

The Case for Distance-Bounded Spatial Approximations

Eleni Tzirita Zacharatou
TU Berlin
eleni.tziritazacharatou@tu-berlin.de

Andreas Kipf
MIT CSAIL
kipf@mit.edu

Ibrahim Sabek
MIT CSAIL
sabek@mit.edu

Varun Pandey
TU Munich
pandey@in.tum.de

Harish Doraiswamy
NYU
harishd@nyu.edu

Volker Markl
TU Berlin and DFKI GmbH
volker.markl@tu-berlin.de

ABSTRACT

Spatial approximations have been traditionally used in spatial databases to accelerate the processing of complex geometric operations. However, approximations are typically only used in a first filtering step to determine a set of candidate spatial objects that may fulfill the query condition. To provide accurate results, the exact geometries of the candidate objects are tested against the query condition, which is typically an expensive operation. Nevertheless, many emerging applications (e.g., visualization tools) require interactive responses, while only needing approximate results. Besides, real-world geospatial data is inherently imprecise, which makes exact data processing unnecessary. Given the uncertainty associated with spatial data and the relaxed precision requirements of many applications, this vision paper advocates for approximate spatial data processing techniques that omit exact geometric tests and provide final answers solely on the basis of fine-grained approximations. Thanks to recent hardware advances, this vision can be realized today. Furthermore, our approximate techniques employ a *distance-based error bound*, i.e., a bound on the maximum spatial distance between false or missing and exact results which is crucial for meaningful analyses. This bound allows to control the precision of the approximation and trade accuracy for performance.

1 INTRODUCTION

There is an explosion in the amount of spatial data being generated and collected today. Billions of GPS-enabled mobile devices, cars, social networks, satellites, sensors, and many other sources produce spatial data constantly. As a result of the ever-increasing data sizes and the computationally-intensive nature of spatial queries, it is hard to provide fast response times, which opposes the interactivity requirements of exploratory applications.

On the bright side, users often do not need exact results. They are instead satisfied with approximate answers, especially if these answers are accompanied by precision guarantees. However, approximate spatial data processing has attracted limited attention [2, 25, 31–33, 35]. There are two different notions of approximation in spatial databases. Synopsis-based techniques aim to accelerate spatial queries by evaluating them on small samples or models of the data [25, 31–33]. Existing techniques in this category are limited to certain types of queries (i.e., range queries, selectivity estimation,

k-means clustering and spatial partitioning). On the other hand, most spatial querying techniques approximate individual spatial objects with simpler geometries such as rectangles or convex polygons to accelerate queries [5, 35]. Unlike synopsis-based techniques, geometric approximations support arbitrary spatial predicates. Our work is related to the latter category, i.e., spatial query processing based on approximations of individual objects, which is orthogonal to sampling techniques that reduce the number of objects to be processed.

Notably, prior work does not give guarantees on the *spatial distance* between false (or missing) and exact results. Consequently, it is hard to interpret the provided approximate results, as the user has no information about how closely these results correspond to the particular region she is interested in. Guaranteeing *distance-based error bounds* is thus crucial. These bounds should be controlled by the user, essentially allowing to trade off between query results accuracy and query execution time.

Motivating Application: Visual Exploration of Mobility Data.

In an effort to enable urban planners to make data-driven decisions, in early 2017 Uber introduced Uber Movement, a visualization platform for the exploration of Uber rides¹. The platform allows users to visualize data of interest at different resolutions over varying time periods. Such visual analyses require interactivity, since high latency reduces the rate at which users make observations, draw generalizations, and generate hypotheses [17]. Furthermore, exact answers are not required, because visualizations are approximate in nature. Moreover, users typically perform “level-of-detail” exploration. They first look at a high level overview, and then zoom into regions of interest for further details [24]. Finally, there is usually uncertainty with respect to spatial coordinates, as GPS positions are typically accurate to within a 4.9 m radius [30]. Similarly, geographical region boundaries are often fuzzy, in the sense that adjacent regions are separated by extended zones (e.g., a street surface) rather than one-dimensional lines. As a result, these zones can be considered to be part of any of the adjacent regions. Overall, the interactivity expected from exploratory applications (visual or not), coupled with the inherent properties of spatial data, necessitate a paradigm shift towards spatial data processing techniques that have approximation at their core.

Hardware Trends. Spatial approximations have been widely used in spatial databases. Recent hardware trends, however, indicate that the time has come to rethink their design and utility. Existing techniques typically use a two-step “filter and refine” strategy [10] where approximations are only employed in a first filtering step

This article is published under a Creative Commons Attribution License (<http://creativecommons.org/licenses/by/3.0/>), which permits distribution and reproduction in any medium as well as allowing derivative works, provided that you attribute the original work to the author(s) and CIDR 2021. 11th Annual Conference on Innovative Data Systems Research (CIDR '21). January 11-15, 2021, Chaminade, USA.

