



Exploiting Scratchpad Memory for Deep Temporal Blocking

A case study for 2D Jacobian 5-point iterative stencil kernel (j2d5pt)

Lingqi Zhang
Tokyo Tech
AIST
Japan

Mohamed Wahib
RIKEN R-CCS
Japan

Peng Chen
AIST
RIKEN R-CCS
Japan

Jintao Meng
SIAT
China

Xiao Wang
ORNL
USA

Toshio Endo
Tokyo Tech
Japan

Satoshi Matsuoka
RIKEN R-CCS
Tokyo Tech
Japan

ABSTRACT

General Purpose Graphics Processing Units (GPGPU) are used in most of the top systems in HPC. The total capacity of scratchpad memory has increased by more than 40 times in the last decade. However, existing optimizations for stencil computations using temporal blocking have not aggressively exploited the large capacity of scratchpad memory. This work uses the 2D Jacobian 5-point iterative stencil as a case study to investigate the use of large scratchpad memory. Unlike existing research that tiles the domain in a thread block fashion, we tile the domain so that each tile is large enough to utilize all available scratchpad memory on the GPU. Consequently, we process several time steps inside a single tile before offloading the result back to global memory. Our evaluation shows that our performance is comparable to state-of-the-art implementations, yet our implementation is much simpler and does not require auto-generation of code.

CCS CONCEPTS

• **Computing methodologies** → **Vector / streaming algorithms; Massively parallel algorithms;**

KEYWORDS

GPGPU, Temporal Blocking, Iterative Stencil Solvers

ACM Reference Format:

Lingqi Zhang, Mohamed Wahib, Peng Chen, Jintao Meng, Xiao Wang, Toshio Endo, and Satoshi Matsuoka. 2023. Exploiting Scratchpad Memory for Deep Temporal Blocking: A case study for 2D Jacobian 5-point iterative stencil kernel (j2d5pt). In *15th Workshop on General Purpose Processing Using GPU (GPGPU '23), February 25, 2023, Montreal, Canada*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3589236.3589242>

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.
GPGPU '23, February 25, 2023, Montreal, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0776-6/23/02...\$15.00
<https://doi.org/10.1145/3589236.3589242>

1 INTRODUCTION

When observing the previous generations of GPUs, Nivida GPUs for instance, there is a clear trend of increase in the cache capacity. Especially the volume of scratchpad memory (or shared memory in CUDA [2]) increased from 720 KB in K20 (2013) to 17.30 MB in A100 (2020). The latest H100 (2023) GPU even pushes max usable shared memory to be 29.83 MB to more than 200 KB per stream multiprocessor(SM).

GPU optimizations that are commonly used in HPC applications were designed mostly assuming that scratchpad memory is not larger than 100 KB per stream multiprocessor [3]. There is a potential in leveraging the untapped scratchpad memory to aggressively optimize for data locality.

In this work, we use a case study kernel commonly used in HPC applications, namely 2D Jacobian 5-point iterative stencil, to fully take advantage of the scratchpad memory for tiling data in an unusual way. More specifically, we run each of the tiles in a serial fashion one after the other while aggressively using the shared memory to run each tile entirely from shared memory. We use device-wide synchronization to resolve the spatial dependency between thread blocks. We demonstrate a new approach to leverage the large capacity of shared memory by proposing a temporal blocking stencil scheme that optimizes for peak data locality, i.e. running the entire problem from shared memory. Our method is much simpler than complex temporal blocking schemes; iterative kernels that use our methods can be manually written, unlike complex temporal schemes that require auto-generation of code.

2 RELATED WORK

Temporal blocking [1, 4] tiles the domain and processes the domain with in combined time steps. Due to space limitations, we mainly review StencilGen [4] and AN5D [1]. Both works used 2.5D or 3.5D tiling and relied on code auto generation for performance optimization. In addition, they relied on overlapped tiling within thread blocks. They did not exploit the inter thread block data exchange pattern. Regarding the usage of scratchpad memory, StencilGen stores all combined time steps in scratchpad memory; AN5D uses scratchpad memory conservatively for double buffer. As a result, in the j2d5pt double-precision kernel. StencilGen and AN5D consumed about 4.32 MB and 0.864 MB scratchpad memory, respectively. So, both AN5D and StencilGen left most of the scratchpad memory untapped, and are overly complex to implement.

