



# A Compiler Framework for Optimizing Dynamic Parallelism on GPUs

Mhd Ghaith Olabi<sup>1</sup>, Juan Gómez Luna<sup>2</sup>, Onur Mutlu<sup>2</sup>, Wen-mei Hwu<sup>3,4</sup>, Izzat El Hajj<sup>1</sup>

<sup>1</sup>American University of Beirut, Lebanon    <sup>2</sup>ETH Zürich, Switzerland    <sup>3</sup>NVIDIA, USA

<sup>4</sup>University of Illinois at Urbana-Champaign, USA

**Abstract**—Dynamic parallelism on GPUs allows GPU threads to dynamically launch other GPU threads. It is useful in applications with nested parallelism, particularly where the amount of nested parallelism is irregular and cannot be predicted beforehand. However, prior works have shown that dynamic parallelism may impose a high performance penalty when a large number of small grids are launched. The large number of launches results in high launch latency due to congestion, and the small grid sizes result in hardware underutilization.

To address this issue, we propose a compiler framework for optimizing the use of dynamic parallelism in applications with nested parallelism. The framework features three key optimizations: thresholding, coarsening, and aggregation. Thresholding involves launching a grid dynamically only if the number of child threads exceeds some threshold, and serializing the child threads in the parent thread otherwise. Coarsening involves executing the work of multiple thread blocks by a single coarsened block to amortize the common work across them. Aggregation involves combining multiple child grids into a single aggregated grid.

Thresholding is sometimes applied manually by programmers in the context of dynamic parallelism. We automate it in the compiler and discuss the challenges associated with doing so. Coarsening is sometimes applied as an optimization in other contexts. We propose to apply coarsening in the context of dynamic parallelism and automate it in the compiler as well. Aggregation has been automated in the compiler by prior work. We enhance aggregation by proposing a new aggregation technique that uses multi-block granularity. We also integrate these three optimizations into an open-source compiler framework to simplify the process of optimizing dynamic parallelism code.

Our evaluation shows that our compiler framework improves the performance of applications with nested parallelism by a geometric mean of  $43.0\times$  over applications that use dynamic parallelism,  $8.7\times$  over applications that do not use dynamic parallelism, and  $3.6\times$  over applications that use dynamic parallelism with aggregation alone as proposed in prior work.

## I. INTRODUCTION

Dynamic parallelism on GPUs allows threads running on the GPU to launch grids of threads to also run on the GPU. It is useful for programming applications with nested parallelism, particularly where the amount of nested parallelism is irregular and cannot be predicted at the beginning of the computation. An example of such an application is graph processing where a thread visiting a vertex may want to perform some work for each of its neighbors. In this case, the *parent* thread visiting the vertex may launch a grid with many *child* threads, one for each neighbor, to work on the neighbors concurrently.

Prior work has shown that using dynamic parallelism in this way imposes a high performance penalty [14, 38, 40]. The

key problem is that when a massive number of small grids are launched, the launch latency is high because of the congestion caused by the large number of launches, and the device is underutilized because the grids are small in size. To address this issue, various hardware and software optimizations have been proposed to reduce the overhead of dynamic parallelism.

Hardware techniques that have been proposed for reducing the overhead of dynamic parallelism include adding thread blocks to existing grids rather than launching new grids [38] or providing a hardware controller that advises programmers on whether a launch is profitable [34]. Further hardware optimizations include locality-aware scheduling of parent and child grids [35, 39]. However, these techniques require hardware changes so they are not available on current GPUs [30], which motivates the need for software techniques.

One category of software techniques is to have threads in the parent grid perform the nested parallel work without performing a dynamic launch [9, 42]. These techniques mitigate the overhead of dynamic parallelism by avoiding it entirely, but require parent threads to be on standby regardless of whether or not there is work available for them to do. Another category of techniques is to consolidate or aggregate the child grids being launched by multiple parent threads into a single grid [14, 24, 25, 41]. These techniques mitigate the overhead of dynamic parallelism by reducing the number of grids launched, hence the congestion, and increasing the sizes of the grids to ensure better hardware utilization.

In this paper, we propose a compiler framework for optimizing the use of dynamic parallelism that features three key optimizations: thresholding, coarsening, and aggregation. The first optimization, thresholding, involves launching a grid dynamically only if the number of child threads exceeds a certain threshold, and serializing the work in the parent thread otherwise. This optimization reduces the number of grids launched, thereby reducing congestion, and ensures that only large grids are launched that properly utilize the hardware. Some prior works assume that thresholding is applied manually by the programmer [24, 25, 34, 41]. However, manual application of thresholding complicates the launch code, requires code duplication, and hurts code readability. We propose to automate the thresholding optimization in the compiler and discuss the challenges associated with doing so.

The second optimization, coarsening, involves combining multiple child thread blocks into a single one. This optimization

reduces the number of child thread blocks that need to be scheduled, and interacts with the aggregation optimization to amortize the overhead of aggregation across the work of multiple original child blocks. While coarsening is a common optimization applied in various contexts [22, 26, 32], we propose to apply it in the context of dynamic parallelism and observe its benefit in combination with aggregation.

The third optimization, aggregation, is similar to prior works [14, 24, 25, 41] that aggregate child grids launched by multiple parent threads into a single grid. However, the granularity of aggregation in prior work has been limited to launches by parent threads in either the same warp, the same thread block, or the entire grid. We further enhance aggregation by proposing a new aggregation technique that uses multi-block granularity. In this technique, launches are aggregated across parent threads in a group of thread blocks as opposed to the two extremes used in prior work of a single thread block or the entire grid.

In summary, we make the following contributions:

- We present a compiler transformation that automates thresholding for dynamic parallelism (Section III).
- We propose to apply coarsening in the context of dynamic parallelism, and present a compiler transformation for doing so (Section IV).
- We propose to apply aggregation at multi-block granularity and present a compiler transformation for doing so (Section V).
- We combine thresholding, coarsening, and aggregation in one open-source compiler framework (Section VI).

Our evaluation (Section VIII) shows that our compiler framework improves the performance of applications with nested parallelism by a geometric mean of  $43.0\times$  over applications that use dynamic parallelism,  $8.7\times$  over applications that do not use dynamic parallelism, and  $3.6\times$  over applications that use dynamic parallelism with aggregation alone as proposed in prior work.

## II. BACKGROUND

### A. Dynamic Parallelism

Fig. 1(a) shows an example of how dynamic parallelism can be used in practice. In this example, parent threads executing on the GPU each discover some nested work that can be parallelized. Each parent thread launches a child grid to perform the nested work in parallel, and each parent thread provides its child grid with a different set of launch configurations and parameters. The amount of nested work may vary across threads. Hence, the child grids have different sizes. One source of inefficiency that may arise when using dynamic parallelism in this way is that a massive number of child grids may be launched, and many of them may be small in size [40]. In this case, the large number of child grid launches causes congestion, and the small size of the child grids causes the device to be underutilized.

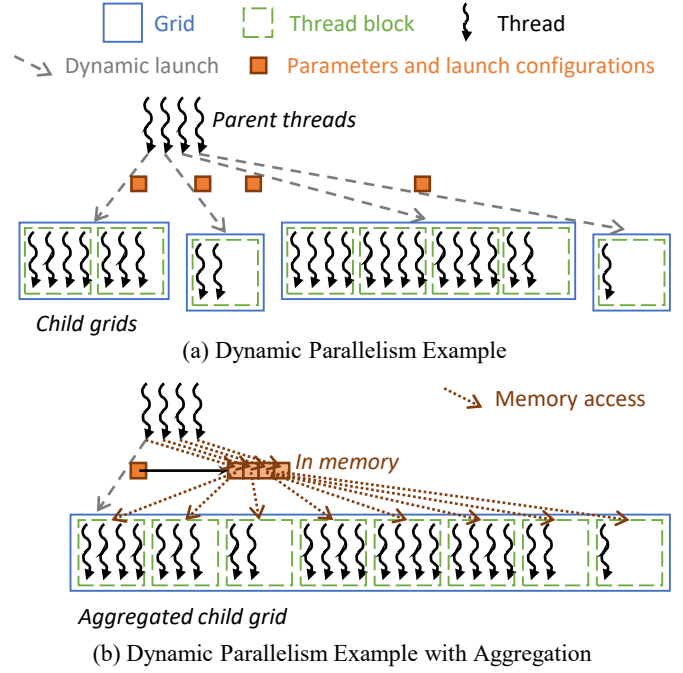


Fig. 1. Background on Dynamic Parallelism

### B. Aggregation

Aggregation is an optimization that consolidates or aggregates child grids that are launched by multiple parent threads into a single grid. Various works [14, 24, 25, 41] apply this optimization either manually or in the compiler. Fig. 1(b) shows how aggregation can be applied to the example in Fig. 1(a). In this example, the parent threads coordinate to find the cumulative size of all their child grids in order to launch a single aggregated grid. In prior work, the scope of coordination has been across parent threads in the same warp, block, or grid. If the scope is across a warp or a block, one of the participating parent threads launches the aggregated grid on behalf of the others. If the scope is across a grid, the aggregated grid is launched from the host. The scope of coordination is referred to as the *aggregation granularity*.

