

Data-Parallel Hashing Techniques for GPU Architectures

Brenton Lessley¹

¹Department of Computer and Information Science, University of Oregon, Eugene, OR, USA

Abstract

Hash tables are one of the most fundamental data structures for effectively storing and accessing sparse data, with widespread usage in domains ranging from computer graphics to machine learning. This study surveys the state-of-the-art research on data-parallel hashing techniques for emerging massively-parallel, many-core GPU architectures. Key factors affecting the performance of different hashing schemes are discovered and used to suggest best practices and pinpoint areas for further research.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [High Performance Computing]: Parallel Algorithms—Hashing

1. Introduction

The problem of searching for elements in a set is a well-studied algorithm in computer science. Canonical methods for this task are primarily based on sorting, spatial partitioning, and hashing [Knu98]. In searching via hashing, an indexable hash table data structure is used for efficient random access and storage of sparse data, enabling fast lookups on average. For many years, numerous theoretical and practical hashing approaches have been introduced and applied to problems in areas such as computer graphics, database processing, machine learning, and scientific visualization, to name a few [Ull72, ML75, KU86, CHM97, Knu98, WSSJ14, WLKC16]. With the emergence of multi-processor CPU systems and thread-based programming, significant research was focused on the design of concurrent, lock-free hashing techniques for single-node, CPU shared-memory [Gre02, Mic02, PGB*05, SS06, GHS*10]. Moreover, studies began to investigate external-memory (off-chip) and multi-node, distributed-memory parallel techniques that could accommodate the oncoming shift towards large-scale data processing [BZ07, CKWT14]. These methods, however, do not demonstrate node-level scalability for the massive number of concurrent threads and parallelism offered by current and emerging many-core architectures, particularly *graphical processing units* (GPUs). GPUs are specifically designed for data-parallel computation, in which the same operation is performed on different data elements in parallel.

CPU-based hashing designs face several notable challenges when ported to GPU architectures:

- *Sufficient parallelism*: Extra instruction- and thread-level parallelism must be exploited to cover GPU global memory latencies and utilize the thousands of smaller GPU compute cores. Data-

parallel design is key to exposing this necessary parallel throughput.

- *Memory accesses*: Traditional pointer-based hash tables induce many random memory accesses that may not be aligned within the same cache line, leading to multiple global memory loads that limit throughput on the GPU.
- *Control flow*: Lock-free hash tables that can be both queried and updated induce heavy thread contention for atomic read-write memory accesses. This effectively serializes the control flow of threads and limits the thread-level parallelism on the GPU.
- *Limited memory*: CPU-based hashing leverages large on-chip caching and shared memory to support random-access memory requests quickly. On the GPU, this fast memory is limited in size and can result in more cache misses and expensive global memory loads.

In this study, we survey the state-of-the-art data-parallel hashing techniques that specifically address the above-mentioned challenges in order to meet the requirements of emerging massively-parallel, many-core GPU architectures. These hashing techniques can be broadly categorized into four groups: open-addressing, perfect hashing, spatial hashing, and separate chaining. Each technique is distinguished by the manner in which it resolves collisions during the hashing procedure.

The remainder of this survey is organized as follows. Section 2 reviews the necessary background material to motivate GPU-based data-parallel hashing. Section 3 surveys the four categories of hashing techniques in detail, with some categories consisting of multiple sub-techniques. Section 4 categorizes and summarizes real-world applications of these hashing techniques at a high-level. Section 5 synthesizes and presents the findings of this survey in terms

of best practices and opportunities for further research. Section 6 concludes the work.

2. Background

The following section reviews concepts that are related to GPU-based data-parallel hashing.

2.1. Scalable Parallelism

Lampert [Lam78] defines *concurrency* as the decomposition of a process into independently-executing events (subprograms or instructions) that do not causally affect each other. *Parallelism* occurs when these events are all executed at the same time and perform roughly the same work. According to Amdahl [Amd67], a program contains both non-parallelizable, or serial, work and parallelizable work. Given P processors (e.g., hardware cores or threads) available to perform parallelizable work, *Amdahl's Law* defines the *speedup* S_P of a program as $S_P \leq T_1/T_P$, where T_1 and T_P are the times to complete the program with a single processor and P processors, respectively. As $P \rightarrow \infty$, $S_\infty \leq \frac{1}{f}$, where f is the fraction of serial work in the program. So, the speedup, or *scalability*, of a program is limited by its inherent serial work, as the number of processors increases. Ideally, a linear speedup is desired, such that P processors achieve a speedup of P ; a speedup proportional to P is said to be scalable.

Often a programmer writes and executes a program without explicit design for parallelism, assuming that the underlying hardware and compiler will automatically deliver a speedup via greater processor cores and transistors, instruction pipelining, vectorization, memory caching, etc [JR15]. While these automatic improvements may benefit *perfectly parallelizable work*, they are not guaranteed to address *imperfectly parallelizable work* that contains data dependencies, synchronization, high latency cache misses, etc [MRR12]. To make this work perfectly parallelizable, the program must be refactored, or redesigned, to expose more *explicit* parallelism that can increase the speedup (S_P). Brent [Bre74] shows that this explicit parallelism should first seek to minimize the *span* of the program, which is the longest chain of tasks that must be executed sequentially in order. Defining T_1 as the total serial work and T_∞ as the span, *Brent's Lemma* relates the work and span as $T_P \leq (T_1 - T_\infty)/P + T_\infty$. This lemma reveals that the perfectly parallelizable work $T_1 - T_\infty$ is scalable with P , while the imperfectly parallelizable span takes time T_∞ regardless of P and is the limiting factor of the scalability of T_P .

A common factor affecting imperfectly parallelizable work and scalability is memory dependencies between parallel (or concurrent) tasks. For example, in a *race condition*, tasks contend for exclusive write access to a single memory location and must *synchronize* their reads to ensure correctness [MRR12]. While some dependencies can be refactored into a perfectly parallelizable form, others still require synchronization (e.g., locks and mutexes) or hardware *atomic* primitives to prevent non-deterministic output. The key to enabling scalability in this scenario is to avoid high contention at any given memory location and prevent *blocking* of tasks, whereby tasks remains idle (sometimes deadlocked) until they can access a lock resource. To enable lock-free progress of work among tasks,

fine-grained atomic primitives are commonly used to efficiently check and increment values at memory locations [Her91, DHM13]. For example, the *compare-and-swap* (CAS) primitive atomically compares the value read at a location to an expected value. If the values are equal, then a new value is set at the location; otherwise, the value doesn't change.

Moreover, programs that have a high ratio of memory accesses to arithmetic computations can incur significant *memory latency*, which is the number of clock or instruction cycles needed to complete a single memory access [PH08]. During this latency period, processors should perform a sufficient amount of parallel work to *hide* the latency and avoid being idle. Given the *bandwidth*, or instructions completed per cycle, of each processor, *Little's Law* specifies the number of parallel instructions needed to hide latency as the bandwidth multiplied by latency [Lit11]. While emerging many-core and massively-threaded architectures provide more available parallelism and higher bandwidth rates, the memory latency rate remains stagnant due to physical limitations [MRR12]. Thus, to exploit this greater throughput and *instruction-level parallelism* (ILP), a program should ideally be decomposed into fine-grained units of computation that perform parallelizable work (*fine-grained parallelism*).

Furthermore, the increase in available parallelism provided by emerging architectures also enables larger workloads and data to be processed in parallel [MRR12, JR15]. Gustafson [Gus88] noted that as a problem size grows, the amount of parallel work increases much faster than the amount of serial work. Thus, a speedup can be achieved by decreasing the serial fraction of the total work. By explicitly parallelizing fine-grained computations that operate on this data, scalable *data-parallelism* can be attained, whereby a single instruction is performed over multiple data elements (SIMD) in parallel (e.g., via a vector instruction), as opposed to over a single scalar data values (SISD). This differs from *task-parallelism*, in which multiple tasks of a program conduct multiple instructions in parallel over the same data elements (MIMD) [PH08]. Task-parallelism only permits a constant speedup and induces *coarse-grained parallelism*, whereby all tasks work in parallel but an individual task could still be executing serial work. By performing inner fine-grained parallelism within outer course-grained parallel tasks, a *nested parallelism* is attained [BS05]. Many recursive and segmented problems (e.g., quicksort and closest pair) can often be refactored into nested-parallel versions [Ble90]. Flynn [Fly72] introduces SIMD, SISD, and MIMD as part of a taxonomy of computer instruction set architectures.

2.2. General-Purpose Computing on GPU (GPGPU)

A *graphical processing unit* (GPU) is a special-purpose architecture that is designed specifically for high-throughput, data-parallel computations that possess a high arithmetic intensity—the ratio of arithmetic operations to memory operations [PH08]. Traditionally used and hard-wired for accelerating computer graphics and image processing calculations, modern GPUs contain many times more execution cores and available instruction-level parallelism (ILP) than a CPU of comparable size [Nvi17b]. This inherent ILP is provided by a group of processors, each of which performs SIMD-like instructions over thousands of independent, parallel threads.

These *stream processors* operate on sets of data, or *streams*, that require similar computation and exhibit the following characteristics [KRD*03]:

- *High Arithmetic Intensity*: High number of arithmetic instructions per memory instruction. The stream processing should be largely compute-bound as opposed to memory bandwidth-bound.
- *High Data-Parallelism*: At each time step, a single instruction can be applied to a large number of streams, and each stream is not dependent on the results of other streams.
- *High Locality of Reference*: As many streams as possible in a set should align their memory accesses to the same segment of memory, minimizing the number of memory transactions to service the streams.

General-purpose GPU (GPGPU) computing leverages the massively-parallel hardware capabilities of the GPU for solving general-purpose problems that are traditionally computed on the CPU (i.e., non-graphics-related calculations). These problems should feature large data sets that can be processed in parallel and satisfy the characteristics of stream processing outlined above. Accordingly, algorithms for solving these problems should be re-designed and optimized for the data-parallel GPU architecture, which has significantly different hardware features and performance goals than a modern CPU architecture [Nvi17a].

Modern GPGPUs with dedicated memory are most-commonly packaged as discrete, programmable devices that can be added onto the motherboard of a compute system and programmed to configure and execute parallel functions [PH08]. The primary market leaders in the design of discrete GPGPUs are Nvidia and Advanced Micro Devices (AMD), with their GeForce and Radeon family of generational devices, respectively. Developed by Nvidia, the CUDA parallel programming library provides an interface to design algorithms for execution on an Nvidia GPU and configure hardware elements [Nvi17b]. For the remainder of this survey, all references to a GPU will be with respect to a modern Nvidia CUDA-enabled GPU, as it is used prevalently in most of the GPU hashing studies.

The following subsections review important features of the GPU architecture and discuss criteria for optimal GPU performance.

2.2.1. SIMT Architecture

A GPU is designed specifically for *Single-Instruction, Multiple Threads* (SIMT) execution, which is a combination of SIMD and simultaneous multi-threading (SMT) execution that was introduced by Nvidia in 2006 as part of the Tesla micro-architecture [Nvi17c]. On the host CPU, a program, or *kernel* function, is written in CUDA C and invoked for execution on the GPU. The kernel is executed N times in parallel by N different CUDA threads, which are dispatched as equally-sized *thread blocks*. The total number of threads is equal to the number of thread blocks times the number of threads per block, both of which are user-defined in the kernel. Thread blocks are required to be independent and can be scheduled in any order to be executed in parallel on one of several independent *streaming multi-processors* (SMs). The number of blocks is typically based on the number of data elements being processed by the kernel or the number of available SMs [Nvi17b]. Since each SM

has limited memory resources available for resident thread blocks, there is a limit to the number of threads per block—typically 1024 threads. Given these memory constraints, all SMs may be occupied at once and some thread blocks will be left inactive. As thread blocks terminate, a dedicated GPU scheduling unit launches new thread blocks onto the vacant SMs.

Each SM chip contains hundreds of ALU (arithmetic logic unit) and SFU (special function unit) compute cores and an interconnection network that provides k -way access to any of the k partitions of off-chip, high-bandwidth global DRAM memory. Memory requests first query a global L2 cache and then only proceed to global memory upon a cache miss. Additionally, a read-only texture memory space is provided to cache global memory data and enable fast loads. On-chip thread management and scheduling units pack each thread block on the SM into one or more smaller logical processing groups known as *warps*—typically 32 threads per warp; these warps compose a *cooperative thread array* (CTA). The thread manager ensures that each CTA is allocated sufficient shared memory space and per-thread registers (user-specified in kernel program). This on-chip shared memory is designed to be low-latency near the compute cores and can be programmed to serve as L1 cache or different ratios thereof (newer generations now include these as separate memory spaces) [Har14].

Finally, each time an instruction is issued, the SM instruction scheduler selects a warp that is ready to execute the next SIMT scalar (register-based) instruction, which is executed independently and in parallel by each active thread in the warp. In particular, the scheduler applies an active mask to the warp to ensure that only active threads issue the instruction; individual threads in a warp may be inactive due to independent branching in the program. A synchronization barrier detects when all threads (and warps) of a CTA have exited and then frees the warp resources and informs the scheduler that these warps are now ready to process new instructions, much like context switching on the CPU. Unlike a CPU, the SM does not perform any branch prediction or speculative execution (e.g., prefetching memory) among warp threads [PH08].

SIMT execution is similar to SIMD, but differs in that SIMT applies one instruction to multiple independent warp threads in parallel, instead of to multiple data lanes. In SIMT, scalar instructions control individual threads, whereas in SIMD, vector instructions control the entire set of data lanes. This detachment from the vector-based processing enables threads of a warp to conduct a form of SMT execution, where each thread behaves more like a heavier-weight CPU thread [PH08]. Each thread has its own set of registers, addressable memory requests, and control flow. Warp threads may take divergent paths to complete an instruction (e.g., via conditional statements) and contribute to starvation as faster-completing threads wait for the slower threads to finish.

The two-level GPU hierarchy of warps within SMs offers massive nested parallelism over data [PH08]. At the outer, SM level of granularity, coarse-grained parallelism is attained by distributing thread blocks onto independent, parallel SMs for execution. Then at the inner, warp level of granularity, fine-grained data and thread parallelism is achieved via the SIMT execution of an instruction among parallel warp threads, each of which operates on an individual data element. The massive data-parallelism and avail-

able compute cores are provided specifically for high-throughput, arithmetically-intense tasks with large amounts of data to be independently processed. If a high-latency memory load is made, then it is expected that the remaining warps and processors will simultaneously perform sufficient work to hide this latency; otherwise, hardware resources remain unused and yield a lower aggregate throughput [Vol16]. The GPU design trades-off lower memory latency and larger cache sizes (such as on a CPU) for increased instruction throughput via the massive parallel multi-threading [PH08].

This architecture description is based on the Nvidia Maxwell micro-architecture, which was released in 2015 [Har14]. While certain quantities of components (e.g., SMs, compute cores, memory sizes, and thread block sizes) change with each new generational release of the Nvidia GPU, the general architectural design and execution model remain constant [Nvi17b]. The CUDA C Programming Guide [Nvi17b] and Nvidia PTX ISA documentation [Nvi17c] contain further details on the GPU architecture, execution and memory models, and CUDA programming.

2.2.2. Optimal Performance Criteria

The following performance strategies are critical for maximizing utilization, memory throughput, and instruction throughput on the GPU [Nvi17a].

Sufficient parallelism: Sufficient instruction-level and thread-level parallelism should be attained to fully hide arithmetic and memory latencies. According to Little's Law, the number of parallel instructions needed to hide a latency (number of cycles needed to perform an instruction) is roughly the latency times the throughput (number of instructions performed per cycle) [Lit11]. During this latency period, threads that are dependent on the output data of other currently-executing threads in a warp (or thread block) are stalled. Thus, this latency can be hidden either by having these threads simultaneously perform additional, non-dependent SIMD instructions in parallel (instruction-level parallelism), or by increasing the number of concurrently running warps and warp threads (thread-level parallelism) [Vol16].

Since each SM has limited memory resources for threads, the number of concurrent warps possible on an SM is a function of several configurable components: allocated shared memory, number of registers per thread, and number of threads per thread block [Nvi17b]. Based on these parameters, the number of parallel thread blocks and warps on an SM can be calculated and used to compute the *occupancy*, or ratio of the number of active warps to the maximum number of warps. In terms of Little's Law, sufficient parallel work can be exploited with either a high occupancy or low occupancy, depending on the amount of work per thread. Based on the specific demands for SM resources, such as shared memory or register usage, by the kernel program, the number of available warps will vary accordingly. Higher occupancy, usually past 50 percent, does not always translate into improved performance [Nvi17a]. For example, a lower occupancy kernel will have more registers available per thread than a higher occupancy kernel, allowing low-latency access to local variables and minimizing register spilling into high-latency local memory.

Memory coalescing: When a warp executes an instruction that

accesses global memory, it *coalesces* the memory accesses of the threads within the warp into one or more memory transactions, or cache lines, depending on the size of the word accessed by each thread and the spatial coherency of the requested memory addresses. To minimize transactions and maximize memory throughput, threads within a warp should coherently access memory addresses that fit within the same cache line or transaction. Otherwise, memory divergence occurs and multiple lines of memory are fetched, each containing many unused words. In the worst case alignment, each of the 32 warp threads accesses successive memory addresses that are multiples of the cache line size, prompting 32 successive load transactions [Nvi17a].

The shared memory available to each thread block can help coalesce or eliminate redundant accesses to global memory [PH08]. The threads of the block (and associated warp) can share their data and coordinate memory accesses to save significant global memory bandwidth. However, it also can act as a constraint on SM occupancy—particularly limiting the number of available registers per thread and warps—and is prone to *bank conflicts*, which occur when two or more threads in a warp access an address in the same bank, or partition, of shared memory [Nvi17b]. Since an SM only contains one hardware bus to each bank, multiple requests to a bank must be serialized. Thus, optimal use of shared memory necessitates that warp threads arrange their accesses to different banks [Nvi17b]. Finally, the read-only texture memory of an SM can be used by a warp to perform fast, non-coalesced lookups of cached global memory, usually in smaller transaction widths.

