

GraphACT: Accelerating GCN Training on CPU-FPGA Heterogeneous Platforms

Hanqing Zeng

University of Southern California
Los Angeles, California
zengh@usc.edu

Viktor Prasanna

University of Southern California
Los Angeles, California
prasanna@usc.edu

ABSTRACT

Graph Convolutional Networks (GCNs) have emerged as the state-of-the-art deep learning model for representation learning on graphs. It is challenging to accelerate training of GCNs, due to (1) substantial and irregular data communication to propagate information within the graph, and (2) intensive computation to propagate information along the neural network layers. To address these challenges, we design a novel accelerator for training GCNs on CPU-FPGA heterogeneous systems, by incorporating multiple algorithm-architecture co-optimizations. We first analyze the computation and communication characteristics of various GCN training algorithms, and select a subgraph-based algorithm that is well suited for hardware execution. To optimize the feature propagation within subgraphs, we propose a light-weight pre-processing step based on a graph theoretic approach. Such pre-processing performed on the CPU significantly reduces the memory access requirements and the computation to be performed on the FPGA. To accelerate the weight update in GCN layers, we propose a systolic array based design for efficient parallelization. We integrate the above optimizations into a complete hardware pipeline, and analyze its load-balance and resource utilization by accurate performance modeling. We evaluate our design on a Xilinx Alveo U200 board hosted by a 40-core Xeon server. On three large graphs, we achieve an order of magnitude training speedup with negligible accuracy loss, compared with state-of-the-art implementation on a multi-core platform.

ACM Reference Format:

Hanqing Zeng and Viktor Prasanna. 2020. GraphACT: Accelerating GCN Training on CPU-FPGA Heterogeneous Platforms. In *Proceedings of the 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '20), February 23–25, 2020, Seaside, CA, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3373087.3375312>

1 INTRODUCTION

Recently, representation learning on graphs has attracted much attention. By extracting structured, low dimensional features from the unstructured, high dimensional graph, many downstream tasks such as node classification [14, 18], link prediction [38], graph classification [32] and clustering [29] can be performed easily and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FPGA '20, February 23–25, 2020, Seaside, CA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7099-8/20/02...\$15.00

<https://doi.org/10.1145/3373087.3375312>

effectively. Among the numerous representation learning methods, Graph Convolutional Networks (GCNs) [18] are capable of learning significantly better features than traditional methods [12, 24]. GCNs have been used to facilitate user recommendation for social networks [31], to identify protein functionality from interaction graphs [14] and to help circuit testability analysis for EDA [20].

Despite the popularity of GCNs, training is still expensive in terms of time and computation resources. To scale GCNs to larger graphs and to enable fast re-training on dynamic graphs, it is critical to develop accelerators. Existing works [14, 34] parallelize training on GPU and multi-core platform. How to accelerate training by exploiting the FPGA hardware is currently not well studied.

Similar to Convolutional Neural Network (CNNs), GCNs are built by iteratively stacking multiple layers. Operations of a *graph convolutional layer* are decomposed into two major steps, to (1) propagate information within the graph and (2) propagate information along the neural network layers. Correspondingly, FPGA accelerators need to address the below challenges to improve performance:

Memory access. The feature propagation (step 1) in a large and sparse graph incurs high volume of irregular memory accesses, both on-chip and off-chip. The memory challenge is unique to the GCN training problem. While CNN accelerators [19, 25, 30, 33, 35, 37, 39] achieve high data reuse and regular accesses by tiling and sliding windows on input tensors, such tensor operations do not generalize to sparse graphs. In addition, while graph processing accelerators [7, 8, 16, 36, 41] optimize data flow and layout for propagation of node labels, such optimizations hardly lead to performance gain when the labels are replaced with long feature vectors.

Computation. The propagation along GCN layers (step 2) involves computationally intensive dense tensor operations on the model weights and graph node features. Fast computation of this step require high consumption of hardware resources.

Load-balance. Degree-imbalance of graph nodes can significantly degrade the performance of feature propagation. Consequently, the overall layer computation (steps 1 and 2) can be load-imbalanced. It is challenging to design on-chip computation modules and the corresponding scheduler to ensure load-balance for arbitrary graphs.

We propose GraphACT, a framework to accelerate GCN training by addressing the three challenges. The main contributions are:

- We propose GraphACT with the following optimizations:
 - **Algorithm selection:** By analyzing various GCN training algorithms, we select a subgraph-based minibatch algorithm to significantly reduce CPU-FPGA communication.
 - **Redundancy reduction:** For each subgraph as a minibatch, we identify and eliminate the recurring aggregation

- operations between common node neighbors. Such online pre-processing on CPU significantly reduces the number of on-chip operations and BRAM accesses on FPGA.
- **Parallelization:** We parallelize the key training steps with optimized on-chip computation modules. Integrating the modules into the processing pipeline, we achieve load-balance on a wide range of target FPGA devices.
 - We develop an accurate performance model for **GraphACT**. The model identifies architectural parameters of the FPGA design, and algorithmic parameters of the minibatch sampler.
 - We evaluate **GraphACT** on a Xilinx Alveo U200 board hosted by a 40-core Xeon server. We achieve an order of magnitude training time speedup with negligible accuracy loss, compared with state-of-the-art multi-core implementation.

2 BACKGROUND AND RELATED WORK

We use bold capital letters (e.g., X) to denote matrices (zero indexed). We use X_u , $X_{:,v}$ and $X_{u,v}$ to denote a row, a column and an element of X , respectively. We use $X[a:b,:]$ to retrieve a sub-matrix by extracting rows of X (from the a^{th} to the $(b-1)^{\text{th}}$ row). Similarly, we use $X[:,a:b]$ to retrieve a sub-matrix by extracting the columns.

We use superscript \top to denote matrix transpose. Superscript within brackets (\cdot) denote index of a GCN layer.

2.1 Graph Convolutional Networks

Graph Convolutional Networks (GCNs) are built by extending the convolution operation defined on grid matrices to unstructured graphs [18]. The input to a GCN is an un-directed, node attributed graph $\mathcal{G}(\mathcal{V}, \mathcal{E}, X)$, where \mathcal{V} and \mathcal{E} denote the set of all nodes and edges. Each node is attributed by a length- f feature vector, so $X \in \mathbb{R}^{|\mathcal{V}| \times f}$ denotes the feature matrix of \mathcal{V} . A GCN embeds each graph node into a f' -dimensional vector. So the output of a GCN is simply an embedding matrix $X' \in \mathbb{R}^{|\mathcal{V}| \times f'}$, generated by information from both the node feature and the node connections.

A GCN consists of multiple graph convolutional layers. Figure 1 visualizes layer $(\ell + 1)$, where $0 \leq \ell \leq L - 1$. The layer has $|\mathcal{V}|$ input nodes, each attributed by a length- $f^{(\ell)}$ vector $(X_u^{(\ell)})^\top$. Note, for the first layer (i.e., $\ell = 0$), $X^{(0)} = X$, and $f^{(0)} = f$. For the last layer, $X^{(L)} = X'$ and $f^{(L)} = f'$. The layer performs the following operations to obtain the length- $f^{(\ell+1)}$ output vectors $(X_u^{(\ell+1)})^\top$:

Feature aggregation. As shown in Figure 1, each node u sends its feature $(X_u^{(\ell)})^\top$ via two types of connection: the self-connection (blue) and the neighbor-connection (green). Neighbor connection is defined by the edges \mathcal{E} . The features propagated to the same destination via neighbor-connection are aggregated by vector mean. Use node 0 as an example. Its input is $(X_0^{(\ell)})^\top = [2, 0, 6]^\top$, and it has neighbors 1, 2, 3. This step produces two vectors for node 0: the blue one $(X_0^{(\ell)})^\top$, and the green one $\frac{1}{3}(X_1^{(\ell)} + X_2^{(\ell)} + X_3^{(\ell)})^\top$.

Weight transformation. Each node independently transforms its two vectors output from the feature aggregation step. In Figure 1, we denote the transform function as $h(\cdot)$, parameterized by the GCN weights. Transformation function for u and v are identical.

Equation 1 defines the forward path of layer $\ell + 1$. D is the (diagonal) degree matrix, where $D_{u,u}$ equals degree of node u , and $D_{u,v}$ equals zero for $u \neq v$. A is the adjacency matrix of \mathcal{G} , where $A_{u,v}$ is 1 if $(u,v) \in \mathcal{E}$, and 0 otherwise. Each layer has two weight matrices: the self-weight $W_\circ^{(\ell+1)}$ and the neighbor-weight $W_\star^{(\ell+1)}$. The “|” operation concatenates two matrices column-wise. The feature aggregation operation corresponds to $D^{-1} \cdot A \cdot X^{(\ell)}$, and the transformation function $h(\cdot)$ applies weights $W_\circ^{(\ell+1)}, W_\star^{(\ell+1)}$, concatenates the intermediate results, and applies ReLU activation.

$$X^{(\ell+1)} = \text{ReLU} \left(X^{(\ell)} \cdot W_\circ^{(\ell+1)} | D^{-1} \cdot A \cdot X^{(\ell)} \cdot W_\star^{(\ell+1)} \right) \quad (1)$$

The output embedding X' (or $X^{(L)}$) of a GCN is often used in downstream applications such as node classification [5, 6, 14]. To classify nodes, we feed X' to a Multi-Layer Perceptron (MLP). A softmax layer [11] converts the MLP outputs to node labels.