In the original code, each parent thread may provide different parameters and launch configurations to its child grid. However, in the transformed code, only one set of parameters and launch configurations can be provided. For this reason, before launching the aggregated grid, the parent threads each store their original parameters and launch configurations in memory, and a pointer to this memory is passed to the aggregated grid. The child threads must then identify who their original parent thread is in order for them to load the right parameters and configurations from memory. To do so, each child thread block executes a search operation. The work done by the parent threads to identify the size of the aggregated grid and store their individual parameters and launch configurations in memory is referred to as the *aggregation logic*. The work done by the child threads to identify their original parent thread and load their original parameters and configurations from memory is referred to as the *disaggregation logic*.

The advantage of aggregation is that it reduces the number of child grids launched, thereby reducing congestion, and it increases the sizes of the child grids to ensure better hardware utilization. The disadvantage of aggregation is that it incurs overhead due to the aggregation and disaggregation logic, and it delays the child grid launches until all parent threads are ready to launch. The choice of aggregation granularity (warp, block, or grid) involves making a trade-off between these advantages and disadvantages. Using a larger granularity reduces congestion and increases utilization, but incurs higher overhead from the aggregation and disaggregation logic and delays the child grids longer before launching them.

### III. THRESHOLDING

#### A. Optimization Overview

Thresholding is an optimization where a child grid is only launched dynamically if the number of child threads exceeds a certain *threshold*. Otherwise, the child threads are executed sequentially by the parent thread. Fig. 2 illustrates how thresholding can be applied to the example in Fig. 1(a). In this example, two of the parent threads have a small number of child threads. The benefit gained from parallelizing these child threads is unlikely to be worth the launch overhead. For this reason, the parent threads instead execute the work of the child threads sequentially.




Fig. 2. Dynamic Parallelism Example with Thresholding

Thresholding is commonly applied manually by programmers [24, 25, 34, 41]. However, manual application of thresholding complicates the launch code, requires code duplication, and hurts code readability. For this reason, we propose to automate thresholding via a compiler transformation (Section III-B) and discuss the challenges associated with doing so (Sections III-C and III-D).

#### B. Code Transformation

Fig. 3 shows an example of how our compiler applies the thresholding transformation. The original code in Fig. 3(a) consists of a parent kernel (lines 04-08) and a child kernel (lines 01-03). The parent kernel calls the child kernel using dynamic parallelism (line 06) and configures it with a grid dimension `gDim` and a block dimension `bDim`.

Fig. 3(b) shows the code after the thresholding transformation is applied. The transformation consists of two key parts: constructing a serial version of the child to be executed by the parent thread (lines 09-15) and applying a threshold to either perform the launch or call the serial version (lines 21-26).

```

01  __global__ child(params) {
02      child body
03  }

04  __global__ parent(...) {
05      ...
06      child <<< gDim, bDim >>> (args);
07      ...
08  }

```

(a) Original Code

```

09  __device__ child_serial(params, dim3 _gDim, dim3 _bDim) {
10      for(_bx = 0; _bx < _gDim.x; ++_bx) {
11          for(_tx = 0; _tx < _bDim.x; ++_tx) {
12              child body // Replace uses of blockIdx.x with _bx,
13          }              // threadIdx.x with _tx, gridDim with
14      }                  // _gDim, and blockDim with _bDim
15  }

16  __global__ child(params) {
17      child body
18  }

19  __global__ parent(...) {
20      ...
21      _threads = ...; // Extracted from gDim expression
22      if(_threads >= _THRESHOLD) {
23          child <<< gDim, bDim >>> (args);
24      } else {
25          child_serial (args, gDim, bDim);
26      }
27      ...
28  }

```

(b) Code after Thresholding Transformation

Fig. 3. Thresholding Code Transformation Example

To construct the serial version, the child kernel is replicated and its attribute is changed to `__device__` so that it becomes a device function (line 09). Two parameters are appended to the parameter list: `_gDim`, which represents the original grid dimension of the parallel version, and `_bDim`, which represents the original block dimension. Loops are inserted around the child body (lines 10-11) to serialize the child threads. The first loop (line 10) iterates over the child thread blocks while the second loop (line 11) iterates over the child threads in each child block, using the bounds passed as parameters. Finally, all uses of the reserved index and dimension variables are replaced with the corresponding loop indices and bounds. The example shows a 1-dimensional child kernel for simplicity, however, if the child kernel is multi-dimensional, loops would be inserted for each dimension.

To apply the threshold, the number of child threads is first identified and stored in the `_threads` variable (line 21). Section III-D discusses how the number of child threads is identified. Next, an if-statement is inserted around the child kernel call (line 22) to ensure that the dynamic launch is

performed only if the number of child threads is greater than or equal to `_THRESHOLD`. Here, `_THRESHOLD` is a macro variable that can be overridden at compile time for tuning purposes. If the threshold is not met, then the device function implementing the serial version is called instead (line 25), thereby serializing the child work in the parent thread.

### C. Non-Transformable Kernels

Not all child kernels are amenable to the kind of transformation described in Section III-B. In particular, there are two kinds of child kernels that we do not transform: (1) child kernels that perform barrier synchronization across threads via `__syncthreads()` or warp-level primitives, and (2) child kernels that use shared memory.

For child kernels that perform barrier synchronization, serializing GPU threads while supporting such synchronization has been done in the literature [12, 18, 20, 21, 33]. The key strategy is to divide the code into regions separated by the barriers and insert loops around each region. However, these techniques target serializing multiple GPU threads in one CPU thread. Extending them to serialize multiple GPU threads in one GPU thread is not practical for two reasons. The first reason is that these techniques perform scalar expansion of all local variables to preserve the state of all threads across barriers. Such a scalar expansion on the GPU would convert all register accesses to memory accesses which would be prohibitively expensive. The second reason for not serializing child threads that perform barrier synchronization is that code that includes barrier synchronizations often implements a parallel algorithm that is not efficient when serialized. For example, a parallel reduction operation uses a reduction tree and leverages barriers to synchronize between levels of the tree. However, reduction trees are not an efficient way to perform sequential reductions. It is more efficient to use a simple reduction loop. In this case, it is better to let the programmer apply the thresholding optimization manually because the best sequential and parallel algorithms are different.

For child kernels that use shared memory, we do not construct a serial version of the kernel because every parent thread would require as much shared memory as an entire child block which would make the shared memory requirements of a parent block too high. Besides, kernels that use shared memory most often use `__syncthreads()` to coordinate access to shared memory across threads. Hence, these kernels will most likely not be transformable anyway for the reasons related to barrier synchronization previously mentioned.

### D. Identifying the Number of Child Threads

The transformation described in Section III-B needs to identify the desired number of child threads in order to compare it with the threshold. Identifying the desired number of child threads is challenging because this information is not what the programmer provides in the kernel call. Instead, the programmer provides the grid dimension (number of blocks) and the block dimension (number of threads per block). The

programmer uses the desired number of child threads in calculating the grid dimension.