¹<https://www.uber.com/newsroom/introducing-uber-movement-2/>

that yields a candidate result set. The subsequent refinement step eliminates false matches by performing exact geometric tests. Most efforts in improving spatial query processing focus only on the first filtering step [12, 19, 20, 22, 29]. In the past decades, the filtering step was in the critical path of the execution, since it was fetching the spatial approximations from disk. As a consequence, the number of slow disk accesses had to be minimized, which led to the design of approximations that sacrifice precision for compactness. Today’s machines, however, have large DRAM sizes that can go up to multiple terabytes and are often equipped with large-capacity Non-Volatile Memory (NVM), making it unnecessary to store the approximations in slow secondary storage devices. As a result of the faster access times provided by DRAM and NVM, the filtering step is no longer in the critical path, and the CPU-intensive refinement step becomes the bottleneck. As recent work shows [8], the main memory-based filtering step takes only a few milliseconds even for billions of points. Therefore, to improve performance, we now need to reduce (or even completely eliminate) the number of costly CPU-based refinements rather than the number of memory accesses. This necessitates the re-design of spatial approximations: we can now easily store more precise (and thus larger) approximations and leverage fast random access storage in exchange for better filtering efficacy and fewer CPU-intensive operations.

With increasing data sizes, the computation of precise approximations becomes more expensive. However, to support exploratory applications where the workload changes dynamically, we need to compute spatial approximations fast and on-the-fly. GPUs, and in particular their native support for *rasterization* make that possible today. The rasterization operation takes as input a geometric primitive (e.g., a polygon) and converts it into a collection of pixels which essentially form a fine-grained uniform grid approximation of the primitive. GPUs perform rasterization at interactive speeds, as they employ highly optimized hardware implementations. This enables us to design techniques that leverage GPUs to compute spatial approximations and evaluate spatial queries in real time.

This vision paper argues that fine-grained grid approximations can form the basis of spatial data processing. We show that these approximations allow to provide distance-based error bounds, enable us to exploit modern hardware, and facilitate further optimizations such as the use of learned indexes. The remainder of this paper outlines our vision of incorporating distance-bounded spatial approximations in different components of a spatial system, highlights individual challenges, and presents promising initial results.

2 APPROXIMATE PROCESSING

In this section, we first present geometric approximations commonly used in spatial data processing. We then describe how we can quantify the error that these approximations introduce. Finally, we discuss the benefits of integrating distance-bounded spatial approximations in different components of a spatial system.

2.1 Geometric Approximations

Spatial objects can have an arbitrarily complex structure. Even worse, different spatial objects can have very different structures (e.g., a point is different from a polygon). To address this challenge, spatial query processing algorithms perform geometric tests (e.g.,




Figure 1: Three example geometric approximations of a polygon: (a) Minimum Bounding Rectangle (MBR), (b) Uniform Raster (UR), (c) Hierarchical Raster (HR).

intersection, containment) on approximations of the geometries [5]. The employed approximations can represent objects with different geometries and retain the objects’ main features. In addition, they have a significantly simpler structure than the actual objects, which reduces computation and storage costs.

The most widely used spatial object approximation is the Minimum Bounding Rectangle (MBR), which is the smallest axis-aligned rectangle that encloses the complete geometry of an object (Figure 1(a)). MBRs are rather rough and inaccurate approximations. Clipped Bounding Rectangles [26] improve the accuracy of MBRs by clipping away empty space that is concentrated around the MBR corners. Brinkhoff et. al. [5] performed a detailed study of different approximations, namely the Rotated Minimum Bounding Rectangle (RMBR), the Minimum Bounding Circle (MBC), the Minimum Bounding Ellipse (MBE), the Convex Hull (CH), and the Minimum Bounding n-Corner (n-C). Raster approximations are another class of approximations that have recently attracted attention as they can provide high approximation accuracy. Raster approximations represent geometric primitives using a set of cells that can be either equi-sized [28, 35] (Uniform Raster, Figure 1(b)) or variable-sized [13, 34] (Hierarchical Raster, Figure 1(c)).

Executing spatial queries on geometric approximations leads to approximate results that are typically further processed to obtain exact answers. However, when the geometric approximation is sufficiently precise and exact answers are not required, approximate query processing techniques can provide final answers solely on the basis of the approximate geometries. In this paper, we advocate for approximate techniques with application-driven accuracy and discuss next how to bound the approximation error.

2.2 Distance Bound

Spatial queries involve predicates that evaluate relations among objects in space (e.g., intersection, containment). Therefore, we argue that it is only natural for *approximate* techniques to provide distance-based error bounds, i.e., *guarantees* on the spatial distance between false (or missing) and exact results. Approximate results without this notion of spatial distance can be misleading and hard to interpret. To illustrate this, consider the example in Figure 2. It shows a set of points corresponding to the pickup location (latitude/longitude) of taxi rides. To optimize its operational planning, the taxi service provider needs to compute the count of trips that originate from within a given region P depicted in the figure. The exact count of taxis is 18. Consider now two approximate results.

