

Solving Large Top-K Graph Eigenproblems with a Memory and Compute-optimized FPGA Design

Francesco Sgherzi*, Alberto Parravicini[†], Marco Siracusa* Marco D. Santambrogio[†]
 Politecnico di Milano, DEIB, Milan, Italy

*{francesco1.sgherzi, marco.siracusa}@mail.polimi.it [†]{alberto.parravicini, marco.santambrogio}@polimi.it,

Abstract—Large-scale eigenvalue computations on sparse matrices are a key component of graph analytics techniques based on spectral methods. In such applications, an exhaustive computation of all eigenvalues and eigenvectors is impractical and unnecessary, as spectral methods can retrieve the relevant properties of enormous graphs using just the eigenvectors associated with the Top-K largest eigenvalues.

In this work, we propose a hardware-optimized algorithm to approximate a solution to the Top-K eigenproblem on sparse matrices representing large graph topologies. We prototype our algorithm through a custom FPGA hardware design that exploits HBM, Systolic Architectures, and mixed-precision arithmetic. We achieve a speedup of 6.22x compared to the highly optimized ARPACK library running on an 80-thread CPU, while keeping high accuracy and 49x better power efficiency.

I. INTRODUCTION

Research in information retrieval and recommender systems has spiked novel interest in spectral methods [1], a class of Machine Learning algorithms able to detect communities in large social and e-commerce graphs, and compute the similarity of graph elements such as users or products [2]. At the core of many spectral methods lies the Top-K eigenproblem for large-scale sparse matrices, i.e. the computation of the eigenvectors associated with the largest eigenvalues (in modulo) of a matrix that stores only non-zero entries (Figure 1). For example, the famous Spectral Clustering algorithm boils down to computing the largest eigenvalues of a sparse matrix representing the graph topology [3]. Despite the rise of theoretical interests for spectral methods, little research has focused on improving the performance and scalability of the Top-K sparse eigenproblem solvers, making them applicable to large-scale graphs.

Most existing high-performance implementations of eigenproblem algorithms operate on dense matrices and are completely unable to process matrices with millions of rows and columns (each encoding, for example, the user’s friends in a social network graph) [4]. Even the highly optimized multi-core implementation of LAPACK requires more than 3 minutes to solve the full eigenproblem on a small graph with $\sim 10^4$ vertices and $\sim 50 \cdot 10^4$ edges on a Xeon 6248, as the eigenproblem complexity scales at least quadratically with the number of vertices in the graph. Many implementations that support sparse matrices, on the other hand, are either forced to compute all the eigenvalues or require an expensive matrix inversion before solving the eigenproblem [5].

The need for high-performance Top-K sparse eigenproblem algorithms goes hand in hand with custom hardware designs

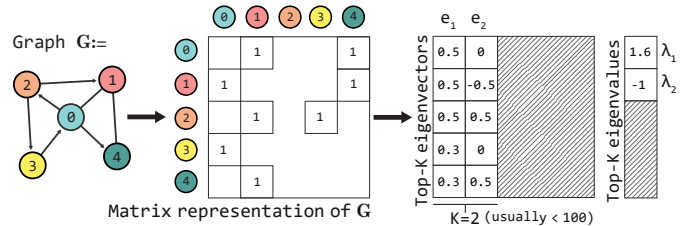


Fig. 1: Top-K eigencomputation of a graph G , represented as a sparse matrix. While G can have millions of vertices, we often need just the Top-K eigenvectors ($K = 2$ in this example).

that can outperform traditional architectures in raw performance and power efficiency, given how applications on-top of Top-K eigenproblem are mostly encountered in data centers. In this work, we tackle both problems by presenting a new Top-K sparse eigensolver whose building blocks are specifically optimized for high-performance hardware designs.

We introduce a novel algorithm to address the Top-K sparse eigenproblem and prototype it through a custom hardware design on an FPGA accelerator card. This is the first FPGA-based Top-K sparse eigensolver, to the best of our knowledge. Our algorithm is a 2-step procedure that combines the Lanczos algorithm (to reduce the problem size) [6] with the Jacobi algorithm (to compute the final eigencomponents) [7], as shown in Figure 2. The Lanczos algorithm, often encountered in Top-K sparse eigensolvers [8], has never been combined with the Jacobi algorithm. Part of the reason lies in their different computational bottlenecks: the Lanczos algorithm demands large memory bandwidth, while the Jacobi algorithm is strongly compute-bound. Our approach exploits the strengths of FPGA accelerator cards and overcomes the limitations of traditional architectures in this class of algorithms.

First, the Lanczos algorithm presents a Sparse Matrix-Vector Multiplication (SpMV) as its main bottleneck, an extremely memory-intensive computation with indirect and fully random memory accesses (Figure 2 (B)). Optimizing SpMV requires high peak memory bandwidth and fine-grained control over memory accesses, without passing through the traditional caching policies of general-purpose architectures. Our hardware design features an iterative dataflow SpMV with multiple Compute Units (CUs), leveraging every High Bandwidth Memory (HBM) channels through a custom memory subsystem that efficiently handles indirect memory accesses.

Then, we introduce a Systolic Array (SA) design for the Jacobi eigenvalue algorithm, a computationally-intensive operation that operates on reduced-size inputs ($K \times K$) (Figure 2 **D**). The Jacobi algorithm maps naturally to a SA that ensures $\mathcal{O}(\log(K))$ convergence, while traditional architectures do not ensure the same degree of performance. CPUs cannot guarantee that all the data are kept in L1 cache and are unlikely to have enough floating-point arithmetic units to parallelize the computation. This results in $\Omega(K^2)$ computational complexity and execution times more than 50 times higher than a FPGA (Section V). Instead, GPUs cannot fill all their Stream Processors, as the input size is much smaller than what is required to utilize the GPU parallelism fully [9].

Moreover, our FPGA-based hardware design employs highly-optimized mixed-precision arithmetic, partially replacing traditional floating-point computations with faster fixed-precision arithmetic. While high numerical accuracy is usually demanded in eigenproblem algorithms, we employ fixed-precision arithmetic in parts of the design that are not critical to the overall accuracy and resort to floating-point arithmetic when required to guarantee precise results.

In summary, we present the following contributions:

- A novel algorithm for approximate resolution of large-scale Top-K sparse eigenproblems (Section III), optimized for custom hardware designs.
- A modular mixed-precision FPGA design for our algorithm that efficiently exploits the available programmable logic and the bandwidth of DDR and HBM (Section IV).
- A performance evaluation of our Top-K eigendecomposition algorithm against the multi-core ARPACK CPU library, showing a speedup of 6.22x and a power efficiency gain of 49x, with a reconstruction error due to mixed-precision arithmetic as good as 10^{-3} (Section V).

II. RELATED WORK

To the best of our knowledge, no prior work optimizes Top-K sparse eigenproblem with custom FPGA hardware designs.

The most well-known large-scale Top-K sparse eigenproblem solver on CPU is the ARPACK library [10], a multi-core Fortran library that is also available in SciPy and MATLAB through thin software wrappers. ARPACK implements the Implicitly Restarted Arnoldi Method (IRAM), a variation of the Lanczos algorithm that supports non-Hermitian matrices. Other sparse eigensolvers provide techniques optimized for specific domains or matrix types, although none is as common as ARPACK [11]–[14].

