



PERKS: a Locality-Optimized Execution Model for Iterative Memory-bound GPU Applications

Lingqi Zhang*
Tokyo Institute of Technology
Japan
zhang.l.ai@m.titech.ac.jp

Mohamed Wahib†
RIKEN Center for Computational
Science
Japan
mohamed.attia@riken.jp

Peng Chen†‡
National Institute of Advanced
Industrial Science and Technology
Japan
chin.hou@aist.go.jp

Jintao Meng
Shenzhen Institutes of Advanced
Technology
China
jt.meng@siat.ac.cn

Xiao Wang†
Oak Ridge National Laboratory
USA
wangx2@ornl.gov

Toshio Endo
Tokyo Institute of Technology
Japan
endo@is.titech.ac.jp

Satoshi Matsuoka§
RIKEN Center for Computational
Science
Japan
matsu@acm.org

ABSTRACT

Iterative memory-bound solvers commonly occur in HPC codes. Typical GPU implementations have a loop on the host side that invokes the GPU kernel as much as time/algorithm steps there are. The termination of each kernel implicitly acts the barrier required after advancing the solution every time step. We propose an execution model for running memory-bound iterative GPU kernels: PERsistent KernelS (PERKS). In this model, the time loop is moved inside persistent kernel, and device-wide barriers are used for synchronization. We then reduce the traffic to device memory by caching subset of the output in each time step in the unused registers and shared memory. PERKS can be generalized to any iterative solver: they largely independent of the solver’s implementation. We explain the design principle of PERKS and demonstrate effectiveness of PERKS for a wide range of iterative 2D/3D stencil benchmarks (geomean speedup of 2.12x for 2D stencils and 1.24x for 3D stencils over state-of-art libraries), and a Krylov subspace conjugate gradient solver (geomean speedup of 4.86x in smaller SpMV datasets from SuiteSparse and 1.43x in larger SpMV datasets over a

state-of-art library). All PERKS-based implementations available at: <https://github.com/neo Zhang307/PERKS>.

CCS CONCEPTS

• **Computing methodologies** → *Concurrent computing methodologies; Parallel computing methodologies*; • **Computer systems organization** → *Parallel architectures*.

KEYWORDS

Persistent Kernel, Iterative Solvers, Memory-bound, GPU

ACM Reference Format:

Lingqi Zhang, Mohamed Wahib, Peng Chen, Jintao Meng, Xiao Wang, Toshio Endo, and Satoshi Matsuoka. 2023. PERKS: a Locality-Optimized Execution Model for Iterative Memory-bound GPU Applications. In *2023 International Conference on Supercomputing (ICS '23)*, June 21–23, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3577193.3593705>

1 INTRODUCTION

GPUs are becoming increasingly prevalent in HPC systems. More than half the systems on the Top500 [1] list include discrete GPUs and seven of the systems in the top ten are GPU-accelerated (November 2022 list). As a result, extensive efforts goes into optimizing iterative methods for GPUs, for instance: iterative stencils [16, 49, 53, 69] used widely in numerical solvers for PDEs, iterative stationary methods for solving systems of linear equations (ex: Jacobi [3, 43], Gauss–Seidel method [18, 28, 43]), iterative Krylov subspace methods for solving systems of linear equations (ex: conjugate gradient [6, 65], BiCG[5, 6], and GMRES[6, 19]).

Although the device memory bandwidth of GPUs has been increasing from generation to generation, the gap between compute and memory is widening. Given that iterative stencils and implicit

*Also with National Institute of Advanced Industrial Science and Technology.

†Corresponding authors

‡Also with RIKEN Center for Computational Science.

§Also with Tokyo Institute of Technology.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ICS '23, June 21–23, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0056-9/23/06...\$15.00
<https://doi.org/10.1145/3577193.3593705>

solvers typically have low arithmetic intensity [49], significant efforts goes into optimizing them for data locality. These included moving the bottleneck from device memory to on-chip scratchpad memory [48] or cache [47], or further pushing the bottleneck to the register files [16, 83]. Those efforts become increasingly effective since the aggregate volume of register files and scratchpad memory capacity are increasing with newer generations of GPUs [38]. In iterative solvers, due to spatial dependencies, a barrier is typically required at the end of each time step (or several time steps when doing temporal blocking [49]). That is to assure that advancing the solution in time step k would only start after all threads finish advancing the solution in time step $k - 1$. Invoking the kernels from the host side in each time step acts as an implicit barrier, where the kernel invocation in time step k would happen after all threads of the kernel invocation at time step $k - 1$ have finished execution. In-between kernel invocations, data stored in registers and scratchpad memory would be wiped out, and the next kernel invocation would start by reading its input from the device memory.

One opportunity to improve the data locality is to extend the lifetime of the kernel across the time steps and take advantage of the large volume of register files and scratchpad memory to reduce traffic to the device memory. In this paper, we propose a generic model for running iterative solvers on GPUs to improve data locality. PERsistent KernelS (PERKS)¹ have the time loop inside them, instead of the host, and use the recently supported device-wide barriers (in CUDA) for synchronization. Next, we identify the cache-able data in the solver: the data that is the output of time step $k - 1$ and input to time step k , as well as the repeatedly loaded constant data. Finally, we use either the scratchpad memory or registers (or both) to cache the data, and reduce the traffic to device memory.

The basic concept and implementation of PERKS is relatively simple, which we argue is essential for encouraging scientists and engineers to adopt PERKS in their iterative solvers implemented for GPUs, and other architectures as well. That being said, a challenging aspect that we address in this paper is a detailed analysis of how and why PERKS is effective. The analysis requires an understanding of the effect of concurrency on performance. More particularly, to gain a deep understanding of why PERKS are effective and the limitations of architectural features, we study the effect of pressure on resources (particularly registers and shared memory). On top of that, we examine the effect of reducing the device occupancy while maintaining high enough concurrency to saturate the device.

It is important to note that PERKS are orthogonal to temporal blocking optimizations. Temporal blocking relies on combining multiple consecutive iterations of the time loop to reduce the memory transactions between them. The dependency along the time dimension is resolved by either: a) redundantly loading and computing cells from adjacent blocks, which limits the effectiveness of temporal blocking to low degrees of temporal blocking [34, 51, 66], or b) using tiling methods of complex geometry (e.g. trapezoidal and hexagonal tiling) along the time dimension and restrict the parallelism due to the dependency between neighboring blocks [11, 30, 56]. In contrast, the execution model of PERKS does not necessitate

the resolution of the dependency along the time dimension since PERKS include an explicit barrier after each time step. This means the PERKS model can be generalized to any iterative solver and can be used on top of any version of the solver. In other words, iterative kernels written as PERKS do not compete with optimized versions of those iterative kernels. For instance, a stencil PERKS does not compete with kernels applying aggressive locality optimizations [14, 16] and out-of-core optimizations [26, 70, 78]; the performance gain from PERKS is added to the performance gain from whatever stencil optimizations are used in the kernel. **As a matter of fact, the more optimized the kernel before it is ported to the PERKS execution model, the higher the speedup that would be gained by PERKS. That is since optimizations to the kernel proportionally increases the overhead of data storing and loading in between iterations, which PERKS aim to reduce.**

The contributions in this paper are as follows:

- We propose PERKS, an execution model that explicitly exploits the large volume of unused on-chip resources to reduce memory traffic in iterative memory-bound applications.
- We introduce the design principles of PERKS and also provide analyses of the potential benefit of PERKS, and how to effectively port iterative solvers to PERKS.
- We implement a wide range of iterative 2D/3D stencil benchmarks and a conjugate gradient solver as PERKS in CUDA. It is important to note that iterative stencils and Krylov subspace solvers are the backbones of numerous scientific and engineering codes. We include an elaborate discussion on implementation details and performance-limiting factors, such as problem sizes, concurrency, and resource contention.
- Our PERKS-based implementation achieves geometric means speedups of 2.12x for 2D stencils and 1.24x for 3D stencils in comparison to several highly optimized state-of-the-art 2D/3D stencil libraries on A100 and V100. PERKS-based conjugate gradient achieves a geometric mean speedup of 2.48x in comparison to the highly GPU-optimized production library Ginkgo [6] for SpMV datasets in SuiteSparse. For smaller datasets, the speedup goes up to 2.86x.

The rest of this paper is organized as follows: Section 2, presents the background and motivation. Section 3 presents the overview of the execution model PERKS. Section 4 shows the implementation of PERKS. In Section 5, we analyze the performance consideration of PERKS. Section 6 displays the evaluated result. In Section 7, we elaborate on the related work. Finally, Section 8 concludes the paper.

2 BACKGROUND AND MOTIVATION

2.1 CUDA Programming Model

CUDA's programming model [61] includes: *threads*, the basic execution unit (32 threads are executed together as a *warp*); *Thread block (TB)*, usually composed of hundreds of threads; *grid*, usually composed of tens of thread blocks.

2.1.1 GPU Memory Hierarchy: On-chip memory in a streaming multiprocessor (SM) includes: shared memory (scratchpad memory), L1 cache, and register file (RF) and. Off-chip memory includes global

¹In this paper, we use PERKS, interchangeably, to refer to our proposed model and, as an abbreviation of PERsistent KernelS

Table 1: Relevant Features of the latest Nvidia GPUs[59]

Feature	P100	V100	A100
Shared memory ²	3.5 MB	7.5 MB	17.29 MB
Register files	14 MB	20MB	27MB
L2 cache	4 MB	6 MB	40 MB
Memory bandwidth	720 GB/s	900 GB/s	1555 GB/s

memory (device memory) and L2 cache. Data in global memory can reside for the entirety of the program, while data in on-chip memory has the lifetime of a kernel. The shared memory is shared among all threads inside a thread block. We summarize the relevant memory features in Table 1.

