

Filtering with Raster Signatures

Leonardo Guerreiro Azevedo¹, Ralf Hartmut Güting², Rafael Brand Rodrigues¹,
Geraldo Zimbrão^{1,3}, Jano Moreira de Souza^{1,3}

¹Computer Science Department, Graduate School of Engineering, Federal
University of Rio de Janeiro,
PO Box 68511, ZIP code: 21945-970, Rio de Janeiro, Brazil

²LG Datenbanksysteme für neue Anwendungen, FB Informatik,
Fernuniversität Hagen, D-58084 Hagen, Germany

³Computer Science Department, Institute of Mathematics,
Federal University of Rio de Janeiro, Rio de Janeiro, Brazil

Email: {azevedo, brand, zimbrao, jano}@cos.ufrj.br, rhg@fernuni-
hagen.de

***Abstract.** Efficient evaluation of spatial queries is an important issue in spatial database. Among spatial operations, spatial join is very useful, and intersection is the most common predicate. However, the exact intersection test of two spatial objects is the most time-consuming and I/O-consuming step in processing spatial joins. On the other hand, the use of approximations can reduce the need for examining the exact geometry of spatial objects in order to find the intersecting ones. This work proposes a new raster approximation (Three-Colors Raster Signature - 3CRS) for representing different data types (polygons, polylines and points) and to be used as filter in the second step of Multi-Step Query Processor. We have also executed experimental tests over real datasets, and the results demonstrated the effectiveness of our approach.*

1. Introduction

The increase of storage capacity and the decrease of hardware costs have made possible for applications to deal with large amount of data, involving Gigabytes, Terabytes and even Petabytes of information. This characteristic is common in Spatial Databases where data usually have high complexity and are available in huge amounts.

Spatial data consists of spatial objects made up of points, lines, regions, rectangles, surfaces, volumes, and even data of higher dimension which includes time (Samet, 1990). Examples of spatial data include cities, rivers, roads, counties, states, crop coverage, mountain ranges etc. It is often desirable to attach spatial with non-spatial attribute information. Examples of non-spatial data are road names, addresses, telephone numbers, city names, etc. Since spatial and non-spatial data are so intimately

connected, it is not surprising that many of the issues that need to be addressed are in fact database issues.

There are numerous applications in spatial database systems area, such as: traffic supervision, flight control, weather forecast, urban planning, route optimization, cartography, agriculture, natural resources administration, coastal monitoring, fire and epidemics control, precision agriculture and intelligent highways (Aronoff, 1989; Tao *et al.*, 2003; Gordon *et al.*, 1994). Each type of application deals with different features, scales and spatiotemporal properties.

Efficient evaluation of spatial queries is an important issue in spatial database. Among spatial operations, spatial join is very useful. Intersection is the most common join predicate. Many works point the exact geometry test as the most time consuming operation regarding both I/O and CPU. Brinkhoff (1994) shows experimental results confirming that the exact geometry test, usually plane-sweep (Boissonnat, 1997; Freiseisen, 1998) is responsible for most of the CPU cost. The I/O cost associated with the exact geometry test is due to the access to the real representation of the spatial objects, which can be very large. Spatial joins have been well studied in the literature, and there are many approaches to processing spatial join operations. Considering points, polylines and polygons as the three data types most common in spatial databases, there are nine classes of different spatial joins. For its usefulness and complexity, the polygon join has been the most investigated, while the point join has been the less investigated because of its similarity with the relational join (Samet, 1990), while there are some proposals for processing polylines and polygon \times polyline joins. In order to improve the efficiency, organization and study of indices and filters for spatial data, Brinkhoff *et al.* (1994) proposed a three-step architecture for spatial join processing, named as Multi-Step Query Processor. The main target of this architecture is to accelerate the most costly step by reducing the number of spatial objects left to be compared. Such a reduction is done by applying filters in previous steps.

This work proposes a new raster approximation suitable to performing spatial joins as a filter in the second step of Multi-Step Query Processor, involving these three common data types (polygon, polyline and point) and the classes of different spatial joins involving them. We propose a raster signature named as Three Color Raster

Signature (3CRS), based on the Four Color Raster Signature (4CRS) proposed by Zimbrao and Souza (1998). The 3CRS signature has the advantages of faster generation time and that can be used for representing polygons, polylines and points (without any specific characteristic). Besides, the same algorithm can be used to evaluate the join predicate involving these three data types. Also, the fast generation time allows generating the signature on the fly. For instance, instead of storing the signature, it can be generated only when it is needed, saving storage space. In order to evaluate the effectiveness of our proposal, we execute spatial joins using 3CRS against the processing without using signatures and the processing using 4CRS. The experimental tests were executed over real data and the results demonstrated the effectiveness of the approach.

