

# FractalCloud: A Fractal-Inspired Architecture for Efficient Large-Scale Point Cloud Processing

Yuzhe Fu<sup>†</sup>, Changchun Zhou<sup>†\*</sup>, Hancheng Ye<sup>†</sup>, Bowen Duan<sup>†</sup>, Qiyu Huang<sup>‡</sup>, Chiyue Wei<sup>†</sup>, Cong Guo<sup>†\*</sup>,  
Hai “Helen” Li<sup>†</sup>, and Yiran Chen<sup>†</sup>

<sup>†</sup>Duke University, <sup>‡</sup>Yale University

{yuzhe.fu, changchun.zhou, cong.guo, hai.li, yiran.chen}@duke.edu

**Abstract**—Three-dimensional (3D) point clouds are increasingly used in applications such as autonomous driving, robotics, and virtual reality (VR). Point-based neural networks (PNNs) have demonstrated strong performance in point cloud analysis, originally targeting small-scale inputs. However, as PNNs evolve to process large-scale point clouds with hundreds of thousands of points, all-to-all computation and global memory access in point cloud processing introduce substantial overhead, causing  $O(n^2)$  computational complexity and memory traffic where  $n$  is the number of points. Existing accelerators, primarily optimized for small-scale workloads, overlook this challenge and scale poorly due to inefficient partitioning and non-parallel architectures. To address these issues, we propose *FractalCloud*, a fractal-inspired hardware architecture for efficient large-scale 3D point cloud processing. *FractalCloud* introduces two key optimizations: (1) a co-designed Fractal method for shape-aware and hardware-friendly partitioning, and (2) block-parallel point operations that decompose and parallelize all point operations. A dedicated hardware design with on-chip fractal and flexible parallelism further enables fully parallel processing within limited memory resources. Implemented in 28 nm technology as a chip layout with a core area of  $1.5 \text{ mm}^2$ , *FractalCloud* achieves  $21.7\times$  speedup and  $27\times$  energy reduction over state-of-the-art accelerators while maintaining network accuracy, demonstrating its scalability and efficiency for PNN inference. The code for *FractalCloud* is available at <https://github.com/Yuzhe-Fu/FractalCloud>.

## I. INTRODUCTION

With the rapid advancement of 3D sensing technologies [1]–[4], such as LiDAR [5] and stereo cameras [6]–[8], point cloud has become prevalent in a wide range of edge applications, including autonomous driving [9], [10], robotics [11], augmented reality [12], and virtual reality [13], [14]. The points in a 3D point cloud are inherently unordered and irregularly distributed, making it challenging for conventional deep neural networks (DNNs) to directly extract meaningful local features [15]–[18]. To overcome this limitation, point-based neural networks (PNNs) [19]–[26] incorporate specialized point operations, such as sampling, neighbor searching, and gathering, to effectively extract learnable local features from unordered point cloud. These tailored operations significantly improve the performance of PNNs in 3D perception tasks. However, these operations often involve irregular memory access patterns and intensive computation, which pose significant challenges to the efficiency of PNNs on edge hardware.

Numerous architecture researchers have proposed various hardware- and software-level optimizations to accelerate PNN

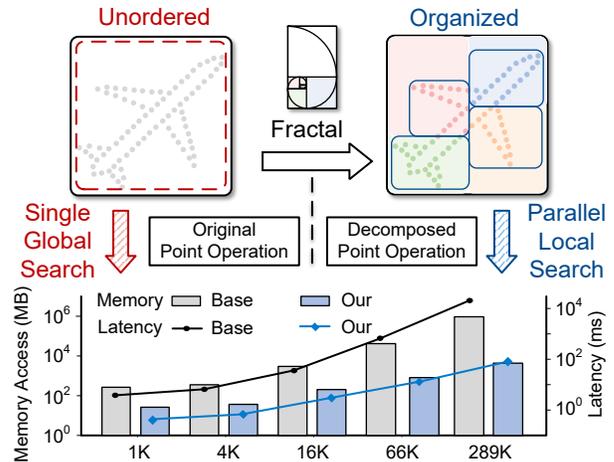


Fig. 1: The memory access (MB) and inferencing latency (ms) of point cloud neural networks in (a) original baseline structure and (b) the proposed *FractalCloud*.

execution, aiming to improve both throughput and energy efficiency. Early works, such as Mesorasi [27] and PointAcc [28], primarily focused on optimizing intensive tensor computations (i.e., DNN layers), which were the main performance bottleneck at the time. As these tensor-level bottlenecks were mitigated, the bottleneck shifted to the irregular memory access from point operations. Consequently, more efforts have focused on optimizing memory access efficiency. One common approach is to partition data into blocks and organize them to alleviate memory access conflicts [29]. Earlier algorithmic work [30], [31] analyzed various partitioning methods (e.g., KD-trees) on conventional datasets without considering hardware optimization or the geometric structures of point-cloud data. In contrast, recent architectural designs such as PNNPU [32] employ a space-uniform partitioning strategy, dividing the 3D space by point coordinates to generate structured memory access patterns. Crescent [29] adopts a density-uniform partitioning scheme based on the KD-tree to achieve more balanced workloads and improved memory locality.

However, as PNN evolves, the scale of point cloud data has expanded from 1K to over 200K in benchmarks [33], [34], with modern LiDAR sensors producing 30K–300K points per frame in real life [5], [35]. This increase introduces new challenges, with point operations emerging as the primary

\*Corresponding Author: Changchun Zhou and Cong Guo

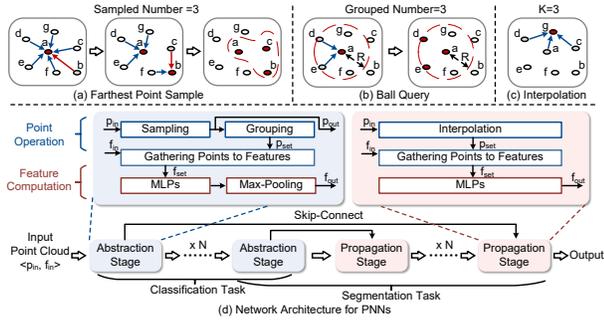


Fig. 2: The workflows for (a) farthest point sampling, (b) ball query, and (c) interpolation. (d) The backbone of PNNs.

bottleneck, as shown in Fig. 4. Conventional point operations require iterative global memory access and all-to-all computation, referred to as global search. The iterative nature and data dependency severely limit decomposition and parallelization, both essential for efficient large-scale processing [36], [37]. Partitioning mitigates global search by localizing computations and memory accesses. However, prior works mainly target small-scale inputs ( $< 4K$  points) and overlook this parallelization issue. As a result, their hardware designs are misaligned with the demands of modern workloads, causing significant performance degradation on large-scale point clouds.

After detailed analysis, we identify two fundamental challenges in state-of-the-art (SOTA) accelerators:

**Inefficient Partitioning.** Existing partitioning strategies exhibit a fundamental trade-off between hardware efficiency and network accuracy. Simple space-uniform partitions [32] are hardware-friendly but ignore highly uneven point distributions, leading to imbalanced blocks and severe accuracy degradation [38], making them unsuitable for large-scale PNNs. Density-aware partitions, such as the KD-tree-based method [29], [39], achieve balanced partitions that preserve accuracy but rely on recursive and exclusive sorting operations. These sorts are inherently sequential and non-decomposable, creating severe hardware inefficiency and becoming a new bottleneck (e.g., 53% latency in Crescent), making them inefficient with increased point cloud.

**Non-Parallel Architecture.** Existing architectures lack optimizations on block-level parallelism. Accelerators without partitioning [27], [28] would make non-decomposable point operations with long latency, as shown in Fig. 1. Even accelerators with partitioning remain constrained by the sequential nature of partitioning methods [29], resulting in block-serial execution. While acceptable for small-scale inputs (e.g., Crescent is only 20% slower than ours), this design assumption fails for large-scale point clouds. As input size  $n$  increases, partitioning latency grows rapidly, and the lack of block-level parallelism causes latency to scale almost  $O(n^2)$ , severely limiting overall performance.

**Main Idea.** To enable efficient large-scale point cloud processing and address the two fundamental challenges, we propose *FractalCloud*, a fractal-inspired hardware architecture.

It organizes point cloud into blocks with a fractal strategy and localizes global point operations into fully parallel block-wise execution, as shown in Fig. 1. Our contributions are summarized as follows:

- **Fractal:** A hardware-driven, shape-aware partitioning method with a DFT-based memory layout and threshold-controlled block division, enabling efficient large-scale point cloud processing.
- **Block-Parallel Point Operations:** Decomposing all point operations from global to local search, enabling fully parallel block-wise processing, and significantly reducing memory accesses and latency.
- **Dedicated Hardware Architecture:** A domain-specific accelerator optimized for fractal-based processing, featuring on-chip fractal, flexible parallelism, and data reuse to accelerate parallelized point operations.
- In 28 nm technology, our *FractalCloud* achieves  $1.5mm^2$  and  $0.58W$ , delivering an average  $21.7\times$  speedup and  $27\times$  energy reduction over state-of-the-art accelerators.

## II. BACKGROUND

In this section, we firstly show the backbone of typical PNNs, then explain the point operations in PNNs.

### A. The Backbone of PNNs

A point cloud consists of two types of information: spatial coordinates  $p$  and feature information  $f$ . As illustrated in Fig. 2(d), a typical PNN architecture comprises two main stages: Abstraction and Propagation. Each stage includes two sequential computational pathways: Point Operations, performed on spatial coordinates to extract learnable local geometric features, and Feature Computation, which processes and aggregates these features to learn high-level semantic representations [20]–[26].

In classification tasks, PNNs consist of the Abstraction stages. Firstly, the spatial coordinates  $p_{in}$  of the input point cloud undergo point operations, where a sampling operation selects representative points. A grouping operation then identifies spatially neighboring points for each sampled point, forming local point sets  $p_{set}$ , which are subsequently transformed into the feature space through a gathering operation, producing Feature Sets  $f_{set}$ .  $f_{set}$  are then processed in the Feature Computation, which applies Multi-Layer Perceptrons (MLPs) to extract hierarchical features.

For more complex segmentation tasks, PNNs introduce additional Propagation stages that are skip-connected to previous Abstraction stages to recover features. The interpolation operations are used to restore the features of points discarded in the previous sampling steps. It leverages the features of sampled points and reconstructs the pre-sampling point cloud information. This process effectively recovers detailed spatial structures and enhances segmentation accuracy.