2.1.2 GPU Device-wide Synchronization: Synchronization in GPUs was limited to groups of threads: thread blocks in CUDA (or a work group in OpenCL). Starting from CUDA 9.0, Nvidia introduced cooperative group APIs [60] that include an API for device-wide synchronization. Before introducing grid-level synchronization, the typical way to introduce device-wide synchronization was to launch sequences of kernels in a single CUDA stream. Zhang et al. [82] conducted a comprehensive study to compare the performance of both methods. The result shows that the latency difference between explicit device-wide synchronization versus implicit synchronization (via repetitive launching of kernels) is negligible in most kernels.

2.2 Iterative Algorithms

In iterative algorithms, the output of time step k is the input of time step $k + 1$. Iterative methods can be expressed as:

$$x^{k+1} = F(x^k) \quad (1)$$

When the domain is mapped out to processing elements, there are two points to consider:

- Spatial dependency necessitates synchronization between time steps, or else advancing the solution in the following time step might use data that has not yet been updated in the previous time step.
- In time step $k + 1$, each thread or thread block needs input from the output of itself in time step k (i.e. temporal dependency). This gives the opportunity for caching data between steps to reduce device memory traffic.

In the following sections, we briefly introduce iterative stencils and Krylov subspace methods. Throughout the paper, we use them as motivation examples, and we use them to report the effectiveness of our proposed methods, given their importance in HPC scientific and engineering codes.

2.2.1 Iterative Stencils. Iterative stencils are widely used in HPC. According to Bastian et al. [33], stencil applications represent 49% of workloads in a wide range of HPC centers. Take 2D Jacobian 5-point stencil (2d5pt) as an example:

$$x(i, j)^{k+1} = N * x(i, j + 1)^k + S * x(i, j - 1)^k + C * x(i, j)^k + W * x(i - 1, j)^k + E * x(i + 1, j)^k \quad (2)$$

²Shared memory is a configurable portion of L1 cache that can be used as a user-managed scratchpad memory

Computation of each point at time step $k + 1$ requires the values of the point itself and its four neighboring points at time step k .

Two blocking methods are widely used to optimize iterative stencils for data locality: *Spatial Blocking* [37, 80] and *Temporal Blocking* [49, 67, 68].

In spatial blocking on GPUs, we split the whole domain into sub-domains, where each thread block can load its sub-domain to the shared memory to improve the data reuse. In the meantime, we require redundant data accesses at the boundary of the thread block to data designated for adjacent thread blocks.

In iterative stencils, each time step depends on the result of the previous time step. One could advance the solution by combining several time steps. The temporal dependency, in this case, is resolved by using a number of halo layers that match the number of combined steps. The amount of data that can be computed depends on the stencil radius (rad) and the number of time steps that are combined (b_t). In overlapped temporal tiling [35, 44, 52], this region can be represented as $2 \times b_t \times rad$ (halo region). Methods based on this kind of blocking are called *overlapped temporal blocking* schemes. Overlapped temporal blocking introduced the overhead of redundant computation that wavefront [10, 47, 79] is aimed to alleviate.

2.2.2 Krylov Subspace Methods. Krylov methods are widely used for large sparse (and dense) linear systems of equations arising in solvers of Partial Differential Equations (PDEs) [7, 29, 64]. Krylov subspace methods can be described as:

$$\kappa_r(A, b) = span\{b, Ab, A^2b, \dots, A^{r-1}b\} \quad (3)$$

Assuming that A is an invertible matrix, it is possible to compute $x = A^{-1}b$ (or solve $Ax = b$) by searching the Krylov subspace without directly computing A^{-1} . Searching the Krylov subspace is a sequence of matrix-vector multiplications, where at each step the approximation of the solution vector x is updated proportionally to the residual error (vector r) from the previous time step.

Conjugate gradient is a main solver in the family of Krylov subspace methods. It is mainly used to solve systems of linear equations for symmetric and positive-definite matrices.

2.3 Motivational Example

We use a motivational example of a double precision 2D 5-point Jacobian (2d5pt) stencil to motivate implementing iterative solvers as PERKS (Readers can refer to Equation 2 and Listing 1 for details of the 2d5pt stencil). **1** Why PERKS: Optimizations for iterative methods focus on a single step to speed up iterative solvers. Single-step optimizations move the performance of the kernel closer to the highest possible attainable performance on the roofline model, yet will not influence the operational intensity. As Figure 1 shows, the optimizations used for the 2D 5-point stencil move the performance vertically at the same operational intensity value of the kernel. Temporal blocking schemes can horizontally move the operational intensity, to the right side of the roofline. Yet resolving the neighborhood dependencies introduces redundancy [35, 44, 66] or hard-to-parallel complex geometrical tile shapes [12, 30, 56], and can cause register pressure [49]. In PERKS, we reduce the unnecessary data access between time steps. The target data traffic to reduce is in-between time steps (i.e., outside the solver) and hence




Figure 1: The roofline model of a 2D 5-point Jacobian stencil kernel (dp), running $T = 20$ time steps, with a domain size of 3072^2 on A100 GPU. Per-time step optimizations only improve the iterative stencil kernel to get closer to the attainable performance. Reducing memory traffic between time steps can increase the attainable performance by increasing the aggregate operation intensity (OI) over all time steps. We also plot different operational intensities for a version of PERKS that reduces the data traffic in-between 20 time steps to 50%, 25%, and 0%.




Figure 2: Performance of a double precision 2D 5-point Jacobian stencil kernel (3072^2) for different Thread Blocks per Streaming Multiprocessor (TB/SM) on A100. Filled regions indicate unused resources. Using one TB/SM and using all unused resources for caching can theoretically provide 1.96x speedup in this situation.

is not subject to the neighborhood dependency issue in temporal blocking schemes. Figure 1 demonstrates how this idea works for a real stencil benchmark running on an A100 GPU for 20 time steps. By caching more of the domain in-between time steps, the operational intensity moves more to the right side of the roofline to be compute-bound. This also demonstrates how PERKS is orthogonal to the per-time step optimizations; PERKS would improve the performance (by moving horizontally on the roofline) regardless of how optimized the baseline algorithm is at its operational intensity. ② The prospect of PERKS: Latency across all operations/instructions in newer generation GPUs has been significantly dropping [39]. As a result, often fewer numbers of warps are enough for CUDA




Figure 3: Runtime (8 time steps) of double precision 2D 5-point Jacobian stencil (3072^2) with different state-of-the-art optimization A100 GPU. PERKS aims to reduce the traffic to/from device memory in-between time steps. SHM [48] uses shared memory. PPCG is a code auto generation tool [74]. NVCC-OPT relies on auto optimization provided by the latest NVCC compiler version. SSAM [16] uses register to improve locality [14, 15, 17]. Stencilgen [68] applies temporal blocking. In the figure we also plot the total runtime of each implementation running 8 time steps, assuming we reduce 50% of the inter time step memory traffic. The results show that the more optimized (i.e. fewer proportion is spent in compute), the more performance improvement we expect from caching.

runtime to hide the latency effectively and hence maintain high performance at low occupancy [76]. In Figure 2, we vary the number of thread blocks per streaming multiprocessor (TB/SM) and plot its performance (left Y-axis). For each TB/SM configuration, we plot on the right Y-axis the unused resources (shared memory and registers). As the figure shows, **even when TB/SM = 5, more than 10.46 MB of shared memory and register files are not in use**. When TB/SM decreases, the performance is slightly fluctuated ($72.74 \rightarrow 68.12$ GCells/s³) while the freed shared memory and registers gradually increase. By reducing TB/SM to its minimum while maintaining enough concurrency to sustain the performance level, the projection from performance gain when caching a subset of the results in unused resources can improve performance by more than 1.96x. Figure 3 uses an alternative perspective, assuming that memory accessing and compute within a solver can not be overlapped. We profiled different 2d5pt state-of-the-art implementations. As Figure 3 shows, the compute part decreases as the more optimized the stencil implementation is. The prospect of PERKS is to reduce this data movement time that dominates the runtime in highly optimized stencil implementations. Note that while temporal-blocking schemes do also reduce the data movement to some extent, they cannot be generalized to all iterative solvers. Additionally, resolving temporal and spatial dependency adds compute overhead and can also lead to increased register pressure that limits the degree of temporal blocking on GPUs [49].

³GCells/s denotes giga-cells updated per second.




Figure 4: Changing a traditional iterative CUDA kernel (time loop on the host) to PERKS: 1) move the time stepping loop from the host to the device, and use CUDA’s grid synchronization API as a device-wide barrier at each time step. This way, we expose the temporal data locality across time steps to thread blocks. We then use the register files and shared memory to reduce traffic to the global memory by caching the domain (or a subset from it) in-between time steps.

3 PERKS: PERSISTENT KERNELS TO IMPROVE LOCALITY

PERsistent Kernels (PERKS) is a generic execution model for running iterative solvers on GPUs to improve data locality by taking advantage of the large capacity of register files and shared memory. As Figure 4 illustrates, in PERKS, we move the time stepping loop from the host to the device, and use CUDA’s grid synchronization API as a device-wide barrier at each time step. This way, we expose the temporal data locality across time steps to thread blocks. We then use the register files and shared memory to reduce traffic to the global memory by caching the domain (or a subset from it) in-between time steps.

3.1 Assumptions and Limitations

The techniques discussed in this paper are based on the following assumptions about the applications.

