Efficient Parallel Zonal Statistics on Large-Scale Global Biodiversity Data on GPUs

Jianting Zhang
Dept. of Computer Science
City College of New York
New York City, NY, 10031

jzhang@cs.ccny.cuny.edu

Simin You
Dept. of Computer Science
CUNY Graduate Center
New York, NY, 10016

syou@gc.cuny.edu

Le Gruenwald
School of Computer Science
University of Oklahoma
Norman, OK 73071

ggruenwald@ou.edu

ABSTRACT

Analyzing how species are distributed on the Earth has been one of the fundamental questions in the intersections of environmental sciences, geosciences and biological sciences. With world-wide data contributions, more than 375 million species occurrence records for nearly 1.5 million species have been deposited into the Global Biodiversity Information Facility (GBIF) data portal. The sheer amounts of point and polygon data and the computation-intensive point-in-polygon tests for zonal statistics for biodiversity studies have imposed significant technical challenges. In this study, we have developed efficient techniques to enable parallel zonal statistics on the global GBIF data completely on GPUs with limited memory capacity. Experiment results have shown that an impressive end-to-end response time under 100 seconds can be achieved for zonal statistics on the 375+ million species records over 15+ thousand global eco-regions with 4+ million vertices on a single Nvidia Quadro 6000 GPU device. The achieved high performance, which is several orders of magnitude faster than reference serial implementations using traditional open source geospatial techniques, not only demonstrates the potential of GPU computing for large scale geospatial processing, but also makes interactive query driven visual exploration of global biodiversity data possible.

1. INTRODUCTION

Quantifying species-environment relationships, i.e., analyzing how species are distributed on the Earth has been one of the fundamental questions studied by biogeographers and ecologists for a long time [1]. Several enabling technologies have made biodiversity data available at much finer scales in the past decade [2], including DNA barcoding for species identification, geo-referring for converting descriptive museum records to geographical coordinates, database technologies for managing species presence locations and related taxonomic and environmental data, and, Geographical Information System (GIS) for species distribution data modeling and analysis. The newly emerging cyberinfrastructure technologies (e.g., metadata,

ontology, Web services and scientific workflow) have made exchanging and sharing species distribution data over the Web much easier. The currently largest species occurrence data repository might be the Global Biodiversity Information Facility (GBIF) which was established by governments in 2001 to encourage free and open access to biodiversity data via the Internet[1]. Through a global network of countries and organizations, as of August 2012, the GBIF data portal has more than 375 million species occurrences records on 1,487,496 species. The majority of the records are geo-referenced which makes it possible to overlay species occurrence records with different types of raster and vector data layers for exploring biodiversity patterns and their relationships with environments and human impacts at global and regional scales.

Given the virtually countless combinations of species taxa, geographical regions and ecosystems [3], many types of exploratory analysis on integrated taxonomic-geographical-environmental data can be investigated [4]. In this study, we will be focusing on a fundamental spatial operation for zone-based point location data summation, i.e., counting the numbers of points that fall within a set of polygons in a zonal dataset. The operation is closely related to point-in-polygon test based spatial joins [5] [6] and is well-supported in several leading GIS software known as Zonal Statistics [7]. While both spatial databases and GIS have exploited optimization techniques, such as indexing and preprocessing, existing designs and implementations are mostly based on serial CPU computing models and usually incur significant delays when processing large scale datasets. Modern commodity personal computers are increasingly equipped with large memory and many-core accelerators, such as Nvidia GPUs that are capable of general computing based on the Compute Unified Device Architecture (CUDA) parallel programming model [8]. Unfortunately, many commercial and open source spatial databases and GIS are optimized for the previous generations of hardware based on outdated cost models and fail to make full use of the computing power provided by modern commodity hardware.

Built on top of our previous research and development efforts on spatial indexing and query processing on GPUs [6] [9], in this study, we aim at accelerating explorations of the GBIF global biodiversity data by designing and implementing efficient data parallel algorithms for high-performance zonal statistics on the hundreds of millions of species occurrences over tens of thousands of complex polygons on commodity GPUs with limited memory capacities. First of all, we have designed a framework to allow efficient use of mapped memory on CPUs as extended GPU memory automatically and support data parallel designs. Second,

we have developed a flexible point indexing technique to index large point datasets that are beyond GPU memory capacity by using mapped memory efficiently through batched processing. Third, we have extended our binary search based spatial filtering algorithm to work with the new point indexing technique. Fourth, a cell-in-polygon test based optimization technique for advanced spatial filtering is developed to allow assigning polygon identifiers to points if the grid cell that the points fall into is tested to be completely within a polygon without performing expensive point-in-polygon tests for individual points. We have performed extensive experiments to demonstrate the efficiency of GPU-based massively data parallel zonal statistics technique and compare it with two reference serial implementations using traditional open source geospatial software packages. The performance of the data parallel framework and techniques as well as the effectiveness of the cell-in-polygon based optimization technique are tested under different experiment settings.

The rest of the paper is arranged as follows. Section 2 introduces background and motivation and briefly reviews related work. Section 3 provides details of the data parallel zonal statistics framework, the point indexing and spatial filtering techniques, and the cell-in-polygon test based optimization technique for advanced spatial filtering. Section 4 presents the experiments and results. Finally Section 5 is the conclusion and future work directions.

2. BACKGROUND, MOTIVATION AND RELATED WORK

Given a point dataset T_O representing species occurrences with two attributes (sp_id, the_geom) and a polygon dataset T_Z representing zones also with two attributes (z_id, the_geom), the basic zonal statistics operation to count the number of occurrences of species in each polygon can be expressed as the following SQL statement:

SELECT COUNT(*) from T_O, T_Z
WHERE ST_WITHIN (T_O.the_geom,T_Z.the_geom)
GROUP BY T_Z.z_id;

Here the_geom attributes in the two datasets represent geometry, i.e., points and polygons, respectively. Advanced zonal statistics operations likely involve species identifiers in additional clauses (such as WHERE and GROUP) to derive the occurrence counts for a single or a group of species. The species occurrence counts can be used to compute abundance and richness measurements for a variety of types of biodiversity studies [10]. The GBIF data portal has provided overview maps of species occurrences for different species groups as well as countries which is very useful in understanding the overall species distribution patterns. However, the maps are mostly for visualization purposes and are limited to a few fixed resolutions up to 0.1 by 0.1 degree which might be too coarse for many scientific inquiries. It is clear that the zonal statistics based data summation operations are closely related to the point-in-polygon test used in the ST_WITHIN function. The function is defined by the Open Geospatial Consortium (OGC) Simple Feature Specification (SFS[2]) and has been implemented in several spatial database systems and GIS, e.g., Java Topology Suit (JTS[3]), Geometry Engine - Open Source (GEOS[4]) and PostGIS/PostgreSQL [11]. The Oracle Spatial development team has recently proposed to build an in-memory R-Tree to speed up

topological relationship query processing for complex regions [12], including point-in-polygon test.

