

TopSort: A High-Performance Two-Phase Sorting Accelerator Optimized on HBM-based FPGAs

Weikang Qiao, *Student Member, IEEE*, Licheng Guo, Zhenman Fang, *Member, IEEE*,
Mau-Chung Frank Chang, *Life Fellow, IEEE* and Jason Cong, *Fellow, IEEE*

Abstract—The emergence of high-bandwidth memory (HBM) brings new opportunities to boost the performance of sorting acceleration on FPGAs, which was conventionally bounded by the available off-chip memory bandwidth. However, it is nontrivial for designers to fully utilize this immense bandwidth. First, the existing sorter designs cannot be directly scaled at the increasing rate of available off-chip bandwidth, as the required on-chip resource usage grows at a much faster rate and would bound the sorting performance in turn. Second, designers need an in-depth understanding of HBM’s characteristics to effectively utilize the HBM bandwidth. To tackle these challenges, we present TopSort, a novel two-phase sorting solution optimized for HBM-based FPGAs. In the first phase, 16 merge trees work in parallel to fully utilize 32 HBM channels’ bandwidth. In the second phase, TopSort reuses the logic from phase one to form a wider merge tree to merge the partially sorted results from phase one. TopSort also adopts HBM-specific optimizations to reduce resource overhead and improve bandwidth utilization. TopSort can sort up to 4 GB data using all 32 HBM channels, with an overall sorting performance of 15.6 GB/s. TopSort is $6.7\times$ and $2.2\times$ faster than state-of-the-art CPU and FPGA sorters.

Index Terms—Sorting, merge sort, hardware acceleration, high-bandwidth memory, memory-centric design, FPGA, floorplan.



1 INTRODUCTION

Sorting is one of the fundamental computation kernels in many big data applications and there have been continuous efforts on designing high-performance sorting accelerators on FPGAs [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]. Most of these sorting accelerators are based on the multi-way merge tree sort algorithm [1], [2], [3], [4], [5], [6], [7], [8], [9], due to its massive data parallelism and regular memory access patterns. Before the maturing of the high-bandwidth memory (HBM) technology, these sorters were usually implemented on DRAM-based FPGAs and bounded by the off-chip memory bandwidth [7], [8]. For example, Bonsai [7] presented the state-of-the-art merge tree based sorter on the AWS F1 datacenter FPGA and achieved an overall sorting performance of 7.1 GB/s with a merge tree’s throughput of 32 GB/s;¹ it could sufficiently scale up the number of parallel merge units until the memory bandwidth becomes the bottleneck.

The emergence of HBM-based FPGAs has brought the potential to further boost the performance of the sorting accelerators by offering a much higher off-chip memory bandwidth. For example, the recent HBM-based datacenter FPGAs report a peak memory bandwidth of about 420

GB/s, which is roughly $6\times$ of that on a similar DRAM-based datacenter FPGA [16], [17]. A natural question is: if we simply port and scale the state-of-the-art merge tree designs onto an HBM-based FPGA, can we get $6\times$ higher merge tree throughput? Unfortunately, the answer is no due to two nontrivial challenges.

On the one hand, with the tremendous off-chip bandwidth increase, the bottleneck of sorting acceleration would shift from off-chip memory to the available on-chip resources. As will be analyzed in Section 3, for a merge tree-based sorter that can output p elements per cycle, the required off-chip bandwidth increases linearly with p , while the required on-chip resources increase much faster at the rate of $\theta(p\log^3(p))$. Unfortunately, for an HBM-based FPGA (e.g., Xilinx Alveo U280 FPGA [18]) that provides roughly $6\times$ more bandwidth than a similar DRAM-based FPGA (e.g., AWS F1 FPGA), it only provides merely $1.2\times$ more on-chip resources. Indeed, our analysis shows that a single merge tree accelerator can not be scaled to use more than $1/4$ of the available HBM bandwidth.

On the other hand, it is nontrivial to fully utilize the HBM bandwidth in an accelerator design. Typically, the HBM is composed of 32 small channels. The accelerator has to rely on multiple memory controllers to access those HBM channels in parallel to maximize the memory bandwidth, at the expense of on-chip resources. Moreover, there could easily be contention between multiple channel accesses due to HBM’s internal channel switching, which would degrade the effective bandwidth. Therefore, it requires delicate memory access control to improve the bandwidth utilization and reduce resource overhead. Finally, the HBM stacks are physically connected to a datacenter FPGA’s bottom die only, making it difficult to spread the resource utilization across multiple FPGA dies to achieve desirable timing closure. To the best of our knowledge, none of the published HBM-

- Weikang Qiao and Mau-Chung Frank Chang are with the Department of Electrical and Computer Engineering, University of California, Los Angeles, CA, 90025.
E-mail: wkqiao2015@ucla.edu, mfchang@ee.ucla.edu
- Licheng Guo and Jason Cong are with the Department of Computer Science, University of California, Los Angeles, CA, 90025.
E-mail: {lguo, cong}@cs.ucla.edu
- Zhenman Fang is with the School of Engineering Science, Simon Fraser University, Burnaby, BC V5A 1S6, Canada.
E-mail: zhenman@sfu.ca

1. The definitions of a merge tree’s throughput and the overall sorting performance are presented in Section 2 and Section 3.

based accelerator designs [19], [20], [21], [22], [23] is able to fully utilize the entire bandwidth of the 32 HBM channels.

In this work, we present TopSort, a high-performance two-phase sorting accelerator specialized for HBM-based FPGAs. TopSort avoids the excessive resource consumption of directly scaling a single giant sorter’s throughput by novelly splitting the complete sorting process into two separate merge phases with smaller sorters and reusing the resources between the two phases. In the first phase, TopSort employs 16 parallel small merge tree kernels, each of which sorts a portion of the input sequence with two HBM channels (one for reading unsorted inputs and the other for writing sorted outputs). In this phase, the effective merge tree throughput is equal to the bandwidth of 16 HBM channels. In the second phase, TopSort merges the sorted results from all HBM channels into one final sorted sequence. To reduce the resource consumption, this phase reuses 4 merge tree kernels from phase one to form a wider merge tree with a $4\times$ higher throughput, since merge trees with different throughput share a similar operation pattern.

To improve the effective HBM bandwidth utilization, we profile the corresponding HBM channel characteristics and carefully optimize the merge tree’s memory access pattern for each phase, including optimizing its data layout and burst access, as well as reducing the usage of HBM controllers. Besides, the novel merge tree reuse architecture allows us to easily improve the design frequency through coarse-grained floorplanning of each separate merge tree kernel. We floorplan TopSort by evenly distributing the design across the entire FPGA logic regions, based on an efficient resource model that considers the accelerator’s design complexity, the available resources of each FPGA die, the limitation of cross-die signals, and the overhead of HBM controllers.

When implemented on the Xilinx Alveo U280 FPGA, TopSort runs at 214 MHz and achieves an overall performance of 15.6 GB/s, which is $6.7\times$ and $2.2\times$ faster than state-of-the-art CPU and FPGA sorters [7], [24]. TopSort can sort up to 4 GB data at a time, which is half of the total HBM capacity. Although this work includes a number of device-specific optimization for the HBM-based U280 FPGAs from Xilinx, the general two-phase reused-based merge sort architecture is applicable to all HBM-based FPGAs across different vendors.

We summarize the contributions of TopSort as below:

1. A novel two-phase merge tree based sorter optimized on HBM-based FPGAs, which fully utilizes the HBM bandwidth and alleviates the on-chip resource bottleneck by intelligently reusing multiple merge trees between two phases.
2. Techniques and insights for HBM-specific optimizations, including the data layout, burst access, and HBM controller optimizations, as well as the floorplanning strategy.
3. Experimental results that demonstrate the superior 15.6 GB/s sorting performance and show TopSort is $6.7\times$ and $2.2\times$ faster than state-of-the-art CPU and FPGA sorters.