One way to identify the desired number of child threads is to multiply the grid dimension with the block dimension. However, this approach gives the total number of threads in the child grid including threads that will not be used. This value is not a representative value to compare with the threshold. For example, consider a kernel with a nested kernel call where the child block dimension is configured to 1024 threads. One of the parent threads identifies 2 units of nested parallel work to be processed by the child. This parent thread will configure its child kernel with 1 block. In this case, multiplying the grid dimension with the block dimension gives 1024 total threads which is much larger than the actual number of threads desired (2 threads). Ideally, the value that should be compared to the threshold is 2, not 1024. Hence, multiplying the grid dimension with the block dimension is not a good approach.

The approach we use to identify the desired number of child threads is based on the observation that programmers usually calculate the grid dimension as a ceiling-division of the desired number of threads by the block dimension. There are arbitrary ways in which a ceiling-division can be expressed so it is not possible to have a static analysis that always determines the desired number of threads with certainty. Instead, we identify the most common patterns used by programmers to calculate ceiling-division and employ a simple static analysis based on these patterns.

Fig. 4 shows common expressions that programmers write for calculating the grid dimension using ceiling-division. Options (a)-(c) use integer arithmetic, while options (d)-(e) convert to floating point and use the `ceil` function. Option (f) is used in multi-dimensional blocks, where the operands to the `dim3` constructor could each be an expression that resembles options (a)-(e). For all options, the expression may be expressed as a whole, or it may be expressed in parts where subexpressions are stored in intermediate variables. Note that `N` and `b` can be arbitrary expressions.

- |                                      |  |
|--------------------------------------|--|
| (a) $(N - 1) / b + 1$                | <code>N</code> : desired number of threads                   |
| (b) $(N + b - 1) / b$                | <code>b</code> : block dimension                             |
| (c) $N / b + (N \% b == 0) ? 0 : 1$  |  |
| (d) <code>ceil((float)N/b)</code>    |  |
| (e) <code>ceil(N/(float)b)</code>    |  |
| (f) <code>dim3(..., ..., ...)</code> | <code>dim3</code> args could be one of the above expressions |

Fig. 4. Common Expressions for Calculating the Grid Dimension

We observe from the examples in Fig. 4 that `N` is usually in the subexpression on the left hand side of the division. Moreover, the subexpression containing `N` may also contain constants such as 1 or `b` (which is usually a constant). Based on this observation, our analysis pass looks for a division operation, takes the subexpression on the left hand side, and removes additions and subtractions of constants, considering the remaining subexpression as the desired number of threads. This analysis is heuristic by nature and is not guaranteed to find



the true desired number of threads. However, using a heuristic is acceptable in this context because the result will only be used to choose whether to serialize or parallelize the work. This choice does not impact program correctness in any way.

The subexpression that is found is assigned to `_threads` in Fig. 3 on line 21. The occurrence of the subexpression in `gDim` is then replaced with `_threads` to ensure that the expression is not duplicated in the code just in case the expression has side effects.

#### IV. COARSENING

##### A. Optimization Overview

Coarsening is an optimization where the work of multiple thread blocks in the original code is assigned to a single thread block. The number of original thread blocks assigned to each coarsened thread block is referred to as the *coarsening factor*. When the GPU hardware is oversubscribed with more thread blocks than it can accommodate simultaneously, the hardware serializes these thread blocks, scheduling a new one whenever an old one has completed. There are multiple advantages of serializing these thread blocks in the code rather than letting the hardware do it. First, it reduces the number of thread blocks that need to be scheduled, and allows some warps in the coarsened block to proceed to executing the work of the next original block before other warps have completed their part in the previous original block. Second, if there is common work across the original thread blocks, that work can be factored out and executed once by the coarsened thread block, allowing its cost to get amortized. The disadvantage of coarsening is that it reduces parallelism, thereby underutilizing the device if the coarsening factor is too high.

Coarsening as an optimization is often applied by programmers in many different contexts [22]. Prior works have also applied coarsening in the compiler [26, 32], but not in the context of dynamic parallelism. We propose to apply coarsening in the context of dynamic parallelism, and automate its application via a compiler transformation.

Fig. 5 illustrates how coarsening can be applied to the child thread blocks in the example in Fig. 1(a). In this example, each coarsened child thread block in the transformed code executes the work of two child thread blocks in the original code. The advantage of applying coarsening in the context




Fig. 5. Dynamic Parallelism Example with Coarsening

```

01  __global__ child(params, _gDim) {
02      for(_bx = blockIdx.x; _bx < _gDim.x; _bx += gridDim.x) {
03          child body // Replace uses of blockIdx.x with _bx
04      }              // and gridDim with _gDim
05  }

06  __global__ parent(...) {
07      ...
08      _cgDim = _gDim = gDim ;
09      _cgDim.x = (_gDim.x + _CFCTOR - 1) / _CFCTOR;
10      child <<< _cgDim, bDim >>>(args, _gDim);
11      ...
12  }

```

Fig. 6. Coarsening Code Transformation Example

of dynamic parallelism is that it reduces the number of child thread blocks that need to be scheduled. More importantly, when applied before aggregation, coarsening also has the advantage of providing child thread blocks that do more work per block, which makes them more capable of amortizing the disaggregation logic overhead.

##### B. Code Transformation

Fig. 6 shows an example of how our compiler applies the coarsening transformation to the original code in Fig. 3(a). The transformation consists of two key parts: coarsening the child kernel (lines 01-05) and modifying the launch configurations to launch the coarsened child (lines 08-10).

To coarsen the child kernel, a parameter `_gDim` is appended to the parameter list (line 01) which represents the original grid dimension without coarsening. A *coarsening loop* is inserted (line 02) that iterates over the work of the original child thread blocks assigned to the coarsened block. Uses of the reserved index and dimension variables are replaced with the corresponding loop indices and bounds. The example shows coarsening in one dimension only for simplicity, however, if the child grid is multi-dimensional, loops would be inserted for each dimension.

To modify the launch configurations to launch the coarsened child, the original grid dimension `gDim` is stored in a variable `_gDim` (line 08). The value is also copied to `_cgDim`, which represents the coarsened grid dimension. The x-dimension of the coarsened grid dimension `_cgDim` is then ceiling-divided by the coarsening factor `_CFCTOR` (line 09). Here, `_CFCTOR` is a macro variable that can be overridden at compile time for tuning purposes. Again, the example shows coarsening in one dimension for simplicity. Finally, the child kernel is configured with the coarsened grid dimension, and the original grid dimension is passed as a parameter (line 10).

#### V. AGGREGATION

Aggregation has been proposed by prior work [14, 24, 25, 41] and has been described in Section II-B. In this paper, we propose a new aggregation granularity, namely, multi-block granularity (Section V-A). We also propose to apply

an aggregation threshold to optimize aggregation at warp and block granularity (Section V-B).

#### A. Multi-block Granularity Aggregation

Recall from Section II-B that prior work has performed aggregation at warp, block, and grid granularity. Using a larger granularity reduces congestion and increases hardware utilization, but incurs higher overhead from the aggregation and disaggregation logic and delays the child grids longer before launching them. The choice of aggregation granularity presents a trade-off between these performance factors. However, there is a wide gap between the block and grid granularity aggregation schemes that leaves a large part of the trade-off space unexplored. A parent thread block typically has hundreds to up to 1,024 parent threads, whereas a parent grid can have many thousands to millions of parent threads. To better explore the trade-off space, we propose an intermediate granularity between block and grid granularity, which is multi-block granularity aggregation.

In multi-block granularity aggregation, we divide a parent grid into groups of blocks, each group having a fixed number of blocks. The threads within the same group of blocks collaborate to aggregate their child grids into a single aggregated grid. In prior work [14], the aggregation logic involves a scan operation on the original grid dimensions and a max operation on the block dimension to identify the aggregated grid configuration. It also involves a barrier synchronization to wait for all participating threads to store their individual configurations and arguments to memory before the aggregated launch is performed. Since we cannot synchronize across multiple thread blocks, we perform the scan and max operations for multi-block granularity using atomic operations similar to what is done at grid granularity in prior work. As for the barrier synchronization, we replace it with a group-wide counter that is atomically incremented by each thread block when it finishes. The last thread block in the group to increment the counter performs the aggregated launch.