Control flow: Control flow instructions (e.g., if, switch, do, for, while) can significantly affect instruction throughput by causing threads of the same warp to diverge and follow different execution paths, or branches. Optimal control flow is realized when all the threads within a warp follow the same execution path [Nvi17a]. This scenario enables SIMD-like processing, whereby all threads complete an instruction simultaneously in lock-step. During *branch divergence* in a warp, the different executions paths, or branches, must be serialized, increasing the total number of instructions executed for the warp. Additionally, the use of atomics and synchronization primitives can also require additional serialized instructions and thread starvation within a warp, particularly during high contention for updating a particular memory location [SO11].

2.3. Data Parallel Primitives

The redesign of serial algorithms for scalable data-parallelism offers platform portability, as increases in processing units and data are accompanied by unrestricted increases in speedup. *Data-parallel primitives* (DPPs) provide a way to explicitly design and program an algorithm for this scalable, platform-portable data-parallelism. DPPs are highly-optimized *building blocks* that are combined together to compose a larger algorithm. The traditional design of this algorithm is thus refactored in terms of DPPs. By providing highly-optimized implementations of each DPP for each platform architecture, an algorithm composed of DPPs can be executed efficiently across multiple platforms. This use of DPPs eliminates the combinatorial (cross-product) programming issue of having to implement a different version of the algorithm for each different architecture.

The early work on DPPs was set forth by Blelloch [Ble90], who proposed a *scan vector model* for parallel computing. In this model, a vector-RAM (V-RAM) machine architecture is composed of a vector memory and a parallel vector processor. The processor executes vector instructions, or *primitives*, that operate on one or more arbitrarily-long vectors of atomic data elements, which are stored in the vector memory. This is equivalent to having as many independent, parallel processors as there are data elements to be processed. Each primitive is classified as either *scan* or *segmented* (per-segment parallel instruction), and must possess a parallel, or *step*, time complexity of $O(\log n)$ and a serial, or *element*, time complexity of $O(n)$, in terms of n data elements; the element complexity is the time needed to simulate the primitive on a serial random access machine (RAM). Several canonical primitives are then introduced and used as building blocks to refactor a variety of data structures and algorithms into data-parallel forms.

The following are examples of DPPs that are commonly-used as building blocks to construct data-parallel algorithms:

- *Map*: Applies an operation on all elements of the input array, storing the result in an output array of the same size, at the same index;
- *Reduce*: Applies an aggregate binary operation (e.g., summation or maximum) on all elements of an input array, yielding a single output value. *ReduceByKey* is a variation that performs segmented *Reduce* on the input array based on unique key, yielding an output value for each key;
- *Gather*: Given an input array of values, reads values into an output array according to an array of indices;
- *Scan*: Calculates partial aggregates, or a prefix sum, for all values in an input array and stores them in an output array of the same size;
- *Scatter*: Writes each value of an input data array into an index in an output array, as specified in the array of indices;
- *Compact*: Applies a unary predicate (e.g., if an input element is greater than zero) on all values in an input array, filtering out all the values which do not satisfy the predicate. Only the remaining elements are copied into an output array of an equal or smaller size;
- *SortByKey*: conducts an in-place segmented *Sort* on the input array, with segments based on a key or unique data value in the input array;
- *Unique*: Ignores duplicate values which are adjacent to each other, copying only unique values from the input array to the output array of the same or lesser size; and
- *Zip*: Binds two arrays of the same size into an output array of pairs, with the first and second components of a pair equal to array values at a given index.

Several other DPPs exist, each meeting the required step and element complexities specified by Blelloch [Ble90]. Cross-platform implementations of a wide variety of DPPs form the basis of several notable open-source libraries.

The Many-Core Visualization Toolkit (VTK-m) [MSU*16] is a platform-portable library that provides a growing set of DPPs and DPP-based algorithms [vtk17]. With a single code base, back-end code generation and runtime support are provided for use on GPUs

and CPUs. Currently, each GPU-based DPP is a modified variant from the Nvidia CUDA Thrust library of parallel algorithms and data structures [Nvi17d], and each CPU-based DPP is adopted from the Intel Thread Building Blocks (TBB) library for scalable data parallel programming [Int17]. VTK-m provides the flexibility to develop custom *device adapter algorithms*, or DPPs, for a new device type. This device can take the form of an emerging architecture or a new parallel programming language (e.g., Thrust and TBB) for which DPPs must be re-optimized. Thus, a DPP can be invoked in the high-level VTK-m user code and executed on any of the devices at runtime. The choice of device is either specified at compile-time by the user, or automatically selected by VTK-m. VTK-m, Thrust, and TBB all employ a generic programming model that provides C++ Standard Template Library (STL)-like interfaces to DPPs and algorithms [PLMS00]. Templated arrays form the primitive data structures over which elements are parallelized and operated on by DPPs. Many of these array types provide additional functionality on top of underlying vector iterators that are inspired by those in the Boost Iterator Library [Boo03].

The CUDA Data Parallel Primitives Library (CUDPP) [cud17] is a library of fundamental DPPs and algorithms written in Nvidia CUDA C [Nvi17b] and designed for high-performance execution on CUDA-compatible GPUs. Each DPP and algorithm incorporated into the library is considered best-in-class and typically published in peer-reviewed literature (e.g., radix sort [MG10, ADMO16], mergesort [SHG09, DTGO12], and cuckoo hashing [ASA*09, AVS*12]). Thus, its data-parallel implementations are constantly updated to reflect the state-of-the-art.

2.4. GPU Searching

The following section reviews canonical approaches for organizing, storing, and searching data on the GPU.

Let $U = \{i\}_{0 \leq i < u}$ be the universe for some arbitrary positive integer u . Then let $S \subset U$ be an unordered set of $n = |S|$ elements, or *keys*, belonging to U . The *search problem* seeks an answer to the *query*: “Is key k a member of S ?” If $k \in S$, then we return its corresponding *value*, which is either k itself or a different value. A *data structure* is built or constructed over S to efficiently facilitate the searching operation. The data structure is implementation-specific and can be as simple as a sorted (ordered) variant of the original set, a hash table, or a tree-based partitioning of the elements.

A generalization of the search task is the *dictionary problem*, which seeks to both modify and query key-value pairs (k, v) in S . A canonical dictionary data structure supports *insert* (k, v) , *delete* (k, v) , *query* (k) , *range* (k_1, k_2) (returns $\{k | k_1 \leq k \leq k_2\}$), and *count* (k_1, k_2) (returns $|\text{range}(k_1, k_2)|$). To support these operations, the dictionary must be dynamic and accommodate incremental or batch updates after construction; this contrasts to a static data structure, which either does not support updates after a one-time build or must be rebuilt after each update. In multi-threaded environments, these structures must also provide concurrency and ensure correctness among mixed, parallel operations that may access the same elements simultaneously.

An extensive body of work has embarked on the redesign of data structures for construction and general computation on the

GPU [OLG*07]. Within the context of searching, these *acceleration* structures include sorted arrays [ASA*09, SGL09, AVS*12, KDB12, LBMC16, LMLC17, ALF*17] and linked lists [YHGT10], hash tables (see section 3), spatial-partitioning trees (e.g., k -d trees [ZHWG08, Kar12, WBS*14], octrees [ZGHG11, Kar12], bounding volume hierarchies (BVH) [LGS*09, Kar12], R-trees [LWL12], and binary indexing trees [KCS*10, SR17]), spatial-partitioning grids (e.g., uniform [LD08, KS09, GTW15] and two-level [KBS11]), skiplists [MCP17], and queues (e.g., binary heap priority [HAP12] and FIFO [CCT12, SF15]). Due to significant architectural differences between the CPU and GPU, search structures cannot simply be “ported” from the CPU to the GPU and maintain optimal performance. On the CPU, these structures can be designed to fit within larger cache, perform recursion, and employ heavier-weight synchronization or hardware atomics. However, during queries, the occurrence of varying paths of pointers (pointer chasing) and dependencies between different phases or levels of the structure both limit the parallel throughput on the GPU. Moreover, these structures ideally should be constructed directly on the GPU, as transfers from the CPU over the PCIe bus induce costly latencies.

For searching an unordered array of elements on the GPU, two canonical data structures exist: the sorted array and the hash table. Both of these data structures are known to be relatively fast to construct on the GPU and are amenable to data-parallel design patterns [ALF*17].

2.4.1. Searching Via Sorting

Given a set of n unordered elements, a canonical searching approach is to first sort the elements in ascending order and then conduct a binary or k -nary search for the query element. This search requires a logarithmic number of comparisons in the worst-case, but is not as amenable to caching as consecutive comparisons are not spatially close in memory for large n . Moreover, on the GPU, an ordered query pattern by threads in a warp can enable memory coalescing during comparisons.

The current version of the CUDA Thrust library [Nvi17d] provides fast and high-throughput data-parallel implementations of mergesort [SHG09] and radix sort [MG10] for arrays of custom (e.g., comparator function) or numerical (i.e., integers and floats) data types, respectively. Similarly, the latest version of the CUDPP library [cud17] includes best-in-class data-parallel algorithms for mergesort [SHG09, DTGO12] and radix sort [MG10, ADMO16], each of which are adapted from published work. Singh et al. [SJC17] survey and compare the large body of recent GPU-based sorting techniques.

A few studies have investigated various factors that affect the performance of data-parallel sort methods within the context of searching [ASA*09, AVS*12, LMLC17]. Kaldewey and Blas introduce a GPU-based p -ary search that first uses p parallel threads to locate a query key within one of p larger segments of a sorted array, and then iteratively repeats the procedure over p smaller segments within the larger segment. This search achieves high memory throughput and is amenable to memory coalescing among the threads [KDB12]. Moreover, the algorithm was also ported to the CPU to leverage the SIMD vector instructions in a fashion similar

to the k -ary search introduced by Schlegel et al. [SGL09]. However, the fixed vector width restricts the degree of parallelism and value of p , which is significantly higher on the GPU.

Inserting or deleting elements into a sorted array is generally not supported and requires inefficient approaches such as appending/removing new elements and re-sorting the larger/smaller array, or first sorting the batch of new insertions and then merging them into the existing sorted array. Ashkiani et al. [ALF*17] present these approaches and the resulting performance for a dynamic sorting-based dictionary data structure, along with setting forth the current challenges of designing dynamic data structures on the GPU.

2.4.2. Searching Via Hashing

Instead of maintaining elements in sorted order and performing a logarithmic number of lookups per query, *hash tables* compactly reorganize the elements such that only a constant number of direct, random-access lookups are needed on average [CSRL01]. More formally, given a universe U of possible keys and an unordered set $S \subseteq U$ of n keys (not necessarily distinct), a *hash function*, $h : U \mapsto H$, maps the keys from S to the range $H = \{j\}_{0 \leq j < m}$ for some arbitrary positive integer $m \geq n$. Defining a memory space over this range of size m specifies a hash table, into which keys are inserted and queried. Thus, the hash table is addressable by the hash function. During an insertion or query operation for a key q , the hash function computes an address $h(q) = r$ into H . If the location $H[r]$ is empty, then q is either inserted into $H[r]$ (for an insertion) or does not exist in H (for a query). If $H[r]$ contains the key q (for a query), then either q or an associated *value* of q is returned[†], indicating success. Otherwise, if multiple distinct keys $q' \neq q$ are hashed to the same address $h(q') = r$, then a situation known as a *hash collision* occurs. These collisions are typically resolved via *separate chaining* (i.e., employing linked lists to store multiple keys at a single address) or *open-addressing* (e.g., when an address is occupied, then store the key at the next empty address).

The occurrence of collisions deteriorates the query performance, as each of the collided keys must be iteratively inspected and compared against the query key. According to the birthday paradox, with a discrete uniform distribution hash function that outputs a value between 1 and 365 for any key, the probability that two random keys hash to the same address in a hash table of size 23 is 50 percent [STKT06]. More generally, for n hash values and a table size of m , the probability $p(n, m)$ of a collision is

$$p(n, m) = \begin{cases} 1 - \prod_{k=1}^{n-1} \left(1 - \frac{k}{m}\right) & n \leq m \\ 1 & n > m \end{cases}$$

$$\approx 1 - \left(\frac{m-1}{m}\right)^{\frac{n(n-1)}{2}}.$$

Thus, for a large number of keys (n) and small hash table (m), hash collisions are inevitable.

[†] In practice, the values should be easily stored and accessible within an auxiliary array or via a custom arrangement within the hash table.

In order to minimize collisions, an initial approach is to use a good quality hash function that is both efficient to compute and distributes keys as evenly as possible throughout the hash table [CSRL01]. One such family of functions are randomly-generated, parameterized functions of the form $h(k) = (a \cdot k + b) \bmod p \bmod |H|$, where p is a large prime number and a and b are randomly-generated constants that bias h from outputting duplicate values [AVS*12]. However, h is a function of the table size, $|H|$. If $|H|$ is too small, then not even the best of hash functions can avoid an increase in collisions. Given the table size, the *load factor* α of the table is defined as $\alpha = n/|H|$, or the percentage of occupied addresses in the hash table, which $|H|$ is typically larger than n . If new keys are inserted into the table and α reaches a maximum threshold, then typically the table is allocated to a larger size and all the keys are *rehashed* into the table.

To avoid collision resolution altogether, a *perfect* hash function can be constructed to hash keys into a hash table without collisions. Each key is mapped to a distinct address in the table. However, composing such a perfect hashing scheme is known to be difficult in general [LH06]. The probability of attaining a perfect hash for n keys in a large table of size m ($m \gg n$) is defined as

$$p(n, m) = (1 - \frac{1}{m}) \cdot (1 - \frac{2}{m}) \cdots (1 - \frac{n-1}{m}) \approx e^{-\frac{n^2}{2m}},$$

which is very small for a large n or small m . Nonetheless, a significant body of research has investigated this approach and is reviewed in this survey, within the context of parallel hashing.

A hash table is *static* if it does not support modification after being constructed; that is, the table is only constructed to handle query operations. Thus, a static hash table also does not support mixed operations and the initial batch of insertions used to construct the table (bulk build) must be completed before the batch of query operations. A hash table that can be updated, or mutated, via insertion and deletion operations post-construction is considered *dynamic*. Denoting the query, insert, and delete operations as q , i , and d , respectively, the *operation distribution* $\Gamma = (q, i, d)$, $q + i + d = 1$ specifies the percentage of each operation that are conducted concurrently in a hashing workload [AFO17]. For example, $\Gamma = (0.7, 0.15, 0.15)$ represents a query-heavy workload that performs 70% queries and 30% updates. Additionally, the percentage q can be split into queried keys that exist in the hash table and those that do not. Often, queries for non-existent keys can present worst-case scenarios for many hash techniques, as a maximum number of searches are conducted until failure [AFO17].

As general data structures, hash tables do not place any special emphasis on the key access patterns over time [CPK16]. However, the patterns that appear in various real-world applications do possess observable structure. For example, geometric tasks may query spatially-close keys in a sequential or coherent pattern, and database tasks may query certain subsets of keys more frequently than others, whereby the hash table serves as a working set or most-recently-used (MRU) table for cache-like accesses [AVS*12, PR13, CPK16]. Moreover, dynamic hash tables do not place special emphasis on the mixture Γ , or pattern, of query and update operations. However, execution time per-

formance may be better or worse for some hashing techniques, depending on the specific Γ , such as query-heavy for key-value stores [ZWY*15] or update-heavy for real-time, interactive spatial hash tables [LH06, ASA*09, GLHL11, NZIS13].

Finally, hash tables offer compact storage for sparse spatial data that contains repeated elements or empty elements that don't need to be computed. For example, instead of storing an entire, mostly-empty voxelized 3D grid, the non-empty voxels can be hashed into a dense hash table [LH06]. Then, every voxel can be queried to determine whether it should be rendered or not, returning a negative result for the empty voxels. Furthermore, a hash table does not have to be one-dimensional. Instead, the data structure can consist of multiple hash tables or *bucketed* partitions that are each addressed by a different hash function.

While collision resolution is straightforward to implement in a serial CPU setting, it does not easily translate to a parallel setting, particularly on massively-threaded, data-parallel GPU architectures. GPU-based hashing[‡] presents several notable challenges:

- Hashing is a memory-bound problem that is not as amenable to the compute-bound and limited-caching design of the GPU, which hides memory latencies via a large arithmetic throughput.
- The random-access nature of hashing can lead to disparate writes and reads by parallel-cooperating threads on the GPU, which performs best when memory accesses are coalesced or spatially coherent.
- The limited memory available on a GPU puts restrictions on the maximum hash table size and number of tables that can reside on device.
- Collision resolution schemes handle varying numbers of keys that are hashed and chained to the same address (separate chaining), or varying numbers of attempts to place a new, collided key into an empty table location (open-addressing). This variance causes some insert and query operations to require more work than others. On a GPU, threads work in groups to execute the same operation on keys in a data-parallel fashion. Thus, a performance bottleneck arises when faster, non-colliding threads wait for slower, colliding threads to finish. Moreover, some threads may insert colliding keys that are unable to find an empty table location, leading to failure during construction of the table.

Searching via the construction and usage of a hash table on the GPU has recently received a breadth of new research, with a variety of different parallel designs and applications, ranging from collision detection to surface rendering to nearest neighbor approximation. The following section covers these GPU-based parallel hashing approaches.

3. Hashing Techniques

We consider four different categories of hashing techniques: open-addressing probing, perfect hashing, spatial hashing, and separate

[‡] In this study, the term *hashing* refers to the entire process from constructing the hash table to the handling of collisions while querying or updating; this is not to be confused with hash function design and computation, or its application to cryptographic protocols and message passing.

chaining. Each category is discussed in a separate subsection and distinguished by its method of handling hash collisions or placement of elements within the hash table.

3.1. Open-addressing Probing

In *open-addressing*, a key is inserted into the hash table by *probing*, or searching, through alternate table locations—the *probe sequence*—until a location is found to place the element [CSRL01]. The determination of where to place the element varies by probing scheme: some schemes probe for the first unused location (*empty slot*), whereas others *evict* the currently-residing key at the probe location (i.e., a collision) and swap in the new key. Each probe location is specified by a hash function unique to the probing scheme. Thus, some probe sequences may be more compact or greater in length than others, depending on the probing method. For a query operation, the locations of the probe sequence are computed and followed to search for the queried key in the table.

Each probing method trades-off different measures of performance with respect to GPU-based hashing. A critical influence on performance is the load factor, which is the percentage of occupied locations in the hash table (subsection 2.4.2). As the load factor increases towards 100 percent, the number of probes needed to insert or query a key increases greatly. Once the table becomes full, probing sequences may continue indefinitely, unless bounded, and lead to insertion failure and possibly a hashing *restart*, whereby the hash table is reconstructed with different hash functions and parameters. Moreover, for threads within a warp on the GPU, variability in the number of probes per thread can induce branch divergence and inefficient SIMD parallelism, as all the threads will need to wait for the worst-case number of probes to execute the next instruction.