Such node classification procedure enables supervised learning of $W_\circ^{(\ell)}, W_\star^{(\ell)}$. Suppose each training node $v \in \mathcal{V}$ is provided with a ground truth label. During training, we calculate cross-entropy loss \mathcal{L} to measure difference between the ground truth and the generated labels. We minimize the loss by updating $W_\circ^{(\ell)}, W_\star^{(\ell)}$ using gradient descent. Gradients with respect to layer ℓ parameters can be calculated from gradients of layer $\ell + 1$ by chain rule:

$$\frac{\partial \mathcal{L}}{\partial W_\circ^{(\ell)}} = \left(X^{(\ell-1)} \right)^\top \cdot \text{mask} \left(\frac{\partial \mathcal{L}}{\partial X^{(\ell)}} \left[:, 0 : \frac{1}{2} f^{(\ell)} \right] \right) \quad (2a)$$

$$\frac{\partial \mathcal{L}}{\partial W_\star^{(\ell)}} = \left(D^{-1} A X^{(\ell-1)} \right)^\top \cdot \text{mask} \left(\frac{\partial \mathcal{L}}{\partial X^{(\ell)}} \left[:, \frac{1}{2} f^{(\ell)} : f^{(\ell)} \right] \right) \quad (2b)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial X^{(\ell-1)}} = & \text{mask} \left(\frac{\partial \mathcal{L}}{\partial X^{(\ell)}} \left[:, 0 : \frac{1}{2} f^{(\ell)} \right] \right) \cdot \left(W_\circ^{(\ell)} \right)^\top \\ & + A D^{-1} \cdot \text{mask} \left(\frac{\partial \mathcal{L}}{\partial X^{(\ell)}} \left[:, \frac{1}{2} f^{(\ell)} : f^{(\ell)} \right] \right) \cdot \left(W_\star^{(\ell)} \right)^\top \end{aligned} \quad (2c)$$

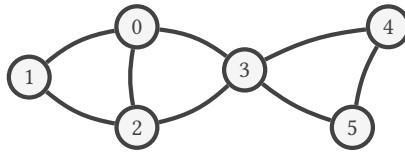
where function $\text{mask}(\cdot)$ corresponds to gradient of $\text{ReLU}(\cdot)$. Both mask and ReLU are light-weight elementwise functions on matrices.

For the node classifier, forward pass of a MLP layer is simply $X_{\text{MLP}}^{\text{out}} = \text{ReLU} \left(W_{\text{MLP}} \cdot X_{\text{MLP}}^{\text{in}} \right)$, where $X_{\text{MLP}}^{\text{in}}$ and $X_{\text{MLP}}^{\text{out}}$ are the input and output features for all \mathcal{V} , and W_{MLP} is the layer weight. Backward pass of a MLP layer performs computation similar to Equations 2a and 2c. For the softmax layer, its forward pass involves computation of exponential function, and its backward pass under cross-entropy loss only involves matrix subtraction [1].

To train GCNs on large graphs, minibatches need to be constructed. Unlike the case of CNN training where images are independent and identically distributed, training samples of GCNs (i.e., graph nodes) depends on each other due to edges. Thus, formulating GCN minibatches is challenging. There have been various techniques to sample minibatches [4–6, 14, 15, 34]. Section 3.1 analyzes the feasibility of hardware implementation for these techniques.

2.2 Deep Learning Training Accelerators

Various FPGA-based accelerators [10, 13, 23, 28, 40] have been proposed to train CNNs. The work in [40] proposes a modular design based on the types of layer operations, and improves performance



Training graph $\mathcal{G}(\mathcal{V}, \mathcal{E}, \mathbf{X})$

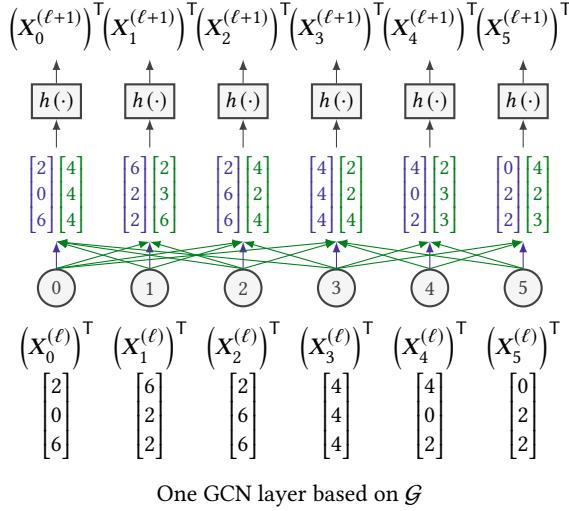


Figure 1: Overview of the GCN model

via reconfiguration. The work in [10] proposes a scalable framework to training CNNs on multi-FPGA clusters. Its partitioning and mapping strategy ensures load-balance. The works in [13, 23] accelerate training by model compression. The reduced model size alleviates the burden on BRAM and thus improves resource utilization.

To accelerate GCN training, the work in [34] proposes parallelization techniques for multi-core platform. It partitions features to increase cache-hit of each core. Although GCNs are an extension of CNNs to graphs, challenges to accelerate GCNs are significantly different. GCNs require both sparse and dense matrix operations. Apart from intensive computation, GCN accelerators need to address issues such as irregular memory access and load-balance.

2.3 Graph Analytics Accelerators

Many accelerator designs [7, 8, 16, 36, 41] have been proposed for traditional graph analytic problems, such as PageRank (PR), single source shortest path (SSSP) and breadth first search (BFS). The works in [7, 41] process input graphs in a two-phased gather-scatter manner, and utilize graph partitioning to improve access locality. The work in [8] extends the above memory optimization to multi-FPGA platforms, and the works in [16, 36] propose optimization specific to HBM and HMC to further boost FPGA performance.

The memory optimizations proposed in the above works do not directly apply to GCN accelerators. First of all, traditional graph analytics often propagate scalars along the graph edges, while GCNs propagate long feature vectors. Thus, although in both cases memory accesses are irregular, the access patterns are very different. Secondly, traditional graph analytics often propagate information

within the full graph, while GCNs propagate within minibatches. Thus, techniques such as graph partitioning may not be effective since minibatch size is much smaller than the full graph size.

3 TRAINING ALGORITHM

We observe from the forward pass (Equation 1) and the backward pass (Equation 2), that GCN training involves three types of matrix product $P \cdot Q$: (1) P being dense and Q being dense, (2) P being binary sparse and Q being dense, and (3) P being diagonal and Q being dense. From acceleration perspective, type (1) operation is computationally intensive and thus requires massive parallelization of DSPs (Section 4.3); type (2) operation incurs large number of irregular BRAM accesses and motivates the graph topology based optimizations (Section 3.2); type (3) operation is equivalent to scaling each row of Q by the diagonal elements of P , and thus its contribution to the total training cost is negligible. Other operations (e.g., mask, ReLU, concatenation, sub-matrix extraction) are light-weight and straightforward to implement, and we do not discuss them in detail.

3.1 Algorithm Selection

For efficient training on large graphs, the first step is to reduce external memory communication via minibatches. Ideally, a minibatch should be sampled such that all data required for gradient calculation (e.g., $\mathbf{X}^{(\ell)}$ of the minibatch nodes) fits in BRAM. Among the numerous algorithms [5, 6, 14, 15, 34], some return minibatches not suitable for hardware execution. We categorize these algorithms and analyze the hardware cost on external memory accesses.

Minibatch by sampling GCN layers [5, 6, 14, 15]. Sampling algorithms in this category traverse GCN layers backward from layer- L outputs to layer-1 inputs to select minibatch nodes. Assume b nodes of layer $\ell + 1$ are selected. The sampler then take $\alpha \cdot b$ nodes of layer ℓ based on the inter-layer connections and/or the features of the $\bar{d} \cdot b$ neighbors in layer- ℓ (where \bar{d} is the average node degree). Specifically, the sampling of [14] does not depend on neighbor features and $10 \leq \alpha \leq 50$ in general. [6] uses the neighbor features as “historical activation” and $\alpha = 2$. [15] computes the node probability by the neighbor features and $\alpha = 1$ on average. [5] does not require neighbor features and $\alpha = 1$. In summary, suppose we sample b_0 output nodes of layer L . Then the sampler reads the neighbor features of each layer (if required) to eventually return $\alpha^L \cdot b_0$ input nodes in layer 1. During training, we read features of the $\alpha^L \cdot b_0$ input nodes of layer 1, and then perform forward propagation to compute output features of the layer- ℓ sampled nodes ($1 \leq \ell \leq L$).

Minibatch by sampling training graph [34]. The algorithm samples from \mathcal{G} instead of the GCN layers. Given a graph sampling algorithm (e.g., multi-dimensional random walk [26]), [34] returns a subgraph $\mathcal{G}_s(\mathcal{V}_s, \mathcal{E}_s)$ induced from $\mathcal{G}(\mathcal{V}, \mathcal{E})$ (where $|\mathcal{V}_s| \ll |\mathcal{V}|$). The minibatch contains the same set of nodes \mathcal{V}_s for all the GCN layers. In other words, for each minibatch, [34] constructs a *complete* L -layer GCN from \mathcal{G}_s , with the layer nodes defined by \mathcal{V}_s and the layer connections defined by \mathcal{E}_s . Note that unlike [6, 15], the graph sampler of [34] requires no information from node features.

Suppose we were to implement the above minibatch training algorithms on FPGA. Since the full $\mathbf{X}^{(\ell)}$ is too large to fit on-chip, we need to read from external memory the following data: (1) layer-1

input features of the minibatch nodes, and (2) layer- ℓ input features of the minibatch neighbor nodes (if required). For ease of analysis, assume the feature sizes of each layer are the same, and let the cost of transferring one feature vector be 1. Also, let the cost of aggregation and computing $h(\cdot)$ for one node (see Figure 1) be 1. Ignore the computation cost of feature aggregation. Table 1 summarizes the ratio between on-chip computation cost and off-chip communication cost, where for the “Value” row, we set $L = 2$, $\bar{d} = 15$ and α as 25, 1, 2, 1 for [14], [5], [6], [15] respectively. Algorithms of low computation-communication ratio (i.e., [6, 14, 15]) impede the development of efficient hardware accelerators. For the remaining two algorithms, the complexity of the sampler of [34] is much lower than that of [5], and the training accuracy of [34] is higher. Thus, we select [34] as the target training algorithm for acceleration.

