

ASH: A Modern Framework for Parallel Spatial Hashing in 3D Perception

Wei Dong, Yixing Lao, Michael Kaess, and Vladlen Koltun

Abstract—We present ASH, a modern and high-performance framework for parallel spatial hashing on GPU. Compared to existing GPU hash map implementations, ASH achieves higher performance, supports richer functionality, and requires fewer lines of code (LoC) when used for implementing spatially varying operations from volumetric geometry reconstruction to differentiable appearance reconstruction. Unlike existing GPU hash maps, the ASH framework provides a versatile tensor interface, hiding low-level details from the users. In addition, by decoupling the internal hashing data structures and key-value data in buffers, we offer direct access to spatially varying data via indices, enabling seamless integration to modern libraries such as PyTorch. To achieve this, we 1) detach stored key-value data from the low-level hash map implementation; 2) bridge the pointer-first low level data structures to index-first high-level tensor interfaces via an index heap; 3) adapt both generic and non-generic integer-only hash map implementations as backends to operate on multi-dimensional keys. We first profile our hash map against state-of-the-art hash maps on synthetic data to show the performance gain from this architecture. We then show that ASH can consistently achieve higher performance on various large-scale 3D perception tasks with fewer LoC by showcasing several applications, including 1) point cloud voxelization, 2) retargetable volumetric scene reconstruction, 3) non-rigid point cloud registration and volumetric deformation, and 4) spatially varying geometry and appearance refinement. ASH and its example applications are open sourced in Open3D (<http://www.open3d.org>).

Index Terms—Parallel hashing, GPU, Volumetric reconstruction, SLAM, Shape-from-shading, Autodiff



1 INTRODUCTION

3D space is only one dimension higher than the 2D image space, yet the additional dimension introduces an unpredictable multiplier to the computational and storage cost. This increased dimension makes 3D perception tasks such as geometric reconstruction and appearance refinement challenging to implement. To reduce complexity, one compromise is to reuse dense data structures and constrain the 3D space by bounding the region of interest, *e.g.*, adopting a 3D array in a bounded space [40]. While this simple approach succeeds at the scale of objects, it cannot meet the demand of room or city scale perception, which is necessary for virtual tour, telepresence, and autonomous driving.

Since the 3D space is generally a collection of 2D surface manifolds, its sparsity can be exploited by *partitioning* to reduce computational cost. The general idea is to split the large 3D space into smaller regions and only proceed with the non-empty ones. There is a plethora of well-established data structures for 3D space partitioning. Examples include trees (Octree [36], KD-tree [4]) and hash maps (spatial hashing [41]). While trees are able to adaptively achieve high precision, they 1) require an initial bounding volume and 2) usually take unbalanced traversal time for a batch of spatial queries and hence 3) are less friendly to batched operations. On the other hand, spatial hashing coupled with a plain array structure is more scalable and parallelizable for reconstruction tasks.

In classical dense SLAM pipelines [41], [47], spatial hashing is used to map 3D coordinates to internal pointers that are only accessible in GPU device code. A pre-allocated memory pool allows dynamic growing hash maps, but adaptive rehashing is generally unavailable. Recent feature-grid encoders [37] apply spatial hashing to map 3D coordinates to features stored at grid points in a *static* bounded

region. While differentiable spatial query is supported, *collisions* are not resolved, limiting its usage to stochastic feature optimization where incorrect key-value mappings are tolerated. A general, user-friendly, collision-free hash map is missing for efficient spatial perception at scale.

The reason for this absence is understandable. A parallel hash map on GPU has to resolve collisions and thread conflicts and preferably organize an optimized memory manager, none of which is trivial to implement. Previous studies have attempted to tackle the problem in one or more aspects, driven by their selected downstream applications. Furthermore, most of the popular parallel GPU hash maps are implemented in C++/CUDA and only expose low-level interfaces. As a result, customized extensions must start from low-level programming. While these designs usually guarantee performance under certain circumstances [2], [3], [41], [47], as of today, they leave a gap from the standpoint of the research community, which prefers to use off-the-shelf libraries for fast prototyping with a high-level scripting language using tensors and automatic differentiation. Our motivation is to bridge this gap to enable researchers to develop sophisticated 3D perception routines with less effort and drive the community towards large-scale 3D perception.

To this end, we design a modern hash map framework with the following major contributions:

- 1) a user-friendly *dynamic*, *generic*¹, and *collision-free* hash map interface that enables tensor I/O, *advanced indexing*,

¹ A dynamic hash map supports insertion and deletion after hash map construction. A generic hash map supports arbitrary dimensional keys and values in various data types.

and *in-place automatic differentiation* when bridged to autodiff engines such as PyTorch;

- 2) an index-first adaptor that supports various state-of-the-art parallel GPU hash map backends and accelerates hash map operations with an improved structure-of-array (SoA) data layout;
- 3) a number of downstream applications that achieve higher performance compared to state-of-the-art implementations with fewer LoC.

Experiments show that ASH achieves better performance with fewer LoC on both synthetic and real-world tasks.

2 RELATED WORK

2.1 Parallel Hash Map

The hash map is a data structure that seeks to map *sparse keys* (e.g. unbounded indices, strings, coordinates) from the set \mathcal{K} to values from the set \mathcal{V} with amortized $O(1)$ access. It has a hash function $h : \mathcal{K} \rightarrow \mathcal{I}_n, k \mapsto h(k)$ that maps the key to the index set $\mathcal{I}_n = \{0, 1, \dots, n - 1\}$ for indexing (or addressing) that is viable on a computer.

Ideally, with a perfect injective hash function h , a hash map can be implemented by $H : \mathcal{K} \rightarrow \mathcal{V}, k \mapsto \mathbf{v}^A[h(k)]$, where \mathbf{v}^A is an array of objects of type \mathcal{V} and $[\cdot]$ is the trivial array element accessor. However, in practice, it is intractable to find an injective map given a sparse key distribution in \mathcal{K} and a constrained index set \mathcal{I}_n of size n due to the computational budget. Therefore, modifications are required to resolve inevitable *collisions*, where $i = h(k_1) = h(k_2), k_1 \neq k_2$. There are two classes of techniques for collision resolution, *open addressing* and *separate chaining*. Open addressing searches for another candidate $j \neq i, j \in \mathcal{I}_n$ via a *probing* algorithm until an empty address is found. The simplest probing, linear probing [28], computes $j = (h(k) + t) \bmod n$ starting from $i = h(k)$, where t is the number of attempts. Separate chaining, on the other hand, maintains multiple entries per mapped index where a linked list is grown at i if $i = h(k_1) = h(k_2), k_1 \neq k_2$.

While hash map implementations are widely available for CPU, their GPU counterparts have only emerged in the recent decade. Most GPU hash maps use open addressing [2], [17], [27], [53], mainly due to simplicity in implementation and capability of handling highly concurrent operations. CUDPP [2] utilizes Cuckoo Hashing [44], while CoherentHash [17] adopts Robin Hood Hashing [6] – both involving advanced probing design. Although being performant when \mathcal{K}, \mathcal{V} are limited to integer sets, these variations cannot be generalized to spatial hashing and only allow static input. Recently, WarpCore [27] proposes to support non-integer \mathcal{V} and dynamic insertion, but the key domain is still limited to at most 64 bits.

There are also a few separate chaining implementations on GPU involving device-side linked lists. SlabHash [3] builds a linked list with a 128-bit *Slab* as the minimal unit, optimized for Single Instruction Multiple Threads (SIMT) warp operations. Although SlabHash allows dynamic insertions, similar to the aforementioned GPU hash maps, only integer \mathcal{K}, \mathcal{V} are supported. stdgpu [52] follows the conventional C++ Standard Library `std::unordered_map` and builds supporting vectors, bitset lock guards, and linked lists from scratch, resulting in a generic, dynamic hash

TABLE 1

Comparison of existing parallel GPU hash maps. ASH preserves the dynamic, generic, and atomic properties, and is extendable to the non-templated high-level Python interfaces.

	Dynamic	Generic	Collision-free	Python
SlabHash [3]	✓	✗	✓	✓
CUDPP [2]	✗	✗	✓	✗
cuDF [53]	✗	✗	✓	✓
WarpCore [27]	✓	✗	✓	✗
stdgpu [52]	✓	✓	✓	✗
InfiniTAM [47]	✓	✗	✓	✗
VoxelHashing [41]	✓	✗	✗	✗
GPURobust [12]	✓	✓	✗	✗
Instant-NGP [37]	✗	✓	✗	✓
ASH	✓	✓	✓	✓

map. With these rich functionalities, however, stdgpu is not optimized for large value sets. In addition, due to its low-level templated design, users have to write device code for simple tasks.

We refer the readers to a comprehensive review of GPU hash maps [33].

2.2 Space Partitioning Structures

3D data is not as simple to organize as 2D images. While a 2D image can be stored in a dense matrix, exploiting sparsity in 3D data is paramount due to the limits in computer memory of the current day.

The most widely used data structures for 3D indexing are arguably trees. A KD-tree [4] recursively sorts k-dimensional data along a selected axis and partitions data at the median point. By nature, a KD-Tree is designed for neighbor search. In 3D, it is mainly used to organize 3D *points* and their *features*. Examples include normal estimation and nearest neighbor association in Iterative Closest Points (ICP) [49], [62] and 3D feature association in global registration [50], [59]. GPU adaptations exist for KD-trees [51], [58], but are not suitable for incrementally changing scenes, as they are usually constructed once and queried repeatedly.

Bounding volume hierarchy (BVH) is another hierarchical representation that organizes primitives such as *objects* and *triangles* in 3D. There are various GPU adaptations [19], [32], [55] mostly targeted at ray tracing and dynamic collision detection. While a parallel construction is possible and deformation of the nodes is allowed, the tree structure typically remains unchanged, assuming a fixed layout.

While KD-trees and BVH split the space *unevenly* by data distribution, an Octree [36], on the other hand, recursively partitions the 3D space *evenly* into 8 subvolumes according to space occupation states. It has been widely used in adaptive 3D mapping [23], [38] for robot navigation. There have been parallel implementations on GPU, from optimized data structures [22] to domain-specific languages [24]. However, these works generally focus on physics simulation within a bounded region of interest where the spatial partition is predefined. While parallel incremental division [57] is possible, an initial bounding region is still required, and the trees are not guaranteed to be balanced.

Spatial hashing is another variation of spatial management with $O(1)$ access time depending on hash maps.

Bundled with small dense 3D arrays, it has been widely used in real-time volumetric scene reconstruction within *unbounded* region of interest. A handful of CPU implementations have achieved real-time performance [20], [30] at the expense of resolution. Similarly, GPU implementations [12], [15], [41], [47] reach high frame rates using GPU-based spatial hashing. However, all of the studies depend on ad hoc GPU hash maps exclusive to these specific systems. Concurrent race conditions have not been fully resolved in several implementations [12], [41], where volumes can be randomly under-allocated. Recent neural feature grids [37] apply spatial hashing in a *bounded* volume to query voxelized feature embeddings. These approaches use simplified hashing designs that are not collision-free, and thus are only compatible with stochastic optimization that tolerates noise from an incorrect query.

2.3 Spatially Varying 3D Representations

A truncated signed distance function (TSDF) [11] is an implicit representation of surfaces, recording point-wise distance to the nearest surface point. It is frequently used for dense scene reconstruction with noisy input. The distribution of surfaces is generally spatially varying and therefore, a proper parameterization is often necessary, either in a discrete [41] or neural [7], [37] form. Non-rigid deformation methods [39], [60] seek to embed point clouds in a deformable grid, where each point is anchored to and deformed by neighbor grids. They are mainly used for animation or non-rigid distortion calibration. Similar to a deformation grid, complex lighting for rendering can be approximated by spatially-varying spherical harmonics (SVSH) [35] placed at a sparse grid. These grids are natural applications of spatial hashing. A comprehensive review of spatially varying representations for real-world scene reconstruction is available [64].

While ad hoc implementations have been introduced for these representations either on CPU or GPU, ASH provides a device-agnostic interface requiring less code written and providing better performance. Table 1 compares various aspects of existing GPU hash maps, either as a standalone data structure (Section 2.1) or embedded in an application (Section 2.2). To the best of our knowledge, ASH is the first implementation that simultaneously supports dynamic insertion, ensures correctness via atomic collision-free operations, allows generic keys, and has a modern tensor interface and Python binding for better usability.

3 OVERVIEW

Before plunging into the details, we first provide a high-level overview of our framework in Fig. 1.

Conventional parallel hash maps reorganize the structure of arrays (SoA) input, *i.e.*, the separated key array and value array, into an array of structures (AoS) where keys and values are *paired*, inserted, and stored. Therefore, array of *pointers to pair structures* (`std::pair` in C++, `thrust::pair` in CUDA, and `tuple` in Python) are returned upon query. Consequently, the operations from insertion and query to in-place value increment require users to write device code and visit AoS at the low-level pointers.

In contrast, ASH sticks to SoA. Fig. 1 shows the workflow of ASH. Instead of *pointers to pairs*, ASH returns *indices* and *masks* arrays that can be directly consumed by tensor libraries such as PyTorch [46] (without memory copy) and NumPy [21] (with GPU to host memory copy). As a result, post-processing functions such as duplicate key removal and in-place modification can be chained with insertion and query in ASH via advanced indexing without writing any device code. As a general and device-agnostic interface for parallel hash maps, our framework is built upon switchable backends with details hidden from the user. Currently, *separate chain* backends are supported, including the generic `stdgpu` [52] backend, and the extended integer-only Slab-Hash [3] backend for arbitrary key-value data types. TBB’s concurrent hash map [31] powers the CPU counterpart with the identical interface to GPU.