The following subsections review research on open-addressing probing for GPU-based hashing, distinguishing each study by its general probing scheme: linear probing, cuckoo hashing, double hashing, multi-level or bucketized probing, and robin hood hashing.

3.1.1. Linear Probing-based Hashing

Linear probing is the most basic method of open-addressing. In this method, a key k first hashes to location $h(k)$ in the hash table. Then, if the location is already occupied, k linearly searches locations $h(k) + 1, h(k) + 2, \dots$ etc. until an empty slot (insertion) or k itself (query) is found. If $h(k)$ is empty, then k is inserted immediately, without probing; otherwise, a worst-case $O(n)$ probes will need to be made to locate k or an empty slot, where n is the size of the hash table. While simple in design, linear probing suffers from *primary clustering*, whereby a cluster, or contiguous block, of locations following $h(k)$ are occupied by keys, reducing nearby empty slots. This occurs because colliding keys at $h(k)$ each successively probe for the next available empty slot after $h(k)$ and insert themselves into it. An improved variant of linear probing is *quadratic probing*, which replaces the linear probe sequence starting at $h(k)$ with successive values of an arbitrary quadratic polynomial: $h(k) + 1^2, h(k) + 2^2, \dots$ etc. This avoids primary clustering, but also introduces a *secondary clustering* effect as a result. For a more than half-full table, both of these probing methods can incur a long probe sequence to find an empty slot, possibly resulting in failure during an insert.

Bordawekar [Bor14] develops an open-addressing approach based on multi-level bounded linear probing, where the hash table has multiple levels to reduce the number of lookups during linear probing. In the first level hash table, each key hashes to a location $h_1(k)$ and then looks for an empty location, via linear probing, within a bounded probe region $P_1 = [h_1(k), h_1(k) + (j - 1)]$, where j is the size of the region. If an empty location is not found, then the key must be inserted into the second-level hash table, which is accomplished by hashing to location $h_2(k)$ and linear probing within another, yet larger, probe region P_2 . This procedure continues for each level, until an empty location is found. In this work, only 2-level and 3-level hash tables are considered; thus, a thread must perform bounded probing on a key for at most three rounds, before declaring failure. To query a key, a thread completes the same hashing and probing procedure. In a data-parallel fashion, each thread within a warp is assigned a key from the bounded probe region and compares this key with the query key, using warp-level voting to communicate success or failure. This continues across warps, for each hash table level.

The initial design goal of this multi-level approach was to bound and reduce the average number of probes per insertion and query, while enabling memory coalescing and cache line coherency among threads (or lanes) within a warp. By using a small, constant number of hash tables and functions, the load factor could be increased beyond the 70 percent of Alcantara et al.'s cuckoo hashing (subsection 3.1.2), without sacrificing performance. However, experimental results reveal that this approach, with both two and three levels (and hash functions), does not perform as fast as cuckoo hashing for the largest batches of key-value pairs (hundreds of millions); for smaller batches, the multi-level approaches are the best performers. This finding is particularly noticeable for querying the keys, suggesting that improved probing and memory coalescing are likely not achieved. Additional details are needed to ascertain whether the ordering of the keys—spatial or random—affect this multi-level approach, or specific reasons why the expected warp-level memory coalescing is not being realized.

Karnagel et al. [KML15] develop a linear probing hashing scheme to perform database group-by and aggregation queries (i.e., a reduce-by-key operation) on the GPU. In this work, a database of records (e.g., customer data) is stored in SSDs and then queried in SQL format from the CPU. Selected columns of the query (e.g., zip code and order total) are transferred and pinned into host memory, from which the GPU fetches the column values in a coalesced, data-parallel fashion via Universal Virtual Addressing (UVA). Then, a hashing procedure begins to compute an aggregate (e.g. average discount) for each unique item, or *group*, in the group-by column (e.g., customer zip code). Each item in this column is initially inserted into a global memory hash table as a key-value pair, where the key is the item ID and the value is a tuple (count, sum). For each item, the ID key is hashed to a location and then probes linearly until its matching key is found. If an empty slot is encountered, then a new value tuple is inserted for the key; otherwise, the count and sum of the tuple are atomically added to the current value tuple for the key. After all threads have completed, the count and sum of each key in the hash table are used to calculate aggregate values. These aggregate values form the output of the query. All GPU

hash table operations and arithmetic computations are performed in data-parallel fashion by blocks of threads.

The primary contribution of this work is a thorough experimental evaluation of several factors affecting hashing performance on the GPU, including guidance on how to decide the hash table size, load factor, and CUDA grid configuration of number of thread blocks and threads per block. Notable experimental findings are the following:

- As the number of groups increases, either the hash table must grow proportionately in size, or the load factor needs to increase. The ideal table size is one that fits within shared L2 cache, with a load factor below 50 percent. If a higher load factor is used, then thread contention and long probe sequences emerge. This contention is due to multiple threads attempting to access the value tuple for the same key (group) and having to synchronize via atomic compare-and-swap updates.
- Hash tables that cannot fit within L2 cache reside in global memory and each thread must make global memory accesses, unless the data is cached. For the data to reside in cache, a cache line load of a small, fixed size (e.g., 128 bytes for L1 and 32 bytes for L2) must be performed. Since linear probing creates a variable number of probes per thread and threads access random, non-coalesced regions of memory, a thread will likely need multiple cache line loads to complete an operation. With thousands of concurrent threads, each invoking cache line loads into a limited size cache, the chances of cache pollution and eviction are high, diminishing the benefits of caching. Thus, to achieve undisturbed cache access to data, fewer threads should be used. The number of threads, however, should also be at least enough to hide the memory latency of the PCIe data transfers from host to device.

Additionally, the simple hash table and probing scheme are only used in order to minimize the number of factors affecting performance and because the approach is mainly PCIe-bandwidth bound, which affords more probes and non-coalesced memory accesses to hide the latency. The authors acknowledge the bounded linear probing approach of Bordawekar [Bor14], but cite the latter reason for using a simpler hashing scheme.

3.1.2. Cuckoo-based Hashing

In *cuckoo hashing*, each key is assigned two locations in the hash table, as specified by primary and secondary hash functions [PR04]. When inserting a new key, its first location is probed with the primary function and the contents of the location are inspected. If the slot is empty, then the key is inserted and the probe sequence ends. Otherwise, a collided key already occupies the slot and the cuckoo eviction procedure begins. First, the occupying key is evicted and hashed to the location specified by its secondary function, where its contents are probed as before. This *eviction chain* continues until either the evicted key is successfully inserted or a maximum chain length is reached. If the eviction is successful, then the new key is finally inserted at its primary location (first probe). Numerous follow-up studies to this canonical approach have introduced cuckoo hashing approaches with more than two hash functions (probes) per key, a separate hash table for each hash function, and other optimizations for concurrent, mixed operations (e.g., simulta-

neous inserts and queries) on the GPU. These studies are surveyed in this subsection.

Alcantara et al. [ASA*09] introduce a data-parallel, dynamic hashing technique based on perfect hashing and cuckoo hashing that supports both hash table construction and querying at real-time, interactive rates. The querying performance of this technique is compared against that of the perfect hashing technique of Lefebvre and Hoppe [LH06] and a standard data-parallel sort plus binary search approach. In this work, a two-phase hashing routine is conducted to insert and query elements, with the goal of maximizing shared-memory usage during cuckoo hashing.

First, elements are hashed into bucket regions within the hash table, following the perfect hashing approach of Fredman et al. [FKS84]. The maximum occupancy of each bucket is the number of threads in a thread block (e.g., 512), such that the entire bucket can fit within shared memory. The hash function aims to coherently map elements into buckets such that:

- Each bucket, on average, maintains a load factor of 80%, and
- Spatially-nearby elements are located within the same bucket, enabling coalescing of memory among threads during queries.

If more than 512 elements hash to a given bucket, then a new hash function is generated and this phase is repeated. Then, within each bucket, cuckoo hashing is performed to insert or query an element, using $i = 3$ different hash functions h_i (i.e., the multiple choices), each corresponding to a sub-table T_i . During construction, each element simultaneously hashes to its location h_1 in T_1 , in a winner-takes-all fashion. If multiple threads hash to the same location, then the winning thread (i.e., the last thread to write) remains and the other threads proceed to hash into location h_2 in T_2 . This continues for T_3 , after which any remaining unplaced elements cycle back to the beginning and hash into h_1 in T_1 again. At this point, if a collision occurs at h_1 , then the current residing element is evicted and added to the batch of unplaced elements. This cuckoo hashing procedure continues until all elements are successfully placed into a sub-table T_i or a maximum number of cycles have occurred.

An observation of this construction routine is that restarts can occur during both phases if either a bucket overflows or the cuckoo hashing reaches the maximum number of cycles within a bucket. While this reconstruction may be viewed as a disadvantage of probing techniques in general, the authors maintain that the occurrence of these restarts are reasonable in practice and fast to compute on massively-parallel GPU architectures. Moreover, this technique makes extensive use of thread atomics to increment and check values in both global and shared memory. While only a fixed number of atomic operations are made each phase, they are still serialized and must handle varying levels of thread contention, both of which are known to degrade performance.

After construction, a query operation is performed by hashing once into a bucket, and then making at most $d = 3$ hashing probes to locate the element within one of the sub-tables T_i of the bucket. Insertions and queries are all conducted in a data-parallel, SIMD fashion. Since each thread warp assigned to a bucket has its own dedicated block of shared memory, the probing and shuffling of elements in the cuckoo hashing can be performed faster locally, as opposed to accessing global memory.

Experimental results for this technique reveal the following:

- For querying elements (voxels in a 3D grid) in a randomized order, this hashing approach outperforms the perfect hashing approach of Lefebvre et al. [LH06] and the data-parallel binary search of radix-sorted elements of Satish et al. [SHG09], particularly above 5 million elements. After this point, the binary searches used in both methods do not scale and become time-prohibitive.
- For querying in a sequential order, the data-parallel binary search demonstrates better performance than this hashing technique, due to more favorable thread branch divergence and memory coalescing among the sorted elements.
- Constructing the hash table of elements in this approach is comparably-fast to radix sorting the elements, with noticeable slowdowns due to more non-coalesced write operations. Moreover, for large numbers of insertions, both approaches are magnitudes faster than constructing the perfect spatial hash table of [LH06], which is initially built on the CPU, rather than the GPU (onto which the table is copied for subsequent querying).

Alcantara et al. [AVS*12] improved upon their original work [ASA*09] by introducing a generalized parallel variant of cuckoo hashing that can vary in the number of hash functions, hash table size, and maximum length of a probe-and-eviction sequence. In [ASA*09], the authors hypothesized that parallel cuckoo hashing within GPU global memory would encounter performance bottlenecks due to the shuffling of elements each iteration and use of global synchronization primitives; thus, shared memory was used extensively in the two-level hashing scheme. However, in this follow-up work, a single-level hash table is constructed entirely in global memory and addressed directly with the cuckoo hash functions, without the first-level bucket hash. The cuckoo hashing dynamics remain the same, except that the probing and evicting of elements occurs over the entire global memory hash table, as opposed to the shared-memory buckets of the two-level approach.

To construct a hash table of N elements, approximately N threads will operate in SIMD parallel fashion to place their elements into empty slots in the global table. A given thread block will not complete until all of its threads have successfully placed their elements; a smaller block size helps minimize the completion time, as the block will likely contain fewer threads with long eviction chains.

The construction (insertion) and query performance of the single-level hash approach is compared against that of Merrill's radix sort plus binary search [MG10] and the authors' previous two-level cuckoo hashing approach. Experimental results reveal the following:

- *Insertions.* For large numbers of insertions, the radix sort [MG10] becomes increasingly faster than both hashing methods, with a much higher throughput of insertions-per-second. For the same size hash table, the single-level hash table is constructed significantly faster than the two-level table, due to shorter eviction chains on average, over all insertion input sizes (the two-level table uses a fixed $1.42N$ space, while the single-level table is variable-sized). Radix sort achieves an upper bound of 775 million memory accesses (read and write) per second, while the single-level hashing only attains 670

million accesses per second. This higher throughput by radix sort is due to its more-localized memory access patterns that enable excellent memory coalescing among threads sharing a memory-bound instruction (up to 70% of the theoretical maximum pin bandwidth on the tested Nvidia GTX 480 GPU, versus 6% of the single-level hashing).

- *Queries: Binary Search vs. Hashing.* For random, unordered queries, binary search probing of the radix sorted elements is much slower than cuckoo hash probing of the elements. This arises from uncoalesced global memory reads and branch divergence for many of the threads, which use the maximum $O(\log N)$ probes. Both cuckoo hashing approaches lookup elements in a worst-case constant number of probes and, thus, perform significantly better than binary searching, despite these probes being largely uncoalesced.
- *Queries: Two-Level vs. Single-Level.* When all queried elements exist in the hash table, the single-level cuckoo hashing makes a smaller average number of probes per query than the two-level approach, leading to faster completion times. However, when a large percentage of the queried elements do not exist in the hash table, the two-level hashing needs fewer worst-case probes before declaring the element as not found. This is because the single-level hashing uses four hash functions, or probes, to lookup an element, whereas the two-level hashing only uses three functions. By setting the number of hash functions to three in the single-level hashing, the authors observe comparable querying performance between the two approaches.

A notable performance observation in this work is that only randomized queries are considered. The authors indicate, as a limitation of their work, that if the elements to be queried are instead ordered (sorted), then binary searching the radix-sorted elements should yield improved thread branch divergence and memory coalescing. This work has since been incorporated into the CUDPP library [cud17] as a best-in-class GPU hash table data structure.

3.1.3. Multi-level and Bucketized Hashing

Bucketized cuckoo hash tables (BCHT) organize groups of entries into *buckets* (or bins), inside which cuckoo hashing is applied [EMM06]. Typically, a single allocated hash table is used and partitioned into bucket regions, each of which may be assigned to a single warp of threads. Thus, the size of each bucket is uniform and proportional to the number of threads in a warp (e.g., 32 threads), the size of the cache line in a warp (e.g., 128 bytes), or the size of the shared memory within the warp's streaming multiprocessor (e.g., less than 50 kilobytes).

As presented in section 3.1.2, the bi-level design of Alcantara et al. [ASA*09] performs bucketized cuckoo hashing by first perfect hashing into buckets that are the size of a thread block's shared memory, and then conducting cuckoo hashing within each bucket. Moreover, as presented in section 3.1.1, the work of Bordawekar [Bor14] develops a multi-level, bounded linear probing scheme. Sections 3.1.1 and 3.1.2 contain additional details on these approaches.

Zhang et al. [ZWY*15] design a modified variant of a BCHT for use in accelerating queries of a GPU-based, in-memory key-value (IMKV) store, whose values reside in host memory. Tradi-

tionally, a bucket is the size of a thread warp or block (e.g., 512 threads) and each thread is assigned to a separate insert or query operation, with varying probe sequence lengths. However, with tens of thousands of threads operating on different buckets (warps) simultaneously, L2 cache (e.g., 1.5MB) contention will be high and likely lead to frequent evictions, which will force threads to perform multiple memory transactions. This work addresses this issue by sizing a bucket to a *processing unit* of threads, which is set as a multiple of the memory transaction (L2 cache line) size (e.g., 32 bytes); the number of threads is the transaction size (in bytes) divided by 4 bytes, assuming each thread accesses a 4 byte (32-bit) key. Thus, a memory transaction services the entire processing unit, enabling coalescing among the threads. In a query operation, a key is first hashed to a bucket and given a key *signature*, after which each thread in the unit compares its assigned key signature in the bucket with the query key signature. Via a warp-wide *ballot* vote primitive, all threads indicate whether they have a match or not. If unsuccessful, the query key is hashed, via its second cuckoo hash function, into an alternative bucket. The same processing unit then searches for this key as before, reporting failure if it is not found. Insert operations are performed via a modified, bucketized cuckoo hashing routine in which a new key is only added into an empty slot, instead of evicting a collided key.

Breslow et al. [BZG*16] introduce an expansion of a BCHT that allows for higher load factors, improved bucket load balancing, and a lower expected number of bucket lookups (less than 2) for both positive and negative queries. In this *Horton table*, a row is maintained for each bucket, which is denoted as either *Type A* or *Type B*. Each key is hashed by its *primary hash function* into the *primary bucket*. If the primary bucket is full, then the key either hashes, via one of its *secondary hash functions*, to a *secondary bucket*—after which we denote the key as a *secondary item*—or replaces a secondary item in the primary bucket. If the key is a secondary item, then it is placed in the secondary bucket that is least full; note that several secondary hash functions (and buckets/rows) can be specified. Then, the filled primary bucket is *promoted* (if not already) to *Type B* and its last stored key is evicted (moved to a secondary bucket) to make room for a compact *remap entry array* that stores an index, or *remap entry*, to the secondary bucket of each secondary item. This important feature permits all secondary items to be efficiently tracked, allowing no more than two probes and hash function evaluations per query. Additional bookkeeping and logic is used to delete keys and handle a cascading effect where an evicted key causes its secondary bucket to convert into *Type B*, which induces another eviction, and so on.

Experimental results of large query sets reveal that most successful lookups occur within the primary buckets, allowing a high load factor with only one hashing probe. Moreover, the performance of the Horton table is compared against that of a baseline BCHT similar to Zhang et al. [ZWY*15]. For all successful queries, the Horton table increases throughput (billions of keys queried per second) over the baseline by 17 to 35 percent. For a set of all unsuccessful queries, the Horton table increases throughput by 73 to 89 percent over the baseline, needing only one hash probe to detect failure. An important note regarding the design and performance of this approach is that only the query operations are conducted in data-parallel fashion on the GPU. The detailed insertion and construc-

tion phase is run on the CPU, which could make reconstruction costly for use other than as a static hash table, which is sufficient for the query-heavy usage of most key-value store systems [ZWY*15].