The first one is computed over the set of black and red points and equals to 22, while the second one is computed over the set of black and violet points and equals to 28. Although the first aggregate result is closer to the exact value, it contains points which are quite far away from the region P that the user is interested in, while it does not include the violet points that are closer to P . We argue that for such exploratory analyses, the second result is more meaningful as it matches more closely the user's region of interest. We further argue that in order to interpret the obtained approximate result, the user needs information about the spatial distance between the data points from which the approximate result was derived and the query geometry. In other words, it is often admissible for the user to compute the result over a region that closely approximates P , as long as she knows how close in space the approximation is.

Formally, a geometry g' ϵ -approximates a geometry g if the Hausdorff distance $d_H(g, g')$ between the two geometries is at most ϵ , where

$$d_H(g, g') = \max \left\{ \max_{p' \in g'} \min_{p \in g} d(p, p'), \max_{p \in g} \min_{p' \in g'} d(p', p) \right\}$$

and $d(p', p)$ denotes the Euclidean distance between two points. Intuitively, this ensures that any false positive (false negative) results that are present (absent) when answering queries using the approximate geometry g' are within a distance ϵ from the boundaries of the original geometry g .

Interestingly enough, not all geometric approximations can be distance-bounded. The Hausdorff distance between an object and its MBR approximation is data dependent: the coordinates of the MBR corner points are the dimension-wise maxima/minima of the bounded object. Consequently, the distance between a corner and the closest point in the object boundary can be very large.

Raster approximations, in contrast, can be distance-bounded. Given ϵ , raster approximations such as the ones shown in Figure 1, can guarantee that $d_H(g, g') \leq \epsilon$ by using a cell side length equal to $\epsilon' = \frac{\epsilon}{\sqrt{2}}$ (i.e., the length of the diagonal of the cell is ϵ) for the cells that are at the boundary of the geometry (shown with violet color). The interior cells that are fully contained in the geometry can have a cell side length larger than ϵ' as they do not contribute to the approximation error. At the boundary, there can be two types of errors, depending on the implementation. If all the cells that overlap even the slightest with the boundary are part of the approximation, then there can only be *false positive* results as the whole cells are considered to be part of the object. We call such a raster approximation *conservative*. In non-conservative raster approximations, the cells that have a small overlap with the boundary can be omitted, which can introduce *false negative* results. Overall, the precision of raster approximations is independent of the geometry they approximate and *tunable*. This property makes them particularly suitable to form the basis of approximate spatial query processing techniques.




Figure 2: Example polygon and points and two approximations of the polygon, MBR (red), and Uniform Raster (violet).




Figure 3: Uniform Raster approximation of points (left) and polygons (right). Figure from [28].

2.3 The Power of Distance-Bounded Raster Approximations

To illustrate the power of distance-bounded raster approximations, consider the example in Figure 3 showing two input data sets, a set of points (left) and a set of polygons (right) approximated with UR. **Indexing.** Figure 3 essentially shows how the data is represented *logically*: geometric objects are approximated by a set of cells, potentially along with additional information that denotes the cells that intersect with the geometry boundaries. Given this representation, a database system needs efficient indexes to store the approximations and enable their fast retrieval. Since approximate query processing eliminates the expensive refinement step, the index lookup performance is crucial because it determines the query performance. Traditional R-tree-based indexes [1] are not applicable as they are designed to index MBRs and are not compatible with raster approximations. At the same time, raster approximations enable new opportunities for a new generation of indexes. Specifically, mapping the cells to a one-dimensional array by enumerating them with a space-filling curve, enables the use of a *learned index* [15]. As we show in Section 3, by learning the position of the cells in the 1D array, the learned index outperforms other spatial index structures. **Optimization.** Section 4 discusses how, by abstracting away from the specific object geometries and providing a *unified* representation for different geometric data types, the raster approximation creates new opportunities in spatial query optimization. That is, the implementation of primitive operations (e.g., intersection tests) on the raster approximation can be *independent* of the geometries and thus re-usable, while it can also leverage modern GPUs.

Execution. Other than enabling efficient access to a single data set, the raster approximation also enables the efficient execution of queries that involve multiple data sets, such as joins. As we show in Section 5, by mapping geometries to sets of cells, we can observe the overlap at the cell level instead of performing geometry-to-geometry comparisons. Each cell can be processed independently, which makes the computation highly parallelizable. Furthermore, aggregations that are distributive or algebraic can be computed very efficiently. The final aggregate can be obtained by combining partial aggregates calculated (in parallel) for each cell.

In the following, we describe how to use distance-bounded raster approximations in various system components in more detail and present initial results.

3 DATA ACCESS

Storage layouts and index structures determine the efficiency of data access. This section shows the details of how we can build high-performing indexes for polygon and point geometries that leverage raster approximations.

