

# Leveraging Recurrent Patterns in Graph Accelerators

Masoud Rahimi, Sébastien Le Beux

Department of Electrical and Computer Engineering, Concordia University, Montreal, Canada  
Email: masoud.rahimi@mail.concordia.ca, sebastien.lebeux@concordia.ca

**Abstract**— Graph accelerators have emerged as a promising solution for processing large-scale sparse graphs, leveraging the in-situ computation of ReRAM-based crossbars to maximize computational efficiency. However, existing designs suffer from memristor access overhead due to the large number of graph partitions. This leads to increased execution time, higher energy consumption, and reduced circuit lifetime. This paper proposes a graph processing method that minimizes memristor write operations by identifying frequent subgraph patterns and assigning them to graph engines, referred to as *static*, allowing most subgraphs to be processed without a need for crossbar reconfiguration. Experimental results show speed up to  $2.38\times$  speedup and  $7.23\times$  energy savings compared to state-of-the-art accelerators. Furthermore, our method extends the circuit lifetime by  $2\times$  compared to state-of-the-art ReRAM graph accelerators.

**Index Terms**—Processing-in-Memory, ReRAM, graph processing, accelerator.

## I. INTRODUCTION

Graphs are a fundamental tool for modeling and analyzing relationships across various application domains, such as social networks [1], recommendation systems [2], [3], computer networks, and bioinformatics [4]. For instance, in social networks, users are modeled as nodes and their connections as edges to analyze interactions and influence. Real-world graphs often contain millions of nodes [5], [6], [7], are typically sparse and characterized by unstructured connectivity. Processing such graphs introduces several challenges, including: i) large memory capacity requirement, ii) random memory accesses caused by irregular connections, and iii) complexity of processing irregular graph structures.

To address these challenges, memory-centric architectures have been proposed to accelerate graph processing [8]. These architectures typically utilize in- or near-memory processing paradigms to reduce memory access [9]. Given the large size and inherent sparsity of graphs, conventional accelerators partition graphs into smaller subgraphs to satisfy limited capacity of local memories [10]. These subgraphs are then mapped onto ReRAM-based crossbars, where computation is performed directly within crossbars [10], [11], [12]. In these methods, subgraph data are stored in adjacency matrix format. To mitigate inefficiencies introduced by large-scale and sparsity of real-world graphs, several methods have been proposed to improve performance. For instance, ReGraphX [13] utilizes relatively small crossbars (e.g.,  $8\times 8$ ) for edge-centric processing in Graph Neural Networks (GNNs), reducing storage of zero entries. Spara [14] introduces a graph reordering and vertex remapping scheme in preprocessing to maximize crossbar utilization. Alternatively, SparseMEM [15] proposes a hierarchical mapping scheme that mitigates local crossbar underutilization and enables core graph computations on compressed graph representations. ReFlip [12] introduces a flipped-mapping

scheme designed to increase parallelism of Matrix-Vector-Multiplication (MVM) operations in crossbars. In this design, feature vectors of a graph are mapped onto the crossbars, while edge data are fed as inputs. Overall, existing methods mainly address graph sparsity to improve graph accelerator performance by reducing memory constraint requirements.

A pattern is defined by the adjacency matrix of a graph, where ‘1’ represents an edge between two vertices, and ‘0’ indicates the absence of an edge. After partitioning a graph, the pattern occurrence refers to the number of subgraphs sharing an identical pattern. Figure 1-a reports pattern occurrence in Wiki-Vote [5], an unweighted graph with 7K vertices and 104K edges. The data are obtained by applying a  $4\times 4$  non-overlapping sliding window to the adjacency matrix, a commonly used partitioning approach [13], [16]. Pattern with all ‘0’ is discarded since it does not involve any processing. As reported in the figure, the most frequent pattern  $P_0$  accounts for 5.9% of subgraphs and, altogether, the 16 most frequent ( $P_0$  to  $P_{15}$ ) account for 86% of subgraphs. In contrast, remaining patterns ( $P_{16}$  to  $P_{809}$ ) collectively account for only 14% of subgraphs. Similar distributions were observed across all studied datasets [5], [6]. By leveraging the observation that few patterns allow processing most subgraphs, graph accelerators can be designed with minimal ReRAM write operations, that are known to be energy consuming, slow, and impact memory cells lifetime.




Figure 1: a) Pattern occurrence in Wiki-vote [5] partitioned with a  $4\times 4$  sliding window on the graph’s adjacency matrix. b) Key observations.

In this paper, for the first time, we leverage pattern occurrence to design graph accelerators. To this end, we develop a method that identifies and ranks patterns based on their frequency in the input graph. The most frequent patterns are then statically mapped to graph engines, which leads to a drastic reduction of memory crossbar writes. The proposed method is fully automated and consists of: i) a window-based partitioning algorithm to preprocess the input graph, ii) a scheduling algorithm capable of executing classical graph algorithms while handling graph engine configuration at run-time, and iii) a design space exploration framework that identifies optimal architectural parameters. Experimental results demonstrate that our architecture achieves up to  $2.38\times$

speedup and  $7.23\times$  energy savings compared to state-of-the-art accelerators, while also increasing circuit lifetime by up to 2x thanks to reduced ReRAM writes.

The rest of the paper is organized as: Section II presents background and related work. Section III presents the proposed method. Section IV presents our experimental setup and results. Finally, Section V concludes the paper.