Hetherington et al. [HOA15] develop a fixed-sized *set-associative* hash table for scaling-up the throughput of key-value storage and accesses. As part of a *MemcachedGPU* caching service, HTTP GET requests are parsed to extract query keys that are hashed to a hash table on the GPU. This table facilitates *k*-way, set-associative hashing with each set (or bucket) entry consisting of an 8-bit key identifier hash and a pointer to the actual memory address (within a dynamically allocated *slab* of memory) at which the key is stored. After hashing to a set, each query key compares itself to the 8-bit hash and, if a positive match, accesses the key in memory and instigates a return package with the associated value, which is stored in CPU memory. If the query key does not exist in the set, then it was previously a *least-recently-used* (LRU) key that must have been evicted from the set by a colliding, more-recent key in the same set (recent usage based on timestamp). Thus, an HTTP SET request can reinsert this key into an empty entry in the set or evict a LRU key that resides in the set. Each set maintains and updates its own local LRU array. Experiments over varying hash table sizes (number of entries) and a query-heavy distribution of key-value requests (95% GET and 5% SET) reveal that MemcachedGPU achieves an acceptable hash table miss rate with 16-way set associativity. In these experiments, each request key is assigned to an individual warp thread for SIMT execution. Unless requests exhibit spatial locality, branch and memory divergence are inevitable in this approach.

Ashkiani et al. [ADMO16] design a set of *multisplit* data-parallel primitives for the GPU that efficiently permute elements into contiguous buckets. While this study is not focused on hashing, it recommends that the multisplit can be used to map elements into the first level of buckets in a multi-level hash table, such as the two-phase hash table of Alcantara et al. [ASA*09]. Moreover, this work contributes a *reduced-bit* radix sort that converges to and exceeds the performance of state-of-the-art radix sort [MG10] as the number of buckets is increased. Thus, if the order of insertions and queries into a bucket-based hash table are non-random and ordered, then this sorting primitive may offer an effective substitution for a bucketing procedure. These primitives have since been incorporated into the CUDPP library [cud17].

3.1.4. Double Hashing

Double hashing first hashes a key k to location $h(k)$ in the hash table and then, if the location is already occupied, computes another independent hash $h'(k)$ that defines the step size to the next probing location [CSRL01]. Thus, the second probe location is $h(k) + i \cdot h'(k)$, where i is the current i -th probe in the probe sequence. This hashing and probing continues until an empty slot (insertion) or k itself (query) is found. Similar to linear and quadratic probing, if $h(k)$ is empty, then k is inserted immediately, without probing. The choice of the second hash function has a large impact on performance, as it dictates the locality of consecutive probes and, thus, the opportunity for memory coalescing among threads on the GPU.

Khorasani et al. [KBGB15] introduce a *stadium hashing* (*Stash*)

technique that builds and stores the hash table in out-of-core host memory, and resolves insert collisions via double hashing until an empty slot is found. In GPU global memory, a compact auxiliary *ticket-board* data structure is maintained to grant read and write accesses to the hash table. For each hash table location, the ticket board maintains a *ticket*, which consists of a single *availability bit* and small number of optional *info bits*. The availability bit indicates whether the location is empty (set to 1) or occupied by a key (set to 0), while the info bits are a small generated signature of the key to help identify the key prior to accessing its value. Within individual thread warps, a shared-memory, *collaborative lanes* (*clStash*) load-balancing scheme is used to ensure that, during insertions, all threads are kept busy, preventing starvation by unsuccessful threads.

Stadium hashing is meant to address three limitations of previous GPU parallel hashing techniques, specifically in regard to the cuckoo hashing approach of Alcantara et al. [AVS*12]:

1. Support for concurrent, mixed insert and query operations. Without proper synchronization, cuckoo hashing encounters a race condition whereby a query probe fails at a location because a concurrently-inserted key hashes to the location and evicts the queried key, yielding a false negative lookup. Stadium hashing avoids this issue by using eviction-free double hashing and granting atomic access to a location via a ticket board ticket with the availability bit set to 1.
2. Reduce host-to-device memory requests for large hash table sizes. In cuckoo hashing, CAS atomics are used to retrieve the content of a memory location, compare the content with a given value, and swap in a new value, if necessary. When a hash table is stored in host memory, the large number of parallel retrieval requests from thousands of GPU threads will turn the hashing into a PCIe bandwidth-bound problem and degrade performance. Stadium hashing uses the GPU ticket board data structure to minimize retrieval requests to the host memory hash table.
3. Efficient use of SIMD hardware. During a cuckoo hashing operation, a thread failing to insert or query a key can cause starvation among the other threads in the thread warp, as they all perform the same instruction in lock-step. Thus, if the other threads complete their operation early, then they will remain idle and contribute to work imbalance. Stadium hashing uses the *clStash* load-balancing routine to maintain a warp-wide, shared memory store of pending operations that early-completing threads can claim to remain busy.

For an out-of-core hash table, the ticket-board with larger ticket sizes (more info bits per key) helps improve the number of operations per second by reducing the number of expensive host memory accesses over the PCIe bus. This improvement is especially significant for unnecessary queries of elements which do not actually reside in the host hash table. In this case, the PCIe bandwidth from GPU to CPU memory is the primary performance bottleneck. However, when the hash table resides in GPU memory, the underutilization of SIMD thread warps becomes the primary bottleneck on performance for low load factors (fewer collisions). The efficiency of warps is shown to improve by using the collaborative lanes *clStash* scheme in combination with the Stadium hashing.

Regarding the experiments and findings in this work, the cuckoo hashing approach of [AVS*12] is specifically designed for hash table construction and querying within GPU memory. Thus, the use of this technique in out-of-core memory should not necessarily be expected to perform optimally, and should be kept in mind when comparing against the out-of-core performance of stadium hashing.

3.1.5. Robin Hood-based Hashing

Robin Hood hashing [Cel86] is an open-addressing technique that resolves hash collisions based on the *age* of the collided keys. The age of a key is the length of the probe sequence, $h^1(k), h^2(k), \dots$, needed to insert the key into an empty slot in the hash table. During a collision at a probe location, the key with the youngest age is evicted and the older key inserted into that location. The evicted key is then robin hood hashed again until it is placed in a new empty location, incrementing its age along the new probe sequence. The idea of this approach is to prevent long probe sequences by favoring keys that are difficult to place. Even in a full table with high load factor, this eviction policy guarantees an expected maximum age of $\Theta(\log n)$ for an insert or query key. However, the worst-case maximum age M may still be higher and worse than the maximum probe sequence length of cuckoo hashing, prompting a table reconstruction in some cases. These maximum M probes will be required during queries for empty keys (those which do not exist in the hash table), unless they are detected and rejected early.

Garcia et al. [GLHL11] introduce a data-parallel robin hood hashing scheme that maintains coherency among thread memory accesses in the hash table. Neighboring threads in a warp are assigned neighboring keys to insert or query from a spatial domain (e.g., pixels in an image or voxels in a volume). By specifying a coherent hash function, both keys will be hashed near each other in the hash table and the threads can access memory in a coalesced fashion, i.e., as part of the same memory transaction. Thus, the sequence of probes for groups of threads will likely also be conducted in a coherent manner, as nearby keys of a young age are evicted and replaced by nearby keys of an older age.

The insertion and query performance of this technique is evaluated on both randomly- and spatially-ordered key sets from a 2D image and 3D volume. For all load factor settings, the existence of coherence in the keys and probe sequences results in significant improvements in construction and querying performance (millions of keys processed per second), as compared to the use of randomly-ordered keys. Moreover, coherent hashing achieves greater throughput than the cuckoo hashing of Alcantara et al. [AVS*12], which employs four hash functions for a maximum probe sequence of length four. For load factors above 0.7, coherent hashing maintains superior performance without failure (hash table reconstruction) during insertions, whereas the cuckoo hashing exhibits an increase in failures.

In absence of coherence in the access patterns, coherent hashing brings little to no benefit compared to the random access robin hood and cuckoo hashing. Thus, this approach is of particular use for applications with spatial coherence in the data. In one of the spatially-coherent experiments, the task is to insert a sparse subset of pixels from an image (e.g., all the non-white pixels) into the hash table, and then query every pixel to reconstruct the image. Since only

non-white pixels are hashed, there will be empty queries for the white pixel keys. Spatial and coherent ordering of keys is attained by applying either a linear row-major, Morton, or bit-reversal function to the spatial location of elements; a non-coherent, randomized order is attained by shuffling the keys.

Coherent hashing has some notable design characteristics that can affect GPU performance. First, upon completing an insert or query operation, a thread sits idle until all threads in its warp have finished as well. The number of threads per warp (192 in this work) and amount of branch divergence due to incoherent ordering of keys are primary factors affecting the warp load balancing. Second, while inserting a key, the hash table is reconstructed if the age, or probe sequence length, of the key exceeds a threshold maximum (15 in this work). Moreover, the hash table is not fully dynamic and is designed to process queries after an initial build phase. Thus, if new keys are inserted after the build phase, then the table is reconstructed entirely, with a larger table size or load factor if necessary.

Zhou et al. [ZGJ*16] design a modified GPU-based robin hood hashing scheme for use in storing and extracting the top- k most similar matches, or results, of a query (e.g. a document of words or multi-dimensional attribute vector). In this similarity search, a two-level Count Priority Queue (c-PQ) data structure hashes potential candidates for the top- k results into an upper-level hash table, as determined by a lower-level *ZipperArray* histogram array of object counts, where *ZipperArray*[i] is the number of objects (e.g., word or attribute value) that appear i times in the query (count of i). An *AuditThreshold* is set as the minimum index (count) of *ZipperArray* whose value (number of objects) is less than k . For an object to be inserted into the hash table, it must have a count greater than or equal to the *AuditThreshold*. As new items are added, object counts and the *ZipperArray* are updated, and the *AuditThreshold* may be increased. During insertion into the hash table, the robin hood hashing scheme of Garcia et al. [GLHL11] is used, with an additional feature that an object with a count smaller than (*AuditThreshold* - 1) can be directly overwritten, and not evicted, during a hash collision. This helps reduce the average age, or probe length, of new insertions, as previously-inserted objects become *expired* due to an increase in the *AuditThreshold*. This modification, along with a lock-free synchronization mechanism, affectively contributes a dynamic hash table variant of [GLHL11].

3.2. Perfect Hashing

Whereas the previous hashing categories employ *imperfect* hash functions that require collision-handling and multiple probes, *perfect hashing* maps each key to a unique address in the hash table, resulting in no collisions and enabling single-probe queries. If the length of the hash table m is equal to the number of keys n , then a perfect hash function over the keys is *minimal* and effectively scatters, or permutes, the keys within the table.

In theory, obtaining a perfect hash function, especially for large sets of keys, is a difficult, low-probability task. Given a universe $U = \{0, 1, \dots, u-1\}$ of possible keys, subset $S \subset U$ of $|S| = n$ keys to be hashed, a hash table of size m , and class H of hash functions $h : U \rightarrow [0, m-1]$, the probability $P(n, m)$ of randomly placing n

keys in $m \geq n$ slots without a collision is

$$P(n, m) = \frac{m(m-1) \cdots (m-n+1)}{m^n}.$$

This can also be stated as the probability of a randomly-chosen hash function $h \in H$ being a perfect hash function over the set S . As a reinterpretation of the classical birthday paradox, only one in ten million hash functions $h \in H$ is a perfect hash function for $n = 31$ keys mapped into $m = 41$ locations. When $m \gg n$, $P(n, m) \approx e^{-\frac{n^2}{2m}}$, which implies that there is a very low probability of attaining a perfect hash when the load factor or occupancy of the hash table is very high; i.e., $m \ll n^2$. Moreover, when $m = n$, $P(n, m) = \frac{n!}{n^n}$, which is the probability of achieving a minimal perfect hash [Meh82]. For larger key set sizes n , such as those seen in practical applications, the minimal perfect hash probability decreases very rapidly and is approximated as e^{-n} .

In practice, a perfect hash function can be described as an imperfect hash function that is then iteratively corrected into a perfect form. One approach to doing this is to construct one or more auxiliary lookup tables that perturb the hash table addresses of collided keys into non-colliding addresses [FHCD92]. These tables are typically significantly more compact than the hash table. Another foundational approach, introduced by Fredman et al. [FKS84], is the use of a multi-level hash table that hashes keys into smaller and smaller buckets—each with a separate hash function—until each key is addressed to a bucket of its own, yielding a collision-free, perfect hash table with constant worst-case lookup time. Moreover, a perfect hash function is most suitable for static hash tables (and key sets), in which no insertions or deletions occur after the construction of the table. If dynamic updates are performed, then the function will inevitably become imperfect—with collisions among relocated keys—and require a reconstruction procedure. Czech et al. [CHM97] survey a rich body of related work investigating additional perfect hashing and minimal perfect hashing schemes (largely non-parallel), each designed for CPU-based storage and computation.

Lefebvre and Hoppe [LH06] introduce a perfect spatial hashing (PSH) approach that is also the first GPU-specific perfect hashing approach. In PSH, a minimal perfect hash function and table are constructed over a sparse set of multi-dimensional spatial data, while simultaneously ensuring locality of reference and coherence among hashed points. Thus, on the GPU, spatially-close points are queried coherently, in parallel, by threads within the same warp. In order to maximize memory coalescing among these threads, points are also coherently accessed within the hash table, as opposed to via a random access pattern, which can necessitate multiple memory load instructions.

In the PSH study, the spatial domain U is a d -dimensional grid with u points (positions), where $d \in \{2, 3\}$. The grid serves as a bounding box over a sparse subset $S \subset U$ of n points that have associated data records $D(p)$, $p \in S$ (e.g., RGB color value for each pixel or voxel); thus, the remaining $u - n$ grid locations are stored in memory without any compute value. The sparse subset D is more compactly stored in a dense hash table H of size $m \geq n$. This table is addressed by a multi-dimensional perfect hash function h that is composed of two imperfect hash functions, h_0 and h_1 , and an *offset*

table Φ that “jitters” h_0 into perfect form. This function is defined as

$$h(p) = h_0(p) + \Phi[h_1(p)] \bmod \bar{m}, \forall p \in S,$$

where h_0 and h_1 perform simple modulo addressing and wrap the points S multiple times over H and Φ , respectively. Due to this modulo, lockstep addressing by h_1 , spatial coherence is preserved for access into Φ . However, the values $\Phi[h_1(p)]$ perturb the coherency of the combined function h . By constructing Φ such that adjacent offset values are locally constant, the coherency of h can be ensured. Note that h is not strictly a minimal perfect hash function (i.e., when $m = n$), since the hash table size m may need to be slightly increased to facilitate the perfect hash represented within Φ . The number of unused table entries $m - n$ is kept small and is considered *near-minimal*. Thus, the sizing of H and the spatial coherency of h are both dependent on the proper construction of Φ .

The construction of the offset table Φ proceeds by first identifying the smallest table size r that yields a perfect hash h . A geometric progression or binary search of r is iteratively conducted, depending on whether faster construction or a compact representation is desired, respectively. For each size r tested, the offset values $\Phi[q], 0 \leq q \leq r$ are assigned via a greedy heuristic procedure that seeks to maximize spatial coherence. Since $r < n$, on average $\lfloor \frac{n}{r} \rfloor$ points, $h_1^{-1}(q) \subset S$, hash to each entry $q \in \Phi$. The entries q are assigned in descending order by size of their point sets $h_1^{-1}(q)$. Then, each q is assigned one of the following two candidate heuristic values:

1. Same offset value as a neighboring entry, $\Phi[q'], \|q - q'\| < 2$.
2. For each $p \in h_1^{-1}(q)$ with a neighboring point $p' \in S, \|p - p'\| = 1$ already hashed in $H[h(p')]$, the offset value that places p in an empty neighboring slot of $H[h(p')]$.

If Φ yields a perfect hash function h with the tested size r , then the construction phase completes; otherwise, the routine is conducted again with another size r .

In a SIMT fashion on the GPU, point queries $q \in U$ are executed in parallel by threads, each computing $h(q)$ and looking up an associated value from $H[h(q)]$, which is mapped to a single point due to the perfect hash h . If q does not exist in H , then a negative query result is returned.

Note that the construction of Φ is an inherently sequential process, since the assignment of offset values depends on earlier offset values or hashed points in H . Moreover, the construction of H and Φ is performed on the CPU in this work, due to the larger memory requirements and presumed usage as a pre-processing step; thus H must be copied into GPU device memory over the PCIe bus. Moreover, H is designed to be static, since insertions or deletions of points after construction destroy the perfect hash and require Φ to be reconstructed.

Bastos and Celes [BC08] implement a GPU-based *link-less* octree data structure by hashing the parent-child (node) relationships into the perfect hash table of Lefebvre et al. [LH06]. Thus, instead of constructing a multi-level tree with pointers over sparse spatial data, only a compact perfect hash table needs to be built; however, updates to this data structure are costly and require the entire

table to be reconstructed, as the perfect hash is intricately data-dependent. Since perfect hashing is collision-free, direct random-access queries can be made on the octree, as opposed to traditional pointer tracing in tree traversals.

Choi et al. [CJC*09] follow-up the work of Bastos and Celes [BC08] with a similar link-less octree design that avoids the need to store extra bitmaps for the *sparsity encoding* of empty grid cells in the sparse spatial domain. This encoding indicates whether a cell contains associated data that is stored within the hash table; if no data is present, then a query operation for the cell can be avoided. While this latter approach significantly reduces storage costs, it executes random-access queries much slower than similar accesses into the benchmark pointer-based octree.

Schneider and Rautek [SR17] denote sparsity encoding as a memory overhead cost for providing *unconstrained access*, or empty cell querying, in the spatial perfect hashing approach of Lefebvre et al. [LH06]. This study proposes a compact, GPU-based *Fenwick tree* data structure that supports unconstrained accesses without additional occupancy storage to denote empty cells.

3.3. Spatial Hashing

The following two subsections present GPU-based spatial hashing techniques for inserting and querying regular grid cells (subsection 3.3.1) and point coordinates (subsection 3.3.2) within a multi-dimensional spatial domain.