Fig. 7 shows an example of how our compiler applies the multi-block granularity aggregation transformation to the original code in Fig. 3(a). The transformation consists of two parts: the aggregation logic in the parent kernel (lines 14-35) and the disaggregation logic in the child kernel (lines 01-11).

For the aggregation logic in the parent kernel, we save the `gDim` and `bDim` expressions in temporary variables to avoid recomputing them every time we use them in case they have side effects (lines 14-15). We then identify the group that the thread block belongs to (line 16) and find the group's memory segment in a pre-allocated memory buffer (line 17). This memory segment is used to store the configurations and arguments that threads in the group pass to their children. Next, each thread that launches a child grid (has a non-zero grid dimension) atomically increments two global counters simultaneously (lines 19-20): (1) `_numParents` to assign an index to the parent thread so that the thread knows where to store its arguments and configuration, and (2) `_sumGDim` to find the total number of child blocks of prior parent threads

```

01 __global__ child(_paramsArray, _gDimScannedArray, _bDimArray) {
02     _parentIdx = binary search in _gDimScannedArray
03     params = _paramsArray[_parentIdx]
04     _gDim = _gDimScannedArray[_parentIdx] - _gDimScannedArray[_parentIdx - 1]
05     _bx = blockIdx.x - _gDimScannedArray[_parentIdx - 1]
06     _bDim = _bDimArray[_parentIdx]
07     if(threadIdx < _bDim) {
08         child body // Replace uses of blockIdx.x with _bx
09                     // and gridDim with _gDim
10     }
11 }

12 __global__ parent(...) {
13     ...
14     _gDim = gDim
15     _bDim = bDim
16     _groupIdx = blockIdx.x / AGG_GRANULARITY
17     find group's memory segments in a pre-allocated buffer based on _groupIdx
18     if(_gDim > 0) {
19         (_parentIdx, _sumPrevGDim) =
20             atomicAdd(&_numParents[_groupIdx], _sumGDim[_groupIdx]), (1, _gDim))
21         _argsArray[_parentIdx] = args
22         _gDimScannedArray[_parentIdx] = _sumPrevGDim + _gDim
23         _bDimArray[_parentIdx] = _bDim
24         atomicMax(&_maxBDim[_groupIdx], _bDim)
25     }
26     __threadfence()
27     __syncthreads()
28     if(threadIdx == launcher thread in block) {
29         _nFinishedBlocks = atomicAdd(&_numFinishedBlocks[_groupIdx], 1) + 1
30         _isLastBlockToFinish = (_nFinishedBlocks == AGG_GRANULARITY)
31         if(!_isLastBlockToFinish) {
32             child <<< _sumGDim[_groupIdx], _maxBDim[_groupIdx] >>>
33                 (_argsArray, _gDimScannedArray, _bDimArray);
34         }
35     }
36     ...
37 }

```

Fig. 7. Multi-block Granularity Aggregation Code Transformation Example

which we use to initialize the scanned array of grid dimensions. The two global counters are incremented simultaneously by treating them as a single 64-bit integer. Each thread then stores its arguments, scanned grid dimension, and block dimension to memory so they can be passed to the child grid (lines 21-23). Each thread also performs an atomic operation to find the maximum block dimension (line 24).

After each thread writes its configuration and arguments to global memory, it performs a fence operation to ensure that the configuration and arguments are visible to its child blocks (line 26). This fence is not necessary in prior work because the visibility of the data is ensured either by the semantics of the dynamic kernel call (for warp and block granularity) or by grid termination (for grid granularity). However, it is necessary here because the launch may be performed by a different thread block than the one that writes the data. Hence, the writing thread block must ensure that its data is visible in global memory before notifying the other blocks that it is ready. A local barrier (line 27) is also needed to ensure that all threads in the block finish storing their data before notifying.

Finally, it is time to perform the aggregation launch. One thread in the block (line 28) atomically increments the group-wide counter (line 29) and checks if its block is the last block in the group to finish (line 30-31). If so, the thread launches the aggregated grid, configuring it with the sum of all grid

dimensions and the maximum of all block dimensions, and passing pointers to the memory arrays containing the original arguments and configurations (lines 32-33).

For the disaggregation logic in the child kernel (lines 01-11), the code remains largely the same as prior work [14]. Each child block performs a binary search through the scanned grid dimension array to identify its original parent thread (line 02). It then loads its parameters and configuration (lines 04-06) and performs the work of the child kernel based on them (lines 07-10).

Note that besides filling the gap in the trade-off space between block and grid granularity aggregation, multi-block granularity aggregation has another advantage over grid granularity aggregation in particular. Grid granularity aggregation requires the CPU to be involved in the aggregated launch whereas the multi-block granularity aggregation runs entirely on the GPU, which frees the CPU to perform other tasks. Hence, multi-block granularity aggregation is more compatible with the asynchronous semantics of kernel calls. If the CPU is needed for performing other tasks, multi-block granularity aggregation may be better for overall execution time even if grid granularity aggregation has better kernel time.

### B. Aggregation Threshold

When thresholding is applied before aggregation, the number of original child grids that participate in the aggregated grid may be substantially reduced. If there is an insufficient number of original child grids participating in the aggregated grid, the benefit of aggregation may not be worth its overhead. To address this issue, we enhance aggregation with another optimization that applies an *aggregation threshold*. The aggregation logic is preceded by an operation to count the number of participating parent threads. If the number of participating parent threads does not meet a certain threshold, the child grids are launched normally by their parent threads instead of being aggregated. Therefore, a child grid may be executed in one of three ways: it may be serialized within its parent, launched directly by its parent, or launched as part of an aggregated grid. Since applying an aggregation threshold requires parent threads to synchronize to count the number of participating threads, it can only be applied at warp and block granularity where barrier synchronization across threads is possible.

## VI. COMPILER FRAMEWORK

We integrate our three optimizations – thresholding, coarsening, and aggregation – into a single compiler framework. For separation of concerns, each optimization is implemented as a separate source-to-source transformation pass that takes a CUDA `.cu` file and generates a `.cu` file. The transformations are independent, meaning that any combination of them could be applied in any order while generating correct code.

Although the optimizations can be applied in any order, we apply them in the following order: thresholding, coarsening, then aggregation, as shown in Fig. 8(a). Thresholding is applied before coarsening because coarsening manipulates the grid dimension which makes it harder to extract the

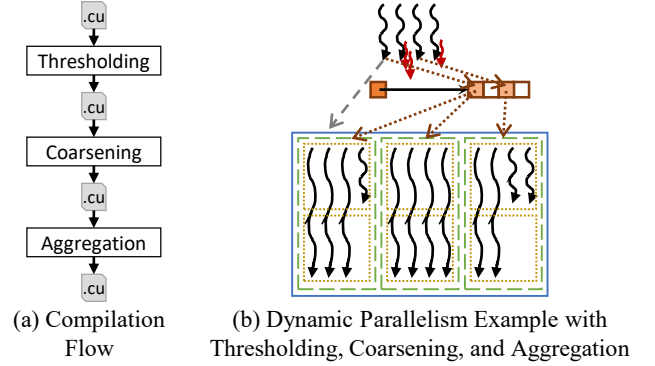


Fig. 8. Combining the three optimizations

number of threads that needs to be compared to the threshold. Thresholding is applied before aggregation because aggregation combines small grids with large grids into a single aggregated grid. It is more difficult to isolate small grids and serialize them after they have been aggregated into larger ones. Coarsening is applied before aggregation because the disaggregation logic should be outside the coarsening loop so that it can be amortized across multiple original child blocks.

Figure 8(b) illustrates the impact of combining all three optimizations on the example in Fig. 1(a). In this example, two of the parent threads have small child grids so the work of the child grids is serialized in the parent threads by the thresholding optimization. The remaining two parent threads have large child grids so they collaborate to aggregate their grids and perform a single launch. Each child block in the aggregated grid searches for its parent thread and obtains the corresponding parameters and configuration from memory. The child block then executes the work of multiple original child blocks because it was coarsened prior to aggregation.