In this paper, we use calligraphic letters to represent sets, and normal lower-case letters for their elements. Normal upper-case letters denote functions. Bold lower-case letters denote vectors of elements or arrays in the programmer’s perspective. Bold upper-case letters are for matrices. For instance, in a hash map, we are interested in key elements $k \in \mathcal{K}$ and their vectorized processing, *e.g.* query $Q(\mathbf{k})$. Specifically, we use \mathcal{I}_n to denote a set of indices $\{0, 1, \dots, n-1\}$, and Θ as the boolean selection $\{0, 1\}$. Given an arbitrary vector \mathbf{x} , we denote $\mathbf{x}(\mathbf{i})$ and $\mathbf{x}(\boldsymbol{\theta})$ as indexing and selection functions applied to \mathbf{x} when $\forall i \in \mathbf{i}, i \in \mathcal{I}_n$ and $\forall \theta \in \boldsymbol{\theta}, \theta \in \Theta$. We use $\langle \mathcal{K}, \mathcal{V} \rangle$ as the key and value sets for a hash map. $h : \mathcal{K} \rightarrow \mathcal{I}_n$ is the internal hash function that converts a key to an index. $H : \mathcal{K} \rightarrow \mathcal{V}$ is the general hash map enclosing h .

4 THE ASH FRAMEWORK

4.1 Classical Hashing

In a hash map $\langle \mathcal{K}, \mathcal{V} \rangle$, since the hash function h cannot be perfect as discussed in Section 2, we have to store keys to verify if collisions happen ($h(k_1) = h(k_2)$ but $k_1 \neq k_2, k_1, k_2 \in \mathcal{K}$).

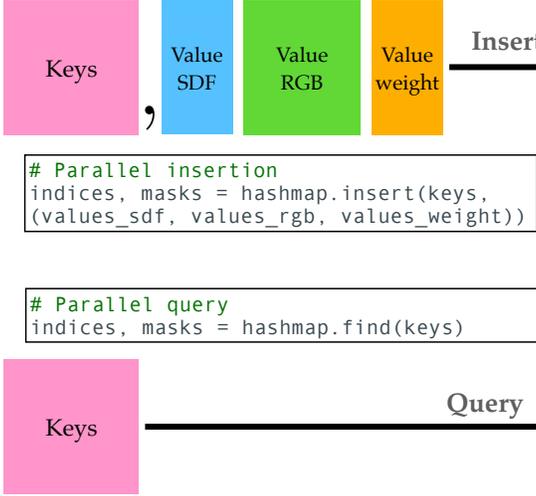
In separate chaining, to resolve hash collisions, the bucket-linked list architecture is used. With n initial *buckets*, we construct the hash function $h : \mathcal{K} \rightarrow \mathcal{I}_n$ where \mathcal{I}_n is defined in Section 2. As shown in Fig. 2, keys with the same hashed index $i = h(k) \in \mathcal{I}_n$ are first aggregated in the i -th bucket, where a linked list grows adaptively to accommodate different keys. A conventional hash map stores key-value pairs as the storage units. Consequently, two keys k_1, k_2 can be distinguished by checking $h(k_1) = h(k_2)$ and $k_1 = k_2$ from the pair in order, and manipulation of the keys and values can be achieved by iterating over such pairs.

With this formulation, assuming a subset $\mathcal{X} \subset \mathcal{K}$ has been inserted into the hash map with associated values $\mathcal{Y} \subset \mathcal{V}$, a query function can be described as

$$Q_{\mathcal{K}, \mathcal{V}} : \mathcal{K} \rightarrow \mathcal{K} \times \mathcal{V} \\ k \mapsto \langle k, v \rangle, \forall k \in \mathcal{X}, \quad (1)$$

where $\langle k, v \rangle$ forms a concrete pair stored in the hash map. This format is common in implementations, *e.g.* in C++ (`std::unordered_map`) and Python (`dict`).

SoA tensor input



SoA tensor output

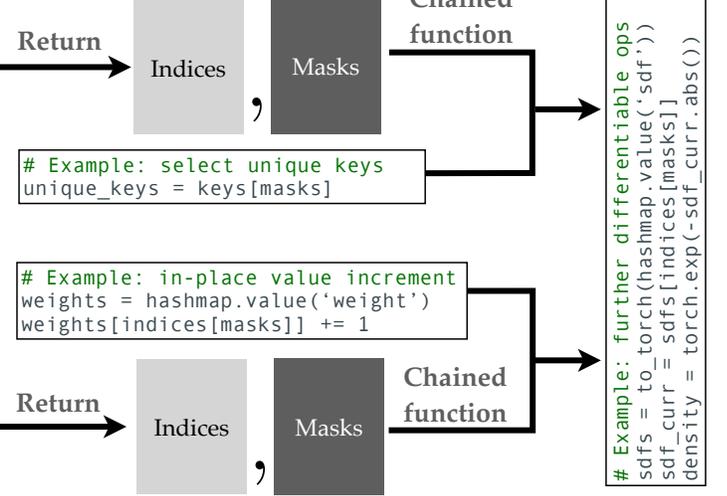


Fig. 1. The interface illustration of ASH. Left: ASH takes a key and/or a list of value tensors as input, with a one-liner interface. For *insertion*, tensors are organized in SoA. Right: buffer indices and masks are returned upon insertion and query, organized together in SoA. Chained functions can be easily applied to hashed data by indexing ASH buffers with indices and masks. Examples include selecting unique keys through insertion, and applying in-place value increment through query. Differentiable ops can be applied in downstream applications.

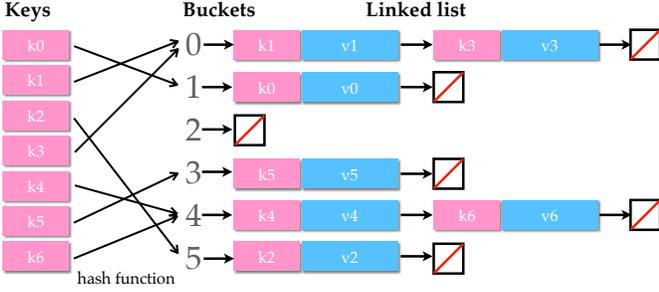


Fig. 2. Illustration of a classical hash map using separate chaining. Keys (left) are put into corresponding buckets (middle) obtained by the hash function h . A linked list (right) is constructed per bucket to store key-value pairs within the same bucket but with unequal keys.

4.2 Function Chaining and Parallel Hashing

The element-wise operation in Eq. 1 can be extended to vectors via parallel device kernels. However, interpretation of the returned iterators of pairs is still required at the low level. In other words, although the parallel version can be implemented efficiently, results are still packed in an AoS instead of SoA:

$$Q_{\mathcal{K},\mathcal{V}}(\mathbf{k}) = \mathbf{array}\{\langle k, v \rangle\}. \quad (2)$$

This forms a barrier when the parallel query is located in a chain of functions. For instance, to apply any function G (e.g., geometry transformation) over the result of a query, the low-level function `second` that selects the value element from a pair $\langle k, v \rangle$ must be provided to dereference the low-level structures and manipulate the keys and values *in-place*. In other words, we have to implement a non-trivial \tilde{G} :

$$\tilde{G}(\mathbf{k}) = (G \circ \text{second} \circ Q_{\mathcal{K},\mathcal{V}})(\mathbf{k}), \quad (3)$$

to force the conversion from AoS to SoA and chain a high-level function G with $Q_{\mathcal{K},\mathcal{V}}$ ². This could be tedious when

2. This can be achieved simply by returning a *copy* of values, but it is not feasible, especially when dealing with large-scale data, e.g. hierarchical voxel grids.

prototyping geometry perception that requires hash map structures since off-the-shelf operations have to be reimplemented in device code.

We reformulate this problem by introducing two affiliate arrays, \mathbf{k}^B and \mathbf{v}^B (note with a superscript B for buffering, they are not the input \mathbf{k}, \mathbf{v}) of capacity $c \geq n$, where n is the number of buckets. These arrays are designed for explicit storage of keys and values, respectively, and serve as buffers to support natural SoA. They are exposed to users for direct access and in-place modification. Now the query function can be rewritten as

$$Q_{\mathcal{K},\mathcal{V}} : \mathcal{K} \rightarrow \mathcal{I}_c, k \mapsto i, \\ \text{s.t. } \mathbf{k}^B(i) = k, \mathbf{v}^B(i) = v, \forall \langle k, v \rangle \in \langle \mathcal{X}, \mathcal{Y} \rangle \quad (4)$$

and this version is ready for parallelization. At this stage, to combine G and $Q_{\mathcal{K},\mathcal{V}}$, we can chain $G \circ \mathbf{v}^B \circ Q$ to manipulate values:

$$G(\mathbf{k}) = G\left(\mathbf{v}^B(Q_{\mathcal{K},\mathcal{V}}(\mathbf{k}))\right), \quad (5)$$

which retains convenient properties such as array vectorization and advanced indexing.

When the input set $\tilde{\mathcal{X}} \not\subseteq \mathcal{X}$ is not fully stored in the hash map, our formulation maintains its effectiveness by a simple masked extension:

$$Q_{\mathcal{K},\mathcal{V}}^{\mathcal{I}} : \mathcal{K} \rightarrow \mathcal{I}_c, Q_{\mathcal{K},\mathcal{V}}^{\Theta} \rightarrow \{0, 1\}, \\ Q_{\mathcal{K},\mathcal{V}}^{\mathcal{I}}(\tilde{k}) = i, Q_{\mathcal{K},\mathcal{V}}^{\Theta} = \theta, \\ \text{s.t. } \theta = 1; \mathbf{k}^B(i) = \tilde{k}, \mathbf{v}^B(i) = \tilde{v}, \text{ if } \tilde{k} \in \mathcal{X}, \quad (6) \\ \theta = 0; i = \text{undefined, otherwise,}$$

which is also ready for parallelization. Now the chaining of functions is given by

$$G(\mathbf{k}) = G\left(\mathbf{v}^B(\mathbf{i}(\theta))\right), \quad (7)$$

$$\mathbf{i} = Q_{\mathcal{K},\mathcal{V}}^{\mathcal{I}}(\mathbf{k}), \theta = Q_{\mathcal{K},\mathcal{V}}^{\Theta}(\mathbf{k}), \quad (8)$$

using advanced indexing with masks. We can also select valid queries with $\mathbf{k}(\theta)$ without visiting \mathbf{k}^B . While our discussion was about the query function, the same applies to insertion.

In essence, by converting the pair-first AoS to an index-first SoA format with the help of array buffers, we can conveniently chain high-level functions over hash map query and insertion. This simple change enables easy development on hash maps and unleashes their potential for fast prototyping and differentiable computation. However, the layout requires fundamental changes to the hash map data structure. With this in mind, we move on to illustrate how the ASH layer converts the AoS in native backends to our SoA layout.

4.3 Generic Backends

We start with converting `stdgpu` [52], a state-of-the-art generic GPU hash map as the backend of ASH. `stdgpu` follows the convention of its CPU counterpart `std::unordered_map` by providing a templated interface. The underlying implementation is a classical bucket-linked list structure with locks to avoid race conditions on GPU. To exploit the power of a generic hash map without reinventing the wheel, we seek to reuse the operations over keys (*i.e.* lock-guarded bucket and linked list operations) and redirect the value mapping to our buffer \mathbf{v}^B .

A dynamic GPU hash map requires dynamic allocation and freeing of keys and values in device kernels. With pre-allocated key buffer \mathbf{k}^B and value buffer \mathbf{v}^B , we maintain an additional *index heap* \mathbf{h} , as shown in Fig. 3. The index heap stores *buffer indices* i pointing to the buffers $\mathbf{k}^B, \mathbf{v}^B$ as a map $P: \mathcal{I}_c \rightarrow \mathcal{I}_c$, where the *heap top* t maintains the currently available buffer index in $\mathbf{h}[t]$. Heap top starts at $t = 0$, and is atomically increased at allocation and decreased at free. With \mathbf{h} and the dynamically changing t , we instantiate a generic hash map with the templated value in `stdgpu` to be $\mathcal{V} = \text{Int32}$, where the values are *buffer indices* i stored in \mathbf{h} to access $\mathbf{k}^B, \mathbf{v}^B$ exposed to the user.

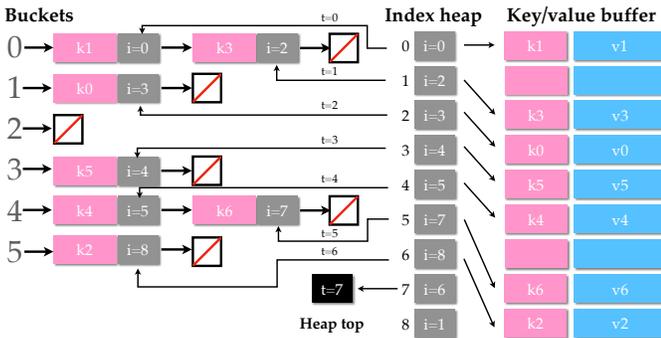


Fig. 3. Illustration of a generic hash map bridged to tensors in ASH. *Buffer indices* i are dynamically provided by the available indices maintained in the index heap at the increasing heap top t (middle), acting as the *values* in the underlying hash map (left). It connects the hash map and the actual key values stored in the buffer (right) by accessing i . The key-bucket correspondences are the same as Fig. 2, omitted for simplicity.

4.3.1 Insertion

The insertion of a $\langle k, v \rangle \in \langle \mathcal{K}, \mathcal{V} \rangle$ pair is now decoupled into two steps, with i) insertion of $\langle k, i \rangle \in \langle \mathcal{K}, \mathcal{I}_c \rangle$ into the hash

map, where i is the buffer index dynamically acquired from the heap top $\mathbf{h}[t]$ and ii) insertion of $\mathbf{k}^B(i) := k, \mathbf{v}^B(i) := v$ into buffers.