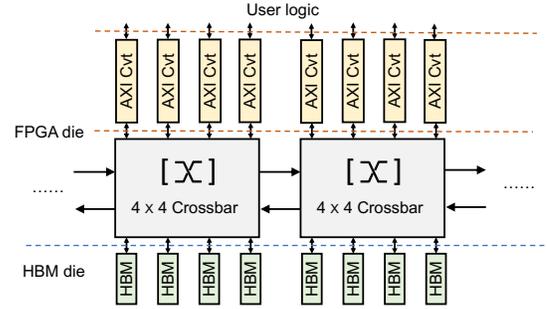


Fig. 1: Connections from user logic to HBM channels. AXI Cvt is short for AXI Rate Converter. Lateral connections between nearby crossbars enable AXI to access HBM channels that are not located in the same group.

2 BACKGROUND REVIEW

The merge tree sorting algorithm is favored for FPGA-based sorters due to its massive data parallelism, less control overhead and regular memory access patterns [1], [2], [3], [4], [5], [6], [7], [8], [9]. In this section, we first introduce the HBM-based FPGAs. Then we give an overview of the hardware merge units and the existing DRAM-based merge tree sorting accelerators.

2.1 HBM-Based FPGAs

HBM achieves higher bandwidth than DDR4 DRAMs by stacking multiple small DRAM dies together and is one of the most promising candidates enabling memory-centric designs [25]. Taking Xilinx U280 board as an example, the board is equipped with 2 HBM stacks and each stack contains 16 pseudo channels [18]. Figure 1 exhibits the connections between the user logic and the HBM channels. Each HBM channel can be accessed through a 256-bit wide AXI interface running at 450 MHz. The vendor tool will by default implement AXI rate converters in the HBM Memory Subsystem (HMSS) to adapt the original 256-bit AXI interfaces to 512-bit AXI interfaces running at 225 MHz to the user logic. In the case of routing congestion which happens if the on-chip resources are over-utilized or the design is not well pipelined, both the HBM-side AXI frequency and the user-side design frequency will be reduced and the available HBM bandwidth will be degraded.

There is no global crossbar to allow 32 AXI interfaces to access the 32 HBM channels at the same time. Instead, the 32 HBM channels are physically bundled into 8 groups and each group contains 4 adjacent channels joined by a built-in 4×4 crossbar, as is shown in Figure 1. The crossbar provides full connectivity within the group. Meanwhile, each AXI interface at the user side can still access any HBM channels outside its group. The data will sequentially traverse through each of the lateral connections until it reaches the crossbar connecting to the target channel.

2.2 Hardware Merge Unit

A hardware merge unit takes two sorted sequences of elements as its inputs and then merges them into one sorted sequence. Specifically, an E -rate hardware merge unit takes E input elements from its two E -element wide inputs and

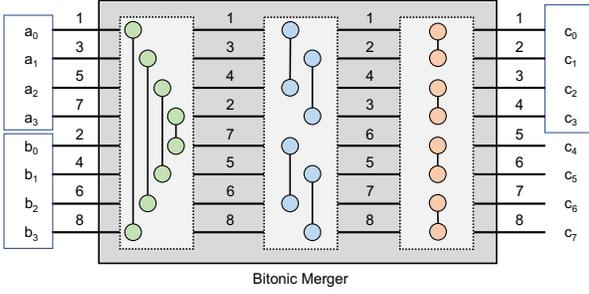


Fig. 2: An example of 4-rate bitonic merge unit: each vertical line that connects two dots is a compare-swap cell and the compare-swap cells in the same box are processed in the same cycle. c_{0-3} will be the outputs after 3 cycles.

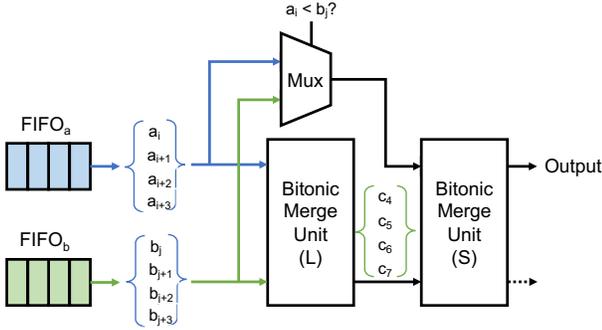


Fig. 3: The topology of a 4-rate MMS merge unit. Registers in the datapath are omitted for simplicity.

outputs E sorted elements every cycle. The *compare-swap cell* is the basic building block for hardware merge unit, which compares two elements' values and swaps them into the correct ordering [1]. A compare-swap cell usually contains a comparator and a 2-input multiplexer, which is suitable to be implemented using the Look-Up Tables (LUTs) on the FPGAs. Using a pipeline of multiple compare-swap cells, designers can develop the hardware merge unit. The bitonic merge unit shown in Figure 2 is one of the most widely used hardware merge units.

If the two sorted input sequences are longer than E , the E -rate merge unit has to be time multiplexed (i.e., executed multiple rounds) with extra control signals to ensure the outputs are in order. For example, in Figure 2, c_{4-7} represent the largest four elements of a_{0-3} and b_{0-3} . In the next cycle, c_{4-7} need to be sent back to the input side of the same bitonic merge unit and merged with either a_{4-7} or b_{4-7} to create the second 4-element sorted outputs. The feedback paths from the output side of the bitonic merge unit to its input side mean the merge unit has to wait for extra cycles before it can do the next merge operation. In other words, the initial interval (II) is equal to the number of pipeline stages in the bitonic merge unit.

To improve the merging performance with an II equal to 1, several optimization techniques have been proposed [11], [12], [13], [14]. [12] proposes an E -rate streaming merge unit called *MMS* that outputs E elements every cycle using two bitonic merge units, as shown in Figure 3. The intuition is

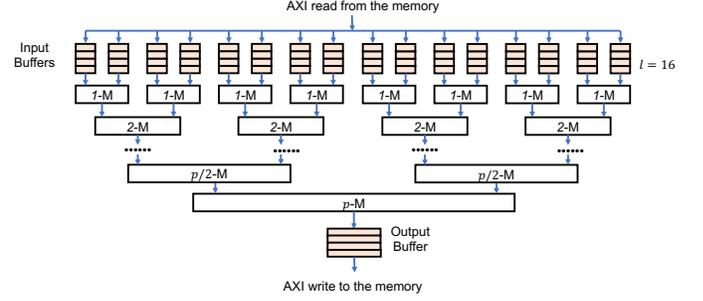


Fig. 4: An example of the merge tree ($p = 16$, $l = 16$), where an E -M box denotes an E -rate hardware merge unit. Note that l can be larger than p , e.g., to make a merge tree ($p = 16$, $l = 32$), one can have another layer of 32 1-M merge units on top of the current leaf layer.

that the larger half outputs of the original bitonic merge unit can be first calculated through the bitonic merge unit (L) and are later fed into the bitonic merge unit (s) with delayed inputs from either sequence a or b . In this work, we use the same topology of MMS to construct the hardware merge unit, as it is free of feedback paths and thus helps alleviate the design routing congestion.

2.3 DRAM-based Merge Tree Sorting Accelerator

Using a combination of various hardware merge units with different rates, we can build a complete binary tree that consumes l unsorted input sequences concurrently at its leaves and outputs p sorted elements at its root every cycle [4], [6], [7], [8]. Figure 4 shows the architecture of a merge tree [7]. Such a merge tree can be uniquely denoted by (p, l) , where p refers to the **merge tree's throughput** and l is **the number of leaves**.

In one sorting *pass*, the input elements are streamed from the off-chip memory into the leaf buffers and then through the merge tree, and the output elements are streamed back into the off-chip memory. Assume there are N unsorted elements initially stored in the off-chip memory, these N elements need to be streamed into the merge tree for multiple passes and after each pass, the size of the partially sorted sequence grows exactly by l times. During the first pass, the merge tree reads these N sub sequences (each containing one element) from the off-chip memory, merges them into N/l sorted sub sequences, each containing l sorted elements, and writes them back to the off-chip memory. In the next pass, the l -element sorted sub sequences are fed into the merge tree again to get N/l^2 sorted sequences of length l^2 . These steps are repeated until the last pass, where l sorted sequences of length N/l are processed by the merge tree to form a complete N -element sorted sequence. One can easily derive that the total number of passes is $\lceil \log_l N \rceil$.