Target Applications: In this paper, we target iterative kernels that are bounded by memory bandwidth. Although execution in a PERKS fashion makes no assumptions on the underlying implementation, optimal PERKS performance can sometimes require minor adaptations to the kernel. The changes, for instance, can be as simple as changing the thread block and grid sizes to reduce over-subscription, or more elaborate as favoring a specific SpMV method from the space of SpMV methods in the case of the conjugate gradient (more details in Section 4.3). Finally, despite not

reporting results for compute-bound iterative kernels, it is important to note that compute-bound iterative kernels could potentially also benefit from becoming PERKS, if the kernel generates memory traffic in-between iterations that CUDA runtime can not effectively overlap with computation.

PERKS in Distributed Computing: PERKS in this paper is demonstrated on a single GPU. In distributed applications that require halo regions (e.g., stencils), PERKS can potentially be used on top of communication/computation overlapping schemes [27, 63]. In overlapping schemes, the boundary points that are computed in a separate kernel would not be cached, while the kernel of the interior points would run as PERKS to cache the data of the interior points. PERKS could also be used with communication-avoiding algorithms (e.g., communication-avoiding Krylov methods [36])

Practicality of PERKS: A wide range of iterative solvers (particularly iterative stencils) can be written as PERKS. However, it should be mentioned that there are applications at which the time stepping loop (on the host) is comprised of different GPU kernel. For instance, in production libraries, conjugate gradient (and Krylov solvers in general) are typically implemented as different kernels corresponding to different steps in the algorithm. In such case the kernels have to be fused [32, 77] first before transforming them to PERKS.

Use of Registers: PERKS uses registers and shared memory for caching data in-between time steps. It should be noted that there are no guarantees that the compiler releases all the registers after the compute portion in each iteration is finished (with Nvidia’s nvcc compiler we did not observe such inefficiency). If such register reuse inefficiency exists, imperfect register reuse by the compiler could result in fewer registers being available for caching and leaves only shared memory to be used for caching. PERKS would not be effective if the target kernel consumes all on-chip resources (both register file and shared memory) even in its minimal occupancy.

Iterative Solvers as PERKS: While this paper’s focus is to demonstrate PERKS model for iterative stencils and Krylov subspace methods (conjugate gradient), the discussion in this section (and paper in general) is applicable to a high degree for other types of iterative solvers. That is since PERKS is not much concerned with the implementation of the solver and only loads/stores the domain (or a subset of it) before/after the solver part in the kernel, under resource constraints. Iterative solvers that use the same flow expressed in Figure 4 can, in principle, be ported to PERKS (with relative ease). Generally speaking, the porting process is as follows: move the time step outside the kernel to be inside the kernel, add grid synchronization to ensure dependency, and store/load a portion of the input or output to cache: either shared memory and/or register (using register arrays). More details on porting kernels to PERKS are in Section 4.1.

4 PORTING SOLVERS TO PERKS

Transforming the existing iterative solvers to PERKS is straightforward. This section first explains briefly how end-users can transform or port their iterative solvers to PERKS. Next, we elaborate on how we implemented memory-bound iterative methods (namely 2D/3D stencils and a conjugate gradient solver) as PERKS.

4.1 Transforming Kernels to PERKS: the End-user Perspective

4.1.1 Identifying the minimal concurrency of the kernel. The end-user can rely on CUDA APIs⁴ [60] to get the max concurrently running parameters. For even better performance, the end-user only needs to reduce the device occupancy to its minimum (while maintaining performance) via manual tuning of the kernel launch parameters or using auto-tuning tools [2, 71, 73].

4.1.2 Porting a Kernel to become PERKS. As Listing 1 shows, PERKS does not modify the computation; the manually written code to move the time loop inside the kernel and load/store to cache is straightforward. Alternatively, though outside this paper’s scope, we point out the possibility of simplifying the process of converting a kernel to PERKS by using source-to-source translation, C++ templates, or domain specific languages.

4.1.3 What to Cache. The end-user can use a profiler, offline, to decide on what data arrays to cache by identifying the arrays that generate the most traffic to/from global memory. In many iterative solvers, profiling is not even needed since the algorithm clearly implies the main data array(s) causing the highest traffic (e.g., the matrix A in conjugate gradient and the discretized domain in stencil applications).

4.1.4 Where to Cache. The end-user would simply use the unused shared memory for caching. For additional performance benefits, advanced users can choose to also cache in registers by manually identifying the adequate number of registers that can be used for caching, without causing register spilling (we provide a Python script to automate this process), or by following the trace of existing on-chip resources management research [45, 75]. We anticipate the possibility of automating this step by source-to-source translation or domain-specific languages so that this step of using on-chip resources could be as easy as adding a persisting range in the domain, similar, in principle, to the method of using L2 cache residency control in A100 [59], except that l2 cache residency control does not guarantee the data is definitely persistent [20].

4.2 Transforming Stencil Kernels to PERKS

4.2.1 Stencil Kernel. We use SHM [48] implementation as baseline. In SHM, 3D stencil implementation uses the standard shared memory implementation where 2D planes (1D planes in 2D stencils) are loaded one after the other in shared memory. Each thread computes the cells in a vertical direction [55, 69]. In our PERKS implementation, before the compute starts, planes that already have the data cached from the previous time step do not load from global memory.

4.2.2 Porting the Stencil Kernel. We do not interfere with compute; only after the computation is finished do we store the results in the registers/shared memory. As Listing 1 shows, after adjusting to handle the input and output of the computation part of the kernel. We exchange halo region (inter thread block dependency data) between time steps. To ensure coalesced memory accesses in the halo region, we transpose the vertical edges of the halo region. Also, we reuse the on-chip resources for caching as soon as the

⁴cudaOccupancyMaxActiveBlocksPerMultiprocessor.

Listing 1: Iterative 2D 5-pt stencil implemented in PERKS.

```

1  __global__ void 2d5pt_PERKS(ptr_in, ptr_out)
2  {
3  ...
4  for(k=0; k<timestep; k++)
5  {
6  ...
7  //Use branch to control the source
8  load(tid, ptr_in,
9  [ptr_in|sm_cache|reg_cache]->sm_in);
10  2d5pt_Compute(sm_in, reg_out);
11  //Use branch to control the destination
12  store(tid, ptr_out,
13  reg_out->[ptr_out|sm_cache|reg_cache]);
14  ...
15  grid.sync();
16  ...
17  }
18  ...
19  }
20
21  __device__ void 2d5pt_Compute(sm_in, reg_out){
22  x = threadIdx.x;
23  t[IPT+2]; //IPT: items per thread
24  for(y=0; y< IPT+2; y++)
25  t[y]=sm_in[x, y+ind_y-1];
26  for(y=0; y< IPT; y++){
27  reg_out[y]=sm_in[x+ind_x-1,y+1+ind_y]*WEST
28  +sm_in[x+ind_x+1,y+1+ind_y]*EAST
29  +t[y-1+1]*SOUTH
30  +t[y+1]*CENTER
31  +t[y+1+1]*NORTH;
32  }
33  }

```

data is consumed. Finally, since the original kernel uses shared memory [55, 69] and registers [16] to optimize stencils, we use the version of the output residing in shared memory or registers at the end of each time step as an already cached output. In this way, we avoid an unnecessary copy to shared memory and registers that we would use for caching.

4.3 Transforming the Conjugate Gradient Solver to PERKS

4.3.1 Conjugate Gradient Kernel. For simplicity and accessibility, we use the Conjugate Gradient (CG) solver implementation that is part of the CUDA SDK samples (*conjugateGradientMulti-BlockCG* [58]). Since the implementation of SpMV in the CG sample is relatively naive, we use the highly optimized merge-based SpMV [54] that is part of the C++ CUB [57] library in the CUDA Toolkit [22], as it fits naturally with the caching scheme in PERKS. We do not discuss the details of merge-based SpMV due to the space limit. The reader can refer to details in [54].

4.3.2 Porting the Conjugate Gradient Kernel. We do not change the implementation or algorithm of the merge-based SpMV since PERKS does not necessitate changes in the underlying algorithm. For merge-based SpMV, we cache the matrix A since it is the largest data array in the solver. To further improve performance, we also cache the residual vector r and the intermediate results. The merge-based SpMV [54] in CUB [57] is composed of two steps: *search* and *compute*. The search step is done twice. The search step first finds the workload for each thread block, and then finds the workload for each thread inside a thread block. The search result for the thread block workloads in global memory is saved since the matrix is static throughout the entire iteration. The second search (thread-level) is conducted in shared memory. Those two steps repeatedly

Listing 2: Iterative Sparse matrix–vector multiplication with merge-based SpMV [54] implemented in PERKS.

```

1  __global__ void Iterative_SpMV_PERKS (...)
2  {
3
4  while (...)
5  {
6      //merge-based SpMV [54] contains:
7      //search: determine the workload range
8      //SpMV_Compute: performs SpMV within range
9
10     //determine current thread block workload
11     get([ cached_tb_range | search(gm_matrix) ]-> tb_range);
12     //load tb_range of matrix to shared memory
13     load(tb_range, [ sm_matrix | gm_matrix ]-> tb_sm_matrix);
14
15     //determine current thread workload
16     get([ cached_t_range | search(tb_matrix) ]-> t_range);
17     SpMV_Compute(tb_sm_matrix, t_range);
18
19     grid.sync();
20
21 }
22
23 }

```

generate intermediate data that we cache, in addition to the matrix A . Listing 2 shows code sample of PERKS based Iterative SpMV, which can be extended to a conjugate gradient solver.