This paper is divided in sections, as follows. Section One is this introduction. Section Two surveys the related literature. In Section Three, we present our raster approximation and its implementation as the second step of the MSQP architecture. Section Four shows the experimental results. Finally, in Section Five we present our conclusions.

2. Related Work

Brinkhoff *et al.* (1994) define spatial join as a subset of the Cartesian product of two sets, A and B , not necessarily distinct, containing, respectively, m and n elements. This resulting subset is composed of elements that meet a given spatial predicate. The overlap of spatial objects is of special interest in practical applications.

Spatial joins may be more efficiently evaluated by means of indexes previously built on each data set and simultaneously traverse these indexes, searching for object intersections. This approach corresponds to a sort-merge approach in relational databases, and it is pointed in (Brinkhoff *et al.*, 1993b) as being quite efficient, especially when the indexes already exist. Typically, an index is composed of two parts: the index structure, that only stores the data keys; and the data structure, that stores the data itself. Thus, a spatial index should store the objects in a spatial structure according to a geometric key (Brinkhoff *et al.*, 1994). Due to its simplicity, MBR (minimum bounding rectangle) is the most popular geometric key. When we use MBRs, the complexity of a spatial object is reduced to four parameters which retain the most

important characteristics of the object: position and extension. Nonetheless, as pointed in (Brinkhoff *et al.*, 1994) and (Zimbrao and Souza, 1998), objects of the real world are very poorly represented by MBRs.

2.1. Architectures for Processing Spatial Joins

There are many approaches to processing spatial join operations. Zhu *et al.* (2000) emphasize that traditional approaches perform the spatial join processing in two steps (Orenstein, 1986; Kothuri and Ravada, 2001). They propose efficient algorithms to be used in the second step. In the two steps approach, presented in Figure 1a, the first step employs a Spatial Access Method (SAM) in order to reduce the search space. The Minimum Bounding Rectangle (MBR) is usually used by SAM methods. This step does not have as output the result of the join operation. Instead, it provides a set of candidate pairs that correspond to a super-set of the solution, and that is sent to the second step. The second step is a refinement step where the pairs resulting from the first step are read from disk and have their geometries processed. This is the most costly step, requiring I/O time to seek and read the spatial objects from disk, and CPU time to compute the exact answer.

Brinkhoff *et al.* (1994) propose a three-step architecture for processing of spatial join named as Multi-Step Query Processor (MSQP), presented in Figure 1.b. In this architecture, another step is included between the first (SAM) and the second (Exact geometry processor) steps. The proposed step consists in comparing the candidate pairs resulted from the first step using a geometric filter. The geometric filter uses a compact and approximated representation of object trying to retain its main characteristics. Examples of such proposals of object representations are: 4CRS (Zimbrao and Souza, 1998), Convex Hull, 5C, RMBR and others found in Brinkhoff *et al.* (1993). As the result of this step we have three possibilities: pairs that belongs to the solution (hit); pairs that do not belong to the solution (false hit); and, pairs that it is not possible to have a conclusive answer (inconclusive comparisons). The later are sent to the third step, the refinement step, where the pairs of objects are read from disk and have their geometries processed.