### B. Point Operations in PNNs

**Sampling.** Farthest Point Sampling (FPS) is the dominant sampling method in both SOTA PNNs [19]–[21] and hardware

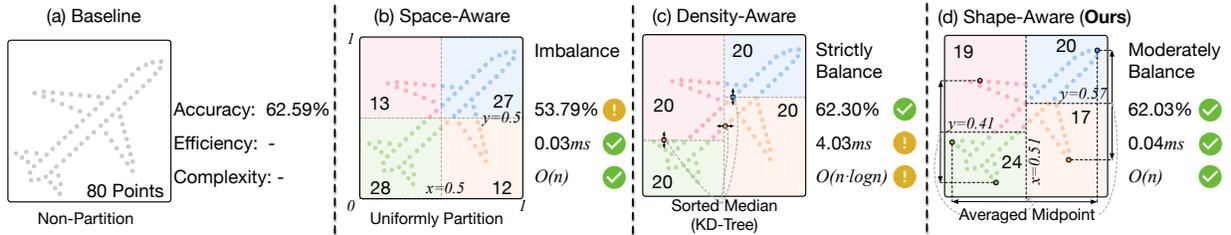


Fig. 3: Comparison of different partitioning strategies on latency and network accuracy for PointNeXt segmenting S3DIS dataset across (a) the original point cloud (PointAcc [28]), (b) uniform partitioning (PNNPU [32]), (c) KD-tree-based partitioning (Crescent [29]), and (d) the proposed Fractal method.

accelerators [28], [40], [41] for its strong representativeness and accuracy for point cloud processing, as shown in Fig. 2 (a). The process starts with a randomly selected initial point as the first sampled point in the sampled set. In each iteration, the algorithm computes the Euclidean distance between each unsampled point and the sampled set, selecting the farthest unsampled point for inclusion of the sampled set.

**Neighbor Searching.** Neighbor searching selects local points around designated center points for grouping and interpolation. PNNs commonly adopt the Ball Query (BQ) method for grouping, which selects up to  $num$  points within a radius  $R$  based on Euclidean distance as shown in Fig. 2(b). For interpolation, the K-Nearest Neighbors (KNN) method is used, retrieving the  $K$  closest points without radius constraints by searching the entire candidate set, as shown in Fig. 2(c).

**Gathering.** The gathering operation follows the neighbor searching operations, such as grouping and interpolation. It leverages the neighbor indices obtained from these operations to retrieve corresponding feature values from the point cloud’s feature space. These gathered features are then passed to the subsequent MLP module for further feature computation.

**Complexity.** As shown in Fig. 2(a–c), original FPS and neighbor searching are iterative with  $O(n)$  iterations, each performing an all-to-all computation over  $n$  points, resulting in a  $O(n^2)$  computation complexity and memory traffic. Gathering also requires frequent global data accesses to retrieve scattered features due to the random layout in original point cloud.

### III. MOTIVATION

#### A. Bottleneck Shift: From Tensor to Point Operation

In recent years, the latency of point-based neural networks (PNNs) has escalated significantly with increasing point cloud sizes and network complexity. As shown in Fig. 4, the primary performance bottleneck has shifted from tensor computations to point operations. Unlike tensor computations, which can be efficiently parallelized into small, localized workloads, point operations often require global search over the entire point cloud, making them inherently difficult to parallelize. For example, sampling techniques such as FPS involve sequential and iterative global traversals to select representative points, while neighbor searching scans the full point set to identify local structures from unordered data. These global operations incur  $O(n^2)$  computational complexity and memory traffic,

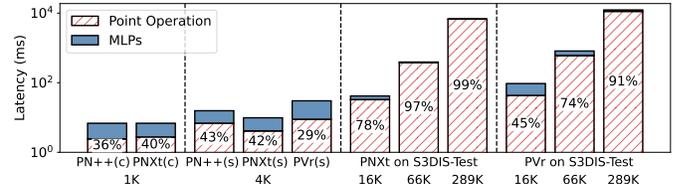


Fig. 4: The latency for PN++, PNxt, and Pvr [20]–[22] inference on GPU for classification and segmentation workloads, with notations defined in Table I, under varying input points.

causing point operations to dominate overall latency—rising from 30% at 1K points to over 90% at 289K points.

This bottleneck shift is also evident in SOTA accelerators [27]–[29], which primarily target small-scale point clouds (less than 4K points), where the cost of point operations is less pronounced. As a result, these designs focus on optimizing MLP computation and memory streaming, while leaving the global-search inefficiencies of point operations unaddressed. When applied to large-scale point clouds, such as the S3DIS dataset [34], existing accelerators suffer from limited on-chip storage, causing 51% average cache misses and frequent DRAM accesses for global search, which further exacerbate latency and expose inefficiencies in point operations.

**Takeaway.** Point operations have emerged as the dominant performance bottleneck (accounting for > 90% of runtime) in large-scale PNNs, underscoring the urgent need for architectural support and algorithmic innovations targeting their global-search nature.

#### B. Algorithm: Accuracy and Efficiency Trade-off

A variety of PNN acceleration studies have proposed partitioning strategies capable of mitigating the global-search bottleneck, as illustrated in Fig. 3. These approaches include the baseline design [28], space-aware partitioning [32], density-aware partitioning [29], and our proposed method—**Fractal**, a shape-aware partitioning scheme.

**Baseline.** As discussed in the previous subsection, early-stage PNN architectures retain full accuracy [27], [28], but suffer from poor efficiency due to irregular and memory-intensive point operations. To address this limitation, researchers have introduced partitioning schemes aimed at enhancing data locality and streaming memory access.

**Space-Uniform Partition.** As shown in Fig. 3(b), space-aware partitioning uniformly divides the 3D space by spatial coordinates [32], [38]. It introduces negligible overhead with a single global traversal, but ignores the non-uniform distribution of real-world point clouds. This leads to highly imbalanced point densities across blocks, weakening local context and causing notable accuracy loss in subsequent point operations ( $> 9\%$  in PointNeXt [38]). Thus, space-uniform partitioning is inadequate for high-accuracy applications.

**Density-Uniform Partition.** To address the workload imbalance and preserve model accuracy, density-aware partitioning (DAP) methods, such as KD-tree-based partition used in Crescent [29], generate balanced partitions by recursively splitting the point cloud based on density. As depicted in Fig. 3(c), KD-tree construction requires recursive median sorting, resulting in high computational complexity [42] and accounting for 53% of Crescent’s latency [29]. Although KD-trees work reasonably well for small-scale point clouds optimization, their performance deteriorates significantly as the input scale increases by 2–3 orders of magnitude. For instance, with 1K points, KD-tree based method is only 20% slower than our proposed *FractalCloud*. However, for large-scale inputs (e.g., 289K points), it becomes nearly  $100\times$  slower than *FractalCloud* due to poor algorithm-hardware alignment, as further discussed in the next subsection.

**Shape-Uniform Partition.** To overcome the limitations of prior methods, we propose the **Shape-Aware Partition** method, inspired by fractal geometry, and refer to it as **Fractal**. As shown in Fig. 3(d), *FractalCloud* performs linear-time traversal to identify marginal points along each dimension and computes axis-wise midpoints for recursive partitioning. This approach leverages object shape information to generate partitions with minimal computational overhead ( $< 0.8\%$  latency in *FractalCloud*) compared to KD-tree methods. Moreover, *FractalCloud* adaptively partitions dense regions into finer-grained blocks to avoid extreme imbalance, achieving moderately balanced partitions with negligible accuracy loss ( $< 0.6\%$ ) and significantly improved processing efficiency.

**Takeaway.** The core insight behind Fractal is that point distributions often align with the object’s geometric shape due to consistent sampling frequency. By being shape-aware, *FractalCloud* achieves an effective balance between efficiency and accuracy, making it well-suited for scalable, high-performance PNN processing.

### C. Architecture: Being Friendly to Hardware

The Fractal method significantly outperforms density-aware partitionings (DAP) in hardware efficiency, primarily due to its structural compatibility with parallel hardware architectures.

**Exclusive Sorter.** Density-aware partitioning (DAP) methods, such as KD-tree, rely on recursive median sorting as shown in Fig. 5 left. Each sort in the KD-tree acts as an exclusive, indivisible functional operation, making it inherently sequential and difficult to parallelize or pipeline efficiently. For example, when processing 1024 points, a 1024-point sorter yields a global order that cannot be directly decomposed

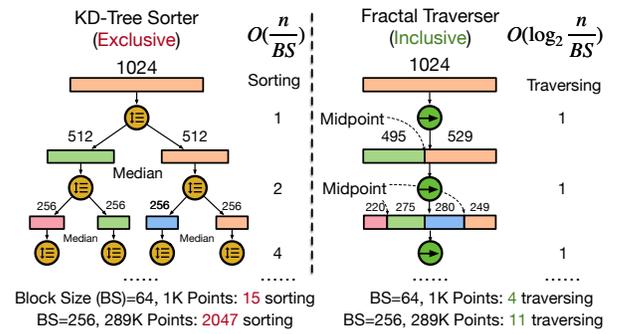


Fig. 5: Workflows for KD-tree and proposed Fractal. BS is the block size, indicating the max points per block.

or reused to execute two subsequent local 512-point sorts in parallel. And all sorts in KD-tree are mandatory as KD-tree’s changing subsets and axes invalidate prior results, requiring re-sorting at each iteration. Moreover, deeper iterations progressively increase sorts with smaller workloads, which mismatches hardware optimized for large sorting and leads to low resource utilization. Consequently, DAP incurs substantial computational overhead, with  $O(n/BS)$  complexity, where  $BS$  is the final block size, and becomes a primary bottleneck (e.g., 53% latency in Crescent [29]).

However, for small-scale inputs, this bottleneck has limited impact. Even with block-serial processing, existing accelerators achieve acceptable performance (only 20% slower than ours). As a result, prior work largely overlooked block-level parallelism. In large-scale scenarios, this assumption breaks down: KD-tree requires 2047 serial sorts for 289K points, and the lack of block-level parallelism further amplifies overall latency. As a result, DAP severely limits the parallelism of point cloud processing, posing a critical bottleneck for high-throughput accelerators where scalable parallel execution is essential.

**Inclusive Traverser.** In contrast, the proposed Fractal adopts a lightweight, linear-time traversal to identify marginal points and determine partitioning midpoints. As shown in Fig. 5 right, only four traversals are required to partition 1024 points, and even for 289K points the process completes in just 11 traversals, a sharp contrast to the thousands of recursive sorts in KD-tree. The traversal unit is inherently simple and highly parallelizable. Analogous to arithmetic components like adders, it can be decomposed into concurrent sub-tasks operating across point dimensions or blocks. This design fundamentally aligns with the requirements of parallel-processing architectures, reducing complexity to  $O(\log_2 n/BS)$  and enabling scalable, low-latency, and hardware-friendly partitioning for large-scale point clouds. Moreover, by producing balanced and independently processable blocks, Fractal provides the essential foundation for block-level parallelism in subsequent point operations.

