Geospatial Filter and Refine Computations on NVidia Bluefield Data Processing Units

Derda Kaymak, Marquette University

Satish Puri, Missouri University of Science and Technology (email: satish.puri@mst.edu)

I. Introduction

Geospatial vector data consists of a collection geometries. A geometry can be a polygon or polyline represented as a sequence of vertices (x, y coordinates). A geometry can be approximated as a minimum bounding rectangle (MBR) as shown in Figure 1.

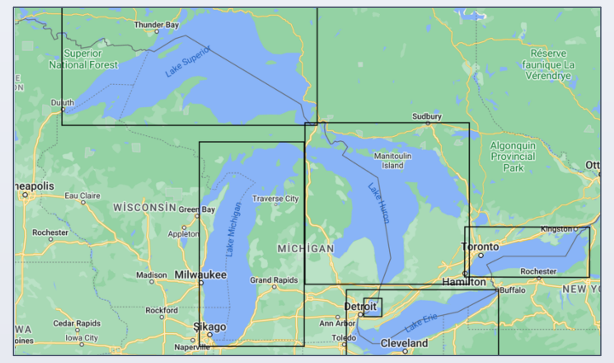

Fig. 1 Polygons representing Great Lakes. Each lake has an associated MBR.

Dataset	Size	Records
Cemetery	56 MB	193 K
Sports	590 MB	1.8 M
Lakes	9 GB	8.4 M

Table I. Geospatial Datasets [1]

Spatial Query: Given a shape, find overlapping or intersecting shapes from the database of shapes.

Spatial Join: Given two input shape datasets, find the overlapping/intersecting pairs of shapes from the input datasets.

II. Objectives

1. Benchmarking of geospatial analytics on emerging hardware like SmartNICs and Data Processing Units (shown in Figure 3).