## II. BACKGROUND AND RELATED WORK

### A. ReRAM-based Crossbars

Resistive Random Access Memory (ReRAM) is a non-volatile emerging memory based on a metal–oxide–metal structure where data are stored by altering the resistance states of individual cells [17]. ReRAM offers fast read speeds and high storage density compared to other non-volatile memories [18], [19]. In contrast, writes are slower than reads and consume more energy [20], [21]. Moreover, ReRAM devices have limited endurance, between  $10^5$  to  $10^8$  write cycles [22], [23]. ReRAM cells are arranged in crossbar arrays [24], enabling in-situ MVM. Each bitline performs a MAC operation expressed as  $\sum_i G_{i,j} \cdot V_i$ , where  $G_{i,j}$  denotes the conductance and  $V_i$  is the input voltage applied to wordlines [10], [25]. Therefore, a crossbar efficiently performs MVM in  $O(1)$  computational complexity.

### B. Graph Partitioning and Mapping to ReRAM

Realistic graph applications typically involve millions of vertices and edges [5], [6], [7]. Processing such large graphs is challenging due to memory capacity limitations. To address this, the graph is partitioned into smaller subgraphs using various methods. For instance, the method detailed in [26] involves duplicating vertices across multiple subgraphs, and distributing edges among them. This allows precise control over the number of vertices within each subgraph. Vertex duplication is particularly effective for compressed graph representations, where only edges are stored. Window-based partitioning is another method for graphs represented in adjacency matrix format. In this approach, a non-overlapping sliding window is applied across the matrix to divide it into smaller submatrices and are directly mapped onto ReRAM-based crossbars [11]. Since submatrices containing only zeros are excluded, using a small window size can significantly reduce storage requirements for sparse graphs [10]. Our method relies on window-based partitioning, using small windows (e.g.,  $4\times 4$ ) to maximize zero submatrices, reducing memory footprint.

To store graph data in crossbars, both compressed and uncompressed formats have been explored. For instance, Spara [14] and Janus [27] utilize Compressed Sparse Row (CSR) to improve memory utilization. This representation enables efficient access to edges of active vertices, allowing simultaneous traversal from a common source. In contrast, GraphR [10] and Remagn [11] store graphs using uncompressed adjacency matrix, enabling in-situ MVM on ReRAM crossbars. In our design, we store input graphs using Coordinate list (COO) format. This ensures efficient storage and sequential edge access, while utilizing adjacency matrix format in local memory to enable in-memory processing on ReRAM.

### C. Graph Accelerators

Prior studies have leveraged in-situ MVM, a fundamental operation in many graph algorithms to accelerate graph processing [10], [12], [13]. For example, GraphR [10] achieves up to two orders of magnitude speedup over a CPU baseline for classical graph algorithms. ReGraphX [13] proposes heterogenous crossbars for on-chip GNN training. However, large-size and sparsity of real-world graphs combined with limited number of graph engines necessitate frequent memory access (i.e., read and write operations) in ReRAM-based crossbars. Consequently, this leads to i) increased execution time, ii) high energy consumption, and iii) reduced circuit lifetime due to limited ReRAM endurance.

Recent works mitigate inefficiencies caused by graph sparsity, as summarized in Table 1. For instance, GraphIte [28] applies graph sparsification by selectively removing edges with minimal impact on vertex updates at the expense of some accuracy loss. Alternatively, SparseMEM [15] proposes a compressed hierarchical mapping analogous to CSR format. In this design, destination vertices and associated weights are stored sequentially within one crossbar, while vertex locations are stored in a separate crossbar. This maximizes crossbar utilization by eliminating zero-valued cells but precludes in-situ MVM operations and requires high-resolution Multi-Level Cell (MLC) ReRAM to store vertex indices. TARE [16] introduces a write-free mapping strategy that partitions a ReRAM crossbar into multiple smaller computing blocks (CBs), each preconfigured with complete sets of possible binary submatrices. While this eliminates runtime write operations, it restricts parallel MVM operations and incurs frequent off-chip memory reads, degrading performance. In contrast, our design minimizes ReRAM writes without excessive off-chip memory access and operates on 1-bit ReRAM crossbars.

Table 1: Comparison of existing graph accelerators

Reference	In-engine Graph Representation	Memory Access (Read/Write)	Required MLC <sup>1</sup> ReRAM	Graph Algorithm
GraphR [10]	Adjacency	High/High	4-bit	Classical
ReFlip [12]	Compressed	High/Low	Variable <sup>2</sup>	GNN
SparseMEM [15]	Compressed	Low/Low	Variable <sup>3</sup>	Classical
TARE [16]	Adjacency	High/Low	1-bit	GNN
Proposed	Adjacency	Low/Low	1-bit	Classical

1: Multi-Level Cell. 2: Depends on the size of input feature vectors. 3: Depends on the number of vertices.

## III. PROPOSED METHOD

In this section, we first present the design flow used to configure the proposed architecture. We then introduce the preprocessing model, which analyzes the input graph and extracts structural patterns. Next, we present the graph processing algorithm. Finally, we detail the graph engine architecture.

### A. Design Flow

Figure 2 presents our design flow, which takes as inputs a graph and an architecture model. The architecture is generic and is characterized by the following parameters: crossbar size ( $C$ ), total number of graph engines ( $T$ ), number of static graph engines ( $N$ ), and number of crossbars per graph engine ( $M$ ).