**Takeaway.** The hardware efficiency of Fractal stems from its elimination of sorting overhead and adoption of fully decomposable traversal logic. With an inclusive, hardware-

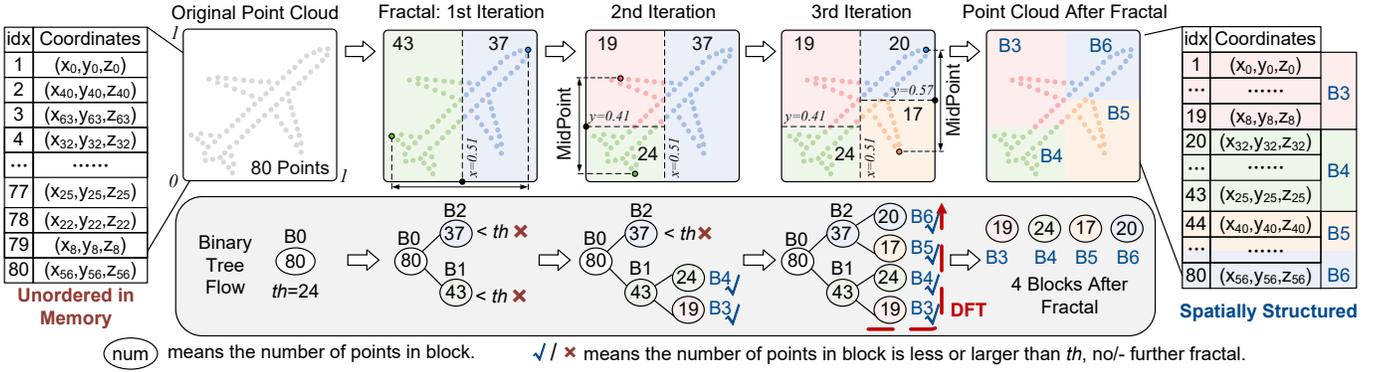


Fig. 6: The workflow of the proposed Fractal method, where  $th$  is the maximum points per block.

friendly traverser, *FractalCloud* unlocks greater parallelism and improved scalability, making it well-suited for modern PNN accelerators targeting real-world, large-scale inputs.

Finally, *FractalCloud* embodies a principled algorithm-hardware co-design philosophy. At the algorithmic level, it exploits object shape to achieve balanced partitioning with minimal complexity. At the architectural level, its traversal mechanism maps naturally to highly parallel hardware structures. Moreover, it establishes a foundation for decomposing and parallelizing all point operations, an aspect overlooked by prior accelerators, as further discussed in the next section. This synergy between algorithmic insight and architectural design allows *FractalCloud* to deliver high efficiency and scalability for large-scale PNN workloads.

#### IV. SYSTEM DESIGN: FRACTAL AND PARALLELIZATION

This section explains the system-level optimizations of *FractalCloud*, which contains two techniques: the proposed Fractal method and the Block-Parallel Point Operations.

##### A. Fractal: A Shape-Aware Partition

Fractal is a hardware-driven, shape-aware partitioning method inspired by fractal geometry, featuring a DFT-based memory layout and threshold-controlled block division, which jointly balance hardware efficiency and network accuracy. The core idea of Fractal is that point distributions in point clouds typically align with object shapes. Based on this, Fractal leverages shape information to adaptively partition point clouds into balanced blocks: dense regions are recursively partitioned into finer blocks, while sparse regions require fewer subdivisions. As shown in Fig. 6, Fractal adopts a binary tree structure for both partitioning and storage. Given a point cloud and a predefined threshold  $th$  (maximum points per block), Fractal iteratively partitions the input into two sub-blocks per step, expanding the binary tree through a recursive process, as illustrated in Alg. 1 (Rows 4–7). Each iteration involves a single traversal to compute the midpoint from the current dimension’s extrema (Row 5), with dimensions cycled across iterations (Row 3) to ensure balanced partitioning and low hardware complexity. Blocks exceeding  $th$  are further partitioned, with smaller  $th$  producing deeper, finer-grained

#### Algorithm 1 Fractal Partitioning

**Require:** Point cloud  $P$ , threshold  $th$ , dimension index  $d$

```

1: function FRACTAL( $P, d$ )
2:   if  $|P| \leq th$  then return  $P$ 
3:   else
4:      $dim \leftarrow (d \bmod 3)$  ▷ Cycle over x, y, z
5:      $mid \leftarrow (\max(P[dim]) + \min(P[dim]))/2$ 
6:      $(P_l, P_r) \leftarrow \text{split}(P, dim, mid)$ 
7:     return Node(FRACTAL( $P_l, d+1$ ), FRACTAL( $P_r, d+1$ ))
8:  $Blocks \leftarrow \text{DFT\_traverse}(\text{FRACTAL}(P, d = 0))$ 

```

partitions. Fractal organizes blocks in memory using depth-first traversal (DFT) (Row 8), while points within each block remain randomly arranged without degrading performance. Before Fractal, the globally random layout of point clouds can cause multiple compute units to access different addresses within the same memory bank, leading to severe bank conflicts [29]. After Fractal, the point cloud is organized in blocks that are distributed across different banks.

Each block is exclusively assigned to a single compute unit with a dedicated memory interface, rendering the intra-block randomness irrelevant to performance. Moreover, the DFT-based organization in Fractal ensures that adjacent memory blocks correspond to spatially adjacent regions, maintaining spatial continuity. This mitigates the control complexity of identifying neighboring blocks in conventional strategies [29], [32], and enables efficient memory access, as accessing neighboring regions simply involves reading consecutive memory blocks, facilitating further hardware optimizations.

Fig. 6 illustrates an example of Fractal with an 80-point input point cloud and a threshold  $th = 24$ . In the first iteration, Fractal computes the midpoint of the x-dimension of input point cloud  $B_0$ , and partitions the input into two blocks:  $B_1$  and  $B_2$ , with (43, 37) points, respectively. Since the number of points in both blocks exceeds  $th$ , they both undergo further fractal. In the second iteration,  $B_1$  is split along the y-dimension into  $B_3$  (19 points) and  $B_4$  (24 points), both meeting the threshold and requiring no further partitioning. The process continues until all blocks contain fewer than  $th$  points, forming a binary tree structure with four leaf nodes, representing the four point blocks after Fractal pro-

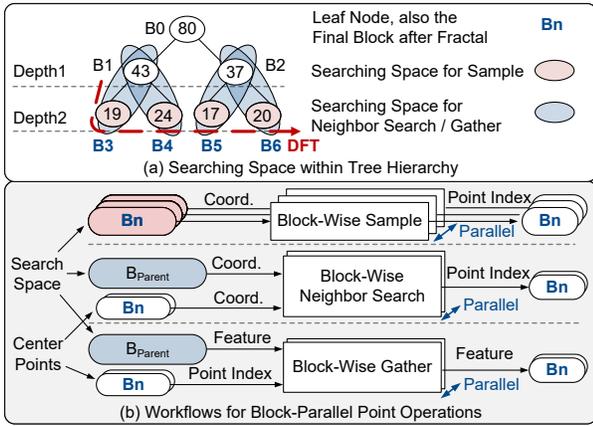


Fig. 7: The (a) search space in tree hierarchy and (b) workflows of block-parallel point operations, including block-wise sampling, neighbor searching, and gathering.

cess. By leveraging shape information to adaptively partition with traversing, Fractal produces balanced blocks with low computational complexity. Its binary tree structure, organized in depth-first memory order, enables streamed memory access. Therefore, Fractal is hardware-friendly and scalable, making it efficient for large-scale point cloud processing.

### B. Block-Parallel Point Operations (BPPO)

After applying Fractal, the point cloud is adaptively partitioned into multiple balanced blocks and stored accordingly. Leveraging this strategy, we propose block-parallel point operations, which decompose and parallelize all sampling, neighbor searching, and gathering in PNNs by reducing their search space from entire point cloud to local blocks, as illustrated in Fig. 7 with the 80-point point cloud in Fig. 6 as an example.

**Block-Wise Sampling.** During sampling, conventional Farthest Point Sampling (FPS) iteratively searches the entire point cloud, incurring high computational overhead. We propose a block-wise sampling method that leverages the tree-structured point cloud storage from Fractal to decompose FPS. In block-wise sampling, instead of performing global searches, FPS is executed independently within each block, limiting the search space and computation to the current node in the tree structure, as shown in Fig. 7. Since leaf nodes are mutually independent, these block-wise operations can run in parallel, enabling inter-block parallelism and simultaneous sampling across multiple regions. The final sampled set is obtained by aggregating the results from all blocks.

A key challenge for applying FPS within blocks while maintaining accuracy is determining the appropriate number of sampled points per block, given the non-uniform point distribution. Prior works [32], [38], [43] rely on multiple hyperparameters and additional prediction steps to estimate sampling counts. In contrast, our block-wise sampling uses a fixed sampling rate across all blocks, avoiding such complexity without sacrificing accuracy. This is enabled by Fractal, which already aligns with point distribution and adaptively partitions

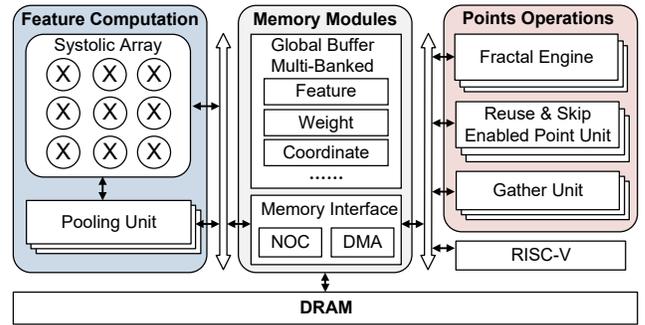


Fig. 8: The overall architecture for *FractalCloud*.

point clouds into moderately balanced blocks, eliminating the need to track point distribution. Although block-wise sampling introduces slight deviations from global search, the impact on accuracy is negligible ( $< 0.2\%$ ) because Fractal preserves global geometry and neural networks tolerate minor variations. Consequently, block-wise sampling removes the global search in standard FPS while maintaining accuracy.

**Block-Wise Neighbor Searching.** Neighbor searching in PNNs consists of two key operations: grouping and interpolation, both require a search space and a center point cloud to identify neighboring points of the center point cloud within the search space. With Fractal, the point cloud is partitioned into blocks, enabling both block-wise grouping and block-wise interpolation, as illustrated in Fig. 7. The center points correspond to leaf nodes in the point cloud tree. In this approach, if the leaf node is at depth 1, its search space is restricted to the same node. For deeper leaf nodes, the search space expands to include their immediate parent nodes, which ensures a broader scope and is sufficient for maintaining network accuracy as shown in Fig. 14.

