Field-Programmable Technology, 2009. FPT 2009. International Conference on*. Citeseer, 2009, pp. 255–262.
- [11] P. Wilson, *Design Recipes for FPGAs: Using Verilog and VHDL*. Newnes, 2015.
- [12] A. Munshi, "The opencl specification," in *Hot Chips 21 Symposium (HCS), 2009 IEEE*. IEEE, 2009, pp. 1–314.
- [13] C. Nugteren, "Cliblast: A tuned opencl BLAS library," *arXiv preprint arXiv:1705.05249*, 2017.
- [14] A. Danalis, G. Marin, C. McCurdy, J. S. Meredith, P. C. Roth, K. Spafford, V. Tipparaju, and J. S. Vetter, "The scalable heterogeneous computing (SHOC) benchmark suite," in *Proceedings of the 3rd Workshop on General-Purpose Computation on Graphics Processing Units*. ACM, 2010, pp. 63–74.
- [15] J. Herdman, W. Gaudin, S. McIntosh-Smith, M. Boulton, D. A. Beckingsale, A. Mallinson, and S. A. Jarvis, "Accelerating hydrocodes with OpenACC, OpenCL and CUDA," in *2012 SC Companion: High Performance Computing, Networking Storage and Analysis*. IEEE, 2012, pp. 465–471.
- [16] J. Cong, Z. Fang, Y. Hao, P. Wei, C. H. Yu, C. Zhang, and P. Zhou, "Best-effort FPGA programming: A few steps can go a long way," *arXiv preprint arXiv:1807.01340*, 2018.
- [17] S. Qin and M. Berekovic, "A comparison of high-level design tools for soc-fpga on disparity map calculation example," *arXiv preprint arXiv:1509.00036*, 2015.
- [18] A. A. Freitas and S. H. Lavington, "Basic concepts on parallel processing," in *Mining Very Large Databases with Parallel Processing*. Springer, 2000, pp. 61–69.
- [19] H. R. Zohouri, A. Podobas, and S. Matsuoka, "Combined spatial and temporal blocking for high-performance stencil computation on FPGAs using OpenCL," in *Proceedings of the 2018 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. ACM, 2018, pp. 153–162.
- [20] M. R. Meswani, L. Carrington, D. Unat, A. Snaveley, S. Baden, and S. Poole, "Modeling and predicting performance of high performance computing applications on hardware accelerators," *The International Journal of High Performance Computing Applications*, vol. 27, no. 2, pp. 89–108, 2013.
- [21] S. Kumar, V. Srinivasan, A. Sharifian, N. Sumner, and A. Shriraman, "Peruse and profit: Estimating the accelerability of loops," in *Proceedings of the 2016 International Conference on Supercomputing*. ACM, 2016, p. 21.
- [22] Intel, "Fog Reference Unit," <https://www.intel.com/content/www/us/en/internet-of-things/fog-reference-design-overview.html>.
- [23] G. Guennebaud, B. Jacob *et al.*, "Eigen," URL: <http://eigen.tuxfamily.org>, 2010.
- [24] M. E. Belviranlı, P. Deng, L. N. Bhuyan, R. Gupta, and Q. Zhu, "Peerwave: Exploiting wavefront parallelism on gpus with peer-sm synchronization," in *Proceedings of the 29th ACM on International Conference on Supercomputing*. ACM, 2015, pp. 25–35.
- [25] R. Bellman, *Dynamic programming*. Courier Corporation, 2013.
- [26] S. B. Needleman and C. D. Wunsch, "A general method applicable to the search for similarities in the amino acid sequence of two proteins," *Journal of molecular biology*, vol. 48, no. 3, pp. 443–453, 1970.
- [27] S. Williams, A. Waterman, and D. Patterson, "Roofline: an insightful visual performance model for multicore architectures," *Communications of the ACM*, vol. 52, no. 4, pp. 65–76, 2009.
- [28] Y. J. Lo, S. Williams, B. Van Straalen, T. J. Ligocki, M. J. Cordery, N. J. Wright, M. W. Hall, and L. Oliker, "Roofline model toolkit: A practical tool for architectural and program analysis," in *International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems*. Springer, 2014, pp. 129–148.
- [29] H. Jia, Y. Zhang, G. Long, J. Xu, S. Yan, and Y. Li, "GPU Roofline: a model for guiding performance optimizations on GPUs," in *European Conference on Parallel Processing*. Springer, 2012, pp. 920–932.
- [30] D. Doerfler, J. Deslippe, S. Williams, L. Oliker, B. Cook, T. Kurth, M. Lobet, T. Malas, J.-L. Vay, and H. Vincenti, "Applying the roofline performance model to the intel xeon phi knights landing processor," in *International Conference on High Performance Computing*. Springer, 2016, pp. 339–353.
- [31] J. L. Henning, "SPEC CPU2000: Measuring CPU performance in the new millennium," *Computer*, vol. 33, no. 7, pp. 28–35, 2000.
- [32] R. Haney, T. Meuse, J. Kepner, and J. Lebak, "The HPEC challenge benchmark suite," in *HPEC 2005 Workshop*, 2005.
- [33] G. Ndu, J. Navaridas, and M. Luján, "CHO: towards a benchmark suite for OpenCL FPGA accelerators," in *Proceedings of the 3rd International Workshop on OpenCL*. ACM, 2015, p. 10.
- [34] S. Maleki, Y. Gao, M. J. Garzar, T. Wong, D. A. Padua *et al.*, "An evaluation of vectorizing compilers," in *Parallel Architectures and Compilation Techniques (PACT), 2011 International Conference on*. IEEE, 2011, pp. 372–382.
- [35] J. He, A. E. Snaveley, R. F. Van der Wijngaart, and M. A. Frumkin, "Automatic recognition of performance idioms in scientific applications," in *Parallel & Distributed Processing Symposium (IPDPS), 2011 IEEE International*. IEEE, 2011, pp. 118–127.
- [36] J. Meng, V. A. Morozov, K. Kumaran, V. Vishwanath, and T. D. Uram, "GROPHECY: GPU performance projection from CPU code skeletons," in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 2011, p. 14.
- [37] G. Chapuis, S. Eidenbenz, and N. Santhi, "Gpu performance prediction through parallel discrete event simulation and common sense," in *Proceedings of the 9th EAI International Conference on Performance Evaluation Methodologies and Tools*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2016, pp. 204–211.
- [38] M. Boyer, J. Meng, and K. Kumaran, "Improving GPU performance prediction with data transfer modeling," in *Parallel and Distributed Processing Symposium Workshops & PhD Forum (IPDPSW), 2013 IEEE 27th International*. IEEE, 2013, pp. 1097–1106.
- [39] S. Biokaghazadeh, "LoopBench," <https://github.com/samanaghazadeh/shoc-fpga/>.