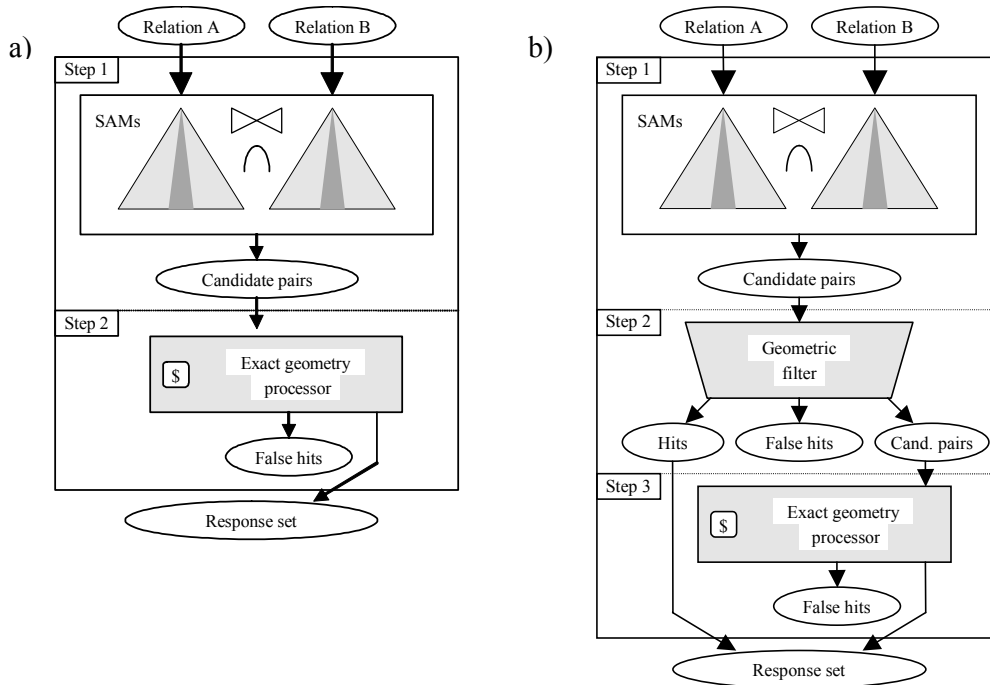


Figure 1. a) Two steps architecture for processing of spatial join; b) Three-step architecture (Brinkhoff et al., 1994).

There are two main advantages in introducing the filter step. First, the approximation size is only a fraction of the spatial object size; therefore it can be stored in the index, together with the object MBR. Second, testing two approximations requires less CPU time than testing two objects. The pair of objects that have a conclusive test (hit or false hit) are not sent to the third step.

3. Three-Color Raster Signature

In this section we will present the characteristics of 3CRS. Section 3.1 presents 3CRS main characteristics. Section 3.2 proposes an algorithm to generate 3CRS. A simple algorithm for changing the resolution of 3CRS is proposed in Section 3.3, and the algorithm to evaluate if two objects overlap using their 3CRS is presented in Section 3.4.

3.1 3CRS Characteristics

The 3CRS is based on 4CRS (Zimbrao and Souza, 1998). 3CRS is a compact and approximated raster representation of objects upon a grid of cells that uses few colors. Each color represents an intersection type between the object and the cell (Table

1). Figure 2 presents an example of a 3CRS representation of a polygon. Actually, 3CRS is a 4CRS where the *Weak* and *Strong* cell types are replaced by an *Inconclusive* type. The *Weak* 4CRS cell type represents that the polygon has an intersection equal or less than 50% with the cell, and the *Strong* type represents an intersection greater than 50% and less than 100%. The 3CRS *Inconclusive* cell type replaces these two types, and it represents that there is a portion of the object within cell, which does not overlap the whole cell. This characteristic allows 3CRS to represent polylines and points in the same way it represents polygons.

Table 1. 3CRS cell types

Cell type	Description
Empty	The cell is not intersected by the object.
Inconclusive	There is a portion of the object within cell, and it does not fulfill the cell.
Full	The cell is fully occupied by the object. This type of cell only exists when representing polygons.




Figure 2. Example of 3CRS representation of a polygon.

When computing a 3CRS it is not required to compute the exact area of polygon within cell, as is done when computing 4CRS signature. In other words, when generating the signature, we do not need to clip the polygon against the cell, and to compute the area of the clipping region, which is very costly. Instead, it is required only to evaluate if the cell is crossed by the object. As result, each cell type is computed fast, and objects that do not have area within cell, such as points and polylines, can be represented using 3CRS. In the case of points, the 3CRS signature is composed by only one inconclusive cell, while regards to polylines, the 3CRS contains *Empty* cells (cells that are not intersected by the polyline) and *Inconclusive* cells (cells that are crossed by the polyline).

3.2 Algorithm to generate 3CRS

The algorithm to generate 3CRS can be divided into three steps:

1. Compute the MBR- 2^n that encloses the object;
2. Follow the object segments, setting as *Inconclusive* the type of the cells they intersect;
3. If the object is a polygon, scan the signature cells, adjusting the type of unmarked cells as *Empty* or *Full* according to if the cell is outside the polygon or inside it, respectively.

In the first step, the MBR- 2^n of the polygon is computed according to the algorithm presented in Zimbrao and Souza (1998), and a grid of empty cells is computed from this MBR.

In the second step, the algorithm scans the object's segments. For each segment s , it goes from the first cell (where an ending point of s is within) to the last cell (where the other ending point of s is within) marking the type of these cells as *Inconclusive*. The algorithm goes from one cell to another according the edge of the cell that is intersected by s (left, right, bottom or top edge). The cell's edge that is intersected by the segment s is identified using the Sutherland-Cohen line clipping algorithm (Newman and Sproull, 1979). The cell adjacent to the current cell's edge is marked as *Inconclusive* and becomes the current cell. The process is repeated until the current cell is equal to the last cell.