Block-wise neighbor searching follows the same DFT order used in Fractal’s block storage, enabling data reuse and minimizing memory access. As illustrated in Fig. 7(a), the DFT starts from  $B_0$ , first processing the leaf node  $B_3$  (depth 2). Its parent node  $B_1$  (depth 1, containing  $B_3$  and  $B_4$ ) is loaded on-chip to expand  $B_3$ ’s searching space. When moving to  $B_4$  (depth 2), its parent node  $B_1$  (depth 1) is already cached, allowing direct data reuse and eliminating redundant memory access. The same process applies to  $B_5$  and  $B_6$ , where parent data are shared across sibling nodes without additional memory overhead. This DFT-based traversal supports a fully streamed block-wise neighbor searching pattern while avoiding additional parent-node access overhead. Moreover, this block-wise strategy relaxes the global-search dependencies of conventional neighbor searching, requiring only relevant blocks to be loaded on-chip. Consequently, all operations can be executed entirely on-chip, even for large-scale point cloud processing.

**Block-Wise Gathering.** Gathering retrieves feature values based on the point indices generated by neighbor searching. Conventional gathering requires a global search, meaning all feature data must be loaded on-chip. When the input size

exceeds the on-chip buffer capacity, it incurs frequent DRAM accesses and substantial latency. In contrast, our block-wise gathering confines the search to relevant blocks, following the same search-space rule as block-wise neighbor searching. Enabled by Fractal, only related blocks need to be kept on-chip, enabling fully on-chip execution with limited buffer capacity and significantly reducing latency and energy consumption.

## V. HARDWARE ARCHITECTURE OF FRACTALCLOUD

This section presents the proposed hardware architecture of *FractalCloud*, including the overall architecture design and detailed explanations of the submodules.

### A. Overall Architecture

*FractalCloud* is a domain-specific accelerator designed to efficiently support PNNs across different input scales and diverse point operations. As illustrated in Fig. 8, the architecture comprises a fractal engine, reuse-and-skip-enabled point units (RSPUs), gathering units, pooling units, a systolic processing element (PE) array, a memory module, and a RISC-V core. The fractal engine implements all mainstream partitioning methods. RSPUs execute all point operations with block-parallel architecture supported. The gathering unit retrieves feature values based on the spatial indices, while the pooling unit performs max pooling operations in PNNs. The systolic PE array, configured as a  $16 \times 16$  array following [27]–[29], handles convolutional and feedforward computations. The memory subsystem integrates a 274KB multi-bank global buffer following [28] and a memory interface with a Network-on-Chip (NoC) for intra-chip communication and a Direct Memory Access (DMA) unit for DRAM interactions. A single-core six-stage RV32IMAC RISC-V processor, extended with Zifencei and additional control instructions, manages module configuration and data transfers. A lightweight configuration module between the RISC-V core and computation modules handles these instructions. The core writes control data into a buffer within this configuration module, which then segments and packages the data based on each computation module’s instruction length and dispatches it for execution.

**Reconfigurability and Scalability.** Above hardware modules provide fundamental support for core operations in PNNs, enabling *FractalCloud* to achieve high reconfigurability and scalability. Each module in *FractalCloud* is equipped with dedicated instructions and configurable registers, allowing the RISC-V CPU to dynamically adjust its functionality and compose different point operations. For example, the Fractal engine supports multiple partitioning strategies, including Fractal, uniform, and KD-tree-based methods, while the RSPUs and gather units can flexibly adapt to various sampling, grouping, and interpolation patterns. This modularity allows *FractalCloud* to efficiently accommodate a wide range of PNN networks with different depths, resolutions, and sampling rates [20]–[24]. Moreover, *FractalCloud* scales efficiently across varying workloads. Large operations, such as large matrix multiplications, are automatically tiled by the compiler based on input size and hardware configurations. It generates

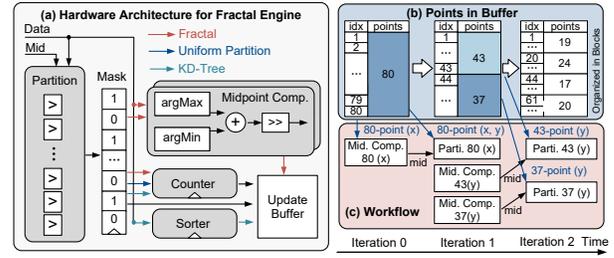


Fig. 9: The (a) hardware architecture of fractal engine. (b) The optimized memory flow and (c) workflow for the Fractal.

a hardware-compatible storage allocation and configures the computing units, enabling seamless scaling with low overhead. Large-scale point operations are decomposed into fine-grained computations via the proposed Fractal method and executed in parallel following the BPO, all scheduled and controlled directly on-chip. With such high reconfigurability and scalability, *FractalCloud* achieves efficient acceleration for PNNs across different workloads, as shown in Fig. 13.

### B. Fractal Engine

The fractal engine efficiently supports the proposed Fractal method as well as mainstream partitioning methods such as uniform and KD-tree, as shown in Fig. 9(a). It comprises partition units, midpoint computation units, counters, and sorter units. Partition units use parallel comparators to partition points into sub-blocks based on the input Mid, while midpoint computation units calculate the midpoint of sub-block via min–max averaging with an addition and right-shift operation. Counters track point distribution after partitioning. And sorter units, based on a merge-sort hardware structure [28], provide the median for KD-tree partitioning. The dataflows of different partitioning methods are illustrated in Fig. 9(a), demonstrating the high hardware reuse of the fractal engine.

We have optimized the workflow for the proposed Fractal, as shown in Fig. 9 (b) and (c) based on the same example of 80-point cloud with a threshold of 24 in Fig. 6. The partition unit and midpoint computing unit operate in a pipelined manner to enhance throughput. In Iteration 0, the engine computes the midpoint along the x-dimension of the input point cloud. In Iteration 1, both x- and y-dimensions are processed: the x-dimension is used by the partition unit to divide points based on the previously computed midpoint, while the y-dimension is used by the midpoint computing unit to determine new midpoints for the resulting sub-blocks. After each iteration, the point cloud is partitioned into two sub-blocks based on mask values and stored in memory. Blocks with more points than the threshold (e.g., the 43-point and 37-point blocks) are further partitioned and passed to the next iteration. By Iteration 2, the entire point cloud is hierarchically partitioned, effectively optimizing the memory structure for subsequent processing.

**Block-Wise Feature Access for Gathering.** The proposed block-wise gather method narrows the search space from the entire point cloud to individual blocks, with each gather unit processing only its assigned block. Fig. 10 illustrates

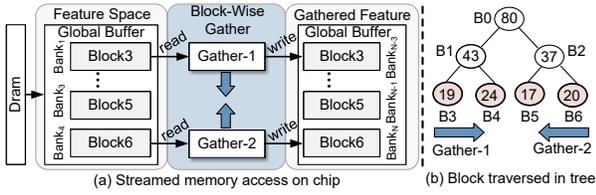


Fig. 10: (a) The memory communication for block-wise gathering and (b) the traversing order of blocks in tree hierarchy.

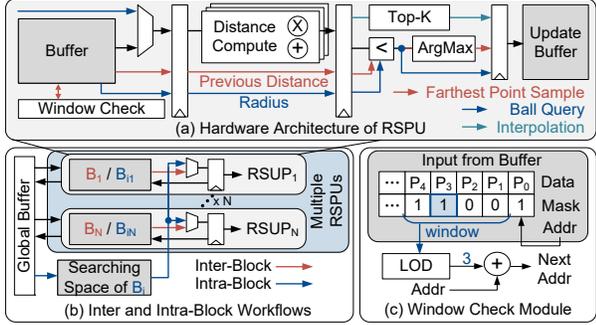


Fig. 11: (a) The hardware design of reuse-and-skip-enabled point unit (RSPU) and (b) inter-block and intra-block workflow for multi-core RSPUs. (c) The hardware design for window check module, with a lowest-one detector (LOD) unit.

the memory access pattern for two gather units, where each block is stored in a separate memory bank. Gather Unit 1 performs a DFT traversal from left to right, processing Block 3 to Block 4, while Gather Unit 2 traverses in the opposite direction, processing Block 6 to Block 5. Both meet at the tree’s midpoint, completing the gathering process. For single-block gathering, the current block and its parent data are read from DRAM, and parent data are shared among all child nodes, increasing data reuse.

Block-wise gathering confines each unit to its designated memory bank, avoiding SRAM conflicts. Because only block-level data is retrieved, all necessary information fits in SRAM, eliminating DRAM lookups during processing. Moreover, thanks to organized data after Fractal, all DRAM accesses are also streamed rather than random, unlike conventional gathering. Detailed evaluations are provided in Section VI-C.

### C. Reuse-and-Skip-Enabled Point Unit (RSPU)

We propose multi-core reuse-and-skip-enabled point units (RSPU) that support both sampling and neighbor searching. It features a flexible block-level parallelism to enable efficient data reuse and redundant computation skipping, providing high hardware efficiency for block-parallel point operations. As shown in Fig. 11(a), the RSPU comprises a windows check module, distance computing units, comparison units, an argmax unit, a top-k unit, and local buffers. RSPU operates in a pipeline manner. The distance computing unit computes Euclidean distances. The comparison unit and argmax unit mainly incorporate comparators, while the argmax unit stores the maximum value from each comparison. The top-k unit

### Algorithm 2 Flexible-Parallel Workflow on Multi-RSPUs

**Require:** Blocks  $\{B_i\}$ , parent block  $B_p$ , mode  $m$ , RSPUs  $\{R_i\}$

- 1: **if**  $m = \text{FPS}$  **then** ▷ Inter-block parallelism
- 2:     **for all**  $R_i$  **in parallel do**
- 3:          $B'_i \leftarrow R_i(B_i, m)$
- 4: **else** ▷ Neighbor Search
- 5:     **for all**  $B_i$  **do**
- 6:         share  $B_p$  data via buffer
- 7:         **for all**  $R_i$  **in parallel do** ▷ Intra-block parallelism
- 8:              $B'_{i,p_i} \leftarrow R_i(p_i, B_p, m)$  ▷  $p_i$ : center points of  $B_i$
- 9: **Update**  $\{B'_i\}$  **in memory**

employs an merge sort mechanism to compare and sort the input data against stored values, ultimately selecting the K smallest values from the input. The local buffer temporarily stores input/output data and mask information indicating whether each point has been sampled, ensuring efficient data management before writing to SRAM. The workflows of an RSPU for FPS, grouping, and interpolation are shown in Fig. 11(a), demonstrating high hardware reuse.