3 SCALABILITY ANALYSIS OF A SINGLE MERGE TREE

In this section, we explain why the single merge tree employed on existing DRAM-based FPGAs cannot be directly scaled to efficiently work on HBM-based FPGAs.

The throughput of a single merge tree can be derived using the analytical model from [7]. First, the merge tree with l leaves takes $\lceil \log_\ell N \rceil$ passes. Second, the merge tree outputs p elements every cycle to the off-chip memory, where p is directly reflected by the off-chip memory bandwidth β_{memory} allocated to the writing operations. Generally, half of the system bandwidth is allocated for writing and the other half is for reading. We define the **overall sorting performance** as the number of bytes of unsorted elements divided by the time it takes to get the completely sorted elements. Assuming there are unlimited on-chip resources, we have the overall sorting performance $\beta_{overall}$ in Equation 1.

$$\beta_{overall} = \frac{\beta_{memory}}{\lceil \log_\ell N \rceil} \quad (1)$$

Clearly, the on-chip resources are finite. Next, we analyze the on-chip resource requirement in scaling the merge tree-based sorter. The merge tree in Figure 4 can be viewed as two sub merge trees whose root rates are $p/2$ each, plus a merge unit whose rate is p . Since a merge unit is a variation of a parallel sorting network, such as bitonic or even-odd sorting network, and such a p -rate merge unit consumes $p \log^2(p)$ number of comparison operators [11], [12], the number of comparators, $L(p)$, required when scaling p can be summarized in Equation 2.

$$L(p) = 2L(p/2) + \theta(p \log^2(p)) \quad (2)$$

Based on this equation, we can derive that the total number of logic elements (i.e., comparators) for the complete binary tree is $\theta(p \log^3(p))$ [26]. Considering that migrating from 4 DRAM channels in [7] to 32 HBM channels increases the available off-chip bandwidth by nearly 8 \times , it is not possible to directly scale a single merge tree's throughput to catch up with the increase of the HBM bandwidth.

In fact, we find that the root throughput of a single merge tree can only be scaled to match 4 HBM channels (for write operations) on the Xilinx Alveo U280 board. Since the read operations of the tree also occupies the same amount of HBM bandwidth, the bandwidth utilized in such a merge tree is at most equal to 8 HBM channels. In other words, such a merge tree only uses 25% of the HBM bandwidth in each pass.

4 TOPSORT METHODOLOGY & ARCHITECTURE

In this section, we present the methodology of TopSort. First, we illustrate the idea of the two-phase sorting. Second, we show how TopSort reuses the logic between two phases to alleviate the resource contention and improve the performance.

4.1 Two Phases in TopSort

Since the required resources grow super-linearly when directly scaling a single merge tree, TopSort chooses to split the N unsorted elements evenly into k parts and have k smaller merge trees work in parallel. Each merge tree sorts N/k elements and is configured to have a throughput that saturates the bandwidth of a single HBM channel. This method has better scalability as its resource consumption grows linearly with the number of trees k . However, the k

sorted sequences still need to be merged into the final sorted sequence. This can be done by streaming the k sequences into another merge tree for the final merge. We refer to the process during which k merge trees work in parallel as phase 1, and the subsequent final merge as phase 2.

If each merge tree in phase 1 has l_{phase1} number of leaves and a single HBM channel's bandwidth is $\beta_{channel}$, the average performance in phase 1 is given in Equation 3. The fraction part is a single merge tree's performance when sorting N/k elements, and there are k such merge trees in parallel.

$$\beta_{phase1} = k \cdot \frac{\beta_{channel}}{\lceil \log_{l_{phase1}}(N/k) \rceil} \quad (3)$$

Let p_{phase2} denote the merge tree's throughput in phase 2. As long as the merge tree has k or more leaves, phase 2 will require only *one* pass, e.g., if phase 2's merge tree has k leaves, then each of the k sorted sequences from phase 1 will be fed into one of the corresponding k leaves once. According to the analysis in Section 3, the merge tree of phase 2 only exploits a portion of the HBM bandwidth due to on-chip resource constraints. Therefore, the performance β_{phase2} of phase 2 is proportional to p_{phase2} .

The overall sorting performance is in Equation 4.

$$\beta_{overall} = \frac{1}{\frac{1}{\beta_{phase1}} + \frac{1}{\beta_{phase2}}} \quad (4)$$

This highlights that the resource balance between the parallel merge trees of phase 1 and the final merge tree of phase 2 plays an important role in the overall sorting performance. For example, we may have 16 merge trees in the first phase to fully utilize all 32 HBM channels and achieve the best performance. But if the remaining resources only allow us to build a final merge tree whose throughput is equal to the bandwidth of a single channel, then the overall sorting performance will be limited by the final merge, no matter how much better the memory bandwidth utilization we achieve in phase 1. One may wonder if we can reprogram the FPGA in phase 2 to have a wider tree for final merging, but the reprogramming overhead takes several seconds [9] and thus is not practical in this case.

4.2 Merge Tree Reuse in TopSort

To resolve the resource contention between phases 1 & 2, TopSort proposes to reuse the merge units in phase 1 to build the final merge tree. This is based on the observation that a merge tree can be decomposed as sub merge trees plus some extra levels of merge units. For example, a merge tree whose throughput is p in Figure 4 consists of four merge trees of throughput $p/4$, plus two $p/2$ -rate merge units and one p -rate merge unit, as shown in Figure 5. In other words, if we want to build a merge tree whose throughput is equal to 4 HBM channels' bandwidth in phase 2, we only need extra resources to build the 3 more merge units. The rest of the resources including the input buffers can be reused from four of the merge trees in phase 1. With this novel merge tree reuse, conceptually, TopSort is able to integrate 16 merge trees to saturate all 32 HBM channels' bandwidth in phase 1 and another merge tree whose throughput saturates 4 HBM channels' bandwidth in phase 2 onto one FPGA.

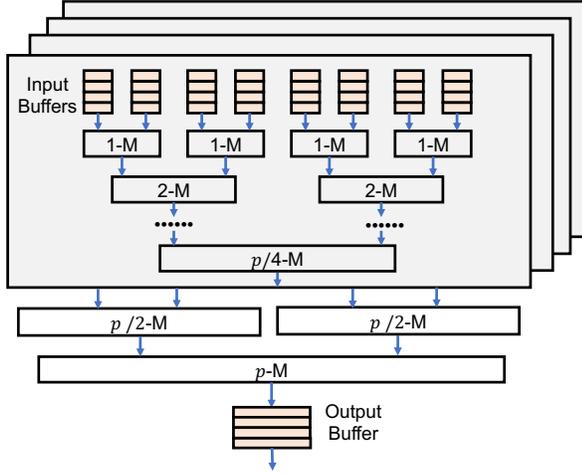


Fig. 5: Merge tree reuse: four $p/4$ -rate merge trees in phase 1 are reused to build a larger p -rate merge trees in phase 2.

Please note this merge tree reuse approach shows a general architecture that could be easily scaled with different hardware configurations. Given the available off-chip bandwidth (e.g., the number of HBM channels) and on-chip resources (e.g., available BRAMs and LUTs), we can easily adapt (either scale up or down) the number of the parallel merge trees in phase 1, the size (i.e., number of tree leaves) of each merge tree in phase 1 and the final merge tree in phase 2.

4.3 Architecture to Support Logic Reuse

The overall architecture of TopSort is shown in Figure 6. There are 32 HBM channels available, so TopSort uses 16 merge tree kernels in phase 1. Given a total of N unsorted elements, we split them into 16 sequences, each of which contains $N/16$ elements and is stored in one HBM channel. Then each of the 16 trees will stream the unsorted sequence from one HBM channel and stream the merged (sorted) sequence to another HBM channel, as described in Section 2.3. We design each tree so that it has 16 leaves and outputs 64 bytes per cycle at its root to saturate the bandwidth of a single HBM channel. The leaf number for each merge tree is chosen to be 16 is because the leaf buffers cannot exceed the available on-chip BRAM limit. Assuming each element is 64-bit, the merge tree will output 8 elements every cycle.