A naive implementation will acquire a buffer index i from \mathbf{h} on every insertion attempt and free it if the insertion fails because the key already exists. However, when running in parallel, `atomicAdd` and `atomicSub` may be conflicting among threads, leading to race conditions. A *two-pass insertion* could resolve the issue: in the first pass, we allocate a batch of indices from \mathbf{h} determined by the input size, attempt insertions, and record results; in the second pass, we free the indices to \mathbf{h} from failed insertions.

We adopt a more efficient *one-pass lazy* insertion. We first attempt to insert $\langle k, -1 \rangle$ with -1 as the dummy index into the backend and observe if it is successful. If not, nothing needs to be done. Otherwise, we capture the returned pointer to the pair, trigger an index i allocation from \mathbf{h} , and directly replace the dummy -1 with i . This significantly reduces the overhead when the key uniqueness is low (*i.e.*, many duplicates exist in the keys to be inserted).

4.3.2 Query

The query operation is relatively simpler. We first look up the buffer index $i \in \mathcal{I}_c$ given k in the backend. If it is a success, we end up with $k = \mathbf{k}^B(i)$, and the target $v = \mathbf{v}^B(i)$ is accessible with i by users.

4.4 Non-generic Backends

While the generic GPU hash map has only recently been available, the research community in parallel computation has been focusing on more controlled setups where both \mathcal{K} and \mathcal{V} are limited to certain dimensions or data types. We seek to generalize this non-generic setup with our index heap and verify their performance in more real-world applications. In this section, we show how ASH can be used to generalize SlabHash [3], a warp-oriented dynamic GPU hash map that only allows insertions and queries to `Int32` data type.

An extension to generic key types is non-trivial for SlabHash since its warp operations only apply to variables with limited word length. Our implementation extends the hash set variation of SlabHash, where only integers as keys are maintained in the backend.

4.4.1 Generalization via Index Heap

The index heap \mathbf{h} is the core to generalizing the SlabHash backend. In brief, a generic key is represented by its associated buffer index i in an integer-only hash set, allocated the same way as discussed in Section 4.3.1. As illustrated in Fig. 4, all the insertions and queries are redirected from the buffer indices to actual keys and values via the index heap. However, the actual implementation involves more complicated changes in design.

Given a generic key k , we first locate the bucket $b = h(k) \in \mathcal{I}_n$. Ideally, we can then allocate a buffer index i at \mathbf{h} 's top t and insert it into the linked list at the bucket b in the integer-only hash set. The accompanying key and value are put in $\mathbf{k}^B(i), \mathbf{v}^B(i)$. During query, we similarly first locate the bucket b then search the key in the linked list by visiting \mathbf{k}^B via the stored index i .

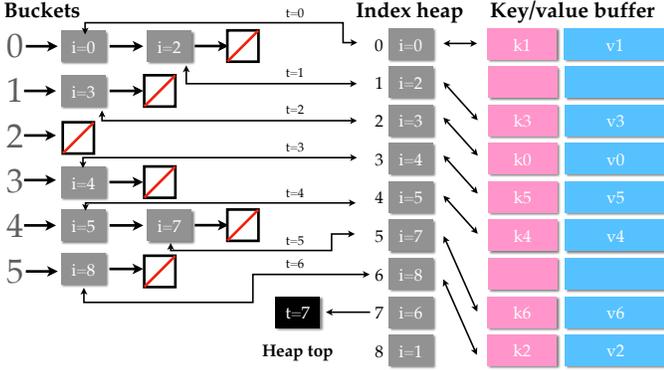


Fig. 4. Illustration of a non-generic hash set enhanced by ASH. Integer *buffer indices* i allocated from the index heap (middle) are inserted as delegate *keys* directly into the hash set (left), associated with *actual keys* in the buffer (right) at i . The key-bucket correspondences are the same as Figs. 2 and 3, omitted for simplicity.

4.4.2 Multi-Pass Insertion

Although query can be applied as mentioned above, lazy insertion mentioned in Section 4.3.1 is problematic in this setup. The main reason is that while the race condition in inserting index i does not occur in warp-oriented insertions, the copy of the *actual* key $k \in \mathcal{K}$ to \mathbf{k}^B requires global memory write. They may not be synchronized among threads, as copying a multi-dimensional key takes several non-atomic instructions. As a result, the insertion of a key k_2 could be accidentally triggered when i) a duplicate $k_1 (= k_2)$'s index $i_1 \in \mathcal{I}_c$ has been inserted but ii) whose actual key k_1 has only been partially copied to the buffer \mathbf{k}^B . This would mistakenly result in $\mathbf{k}^B(i_1) = k_1 \neq k_2$ followed by the unexpected insertion of k_2 when unsynchronized. In practice, with more than 1 million keys to be inserted in parallel, these kinds of conflicts happen with probability as low as $\leq 0.1\%$. To resolve conflicts, we split insertion into three passes:

- *Pass 1*: batch insert *all* keys \mathbf{k} to \mathbf{k}^B by directly copying all candidates via batch allocated corresponding indices i from \mathbf{h} ;
- *Pass 2*: perform parallel hashing with indices i from pass 1. In this pass, keys are read-only in global buffers and hence do not face race conditions. Successful insertions are marked in a mask array.
- *Pass 3*: batch insert values to \mathbf{v}^B with successful masks, and free the rest to \mathbf{h} .

While there is overhead due to the multi-pass operation, it is still practical for a dynamic hash map. First, keys are relatively inexpensive to copy, especially for spatial coordinates, while the more expensive copying of values is done without redundancy. Second, a dynamic hash map generally reserves sufficient memory for further growth so that the *all* key insertion would not exceed the buffer capacity.

4.5 Reshaping and Memory Management

While buffers are represented as fixed-size arrays, growth of storage is needed to accommodate the accumulated input data, which can exceed the hash map's capacity, *e.g.* 3D points from an RGB-D stream. This triggers reshaping, where we adopt the conventional $\times 2$ strategy to double the

buffer size as common in the C++ Standard Library, collect all the active keys and values, and batch insert them into the enlarged buffer.

In dynamic insertions, there can be frequent free and allocation of small memory blobs that are adjacent and mergeable. In view of this, we implement another tree-structured global GPU manager similar to PyTorch [46].

4.6 Dispatch Routines

To enable bindings to non-templated languages, *e.g.* Python, the tensor interface is non-templated so that it can take data types and shapes as arguments. In the context of spatial hashing, we support arbitrary dimensional keys by expanding the dispatcher macros in C++. Float types have undetermined precision behaviors on GPU. Therefore, a conversion to the integers given the desired precision is recommended to use the hash map.

We also additionally dispatch values by their element byte sizes into intrinsically supported vectors: `int`, `int2`, `int3`, and `int4`. This adaptation accelerates trivially copiable value objects such as `int3`, and supports non-trivially copiable value blocks (*e.g.* an 8^3 array pointed to a void pointer). This improves the insertion of large value chunks by a factor of 10 approximately.

4.7 Multi-value Hash Map and Hash Set

ASH supports multi-value hash maps that store values organized in SoA, as well as hash sets with only keys and no values.

Multi-value hash maps. Various applications in 3D processing require mapping coordinates to several properties. For instance, a 3D coordinate can be mapped to a normal, a color, and a label in a point cloud. While the mapped values can be packed as an array of structures (*i.e.*, AoS) to fit a hash map, code complexity could increase since structure-level functions have to be implemented. We generalize the hash map's functionality by extending the *single* value buffer \mathbf{v}^B to an *array of value buffers* $\{\mathbf{v}_i^B\}$ and applying loops over properties per index during an insertion. This simple change supports the storage of complex value types in SoA that allows easy vectorized query and indexing.

Hash set. A hash set, on the other hand, is a simplified hash map – an unordered set that stores unique keys. It is generally useful in maintaining a set by rejecting duplicates, such as in point cloud voxelization. By removing \mathbf{v}^B and ignoring value insertion, a hash map becomes a hash set.

5 EXPERIMENTS

We start with synthetic experiments to show that ASH, with its optimized memory layout, increases performance while improving usability. All experiments in this section are conducted on a laptop with an Intel i7-6700HQ CPU and an Nvidia GTX 1070 GPU. In all experiments, we assume the hash map capacity is equivalent to the number of input keys (regardless of duplicates). Each reported time is an average of 10 trials.

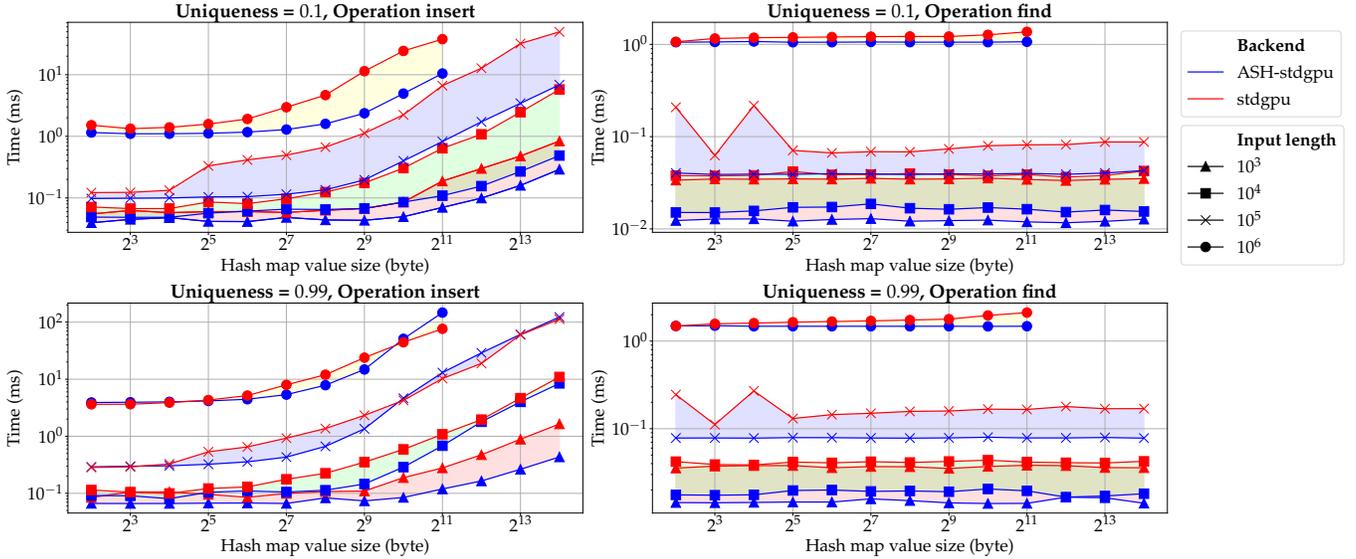


Fig. 5. Hash map performance comparison between ASH-stdgpu and the vanilla stdgpu with 3D integer keys. Each curve shows the average operation time (y -axis) with varying hash map value sizes in bytes (x -axis), given a controlled backend, input length, and input key uniqueness ratio. Lower is better. Further factors are denoted by the legends on the right. ASH-stdgpu runs consistently faster than the vanilla stdgpu.

5.1 Spatial Hashing with Generic Backend

The first experiment is the performance comparison between vanilla stdgpu and ASH with stdgpu backend (ASH-stdgpu). For fairness, we extend the examples of stdgpu such that an array of iterators and masks are returned for in-place manipulations. The number of buckets and the load factor are determined internally by stdgpu.

Setup 1. We test randomly generated 3D spatial coordinates mapped to float value blocks of varying sizes. The key \mathcal{K} , value \mathcal{V} , capacity c , and uniqueness ρ are chosen as follows:

$$\begin{aligned} \mathcal{K} &= \{\text{Tensor}((3), \text{Int32})\}, \\ \mathcal{V} &= \{\text{Tensor}((2^j), \text{Float32}) \mid j = 0, 1, \dots, 12\}, \\ c &= \{10^j \mid j = 3, 4, 5, 6\}, \\ \rho &= \{0.1, 0.99\}, \end{aligned}$$

where ρ indicates the ratio of the unique number of keys to the total number of keys being inserted or queried.

Fig. 5 illustrates the comparison between vanilla stdgpu and ASH-stdgpu. For *insert* operation, ASH-stdgpu is significantly faster than stdgpu when $\rho = 0.1$ is low, and the performance gain increases when the value byte size increases. This is mainly due to the SoA memory layout and the lazy insertion mechanism, where a lightweight integer $i \in \mathcal{I}_c$ is inserted in an attempt instead of the actual value $v \in \mathcal{V}$. At a high input uniqueness $\rho = 0.99$, ASH-stdgpu maintains the performance advantage with low and medium value sizes, and its performance is comparable to stdgpu with a large value size. This indicates that our dispatch pattern in copying values helps in a high throughput scenario. For *find* operation, ASH-stdgpu is consistently faster than vanilla stdgpu, under both high and low key uniqueness settings.

In addition to *insert* and *find*, we introduce a new *activate* operation. It “activates” the input keys by inserting them into the hash map and obtaining the associated buffer indices. This is especially useful when we can pre-determine and apply the element-wise initialization. Examples include

the TSDF voxel blocks (zeros) and multi-layer perceptrons (random initializations). The *activate* operation is absent in most existing hash maps and is only available as hard-coded functions [12], [41], [47].

With the *activate* operation, we conduct ablation studies to compare the insertion time of merely the keys versus the insert time of both the keys and values. Fig. 6 compares the runtime between *insert* and *activate* in ASH-stdgpu. The key, value, capacity, and uniqueness choices are the same as in Setup 1. We observe that while the insertion time increases as the value size increases, the activation time remains stable.

5.2 Integer Hashing with Non-Generic Backend

Next, we compare ASH based on the SlabHash backend (ASH-slab) with the vanilla SlabHash. Since SlabHash only supports integers as keys and values, we limit our ASH-slab backend to the same integer types here. The number of buckets is $2 \times$ capacity (load factor is approx. 0.5), since it is empirically the best factor when ASH-slab is applied to non-generic and generic tasks. Since vanilla SlabHash only supports data I/O from the host, we include the data transfer time between host and device when measuring the performance of ASH-slab.