Point-in-polygon test has been extensively investigated by the computational geometry and spatial database research communities. While computational geometry research usually focuses on a single point and polygon pair, spatial database research addresses the overall efficiency on testing a large set of points and polygons that can be abstracted as a special type of spatial joins. Spatial joins are typically divided into two phases, i.e., the filtering phase and the refinement phase [5]. The filtering phase utilizes some pre-built or on-the-fly constructed spatial indices to pair subsets of points and subsets of polygons for further refinements. The refinement phase in the point-in-polygon test based spatial joins applies computational geometry algorithms to determine whether a point is within a polygon for paired points and polygons. Obviously, building indices incurs additional overheads but can significantly reduce the number of required point-in-polygon tests and improve spatial join efficiency.

In our previous works, we have extensively investigated the potentials of GPU-based spatial indexing and spatial joins and many of them are based on data parallel designs. For a brief review of these works, we refer the reader to our ACM SIGSPATIAL SPECIAL paper [13]. In particular, a parallel binary search based spatial join framework [9] is proposed for joining indexed point data (using quadtrees or grid-files) and indexed polyline or polygon data (using grid-files). Several applications that demonstrate the effectiveness and efficiency of the indexing and spatial join techniques have been reported, including point-in-polygon test based spatial association between taxi pickup/drop-off locations and census tracks in the New York City (NYC) [6], point-to-polyline nearest neighbor search based spatial associations between taxi pickup/drop-off locations and street network in NYC [9] and Hausdorff distance based trajectory similarity queries in Beijing [14].

Zonal statistics on GBIF species occurrence data is conceptually similar to the point-in-polygon test based spatial join which may suggest that we can simply apply the techniques we have developed in [6] to this new dataset. However, first of all, the number of species occurrences in the dataset (375+ million) is more than two times larger than the number of taxi pickup locations we have processed previously (~170 million) and it is impossible to index all species occurrences on GPUs completely due to their memory capacity limit. Second, the polyline/polygon/trajectory data that we have used in our previous applications are considerably simpler than the World Wild Fund (WWF) ecoregion polygon data[5]. The average number of vertices per polygon in the WWF dataset (279) is nearly three times as large as that of the NYC census block dataset (108). Third, comparing with taxi pickup locations that are mostly clustered in major street intersections, the distributions of species occurrences are much more dispersed which is likely to cause significant execution flow divergences on GPUs, a typical problem in degrading GPU computing performance [8]. As such, effective optimization techniques are keys to achieving high performance for large datasets at the scale in order to support interactive visual explorations. Finally, perhaps more importantly, while the recent Nvidia GPUs set the GPU memory capacity (24 GB) to a new level, from a research perspective, it is crucial to develop a flexible framework to support zonal statistics

and other types of geospatial processing on large datasets that exceed GPU memory capacity limit.

3. Efficient Zonal Statistics on GPUs

We propose to follow the GPU-based spatial join framework we have developed previously [9] and reuse existing components, e.g., point-in-polygon test GPU routine presented in [6], whereas possible. Our new contributions in this study are four-fold: 1) a framework to allow efficient use of mapped CPU memory as extended GPU memory automatically and support data parallel designs, (2) a flexible and efficient point indexing technique to index large point datasets that are beyond GPU memory capacity, (3) an extended binary search based spatial filtering algorithm to work with the new point indexing technique, and (4) a cell-in-polygon test based optimization technique for advanced spatial filtering. *The four new designs are highlighted and numbered in Fig. 1.* We next introduce our data parallel framework as the motherboard for relevant techniques before the design and development details are presented in the following subsections.

3.1 Data Parallel Framework

The data parallel framework for high-performance zonal statistics is shown in Fig. 1. Note that we use solid arrows to show data processing steps and dashed arrows to show the correspondences among data used in different components in the

framework. Following our previous studies [13], the point coordinates and polygon vertices are stored as arrays with each element has a fixed length, instead of storing them as objects that may have variable lengths. Although not shown in Fig. 1 due to space limit, a polygon index array is constructed to store the first vertex positions of polygons to efficiently access polygon vertex arrays on both GPUs (for coalesced memory accesses) and CPUs (for cache-friendly memory accesses). Since the GPU-based zonal statistics technique is built on top of the point-in-polygon test based spatial joins, we reuse the relevant data parallel designs presented in [9] including sort-based point indexing, grid-file based polygon MBB (Minimum Bounding Box) rasterization and indexing, and, binary search based spatial filtering and nested-loop based spatial refinement. The GPU-based point-in-polygon test technique [6] is plugged into spatial refinement to implement the required zonal statistics functionality. These designs are extended when necessary and will be described in their respective subsections next. As both the previous implementations and the implementations for new extensions can be realized using either data parallel primitives supported by parallel libraries (e.g., Thrust[6] that comes with CUDA SDK) or nested loops with regular data access patterns and can be efficiently realized by using native GPU programming languages (e.g. CUDA), we consider both the new designs for individual components and the overall framework data parallel [15].