It is important to emphasize some aspects about this step:

- There is no specific order to evaluate polygon segments;
- It is possible to get a cell that is indexed in the signature array in a constant time using a mod operation. Hence one can compute the cell for a point in constant time;
- Given the line equation, one can go along the line and mark intersected cells in a constant time per cell. Hence for a line segment, the time required will be linear in the number of marked cells;

- Usually a line segment is short and will mark only a constant number of cells, hence requiring $O(1)$ time per line segment;
- Finally, this step compute all partial cells (*Inconclusive* cells), and it needs $O(n)$ time if all segments are sufficiently short.

During the segment evaluation, if the object from which the signature is being computed is a polygon, a matrix of border intersection is used to store a value for each cell, representing a type of intersection between the segment and the left and right borders of the cell. This value is based on the *InsideAbove* flag of the segment. The *InsideAbove* flag of a segment is true when the area inside the polygon lies above the segment; or, if the segment is a vertical line, it indicates that the area inside the polygon is on the left of the segment. The value of 1 is added to the cell value if the *InsideAbove* attribute of the segment is true. Otherwise, 1 is deducted from the cell value. Figure 3 3 shows an example of matrix values. The inconclusive cells of the signature are represented as dark gray cells.

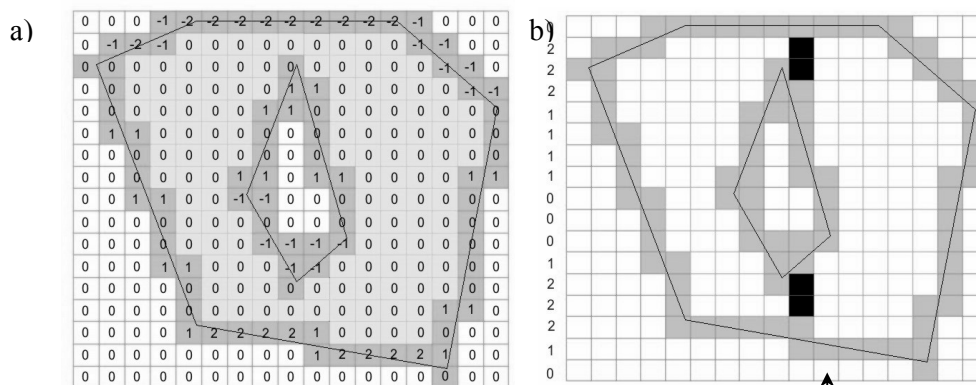


Figure 3. (a) Matrix of border intersection (b) example of the third step execution

The third step (only used when computing 3CRS from polygons) is responsible for marking the cells that are inside the polygon as *Full* cells. The algorithm goes from bottom to top, following each column of the *matrix*, summing the value of the matrix of border intersection corresponding to each cell. A variable named *counter* is used to store this sum. For each cell which type was not already set *Inconclusive*, we evaluate the current value of *counter* variable. Since the maximum number of cells of the grid is a fixed value K , then this step requires $O(K)$ time. Figure 3.b shows an example of the third step execution for the column pointed by the arrow. This example uses the matrix

presented in Figure 3.a. The inconclusive cells are already marked (they are represented as the gray cells). On the left of the grid, the numbers correspond to the values of the *counter*, as the algorithm goes through each row. The counter starts with zero. In the first row (from bottom to top), the polygon do not intersects the sides of the cell, so the value of counter stays zero. The cell is not marked as *Inconclusive*, so it stays as *Empty*, because counter is equal to zero. On the second row, the correspondent cell of the *matrix* has the value 1 (as calculated on the previous step of the algorithm, because one segment of the polygon intersects the right side of the cell and the *InsideAbove* flag of that segment is true), so we increase in one the value of counter. The cell was marked before as *Inconclusive*, so its type does not change. The same occurs in the third row. On the fourth row, the value of counter is two, and the cell is not *Inconclusive*. Hence the cell is marked as *Full*. The algorithm continues to the end of the column (and lines), until all the cells are marked.

3.3 Intersection test using 3CRS

When evaluating two 3CRS signatures of polygons, it is essential that both of them have the same cell size. If it does not apply, it is imperative to perform a change of scale. Whenever a change of scale is necessary, it is accomplished through the grouping of 2^m cells, having in mind that the coordinates of the beginning of each cell are proportional to the length of its side. An algorithm for scale change is presented in Zimbrao and Souza (1998).

After performing the scale changes (if it is required), the cells of two 3CRS that overlap each other are evaluated. Only the cells that are inside the intersection MBR of the signatures are processed. The result of the comparison between two cells is presented in Table 2. Note that there is only one *Inconclusive* result, and it occurs when comparing two *Inconclusive* cells.