Figure 2: Design flow and architecture overview: a) input graph preprocessed based on architectural parameters; b) architecture initialized; and c) preconfigured architecture with static and dynamic graph engines.

During preprocessing (shown in Figure 2 upper part); ① the graph is partitioned into subgraphs sized by crossbar size. Patterns are then identified per subgraph ② and ranked by frequency across the graph ③. In ④, the most frequent patterns are assigned to static graph engines, which are configured once. The least frequent patterns are assigned to dynamic graph engines, reconfigured at runtime. Section III.B details our preprocessing algorithm. The architecture (depicted at the bottom of Figure 2) is composed of multiple graph engines to enable parallel processing. To maximize graph engine activations, most engines are implemented as static and configured once during the initialization phase. In contrast, dynamic engines are reconfigured as needed at runtime.

### B. Graph Preprocessing Algorithm

Algorithm 1 outlines preprocessing where the inputs are graph  $G$  and architectural parameters. Outputs are configuration table  $CT$ , and subgraph table  $ST$ . The algorithm: i) partition input graphs into smaller subgraphs based on the crossbar size  $C$  (Line 4), ii) extract patterns from subgraphs and rank them by frequency (Lines 5–12). iii) assign most frequent patterns to static engines and map remaining patterns to dynamic engines (Lines 13–19).

Figure 3 illustrates preprocessing of a graph with six vertices using three graph engines, each with a single  $2 \times 2$  crossbar. The graph is converted to an adjacency matrix, partitioned into nine subgraphs ( $S_0$  to  $S_8$ ), then patterns are extracted (Figure 3-b) and ranked by frequency (Figure 3-c). For example, pattern  $P_0$  has the highest frequency, appearing in  $S_0$ ,  $S_4$ , and  $S_6$ . In contrast,  $P_2$  and  $P_3$  each appear once, in  $S_3$  and  $S_5$ , respectively. As mentioned earlier, subgraphs with no edges ( $S_5$  and  $S_8$ ), corresponding to patterns consisting entirely of zeros, are excluded from the ranking.

Following pattern identification and ranking, patterns are mapped to graph engines based on their frequency (Figure 3-d). Most frequent patterns,  $P_0$  and  $P_1$ , are assigned to static engines

$GE_0$  and  $GE_1$ , respectively, while less frequent patterns,  $P_2$  and  $P_3$ , are allocated to the dynamic engine  $GE_2$ . Finally, patterns and subgraphs are stored in the main memory (Figure 3-e). The configuration table stores the pattern data, represented in COO format, along with their associated graph engines. The subgraph table stores vertex data for each subgraph along with its corresponding pattern. Since all subgraphs have same number of vertices, only the starting source and destination vertices are recorded. For instance, for  $S_1$ , only vertices  $V_0$  and  $V_2$  are stored, thereby reducing storage overhead in main memory.

---

#### Algorithm 1: Preprocessing

**Input:** Graph  $G$ , crossbar size  $C$ , No. of static graph engines  $N$ , No. of crossbars per graph engine  $M$

**Output:** Graph engine configuration table  $CT$ , subgraph table  $ST$

```

1 Let Partitioning(Graph, C): partition Graph into subgraphs with C vertices.
2 Procedure Preprocess( $G, C, N, M$ ):
3   patterns  $\leftarrow \{\}$ 
4   // partition the input graph
   subgraphs  $\leftarrow$  Partitioning( $G, C$ )
5   // identify & rank patterns
   foreach  $s \in$  subgraphs do
6      $p \leftarrow$  GetPattern( $s$ )
7     // increment frequency for an existing pattern
     if  $p \in$  patterns then
8       patterns[ $p$ ]  $\leftarrow$  patterns[ $p$ ] + 1
9     else
10      patterns.add( $p, 1$ )
11      ST.add( $s, p$ )
12   // sort patterns by frequency
   patterns.sort()
13   // assign patterns to graph engines
    $i \leftarrow 0$ 
14   foreach  $p \in$  patterns do
15     if  $i < N * M$  then
16       CT.add( $p, "static"$ )
17        $i \leftarrow i + 1$ 
18     else
19       CT.add( $p, "dynamic"$ )
20   return CT, ST

```

---

In case graph engines contain multiple crossbars (i.e.,  $M > 1$ ), patterns assigned to static engines are evenly distributed across their crossbars (function *FindGE* in algorithm 1). This distribution balances pattern load among static engines, improving overall utilization. Furthermore, since patterns with a single edge are more frequent (due to power-law degree distribution [29]), the row address is also stored in the configuration table. This eliminates iteration over all crossbar rows, thereby reducing ReRAM reads in static engines.

### C. Graph Processing

For runtime processing, we use the streaming-apply execution model detailed in [10]. In this model, subgraphs sharing same source vertices (row-major order) or destination vertices (column-major order) are processed together within an iteration. The table in Figure 3-e shows subgraphs in column-major order. For example,  $S_0$ ,  $S_3$ , and  $S_6$  are grouped together since they share same destination vertices  $V_0$  and  $V_1$ . Our design supports both row- and column-major execution models, which can be selected for different graph algorithms. Furthermore, execution order of subgraphs can be further optimized to maximize engine activations, which we leave as future work.