**Flexible-Parallel Workflow of RSPUs.** Building on the RSPU design and BPPO, we propose a flexible-parallel workflow that exploits both intra-block and inter-block parallelism to accelerate point operations on multi-core RSPUs. The intra-block parallelism (blue lines in Fig. 11(b), Row 6-8 in Alg. 2) targets neighbor searching, with multiple RSPUs concurrently process different center points within a block while reusing search-space data from parent blocks, thereby saving memory accesses. The inter-block parallelism (red lines in Fig. 11(b), Rows 2–3 in Alg. 2) enables independent FPS execution on different blocks, achieving concurrent processing and avoiding redundant computations through a window-check mechanism discussed later.

**Data Reuse of Coordinates.** As shown by the blue lines in Fig. 11(b), intra-block workflows allow multiple RSPUs to perform neighbor searching concurrently with high data reusing. Specifically, each RSPU computes different center points within the same block while sharing the search space data from a dedicated buffer. This approach enhances data reuse by the number of RSPUs times, reducing redundant memory accesses for neighbor searching operations.

**Computation Skipping Enabled by Window Check.** In standard FPS, each iteration performs a global traversal over all input points [28]. However, already sampled points won’t be sampled again, leading to redundant computations. To address this inefficiency, we introduce a window-check module that filters out sampled points (0s in Fig.11(c)) before each iteration, ensuring only valid candidates (1s in Fig.11(c)) participate in distance updates. As shown in Fig.11(c), the module has a mask window input, indicating the sampling status of upcoming points. A lowest-one detector (LOD), implemented with a priority encoder, identifies the nearest ‘1’ to locate the next valid point and update the next memory address accordingly. This mechanism effectively bypasses redundant memory accesses and computations associated with already sampled points. Leveraging inter-block parallelism,

TABLE I: Evaluated Networks and Datasets.

Model	Notation	Task	Dataset	Scene
PointNet++	PN++ (c)	Classification	ModelNet40	Object
PointNeXt	PNXt (c)			
PointNet++	PN++ (ps)	Part Segmentation	ShapeNet	Object
PointNeXt	PNXt (ps)			
PointNet++	PN++ (s)	Segmentation	S3DIS	Indoor
PointNeXt	PNXt (s)			
PointVector	PVr (s)			

where each FPS operates independently on its assigned RSPU, the window-check mechanism enables concurrent redundancy skipping across all parallel FPS operations.

## VI. EVALUATION

### A. Experimental Setup

**Software Benchmarks.** We evaluate the proposed *FractalCloud* using three advanced point cloud neural networks, PointNet++ [22], PointNeXt [21], and PointVector [20], across three fundamental tasks: classification, part segmentation, and semantic segmentation, as summarized in Table I. To assess performance, we define two input scales: small-scale inputs (ranging from 1K to 4K points from ModelNet40 [33], ShapeNet [44], and S3DIS [34] datasets), and large-scale inputs (ranging from 8K to 289K points from the S3DIS dataset). For accuracy benchmarking, we retrain and evaluate publicly available models from original repositories using PyTorch [45], ensuring results comparable to those reported in the original papers. Classification performance is measured using overall accuracy (OA), whereas segmentation performance is evaluated using the mean Intersection-over-Union (mIoU) metric on the Area-5 test of the S3DIS dataset.

**Hardware Baselines.** To evaluate the proposed *FractalCloud*, we employ two categories of hardware baselines. The first is a server-level GPU platform equipped with an NVIDIA TITAN RTX GPU, where all PNN implementations are CUDA-optimized using the Openpoints Library [46]. The GPU is warmed up before measurement, and latency is obtained by averaging 20 synchronized inference runs measured with `time()`. Energy consumption of GPU is measured using the NVIDIA Management Library [47]. The second category includes state-of-the-art (SOTA) PNN accelerators: Mesorasi [27], PointAcc [28], and Crescent [29]. For a fair comparison, since Mesorasi and Crescent utilize the delayed-aggregation technique to reduce the latency of MLPs, we also integrate it into *FractalCloud*. Moreover, as Mesorasi and Crescent do not support the FPS method in the sampling step, which is critical for large-scale point cloud processing, we equip them with the FPS engine from PointAcc. Key hardware specifications (e.g., frequency, cores, and data bit-width) are aligned with all SOTA accelerators to ensure fair comparisons, and all accelerators are evaluated under DDR4 memory for consistency.

**Hardware Implementation.** The proposed *FractalCloud* is implemented in Verilog HDL. We synthesize the hardware design using the Synopsys Design Compiler [48], and

TABLE II: Evaluated Hardware Accelerators.

Accelerator	Mesorasi [27]	PointAcc [28]	Crescent [29]	<i>FractalCloud</i>
Cores	16x16	16x16	16x16	16x16
SRAM (KB)	1624	274	1622.8	274
Frequency	1GHz	1GHz	1GHz	1GHz
Area (mm <sup>2</sup> )	4.59	1.91	4.75	1.5
DRAM	DDR4-2133	DDR4-2133	DDR4-2133	DDR4-2133
Bandwidth	17GB/s	17GB/s	17GB/s	17GB/s
Technology	28nm	28nm	28nm	28nm
Peak Performance	512 GOPS	512 GOPS	512 GOPS	512 GOPS

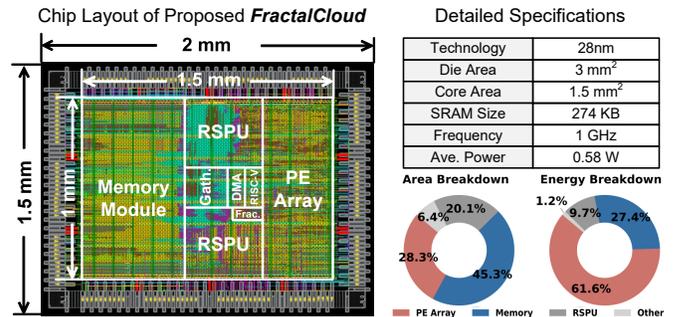


Fig. 12: Chip layout, specifications, and breakdown of on-chip area and energy consumption for *FractalCloud*.

place&route using Cadence Innovus [49] in TSMC 28 nm technology. Area for all modules are derived from post-layout reports, with SRAM generated by the memory compiler provided by TSMC. Power is measured through post-layout simulations using the real network data, ensuring more accurate estimation than prior synthesis-based and CACTI-based works [28], [29], [50]. We use 16-bit half-precision floating-point arithmetic for all computations to align with all SOTA works and preserve network accuracy. For end-to-end performance evaluation, we compare *FractalCloud* with SOTA baselines using a cycle-level accurate simulator based on the open-source framework [51]. The simulator is verified through RTL simulation. We do not simulate Mesorasi directly, as Crescent is an improved version of Mesorasi. Instead, the results for Mesorasi are sourced and scaled from the original paper. We use DRAMsim3 [52] to estimate the off-chip DRAM power consumption. To ensure a fair comparison, all designs are scaled to the 28nm technology node using the scaling method in [53]. The total core area of proposed *FractalCloud* is 1.5mm<sup>2</sup>, with an average power consumption of 0.58W, as shown in Fig. 12. The hardware configurations of *FractalCloud* and SOTA works are detailed in Table II.

### B. Overall Results

**Network Accuracy.** We evaluate our proposed methods using PointNet++, PointNeXt, and PointVector models for classification and segmentation tasks, respectively, by replacing the original point operations with our proposed techniques and retraining networks, as shown in Fig. 14. The threshold for the proposed Fractal method is set to 64 for classification and 256

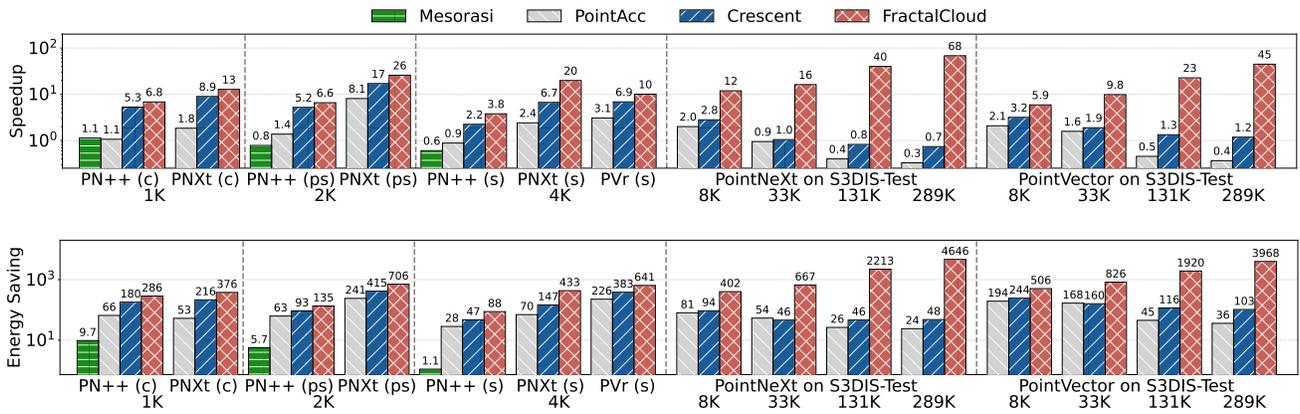


Fig. 13: Comparison of speedup and energy saving achieved by Mesorasi, PointAcc, Crescent, and proposed *FractalCloud*, normalized to GPU performance. Higher is better.

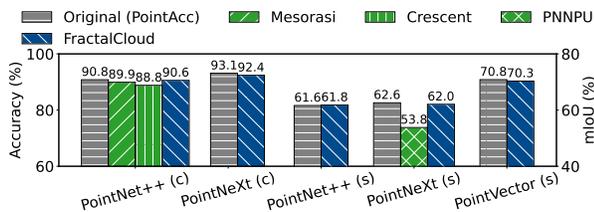


Fig. 14: Network accuracy comparison among: Original (PointAcc [28]), Mesorasi [27], Crescent [29], PNNPU [32], and *FractalCloud* on different workloads. PointAcc is lossless.

for segmentation. PointAcc [28] preserves network accuracy but offers limited hardware improvement. Other prior works exhibit greater accuracy loss on the same tasks compared to *FractalCloud*. Based on the same baseline, Mesorasi and Crescent have 0.9% [27] and 2% [29] loss for PointNet++ (c), respectively, due to network structure modifications and approximate memory access for improved hardware efficiency. PNNPU has 8.8% loss for PointNetx (s) [32], [38] due to degraded feature representations after partitioning. In contrast, *FractalCloud* achieves the best overall hardware performance, as shown in Fig. 13, while maintaining accuracy within 0.7% loss across all models and input scales, demonstrating the best trade-off between hardware efficiency and model accuracy.