Fig. 1 Data Parallel Framework for Efficient Zonal Statistics on GPUs

Our new flexible data parallel framework utilizes the Unified Virtual Addressing (UVA) feature that is available in newer generations of Nvidia GPUs [8], which include both Fermi, Kepler and Maxwell based Nvidia GPUs, to allocate chunks of CPU memory and make them accessible to both CPUs and GPUs. We term such CPU memory chunks as GPU mapped memory on CPUs, or simply GPU mapped memory when there is no confusion. Using GPU mapped memory virtually extends GPU memory capacity by using CPU memory which can be two orders of magnitude larger (1-6 GB vs. 100-1000 GB). However, in a way similar to using disks as virtual CPU memory [16], using GPU mapped memory in a naive way may perform poorly. For example, our experiments show that simply applying the parallel sort primitive on GPUs (which is based on the radix sort algorithm) for point indexing using mapped memory can result in a much inferior performance. Our data parallel framework allows GPU mapped memory to be effectively utilized for scalability without significant degrading the overall performance when applied to larger scale data.

3.2 Flexible Point Indexing on GPUs using Batched Processing

As reported in [9], the Flatly Structured Grid-file (FSG) approach is much simpler than the Multi-Level Quadrant (MLQ) based approach for point indexing from both design and implementation perspectives. The load balancing guarantee of the MLQ approach is not instrumental for large scale data to achieve good performance when the number of (point quadrants, polygon) pairs after spatial filtering is much larger than the number of parallel processing units. The multi-core CPU implementation (using 8 Intel Xeon E5405 CPU cores) of the FSG approach actually has achieved much better performance than the MLQ approach on GPUs (using Nvidia Quadro 6000) despite that the GPU can achieve a much higher sorting rate which is a key to the performance of both MLQ and FSG implementations. The results suggest that the FSG approach is superior to the MLQ approach for spatially joining large scale datasets. Therefore, the FSG approach is adopted in this study for point indexing. While it is interesting to implement the FSG approach on GPUs, it has a much larger memory footprint which limits the number of species occurrence point records to about 100 million when each record has a length of 12 bytes, i.e., 4-byte float for x/y coordinate and 4 byte integer for taxon identifier. This is also the reason that we were forced to index 170 million taxi trip records on multi-core CPUs as reported in [9]. The scalability issue of the existing point indexing technique has motivated us to develop a more flexible parallel design for the FSG approach on GPUs.

Given a point dataset with N records where each record includes a longitude and latitude pair (optionally with some other attributes such as taxon identifier in the GBIF dataset), the dataset is stored as an array of records in a CPU memory block which is mapped by a GPU device through the UVA mechanism [8]. Both the CPU and the GPU in a computing node can access the memory block, not only for point indexing but also for point-in-polygon test in spatial refinement. When GPUs access the mapped memory in CPUs, as illustrated at the top of Fig. 1, they are required to transfer data in small units from the mapped memory in CPUs to their processors that need the data through a PCI-E bus dynamically. This is quite different from the conventional way that transfers data from CPUs to GPUs in large chunks before they are processed by GPUs. Clearly the flexibility of being able to utilize larger CPU memory is at the cost of lower

efficiency in data transfer, in a way very similar to virtual memory in traditional CPU computing and buffer management in relational database systems.

One might attempt to apply the FSG design to GPU mapped memory to minimize the effort of reimplementation which can be costly. However, this will not work for two reasons. First, while the inputs and outputs of the FSG design can use GPU mapped memory, the implementations of many parallel primitives used in the design (including *sort* in Thrust which is used by FSG) may use temporal GPU memory storage for intermediate results which is typically proportional to input sizes. The required temporal memory footprints are likely to exceed GPU memory capacity for large scale data and the process will fail due to out of memory. For example, the Nvidia Quadro 6000 GPU can only sort about 200 million records (including longitude/latitude and taxon identifier) which is well below our goal for a flexible solution. Second, even if little intermediate results are produced and the GPU is free from the memory capacity problem after putting both inputs and outputs in GPU mapped memory in CPUs, excessive accesses to the mapped host memory in an uncoordinated manner may significantly degrade performance and make GPU implementations unattractive. For example, sorting a subset of 125 million GBIF point data records in a Quadro 6000 GPU using mapped memory needs 23.867 seconds while only 0.683 second is required if the sorting is done completely in GPU memory. This represents a 34.7X slowdown which is not surprising, given that the underlying radix sort algorithm requires significant amount of data movements and PCI-E bus bandwidths are about 1-2 orders of magnitude slower than GPU memory bandwidths.

Our solution is to partition the input point data array into chunks and process the chunks in batches. While we refer to [9] for the detailed design of the original (i.e., single-chunk) FSG algorithm and a multi-core CPU implementation for references, we next briefly repeat the key ideas of the single-chunk FSG design before presenting details of the multi-chunk FSG algorithm and its GPU implementation for the purpose of being self-contained. As discussed earlier, the FSG algorithm for point indexing actually is much simpler than the MLQ algorithm presented in [6] and requires a simple chaining of only four parallel primitives, i.e., *transform*, *sort*, *reduce* (by key) and *scan*. The transform primitive derives a cell identifier for each point based on its (longitude, latitude) pair. Given a grid cell size, row-major ordering is used to compute the cell identifier for easy calculation. The next step is to sort the points based on their cell identifiers to put all points that fall within a grid cell close to each other. Clearly, points within a grid cell are not sorted for performance concerns. A *reduce* (by key) primitive is used to count the numbers of points within all grid cells which are subsequently used to compute the positions of the first points among the points that are within the corresponding grid cells. As shown in the middle part of Fig. 2, given an input array *PntRec*, four arrays will be in the output list. In addition to the sorted *PntRec* array, we also have *PntCID* that stores grid cell identifiers, *PntLen* array that stores the numbers of points in cells and *PntPos* array that stores the positions of first points among the points in a grid cell in the sorted *PntRec* array.

When there are multiple chunks in a point data array, thanks to our data parallel design, each chunk can be processed independently, either using a single GPU where the chunks are processed sequentially, or using multiple GPUs where the chunks

are processed in parallel, or in a way that combines the two options. For GPUs with smaller memory capacities, we can simply decrease batch sizes and make the technique flexible. The performance will degrade gracefully for smaller GPU memory capacities but the tradeoff can be justified in this case. The design is similar to the mapping phase in the MapReduce computing model [17] in the sense that chunks are processed independently and no communications are required among chunks in this step.