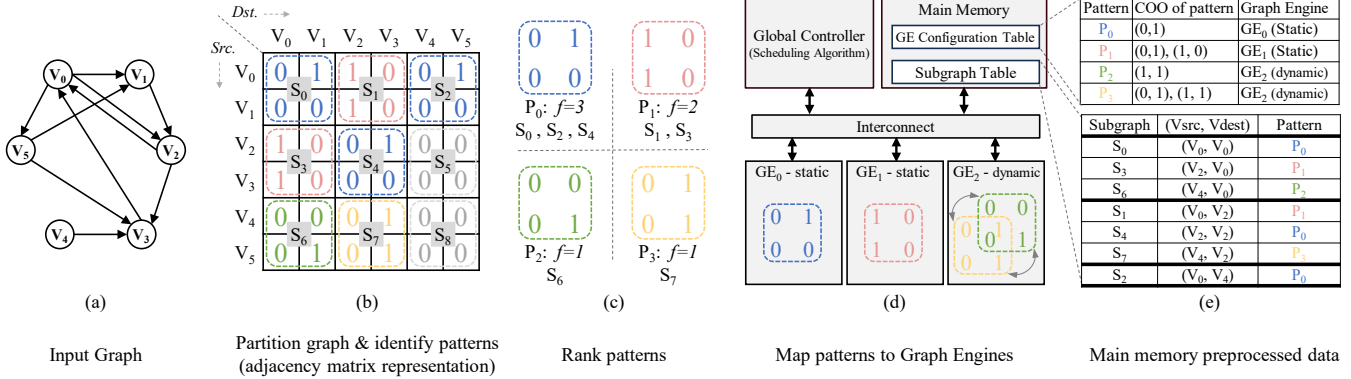


Figure 3: Graph preprocessing: a) Input graph, b) partition and identify patterns in adjacency matrix representation, c) rank patterns, d) map to two static and one dynamic graph engines, and e) graph engine configuration and subgraph tables (column-major order) in main memory.

Algorithm 2 details our graph processing. Static graph engines are pre-configured with assigned patterns. Subsequently, subgraphs are processed over multiple iterations, where in each iteration, a batch (a group of subgraphs) is fetched from the subgraph table ( $ST$ ) (Line 9) and executed in parallel (Line 10). Batch size is determined by available graph engines. For each subgraph in a batch, if the pattern is assigned to a static engine, only vertex data is transferred (Line 12). Otherwise, a dynamic engine is selected based on the replacement policy ( $FindGE$  function) and then re-configured with the corresponding pattern (Line 14-15). Once all subgraphs in a batch are processed ( $Allocate$  function), the updated vertex data is aggregated (Line 17).

#### Algorithm 2: Graph Processing & Scheduling

**Input:** Graph engine configuration table  $CT$ , subgraph table  $ST$

**Output:** Processed Vertex Data  $Results$

```

1 Let Configure(GE, Config) configure graph engine GE with Config.
2 Let Allocate(GE, VertexData) allocate VertexData to graph engine GE.
3 Let Aggregate(VertexData) aggregate VertexData from the graph engines.
4 Procedure Process(CT, ST):
5   Results ← {}
6   // initialization: configure static GEs
7   foreach pattern p ∈ CT do
8     if p is static then
9       Configure(p.GE, p.data)
10  // iteratively process subgraphs
11  foreach batch of subgraphs with same dest. vertices ∈ ST do
12    parallelforeach subgraph s, pattern p ∈ batch do
13      if p is static then
14        // process subgraph with initial vertex data
15        r ← Allocate(p.GE, s.vertex_data)
16        // process subgraph in dynamic GE
17      else
18        ge ← findGE(p.data)
19        Configure(ge, p.data)
20        r ← Allocate(ge, s.vertex_data)
21        // aggregate vertex data from GEs
22        Results ← Aggregate(r)
23  return Results

```

#### D. Graph Engine

Figure 4 illustrates the architecture, with two ReRAM crossbars:  $CB_0$  (static) and  $CB_1$  (dynamic). Graph execution is managed by the control unit, which handles crossbar configuration, data movement, and scheduling. The input register holds initial vertex data, while the output register stores the processed results. Analog

outputs from crossbar bitlines are sampled by sample-and-hold (S/H) circuits and digitized by analog-to-digital converters (ADCs), which are shared across multiple bitlines to reduce area and power overhead. The driver applies input voltages during MVM operations and configures crossbars as needed. Computations not natively supported by in-situ MVM are offloaded to a lightweight arithmetic logic unit (ALU). Note that the architecture shown is generic, however our method requires only static or dynamic crossbars within an engine. Furthermore, we assume same crossbar size for both static and dynamic engines. We leave exploration of non-uniform crossbar configurations to future work.