Dimensionality Reduction. While raster cells could be indexed using spatial data structures such as a Quadtree, a *linearization* step can simplify the indexing problem significantly. A common approach is to map 2D cells into a 1D domain by enumerating them with a space-filling curve, such as the Hilbert or Z curve. As we will show, we can achieve much higher lookup performance with linearized cells, even compared to well-tuned 2D spatial indexes [18].

Polygon Indexing. Given the *logical* representation of polygons as a collection of linearized hierarchical cells, a database system can use different *physical* representations to store these cells, such as a B+-tree or a sorted array. Adaptive Cell Trie (ACT) [13, 14] is a recently proposed radix tree data structure for indexing linearized cells of hierarchical raster approximations. A radix tree has a clear advantage over a B+-tree or a sorted array in this setting. That is, matching cells can be found in any level of the tree, and larger cells are indexed closer to the root. Hence, larger cells are likely to be found sooner during the tree traversal. In addition, the radix tree offers implicit prefix compression as keys are not stored explicitly.

To index a set of polygons in ACT, we first perform a hierarchical raster approximation of the polygons that conforms to a user-defined distance bound (Section 2.2). ACT uses the IDs of the linearized cells to build the radix tree. To find a matching polygon for a query point, we first transform the query point to a cell on the most fine-grained grid level. Then, we traverse the radix tree with the query cell of this point and retrieve the ID of the matching polygon (if such a polygon exists).

Point Indexing. Like polygons, points are traditionally indexed with spatial data structures such as R-trees. Here, we propose to apply the same linearization for mapping 2D points to 1D cell identifiers. This again simplifies the indexing problem potentially leading to large speedups as we will demonstrate. We store the resulting 1D cell identifiers (corresponding to 2D points) in a data structure such as a B+-tree or simply in a *sorted* array.

To query the points with a polygon, we first approximate the query polygon using a hierarchical raster approximation, which yields a set of non-overlapping variable-sized cells that we call *query cells*. Then, for each cell, we perform a binary search on the sorted array to get the qualifying points. For aggregation queries (e.g., COUNT, SUM), one can pre-compute a prefix sum array and simply perform a lower and an upper bound lookup with the query cell’s boundaries [11]. By subtracting the lower bound from the upper bound, we can compute the aggregate value. In this setting, the time for computing both lower and upper bounds (essentially a binary search each) really matters. Therefore, we also explore using a learned index to speed up these searches.

We employ RadixSpline (RS) as a learned index [15]. RS consists of two main components: i) a set of spline points, and ii) a radix table to quickly determine the spline points to be examined for a lookup key (i.e., the query cell in our case). At lookup time, we first consult the radix table to determine an initial range of spline points. Next, this range is searched over to determine the spline points




Figure 4: Data access efficiency. (a) Point-polygon containment query performance. (b) Impact of the precision of the raster approximation on the number of qualifying points.

surrounding the lookup key. Finally, we use linear interpolation to predict the position of the lookup key in the sorted array. Building RS requires only one pass over the data, and is thus efficient.

Performance. We experimentally compare the performance of our proposed RS-based index with binary search (BS) and other four spatial indexes, namely, R*-tree [1] from Boost Geometry [4], Quadtree [9], STR-packed R-tree [16], and Kd-tree [3]. The spatial indexes act as baselines for filtering based on the MBR approximation. In our experiment, we use 39,200 polygons corresponding to the NYC Census regions (query polygons) and 1.2B points from the NYC taxi data set (years 2009 to 2016) [27]. We implemented the Quadtree, the STR-packed R-tree, and the Kd-tree baselines based on recent research [18]. For the Boost R*-tree, we chose the bulk-loading mode and manually optimized the number of elements per node. For the RadixSpline, we have set the number of radix bits to 25 and the spline error to 32. This experiment was run single-threaded on a two-socket Arch Linux 5.7.4 machine with an Intel Xeon Gold 6230 Processor CPU (2.10 GHz, 10 cores, 3.90 GHz turbo) and 256 GB DDR3 RAM.

Figure 4(a) shows the cumulative query time to find the total number of points inside the query polygons, while varying the precision of the raster approximation (i.e., number of approximating cells per query polygon). We compared the results of three RS-based index variations, corresponding to three precision levels (32, 128, and 512 cells per polygon), with binary search at the highest precision level used (i.e., 512) and the other four spatial baselines. Note that the spatial baselines use MBR filtering, and hence they are agnostic to changing the precision level. Clearly, the three RS-based variations outperform both Boost R*-tree and BS baselines (at least 10× and 35% better than Boost R*-tree and BS, respectively). For the Quadtree, STR-packed R-tree, and Kd-tree baselines, the RS-based variations are still either better or very close to them in terms of query time. However, as shown in Figure 4(b), RS-based variations are significantly better in terms of finding the tightest number of qualifying points compared to the exact number (precision level of 512 is almost similar to the exact case). Thus, in summary, our proposed RS-based index hits a sweet spot in the trade-off between precision and query time compared to all other baselines.