Setup 2. We test random scalar integer values mapped to scalar float values. The key \mathcal{K} , value \mathcal{V} , capacity c , and uniqueness ρ are chosen as follows:

$$\begin{aligned} \mathcal{K} &= \{\text{Tensor}((1), \text{Int32})\}, \\ \mathcal{V} &= \{\text{Tensor}((1), \text{Float32})\}, \\ c &= \{10^j \mid j = 3, 4, 5, 6\}, \\ \rho &= \{0.1, 0.2, \dots, 0.9, 0.99\}, \end{aligned}$$

As shown in Fig. 7, although ASH-slab does not make use of the non-blocking warp-oriented operations in SlabHash in order to enable support for generic key and value types, our *insert* is still comparable to vanilla SlabHash

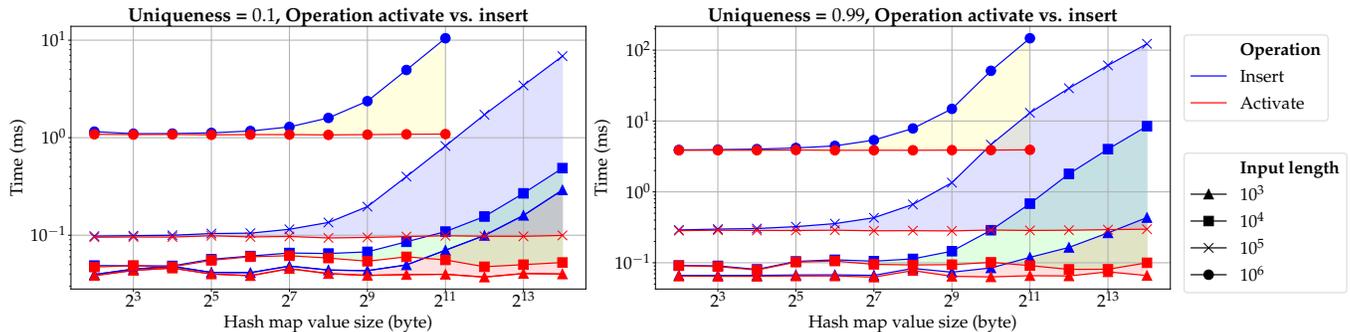


Fig. 6. Study of the *activate* operation introduced in ASH against *insert* with 3D integer keys on the ASH-stdgpu backend. Each curve shows the average operation time (*y*-axis) with varying hash map value sizes in bytes (*x*-axis), given an input length and input key uniqueness ratio. *Lower is better*. *Activate* keeps a stable runtime in the tasks that do not require explicit value insertion, while *insert* time increases corresponding to the hash map value size.

which is only optimized for integers. The drop in performance of ASH-slab when ρ increases is an expected indication that the overhead of multi-pass insertion increases correspondingly.

It is worth mentioning that with an improved global memory manager, the construction of an ASH-slab hash map takes less than 1ms under all circumstances, while the vanilla SlabHash constantly takes 30ms for the redundant slab memory manager. In practice, where the hash map is constructed and used once (*e.g.* voxelization), ASH-slab is a more practical solution.

5.3 Ablation Between Backends

We now conduct an ablation study with different backends in ASH, namely ASH-stdgpu and ASH-slab, with arbitrary input key-value types beyond integers. The experimental setup follows Section 5.1.

In Fig. 8, we can see that ASH-stdgpu outperforms ASH-slab in most circumstances with the 3D coordinate keys and varying length values that are common in real-world applications. While warp-oriented operations heavily used in SlabHash enjoy the benefits of intrinsic acceleration, they sacrifice the granularity of operations. Threads can only move on to the next task once all the operations in a warp (of 32) are finished. As a result, early termination when an insertion failure occurs are less likely in a warp-oriented hash map. If the data layout is not well-distributed for the intrinsic operations (*e.g.*, low-uniqueness input, keys with long word width), the performance drop could be significant.

This observation is more apparent in insertion under varying input densities. With a relatively small value size and a high uniqueness, ASH-slab performs better. When the uniqueness is low, however, each thread in ASH-slab still has to finish a similar workload before termination, while ASH-stdgpu can reject many failure insertions early and move on to the following workloads. As of now, ASH-slab is suitable for the voxel downsampling application, while ASH-stdgpu is better for other tasks. Therefore, we set stdgpu as the default backend for ASH in the remaining sections.

TABLE 2
Comparison of the complexity of coding (top) and LoC (bottom) of each operation among the implementations. Unlike stdgpu and SlabHash, ASH does not require that users write device code or use a CUDA compiler. It requires few LoC for construction, query, and insertion.

	stdgpu	SlabHash	ASH
Device code free?	✗	✓	✓
CUDA compiler free?	✗	✗	✓
Construct	3	9	3
Find	22	2	2
Insert	27	1	2

5.4 Code Complexity

We now study the usage at the user end. First of all, the ASH framework, regardless of the backend used, is already compiled as a library. A C++ developer can easily include the header and build the example directly with a CPU compiler and link to the precompiled library with a light tensor engine. An equivalent Python interface is provided via pybind [26] as shown in Fig. 1.

In comparison, to use SlabHash’s interface with an input array from host memory, a CUDA compiler is required, along with manual bucket-linked list chain configurations. For further performance improvement, detailed memory management has to be done manually via `cudaMalloc`, `cudaMemcpy` and `cudaFree`. stdgpu provides a built-in memory manager but requires writing device and host functions. In a query operation, the found values are returned *by-copy* for SlabHash, so in-place modification requires further modification of the library. stdgpu exposes iterators in an AoS fashion, therefore the device code needs to be implemented to reinterpret an array of iterators and masks for further operations.

The compilation complexity and the interface LoC required for the same functionality in C++ are listed in Table 2.

6 APPLICATIONS

We now demonstrate a number of applications and ready-to-use systems in 3D perception to demonstrate the power of ASH with fewer LoC and better performance. The presented applications include:

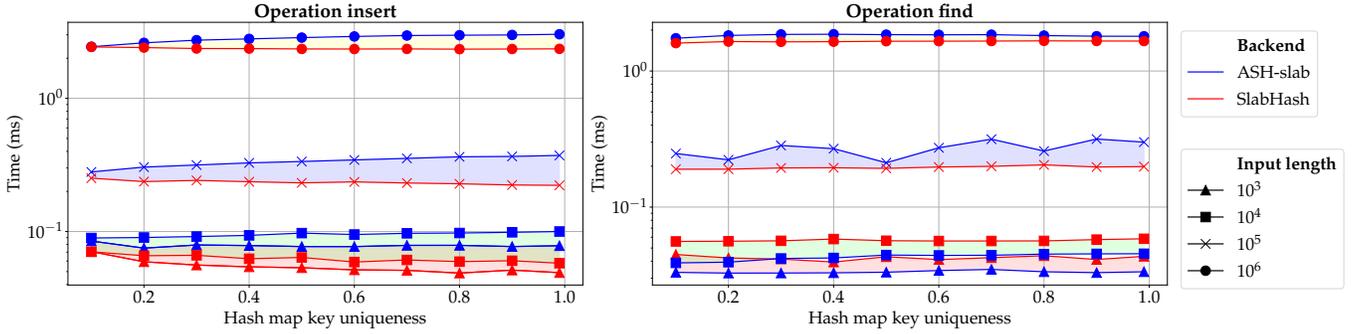


Fig. 7. Hash map performance comparison between ASH-slab and the vanilla SlabHash with 1D integer keys and values. Each curve shows the average operation time (y -axis) with varying input key uniqueness ratio (x -axis), given an input length. *Lower is better*. ASH-slab retains a comparable performance for integers while supporting generalization to arbitrary dimensional keys and values of various data types.

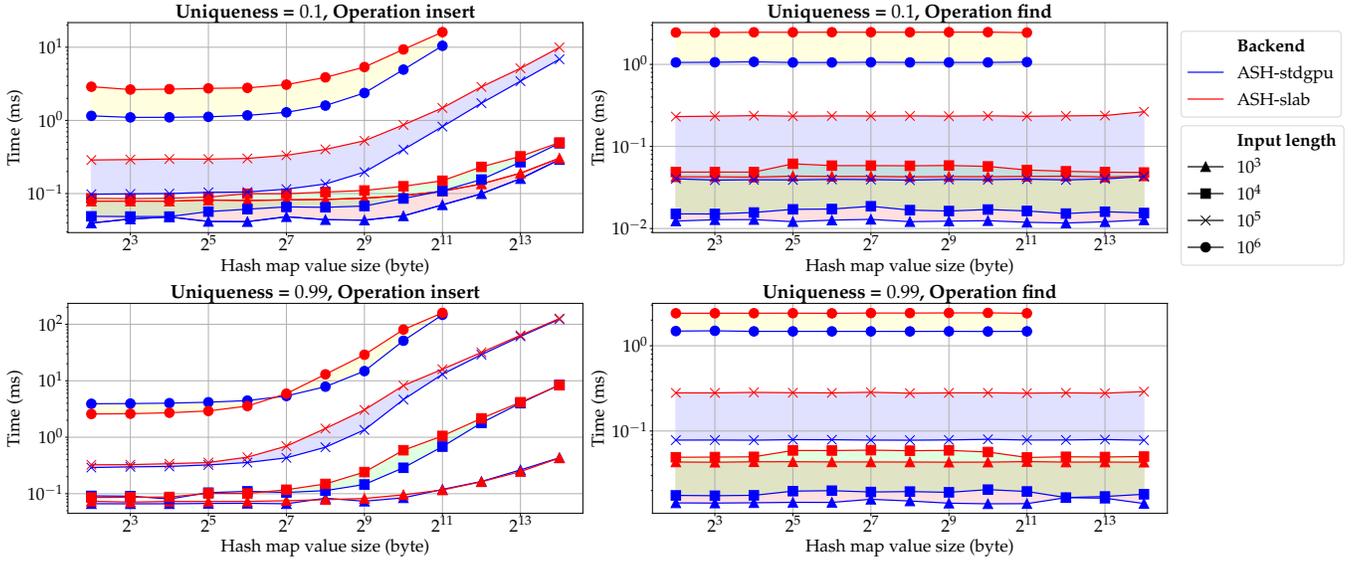


Fig. 8. Ablation study of the hash map performance with 3D integer keys over different backends. Each curve shows the average operation time (y -axis) with varying hash map value sizes in bytes (x -axis), given a controlled backend, input length, and input key uniqueness ratio. *Lower is better*. ASH-stdgpu outperforms ASH-slab in most circumstances with the 3D integer keys and varying length values, which are common in real-world applications.

- 1) Point cloud voxelization;
- 2) Retargetable volumetric reconstruction;
- 3) Non-rigid registration and deformation;
- 4) Joint geometry and appearance refinement.

The first two experiments are conducted on an Intel i7-6700HQ CPU and an Nvidia GeForce GTX 1070 GPU for indoor scenes. Outdoor scene experiments are run on an Intel i7-11700 CPU and an Nvidia RTX 3060 GPU. The rest are done on an Intel i7-7700 CPU and an Nvidia GeForce GTX 1080Ti GPU.

6.1 Point Cloud Voxelization

Setup 3. In voxelization, a hash map maps a point cloud’s discretized coordinates to its natural array indices, and the hash map capacity is the point cloud size, typically ranging from 10^5 to 10^7

$$\mathcal{K} = \{\text{Tensor}((3), \text{Int}32)\},$$

$$\mathcal{V} = \{\text{Tensor}((1), \text{Int}32)\}.$$

Voxelization is a core operation for discretizing coordinates. It is essential for sparse convolution [9], [10] at

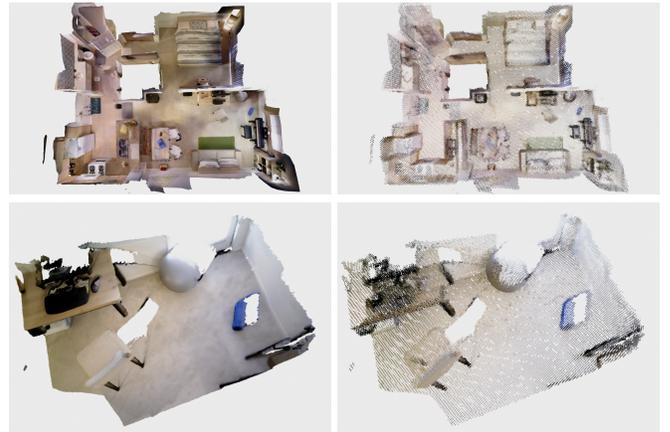


Fig. 9. Visualization of point cloud voxelization. Top: scene-level large-scale inputs. Bottom: fragment-level small-scale inputs. Left: original point clouds. Right: voxelized point clouds.

the quantization preprocessing stage, and is often used to generate point cloud “pyramids” [62] in coarse-to-fine 3D registration for improved speed and robustness.

Voxelization is a natural task for parallel hashing, as the essence of the operation is to discard duplicates at grid points. To achieve this, we first discretize the input by converting the coordinates described by the continuous meter metric to the voxel units. Then a simple hash set insertion eliminates the duplicates and corresponds them to the remaining unique coordinates. The returned indices can be reused for tracing other properties such as colors and point normals associated with the input. The python code for voxelization can be found in Listing 1.

Listing 1. Voxelization

```

1 import open3d.core as o3c
2
3 def voxelize(pcd, voxel_size):
4     xyz = pcd.point.positions
5     N = len(xyz)
6     hashset = o3c.HashSet(N, o3c.int32, (3,))
7     xyz_int = (xyz / voxel_size).floor().to(o3c.int32)
8     _, mask = hashset.insert(xyz_int)
9     # Return points with indices
10    return xyz[mask], o3c.arange(N)[mask]

```

We compare voxelization implemented in ASH with two popular implementations, MinkowskiEngine [10] on CUDA and Open3D [62] on CPU. Our experiments are conducted on a large scene input with 8×10^6 points which is typical for scene perception, and a small fragment of the scene with 5×10^5 points which is typical for an RGB-D input frame, as shown in Fig. 9.