We further evaluate individual point operations to explain how *FractalCloud* maintains high accuracy. Our findings indicate that directly applying the proposed block-wise sampling and block-wise interpolation preserves network accuracy, with less than 0.2% accuracy loss. This is primarily because the blocks after Fractal align with the overall geometry of the point cloud, thereby maintaining the representational capacity of the points. Block-wise grouping introduces slight accuracy degradation, primarily due to numerical differences between local and original global searches. However, such errors can be mitigated through retraining as evidenced by Fig. 14, since the extended search space in the block-wise grouping provides sufficient candidate points for neighbor selection. Block-wise

gathering has no impact on network accuracy, as it does not modify original feature values. After retraining, all networks recover their accuracy, achieving performance almost identical to the original models.

**Speedup.** Fig. 13 presents the inference speedup achieved by *FractalCloud* compared to state-of-the-art accelerators, using GPU performance as the baseline. In the small-scale input scenario, *FractalCloud* achieves speedups of 6.9 $\times$ , 7.6 $\times$ , 2.7 $\times$ , and 19.4 $\times$  over Mesorasi, PointAcc, Crescent, and GPU, respectively. For large-scale inputs, where point operations dominate latency, *FractalCloud* demonstrates even more substantial improvements, with 63.4 $\times$ , 27.8 $\times$ , and 27.4 $\times$  speedup over PointAcc, Crescent, and GPU. This performance gain primarily stems from three key optimizations. First, the Fractal method partitions the point cloud into multiple blocks with low hardware complexity, enabling balanced memory patterns for further optimization. Second, block-parallel point operations decompose and parallelize all point operations from global search to local search. Third, the optimized hardware architecture enables fully parallel execution of both partitioning and point operations. Together, these optimizations ensure an average speedup of 21.7 $\times$  compared to SOTA accelerators.

**Energy Saving.** Fig. 13 shows the energy savings of *FractalCloud* under the same settings. In small-scale input scenarios, *FractalCloud* achieves energy reductions of 44.3 $\times$ , 4.1 $\times$ , 1.9 $\times$ , and 380 $\times$  compared to Mesorasi, PointAcc, Crescent, and GPU, respectively. The improvements become more pronounced in large-scale inputs, with energy savings of 56.8 $\times$ , 28.2 $\times$ , and 1893 $\times$  over PointAcc, Crescent, and GPU. Similar to the speedup results, these energy savings are rooted in *FractalCloud*'s three optimizations. Mesorasi suffers from significant energy overhead from point operations, accounting for approximately 90% of total energy consumption, due to the lack of optimization. PointAcc faces energy bottlenecks from excessive DRAM access caused by the global search problems and its limited 274KB on-chip buffer. Crescent mitigates DRAM energy consumption with a larger buffer and streamed memory access, but at the cost of increased SRAM energy

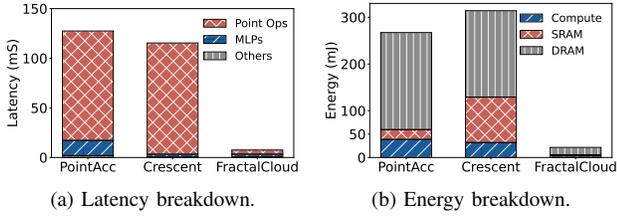


Fig. 15: (a) Latency breakdown and (b) energy breakdown for *FractalCloud* and SOTA accelerators evaluated on PointNeXt segmenting S3DIS dataset.

usage, leading to only marginal gains compared to PointAcc. As a result, *FractalCloud* achieves an average energy saving of 27× compared to SOTA accelerators.

**Latency and Energy Breakdown Analysis.** Fig. 15 presents a detailed breakdown of latency and energy consumption among PointAcc, Crescent, and *FractalCloud* when executing PointNeXt on the S3DIS test dataset with 33K input points. In PointAcc, point operations rely on global search, requiring access and computation over the entire data. The small on-chip buffer cannot accommodate large-scale point cloud data, causing around 41% of data to be retrieved from DRAM during point operations. These frequent off-chip accesses significantly increase both latency and energy consumption. Crescent reduces MLP latency by adopting the delay-aggregation method from [27], but this increases the search space for gather operations, worsening the latency. To compensate, Crescent employs an improved KD-tree for memory streaming and a large 1622.8KB on-chip buffer. While this mitigates the problem of frequent DRAM access, it significantly increases SRAM power usage, and even causes a higher total energy consumption than PointAcc. As a result, Crescent achieves 1.1× speedup but 17% more energy consumption compared to PointAcc. In contrast, *FractalCloud* leverages Fractal to organize the point cloud, providing efficient partitioning for further optimization. BPPO reduces the search space of point operations to blocks, while RSPUs facilitate flexible and parallel hardware execution of these operations. Together, with 33K input points, these optimizations enable *FractalCloud* to achieve an average of 16.2× latency reduction, 8.5× computation-related energy savings, and 14.7× memory-related energy savings, compared with PointAcc and Crescent, all within a limited 274KB buffer.

### C. Ablation Study on Proposed Techniques

This section is organized as follows: We first compare the proposed Fractal with SOTA partitioning approaches and analyze the effect of different threshold settings. We then evaluate each point operation in BPPO to quantify each contribution. Finally, we assess the effectiveness of the proposed RSPU.

**Ablation Study on Fractal.** We compare Fractal with three representative partitioning techniques: uniform [32], [43], KD-tree [29], [42], and Octree partitionings [41], [54], [55], as shown in Fig. 16. Uniform partitioning offers minimal complexity but overlooks point cloud distribution, causing large

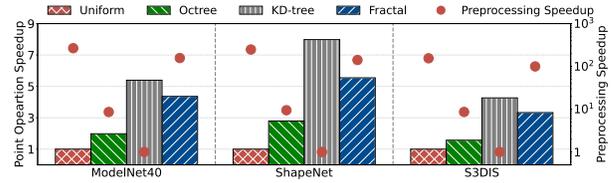


Fig. 16: Speedup comparisons of point operation (bars, with uniform as baseline) and partitioning (dots, with KD-tree as baseline) among uniform, octree, KD-tree, and Fractal.

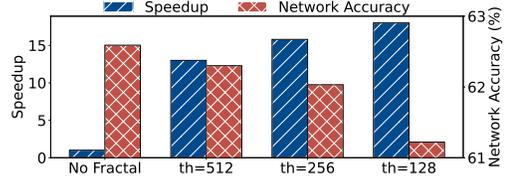


Fig. 17: Impact of threshold selection in Fractal on hardware speedup and network accuracy for PointNeXt segmenting S3DIS dataset.

variance within blocks and reducing point operation efficiency, with 8.8% accuracy loss in large-scale tasks [38]. KD-tree achieves balanced blocks via recursive median selection but suffers from recursive and exclusive sorts, contributing up to 71% of total latency when applied on *FractalCloud*. Octree, as a uniform-based extension with dynamic subdivision, slightly reduces imbalance with increased control complexity [41], but still incurs 3% accuracy loss due to residual uniform partitioning [38], [55]. Fractal addresses these limitations by adaptively partitioning points according to shape information, preserving geometric features while remaining hardware-friendly. It achieves 133× faster partitioning than KD-tree, 14.9× faster than Octree, and improves point operation performance by 2.1× over Octree and 4.4× over uniform partitioning. Additionally, Fractal incurs only 1% area and 0.2% energy overhead in end-to-end inference.

Threshold selection in Fractal trades off hardware speedup and accuracy, as shown in Fig. 17. Overall, lower thresholds ( $th$ ) increase speedup by reducing points per block but may impair accuracy due to insufficient neighbors for feature extraction. Under extreme cases, excessive thresholds ( $th=4k$ ) preserve accuracy but increase per-block computation complexity and memory traffic, with only 4.6× speedup. Overly small thresholds ( $th=8$ ) over-partition the input, disrupting geometry, and causing random-like FPS and more errors in neighbor searching, causing over 8% accuracy loss despite a 21× speedup. Through a greedy design-space exploration,  $th=256/64$  is chosen for large/small-scale inputs, achieving an optimal trade-off between speedup and accuracy.

**Ablation Study on Block-Parallel Points Operations (BPPO).** Fig. 18 presents the incremental speedup and energy savings of *FractalCloud* during PointNeXt inference on the S3DIS dataset with 289K points. The *Baseline* denotes *FractalCloud* without optimizations, while *Baseline (Meso)* applies Mesorasi’s delayed-aggregation [27] for MLPs, offering only

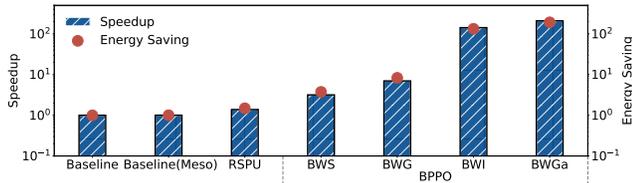


Fig. 18: Comparison of speedup and energy savings after applying RSPU and each optimized point operation in BPPO.

a marginal 1.004 $\times$  gain. Adding RSPU eliminates redundant computations and improves data reuse in hardware, yielding 1.37 $\times$  speedup and 1.48 $\times$  energy savings over *Baseline (Meso)*. BPPO further mitigates the global-search latency with parallel processing. Now sampling dominates 57% of latency. Block-wise sampling (BWS) limits computation to local blocks, achieving 2.3 $\times$  speedup and 2.5 $\times$  energy savings. Grouping and interpolation then dominate latency (55% and 43%). Block-wise grouping (BWG) and interpolation (BWI) restrict searches to relevant blocks, providing 2.2 $\times$  and 20 $\times$  speedup and 2.2 $\times$  and 16 $\times$  energy savings, respectively. Finally, block-wise gathering (BWGa) alleviates DRAM fetches by enabling fully on-chip retrievals, adding 1.5 $\times$  speedup and 1.4 $\times$  energy savings. Overall, *FractalCloud* achieves 209 $\times$  speedup and 192 $\times$  energy savings over both *Baseline* and *Baseline (Meso)*.

**Ablation Study on Reuse-and-Skip-Enabled Point Unit (RSPU).** To efficiently accelerate point operations, we propose the reuse-and-skip-enabled point unit (RSPU), which supports block-parallel computing and integrates data reuse for neighbor searching and redundant computation skipping for sampling. Existing SOTA designs [28], [40] mainly adopt point-level parallelism, where multiple input points are processed simultaneously to compute a single output point. Although this improves throughput, it requires complex hardware components (e.g., N-sorters, N-mergers) and cannot eliminate redundant computations in iterative FPS. Moreover, processing only one center point per iteration limits data reuse in neighbor searching, as search-space data cannot be effectively shared. In contrast, our RSPU employs a window-check mechanism to skip redundant FPS computations, achieving 3.6 $\times$  speedup and 3.4 $\times$  memory-access reduction over PointAcc. For neighbor searching, intra-block parallelism enables concurrent processing of multiple center points with the shared searching space, significantly improving data reuse with a 7.6 $\times$  memory-access reduction. Overall, these optimizations improve PNN inference by 1.37 $\times$  speedup and 1.48 $\times$  energy efficiency.