3.3.1. Grid-based Spatial Hashing

Most real-world use cases of searching require a data structure that can support lookups of geometric primitives — e.g., point coordinates, polygonal shapes, and voxels — that exist within a multi-dimensional *spatial domain*, such as \mathbb{R}^2 , \mathbb{R}^3 , or \mathbb{R}^n . One approach is to explicitly compute a bounding box over the domain and then recursively subdivide it into smaller and smaller regions, or cells, which contain a group of primitives or a subset of the spatial domain. This subdivision hierarchy can be represented by a grid (e.g., uniform and two-level) or tree (e.g., k -d tree, octree, or bounding volume hierarchy) data structure (see Subsection 2.4) that conducts a query operation by traversing a path through the hierarchy until the queried primitive is found. While these structures are designed for fast, highly-parallel usage, they typically do not exhibit fast reconstruction rates due to complex spatial hierarchies, and may contain deep tree structures that are conducive to thread branch divergence during parallel query traversals. These attributes are particularly important to real-time, interactive applications, such as surface reconstruction and rendering, that make frequent updates and queries to the acceleration structure.

An alternative approach that addresses these limitations is to perform *spatial hashing* over the primitives, whereby the multi-dimensional domain is projected, or compressed, to a single dimension in the form of a hash table data structure. Instead of computing a bounding box over the spatial domain and explicitly storing the entire space, spatial hashing assumes an implicit, infinite regular grid over the domain and maps each positional primitive (e.g., a point coordinate) to a uniformly-sized and axis-aligned cell

within the grid. Each cell is uniquely addressed by unit coordinates and contains a user-specified number of primitives within its bounds [SLM04]. These coordinates are used by the hash function to hash the cell into the hash table. Two or more cells may hash to the same address, resulting in collisions that must be resolved. To query a primitive (or cell), the primitive is mapped to its cell and the cell is hashed to an address in the hash table. From this address, the cell is searched, using more than one probe if a collision occurs. Typically, to exploit sparsity, only non-empty cells that contain computable primitive data (e.g., pixel intensity, RGB, or density) are inserted into the hash table. A query of an empty cell will return a negative result, as it doesn't exist in the table.

This canonical grid-based *voxel hashing* approach was introduced by [THM*03] as a CPU-based search structure for detecting colliding 3D tetrahedral cells in \mathbb{R}^3 domain space. Several follow-up studies have since introduced GPU-based spatial hashing techniques based off of this approach, and they are surveyed as follows.

Nießner et al. [NZIS13] extend the approach of Teschner et al. [THM*03] with more sophisticated collision handling and a 3D voxel hashing scheme that is designed particularly for fast, real-time hash table updates on the GPU. An infinite uniform grid subdivides the world space into voxel blocks, each of which consists of 8^3 voxels. The world coordinates of each voxel block are hashed as an address into a bucketed hash table. During an insertion, a block probes linearly through its assigned bucket for the first empty slot that it can occupy. If a free slot is found, then the block is inserted. Otherwise, if the bucket is already full, then *overflow* occurs and a linked list entry in the last slot points to the next free slot in another bucket of the hash table. The block then probes along this overflow chain to find the next empty slot. Due to interconnection among buckets, each hash entry contains an offset pointer to its neighboring bucket entry, which may not be adjacent in the table. A query operation conducts similar probing to find a particular block within the hash table. Additionally, lightweight GPU atomic primitives are used to coordinate data-parallel insertions and deletions of blocks, each assigned to an individual thread. While an entire bucket is locked for writing during an insertion into the bucket, no degradation in performance is observed for a high-throughput, real-time 3D scene reconstruction experiment. Moreover, by using a larger hash table size, the number of collision is kept minimal and prevents bucket overflows into other disparate buckets, which can cause uncoalesced memory accesses among warp threads.

Kähler et al. [KPVM16] introduce a GPU-based hierarchical voxel block hashing technique that uses multiple hash tables in a hierarchy to store finer and finer resolutions of grid discretization for voxel blocks (cells). Initially, each block is hashed to an entry in a first-level hash table of *coarse resolution*. Then, if the voxels within this block are represented at a *finer resolution*—as indicated by a flag in the entry of each hash entry—the block is hashed again with a different hash function into a second-level hash table. This hierarchical hashing continues until an entry is reached that contains a pointer to the *voxel block array*, which stores the actual, individual block data. Atomic voxel block *splitting* and *merging* operations are supported to enable the addition or removal of hash table entries for finer or coarser resolutions, respectively. On scene reconstruction experiments with signed distance function (SDF) values

for roughly 20 million points, this adaptive hierarchical representation, with 3 resolution levels and base voxel size of 2 mm, attains greater accuracy than a fixed representation with 8 mm voxel sizes.

Note that the hierarchical voxel hashing of Kähler et al. [KPVM16] is inspired by the general approaches of Eitz and Lixu [EL07] and Pouchol et al. [PACT09], which themselves expanded upon the original spatial hashing of Teschner et al. [THM*03]. These studies are each CPU-based and use real-time collision detection as a motivating application.

Chentanez et al. [CMM16] introduce a GPU-based spatial hashing variant of Teschner et al. [THM*03] for detecting and deleting overlapping triangles on the surface of a 3D mesh volume, as vertices are advected (i.e., mesh refinement). In this work, the 3D bounding cells of triangles are inserted into a specially-arranged hash table using the coordinate-based hash function from [THM*03]. The hash table consists of n buckets each with m available slots ($n \cdot m$ entries), and the first n entries of the table are reserved to store counts of the number of slots $j \leq m$ that are occupied in each bucket. Thus, the total allocated size of the table is $n(1+m)$. During an insertion of a cell k into bucket $h(k) = b$, the thread assigned to cell k first checks the occupancy count value for bucket b . If b has open slots, then k is inserted into the first available slot and the count for b is atomically incremented. Otherwise, the thread examines the count for the next bucket $b+1$ and inserts k into the first open slot of $b+1$, if possible, so on and so forth until k is successfully inserted. This is a modified collision resolution scheme whereby a bucket collision only occurs when the bucket is full and subsequent buckets are then linearly-probed for one that has an empty slot. During a cell query, the same linear probing over buckets is performed, beginning with the bucket to which the cell is hashed.

Note that, in this approach, thousands of other parallel threads are executing the same operation on different triangle cells, likely resulting in high contention for atomic writes for the bucket count values and worst-case linear probing sequences that induce branch divergence within warps. The extent of such divergence depends on the size m of each bucket and whether locality of reference is maintained among bucket entries when hashing spatially-nearby cells. These performance effects are not explored in the experimental findings of this approach.

Tumblin et al. [TAHR15] expand upon traditional perfect spatial hashing (PSH) with a *compact spatial hashing* (CSH) variant that compacts a perfect hash table when it becomes too sparse, saving unused memory on the GPU. As a larger number of keys need to be hashed, a sufficiently large hash table must be allocated to construct a perfect hash among the keys. Often, this large table still contains many empty locations, resulting in a low occupancy and high *compressibility*, which is the ratio of available table locations to occupied locations. A *compression* function randomizes the original hash locations of each key and fits them within a smaller, compact hash table of size proportional to the number of keys divided by a desired load factor. Since PSH is collision-free, this compaction inevitably induces collisions, which are handled in this work by a canonical quadratic probing open-addressing method in parallel. The goal of the compression function, thus, is to reduce the occurrence of collisions via random scattering of keys. However, this

random re-assignment of hash locations disrupts any spatial locality that existed among the keys, preventing warp-level memory coalescing during accesses. Experimental results for an adaptive mesh refinement (AMR) task show that as the perfect hash table reaches 20 to 40 times the size of the compact hash table, the CSH becomes the faster method. Thus, the exceedingly larger memory of PSH offsets the extra costs (e.g., thread divergence and uncoalesced memory) of resolving collisions and querying in CSH.

Duan et al. [DLN*17] present an *exclusive grouped spatial hashing* (EGSH) scheme that is optimized to compactly represent multi-dimensional domains that contain repetitive data (e.g., duplicate RGB or density values). The goal of this approach is to identify all groups of points that share the same data values and then, for each group, compress its points into a single group-wide value, avoiding the unnecessary storage of duplicates, which are significantly prevalent in some domains. This grouped hashing is performed over multiple iterations using multi-level hash tables until each group has been exclusively hashed into a unique table location. The logistics of this approach are discussed as follows:

- In the first iteration, all points in the *input domain* are hashed into a hash table of size equal to the number of points. Then, at each non-empty hash table location, collided points are binned into disjoint groups based on their corresponding data values. The data value of the group with the most points is set as the exclusive value in this hash table location, replacing and compressing the many repetitive points of the group. For subsequent querying, the hash table ID (iteration index) of each compressed point is stored in a global, persistent *coverage table*. The remaining uncompressed points of the other groups are then advanced as the input domain to another iteration of exclusive group hashing. In the next iteration, all the uncompressed points (among all hash table locations from the previous iteration) are hashed into a smaller hash table of size approximately equal to the number of groups from the previous iteration. The grouping and compression routine is then conducted as before, and the hashing iteratively continues until all points are compressed.
- The compression of repetitive elements contrasts with other hashing approaches covered in this survey, which hash repetitive keys into separate, and possibly disparate (depending on load factor and table size) addresses of the table upon collision.

Experiments on the GPU reveal that after several iterations of EGSH, the input domain becomes very sparse and has a rapid reduction in the amount of repetitive data (uncompressed groups). Both of these traits are highly suitable for the perfect spatial hashing (PSH) of Lefebvre and Hoppe [LH06], which similarly provides constant-time random accesses. Thus, an optimized variant of EGSH performs exclusive grouped hashing for a small number, k , of iterations—generating k levels of hash tables—and then applies the PSH on the remaining uncompressed input domain. In this work, $k = 6$ is used for a set of 2D and 3D grid-based input textures, all of which possess a high ratio of points with repetitive data values. The results of a comparison between optimized EGSH and PSH on these textures reveal that both schemes have similar constant access times, while the EGSH memory cost is typically less than half of the PSH memory cost. Takeaways of this study are that PSH achieves best performance on sparse, slightly repetitive

domains, as opposed to sparse, highly repetitive or dense domains. Contrarily, EGSH attains the smallest memory savings and construction time for input domains with highly repetitive data.

A few important notes regarding this EGSH work are the following:

- Thread- and warp-level GPU performance findings are not provided. Only high-level runtimes and memory usage are analyzed. Moreover, the ESGH multi-level hash tables appear to be constructed on the CPU and copied over the PCIe bus to the GPU for subsequent querying, much like that of the PSH hash table construction. This is notable, since, during construction, the search for the group with the maximal number of elements at each hash table location is the most time-consuming task and may not be optimally parallelized on the CPU.
- When each group has only one point, ESGH degenerates into PSH, whereby all points are hashed to unique locations of a single table. When the groups contain multiple repetitive points, multiple sub-tables are needed to complete the ESGH.

3.3.2. Locality-Sensitive Hashing

Much like grid-based spatial hashing, locality-sensitive hashing (LSH) reduces the dimensionality of high-dimensional data via a projection to a 1D hash table. LSH hashes input elements so that similar elements map to the same buckets with high probability, with the number of buckets in the hash table being much smaller than the universe of possible input elements. LSH differs from the other hashing approaches covered in this survey because it aims to maximize the probability of collisions between similar items [IM98]. Similar to other approaches, LSH employs a collision resolution scheme to relocate collided elements that are inserted into the hash table. During a query operation, LSH is well-suited for returning the k approximate nearest neighbors (k ANN) of the query element q , since these neighbors will likely reside in the same bucket as q [IM98, DIIM04]. While performant GPU-based, brute-force approaches exist to find the *exact* k nearest neighbors [KZ09, LA15], a large body of recent research has investigated the design and performance of LSH for k ANN.

More formally, canonical LSH proceeds as follows, beginning with the construction of the LSH hash tables. Given a set of D -dimensional points $S \subset \mathbb{R}^D$, $M < D$ different hash functions $h_i(\mathbf{p}) = \lfloor (\mathbf{a}_i \cdot \mathbf{p} + b_i) / W \rfloor$ are used to project each point $\mathbf{p} \in S$ to a cell within a \mathbb{Z}^M lattice space, the size of which is determined by M and W ($\mathbf{a}_i \in \mathbb{R}^D$ and $b_i \in [0, W)$ are randomly generated). This cell location, or *LSH projection*, is specified as a M -dimensional vector $H_j(\mathbf{p}) = \langle h_1(\mathbf{p}), h_2(\mathbf{p}), \dots, h_M(\mathbf{p}) \rangle$ and then mapped into a single value $g_j(H_j(\mathbf{p}))$ known as the *LSH code*. This code represents the bucket index of the hash table G_j into which \mathbf{p} is then inserted. To decrease the collision probability of false neighbors, \mathbf{p} is projected and hashed into $j = L$ different and independent hash tables, with each instance of H_j and G_j being randomly generated. During a query for arbitrary point $\mathbf{v} \in \mathbb{R}^D$, \mathbf{v} first computes its L different LSH codes into the L different hash table buckets. Then, a *candidate set* of nearest neighbors of \mathbf{v} is composed of the union $A(\mathbf{v})$ of all points hashed into the same L buckets as \mathbf{v} . A *short-list search* over $A(\mathbf{v})$ calculates the subset $k \subset A(\mathbf{v})$ of k neighbor points that are closest in distance to \mathbf{v} . This short-list is typically implemented

with a max-heap data structure and requires the most computation with LSH [IM98, DIIM04, PLM10].

Two surveys led by J. Wang et al. [WSSJ14, WLKC16] review additional CPU-based techniques related to LSH. The following studies exclusively focus on data-parallel, GPU-based LSH approaches.

Pan et al. [PLM10] design a data-parallel GPU-based variant of the canonical LSH method that simulates the L different hash tables with a single cuckoo hash table of Alcantara et al. [ASA*09]. In SIMT parallel fashion, threads execute insert and query operations on their assigned D -dimensional points. During an insertion, all of the LSH codes of points are data-parallel radix sorted in a linear array. Then, the sorted codes are partitioned into *buckets*, based on unique code values. A data-parallel prefix-sum scan identifies the starting and ending indices of each bucket. The LSH code and its bucket interval define a key-value pair that is inserted into the cuckoo hash table, or *indexing table*. Multiple points may have the same LSH code/key and thus collide at the same index table location. These collisions are resolved via a set of $L \geq 3$ hash functions that define the probe sequence needed to relocate points upon eviction. Note that the L functions simulate hashing a point into the L separate hash tables (or buckets) of traditional LSH. Finally, a query operation on a point computes the LSH code of the point, probes for the corresponding key in the indexing table, and then extracts the associated bucket interval value. The points within this bucket define the candidate set of neighbor points. A data-parallel search over a max-heap of these points returns the k ANN. Experiments on real-time motion planning data reveal that this approach is both faster and more accurate than comparable tree-based k ANN approaches. Also, the accuracy, or spatial locality, of the nearest neighbor hashing increases as the number of cuckoo hash functions L increases.

Pan and Manocha [PM11] follow-up their original GPU-based LSH approach [PLM10] with *bi-level* LSH that adds the following four components:

1. *Random projection (RP) tree*: In the first level, an RP tree is constructed in parallel over the input data points by iteratively partitioning the points into smaller and smaller clusters until a desired tree depth is reached, with leaf nodes containing small subsets of likely spatially-similar points. The tree is similar to a k -D tree, but splits the points along random directions instead of along coordinate directions [DF08]. This addition to the LSH helps generate more compact and accurate LSH codes.
2. *Hierarchical LSH table*: In the second level, an LSH table is constructed for each different RP tree leaf node and its subset of points. Unlike the previous LSH approach, two additional steps are performed prior to computing LSH codes. First, each point in the leaf node is projected into a more compact E_8 lattice space that produces more accurate projections for high-dimensional data. Then, these LSH projections are mapped to 1D *Morton curve* values that preserve neighborhood relationships. These values are efficiently constructed on the GPU, via the approach of Lauterbach et al. [LGS*09], and serve as LSH codes.
3. *Bi-level LSH code*: A modified Bi-level LSH code for a point \mathbf{v} is specified by the pair $\tilde{H} = (\text{RP-tree}(\mathbf{v}), H(\mathbf{v}))$, where $\text{RP-tree}(\mathbf{v})$ is the index of the leaf node to which \mathbf{v} belongs and $H(\mathbf{v})$ is the

Morton curve value (or LSH code). These bi-level codes are then data-parallel radix sorted and bucketed, producing the bucket intervals. Similar to the previous approach, the LSH codes $H(\mathbf{v})$ and their corresponding bucket intervals form key-value pairs for cuckoo indexing table.

4. *Work queue*: Instead of extracting the k ANN of a query point \mathbf{v} from a global memory max-heap of size k , a global memory work-queue is used to perform a clustered sort that arranges the candidate set of neighbors in ascending order of distance from \mathbf{v} . This sort works in data-parallel across multiple queries and candidate sets, and employs radix sorting, which can benefit from the high-speed GPU shared memory. Moreover, this queuing approach increases parallel throughput and avoids thread branch divergence inherent in the max-heap tree traversals.

A set of experiments on an image dataset with nearly 2 million images, each represented as a 512-dimensional point, demonstrate that the Bi-level LSH can provide higher quality ANN results than the previous LSH method, given the same computational budget. Each of the algorithmic improvements discussed above attain accelerated GPU performance over the original methods.

Gieseke et al. [GHO14] introduce a *buffer k -d tree* data structure for massively-parallel ANN search on the GPU. While hashing is not used in this study, the authors state a weakness of the Pan and Manocha [PM11] approach is that it possibly yields inexact answers, as compared to those of a serial benchmark. While a reason was not provided, this inaccuracy may be due to either the RP-tree spatial partitioning of points or the hierarchical LSH code calculation, which involves consecutive mappings to lower-dimensional spaces.

Lukač and Žalik [LŽ15] implement a GPU-optimized variant of *multi-probe LSH* (MLSH) that was originally introduced by Lv et al. [LJW*07]. In this approach, L hash tables are constructed, one at a time, on the GPU using the unique LSH codes of projected multi-dimensional points as bucket indices (the LSH code is a function of K different LSH projections to a line). The points within each bucket are sorted in ascending order by ID using the data-parallel radix sort of Merrill et al. [MG10]. During point queries, the candidate sets are composed using *query-directed probing* to first visit buckets that possess a high success probability of containing nearest neighbors. Given the properties of LSH, these neighbors are likely to be in buckets that only differ slightly in distance in the table. A *scoring criteria* assigns a threshold for each bucket, determining whether it should be probed, based on its likelihood of containing a nearest neighbor of the query point. In this work, a simplified multi-probe scheme assigns a scoring criteria to the immediate left and right buckets of the bucket into which the query point is hashed. Thus, the points of at most 3 buckets combine to form the candidate set. In order to quickly extract the k ANN for the query point, a *deterministic skip-list* (DSL) search structure is built in global memory. This structure arranges the candidate set points in multiple levels of sorted linked lists, or subsequences, each of increasing size and in order of distance from the query point [MPS92, MCP17]. The resulting k ANN is copied back to host CPU memory, and then the LSH procedure is iteratively repeated for the remaining $L - 1$ hash tables, after which L different sets of

k ANN reside on the host. From these $L \cdot k$ points, a final k ANN is determined.

