

Towards an End-to-End Processing-in-DRAM Acceleration of Spectral Library Search

Tianyun Zhang, Franz Franchetti

Task 3135.019

Department of Electrical and Computer Engineering, College of Engineering, Carnegie Mellon University

Introduction

- Overall objective:** Explore processing-in-memory (PIM) potential to accelerate spectral library searches by integrating hardware and algorithm co-design across the mass spectrometry (MS) data analysis pipeline
- Application motivation:** Modern workloads (e.g. mass spectrometry) are increasingly **data-intensive**
- Hardware motivation:** Computer architectures are designed for **infrequent memory access**
 - Data transfers between the memory and the processor over a narrow memory channel incur high **energy and latency costs**

We make the following contributions:

- Adaptation of an approximate sum of absolute differences algorithm for PIM-based spectral matching
- A preliminary approach to integrate PIM into mass spectrometry workflows
- Early performance assessments via PIM simulation
- Evaluation of algorithm robustness under noise and low precision representations

DRAM + Processing-in-memory (PIM)

- Processing-in-memory (PIM)**, a.k.a. **near-data-processing (NDP)** reduces data movement by adding compute capability to memory

PIM types

1. Processing-using-memory (PUM):


- Use inherent circuit-level properties of memory cells to perform logic within memory arrays

2. Processing-near-memory (PNM):

- 3D-stacked memory with a logic layer
- Take advantage of the high bandwidth communication over vertical interconnects between layers


DRAM for PIM

- DRAM-based PIM can benefit from simultaneous access to multiple DRAM arrays for parallel data processing




Mass Spectrometry (MS) and Spectral Library Search

- Mass spectrometry (MS)** analyzes the proteome of protein samples



- Spectral library search** is used to interpret MS/MS data



Minimum Sum of Absolute Differences

- DSP-inspired algorithm: sum of absolute differences
- Calculate a similarity score between the experimental spectrum $S = (s_1, s_2, \dots, s_n)$ and each library spectrum $L_p = (\ell_{p1}, \ell_{p2}, \dots, \ell_{pn})$ for $p=1, 2, \dots, m$.
- Assuming m library spectra and n elements in each spectrum, the basic algorithm seeks L_k such that:

$$k = \arg \min_p \sum_{j=1}^n |s_j - \ell_{pj}|$$

Pre-processing and Pre-filtering

Discretization


- Divide the m/z range of spectra into bins and sum intensity values in the same bin

Approximation

- Fixed-point for floating-point values with reduced bit-precision enables program execution on PIM hardware


Key insights:

- Pre-filter using PIM to decrease the search space of spectral library search
- Co-design application with hardware
- Low-precision approach aligns with MS error tolerance



PIM Simulation

```
for (int idx=0; idx < subvecLen; idx++) {  
    pimSub(obj1, obj2, obj3);  
    pimAbs(obj3, obj3);  
    for (int i=idx; i+subvecLen-1 < vecLen; i+=subvecLen) {  
        pimRedSumRanged(obj3, i, i+subvecLen-1, &sumAbsDiff);  
    }  
    pimRotateElementsRight(obj2);  
}
```



- PIMEval** is a performance and energy simulation framework for various **DRAM-based PIM** architectures: subarray-level bit-serial, subarray-level bit-parallel, bank-level
- We model **subarray-level bit-serial PIM** with single/dual rank DIMM with 8 chips, 16 banks per chip, and 32 subarrays per bank of 8192×8192 cells

Key insights:

- PIM performance benefits scale with increasing problem size; small datasets favor CPU due to lower overhead

Algorithm Robustness Evaluation




TABLE I: Algorithm Robustness under Varying Degradation

Experiment Condition	Degradation (Rem/Jit/Noise)	AUC (32-bit)	AUC (16-bit)	AUC (8-bit)
No-Noise Control	0% / ± 0 / 0%	1.0000	1.0000	1.0000
Moderate Noise	10% / ± 1 / 10%	0.9892	0.9681	0.9732
High-Noise Test	25% / ± 2 / 25%	0.7819	0.7583	0.7606

- Experiments show algorithm is **robust** in handling **noise, variations, and degradations** inherent to real-world bioinformatics workflows
- 8-bit precision delivers accuracy comparable to that of higher bit-precisions

Conclusion and Next Steps

- Proof of concept for hardware-algorithm co-designed pipeline for DRAM-based PIM architecture acceleration of spectral library search
- Next steps are to estimate analytical scaling results, observe performance on larger workloads in simulation, and use larger datasets.

References

- O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "A modern primer on processing in memory," in Emerging Computing: From Devices to Systems. Springer, 2023, pp. 171–243.
- J. Kang, W. Xu, W. Bittremieux, N. Moshiri, and T. Rosing, "Accelerating open modification spectral library searching on tensor core in high-dimensional space," Bioinformatics, vol. 39, no. 7, p. btad404, 2023.
- V. R. Pagala, A. A. High, X. Wang, H. Tan, K. Kodali, A. Mishra, K. Kavdia, Y. Xu, Z. Wu, and J. Peng, "Quantitative protein analysis by mass spectrometry," Protein-protein interactions: Methods and applications, pp. 281–305, 2015.
- J. M. Chick, D. Kolipakkam, D. P. Nusinow, B. Zhai, R. Rad, E. L. Huttlin, and S. P. Gygi, "A mass-tolerant database search identifies a large proportion of unassigned spectra in shotgun proteomics as modified peptides," Nature biotechnology, vol. 33, no. 7, pp. 743–749, 2015.
- F. A. Siddique, D. Guo, Z. Fan, M. Gholamrezaei, M. Baradaran, A. Ahmed, H. Abbot, K. Durrer, K. Nandagopal, E. Ermovick et al., "Architectural modeling and benchmarking for digital dram pim," <https://www.thermofisher.com/us/en/home/industrial/mass-spectrometry/liquid-chromatography-mass-spectrometry-lc-ms/lc-ms-systems.html>

Acknowledgements

This work was supported in part by the Jump 2.0 PRISM center, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

Carnegie Mellon University



Semiconductor Research Corporation