Merge-based SpMV originally uses small thread blocks, i.e., 64 threads per TB. This introduces a high volume of concurrently running thread blocks per streaming multiprocessor. To reduce the device occupancy while maintaining performance, we increased the TB size to 128 and slightly changed the memory access order to accommodate the larger TB size.

4.4 PERKS and CUDA Considerations

4.4.1 Restrictions of Synchronization APIs. PERKS relies on cooperative groups APIs [60] (supported since CUDA 9.0). Currently, the APIs do not allow over-subscription, i.e., one needs to explicitly assign workload to blocks and threads to expose enough parallelism to the device. However, it is worth mentioning that this API does not limit the flexibility, as different kernels can still run concurrently in a single GPU, as long as they as a whole doesn't exceed the hardware limitation.

4.4.2 New Features in Nvidia Ampere. The Nvidia Ampere generation of GPUs introduced two new features that have the potential to improve the performance of PERKS. Namely, asynchronous copy for shared memory and L2 cache residency control [59]. When testing asynchronous copy to cache in PERKS, we did not observe noticeable performance difference. For L2 cache residency control, we experimented with setting the input and halo region to be persistent in stencils. We observed a 8% slowdown and no change in performance, respectively. Accordingly, we do not include those new features in our PERKS implementations.

4.4.3 Register pressure in PERKS. One concern with PERKS is that kernels might run into register pressure if the compiler is not optimally reusing registers for different time steps, potentially affecting concurrency and penalizing performance. To illustrate this issue, take a high register-pressure 2D 25-point double precision Jacobi stencil as an example. The shared memory optimized baseline version (SHM) uses 78 registers per thread, yet the PERKS version

uses 112 registers⁵. Similar behavior is also observed in other stencil benchmarks. Reducing the occupancy while maintaining the concurrency –as mentioned in the previous section– reduces the impact of this compiler's inefficiency in register reuse in all the benchmarks we report in the results section. In the above example, at worst, 48 registers among the maximum available 178 registers per thread could not be used for caching data; it neither harms concurrency nor triggers register spilling.

5 PERFORMANCE ANALYSIS

In this section, we propose a performance model that serves the following purposes. First, we propose a projection of achievable performance that we compare with measured results to detect abnormal behavior or implementation shortcomings. We relied on this projection in the analysis of our PERKS implementation quality (Section 5.2). Second, we identify the bounds on reducing concurrency before performance regression and use concurrency to explain potential optimizations for further performance improvement (Section 5.4). It is worth mentioning that the concurrency analysis is not a requirement for porting kernels to PERKS; we use the analysis to understand the feasibility of PERKS in practice, and address its implication on performance.

5.1 Overview

This performance model relies on three performance attributes: a) measured performance \mathbb{M} of our PERKS implementation, b) the projected peak performance \mathbb{P} achievable on a given GPU, and c) the efficiency function $\mathfrak{E}()$ describing the efficiency of the given kernel running on the device. More specifically, $\mathfrak{E}()$ is a function of the concurrency exposed by the software \mathbb{C}_{sw} and the concurrency required by the hardware \mathbb{C}_{hw} . The relation of measured performance to projected peak performance becomes:

$$\mathbb{M} = \mathbb{P} \times \mathfrak{E}(\mathbb{C}_{sw}, \mathbb{C}_{hw}) \quad (4)$$

We discuss projected peak performance in the following section. A detailed discussion of the efficiency and concurrency functions is in Section 5.3.

5.2 Projecting Peak Achievable Performance

We rely on the figure of merit as the performance metric in this analysis. In stencils, we use the giga-cells updated per second (GCell-s/s) [16, 49]. Given the memory-bound nature of the conjugate gradient solver, we directly use sustained memory bandwidth as a metric, following other works on conjugate gradient [6]. Due to space limitations, this section mainly focuses on stencils to explain the performance analysis. Without loss of generality, the analysis is applicable to other cases (ex: conjugate gradient) by adjusting the performance metric and code concurrency accordingly.

We use a simple performance model inspired by the roofline model [42, 62]. The model's utility is to project the upper bound on performance based on the reduction of global memory traffic. This model, in turn, helps us in this paper to identify performance gaps in our PERKS implementation and later inspect the reasons for those gaps.

⁵We gathered the number of registers used by finding the maximum number of registers available as cache before spilling with "`__launch_bounds__`" instruction. Register spilled can be indicated by "`-Xptxas -v -dlcm=cg`" flag.

In a kernel implemented as PERKS, the bottleneck could either be the global memory bandwidth or the shared memory bandwidth (if the PERKS caching scheme moves the bottleneck to become the shared memory bandwidth). We don't assume the registers to be a bottleneck since we assume that as long as we ensure that no register spilling occurs, we avoid register pressure.

We assume a total domain of size \mathbb{D} bytes, the cached portion to be \mathbb{D}_{cache} bytes, and the uncached portion to be $\mathbb{D}_{uncache} = \mathbb{D} - \mathbb{D}_{cache}$ bytes. The cached portion of the domain data would be divided between registers and shared memory (since we cache in both registers and shared memory): $\mathbb{D}_{cache} = \mathbb{D}_{cache}^{sm} + \mathbb{D}_{cache}^{reg}$. For N time steps, assuming the number of bytes stored to global memory in each time step is \mathbb{S}_{gm} and the number of bytes loaded is \mathbb{L}_{gm} , the total global memory bytes accessed \mathbb{A}_{gm} becomes:

$$\mathbb{A}_{gm}(\mathbb{D}) = N \cdot (\mathbb{L}_{gm} + \mathbb{S}_{gm}) = 2 \cdot N \cdot \mathbb{D}_{uncache} + 2 \cdot \mathbb{D}_{cache} \quad (5)$$

When the kernel is bounded by global memory bandwidth, i.e., the volume of cached data does not move the bottleneck from global memory to shared memory, for the global memory bandwidth of \mathbb{B}_{gm} and data type size of $\mathbb{S}(type)$, the time $\mathbb{T}_{gm}(\mathbb{D})$ for accessing the global memory becomes:

$$\mathbb{T}_{gm}(\mathbb{D}) = \mathbb{A}_{gm}(\mathbb{D}) \cdot \mathbb{S}(type) / \mathbb{B}_{gm} \quad (6)$$

In the case when the kernel is bounded by shared memory bandwidth, i.e., the volume of data cached in shared memory moves the bottleneck to be the shared memory bandwidth, the total shared memory (in bytes) accessed \mathbb{A}_{sm} becomes:

$$\mathbb{A}_{sm}(\mathbb{D}_{cache}^{sm}) = N \cdot (\mathbb{L}_{sm} + \mathbb{S}_{sm}) = 2 \cdot (N - 1) \cdot \mathbb{D}_{cache}^{sm} \quad (7)$$

Assuming $\mathbb{A}_{sm}(\text{KERNEL})$ to be the shared memory originally used by the kernel, e.g., shared memory used in the baseline implementation of a stencil kernel to improve the locality, and \mathbb{B}_{sm} to be the shared memory bandwidth, the time $\mathbb{T}_{sm}(\mathbb{D})$ for accessing the shared memory becomes:

$$\mathbb{T}_{sm}(\mathbb{D}_{cache}^{sm}) = \{(\mathbb{A}_{sm}(\mathbb{D}_{cache}^{sm}) + \mathbb{A}_{sm}(\text{KERNEL})) \cdot \mathbb{S}(type)\} / \mathbb{B}_{sm} \quad (8)$$

The projected best-case total time required for the PERKS kernel may be written as:

$$\mathbb{T}_{PERKS} = \max(\mathbb{T}_{gm}(\mathbb{D}), \mathbb{T}_{sm}(\mathbb{D}_{cache}^{sm})) \quad (9)$$

Accordingly, the projected peak performance (\mathbb{P} in Equation 4) for the N time steps can be expressed as:

$$\mathbb{P} = \mathbb{D} \cdot N / \mathbb{T}_{PERKS} \quad (10)$$

We give an example of computing $N = 1000$ time-steps of a single precision 2D 5-point Jacobi stencil on A100. We use the domain size $\mathbb{D} = 3072^2$; the total cache-able region is $\mathbb{D}_{cache} = 3072 \cdot 2448$ leading to $\mathbb{T}_{gm}(\mathbb{D}) = 9900.70$ us. The total number of bytes for the halo accesses is $\mathbb{A}(\mathfrak{H}(\mathbb{D}_{cache})) = 1000 \cdot 2 \cdot 216 \cdot (136 \cdot 2 + 256 \cdot 2)$. Thus $\mathbb{T}_{gm}^{\mathfrak{H}}(\mathbb{D}_{cache}) = 871.22$ us. So $\mathbb{P}_{PERKS} = 3072^2 \cdot 1000 / \mathbb{T}_{PERKS} = 876.09$ GCells/s.