Table 1: Computation-communication ratio of various training algorithms. Let $B = \sum_{\ell=0}^{L-1} \alpha^\ell b_0$, and $B' = L \cdot b_0$.

	[14]	[5]	[6]	[15]	[34]
Expression	$\frac{B}{\alpha^L b_0}$	$\frac{B}{\alpha^L b_0}$	$\frac{B}{\alpha^L b_0 + \bar{d}B}$	$\frac{B}{\alpha^L b_0 + \bar{d}B}$	$\frac{B'}{b_0}$
Value	0.04	2	0.06	0.06	2

Remark on notation. Since we select the algorithm of [34], in the following sections, we mainly focus on GCN built on a subgraph \mathcal{G}_s . A symbol with an “s” subscript means it is related with a subgraph. Thus, forward and backward pass of the GCN can be updated by simply replacing $A, D, X^{(\cdot)}$ with $A_s, D_s, X_s^{(\cdot)}$ in Equations 1 and 2.

3.2 Redundancy Reduction

Before designing the hardware architecture, we first optimize the redundancy in minibatch training from an algorithm perspective. Take Figure 2a as an example. Suppose the graph sampler returns a subgraph \mathcal{G}_s (of \mathcal{G} in Figure 1) as the minibatch. In each GCN layer, nodes 1, 2, 3 and 4 aggregate neighbor features, with node 0 calculating $\frac{1}{3} (X_1^\top + X_2^\top + X_3^\top)$, node 1 calculating $\frac{1}{2} (X_0^\top + X_2^\top)$, node 2 calculating $\frac{1}{3} (X_0^\top + X_1^\top + X_3^\top)$ and node 3 calculating $\frac{1}{2} (X_0^\top + X_2^\top)$ (omitting superscript (ℓ) for simplicity). This corresponds to 6 vector additions and 10 vector reads in total. Observe that the vector pair (X_0^\top, X_2^\top) appears in the aggregation of both nodes 1 and 3.

Similarly, the pair (X_1^\top, X_3^\top) is aggregated by both nodes 0 and 2. Thus, we can perform pre-processing to compute the partial sum $X_0^\top + X_2^\top$ and $X_1^\top + X_3^\top$. The aggregation of the four nodes after pre-processing requires only 2 additions and 4 reads. In general, even considering the pre-processing cost, we can still significantly reduce both computation and communication of feature aggregation. The redundancy reduction helps us achieve perfect load-balance of the *complete pipeline* for a wide range of FPGAs (Section 5.3).

The key to reduce redundancy is to identify sets of nodes appearing frequently in the neighbor lists of $v \in \mathcal{V}_s$. To simplify the problem, we aim at finding such sets of size 2 – we identify *common pairs of neighbors*. We first construct an un-directed *aggregation graph* \mathcal{G}_a from \mathcal{G}_s by Algorithm 1. Each edge (u, v) in \mathcal{G}_a represents potential vector sum operations between u and v . The attribute

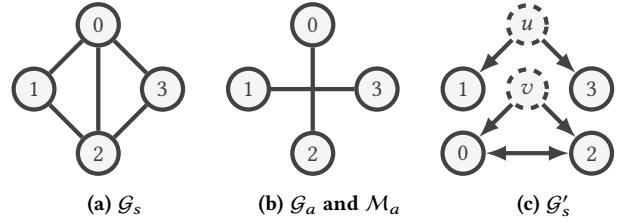


Figure 2: Example on reducing redundancy in feature propagation, using a graph theoretical approach.

of (u, w) consists of a set of nodes $\{v_1, \dots, v_n\}$, meaning that the neighbor list of v_i ($1 \leq i \leq n$) contains both u and w . The weight of an edge is simply the size of its attribute set (we remove edges of weight 1). In Algorithm 1, \mathcal{W}_a and \mathcal{D}_a can be implemented as a hash table. And $\mathcal{W}_a[(u, w)]$ (or $\mathcal{D}_a[(u, w)]$) returns the value corresponding to the key (u, w) . Figure 2b shows the aggregation graph \mathcal{G}_a for our above example, where edges have weight 2.

Algorithm 1 Construction of aggregation graph \mathcal{G}_a

Input: Subgraph \mathcal{G}_s ($\mathcal{V}_s, \mathcal{E}_s$)

Output: Aggregator graph \mathcal{G}_a

```

1:  $\mathcal{W}_a \leftarrow \mathcal{V}_s$ 
2:  $\mathcal{W}_a \leftarrow \emptyset$   $\triangleright$  Key-value pairs mapping edges to edge weights
3:  $\mathcal{D}_a \leftarrow \emptyset$   $\triangleright$  Key-value pairs mapping edges to edge attributes
4: for  $v \in \mathcal{V}_s$  do
5:   for  $u \in \text{neigh}(v)$  do
6:     for  $w \in \text{neigh}(v) \setminus u$  do
7:       if Key  $(u, w)$  not in  $\mathcal{W}_a$  then
8:         Add key-value pair  $((u, w), 1)$  to  $\mathcal{W}_a$ 
9:       else
10:         $\mathcal{W}_a[(u, w)] \leftarrow \mathcal{W}_a[(u, w)] + 1$ 
11:       if Key  $(u, w)$  not in  $\mathcal{D}_a$  then
12:         Add key-value pair  $((u, w), \{v\})$  to  $\mathcal{D}_a$ 
13:       else
14:          $\mathcal{D}_a[(u, w)] \leftarrow \mathcal{D}_a[(u, w)] \cup \{v\}$ 
15: Remove key-value pairs of weight-1 edges in  $\mathcal{W}_a$  and  $\mathcal{D}_a$ 
16:  $\mathcal{G}_a \leftarrow$  Un-directed graph based on  $\mathcal{W}_a$  and  $\mathcal{D}_a$ 

```

Intuitively, the next step upon getting \mathcal{G}_a is to extract the edges with largest weights, so that pre-computation of the corresponding vector sums reduces the most redundancy. However, there is one subtlety to notice. Suppose two edges of \mathcal{G}_a , (u_i, u_j) and (u_j, u_k) , have large weights, and $\{v_1, \dots, v_m\}$, the intersection of their attributes, is non-empty. Consider the aggregation of nodes $\{v_1, \dots, v_m\}$ after pre-computing $\mathbf{x}' = X_i^\top + X_j^\top$ and $\mathbf{x}'' = X_j^\top + X_k^\top$. By replacing the pair (X_i^\top, X_j^\top) with \mathbf{x}' , the other pair (X_j^\top, X_k^\top) disappears in aggregation. So \mathbf{x}'' does not help reduce redundancy. To avoid such useless pre-computation on \mathbf{x}'' , a good solution is to find a *maximum weight matching* of \mathcal{G}_a , so that the selected pairs imply high redundancy, and share no common nodes.

THEOREM 3.1. *For feature aggregation of each layer, number of vector reads and additions can decrease by at least $\sum_{e \in \mathcal{M}_a^*} (\mathcal{W}_a[e] - 2)$ and $\sum_{e \in \mathcal{M}_a^*} (\mathcal{W}_a[e] - 1)$. \mathcal{M}_a^* is maximum weight matching of \mathcal{G}_a .*

PROOF. We first consider the reduction in number of reads. Since edges in a matching are disjoint, for each $(u, v) \in \mathcal{M}_a^*$, accessing $X_u^\top + X_v^\top$ instead of X_u^\top and X_v^\top saves $(2 - 1) \cdot \mathcal{W}_a[(u, v)]$ number of reads. The pre-computation of $X_u^\top + X_v^\top$ consumes 2 reads. In sum, the total reduction is $\sum (\mathcal{W}_a[e] - 2)$. The proof for reduction in number of additions can be similarly derived. \square

Note that Theorem 3.1 considers the total cost including any pre-computation cost. Also, although polynomial complexity algorithm [22] exists for computing the maximum weight matching, it is still too expensive in practice. Thus, we propose a greedy approach (Algorithm 2) to quickly compute a good matching \mathcal{M}_a .

Algorithm 2 Greedy approach to find a good matching \mathcal{M}_a

Input: Aggregation graph \mathcal{G}_a ; threshold θ
Output: Matching \mathcal{M}_a

- 1: $\mathcal{M}_a \leftarrow \emptyset$
- 2: $\mathbb{H} \leftarrow$ Max-heap of edges \mathcal{E}_a , ordered by edge weights
- 3: $\mathcal{S} \leftarrow$ Length $|\mathcal{V}_a|$ vector of values True
- 4: **while** $\text{maxWeight}(\mathbb{H}) > \theta$ **do**
- 5: $(u, v) \leftarrow \text{extractMaxEdge}(\mathbb{H})$
- 6: **if** $\neg(\mathcal{S}_u \wedge \mathcal{S}_v)$ **then** ► (u, v) violates edge disjointness
 continue
- 7: $\mathcal{S}_u \leftarrow \text{False}; \mathcal{S}_v \leftarrow \text{False}$
- 8: $\mathcal{M}_a \leftarrow \mathcal{M}_a \cup \{(u, v)\}$

The next step after obtaining \mathcal{M}_a is to update the original \mathcal{G}_s to \mathcal{G}'_s (following Algorithm 3), so that the pre-computed sums can propagate in \mathcal{G}'_s . The idea is to merge each pair of nodes in \mathcal{M}_a into a new node, whose feature vector is returned by pre-computation. Compared with \mathcal{G}_s , the updated \mathcal{G}'_s has more nodes ($|\mathcal{V}'_s| = |\mathcal{V}_s| + |\mathcal{M}_a|$), but less edges. Note that \mathcal{G}'_s is directed, even when \mathcal{G}_s is un-directed. The example \mathcal{G}'_s is shown in Figure 2c.