4 QUERY OPTIMIZATION

Existing approaches for spatial query processing are tied to specific geometric data representations and closely follow the relational




Figure 5: The blend and mask operators applied on rasterized canvases. The different colors are used for illustrative purposes to denote the information stored in each pixel of the rasterized canvas; the grey color denotes empty pixels.

model for query optimization [23]. They use operators that are tightly coupled to specific geometric types and query classes. Let us consider again the selection query from Figure 2. As mentioned earlier, this query is typically implemented as a *single* operator that uses two phases: filtering and refinement. While the filtering phase relies on MBRs and is thus generic, the refinement phase depends on the geometric type and operation. In this example, the refinement is specific to the input being points, and the performed operation is a point-in-polygon test. If the input changes from taxi pickup locations to restaurants represented by polygons, then a different implementation is required, since a polygon-intersect-polygon test must be performed instead. The use of such large monolithic operators limits the set of options over which optimization can be performed, as the operators cannot be reused across query classes.

To overcome these limitations, and to exploit modern GPUs, a GPU-friendly spatial data model and algebra was introduced in [6], which proposes a uniform data representation called *canvas* and a small set of simple parallelizable operators. These operators include common computer graphics operations: *blend*, *mask*, and *affine transformations*. More importantly, these operators are sufficient to realize common spatial query classes without being tied to specific geometries. For instance, both point-polygon and polygon-polygon intersection tests boil down to applying a combination of the above operations on the canvas. We propose to adapt the canvas model to support distance-bounded approximate queries: the canvas now simply becomes a rasterized image, where the pixel size depends on the required bound. The GPU-amenable operators work directly on such a *rasterized canvas*—in fact, the implementation of these operators now becomes straightforward since boundary conditions [6] need not to be taken care of. Figure 5 illustrates examples for the blend and the mask operators. The *blend* binary operator merges two rasterized canvases into one. The blend function \odot defines how the merge is performed. The *mask* operator filters pixels of the rasterized canvas to retain only those pixels that satisfy the condition specified by M . There are two ways to generate a rasterized canvas: by rendering the data directly on the GPU, or through the use of indexes (e.g., using ACT described in Section 3).

The rasterized canvas along with the proposed set of operators enable the creation of multiple alternative plans to realize any given ad-hoc query, thereby adding flexibility in the optimization process. Furthermore, each operator can have multiple implementations and indexes can be reused across operators, which provides a wider set of options for the optimizer. Thus, the optimizer can choose different query plans based on the query parameters, the distance bound (i.e., the resolution of the rasterized canvas), and the estimated selectivity. As an example of the potential gains that our proposed

model provides, we show in Section 5.2 how the model allows for an alternate plan for an approximate spatial aggregation query that performs significantly faster than traditional approaches.

5 QUERY EXECUTION

This section highlights the benefits of distance-bounded raster approximations in query evaluation. As a representative example, we focus on the evaluation of spatial aggregation queries defined as follows in SQL-like notation:

```
SELECT AGG( $a_i$ ) FROM P, R
WHERE P.loc INSIDE R.geometry [AND filterCondition]*
GROUP BY R.id
```

Given a set of points of the form $P(loc, a_1, a_2, \dots)$, where loc and a_i are the location and attributes of the point, and a set of regions $R(id, geometry)$, this query performs an aggregation (AGG) over the result of the join between P and R . The geometry of a region can be any *arbitrary polygon*. Functions such as COUNT(*) or AVG(a_i) are commonly used for AGG.

This query typically uses point-in-polygon (PIP) tests to identify polygons that contain each of the points. Note that each PIP test requires time linear with respect to the size of the polygon. Since real-world polygonal regions often consist of hundreds of vertices, these tests are computationally intensive. This challenge is compounded by the fact that data sets can have hundreds of millions, or even billions of points, requiring a large number of PIP tests to be performed.

Existing systems typically evaluate spatial aggregation queries by performing a spatial join of the points and the polygons, followed by the aggregation of the join results. To reduce the number of PIP tests, the join is first solved using MBR approximations. As we show next, our evaluation strategies that are based on raster approximations, outperform the above approach significantly.

5.1 Main-Memory Join

Using our ACT index (Section 3), we can evaluate the query with an index-nested loop join: we simply index the polygons with ACT, and query the radix tree for every point. We combine the join with the aggregation to avoid materializing the join result. Given that ACT employs a fine-grained distance-bounded HR approximation, we omit the PIP tests and provide approximate results.