To evaluate the performance, we vary the parameter *voxel size* from 5mm to 5cm, which is typical in the spectrum of voxelization applications, from dense reconstruction to feature extraction. In Fig. 10 we can see that our implementation outperforms baselines consistently for inputs at both scales. Meanwhile, in measuring the LoC written in C++ (the Python wrappers are one-liners) required for the functionality given the hash map interface, we observe that ASH requires only 28 LoC, while MinkowskiEngine and Open3D on CPU take 71 and 72 LoC respectively.

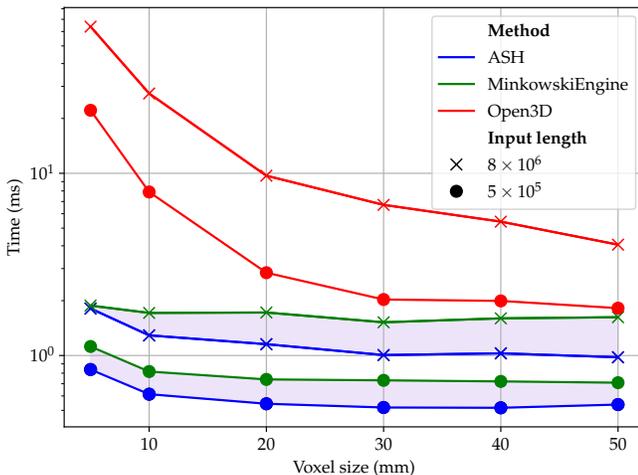


Fig. 10. Performance comparison of voxelization. Each curve shows the run time (*y*-axis) over the varying voxel size (*x*-axis). Lower is better. ASH is consistently faster than Open3D’s default voxelizer (CPU) and MinkowskiEngine (CUDA).

6.2 Retargetable Volumetric Reconstruction

6.2.1 Truncated Signed Distance Function

Scene representation with truncated signed distance function (TSDF) from a sequence of 3D input has been introduced [11] and adapted to RGB-D [40]. It takes a sequence of depth images $\{D^j\}$ with their poses $\{\mathbf{T}^j \in \mathbf{SE}(3)\}$ as input, and seeks to optimize the signed distance, an implicit function value d per point at $\mathbf{x} \in \mathbb{R}^3$. The signed distance measured for frame j is given by³

$$[u, v, r]^\top = \Pi(\mathbf{T}^{j-1} \mathbf{x}), \quad (9)$$

$$d^j = D^j(u, v) - r, \quad (10)$$

where Π projects the 3D point to 2D with a range reading after a rigid transformation. To reject outliers, a truncate function $\Psi_\mu(d) = \text{clamp}(d, -\mu, \mu)$ is applied to d^j . There are multiple variations of Ψ and the definition of signed distance d^j [5]. For this paper, we follow the convention in KinectFusion [40].

With a sequence of inputs, per-point signed distance can be estimated in least squares with a closed-form solution

$$d = \arg \min_t \sum_j w^j \|t - d^j\|^2, d = \frac{\sum_j w^j d^j}{\sum_j w^j}, \quad (11)$$

where w^j is the selected weight depending on view angles and distances [5]. In other words, with a sequence of depth inputs and their poses, we can measure TSDF at any point in 3D within the camera frustums. We can also rewrite Eq. 11 incrementally:

$$d := \frac{w \cdot d + w^j \cdot d^j}{w + w^j}, w := w + w^j, \quad (12)$$

where w is the accumulated weight paired with d .

Equipped with a projection model Π that converts a point to signed distance, TSDF reconstruction can be generalized to imaging LiDARs [13] for larger scale scenes.

6.2.2 Spatially Hashed TSDF Blocks

Setup 4. In a scene represented by a volumetric TSDF grid, a hash map maps the coarse voxel blocks’ coordinates to the TSDF data structure of the voxel block, and the hash map capacity is typically 10^3 to 10^5 for small to large-scale indoor scenes:

$$\mathcal{K} = \{\text{Tensor}((3), \text{Int32})\},$$

$$\mathcal{V} = \{\text{Tensor}((\ell^3), \text{Float32}), \text{Tensor}((\ell^3), \text{Float32})\},$$

where ℓ is the voxel block resolution, which is set to 8 or 16.

While recent neural representations utilize multi-layer perceptrons to approximate the TSDF in continuous space [7], classical approaches use discretized voxels. Such representations have a long history and are ready for real-world, real-time applications. They can also provide data for training neural representations.

The state-of-the-art volumetric discretization for TSDF reconstruction is spatial hashing, where points are allocated around surfaces on-demand at a voxel resolution of around 5mm. While it is possible to hash high-resolution voxels directly, the access pattern could be less cache-friendly,

3. Details including depth masking and projective pinhole camera model are omitted for clarity and could be found in KinectFusion [40].

as the neighbor voxels are scattered in the hash map. A hierarchical structure is a better layout, where small dense voxel grids (e.g. in the shape of 8^3 or 16^3) are the minimal unit in a hash map; detailed access can be redirected to simple 3D indexing. In other words, a voxel can be indexed by a coupled hash map lookup and a direct local addressing

$$\mathbf{x}_{\text{block}} = \lfloor \mathbf{x} / (s\ell) \rfloor, \quad (13)$$

$$\mathbf{x}_{\text{voxel}} = \lfloor (\mathbf{x} - \mathbf{x}_{\text{block}} \cdot s\ell) / s \rfloor, \quad (14)$$

where s is the voxel size and ℓ is the voxel block resolution as described in Setup 4.

While previous implementations [12], [41], [47] have achieved remarkable performance, modularized designs are missing. Geometry processing and hash map operations were coupled due to the absence of a well-designed parallel GPU hash map. One deficiency of this design is *unsafe* parallel insertion, where the capacity of a hash map can be exceeded. Another is ad hoc recurring low-level linked list access in geometry processing kernels that cause high code redundancy. Our implementation demonstrates the first modularized pipeline where safe hash map operations are used without any ad hoc modifications.

6.2.3 Voxel Block Allocation and TSDF Integration

Setup 5. For an input depth image, a hash map maps the unprojected point coordinates to the active indices as described in Setup 4, and the capacity of a hash map is typically 10^5 to 10^6 , with 10^2 to 10^3 valid entries:

$$\mathcal{K} = \{\text{Tensor}((3), \text{Int32})\},$$

$$\mathcal{V} = \{\text{Tensor}((1), \text{Int32})\}.$$

In the modularized design, we first introduce a double hash map structure for voxel block allocation and TSDF estimation. Voxel block allocation identifies points from $\{D^j\}$ as surfaces and computes coordinates with Eq. 13. Intuitively, they can be directly inserted to the *global* hash map described in Setup 4. This is achievable in an ad hoc implementation where the core of the hash map is modified at the device code level, and *unsafe* insertion is allowed [41], [47]. However, in a modularized and *safe* setup, this could lead to problems. A VGA resolution depth input contains $640 \times 320 \approx 3 \times 10^5$ points and easily exceeds the empirical *global* hash map capacity. As we have mentioned, rehashing will be triggered under such circumstances, which is both time and memory consuming, especially for a hash map with memory-demanding voxel blocks as values.

To address this issue without changing the low-level implementation and sacrificing safety, we introduce a second hash map, the *local* hash map from Setup 5. This hash map is similar to the one used in *voxelization*: it maps discretized 3D coordinates unprojected from depths to integer indices. With this setup, a larger input capacity is acceptable, as the *local* hash map is lightweight and can be cleared or constructed from scratch per iteration.

There are two main benefits to using a *local* hash map: it converts the input from the 10^5 raw point scale to the 10^3 voxel block scale, which is safe for the *global* hash map without rehashing; as a byproduct, it keeps track of the active voxel blocks for the current frame j , which can be directly used in the following TSDF integration and ray

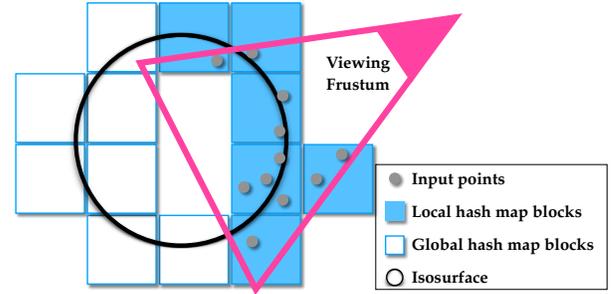


Fig. 11. Illustration of local and global hash maps iteratively used in real-time reconstruction. The local hash map activates voxel blocks enclosing points observed in the viewing frustum. The global hash map accumulates such activated blocks and maintains all the blocks around the isosurface.

casting. The local and global hash maps can be connected through indices, where a query of coordinate in the local map is redirected to the global map in-place. Fig. 11 shows the roles of the two hash maps. Listing 2 details the construction and interaction between the two hash maps.

Listing 2. Double hash map allocation

```

1 import open3d.core as o3c
2
3 # Map block coords to actual storage
4 global_hashmap = o3c.HashMap(
5     global_capacity,
6     key_dtype=o3c.int32,
7     key_shape=(3,),
8     # Float 1-channel TSDF and 3-channel color
9     value_dtypes=(o3c.float32, o3c.float32),
10    values_shapes=((8, 8, 8, 1), (8, 8, 8, 3)))
11 # Map block coords to global hashmap indices
12 local_hashmap = o3c.HashMap(
13     local_capacity,
14     key_dtype=o3c.int32,
15     key_shape=(3,)
16     # Index in global hash map
17     value_dtypes=(o3c.int32),
18     value_shapes=((1,)))
19 # Discretize and insert to local map
20 xyz_int = (xyz / block_size).floor().to(o3c.int32)
21 i_local, mask = local_hashmap.activate(xyz_int)
22 # Remove duplicates
23 xyz_int = xyz_int[mask]
24 i_local = i_local[mask]
25 # Activate and query in the global map
26 global_hashmap.activate(xyz_int)
27 i_global, mask = global_hashmap.find(xyz_int)
28 # Associate local and global maps via indices
29 local_v = local_hashmap.value()
30 local_v[i_local] = i_global

```

By accessing the global hash map’s TSDF and colors through returned indices, TSDF integration can then be implemented following Eq. 12 in a pure geometry function, either in a low-level GPU kernel or a high-level vectorized Python script. Spatial hash map is detached from the core geometric computation, providing more flexibility in performance optimization.

6.2.4 Surface Extraction

A volumetric scene reconstruction is not usable for most software and solutions until the results are exported to point clouds or triangle meshes. Hence we implement a variation of Marching Cubes [34] that extracts vertices with triangle faces at zero crossings in a volume. In a spatially

hashed implementation, boundary voxels in voxel blocks are harder to process since queries of neighbor voxel blocks are frequent, and shared vertices among triangles are hard to track. One common approach is to simply visit vertices at isosurfaces and disregard duplicates, but this usually results in a heavily redundant mesh [30], [47], or time-consuming post-processing to merge vertices [41]. Another method is to introduce an assistant volumetric data structure to atomically record the vertex-voxel association, but the implementations are over-complex and require frequent low-level hash map queries coupled with surface extraction [12], [14].

Now that we have a unified hash map interface, we simplify the voxel block neighbor search routine [12] and set up a 1 -radius neighbor lookup table in advance, as described in Listing 3. Surface extraction is then detached from hash map access and can be optimized separately. As a low-hanging fruit, point cloud extraction is implemented with the same routine by ignoring the triangle generation step. In fact, surface extraction of a median-scale scene shown in Fig. 13 takes less than 100ms, making interactive surface updates possible in a real-time reconstruction system.

Listing 3. Radius nearest neighbor search in 3D

```

1  import open3d.core as o3c
2
3  def radius_nns(xyz, r):
4      N = len(xyz)
5      hashset = o3c.HashSet(N, o3c.int32, (3,))
6      hashset.insert(xyz)
7      # Get offset tensors
8      # ([-r, -r, -r], ..., [r, r, r])
9      offsets = enumerate_radius_offsets(r)
10     # Collect neighbors
11     xyz_query = xyz.clone()
12     for offset in offsets:
13         xyz_query.append(xyz + offset, 0)
14     # Query
15     indices, masks = hashset.find(xyz_query)
16     # Reshape to get neighbor indices for each point
17     indices = indices.view(N, -1)
18     masks = indices.view(N, -1)

```

6.2.5 Ray Casting

Another way to interpret a volume is through ray casting or ray marching. Given camera intrinsics and extrinsics, ray casting renders depth and color images by marching rays in the spatially hashed volumes, querying color and TSDF values, and finding zero-crossing interfaces. It allows rendering at known viewpoints, synthesizing novel views, and estimating camera poses.

Various accelerations can speed up ray casting. Adaptive spherical ray casting and a precomputed min-max range estimate [47] will constrain the search range and boost performance. The latter can be conducted by simply projecting the active keys collected in Listing 2 without the involvement of hash maps. In addition, we can squeeze more from our double-hash-map architecture. Conventional ray marching applies query in the *global* hash map [12], [41], [47]. Since the *local* hash map is directly associated with the *global* hash map with shared active indices, we can replace the *global* hash map with the *local* one accompanied with the v^B buffer in the *global* hash map. With such a simple change, we can now query the more compact *local*

hash map, and access the *global* hash map in-place without touching the geometric computations. Since out-of-frustum voxel blocks are ignored, this operation slightly sacrifices rendering completeness at image boundaries, as shown in Fig. 12. However, it is able to boost speed by a factor of 5, and is useful for real-time systems such as dense SLAM.

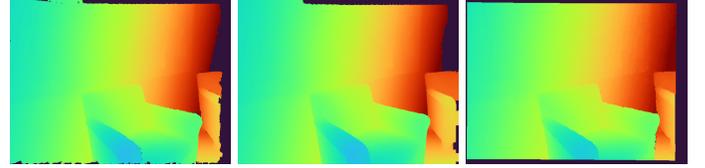


Fig. 12. Visualization of volumetric ray casting. From left to right: rendered depth from *local*, *global* hash maps, and input ground truth depth. Note the difference at boundaries.