On GPUs, the cuSOLVER [15] library by Nvidia provides a simple eigensolver based on the shift-inverse method that retrieves only the largest eigenvalue and its eigenvector (i.e. $K = 1$), which is significantly more limited than the general Top-K eigenproblem. The nvGRAPH library [16], also developed by Nvidia, provides an implementation of spectral clustering at whose core lies the Lanczos algorithm. However, the implementation of the inner Lanczos algorithm is not publicly available. To the best of our knowledge, there is no publicly available GPU implementation of the Lanczos algorithm that

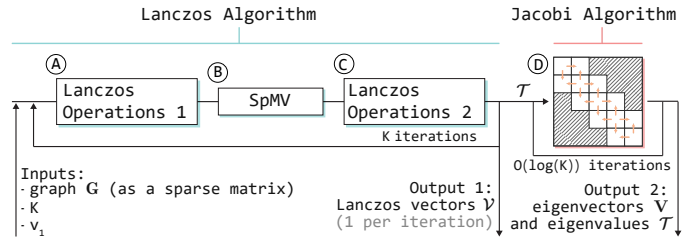


Fig. 2: Steps of our novel Top-K sparse eigenproblem solver, which combines the Lanczos algorithm with a Systolic Array formulation for the Jacobi eigenvalue algorithm.

can solve large scale sparse eigenproblems required by spectral methods. The MAGMA library [17] solves the Top-K sparse eigenproblem through the alternative LOBPCG algorithm [5], which requires multiple iterations (each containing at least one SpMV) to compute even a single eigenvector, to the contrary of the Lanczos algorithm. Other GPU-based Top-K eigensolvers are domain-specific, do not support large-scale inputs, or do not leverage features of modern GPUs such as HBM memory or mixed-precision arithmetic [18], [19]. Eigensolvers for dense matrices are more common on GPUs, as they easily exploit the enormous memory bandwidth of these architectures: Myllykoski et al. [4] focus on accelerating the case of dense high-dimensional matrices (around 10^5 rows) while Cosnauu [9] operates on multiple small input matrices. Clearly, none of the techniques that operate on dense matrices can easily scale to matrices with millions of rows as simply storing them requires terabytes of memory.

Specialized hardware designs for eigensolvers are limited to resolving the full eigenproblem on small dense matrices, through the QR-Householder Decomposition and Jacobi eigenvalue algorithm. Most formulations of the Jacobi algorithm [20], [21] leverage Systolic Array, a major building block of high performance domain-specific architectures from their inception [22] to more recent results [23]–[25]. However, hardware designs of the Jacobi algorithm based on SA cannot scale to large matrices, as the resource utilization scales linearly with the size of the matrix. Implementations of the *QR-Householder* algorithm face similar problems [26], [27] as they also leverage systolic architectures, although research about resource-efficient designs do exist [28].

III. SOLVING THE TOP-K SPARSE EIGENPROBLEM

Algorithms like Spectral Clustering contain as their core step a Top-K sparse eigenproblem, i.e. finding eigenvalues and eigenvectors of sparse matrices representing, for instance, graph topologies with millions of vertices and edges.

Given a sparse square matrix $M \in \mathbb{R}^{n \times n}$ and an integer $K \ll n$ the goal of the Top-K sparse eigenproblem is to find the K eigenvalues with the highest magnitude, and their associated eigenvectors. This is equivalent to computing the approximate decomposition $M \approx M_K = Q_K \Lambda_K Q_K^T$, with $Q_K \in \mathbb{R}^{n \times K}$ and $\Lambda_K \in \mathbb{R}^{K \times K}$. Q_K contains the eigenvectors, while Λ_K is a diagonal matrix containing the

Algorithm 1 Lanczos algorithm for the Top-K eigenvectors

```
1: function LANCZOS( $\mathbf{M}, K, v_1$ )
2:    $\beta_1 \leftarrow 0, v_0 \leftarrow \mathbf{0}_N$   $\triangleright$  Initialization
3:   for  $i$  in  $1, K$  do  $\triangleright$  Main Lanczos loop
4:     if  $i > 1$  then
5:        $\beta_i \leftarrow \|w'_{i-1}\|_2$ 
6:        $v_i \leftarrow w'_{i-1}/\beta_i$   $\triangleright$  Compute new Lanczos vector
7:        $w_i \leftarrow \mathbf{M}v_i$   $\triangleright$  Sparse matrix-vector multiplication
8:        $\alpha_i \leftarrow w_i v_i$ 
9:        $w'_i \leftarrow w_i - \alpha_i v_i - \beta_i v_{i-1}$ 
10:      Orthogonalize  $w'_i$  with respect to  $\mathcal{V}$ 
11:    $\triangleright$  Tridiagonal matrix  $\mathcal{T}$  and Lanczos vectors  $\mathcal{V}$ 
12:   Output  $\{\mathcal{T} = [\alpha_1, \dots, \alpha_K], [\beta_1, \dots, \beta_{K-1}]\}$ 
13:   Output  $\mathcal{V} = [v_1, \dots, v_K]$ 
```

eigenvalues. Indeed, computing all the n eigenvalues of the matrix is intractable for large matrices and redundant for many applications that require only a handful of eigencomponents. For example, Spectral Clustering and many of its variations rely only on the Top-K eigenvectors, with K rarely above ~ 10 .

In this work, we propose a novel algorithm to solve the Top-K sparse eigenproblem, combining the Lanczos algorithm and the Jacobi eigenvalue algorithm. Our technique is particularly suited for highly optimized and modular hardware designs. The first phase leverages the Lanczos algorithm, taking as input the original matrix \mathbf{M} , the number of desired eigencomponents K and an $L2$ -normalized random vector $v_1 \in \mathbb{R}^n$, initialized with values equal to $1/n^2$. The Lanczos algorithm outputs a $K \times K$ symmetric tridiagonal matrix \mathcal{T} (Figure 3) and a set of orthogonal Lanczos vectors $\mathcal{V} \in \mathbb{R}^{K \times n}$. As second step, we apply the Jacobi eigenvalue algorithm to \mathcal{T} . This algorithm transforms \mathcal{T} into a diagonal matrix containing its eigenvalues, and returns a matrix \mathbf{V} with the eigenvectors of \mathcal{T} . Each eigenvalue λ of \mathcal{T} is also an eigenvalue of the original matrix \mathbf{M} . Moreover, if x is the eigenvector of \mathcal{T} associated to λ , then $\mathcal{V}x$ is the eigenvector of \mathbf{M} associated to λ . \mathbf{M}_K can be obtained as $\mathbf{M}_K = (\mathcal{V}\mathbf{V})\mathcal{T}(\mathcal{V}\mathbf{V})^T$, although many applications in spectral analysis only require the Top-K eigenvalues and eigenvectors of \mathbf{M} instead of retrieving \mathbf{M}_K .

A. The Lanczos Algorithm

The Lanczos algorithm retrieves the Top-K eigencomponents of a matrix and is often employed as a building block

$$\begin{bmatrix} \alpha_1 & \beta_1 & 0 & 0 & 0 \\ \beta_1 & \alpha_2 & \beta_2 & 0 & 0 \\ 0 & \beta_2 & \alpha_3 & \beta_3 & 0 \\ 0 & 0 & \beta_3 & \alpha_4 & \beta_4 \\ 0 & 0 & 0 & \beta_4 & \alpha_5 \end{bmatrix}$$

Fig. 3: Example of (5×5) tridiagonal matrix, obtained as output of the Lanczos algorithm for $K = 5$.