The compiler transformations are implemented as source-to-source transformation passes in Clang [23]. The thresholding and coarsening transformations and their supporting analyses are implemented from scratch. The aggregation transformations are implemented by modifying the implementations provided by one of the prior works [14]. We leverage this work because it is open source, but our techniques can also be applied to any of the prior works that perform aggregation [14, 41].

## VII. METHODOLOGY

We evaluate our compiler framework on a system with a Volta V100 GPU [28] with 32GB of device memory and an AMD EPYC 7551P CPU [3] with 64GB of main memory. Table I shows the benchmarks and datasets used in our evaluation. We set the pending launch count appropriately to avoid overflowing the launch buffer pool [30]. We compile the benchmarks with per-thread default streams enabled to ensure that launches from the same block are not bottlenecked on the same default stream. We use larger datasets than prior work [14] does because we evaluate on a larger GPU. However, our evaluation on smaller datasets shows similar trends. Note

TABLE I  
BENCHMARKS AND DATASETS

Benchmark	Description	Dataset Used
BFS	Breadth First Search [10]	KRON, CNR
BT	Bezier Tessellation [29]	T0032-C16, T2048-C64
MSTF	Minimum Spanning Tree (find kernel) [8]	KRON, CNR
MSTV	Minimum Spanning Tree (verify kernel) [8]	KRON, CNR
SP	Survey Propagation [8]	RAND-3, 5-SAT
SSSP	Single Source Shortest Path [8]	KRON, CNR
TC	Triangle Counting [27]	KRON, CNR

Dataset	Description
KRON	kron_g500-simple-logn16, 65,536 vertices, 2,456,071 edges [31]
CNR	cnr-2000, 325,557 vertices, 2,738,969 edges [7]
T0032-C16	Max Tessellation: 32, Curvature: 16, Lines: 20,000 [29]
T2048-C64	Max Tessellation: 2048, Curvature: 64, Lines: 20,000 [29]
RAND-3	random-42000-10000-3, 10,000 literals [8]
5-SAT	5-SATISFIABLE, 117,296 literals [5]

that for TC, we use parts of the graphs in Table I due to memory constraints.

We compare our results to three different baselines. The *No CDP* versions of benchmarks are the original versions cited in Table I that do not use CUDA Dynamic Parallelism. The *CDP* versions use CUDA Dynamic Parallelism and are obtained from prior work [14]. The *KLAP (CDP+A)* versions perform aggregation only and are also obtained from prior work [14]. We use KLAP as a baseline from among prior works because it is open source and because we build on it in our compiler framework.

We report the performance of code versions generated by our compiler from the CDP version for multiple combinations of optimizations. We indicate the combinations of optimizations applied as follows: *T* for thresholding, *C* for coarsening, and *A* for aggregation. The only exception is the TC benchmark with CDP+T because the original benchmark already applies dynamic parallelism with thresholding.

For each combination of optimizations, we tune the relevant parameters and report results for the best configuration. The tuned parameters are the launch threshold, coarsening factor, and aggregation granularity. The threshold is not tuned beyond the largest dynamic launch size to ensure that at least one dynamic launch is performed. We use an exhaustive search to perform tuning to show the maximum potential of the optimizations and to present a complete view of the design space. However, such an exhaustive search is unnecessary in practice as we discuss in Section VIII-C.

To extract the breakdown of execution time in Section VIII-B, we incrementally deactivate portions of the code and calculate the time difference. For benchmarks with multiple iterations, we report the time for the longest running iteration. For BT, the aggregated `cudaMalloc` in the parent is considered part of the parent work because it is not affected by thresholding and coarsening.

## VIII. EVALUATION

### A. Performance

Fig. 9 shows the performance results for each benchmark and dataset with all combinations of optimizations applied. Performance is reported as speedup over the CDP version.

Applying CDP alone leads to a performance degradation in almost all cases compared to not applying CDP. Aggregation alone recovers from this degradation and substantially improves performance, where CDP+A is  $12.1\times$  faster than CDP and  $2.4\times$  faster than No CDP. These observations are consistent with those made in prior works [14, 24, 25, 41], which note that the large number of launches causes congestion and aggregating the grids mitigates this overhead.

Thresholding alone gives substantial speedup, where CDP+T is  $13.4\times$  (geomean) faster than CDP. Thresholding also gives speedup in the presence of aggregation, where CDP+T+A is  $2.9\times$  (geomean) faster than CDP+A, and CDP+T+C+A is also  $3.1\times$  (geomean) faster than CDP+C+A. The incremental benefit of thresholding with aggregation is not as pronounced as without aggregation because the benefit of thresholding is reducing the number of launches which aggregation also does. Nevertheless, the speedup is still significant.

Coarsening without aggregation gives modest speedup, where CDP+C is  $1.01\times$  (geomean) faster than CDP and CDP+T+C is  $1.09\times$  (geomean) faster than CDP+T. On the other hand, coarsening with aggregation gives more speedup, where CDP+C+A is  $1.16\times$  (geomean) faster than CDP+A, and CDP+T+C+A is  $1.22\times$  (geomean) faster than CDP+T+A. Notice how coarsening is synergistic with aggregation (its speedup in the presence of aggregation is greater than its speedup in the absence of aggregation). The reason is that in the presence of aggregation, coarsening helps amortize the disaggregation logic across more work in the coarsened child block in contrast with much less work in each original child block. Although the benefit of coarsening is much smaller




Fig. 9. Performance (higher is better)



than that of thresholding and aggregation, it is still substantial. A speedup of  $1.22\times$  is significant considering that it comes from a compiler optimization that requires no programmer intervention or additional architecture support.

Ultimately, our compiler framework as a whole substantially improves performance with the three optimizations combined, where CDP+T+C+A is  $43.0\times$  (geomean) faster than CDP. CDP+T+C+A is also  $8.7\times$  (geomean) faster than No CDP, showing that dynamic parallelism is a powerful programming feature when combined with the right optimizations. Moreover, CDP+T+C+A is  $3.6\times$  (geomean) faster than KLAP (CDP+A), showing that our framework provides substantial speedup over prior works that perform aggregation alone [14].

### B. Breakdown of Execution Time

Fig. 10 shows how the execution time is spent for each benchmark and dataset. We use KLAP (CDP+A) as baseline instead of CDP because prior work [14] has already shown how aggregation affects the execution time breakdown, and including CDP makes the figure illegible. We compare with CDP+T+A and CDP+T+C+A to show the incremental impact of thresholding and coarsening, respectively.

The first observation is that thresholding increases parent work and decreases child work. Thresholding serializes child work in parent threads so it is natural for parent threads to have more work to do and child threads to have less.

The second observation is that thresholding decreases the overhead from aggregation, launching, and disaggregation. Thresholding results in fewer parent threads launching and fewer child threads being launched. The aggregation overhead decreases because fewer parent threads participate, the launch overhead decreases because there are fewer launches, and the disaggregation overhead decreases because there are fewer child threads searching for their parents.

The third observation is that coarsening decreases the launch overhead. Since coarsening reduces the number of child blocks that need to be scheduled, the launch overhead decreases.

The fourth observation is that coarsening decreases the disaggregation overhead. Coarsening amortizes the disaggregation logic across multiple child thread blocks rather than having each child thread block perform its own disaggregation.

The final observation is that coarsening decreases parent work and increases child work for some benchmarks, namely BFS and SSSP. This result is unintuitive because coarsening does not change the work that the parent does. The reason is that for these two benchmarks, after coarsening is applied, the reduction in launch and disaggregation overhead results in a lower optimal threshold to be found. The lower threshold results in more work being offloaded from parent to child. This observation demonstrates an important interaction between thresholding and coarsening.

### C. Impact of Threshold and Aggregation Granularity

Fig. 11 shows how performance varies for different threshold values and aggregation granularity while keeping the coarsening factor constant at the best coarsening factor found. For space constraints, we only show the results for one dataset.

The first observation is that for most benchmarks (all except SP), as the threshold increases initially, performance also improves. Increasing the threshold initially results in more of the small child grids getting serialized in their parent threads, which reduces the number of launches, hence the congestion, and avoids launching small grids that underutilize the device.