6.2.6 Retargetable Reconstruction System

Thanks to the design where spatial hashing and geometry processing are detached, the spatially hashed volumetric representation can be reused for multiple purposes, from posed RGB-D surface reconstruction, dense RGB-D SLAM, to large-scale LiDAR surface reconstruction, with minimal modifications. For RGB-D input, a handful of existing re-



Fig. 13. Visualization of triangle mesh extracted from the real-time dense SLAM system on scene *lounge* and *copyroom* in the *fast* mode. Rendered with Mitsuba 2 [42].

construction systems [12], [41], [47] run on GPU. We first retarget our system to *fast* SLAM setup for fair comparisons. In this setup, we use a less aggressive ray-based allocation strategy [47] and store only weighted TSDF, consistent with the baselines. Camera poses are estimated in real-time via

frame-to-model alignment [41] between input frames and ray-casting rendering through the double-hash-map. We compare the performance of this setup against the state-of-the-art implementations with the same parameters: voxel size is 5.8mm, voxel block resolution is $\ell = 8$, TSDF truncation distance is 4cm, min/max acceptable range of depth scanning is 0.2m and 3m. Performance is profiled on the *lounge* scene shown in Fig. 13. Breakdown analysis of runtime and LoC⁴ is shown in Fig. 16. In most comparisons, we can see a significant performance gain, with fewer LoC to write thanks to the elimination of redundant hash map look-ups in geometry kernels.

Fast mode runs far beyond framerate and reconstructs pure geometry. In addition, we implement a *quality* mode for richer volumetric information including color, and reduce noise by integrating depth into 16^3 TSDF voxel blocks from a 1-radius-neighbor allocation [62]. The introduction of color requires double memory and triple computation cost in trilinear interpolation for ray casting and surface extraction, thus the *quality* mode is around $3\times$ slower. However, it still runs in real-time, and provides a better user experience. Fig. 14 shows the interactive reconstruction system in the *quality* mode. The system runs at 30Hz on a mid-end laptop, providing incremental volumetric reconstruction and interactive point cloud extraction and realistic rendering. Note that to retarget from a *fast* system to a user-friendly *quality* application, we only need to change the block allocation function and several parameters, in total several dozen of lines, without re-writing the core.

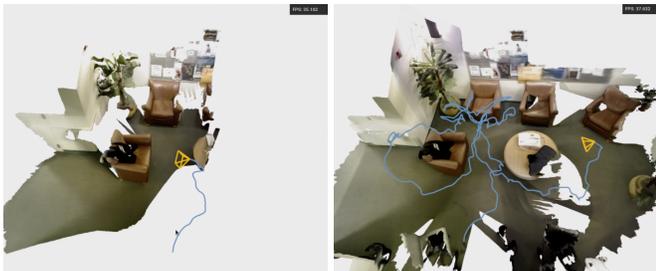


Fig. 14. Visualization of the real-time dense SLAM system in the *quality* mode with colored and interactive surface reconstruction. Viewpoints can be changed by users to visualize the incremental reconstruction of the scene.

The system can also be adapted to LiDAR point clouds. By simply replacing the conventional pinhole camera model with a customized spherical projection model [13], we can reconstruct large-scale scenes from LiDAR data. Fig. 15 shows the fast city-scale scene reconstruction from posed KITTI LiDAR point clouds [18]. At such a scale, to the best of our knowledge, GPU accelerated volumetric reconstruction has not been achieved before. Thus we compare our approach against efficient CPU LiDAR volumetric reconstruction pipelines [23], [43] in the Ouster imaging LiDAR dataset [13]. The retargeted reconstruction system results in significant performance gain, better reconstruction quality [13], at a small developing cost in LoC, as seen in Table 3.

To conclude this subsection, we presented a retargetable volumetric reconstruction system with a modular design,

4. All the code are reformatted with clang-format with a modified Google style.

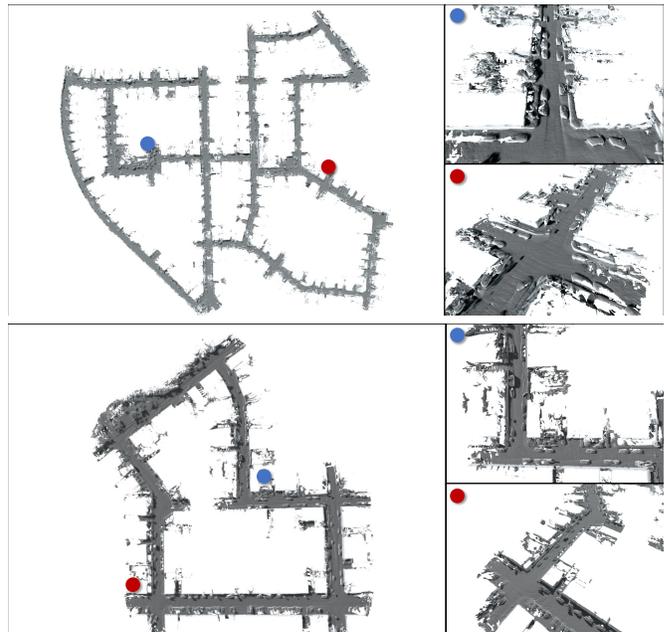


Fig. 15. City-scale volumetric TSDF reconstruction of KITTI LiDAR sequence 00 and 07 with a voxel size of 0.5m.

TABLE 3
Performance and LoC (top) and reconstruction quality (bottom) comparison between ASH, OctoMap [23], and Voxblox [43] on the Ouster LiDAR dataset [13]. ASH is faster with fewer LoC, and produces a better reconstruction.

Scenario	OctoMap		Voxblox		ASH	
	Time (ms)	LoC	Time (ms)	LoC	Time (ms)	LoC
Indoor	676.20	25371	1003.26	13725	22.29	2766
Outdoor	1167.68	-	1002.22	-	58.55	-
Reconstruction quality						
F-score (\uparrow)	92.06		97.98		98.34	

separating hash maps and geometric operations. With minimal changes in geometry functions, the core volumetric representation can be used for RGB-D and LiDAR scene reconstruction, and interactive dense SLAM. Our system is faster, requires fewer LoC, and supports an easy switch between speed and fidelity.

6.3 Non-Rigid Volumetric Deformation

While fast online volumetric reconstruction is useful in exploration and visualization, offline reconstruction systems [8] are sometimes preferred when a higher quality is required for design and evaluation.

State-of-the-art offline systems adopt divide-and-conquer. Long input sequences are split into smaller subsets, each yielding a submap point cloud \mathcal{M}^j reconstruction with less drift. In this setup, we can simply reuse the integration and surface extraction components in the previous subsection [8], [12]. A global pose graph is then constructed and optimized after robust registration [9], [56] of submaps. For the details, we refer the readers to state-of-the-art RGBD reconstruction systems [8], [12].

However, issues still persist in challenging scenes, *e.g.*, heavy misalignment due to the strong simulated noise in the Augmented ICL dataset [8], and the artifacts presented

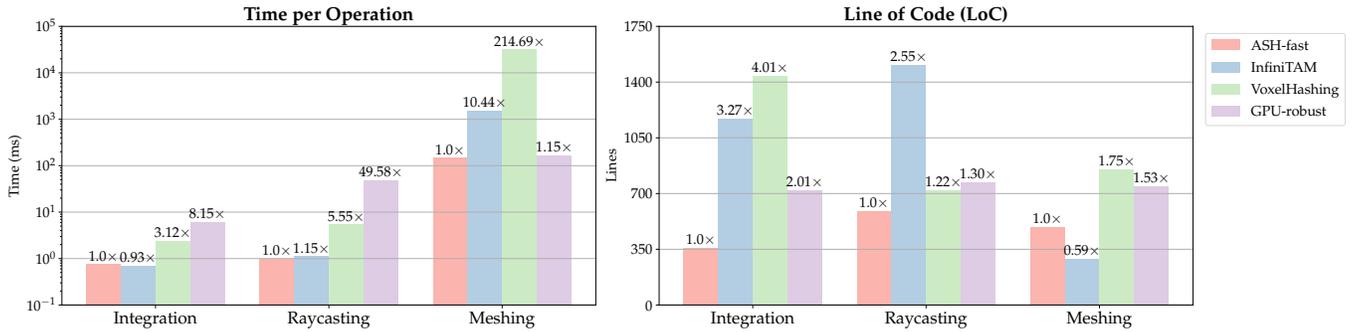


Fig. 16. Performance and LoC comparison of our real-time dense SLAM pipeline in the *fast* mode (ASH-fast) against state-of-the-art implementations: InfiniTAM [47], VoxelHashing [41], GPU-robust [14]. Evaluated on the *lounge* scene [8]. Left: detailed comparison of separating modules. Right: corresponding LoC comparison. *Lower is better*. Note the meshing LoC in InfiniTAM is significantly fewer since the implementation is over-simplified and requires further postprocessing. ASH-fast achieves a consistent fast speed with fewer LoC.

in the large-scale indoor RGBD LiDAR dataset [45], as shown in Fig. 18. To deal with this, non-rigid volumetric deformation is presented in Simultaneous Localization and Calibration (SLAC) [60], [61].

SLAC attempts to minimize the distance between correspondences from different submaps by optimizing a combination of rigid transformations and non-rigid deformations. While the rigid transformations are simply the submap poses $\{\mathbf{T}^j\}$, deformation is parameterized by a control grid \mathbf{c} . In \mathbf{c} , each grid point \mathbf{u} stores a local Euclidean offset $\mathbf{c}_\mathbf{u} \in \mathbb{R}^3$, and the accompanying function $C_\mathbf{c}(\mathbf{x}) = \mathbf{x} + \sum_{\mathbf{u} \in \mathbf{N}_\mathbf{x}} w_\mathbf{u}(\mathbf{x}) \mathbf{c}_\mathbf{u}$ deforms a point $\mathbf{x} \in \mathbb{R}^3$ by applying interpolated neighbor grid offsets, where $w_\mathbf{u}(\mathbf{x})$ is the interpolation ratio. The loss function is then parameterized over $\{\mathbf{T}^j\}$ and \mathbf{c} with as-rigid-as-possible regularizers:

$$\begin{aligned} \min_{\mathbf{c}, \mathbf{T}} \quad & \sum_{\mathbf{p} \in \mathcal{M}^i, \mathbf{q} \in \mathcal{M}^j} \|(\mathbf{T}^i C_\mathbf{c}(\mathbf{p}) - \mathbf{T}^j C_\mathbf{c}(\mathbf{q}))\|^2 \\ & + \lambda \sum_{\mathbf{u}, \mathbf{v} \in \mathbf{N}_\mathbf{u}} \|\mathbf{c}_\mathbf{u} - \mathbf{R}_\mathbf{v}^{\mathbf{c}_\mathbf{v}} \mathbf{u}\|^2, \end{aligned} \quad (15)$$

where $\mathbf{p} \in \mathcal{M}^i, \mathbf{q} \in \mathcal{M}^j$ are corresponding 3D points between submaps obtained by nearest neighbor search. $\mathbf{R}_\mathbf{v}^{\mathbf{c}_\mathbf{v}}$ is the rigid rotation that minimizes $\|\mathbf{R}_\mathbf{v}^{\mathbf{c}_\mathbf{v}}(\mathbf{v} - \mathbf{u}) - (\mathbf{c}_\mathbf{v} - \mathbf{c}_\mathbf{u})\|$ locally, where \mathbf{v} is a 1-ring neighbor $\mathbf{N}_\mathbf{u}$ of \mathbf{u} . It controls local distortions in the as-rigid-as-possible regularizer.

This problem formulation is complicated to realize in code, and in the original implementation, the deformation grid is a simplified dense 3D array where points out-of-bound are discarded during optimization. As of today, SLAC has never been reproduced apart from the original implementation. We observe that similar operations for TSDF grids can be applied here by ASH to generate a spatially hashed control grid.

Setup 6. A volumetric deformation hash map maps grid coordinates to position offsets, and the capacity of the hash map is typically 10^3 to 10^4 :

$$\begin{aligned} \mathcal{K} &= \{\text{Tensor}((3), \text{Int}32)\}, \\ \mathcal{V} &= \{\text{Tensor}((3), \text{Float}32)\}. \end{aligned}$$

Equipped with ASH, the non-rigid deformation can be written in several lines which results in a significant drop in LoC. We first voxelize the input point cloud with the deformation grid size following Listing 1. Then, instead

of the 1-radius nearest neighbors (3^3 entries in Listing 3), we look for 1-cube nearest neighbors (2^3), where a point is enclosed in a cube formed by grid points. The interpolation ratio can be computed jointly. We also adapt the 1-radius neighbor search to 1-ring neighbors for the regularizer. The embedding of a submap point cloud is visualized in Fig. 17, where the edges indicate the association between points to grids, and the colors show the interpolation ratio. Note that this visualization is also made easy thanks to the simple interface of ASH.

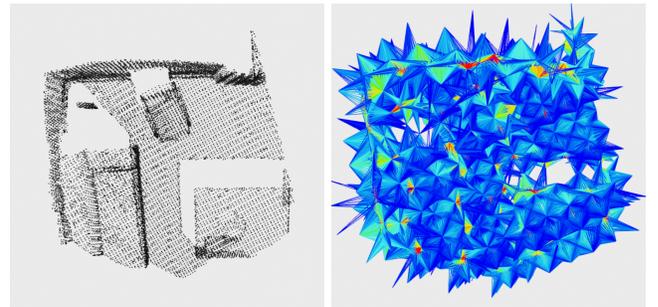


Fig. 17. Visualization of a point cloud and its embedding in the volumetric deformation grids. Left: original point cloud. Right: embedding graph connecting the input points and associated deformation grid points. Each edge's color indicates the interpolation weight: blue shows a lower weight (closer to 0), while red shows a higher weight (closer to 1).