Algorithm 2 Jacobi eigenvalue algorithm with Systolic Arrays

```
1: function JACOBI( $\mathcal{T}$ )
2:    $\mathbf{V} \leftarrow \mathbb{1}_K$   $\triangleright$  Identity matrix of size  $K \times K$ 
3:   repeat
4:     for  $i$  in  $1, K/2$  do  $\triangleright$  Diagonal CU
5:        $p_{ii} \leftarrow \mathcal{T}[2i : 2i + 1, 2i : 2i + 1]$ 
6:        $\theta_i \leftarrow \frac{1}{2} \arctan \frac{2\beta}{\alpha - \delta}$ 
7:       Rotate  $p_{ii}$   $\triangleright$  Full equation in Figure 4a
8:       Propagate  $c_i$  and  $s_i$ 
9:     for  $j$  in  $1, K/2 - 2$  do  $\triangleright$  Offdiagonal CU
10:       $i \leftarrow j + 1$ 
11:      Receive  $c_i, c_j, s_i, s_j$  from  $p_{ii}, p_{jj}$ 
12:       $p_{ij} \leftarrow \mathcal{T}[2i : 2i + 1, 2j : 2j + 1]$ 
13:      Rotate  $p_{ij}$   $\triangleright$  Full equation in Figure 4b
14:     for  $i$  in  $1, K/2$  do  $\triangleright$  Eigenvector CU
15:       for  $j$  in  $1, K/2$  do
16:          $v_{ij} \leftarrow \mathbf{V}[2i : 2i + 1, 2j : 2j + 1]$ 
17:         Receive  $c_j, s_j$  from  $p_{jj}$ 
18:         Rotate  $v_{ij}$   $\triangleright$  Full equation in Figure 4c
19:       Permute rows and columns of  $\mathcal{T}$  and  $\mathbf{V}$   $\triangleright$  Figure 5
20:     until  $\mathcal{T}$  becomes diagonal
21:   Output  $\mathcal{T}$   $\triangleright$  Eigenvalues of the input  $\mathcal{T}$ 
22:   Output  $\mathbf{V}$   $\triangleright$  Eigenvectors of the input  $\mathcal{T}$ 
```

of large-scale eigenproblem algorithms [10], [29], [30]. The $K \times K$ output tridiagonal matrix \mathcal{T} is significantly smaller than the input ($K \ll n$) and also simpler in structure, as elements outside of the band enclosing the main diagonal and the ones immediately above and below are zero. Pseudo-code of the algorithm is provided in Algorithm 1. For each of the K iterations, it computes a Lanczos vector v_i by normalizing w'_{i-1} , obtained at the previous iteration (line 6 and Figure 2A). From v_i , we obtain w'_i by projecting the matrix \mathbf{M} into v_i (line 7 and Figure 2B), followed by an orthogonalization (lines 8 to 10 and Figure 2C). The algorithm is highly efficient as each vector v_i is computed in a single iteration, and $K \ll n$.

The Lanczos algorithm is particularly efficient on sparse matrices, as its most expensive operation is an iterative SpMV, bounding its computational complexity to $\mathcal{O}(K \cdot E)$, with E being the number of non zero elements of \mathbf{M} . In our hardware design, we optimize the memory-intensive SpMV computation through multiple independent CUs, so that we can take advantage of all the available 32 HBM channels of a Xilinx Alveo U280 FPGA accelerator card (Section IV-B).

This algorithm is prone to numerical instability as the Lanczos vectors \mathcal{V} can quickly lose pairwise orthogonality if K is very large. To prevent instability, we normalize the input matrix in *Frobenius norm* as eigencomponents are invariant to constant scaling: values of the matrix are in the range $(-1, 1)$, which implies that eigenvalues and eigenvectors are also in the range $(-1, 1)$. This property enables the use of fixed-point arithmetic to improve performance and reduce resource

$$\begin{bmatrix} c_i & s_i \\ -s_i & c_i \end{bmatrix} \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} \begin{bmatrix} c_i & -s_i \\ s_i & c_i \end{bmatrix} = \begin{bmatrix} \alpha' & 0 \\ 0 & \delta' \end{bmatrix}$$

(a) Operations for the Diagonal Processor p_{ii} (Figure 5A).

$$\begin{bmatrix} c_i & s_i \\ -s_i & c_i \end{bmatrix} \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} \begin{bmatrix} c_j & -s_j \\ s_j & c_j \end{bmatrix} = \begin{bmatrix} \alpha' & \beta' \\ \gamma' & \delta' \end{bmatrix}$$

(b) Operations for the Offdiagonal Processor p_{ij} (Figure 5C).

$$\begin{bmatrix} w & x \\ y & z \end{bmatrix} \begin{bmatrix} c_j & -s_j \\ s_j & c_j \end{bmatrix} = \begin{bmatrix} w' & x' \\ y' & z' \end{bmatrix}$$

(c) Operations for the Eigenvector Processor p_{ij} (Figure 5D).

Fig. 4: Operations performed by different processors in the Jacobi eigenvalue Systolic Array architecture. Values c_i and s_i indicate $\cos(\theta_i)$ and $\sin(\theta_i)$, with $\theta_i = \frac{1}{2} \arctan \frac{2\beta}{\alpha-\delta}$.

usage (Section V-C). We further improve numerical stability by adopting a version of the algorithm that reorders operations [31] and reorthogonalizes Lanczos vectors in each iteration [32]. Reorthogonalization (Algorithm 1, line 10) requires $K^2/2$ more operations of cost $\mathcal{O}(n)$, increasing complexity to $\mathcal{O}(K(E + nK^2/2))$. We also introduce the option of performing reorthogonalization every 2 iterations, for a lower overhead of $\mathcal{O}(n(K/2)^2/2)$, with negligible accuracy loss (Section V-C). In practice, execution time is usually dominated by SpMV making reorthogonalization a viable option.

B. The Jacobi Eigenvalue Algorithm

The Jacobi eigenvalue algorithm computes the eigenvalues and eigenvectors of a dense symmetric real matrix. It is an iterative procedure that performs rotations on square submatrices. Each iteration is highly computationally-intensive as it contains $\Omega(K^2)$ trigonometric operations. However, this algorithm is particularly well suited to solve eigenproblems on small tridiagonal matrices. As many matrix values are zero and cannot introduce data-dependencies in rotations, it is possible to parallelize the entire computation at hardware-level.

The Jacobi eigenvalue algorithm has sought many formulations to improve either its parallelism or its resource utilization. The best-known formulation of the algorithm was proposed by Brent and Luk [33] and has been the standard for implementing the algorithm on FPGA to this day [20], [21]. Our design improves this formulation with a more resource-efficient procedure for interchanging rows and columns, and its structure is shown in Figure 5.

We employ a SA design that maps the input matrix as 2×2 submatrices to $K^2/4$ adjacent *processors* (or CU) (Figure 5, (A)). The systolic architecture propagates the rotation angles (B) and the values stored in each processors (E).

Starting from \mathcal{T} , the algorithm set to zero $K/2$ off-diagonal entries per iteration by using rotations. Diagonal processors annihilate β and γ components (Algorithm 2, line 7) with a rotation of angle θ . This angle is propagated (line 8) to the off-diagonal processor (line 9), and to the eigenvector processor that applies the same rotation to the identity matrix (line 14).