The second observation is that for some benchmarks (e.g., BFS, BT, MSTF, SSSP), increasing the threshold too much causes performance to degrade again. Increasing the threshold too much results in large child grids getting serialized within their parent threads, which reduces parallelism and causes high control divergence.

The third observation is that different benchmarks perform best with different levels of aggregation granularity. For example, SP and TC perform best with grid granularity, BFS and SSSP perform best with multi-block granularity, BT and MSTF perform best with block granularity, and MSTV performs best without aggregation. The fact that multi-block granularity aggregation performs best in some of the cases reflects the importance of providing an intermediate granularity between grid and block granularity.

As mentioned in Section VII, we use an exhaustive search to perform tuning to show the maximum potential of the optimizations and to present a complete view of the design space. However, from our experience, such a broad search is




Fig. 10. Breakdown of Execution Time (lower is better)

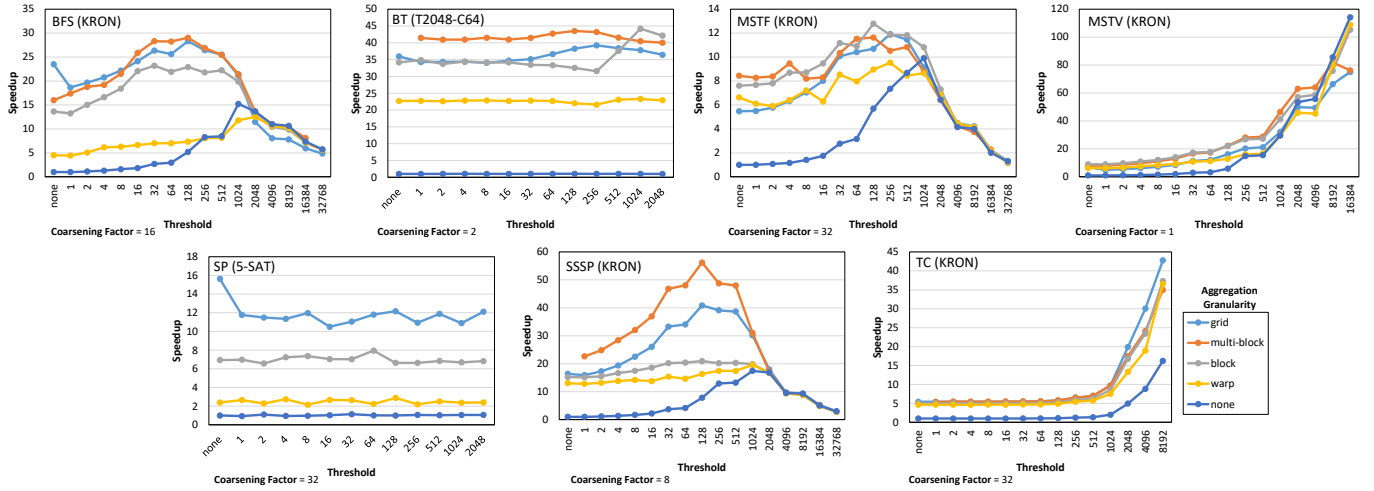


Fig. 11. Impact of Threshold and Aggregation Granularity

unnecessary. First, the best threshold is typically the one that allows approximately 6,000-8,000 child grid launches. Second, performance is not very sensitive to the coarsening factor provided that it is sufficiently large ( $>8$ ) so the coarsening factor does not need to be searched with a broad range or with high resolution. Third, aggregation at warp granularity is never favorable. With these observations in mind, users can typically find a combination of parameters that is very close to the best with less than ten runs. Moreover, the compiler framework exposes these parameters in a configurable manner to make it easy for users to leverage off-the-shelf autotuners [4]. Such tuning is worthwhile for kernels that run repeatedly on similar datasets. On the other hand, if a user cannot afford to tune the kernel, it is not necessary for the user to find the best parameters to benefit from the optimizations. Selecting a sub-optimal set of parameters would still yield a speedup, just not the maximum possible speedup. For example, if we fix the threshold to 128 for all benchmarks and datasets, then CDP+T+C+A will have a geomean speedup of  $1.9\times$  over CDP+C+A, as opposed to  $3.1\times$  when the best threshold is used. Hence, our optimizations are still useful even if the best parameters are not searched for or found.

#### D. Workloads with Low Nested Parallelism

Dynamic parallelism is useful when the amount of nested parallelism is high such that the launch overhead is worth the parallel work extracted. If the amount of nested parallelism is low for all parent threads, the benefit of dynamic parallelism is limited. To demonstrate the impact of dynamic parallelism on applications with a small amount of nested parallelism, we evaluate the graph benchmarks on a road graph (USA-road-NY [11]). The graph has 264,346 vertices, 730,100 edges, an average degree of 3, and a maximum degree of 8. Hence, each vertex has a small number of outgoing edges so processing the graph has a small amount of nested parallelism.

Fig. 12 shows the performance results for each graph benchmark on the road graph with all combinations of optimizations




Fig. 12. Performance of Graph Benchmarks on Road Graphs (higher is better)

applied. It is clear that CDP versions perform substantially more poorly on this graph relative to the No CDP versions. Our proposed optimizations are able to recover much of the performance degradation, but not all of it. In fact, in this experiment, we tune the threshold beyond the largest launch size such that CDP+T degenerates to serializing all child threads like No CDP. However, CDP+T still cannot recover fully. The reason is that the mere existence of a dynamic launch that is guarded by a condition that is always false such that the launch is never performed. Upon inspecting the assembly code of the two kernels, we observe that a large number of additional instructions are generated besides the instructions for performing the launch. Upon profiling the execution of the two kernels, we observe that a large number of additional instructions are executed even though the launch is never performed.

Note that the SP benchmark on the RAND-3 dataset performs poorly in Fig. 9 also because the amount of nested parallelism is low (all child grids have fewer than 32 threads). These results show the importance of being aware of the application and dataset when choosing whether or not to apply dynamic parallelism and its associated optimizations.

## IX. RELATED WORK

To our knowledge, our work is the first to provide a compiler framework for optimizing dynamic parallelism code that combines the thresholding, coarsening, and aggregation optimizations together.

Several benchmarking efforts [16, 36, 40] observe the inefficiencies of dynamic parallelism caused by high launch overhead and hardware underutilization, which has motivated various hardware and software optimizations.

Many hardware optimizations have been proposed for mitigating the overhead of dynamic parallelism. Dynamic Thread Block Launch (DTBL) [37, 38] proposes hardware support for lightweight dynamic launching of thread blocks rather than heavyweight dynamic launching of entire grids. The dynamically launched thread blocks are essentially added to existing grids on the fly by the hardware. LaPerm [39] extends DTBL with a locality-aware scheduler. SPAWN [34] is a hardware controller that advises programmers whether or not a dynamic launch is profitable. LASER [35] enhances dynamic parallelism with locality-aware scheduling. Hardware optimizations are promising for future generations of GPUs, but they are not available on current GPUs, which motivates the need for software optimizations. Moreover, hardware optimizations, if implemented, are potentially synergistic with the software optimizations we propose.

Many compiler/software optimizations have been proposed to improve dynamic parallelism performance or to provide an alternative to dynamic parallelism. CUDA-NP [42] is a compiler approach that enables annotation of parallel loops in the code with directives, instead of using dynamic parallelism. The compiler transforms the kernel to launch excess threads for each original thread. Using control flow, the excess threads are activated whenever a parallel loop is encountered. A later work [19] improves this technique by having multiple original threads share the excess threads for better load balance and less control divergence. Another approach, Free Launch [9], eliminates launches of child grids by applying transformations that reuse parent threads to execute child threads either sequentially or in parallel. These approaches mitigate the overhead of dynamic parallelism by avoiding it entirely, but require threads to be on standby regardless of whether or not there is work available for them to do.

Li et al. [24, 25], Wu et al. [41], and KLAP [13, 14] are aggregation techniques where multiple child grids are combined into a single aggregated grid to reduce the launch overhead. Zhang et al. [43, 44] further enhance this approach by grouping together child grids with similar optimal configurations rather than placing all child grids in the same aggregated grid. We leverage one of these works, KLAP [14], as the aggregation component in our flow.