Table 2. Possible results when comparing two cells.

	Empty	Inconclusive	Full
Empty	No	No	No
Inconclusive	No	Perhaps	Yes
Full	No	Yes	Yes

Figure 4 presents a proposal of algorithm for 3CRS comparison based on the comparison results of pair of cells presented Table 22. If all cell comparisons result in “No”, then there is no intersection between the polygons. On the other hand, if a “YES” result is found, it means that the polygons intersect, and the comparison can stop. During the comparison, if there is a “PERHAPS” result and no “YES” result, then it is not possible to ensure that the polygons intersect or that they do not has intersection. In this case, it is needed to execute the refinement step, seeking and reading the polygons’ exact representations from disk and executing the exact test.

```

algorithm hasIntersection(signat3CRS1, signat3CRS2)
  interMBR = intersectionMBR(signat3CRS1, signat3CRS2);
  if (signat3CRS1.lengthOfCellSide <
      signat3CRS2.lengthOfCellSide)
    s3CRS = changeScale(signat3CRS1, signat3CRS2.lengthOfCellSide);
    b3CRS = signat3CRS2;
  else
    if (signat3CRS1.lengthOfCellSide >
        signat3CRS2.lengthOfCellSide)
      b3CRS = signat3CRS1;
      s3CRS = changeScale(signat3CRS2, signat3CRS1.lengthOfCellSide);
    else
      s3CRS = signat3CRS1;
      b3CRS = signat3CRS2;
  result = NO;
  for each b3CRS cell b that is inside interMBR do
    for each s3CRS cell s that intersects cell b do
      if b.type == EMPTY or s.type == EMPTY
        continue;
      if b.type == INCONCLUSIVE or s.type == INCONCLUSIVE
        result = PERHAPS;
      if ( (b.type == FULL) and
          (s.type == FULL or s.type == INCONCLUSIVE) ) or
          ( (s.type == FULL) and
            (b.type == FULL or b.type == INCONCLUSIVE) )
        return YES;
  return result;

```

Figure 4. Algorithm for 3CRS comparison.

4. Experimental Tests

This section is dedicated to presenting the experimental results concerning to the evaluation of use of 3CRS signature in query processing. We evaluated the use of 3CRS as a filter in the second step of MSQP (Brinkhoff *et al.*, 1994) against the use of 4CRS, and against the processing without a filter step, the architecture of two steps (Orenstein, 1986; Kothuri and Ravada, 2001). We evaluated the intersection join of set of polygons.

4.1. Experimental Data Sets

The polygon real datasets used in the experiments consist of township boundaries, census block-group, geologic map and hydrographic map from Iowa

(USA), available online at “<http://www.igsb.uiowa.edu/nrgis/gishome.htm>”, and Brazilian municipalities (IBGE, 1996). In order to simulate large datasets, the Iowa datasets were replicated six times, in the same way as suggested by Brinkhoff *et al.* (1994). The original polygons were shifted by random displacements of x and y coordinates. In the case of Brazilian municipalities, we performed one replication (named Brazilian municipalities’), so that we could execute the test of Brazilian municipalities against Brazilian municipalities’. Data characteristics are presented in Table3.

Table 3. Test datasets

Datasets		Size (KB)	# pol.	# segments	Avg. # segments
Iowa	Census block group	38,824	17,844	1,764,588	98
	Topography	61,748	20,070	3,780,552	188
	Hydrologic map	6,904	2,544	475,434	186
	Township boundaries	25,288	12,216	1,059,438	86
	Geologic maps	21,856	9,984	640,428	64
Brazil	Municipalities	9,840	4,645	399,002	85
	Municipalities’	9,840	4,645	399,002	85
Average		24,757	10,278	1,216,921	118

4.2. Test Environment and R*-tree characteristics

Tests were executed on a PC Athlon XP 1600+ 1.4 GHz with 256 MB of RAM. A page size of 2,048 bytes for I/O operations was defined.

The R*-tree (Beckmann *et al.*, 1990) was chosen as a spatial access method in order to reduce the search space. In other words, the R*-Tree was used to take account only the objects that have at least MBR intersection and not all of them. That choice was due to the wide use of R*-Tree, as well as, to the successful results found in the literature. The access methods traditionally used employ the object’s Minimum Bounding Rectangle (MBR), and the access methods execution returns what is called a set of candidates, since it contains all pairs of polygons that belong to the answer plus other pairs that have only MBR intersection. In the same way as Brinkhoff *et al.* (1994) and Zimbrão and Souza (1998) do, for our tests we generated R*-Trees that store the 4CRS signatures as part of the polygons’ keys. This means that they were stored in the leaf nodes of the R*-Tree index.

