

UWB-GCN: Accelerating Graph Convolutional Networks through Runtime Workload Rebalancing

Tong Geng^{†‡}, Ang Li[‡], Tianqi Wang[†], Chunshu Wu[†], Yanfei Li[¶],
Antonino Tumeo[‡], Shuai Che[§], Steve Reinhardt[§] and Martin Herbordt[†]

[†]Boston University

[‡]Pacific Northwest National Laboratory

[§]Microsoft Research

[¶]Zhejiang University

Abstract—The recent development of deep learning has mostly been focusing on Euclidean data, such as images, videos, and audios. However, most real-world information and relationships are often expressed in graphs. Graph convolutional networks (GCNs) appear as a promising approach to efficiently learn from graph data structures, showing advantages in several practical applications such as social network analysis, knowledge discovery, 3D modeling, and motion capturing. However, practical graphs are often extremely large and unbalanced, posing significant performance demand and design challenges on the hardware dedicated to GCN inference.

In this paper, we propose an architecture design called Ultra-Workload-Balanced-GCN (UWB-GCN) to accelerate graph convolutional network inference. To tackle the major performance bottleneck of workload imbalance, we propose two techniques: dynamic local sharing and dynamic remote switching, both of which rely on hardware flexibility to achieve performance auto-tuning with negligible area or delay overhead. Specifically, UWB-GCN is able to effectively profile the sparse graph pattern while continuously adjusting the workload distribution among parallel processing elements (PEs). After converging, the ideal configuration is reused for the remaining iterations. To the best of our knowledge, this is the first accelerator design targeted to GCNs and the first work that auto-tunes workload balance in accelerator at runtime through hardware, rather than software, approaches. Our methods can achieve near-ideal workload balance in processing sparse matrices. Experimental results show that UWB-GCN can finish the inference of the Nell graph (66K vertices, 266K edges) in 8.1ms, corresponding to 199 \times , 16 \times , and 7.5 \times respectively, compared to the CPU, GPU, and the baseline GCN design without workload rebalancing.

I. INTRODUCTION

Deep learning and its use in a wide range of applications, from image classification to video processing to speech recognition and natural language processing, have given rise to paradigms such as *Convolutional-Neural-Networks* (CNNs) [17] and *Long-Short-Term-Memory* (LSTM) [14]. These paradigms typically use data represented in the Euclidean space and can efficiently extract latent information from Euclidean data such as images, videos, audios and texts [25].

While deep learning has achieved significant success in tasks with Euclidean data, there is an increasing number

of applications that use data generated from non-Euclidean domains and are represented as graphs with complex relationships and inter-dependency between objects, for which most existing deep learning algorithms may fall short. For instance, in E-commerce, a graph-based learning system exploits the interactions between users and products [3], [26] to make highly accurate recommendations. In chemistry, molecules are modeled as graphs for bioactivity identification in drug research [10]. In a citation network, papers can be categorized into different groups while linked to each other via citations [16], [24]. In each of these cases, the graph has various numbers of unordered nodes and each node has a different number of neighbors, leading to massive data dependencies among nodes.

The irregularity in graph data imposes significant challenges on existing machine learning algorithms and makes critical feature extraction operations, such as convolutions, not directly applicable. As a result, Graph Neural Networks have been proposed, in various forms, to extend deep learning approaches to graph data [7], [11], [20], [21], [23], [25]. Among these, the *Graph Convolutional Network* (GCN), an approach that marries some ideas of CNNs to the distinct needs of graph data processing, has demonstrated significant potential [4], [8], [16].

With the development of GCNs, their acceleration becomes an urgent issue. GCNs have been mainly adopted in two contexts: (1) big data and data mining, where social media and E-commerce companies use GCNs to identify user preference and where police officers use GCNs to learn about the social relationships of a criminal suspect; and (2) embedded devices, where GCNs are used in real-time motion capture and 3D modeling in animation and sports game development. Both contexts need high-performance GCN inference. For big data mining, it requires high throughput, especially on events like “Black Friday”, when millions of people shop on Amazon at the same time. To advertise the correct products to every customer, a very large graph needs to be continuously evaluated. For real-time motion capture and modeling, extremely short-latency inference is required. Although most of the existing




Fig. 1. Adjacency matrix non-zero distribution imbalance of the CORA and Pubmed dataset

GCNs have two or three layers, the current trend shows that GCNs are becoming deeper (as is what happened with CNNs). A GCN network with 152 layers has been proposed recently, and its efficiency has been demonstrated in the task of cloud semantic segmentation [19]. Strict timing requirement makes the acceleration of GCNs a critical area of research.

There has so far been relatively little work in GCN acceleration. The major kernel of GCN inference is *Sparse-Dense-Matrix-Multiplication* (SPMM). In the past few years, many efficient accelerators have been proposed for sparse CNN [12], [15], [28], as summarized in Section 6. These accelerators benefit from the following three features of sparse CNNs: (i) the non-zeros among CNN input and output channels are roughly balanced, leading to an easier workload distribution among parallel PEs; (ii) the sparsity of CNNs is relatively low, so typically the matrices can still be processed using a dense format; (iii) the sparse matrices or tensors in CNNs are relatively small in size. However, for this new GCN problem, these three assumptions are not held any more. Particularly, (a) the distribution of non-zeros can be extremely unbalanced and clustered (see the two examples in Figure 1) in sparse matrices for describing graphs (e.g., a graph adjacency matrix), given that real-world graphs often follow the power-law distribution. This leads to great challenges in workload distribution balancing, particularly on hardware; (b) the sparse matrices in GCN can be significantly sparser than in CNNs, so the hardware design need to efficiently handle sparse data formats (e.g., *CSR*, *CSC*, *COO*, etc.), which is much more difficult than processing it in software; (c) the imbalance issue is significantly exacerbated given the huge size of real-world graphs. For example, the *Reddit* graph (233K nodes and 12M links) is significantly larger than the typical images for sparse CNNs. Consequently, providing architectural support for efficient GCN acceleration with balanced workload distribution among large number of PEs becomes a very challenging task.

In this work, we propose UWB-GCN, an architecture for accelerating GCN inference through two dynamic strategies for workload balancing: *local sharing* and *remote switching*. We first propose a baseline architecture design and then present our online workload rebalancing techniques, which monitor and adjust the workload distribution by dynamically reconfiguring the task distribution network. The ideal configuration is reused for later iterations, forming a hardware performance




Fig. 2. Local and remote non-zero imbalance in a sparse matrix. We only show remote imbalance among rows since column remote imbalance can be easily handled by task queuing technology [12].

auto-tuning paradigm. We implement UWB-GCN in Verilog-HDL and evaluate it on a Xilinx VCU-118 FPGA. Compared with the baseline design, UWB-GCN enhances PE utilization rate, on average, from 60.4% to 92.6%, leading to $2.7\times$ speedups. This paper thus makes the following contributions:

- We propose UWB-SPMM, an SPMM engine, for hardware based sparse-dense-matrix-multiplication, on matrices with drastically imbalanced non-zero distribution. Particularly, we propose a novel hardware-based performance auto-tuning paradigm for workload rebalancing.
- We propose UWB-GCN, a GCN accelerator based on the UWB-SPMM engine that can significantly accelerate GCN inference with negligible overhead. Evaluation results show that UWB-GCN can provide, on average, $246.7\times$, $78.9\times$, $2.8\times$ speedups compared with CPU, GPU, and the baseline design without workload rebalancing.

II. BACKGROUND

In this section we briefly introduce GCNs, showing their differences from traditional network models like CNNs and the challenges on hardware design arising from these differences.