KLAP [14] also includes another dynamic parallelism optimization, promotion, which targets a specific pattern where a single-block kernel calls itself recursively. Our optimizations are not applicable to this pattern. Thresholding is not applicable because all child grids have the same size. Coarsening is not

applicable because a child grid has only one block. Aggregation is not applicable because only one thread per parent grid performs a launch.

Various frameworks have been proposed to optimize the execution of applications with irregular parallelism on GPUs. Wireframe [1], Juggler [6], ATA [17], and BlockMaestro [2] facilitate the execution of irregular parallel workloads where data-dependences between thread blocks need to be enforced. VersaPipe [45] facilitates the extraction of pipeline parallelism from different GPU kernels. NestGPU [15] optimizes the execution of nested SQL queries on GPUs while avoiding the use of dynamic parallelism due to its inefficiency. Our work focuses on optimizing GPU applications with nested parallelism expressed using dynamic parallelism.

## X. CONCLUSION

We present an open-source compiler framework for optimizing the use of dynamic parallelism in applications with nested parallelism. The framework includes three key optimizations: thresholding, coarsening, and aggregation. Our evaluation shows that our compiler framework substantially improves performance of applications with nested parallelism that use dynamic parallelism, compared to when dynamic parallelism is not used or when it is used with aggregation only like in prior work.

## ACKNOWLEDGMENTS

This work is supported by the University Research Board of the American University of Beirut (URB-AUB-103782-25509).

## APPENDIX

### A. Abstract

Our artifact is a compiler for optimizing applications that use dynamic parallelism following the workflow illustrated in Figure 8(a). We have implemented the compiler in Clang [23] and have made the compiler code publicly available. Since building the compiler requires building Clang/LLVM which can be time and resource consuming, we provide pre-built binaries of the compiler in a Docker image, along with the required dependences and the benchmarks/datasets on which the compiler has been evaluated. Reviewers can use the compiler binaries to transform the benchmark CUDA code with our optimizations, then compile and run the code on a CUDA-capable GPU to verify the timing/speedup results reported in Section VIII. Scripts are provided to automate this process.

### B. Artifact Check-list (Meta-information)

- **Program:** The benchmarks used are listed in Table I. The benchmark code is obtained from prior work [14]<sup>1</sup>. We also include copies of the benchmark code in the artifact.
- **Compilation:** The CUDA code for the benchmarks before and after transformation requires NVCC to be compiled, which has been included in the Docker image in the artifact.
- **Transformations:** The software transformations are implemented in Clang [23]. The code for these transformations is in

<sup>1</sup>Prior work's code is available here: <https://github.com/illinois-impact/klap>

the artifact, and a pre-compiled binary from this code has been included in the Docker image in the artifact.

- **Binaries:** We include x86-x64 Linux binaries for our pre-compiled Clang-based compiler passes in the Docker image in the artifact.
- **Data set:** The datasets are listed in Table I and have been included in the artifact.
- **Run-time environment:** The Docker image in the artifact builds x86-x64 Linux binaries and includes all the dependences.
- **Hardware:** The evaluation requires a CUDA-capable GPU that can execute dynamic parallelism code. Our evaluation used a V100 GPU with 32GB of memory. We recommend having at least 16GB of GPU memory to support the datasets.
- **Metrics:** The metric reported is execution time/speedup.
- **Experiments:** Two scripts are provided in the artifact: one script that tests the best configuration (threshold value, coarsening factor, aggregation granularity) for each combination of optimizations (thresholding, coarsening, aggregation) to verify Figures 9 and 12; and one script that exhaustively tests all possible configurations for each combination of optimizations to verify Figure 11. In our evaluation, ten runs are used and an average is taken, but there is little variation across runs for most benchmarks/datasets. We do not include a script for reproducing the results in Figure 10 because the process of collecting these results (described in Section VII) is manual and difficult to automate.
- **Output:** Running the binary of a single benchmark outputs the time it takes for that benchmark to run. Running the experiment scripts outputs a CSV file with the execution time of each benchmark/dataset for each combination of compiler optimizations/configurations used.
- **How much disk space required (approximately)?:** < 4GB
- **How much time is needed to prepare the workflow (approximately)?:** < 1hr
- **How much time is needed to complete the experiments (approximately)?:** A single benchmark compilation and run should finish in < 1min. Running the script for the best configurations for all benchmarks should take < 1hr. Running the script for the exhaustive search can take up to 24hrs per benchmark.
- **Publicly available?:** Yes.
- **Code licenses (if publicly available)?:** MIT License.
- **Archived (provide DOI)?:** Yes, at the following link: <https://doi.org/10.6084/m9.figshare.17048447.v1>.

## C. Description

1) *How to Access:* The artifact can be downloaded at the following link: <https://doi.org/10.6084/m9.figshare.17048447.v1>.

The artifact consists of the code for the compiler passes, the benchmark code, the datasets, and a Docker image which includes the binaries of the compiler passes and the dependences required to execute the compiler passes, compile the transformed benchmark code, and execute the benchmarks.

2) *Hardware Dependencies:* Running the binaries depends on having a CUDA capable devices. Our evaluation used a V100 GPU with 32GB of memory. We recommend having at least 16GB of GPU memory to support the datasets.

3) *Software Dependencies:* We ran our tests on an environment with CUDA-9.1 installed. Software dependencies have been packaged in the Docker image in the artifact.

4) *Datasets:* The datasets are listed in Table I and have been included in the artifact.

## D. Installation

The artifact contains a README file with installation instructions and a convenience script for handling the installation.

## E. Experiment Workflow

To run the experiments, run the Docker image: `./run.sh`.

In order to compile all the binaries with default parameters, from inside the Docker container, change directory to the relevant benchmark inside the `test/` directory and run `make all`.

In order to run the benchmark with the best configurations, run `bestcombination.sh`. The parameters used here are the best combination of parameters we found for each benchmark and data set when we exhaustively searched the space.

In order to perform the exhaustive search, run `sweep.sh` which is available inside every benchmark directory.

## F. Evaluation and Expected Result

Running `bestcombination.sh` for each benchmark should provide the execution times used to report the speedups in Figures 9 and 12. Running `sweep.sh` for each benchmark should provide the execution times used to report the speedups in Figure 11. In our evaluation, ten runs are used and an average is taken, but there is little variation across runs for most benchmarks/datasets.

## G. Experiment Customization

The experiments can be customized by running each benchmark with the desired parameters. The parameters can be updated in the provided Makefile for each benchmark.

## REFERENCES

- [1] A. A. Abdolrashidi, D. Tripathy, M. E. Belviranli, L. N. Bhuyan, and D. Wong, "Wireframe: supporting data-dependent parallelism through dependency graph execution in GPUs," in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 2017, pp. 600–611.
- [2] A. Abdolrashidi, H. A. Esfeden, A. Jahanshahi, K. Singh, N. Abu-Ghazaleh, and D. Wong, "BlockMaestro: Enabling programmer-transparent task-based execution in GPU systems," in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2021, pp. 333–346.
- [3] AMD, "AMD EPYC 7551P." [Online]. Available: <https://www.amd.com/en/products/cpu/amd-epyc-7551p>
- [4] J. Ansel, S. Kamil, K. Veeramachaneni, J. Ragan-Kelley, J. Bosboom, U.-M. O'Reilly, and S. Amarasinghe, "OpenTuner: An extensible framework for program autotuning," in *Proceedings of the 23rd International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2014, pp. 303–316.
- [5] A. Belov, D. Diepold, M. J. Heule, and M. Järvisalo, "Proceedings of SAT Competition 2014," 2014.
- [6] M. E. Belviranli, S. Lee, J. S. Vetter, and L. N. Bhuyan, "Juggler: a dependence-aware task-based execution framework for GPUs," in *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. ACM, 2018, pp. 54–67.
- [7] P. Boldi and S. Vigna, "The WebGraph framework I: Compression techniques," in *Proc. of the Thirteenth International World Wide Web Conference (WWW 2004)*. Manhattan, USA: ACM Press, 2004, pp. 595–601.
- [8] M. Burtcher, R. Nasre, and K. Pingali, "A quantitative study of irregular programs on GPUs," in *2012 IEEE International Symposium on Workload Characterization (IISWC)*, 2012, pp. 141–151.