2. Use DPUs as an offload target for geospatial analytics.

3. Design space exploration of spatial analytics based on filter and refine strategy.

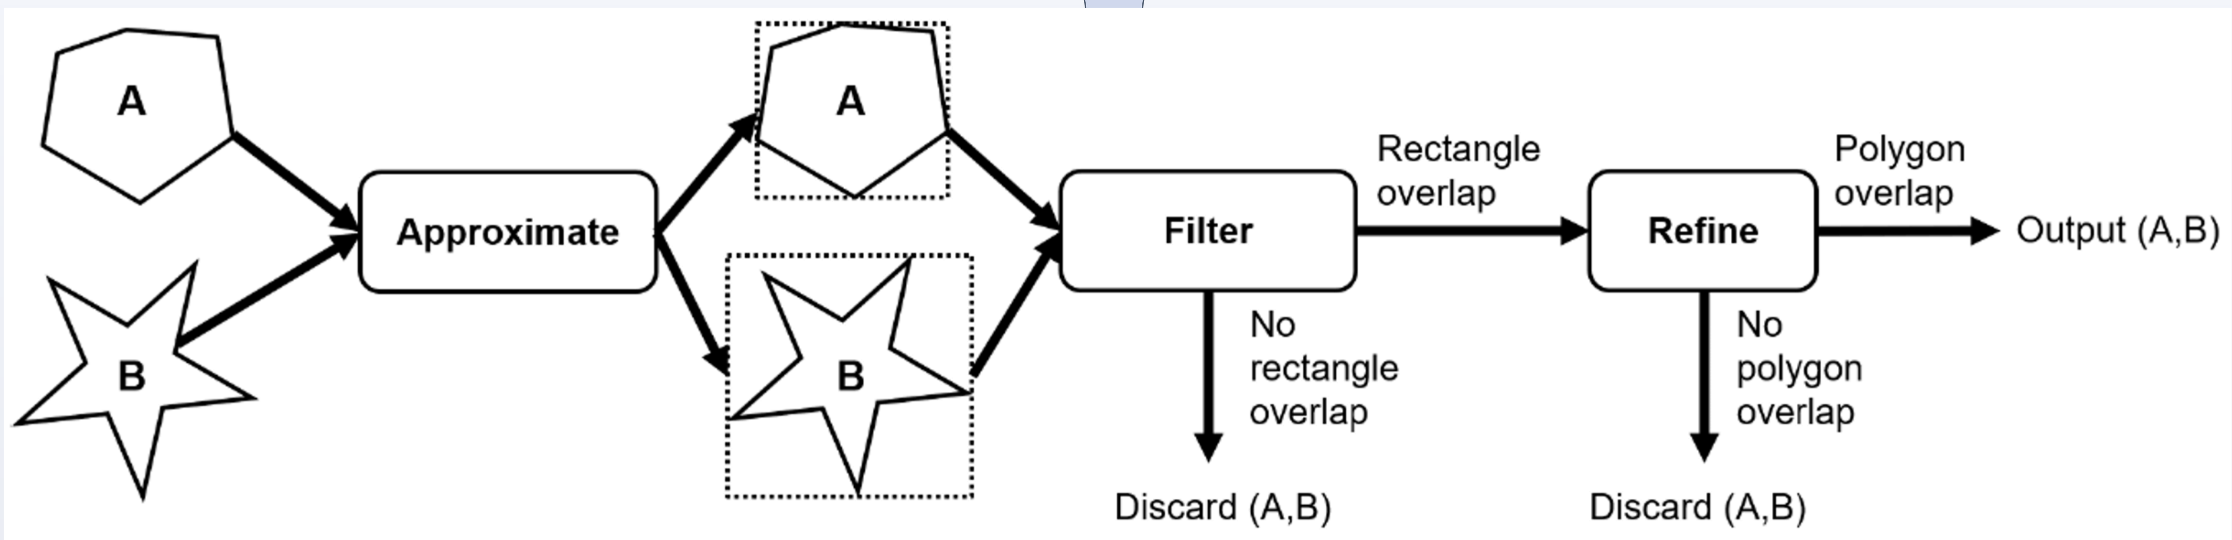

Fig. 2 Filter and refine for two input polygons using their rectangles. Only those pairs (A,B) will pass the filter if their rectangle approximations overlap. Spatial join processes millions of such pairs.

III. Experimental Setup

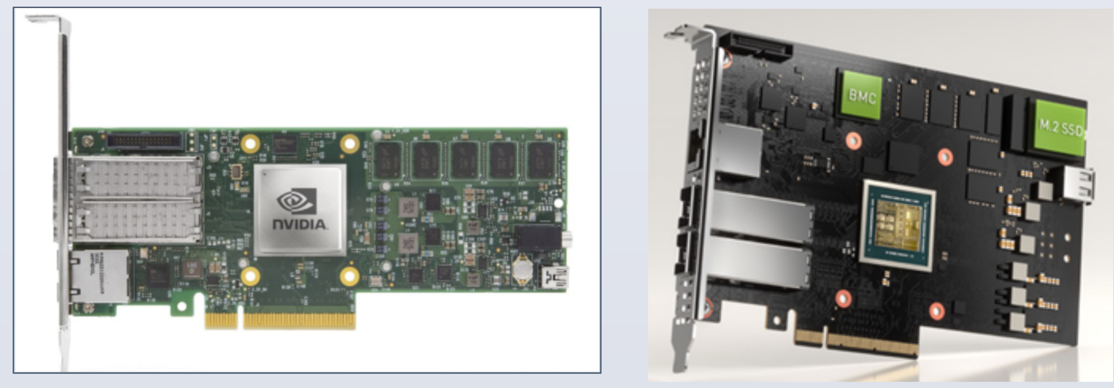

Fig. 3 NVIDIA BlueField-2 and BlueField-3 DPU [2]

Hardware details:

We used Thor cluster which has multi-core Intel CPUs and Bluefield DPUs [5].

Dual Socket Intel Xeon 16-core CPUs E5-2697A V4 @ 2.60 GHz (Broadwell).