A. Graph Convolutional Network

In this paper, we focus on spectral-based graph convolutional networks [16], [25] since it is one of the most fundamental and widely used GCN structures with numeric variants [8], [13], [16]. Please refer to Section 6 for the history and alternative types of GCNs. Eq. 1 shows the layer-wise forward propagation of a multi-layer spectral GCN:

$$X^{(l+1)} = \sigma(AX^{(l)}W^{(l)}) \quad (1)$$

A is the graph adjacency matrix with each row delineating the connection of a vertex with other vertices. $X^{(l)}$ is the matrix of input features in layer- l ; each column of X represents a feature while each row denotes a node. W^l is the weight matrix of layer- l . $\sigma(\cdot)$ denotes the non-linear activation function, e.g., *ReLU* [17]. In general, A needs to be normalized via $\tilde{A} = D^{-\frac{1}{2}} \times (A + I) \times D^{-\frac{1}{2}}$ where I is the identity matrix, and $D_{ii} = \sum A_{ij}$. The reason is that without normalization, multiplying the feature vector $X^{(l)}$ by A will change its




Fig. 3. The structure of a single layer in GCN.

scale – those nodes with more neighbors tend to have larger values under feature extraction. Note that during both training and inference of GCN, \tilde{A} remains constant. Since \tilde{A} can be computed from A offline, in the remainder of this paper, we use A to denote the normalized \tilde{A} . In general, A is multiplied only once per layer. However, when multi-hop neighboring information is to be collected, A can be multiplied twice or more (i.e., A^2 , A^3 , etc.).

Eq. 1 is essentially derived from graph signal processing theory: convolutions on a graph can be converted to a multiplication of signal $x \in \mathbb{R}^N$ (i.e., a scalar for each node) and a filter $g \in \mathbb{R}^N$ in the frequency domain via the Fourier transform:

$$\text{CONV}(g, x) = \mathcal{F}^{-1}(\mathcal{F}(x) \odot \mathcal{F}(w)) = U(U^T x \odot U^T g) \quad (2)$$

where \odot denotes the Hadamard product. U is collection of eigenvectors for the normalized graph Laplacian $\mathcal{L} = I_N - D^{-\frac{1}{2}} A D^{-\frac{1}{2}} = U \Lambda U$. The diagonal matrix Λ comprises the corresponding eigenvalues. If a frequency domain filter $g_w = \text{diag}(W)$ is defined, then Eq. 2 can be simplified [4] as:

$$\text{CONV}(g_w, x) = U g_w U^T x \quad (3)$$

Eq. 3 can be further simplified by defining the filter as the Chebyshev polynomials of the diagonal matrix Λ [8], [16] to obtain Eq. 1.

Figure 3 illustrates the structure and compute flow for a layer of a GCN. By multiplying A and $X^{(l)}$, we are integrating information from connected neighboring nodes. By multiplying $A X^{(l)}$ with $W^{(l)}$, and going through the non-linear activation function $\sigma(\cdot)$, we obtain the input features for the next layer. After multiple layers, the GCN is able to extract very high-level abstracted features for various learning purposes.

B. GCN Matrices Profiling

To leverage the specific characteristics of the matrices for performance improvement, we profile the sparsity and dimensions of A , X , and W for a 2-layer GCN using 5 publicly available datasets that have been widely evaluated for GCNs [16]. The profiling results are listed in Table I. As can be seen, A is quite sparse (sparsity $\geq 99\%$). For the most of datasets, the distribution of input features for the first layer X_1 is also very sparse (sparsity $\geq 90\%$), as these are raw features obtained directly from the graph. X_2 becomes much denser (sparsity $< 50\%$). W s are, in general, dense matrices.

The dimensions of the matrices in GCNs depend on the dataset and can range from thousands to millions or even more. Therefore, A can be extremely large and have been saved in a

TABLE I

SPARSITY AND DIMENSIONS OF MATRICES IN A 2-LAYER GCN FOR THE 5 MOST WIDELY EVALUATED GCN GRAPH DATASETS. F1, F2, F3 REFER TO INPUT FEATURES OF LAYER-1,2,3. DIM IS SHORT FOR DIMENSION. DENSE IS SHORT FOR DENSITY.

		CORA	CITSEER	PUBMED	NELL	REDDIT
Dense	A	0.18%	0.11%	0.028%	0.0073%	0.043%
	W	100%	100%	100%	100%	100%
	X1	1.27%	0.85%	10.0%	0.011%	51.6%
	X2	78.0%	89.1%	77.6%	86.4%	60.0%
Dim	Node	2708	3327	19717	65755	232965
	F1	1433	3703	500	61278	602
	F2	16	16	16	64	64
	F3	7	6	3	186	41

sparse data format. Different from CNNs, where the number of features per layer are roughly similar or increasing, in GCNs the number of features often reduces drastically by layer. There are typically thousands of features in the first layer, but only a few dozen in the second.

These observations here are essentially general and widely applicable to other datasets for GCNs. A being large and sparse is due to the scale and fundamental properties (e.g., power-law) of real-world graphs. For CNNs, the size of feature-maps decreases with layers for higher abstraction, while the channels increases for stronger abstracting capability. In GCNs, however, the structure of the graph (i.e., A) keeps constant for each layer, while the number of feature channels decreases for the aggregation and abstraction of features. This overwhelming information aggregation explains why sparsity of X drops significantly with layers.

III. GCN BASELINE ARCHITECTURE

We propose a baseline architecture design for GCN. Although we call it baseline, it is the first architectural design specially for GCN, to the best of our knowledge. The baseline design shares some similarity with existing sparse CNN accelerator designs, but in addition needs to support ultra high sparsity and large dimensions for the input matrices. In the next section, we show how to achieve near optimal workload balancing on top of the baseline design.

A. Matrix Computation Order

To compute $A X W$, there are two alternative computation orders: $(A \times X) \times W$ and $A \times (X \times W)$. The choice dictates the volume of non-zero multiplications. Based on our profiling observations, A is ultra sparse and large, X is general sparse and usually large in columns, and W is small and dense. Since multiplying A and X leads to a very large dense matrix, then multiplying another dense matrix brings significant computation workload and delay. Alternatively, for $A \times (X \times W)$, both are sparse-dense matrix multiplications¹; the scale of computation is thus drastically smaller. Table II lists the amount of computation for four datasets following the two approaches. Since the difference is obviously huge, in our design we first perform $X \times W$ and then left multiply with A .

¹Sparse-matrix-dense-matrix multiplication is known as SPMM; Sparse-matrix-sparse-matrix multiplication is known as SPGEMM.

TABLE II
OPERATIONS REQUIRED UNDER DIFFERENT EXECUTION ORDERS

Layer	Order	CORA	CITeseer	PUBMED	NELL	REDDIT
Layer1	$(A \times X) \times W$	62.3M	197.5M	163.2M	257G	16.3G
	$A \times (X \times W)$	999.7K	1.87M	17.5M	47M	6.1G
Layer2	$(A \times X) \times W$	468.2K	493.0K	2.3M	800M	764.3M
	$A \times (X \times W)$	329.3K	357.6K	1.06M	735M	530.3M
ALL	$(A \times X) \times W$	62.8M	198.0M	165.5M	258G	17.1G
	$A \times (X \times W)$	1.33M	2.23M	18.6M	782M	6.6G




Fig. 4. Compressed-Sparse-Column sparse matrix format.

B. Baseline SPMM

Take $A \times B = C$ as an example, if A is in size $(m \times n)$, B is in size $(n \times k)$, C is in size $(m \times k)$, we can reformulate C as:

$$C = \left[\sum_{j=0}^{n-1} A_j b_{(j,0)}, \sum_{j=0}^{n-1} A_j b_{(j,1)}, \dots, \sum_{j=0}^{n-1} A_j b_{(j,k-1)} \right] \quad (4)$$

where A_j is the j th column of A . $b_{j,k}$ is an element of B at row- j and column- k . In other words, by broadcasting the j th element from column- k of W , to the entire column- j of A , we can obtain a partial resulting column of C . Essentially, B is processed in a streaming fashion: each element $b_{(j,k)}$ finishes all computation it involves at once, and then aborts completely. In this way, we reuse the entire sparse matrix A for each column of C (k times in total). Such a design brings additional advantages when A and C are stored in *Compressed-Sparse-Column* (CSC) format (see Figure 4). Further benefit with this design is that it provides opportunity to pipeline multiple SPMM operations, as will be discussed later. Since a complete resulting element of C requires an entire corresponding row of A , to avoid expensive parallel reduction in hardware, we partition A and C along the rows and assign them to the PEs. Figure 5 depicts the procedure of calculating C . The columns of A and elements of B in the same color are to be multiplied, and stored as partial results in C with the same color.