With this parameterization, we reproduce SLAC after rewriting the non-linear least squares solver and jointly optimizing the grid points and submap poses given the correspondences. In addition, the hash map can be saved and loaded from the disk for further processing, including deformed TSDF integration that reconstructs the scene from the deformed input depth images embedded in the grids. Experiments show that with a modularized design and a spatial hash map, we can reproduce SLAC by reducing artifacts after optimization, as shown in Fig. 18.

We can see a gain in performance with fewer LoC in Table 4 in the *livingroom 1* scene with heavy simulated noise⁵. Note while the hash map generalizes deformation grids from bounded to unbounded scenes, the LoC and time contributing to the core non-linear least squares op-

⁵ To control the experiment, we use the initial submap pose graph from the baseline implementation.

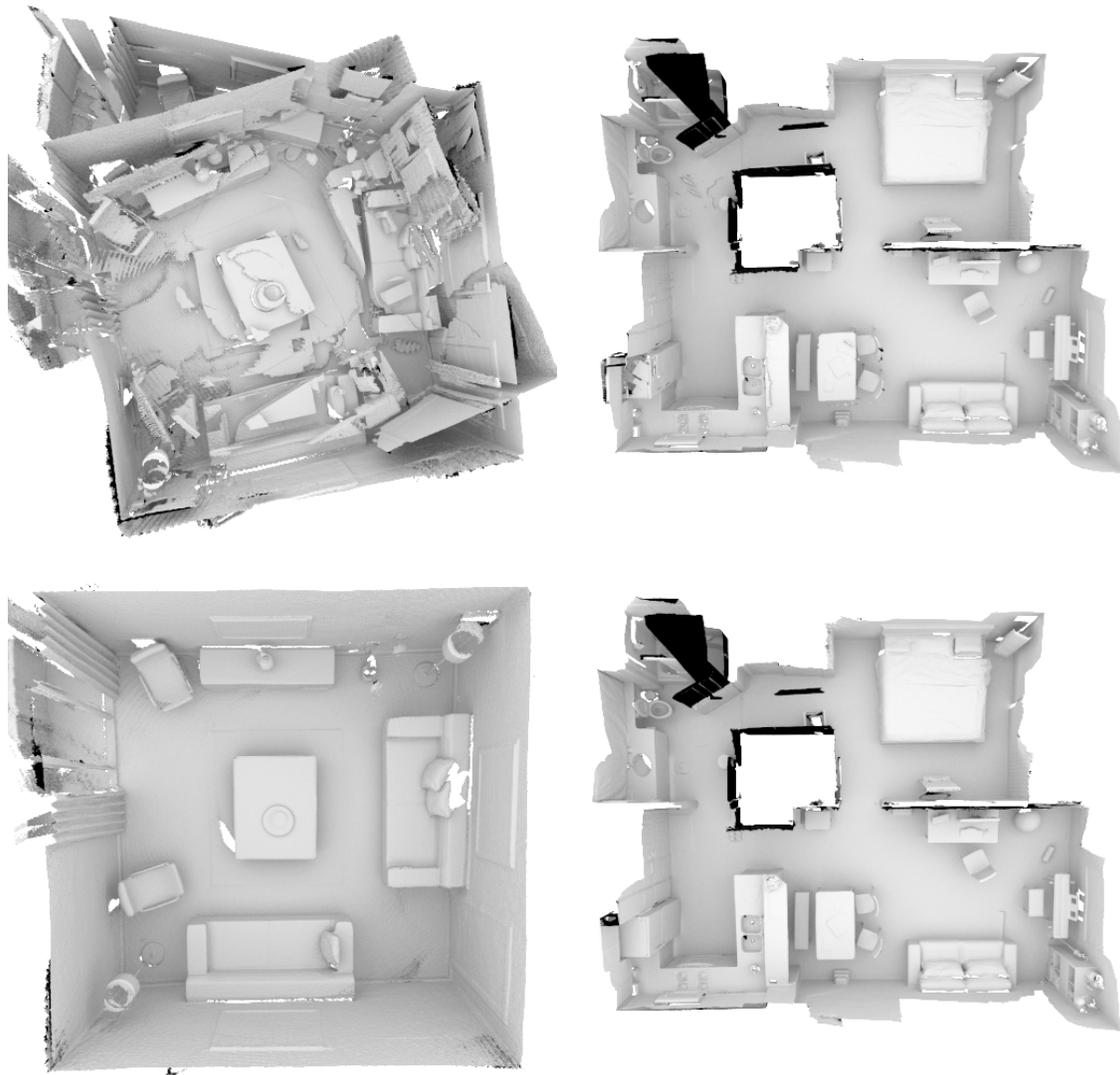


Fig. 18. Visualization of scene reconstructions before and after ASH-SLAC. First row: before ASH-SLAC. Second row: after ASH-SLAC. Left: *livingroom-1* from Augmented ICL [8]. Right: *apartment* from Indoor LiDAR RGBD [45]. Artifacts are eliminated by global pose adjustment and local deformation via deformable TSDF integration. Rendered with Mitsuba 2 [42].

timization are slightly reduced. Meanwhile, the deformation and integration speed *per frame* is significantly faster ($11.8\times$), which is critical for large-scale ($\geq 30K$ frames) sequences. While being faster and easier to develop, our system achieves a higher reconstruction quality in terms of precision, recall, and F-score with a distance threshold $\tau = 20\text{mm}$ [45].

TABLE 4

Performance and LoC (top) and reconstruction quality (bottom) comparison between ASH-SLAC and the original implementation [60] on the *livingroom-1* scene [8]. ASH-SLAC is faster with fewer LoC, and produces a better reconstruction.

Operation	Original SLAC		ASH-SLAC	
	Time (ms)	LoC	Time (ms)	LoC
Non-rigid optim.	2041.1	1585	1982.1	1535
Deformed integration	125.38	944	10.62	446
Reconstruction quality				
Precision (\uparrow)	29.19		36.10	
Recall (\uparrow)	51.44		61.34	
F-score (\uparrow)	37.24		45.45	

6.4 Joint Geometry and Appearance Refinement

SLAC reduces artifacts for large-scale scenes. For small-scale objects, while volumetric reconstruction outputs smooth surfaces, fine details are often impaired due to the weight averaging of the TSDF.

Shape-from-Shading (SfS) refines details by jointly optimizing volumetric TSDF functions given the initial geometry and appearance [63]. It takes a reconstructed volumetric TSDF grid \mathbf{d}^0 with a set of high-resolution key frame RGB images I^j and their poses \mathbf{T}^j as input, and outputs jointly optimized TSDF \mathbf{d} and albedo \mathbf{a} through an image formation model

$$\begin{aligned}
 \min_{\mathbf{a}, \mathbf{d}} \sum_{\mathbf{x}, j} & \|\nabla B(\mathbf{x}) - \nabla I^j(\Pi(\mathbf{T}^{j-1} \mathbf{x}')\|)^2 \\
 & + \lambda_{\text{smooth}} \sum_{\mathbf{x}} \|\Delta \mathbf{d}_{\mathbf{x}}\|^2 + \lambda_{\text{init}} \sum_{\mathbf{x}} \|\mathbf{d}_{\mathbf{x}} - \mathbf{d}_{\mathbf{x}}^0\|^2 \\
 & + \lambda_{\text{chrome}} \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{N}_{\mathbf{x}}} w(\mathbf{x}, \mathbf{y}) \|\mathbf{a}_{\mathbf{x}} - \mathbf{a}_{\mathbf{y}}\|^2, \quad (16)
 \end{aligned}$$

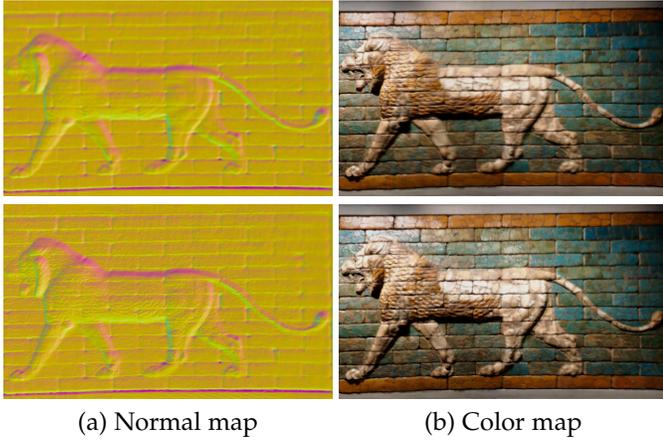


Fig. 19. Appearance and geometry refinement before and after ASH-Intrinsic3D on *lion* [63]. First row: initial reconstruction from volumetric integration. Second row: refined reconstruction after optimization.

where the estimated voxel-wise appearance is computed by $B(\mathbf{x}) = \mathbf{a}_x \text{SH}(\mathbf{n}_x)$ (SH stands for spherical harmonics), and associated with the closest surface point

$$\mathbf{x}' = \mathbf{x} - \mathbf{d}_x \mathbf{n}_x, \quad \mathbf{n}_x = \frac{\nabla_{\mathbf{x}}(\mathbf{d})}{\|\nabla_{\mathbf{x}}(\mathbf{d})\|}, \quad (17)$$

which is projected to image I^j through Π after a rigid transformation \mathbf{T}^{j-1} . Here the voxel-wise gradient is directly derived from \mathbf{d} with a finite difference

$$\nabla_{\mathbf{x}}(\mathbf{d}) = \frac{\mathbf{d}_{\mathbf{x}+\delta} - \mathbf{d}_{\mathbf{x}-\delta}}{2\delta}. \quad (18)$$

Similar to SLAC, we use $\mathbf{d}_x, \mathbf{a}_x$ to access TSDF and albedo values at grid point $\mathbf{x} \in \mathbb{R}^3$. $\lambda_{\text{smooth}}, \lambda_{\text{init}}, \lambda_{\text{chrome}}$ are coefficients for regularizing smoothness through the Laplacian, stability, and piece-wise albedo constancy via a weighted chromaticity regularizer $w(\mathbf{x}, \mathbf{y})$, respectively [63].

While the image formation model is straightforward, similar to SLAC, the underlying data structure used in implementing the model can be complex mainly because of the prevalent nearest neighbor search in normal computation and neighbor voxel regularizers. As a result, to enable such a system without a modern hash map, one has to rely on low-level C++ implementation and is consequently limited to the low-level Ceres solver [1] for autodiff in optimization. Further, the spatially hashed voxels have to be bounded to reduce computation cost [35].

Now equipped with ASH, we provide a simplified solution that is built upon the hash map and advanced indexing. Unlike SLAC which requires time-consuming deformable TSDF re-integration for final scene reconstruction, Sfs allows reusing accelerated surface extraction from the TSDF grids without further optimization. Therefore, we implement the Sfs pipeline in *pure Python* as an example of fast prototyping of a differentiable rendering pipeline. Running on GPU, we lift the constraint of a user-defined bounding box and optimize the full reconstructed surface.

Without the requirement of extreme performance, we drop the hierarchical volumetric layout and use the simple voxel-based hash map:

Setup 7. A voxel indexer is given by a hash set $\mathcal{K} = \{\text{Tensor}((3), \text{Int32})\}$. The typical capacity is 10^9 to 10^{10} .

With this setup, we can reuse the code in SLAC to look up the 1-ring neighbors for normal estimation and Laplacian regularization. There is, however, another lookup required since we are minimizing the difference of appearance *gradient* in Eq. 16: we need to find the 1-ring neighbors that also *have* 1-ring neighbors. In other words, we have to find the intersection of two sets. While NumPy provides the functionality for 1D arrays through *ordered* sorting, our hash map allows *unordered* intersection that can be generalized to multi-dimensional inputs:

Setup 8. With two input sets $\mathbf{k}_1 \subset \mathcal{K}, \mathbf{k}_2 \subset \mathcal{K}$, the intersection $\mathbf{k}_1 \cap \mathbf{k}_2$ is given by the following operations: initialize a hash set with \mathbf{k}_1 ; query \mathbf{k}_2 and obtain success mask θ ; return $\mathbf{k}_2(\theta)$.

After data association is found and SH parameters are estimated in a preprocessing step, all the terms in Eq. 16 are converted to a trivial combination of indexing and arithmetic operations. We can take advantage of PyTorch’s autodiff, and backpropagate the gradient through the built-in differentiable index layer. ADAM [29] with an initial learning rate 10^{-3} is used. Thus the core volumetric Sfs pipeline [63] is reproduced in pure Python.

An extension can be easily implemented by introducing spatially varying lights [35], wrapped up with a hash map.

Setup 9. Spatially varying spherical harmonics (SVSH) (bands = 3) can be described by a hash map that maps lighting subvolume coordinates to the corresponding coefficients:

$$\begin{aligned} \mathcal{K} &= \{\text{Tensor}((3), \text{Int32})\}, \\ \mathcal{V} &= \{\text{Tensor}((9), \text{Float32})\}. \end{aligned}$$

The embedding of an active voxel in an SVSH map is identical to SLAC, with 1-cube neighbors for the data term and 1-ring neighbors for the regularizer. Further description is omitted here as the formulation and implementation are similar to Eq. 16 [35].

Having both Sfs and SVSH optimization implemented⁶ [35], we show the results on the scene *lion* in Fig. 19. Without voxel grid upsampling, both the geometry and appearance details are sharper. Regarding performance and code complexity, we show in Table 5 that our code is much shorter in pure Python, and $150\times$ faster per iteration thanks to the CUDA autodiff engine in PyTorch. Note that Ceres is a 2nd-order optimizer on CPU that empirically converges faster than the 1st-order ADAM optimizer. In practice, however, we found that in 50 iterations ADAM converges well against the preset 10 iterations for the non-linear least squares solver. Thus the total optimization performance of our implementation is still $30\times$ faster with more voxels to process (remember that we do not require an additional bounding box).

