

Efficient Join-Index-Based Join Processing: A Clustering Approach

Shashi Shekhar, Chang-tien Lu, Sanjay Chawla
Computer Science Department, University of Minnesota
200 Union Street SE, Minneapolis, MN-55455
[shekhar,ctl, chawla]@cs.umn.edu TEL:(612) 6248307 FAX:(612)6250572
<http://www.cs.umn.edu/Research/shashi-group>

Sivakumar Ravada
Spatial Data Product Division
One, Oracle Drive, Nashua, NH 03062
sravada@us.oracle.com

November 12, 1999

Abstract

A Join Index is a data structure used for processing join queries in databases. Join indices use pre-computation techniques to speed up online query processing and are useful for data-sets which are updated infrequently. The cost of join computation using a join-index with limited buffer space depends primarily on the page-access sequence used to fetch the pages of the base relations. Given the join-index, we introduce a suite of methods based on clustering to compute the joins. We derive upper-bounds on the lengths of the page-access sequences. Experimental results with Sequoia 2000 data sets show that the clustering method outperforms the existing methods based on sorting and online-clustering heuristics.

Acronym	Full form	Definition section/page
OPAS-FB	Optimal Page-Access Sequence with Fixed Buffer	Section 1
PCG	Page-Connectivity Graph	Section 1
AC	Asymmetric Clustering-based heuristic	Section 2
Sorting	Sorting heuristic	Section 2
SC	Symmetric Clustering based heuristic	Section 3
FP	Fotouhi and Pramanik's heuristic	Section 3
OM	Omiecinski's heuristic	Section 3
Chan	Chan's heuristic	Section 3
B-diagonal entry $M[i, j]$	$ i - j \leq \lfloor \frac{B}{2} \rfloor$	Section 3

Table 1: Table of Acronyms

Keywords: Optimal Page Access Sequence, Join Index, Join Processing, Spatial Join.

1 Introduction

The join operation is a fundamental operation in relational databases, and substantial work has been done in optimizing join operations [11, 27]. A join index [30, 34] is a special data structure that facilitates rapid join-query processing. For data sets which are updated infrequently, the join index can be particularly useful.

The join-index is typically represented as a bi-partite graph between the pages of incumbent relations or their surrogates. When the number of buffer pages is fixed, the join-computation problem is transformed into determining a page-access sequence such that the join can be computed with the minimum number of redundant page accesses. This problem has been shown to be NP-hard [26, 29], and consequently, it is unlikely that a polynomial time solution exists for this problem. Solutions in the literature use a clustering method that groups pages in one or both tables involved in the join to reduce total page accesses. Available heuristics either group the pages of a single table via sorting [34] or use incremental clustering methods [5, 7, 28].

Our Contribution: We introduce two new heuristics for this problem. One heuristic uses the clustering method to group the pages in one table, generalizing the sorting-based heuristic for joins. The other heuristic uses clustering for the pages of both tables. The former generalizes and outperforms the sorting heuristic, while the latter generalizes and outperforms the incremental clustering methods for joins. We provide a formal characterization of an upper bound on the number of redundant I/Os performed by our approaches. Experiments with the Sequoia 2000 [33] data-set show that both heuristics outperform other methods when the memory size is relatively small. The proposed approaches are useful for computing joins, given join-indices for large database, where the size of memory is small compared to the sizes of the individual datasets.