Workload Mapping: In the baseline design, with the assumption that non-zeros of a sparse matrix are evenly distributed among the rows, we adopt a direct and static mapping from matrix rows to PEs. For example, in Figure 6, each two rows of A are mapped to a separated PE; each PE processes three non-zeros of A eventually.

C. Baseline Architecture Design

The objective here is an efficient architecture design for sparse-matrix-dense-matrix multiplication (SPMM): $A \times B = C$, given A can be extremely sparse (e.g., sparsity $\geq 99\%$) or generally sparse (e.g., sparsity $\approx 50\%$). Figure 7 illustrates our baseline design, comprising the modules of *sparse-matrix-memory* (SPMMeM), *dense-column-memory* (DCM), *task-distributor & Queue* (TDQ), *PE-array*, and an *accumulation-*




Fig. 5. SPMM computation approach.




Fig. 6. Partitioning the sparse matrix rows among 8 PEs.

buffers-array (ACC). SPMMeM buffers the input sparse matrix A . DCM buffers the input dense matrix B . TDQ is for task distribution to the PEs. PE-array is for concurrent multiplication. Finally, ACC buffers the partial results of the resulting matrix C for accumulation.

Depending on the sparsity and storage format of A , we have two alternative designs for TDQ.

TDQ-1 (left of Figure 7) is used when A is general sparse and stored in dense format. We perform the direct row partition as discussed, and map to the input buffer of a PE (see Figure 6). Each cycle, $NPE/(1 - Sparsity)$ data are forwarded to a PE given evenly-distributed non-zeros. As one PE may account for more than a single row of A , we allocate multiple task queues (TQs) per PE. As shown in Figure 6-(B), in each cycle a PE can receive up to 4 non-zero elements. We have four queues to buffer these non-zeros from different rows of A . Each cycle, an arbitrator selects a non-empty queue, pops an element, checks for Read-after-Write (RaW) hazard (discussed later), and forwards it to the PE for processing.

TDQ-2 (right of Figure 7) is used when A is ultra-sparse and stored in CSC sparse format. Since in CSC the non-zeros are continuously stored in a dense array, if we can directly process the dense array, we gain from avoiding all the zeros. However, we suffer from the overhead of navigating to the correct PE as the indices are no longer continuous and essentially stored in another index array. We use a multi-stage Omega-network for routing the non-zero data stream to the correct PE according to their row indices from the index array. Each router in the *Omega-network* has a local buffer in case the buffer of the next stage is saturated. Our design attempts to balance the data forwarding rate and the processing capability of the PEs. This is achieved when non-zero elements are distributed evenly

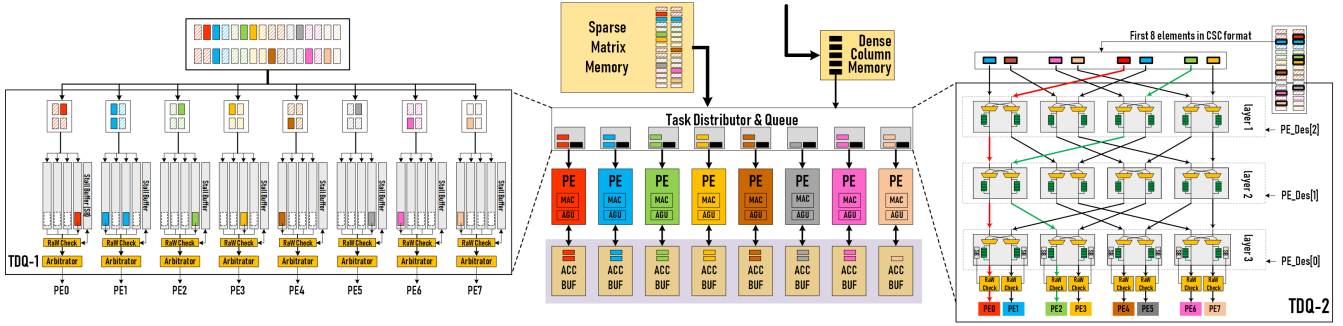


Fig. 7. Architecture design for the baseline SPMM engine.

among rows. Compared with a global crossbar network, the Omega-network design incurs much less area and hardware complexity; this is especially the case when we have a large number of PEs. Meanwhile, TDQ also accepts streaming data from a particular column of dense matrix B in DCM.

PEs fetch present partial results of C from ACC, perform the new multiplication task, add to the partial results, and save back to ACC. Each PE is coupled with a bank of ACC to store the rows of C it accounts for. A PE features two units: a *multiply-accumulate-unit* (MAC), and an *address-generation-unit* (AGU) for result address generation and forwarding. Since C is roughly a dense matrix and stored in dense format, the rows of C are statically partitioned among ACC buffers. Synchronization is only needed when an entire column of the resulting matrix C is completely calculated. Consequently, the imbalanced distribution of non-zeros across columns does not bring any performance issues.

An important issue here is the *Read-after-Write* (RaW) hazard. Since the computations are all floating-point, the pipelined MAC unit usually takes several cycles to process, but can still accept new tasks while processing. If the new task tries to accumulate to the same partial result of C (i.e., from the same row of A), it actually fetches a stale partial result from ACC, and a RaW hazard occurs. To avoid this hazard, we implement a stall buffer of size T , where T is the delay of the MAC units. We track the row indices currently being processed by the MAC and check whether the current element is targeting the same row in the *RaW-check-unit* (see Figure 7). If so, we buffer that job and delay for a few cycles until the hazard is resolved. This is similar to the role of the scoreboard for register RaW hazards in processor design.

Overall, for each layer of GCN, we first execute SPMM on $X \times W$. Since X is general sparse and stored in dense format, we use TDQ-1. The result of XW is dense. We then compute $A \times (XW)$. Again, this is SPMM. However, as A is ultra-sparse, and stored in CSC format, we use TDQ-2. The result is relatively dense. But after the activation function *ReLU*, a large portion of entries become zero, and we again have a general sparse matrix for the input feature matrix of the next layer.

To further improve the performance, we exploit the parallelism between the two consecutive SPMMs (i.e., $X \times W$ and $A \times (XW)$). This is based on the observation that when




Fig. 8. Exploiting extra parallelism across consecutive SPMM computation through pipelining.

a column of (XW) has finished computing, and A is constant and ready, we can already start the multiplication of A with that column, without the need to wait for the entire XW ; This is shown in Figure 8. This design brings two major benefits: (i) We gain extra parallelism and reduce the overall delay through coarse-grained pipelining and (ii) Instead of looking for large off-chip storage to cache the resulting XW matrix, we only need to buffer a single column of XW on-chip. Such a pattern can be reused within a GCN layer if left-multiplied by other sparse matrices, e.g., some GCNs collect information from 2-hop neighbors, so the layer formulation becomes $A \times (A \times (X \times W))$; the three multiplications can be pipelined.

D. Workload Balance Problem

The baseline architecture works well when non-zeros are evenly distributed among the rows of A . However, when this assumption does not hold, which is very likely for power-law graphs, then the performance of the baseline architecture degrades significantly due to PE workload imbalance.

Figure 9-(A) and (B) illustrate two types of workload imbalance: *local imbalance* and *remote imbalance*, and also the histogram when mapping to the baseline architecture with 8 PEs. Note that both types lead to significant performance degradation: the delay increases from the expected 2 cycles to 5 and 7 cycles, respectively.

This imbalance issue is unique for GCNs and has not been faced or resolved by existing works of sparse-CNNs [2], [12], [15], [28]. The reason is that non-zeros in those sparse matrices are more or less evenly distributed. However, when dealing with huge and ultra sparse matrices such as the adjacency matrix of a social-network graph following a power-law distribution, the condition is totally different. Efficiently handling of this unique workload balance problem from this new GCN application is the major research problem for this




Fig. 9. Local and remote workload imbalance among 8 PEs with 75% sparsity. work. Typically, when dealing with sparse data structures such as sparse matrices/tensors, trees and graphs, etc., to achieve workload balance, the software approach is to first profile the structure through, for example, symbolic analysis, in a preprocessing stage, and then use the sampled information to guide the partition strategy later for real processing. In this work, we show how to dynamically adjust hardware configuration for continuous workload rebalancing. Our design can be applied to a variety of specialized accelerators for processing sparse data structures.