The tests (Table 3) can be described according to the concepts presented in Sub-Section 2.1 (Architectures for Processing Spatial Joins). The experimental tests using 3CRS and 4CRS were executed according to the MSQP (architecture of three steps), while the test without using a filter step can be described as the architecture of two steps.

Table 3. Joins executed to test the algorithm that computes the approximate area of polygon x polygon intersection

Labels	Dataset 1	Dataset 2
Join-1	Geological map	Township boundaries
Join-2	Geological map	Census block
Join-3	Township boundaries	Census block
Join-4	Brazilian municipalities	Brazilian municipalities'

In the filter step, in the case of 4CRS we use the algorithm proposed by Zimbrao and Souza (1998). For evaluating the overlapping of two 3CRS signatures, we employed the algorithm presented in Sub-section 3.3.

In the third step, we used a variant of plane sweep (Preparata, 1998) for detecting intersection of polygon's contours, and a test implemented by us for verifying if the polygon contains the other polygon.

4.3. Experimental Results

In order to generate the raster signatures, we have to choose the maximum number of grid cells (Zimbrao and Souza, 1998). Intuitively, the larger the number of cells, the closer is the approximation to the original polygon. However, processing raster signatures that have large sizes could produce high I/O and CPU costs. To evaluate the effects of different choices, we executed experimental tests using different maximum numbers of cells, such as: 250, 500, 1000 and 1500. We executed the following evaluations: storage requirements; number of pair identified in the second step which represent hits (pairs of objects that intersect each other) and false hits (pairs of objects that do not have intersection), and those are identified without executing the refinement step; CPU costs; and, I/O costs.

Despite of the 3CRS be based on the 4CRS, it presents some performance differences compared to 4CRS, in terms of store requirements, generation time, execution time and number of inconclusive hits.

Related to storage requirements, first, it is important to emphasize that signatures with 3 or 4 colors (3CRS or 4CRS) have same storage requirements, since it is required 2 bits to store the color of each cell. Therefore, both of them spend the same space to be stored. The experiments demonstrated that signatures with maximum number of cells equal to 250 have smaller storage requirements. Figure 5.a presents the size of the signature in comparison with the size of the original dataset. Signatures of maximum number of cells equal to 1500 have more storage requirements.

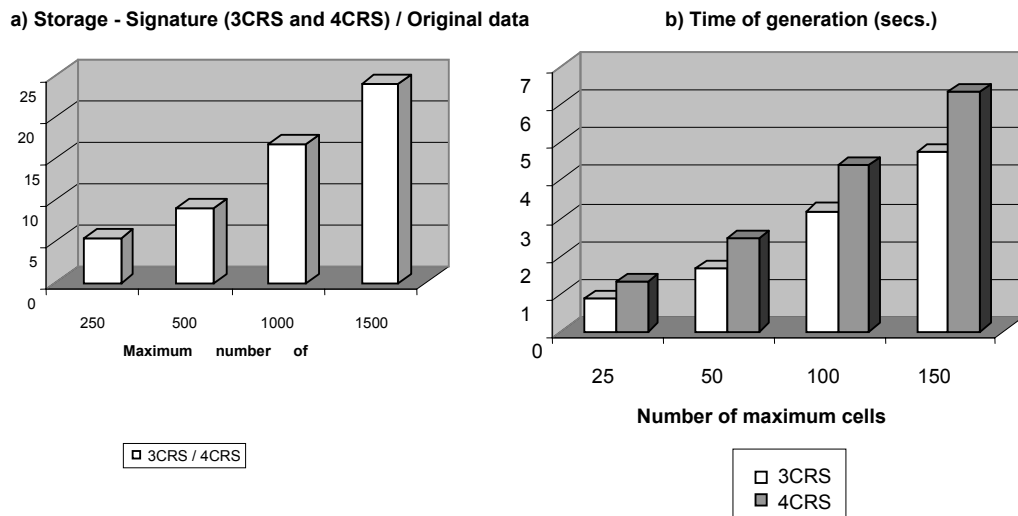


Figure 5. a) Size of the signatures (3CRS and 4CRS) related to the real data. b) Generation time for 3CRS and 4CRS

Related to generation time, 3CRS can be generated faster than 4CRS signatures, as presented in Figure 5.b. As it is expected, when the maximum number of cells increases, the time to generate the signatures also increases.

Table 4 presents the 4CRS and 3CRS signatures characteristics for the maximum number of cells equals to 500. To store 4CRS or 3CRS signatures of maximum number of cells equal to 500 it is only needed, on average, 2.85% of the space needed to store the real datasets. 3CRS can be generated in approximated 70% of the time needed for generating 4CRS, as presented in Table 4.