While it seems that we will need to rearrange the sorted point array in multiple chunks to proceed to spatial filtering, our design avoids such data movements (which could be expensive for hundreds of millions of records) by only manipulating the three arrays at the grid cell level, i.e., *PntCID, PntLen and PntPos* arrays. Since the number of grid cells for indexing is typically much smaller than the number of point records, the costs for manipulating such arrays are much lower. This is the key to the scalability and efficiency of our new design for indexing point data. The steps are illustrated in the lower part of Fig. 2. First of all, the total number of points in each chunk is collected for all chunks and stored in the *CLen* array. Similar to computing the *PntPos* array from the *PntLen* array by using a *scan* (prefix-sum) primitive, we can compute the *CPos* array from the *CLen* array. Note that the lengths of the *CLen* and *CPos* arrays are the same as the number of chunks which are typically very small and the costs of this step are negligible. Next, the value of each *CPos* array element is added back to all the elements in the *PntPos* array

within each chunk (bottom part of Fig. 2), so that the elements in the *PntPos* array correctly index points in grid cells after concatenating the *PntRec*, *PntLen* and *PntPos* arrays in all chunks. Again, since all the steps are implemented using parallel primitives, the design is highly data parallel and can be implemented on top of parallel libraries that support these fundamental primitives in a straightforward manner. Experiments on the GBIF point data shows that, about 1/3 of the total processing time is spent on transferring data between GPUs and CPUs while the rest 2/3 of the time is spent on sorting for all batches. The runtimes of the rest of the steps (including *transform*, *reduce* and *scans*) are relatively insignificant. Given that GPUs have excellent performance on sorting [18], the new design, termed as Multi-Chunked FSG for point indexing, is expected to be not only flexible but also highly efficient.

3.3 Extending Spatial Filtering to Support Chunked Point Indexing

The binary search based spatial filtering design and its GPU-based implementation [6] [9] does not allow duplicated cell identifiers which means that the technique will not work for the multi-chunked point indices using the technique presented in Section 3.2. For a grid cell appears in *K* chunks, there will be *K* duplicated cell identifiers in the *PntCID* array. We next present details on how binary search based spatial filtering can be extended to work with the Multi-Chunk FSG approach for point indexing.
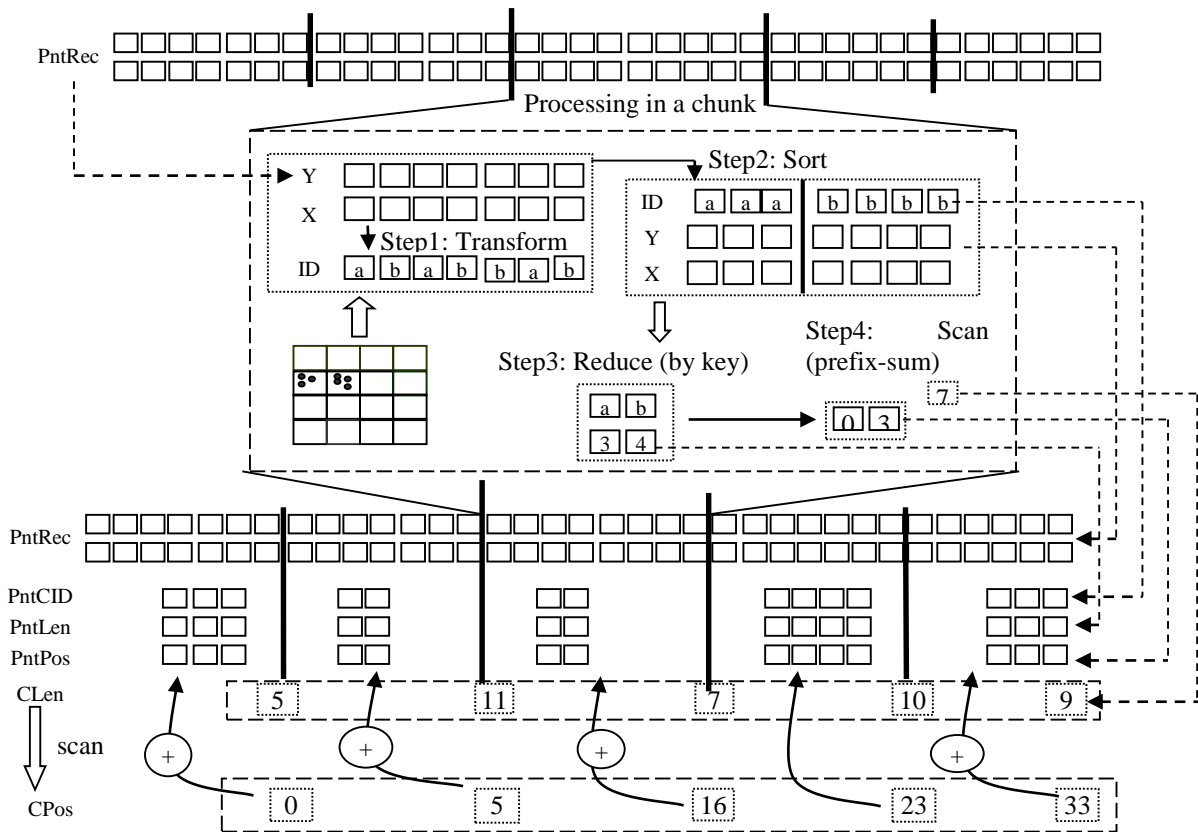