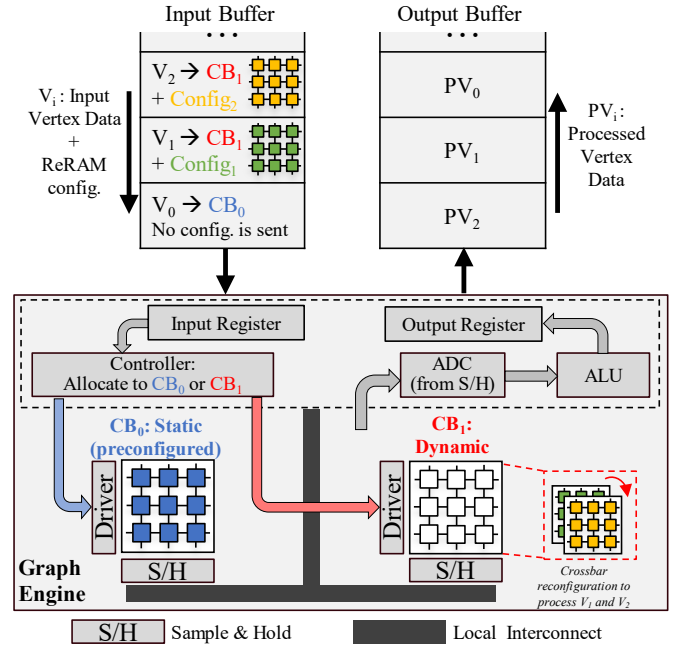


Figure 4: Graph engine architecture composed of two crossbars.

At runtime, the vertex data ( $V_i$ ) of a given subgraph is loaded from the input buffer. For dynamic crossbars (e.g.,  $CB_1$  in Figure 4), the crossbar configuration (edge data labeled as  $Config_i$ ) is also transferred via the input buffer. In contrast, for static crossbars (e.g.,  $CB_0$ ), no configuration data is required. After processing,

updated vertex data (PV<sub>i</sub>) is written back through an output buffer. Both input and output buffers operate as FIFOs, where each input entry is paired with a corresponding output entry, enabling pipeline processing of multiple subgraphs.

Our architecture supports a range of graph algorithms such as Breadth-First Search (BFS), Single-Source Shortest Path (SSSP), and PageRank that follow the vertex programming model described in [10]. In this model, graph processing is divided into two phases: i) *edge computation*, where a value is computed for each outgoing edge of an active vertex using in-situ MVM operation of ReRAM crossbars; and ii) *reduce and apply*, where each vertex property (i.e., vertex data) is updated by applying a reduction function over all incoming edge values using the ALU.

#### IV. EVALUATION

In this section, we first present the experimental setup used to evaluate our design. We then analyze activity of graph engines to demonstrate the functionality of our method and identify optimal architectural parameters. Finally, we compare our design performance against state-of-the-art accelerators.

##### A. Experimental Setup

We develop a system-level simulator to evaluate the design performance. It estimates the execution time and energy consumption by monitoring memory access performed by the graph engines during processing. To determine required number of clock cycles, the simulator models the parallel operation of graph engines. Several preprocessing parameters are provided to the simulator, including crossbar size, total number of graph engines, number of static and dynamic engines, and number of crossbars per engine.

We evaluate our design using datasets sourced from various domains, as listed in Table 2 [5]. These benchmarks are graph datasets with varying sizes and sparsity levels, enabling a comprehensive analysis. The average degree of a graph represents the average number of edges per vertex, indicating the graph’s connectivity level. For benchmarking, since input graphs are unweighted, we use BFS as the baseline graph algorithm. We also use column-major order as our baseline execution model [10].

Table 2: Graph datasets

Name	#vertices	#edges	Average Deg.	Sparsity (%)	Domain
web-Google (WG)	875K	5.1M	12	99.999	Web
Amazon302 (AZ)	262K	1.2M	9	99.998	Recom.
Slashdot0902 (SD)	82K	948K	23	99.985	Social
soc-Epinions1 (EP)	76K	509K	13	99.991	Social
p2p-gnutella31 (PG)	5K	148K	5	99.996	Network
Wiki-vote (WV)	7K	104K	29	99.795	Social

Unless stated otherwise, we assume 32 graph engines containing 4×4 crossbars. We consider single-level cell (i.e., 1-bit) ReRAM crossbars, modeled in 32nm node technology, with capacity of 32KB and 8 bits data width. ReRAM parameters are obtained from NVSim [30]. We use CACTI-6.5 [31] at 32nm to simulate I/O buffers and main memory. Furthermore, ADC parameters are

derived from [32], assuming a 32 nm technology node with 8-bit resolution. Table 3 summarizes the simulation specifications.

Table 3: Simulation specifications

		Latency	Energy
4x4 ReRAM crossbar 32KB $V_{SET} = V_{RESET} = 2V$	Per bit Read	1.3ns	1.1pJ
	Per bit Write	20.2ns	4.9pJ
	Sense Amplifier	1ns	1pJ
SRAM buffer 32KB	Per access	0.31ns	29pJ
ADC 8-bit resolution	Per access	1ns	2pJ

In our evaluation, we consider both graph engine performance and main memory access. Since baseline designs perform better with large crossbars, we assume a 128×128 crossbar configuration with the same capacity and peripheral circuitry across designs. To compare with graph neural network designs such as TARE [16], we consider only its mapping scheme and adapt it for classical graph algorithms.

##### B. Graph Engine Activity

Analyzing graph engine activity characterizes resource utilization, workload distribution, and potential performance bottlenecks. For this purpose, we assume a design with 6 graph engines including 4 static and 2 dynamic, each containing 4 crossbars. Figure 5 illustrates the crossbar read and write activity during processing of Wiki-Vote. Activity levels (scale of 0 to 100) are computed by aggregating number of crossbar reads/writes over a sliding window of iterations. We observe non-uniform read and write activity across iterations. This is caused by the execution model, where the number of subgraphs processed in parallel varies across iterations. Analysis of other datasets shows this behavior is strongly graph-dependent. Since most subgraphs are processed by static engines, their read activity is significantly higher than dynamic engines. This variation motivates us to explore workload balance between static and dynamic engines to achieve higher speedup.