#### D. Discussion

**Imbalance effect in Fractal.** Outliers, irregular shapes, and different scenes in point cloud may cause partial imbalance in Fractal, but their impact on performance and memory load is minimal. Fractal enforces a threshold  $th$  that limits the number of points per block. If an imbalance occurs, recursive partitioning continues until all blocks meet this constraint. Even under extreme shapes, such as two distant dense regions, the maximum imbalance among blocks after Fractal is strictly

bounded by  $th$ , while spatial partitioning methods can reach the full input size in the worst case. Since both latency and memory load are dominated by the largest block, Fractal effectively constrains hardware imbalance. Moreover, Fractal cyclically utilizes all three spatial dimensions during partitioning, avoiding extreme cases in scene-level point clouds [34] where coplanar points make one dimension non-splittable. Although outliers may lead to a few underfilled blocks, they represent only 0.5-2.5% of points in S3DIS and are further mitigated by the threshold mechanism. Therefore, experiments on S3DIS show that, compared with a strictly balanced case, partial imbalance in Fractal increases overall latency by only 3.0% and 2.8% for PointNeXt and PointVector, respectively, confirming its negligible impact.

**Asymptotic Limitation.** *FractalCloud* can also accelerates extremely large-scale ( $> 500K$ ) point clouds, achieving an asymptotic speedup of 105.7 $\times$  over GPU at 1M points on S3DIS with PointNeXt. But the accuracy may degrade if the block size is fixed, consistent with the trend shown in Fig. 17, where larger block sizes can better preserve accuracy. This reflects a performance-accuracy tradeoff rather than a limitation of the Fractal itself. On the hardware side, the maximum supported input size mainly depends on DRAM capacity and the network configuration (e.g., a 24GB DRAM can handle 3M-point PointNeXt). In practice, however, existing PNNs rarely exceed 300K points, which already fully cover the needs of edge-level devices [5] targeted by our design.

**Potential Adaptations.** While *FractalCloud* focuses on 3D point clouds, Fractal can also bring benefits to other data domains. In general, Fractal is applicable when (i) data preserve organized structures (e.g., graphs, tables), rather than being purely random, and (ii) local coherence exists among elements. For example, Fractal can exploit spatial locality in dynamic graphs to accelerate their construction and updates in DGCNN [24], [56], and may benefit high-dimensional tabular data [57] via locality-aware feature grouping and selection.

## VII. CONCLUSION

The rapid growth of large-scale 3D point clouds presents new challenges and opportunities for efficient hardware acceleration of point cloud neural networks. To meet the emerging demands, we propose *FractalCloud*, a fractal-inspired accelerator for efficient large-scale 3D point cloud processing. Specifically, *FractalCloud* introduces the Fractal method to achieve balanced partitioning with low hardware complexity and maintained network accuracy. Block-parallel point operations are developed, decomposing and parallelizing all point operations from global search to local search. Moreover, a dedicated hardware design integrates on-chip fractal, flexible parallel processing, and optimized data reuse with redundancy elimination, maximizing the efficiency of our fractal approach. Implemented as a chip layout in 28nm technology, *FractalCloud* achieves an average speedup of 21.7 $\times$  and energy reduction of 27 $\times$  compared to state-of-the-art PNN accelerators.

## ACKNOWLEDGMENTS

This work was supported in part by NSF grants 2328805 and 2112562, by ARO W911NF-23-2-0224, as well as the Beyond the Horizon Initiative through the Pratt School of Engineering, Duke University. The authors would like to thank the anonymous reviewers for their constructive feedback, which helped improve the quality of this work. We also extend our gratitude to Junyao Zhang, Haoxuan Shan, Mingyuan Ma, and Jonathan Ku for their technical support and insightful discussions.

## APPENDIX

### A. Abstract

Our artifact contains the algorithmic implementation of *FractalCloud*, enabling the evaluation of network performance under the proposed fractal-inspired point cloud processing architecture. The artifact includes the PyTorch implementation of the Fractal operation and Block-Parallel Point Operations, scripts to reproduce classification and segmentation results on ModelNet40 and S3DIS, and pretrained models for PointNet++, PointNeXt-S, and PointVector-L. The purpose of this artifact is to validate algorithmic correctness and network performance as reported in the paper.

### B. Artifact check-list (meta-information)

- **Algorithm:** Fractal operation; block-parallel point operations; Point cloud neural networks
- **Program:** Python (PyTorch implementation)
- **Compilation:** N/A (Conda environment and Docker)
- **Transformations:** N/A
- **Binary:** Docker Image
- **Model:** PointNet++, PointNeXt-S, PointVector-L
- **Data set:** ModelNet40, S3DIS
- **Run-time environment:** Ubuntu 24.04.3 LTS (Host), Docker container with Ubuntu 16.04.7 LTS, Python 3.7.16, PyTorch 1.10.1, CUDA 11.3
- **Hardware:** x86-64 server with one NVIDIA TITAN RTX (24GB)
- **Run-time state:** Standard inference with pretrained models
- **Execution:** Python scripts
- **Metrics:** Classification accuracy, segmentation mIoU
- **Output:** Reported accuracy and segmentation metrics
- **Experiments:** Classification on ModelNet40; segmentation on S3DIS
- **How much disk space required (approximately)?:** ~96GB total (including Docker, dependencies, datasets, pretrained models, and cache)
- **How much time is needed to prepare workflow (approximately)?:** ~20 minutes (Docker)
- **How much time is needed to complete experiments (approximately)?:** 5–10 minutes for classification; 2–14 hours for segmentation
- **Publicly available?:** Yes (<https://github.com/Yuzhe-Fu/FractalCloud>)
- **Code licenses (if publicly available)?:** MIT license
- **Data licenses (if publicly available)?:** Original dataset licenses
- **Workflow automation framework used?:** N/A
- **Archived (provide DOI)?:** <https://zenodo.org/records/17851585>

### C. Description

1) *How to access:* The source code is publicly available at: <https://github.com/Yuzhe-Fu/FractalCloud>. The artifact sources are also archived at Zenodo: <https://zenodo.org/records/17851585>. Pretrained weights for all evaluated models (PointNet++, PointNeXt-S, and PointVector-L) and inference logs for evaluated workflows are available at: <https://huggingface.co/YuzheFu/FractalCloud>.

2) *Hardware dependencies:* A server equipped with a GPU is required for evaluation. The artifact has been tested with an NVIDIA TITAN RTX GPU (24GB).

3) *Software dependencies:* The artifact relies on the following software components:

- Host OS: Ubuntu 24.04.3 LTS
- Docker (optional): Docker 28.3.3
- Docker Image (provided): Ubuntu 16.04.7 LTS with all required dependencies preinstalled
- Python 3.7.16
- Pytorch 1.10.1
- CUDA 11.x (only required for local installation without the provided Docker image)
- Conda 4.9.2

4) *Data sets:* The evaluated datasets are:

- ModelNet40 (classification)
- S3DIS (segmentation)

5) *Models:*

- PointNet++
- PointNeXt-S
- PointVector-L

Pretrained weights and inference logs (exact outputs reported in the paper) are available at <https://huggingface.co/YuzheFu/FractalCloud> for validating reproduced results.

### D. Installation

Installation can be done through the provided Docker image or via running local setup scripts. A detailed step-by-step installation guide is provided in the project README file.

### E. Experiment workflow

Example execution commands to reproduce classification and segmentation results are included in the README. Running the provided scripts will produce network accuracies matching the results reported in the paper.

### F. Evaluation and expected results

The artifact is expected to reproduce the network accuracy metrics reported in the paper. Specifically, it validates the numerical consistency of the fractal implementation and its impact on network performance across multiple point cloud models.

### G. Experiment customization

Users can modify experiment configurations, model architectures, and dataset paths by editing the script arguments and YAML configuration files as described in the project README.

## H. Notes

N/A

## I. Methodology

Submission, reviewing and badging methodology:

- <https://www.acm.org/publications/policies/artifact-review-and-badging-current>
- <https://cTuning.org/ae>