IV. UWB-GCN ARCHITECTURE DESIGN

We treat the two types of imbalance problem (shown in Figure 9) separately. For local imbalance, we propose *dynamic local sharing*; for remote imbalance, we propose *dynamic remote switching*. Both of them are dynamic techniques that measure and adjust for a better task distribution configuration each time a column of the dense input matrix is processed. After several columns, the optimal configuration best matching the non-zero structure of the sparse input matrix is obtained. This configuration is then reused for the processing of the remaining columns of the dense matrix.

Their difference is the granularity. Figure 10 describes our design flow. We use heat-map to represent the utilization of different PEs (from blue 0% to red 200%). Initially, we employ equal partitioning for the baseline design. Some of the PEs (e.g., PE2, PE7 and PE8) are over-utilized while some (PE1, PE4 and PE6) are under-utilized. The ultimate purpose of the design is to balance the colors (i.e., utilization) by adjusting or exchanging the workloads of PEs (i.e., area in Figure 10). We first employ local balancing (right arrow) by averaging out some of the overloaded work to neighbors, improving the situation. However, the offloaded work needs to be returned for aggregation after processing. Due to chip area and design complexity restrictions, we may exchange workload between direct neighbors ($2 \leftarrow 3 \rightarrow 4$), 2-hop neighbors ($1 \leftarrow 2 \leftarrow 3 \rightarrow 4 \rightarrow 5$), or even 3-hop neighbors ($0 \leftarrow 1 \leftarrow 2 \leftarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 6$), but not all of them. In case non-zeros are clustered in a region across several PEs, local strategy may not work.

To allow an overloaded PE to exchange data with a remote under-loaded PE, we propose remote PE switching (down-arrow in Figure 10). By interchanging workloads between remote over-utilized and under-utilized PEs, followed by another round of local sharing (lower right-arrow), we can significantly improve load balancing. Note, this is what happens in the

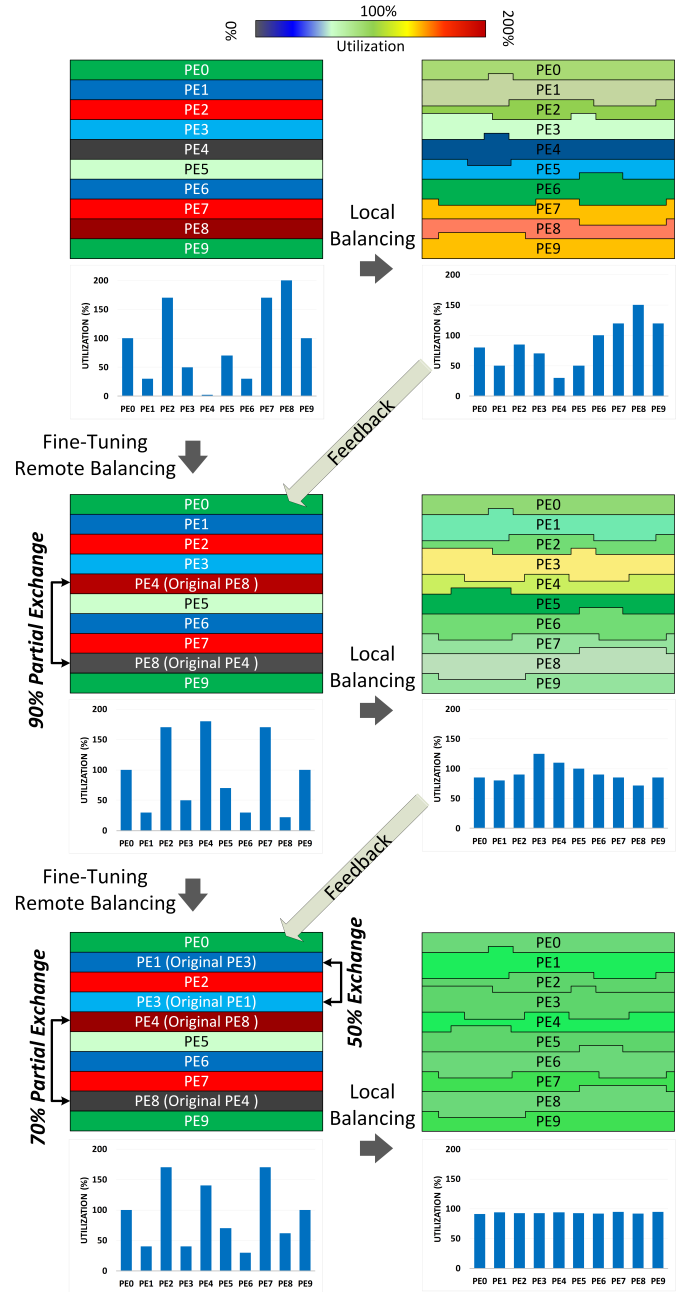


Fig. 10. Neighbor PE workload sharing and remote PE workload switching. processing of one column of the dense matrix B . Our accelerator can remember this plan and incrementally adjust it when processing the next column, for the same sparse matrix A is reused every time. After several rounds, the configuration best matching the sparse structure of A is obtained, and we use it for the remaining rounds. In the following, we discuss how to realize this strategy in hardware.

A. Dynamic Local Sharing

We need to estimate the PE utilization difference before adjusting the workload. This is achieved by comparing the number of pending tasks in its task queue (TQ) with the




Fig. 11. Architecture design for local workload sharing.

neighboring PEs. Figure 11 illustrates our design on how to realize 1-hop local sharing for TDQ-1 and TDQ-2 respectively.

TDQ-1: Before a new task is pushed into a PE’s TQ, it compares the number of pending tasks (by checking a waiting task counter) with the TQs of neighboring PEs. The task is then forwarded to the TQ with fewer pending tasks. If forwarded to a neighbor, the result needs to be returned to the ACC buffer of its original PE for accumulation after the multiplication, as shown in Figure 11-(B). The valid return address is calculated in the AGU unit of a PE.

TDQ-2: For an Omega network, it is the final layer of the multi-stage network accounting for forwarding between neighboring PEs. For example, in Figure 11-(C), two PEs share the same final-layer switch. Let us call them a group here. In Figure 12, a group has four PEs sharing the same final-layer switch. Therefore, we focus on the TQs of the final layer. After figuring out the pending task condition, we know the proper destination PE id. We adjust the address tag of the task before it is pushed into the TQs of the final layer. To enable PEs on the group border (e.g., the leftmost or rightmost PEs) to communicate with their out-of-the-group neighbors, we add extra links in the final layer, as shown in Figure 11-(D). Note, Figure 11-(D) just shows the condition of sharing among 1-hop neighbors. By considering more hop neighbors, we obtain a more balanced design at the cost of higher hardware complexity and area. This is a design trade-off. We discuss it in more detail in the evaluation.

B. Dynamic Remote Switching

For remote switching, we also adopt the multi-round auto-tuning approach. The idea is to find the most over-utilized PE and the most under-utilized PE per round (i.e., for a column of B), and switch a part or all of their workloads. The

percentage to be switched depends on their utilization gap. The architecture design is shown in Figure 12.

First, we discuss how to identify the most over-utilized (hotspot) PE and the most under-utilized (coldspot) PE. This is achieved by the *PE Status Monitor* (PESM). As previously mentioned, each TQ of a PE has a counter to track the number of pending tasks, which can trigger an “empty” signal when reaching zero. All the counters are connected to a multi-stage MUX-tree for integration, and a single signal Ψ is the output. Right after the jobs of the current round are dispatched, we start to monitor Ψ . When Ψ triggers, we know some PEs become idle. By voting at each level, the Mux-tree is able to identify the PE group with the highest number of “empty” signals triggered, i.e., the coldspot. When all PEs have triggered the empty signal, we record the present Ψ , which is the last PE to abort, i.e., the hotspot.