3.4. Separate Chaining

Separate chaining is a classic collision resolution technique that uses a linked list or node-based data structure to store multiple collided keys at a single hash table entry. Each hash table entry contains a pointer, or memory address, to a head node of a linked list, or *chain*. Each node in the linked list consists of a key, associated value (optional), and a pointer to the next node in the list, if any. If a single key hashes to an entry, then the linked list consists of a single node with a null pointer to the non-existent next node. Otherwise, if multiple keys collide and hash to the same location, then the linked list forms a *chain* of these keys, each represented by a separate node in the list. During a query operation, a key hashes to an entry in the table and then iterates through each of the nodes of the chain referenced at the entry, searching for a matching key. This iterative search is similar in nature to open-addressing linear probing (refer to subsection 3.1.1), where a key hashes to an initial table entry and then probes each subsequent entry until a matching key is found. Both techniques can result in degenerate, worst-case queries that require a non-constant number of probes. Unlike separate chaining, linear probing is prone to primary clustering of collided keys and performs *lazy deletion* of keys that renders unoccupied table entries heavily fragmented and may require re-hashing or compaction. However, linear probing is more amenable to caching, as probes are conducted within a contiguous block of memory, instead of over the scattered memory of linked list nodes.

Moreover, separate chaining is a form of *open hashing* in which keys and values are stored in allocated linked lists outside of the hash table and then referenced by head node pointers that are stored inside the table. Contrarily, open addressing collision resolution follows *closed hashing*, whereby each hashed key (and value) is inserted directly into the hash table.

In the context of parallel hashing, separate chaining must synchronize collisions during key insertions to ensure that the linked list data structures are properly allocated and constructed. Moreover, a dynamic memory allocation scheme must ensure that concurrent threads conducting insert operations properly synchronize their requests for new available blocks of memory. Similar design challenges exist for the deletion of keys, and the simultaneous execution of queries by threads must avoid reader-writer race conditions that result in faulty memory accesses to incorrect or deallocated nodes (keys).

A large body of research has investigated *concurrent* hash tables for separate chaining on multi- and many-core CPU systems [Gre02, Mic02, PGB*05, SS06, GHS*10]. Each of these hash tables is designed to support dynamic[§] updates and resizing with lock-based methods (e.g., mutexes or spin-locks) or lock-free (non-blocking) hardware atomics, such as compare-and-swap (CAS). Since the majority of these hash tables are linked list-based data structures, they are not designed for scalable, high-performance on massively-parallel GPU architectures. In particular, when ported to the GPU, the performance of these approaches may degrade due to several reasons:

- Lock-based methods induce substantial thread contention during blocking operations for shared resources and are not scalable with increasing numbers of concurrent threads. This contention creates starvation for blocked threads and warp underutilization, since each thread must wait for its other warp threads to finish acquiring and releasing the lock. Moreover, lock-free hardware atomic primitives prevent deadlock, but still neglect the sensitivity of GPUs to global memory accesses and thread branch divergence.
- Lack of coalescing among memory accesses due to the scattering of linked list node pointers in memory and random addressing of keys by threads within the same warp, which can lead to additional global memory transactions (cache line loads).
- Dynamic memory management and pointer chasing required for the linked lists on the GPU is challenging for traditional CPU-based synchronization schemes, due to the massive thread parallelism. This performance challenge is similarly observed in pointer-heavy spatial search tree structures that are ported to the GPU.

The following studies introduce GPU-based separate chaining hashing approaches that attempt to address these performance challenges.

Moazeni and Sarrafzadeh [MS12] and Misra and Chaudhuri [MC12] deploy some of the earliest lock-free, separate chaining-based hash tables on a GPU architecture. Using CUDA atomic CAS operations (atomicCAS and atomicInc), both approaches support batches of concurrent query and insert operations, with only [MC12] also supporting delete operations. [MS12] achieves a significant execution time speedup for queries over counterpart lock-based and OpenMP-based CPU implementations. However, the lock-free table only attains significantly higher throughput (operations per second) than the OpenMP implementation for query-heavy batches (80% queries and 20% inserts). Additionally, this work does not focus on larger, scalable batch sizes and provides little analysis regarding thread- or warp-level performance. [MC12] demonstrates that a GPU lock-free hash table leverages a much higher degree of concurrency and throughput than a CPU implementation for both query-heavy and update-heavy workload batches. This performance increase is largely due to spreading the thread contention and atomic comparisons over multiple different hash locations, as threads work in SIMT data-parallel fashion to conduct mixed operations at random locations.

Stuart and Owens [SO11] and newer versions of the Nvidia CUDA C library [Nvi17b] both introduce new efficient synchronization and atomic primitives (e.g., warp-level and share memory atomics) for CUDA-compatible GPUs. These primitives likely satisfy the inefficiencies of atomics for pointer-based data structures cited in Misra and Chaudhuri [MC12].

Steinberger et al. [SKKS12] design *ScatterAlloc*, an efficient GPU-based dynamic memory allocator that is significantly faster

[§] Some implementations are aware of future insertions at compile-time and preallocate sufficiently-large additional memory. These hash tables are *semi-dynamic* since they do not dynamically allocate new memory at runtime for unknown insertions.

than the built-in CUDA toolkit allocator and the first-published GPU allocator, *XMalloc*, of Huang et al. [HRJ*10]. This scheme maintains a linked-list memory pool of *super blocks* and organizes the blocks into larger fixed-size *pages*, which are addressed via a hash function. For usage in separate chaining hashing on the GPU, linked list collision chains can be dynamically allocated or deallocated as super blocks to large numbers of threads in parallel, as part of update operations (e.g., insert or delete). Due to the hash-based addressing of available memory pages, threads can minimize contention for the same block of memory and scatter their block assignments for efficient random access (with a possible tradeoff of memory fragmentation). Vinkler and Havran [VH15] survey and experimentally compare existing GPU dynamic memory allocation schemes. The performance of each scheme varies across different criteria, including fragmentation of available memory blocks, per-block thread contention for atomic allocation requests, size and coalescing of requested memory by inter-warp threads, uniformity of the number of allocation requests per inter-warp thread, and dependence on the number of user-specified registers available to threads in each SM of the GPU.

Ashkiani et al. [AFO17] propose a dynamic *slab hash* table on the GPU that is built upon an array of linked-lists, or *slab lists*, each of which represent a chain of one or more *slabs*, or memory units, that store collided keys. Each slab of memory is roughly the size of a warp memory transaction width (128 bytes), or the number of warp threads (32) times the size of a key (4 bytes). Thus, each warp is aligned to perform operations over the keys stored in a single slab. As part of a novel *work-cooperative work sharing* (WCWS) strategy, each warp maintains a *work queue* that stores all the arbitrary operations assigned to the different threads in the warp. In a round-robin fashion, each batch of the same operation type in the queue is fully and cooperatively executed by the threads. For a given operation type, all threads perform a warp-wide ballot instruction to denote the *active* threads that were assigned this operation. For each active thread, the entire warp cooperates to execute the active thread's operation and its corresponding key. If the operation is a query for a key q , then the warp hashes q to a slab list $b_i = h(q)$ in the slab hash table H . The first warp-sized slab, b_{i0} , of the slab list at $H[b_i]$ is loaded from global memory via a single memory transaction. This slab memory unit contains the same number of keys as threads in the warp. So, each warp thread then compares its corresponding key k with the query key q . If any thread has $k = q$, then a successful result is returned. Otherwise, the warp follows a *next* pointer stored in b_{i0} to load the next connected slab b_{i1} , in which q is cooperatively searched again. This search ends when either q is found or the last slab in b_i has been searched.

An insert operation proceeds similarly, except the threads search for an empty slab spot into which a new key can be atomically inserted. If no empty spot is found in any of the slabs of the slab list, then a new slab must be atomically and dynamically allocated (since other warps may also be trying to allocate). This allocation is efficiently performed via a novel warp-synchronous *SlabAlloc* allocator (see [AFO17] for further details).

This warp-cooperative approach differs from previous GPU separate chaining in which the threads of a warp execute a SIMT query or update operation for different keys, each of which likely require

a random, uncoalesced memory access. WCWS ensures memory coalescing for each operation by perfectly aligning the threads of a warp with the keys of a slab, both of the same size. Thus, the same block of cache-aligned global memory is loaded in a single transaction for every operation by the warp, exposing increased throughput (millions of operations per second). Moreover, while being inserted, keys are always stored at contiguous addresses within a slab memory unit. This contrasts with traditional linked list storage in which keys are inserted as new nodes at random memory locations.

The performance of the dynamic slab hash table is compared to the static cuckoo hash table of Alcantara et al. [ASA*09]—which must be rebuilt upon updates—and the semi-dynamic lock-free hash table of Misra and Chaudhuri [MC12]. For static bulk builds, cuckoo hashing consistently achieves a higher throughput of insertions per second, while slab hashing achieves higher query throughput only when the average number of slabs per slab list is less than 1 (i.e., approximately a single “node” list). Over all configurations, cuckoo hashing attains the better query throughput. In the best case scenario, it only makes a single atomic comparison for an insertion and a single random memory access for a query; contrarily, in the best case, slab hashing requires both a memory access (to load a slab) and an atomic comparison for an insertion. For dynamic updates, slab hashing significantly outperforms cuckoo hashing, in terms of execution time, as the number of inserted keys increases. This is due to the rebuilding of the static cuckoo hash table each time a new batch is inserted. Additionally, slab hashing significantly outperforms lock-free hashing across different distributions of mixture operations and increasing numbers of slab lists (i.e., the size of the hash table).

4. Hashing Applications

The following section highlights a variety of real-world applications of the GPU-based hashing techniques presented in this survey. These applications can be broadly divided into six categories, many falling within the domains of computer graphics and database processing. The majority of the studies cited within each application area also introduce a novel hashing technique and are discussed in section 3; the remaining studies strictly apply one of the hashing techniques.

Collision detection: Teschner et al. [THM*03] and Eitz and Lixu [EL07] use spatial hashing to detect real-time intersections between deformable objects in a scene and tetrahedral cells in 3D mesh volumes. Lefebvre and Hoppe [LH06] use perfect spatial hashing to detect collisions among surfaces of 3D objects. Pouchol et al. [PACT09] use spatial hashing to model the interaction between solid objects (e.g., spheres) and fluid. Choi et al. [CJC*09] use perfect spatial hashing to detect interference between characters and obstacles in a free space mapped virtual environment. Chentanez et al. [CMM16] use spatial hashing to detect and delete overlapping, or intersecting, triangles on the surface of 3D mesh volumes.

Adaptive mesh refinement (AMR): Tumblin et al. [TAHR15] use compact perfect hashing to search for neighboring cells in cell-based AMR for a shallow-water hydrodynamics simulation (e.g., AMR at the boundary of a water wave). Chentanez et al. [CMM16]

use spatial hashing to perform AMR on 3D mesh volumes, as vertices are advected in real-time.

Surface rendering: Lefebvre and Hoppe [LH06] use perfect spatial hashing to render the color surfaces of 3D volumetric textures. Alcantara et al. [ASA*09, AVS*12] (open-addressing cuckoo hashing), Garcia et al. [GLHL11] (open-addressing robin hood hashing), Nießner et al. [NZIS13] (spatial hashing), and Duan et al. [DLN*17] (spatial hashing) all perform real-time surface rendering and reconstruction of 3D volumes within voxelized grids. Bastos and Celes [BC08] use perfect hashing to perform isosurface rendering and morphing of adaptively sampled distance fields (ADFs). Kähler et al. [KPVM16] use spatial hashing to render voxelized 3D scene models of signed distance fields (SDFs).

Interactive drawing and painting: Lefebvre and Hoppe [LH06] use perfect spatial hashing to interactively paint over 3D volumetric textures. Garcia et al. [GLHL11] use open-addressing robin hood hashing to interactively draw on 2D surfaces, such as an atlas. Eyyurekli and Breen [EB11] use spatial hashing to interactively edit and draw over 3D level-set surfaces.

Database processing: Hetherington et al. [HOA15] and Choudhury et al. [CPK16] use open-addressing cuckoo hashing to cache most-recently used, or working set, queries in a key-value store. Karnagel et al. [KML15] use open-addressing linear probing to perform group-by and aggregation queries from a key-value store. Zhang et al. [ZWY*15] and Breslow et al. [BZG*16] use open addressing bucketized cuckoo hashing to accelerate queries and updates in key-value stores.

Similarity search: Zhou et al. [ZGJ*16] use open-addressing robin hood hashing to extract the top- k most similar matches for query records in real-world document and relational datasets. Alcantara et al. [ASA*09] use open-addressing cuckoo hashing to perform geometric hashing, which is a form of 2D image matching. Pan et al. [PLM10], Pan and Manocha [PM11], and Lukač and Žalik [LŽ15] each use locality-sensitive hashing to find the k approximate nearest neighbors (k ANN) of query points within multi-dimensional record sets. Pouchol et al. [PACT09] use spatial hashing to perform particle neighbor search within fluid and solid interaction simulations. Todd et al. [TTD*16] use multi-level bucketized hashing to identify genes with similar k -motifs, or DNA subsequences of length k .

5. Analysis and Future Work

This section analyzes the findings of the surveyed hashing techniques and identifies opportunities for future work. Table 1 enumerates a set of 17 hashing use case attributes and suggests the most-suitable or performant hashing technique(s) for each attribute. Due to the large number of possible subsets of use case attributes, a technique is only suggested for a single attribute. A practitioner can consult the table for a set of desired attributes, identify overlapping suggested techniques, and then investigate the suitability of these techniques for a specific task. Table 2 evaluates the most-suitable hashing techniques from Table 1 based on their ability to address optimal GPU performance criteria and utilize performant GPU hardware features. This evaluation assesses performance as it pertains to arbitrary access patterns for insertions and queries.

Thus, special cases such as empty queries or ordered accesses are not considered unless a technique is specifically designed to perform well for such cases; for example, *CoherentHash* [GLHL11] achieves best-in-class throughput and memory coalescing among open-addressing techniques only when coherence exists among input elements and their hash table locations. The GPU performance criteria and hardware features are described as follows:

- **Sufficient Parallelism:** The hashing technique experimentally demonstrates a sufficient throughput of insertion and query operations (millions per second) to hide global memory access latency.
- **Memory Coalescing:** All the threads in a warp access addresses within the same fetched cache line of contiguous memory. These memory requests are necessary to execute the given SIMT instruction.
- **Control Flow:** All the threads in a warp follow the same execution path for a SIMT instruction.
- **CPU↔GPU Data Transfers:** The hash table is constructed and/or stored in CPU memory and then accessed from or copied onto the GPU via the interconnection bus (e.g., PCI-e); thus, the hashing experiences data transfer bandwidth latency.
- **Shared Memory:** Per-thread-block GPU memory space that is smaller in size than global DRAM memory, but offers faster memory accesses.
- **Atomic Operations:** Lightweight hardware atomic functions, such as compare-and-swap (CAS), that guard and manage hash table entries during parallel insertions, probing evictions (e.g., in cuckoo hashing), and deletions.
- **Warp-wide Voting:** Lightweight functions used by all the threads in a warp to communicate data and perform collaborative execution, such as when all warp threads query the hash table for the same key.

For arbitrary, random access patterns, *CuckooHash2* cuckoo hashing [AVS*12] offers best-in-class throughput performance among the surveyed hashing techniques (subsection 3.1.2). This is due to the small constant number of probes necessary in both the best- and worst-case scenarios. In the worst-case insertion scenario of not finding an empty slot, the cuckoo hash table demonstrates fast reconstruction rates. In the presence of spatially-ordered access patterns, the *CoherentHash* robin hood hashing [GLHL11] achieves greater throughput than cuckoo hashing and is robust to higher load factors (subsection 3.1.5).

In the ideal, “fast-path,” scenario, an open-addressing technique only requires a single atomic CAS operation for an insertion and a single random global memory access for a query. However, in a typical scenario, a variable number of probes are needed to insert and query a key, often spanning non-contiguous regions of memory. This induces non-coalesced memory accesses and control flow divergence among threads of a warp. Thus, most of the open-addressing techniques assessed in Table 2 cannot guarantee to attain memory coalescing and control flow.

The combination of radix sorting and binary searching is a very effective alternative to searching via hashing when access patterns are ordered or the data is already in near-sorted order prior to sorting. However, for interactive use, this approach naively requires a

Use Case Attribute/Hashing Category	Open-Addressing	Perfect Hashing	Spatial Hashing	Separate Chaining
Access Patterns:				
<ul style="list-style-type: none"> – Ordered queries – Random insertions and queries – Duplicate insertions and queries – Query-heavy operation mix – Update-heavy operation mix – Unsuccessful (empty) queries 	<i>CoherentHash</i> [GLHL11] <i>CuckooHash2</i> [AVS*12] <i>HortonHash</i> [BZG*16]; <i>MemcachedGPU</i> [HOA15]	<i>PerfectHash</i> [LH06] <i>PerfectHash</i> [LH06]	<i>EGSH</i> [DLN*17]	<i>SlabHash</i> [AFO17]
Data Type:				
<ul style="list-style-type: none"> – Grid-based spatial primitives – Integer or index-based – Multi-dimensional attribute vector 	<i>CuckooHash2</i> [AVS*12]		<i>VoxelHash</i> [NZIS13] <i>BiLevelLSH</i> [PM11]	
Hash Table:				
<ul style="list-style-type: none"> – Collision-free – Fast construction – Dynamic – Low occupancy – High occupancy; maximum load 	<i>CuckooHash2</i> [AVS*12] <i>CoherentHash</i> [GLHL11]	<i>PerfectHash</i> [LH06]	<i>CompactHash</i> [TAHR15]	<i>SlabHash</i> [AFO17]
Hardware:				
<ul style="list-style-type: none"> – Use of CPU memory (PCIe bound) – Use of GPU shared memory – Efficient use of atomics 	<i>StadiumHash</i> [KBGB15]; <i>Horton-Hash</i> [BZG*16] <i>CuckooHash1</i> [ASA*09] <i>CuckooHash2</i> [AVS*12]		<i>BiLevelLSH</i> [PM11]	

Table 1: Suggested hashing technique(s) for different use case attributes. For each attribute, the most suitable or best-performing technique from one or more of the four hashing categories is denoted. Additional details regarding a technique can be found within the section of its encompassing hashing category.

re-sort of a larger array each time new data is added. Additional research is needed to investigate more-efficient data-parallel schemes for accommodating dynamic data.