Bluefield-2 DPU has Arm Cortex-A72 processor with 8 cores (2.4 GHz) and 16 GB memory. Bluefield-3 DPU has Armv8.2+ A78 Hercules processor with 16 cores (2.1 GHz) and 16 GB memory.

MPI library was used for internode communication. gRPC library was used CPU-DPU communication. GEOS library was used for computational geometry algorithms.

IV. Experimental Results

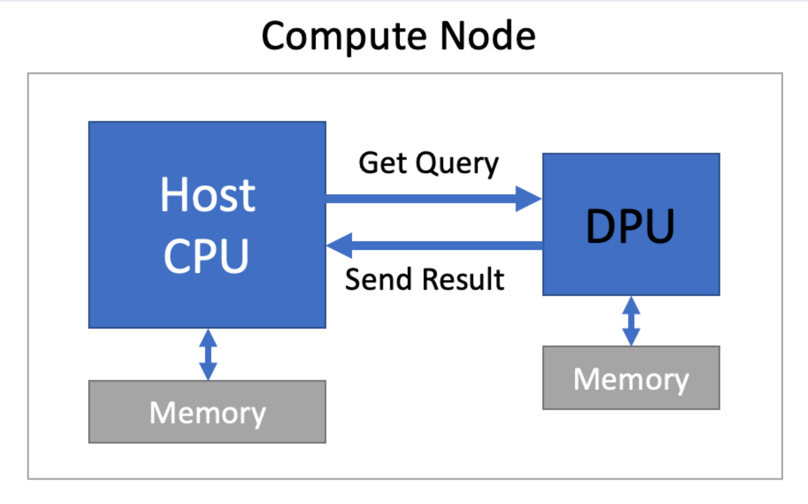

Fig. 4 Query processing offload on a heterogeneous compute node. Communication using with gRPC.

Experiment #1: The base layer shape data is divided into 128 chunks and distributed among processes with round robin scheduling. Cores of a single node were used in this experiment. CPU performs the same geospatial operation approximately 2.7 times faster than the BlueField-2 DPU regardless of the number of processes. CPU is 1.9 times faster for a single MPI process compared to CPU. With multi-threaded code on multiple cores, the performance of BlueField-3 increases more than that of the CPU. This may be an indication that the DPU can be more scalable than the CPU (Fig. 6). has 16 cores and Bluefield-2 has 8 cores.

Offload Experiment: We conducted experiments to use DPU for offloading computations from CPU. When the CPU and DPU were used together with dynamic load balancing, the processing time significantly improved. While the spatial join computation took about 76 seconds on BlueField-2 and approximately 28 seconds on the CPU, using both together reduced the processing time to around 16 seconds.