Algorithm 3 Construction of the update subgraph \mathcal{G}'_s

Input: Original subgraph \mathcal{G}_s ; Matching \mathcal{M}_a ; Edge attribute \mathcal{D}_a
Output: Updated (directed) subgraph \mathcal{G}'_s ($\mathcal{V}'_s, \mathcal{E}'_s$)

- 1: $\mathcal{V}'_s \leftarrow \mathcal{V}_s; \mathcal{E}'_s \leftarrow \mathcal{E}_s$
- 2: **for** $(u, v) \in \mathcal{M}_a$ **do**
- 3: Assign a new node v' corresponding to the edge (u, v)
- 4: $\mathcal{V}'_s \leftarrow \mathcal{V}'_s \cup \{v'\}$
- 5: **for** $w \in \mathcal{D}_a[(u, v)]$ **do**
- 6: Remove (u, w) from \mathcal{E}'_s
- 7: Replace (v, w) with (v', w) in \mathcal{E}'_s

Complexity analysis. Complexity of Algorithms 1, 2 and 3 are low compared with the feature aggregation. Complexity of Algorithm 1 is $O(|\mathcal{E}_a|) = O(\sum_{v \in \mathcal{V}_s} d_v^2)$, where d_v is v 's degree. Although in the worst case, $O(|\mathcal{E}_a|) = O(|\mathcal{V}_s| \cdot d_{\max}^2)$, for the graphs in practice, we may assume $O(|\mathcal{E}_a|) = O(|\mathcal{V}_s| \cdot \bar{d}^2) = O(|\mathcal{E}_s| \cdot \bar{d})$, where d_{\max} and \bar{d} are the max and average degree of \mathcal{G}_s . Complexity of Algorithm 2 is $O(|\mathcal{E}_a| + N \log |\mathcal{E}_a|)$, where the first term counts for line 2, and the second term is for the loop from line 4 to 9. Number of times max-edge is extracted from heap, N , depends on

the threshold θ . Typically, $N \ll |\mathcal{E}_a|, |\mathcal{V}_s| < 5000, \bar{d} < 20$. Overall complexity of Algorithms 1 and 2 is much less than the complexity of single layer feature aggregation (i.e., $O(|\mathcal{E}_s| \cdot f)$, where the feature length f is to the order of 10^2 to 10^3). For Algorithm 3, lines 6 and 7 each takes at most d_w operations if we simply scan the neighbor list of w . In return, we save one vector sum operation for u and v . Considering $\bar{d} \ll f$, overhead of Algorithm 3 is thus negligible compared with the benefit in redundancy reduction.

It is worth noticing that (1) one time transformation from \mathcal{G}_s to \mathcal{G}'_s benefits $3L$ number of feature aggregation in a L -layer GCN, and (2) such transformation, which only involves integer operations, reduces floating point arithmetics during feature aggregation. These two observations further justify the cost of Algorithms 1, 2 and 3.

Iterative redundancy reduction. After obtaining \mathcal{G}'_s , redundancy still exists when aggregating features of \mathcal{G}'_s . Viewing \mathcal{G}'_s as the new \mathcal{G}_s , we can again apply Algorithms 1, 2 and 3 to obtain a \mathcal{G}''_s . This process can continue until few edges can be reduced (i.e., the matching is small). Note that although the training graph (and thus \mathcal{G}_s) is often un-directed, Algorithms 1, 2 and 3 still apply when \mathcal{G}_s is directed. Therefore, iterative redundancy reduction is feasible.

For sake of notation, define one *round* of redundancy reduction as one invocation of Algorithms 1, 2 and 3. Define the subgraph output by the last round as $\mathcal{G}^\#$, and its adjacency matrix as $A^\#$. Define the set of matchings in all rounds as $\mathbb{M}_a = \{\mathcal{M}_a, \mathcal{M}'_a, \mathcal{M}''_a, \dots\}$. Define $\gamma_{\text{add}} := \frac{\text{numAdd}(\mathcal{G}^\#) + \sum_{M \in \mathbb{M}_a} |\mathcal{M}|}{\text{numAdd}(\mathcal{G}_s)}$ as the redundancy reduction rate for additions, where $\text{numAdd}(\mathcal{G}) = \sum_{v \in \mathcal{V}; d_v \geq 1} (d_v - 1)$ denotes the number of additions to aggregate by traversing \mathcal{G} 's neighbor lists. Define $\gamma_{\text{read}} := \frac{\text{numRead}(\mathcal{G}^\#) + \sum_{M \in \mathbb{M}_a} 2|\mathcal{M}|}{\text{numRead}(\mathcal{G}_s)}$ as the redundancy reduction rate for BRAM reads, where $\text{numRead}(\mathcal{G}) = \sum_{v \in \mathcal{V}} d_v = \bar{d} |\mathcal{V}|$ denotes number of reads to aggregate features.

4 ACCELERATOR DESIGN

4.1 Overview

We first partition the workload between FPGA and CPU. We let FPGA execute the computation intensive operations, and leave the communication intensive parts to CPU. A natural partition is: CPU performs graph sampling, and then converts \mathcal{G}_s to $\mathcal{G}^\#$; FPGA performs the forward and backward pass defined by Equations 1 and 2, where feature propagation of each layer is based on $\mathcal{G}^\#$. Notice that under supervised learning, the last graph convolutional layer is followed by a node classifier (see Section 2.1). The softmax layer at the end of the classifier and the cross-entropy loss require computation of exponential and logarithmic functions. Since softmax and loss contribute to a negligible portion of the total workload and their accurate calculation requires complicated hardware [9, 21], we assign their computation to CPU. Section 4.4 describes the scheduling.

To improve overall training throughput, we need to (1) reduce the overhead in external memory access, and (2) increase the utilization of the on-chip resources. The first goal can be achieved by setting the minibatch parameters so that the size of \mathcal{G}_s is appropriate for the BRAM capacity. Ideally, once receiving the initial node features (i.e., $X_s^{(0)}$ of \mathcal{V}_s) and the connection information (i.e., $D_s, A_s^\#, \mathbb{M}_a$), the FPGA should propagate forward from the first graph convolutional layer to the last classifier MLP layer without

accessing external memory. Similarly, once receiving the gradient with respect to softmax, FPGA should propagate backward without accessing external memory. Thus, CPU needs to communicate:

- To FPGA: $X_s^{(0)}$, D_s , $A_s^\#$, \mathbb{M}_a and $\frac{\partial \mathcal{L}}{\partial X_{\text{MLP}}^{\text{out}}}$
- From FPGA: $X_{\text{MLP}}^{\text{out}}$

And on-chip BRAM needs to store:

- Node features: $\bigcup_{0 \leq \ell \leq L} \{X_s^{(\ell)}\}$
- Subgraph topological data: D_s , $A_s^\#$ and \mathbb{M}_a
- Pre-computed vector sum for pairs in \mathbb{M}_a
- Intermediate feature aggregation results
- Model weights: $\bigcup_{1 \leq \ell \leq L} \{W_\circ^{(\ell)}, W_\star^{(\ell)}\}$ and W_{MLP}
- Gradient information: $\frac{\partial \mathcal{L}}{\partial X_s^{(\ell)}}$ and $\frac{\partial \mathcal{L}}{\partial X_s^{(\ell-1)}}$, for any $2 \leq \ell \leq L$
- Optimizer specific auxiliary data
- Other data (i.e., tile buffer in Section 4.3)

Note that we only need to store gradients with respect to activations of consecutive layers. When calculating $\frac{\partial \mathcal{L}}{\partial W_\circ^{(\ell)}}$ and $\frac{\partial \mathcal{L}}{\partial W_\star^{(\ell)}}$, the data $\frac{\partial \mathcal{L}}{\partial X_s^{(\ell)}}$ and $\frac{\partial \mathcal{L}}{\partial X_s^{(\ell+1)}}$ are stored on chip. When calculating $\frac{\partial \mathcal{L}}{\partial X_s^{(\ell-1)}}$, we overwrite $\frac{\partial \mathcal{L}}{\partial X_s^{(\ell+1)}}$ with the newly generated $\frac{\partial \mathcal{L}}{\partial X_s^{(\ell-1)}}$. Also note that the optimizer specific auxiliary data depends on the optimizer used. For vanilla gradient descent, no auxiliary data is needed. For gradient descent with momentum [27], we need to store the gradient with respect to model weights for the previous minibatch. For Adam optimizer [17], we need to store more gradient related data (e.g., first moment estimate, second moment estimate, etc.). No matter what optimizer is used, the size of the auxiliary data is comparable to the size of model weights.

Regarding the processing pipeline on-chip, there are two main computation modules to perform feature aggregation and weight transformation. In Figure 3, the Feature Aggregation module consists of a 1D accumulator array to sum the node vectors. The Weight Transformation module consists of a 2D systolic array to compute the dense matrix product between X and W . The two modules are reused for the computation of the L graph convolutional layers, and the Weight Transformation module is also reused for the MLP layer. The various BRAM buffers stores the data listed above. During forward pass, when computing layer ℓ , the Feature Aggregation and Weight Transformation modules read from the buffer of $X_s^{(\ell-1)}$, and write into $X_s^{(\ell)}$. In backward pass, for layer ℓ , the two computation modules read the buffer of $X_s^{(\ell-1)}$, and read / write into the two gradient buffers in a ping-pong fashion.

The $X_s^{(\ell)}$ and $\frac{\partial \mathcal{L}}{\partial X_s^{(\ell)}}$ buffers use feature major data layout. So each cycle, a buffer can provide the full feature vector of one node. For a BRAM block ($36\text{bits} \times 1\text{K}$) of the Xilinx Alveo board we use, we store feature values (32-bit floating point) of 1K different nodes.