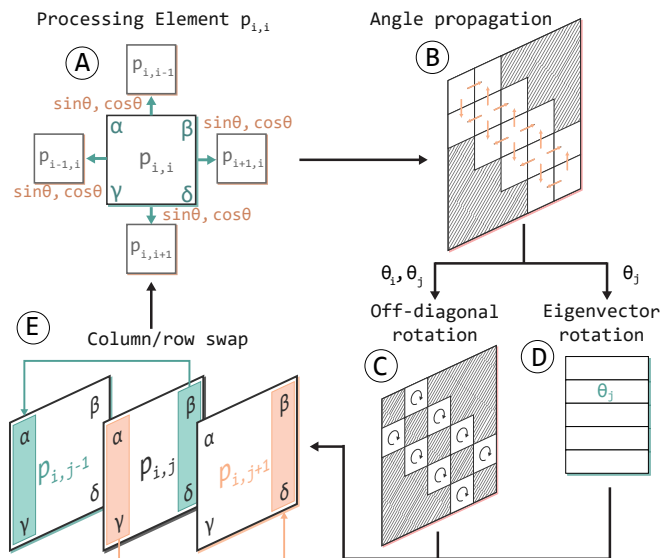


Fig. 5: Steps of the Jacobi eigenvalues computation using Systolic Arrays. Each Processing Element (PE) p_{ij} holds 4 values α , β , γ , δ , and $\theta = \frac{1}{2} \arctan \frac{2\beta}{\alpha-\delta}$.

To ensure convergence, diagonal CUs are fed non-zero elements at each iteration in the β and γ position. New non-zero elements are provided to the diagonal CUs by swapping rows and columns, since eigencomponents are invariant to linear combinations. We improve the swap procedures of Brent and Luk [33] by swapping vectors *in reverse*, obtaining the same results with fewer resources (Section IV-C2).

The SA formulation allows performing each iteration of the algorithm in constant time, enabling complexity equal to the number of iterations, $\mathcal{O}(\log(K))$, instead of having cost above $\Omega(K^2 \cdot \log(K))$ due to the matrix multiplications [33].

IV. THE PROPOSED HARDWARE DESIGN

This section presents our custom FPGA-based hardware design for the Top-K sparse eigenproblem algorithm previously introduced. The logical division between the Lanczos and Jacobi algorithms is also present in the hardware implementation. Our hardware design is composed of two macro-areas that are mapped to separate reconfigurable Super Logic Regions (SLRs) of the FPGA, to provide more efficient resource utilization and higher flexibility in terms of clock frequency, memory interfaces, and reconfigurability. Figure 6 shows a high level view of our FPGA design. We prototyped our hardware design on an Alveo U280 accelerator card with HBM2 and DDR4 memory. The Lanczos algorithm, being a memory-intensive computation, is mapped to SLR0, which provides direct access to all the HBM2 memory interfaces on the accelerator card. SLR1 and SLR2 hosts different replicas of the IP core implementing the Jacobi algorithm, optimized for different numbers of eigenvectors K .

A. Lanczos Hardware Design

The left part of Figure 6 highlights the Lanczos algorithm hardware design components. Partitions of the sparse input

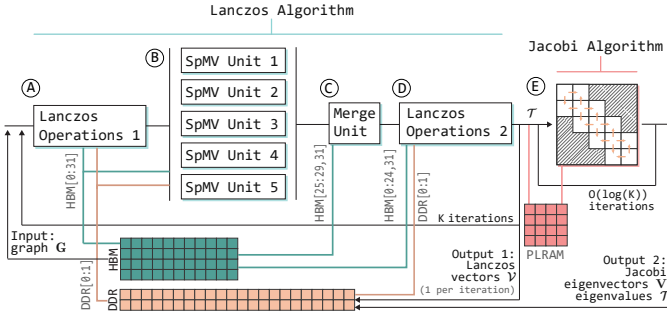


Fig. 6: High-level architecture of our Top-K Sparse Eigencomputation FPGA design. We highlight interconnections between FPGA computational units and the FPGA board memory.

matrix are read from HBM (A) and distributed to 5 parallel SpMV CUs (B) (Algorithm 1, line 7). Partial results from every partition are merged (C) into a single vector to be used by the remaining linear operations (D) (lines 5, 6, 8, 9). Operations are then repeated K times to produce the $3 \cdot K - 2$ values in the tridiagonal matrix T and the K Lanczos vectors in V , stored in DDR memory.

B. SpMV Hardware Design

The biggest bottleneck in the Lanczos algorithm is an iterative SpMV computation (Algorithm 1, line 7). While other computations in the Lanczos algorithm are relatively straightforward to optimize and parallelize, SpMV is well-known for being a complex, memory-intensive computation that presents indirect and random memory accesses [34]. Although significant research has been made into developing high-performance SpMV implementations on FPGA [35]–[38], the Lanczos algorithm introduces circumstances that prevents us from using an out-of-the-box FPGA SpMV implementation. Our SpMV design must perform multiple iterations without communication from device to host, as data-transfer and synchronizations would hinder performance. Then, the SpMV must be easily partitioned and replicated to provide flexibility over the hardware resources. Finally, we require access to multiple HBM channels to maximize the overall memory bandwidth achieved in the computation.

Our final SpMV design extends and improves the one recently proposed by Parravicini et al. [39] in the context of graph ranking algorithms, which are also variations of the power iteration method as in the case of the Lanczos algorithm. Below we introduce how we leveraged HBM memory in our SpMV design to provide better scalability and performance.

1) *SpMV Dataflow Architecture*: As SpMV is an extremely memory-intensive computation, a good SpMV implementation should make efficient use of the memory bandwidth made available by the underlying hardware. Figure 7 shows the structure of one of our SpMV CUs. We employ a streaming dataflow SpMV design that reads the input sparse matrix stored using the Coordinate (COO) format. In the COO format, non-zero entries of the matrix are stored using 3 32-bits values: the row and column index in the matrix and the value

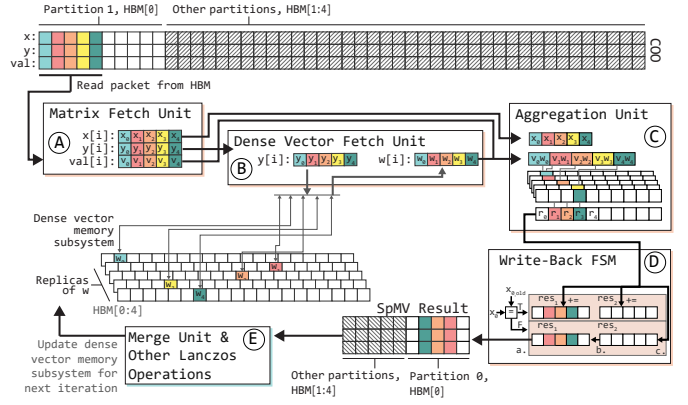


Fig. 7: Architecture of one iterative SpMV CU. Each CU processes a portion of the input matrix through a 4-stage dataflow design, and results are replicated on the *dense vector memory subsystem* after each iteration.

itself. Compared to other sparse matrix data-layouts, such as Compressed Sparse Row (CSR), the COO format does not present indirect data accesses that can severely reduce the opportunities for a pipelined design. The *Matrix Fetch Unit* in each CU is connected to a single HBM channel and reads, for each clock cycle, a packet of 512 bits containing 5 non-zero matrix entries (A). Memory transactions happen in continuous bursts of maximum AXI4 length (256 beats): each CU reads the matrix at the maximum bandwidth offered by the HBM channel (14.37 GB/s, for a total of 71.87 GB/s using 5 CU). For each of the 5 non-zero values in each COO packet, the *Dense Vector Fetch Unit* performs a random access to the SpMV dense vector (B). This step is critical to the overall SpMV performance: compared to [39], we leverage HBM instead of UltraRAM (URAM), achieving better scalability and performance. We detail our *Dense Vector Memory Subsystem* below and in Figure 8. The *Aggregation Unit* sums results within a single data-packet that refers to the same matrix column (C). A *Write-Back Finite-State Machine* stores results of each CU to HBM (D). Each write-transaction is a 512-bits data-packet containing up to 15 values, each referring to a single matrix row. Compared to [39] we reduce the number of write transactions by 3 times the average number of non-zeros per row. As such, we can store results through the same HBM channels of the dense vector with no detriment to performance.