Table 4. Raster signatures' characteristics with maximum number of cells equal to 500

Datasets		Data size (KB)	4CRS size (KB)	4CRS / data (%)	Gener. time 3CRS (sec.)	Gener. time 4CRS (sec.)	3CRS / 4CRS (%)
Iowa	Census block group	38,824	1163	3.00	14.83	20.47	71.40
	Hydrologic map	6,904	169	2.45	2.70	3.97	68.02
	Topography map	61,748	1455	2.36	33.71	50.67	66.54
	Township boundaries	25,288	838	3.31	10.18	14.53	70.09
	Geologic maps	21,856	676	3.09	8.43	11.93	70.64
Average		29.050	827	2.85	13.97	20.37	69.34

Signatures of 250 maximum number of cells are processed faster, since there is less cells to process, while signatures of 1500 maximum cells are processed slower, since there are more cells to evaluate. Figure 6 shows the average time to execute the first two steps of the architecture, using 3CRS and 4CRS. In both cases, the bigger is the maximum number of cells, the bigger is the processing time.




Figure 6. Average time to execute the first two steps of the signature, using 3CRS and 4CRS

The number of inconclusive pairs that goes from the second step to the third step is bigger when 250 is chosen as the maximum number of cells, as demonstrated in Figure 7. In other words, signatures with maximum number of cells equal to 250 identify less hits and false hits. On the other hand, when the maximum number of cells is 1500, the number of inconclusive pairs is the smallest. The 3CRS and the 4CRS have

almost the same results. 3CRS generates only a few more inconclusive pairs to the third step.

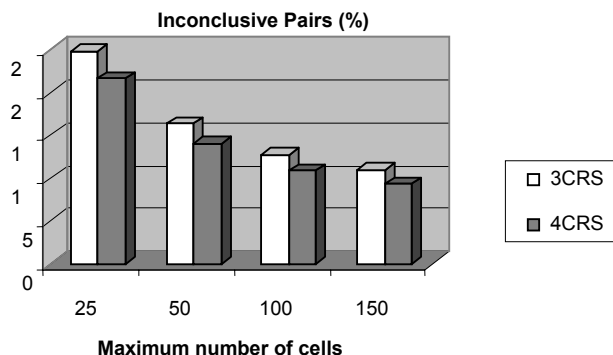


Figure 7. Inconclusive pairs (%) for the 3CRS and 4CRS signatures

Despite of the differences between 3CRS and 4CRS, both presents significant gain over the 2-step architecture, in terms of total execution time and disk access.

The total execution time is presented in Figure 8. Notice that both 3CRS and 4CRS use much less time than the 2-step architecture (more than 50% of reduction). The time used in 3CRS is slightly bigger than 4CRS, because the 3CRS test generates more inconclusive pairs, which leads to more exact and slower tests.




Figure 8. Comparison of the total execution time (secs.) between 3CRS, 4CRS and 2-step architecture

Figure 9 shows the relation between the disk accesses using 3-step architecture (3CRS and 4CRS) and 2-step architecture. It shows the ration in percentage: number of disk accesses used by the 3-step architecture divided by number of disk accesses used by the 2-step architecture. Notice that the 4CRS uses slightly less disk accesses than the 3CRS. Again, it is because the 3CRS generates more inconclusive pairs, which leads to more exact test (the refinement step). Therefore, the algorithm needs to access the disk more often, to seek and read objects from disk and execute the exact test. However, it is

important to emphasize that the difference between the number of disk accesses using 3CRS and 4CRS is too small.

Maximum number of cells	3CRS	4CRS
250	52.11	48.50
500	51.00	48.18
1000	59,64224	57.41
1500	71,07493	69.13




Figure 9. Comparison of the number of disk accesses (average) of the 3-step join over the 2-step join

5. Conclusion

Spatial operations are very costly. The literature presents different approaches for processing spatial operations. This work proposed a new raster signature, the 3CRS, based on the 4CRS (Zimbrao and Souza, 1998), that can be used as a second filter step in the Multi-Step Query Processor of Brinkhoff *et al.* (1994). 3CRS has the good performance of the 4CRS when compared against the 2-step processing. Moreover, 3CRS has faster generation time than 4CRS and it is also more flexible, since it can be used to represent different spatial data types, such as polygons, polylines and points. Due to its faster generation time, 3CRS can also be calculated on-the-fly, for example when the optimizer decides to use it. Another advantage is that the 3CRS can be computed based on others 3CRS signatures, as a result of a previous operation, like a join.