4.2 Feature Aggregation Module

This module performs feature aggregation in three steps. The first pre-computation step calculates the vector sum of node pairs in \mathbb{M}_a , and stores the results in the buffer for $X_{\mathbb{M}}$. The second step computes $A_s^\# \cdot X_s^{(\ell)}$ by reading $X_{\mathbb{M}}$ and $X_s^{(\ell)}$. The final step applies the scaling coefficient based on D_s . The aggregated features are written

into a temporary buffer to be consumed by the Weight Transformation module. For the below discussion, we ignore step 3 since its complexity is low (i.e., $\mathcal{O}(|\mathcal{V}_s|)$) compared with the other steps (i.e., $\mathcal{O}(|\mathcal{E}_s|)$). Regarding steps 1 and 2, since the feature length is large, we explore parallelism along the feature dimension. In this case, a simple 1D accumulator array of size $P_{\text{agg}} = \max_{0 \leq \ell \leq L-1} f^{(\ell)}$ is sufficient. During pre-computation, pairs of node indices are read sequentially from \mathbb{M}_a . Vector sum of each pair consumes two cycles. Similarly, during propagation of $A_s^\#$, indices in the neighbor lists are read sequentially. Note that even though \mathbb{M}_a contains pairs of multiple rounds (see last paragraph of Section 3.2), as long as the pre-computation on the earlier rounds finishes before that of later rounds, the memory controller for the accumulator array is straightforward. For example, assume the original subgraph nodes are indexed continuously from 1 to $|\mathcal{V}_s|$ and the new nodes generated in the first round are indexed from $|\mathcal{V}_s| + 1$ to $|\mathcal{V}_s| + |\mathcal{M}_a|$. The matching in the second round, \mathcal{M}'_a , can only contain indices from 1 to $|\mathcal{V}_s| + |\mathcal{M}_a|$. If the second round starts after the first round has finished, all features required by \mathcal{M}'_a are ready. So the accumulator array can directly read $X_s^{(\ell)}$ and $X_{\mathbb{M}}$, and continue filling $X_{\mathbb{M}}$ with new features indexed from $|\mathcal{V}_s| + |\mathcal{M}_a| + 1$ to $|\mathcal{V}_s| + |\mathcal{M}_a| + |\mathcal{M}'_a|$.

Remark. To further increase parallelism, we can aggregate vectors of multiple neighbors in the same cycle. Then the accumulator array would be replaced by an accumulator tree, and the buffer would be further partitioned to reduce bank conflicts in BRAM. In this case, challenges such as load-balance would emerge. Fortunately, for most target FPGA devices, parallelism of just $\max f^{(\ell)}$ is sufficient. We discuss this in detail in Section 5.3 and 5.4. Also, we evaluate the storage overhead due to $X_{\mathbb{M}}$ in Section 6.2.

4.3 Weight Transformation Module

This module performs weight transformation of either a GCN layer (i.e., $h(\cdot)$ function in Figure 1) or a MLP layer. The main operation is multiplication of dense matrices. We use a 2D systolic array to execute the blocked matrix multiplication algorithm. Let the dimension of the systolic array be P_{sys} (where $P_{\text{sys}} \ll f^{(\ell)}$). So the computation parallelism of this module is P_{sys}^2 , and each cycle $2P_{\text{sys}}$ data are streamed in the array. Next we specify the data access pattern. Suppose we compute $X \cdot W$, where $X \in \mathbb{R}^{|\mathcal{V}_s| \times f}$ and $W \in \mathbb{R}^{f \times f'}$. Then the X buffer is depth-wise partitioned to tiles of layout $P_{\text{sys}} \times f$, and the W buffer is width-wise partitioned to tiles of layout $f \times P_{\text{sys}}$. There are $\frac{|\mathcal{V}_s|}{P_{\text{sys}}} \times \frac{f'}{P_{\text{sys}}}$ pairs of tiles. For each pair, the systolic array reads a diagonal (or anti-diagonal if matrix is transposed) of the two tiles per cycle. Thus, computing a pair of tiles consumes $f + P_{\text{sys}} - 1$ cycles, and computing the full dot product takes $\frac{|\mathcal{V}_s|}{P_{\text{sys}}} \cdot \frac{f'}{P_{\text{sys}}} \cdot (f + P_{\text{sys}} - 1) \approx \frac{1}{P_{\text{sys}}^2} \cdot |\mathcal{V}_s| \cdot f \cdot f'$ cycles. To perform ReLU in the forward pass, we only need to add a comparator in each PE of the systolic array. When the results for a pair of tiles are ready, the PEs clip negative values to zero and append a status bit to indicate whether or not the clipping occurs. The `mask` function in the backward pass can thus be implemented in a simple way based on the `True` or `False` of the status bit.

Interaction with Feature Aggregation module. When operating on the neighbor weight W_\star , Weight Transformation module reads the

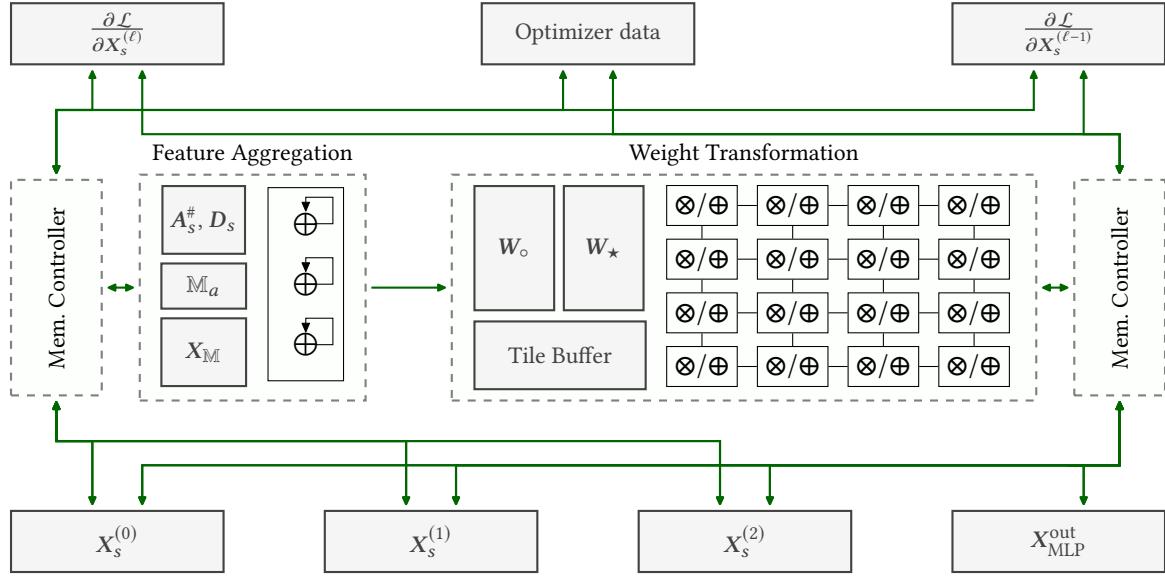


Figure 3: Overview of the processing pipeline on FPGA ($L = 2$)

feature from the buffer filled by the Feature Aggregation module. When operating on the self weight W_o , conflicts may occur since both modules read from the X buffer. We add a small tile buffer in the Weight Transformation module to *completely avoid read conflicts*. The tile buffer of size $P_{sys} \times f$ stores a tile of X , and is filled in f cycles. Data in the tile buffer stay for $(f + P_{sys} - 1) \cdot \frac{f'}{P_{sys}}$ cycles to enumerate all tiles of W . Read conflicts can only happen during the filling of the tile buffer. And we simply stall the Feature Aggregation module in this short period of time. Since $f' \gg P_{sys}$, the pipeline stall has negligible impact on performance. The above analysis is based on the forward pass operation. In the backward pass, we swap the dimensions of $|\mathcal{V}_s|$ and f , and the same conclusion holds.

4.4 Scheduling

Scheduling between CPU and FPGA. CPU samples the subgraph \mathcal{G}_s , transforms it to $\mathcal{G}_s^{\#}$, and calculates softmax, loss and the corresponding gradients. These are shown in light blue blocks of Figure 4. FPGA handles majority of the computation workload in the forward and backward pass, as shown in light green. Communication between CPU and FPGA, as specified in Section 4.1, is in dark blue. Notice that the subgraphs are independently sampled for each minibatch. Thus, CPU can prepare \mathcal{G}_s and $\mathcal{G}_s^{\#}$ for the next minibatch, while simultaneously FPGA is training the current one. This explains the overlap between step 7 and 2. Parallelism can be explored by multiple cores of the CPU. The C cores can process subgraphs for the next C minibatches in parallel without any dependency.

Scheduling of FPGA modules. In Figure 4, we only show the scheduling of the GCN forward pass. Scheduling of MLP and the backward pass can be analogously designed. Computation on neighbor weights depends on the aggregated feature and computation on self weights has no dependency. Thus, we overlap the operations of a and b. The avoidance of read conflicts between a and b is discussed

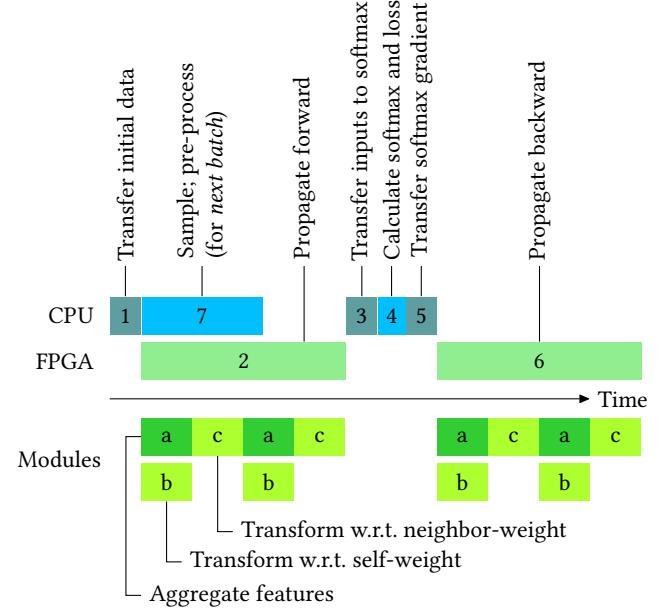


Figure 4: Per-minibatch scheduling between CPU and FPGA (up), and between FPGA computation modules (down).

in Section 4.3. During operation c, Feature Aggregation module is idle. We analyze its impact on DSP utilization in Section 5.4.