Fig. 2 Parallel Design for Indexing Point Data in Chunks

First, the *PntLen* and *PntPos* arrays derived from the Multi-Chunked FSG point indexing approach are sorted by using the *PntCID* array as keys to make the same cell identifiers appear next to each other in the *PntCID* array. Note that the positions of the elements in the *PntLen* and *PntPos* arrays are changed according to the key-value based sorting. Next, as shown in Fig. 3, for each of the elements in the MID array, our spatial filtering algorithm binary searches the *PntCID* array by using the corresponding element in the *MC* array as the key. Recall that the MID array and the MC array store the correspondences between polygon MBB identifiers and cell identifiers of rasterized polygon MBBs [9]. The key extension is to match cell identifiers in the MC array and the sorted PntCID array by using three parallel primitives, i.e., *binary_search*, *lower_bound* and *upper_bound*, as a bundle for binary searches. While the *lower_bound* and *upper_bound* primitives return the first and the last positions where values could be inserted without violating the ordering during binary searching, the *binary_search* primitive returns the result of whether the values being searched are or are not in the array being searched. The resulting position vectors from the *lower_bound* and *upper_bound* primitives need to be filtered out by the resulting boolean vector from the *binary_search* primitive to eliminate unsuccessful searches while keeping the upper bounds and lower bounds of successful searches. Note that it is not necessary to use the *upper_bound* primitive if the cell identifiers in the *PntCID* array are guaranteed to be unique, which is the case if the point dataset is not chunked. This is exactly the original FSG design for spatial filtering presented in [9]. Finally, for each matched (MID$_i$, lower_bound$_i$, upper_bound$_i$) triple, we can use MID$_i$ and lower_bound$_i$ and upper_bound$_i$ values to access the polygon vertex arrays and point coordinate arrays as follows. Assuming the arrays that store the vertex positions and the numbers of polygon vertices are *PlyPos* and *PlyLen*, respectively, then the polygon vertices will be at the positions *PlyPos*[**idx**(MID$_i$)] .. *PlyPos*[**idx**(MID$_i$)+1]-1 with *PlyLen*[i] vertices. Function **idx**(i) maps polygon identifier i to an index in the *PlyPos* or *PlyLen* array, which can be as simple as **idx**(i)=i. Similarly points that fall within the grid cell whose identifier is being searched are distributed in upper_bound$_i$ - lower_bound$_i$ blocks. Note that here blocks are combinations of chunks and grid cells, i.e., a block of points are within a grid cell in a chunk. For each j= lower_bound$_i$ .. upper_bound$_i$, the starting position and

number of points in these blocks are recorded in *PntPos*[j] and *PntLen*[j], respectively. They can be used to access the *PntRec* array to retrieve point coordinates or other information for further processing. While supporting multiple data point chunks has added significant complexity to our original spatial filtering design, it eliminates the need to actually sort point records across multiple chunks as it would have been done for a single chunk. We note that data movements are typically expensive in various sorting implementations on both CPUs and GPUs and should be avoided as much as possible for large scale data.

To better illustrate our extended design, an example is provided in Fig. 3. In the top part of the figure, after binary searching each cell identifier in the *MC* array from the *PntCID* array, while there are two matched cell identifiers in the *PntCID* array (at positions 1 and 2 and shaded with light and dark gray colors, respectively) are paired with cell identifier 2 in the *MC* array, there is only one match for cell identifiers 6 and 8, respectively, and there is no match for cell identifiers 5, 4 and 1. As shown in the bottom part of Fig. 3, the three points in the first chunk and the four points in the second chunk in grid cell #2 can be accessed by combining the corresponding elements in the *PntPos* and the *PntLen* arrays. The point data records are colored in light and dark gray in the same way as the two matched elements in the *PntCID, Pntlen and PntPos* arrays are colored.

3.4 Parallel Cell-in-Polygon Test for Optimization

The tradeoffs between spatial filtering and spatial refinement in spatial joins are well studied in spatial databases [5]. In our FSG approach, clearly, using a high resolution grid for point/polygon indexing will increase the amount of workload in indexing and spatial filtering but is likely to reduce the workload in the final spatial refinement phase. However, for heavily clustered regions, the numbers of points that fall within some grid cells are likely to be large. Assuming that there are K points in a grid cell, directly applying the point-in-polygon test would require O(K) tests, each requires O(V) operations where V is the number of vertices in the polygon to be tested. When K is large in such grid cells, directly performing point-in-polygon test can be very expensive.


Fig. 3 Data Parallel Design for Spatial Filtering with Chunked Point Indexing

By observing that if a grid cell is completely inside or outside a polygon, we can directly assign the test results to all points in the grid cell without requiring any point-in-polygon test. Although a cell-in-polygon test is generally more expensive than a point-in-polygon test, when K is large, the optimization is likely to be beneficial. From a probabilistic perspective, if the probability that the grid cell is completely within or outside of a polygon is high, the overall computing cost can be significantly decreased by performing a single cell-in-polygon test instead of multiple point-in-polyline tests. We consider this optimization technique as part of spatial filtering and refer it as advanced spatial filtering in this study.

Several well-established computational geometry principles can be used to test the relationships between a rectangle (including a squared grid cell) and a polygon. Motivated by the procedure used in [19], we have used the following two steps to determine whether a grid cell intersects, is within, or, is outside of a polygon. Note that multi-rings are allowed in our technique by separating rings with the origin of the underlying coordinate system. Our technique extends the work in [19] that only supports single-ring polygons and the extension is necessary for WWF ecoregion data as polygons in this dataset are complex and many of them have multiple rings. As shown in Fig 4A, the first step for cell-in-polygon test is to check whether any of the grid cell's four edges intersects with any of the polygon edges, or, whether any of the polygon's vertices is within the cell, to determine whether the grid cell intersects with the polygon. If the grid cell does not intersect with the polygon, then it is either completely inside (Fig. 4B) or completely outside the polygon (Fig. 4C). We subsequently test whether any of the cell's corners is within the polygon. If the test is true then the grid cell is inside the polygon; otherwise the grid cell is outside of the polygon.


Fig. 4 Three Cases in Cell-in-Polygon Tests

4. EXPERIMENTS AND RESULTS

4.1 Data and Experiment Setup

The GBIF global species occurrence dataset has 375+ million species occurrences records as of 08/02/2012 when we obtained the dataset. Our preprocessing results have shown that the dataset contains 1,487,496 species, 168,280 genus, 1,142 families in 262 classes, 109 phyla and 9 kingdoms. The majority (95.7%) of the records is related to animals and plants. A large portion (74.1%) is geo-referenced (with latitude/longitude coordinates at different accuracy levels) and can be associated with terrestrial eco-regions. The WWF ecoregion dataset comes in ESRI shapefile format[7] and has 14,458 polygons, 16,838 rings and 4,028,622 points. The ecoregion data volume is relatively small when compared to today's CPU memory capacities. However, the raw GBIF species occurrence data we received is in the form of a relational database dump with 35 columns and has a total data volume of 180 GB. Many of these columns use the variable character type which makes random accesses very difficult. We have extracted individual columns and converted them into binary format for further processing. In this study, we

primarily focus on three attributes, i.e., latitude, longitude and taxon identifier. The total data volume of the three columns is about 4.2 GB for the 375 million point data records. As the total data volume of the three attributes is less than 1/3 of the CPU memory in our experiment system (16 GB), hereafter we assume that all data involved are memory-resident.