If data will be updated at run-time, then *SlabHash* [AFO17] offers best-in-class dynamic hashing, achieving a significant increase in throughput over cuckoo hashing, which must be reconstructed

after each batch of updates (section 3.4). Moreover, as seen in Table 2, this technique addresses each of the criteria for optimal GPU performance. Further research is needed to compare the performance of slab hashing with that of *CoherentHash* robin hood hashing [GLHL11] in the presence of coherent access patterns.

When data must be stored and accessed off-device in CPU

Hashing Technique	GPU Performance Criteria				GPU Hardware Features		
	Sufficient Parallelism	Memory Coalescing	Control Flow	CPU↔GPU Data Transfers	Shared Memory	Atomic Operations	Warp-wide Voting
Open-addressing:							
– <i>CoherentHash</i> [GLHL11]	✓	✓	×	×	×	✓	×
– <i>CuckooHash1</i> [ASA*09]	✓	×	×	×	✓	✓	×
– <i>CuckooHash2</i> [AVS*12]	✓	×	×	×	×	✓	×
– <i>HortonHash</i> [BZG*16]	✓	×	✓	✓	×	×	×
– <i>MemcachedGPU</i> [HOA15]	✓	×	×	✓	✓	✓	×
– <i>StadiumHash</i> [KBGB15]	✓	×	×	✓	✓	✓	✓
Perfect Hashing:							
– <i>PerfectHash</i> [LH06]	✓	✓	✓	✓	×	×	×
Spatial Hashing:							
– <i>BiLevelLSH</i> [PM11]	✓	×	×	×	✓	✓	×
– <i>EGSH</i> [DLN*17]	✓	×	✓	✓	×	×	×
– <i>VoxelHash</i> [NZIS13]	✓	×	×	×	×	✓	×
Separate Chaining:							
– <i>SlabHash</i> [AFO17]	✓	✓	✓	×	×	✓	✓

Table 2: Select hashing techniques and their ability to address GPU criteria for optimal performance and utilize performant GPU hardware features. The techniques are grouped by category and represent the subset of techniques that are identified as highly-suitable for different use-case attributes in Table 1.

memory, the use of ticketing, or key bit signatures, is beneficial to save expensive accesses for obvious non-matches during probing/querying. Future hashing approaches should assess the performance benefits of ticketing even when off-device accesses do not occur. Maintaining the ticketing structure in shared memory appears to be particularly beneficial, as demonstrated by the *StadiumHash* open-addressing technique [KBGB15].

Regardless of the data use case, shared memory should be leveraged as much as possible to perform warp operations and faster memory accesses (not necessarily coalesced). This is facilitated by sizing buckets to the size of a thread block, such as in *CuckooHash1* cuckoo hashing [ASA*09]. If data must be accessed outside of shared memory, warps should be modeled as collaborative processing units the size of a memory transaction. Each thread is assigned to an entry within the loaded cache line and all threads then compare their entries (possibly empty) to the query or insert key via a warp-wide voting function. *CuckooHash1* [ASA*09], *StadiumHash* [KBGB15], and *SlabHash* [AFO17] make particularly good use of shared memory and warp-wide voting (Table 2).

Fast hash table construction enables larger load factors, acceptance of insertion failure, and dynamic usage in interactive applications. CPU-constructed hash tables face two bottlenecks: slower construction on the CPU and copying over the PCIe bus to the GPU. Both bottlenecks render these tables infeasible for updates or reconstructions. From Table 2, the *HortonHash* [BZG*16], *PerfectHash* [LH06], and *EGSH* [DLN*17] techniques are bandwidth-bound by the transfer of the hash table from CPU to GPU prior to querying. Additionally, *MemcachedGPU* [HOA15] and *StadiumHash* [KBGB15] must service data transfers during querying, as hash table data resides on the CPU. Further research is needed

to redesign CPU-constructed hash tables for efficient data-parallel construction on the GPU.

Perfect hashing (section 3.2), *PerfectHash* [LH06], avoids collision resolution, but is not well-suited for updates, since the hash table must be reconstructed on the CPU and remain PCIe bandwidth-bound. A trade-off arises: either use multiple separate hash tables (and multiple probes), or use a single addressable hash table and construct the offset table, which is the primary bottleneck during construction. Further research towards constructing the offset table in data-parallel on the GPU is needed to make perfect hashing a more dynamic, interactive solution.

Compact spatial hashing (subsection 3.3.1), *CompactHash* [TAHR15], offers the useful feature of downsizing a perfect hash table that contains a significant number of unused entries, which arises often in spatial hashing. This comes with the trade-off of new hash collisions that must be resolved. Further research should assess the viability of this approach for other types of hash tables and varying load factors.

The *BiLevelLSH* [PM11] locality-sensitive hashing technique takes advantage of fast on-device data-parallel operations to sort key-value pairs and hash them into a cuckoo hash table. Further work is needed to design a dynamic variant that supports updates to the hash table and sorted key-values. Moreover, future research should investigate the use of LSH for *approximate* surface rendering and reconstruction tasks. For instance, instead of querying the data to render for each point in a grid, select points can be queried and return, in a single operation, the approximate data for an entire k -point bounding box in the form of the k ANN.

Finally, prospective avenues for future research exist for a *Hash-Fight* technique that is introduced by Lessley et al. [LBMC16, LMLC17] as part of a platform-portable, GPU-compatible hash-

ing approach. This approach employs an iterative winner-takes-all collision resolution technique that does not use hardware atomic primitives to synchronize writes to the hash table. Instead, race conditions are a fundamental and non-detrimental feature of resolving collisions. However, *HashFight* does not maintain a persistent hash table, but rather reconstructs a new, smaller-sized table at the beginning of each iteration. Thus, additional work is needed to expand the technique to support query and insert operations, with accompanying throughput analyses. Then, the build speed of *HashFight* can be compared against the build times of the best-in-class static cuckoo and robin hood hashing techniques, particularly *CuckooHash2* [AVS*12] and *CoherentHash* [GLHL11].

6. Conclusion

This paper provides a survey of parallel hashing techniques for GPU architectures. These techniques are categorized according to the method of collision resolution: open-addressing, perfect hashing, spatial hashing, and separate chaining. Each of the surveyed studies offer various design choices and patterns that help inform a more-general set of best practices for performant hashing on the GPU. These best practices and the most-suitable hashing techniques for different use-case factors are analyzed and used to reveal opportunities for future research.

References

- [ADMO16] ASHKIANI S., DAVIDSON A., MEYER U., OWENS J. D.: Gpu multisplit. In *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (New York, NY, USA, 2016), PPOPP '16, ACM, pp. 12:1–12:13. URL: <http://doi.acm.org/10.1145/2851141.2851169>, doi:10.1145/2851141.2851169. 5, 6, 11
- [AFO17] ASHKIANI S., FARACH-COLTON M., OWENS J. D.: A Dynamic Hash Table for the GPU. *ArXiv e-prints* (Oct. 2017). arXiv:1710.11246. 7, 19, 21, 22
- [ALF*17] ASHKIANI S., LI S., FARACH-COLTON M., AMENTA N., OWENS J. D.: GPU LSM: A dynamic dictionary data structure for the GPU. *CoRR abs/1707.05354* (2017). URL: <http://arxiv.org/abs/1707.05354>, arXiv:1707.05354. 6
- [Amd67] AMDAHL G. M.: Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, Spring Joint Computer Conference* (New York, NY, USA, 1967), AFIPS '67 (Spring), ACM, pp. 483–485. 2
- [ASA*09] ALCANTARA D. A., SHARF A., ABBASINEJAD F., SENGUPTA S., MITZENMACHER M., OWENS J. D., AMENTA N.: Real-time parallel hashing on the gpu. In *ACM SIGGRAPH Asia 2009 Papers* (New York, NY, USA, 2009), SIGGRAPH Asia '09, ACM, pp. 154:1–154:9. 5, 6, 7, 9, 10, 11, 17, 19, 20, 21, 22
- [AVS*12] ALCANTARA D. A., VOLKOV V., SENGUPTA S., MITZENMACHER M., OWENS J. D., AMENTA N.: Chapter 4 - Building an efficient hash table on the {GPU}. In *{GPU} Computing Gems Jade Edition*, Hwu W.-m. W., (Ed.), Applications of GPU Computing Series. Morgan Kaufmann, Boston, 2012, pp. 39 – 53. 5, 6, 7, 10, 12, 20, 21, 22, 23
- [BC08] BASTOS T., CELES W.: Gpu-accelerated adaptively sampled distance fields. In *2008 IEEE International Conference on Shape Modeling and Applications* (June 2008), pp. 171–178. 14, 20
- [Ble90] BLELLOCH G. E.: *Vector models for data-parallel computing*, vol. 75. MIT press Cambridge, 1990. 2, 5
- [Boo03] BOOST C++ LIBRARIES: Boost.Iterator Library, 2003. http://www.boost.org/doc/libs/1_65_1/libs/iterator/doc/index.html. 5
- [Bor14] BORDAWEKAR R.: Evaluation of parallel hashing techniques. In *GPU Technology Conference* (Mar. 2014). 8, 9, 10
- [Bre74] BRENT R. P.: The parallel evaluation of general arithmetic expressions. *J. ACM* 21, 2 (Apr. 1974), 201–206. 2
- [BS05] BLIKBERG R., SÄYREVIK T.: Load balancing and openmp implementation of nested parallelism. *Parallel Computing* 31, 10 (2005), 984 – 998. OpenMP. 2
- [BZ07] BOTELHO F. C., ZIVIANI N.: External perfect hashing for very large key sets. In *Proceedings of the Sixteenth ACM Conference on Conference on Information and Knowledge Management* (New York, NY, USA, 2007), CIKM '07, ACM, pp. 653–662. 1
- [BZG*16] BRESLOW A. D., ZHANG D. P., GREATHOUSE J. L., JAYASENA N., TULLSEN D. M.: Horton tables: Fast hash tables for in-memory data-intensive computing. In *Proceedings of the 2016 USENIX Conference on Usenix Annual Technical Conference* (Berkeley, CA, USA, 2016), USENIX ATC '16, USENIX Association, pp. 281–294. URL: <http://dl.acm.org/citation.cfm?id=3026959.3026986>. 11, 20, 21, 22
- [CCT12] CEDERMAN D., CHATTERJEE B., TSIGAS P.: Understanding the performance of concurrent data structures on graphics processors. In *Proceedings of the 18th International Conference on European Parallel Processing* (August 2012), Euro-Par 2012, pp. 883–894. 6
- [Cel86] CELIS P.: *Robin Hood Hashing*. PhD thesis, Waterloo, Ont., Canada, Canada, 1986. 12
- [CHM97] CZECH Z. J., HAVAS G., MAJEWSKI B. S.: Perfect hashing. *Theoretical Computer Science* 182, 1 (1997), 1 – 143. 1, 13
- [CJC*09] CHOI M. G., JU E., CHANG J.-W., LEE J., KIM Y. J.: Link-less octree using multi-level perfect hashing. *Comput. Graph. Forum* 28 (2009), 1773–1780. 14, 19
- [CKWT14] CHENG L., KOTOULOS S., WARD T. E., THEODOROPOULOS G.: Design and evaluation of parallel hashing over large-scale data. In *2014 21st International Conference on High Performance Computing (HiPC)* (Dec 2014), pp. 1–10. 1
- [CMM16] CHENTANEZ N., MÜLLER M., MACKLIN M.: GPU accelerated grid-free surface tracking. *Computers & Graphics* 57, Supplement C (2016), 1 – 11. URL: <http://www.sciencedirect.com/science/article/pii/S0097849316300152>, doi:https://doi.org/10.1016/j.cag.2016.03.002. 15, 19
- [CPK16] CHOUDHURY Z., PURINI S., KRISHNA S. R.: A hybrid cpu+gpu working-set dictionary. In *2016 15th International Symposium on Parallel and Distributed Computing (ISPDC)* (July 2016), pp. 56–63. 7, 20
- [CSRL01] CORMEN T. H., STEIN C., RIVEST R. L., LEISERSON C. E.: *Introduction to Algorithms*, 2nd ed. McGraw-Hill Higher Education, 2001. 6, 7, 8, 11
- [cud17] CUDA Data Parallel Primitives Library. <http://cudpp.github.io>, Nov. 2017. 5, 6, 10, 11
- [DF08] DASGUPTA S., FREUND Y.: Random projection trees and low dimensional manifolds. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (New York, NY, USA, 2008), STOC '08, ACM, pp. 537–546. 17
- [DHM13] DICE D., HENDLER D., MIRSKY I.: Lightweight contention management for efficient compare-and-swap operations. In *Proceedings of the 19th International Conference on Parallel Processing* (Berlin, Heidelberg, 2013), Euro-Par '13, Springer-Verlag, pp. 595–606. 2
- [DIIM04] DATAR M., IMMORLICA N., INDYK P., MIRROKNI V. S.: Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the Twentieth Annual Symposium on Computational Geometry* (New York, NY, USA, 2004), SCG '04, ACM, pp. 253–262. 16, 17

- [DLN*17] DUAN W., LUO J., NI G., TANG B., HU Q., GAO Y.: Exclusive grouped spatial hashing. *Computers & Graphics* (2017). doi:<https://doi.org/10.1016/j.cag.2017.08.012>. 16, 20, 21, 22
- [DTGO12] DAVIDSON A., TARJAN D., GARLAND M., OWENS J. D.: Efficient parallel merge sort for fixed and variable length keys. In *Innovative Parallel Computing* (May 2012), p. 9. 5, 6
- [EB11] EYIYUREKLI M., BREEN D. E.: Data structures for interactive high resolution level-set surface editing. In *Proceedings of Graphics Interface 2011* (School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2011), GI '11, Canadian Human-Computer Communications Society, pp. 95–102. 20
- [EL07] EITZ M., LIXU G.: Hierarchical spatial hashing for real-time collision detection. In *Shape Modeling and Applications, 2007. SMI '07. IEEE International Conference on* (June 2007), pp. 61–70. doi:10.1109/SMI.2007.18. 15, 19
- [EMM06] ERLINGSSON Å., MANASSE M., MCSHERRY F.: A cool and practical alternative to traditional hash tables. In *7th Workshop on Distributed Data and Structures (WDAS'06)* (Santa Clara, CA, January 2006), pp. 1–6. 10
- [FHCD92] FOX E. A., HEATH L. S., CHEN Q. F., DAOUD A. M.: Practical minimal perfect hash functions for large databases. *Commun. ACM* 35, 1 (Jan. 1992), 105–121. 13
- [FKS84] FREDMAN M. L., KOMLÓS J., SZEMERÉDI E.: Storing a sparse table with $O(1)$ worst case access time. *J. ACM* 31, 3 (June 1984), 538–544. 9, 13
- [Fly72] FLYNN M. J.: Some computer organizations and their effectiveness. *IEEE Trans. Comput.* 21, 9 (Sept. 1972), 948–960. 2
- [GHOI14] GIESEKE F., HEINERMANN J., OANCEA C., IGEL C.: Buffer k-d trees: Processing massive nearest neighbor queries on GPUs. 172–180. 17
- [GHS*10] GOODMAN E. L., HAGLIN D. J., SCHERRER C., CHAVARRÍA-MIRANDA D., MOGILL J., FEO J.: Hashing strategies for the cray xmt. In *2010 IEEE International Symposium on Parallel Distributed Processing, Workshops and Phd Forum (IPDPSW)* (April 2010), pp. 1–8. 1, 18
- [GLHL11] GARCÍA I., LEFEBVRE S., HORNUS S., LASRAM A.: Coherent parallel hashing. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 161:1–161:8. 7, 12, 13, 20, 21, 22, 23
- [Gre02] GREENWALD M.: Two-handed emulation: How to build non-blocking implementations of complex data-structures using dcas. In *Proceedings of the Twenty-first Annual Symposium on Principles of Distributed Computing* (New York, NY, USA, 2002), PODC '02, ACM, pp. 260–269. 1, 18
- [GTW15] GAO H., TANG J., WU G.: Parallel surface reconstruction on gpu. In *Proceedings of the 7th International Conference on Internet Multimedia Computing and Service* (New York, NY, USA, 2015), ICIMCS '15, ACM, pp. 54:1–54:5. 6
- [Gus88] GUSTAFSON J. L.: Reevaluating amdahl's law. *Commun. ACM* 31, 5 (May 1988), 532–533. 2
- [HAP12] HE X., AGARWAL D., PRASAD S. K.: Design and implementation of a parallel priority queue on many-core architectures. In *2012 19th International Conference on High Performance Computing* (Dec 2012), pp. 1–10. 6
- [Har14] HARRIS M.: Maxwell: The Most Advanced CUDA GPU Ever Made. <https://devblogs.nvidia.com/parallelforall/maxwell-most-advanced-cuda-gpu-ever-made/>, 2014. 3, 4
- [Her91] HERLIHY M.: Wait-free synchronization. *ACM Trans. Program. Lang. Syst.* 13, 1 (Jan. 1991), 124–149. 2
- [HOA15] HETHERINGTON T. H., O'CONNOR M., AAMODT T. M.: Memcachedgpu: Scaling-up scale-out key-value stores. In *Proceedings of the Sixth ACM Symposium on Cloud Computing* (New York, NY, USA, 2015), SoCC '15, ACM, pp. 43–57. 11, 20, 21, 22
- [HRJ*10] HUANG X., RODRIGUES C. I., JONES S., BUCK I., M. HWU W.: XMalloc: A scalable lock-free dynamic memory allocator for many-core machines. In *2010 10th IEEE International Conference on Computer and Information Technology* (June 2010), pp. 1134–1139. 19
- [IM98] INDYK P., MOTWANI R.: Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing* (New York, NY, USA, 1998), STOC '98, ACM, pp. 604–613. 16, 17
- [Int17] INTEL CORPORATION: Introducing the Intel Threading Building Blocks, May 2017. <https://software.intel.com/en-us/node/506042>. 5
- [JR15] JEFFERS J., REINDERS J.: *High Performance Parallelism Pearls Volume Two: Multicore and Many-core Programming Approaches*, 1st ed., vol. 2. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2015. 2
- [Kar12] KARRAS T.: Maximizing Parallelism in the Construction of BVHs, Octrees, and k-d Trees. In *Eurographics/ ACM SIGGRAPH Symposium on High Performance Graphics* (2012), Dachsbacher C., Munkberg J., Pantaleoni J., (Eds.), The Eurographics Association, pp. 33–37. 6
- [KGBG15] KHORASANI F., BELVIRANLI M. E., GUPTA R., BHUYAN L. N.: Stadium hashing: Scalable and flexible hashing on gpus. In *2015 International Conference on Parallel Architecture and Compilation (PACT)* (Oct 2015), pp. 63–74. 11, 21, 22
- [KBS11] KALOJANOV J., BILLETER M., SLUSALLEK P.: Two-level grids for ray tracing on gpus. *Computer Graphics Forum* 30, 2 (2011), 307–314. 6
- [KCS*10] KIM C., CHHUGANI J., SATISH N., SEDLAR E., NGUYEN A. D., KALDEWEY T., LEE V. W., BRANDT S. A., DUBEY P.: Fast: Fast architecture sensitive tree search on modern cpus and gpus. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data* (New York, NY, USA, 2010), SIGMOD '10, ACM, pp. 339–350. 6
- [KDB12] KALDEWEY T., DI BLAS A.: Large-scale gpu search. 3–14. 6
- [KML15] KARNAGEL T., MUELLER R., LOHMAN G. M.: Optimizing gpu-accelerated group-by and aggregation. In *ADMS@VLDB* (2015). 8, 20
- [Knu98] KNUTH D. E.: *The Art of Computer Programming, Volume 3: (2nd Ed.) Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1998. 1
- [KPV16] KÄHLER O., PRISACARIU V., VALENTIN J., MURRAY D.: Hierarchical voxel block hashing for efficient integration of depth images. *IEEE Robotics and Automation Letters* 1, 1 (Jan 2016), 192–197. doi:10.1109/LRA.2015.2512958. 15, 20
- [KRD*03] KAPASI U. J., RIXNER S., DALLY W. J., KHAILANY B., AHN J. H., MATTSON P., OWENS J. D.: Programmable stream processors. *Computer* 36, 8 (Aug. 2003), 54–62. 3
- [KS09] KALOJANOV J., SLUSALLEK P.: A parallel algorithm for construction of uniform grids. In *Proceedings of the Conference on High Performance Graphics 2009* (New York, NY, USA, 2009), HPG '09, ACM, pp. 23–28. 6
- [KU86] KARLIN A. R., UPFAL E.: Parallel hashing—an efficient implementation of shared memory. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing* (New York, NY, USA, 1986), STOC '86, ACM, pp. 160–168. 1
- [KZ09] KUANG Q., ZHAO L.: A practical GPU based KNN algorithm. In *Proceedings of the Second Symposium on International Computer Science and Computational Technology (ISCSCT '09)* (Dec. 2009), Academy Publisher, pp. 151–155. 16
- [LA15] LI S., AMENTA N.: Brute-force k-nearest neighbors search on the gpu. In *Proceedings of the 8th International Conference on Similarity Search and Applications - Volume 9371* (New York, NY, USA, 2015), SISAP 2015, Springer-Verlag New York, Inc., pp. 259–270. 16