5 PERFORMANCE ANALYSIS

To simplify notation, we assume $f^{(\ell)} = f$, $\forall 0 \leq \ell \leq L$. So the weight matrices $W_o^{(\ell)} \in \mathbb{R}^{f \times \frac{1}{2}f}$ and $W_{\star}^{(\ell)} \in \mathbb{R}^{f \times \frac{1}{2}f}$, where the

$\frac{1}{2}$ factor is due to the concatenation in the forward pass. For the classifier, we assume a single layer MLP, with $W_{\text{MLP}} \in \mathbb{R}^{f \times f}$.

Regarding FPGA resources, we assume accumulators and multipliers are implemented by DSPs and have the same hardware cost. We assume the target FPGA can implement R_{DSP} number of accumulators / multipliers, store R_{BRAM} words, and communicate R_{BW} words with external memory per cycle. Here a word represents an element of the feature vectors or the weight matrices.

Training performance depends on the parameters related to:

- Minibatch and GCN: $|\mathcal{V}_s|$, \bar{d} and f
- Redundancy reduction: γ_{add} , γ_{read} and $\sum_{\mathcal{M} \in \mathbb{M}_a} |\mathcal{M}|$
- FPGA architecture: P_{agg} , P_{sys}

We also use the fact that $\bar{d} \ll f \ll |\mathcal{V}_s|$ to simplify analysis.

5.1 Computation

Each graph convolutional layer performs feature aggregation 3 times and product on weights 6 times. Two of the feature aggregation operate on length- f vectors, and the other one on length- $\frac{1}{2}f$ vectors. Since feature aggregation is parallelized along feature dimension only, it takes exactly $\gamma_{\text{read}} \cdot |\mathcal{V}_s| \cdot \bar{d} \cdot f / P_{\text{agg}}$ cycles to aggregate length- f features. All six products on weights have the same complexity, and each takes $\frac{1}{2} \cdot |\mathcal{V}_s| \cdot f^2 / P_{\text{sys}}^2$ cycles. By the schedule in Figure 4, to hide the feature aggregation time, we have:

$$\gamma_{\text{read}} \cdot |\mathcal{V}_s| \cdot \bar{d} \cdot f \cdot \frac{1}{P_{\text{agg}}} = \left(1 - \frac{2P_{\text{sys}}}{f}\right) \cdot \frac{1}{2} \cdot |\mathcal{V}_s| \cdot f^2 \cdot \frac{1}{P_{\text{sys}}^2} \quad (3a)$$

$$P_{\text{agg}} + 2P_{\text{sys}}^2 = R_{\text{DSP}} \quad (3b)$$

$$P_{\text{agg}} \leq f \quad (3c)$$

where factor $1 - \frac{2P_{\text{sys}}}{f}$ is due to the pipeline stall analyzed in Section 4.3. Solving Equations 3a and 3b under the constraint of 3c gives the architectural parameters P_{agg}^* , P_{sys}^* . Under reasonable values of f , $|\mathcal{V}_s|$, \bar{d} , γ_{read} and R_{DSP} , the solutions P_{agg}^* and P_{sys}^* always exist (see also Section 5.3). Total FPGA cycles* to complete one minibatch is:

$$T_{\text{batch}} = (3L + 3) \cdot |\mathcal{V}_s| \cdot f^2 \cdot \frac{1}{(P_{\text{sys}}^*)^2} \quad (4)$$

5.2 Communication

On-chip storage. All data listed in Section 4.1 have to fit on-chip. Index data $A_s^{\#}$ and \mathbb{M}_a and coefficient D are negligible compared with feature data (since $\bar{d} \ll f$). Size of the buffer for $X_{\mathbb{M}}$ is $f \cdot \sum_{\mathcal{M} \in \mathbb{M}_a} |\mathcal{M}|$. Size of the buffer between Feature Aggregation and Weight Transformation modules is $f \cdot |\mathcal{V}_s|$. Size of the tile buffer in Weight Transformation module is $P_{\text{sys}} \cdot |\mathcal{V}_s|$. Weights for each layer takes f^2 . Under gradient descent with momentum, the optimizer requires additional f^2 storage per layer for the auxiliary data. Thus,

$$(L + 5) f |\mathcal{V}_s| + f \sum_{\mathcal{M} \in \mathbb{M}_a} |\mathcal{M}| + P_{\text{sys}} |\mathcal{V}_s| + 2(L + 1) f^2 \leq R_{\text{BRAM}} \quad (5)$$

*We ignore the number of cycles to apply gradients to weights (elementwise operation), since it is negligible compared to the number of cycles to calculate gradients.

By adjusting θ and number of rounds (see Algorithm 2 and Section 3.2), we control $\sum |\mathcal{M}|$ to meet a pre-defined budget (say $|\mathcal{V}_s|$). Then, we tune the graph sampler so that $|\mathcal{V}_s|$ satisfies inequality 5.

Off-chip accesses. Ignoring D_s , $A_s^{\#}$ and \mathbb{M}_a , CPU-FPGA data transfer include the initial features and the MLP output features and gradients. Thus, β , the ratio between on-chip computation and off-chip communication, is lower bounded[†] by f , indicating high reuse of on-chip data, and low overhead in external memory access.

5.3 Load-Balance

If feature aggregation is parallelized along feature dimension only, the full FPGA pipeline is *perfectly* load-balanced, regardless of the graph connection. However, if we would have to aggregate features of multiple nodes in parallel, BRAM access conflicts would cause load-imbalance of the module and degrade training performance.

Load-imbalance is more likely to happen on FPGAs with more DSP resources. The threshold \hat{R}_{DSP} for a design to become load-imbanced can be derived by plugging $P_{\text{agg}} = f$ into Equations 3a and 3b. After simplifying the expression, we have:

$$\hat{R}_{\text{DSP}} > \left(\frac{f}{\bar{d} \gamma_{\text{read}}} \right)^2 \cdot \left(1 + \bar{d} \gamma_{\text{read}} - \sqrt{1 + 2\gamma_{\text{read}} \bar{d}} \right) \quad (6)$$

Feature dimension is often set to 256, 512 or 1024 in the literature. Subgraph degree is often less than 20, so we set $\bar{d} = 15$. The ratio γ_{read} depends on \mathcal{G}_s , and we set it to 0.7. With such parameter values, $\hat{R}_{\text{DSP}} > 4048$, meaning that there have to be at least 4048 multipliers / accumulators on chip for our FPGA pipeline to become load-imbanced. Even the largest FPGAs in the market (e.g., Xilinx UltraScale+, Intel Stratix-10) does not have such high DSP capacity (considering that data in single precision floating point are necessary to achieve high training accuracy). Note that γ_{read} helps keep the threshold \hat{R}_{DSP} high. If $\gamma_{\text{read}} = 1$, \hat{R}_{DSP} is only 3038.

5.4 DSP Utilization

Since load-balance is achieved and BRAM conflicts are eliminated, we can analytically derive the DSP utilization for any \mathcal{G}_s . From the timeline of Figure 4, the systolic array is idle when CPU is communicating with FPGA (ignoring the CPU computation time on softmax and loss). The accumulator array of Feature Aggregation module is idle during the stall and during computation on neighbor weights. Using the fact that computation-communication ratio $\beta > f$ (see Section 5.2), we derive the overall DSP utilization μ as:

$$\mu > \mu' \cdot \frac{1}{1 + \mu' \cdot \frac{R_{\text{DSP}}}{f \cdot R_{\text{BW}}}} \quad (7)$$

where $\mu' = \frac{1}{R_{\text{DSP}}} \cdot \left(2 \left(P_{\text{sys}}^* \right)^2 + \frac{5}{12} P_{\text{arr}}^* \left(1 - \frac{2P_{\text{sys}}^*}{f} \right) \right)$, and P_{arr}^* , P_{sys}^* satisfy Equations 3a and 3b. The $\frac{5}{12}$ factor is due to the aggregation of length- $\frac{1}{2}f$ features in the backward pass. If we can reduce the redundancy up to $\gamma_{\text{read}} > 1 - \frac{2P_{\text{sys}}^*}{f} > 1 - \frac{\sqrt{2R_{\text{DSP}}}}{f}$, we can then lower-bound μ' by $\frac{1}{1+d/f}$. With the typical parameter values specified in

[†]To simplify expression, we only consider the computation of the systolic array on graph convolutional layers. Therefore, f is just a lower bound on the ratio.

Section 5.3, $\mu' > 95\%$. For most FPGA platforms, $\frac{R_{\text{DSP}}}{f \cdot R_{\text{BW}}} \ll 1$, so overall DSP utilization μ is close to 1. For example, on Xilinx Alveo U200 board connected to CPU via PCIe 3.0 x16, we have $\mu > 93\%$.

6 EXPERIMENTS

We evaluate on a Xilinx Alveo U200 board hosted by a 40-core Xeon server (E5-2698 v4 @2.2GHz, hyper-threaded). CPU-FPGA communication is via PCIe 3.0 x16 (peak bandwidth: 15.8GB/s). The UltraScale+ XCU200 FPGA on Alveo has 5867 DSP slices, 1766 36Kb block RAMs and 800 288Kb Ultra RAMs[‡]. For Float32, on-chip RAM can store 8166K words, and DSPs can implement 1173 accumulators or 1955 multipliers. We use Vivado 2018.3 for synthesis.

Table 2: Dataset statistics

Dataset	Nodes	Edges	Attribute	Classes	Train/Val/Test
PPI	14,755	225,270	50	121	0.66/0.12/0.22
Reddit	232,965	11,606,919	602	41	0.66/0.10/0.24
Yelp	716,847	6,977,410	300	100	0.75/0.15/0.10

We use three standard, large benchmark graphs [5, 6, 14, 15, 34] to evaluate the training speed and accuracy. In Table 2, “Attribute” specifies the initial feature length (i.e., $f^{(0)} = |X_v^{(0)}|$). “Classes” specifies total number of node classes. “Train/Val/Test” splitting follows [34]. For all the datasets, following [5, 6, 14, 15, 34], the GCN has two graph convolutional layers ($L = 2$) and one MLP layer in the classifier. The hidden dimension is $f^{(\ell)} = 256$, where $\ell = 1, 2$.