Performance. We experimentally compare the performance of our approximate join with exact joins using the Boost [4] R*-tree [1] and Google’s S2ShapeIndex (SI)², all implemented in C++. ACT uses HR polygonal approximations satisfying a 4m distance bound. The R-tree indexes the polygons’ MBRs, while, similarly to ACT, SI uses HR approximations. However, SI’s approximation is not distance-bounded and SI does not support approximate evaluation. We use 1.2B points from the NYC taxi data set [27] and three NYC polygon data sets: Boroughs (5), Neighborhoods (289), and Census (39,200). This experiment was run single-threaded on a machine with 14-core Intel Xeon E5-2680 v4 CPUs and 256 GB DDR4 RAM. Figure 6 shows that the ACT-based approximate join significantly outperforms other approaches. Compared to the R*-tree, it brings over two orders of magnitude improvement for Boroughs, and over one order

²<https://s2geometry.io/devguide/s2shapeindex>

otherwise, while it is over one order of magnitude faster than SI in all cases. The low performance of the R*-tree for Boroughs is due to the fact that Boroughs are more complex polygons than Neighborhoods and Census and thus PIP tests are more expensive. Specifically, Boroughs have 663 vertices per polygon on average, while Neighborhoods have 30.6 and Census 13.6. Therefore, reducing the number of PIP tests by approximating the polygons more closely (as SI does) or completely eliminating them by using distance-bounded fine-grained approximations (like ACT) has a significant impact on performance. On the contrary, in the case of Census, which are the simplest polygons, ACT brings the least improvement. Experiments on other data sets [13] also confirm the above findings.

Overall, ACT trades memory consumption for approximation accuracy, which in turn enables approximate evaluation and leads to higher performance. Therefore, ACT has higher space consumption compared to the other approaches. For example, the HR approximation of the Neighborhood polygons consists of 13.2M cells, which are represented using 64-bit integer IDs. The total size of ACT is 143 MB. In contrast, SI that uses a coarser-grained HR approximation occupies 1.2 MB, while the R*-tree that approximates polygons even more coarsely using MBRs occupies only 27.9 KB.

5.2 GPU Join

Section 4 outlined the use of a rasterized canvas model for executing spatial queries on GPUs. Here we show the gains that the proposed model brings in the evaluation of spatial aggregation queries. In fact, the query can be realized by simply combining a small set of operators from our query algebra on top of the rasterized canvas model. This is exactly what our recently proposed algorithm, Bounded Raster Join [7, 28] (BRJ), does. Intuitively, BRJ takes as input a uniform representation of the points and polygons on rasterized canvases. It then merges (using the *blend* operator) all the points into a single canvas that maintains partial aggregates, i.e., each canvas pixel keeps the aggregate of all points falling in that pixel. Then, it joins this canvas with the set of polygon canvases (by composing the *blend* and *mask* operators) to identify points that intersect with the polygons, and finally merges the results (using a combination of *transformations* and *blending*) to compute the final aggregates. That is, it combines the aggregates from the individual pixels that fall within a polygon to generate the aggregation for that polygon. The precise query plan can be found in [6]. The above operations are natively supported by the graphics pipeline, leading to orders of magnitude speedup over typical evaluation strategies on CPUs without requiring any pre-computation [28].

Performance. We implemented BRJ using C++ and OpenGL. We create the canvases on-the-fly by simply rendering the geometries onto an off-screen buffer and store the aggregates in the buffer's color channels (r,g,b,a). We experimentally compare BRJ with an




Figure 6: Main-memory join.

accurate GPU Baseline that follows the traditional index-based evaluation strategy of first filtering the polygons with a grid index (with 1024² cells) and then performing PIP tests. This experiment was run on a machine with an Intel Core i7 Quad-Core CPU, 16 GB RAM, and an NVIDIA GTX 1060 mobile GPU with 6 GB of memory, out of which we use only 3 GB. We join 600M points of the NYC taxi data set [27] (transferred in batches to the GPU) with 260 NYC neighborhood polygonal regions (some of the regions are multi-polygons) and count the number of points in each region. Figure 7 shows that there is a trade-off between the accuracy and the query time. For a distance bound of 10m, BRJ is about 8.5× faster than the baseline, while for 1m it becomes slower. This is because lower bounds require smaller pixel sizes, and hence increasing the canvas resolution. When this resolution becomes higher than what the GPU supports, BRJ needs to divide the rasterized canvas and perform multiple aggregations, one for each subdivision. We note, however, that with a distance bound of 10m we get close to accurate counts: over all the polygons, the median error is only about 0.15%. BRJ can therefore provide a significant speedup with only a small accuracy loss. The accuracy-time trade-off has a similar behavior for larger inputs as well as for other data sets [28].




Figure 7: Bounded Raster Join (GPU). Impact of the distance bound on performance.

6 DISCUSSION

Synopsis-based Approximate Spatial Query Processing. Approximate Query Processing (AQP) typically refers to extracting small data synopses (e.g., samples) from large spatial data sets, and performing accurate evaluation on top of those samples, yielding approximate answers due to the initial data reduction [25]. Prior work in that direction [32, 33] does not provide support for arbitrary spatial queries such as joins and group-by predicates. Furthermore, most existing methods do not provide any accuracy guarantees and do not have the notion of distance bounds. Initial efforts to provide such guarantees [32] focus on the selectivity estimation problem and only provide bounds on the relative error between the actual and the estimated selectivity.