Figure 5: Graph Engine (GE) activity during processing of Wiki-Vote. GE<sub>1</sub> to GE<sub>4</sub> are static, while GE<sub>5</sub> and GE<sub>6</sub> are dynamic. Activity levels range from 0 (no activity) to 100 (maximum activity).

We now study the trade-off between number of static and dynamic graph engines, keeping the total number fixed. Figure 6 shows speedup achieved for different allocations of static engines, normalized to the baseline configuration with no static engines.

For simplicity we assume a single crossbar per engine. For the sake of brevity, results are shown for three representative datasets. Across all three datasets, assigning 16 static engines results in highest speedup. For instance, in EP, utilizing 16 static engines achieves a 1.6 $\times$  speedup. Given a 4 $\times$ 4 crossbar, most subgraphs contain only a single edge due to the power-law degree distribution [29]. Therefore, at least 16 static engines are required to ensure all such subgraphs are processed by static engines. In contrast, assigning most engines as static leads to longer execution time because few dynamic engines limit parallelism and increase the iteration count. We observed a similar trend across other datasets. In the following, we use the activity profile to quantify total memory access and evaluate system performance.

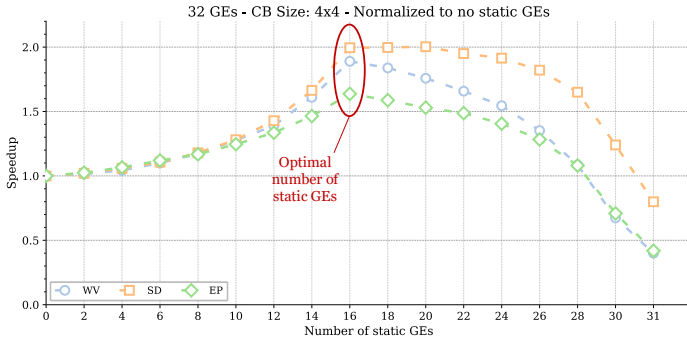


Figure 6: Speedup wrt. to the number of static graph engines. A total of 32 graph engines, each containing a 4 $\times$ 4 crossbar. Results are normalized to a system with no static engines.

### C. Performance Comparison

In the following, we compare energy consumption and speedup of our design with three baseline approaches, GraphR [10], SparseMEM [15], and TARE [16]. For fair comparison, we assume equal number of graph engines and identical memory capacity across all designs.

1) *Energy*: Table 4 reports the energy consumption of our proposed design and state-of-the-art methods for executing the BFS algorithm on each dataset. Overall, our design is three orders of magnitude more energy efficient than GraphR and achieves 7.23 $\times$  and 2.3 $\times$  energy efficiency compared to SparseMEM and TARE, respectively. The improvement over GraphR and SparseMEM is explained by fewer ReRAM writes. For TARE, although the write-free engine suppresses ReRAM writes, frequent off-chip memory accesses lead to higher overall energy consumption.

Table 4: Total energy consumption of BFS across evaluated datasets.

Dataset	GraphR [10]	SparseMEM [15]	TARE [16]	Proposed
WV	3.3 mJ	23 $\mu$ J	24 $\mu$ J	5.9 $\mu$ J
PG	60 mJ	55 $\mu$ J	30 $\mu$ J	7.1 $\mu$ J
EP	53 mJ	182 $\mu$ J	35 $\mu$ J	26 $\mu$ J
SD	110 mJ	260 $\mu$ J	50 $\mu$ J	48 $\mu$ J
AZ	460 mJ	688 $\mu$ J	79 $\mu$ J	54 $\mu$ J
WG	4.1 J	2.12 mJ	470 $\mu$ J	318 $\mu$ J

2) *Speedup*: Figure 7 compares speedup of the four designs. On average, our design achieves three orders of magnitude higher

speedup than GraphR, and 2.38 $\times$  and 1.27 $\times$  over SparseMEM and TARE, respectively. The improvement over GraphR results from significantly reduced crossbar access, as GraphR performance is constrained by sparse subgraph mapping. Although SparseMEM achieves better crossbar utilization, its execution time is higher due to decompression of graph data in graph engines. Finally, while TARE minimizes ReRAM writes by configuring crossbars only once, it limits parallel execution of MVM operations, resulting in more iterations.

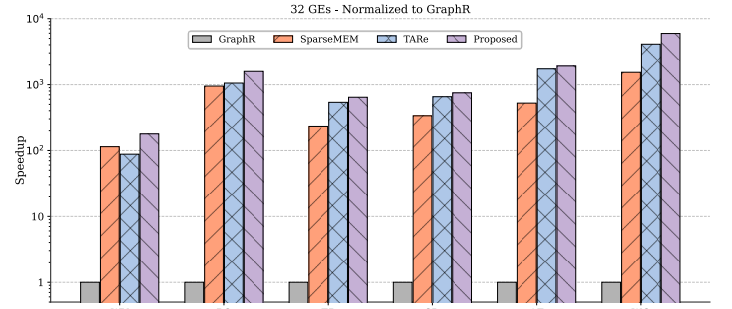


Figure 7: Speedup of the proposed design compared to state-of-the-art. Results are obtained for BFS algorithm and normalized to GraphR.