In phase 2, TopSort reuses 4 merge trees from phase 1 and adds two levels of merge units to form a wider merge tree. This merge tree has 64 leaves and outputs 256-byte elements per cycle. Since there are 16 sequences sitting in 16 HBM channels after phase 1, we split each sequence into 4 segmented sub sequences. Then each of the 64 sub sequences is fed into one of the 64 leaves for one pass to get the final sorted results.

In phase 2, the input buffers and each merge unit of the four reused merge trees stay exactly the same as phase 1. We only need to change the memory addresses and the length of the sequences that the AXI interfaces read from each channel for each of the 64 leaves. The change of the write behavior of phase 2 is discussed in Section 4.5.

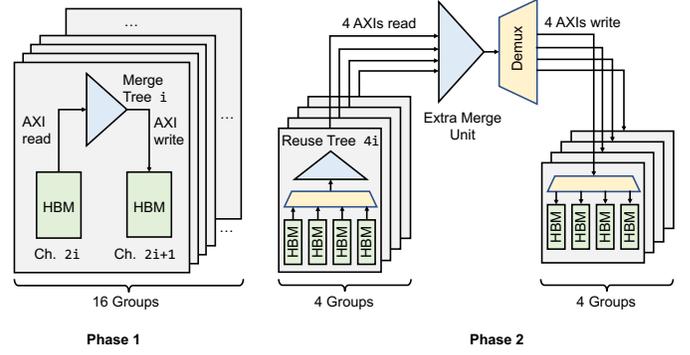


Fig. 6: The overall architecture of the two phases in TopSort. The merge trees rely on AXI interfaces to access HBM channels. The role of demux is covered in Section 4.5.

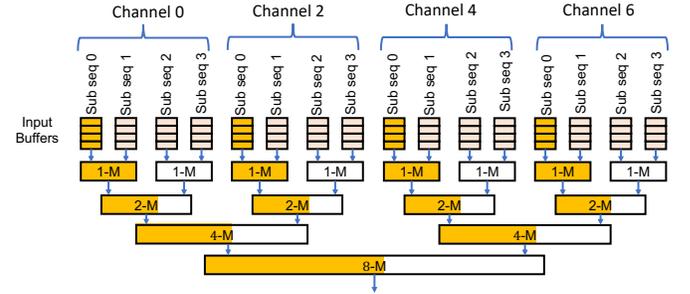


Fig. 7: The active rate of one reused tree in phase 2, if the sequence in each channel is completely sorted. Initially, only the leaves feeding sub sequence 0 from each channel are active, then are the leaves feeding sub sequences 1, and so on. Leaves marked as yellow are actively feeding. Merge units partially marked mean they are idle half of the time.

4.4 Tuning Sorted Sequence Size from Phase 1

One hidden requirement for a merge tree to output p elements per cycle is that, on average, it has p leaves out of its total l leaves providing elements every cycle. Section 4.3 mentions that each of the 16 sequences in 16 HBM channels after phase 1 needs to be divided into 4 segmented sub sequences so that there are 64 sub sequences fed into the 64 leaves of the tree in phase 2. If each sequence is totally sorted after phase 1, then the 4 sub sequences have the relation that elements in sub sequence 0 are always smaller than elements in sub sequence 1, and so on. When such 4 sub sequences are fed into 4 leaves, only 1 of the 4 leaves will feed the element to the merge tree at any cycle. For each of the reused merge trees that has 16 leaves, this means only 4 leaves are providing elements into the tree per cycle, as shown in Figure 7. As a result, each merge tree is idle half of the time. Although each tree can output 8 elements per cycle, the average throughput of the tree is 4 elements per cycle, which is below our expectation.

To solve this issue, we need to tune the sorted sequence size from phase 1. Instead of completely sorting the $N/16$ -element sequence in each channel, we want 4 $N/64$ -element sorted sub sequences from phase 1. This is done by controlling the number of elements that go into the leaves in the last pass of phase 1. Since the sorted sequence size always grows by the number of leaves l , we deliver $N/1024$ elements into

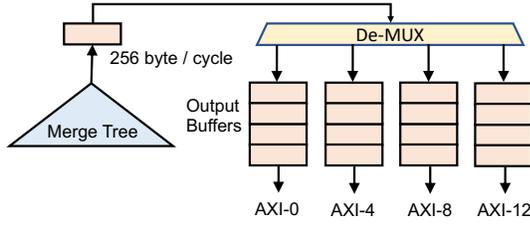


Fig. 8: TopSort’s write behavior of phase 2. Writes from 4 AXI interfaces are used to match the tree’s throughput.

each leaf. At the tree root side, we get $N/64$ -element sorted sequence. We repeat this process 4 times for each merge tree in phase 1 to get 4 $N/64$ -element sorted sub sequences.

4.5 Memory Write Pattern in Phase 2

Since the $4 \times$ wider merge tree of phase 2 outputs 256-byte sorted elements per cycle, there are 4 AXI interfaces writing to the HBM channels in parallel. The choice of the specific AXI interfaces will be explained in Section 5.2. The on-chip memory is not big enough to hold all the continuously sorted elements before they are written back to entirely fill one HBM channel. Therefore, these continuously sorted elements have to be split into batches and then written back to separate HBM channels through 4 AXI interfaces. Figure 8 depicts the details of such behavior. In our implementation, we choose the batch size to be 4 KB, e.g., we write the first 4 KB sorted elements into the first output buffer, then write the second 4 KB sorted elements into the second output buffer, and so on. The elements in the four buffers will then be written to HBM channels via the four AXI interfaces. As a result, although the elements are completely sorted, one needs to pick the 4 KB batches channel by channel to form the final sorted results.

5 HBM-SPECIFIC OPTIMIZATIONS

In this section, we introduce the design choices in TopSort that are coupled with specific HBM characteristics. These HBM-related optimizations are necessary to achieve a full bandwidth usage. First, we minimize the number of AXI interfaces to reduce the resource overhead. Second, we propose a dedicated data layout with spatial locality to prevent bandwidth degradation. In addition, we properly set different burst sizes for different merge trees to strike a balance between the area overhead and the bandwidth. Finally, we floorplan and pipeline our design based on the architecture of the HBM platform for timing optimization.

5.1 Avoiding Unnecessary AXI Interfaces

According to our measurement, an AXI rate converter in Figure 1 for an HBM channel requires about 5K LUTs and about 6K Flip-Flops (FFs) on the Xilinx U280 FPGA. As a result, the naive choice of instantiating one AXI interface for each of the 32 HBM channels will result in about 320K LUTs, which is more than 40% of the available LUTs on the bottom die where the HBM resides. Such high resource overhead will squeeze the space for user logic and cause severe routing issues.

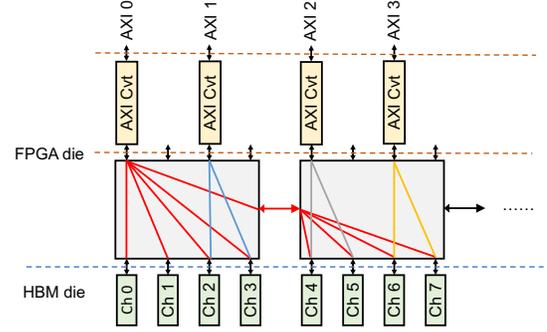


Fig. 9: Usage of the first 4 AXI interfaces. The same behavior is repeated for rest of the 12 AXI interfaces.

To address this issue, we utilize the internal crossbar among HBM channels. For each merge tree of phase 1, we only need one AXI module to read from and write to two adjacent HBM channels. This way halves the AXI area overhead and effectively reduces the routing congestion in the bottom die of the FPGA.