We have empirically set the data grid resolution to 1 arc-minute (approximately 2 kilometers around the equator) primarily because this might be the finest resolution for global biodiversity studies and it may already be beyond the accuracy of some species occurrence records. The width and height of the resulting grid are 21,600 and 10,800, respectively. The gridded coordinates of a point location can be easily stored as a 2-byte short integer along both longitude and latitude dimensions. As the indexing grid resolutions are allowed to be coarser than the data grid resolution, we have chosen three grid resolutions for spatial indexing, i.e., 2^n*2^n for n=13, 14 and 15, to investigate how various performance measurements change with indexing grid resolutions.

All experiments are performed on a Dell Precision T5400 workstation equipped with 16 GB memory and a 500 GB 7200 RPM hard drive. The workstation has dual quad-core Intel E5405 CPUs (8 cores in total) running at 2.00 GHZ and with 6MB L2 cache per core pair, 128 KB L1 cache per core and 12.8 GB/s memory bandwidth per CPU. The workstation is also equipped with an Nvidia Quadro 6000 GPU device with 448 CUDA cores (1.15 GHz), 6 GB GDDR5 memory and 144 GB/s memory bandwidth. The sustainable disk I/O speed is about 100 MB/s while the theoretical data transfer speed between the CPU and the GPU is 8 GB/s through PCI-E. The relevant software installed on the workstation are Nvidia CDUA SDK 5.0 (with Thrust library 1.6), g++ 4.6.3 and Intel TBB 4.1. All programs, including the two serial implementations using traditional technologies (Section 4.3), are optimized with -O3 option during compilations for fair comparisons.

4.2 Overall Results

The runtimes of the four components in our GPU-based zonal statistics technique, i.e., point indexing, polygon MBB indexing, spatial filtering and spatial refinement, under the three grid resolutions are measured. We do not plot polygon MBB indexing runtimes in Fig. 5 because they are negligible (51, 197 and 787 milliseconds for the three grid cell levels, respectively) when compared to others which are plotted. Note that the spatial filtering runtimes are measured with the optimization technique described in Section 3.4.

From Fig. 5 we can see that the runtimes of spatial filtering and spatial refinement dominate the overall runtimes under all the three grid resolutions. From an application perspective, the most significant conclusion we can draw from the experiment results is that, zonal statistics on the 375+ million species occurrences over the 15 thousand complex ecoregion polygons based on point-in-polygon test spatial relationship can be completed on a commodity workstation equipped with a single GPU device in the order of 100 seconds.

Our data parallel designs make it relatively easy to implement the designs in multiple parallel hardware platforms. For demonstration and comparison purposes, we have also implemented the designs on multi-core CPUs. To minimize the additional implementation efforts, since the Thrust parallel library also provides interfaces to the Intel Threading Building Block

(TBB[8]) library that is known to be efficient on multi-core CPUs, we recompile our GPU-based Thrust code to use TBB and link it with the TBB runtime library to utilize multi-core CPUs in a way similar to the work reported in [9] for point-to-polyline nearest neighbor search based spatial joins, but with two exceptions. The first exception is on point indexing where we have found that the GNU parallel mode library[9] is more efficient for multi-core CPU based sorting and we use it instead for fair comparisons. The second exception is related to the native CUDA implementation of the point-in-polygon test module as reported in [6]. For fair comparisons, we have implemented the point-in-polygon test module using the native TBB programming model by assigning a range of (polygon, block) pairs as a task and letting a single CPU core loop through all the points in the polygon for point-in-polygon test.

As expected, the GPU-based implementations are significantly faster than their peer multi-core CPU implementations with speedups ranging from 2.7X to 4.7X for the three major components (point indexing, spatial filtering and spatial refinement) under the three grid resolutions, as shown in Fig. 6. The speedups are higher for spatial filtering and spatial refinement as they are more computing intensive and can better use GPU's massive floating point computing power. Please note that the CPU performance is measured when all the 8 cores are fully utilized and the multi-core CPU implementations have been optimized as much as possible for fair comparisons. Our results agree with the rigorous performance analysis on quite a few non-geospatial benchmarks reported in [20] when comparing the performance of GPUs and multi-core CPUs. The comparisons also suggest that our data parallel designs can achieve high efficiency on both GPUs and multi-core CPUs by using parallel primitives that are optimized for the respective hardware platforms. As such, they are less likely to depend on the programming skills of individual programmers and are more preferable from a software development perspective.

After comparing with the multi-core CPU implementations based on our data parallel designs, we would like to comment on the relationships between filtering and refinement using different grid resolutions in our GPU-based implementation as observed in the experiments. First of all, from Fig. 5, it is easy to see that the filtering runtimes increase with grid resolutions while the refinement runtimes decrease with grid resolutions for both CPU and GPU implementations. This is expected as using finer resolution grid for filtering reduces false positives and requires fewer point-in-polygon tests in the refinement phase. Since cell-in-polygon test is used in the filtering phase as an optimization technique, which is also computation intensive, the runtimes in the filtering phase are comparable with the runtimes in the refinement phase, although the computing workload for the basic spatial filtering design can be quite light [9]. While the runtime of spatial filtering is about 1/5 of the runtime of spatial refinement at the grid level 13, the ratio quickly increases to 1.6 at the grid level 15. The totals of the filtering and refinement runtimes (and hence the end-to-end runtimes) are minimized at the grid level 14. The results indicate that choosing a proper grid level is important in improving the system performance and we leave a more comprehensive investigation for future work.
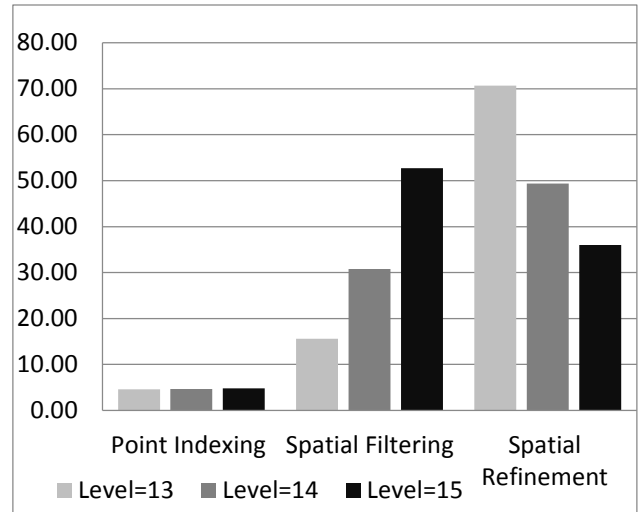