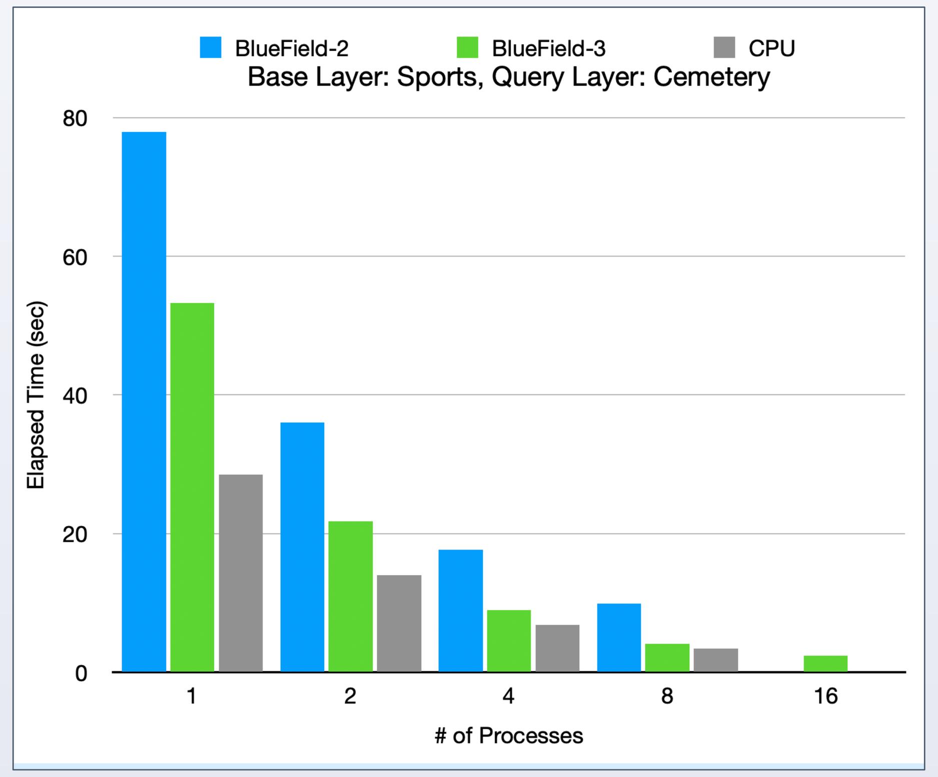

Fig. 6 Intersects performance of single node using data divided into 128 partitions.

Experiment #2: The base layer is divided into 128 chunks and different dataset pairs for base and query layers are tested. We see that the CPU spatial join performance for the Sports-Cemetery file pair is about 2.7 times better than that of the BlueField-2 and 1.9 times better than BlueField-3, but this ratio increases slightly as the size of the data pairs used increases (Fig. 7).

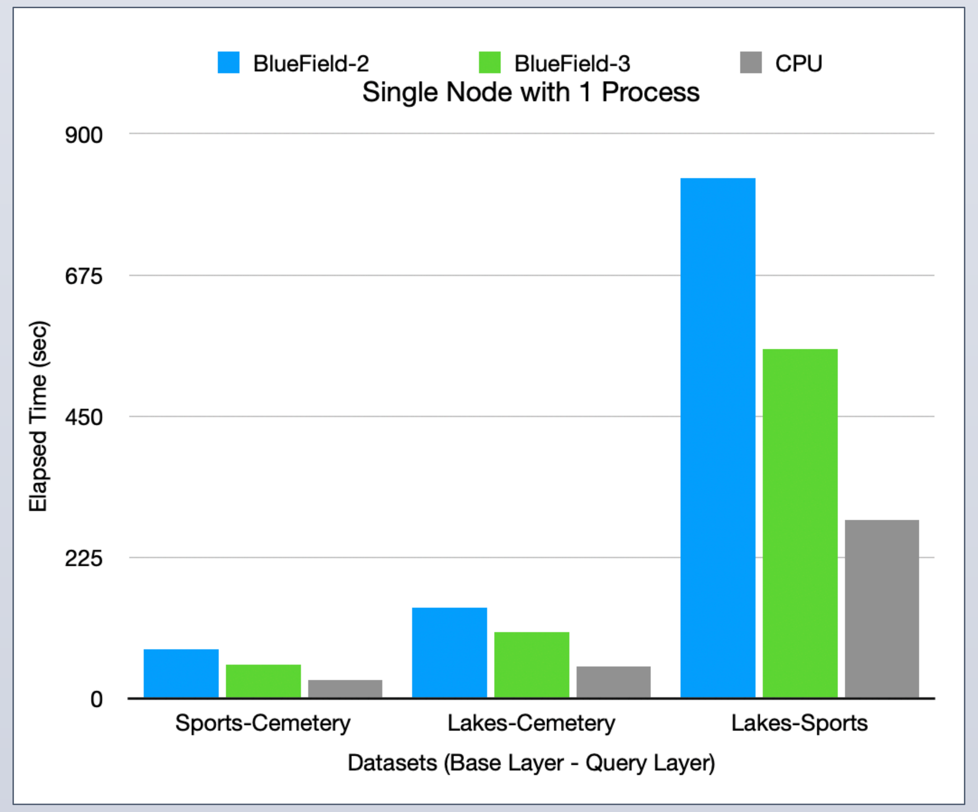

Fig. 7 Performance comparison among two generations of DPUs and CPU using data divided among MPI processes in a single node environment.