5.2 Data Layout Optimization

Although each AXI interface from the user side can access any of the 32 HBM channels, non-local data accesses could cause contention over the lateral connections between channel groups, thus leading to bandwidth decrease. To access an HBM channel outside the group, the data must occupy and traverse through each of the lateral connections until it reaches the destination. Two memory accesses requiring the same lateral connection will cause a conflict and one will be blocked.

Therefore, we must carefully arrange for the data placement to minimize out-of-group channel accesses and avoid conflicts in lateral connections. Table 1 summarizes the target HBM channels for each AXI interface. In the first phase, merge tree i uses AXI- i to only read/write a pair of nearby HBM channels $2i$ and $2i + 1$, where i ranges from 0 to 15. In the second phase, the merge tree $4i$ uses AXI- $4i$ to read/write between channels $8i+2j$ and channels $8i+2j+1$, where i & j both range from 0 to 3.

TABLE 1: Correspondence between each user side AXI and the HBM channels it accesses in TopSort, $i=0-3$.

AXI NO.	HBM Channel NO.
AXI- $4i$	$8i - (8i+7)$
AXI- $(4i + 1)$	$(8i+2) - (8i+3)$
AXI- $(4i + 2)$	$(8i+4) - (8i+5)$
AXI- $(4i + 3)$	$(8i+6) - (8i+7)$

Figure 9 illustrates the access pattern of the first four AXI interfaces in different colors. AXI-1, 2 and 3 will only access the two channels within its crossbar group. Although the AXI- $4i$ in the second phase needs to access channels outside of its local group, our data layout guarantees that different AXI interfaces will not cause conflicts in lateral connections.

5.3 Burst Size Optimization

For TopSort, each of the AXI interfaces will always access the HBM in burst mode. We need to properly choose the

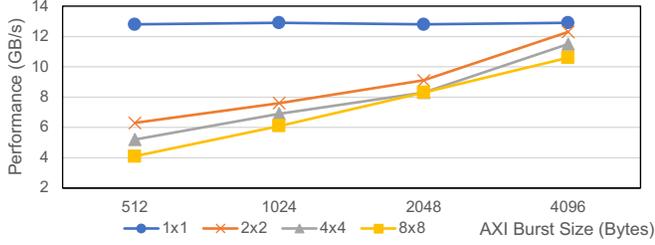


Fig. 10: HBM channel performance of inter-crossbar & intra-crossbar behaviors when varying the AXI burst size. The number of outstanding bursts is fixed at 32.

burst sizes because a large burst requires excessive buffers while a small burst may not fully utilize the bandwidth.

In order to select proper burst sizes, we profile the HBM performance based on the memory access pattern of each AXI interface. As shown in Fig. 9, In phase one, each AXI reads from one channel and writes to the adjacent channel; In phase two, AXI-0 reads from four even/odd channels in a round-robin way and writes to the other four nearby odd/even channels.

We define reading from m HBM channels and writing to their nearby m channels as pattern $m \times m$. The patterns of phase 1 and phase 2 are thus 1×1 and 4×4 by this definition. For completeness, we also perform tests for patterns 2×2 and 8×8 . Then we measured the achieved bandwidth when using different AXI burst sizes as seen in Figure 10.

Based on our experiments, any AXI burst size from 512 B to 4 KB can achieve the peak HBM bandwidth for the 1×1 pattern, which is used by merge trees in phase 1. However, for the 4×4 inter-crossbar memory access pattern used in phase 2, the AXI burst size needs to be maximized to 4 KB to achieve the best bandwidth. Therefore, we set the burst size to be 1 KB for the 12 merge trees that are not reused in phase 1, and 4 KB for the 4 merge trees that are reused in phase 2.

Minimizing the burst sizes under the bandwidth constraints could significantly save resources because not all buffers are implemented in BRAM. For each merge tree, each leaf node requires a 512-bit wide input buffer (see Figure 4) to hold the read bursts from AXI interfaces. When implemented using BRAM, each buffer consumes at least 7.5 BRAMs on Xilinx FPGAs. Since we have 16 merge trees and each tree has 16 leaves, the FPGA does not have enough BRAMs to implement all the buffers. So, about 3/4 of the buffers are implemented using LUT-based shift registers. While BRAM-based buffers are insensitive to the burst size, the area of a LUT-based buffer is directly proportional to the buffer depth. A single LUT can implement a 1-bit shift register with a depth of 32 and 512 LUTs can hold two AXI bursts of 1 KB. Reducing the bursts size from 4 KB to 1 KB directly reduces the LUTs consumption by $4 \times$. Further reducing the burst size from 1 KB to 512 B does not save more LUTs because 2 bursts of 512 B still require 512 LUTs. The actual sorting results with different burst sizes in TopSort is further discussed in Section 6.5.

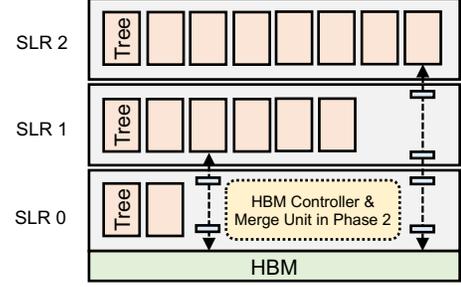


Fig. 11: The floorplan of TopSort on the Xilinx U280 FPGA. Each Super Logic Region (SLR) is a single FPGA die slice in the Stacked Silicon Interconnect (SSI) FPGA device

5.4 Floorplanning Optimization

Thanks to the parallel sorting in the first phase and merge tree reuse in the second phase, we are able to separate each of the merge tree kernels independently and apply a simple yet efficient resource model to better distribute these kernels across the three FPGA dies, as shown in Figure 11. Please note that the merge tree kernels in the first phase only expose simple AXI interfaces and each of the merge tree has the same size. This feature allows us to easily implement a coarse-grained floorplanning with the granularity of the separate merge tree kernel. In contrast, if we simply scale the single merge tree and want to make use of all of the FPGA dies, it would require significantly more engineering efforts to manually split the design logics and fit them into each die.

Considering that the sorted output from the merge tree of phase 2 will be directly written to the HBM channels, a natural choice is to place the extra merge units on the bottom die to minimize the signal crossing. After that, we will migrate as many merge tree kernels to the mid and top dies, to save more spaces in the bottom die to alleviate the routing issues.

We derive a simple resource model to solve this floorplanning problem. Let L be the resource consumption of each merge tree, u_i be the number of merge trees and a_i be the available resources on the i -th die, w be the width of an AXI interface, and W be the available signals crossing two neighbor dies. Then the problem can be formulated as *maximizing* $(u_1 + u_2)$ under the following constraints:

$$\begin{cases} (u_1 + u_2) \cdot w \leq W & (1) \\ u_1 \cdot L \leq a_1 & (2) \\ u_2 \cdot L \leq a_2 & (3) \end{cases}$$

where constraint (1) limits the AXI connections from all merge trees in both die 1 and 2 to die 0 do not exceed the available amount of signals crossing die 0 and 1, and constraints (2) and (3) make sure the merge trees placed on die 1 and 2 do not exceed each die's available resource.

By solving the above problem, we place eight merge trees on the top die and six merge trees on the middle die. There are another two merge trees remaining on the bottom die; we select these two merge trees as merge tree 0 and merge tree 8. The detailed layout are presented in Section 6.3.

Finally, to improve the timing, the signals crossing dies need to be pipelined with registers [27]. In general, we add 2 stages of pipeline registers when a signal crosses one die and add 4 stages of pipeline registers when it crosses two dies.

6 EXPERIMENTAL RESULTS

6.1 Experimental Setup

We perform all of the experiments on the Xilinx U280 FPGA board. The kernel is developed using System Verilog and is synthesized and implemented using Xilinx Vitis 2020.2. We use the pblock method [28] to place the design modules during the floorplanning. The input data are in a key-value pair format, each with a 32-bit key and a 32-bit value. The keys used in the experiments have a uniform distribution. To ensure the keys are fully randomized, we first generate N records whose keys are incremented from 1 to N , then we use the `random_shuffle()` function from the python library to shuffle the records. The sorted results are validated by checking whether their keys are from 1 to N .