## REFERENCES

- [1] A. Rani, D. Ortiz-Arroyo, and P. Durdevic, "Advancements in point cloud-based 3d defect classification and segmentation for industrial systems: A comprehensive survey," *Information Fusion*, vol. 112, p. 102575, 2024.
- [2] P. Guo, H. Zhu, H. Ye, T. Li, and T. Chen, "Revisiting 3d visual grounding with context-aware feature aggregation," *Neurocomputing*, vol. 601, p. 128195, 2024.
- [3] X. Yan\*, R. Chen\*, B. Zhang, H. Ye, R. Xia, J. Yuan, H. Zhou, X. Cai, B. Shi, W. Shao *et al.*, "Spot: Scalable 3d pre-training via occupancy prediction for learning transferable 3d representations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [4] C. Zhou, M. Liu, S. Qiu, X. Cao, Y. Fu, Y. He, and H. Jiao, "Sagitta: An energy-efficient sparse 3d-cnn accelerator for real-time 3-d understanding," *IEEE Internet of Things Journal*, vol. 10, no. 23, pp. 20 703–20 717, 2023.
- [5] R. Roriz, H. Silva, F. Dias, and T. Gomes, "A survey on data compression techniques for automotive lidar point clouds," *Sensors*, vol. 24, no. 10, p. 3185, 2024.
- [6] P. Dong, Z. Chen, K. Li, L. Chen, K.-T. Cheng, and F. An, "A 1920×1080 129fps 4.3 pj/pixel stereo-matching processor for pico aerial vehicles," in *ESSCIRC 2023-IEEE 49th European Solid State Circuits Conference (ESSCIRC)*. IEEE, 2023, pp. 345–348.
- [7] P. Dong, Z. Chen, Z. Li, Y. Fu, L. Chen, and F. An, "A 4.29 nj/pixel stereo depth coprocessor with pixel level pipeline and region optimized semi-global matching for iot application," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 69, no. 1, pp. 334–346, 2021.
- [8] P. Dong, Z. Li, Z. Chen, R. Yao, H. Deng, W. Zhang, Y. Zhang, L. Chen, C. Wang, and F. An, "A 139 fps pixel-level pipelined binocular stereo vision accelerator with region-optimized semi-global matching," in *2021 IEEE Asian Solid-State Circuits Conference (A-SSCC)*. IEEE, 2021, pp. 1–3.
- [9] L. Liu, B. Su, J. Jiang, G. Wu, C. Guo, C. Xu, and H. F. Yang, "Towards accurate and efficient 3d object detection for autonomous driving: A mixture of experts computing system on edge," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 25 903–25 913.
- [10] Y. Cui, R. Chen, W. Chu, L. Chen, D. Tian, Y. Li, and D. Cao, "Deep learning for image and point cloud fusion in autonomous driving: A review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 2, pp. 722–739, 2021.
- [11] L. Kästner, V. C. Frasinianu, and J. Lambrecht, "A 3d-deep-learning-based augmented reality calibration method for robotic environments using depth sensor data," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1135–1141.
- [12] Y. Wang, S. Zhang, B. Wan, W. He, and X. Bai, "Point cloud and visual feature-based tracking method for an augmented reality-aided mechanical assembly system," *The International Journal of Advanced Manufacturing Technology*, vol. 99, pp. 2341–2352, 2018.
- [13] J. D. Stets, Y. Sun, W. Corning, and S. W. Greenwald, "Visualization and labeling of point clouds in virtual reality," in *SIGGRAPH Asia 2017 Posters*, 2017, pp. 1–2.
- [14] S. Li, Y. Zhao, C. Li, B. Guo, J. Zhang, W. Zhu, Z. Ye, C. Wan, and Y. C. Lin, "Fusion-3d: Integrated acceleration for instant 3d reconstruction and real-time rendering," in *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2024, pp. 78–91.
- [15] W. Ali, S. Abdelkarim, M. Zidan, M. Zahran, and A. El Sallab, "Yolo3d: End-to-end real-time 3d oriented object bounding box detection from lidar point cloud," in *Proceedings of the European conference on computer vision (ECCV) workshops*, 2018, pp. 0–0.
- [16] B. Yang, W. Luo, and R. Urtaşun, "Pixor: Real-time 3d object detection from point clouds," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 7652–7660.
- [17] T. Zhang, M. Ma, F. Yan, H. Li, and Y. Chen, "Pids: Joint point interaction-dimension search for 3d point cloud," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 1298–1307.
- [18] C. Choy, J. Gwak, and S. Savarese, "4d spatio-temporal convnets: Minkowski convolutional neural networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 3075–3084.
- [19] Q. He, J. Zhang, J. Peng, H. He, X. Li, Y. Wang, and C. Wang, "Pointnrvkv: Efficient rvkv-like model for hierarchical point cloud learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 3, 2025, pp. 3410–3418.
- [20] X. Deng, W. Zhang, Q. Ding, and X. Zhang, "Pointvector: A vector representation in point cloud analysis," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 9455–9465.
- [21] G. Qian, Y. Li, H. Peng, J. Mai, H. Hammoud, M. Elhoseiny, and B. Ghanem, "Pointnext: Revisiting pointnet++ with improved training and scaling strategies," *Advances in neural information processing systems*, vol. 35, pp. 23 192–23 204, 2022.
- [22] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *Advances in neural information processing systems*, vol. 30, 2017.
- [23] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660.
- [24] A. V. Phan, M. Le Nguyen, Y. L. H. Nguyen, and L. T. Bui, "Dgcnn: A convolutional neural network over large-scale labeled graphs," *Neural Networks*, vol. 108, pp. 533–543, 2018.
- [25] H. Zhao, L. Jiang, C.-W. Fu, and J. Jia, "Pointweb: Enhancing local neighborhood features for point cloud processing," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5565–5573.
- [26] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *ACM Transactions on Graphics (tog)*, vol. 38, no. 5, pp. 1–12, 2019.
- [27] Y. Feng, B. Tian, T. Xu, P. Whatmough, and Y. Zhu, "Mesorasi: Architecture support for point cloud analytics via delayed-aggregation," in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2020, pp. 1037–1050.
- [28] Y. Lin, Z. Zhang, H. Tang, H. Wang, and S. Han, "Pointacc: Efficient point cloud accelerator," in *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, 2021, pp. 449–461.
- [29] Y. Feng, G. Hammonds, Y. Gan, and Y. Zhu, "Crescent: taming memory irregularities for accelerating deep point cloud analytics," in *Proceedings of the 49th Annual International Symposium on Computer Architecture*, 2022, pp. 962–977.
- [30] D. A. Talbert and D. Fisher, "An empirical analysis of techniques for constructing and searching k-dimensional trees," in *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2000, pp. 26–33.
- [31] A. W. Moore, "Efficient memory-based learning for robot control," University of Cambridge, Computer Laboratory, Tech. Rep., 1990.
- [32] S. Kim, J. Lee, D. Im, and H.-J. Yoo, "Pnnp: A 11.9 tops/w high-speed 3d point cloud-based neural network processor with block-based point processing for regular dram access," in *2021 Symposium on VLSI Circuits*. IEEE, 2021, pp. 1–2.
- [33] J. Sun, Q. Zhang, B. Kaikhura, Z. Yu, C. Xiao, and Z. M. Mao, "Modelnet40-c: A robustness benchmark for 3d point cloud recognition under corruption," in *ICLR 2022 Workshop on Socially Responsible Machine Learning*, vol. 7, 2022.
- [34] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese, "3d semantic parsing of large-scale indoor spaces," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1534–1543.
- [35] S. Li, C. Li, W. Zhu, B. Yu, Y. Zhao, C. Wan, H. You, H. Shi, and Y. Lin, "Instant-3d: Instant neural radiance field training towards on-device ar/vr 3d reconstruction," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, 2023, pp. 1–13.

- [36] C. Guo, B. Y. Hsueh, J. Leng, Y. Qiu, Y. Guan, Z. Wang, X. Jia, X. Li, M. Guo, and Y. Zhu, "Accelerating sparse dnn models without hardware-support via tile-wise sparsity," in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2020, pp. 1–15.
- [37] C. Guo, F. Xue, J. Leng, Y. Qiu, Y. Guan, W. Cui, Q. Chen, and M. Guo, "Accelerating sparse dnns based on tiled gemm," *IEEE Transactions on Computers*, vol. 73, no. 5, pp. 1275–1289, 2024.
- [38] C. Zhou, Y. Fu, Y. Ma, E. Han, Y. He, and H. Jiao, "Adjustable multi-stream block-wise farthest point sampling acceleration in point cloud analysis," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 71, no. 7, pp. 3523–3527, 2024.
- [39] R. Pinkham, S. Zeng, and Z. Zhang, "Quicknnc: Memory and performance optimization of kd tree based nearest neighbor search for 3d point clouds," in *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2020, pp. 180–192.
- [40] H. Yoon and J.-J. Kim, "Efficient sampling and grouping acceleration for point cloud deep learning via single coordinate comparison," in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. IEEE, 2023, pp. 1–9.
- [41] Y. Gao, C. Jiang, W. Piard, X. Chen, B. Patel, and H. Lam, "Hgpnc: A heterogeneous architecture for e2e embedded point cloud inference," in *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2024, pp. 1588–1600.
- [42] T. Xu, B. Tian, and Y. Zhu, "Tigris: Architecture and algorithms for 3d perception in point clouds," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 629–642.
- [43] C. Zhou, Y. Fu, M. Liu, S. Qiu, G. Li, Y. He, and H. Jiao, "An energy-efficient 3d point cloud neural network accelerator with efficient filter pruning, mlp fusion, and dual-stream sampling," in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. IEEE, 2023, pp. 1–9.
- [44] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "Shapenet: An information-rich 3d model repository," *arXiv preprint arXiv:1512.03012*, 2015.
- [45] S. Imambi, K. B. Prakash, and G. Kanagachidambaresan, "Pytorch," *Programming with TensorFlow: solution for edge computing applications*, pp. 87–104, 2021.
- [46] Qian, Guocheng. (2022) Openpoints. [Online]. Available: <https://github.com/guochengqian/openpoints>
- [47] NVIDIA. (2025) Nvidia management library - nvidia. [Online]. Available: <https://developer.nvidia.com/management-library-nvml>
- [48] Synopsys. (2025) Design compiler - synopsys. [Online]. Available: <https://www.synopsys.com/implementation-and-signoff/rtl-synthesis-test/dc-ultra.html>
- [49] Cadence. (2025) Innovus implementation system. [Online]. Available: [https://www.cadence.com/en\\_US/home/tools/digital-design-and-signoff/soc-implementation-and-floorplanning/innovus-implementation-system.html](https://www.cadence.com/en_US/home/tools/digital-design-and-signoff/soc-implementation-and-floorplanning/innovus-implementation-system.html)
- [50] N. Muralimanohar, R. Balasubramonian, and N. P. Jouppi, "Cacti 6.0: A tool to model large caches," *HP laboratories*, vol. 27, p. 28, 2009.
- [51] C. Guo, J. Tang, W. Hu, J. Leng, C. Zhang, F. Yang, Y. Liu, M. Guo, and Y. Zhu, "Olive: Accelerating large language models via hardware-friendly outlier-victim pair quantization," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, 2023, pp. 1–15.
- [52] S. Li, Z. Yang, D. Reddy, A. Srivastava, and B. Jacob, "Dramsim3: A cycle-accurate, thermal-capable dram simulator," *IEEE Computer Architecture Letters*, vol. 19, no. 2, pp. 106–109, 2020.
- [53] Y. Fu, C. Zhou, T. Huang, E. Han, Y. He, and H. Jiao, "Softact: A high-precision softmax architecture for transformers supporting nonlinear functions," *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.
- [54] F. Chen, R. Ying, J. Xue, F. Wen, and P. Liu, "Parallelnnc: A parallel octree-based nearest neighbor search accelerator for 3d point clouds," in *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2023, pp. 403–414.
- [55] C. Zhou, T. Huang, Y. Ma, Y. Fu, X. Song, S. Qiu, J. Sun, M. Liu, G. Li, Y. He, Y. Yang, and H. Jiao, "23.4 nebula: A 28nm 109.8 tops/w 3d pnn accelerator featuring adaptive partition, multi-skipping, and block-wise aggregation," in *2025 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 68. IEEE, 2025, pp. 412–414.
- [56] Y. Wang and J. M. Solomon, "Object dgcnn: 3d object detection using dynamic graphs," *Advances in Neural Information Processing Systems*, vol. 34, pp. 20745–20758, 2021.
- [57] N. Ukey, Z. Yang, B. Li, G. Zhang, Y. Hu, and W. Zhang, "Survey on exact knn queries over high-dimensional data space," *Sensors*, vol. 23, no. 2, p. 629, 2023.