The experiments that we executed to evaluate the new signature demonstrated the effectiveness of 3CRS. It only needs, on average, 2.85% of the storage of the real datasets to store 3CRS signature. It has a significant reduction of the generation time (30% of reduction, in average) related to 4CRS. The time and number of disk accesses to process the queries where much smaller then the time to execute the queries without signature. Related to 4CRS, the processing using 3CRS produces a small growth of the inconclusive answer. In other words, the increasing of total processing time/number of disk accesses is too small, and motivates the use of 3CRS in processing spatial joins.

As future works, we intend to develop a different storage mechanism to reduce the storage requirements. We also plan to evaluate the use of 3CRS to representing polylines and points. Another future work is to implement this signature in SECONDO (Güting *et al.*, 2005).

6. References

- Azevedo, L. G., Monteiro, R. S., Zimbrão, G., Souza, J. M.: Approximate Spatial Query Processing Using Raster Signatures. In Proc. of VI Brazilian Symposium on GeoInformatics, Campos do Jordao, Brazil (2004).
- Aronoff, S. (1989), “Geographic Information Systems”, 1 ed. Ottawa, Canada, WDL Publications.
- Beckmann, N., Kriegel, H. P., Schneider, R., et al., 1990, “The R*-tree: An Efficient and Robust Access Method for Points and Rectangles”. In Proceedings of the 1990 ACM SIGMOD International Conference on Management of Data, pp. 322-331, Atlantic City, NJ, USA, May.
- Boissonnat, J. D., Preparata, F. P. (1997) “Robust Plane Sweep for Intersecting Segments”, SIAM Journal on Computing
- Brazilian Institute of Geography and Statistics Fundação Instituto Brasileiro de Geografia e Estatística – IBGE (1996): “Malha Municipal Digital do Brasil - 1994”, Rio de Janeiro.
- Brinkhoff, T., Kriegel, H. P., Schneider, R. (1993a), “Comparison of Approximations of Complex Objects Used for Approximation-based Query Processing in Spatial Database Systems”. In Procs. of Int. Conf. on Data Engineering, Vienna, Austria, 40-49.
- Brinkhoff, T., Kriegel, H. P., Seeger, B. (1993b) “Efficient processing of Spatial Joins Using R-Trees”. In Proceedings of the 1993 ACM-SIGMOD Conference, Washington, DC, USA.
- Brinkhoff, T., Kriegel, H. P., Schneider, R., Seeger, B. (1994), “Multi-step Processing of Spatial Joins”. In Proc. of ACM-SIGMOD Int. Conference on Management of Data, Minneapolis, MN 197-208.
- Freiseisen W., Pau P. (1998), “A generic plane-sweep for intersecting line segments,”RISC-Linz Report, pp. Series No. 98–18
- Gordon, S. R., et al. (1994), “Final Report on Status of Spatial/Map Databases”. Technical Report of Oak Ridge National Laboratory, June 1994.
- Güting, R.H., Almeida, V., Ansorge, D., Behr, T., Ding, Z., Höse, T., Hoffmann, F., Spiekermann, M. (2005), “Secondo: An Extensible DBMS Platform for Research Prototyping and Teaching”, 21st Intl. Conf. on Data Engineering, ICDE, Tokyo, Japan

- Kothuri, R. K., Ravada, S. (2001) "Efficient Processing of Large Spatial Queries Using Interior Approximations". Proc. of the Int. Symposium on Spatial and Temporal Databases, Los Angeles, CA 404-424.
- Newman, W. M., Sproull, R. F. (1979), "Principles of Interactive Computer Graphics", 2 ed. New York, McGraw-Hill Book Company.
- Orenstein, J. A. (1986), "Spatial query processing in an object-oriented database system". In Proc. of ACM SIGMOD Int. Conf. on Management of Data, Washington, DC 326-336.
- Preparata, F. P., Shamos M. I.: "Computational Geometry", Springer, 1998.
- Samet, H. (1990), "The Design and Analysis of Spatial Data Structure", 1 ed., Boston, Massachusetts, Addison-Wesley Publishing Company.
- Tao, Y., Sun, J., Papadias, D. (2003), "Selectivity estimation for predictive spatio-temporal queries". In: Proceedings of the 19th International Conference on Data Engineering, pp. 417-428, Bangalore, India.
- Zimbrão, G., Souza, J. M. (1998). "A Raster Approximation For Processing of Spatial Joins". In: Proceedings of VLDB'98 - 24th International Conference on Very Large Databases, New York City, NY, USA, Aug 1998.
- Zhu, H., Su, J., Ibarra, O. H. (2000), "Toward Spatial Joins for Polygons". In Proc. Int. Conf. of SSDBM, Berlin, Germany 2431-246.