5.3 Concurrency and Micro-benchmarks

Reducing device occupancy increases the availability of resources to be used for caching in PERKS (as illustrated earlier in Figure 2). On the contrary, reducing occupancy can lead to lower device utilization. To effectively implement PERKS, one has to reduce the

Table 2: Concurrency analysis of global memory accesses of a single precision 2D-5point Jacobi stencil kernel running on A100 (1000 time-steps on 3072² domain)

$\frac{TB}{SM}$	Used Reg. /SM	Unused Reg. /SM	GM Load op/SM	GM Store op/SM	Measured GCells/s
1	32KB	224KB	2580	2048	94.75
2	64KB	192KB	5160	4096	133.24
8	256KB	0KB	20640	16384	138.29

occupancy as much as possible without scarifying performance. Inspired by the findings of Volkov [76], we assume that the efficiency function \mathfrak{E} reaches its peak point when the code provides enough concurrency to saturate the device (irrespective of the occupancy):

$$\mathfrak{E}(\mathbb{C}_{sw}, \mathbb{C}_{hw}) = 100\%, \text{ if } \forall \mathbb{C}_{sw} \geq \mathbb{C}_{hw} \quad (11)$$

Where $\mathbb{C}_{sw}(\text{OP})$ is the minimum number of concurrently executable instructions of the operation OP exposed by the launched kernel, and $\mathbb{C}_{hw}(\text{OP})$ is the maximum numbers of instructions of the operation OP that the device is capable of handling concurrently. Because this paper mainly focuses on memory bound applications, the OP referred to in this paper are limited to data access operations, i.e. global memory load/store $\mathbb{C}(\text{GM})$, shared memory load/store $\mathbb{C}(\text{SM})$, and L2 cache load/store $\mathbb{C}(\text{L2})$.

5.3.1 Measuring \mathbb{C}_{sw}^{SM} . $\mathbb{C}_{sw}^{SM}(\text{OP})$, the kernel concurrency at the Streaming Multi-processor (SM) level, can be computed based on the concurrency exposed by the threads of a thread block $\mathbb{C}_{sw}^{TB}(\text{OP})$ and number of concurrently running thread blocks per SM TB/SM : $\mathbb{C}_{sw}^{SM}(\text{OP}) = \mathbb{C}_{sw}^{TB}(\text{OP}) \cdot \text{TB}/\text{SM}$.

5.3.2 Measuring \mathbb{C}_{hw} . According to Little's Law [46], the hardware concurrency \mathbb{C}_{hw} can be determined by the throughput THR and latency \mathbb{L} [76]:

$$\mathbb{C}_{hw}(\text{OP}) = \text{THR}(\text{OP}) \cdot \mathbb{L}(\text{OP}) \quad (12)$$

The throughput THR for data access operations are available in CUDA documentation [21, 61]. We measure the latency \mathbb{L} with commonly used microbenchmarks [50, 81, 82].

5.4 Concurrency Analysis

In this section, we briefly describe how we analyze the concurrency to reduce the occupancy of the original kernel in order to release resources for caching while sustaining performance. We conduct a static analysis to extract the data movement operations in the kernel. Note that we account for any barriers in the original kernels that could impact the concurrency of operations, i.e., we do not combine operators/instructions from before and after the barrier when we count the operators. Finally, we apply a simple model (Equation 11) to identify the least occupancy we could drop to before the concurrency starts to drop. The results summarized in Table 2 show that for a 2D 5-point Jacobi stencil kernel kernel, we could reduce the original occupancy to 1/4th while maintaining performance.

To understand the gap between the performances at 1 vs. 8 TB/SM (94.75/138.29 = 68.52%), we inspect the efficiency function $\mathfrak{E}()$. The number of concurrent global memory accesses and shared memory accesses in the 2D 5-point Jacobi stencil kernel are enough to saturate A100 when $\text{TB}/\text{SM}=1$. Accordingly, we get $\mathfrak{E}(\mathbb{C}_{sw=j2d5pt, \text{TB}/\text{SM} \geq 1, \mathbb{C}_{hw=A100}) = 1$, which would indicate that

Table 3: Stencil benchmarks and domain sizes we use. A detailed description of the stencil benchmarks can be found in [67, 83]. For fairness, the domain sizes we use are the minimum domain sizes that would saturate the device for different stencil benchmarks. The minimum domain sizes are identified by running each benchmark at different domain sizes and using the domain size after which the performance (in GCells/s) cease to improve

Type (Stencil Order, FLOPs/Cell)	2d stencils		Type (Stencil Order, FLOPs/Cell)	3d stencils	
	A100	V100		A100	V100
2d5pt (1,10)	2304 × 2304	2048 × 1280	3d7pt (1,14)	256 × 288 × 256	128 × 128 × 128
2ds9pt (2,18)	2304 × 2304	2048 × 1280	3d13pt (2,26)	256 × 288 × 256	256 × 320 × 256
2d13pt (3,26)	4608 × 3072	2048 × 2048	3d17pt (1,34)	256 × 288 × 256	160 × 160 × 256
2d17pt (4,34)	3072 × 2304	4096 × 2560	3d27pt (1,54)	256 × 288 × 256	160 × 160 × 256
2d9pt (1,18)	2304 × 2304	2048 × 1280	poisson (1,38)	256 × 288 × 256	160 × 160 × 256
2d25pt (2,50)	4608 × 3072	2048 × 1280	—	—	—

Table 4: Datasets for the Conjugate Gradient (CG) solver (from SuiteSparse [25])

Code	Name [25]	Rows	NNZ	Code	Name [25]	Rows	NNZ	Code	Name [25]	Rows	NNZ
D1	Trefethen_2000	2,000	41,906	D8	finan512	74,752	596,992	D15	crankseg_1	52,804	10,614,210
D2	msc01440	1,440	46,270	D9	cbuckle	13,681	676,515	D16	bmwcra_1	148,770	10,644,002
D3	fv1	9,604	85,264	D10	G2_circuit	150,102	726,674	D17	hood	220,542	10,768,436
D4	msc04515	4,515	97,707	D11	thermomech_dM	204,316	1,423,116	D18	BenElechi1	245,874	13,150,496
D5	Muu	7,102	170,134	D12	ecology2	999,999	4,995,991	D19	crankseg_2	63,838	14,148,858
D6	crystm02	13,965	322,905	D13	tmt_sym	726,713	5,080,961	D20	af_1_k101	503,625	17,550,675
D7	shallow_water2	81,920	327,680	D14	consph	83,334	6,010,480	—	—	—	—

the observed gap in performance is not due to a drop in concurrency we did not model. While this confirms the effectiveness of the concurrency analysis (i.e., since the concurrency analysis resonates with the empirical measurements in Table 2), it does not uncover the source of the performance gap. Investigative profiling revealed that the concurrency for accesses in L2 cache, not global memory, is impacted by reducing occupancy on A100 in specific to the level that affects performance notably. More particularly, access to global memory for the halo region garners a high L2 cache hit rate. This effectively means that higher concurrency is necessary to saturate the L2 cache when hit rates are high. To confirm, we manually doubled the concurrency $C_{s,w}^{TB}$: the performance increased to 123.94 GCells/s with TB/SM=1 (from 68.52% up to 89.6%).

6 EVALUATION

6.1 Hardware and Software Setup

The experimental results presented here are evaluated on the two latest generations of Nvidia GPUs: Volta V100 and Ampere A100 with CUDA 11.5 and driver version 495.29.05.

We run each evaluation ten times for all iterative stencils and conjugate gradient experiments. All experimental results reported are done in double precision.

6.2 Benchmarks and Datasets

6.2.1 Stencil Benchmarks. To evaluate the performance of PERKS stencils, we perform a wide set of experiments on various 2D/3D stencil benchmarks (listed in Table 3).

We compare PERKS (w/ SHM [48] as base) with a wide range of state-of-the-art stencil implementations/libraries: PPCG [74], Bricks [83], SSAM [16], STENCILGEN [68] and SHM [48]. The implementations/libraries represent different classes of stencil optimization approaches: code auto-generation (PPCG), vector-level data reuse (Bricks), shared memory optimization (SHM), accumulated summations optimization (SSAM), and temporal blocking

(STENCILGEN). For a fair comparison, when comparing with STENCILGEN [68], SSAM [16], and Bricks [83], we use the default setting (including the default domain size) in their papers, except that we adjusted the build system to add the latest GPU (A100). We use SSAM’s setting to evaluate PPCG [74].

We use the test data provided by STENCILGEN [68]. We tested three PERKS (SHM) implementations: PERKS_SM that only uses shared memory to cache data; PERKS_REG that only uses registers to cache data; and PERKS_MIX that uses both shared memory and registers to cache data. Due to space limitations, we report only the peak performance among those three PERKS variants.

6.2.2 Conjugate Gradient Datasets. The datasets for conjugate gradient come from the SuiteSparse Matrix Collection [25]. We selected symmetric positive definite matrices that can converge in a CG solver. The details of the selected datasets are listed in Table 4.

We compare the performance of PERKS (CG Solver) with Ginkgo library [6], a widely used library heavily optimized for GPUs (including A100). We run 10,000 time steps in our performance evaluation (similar to Ginkgo’s basic setting [6]). We report the speedup per time step and the measured sustained bandwidth achieved by Ginkgo. For PERKS (CG Solver), we run different variants that implement caching the vector r or the matrix A plus the additional caching of TB-level search result and thread-level search result policies. We only report the best-performing variant for each dataset.

6.3 Sizes of Domains and Problems

PERKS intuitively favors smaller domain/problem sizes. However, for a fair evaluation of PERKS, we can not choose arbitrarily small domain sizes; we need domain/input sizes that fully utilize the compute capability of the device. Similar to [24], we conducted an elaborate set of experiments for every individual stencil benchmark to identify the minimum domain size that would fully utilize the device. Note that domain/problem sizes that are beyond domain/problem sizes that could fully utilize the device are effectively




Figure 5: Comparison of PERKS(SHM) [48] over a wide range of stencil libraries in a wide range of 2D/3D stencil benchmarks on A100 (top) and V100 (bottom) GPUs. (a) & (c) in the left report the performance; (b) & (d) in the right report the geometric mean speedup of PERKS(SHM) over other state-of-the-art implementations.