We compare with the baseline [34] since GraphACT uses the same minibatch algorithm as [34]. For [34], we run the public implementation[§] by Python (version 3.6.7) and Tensorflow (version 1.12). We run [34] on both the CPU and GPU platforms. Table 4 shows the hardware specification of the hardware. Parallelization of [34] on CPU is via the Intel MKL library. Thus, we add the `--mkl` flag to run [34] on CPU. Parallelization of [34] on GPU is directly via Tensorflow. The implementation of GraphACT is open-sourced[¶].

6.1 Training Accuracy

Under the same hyper-parameters, the accuracy and convergence of GraphACT are identical [34]. However, since we set the subgraph size under the BRAM constraint (Equation 5), it is necessary to evaluate accuracy under our settings. Table 3 summarizes the comparison, where except the subgraph size, all hyper-parameters of Para-GCN and GraphACT are the same. The popular GraphSAGE does not sample subgraphs, and we use its default minibatch size of 512. Comparing with [14], we achieve much higher accuracy on PPI, and comparable accuracy on Reddit and Yelp after similar number of epochs. Comparing with [34], reducing the subgraph size $|\mathcal{V}_s|$ has negligible impact on the overall accuracy and convergence rate.

6.2 Redundancy Reduction

After transforming \mathcal{G}_s to $\mathcal{G}_s^\#$, redundancy in computation and communication are reduced at the cost of extra on-chip buffer to store

the vector sums of \mathbb{M}_a . Figure 5 shows such tradeoff, as well as the effectiveness of redundancy reduction. Each “ \times ” marker corresponds to one round of redundancy reduction, and we run for 5 rounds ($\theta = 2$, Algorithm 2). We observe (1) number of reads and additions can be reduced to as low as 60%, at the cost of a buffer less than $2|\mathcal{V}_s|$, (2) the γ_{read} value used in Section 5.3 and 5.4 (to analyze load-balance and DSP utilization) is realistic, and (3) amount of redundancy depends on the graph topology — while our algorithm is effective on many graphs (e.g., PPI, Yelp), exceptions exist (e.g., Reddit). By Figure 5, in our implementation, we set a budget B such that $\sum_{M \in \mathbb{M}_a} |\mathcal{M}| < B \cdot |\mathcal{V}_s|$. For PPI, Reddit, Yelp, $B = 2, 0.5, 0.5$.

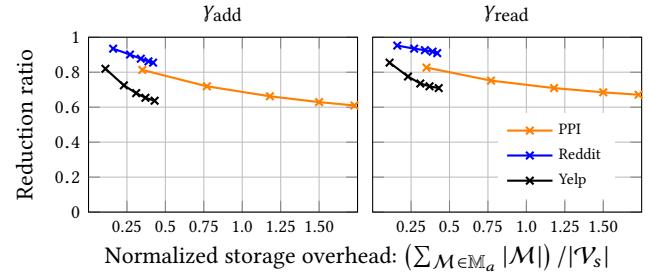


Figure 5: Reduced redundancy vs. Storage overhead

6.3 Comparison with State-of-the-Art

Table 4 shows the training speed comparison with the state-of-the-art [34]. All implementations use single precision floating point for weights and features. In our design, the DSP and BRAM resources are heavily utilized. We set $P_{\text{sys}} = 24$ and $P_{\text{arr}} = 128$ for all datasets. For the multi-core implementation, “L3” shows the aggregated L3-cache size and “Off-chip bandwidth” shows the peak CPU-main memory data transfer speed. For the GPU implementation, “Off-chip bandwidth” is the same as the proposed implementation since the GPU is connected via PCIe 3.0 x16. “Total convergence time” includes the time for graph sampling, pre-processing (for redundancy reduction), GCN forward/backward propagation and data transfer from/to main memory. It excludes the time for initial data loading from the disk and Tensorflow initialization (for allocating device resources and building computation graphs).

Comparing with the CPU baseline, we achieve 12 \times to 15 \times speedup. Apart from the overhead of the Python language, inefficiency of the CPU baseline is mainly due subgraph feature aggregation operation. We observe that although feature aggregation requires less than 10% of the multiplication/addition of weight transformation, the CPU spends about equal amount of time on the two operations (see also Figure 3.D of [34]). Comparing with the GPU baseline, our design converges slightly faster (1.1 \times to 1.5 \times). The theoretical peak performance of the GPU (9.3 TFLOPS for Float32 [2]) is much higher than that of the FPGA (1.8 TFLOPS for Float32 [3][¶]). Inefficiency of the GPU baseline is mainly due to the sub-optimal Tensorflow implementation of sparse matrix multiplication. Note that the CPU in our proposed design only executes light weight operations (Section 3.2 and 4.4). Thus, the training time improvement mainly comes from the highly optimized FPGA architecture.

[‡]In this paper, we refer to block RAM and Ultra RAM by a general term “BRAM”.

[§]Baseline code: <https://github.com/ZimpleX/gcn-ipdps19>, commit a6f531.

[¶]GraphACT code: <https://github.com/GraphSAINT/GraphACT>

[¶]Each Float32 multiplier consumes 3 DSPs. The max DSP frequency is 775MHz.

Table 3: Comparison of convergence and test set accuracy (F1-micro score)

Method	PPI			Reddit			Yelp		
	Accuracy	Epochs	Subgraph size	Accuracy	Epochs	Subgraph size	Accuracy	Epochs	Subgraph size
GraphSAGE [14]	0.603	71	–	0.953	6	–	0.593	10	–
Para-GCN [34]	0.685	120	5500	0.957	5	8000	0.606	8	4000
GraphACT	0.678	104	4000	0.952	5	2750	0.598	7	2750

Table 4: Comparison of resources and training time with state-of-the-art implementations

	Xilinx Alveo U200			Xeon E5-2698 v4 server			NVIDIA Tesla P100		
	PPI	Reddit	Yelp	PPI	Reddit	Yelp	PPI	Reddit	Yelp
Data type	Float32	Float32	Float32	Float32	Float32	Float32	Float32	Float32	Float32
Frequency (GHz)	0.2	0.2	0.2	2.2	2.2	2.2	1.2	1.2	1.2
DSP slices / CPU cores / CUDA cores	5632 (96%)	5632 (96%)	5632 (96%)	40	40	40	3584	3584	3584
BRAM / L3 / HBM2 (MB)	34.6 (96%)	32.8 (91%)	27.3 (75%)	102	102	102	16280	16280	16280
Off-chip bandwidth (GB/sec)	15.8	15.8	15.8	136.6	136.6	136.6	15.8	15.8	15.8
Total convergence time (sec)	9.6	7.6	23.4	151.4	95.5	359.4	10.6	11.4	30.4

Although the baseline training time may be further improved by re-implementation using C++/CUDA, inefficiency in CPU and GPU due to sparse feature aggregation may not be easily eliminated.

7 DISCUSSION

This work proposed several hardware-software optimizations for GCN training on CPU-FPGA heterogeneous platforms. We discuss these optimizations and their applicability to various platforms.

Design challenges. The key to accelerating GCN training is to address the challenges of memory access and load-balance. The solutions for these differ on various platforms. Regarding memory access, the solution on FPGA must optimize both on-chip and off-chip accesses. For data in BRAMs, we need to increase their reuse so as to reduce off-chip communication. We also need to reduce bank access conflicts to reduce pipeline stalls. In **GraphACT**, we reduce off-chip communication by setting the subgraph size based on the BRAM capacity (Section 5.2). We eliminate on-chip access conflicts by appropriately parallelizing the feature aggregation operation (Section 4.2) and buffering data tiles between computation modules (Section 4.3). On the other hand, for CPU and GPU, the critical issue is to enhance data access locality since the memory hierarchy is not explicitly exposed to the programmer. Considering that the full graph fits in the CPU main memory or GPU global memory (sizes of which range from 10^1 to 10^3 GB), the memory access challenge on CPU and GPU may be addressed by software node-reordering or by hyper-threading. Regarding load-balance, CPU and GPU can decouple the operations of feature propagation and weight transformation, and then adopt separate strategies to balance them. However, on FPGA, since the two operations are processed by a single pipeline, an integrated strategy needs to simultaneously balance the very different operations (one on dense tensors and the other on sparse tensors). While memory access and load-balance are challenging to optimize, a carefully designed FPGA pipeline can

achieve much higher efficiency than the CPU or GPU solutions. This is due to the ability of FPGA to customize computation modules and to control memory accesses in an explicit and fine-grained manner.

Remark on redundancy reduction. Although redundancy reduction is an algorithm-level optimization independent of the platform, it may not be directly applicable to CPU to improve performance. Recall that redundancy reduction (Section 3.2) constructs a subgraph $\mathcal{G}_s^\#$ with less edges yet *more* nodes than \mathcal{G}_s . Thus, feature aggregation on $\mathcal{G}_s^\#$ may have worse data locality than \mathcal{G}_s . On a CPU, even though the L3-cache can hold the node features of $\mathcal{G}_s^\#$, the overall feature aggregation performance may not benefit from the reduced computation load due to the potential increase in L1- or L2-cache misses. On the other hand, data access locality of **GraphACT** does not affect FPGA performance, as long as data of one minibatch completely fits in BRAM. Under the proposed architecture and data layout (Section 4), BRAM will always provide the requested node feature of $\mathcal{G}_s^\#$ to the Feature Aggregation module within one cycle.

8 CONCLUSION

We accelerated GCN training on CPU-FPGA heterogeneous platforms. By multiple software-hardware co-optimizations, we achieved conflict free BRAM access, load-balance and high DSP utilization.