Having identified a hotspot and coldspot PE-tuple with id T_j for the current round i , to avoid thrashing, we only exchange a portion of the workload between them. The number of jobs (i.e., rows of A) to be switched in the i -th round (i.e., a column of B) N_i is calculated via:

$$N_i = \begin{cases} 0 & \text{if } i = 1 \\ N_{i-1} + G_i/G_1 \times (R/2) & \text{otherwise} \end{cases} \quad (5)$$

where G_i is the largest workload gap (i.e., the workload difference between hot-spot and cold-spot PEs) for the i -th round, R is the initial workload under equal partition. In the current design, we track each T_j tuple for two rounds. In other words, in the PE Status Monitor in Figure 12, we have two slots for tracking the current PE-tuple, and the PE-tuple from the previous round. We ensure there is no conflict here. Each PE-tuple being tracked is updated per round according to Eq 5. In this way, the workload switching ratio for each tracked PE-tuple is adjusted for two or several rounds and high-likely to converge. The number of rounds we can track depends on the size of the tracking window in the PESH, and is a design tradeoff between area and performance. Calculating Eq 5 is conducted in the *Utilization Gap Tracker* in Figure 12. To reduce the hardware cost of division and multiplication in calculating $G_i/G_1 \times (R/2)$, we also design a hardware-efficient approximation approach for computing Eq 5, which will not be discussed in detail here due to space limitation.

Knowing how many rows are to be switched between remote PEs, we use a *Shuffling Lookup Table* (SLT) to determine which rows are to be interchanged between the PE-tuple. The IDs of these rows are forwarded to the *Remote Balancing Control Register* (RBCR). In the next round, the destination PE of these rows are updated in the *Shuffling Switches* (SS).

V. EVALUATION

We evaluate the baseline and UWB-GCN designs, and compare them with the same GCN networks on other platforms such as CPU and GPU.

A. Evaluation Configuration

To evaluate UWB-GCN, we implement the RTL of the baseline architecture and UWB-GCN in *Verilog-HDL*. We

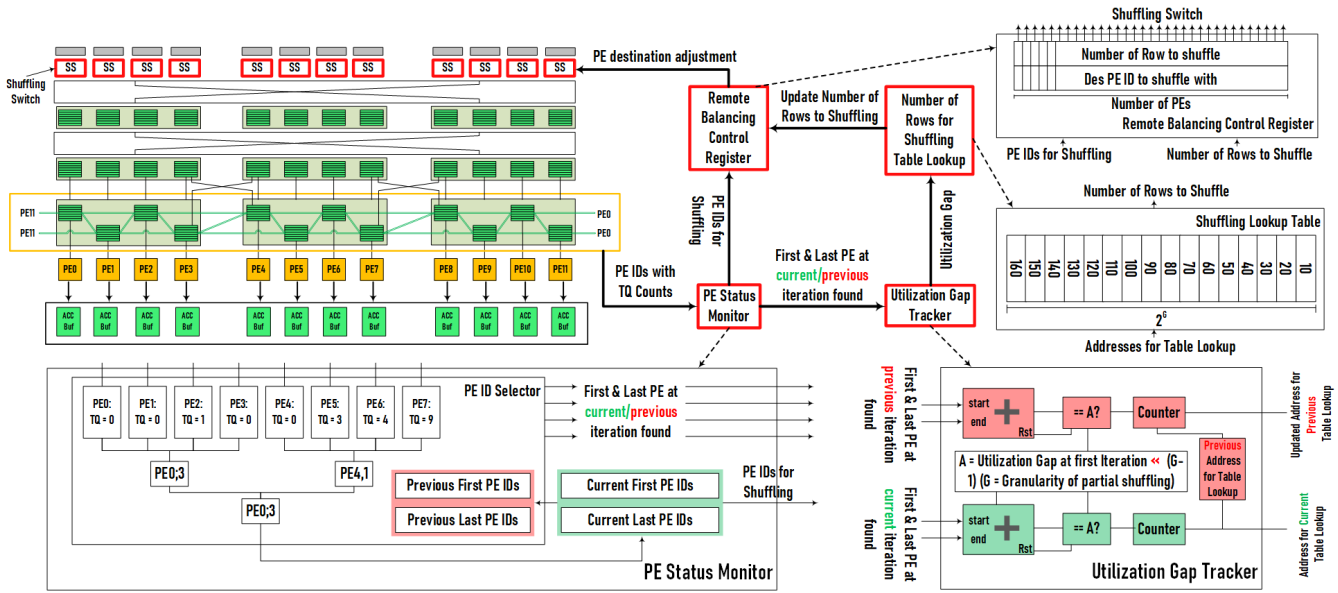


Fig. 12. UWB-SPMM overall architecture design with local workload sharing and remote workload switching. Modules in red boxes are used for remote workload switching.




Fig. 13. Number of non-zeros per row in the adjacency matrices of *Citeseer*, *Nell* and *Reddit*.

measure the PE utilization, performance, energy efficiency, and hardware resources consumption on a *Xilinx Virtex UltraScale+ VCU118* FPGA board. Note, we only use FPGA as an evaluation platform to demonstrate the performance and efficiency of UWB-GCN. Our design is a general architecture design that does not leverage any FPGA-specific features or units.

We allocate a counter to each PE for tracking the number of idle cycles for utilization measurement. The number of operation-cycles (i.e., execution delay) are measured by a cycle-accurate hardware counter. The counter triggers when the first data is forwarded to UWB-GCN and stops when the last output feature is received. The hardware consumption and operation frequency are reported by the *Vivado Design Suite-2019.1* after synthesis and implementation. The board-level power consumption is measured by a power meter. To perform cross-platform comparison, we implement the reference GCN networks in *PyTorch* and run them on a high-end server CPU *Intel Xeon E5-2698-V4* and a *NVIDIA Tesla P100 GPU*. *PyTorch* calls *cuSPARSE* library in the SPMM computation.

The datasets for evaluation are *Cora*, *Citeseer*, *Pubmed*, *Nell* and *Reddit*, which are the five most widely used public available datasets in GCN research. Their characteristics are listed in Table I.

B. UWB-GCN Evaluation

The efficiency of our design is evaluated by comparing the performance, hardware resources consumption, PE utilization of the baseline design (i.e., *Base*) with the 4 different design choices of UWB-GCNs: (i) 1-hop local sharing (i.e., *Design(A)*), (ii) 2-hop local sharing (i.e., *Design(B)*), (iii) 1-hop local sharing + remote switching (i.e., *Design(C)*), (iv) 2-hop local sharing + remote switching (i.e., *Design(D)*). The only exception is for *Nell*, where we use 2-hop and 3-hop local sharing; we explain the reason later.

Figure 13 illustrates the non-zero element distribution in the adjacency matrices of *Citeseer*, *Nell* and *Reddit*. The distributions of adjacent matrices of *Cora* and *Pubmed* have already been shown in Figure 1. As can be seen, the distribution of non-zeros in the graph adjacency matrices *A* are extremely unbalanced, confirming the importance of workload rebalancing.

Figure 14-(A-E) compare the overall GCN inference delay and average utilization of PEs for the five designs over the five datasets, respectively. The lines show overall PE utilization. The bars show the break down of the delay cycles according to the GCN layers. We also mark the latency lower bound when assuming full PE utilization. For *Cora*, *Citeseer* and *Pubmed*, using 1-hop and 2-hop local sharing can improve PE utilization from 53%, 71% and 69%, to 83%, 83% and 93%, respectively, leading to $1.93\times$, $1.25\times$ and $1.56\times$ performance improvement. Enabling remote switching can further promote PE utilization to 90%, 89% and 96%, respectively, bringing performance gain by $2.12\times$, $1.40\times$ and $1.62\times$. After analysis, we found the remaining 4-10% utilization gap is due to PE under-utilization in the auto-tuning phase, e.g., in the first iteration. For *Nell*, as shown in Figure 13, the non-zeros are quite clustered. In this case, one or two PEs are extremely over-utilized in the baseline design, leading to only 13%