Compared to the original SpMV design in [39], we support multiple SpMV CUs that operate on partitions on the COO input matrix, created by assigning an equal number of rows to each CU. We employ up to 5 SpMV CUs (Figure 6). While in principle it is possible to place more CUs, our current design is limited by the hardened AXI switch in the Alveo U280 that prevents the use of more than 32 AXI master channels to HBM, which we fully employ [40]. Each SpMV CU compute a portion of the output vector: partial results are aggregated by the *Merge Unit* (Figure 6 (C)) and replicated across HBM channels to use them in the following iteration.

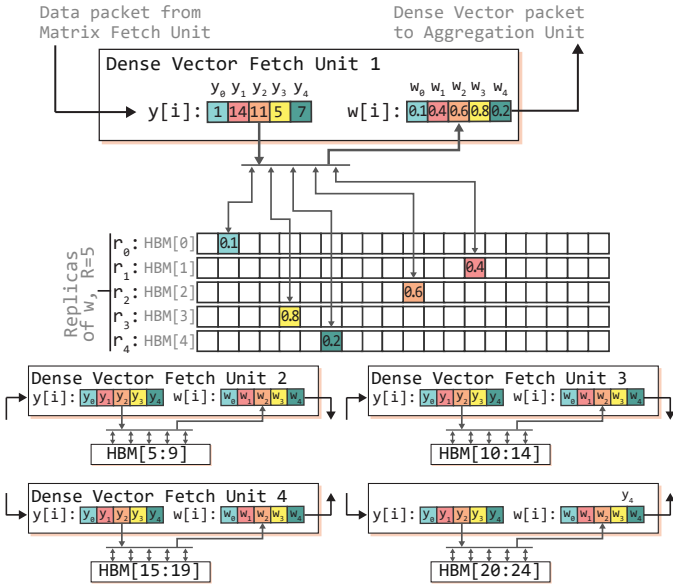


Fig. 8: Dense vector memory subsystem of our SpMV FPGA design. Index y_i accesses replica r_i , guaranteeing a pipelined design with 5 random vector accesses per clock cycle.

2) *SpMV Dense Vector Memory Subsystem*: Each SpMV CU processes 5 non-zero matrix entries per clock cycle, and for each non-zero entry it must perform a random access on a dense vector of size n (in our case, the Lanczos vector v_i at iteration i). As each AXI master channel can handle only one read transaction per cycle, we need to replicate the dense vector 5 times, similarly to [41]. The hardened AXI switch in the Alveo U280 renders highly inefficient to attach multiple AXI master channels to the same HBM bank: only 32 AXI master channels are available, and small memory transactions (32 bits) have the same performance as larger transactions, preventing sustained bandwidth sharing [42]–[44]. We solve the issue by leveraging the abundant HBM memory on the Alveo U280, and replicating the dense vector 5 times for each CU, as in Figure 8. A more flexible AXI switch would enable multiple 32-bits read transactions on the same HBM channel in a single clock cycle, reducing the demand for data replication. Compared to [39], our HBM-based memory subsystem marks a significant improvement, as we avoid URAM to store the intermediate dense vector and results. Instead of being limited by the FPGA’s 90 MB of URAM, we store the dense vector using individual HBM banks with 250 MB of capacity, allowing computations on matrices with up to 62.4 million rows. Moreover, high URAM consumption significantly limits the maximum attainable frequency, while we do not incur in this limitation (Table I).

C. Jacobi Systolic Array Design

The Jacobi eigenvalue algorithm is very computationally intensive. Although it processes a small input of size $K \times K$, unoptimized implementations still require a significant amount of time due to a large number of dense matrix multiplications.

Moreover, its convergence rate is implementation-dependent and as high as $\mathcal{O}(K^2)$. By adopting a SA-based design, we overcome both issues. By parallelizing the computation through a SA formulation and performing rotations concurrently, we decrease the number of iterations for convergence to $\mathcal{O}(\log(K))$. Rotations, equivalent to multiplications on 2×2 submatrices, are unrolled and performed in constant time.

Our design for the Jacobi algorithm is optimized to compute up to K eigenvalues. While it can compute a lower amount of eigenvalues without a reconfiguration, we place in the same FPGA bitstream multiple *Jacobi cores* optimized for specific K (4, 8, 16, etc.). We can configure both SLR1 and SLR2 with Jacobi cores to fully utilize the FPGA resources and opening the doors for independent optimization on specific values of K by reconfiguring individual SLRs. The *Lanczos Core* on SLR transfers only the $3K - 2$ values of \mathcal{T} to the Jacobi cores on SLR1 and SLR2. We prevent inefficiencies related to inter-SLR communication by moving data through PLRAM, while also avoiding the long read-write latency of DDR and HBM.

In practice, the systolic formulation cannot scale beyond very small matrices ($K \approx 32$) due to the large number of resources required for trigonometric operations in each CU. While resource utilization has prevented widespread adoption of the Jacobi algorithm for general eigenproblem resolution, it is not a limitation for our use case, as we apply the Jacobi eigenvalue algorithm on small $K \times K$ inputs by design.

On CPU, approaches such as QR factorization are more common [45], because efficient systolic array formulations of the Jacobi algorithm require full control over cache eviction policies. Moreover, even modern CPUs lack enough floating-point arithmetic units to perform the operations required for an iteration at once: even for a small K such as $K = 8$, the Jacobi algorithm computes 16 trigonometric operations and about 800 floating-point multiplications per iteration.

Instead, we leverage the abundant hardware resources of our FPGA platform to perform all these operations concurrently, making it the optimal choice for our Jacobi SA design.

1) *Diagonal And Offdiagonal CU*: Diagonal CU (Algorithm 2, line 4) annihilate elements immediately outside the diagonal via a matrix rotation. Although the rotation angle is arbitrary, the fastest convergence is achieved by setting $\theta = \frac{1}{2} \arctan \frac{2\beta}{\alpha - \delta}$, which eliminates the β and γ components (Figure 4a). We efficiently compute the components of the rotation matrix via Taylor series expansion. Even an order-3 approximation provides excellent accuracy ($\sim 10^{-6}$ at $\pm \frac{\pi}{4}$), using significantly fewer DSPs and BRAMs than the CORDIC core. Rotation on the diagonal (Figure 5 (A)) are performed by $K/2$ parallel cores, propagating rotation values (B) in constant time to the Offdiagonal CU (Algorithm 2, line 9). As each CUs holds only 4 elements, matrix multiplications are fully unrolled and performed in constant time. Eigenvectors (Algorithm 2, line 14) (D) are computed in parallel to the rotation of the Offdiagonal CU (C) as they only require rotation values.