Figure 6: Comparison of PERKS(CG Solver, CUDA Sample [58]+CUB [54]) over Ginkgo [6] on A100 (top) and V100 (bottom) GPUs in datasets (D1 to D20) listed in Table 4. (a) & (c) in the left report the performance (sustained memory bandwidth); (b) & (d) in the right report the geometric mean speedup of PERKS(CG Solver) over Ginkgo.

serialized by the device once we go beyond peak concurrency sustainable by the device. Table 3 summarizes the domain sizes for the different stencil benchmarks that would achieve a fair performance in the SHM implementation.

For conjugate gradient experiments, we include datasets from SuiteSparse that cover a wide range of problem sizes: from strong-scaling small dataset sizes that would fit in L2 cache and up to large dataset sizes typically reported by libraries for a single GPU of the same generations we use (Ginkgo [6, 8] and MAGMA [9, 72]).

6.4 Iterative 2D/3D Stencils

Figure 5 compares the performance of SHM [48] and PERKS based SHM with a wide range of state-of-the-art stencil implementations/libraries. The performance of SHM is comparable to state-of-the-art implementations, i.e. Bricks [83] and SSAM [16], across all stencil benchmarks. Applying PERKS consistently speedup SHMs: in comparison to SHM, PERKS (SHM) achieves a the geometric mean speedup of 1.95x in A100 and 2.31x in V100, for 2D stencils. The geometric mean speedup for 3D stencils is 1.13x for A100 and 1.36x for V100.

In comparison to the best state-of-the-art spatial blocking implementations/libraries, SSAM, PERKS (SHM) achieves a geometric mean speedup of 1.91x and 2.11x for 2D stencils in A100 and V100, respectively. The geometric mean speedup for 3D stencils is 1.15x for A100 and 1.14x for V100. In comparison to the state-of-the-art temporal blocking implementation of STENCILGEN, PERKS (SHM) achieves a geometric mean speedup of 1.28x (A100) and 1.02x (V100).

6.5 Conjugate Gradient

Figure 6 compares PERKS (CG Solver) to Ginkgo. We observe significantly higher performance advantage when the input size is within the L2 cache capacity. This phenomenon implies that PERKS (CG Solver) automatically benefits from the large L2 cache, possibly because the constant matrix can reside in the L2 cache, saving memory traffic to the global memory. When the input is less than the L2 capacity, PERKS running on A100 achieves a geometric mean of 4.49x speedups; in V100, PERKS achieves 5.50x speedups. When the input matrix exceeds the L2 cache capacity, PERKS running on A100 achieves a geometric mean of 1.18x speedup. On V100, PERKS

achieves a geometric mean of 1.64x (double) speedups. It is important to remember that the Ginkgo library that we use as baseline is among the top performing libraries in CG solvers, emphasizing GPU optimizations [6].

Note that regardless of whether we stay within the L2 cache capacity or exceed it, we are still caching the domain using one of the caching policies we described earlier.

7 RELATED WORK

The concept of persistent threads and persistent kernels dates back to the introduction of CUDA [4, 31]. The main motivation for persistence at the time was load imbalance issues with the runtime warp scheduler [4, 13]. Later research focused on using persistent kernels to overcome the kernel invocation overhead (which was high at the time). GPUrdma [23] and GPU-Ether [40] expanded on the concept of persistent kernels to reduce the latency of network communication.

As on-chip resources increased, researchers began to capitalize on data reuse in persistent kernel. Most of them focused on specific applications, GPUrdma [23] proposed to keep the constant matrix in shared memory. Khorasani et al. [41] proposed to keep parameters in registers. Zhu et al. [84] proposed a sparse persistent implementation of recurrent neural networks. To our knowledge, this work is the first to propose a generic and methodological blueprint for accelerating memory-bound iterative applications using persistent kernels.

8 CONCLUSION

We propose a persistent kernel execution model for iterative applications. We enhance performance by moving the time loop to the kernel and cache the intermediate output of each time step with unused on-chip resources. We show a notable performance improvement for iterative 2D/3D stencils and a conjugate gradient solver for both V100 and A100 over highly optimized baselines.

ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI under Grant Numbers JP22H03600 and JP21K17750. This work was supported by JST, PRESTO Grant Number JPMJPR200MA, Japan. This paper is based on results obtained from JPNP20006 project, commissioned by the New Energy and Industrial Technology Development Organization (NEDO). This manuscript has been co-authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The publisher acknowledges the US government license to provide public access under the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan/>). The authors wish to express their sincere appreciation to Jens Domke, Aleksandr Drozd, Emil Vatai and other RIKEN R-CCS colleagues for their invaluable advice and guidance throughout the course of this research. They also wish to thank Dr. Zhao Tuowen from the SambaNova for the helpful discussions. Finally, the first author would also like to express his gratitude to RIKEN R-CCS for offering the opportunity to undertake this research in an intern program.

REFERENCES

- [1] 2022. TOP500. <https://www.top500.org/lists/top500/2022/06/highs/> [Online; accessed 27-Mar-2021].
- [2] Laksono Adhianto, Sinchan Banerjee, Mike Fagan, Mark Krentel, Gabriel Marin, John Mellor-Crummey, and Nathan R Tallent. 2010. HPCToolkit: Tools for performance analysis of optimized parallel programs. *Concurrency and Computation: Practice and Experience* 22, 6 (2010), 685–701.
- [3] Abal-Kassim Cheik Ahamed and Frédéric Magoulès. 2017. Efficient implementation of Jacobi iterative method for large sparse linear systems on graphic processing units. *The Journal of Supercomputing* 73, 8 (2017), 3411–3432.
- [4] Timo Aila and Samuli Laine. 2009. Understanding the efficiency of ray traversal on GPUs. In *Proceedings of the conference on high performance graphics 2009*. 145–149.
- [5] José I. Aliaga, Joaquín Pérez, and Enrique S. Quintana-Ortí. 2015. Systematic Fusion of CUDA Kernels for Iterative Sparse Linear System Solvers. In *Euro-Par 2015: Parallel Processing*, Jesper Larsson Träff, Sascha Hunold, and Francesco Versaci (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 675–686.
- [6] Hartwig Anzt, Terry Cojean, Goran Flegar, Fritz Göbel, Thomas Grützmacher, Pratik Nayak, Tobias Ribizel, Yuhsiang Mike Tsai, and Enrique S. Quintana-Ortí. 2022. Ginkgo: A Modern Linear Operator Algebra Framework for High Performance Computing. *ACM Trans. Math. Softw.* 48, 1, Article 2 (feb 2022), 33 pages. <https://doi.org/10.1145/3480935>
- [7] Hartwig Anzt, Mark Gates, Jack Dongarra, Moritz Kreutzer, Gerhard Wellein, and Martin Köhler. 2017. Preconditioned Krylov solvers on GPUs. *Parallel Comput.* 68 (2017), 32–44.
- [8] Hartwig Anzt, Yuhsiang M. Tsai, Ahmad Abdelfattah, Terry Cojean, and Jack Dongarra. 2020. Evaluating the Performance of NVIDIA's A100 Ampere GPU for Sparse and Batched Computations. In *2020 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*. 26–38. <https://doi.org/10.1109/PMBS51919.2020.00009>
- [9] N. Beams, A. Abdelfattah, S. Tomov, J. Dongarra, T. Kolev, and Y. Dudouit. 2020. High-Order Finite Element Method using Standard and Device-Level Batch GEMM on GPUs. In *2020 IEEE/ACM 11th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (Scala)*. IEEE Computer Society, Los Alamitos, CA, USA, 53–60. <https://doi.org/10.1109/Scala51936.2020.00012>
- [10] Mehmet E Belviranlı, Peng Deng, Laxmi N Bhuyan, Rajiv Gupta, and Qi Zhu. 2015. Peerwave: Exploiting wavefront parallelism on gpus with peer-sm synchronization. In *Proceedings of the 29th ACM on International Conference on Supercomputing*. 25–35.
- [11] U. Bondhugula, V. Bandishti, and I. Pananilath. 2017. Diamond Tiling: Tiling Techniques to Maximize Parallelism for Stencil Computations. *IEEE Transactions on Parallel and Distributed Systems* 28, 5 (2017), 1285–1298. <https://doi.org/10.1109/TPDS.2016.2615094>
- [12] Uday Bondhugula, Vinayaka Bandishti, and Irshad Pananilath. 2017. Diamond Tiling: Tiling Techniques to Maximize Parallelism for Stencil Computations. *IEEE Trans. Parallel Distrib. Syst.* 28, 5 (May 2017), 1285–1298. <https://doi.org/10.1109/TPDS.2016.2615094>
- [13] Long Chen, Oreste Villa, Sriram Krishnamoorthy, and Guang R Gao. 2010. Dynamic load balancing on single- and multi-GPU systems. In *2010 IEEE International Symposium on Parallel & Distributed Processing (IPDPS)*. IEEE, 1–12.
- [14] Peng Chen, Mohamed Wahib, Shinichiro Takizawa, Ryousei Takano, and Satoshi Matsuoka. 2018. Efficient Algorithms for the Summed Area Tables Primitive on GPUs. In *2018 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 482–493.
- [15] Peng Chen, Mohamed Wahib, Shinichiro Takizawa, Ryousei Takano, and Satoshi Matsuoka. 2019. IFDK: A Scalable Framework for Instant High-Resolution Image Reconstruction. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (Denver, Colorado) (SC '19)*. Association for Computing Machinery, New York, NY, USA, Article 84, 24 pages. <https://doi.org/10.1145/3295500.3356163>
- [16] Peng Chen, Mohamed Wahib, Shinichiro Takizawa, Ryousei Takano, and Satoshi Matsuoka. 2019. A versatile software systolic execution model for GPU memory-bound kernels. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–81.
- [17] Peng Chen, Mohamed Wahib, Xiao Wang, Takahiro Hirofuchi, Hirotaaka Ogawa, Ander Biguri, Richard Boardman, Thomas Blumensath, and Satoshi Matsuoka. 2021. Scalable FBP Decomposition for Cone-Beam CT Reconstruction. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (St. Louis, Missouri) (SC '21)*. Association for Computing Machinery, New York, NY, USA, Article 9, 16 pages. <https://doi.org/10.1145/3458817.3476139>
- [18] Hadrien Courtécuisse and Jérémie Allard. 2009. Parallel dense gauss-seidel algorithm on many-core processors. In *2009 11th IEEE International Conference on High Performance Computing and Communications*. IEEE, 139–147.
- [19] Raphaël Couturier and Stéphane Domas. 2012. Sparse systems solving on GPUs with GMRES. *The Journal of Supercomputing* 59, 3 (2012), 1504–1516.