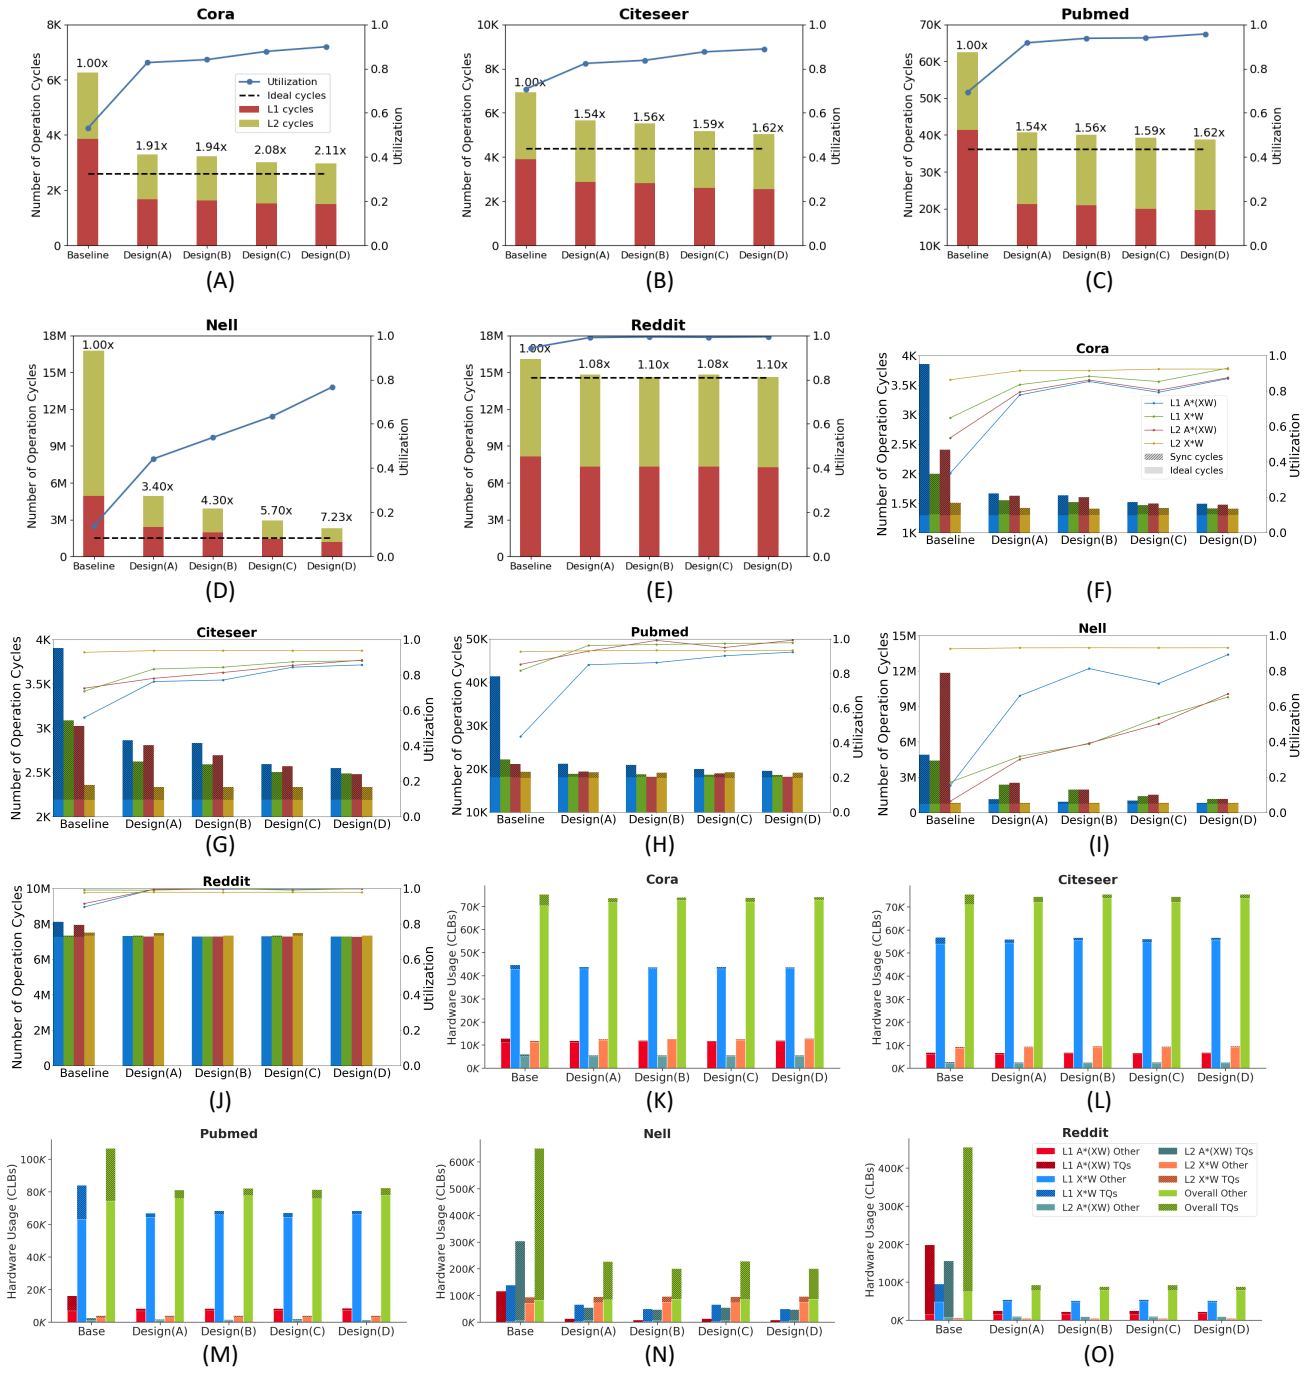


Fig. 14. A-E: Overall performance and PE utilization of the proposed five designs; F-J: Per-SPMM performance and PE utilization of the proposed five designs; K-O: Hardware resources consumption normalized to the number of CLBs of the five designs. The legends are shown in (A), (F) and (O).

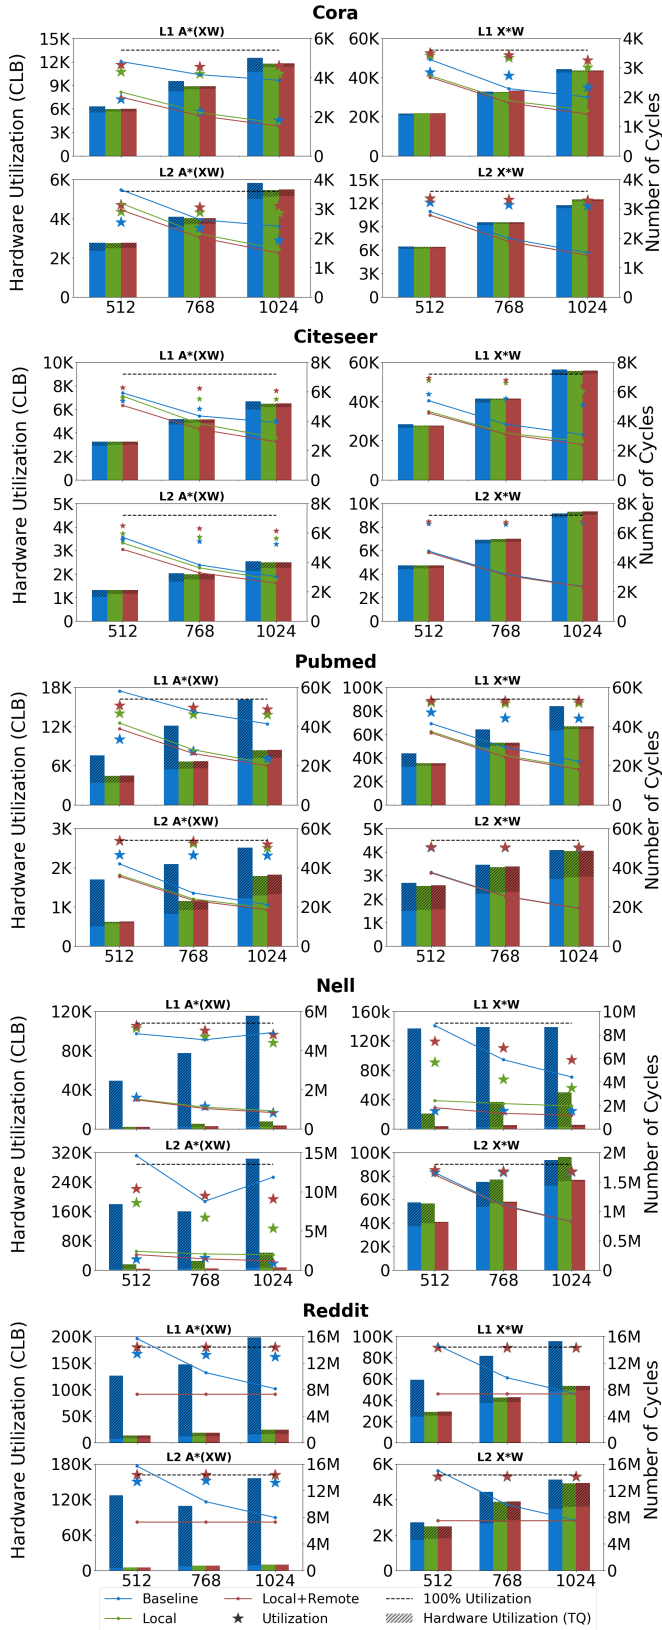


Fig. 15. Utilization, performance and hardware resources consumption with different number of PEs.