2) *Row/Column Interchange*: Each CU has 8 connections to propagate input and output values of $\alpha, \beta, \gamma, \delta$ values to adjacent processors, in addition to communicating the rotation

TABLE I: Resource usage and clock frequency in our FPGA hardware design, divided by algorithm.

Algorithm	SLR	LUT	FF	BRAM	URAM	DSP	Clock (MHz)
Lanczos	SLR0	42%	13%	15%	0%	16%	225
Jacobi	SLR1	40%	42%	0%	0%	68%	225
Jacobi	SLR2	15%	17%	0%	0%	34%	225
Available		1097419	2180971	1812	960	9020	

TABLE II: Matrices/graphs in the evaluation, sorted by number of edges/non-zero entries (in millions). For each matrix, we report the memory footprint when stored as COO.

ID	Name	Rows (M)	Non-zeros (M)	Sparsity (%)	Size (GB)
WB-TA	wiki-Talk	2.39	5.02	8.79×10^{-4}	0.06 GB
WB-GO	web-Google	0.91	5.11	6.17×10^{-4}	0.07 GB
WB-BE	web-Berkstan	0.69	7.60	1.60×10^{-3}	0.10 GB
FL	Flickr	0.82	9.84	1.46×10^{-3}	0.13 GB
IT	italy_osm	6.69	14.02	3.13×10^{-5}	0.18 GB
PA	patents	3.77	14.97	1.05×10^{-4}	0.19 GB
VL3	venturiLevel3	4.02	16.10	9.96×10^{-5}	0.21 GB
DE	germany_osm	11.54	24.73	1.86×10^{-5}	0.32 GB
ASIA	asia_osm	11.95	25.42	1.78×10^{-5}	0.33 GB
RC	road_central	14.08	33.87	1.71×10^{-5}	0.43 GB
WK	Wikipedia	3.56	45.00	3.55×10^{-4}	0.60 GB
HT	hugetrace-00020	16.00	47.80	1.87×10^{-5}	0.61 GB
WB	wb-edu	9.84	57.15	5.90×10^{-5}	0.73 GB

value θ . As shown in Figure 5E, each processor $p_{i,j}$ with i and $j \neq (1, K/2)$ propagates its α and γ values to the β and δ slots of $p_{i,j+1}$ and its β and δ values to the α and γ slots of $p_{i,j-1}$. Processors in the first column ($p_{i,1}$) propagate β and δ to the α and γ slots of $p_{i,2}$. Processors $p_{i,K/2}$ propagate β and δ to their own α and γ slots. Operations for the column interchange are symmetrical. As α and γ of $p_{i,1}$ are never propagated, more swaps are performed towards lower indices than higher indices. These additional swaps require K temporary vectors to store rows that would be overwritten by the swaps. To avoid wasting resources for these temporary vectors, we execute operations in *reverse*, from $K/2$ to 1. As row/column swaps do not introduce additional data dependencies, we perform them in a single clock cycle using FFs.

V. EXPERIMENTAL EVALUATION

To prove that our custom FPGA design is suitable for solving large-scale Top-K sparse eigenproblems, we compare it against the popular ARPACK library, measuring how it compares in terms of execution time, power efficiency, and accuracy. The multi-threaded ARPACK library [10], a Top-K sparse eigensolver that employs IRAM, runs on two Intel Xeon Gold 6248 (80 threads in total) and 384 GB of DRAM using single-precision floating-point arithmetic. Our eigensolver is prototyped on a Xilinx Alveo U280 accelerator card equipped with 8 GB of HBM2 memory, 32GB of DDR4 memory, and a `xcu280-fsvh2892-2L-e` FPGA whose resources are reported in Table I. Results are averaged over 20 runs.

Tests are carried out using a collection of large sparse matrices representing graph topologies, each containing millions of rows and non-zero entries (Table II). All test matrices come

from the SuiteSparse collection [46]. While our evaluation is focused on sparse matrices representing graphs, our Top-K sparse eigenproblem FPGA design is applicable to other domains such as image analysis [47]–[49].

Resource utilization and clock frequency of our design are reported in Table I. The Lanczos algorithm and Jacobi algorithm have similar utilization, with around 20% LUT utilization each (50% of the available LUTs in each SLR). Although the SA architecture of the Jacobi algorithm processes small $K \times K$ inputs, it requires the computation of many trigonometric operations and multiplications (16 and > 800 for $K=8$) in each iteration. Resource utilization of the Jacobi algorithm scales quadratically with the number of eigenvalues K , while the Lanczos algorithm is not affected.

A. Execution Time

We measure the execution time speedup of the FPGA-based hardware design implementing our Top-K sparse eigenproblem solver against the CPU baseline and report results in Figure 9. We are always faster than the baseline, with a geometric mean speedup of 6.22x, up to 64x for specific graphs. The speedup is mostly unaffected by K , showing how our design can efficiently compute many eigenvalues at once. Figure 10A shows how the time required by our FPGA design to process a single matrix value is unaffected by the overall graph size, while the CPU behavior is drastically more unpredictable.

We estimate that the Lanczos dominates the overall execution time due to the SpMV computations, taking more than 99% of the execution time. However, optimizing the Jacobi algorithm with a SA design is still worth the effort, compared to running this step on CPU. Figure 10B shows the speedup of our Jacobi SA design compared to an optimized C++ CPU implementation: the execution time on CPU grows quadratically due to repeated matrix multiplications, becoming a non-negligible part of the execution time for large K .

Our hardware design synthesized at 225 Mhz on the Alveo U280 accelerator card. A clock frequency beyond 225 Mhz does not significantly improve performance as SpMV represents the main computational bottleneck in the computation, and its performance is bound by HBM bandwidth [43]. Each SpMV CU processes data at the maximum bandwidth offered by the HBM channel from which it reads the matrix (14.37 GB/s, for a total of 71.87 GB/s using 5 CU).

B. Power Efficiency

We measured via an external power meter that our FPGA design consumes about 38W during execution, plus 40W for the host server. The CPU implementation consumes around 300W during execution. Our FPGA design provides 49x higher Performance/Watt ratio (24x if accounting for the FPGA host machine): we provide higher performance without sacrificing power efficiency, making our design suitable for repeated computations typical of data center applications.

C. Accuracy Analysis of the Approximate Eigencomputation

The Lanczos algorithm is known to suffer from numerical instability [31]. To limit this phenomenon, we reorganize the

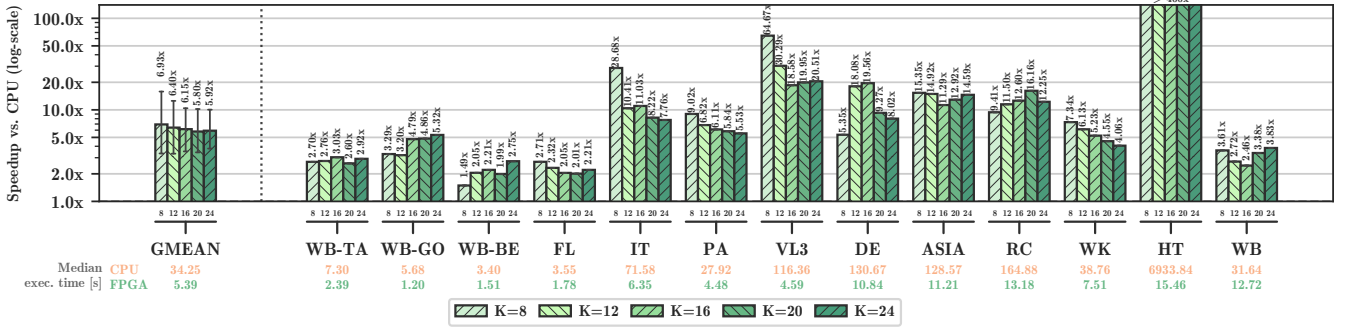


Fig. 9: Speedup (higher is better) of our Top-K sparse eigensolver vs. the ARPACK multi-core CPU library, for increasing number of eigenvalues K . Geomean excludes the outlier graph HT, where the speedup of our FPGA design exceeds 400x.