We also evaluate reconstruction quality in Table 5. We render our optimized mesh given the keyframe camera extrinsic and intrinsic parameters and compute RMSE against the raw input images. For the baseline [35], we follow a similar procedure and render the optimized mesh (not upsampled for fairness) given refined camera parameters. We use the same mask given by the baseline to ensure the

6. Pose optimization and voxel upsampling are disabled at current.

same region of interest. The results show that our implementation produces improved RMSE despite the simplified development.

TABLE 5

Performance per epoch and LoC (top), and rendering quality (bottom) comparison between ASH-Intrinsic3D and the original implementation [35] on the *lion* scene [35]. ASH-Intrinsic3D is faster with fewer LoC, and results in comparable rendering from the refined reconstruction.

Operation	Original-Intrinsic3D		ASH-Intrinsic3D	
	Time (s)	LoC	Time (s)	LoC
SVSH optim.	0.503	605	0.092	254
Joint optim.	147.323	7399	0.916	1416
Rendering quality				
RMSE mean (\downarrow)	0.677		0.627	
RMSE std (\downarrow)	0.095		0.120	

7 CONCLUSION AND FUTURE WORK

We presented ASH, a performant and easy-to-use framework for spatial hashing. Both synthetic and real-world experiments demonstrate the power of the framework. With ASH, users can achieve the same or better performance in 3D perception tasks while writing less code.

There are various avenues for future work. At the architecture level, we seek to introduce the open address variation [2], [27] of parallel hash maps for flexibility and potential high performance static hash maps. At the low level, we plan to further optimize the GPU backend, and accelerate the CPU counterpart, potentially with cache level optimization and code generation [24], [25]. We also plan to apply ASH to sparse convolution [10], [54] and neural rendering [16], [48], where spatially varying parameterizations are exploited.

ASH accelerates a variety of 3D perception workloads. We hope that the presented framework will serve both research and production applications.

ACKNOWLEDGEMENT

We thank Patrick Stotko for adapting stdgpu to our framework, Qian-Yi Zhou for the valuable comments and detailed math derivations that helped us implement SLAC, and Kwonyoung Ryu and Jaesik Park for providing the imaging LiDAR dataset.

REFERENCES

- [1] S. Agarwal, K. Mierle, and Others. Ceres solver. <http://ceres-solver.org>. 16
- [2] D. A. Alcantara, A. Sharf, F. Abbasnejad, S. Sengupta, M. Mitzenmacher, J. D. Owens, and N. Amenta. Real-time parallel hashing on the GPU. In *ACM SIGGRAPH Asia*. 2009. 1, 2, 17
- [3] S. Ashkiani, M. Farach-Colton, and J. D. Owens. A dynamic hash table for the GPU. In *IEEE IPDPS*, 2018. 1, 2, 3, 5
- [4] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975. 1, 2
- [5] E. Bylow, J. Sturm, C. Kerl, F. Kahl, and D. Cremers. Direct camera pose tracking and 3D reconstruction with signed distance functions. In *RSS*, 2013. 10
- [6] P. Celis, P.-A. Larson, and J. I. Munro. Robin hood hashing. In *Annual Symposium on Foundations of Computer Science*, 1985. 2
- [7] R. Chabra, J. E. Lenssen, E. Ilg, T. Schmidt, J. Straub, S. Lovegrove, and R. Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *ECCV*, 2020. 3, 10
- [8] S. Choi, Q. Zhou, and V. Koltun. Robust reconstruction of indoor scenes. In *CVPR*, 2015. 13, 14, 15
- [9] C. Choy, W. Dong, and V. Koltun. Deep global registration. In *CVPR*, 2020. 9, 13
- [10] C. Choy, J. Gwak, and S. Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, 2019. 9, 10, 17
- [11] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *ACM SIGGRAPH*, 1996. 3, 10
- [12] W. Dong, J. Park, Y. Yang, and M. Kaess. GPU accelerated robust scene reconstruction. In *IROS*, 2019. 2, 3, 7, 11, 12, 13
- [13] W. Dong, K. Ryu, M. Kaess, and J. Park. Revisiting lidar registration and reconstruction: A range image perspective. In *arXiv preprint arXiv:2112.02779*, 2021. 10, 13
- [14] W. Dong, J. Shi, W. Tang, X. Wang, and H. Zha. An efficient volumetric mesh representation for real-time scene reconstruction using spatial hashing. In *ICRA*, 2018. 12, 14
- [15] W. Dong, Q. Wang, X. Wang, and H. Zha. PSDF fusion: Probabilistic signed distance function for on-the-fly 3D data fusion and scene reconstruction. In *ECCV*, pages 701–717, 2018. 3
- [16] S. J. Garbin, M. Kowalski, M. Johnson, J. Shotton, and J. Valentin. Fastnerf: High-fidelity neural rendering at 200fps. *arXiv preprint arXiv:2103.10380*, 2021. 17
- [17] I. García, S. Lefebvre, S. Hornus, and A. Lasram. Coherent parallel hashing. *ACM Transactions on Graphics*, 30(6):1–8, 2011. 2
- [18] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, pages 3354–3361. IEEE, 2012. 13
- [19] J. Gunther, S. Popov, H.-P. Seidel, and P. Slusallek. Real time ray tracing on GPU with bvh-based packet traversal. In *Symposium on Interactive Ray Tracing*, 2007. 2
- [20] L. Han and L. Fang. Flashfusion: Real-time globally consistent dense 3D reconstruction using cpu computing. In *RSS*, 2018. 3
- [21] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, Sept. 2020. 3
- [22] R. K. Hoetzlein. Gvdb: Raytracing sparse voxel database structures on the GPU. In *Proceedings of High Performance Graphics*. 2016. 2
- [23] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard. Octomap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*, 34(3):189–206, 2013. 2, 13
- [24] Y. Hu, T.-M. Li, L. Anderson, J. Ragan-Kelley, and F. Durand. Taichi: a language for high-performance computation on spatially sparse data structures. *ACM Transactions on Graphics*, 38(6):1–16, 2019. 2, 17
- [25] W. Jakob. Enoki: structured vectorization and differentiation on modern processor architectures, 2019. <https://github.com/mitsuba-renderer/enoki>. 17
- [26] W. Jakob, J. Rhineland, and D. Moldovan. pybind11 – seamless operability between c++11 and python, 2017. <https://github.com/pybind/pybind11>. 8
- [27] D. Jünger, R. Kobus, A. Müller, C. Hundt, K. Xu, W. Liu, and B. Schmidt. Warpcore: A library for fast hash tables on GPUs. *arXiv preprint arXiv:2009.07914*, 2020. 2, 17
- [28] T. Karnagel, R. Müller, and G. M. Lohman. Optimizing GPU-accelerated group-by and aggregation. *VLDB*, 8:20, 2015. 2
- [29] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 16
- [30] M. Klingensmith, I. Dryanovskii, S. S. Srinivasa, and J. Xiao. Chisel: Real time large scale 3D reconstruction onboard a mobile device using spatially hashed signed distance fields. In *RSS*, 2015. 3, 12
- [31] A. Kukanov and M. J. Voss. The foundations for scalable multi-core software in intel threading building blocks. *Intel Technology Journal*, 11(4), 2007. 3
- [32] C. Lauterbach, M. Garland, S. Sengupta, D. Luebke, and D. Manocha. Fast bvh construction on GPUs. In *Computer Graphics Forum*, volume 28, pages 375–384, 2009. 2
- [33] B. Lessley and H. Childs. Data-parallel hashing techniques for GPU architectures. *IEEE Transactions on Parallel and Distributed*

- Systems*, 31(1):237–250, 2019. 2
- [34] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH*, 21(4):163–169, 1987. 11
- [35] R. Maier, K. Kim, D. Cremers, J. Kautz, and M. Nießner. Intrinsic3D: High-quality 3D reconstruction by joint appearance and geometry optimization with spatially-varying lighting. In *ICCV*, 2017. 3, 16, 17
- [36] D. Meagher. Geometric modeling using octree encoding. *Computer Graphics and Image Processing*, 19(2):129–147, 1982. 1, 2
- [37] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv:2201.05989*, Jan. 2022. 1, 2, 3
- [38] K. Museth, J. Lait, J. Johanson, J. Budsberg, R. Henderson, M. Alden, P. Cucka, D. Hill, and A. Pearce. Openvdb: an open-source data structure and toolkit for high-resolution volumes. In *ACM SIGGRAPH Courses*, pages 1–1. 2013. 2
- [39] R. A. Newcombe, D. Fox, and S. M. Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *CVPR*, 2015. 3
- [40] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *ISMAR*, 2011. 1, 10
- [41] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger. Real-time 3D reconstruction at scale using voxel hashing. *ACM Transactions on Graphics*, 32(6):1–11, 2013. 1, 2, 3, 7, 11, 12, 13, 14
- [42] M. Nimier-David, D. Vicini, T. Zeltner, and W. Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics*, 38(6):1–17, 2019. 12, 15
- [43] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto. Voxblox: Incremental 3D Euclidean signed distance fields for on-board MAV planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1366–1373. IEEE, 2017. 13
- [44] R. Pagh and F. F. Rodler. Cuckoo hashing. *Journal of Algorithms*, 51(2):122–144, 2004. 2
- [45] J. Park, Q.-Y. Zhou, and V. Koltun. Colored point cloud registration revisited. In *ICCV*, 2017. 14, 15
- [46] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshain, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*. 2019. 3, 6
- [47] V. A. Prisacariu, O. Kähler, S. Golodetz, M. Sapienza, T. Cavallari, P. H. Torr, and D. W. Murray. Infinitam v3: A framework for large-scale 3D reconstruction with loop closure. *arXiv preprint arXiv:1708.00783*, 2017. 1, 2, 3, 7, 11, 12, 14
- [48] C. Reiser, S. Peng, Y. Liao, and A. Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. *arXiv preprint arXiv:2103.13744*, 2021. 17
- [49] S. Rusinkiewicz and M. Levoy. Efficient variants of the icp algorithm. In *3DV*, 2001. 2
- [50] R. B. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (FPFH) for 3D registration. In *ICRA*, 2009. 2
- [51] R. B. Rusu and S. Cousins. 3D is here: Point cloud library (PCL). In *ICRA*, 2011. 2
- [52] P. Stotko. stdgpu: Efficient STL-like Data Structures on the GPU. *arXiv:1908.05936*, 2019. 2, 3, 5
- [53] R. D. Team. *RAPIDS: Collection of Libraries for End to End GPU Data Science*, 2018. 2
- [54] B. Ummenhofer, L. Prantl, N. Thuerey, and V. Koltun. Lagrangian fluid simulation with continuous convolutions. In *ICLR*, 2020. 17
- [55] I. Wald, S. Boulos, and P. Shirley. Ray tracing deformable scenes using dynamic bounding volume hierarchies. *ACM Transactions on Graphics*, 26(1):6–es, 2007. 2
- [56] H. Yang, W. Dong, L. Carlone, and V. Koltun. Self-supervised geometric perception. In *CVPR*, pages 14350–14361, 2021. 13
- [57] M. Zeng, F. Zhao, J. Zheng, and X. Liu. Octree-based fusion for realtime 3D reconstruction. *Graphical Models*, 75(3):126–136, 2013. 2
- [58] K. Zhou, Q. Hou, R. Wang, and B. Guo. Real-time kd-tree construction on graphics hardware. *ACM Transactions on Graphics*, 27(5):1–11, 2008. 2
- [59] Q. Zhou, J. Park, and V. Koltun. Fast global registration. In *ECCV*, 2016. 2
- [60] Q.-Y. Zhou and V. Koltun. Simultaneous localization and calibration: Self-calibration of consumer depth cameras. In *CVPR*, 2014. 3, 14, 15
- [61] Q.-Y. Zhou, S. Miller, and V. Koltun. Elastic fragments for dense scene reconstruction. In *ICCV*, 2013. 14
- [62] Q.-Y. Zhou, J. Park, and V. Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018. 2, 9, 10, 13
- [63] M. Zollhöfer, A. Dai, M. Innmann, C. Wu, M. Stamminger, C. Theobalt, and M. Nießner. Shading-based refinement on volumetric signed distance functions. *ACM Transactions on Graphics*, 34(4):1–14, 2015. 15, 16
- [64] M. Zollhöfer, P. Stotko, A. Görlitz, C. Theobalt, M. Nießner, R. Klein, and A. Kolb. State of the art on 3D reconstruction with RGB-D Cameras. In *Computer Graphics Forum*, volume 37, pages 625–652, 2018. 3



Wei Dong is a PhD student at Carnegie Mellon University since 2018, advised by Michael Kaess. Prior to PhD, he received his bachelor's degree (2015) and master's degree (2018) in Computer Science from Peking University, advised by Hongbin Zha.



Yixing Lao is a Research Engineer at Intel Intelligent Systems Lab. He received his master's degree in Computer Science from University of California, San Diego in 2016, and bachelor's degree in Computer Engineering from the University of Hong Kong in 2013.



Michael Kaess is an Associate Professor in the Robotics Institute at Carnegie Mellon University and the Director of the Robot Perception Lab. He was a Research Scientist and a Postdoctoral Associate at the Massachusetts Institute of Technology. In 2008, he received the PhD degree in Computer Science from the Georgia Institute of Technology. He was Associate Editor for IEEE Transaction on Robotics and is currently Associate Editor for IEEE Robotics and Automation Letters. He is a Senior Member of IEEE.



Vladlen Koltun is a Distinguished Scientist at Apple. Previously he was the Chief Scientist for Intelligent Systems at Intel, where he built an international research lab, based on four continents, that produced high-impact results in robotics, computer vision, image synthesis, machine learning, and other areas. He has mentored more than 50 PhD students, postdocs, research scientists, and PhD student interns, many of whom are now successful research leaders.