overall utilization. In this case, even 2-hop local sharing is still insufficient to rebalance the workload. Therefore, for the *Nell* dataset only, we use 2-hop and 3-hop local sharing (rather than 1-hop and 2-hop) in our evaluation. Experiment results show that 2-hop and 3-hop local sharing can enhance PE utilization from 13% to 44% and 53%, bringing $3.4\times$, $4.3\times$ performance improvement. With remote switching enabled, the utilization further increases to 68% and 82%, leading to $5.9\times$ and $7.5\times$ performance gain. For *Reddit*, through local sharing, the utilization has already achieved 99% (from 92% in the baseline). In UWB-GCN, hardware resources are allocated to different layers in proportion to their number of arithmetic operations. Thus, when perfect utilization is achieved, the different layers have almost the same execution time. As shown in Figure 14(A-E), the green and red bars have similar length at Design(D), while their lengths vary significantly at Baseline.

Figure 14(F-J) further break down the cycles of SPMM for the five designs over the five datasets, respectively. The shading area of the bars represent the "Sync" cycles due to workload imbalance (i.e., the waiting cycles at the barrier); the non-shading area represent the "Ideal" cycles assuming perfect workload balance. The bars in different color represent the cycles of the four SPMM operations (i.e., $A \times (XW)$ of Layer-1, $X \times W$ of Layer-1, $A \times (XW)$ of Layer-2, and $X \times W$ of Layer-2) in the two-layer GCNs [16], [25]. The curves show corresponding PE utilization.

Comparing among the datasets, for *Cora*, *Citeseer* and *Pubmed*, the imbalance mainly occurs in the $A \times (XW)$ SPMM operation of the input layer, which are significantly mitigated by the rebalancing techniques of UWB-GCN. For *Nell*, the imbalance mainly occurs in the $A \times (XW)$ SPMM of the hidden layer, which is also diminished by UWB-GCN's rebalancing techniques. *Reddit* by itself is already very balanced. Comparing among the SPMM operations, the utilization improves significantly for $A \times (XW)$ of Layer-1, $X \times W$ of Layer-1, and $A \times (XW)$ of Layer-2. For $X \times W$ of Layer-2, although X is sparse after filtered by the *ReLU* activation function of Layer-1, its sparsity is much lower compared to X in Layer-1, so the utilization is also high for the baseline (except *Cora*).

Figure 14(K-O) compare the overall hardware resources consumption of the five designs over the five datasets. The hardware resources cost is normalized to the number of *Configurable Logic Blocks* (CLBs) used in the design, which is the basic components of FPGA. In an ASIC design, this can be normalized to the number of transistors. The red area represents the CLB consumption for the TQs of the TDQ modules. The idea here is that, if the task distribution is more unbalanced, the TQs require more slots of the buffering queue. Therefore, by introducing the rebalancing techniques of UWB-GCN, the area cost of TQs should be reduced. This is especially the case for *Pubmed*, *Nell* and *Reddit*. The green area represents other hardware modules excluding the TQs. For this part, it keeps almost unchanged across the five datasets, which means the area overhead from the rebalancing logic of UWB-GCN is very small – only 2.1%, 3.9% and 1.9%

of the whole baseline-design area for 1-hop local-sharing, 2-hop local sharing, and remote switching designs, respectively (for *Nell*, it is 2-hop and 3-hop). Combining the two parts, the UWB-GCN design can essentially reduce hardware resources consumption when comparing to the baseline design, largely due to dramatically reduced per-PE TQ size under more balanced workload (e.g., For *Nell*, the TQ depth for $A \times (XW)$ of Layer-1 is 65128 in the baseline, which reduces to 2675 in Design(D)).

C. Scalability of UWB-GCN

We evaluate the scalability of UWB-GCN by increasing the number of PEs from 512 to 768 to 1024 for the baseline, local sharing, local sharing plus remote switching. For *Cora*, *Citeseer*, *Pubmed* and *Reddit*, we adopt 1-hop local sharing; for *Nell*, we adopt 3-hop local sharing. We show the performance, PE utilization and hardware consumption for the four SPMV operations of the five datasets, respectively, in Figure 15. The bars represent the hardware consumption (normalized to the number of CLBs). The lines represent the performance. The stars represent PE utilization. The dotted lines mark the full utilization (100%) upper bound.

For the baseline design, the PE utilization drops with more PEs due to increased unbalancing degree, since more PEs for partitioning means less rows per PE, highlighting the unbalance degree among PEs. In other words, PEs have less opportunities to absorb inter-row imbalance due to less number of rows to averaging out the imbalance. In contrast, GCNs with both local sharing and remote switching show relatively stable and high PE utilization. The PE utilizations with only local sharing scales better than the baseline, but worse than with both local sharing and remote switching. Overall, by introducing the rebalancing techniques, the performance of UWB-GCN scales almost in linear with the number of PEs, much better than the baseline design.

D. Cross-platform Comparison

Table III presents the cross-platform evaluation of our UWB-GCN design. We compare the inference latency (*milliseconds*), energy efficiency (normalized to *Graph Inference/kJ*) and operation frequency (*MHz*) of UWB-GCNs with implementations of GCNs in a high-end Intel CPU, a Pascal architecture based NVIDIA Tesla-P100 GPU, the baseline design without workload balancing, and the reproduced EIE reference implementation [12] (but tweaked for GCN processing). For UWB-GCN, we use the design choice-(d) for the comparison here. We use 1024 PEs in the UWB-GCN design here. As we can see, despite running at a relatively low frequency, our design achieves $248.2\times$, $78.9\times$, $2.8\times$ and $2.7\times$ speedups on average, over the high-end CPU, GPU, the baseline design without workload balancing, and the reference EIE design, respectively, across the five GCN graph datasets. Our design also achieves $870\times$, $618\times$, $2.7\times$ and $2.7\times$ better energy efficiency, respectively.

VI. RELATED WORK

The objects of early studies integrating graph and neural networks are known as graph neural networks or GNNs, which were first proposed by Gori et al. [11], and developed further by Micheli [21] and Scarselli et al. [23]. In GNNs, the representation of a target node is inferred by iteratively propagating neighbor information through recurrent neural architectures until they reach a stable fixed point [25]. The whole process is very computation intensive [7].

More recently, inspired by the huge success of convolutional neural networks (CNN) in extracting local features from images or videos, graph convolutional networks (GCNs) have emerged as an alternative approach in addressing graph data. In 2013, Bruna et al. [4] proposed a design for graph convolutional networks based on spectral graph theory; this was followed by a number of variants [8], [13], [16]. Since then other types of GCNs have been proposed, including spatial-based GCNs [7], [9], graph attention networks [1], and graph generative networks [27].