We plan to extend **GraphACT** to accelerate inference, where the GCN operates on large, un-sampled graphs. The redundancy reduction and the computation modules of **GraphACT** can be applied to achieve high DSP utilization. In addition, since the inference graph is much larger than the training subgraphs, we need to optimize off-chip communication by partitioning and scheduling algorithms.

ACKNOWLEDGMENTS

This work was supported by US NSF under grant No. CNS-1643351, Intel Strategic Research Alliance and the Defense Advanced Research Projects Agency under contract No. FA8750-17-C-0086.

REFERENCES

- [1] [n. d.]. Gradient of softmax. <https://deepnotes.io/softmax-crossentropy>. ([n. d.]). Accessed: 2019-09-08.
- [2] [n. d.]. NVIDIA Tesla P100 peak performance. <https://images.nvidia.com/content/tesla/pdf/nvidia-tesla-p100-PCIe-datasheet.pdf>. ([n. d.]). Accessed: 2019-11-30.
- [3] [n. d.]. Xilinx Alveo U200 peak performance. <https://www.xilinx.com/products/boards-and-kits/alveo/u200.html#specifications>. ([n. d.]). Accessed: 2019-11-30.
- [4] Anonymous. 2020. GraphSAINT: Graph Sampling Based Inductive Learning Method. In *Submitted to International Conference on Learning Representations*. <https://openreview.net/forum?id=Bje8pkHFwS> under review.
- [5] Jie Chen, Tengfei Ma, and Cao Xiao. 2018. FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling. In *International Conference on Learning Representations (ICLR)*.
- [6] Jianfei Chen, Jun Zhu, and Le Song. 2018. Stochastic Training of Graph Convolutional Networks with Variance Reduction.. In *ICML*. 941–949.
- [7] Guohua Dai, Yuze Chi, Yu Wang, and Huazhong Yang. 2016. Fppg: Graph processing framework on fpga a case study of breadth-first search. In *Proceedings of the 2016 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. ACM, 105–110.
- [8] Guohua Dai, Tianhao Huang, Yuze Chi, Ningyi Xu, Yu Wang, and Huazhong Yang. 2017. Foregraph: Exploring large-scale graph processing on multi-fpga architecture. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. ACM, 217–226.
- [9] Florent De Dinechin, Pedro Echeverria, Marisa López-Vallejo, and Bogdan Pasca. 2013. Floating-point exponentiation units for reconfigurable computing. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)* 6, 1 (2013), 4.
- [10] Tong Geng, Tianqi Wang, Ahmed Sanaullah, Chen Yang, Rui Xu, Rushi Patel, and Martin Herbordt. 2018. FPDeep: Acceleration and load balancing of CNN training on FPGA clusters. In *2018 IEEE 26th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 81–84.
- [11] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- [12] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 855–864.
- [13] Kaiyuan Guo, Shuang Liang, Jincheng Yu, Xuefei Ning, Wenshuo Li, Yu Wang, and Huazhong Yang. 2019. Compressed CNN Training with FPGA-based Accelerator. In *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '19)*. ACM, New York, NY, USA, 189–189. <https://doi.org/10.1145/3289602.3293977>
- [14] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems 30*. 1024–1034.
- [15] Wenbing Huang, Tong Zhang, Yu Rong, and Junzhou Huang. 2018. Adaptive sampling towards fast graph representation learning. In *Advances in Neural Information Processing Systems*. 4558–4567.
- [16] Soroosh Khoram, Jiliang Zhang, Maxwell Strange, and Jing Li. 2018. Accelerating Graph Analytics by Co-Optimizing Storage and Access on an FPGA-HMC Platform. In *Proceedings of the 2018 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '18)*. ACM, New York, NY, USA, 239–248. <https://doi.org/10.1145/3174243.3174260>
- [17] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [18] Thomas N. Kipf and Max Welling. 2016. Semi-Supervised Classification with Graph Convolutional Networks. *CoRR* abs/1609.02907 (2016). arXiv:1609.02907 <http://arxiv.org/abs/1609.02907>
- [19] Yufei Ma and et al. 2017. Optimizing Loop Operation and Dataflow in FPGA Acceleration of Deep Convolutional Neural Networks. In *Proceedings of the 2017 ACM/SIGDA Intl. Symposium on Field-Programmable Gate Arrays (FPGA '17)*.
- [20] Yuzhe Ma, Haoxing Ren, Brucek Khailany, Harbinder Sikka, Lijuan Luo, Karthikeyan Natarajan, and Bei Yu. 2019. High Performance Graph Convolutional Networks with Applications in Testability Analysis. In *Proceedings of the 56th Annual Design Automation Conference 2019 (DAC '19)*. ACM, New York, NY, USA, Article 18, 6 pages. <https://doi.org/10.1145/3316781.3317838>
- [21] AM Mansour, AM El-Sawy, MS Aziz, and AT Sayed. 2015. A new hardware implementation of base 2 logarithm for FPGA. (2015).
- [22] S. Micali and V. V. Vazirani. 1980. An $O(v|v| c |E|)$ algorithm for finding maximum matching in general graphs. In *21st Annual Symposium on Foundations of Computer Science (sfcs 1980)*, 17–27. <https://doi.org/10.1109/SFCS.1980.12>
- [23] Hiroki Nakahara, Akira Jinguji, Masayuki Shimoda, and Shimpei Sato. 2019. An FPGA-based Fine Tuning Accelerator for a Sparse CNN. In *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. ACM, 186–186.
- [24] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 701–710.
- [25] R. Rajat, H. Zeng, and V. Prasanna. 2019. A Flexible Design Automation Tool for Accelerating Quantized Spectral CNNs. In *2019 29th International Conference on Field-Programmable Logic and Applications (FPL)*. 144–150. <https://doi.org/10.1109/FPL.2019.00031>
- [26] Bruno Ribeiro and Don Towsley. 2010. Estimating and sampling graphs with multidimensional random walks. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. ACM, 390–403.
- [27] David E Rumelhart, Geoffrey E Hinton, Ronald J Williams, et al. [n. d.]. Learning representations by back-propagating errors. *Cognitive modeling* 5, 3 ([n. d.]), 1.
- [28] Shreyas Kolala Venkataramanaiah, Yufei Ma, Shihui Yin, Eriko Nurvitadhi, Aravind Dasu, Yu Cao, and Jae-sun Seo. 2019. Automatic Compiler Based FPGA Accelerator for CNN Training. *arXiv preprint arXiv:1908.06724* (2019).
- [29] Chun Wang, Shirui Pan, Guodong Long, Xingquan Zhu, and Jing Jiang. 2017. Mgae: Marginalized graph autoencoder for graph clustering. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. ACM, 889–898.
- [30] Xuechao Wei, Cody Hao Yu, Peng Zhang, Youxiang Chen, Yuxin Wang, Han Hu, Yun Liang, and Jason Cong. 2017. Automated Systolic Array Architecture Synthesis for High Throughput CNN Inference on FPGAs. In *Proceedings of the 54th Annual Design Automation Conference 2017 (DAC '17)*. ACM, New York, NY, USA, Article 29, 6 pages. <https://doi.org/10.1145/3061639.3062207>
- [31] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. 2018. Graph Convolutional Neural Networks for Web-Scale Recommender Systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '18)*. 10.
- [32] Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L. Hamilton, and Jure Leskovec. 2018. Hierarchical Graph Representation Learning with Differentiable Pooling. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS'18)*. Curran Associates Inc., USA, 4805–4815. <http://dl.acm.org/citation.cfm?id=3327345.3327389>
- [33] Hanqing Zeng, Ren Chen, Chi Zhang, and Viktor Prasanna. 2018. A Framework for Generating High Throughput CNN Implementations on FPGAs. In *Proceedings of the 2018 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '18)*. ACM, New York, NY, USA, 10.
- [34] H. Zeng, H. Zhou, A. Srivastava, R. Kannan, and V. Prasanna. 2019. Accurate, Efficient and Scalable Graph Embedding. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 462–471. <https://doi.org/10.1109/IPDPS.2019.00056>
- [35] Chen Zhang, Guangyu Sun, Zhenman Fang, Peipei Zhou, Peichen Pan, and Jason Cong. 2018. Caffeine: Towards uniformed representation and acceleration for deep convolutional neural networks. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2018).
- [36] Jiliang Zhang, Soroosh Khoram, and Jing Li. 2017. Boosting the Performance of FPGA-based Graph Processor Using Hybrid Memory Cube: A Case for Breadth First Search. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '17)*. ACM, New York, NY, USA, 207–216. <https://doi.org/10.1145/3020078.3021737>
- [37] J. Zhang, W. Zhang, G. Luo, X. Wei, Y. Liang, and J. Cong. 2019. Frequency Improvement of Systolic Array-Based CNNs on FPGAs. In *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*. 1–4. <https://doi.org/10.1109/ISCAS.2019.8720271>
- [38] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*. 5165–5175.
- [39] Xiaofan Zhang, Junsong Wang, Chao Zhu, Yonghua Lin, Jinjun Xiong, Wen-mei Hwu, and Deming Chen. 2018. DNNBuilder: an automated tool for building high-performance DNN hardware accelerators for FPGAs. In *Proceedings of the International Conference on Computer-Aided Design*. ACM, 56.
- [40] Wenlai Zhao, Haohuan Fu, Wayne Luk, Teng Yu, Shaojun Wang, Bo Feng, Yuchun Ma, and Guangwen Yang. 2016. F-CNN: An FPGA-based framework for training Convolutional Neural Networks. In *2016 IEEE 27th International Conference on Application-specific Systems, Architectures and Processors (ASAP)*. IEEE, 107–114.
- [41] S. Zhou, C. Chelmis, and V. K. Prasanna. 2016. High-Throughput and Energy-Efficient Graph Processing on FPGA. In *2016 IEEE 24th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. 103–110. <https://doi.org/10.1109/FCCM.2016.35>