Listing 1: Pseudo code for j2d5pt stencil kernel function

```

1 //kernel function don't assume whether ptr_in and ptr_out is in
2 device memory or scratchpad memory
3 __device__ void j2d5pt(ptr_in, ptr_out, loc_x, loc_y){
4     x = threadIdx.x;
5     t[ILP+2]; //ILP: instruction level parallel
6     for(y=0; y< ILP+2; y++){
7         t[y]=ptr_in[x, y+ind_y-1];
8         for( y=0; y< ILP; y++){
9             result[y]=ptr_in[x+loc_x-1,y+1+loc_y]*W
10            +ptr_in[x+loc_x+1,y+1+loc_y]*E
11            +t[y-1]*S
12            +t[y+1]*C
13            +t[y+1+1]*N;
14        }
15    }
16    for( y=0; y< ILP; y++){
17        ptr_out[x+loc_x, y+loc_y] = result[y];
18    }

```




Figure 1: How DTB processes the tiles. DTB loads the tile to populate the scratchpad memory with the input, processes T time steps, and then stores the results to the output address. DTB processes tiles in a serial order.

3 DEEP TEMPORAL BLOCKING (DTB)

3.1 Basic function

Listing 1 shows the base kernel function we used in this case study. We only modified the input and output pointer location to use scratchpad memory. In this kernel, we move the time loop from the host side to be inside the kernel. Next, we tile the domain of the problem spatially and run the tiles in a serial fashion. For each tile, we run it entirely to completion, over all its time steps, before we start on the next tile.

3.2 Dependency Between Thread Blocks

We use the CUDA grid-level barrier to ensure that each thread block can exchange the halo region correctly. We use the bulk synchronous parallel (BSP) model.

3.3 Processing the Tiles in Order

After we load a tile into the scratchpad memory, we process the tile for several time steps (temporal blocking) before moving to the next tiling. Figure 1 shows the process.

4 EVALUATION

We compare DTB with StencilGen [4] and AN5D [1], the state-of-the-art implementations for temporal blocking for stencils (a




Figure 2: Comparing the performance of DTB with other state-of-the-art temporal blocking implementations (SOTAs), i.e., StencilGen [4] and AN5D [1]. The temporal blocking depth (number of time steps) is marked inside the parentheses. DTB runs a 8592×8328 domain. DTB_pruned, StencilGen, and AN5D run 8192^2 domain size. We use the valid domain to evaluate the performance. DTB shows comparable performance with other SOTAs.

domain size of 8192^2). We used 8592×8328 to run the DTB. We also report a pruned version that considers 8192^2 as a valid domain size. Figure 2 shows the result: the performance of DTB is comparable to that of state-of-the-art temporal blocking implementations (SOTAs).

5 CONCLUSION

In this work, we discuss a case study on the use of scratchpad memory for DTB on the j2d5pt stencil. Instead of applying a complex temporal blocking implementation, we just tile the domain so that each tile fully occupies the scratchpad memory. Evaluation shows that DTB is compatible with other SOTAs. We anticipate that DTB could perform even better on a larger scratchpad memory architecture, which would be explored in future work.

ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI under Grant Number JP21K17750. This paper is based on results obtained from a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO). This research used resources at the Oak Ridge Leadership Computing Facility, a DOE Office of Science User Facility operated by the Oak Ridge National Laboratory. 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. 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] 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>
- [2] Nvidia. 2022. CUDA Programming guide. <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>
- [3] Prashant Singh Rawat. 2018. *Optimization of stencil computations on GPUs*. Ph.D. Dissertation. The Ohio State University.
- [4] 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.