Experiment #3: Tests were conducted on BlueField-3 using different number of processes on a single node and using multiple nodes (Fig. 8).

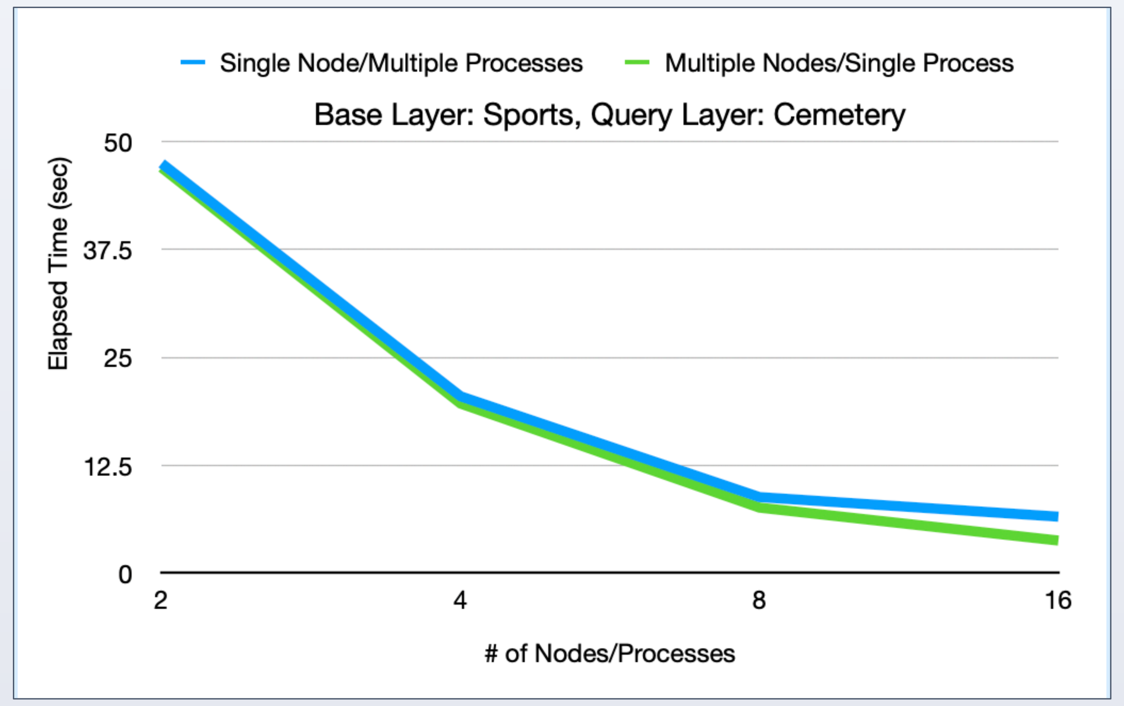

Fig. 8 Performance difference between multiple nodes and multiple processes on BlueField-3 DPU(s).

V. Conclusion

We show benchmarking results of spatial join computations on current generation DPUs. We also show that offloading filter and refine tasks to DPUs can lead to performance improvement in spatial query processing.

V. Acknowledgement

This work is supported in part by National Science Foundation CAREER Grant No. 2344578. We thank the HPC-AI Advisory Council testbed for BlueField testing and evaluation.

VI. References

1. *Suresh, Kaushik Kandadi, et al. "A Novel Framework for Efficient Offloading of Communication Operations to Bluefield SmartNICs." 2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS). IEEE, 2023.*
2. *Karamati, Sara, et al. ""Smarter" NICs for faster molecular dynamics: a case study." 2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS). IEEE, 2022.*
3. *Thor Cluster in HPCAdvisory Council, https://hpcadvisorycouncil.atlassian.net/wiki/spaces/HPCWORKS/pages/7864401/Thor*