- [9] G. Chen and X. Shen, "Free launch: optimizing GPU dynamic kernel launches through thread reuse," in *Proceedings of the 48th International Symposium on Microarchitecture*. ACM, 2015, pp. 407–419.
- [10] 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*, ser. GPGPU-3. New York, NY, USA: Association for Computing Machinery, 2010, p. 63–74. [Online]. Available: <https://doi.org/10.1145/1735688.1735702>
- [11] DIMACS, "9th DIMACS implementation challenge - shortest paths," 2006.
- [12] I. El Hajj, "Dynamic loop vectorization for executing OpenCL kernels on CPUs," Master's thesis, University of Illinois, 2014.
- [13] —, "Techniques for optimizing dynamic parallelism on graphics processing units," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2018.
- [14] I. El Hajj, J. Gómez-Luna, C. Li, L.-W. Chang, D. Milojicic, and W.-m. Hwu, "KLAP: Kernel launch aggregation and promotion for optimizing dynamic parallelism," in *Microarchitecture (MICRO), 2016 49th Annual IEEE/ACM International Symposium on*. IEEE, 2016, pp. 1–12.
- [15] S. Floratos, M. Xiao, H. Wang, C. Guo, Y. Yuan, R. Lee, and X. Zhang, "NestGPU: Nested query processing on GPU," in *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 2021, pp. 1008–1019.
- [16] I. Harb and W.-C. Feng, "Characterizing performance and power towards efficient synchronization of GPU kernels," in *Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), 2016 IEEE 24th International Symposium on*. IEEE, 2016, pp. 451–456.
- [17] A. E. Helal, A. M. Aji, M. L. Chu, B. M. Beckmann, and W.-c. Feng, "Adaptive task aggregation for high-performance sparse solvers on GPUs," in *2019 28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 2019, pp. 324–336.
- [18] R. Karrenberg and S. Hack, "Improving performance of OpenCL on CPUs," in *International Conference on Compiler Construction*. Springer, 2012, pp. 1–20.
- [19] F. Khorasani, B. Rowe, R. Gupta, and L. N. Bhuyan, "Eliminating intra-warp load imbalance in irregular nested patterns via collaborative task engagement," in *Parallel and Distributed Processing Symposium, 2016 IEEE International*. IEEE, 2016, pp. 524–533.
- [20] H.-S. Kim, I. El Hajj, J. Stratton, S. Lumetta, and W.-M. Hwu, "Locality-centric thread scheduling for bulk-synchronous programming models on CPU architectures," in *Proceedings of the 13th Annual IEEE/ACM International Symposium on Code Generation and Optimization*. IEEE Computer Society, 2015, pp. 257–268.
- [21] H.-S. Kim, I. El Hajj, J. A. Stratton, and W.-M. Hwu, "Multi-tier dynamic vectorization for translating GPU optimizations into CPU performance," Center for Reliable and High-Performance Computing, Tech. Rep., 2014.
- [22] D. B. Kirk and W.-m. Hwu, *Programming Massively Parallel Processors: A Hands-On Approach*. Morgan Kaufmann, 2016.
- [23] C. Lattner and V. Adve, "LLVM: A compilation framework for lifelong program analysis & transformation," in *International Symposium on Code Generation and Optimization, 2004. CGO 2004*. IEEE, 2004, pp. 75–86.
- [24] D. Li, H. Wu, and M. Becchi, "Exploiting dynamic parallelism to efficiently support irregular nested loops on GPUs," in *Proceedings of the 2015 International Workshop on Code Optimisation for Multi and Many Cores*. ACM, 2015, p. 5.
- [25] —, "Nested parallelism on GPU: Exploring parallelization templates for irregular loops and recursive computations," in *Parallel Processing (ICPP), 2015 44th International Conference on*. IEEE, 2015, pp. 979–988.
- [26] A. Magni, C. Dubach, and M. O'Boyle, "Automatic optimization of thread-coarsening for graphics processors," in *Proceedings of the 23rd International Conference on Parallel Architectures and Compilation Techniques (PACT), 2014*, pp. 455–466.
- [27] V. S. Mailthody, K. Date, Z. Qureshi, C. Pearson, R. Nagi, J. Xiong, and W.-m. Hwu, "Collaborative (CPU + GPU) algorithms for triangle counting and truss decomposition," in *2018 IEEE High Performance Extreme Computing Conference (HPEC)*, 2018.
- [28] NVIDIA, "NVIDIA Tesla V100 GPU architecture: The world's most advanced data center GPU. version WP-08608-001\_v1." [Online]. Available: <https://nvidia.com/v100>
- [29] —, "CUDA samples v. 7.5," 2015.
- [30] —, "CUDA programming guide v. 11.2," 2021.
- [31] R. A. Rossi and N. K. Ahmed, "The network data repository with interactive graph analytics and visualization," in *AAAI*, 2015. [Online]. Available: <https://networkrepository.com>
- [32] N. Stawinoga and T. Field, "Predictable thread coarsening," *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 15, no. 2, pp. 1–26, 2018.
- [33] J. A. Stratton, S. S. Stone, and W.-m. Hwu, "MCUDA: An efficient implementation of CUDA kernels for multi-core CPUs," in *International Workshop on Languages and Compilers for Parallel Computing*. Springer, 2008, pp. 16–30.
- [34] X. Tang, A. Pattnaik, H. Jiang, O. Kayiran, A. Jog, S. Pai, M. Ibrahim, M. T. Kandemir, and C. R. Das, "Controlled kernel launch for dynamic parallelism in GPUs," in *High Performance Computer Architecture (HPCA), 2017 IEEE International Symposium on*. IEEE, 2017, pp. 649–660.
- [35] X. Tang, A. Pattnaik, O. Kayiran, A. Jog, M. T. Kandemir, and C. Das, "Quantifying data locality in dynamic parallelism in GPUs," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 2, no. 3, pp. 1–24, 2018.
- [36] Y. Ukidave, F. N. Paravecino, L. Yu, C. Kalra, A. Momeni, Z. Chen, N. Materise, B. Daley, P. Mistry, and D. Kaeli, "NUPAR: A benchmark suite for modern GPU architectures," in *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '15. ACM, 2015, pp. 253–264.
- [37] J. Wang, "Acceleration and optimization of dynamic parallelism for irregular applications on GPUs," Ph.D. dissertation, Georgia Institute of Technology, 2016.
- [38] J. Wang, N. Rubin, A. Sidelnik, and S. Yalamanchili, "Dynamic thread block launch: A lightweight execution mechanism to support irregular applications on GPUs," in *Proceedings of the 42nd Annual International Symposium on Computer Architecture*. ACM, 2015, pp. 528–540.
- [39] —, "Laperm: Locality aware scheduler for dynamic parallelism on GPUs," in *The 43rd International Symposium on Computer Architecture (ISCA)*, June 2016.
- [40] J. Wang and S. Yalamanchili, "Characterization and analysis of dynamic parallelism in unstructured GPU applications," in *Workload Characterization (IISWC), 2014 IEEE International Symposium on*. IEEE, 2014, pp. 51–60.
- [41] H. Wu, D. Li, and M. Becchi, "Compiler-assisted workload consolidation for efficient dynamic parallelism on GPU," in *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2016, pp. 534–543.
- [42] Y. Yang and H. Zhou, "CUDA-NP: Realizing nested thread-level parallelism in GPGPU applications," in *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP '14. Association for Computing Machinery, 2014, p. 93–106.
- [43] J. Zhang, "Transforming and optimizing irregular applications for parallel architectures," Ph.D. dissertation, Virginia Tech, 2018.
- [44] J. Zhang, A. M. Aji, M. L. Chu, H. Wang, and W.-c. Feng, "Taming irregular applications via advanced dynamic parallelism on GPUs," in *Proceedings of the ACM International Conference on Computing Frontiers*, 2018.
- [45] Z. Zheng, C. Oh, J. Zhai, X. Shen, Y. Yi, and W. Chen, "VersaPipe: A versatile programming framework for pipelined computing on GPU," in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 2017, pp. 587–599.