6.2 TopSort Configuration & Resource Utilization

We implement 16 merge trees in total, each with 16 leaves and outputs 8 elements (each 8 bytes) per cycle. In phase 1, TopSort uses all 16 merge trees. In phase 2, TopSort reuses 4 of the merge trees to form a wider merge tree that can sort 32 elements (each 8 bytes) per cycle.

The resources utilization of the dynamic region² is listed in Table 2, including our TopSort kernel and HBM controllers. TopSort consumes 54.6% LUTs, 35.5% FFs and 56% BRAMs.

Table 3 shows the resource breakdown of each type of the merge trees. For those trees that are not reused (e.g., merge tree 1), each of them takes less than 3% of the available resources. This is consistent with the observation in Section 3 that it takes significantly less resources for a single merge tree to saturate one HBM channel than multiple HBM channels. Meanwhile, a reused tree (e.g., merge tree 4) requires more LUTs because its input buffers need to accommodate larger data bursts. As for the extra 3 merge units used in phase 2, they takes much more resources than a single merge tree, which validates the theoretical analysis that directly scaling the tree throughput requires the resources to grow superlinearly.

TABLE 2: Resources utilization and percentage of the TopSort kernel and HBM controllers. DSPs are used to calculate the addresses of AXI read and write.

Components	LUTs	Flip-Flops	BRAMs	DSPs
Available	1,066,848	2,144,089	1,487	8,484
TopSort kernel	582,244	750,505	834	84
Percentage	54.6%	35.5%	56.0%	1.0%
HBM Controllers	87,848	113,814	4	0
Percentage	8.2%	5.3%	0.2%	0.0%

². The FPGA is partitioned into two regions, with the dynamic region reserved for user logic and the static region used for infrastructure IPs.

TABLE 3: Resources utilization of individual merge trees. Here we pick one merge tree from each type since they have slightly different resource consumption.

Components	LUTs	Flip-Flops	BRAMs	DSPs
Merge tree 1	28,788	39,905	37.5	2
Percentage	2.7%	1.9%	2.5%	0.02%
Merge tree 4	39,195	39,459	60	10
Percentage	3.7%	1.8%	4.0%	0.1%
Extra logic of phase 2	60,239	102,864	144	20
Percentage	5.6%	4.8%	9.7%	0.2%

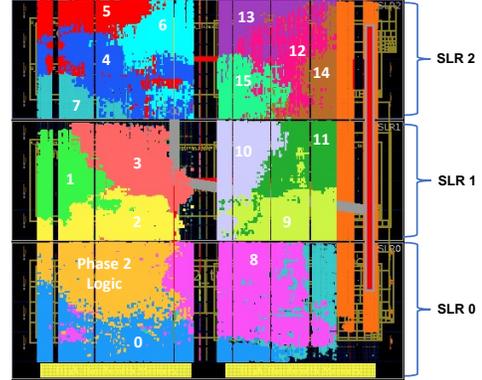


Fig. 12: Layout of TopSort implemented on the Xilinx U280 FPGA board. Number 0-15 label the 16 merge trees of phase 1. The orange part in the bottom die represents the extra logic used to form the wider merge tree of phase 2.

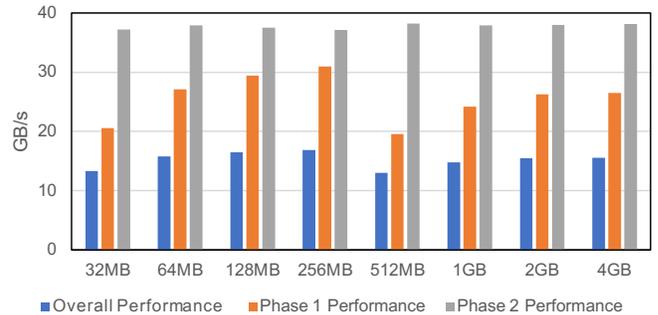


Fig. 13: Sorting performance with different data sizes. The overall performance is calculated through dividing the data size by the total sorting execution time.

6.3 Design Layout & Frequency

Figure 12 shows the final placements of all merge trees and the extra phase 2 logic. Our floorplanning and pipelining effectively reduces local routing congestion, so that TopSort could achieve 214 MHz in user clock and 414 MHz in the HBM control clock. Without our floorplanning, Vivado failed in routing even with the highest optimization level.

6.4 Overall Sorting Performance

Figure 13 presents the performance of sorting data ranging from 32 MB to 4 GB in size. The maximal data size that we can handle is half of the HBM capacity, which is 4 GB in the Xilinx U280 FPGA. This is because half of the HBM channels are used for reading and the other half for writing.

TABLE 4: Comparison with existing FPGA based in-memory sorters.

	Algorithm	Off-Chip Memory Used	Merge Tree Throughput	Design Frequency	Overall Sorting Perf.
FPGA'20 [15]	Sample sort	4 DDR4 channels	N/A	250 MHz	4.3 GB/s
FPL'20 [8]	Merge tree sort	1 DDR4 channel	4 GB/s	214 MHz	1.9 GB/s ³
ISCA'20 [7]	Merge tree sort	4 DDR4 channels	32 GB/s	250 MHz	7.1 GB/s
TopSort	Merge tree sort	32 HBM channels	219 GB/s (phase 1) 55 GB/s (phase 2)	214 MHz	15.6 GB/s

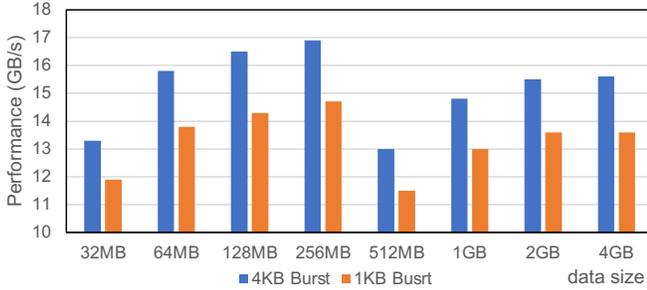


Fig. 14: Overall sorting performance with the same AXI burst size in phase one but different AXI burst sizes in phase two.

When sorting 4 GB of data, TopSort achieves an overall performance of 15.6 GB/s. Please note that 4 GB is half of the HBM capacity and the maximal size that we are able to sort. Sorting 4GB data has a performance of 26.5 GB/s in the first phase and 38 GB/s in the second phase.

Figure 13 also shows a performance drop when the data size is increased from 256 MB to 512 MB. This is because sorting 512MB data takes 6 passes in phase 1 while sorting 256 MB data takes 5 passes. Besides, sorting 512 MB to 4 GB data requires the same number of passes in phase 1, but the sorting performance still varies, which is similar for sorting 32 MB to 256 MB data. This is because some input leaves of the merge trees may not be fully active in the last pass of phase 1, similar to what is illustrated in Figure 7.

Based on our measurement, we are able to utilize at least 318 GB/s of HBM bandwidth in the first phase, where we run 6 passes and each pass reads and writes to HBM channels simultaneously. Thus the average HBM bandwidth utilized is at least $26.5 \times 6 \times 2 = 318$ GB/s. In fact, the used bandwidth is even higher because we have not account for the idle time of the merge units, which needs to be reset after merging two sorted sequences. In the initial pass, the length of the sorted sequences is 1, so the merge units have to be reset frequently.

6.5 Sorting Performance with Different Burst Sizes

Figure 14 compares the performance of two burst size choices. Both tests set the burst size as 1 KB in phase 1. In the second phase, we set the burst sizes to be 1 KB and 4 KB. The results show that reading in 4KB bursts in phase 2 gives better performance, which matches the profiling results in Section 5.3 that the designers need to use 4KB AXI bursts to maximize the HBM channels' bandwidth when they use one AXI to access multiple HBM channels.