The above line of work is orthogonal to what we propose in this paper. We focus on approximations in space, i.e., approximations of individual object geometries, and on tunable distance bounds that control the spatial accuracy of the approximations.

Result Range Estimation. Rather than providing only an approximate result, we can use the raster approximation to provide a result range based on the key insight that errors happen only at the boundary cells. Therefore, by counting the number of results contained in these cells we can get loose bounds on the result range. For example, let us assume that we have a *conservative* raster approximation, i.e., we can only have false positives at the boundary, and let α be the approximate count of points within a polygon. Let C be the set of

cells at the boundary and ϵ be the partial count computed over C . Then, we know that the result falls in the interval $[\alpha - \epsilon, \alpha]$ with 100% confidence. In the above calculation, we assume that all the results at the boundary are false positives, which is the worst case. By making some assumptions about the distribution of points at the boundary, we can obtain a tighter interval.

Higher-Dimensional Data. Even though this paper focuses on 2D primitives, the proposed distance-bounded approximation can be directly extended to support 3D primitives. However, the proposed operators do not have a straightforward GPU implementation over 3D data. In our future work, we plan to investigate extensions to our techniques to handle 3D data.

GPU Rasterization vs. Ray Tracing. This work shows the benefits of using the GPU rasterization pipeline in spatial data processing. Given that spatial databases rely on the same primitive types (geometric objects) and operations that are similar to the ones used in graphics (e.g., spatial selections), we expect further opportunities to exploit advanced graphics techniques and hardware in the design of spatial systems. In future work, it will be particularly interesting to explore the use of native GPU ray tracing, recently introduced by RTX GPUs from Nvidia [21]. Ray tracing can be, for example, used to support 3D spatial queries.

7 CONCLUSION

Changes in applications requirements and hardware have been the main driving forces in rethinking the role of geometric approximations in spatial data management. This paper shows that distance-bounded raster approximations can enable a wider set of optimization options and can form the basis of approximate spatial query processing techniques that take better advantage of modern hardware and improve performance. Our experiments demonstrate that raster approximations can be indexed efficiently and can provide a sweet spot in the trade-off between precision and query time. In doing so, we set the stage for new spatial systems that employ distance-bounded raster approximations at their core.

ACKNOWLEDGMENTS

This work was partially supported by the German Ministry for Education and Research as BIFOLD - Berlin Institute for the Foundations of Learning and Data (ref. 01IS18025A and ref 01IS18037A). This research was further supported by Google, Intel, and Microsoft as part of the MIT Data Systems and AI Lab (DSAIL) at MIT, NSF IIS 1900933, DARPA Award 16-43-D3M-FP040. Ibrahim Sabek was supported by the NSF, under grant #2030859 to the Computing Research Association for the CIFellows Project. Harish Doraiswamy was supported in part by the NYU Moore Sloan Data Science Environment and the NSF award CCF-1533564.