### D. Circuit Lifetime Analysis

ReRAM cells have limited write endurance ( $\sim 10^8$  cycles [23]), and exceeding this limit can cause write-induced faults [17], [33]. We model circuit lifetime as  $\frac{E}{w} \times T$ , where  $E$  is cell endurance,  $w$  is the maximum number of write operations per cell, and  $T$  is the execution interval. We assume graph engines are not used once a crossbar reaches maximum writes, allowing remaining engines to continue operation. Furthermore, static engines are excluded, as they are configured only once. Considering 128 graph engines and executing Wiki-Vote once per hour ( $T = 1h$ ), our design can operate for over 10 years, which is two orders of magnitude longer than GraphR and 2 $\times$  longer than SparseMEM. The improvements are due to reduced number of ReRAM writes in our design.

## V. CONCLUSION

In this paper, we presented a graph accelerator that leverages recurring graph patterns to mitigate frequent memory access. We also propose a method to find the best number of static graph engines for a given application and to explore design tradeoffs. Our results show the proposed accelerator achieves 1.27 $\times$  speedup and 2.3 $\times$  reduction in energy consumption compared to state-of-the-art. As opposed to state-of-the-art graph accelerators that require expensive and error-prone, large ReRAM crossbars, our architecture performs better with smaller, cost-effective crossbars, e.g., 4 $\times$ 4 or 8 $\times$ 8. Our architecture thus has potential for deployment not only in HPC data centers, where graph accelerators are most needed, but also in high-performance embedded systems. Furthermore, static graph engines intrinsically do not require any ReRAM write operations during application execution, which significantly increases architecture lifetime. By leveraging graph remapping on graph engines, we plan to further enhance architecture reliability, which we will investigate in our future work.