6.6 Comparison with State-of-the-art CPU & FPGA Sorters

TopSort is an in-memory sorter [29], which stores the entire input in main memory such as an HBM or a DRAM. An in-memory sorter has much higher overall sorting performance than an in-storage sorter, which requires a second-level storage such as SSD to sort a larger dataset and thus is bounded by the slow I/O [5], [9]. In fairness, we compare TopSort against several CPU and FPGA in-memory sorters here. The state-of-the-art CPU in-memory sorter in [24] can sort 4GB of data in 1.7 second, achieving an overall sorting performance of about 2.3 GB/s. Compared to it, TopSort achieves $6.7\times$ speedup.

Table 4 lists the existing FPGA in-memory sorters. The best-performing DRAM-based FPGA sorter in [7] can fully utilize 4 DDR4 DRAM channels on the AWS F1 datacenter FPGA. When sorting 4 GB data, it achieves an overall sorting performance of 7.1 GB/s. Comparing to it, TopSort achieves $2.2\times$ speedup. Besides, we also list the absolute merge tree throughput, which is the number of elements (p) that the merge tree can output per cycle multiplied by the design frequency and the element's width in bytes, to show the utilization of the off-chip memory bandwidth of different merge tree designs. TopSort has an entire tree throughput of 219 GB/s in phase 1 and saturates the 420 GB/s HBM bandwidth (HBM is half-duplex so half of its bandwidth is for reading data and the other half is for writing data). In phase 2, TopSort still has $1.7\times$ and $14\times$ higher merge tree throughput than that of [7], [8], respectively.

One may find that the overall sorting performance improvement of TopSort over [7] is much less than the $6\times$ off-chip bandwidth increase. This is because the on-chip resources of the Xilinx U280 FPGA is merely enriched by about 20% compared to the AWS F1 FPGA. As analyzed in Section 2, the on-chip resources become the new bottleneck to linearly scale the overall sorting performance. While TopSort could address this issue to some extent, its overall performance is still bound by the merge tree in phase 2. That is why only $2.2\times$ speedup is reported.

6.7 Comparison with Scaled Single Merge Tree

We also compare TopSort to a single giant merge tree, which is scaled to output 32 elements per cycle and has a throughput equal to writing 4 HBM channels, i.e., the same as the merge tree throughput of phase two in TopSort. Unfortunately, we are not able to route this single merge tree, as it requires significantly more engineering efforts of manual placement optimization compared to TopSort. The

3. [8] reported $49\times$ speedup over the `C++ std::sort()` implementation on an ARM Cortex A53 core. We re-implement `C++ std::sort()` on a more advanced Intel Xeon E5-2680 core and multiply the performance by $49\times$ to get the listed value.

reason why it is easier to floorplan TopSort than to floorplan single merge tree is explained in Section 5.4. Nonetheless, we can still estimate its performance, assuming it achieved the same design frequency as TopSort.

The best estimation is that the single merge tree has the same number of leaves as the sum of the 16 merge tree leaves in phase one of TopSort, which is 256. This means sorting 4 GB data still requires 4 passes. Since the throughput of the merge tree is the same as TopSort’s throughput in phase 2, which is 38 GB/s, the overall sorting performance of the single merge tree is thus at most $38/4 = 9.5$ GB/s, much less than the 15.6 GB/s achieved in TopSort.

7 RELATED WORK

Merge Tree Acceleration. The existing works involved in merge tree acceleration are split into two directions. One direction focuses on optimizing the resource consumption of the merge unit itself, at the center of which is a variation of parallel sorting networks such as bitonic or even-odd sorting networks [11], [12], [13], [14]. TopSort benefits from these works in that using efficient merge units may always reduce the on-chip resource consumption of the merge tree. In this work, we adopt the merge units from [12], which contains two bitonic mergers, but rewrite the whole control logic to make the merge unit has the pure streaming behavior. The other direction investigates how to design efficient merge trees that interact with various memory layers such as DRAM and storage [5], [6], [7], [8], [9]. These works all scale the single tree to match the off-chip bandwidth. Our work addresses a different challenge where the bottleneck of merge tree sorting is shifted to limited on-chip resources.

HBM-Specific Optimizations. As for HBM-specific optimizations, [22] presents an HLS design that applies a similar floorplanning strategy and achieves a design frequency of 237 MHz when using 18 HBM channels. The majority of the recent HBM-based accelerators [19], [20], [23] are HLS designs but are not able to get more than 190 MHz and use more than 28 HBM channels. [21] implements a hash join accelerator that uses 32 HBM channels while running at 250 MHz and its random memory accesses are through 256-bit wide AXI interfaces. However, using a 256-bit wide AXI interface can only utilize half of the available HBM bandwidth at most. Compared to it, our work relies on continuous memory reads and writes, which requires the AXI interface to be 512-bit wide to maximize the memory performance and consumes significantly more on-chip resources for the AXI rate converters as well as the AXI burst buffers. Besides, we need to carefully optimize the data layout to avoid access conflicts in the lateral connections of the built-in AXI crossbars, which is never revealed in those applications with random memory access patterns. We believe that our work provides valuable insights on how to optimize HBM-based accelerator designs, especially HLS-based designs that are struggling to fully utilize the HBM bandwidth with a high design frequency.

8 CONCLUSION

In this work, we present TopSort, a high-performance two-phase sorting accelerator specialized for HBM-based FPGAs. Our analysis shows that the sorter performance on

HBM-based FPGAs is bounded by the limited on-chip resources. To achieve better performance, TopSort proposes a novel two-level sorting solution with smaller merge trees. In the first phase, TopSort can fully utilize all the HBM bandwidth. In the second phase, TopSort reuses the logic from the first phase to avoid the resource contention. Moreover, TopSort adopts several HBM-specific optimizations to further reduce the resource overhead and improve the bandwidth utilization. Finally, it also employs coarse-grained floorplanning to achieve better time closure. TopSort is the first HBM-based FPGA accelerator that can fully utilize all the HBM channel bandwidth and achieves $6.7\times$ and $2.2\times$ speedup over state-of-the-art CPU sorter and DRAM-based FPGA sorter.

ACKNOWLEDGMENT

This work is partially supported by CRISP, one of six JUMP centers and the CDSC industrial partners, including Samsung and Siemens Mentor Graphics.

REFERENCES

- [1] 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 (FPGA)*, 2011.
- [2] J. Casper and K. Olukotun, “Hardware Acceleration of Database Operations,” in *Proceedings of the 22th ACM/SIGDA international symposium on Field Programmable Gate Arrays (FPGA)*, 2014.
- [3] J. Matai, D. Richmond, D. Lee, Z. Blair, Q. Wu, A. Abazari, and R. Kastner, “Resolve: Generation of High-Performance Sorting Architectures from High-Level Synthesis,” in *Proceedings of the 24th ACM/SIGDA international symposium on Field-Programmable Gate Arrays (FPGA)*, 2016.
- [4] W. Song, D. Koch, M. Lujan, and J. Garside, “Parallel Hardware Merge Sorter,” in *Proceedings of the 24th IEEE international symposium on Field-programmable Custom Computing Machines (FCCM)*, 2016.
- [5] S.-W. Jun, S. Xu, and Arvind, “Terabyte Sort on FPGA-Accelerated Flash Storage,” in *Proceedings of the 25th IEEE international symposium on Field-programmable Custom Computing Machines (FCCM)*, 2017.
- [6] K. Manev and D. Koch, “Large Utility Sorting on FPGAs,” in *Proceedings of the 2018 International Conference on Field-Programmable Technology (FPT)*, 2018.
- [7] N. Samardzic, W. Qiao, V. Aggarwal, M.-C. F. Chang, and J. Cong, “Bonsai: High-performance adaptive merge tree sorting,” in *Proceedings of the 47th ACM/IEEE international symposium on computer architecture (ISCA)*, 2020.
- [8] P. Papaphilippou, C. Brooks, and W. Luk, “An Adaptable High-Throughput FPGA Merge Sorter for Accelerating Database Analytics,” in *Proceedings of the 30th International Conference on Field-Programmable Logic and Applications (FPL)*, 2020.
- [9] W. Qiao, J. Oh, L. Guo, M.-C. F. Chang, and J. Cong, “Fans: Fpga-accelerated near-storage sorting,” in *Proceedings of the 29th IEEE international symposium on Field-programmable custom computing machines (FCCM)*, 2021.
- [10] R. Chen, S. Siritiyal, and V. Prasanna, “Energy and Memory Efficient Mapping of Bitonic Sorting on FPGA,” in *Proceedings of the 23th ACM/SIGDA international symposium on Field Programmable Gate Arrays (FPGA)*, 2015.
- [11] S. Mashimo, T. V. Chu, and K. Kise, “High-Performance Hardware Merge Sorter,” in *Proceedings of the 25th IEEE international symposium on Field-programmable Custom Computing Machines (FCCM)*, 2017.
- [12] M. Saitoh, E. A. Elsayed, T. V. Chu, S. Mashimo, and K. Kise, “A high-performance and cost-effective hardware merge sorter without feedback datapath,” in *Proceedings of the 26th IEEE international symposium on Field-programmable custom computing machines (FCCM)*, 2018.