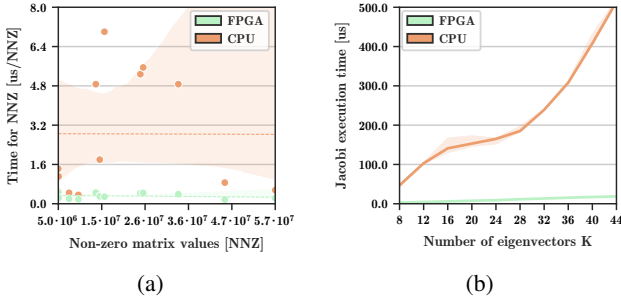


Fig. 10: (a) Relation between number of matrix non-zero values and time to process a single value. (b) Speedup vs. CPU of our Systolic Array architecture for the Jacobi algorithm.

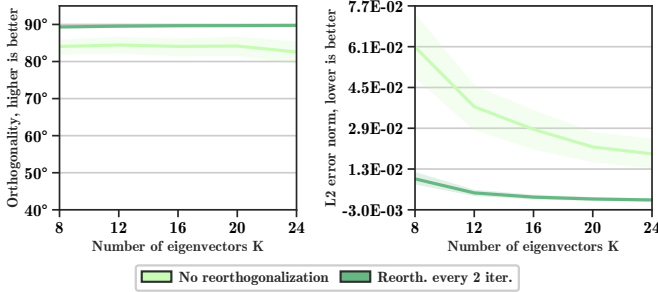


Fig. 11: Accuracy of our Top-K sparse eigensolver, in terms of orthogonality and reconstruction error, for increasing K .

algorithm’s operations as in [31] and apply reorthogonalization as in [32]. To assess the stability of our design, we measure the eigenvectors’ pairwise orthogonality and the eigenvector error norm. Eigenvectors must form an orthonormal basis and be pairwise orthogonal, i.e. their dot product is 0, or equivalently their angle is $\pi/2$. For each pair of eigenvectors, we measure the average angle that occurs. Then, if λ is an eigenvalue of M and v is its associated eigenvector, it must hold $Mv = \lambda v$. By measuring the L2 norm of $Mv - \lambda v$ for all v we evaluate how precise the eigenvector computation is. Results are reported, for increasing K , in Figure 11, aggregated on all graphs due to space constraints. Accuracy is

excellent if reorthogonalization is applied every two iterations, but even without this procedure results are satisfactory. Despite using fixed-precision arithmetic in the Lanczos algorithm, the average reconstruction error is below 10^{-3} , and the average orthogonality is > 89.9 degrees, when applying reorthogonalization every two iterations. Orthogonality is not affected by K , while the average reconstruction error improves as K increases. Spectral methods in machine learning applications use eigenvectors to capture the most important features of their input and do not usually require the same degree of precision as other engineering applications. Reorthogonalization adds an overhead up to $\mathcal{O}(nK^2/2)$ to the algorithm compared to Figure 9. On large graphs this overhead is negligible compared to SpMV, and is a viable option in applications where maximum accuracy is necessary. Still, our hardware design can provide excellent accuracy while being significantly faster than a highly optimized CPU implementation.

VI. CONCLUSION

The computation of the Top-K eigenvalues and eigenvectors on large graphs represented as sparse matrices is critical in spectral methods, a class of powerful Machine Learning algorithms that can extract useful features from graphs. We solve the Top-K sparse eigenproblem with a new algorithm that is optimized for reconfigurable hardware designs: in the first part of the computation, we exploit the enormous bandwidth of HBM through the Lanczos algorithm, while in the second part, we introduce a systolic array architecture that efficiently parallelizes the compute-intensive Jacobi eigenvalue algorithm. Compared to the popular ARPACK CPU library, we achieve a geomean speedup of 6.22x on 13 graphs with millions of vertices, raising the bar for high-performance Top-K sparse eigensolvers at a large scale.

As future work, we will extend our hardware design to support non-Hermitian matrices through the Implicitly Restarted Arnoldi Method. We will also investigate heterogeneous implementations that combine the abundant memory bandwidth of GPUs for high-performance SpMV with our systolic array FPGA design for the Jacobi eigenvalue.