- [20] Nvidia CUDA. 2021. CUDA C Programming Guide. <https://docs.nvidia.com/cuda/cuda-c-programming-guide> [Online; accessed 3-Jan-2023].
- [21] Nvidia CUDA. 2021. NVIDIA A100 Tensor Core GPU Architecture. <https://resources.nvidia.com/en-us-genomics-ep/ampere-architecture-white-paper> [Online; accessed 20-July-2021].
- [22] NVIDIA CUDA. 2023. CUDA Toolkit Documentation. *NVIDIA Developer Zone*. <http://docs.nvidia.com/cuda/index.html> (2023).
- [23] Feras Daoud, Amir Wataf, and Mark Silberstein. 2016. GPUrdma: GPU-side library for high performance networking from GPU kernels. In *Proceedings of the 6th International Workshop on Runtime and Operating Systems for Supercomputers*. 1–8.
- [24] Philip Roth David Eberius, David Rogers. 2022. Understanding Strong Scaling on GPUs Using Empirical Performance Saturation Size. In *The International Conference for High Performance Computing, Networking, Storage, and Analysis (International Workshop on Performance Portability and Productivity (P3HPC))*.
- [25] Timothy A Davis and Yifan Hu. 2011. The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software (TOMS)* 38, 1 (2011), 1–25.
- [26] Toshio Endo. 2018. Applying recursive temporal blocking for stencil computations to deeper memory hierarchy. In *2018 IEEE 7th Non-Volatile Memory Systems and Applications Symposium (NVMISA)*. IEEE, 19–24.
- [27] Toshio Endo, Yuki Takasaki, and Satoshi Matsuoka. 2015. Realizing Extremely Large-Scale Stencil Applications on GPU Supercomputers. In *2015 IEEE 21st International Conference on Parallel and Distributed Systems (ICPADS)*. 625–632. <https://doi.org/10.1109/ICPADS.2015.84>
- [28] Marco Fratarcangeli, Valentina Tibaldo, and Fabio Pellacini. 2016. Vivace: A practical gauss-seidel method for stable soft body dynamics. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 1–9.
- [29] Andreas Frommer, Kathryn Lund, and Daniel B Szyld. 2017. Block Krylov subspace methods for functions of matrices. (2017).
- [30] Tobias Grosser, Albert Cohen, Paul H. J. Kelly, J. Ramanujam, P. Sadayappan, and Sven Verdoolaege. 2013. Split Tiling for GPUs: Automatic Parallelization Using Trapezoidal Tiles. In *Proceedings of the 6th Workshop on General Purpose Processor Using Graphics Processing Units* (Houston, Texas, USA) (*GPGPU-6*). Association for Computing Machinery, New York, NY, USA, 24–31. <https://doi.org/10.1145/2458523.2458526>
- [31] Kshitij Gupta, Jeff A Stuart, and John D Owens. 2012. *A study of persistent threads style GPU programming for GPGPU workloads*. IEEE.
- [32] Tobias Gysi, Tobias Grosser, and Torsten Hoefler. 2015. MODESTO: Data-centric Analytic Optimization of Complex Stencil Programs on Heterogeneous Architectures. In *Proceedings of the 29th ACM International Conference on Supercomputing, ICS'15, Newport Beach/Irvine, CA, USA, June 08 - 11, 2015*. 177–186. <https://doi.org/10.1145/2751205.2751223>
- [33] Bastian Hagedorn, Larisa Stoltzfus, Michel Steuwer, Sergei Gorlatch, and Christophe Dubach. 2018. High performance stencil code generation with lift. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization*. 100–112.
- [34] Justin Holewinski, Louis-Noël Pouchet, and P. Sadayappan. 2012. High-Performance Code Generation for Stencil Computations on GPU Architectures. In *Proceedings of the 26th ACM International Conference on Supercomputing* (San Servolo Island, Venice, Italy) (*ICS '12*). Association for Computing Machinery, New York, NY, USA, 311–320. <https://doi.org/10.1145/2304576.2304619>
- [35] Justin Holewinski, Louis-Noël Pouchet, and P. Sadayappan. 2012. High-performance Code Generation for Stencil Computations on GPU Architectures. In *Proceedings of the 26th ACM International Conference on Supercomputing* (San Servolo Island, Venice, Italy) (*ICS '12*). ACM, New York, NY, USA, 311–320. <https://doi.org/10.1145/2304576.2304619>
- [36] Yasuhiro Idomura, Takuya Ina, Yussuf Ali, and Toshiyuki Imamura. 2020. Acceleration of fusion plasma turbulence simulations using the mixed-precision communication-avoiding krylov method. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2020, Virtual Event / Atlanta, Georgia, USA, November 9-19, 2020*. 93. <https://doi.org/10.1109/SC41405.2020.00097>
- [37] F. Irigoien and R. Triolet. 1988. Supernode Partitioning. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (*POPL '88*). ACM, New York, NY, USA, 319–329. <https://doi.org/10.1145/73560.73588>
- [38] Zhe Jia, Marco Maggioni, Jeffrey Smith, and Daniele Paolo Scarpazza. 2019. Dissecting the NVidia Turing T4 GPU via Microbenchmarking. arXiv:1903.07486 [cs.DC]
- [39] Zhe Jia, Marco Maggioni, Benjamin Staiger, and Daniele Paolo Scarpazza. 2018. Dissecting the NVidia Volta GPU Architecture via Microbenchmarking. *CoRR* abs/1804.06826 (2018). arXiv:1804.06826 <http://arxiv.org/abs/1804.06826>
- [40] Changue Jung, Suhwan Kim, Ikjun Yeom, Honguk Woo, and Younghoon Kim. 2021. GPU-Ether: GPU-native Packet I/O for GPU Applications on Commodity Ethernet. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 1–10.
- [41] Farzad Khorasani, Hodjat Asghari Esfeden, Nael Abu-Ghazaleh, and Vivek Sarkar. 2018. In-Register Parameter Caching for Dynamic Neural Nets with Virtual Persistent Processor Specialization. In *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 377–389. <https://doi.org/10.1109/MICRO.2018.00038>
- [42] Ki-Hwan Kim, KyoungHo Kim, and Q-Han Park. 2011. Performance analysis and optimization of three-dimensional FDTD on GPU using roofline model. *Computer Physics Communications* 182, 6 (2011), 1201–1207.
- [43] Aleksandr Kochurov and Dimitrii Golovashkin. 2015. GPU implementation of Jacobi Method and Gauss-Seidel Method for Data Arrays that Exceed GPU-dedicated Memory Size. *Journal of Mathematical Modelling and Algorithms in Operations Research* 14, 4 (2015), 393–405.
- [44] Sriram Krishnamoorthy, Muthu Baskaran, Uday Bondhugula, J. Ramanujam, Atanas Rountev, and P Sadayappan. 2007. Effective Automatic Parallelization of Stencil Computations. In *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation* (San Diego, California, USA) (*PLDI '07*). ACM, New York, NY, USA, 235–244. <https://doi.org/10.1145/1250734.1250761>
- [45] Chao Li, Yi Yang, Zhen Lin, and Huiyang Zhou. 2015. Automatic data placement into GPU on-chip memory resources. In *2015 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. IEEE, 23–33.
- [46] John D. C. Little. 1961. A Proof for the Queuing Formula: $L = \lambda W$. *Operations Research* 9 (1961), 383–387.
- [47] Tareq Malas, Georg Hager, Hatem Ltaief, Holger Stengel, Gerhard Wellein, and David Keyes. 2015. Multicore-optimized wavefront diamond blocking for optimizing stencil updates. *SIAM Journal on Scientific Computing* 37, 4 (2015), C439–C464.
- [48] Naoya Maruyama and Takayuki Aoki. 2014. Optimizing Stencil Computations for NVIDIA Kepler GPUs. In *Proceedings of the 1st International Workshop on High-Performance Stencil Computations*, Armin Größlinger and Harald Köstler (Eds.). Vienna, Austria, 89–95.
- [49] Kazuaki Matsumura, Hamid Reza Zohouri, Mohamed Wahib, Toshio Endo, and Satoshi Matsuoka. 2020. AN5D: automated stencil framework for high-degree temporal blocking on GPUs. In *CGO '20: 18th ACM/IEEE International Symposium on Code Generation and Optimization*, San Diego, CA, USA, February, 2020. 199–211. <https://doi.org/10.1145/3368826.3377904>
- [50] Xinxin Mei, Kaiyong Zhao, Chengjian Liu, and Xiaowen Chu. 2014. Benchmarking the memory hierarchy of modern GPUs. In *IFIP International Conference on Network and Parallel Computing*. Springer, 144–156.
- [51] Jiayuan Meng and Kevin Skadron. 2009. Performance Modeling and Automatic Ghost Zone Optimization for Iterative Stencil Loops on GPUs. In *Proceedings of the 23rd International Conference on Supercomputing* (Yorktown Heights, NY, USA) (*ICS '09*). Association for Computing Machinery, New York, NY, USA, 256–265. <https://doi.org/10.1145/1542275.1542313>
- [52] Jiayuan Meng and Kevin Skadron. 2009. Performance Modeling and Automatic Ghost Zone Optimization for Iterative Stencil Loops on GPUs. In *Proceedings of the 23rd International Conference on Supercomputing* (Yorktown Heights, NY, USA) (*ICS '09*). ACM, New York, NY, USA, 256–265. <https://doi.org/10.1145/1542275.1542313>
- [53] Jiayuan Meng and Kevin Skadron. 2011. A performance study for iterative stencil loops on GPUs with ghost zone optimizations. *International Journal of Parallel Programming* 39, 1 (2011), 115–142.
- [54] Duane Merrill and Michael Garland. 2016. Merge-based parallel sparse matrix-vector multiplication. In *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 678–689.
- [55] Paulius Micikevicius. 2009. 3D Finite Difference Computation on GPUs Using CUDA. In *Proceedings of 2nd Workshop on General Purpose Processing on Graphics Processing Units* (Washington, D.C., USA) (*GPGPU-2*). Association for Computing Machinery, New York, NY, USA, 79–84. <https://doi.org/10.1145/1513895.1513905>
- [56] Takayuki Muranushi and Junichiro Makino. 2015. Optimal Temporal Blocking for Stencil Computation. *Procedia Computer Science* 51 (2015), 1303–1312. <https://doi.org/10.1016/j.procs.2015.05.315> International Conference On Computational Science, ICCS 2015.
- [57] Nvidia. 2021. CUB Library. <https://nvlabs.github.io/cub>
- [58] Nvidia. 2021. NVIDIA CUDA Sample. <https://docs.nvidia.com/cuda/cuda-samples/index.html>
- [59] Nvidia. 2023. NVIDIA A100 Tensor Core GPU Architecture. <https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/nvidia-ampere-architecture-whitepaper.pdf>
- [60] Nvidia. 2023. NVIDIA CUDA Runtime API. <https://docs.nvidia.com/cuda/cuda-runtime-api/index.html>
- [61] Nvidia. 2023. Programming guide. <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>
- [62] Georg Ofenbeck, Ruedi Steinmann, Victoria Caparros, Daniele G Spampinato, and Markus Püschel. 2014. Applying the roofline model. In *2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 76–85.