Fig. 5 Plots of runtimes (s) of Point Indexing, Spatial Filtering and Spatial Refinement on GPUs using three grid resolutions
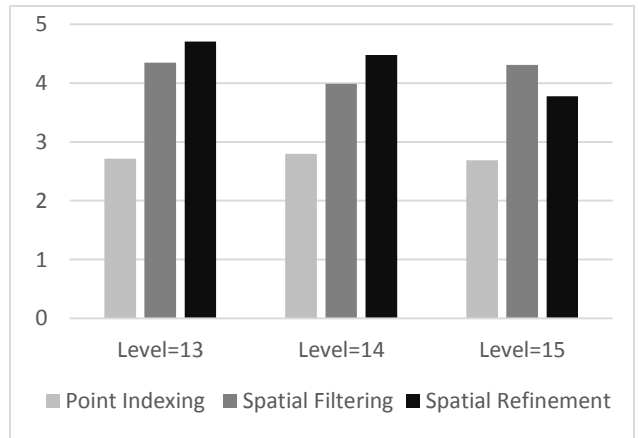

Fig. 6 Plots of GPU over multi-core CPU speedups

4.3 Comparisons with alternatives using traditional technologies

It is not our intention to directly compare our memory-resident massively data parallel technique with serial implementations using traditional geospatial software packages that are designed for uniprocessors and disk-resident systems. This is because the two techniques are developed for different applications targeting at different hardware. Nevertheless, we report the performance comparisons with two serial implementations using libspatialindex[10] for R-Tree based polygon indexing and GDAL[11] (through GEOS) for point-in-polygon tests for reference purposes. The comparisons can also help understand the level of performance that our technique has achieved due to data parallel designs and optimized implementations on GPUs.

The major difference between the two serial implementations is the following: the second implementation incorporates an optimization heuristic in hope to improve the overall performance, while the first serial implementation simply queries polygon MBBs that intersect with each and every point before performing point-in-polygon test between the point and the polygons whose MBBs intersect with the querying point. Given that querying the polygon R-Tree for 375+ million points can be expensive when traversing the polygon R-Tree individually, the heuristic is to locate all the MBBs in the polygon R-Tree leaf nodes that intersect with grid cells of groups of points where only a single R-Tree query is needed for the groups of points within the grid cells. The second implementation clearly requires grid-based indexing of points but can potentially save R-Tree query time as the number of accesses to R-Tree nodes can be significantly reduced through point grouping. Although it is possible to use R-Tree to index points by treating each point as a degenerated MBB, the high index construction cost has led us to decide to either not index the point data (implementation 1) or re-use the results of our grid file based indexing (implementation 2).

We believe the first serial implementation represents a reasonably efficient implementation by applying spatial filtering before refinement as a typically trained geospatial programmer would do. We also expect the second serial implementation to be more efficient by incorporating the optimization heuristic. However, the results are quite the opposite as detailed below. The code for the two serial implementations and the three subsets of point data are publically available online[12] and we encourage interested readers to cross-examine the implementations, validate the experiment results and make independent comparisons.

First of all, neither implementation is as efficient as we have expected. It takes 18.77 hours to process a subset of approximately 10 million point records with a throughput in the order of 139 points per second. Additional experiments using two smaller subsets of species with 279,808 and 746,302 points result in similar performance, i.e., 138 points per second for both smaller datasets. By using a linear extrapolation, it would take 600+ hours to complete the 375 million points using the serial implementations, although the implementation does exhibit excellent scalability and is suitable for MapReduce/Hadoop systems. However, the performance is 4-5 orders of magnitude slower (138 points per second) than our GPU based implementation (375 million points in about 100 seconds) which is inferior from both the usability and monetary cost perspectives.

Second, the experiments show that the optimization heuristic employed in the second serial implementation is largely ineffective. While measured accesses to the polygon R-Tree has been dramatically reduced by the optimization, the runtimes do not get improved noticeably. Further investigations have revealed that the polygon R-Tree is fairly small (a few megabytes) and can be completely cached in memory which makes reducing accesses to R-Tree insignificant as in disk-resident cases. Since querying cell boundaries against the polygon R-Tree will inevitably cause more false positives when compared with directly querying points and point-in-polygon tests, which is much more expensive than accessing memory-resident R-Tree nodes, the heuristic does not work as expected. Since point data are also made memory resident in both serial implementations, we can conclude that the low performance of the serial implementations is largely unrelated to disk I/Os in our experiments.

While we are still in the process of fully understanding the 4-5 orders of magnitude of performance differences, we believe that excessive memory allocation/deallocation to accommodate for low memory capacities, library overheads for generality (e.g., object-oriented abstractions) and mismatches between traditional data structures and algorithms with modern hardware architectures (e.g., cache unfriendliness in depth-first tree traversals) are among the factors that contribute to the low performance of the two serial implementations by using traditional geospatial techniques. Furthermore, it is interesting to observe that, even assuming that our multi-core CPU-based implementations have achieved perfect scalability (8X for 8 cores), the performance of the corresponding serial implementations of our data parallel designs (by multiplying the number of cores with the measured runtimes) is still about three orders of magnitude faster than using traditional technologies. This may suggest that there is a huge room to improve traditional spatial data processing technologies by adopting data parallel designs and hardware architecture aware implementations. We leave this interesting interdisciplinary research topic for our future work.