To the best of our knowledge, this is the first accelerator design focusing on GCNs. There have been many efforts on accelerating sparse CNNs [2], [12], [15], [18], [22], [28]. We briefly summarize these studies and explain why these solutions fall short when applied to GCNs. Kung et al. condense the sparse parameter matrix through column grouping [18]. In case of conflict, only the most significant parameters are kept, others are discarded. Essentially, some accuracy is traded-off for performance. Kim et al. [15] mitigate the workload imbalance problem of sparse CNNs by using information from design-time profiling. Han et al. [12] propose EIE, an SPMV accelerator that forwards non-zeros to PEs in column-major order; this is similar to our baseline design with TDQ-1. However, they only focus on SPMV and do not address the workload imbalance issue among the rows of the sparse matrix. Zhang et al. [28] rely on different indexing methods to identify and select non-zeros. However, these techniques do not function well when the matrix becomes ultra-sparse, as in GCNs. Albericio et al. [2] extend the DaDianNao design [6] by enabling skipping of zeros. However, they also focus on sparse-CNNs and do not address the workload imbalance issue. The reason these studies do not touch on the workload imbalance issue is partially because, compared with GCNs that process graphs, the impact of workload imbalance for sparse-CNNs is much less significant. Chen et al. [5] propose Eyeriss. Rather than skipping zeros, Eyeriss saves energy by power-gating computations with zeros involved. Finally, Zhuo and Prasanna [29] present an SPMV design for FPGAs. They use the CSR format, which can be applied to various sparse matrices. However, this design still suffers from irregular sparse structures and the workload imbalance problem.

VII. CONCLUSION

In this paper, we propose an architecture design called Ultra-Workload-Balanced-GCN to accelerate graph convolutional network inference. To tackle the major performance issues from workload imbalance, we propose dynamic local workload

TABLE III
CROSS-PLATFORM EVALUATIONS OF UWB-GCNs. THE CPU AND GPU DESIGNS ARE IMPLEMENTED IN PYTORCH. THE EIE-LIKE DESIGN IS A HOMEMADE REFERENCE DESIGN ACCORDING TO THE EIE ARCHITECTURE [12] ON THE SAME VCU118 FPGA.

	Network	Cora	Citeseer	Pubmed	Nell	Reddit
Intel-Xeon E5-2698V4	Freq(MHz)	2.2-3.6 GHz				
	Latency (ms)	3.90	4.33	34.15	1.61E3	1.08E4
	Energy (Graph Inference/kJ)	1.90E3	1.71E3	216.9	4.61	0.69
GPU NVIDIA Tesla-P100	Freq(MHz)	1328-1481 MHz				
	Latency (ms)	1.78	2.09	7.71	130.65	2.43E3
	Energy (Graph Inference/kJ)	1.87E3	1.59E3	432.3	25.51	1.37
EIE-like: VCU118 FPGA	Freq(MHz)	285 MHz				
	Latency (ms)	0.022	0.024	0.22	59.1	56.3
	Energy (Graph Inference/kJ)	1.19E6	1.11E6	1.20E5	438.2	452.1
Baseline: VCU118 FPGA	Freq(MHz)	275 MHz				
	Latency (ms)	0.023	0.025	0.23	61.0	58.9
	Energy (Graph Inference/kJ)	1.21E6	1.09E6	1.16E5	433.3	447.0
UWB-GCN: VCU118 FPGA	Freq(MHz)	275 MHz				
	Latency (ms)	0.011	0.018	0.14	8.1	53.2
	Energy (Graph Inference/kJ)	2.40E6	1.46E6	1.81E5	3.16E3	497.3

sharing and remote workload switching techniques. They rely on hardware flexibility to realize performance auto-tuning with area and delay overhead. This is the first accelerator design for GCNs that relies on hardware auto-tuning to achieve workload rebalancing for sparse matrix computations. We conduct RTL design and experiments on a Xilinx VCU-118 FPGA, which show that our design can achieve $246.7\times$, $78.9\times$, $2.8\times$ speedups on average over the high-end CPU, GPU, and the baseline design without workload rebalancing, across the five widely used GCN graph datasets.

VIII. ACKNOWLEDGEMENT

This research was supported by the DMC-CFA project under PNNL's Laboratory Directed Research and Development Program. This research was supported by the U.S. DOE Office of Science, Office of Advanced Scientific Computing Research, under award 66150: "CENATE - Center for Advanced Architecture Evaluation". This research was supported by the High Performance Data Analytics (HPDA) program at PNNL.

REFERENCES

- [1] Sami Abu-El-Haija, Bryan Perozzi, Rami Al-Rfou, and Alexander A Alemi. Watch your step: Learning node embeddings via graph attention. In *Advances in Neural Information Processing Systems*, pages 9180–9190, 2018.
- [2] Jorge Albericio, Patrick Judd, Tayler Hetherington, Tor Aamodt, Natalie Enright Jerger, and Andreas Moshovos. Cnvlutin: Ineffectual-neuron-free deep neural network computing. *ACM SIGARCH Computer Architecture News*, 44(3):1–13, 2016.
- [3] Rianne van den Berg, Thomas N Kipf, and Max Welling. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*, 2017.
- [4] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- [5] Yu-Hsin Chen, Tushar Krishna, Joel S Emer, and Vivienne Sze. Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE Journal of Solid-State Circuits*, 52(1):127–138, 2016.
- [6] Yunji Chen, Tao Luo, Shaoli Liu, Shijin Zhang, Liqiang He, Jia Wang, Ling Li, Tianshi Chen, Zhiwei Xu, Ninghui Sun, et al. Dadiannao: A machine-learning supercomputer. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 609–622. IEEE Computer Society, 2014.
- [7] Hanjun Dai, Zornitsa Kozareva, Bo Dai, Alex Smola, and Le Song. Learning steady-states of iterative algorithms over graphs. In *International Conference on Machine Learning*, pages 1114–1122, 2018.
- [8] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pages 3844–3852, 2016.
- [9] Hongyang Gao, Zhengyang Wang, and Shuiwang Ji. Large-scale learnable graph convolutional networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1416–1424. ACM, 2018.
- [10] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1263–1272. JMLR. org, 2017.
- [11] Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pages 729–734. IEEE, 2005.
- [12] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A Horowitz, and William J Dally. Eie: efficient inference engine on compressed deep neural network. In *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, pages 243–254. IEEE, 2016.
- [13] Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [15] Dongyoung Kim, Junwhan Ahn, and Sungjoo Yoo. A novel zero weight/activation-aware hardware architecture of convolutional neural network. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017*, pages 1462–1467. IEEE, 2017.
- [16] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [18] HT Kung, Bradley McDanel, and Sai Qian Zhang. Packing sparse convolutional neural networks for efficient systolic array implementations: Column combining under joint optimization. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 821–834. ACM, 2019.
- [19] Guohao Li, Matthias Müller, Ali Thabet, and Bernard Ghanem. Can gens go as deep as cnns? *arXiv preprint arXiv:1904.03751*, 2019.
- [20] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- [21] Alessio Micheli. Neural network for graphs: A contextual constructive approach. *IEEE Transactions on Neural Networks*, 20(3):498–511, 2009.

- [22] Angshuman Parashar, Minsoo Rhu, Anurag Mukkara, Antonio Puglielli, Rangharajan Venkatesan, Brucek Khailany, Joel Emer, Stephen W Keckler, and William J Dally. Senn: An accelerator for compressed-sparse convolutional neural networks. In *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*, pages 27–40. IEEE, 2017.
- [23] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.
- [24] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [25] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. A comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*, 2019.
- [26] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 974–983. ACM, 2018.
- [27] Jiaxuan You, Rex Ying, Xiang Ren, William L Hamilton, and Jure Leskovec. Graphrnn: A deep generative model for graphs. *arXiv preprint arXiv:1802.08773*, 2018.
- [28] Shijin Zhang, Zidong Du, Lei Zhang, Huiying Lan, Shaoli Liu, Ling Li, Qi Guo, Tianshi Chen, and Yunji Chen. Cambricon-x: An accelerator for sparse neural networks. In *The 49th Annual IEEE/ACM International Symposium on Microarchitecture*, page 20. IEEE Press, 2016.
- [29] Ling Zhuo and Viktor K Prasanna. Sparse matrix-vector multiplication on fpgas. In *Proceedings of the 2005 ACM/SIGDA 13th international symposium on Field-programmable gate arrays*, pages 63–74. ACM, 2005.