## VI. REFERENCES

- [1] E. Agichtein, C. Castillo, D. Donato, A. Gionis, and G. Mishne, "Finding High-Quality Content in Social Media," Feb. 2008, pp. 183–194. doi: 10.1145/1341531.1341557.
- [2] G. Linden, B. Smith, and J. York, "Amazon.com recommendations: item-to-item collaborative filtering," *IEEE Internet Computing*, vol. 7, no. 1, pp. 76–80, 2003, doi: 10.1109/MIC.2003.1167344.
- [3] J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen, "Collaborative Filtering Recommender Systems," in *The Adaptive Web: Methods and Strategies of Web Personalization*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 291–324. doi: 10.1007/978-3-540-72079-9\_9.
- [4] A. Conesa, S. Götz, J. M. García-Gómez, J. Terol, M. Talón, and M. Robles, "Blast2GO: a universal tool for annotation, visualization and analysis in functional genomics research," *Bioinformatics*, vol. 21, no. 18, pp. 3674–3676, Aug. 2005, doi: 10.1093/bioinformatics/bti610.
- [5] J. Leskovec and A. Krevl, "SNAP Datasets: Stanford Large Network Dataset Collection." Jun. 2014. [Online]. Available: <http://snap.stanford.edu/data>
- [6] R. A. Rossi and N. K. Ahmed, "The Network Data Repository with Interactive Graph Analytics and Visualization," in *AAAI*, 2015. [Online]. Available: <https://networkrepository.com>
- [7] J. Kunegis, "KONECT: the Koblenz network collection," in *WWW Companion - Proceedings of the 22nd International Conference on World Wide Web*, May 2013, pp. 1343–1350. doi: 10.1145/2487788.2488173.
- [8] C.-Y. Gui *et al.*, "A Survey on Graph Processing Accelerators: Challenges and Opportunities," *Journal of Computer Science and Technology*, vol. 34, no. 2, pp. 339–371, Mar. 2019, doi: 10.1007/s11390-019-1914-z.
- [9] C. Zhang, H. Sun, S. Li, Y. Wang, H. Chen, and H. Liu, "A Survey of Memory-Centric Energy Efficient Computer Architecture," *IEEE Transactions on Parallel and Distributed Systems*, vol. 34, no. 10, pp. 2657–2670, 2023, doi: 10.1109/TPDS.2023.3297595.
- [10] L. Song, Y. Zhuo, X. Qian, H. Li, and Y. Chen, "GraphR: Accelerating Graph Processing Using ReRAM," in *IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2018, pp. 531–543. doi: 10.1109/HPCA.2018.00052.
- [11] A. I. Arka, B. K. Joardar, J. R. Doppa, P. P. Pande, and K. Chakrabarty, "Performance and Accuracy Tradeoffs for Training Graph Neural Networks on ReRAM-Based Architectures," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 29, no. 10, pp. 1743–1756, 2021, doi: 10.1109/TVLSI.2021.3110721.
- [12] Y. Huang *et al.*, "Accelerating Graph Convolutional Networks Using Crossbar-based Processing-In-Memory Architectures," in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2022, pp. 1029–1042. doi: 10.1109/HPCA53966.2022.00079.
- [13] A. I. Arka, J. R. Doppa, P. P. Pande, B. K. Joardar, and K. Chakrabarty, "ReGraphX: NoC-enabled 3D Heterogeneous ReRAM Architecture for Training Graph Neural Networks," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2021, pp. 1667–1672. doi: 10.23919/DATE51398.2021.9473949.
- [14] L. Zheng *et al.*, "Spara: An Energy-Efficient ReRAM-Based Accelerator for Sparse Graph Analytics Applications," in *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2020, pp. 696–707. doi: 10.1109/IPDPS47924.2020.00077.
- [15] M. Zahedi, G. Custers, T. Shahroodi, G. Gaydadjev, S. Wong, and S. Hamdioui, "SparseMEM: Energy-efficient Design for In-memory Sparse-based Graph Processing," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2023. doi: 10.23919/DATE56975.2023.10137303.
- [16] Y. He, B. Li, Y. Wang, C. Liu, H. Li, and X. Li, "A Task-Adaptive In-Situ ReRAM Computing for Graph Convolutional Networks," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 43, no. 9, pp. 2635–2646, 2024, doi: 10.1109/TCAD.2024.3375251.
- [17] H.-S. P. Wong *et al.*, "Metal-Oxide RRAM," *Proceedings of the IEEE*, vol. 100, no. 6, pp. 1951–1970, 2012, doi: 10.1109/JPROC.2012.2190369.
- [18] R. Fackenthal *et al.*, "19.7 A 16Gb ReRAM with 200MB/s write and 1GB/s read in 27nm technology," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 2014, pp. 338–339. doi: 10.1109/ISSCC.2014.6757460.
- [19] P. Jangra and M. Duhan, "Performance-based comparative study of existing and emerging non-volatile memories: a review," *Journal of Optics*, vol. 52, no. 4, pp. 2395–2409, Dec. 2023, doi: 10.1007/s12596-022-01058-w.
- [20] P. Chi *et al.*, "PRIME: A Novel Processing-in-Memory Architecture for Neural Network Computation in ReRAM-Based Main Memory," in *ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, 2016, pp. 27–39. doi: 10.1109/ISCA.2016.13.
- [21] A. Levy *et al.*, "EMBER: Efficient Multiple-Bits-Per-Cell Embedded RRAM Macro for High-Density Digital Storage," *IEEE Journal of Solid-State Circuits*, vol. 59, no. 7, pp. 2081–2092, 2024, doi: 10.1109/JSSC.2024.3387566.
- [22] E. Kondratyuk *et al.*, "Automated Testing Algorithm for the Improvement of 1T1R ReRAM Endurance," *IEEE Transactions on Electron Devices*, vol. 68, no. 10, pp. 4891–4896, 2021, doi: 10.1109/TED.2021.3105089.
- [23] H.-D. Kim, H.-M. An, Y. Seo, and T. G. Kim, "Transparent Resistive Switching Memory Using ITO/AlN/ITO Capacitors," *IEEE Electron Device Letters*, vol. 32, no. 8, pp. 1125–1127, 2011, doi: 10.1109/LED.2011.2158056.
- [24] C. Xu *et al.*, "Overcoming the challenges of crossbar resistive memory architectures," in *IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*, 2015, pp. 476–488. doi: 10.1109/HPCA.2015.7056056.
- [25] M. Hu *et al.*, "Dot-product engine for neuromorphic computing: Programming 1T1M crossbar to accelerate matrix-vector multiplication," in *53rd ACM/EDAC/IEEE Design Automation Conference (DAC)*, 2016, pp. 1–6. doi: 10.1145/2897937.2898010.
- [26] C. Nie, G. Chen, W. Zhang, and Z. He, "GIM: Versatile GNN Acceleration with Reconfigurable Processing-in-Memory," in *IEEE 41st International Conference on Computer Design (ICCD)*, 2023, pp. 499–506. doi: 10.1109/ICCD58817.2023.00083.
- [27] X. Li *et al.*, "Janus: A Flexible Processing-in-Memory Graph Accelerator Toward Sparsity," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 43, no. 12, pp. 4813–4826, 2024, doi: 10.1109/TCAD.2024.3405395.
- [28] D. Choudhury, A. Kalyanaraman, and P. Pande, "GraphIte: Accelerating Iterative Graph Algorithms on ReRAM Architectures via Approximate Computing," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2023. doi: 10.23919/DATE56975.2023.10137001.
- [29] W. Aiello, F. Chung, and L. Lu, "Random Evolution in Massive Graphs," in *Handbook of Massive Data Sets*, J. Abello, P. M. Pardalos, and M. G. C. Resende, Eds., Boston, MA: Springer US, 2002, pp. 97–122. doi: 10.1007/978-1-4615-0005-6\_4.
- [30] X. Dong, C. Xu, Y. Xie, and N. P. Jouppi, "NVSim: A Circuit-Level Performance, Energy, and Area Model for Emerging Nonvolatile Memory," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 31, no. 7, pp. 994–1007, 2012, doi: 10.1109/TCAD.2012.2185930.
- [31] N. Muralimanohar, R. Balasubramonian, and N. Jouppi, "Cacti 6.0: A tool to model large caches," *HP Laboratories*, Jan. 2009.
- [32] L. Kull *et al.*, "A 3.1 mW 8b 1.2 GS/s Single-Channel Asynchronous SAR ADC With Alternate Comparators for Enhanced Speed in 32 nm Digital SOI CMOS," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 12, pp. 3049–3058, 2013, doi: 10.1109/JSSC.2013.2279571.
- [33] C.-Y. Chen *et al.*, "RRAM Defect Modeling and Failure Analysis Based on March Test and a Novel Squeeze-Search Scheme," *IEEE Transactions on Computers*, vol. 64, no. 1, pp. 180–190, 2015, doi: 10.1109/TC.2014.12.