REFERENCES

- [1] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger. The R*-tree: An Efficient and Robust Access Method for Points and Rectangles. In *Proc. SIGMOD*, pages 322–331, 1990.
- [2] A. Belussi, B. Catania, and S. Migliorini. Approximate queries for spatial data. In *Advanced Query Processing: Volume 1: Issues and Trends*, pages 83–127. Springer Berlin Heidelberg, 2013.
- [3] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM (CACM)*, 18(9):509–517, 1975.
- [4] Boost Geometry. <https://github.com/boostorg/geometry/>.
- [5] T. Brinkhoff, H.-P. Kriegel, and R. Schneider. Comparison of approximations of complex objects used for approximation-based query processing in spatial database systems. In *Proc. ICDE*, pages 40–49, 1993.
- [6] H. Doraiswamy and J. Freire. A gpu-friendly geometric data model and algebra for spatial queries: Extended version. *arXiv:2004.03630 [cs.DB]*, 2020.
- [7] H. Doraiswamy, E. Tzirita Zacharatos, F. Miranda, M. Lage, A. Ailamaki, C. T. Silva, and J. Freire. Interactive Visual Exploration of Spatio-Temporal Urban Data Sets Using Urbane. In *Proc. SIGMOD*, pages 1693–1696, 2018.
- [8] H. Doraiswamy, H. T. Vo, C. T. Silva, and J. Freire. A GPU-based index to support interactive spatio-temporal queries over historical data. In *Proc. ICDE*, pages 1086–1097, 2016.
- [9] R. A. Finkel and J. L. Bentley. Quad trees: A data structure for retrieval on composite keys. *Acta Inf.*, 4:1–9, 1974.
- [10] R. H. Güting. An introduction to spatial database systems. *The VLDB Journal*, 3(4):357–399, Oct. 1994.
- [11] C. Ho, R. Agrawal, N. Megiddo, and R. Srikant. Range queries in OLAP data cubes. In *Proc. SIGMOD*, pages 73–88, 1997.
- [12] E. H. Jacox and H. Samet. Spatial join techniques. *ACM Trans. Database Syst.*, 32(1):7–es, Mar. 2007.
- [13] A. Kipf, H. Lang, V. Pandey, R. A. Persa, C. Anneser, E. Tzirita Zacharatos, H. Doraiswamy, P. A. Boncz, T. Neumann, and A. Kemper. Adaptive main-memory indexing for high-performance point-polygon joins. In *Proc. EDBT*, pages 347–358, 2020.
- [14] A. Kipf, H. Lang, V. Pandey, R. A. Persa, P. A. Boncz, T. Neumann, and A. Kemper. Approximate geospatial joins with precision guarantees. In *Proc. ICDE*, pages 1360–1363, 2018.
- [15] A. Kipf, R. Marcus, A. van Renen, M. Stoian, A. Kemper, T. Kraska, and T. Neumann. RadixSpline: a single-pass learned index. In *Proc. aIDM@SIGMOD*, pages 5:1–5:5, 2020.
- [16] S. T. Leutenegger, M. A. López, and J. M. Edgington. STR: A simple and efficient algorithm for r-tree packing. In *Proc. ICDE*, pages 497–506, 1997.
- [17] Z. Liu and J. Heer. The Effects of Interactive Latency on Exploratory Visual Analysis. *Proc. TVCG*, 20(12):2122–2131, 2014.
- [18] V. Pandey, A. van Renen, A. Kipf, I. Sabek, J. Ding, and A. Kemper. The Case for Learned Spatial Indexes. In *AIDB Workshop @ VLDB*, 2020.
- [19] M. Pavlovic, D. Sidlauskas, T. Heinis, and A. Ailamaki. QUASII: QUery-Aware Spatial Incremental Index. In *Proc. EDBT*, pages 325–336, 2018.
- [20] M. Pavlovic, E. Tzirita Zacharatos, D. Sidlauskas, T. Heinis, and A. Ailamaki. Space odyssey: efficient exploration of scientific data. In *Proc. ExploreDB@SIGMOD*, pages 12–18, 2016.
- [21] Nvidia Ray Tracing. <https://developer.nvidia.com/rtx/raytracing>.
- [22] I. Sabek and M. F. Mokbel. On Spatial Joins in MapReduce. In *Proc. SIGSPATIAL*, pages 21:1–21:10, 2017.
- [23] H. Samet and W. G. Aref. Spatial data models and query processing. In W. Kim, editor, *Modern Database Systems*, pages 338–360. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 1995.
- [24] B. Shneiderman. The eyes have it: a task by data type taxonomy for information visualizations. In *Proc. VL/HCC*, pages 336–343, 1996.
- [25] A. B. Siddique, A. Eldawy, and V. Hristidis. Comparing synopsis techniques for approximate spatial data analysis. *PVLDB*, 12(11):1583–1596, 2019.
- [26] D. Sidlauskas, S. Chester, E. Tzirita Zacharatos, and A. Ailamaki. Improving spatial data processing by clipping minimum bounding boxes. In *Proc. ICDE*, pages 425–436, 2018.
- [27] TLC Trip Record Data. <https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>.
- [28] E. Tzirita Zacharatos, H. Doraiswamy, A. Ailamaki, C. T. Silva, and J. Freire. GPU rasterization for real-time spatial aggregation over arbitrary polygons. *PVLDB*, 11(3):352–365, 2017.
- [29] E. Tzirita Zacharatos, D. Sidlauskas, F. Tauheed, T. Heinis, and A. Ailamaki. Efficient bundled spatial range queries. In *Proc. SIGSPATIAL*, pages 139–148, 2019.
- [30] F. van Diggelen and P. Enge. The world's first GPS MOOC and worldwide laboratory using smartphones. In *Proc. ION GNSS+*, pages 361–369, 2015.
- [31] D. Vorona, A. Kipf, T. Neumann, and A. Kemper. DeepSPACE: Approximate geospatial query processing with deep learning. In *Proc. SIGSPATIAL*, pages 500–503, 2019.
- [32] T. Vu and A. Eldawy. DeepSampling: Selectivity Estimation with Predicted Error and Response Time. In *Proc. DeepSpatial@SIGKDD*, 2020.
- [33] L. Wang, R. Christensen, F. Li, and K. Yi. Spatial Online Sampling and Aggregation. *PVLDB*, 9(3):84–95, 2015.
- [34] C. Winter, A. Kipf, C. Anneser, E. Tzirita Zacharatos, T. Neumann, and A. Kemper. Geoblocks: A query-cache accelerated data structure for spatial aggregation over polygons. *arXiv:1908.07753 [cs.DB]*, 2020.
- [35] G. Zimbrão and J. M. d. Souza. A Raster Approximation For Processing of Spatial Joins. In *VLDB*, pages 558–569, 1998.