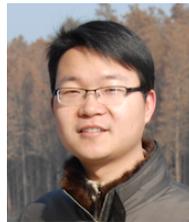
- [13] P. Papaphilippou, C. Brooks, and W. Luk, "FLiMS: Fast Lightweight Merge Sorter," in *Proceedings of the 2018 International Conference on Field-Programmable Technology (FPT)*, 2018.
- [14] E. A. Elsayed and K. Kise, "Towards an Efficient Hardware Architecture for Odd-even Based Merge Sorter," in *Proceedings of the 13th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoc)*, 2019.
- [15] H. Chen, S. Madaminov, M. Ferdman, and P. Milder, "FPGA-Accelerated Samplesort for Large Data Sets," in *Proceedings of the 28th ACM/SIGDA international symposium on Field Programmable Gate Arrays (FPGA)*, 2020.
- [16] A. Lu, Z. Fang, W. Liu, and L. Shannon, "Demystifying the memory system of modern datacenter fpgas for software programmers through microbenchmarking," in *2021 International Symposium on Field-Programmable Gate Arrays*, ser. FPGA '21, 2021, p. 105–115.
- [17] Y. Choi, Y. Chi, W. Qiao, N. Samardzic, and J. Cong, "Hbm connect: high-performance hls interconnect for fpga hbm," in *Proceedings of the 29th ACM/SIGDA international symposium on Field programmable gate arrays (FPGA)*, 2021.
- [18] Xilinx. [Online]. Available: https://www.xilinx.com/content/dam/xilinx/support/documents/data_sheets/ds963-u280.pdf
- [19] Y. Hu, Y. Du, E. Ustun, and Z. Zhang, "Graphlily: Accelerating graph linear algebra on hbm-equipped fpgas," in *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. IEEE, 2021, pp. 1–9.
- [20] L. Song, Y. Chi, L. Guo, and J. Cong, "Serpens: A high bandwidth memory based accelerator for general-purpose sparse matrix-vector multiplication," *arXiv preprint arXiv:2111.12555*, 2021.
- [21] J. Wirth, J. A. Hofmann, L. Thostrup, C. Binnig, and A. Koch, "Scalable and flexible high-performance in-network processing of hash joins in distributed databases," in *Proceedings of the 2021 International Conference on Field-Programmable Technology (FPT)*, 2021, pp. 1–9.
- [22] Y. Du, Y. Hu, Z. Zhou, and Z. Zhang, "High-performance sparse linear algebra on hbm-equipped fpgas using hls: A case study on spmv," 2022.
- [23] Y. Chi, L. Guo, and J. Cong, "Accelerating sssp for power-law graphs," in *Proceedings of the 2022 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, 2022.
- [24] M. Cho, D. Brand, and R. Bordawekar, "PARADIS: An efficient parallel algorithm for in-place radix sort," in *Very Large Data Bases (VLDB)*, 2015.
- [25] M. Gao, G. Ayers, and C. Kozyrakis, "Practical near-data processing for in-memory analytics frameworks," in *2015 International Conference on Parallel Architecture and Compilation (PACT)*, 2015, pp. 113–124.
- [26] J. L. Bentley, D. Haken, and J. B. Saxe, "A general method for solving divide-and-conquer recurrences," *ACM SIGACT News*, vol. 12, no. 3, pp. 36–44, 1980.
- [27] L. Guo, Y. Chi, J. Wang, J. Lau, W. Qiao, E. Ustun, Z. Zhang, and J. Cong, "Autobridge: Coupling coarse-grained floorplanning and pipelining for high-frequency hls design on multi-die fpgas," in *Proceedings of the 29th ACM/SIGDA international symposium on Field programmable gate arrays (FPGA)*, 2021.
- [28] Xilinx. [Online]. Available: https://www.xilinx.com/support/documents/sw_manuals/xilinx2020_2/ug938-vivado-design-analysis-closure-tutorial.pdf
- [29] D. E. Knuth, *Art of Computer Programming: Sorting and Searching*, 2nd ed. Addison-Wesley Professional, 1998.



Weikang Qiao received his B.S. degree in Information and Communication Engineering from Zhejiang University and his M.S. degree in Electrical Engineering from UCLA. Currently, he is a fifth-year Ph.D. student in the Electrical and Computer Engineering department at UCLA. His research focuses on customized accelerator architecture designs and performance modeling across various memory hierarchies, such as DRAM, High-bandwidth Memory (HBM) and SSDs. He is an IEEE student member.



Licheng Guo received his B.S. degree in Electrical Engineering from Zhejiang University in 2018. Currently he is a 4th-year Ph.D student in the UCLA CS department. His research focus on co-optimizing HLS compilers (from C++ to RTL) and physical design tools (from RTL to hardware) to improve the circuit maximal frequency and reduce the compilation time.



Zhenman Fang received his PhD degree in Computer Science from Fudan University, China in 2014. He did his postdoc at UCLA from 2014 to 2017, and worked as a Staff Software Engineer at Xilinx, San Jose, from 2017 to 2019. Currently, Zhenman is an Assistant Professor in School of Engineering Science, Simon Fraser University, Canada. His recent research focuses on customizable computing with specialized hardware acceleration, including emerging application characterization and acceleration, novel accelerator-rich and near-data computing architecture designs, and corresponding programming, runtime, and tool support. He is a member of the IEEE and ACM.



Mau-Chung Frank Chang is the Wintek Chair in Electrical Engineering and Distinguished Professor at UCLA. He pioneered the development of the world's first multi-gigabit/sec data converters in heterojunction bipolar technologies; first mm-Wave Radio-on-Chip with Digitally Controlled on-chip Artificial Dielectric (DiCAD) for broadband frequency tuning and on-chip sensing/actuating with cautious feedback control for achieving self-diagnosis and self-healing capabilities. He realized the first CMOS frequency synthesizers up to terahertz spectra and demonstrated tri-color and 3-dimensional CMOS active imagers at the (sub)-mmWave spectra based on a time-encoded digital architecture. He is a Member of the US National Academy of Engineering, a Fellow of US National Academy of Inventors, an Academician of Academia Sinica of Taiwan, and a Lifetime Fellow of IEEE. He was honored with the IEEE David Sarnoff Award in 2006 for developing and commercializing GaAs HBT and BiFET power amplifiers, which dominated smartphones transmitter worldwide production throughout the past 2.5 decades.



Jason Cong received his B.S. degree in computer science from Peking University in 1985, his M.S. and Ph. D. degrees in computer science from the University of Illinois at Urbana-Champaign in 1987 and 1990, respectively. Currently, he is the Volgenau Chair for Engineering Excellence (and former department chair) at the UCLA Computer Science Department, with joint appointment from the Electrical Engineering Department. Dr. Cong's research interests include novel architectures and compilation for customizable computing and quantum computing. He has over 500 publications in these areas, including 16 best paper awards, three 10-Year Most Influential Paper Awards, and three papers in the FPGA and Reconfigurable Computing Hall of Fame. He was elected to an IEEE Fellow in 2000, an ACM Fellow in 2008, a member of the National Academy of Engineering in 2017, and a Fellow of the National Academy of Inventors in 2020.