5. CONCLUSION AND FUTURE WORK

In this study, we have significantly extended our previous techniques for point-in-polygon test based spatial joins on GPUs for large scale data. The integrated flexible designs and their GPU implementations have successfully performed zonal statistics on 375+ million global species occurrence records over 15 thousand complex ecoregions in facilitating exploring global biodiversity explorations. We have developed a flexible data parallel framework by using GPU mapped memory in CPUs for large-scale data that may exceed GPU memory capacity. We have extended our point data indexing and binary search based spatial filtering designs to accommodate multi-chunked point data indexing while achieving much higher efficiency when compared with using mapped memory naively. Including the cell-in-polygon based optimization, the combined improvements have reduced the total runtime to about 100 seconds using a single GPU device. The performance is several orders of magnitude faster than two reference serial implementations using traditional open source geospatial techniques. The realized high performance on top of flexible designs is not only significant for practical applications in exploring increasingly larger global biodiversity data but also suggests that there are huge rooms to improve the performance of traditional geospatial technologies on modern parallel hardware.

For future work, first of all, we would like to integrate our technique with data management and visualization frontends for practical applications. Second, since the flexible data parallel framework is also applicable to other types of spatial processing, it is thus interesting to examine its scalability in additional applications with larger scale data. For example, spatially and temporally associating 2.7 billion GPS points deposited into Openstreetmap Planet[13] with global road networks by using the point-to-polyline nearest neighbor search based spatial joins [9]. Third, our new data parallel framework allows integrating multi-core CPUs and multi-GPUs as well as other types of hardware accelerators that share the same address space to synergistically process large scale data by assigning chunks of array elements to multiple processors in a straightforward manner. We plan to materialize the design which essentially allows heterogeneous computing and implement it on a hybrid CPU-GPU system for

performance evaluation using the GBIF data and the Openstreetmap Planet GPS location data. Finally, while it is certainly a challenging task that requires significant effort, we plan to investigate the mismatches between the designs and implementations of traditional geospatial processing software packages and the new generation of parallel hardware in a systematic manner. The findings may not only lead to improved performance but also may provide new insights on how to make better use of commodity parallel hardware and enable larger scale geospatial processing with higher efficiency and better scalability.

6. REFERENCES

[1] C. Cox and P. Moore, Biogeography: An Ecological and Evolutionary Approach (7th Ed.), Wiley, 2005.

[2] F. A. Bisby, "The quiet revolution: Biodiversity informatics and the internet," Science, vol. 289 (5488), pp. 2309-2312, 2000.

[3] J. Zhang, "A high-performance web-based information system for publishing large-scale species range maps in support of biodiversity studies," Ecological Informatics, vol. 8, pp. 68-77, 2012.

[4] J. Zhang and L. Gruenwald:, "Embedding and extending GIS for exploratory analysis of large-scale species distribution data," in ACM-GIS Conference, 2008.

[5] E. H. Jacox and H. Samet, "Spatial Join Techniques," ACM Trans. Database Syst., vol. 32, no. 1, p. Article #7, 2007.

[6] J. Zhang and S. You, "Speeding up large-scale point-in-polygon test based spatial join on GPUs," in Proc. ACM BigSpatial'12, 23-32, 2012.

[7] D. Theobald, GIS Concepts and ArcGIS Methods, 2nd Ed., Conservation Planning Technologies, Inc., 2005.

[8] D. B. Kirk and W.-m. W. Hwu, Programming Massively Parallel Processors: A Hands-on Approach, 2nd ed., Morgan Kaufmann, 2012.

[9] J. Zhang, S. You and L. Gruenwald, "Parallel Online Spatial and Temporal Aggregations on Multi-core CPUs and Many-Core GPUs," Information Systems, vol. 44, p. 134–154, 2014.

[10] C. Ricotta, "Through the jungle of biological diversity," Acta Biotheoretica, vol. 53, pp. 29-38, 2005.

[11] R. Obe and L. Hsu, PostGIS in Action, Manning Publications, 2011.

[12] Y. Hu, S. Ravada, R. Anderson and B. Bamba, "Topological relationship query processing for complex regions in Oracle spatial," in Proc. ACM-GIS'12, 2012.

[13] J. Zhang, S. You and L. Gruenwald, "Large-Scale Spatial Data Processing on GPUs and GPU-Accelerated Clusters," ACM SIGSPATIAL Special, vol. 6, no. 3, pp. 27-34, 2014.

[14] J. Zhang, S. You and L. Gruenwald, "U2STRA: High-performance Data Management of Ubiquitous Urban Sensing Trajectories on GPGPUs," in Proc. ACM CDMW '12, 2012.

[15] M. McCool, J. Reinders and A. Robison, Structured Parallel Programming: Patterns for Efficient Computation, Morgan Kaufmann, 2012.

[16] L. Hennessy and D. A. Patterson, Computer Architecture: A Quantitative Approach, 5th edition, Morgan Kaufmann, 2011.

[17] J. Dean and S. Ghemawat, "MapReduce: a flexible data processing tool," Communications of the ACM, vol. 53, no. 1, pp. 72-77, 2010.

[18] D. Merrill and A. S. Grimshaw, "High performance and scalable radix sorting: a case study of implementing dynamic parallelism for GPU computing," Parallel Processing Letters, vol. 21, no. 2, pp. 245-272, 2011.

[19] K. Wang, Y. Huai, R. Lee, F. Wang, X. Zhang and J. H. Saltz, "Accelerating Pathology Image Data Cross-comparison on CPU-GPU Hybrid Systems," Proc. VLDB Endow., vol. 5, no. 11, pp. 1543--1554, 2012.

[20] V. W. Lee, C. Kim, et al, "Debunking the 100X GPU vs. CPU myth: an evaluation of throughput computing on CPU and GPU," in Proc. ACM/IEEE ISCA'10, 2010.

---

1 http://www.gbif.org/

2 http://www.opengeospatial.org/standards/sfs

3 http://www.vividsolutions.com/jts/JTSHome.htm

4 http://trac.osgeo.org/geos/

5 http://worldwildlife.org/biomes

6 https://thrust.github.io/

7 https://www.esri.com/library/whitepapers/pdfs/shapefile.pdf

8 https://www.threadingbuildingblocks.org/

9 http://gcc.gnu.org/onlinedocs/libstdc++/manual/parallel_mode.html

10 http://libspatialindex.github.io/

11 http://www.gdal.org/

12 http://www-cs.ccny.cuny.edu/~jzhang/zs_gbif.html.

13 http://wiki.openstreetmap.org/wiki/Planet.osm