REFERENCES

- [1] Q. Zhou, X. Chen, and J. Zhang, "Spectral clustering-based matrix completion method for top-n recommendation," in *Proceedings of the 2019 5th International Conference on Computing and Data Engineering*, ser. ICCDE' 19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1–6.
- [2] F. Tang, C. Wang, J. Su, and Y. Wang, "Spectral clustering-based community detection using graph distance and node attributes," *Computational Statistics*, vol. 35, no. 1, pp. 69–94, 2020.
- [3] A. Y. Ng, M. I. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Advances in neural information processing systems*, 2002, pp. 849–856.
- [4] M. Myllykoski and C. C. Kjelgaard Mikkelsen, "Task-based, gpu-accelerated and robust library for solving dense nonsymmetric eigenvalue problems," *Concurrency and Computation: Practice and Experience*, vol. n/a, no. n/a, p. e5915. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.5915>
- [5] A. V. Knyazev, "Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method," *SIAM journal on scientific computing*, vol. 23, no. 2, pp. 517–541, 2001.
- [6] C. Lanczos, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, 1950.
- [7] H. Rutishauser, "The jacobi method for real symmetric matrices," *Numerische Mathematik*, vol. 9, no. 1, pp. 1–10, Nov 1966. [Online]. Available: <https://doi.org/10.1007/BF02165223>
- [8] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [9] A. Cosnau, "Computation on gpu of eigenvalues and eigenvectors of a large number of small hermitian matrices," *Procedia Computer Science*, vol. 29, pp. 800–810, 12 2014.
- [10] R. B. Lehoucq, D. C. Sorensen, and C. Yang, *ARPACK users' guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods*. SIAM, 1998.
- [11] H. M. Aktulga, C. Yang, E. G. Ng, P. Maris, and J. P. Vary, "Topology-aware mappings for large-scale eigenvalue problems," in *European Conference on Parallel Processing*. Springer, 2012, pp. 830–842.
- [12] V. Hernandez, J. Roman, A. Tomas, and V. Vidal, "A survey of software for sparse eigenvalue problems," *Universitat Politècnica De Valencia, SLEPs technical report STR-6*, 2009.
- [13] K. J. Maschhoff and D. C. Sorensen, "P_arpack: An efficient portable large scale eigenvalue package for distributed memory parallel architectures," in *International workshop on applied parallel computing*. Springer, 1996, pp. 478–486.
- [14] D. Lee, T. Hoshi, T. Sogabe, Y. Miyatake, and S.-L. Zhang, "Solution of the k-th eigenvalue problem in large-scale electronic structure calculations," *Journal of Computational Physics*, vol. 371, 2018.
- [15] Nvidia, "cusolver," 2020. [Online]. Available: docs.nvidia.com/cuda/cusolver/index.html
- [16] —, "nvgraph," 2019. [Online]. Available: docs.nvidia.com/cuda/nvgraph/index.html
- [17] MAGMA, "magma_slobpcg," 2020. [Online]. Available: icl.cs.utk.edu/projectsfiles/magma/doxygen/group__magma_sparse__ssyev.html
- [18] N. M. Evstigneev, "Implementation of implicitly restarted arnoldi method on multigpu architecture with application to fluid dynamics problems," in *International Conference on Parallel Computational Technologies*. Springer, 2017, pp. 301–316.
- [19] J. Dubois, C. Calvin, and S. Petiton, "Accelerating the explicitly restarted arnoldi method with gpus using an autotuned matrix vector product," *SIAM Journal on Scientific Computing*, 2011.
- [20] K. D. Gupta, M. Wajid, R. Muzammil, and S. J. Arif, "Hardware architecture for eigenvalues computation using the modified jacobi algorithm on fpga," in *2019 5th International Conference on Signal Processing, Computing and Control (ISPCC)*, 2019, pp. 243–246.
- [21] A. Burger, P. Urban, J. Boubin, and G. Schiele, "An architecture for solving the eigenvalue problem on embedded fpgas," in *Architecture of Computing Systems – ARCS 2020*, A. Brinkmann, W. Karl, S. Lankes, S. Tomforde, T. Pionteck, and C. Trinitis, Eds. Cham: Springer International Publishing, 2020, pp. 32–43.
- [22] H.-T. Kung, "Why systolic architectures?" *IEEE computer*, 1982.
- [23] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers *et al.*, "In-datacenter performance analysis of a tensor processing unit," in *Proceedings of the 44th Annual International Symposium on Computer Architecture*, 2017, pp. 1–12.
- [24] M. Langhammer, S. Gribok, and G. Baeckler, "High density 8-bit multiplier systolic arrays for fpga," in *2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 2020, pp. 84–92.
- [25] B. Asgari, R. Hadidi, and H. Kim, "Proposing a fast and scalable systolic array for matrix multiplication," in *2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 2020, pp. 204–204.
- [26] J. Vázquez-Castillo, A. Castillo-Atoche, R. Carrasco-Alvarez, O. Longoria-Gandara, and J. Ortegón-Aguilar, "Fpga-based hardware matrix inversion architecture using hybrid piecewise polynomial approximation systolic cells," *Electronics*, vol. 9, no. 1, p. 182, 2020.
- [27] S. Aslan, S. Niu, and J. Saniie, "Fpga implementation of fast qr decomposition based on givens rotation," 08 2012, pp. 470–473.
- [28] P. Desai, S. Aslan, and J. Saniie, "Fpga implementation of gram-schmidt qr decomposition using high level synthesis," in *2017 IEEE International Conference on Electro Information Technology (EIT)*, 2017.
- [29] G. H. Golub and R. Underwood, "The block lanczos method for computing eigenvalues," in *Mathematical software*. Elsevier, 1977.
- [30] D. Calvetti, L. Reichel, and D. C. Sorensen, "An implicitly restarted lanczos method for large symmetric eigenvalue problems," *Electronic Transactions on Numerical Analysis*, vol. 2, no. 1, p. 21, 1994.
- [31] C. C. Paige, "Computational variants of the lanczos method for the eigenproblem," *IMA Journal of Applied Mathematics*, 1972.
- [32] B. N. Parlett, *The Symmetric Eigenvalue Problem*. Philadelphia: Society for Industrial and Applied Mathematics, 1998.
- [33] R. P. Brent, F. T. Luk *et al.*, "The solution of singular-value and eigenvalue problems on systolic arrays," in *Mathematical Programming and Numerical Analysis Workshop*. Centre for Mathematics and its Applications, Mathematical Sciences Institute . . . , 1984, pp. 38–64.
- [34] T. Nguyen, S. Williams, M. Siracusa, C. MacLean, and N. W. D. Doerfler, "The performance and energy efficiency potential of fpgas in scientific computing," in *IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, 2020.
- [35] F. Sadi, J. Sweeney, T. M. Low, J. C. Hoe, L. Pileggi, and F. Franchetti, "Efficient spmv operation for large and highly sparse matrices using scalable multi-way merge parallelization," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*. New York, NY, USA: Association for Computing Machinery, 2019.
- [36] P. Grigoras, P. Burovskiy, E. Hung, and W. Luk, "Accelerating spmv on fpgas by compressing nonzero values," in *2015 IEEE 23rd Annual International Symposium on Field-Programmable Custom Computing Machines*. IEEE, 2015, pp. 64–67.
- [37] Y. Umuroglu and M. Jahre, "A vector caching scheme for streaming fpga spmv accelerators," in *International Symposium on Applied Reconfigurable Computing*. Springer, 2015, pp. 15–26.
- [38] A. K. Jain, H. Omidian, H. Fraisse, M. Benipal, L. Liu, and D. Gaitonde, "A domain-specific architecture for accelerating sparse matrix vector multiplication on fpgas," in *2020 30th International Conference on Field-Programmable Logic and Applications (FPL)*. IEEE, 2020.
- [39] A. Parravicini, F. Sgherzi, and M. D. Santambrogio, "A reduced-precision streaming spmv architecture for personalized pagerank on fpga," *arXiv preprint arXiv:2009.10443*, 2020.
- [40] Xilinx, "Axi high bandwidth memory controller v1.0 logicore ip product guide," 2020. [Online]. Available: www.xilinx.com/support/documentation/ip_documentation/hbm/v1_0/pg276-axi-hbm.pdf
- [41] S. Kestur, J. D. Davis, and E. S. Chung, "Towards a universal fpga matrix-vector multiplication architecture," in *2012 IEEE 20th International Symposium on Field-Programmable Custom Computing Machines*. IEEE, 2012, pp. 9–16.
- [42] Y.-k. Choi, Y. Chi, J. Wang, L. Guo, and J. Cong, "When hls meets fpga hbm: Benchmarking and bandwidth optimization," *arXiv preprint arXiv:2010.06075*, 2020.
- [43] A. Lu, Z. Fang, W. Liu, and L. Shannon, "Demystifying the memory system of modern datacenter fpgas for software programmers

through microbenchmarking,” in *Proc. ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays*, 2021.

- [44] Z. Wang, H. Huang, J. Zhang, and G. Alonso, “Shuhai: Benchmarking high bandwidth memory on fpgas,” in *2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 2020, pp. 111–119.
- [45] A. Buttari, J. Langou, J. Kurzak, and J. Dongarra, “Parallel tiled qr factorization for multicore architectures,” *Concurrency and Computation: Practice and Experience*, vol. 20, no. 13, pp. 1573–1590, 2008.
- [46] T. A. Davis and Y. Hu, “The university of florida sparse matrix collection,” *ACM Transactions on Mathematical Software (TOMS)*, vol. 38, no. 1, pp. 1–25, 2011. [Online]. Available: sparse.tamu.edu
- [47] P. Perona and W. Freeman, “A factorization approach to grouping,” in *European Conference on Computer Vision*. Springer, 1998.
- [48] J. Shi and J. Malik, “Normalized cuts and image segmentation,” *IEEE Transactions on pattern analysis and machine intelligence*, 2000.
- [49] F. Tung, A. Wong, and D. A. Clausi, “Enabling scalable spectral clustering for image segmentation,” *Pattern Recognition*, vol. 43, no. 12, pp. 4069–4076, 2010.