- [Lam78] LAMPORT L.: Time, clocks, and the ordering of events in a distributed system. *Commun. ACM* 21, 7 (July 1978), 558–565. 2
- [LBMC16] LESSLEY B., BINYAHIB R., MAYNARD R., CHILDS H.: External Facelist Calculation with Data-Parallel Primitives. In *Proceedings of EuroGraphics Symposium on Parallel Graphics and Visualization (EGPGV)* (Groningen, The Netherlands, June 2016), pp. 10–20. 6, 22
- [LD08] LAGAE A., DUTRÉ P.: Compact, fast and robust grids for ray tracing. In *ACM SIGGRAPH 2008 Talks* (New York, NY, USA, 2008), SIGGRAPH '08, ACM, pp. 20:1–20:1. 6
- [LGS*09] LAUTERBACH C., GARLAND M., SENGUPTA S., LUEBKE D., MANOCHA D.: Fast bhv construction on gpus. *Computer Graphics Forum* 28, 2 (2009), 375–384. 6, 17
- [LH06] LEFEBVRE S., HOPPE H.: Perfect spatial hashing. In *ACM SIGGRAPH 2006 Papers* (New York, NY, USA, 2006), SIGGRAPH '06, ACM, pp. 579–588. 7, 9, 10, 13, 14, 16, 19, 20, 21, 22
- [Lit11] LITTLE J. D. C.: Or forum—little’s law as viewed on its 50th anniversary. *Oper. Res.* 59, 3 (May 2011), 536–549. URL: <http://dx.doi.org/10.1287/opre.1110.0940>, doi:10.1287/opre.1110.0940. 2, 4
- [LJW*07] LV Q., JOSEPHSON W., WANG Z., CHARIKAR M., LI K.: Multi-probe lsh: Efficient indexing for high-dimensional similarity search. In *Proceedings of the 33rd International Conference on Very Large Data Bases* (2007), VLDB '07, VLDB Endowment, pp. 950–961. 17
- [LMLC17] LESSLEY B., MORELAND K., LARSEN M., CHILDS H.: Techniques for Data-Parallel Searching for Duplicate Elements. In *Proceedings of IEEE Symposium on Large Data Analysis and Visualization (LDAV)* (Phoenix, AZ, Oct. 2017), pp. 1–5. 6, 22
- [LWL12] LUO L., WONG M. D. F., LEONG L.: Parallel implementation of r-trees on the gpu. In *17th Asia and South Pacific Design Automation Conference* (Jan 2012), pp. 353–358. doi:10.1109/ASPAC.2012.6164973. 6
- [LŽ15] LUKAČ N., ŽALIK B.: *Fast Approximate k-Nearest Neighbours Search Using GPGPU*. Springer Singapore, Singapore, 2015, pp. 221–234. URL: https://doi.org/10.1007/978-981-287-134-3_14, doi:10.1007/978-981-287-134-3_14. 17, 20
- [MC12] MISRA P., CHAUDHURI M.: Performance evaluation of concurrent lock-free data structures on gpus. In *Proceedings of the 2012 IEEE 18th International Conference on Parallel and Distributed Systems* (Washington, DC, USA, 2012), ICPADS '12, IEEE Computer Society, pp. 53–60. 18, 19
- [MCP17] MOSCOVICI N., COHEN N., PETRANK E.: A GPU-friendly skiplist algorithm. In *2017 26th International Conference on Parallel Architectures and Compilation Techniques (PACT)* (Sept 2017), pp. 246–259. 6, 17
- [Meh82] MEHLHORN K.: On the program size of perfect and universal hash functions. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)* (Nov 1982), pp. 170–175. doi:10.1109/SFCS.1982.80. 13
- [MG10] MERRILL D. G., GRIMSHAW A. S.: Revisiting sorting for gpgpu stream architectures. In *Proceedings of the 19th International Conference on Parallel Architectures and Compilation Techniques* (New York, NY, USA, 2010), PACT '10, ACM, pp. 545–546. URL: <http://doi.acm.org/10.1145/1854273.1854344>, doi:10.1145/1854273.1854344. 5, 6, 10, 11, 17
- [Mic02] MICHAEL M. M.: High performance dynamic lock-free hash tables and list-based sets. In *Proceedings of the Fourteenth Annual ACM Symposium on Parallel Algorithms and Architectures* (New York, NY, USA, 2002), SPAA '02, ACM, pp. 73–82. URL: <http://doi.acm.org/10.1145/564870.564881>, doi:10.1145/564870.564881. 1, 18
- [ML75] MAURER W. D., LEWIS T. G.: Hash table methods. *ACM Comput. Surv.* 7, 1 (Mar. 1975), 5–19. 1
- [MPS92] MUNRO J. I., PAPADAKIS T., SEDGEWICK R.: Deterministic skip lists. In *Proceedings of the Third Annual ACM-SIAM Symposium on Discrete Algorithms* (Philadelphia, PA, USA, 1992), SODA '92, Society for Industrial and Applied Mathematics, pp. 367–375. 17
- [MRR12] MCCOOL M., REINDERS J., ROBISON A.: *Structured Parallel Programming: Patterns for Efficient Computation*, 1st ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2012. 2
- [MS12] MOAZENI M., SARRAFZADEH M.: Lock-free hash table on graphics processors. In *2012 Symposium on Application Accelerators in High Performance Computing* (July 2012), pp. 133–136. doi:10.1109/SAHPC.2012.25. 18
- [MSU*16] MORELAND K., SEWELL C., USHER W., LO L., MEREDITH J., PUGMIRE D., KRESS J., SCHROOTS H., MA K.-L., CHILDS H., LARSEN M., CHEN C.-M., MAYNARD R., GEVECI B.: VTK-m: Accelerating the Visualization Toolkit for Massively Threaded Architectures. *IEEE Computer Graphics and Applications (CG&A)* 36, 3 (May/June 2016), 48–58. 5
- [Nvi17a] NVIDIA CORPORATION: CUDA C Best Practices Guide. <http://docs.nvidia.com/cuda/cuda-c-best-practices-guide/index.html>, 2017. 3, 4
- [Nvi17b] NVIDIA CORPORATION: CUDA C Programming Guide. <http://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>, 2017. 2, 3, 4, 5, 18
- [Nvi17c] NVIDIA CORPORATION: Parallel Thread Execution ISA Version 6.0. <http://docs.nvidia.com/cuda/parallel-thread-execution/index.html>, 2017. 3, 4
- [Nvi17d] NVIDIA CORPORATION: Thrust, Nov. 2017. <http://thrust.github.io>. 5, 6
- [NZIS13] NIESSNER M., ZOLLHÖFER M., IZADI S., STAMMINGER M.: Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (TOG)* (2013). 7, 15, 20, 21, 22
- [OLG*07] OWENS J. D., LUEBKE D., GOVINDARAJU N., HARRIS M., KRÄUGER J., LEFOHN A. E., PURCELL T. J.: A survey of general-purpose computation on graphics hardware. *Computer Graphics Forum* 26, 1 (2007), 80–113. doi:10.1111/j.1467-8659.2007.01012.x. 6
- [PACT09] POUCHOL M., AHMAD A., CRESPIAN B., TERRAZ O.: A hierarchical hashing scheme for nearest neighbor search and broad-phase collision detection. *Journal of Graphics, GPU, and Game Tools* 14, 2 (2009), 45–59. doi:10.1080/2151237X.2009.10129281. 15, 19, 20
- [PGB*05] PEIERLS T., GOETZ B., BLOCH J., BOWBEER J., LEA D., HOLMES D.: *Java Concurrency in Practice*. Addison-Wesley Professional, 2005. 1, 18
- [PH08] PATTERSON D. A., HENNESSY J. L.: *Computer Organization and Design, Fourth Edition, Fourth Edition: The Hardware/Software Interface (The Morgan Kaufmann Series in Computer Architecture and Design)*, 4th ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008. 2, 3, 4
- [PLM10] PAN J., LAUTERBACH C., MANOCHA D.: Efficient nearest-neighbor computation for gpu-based motion planning. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Oct 2010), pp. 2243–2248. doi:10.1109/IROS.2010.5651449. 17, 20
- [PLMS00] PLAUGER P., LEE M., MUSSER D., STEPANOV A. A.: *C++ Standard Template Library*, 1st ed. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2000. 5
- [PM11] PAN J., MANOCHA D.: Fast gpu-based locality sensitive hashing for k-nearest neighbor computation. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems* (New York, NY, USA, 2011), GIS '11, ACM, pp. 211–220. 17, 20, 21, 22

- [PR04] PAGH R., RODLER F. F.: Cuckoo hashing. *J. Algorithms* 51, 2 (May 2004), 122–144. 9
- [PR13] POLYCHRONIOU O., ROSS K. A.: High throughput heavy hitter aggregation for modern simd processors. In *Proceedings of the Ninth International Workshop on Data Management on New Hardware* (New York, NY, USA, 2013), DaMoN '13, ACM, pp. 6:1–6:6. 7
- [SF15] SCOGLAND T. R., FENG W.-C.: Design and evaluation of scalable concurrent queues for many-core architectures. In *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering* (New York, NY, USA, 2015), ICPE '15, ACM, pp. 63–74. 6
- [SGL09] SCHLEGEL B., GEMULLA R., LEHNER W.: K-ary search on modern processors. In *Proceedings of the Fifth International Workshop on Data Management on New Hardware* (New York, NY, USA, 2009), DaMoN '09, ACM, pp. 52–60. 6
- [SHG09] SATISH N., HARRIS M., GARLAND M.: Designing efficient sorting algorithms for manycore gpus. In *Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Processing* (Washington, DC, USA, 2009), IPDPS '09, IEEE Computer Society, pp. 1–10. URL: <https://doi.org/10.1109/IPDPS.2009.5161005>, doi:10.1109/IPDPS.2009.5161005. 5, 6, 10
- [SJC17] SINGH D. P., JOSHI I., CHOUDHARY J.: Survey of gpu based sorting algorithms. *International Journal of Parallel Programming* (Apr 2017). 6
- [SKKS12] STEINBERGER M., KENZEL M., KAINZ B., SCHMALSTIEG D.: ScatterAlloc: Massively parallel dynamic memory allocation for the GPU. In *2012 Innovative Parallel Computing (InPar)* (May 2012), pp. 1–10. 18
- [SLM04] SCHROEDER W. J., LORENSEN B., MARTIN K.: *The Visualization Toolkit: An object-oriented approach to 3D graphics*. Kitware, 2004. 15
- [SO11] STUART J. A., OWENS J. D.: Efficient synchronization primitives for GPUs. *CoRR abs/1110.4623*, 1110.4623v1 (Oct. 2011). [arXiv:1110.4623v1](https://arxiv.org/abs/1110.4623v1). 4, 18
- [SR17] SCHNEIDER J., RAUTEK P.: A versatile and efficient gpu data structure for spatial indexing. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 911–920. 6, 14
- [SS06] SHALEV O., SHAVIT N.: Split-ordered lists: Lock-free extensible hash tables. *J. ACM* 53, 3 (May 2006), 379–405. 1, 18
- [STKT06] SUZUKI K., TONIEN D., KUROSAWA K., TOYOTA K.: *Birth-day Paradox for Multi-collisions*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 29–40. 6
- [TAHR15] TUMBLIN R., AHRENS P., HARTSE S., ROBEY R. W.: Parallel compact hash algorithms for computational meshes. *SIAM Journal on Scientific Computing* 37, 1 (2015), C31–C53. doi:10.1137/13093371X. 15, 19, 21, 22
- [THM*03] TESCHNER M., HEIDELBERGER B., MUELLER M., POMERANETS D., GROSS M.: Optimized spatial hashing for collision detection of deformable objects. *Proceedings of Vision, Modeling, Visualization (VMV 2003)* (2003), 47–54. 15, 19
- [TTD*16] TODD A., TRUONG H., DETERS J., LONG J., CONANT G., BECCHI M.: Parallel gene upstream comparison via multi-level hash tables on gpu. In *2016 IEEE 22nd International Conference on Parallel and Distributed Systems (ICPADS)* (Dec 2016), pp. 1049–1058. 20
- [ULL72] ULLMAN J. D.: A note on the efficiency of hashing functions. *J. ACM* 19, 3 (July 1972), 569–575. 1
- [VH15] VINKLER M., HAVRAN V.: Register efficient dynamic memory allocator for gpus. *Computer Graphics Forum*, 8 (2015), 143–154. 19
- [Vol16] VOLKOV V.: *Understanding Latency Hiding on GPUs*. PhD thesis, EECS Department, University of California, Berkeley, Aug 2016. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-143.html>. 4
- [vtk17] VTK-m. <https://gitlab.kitware.com/vtk/vtk-m>, Nov. 2017. 5
- [WBS*14] WIDANAGAMAACHCHI W., BREMER P. T., SEWELL C., LO L. T., AHRENS J., PASCUCCIK V.: Data-parallel halo finding with variable linking lengths. In *2014 IEEE 4th Symposium on Large Data Analysis and Visualization (LDAV)* (Nov 2014), pp. 27–34. 6
- [WLKC16] WANG J., LIU W., KUMAR S., CHANG S. F.: Learning to hash for indexing big data—A survey. *Proceedings of the IEEE* 104, 1 (Jan 2016), 34–57. doi:10.1109/JPROC.2015.2487976. 1, 17
- [WSSJ14] WANG J., SHEN H. T., SONG J., JI J.: Hashing for similarity search: A survey. *CoRR abs/1408.2927* (2014). URL: <http://arxiv.org/abs/1408.2927>, arXiv:1408.2927. 1, 17
- [YHGT10] YANG J. C., HENSLEY J., GRÜN H., THIBIEROZ N.: Real-time concurrent linked list construction on the gpu. In *Proceedings of the 21st Eurographics Conference on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2010), EGSR'10, Eurographics Association, pp. 1297–1304. 6
- [ZGHG11] ZHOU K., GONG M., HUANG X., GUO B.: Data-parallel octrees for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics* 17, 5 (May 2011), 669–681. 6
- [ZGJ*16] ZHOU J., GUO Q., JAGADISH H. V., LUAN W., TUNG A. K. H., YANG Y., ZHENG Y.: Generic inverted index on the GPU. *CoRR abs/1603.08390* (2016). URL: <http://arxiv.org/abs/1603.08390>, arXiv:1603.08390. 13, 20
- [ZHWG08] ZHOU K., HOU Q., WANG R., GUO B.: Real-time kd-tree construction on graphics hardware. In *ACM SIGGRAPH Asia 2008 Papers* (New York, NY, USA, 2008), SIGGRAPH Asia '08, ACM, pp. 126:1–126:11. 6
- [ZWY*15] ZHANG K., WANG K., YUAN Y., GUO L., LEE R., ZHANG X.: Mega-kv: A case for gpus to maximize the throughput of in-memory key-value stores. *Proc. VLDB Endow.* 8, 11 (July 2015), 1226–1237. URL: <http://dx.doi.org/10.14778/2809974.2809984>, doi:10.14778/2809974.2809984. 7, 10, 11, 20