- [63] Carl Pearson, Mert Hidayetoğlu, Mohammad Almasri, Omer Anjum, I-Hsin Chung, Jinjun Xiong, and Wen-Mei W. Hwu. 2020. Node-Aware Stencil Communication for Heterogeneous Supercomputers. In *2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. 796–805. <https://doi.org/10.1109/IPDPSW50202.2020.00136>
- [64] John W Pearson and Jennifer Pestana. 2020. Preconditioners for Krylov subspace methods: An overview. *GAMM-Mitteilungen* 43, 4 (2020), e202000015.
- [65] Everett Phillips and Massimiliano Fatica. 2014. A CUDA implementation of the High Performance Conjugate Gradient benchmark. In *International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems*. Springer, 68–84.
- [66] Prashant Rawat, Martin Kong, Tom Henretty, Justin Holewinski, Kevin Stock, Louis-Noël Pouchet, J. Ramanujam, Atanas Rountev, and P. Sadayappan. 2015. SDSLc: A Multi-Target Domain-Specific Compiler for Stencil Computations. In *Proceedings of the 5th International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing (Austin, Texas) (WOLFHPC '15)*. Association for Computing Machinery, New York, NY, USA, Article 6, 10 pages. <https://doi.org/10.1145/2830018.2830025>
- [67] Prashant Singh Rawat, Changwan Hong, Mahesh Ravishankar, Vinod Grover, Louis-Noël Pouchet, and P Sadayappan. 2016. Effective resource management for enhancing performance of 2D and 3D stencils on GPUs. In *Proceedings of the 9th Annual Workshop on General Purpose Processing using Graphics Processing Unit*. 92–102.
- [68] Prashant Singh Rawat, Miheer Vaidya, Aravind Sukumaran-Rajam, Mahesh Ravishankar, Vinod Grover, Atanas Rountev, Louis-Noël Pouchet, and P Sadayappan. 2018. Domain-specific optimization and generation of high-performance GPU code for stencil computations. *Proc. IEEE* 106, 11 (2018), 1902–1920.
- [69] Prashant Singh Rawat, Miheer Vaidya, Aravind Sukumaran-Rajam, Atanas Rountev, Louis-Noël Pouchet, and P. Sadayappan. 2019. On Optimizing Complex Stencils on GPUs. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 641–652. <https://doi.org/10.1109/IPDPS.2019.00073>
- [70] Amir Hossein Nodehi Sabet, Zhijia Zhao, and Rajiv Gupta. 2020. Subway: Minimizing data transfer during out-of-GPU-memory graph processing. In *Proceedings of the Fifteenth European Conference on Computer Systems*. 1–16.
- [71] Sameer S Shende and Allen D Malony. 2006. The TAU parallel performance system. *The International Journal of High Performance Computing Applications* 20, 2 (2006), 287–311.
- [72] Stanimire Tomov, Rajib Nath, Hatem Ltaief, and Jack Dongarra. 2010. Dense Linear Algebra Solvers for Multicore with GPU Accelerators. In *Proc. of the IEEE IPDPS'10*. IEEE Computer Society, Atlanta, GA, 1–8. DOI: 10.1109/IPDPSW.2010.5470941.
- [73] Ben van Werkhoven. 2019. Kernel Tuner: A search-optimizing GPU code auto-tuner. *Future Generation Computer Systems* 90 (2019), 347–358. <https://doi.org/10.1016/j.future.2018.08.004>
- [74] Sven Verdoolaege, Juan Carlos Juega, Albert Cohen, Jose Ignacio Gomez, Christian Tenllado, and Francky Catthoor. 2013. Polyhedral parallel code generation for CUDA. *ACM Transactions on Architecture and Code Optimization (TACO)* 9, 4 (2013), 1–23.
- [75] Nandita Vijaykumar, Kevin Hsieh, Gennady Pekhimenko, Samira Khan, Ashish Shrestha, Saugata Ghose, Adwait Jog, Phillip B. Gibbons, and Onur Mutlu. 2016. Zorua: A holistic approach to resource virtualization in GPUs. In *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 1–14. <https://doi.org/10.1109/MICRO.2016.7783718>
- [76] Vasily Volkov. 2010. Better performance at lower occupancy. In *Proceedings of the GPU technology conference, GTC*, Vol. 10. San Jose, CA, 16.
- [77] Mohamed Wahib and Naoya Maruyama. 2014. Scalable Kernel Fusion for Memory-Bound GPU Applications. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2014, New Orleans, LA, USA, November 16-21, 2014*. 191–202. <https://doi.org/10.1109/SC.2014.21>
- [78] Mohamed Wahib, Haoyu Zhang, Truong Thao Nguyen, Aleksandr Drozd, Jens Domke, Lingqi Zhang, Ryousei Takano, and Satoshi Matsuoka. 2020. Scaling Distributed Deep Learning Workloads beyond the Memory Capacity with KARMA. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (Atlanta, Georgia) (SC '20)*. IEEE Press, Article 19, 15 pages.
- [79] Gerhard Wellein, Georg Hager, Thomas Zeiser, Markus Wittmann, and Holger Fehske. 2009. Efficient Temporal Blocking for Stencil Computations by Multicore-Aware Wavefront Parallelization. In *2009 33rd Annual IEEE International Computer Software and Applications Conference*, Vol. 1. 579–586. <https://doi.org/10.1109/COMPSAC.2009.82>
- [80] M. Wolfe. 1989. More Iteration Space Tiling. In *Proceedings of the 1989 ACM/IEEE Conference on Supercomputing (Reno, Nevada, USA) (Supercomputing '89)*. ACM, New York, NY, USA, 655–664. <https://doi.org/10.1145/76263.76337>
- [81] Henry Wong, Misel-Myrto Papadopoulou, Maryam Sadooghi-Alvandi, and Andreas Moshovos. 2010. Demystifying GPU microarchitecture through microbenchmarking. In *2010 IEEE International Symposium on Performance Analysis of Systems & Software (ISPASS)*. 235–246. <https://doi.org/10.1109/ISPASS.2010.5452013>
- [82] L. Zhang, M. Wahib, H. Zhang, and S. Matsuoka. 2020. A Study of Single and Multi-device Synchronization Methods in Nvidia GPUs. In *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE Computer Society, Los Alamitos, CA, USA, 483–493. <https://doi.org/10.1109/IPDPS47924.2020.00057>
- [83] Tuowen Zhao, Protonu Basu, Samuel Williams, Mary Hall, and Hans Johansen. 2019. Exploiting reuse and vectorization in blocked stencil computations on CPUs and GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–44.
- [84] Feiwen Zhu, Jeff Pool, Michael Andersch, Jeremy Appleyard, and Fung Xie. 2018. Sparse Persistent RNNs: Squeezing Large Recurrent Networks On-Chip. *CoRR abs/1804.10223* (2018). arXiv:1804.10223 <http://arxiv.org/abs/1804.10223>