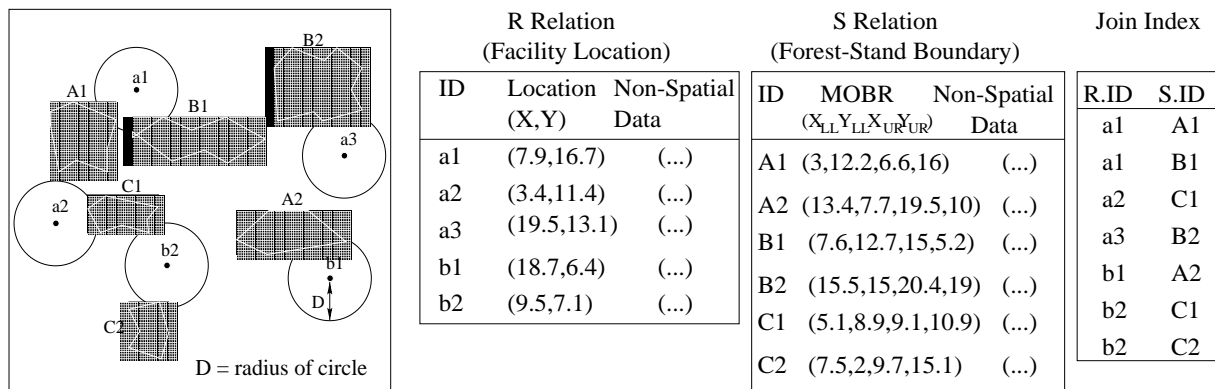
1.1 Join Index: Basic Concept

Consider a database with two relations, Facility and Forest Stand. Facility has a point attribute representing its location, and Forest Stand has rectangle attributes that represent its extent by a bounding box. The polygon representing its extent may be stored separately. A point has the x and y coordinates on the map. A rectangle is represented by points that represent the bottom left and top right corners.

In Figure 1(a), points a_1, a_2, a_3, b_1, b_2 represent facility locations, and polygons $A_1, A_2, B_1, B_2, C_1, C_2$ are the bounding boxes that represent the limits of the forest stands. The circle around each location shows the area within distance D from a facility. The rectangle around each forest boundary represents the Minimal Orthogonal Bounding Rectangle(MOBR) for each forest stand. Figure 1(b) shows two relations, R and S , for this data set. Relation R represents facilities via the attributes of a unique identifier, $R.ID$, the location (x,y coordinates), and other non-spatial attributes. Relation S represents the forest stands via a unique identifier, $S.ID$, the MOBR and non-spatial attributes. $MOBR(X_{LL}, Y_{LL}, X_{UR}, Y_{UR})$, is represented via the coordinates of the lower-left corner point (X_{LL}, Y_{LL}) and the upper right corner point (X_{UR}, Y_{UR}) . Now, consider the following query: **Q**: "Find all forest stands which are within a distance D from each facility". This query will require a join on the Facility and Forest Stand relations, based on their spatial attributes.

A spatial join algorithm [2, 3, 4, 12, 25] may be used to find the pairs (Facility, Forest-stand) which satisfy query **Q**. Alternatively, a join-index may be used to materialize a subset of the result to speed

up processing for the future occurrence of Q if there are few updates to the spatial data. Figure 1(b) shows a join index with two columns. Each tuple in the join index represents a tuple in the table $JOIN(R, S, distance(R.Location, S.MOBR) < D)$. In general, the tuples in the join index may also contain pointers to the pages of R and S where the relevant tuples of R and S reside. We omit the pointer information to simplify the diagrams in this paper.



(a) Spatial Attribute of R and S

(b) R and S Relation Table and Join Index

Figure 1: Constructing a Join Index from two relations

A join index describes a relationship between the objects of two relations. Assume that each tuple of a relation has a surrogate (a system-defined identifier for tuples, pages, etc.) which uniquely identifies that tuple. A join index is a sequence of pairs of surrogates, where each pair of surrogates identifies the result-tuple of a join. The tuples participating in the join result are given by their surrogates. Formally, let R and S be two relations. Then consider the join of R and S on attributes A of R and B of S . Then the join index is an abstraction of the join of the relations. If F defines the join predicate, then the join index is given by the set $JI = \{(r_i, s_j) | F(r_i.A, s_j.B) \text{ is true for } r_i \in R \text{ and } s_j \in S\}$, where r_i and s_j are surrogates of the i th tuple in R and the j th tuple in S , respectively. For example, consider the Facility and Forest Stand relational tables shown in Figure 1. The Facility relation is joined with the Forest Stand relation on the spatial attributes of each relation. The join-index for this join contains the tuple IDs which match the spatial join predicate.

A join index can be described by a bipartite graph $G = (V_1, V_2, E)$, where V_1 contains the tuple IDs of relation R , and V_2 contains the tuple IDs of relation S . Edge set E contains an edge (v_r, v_s) for $v_r \in R$ and $v_s \in S$, if there is a tuple corresponding to (v_r, v_s) in the join index. The bipartite graph models all of the related tuples as connected vertices in the graph. In a graph, the edges connected to a node are called the incident edges of that node, and the number of edges incident on a node is called the degree of that node.

1.2 Page-Connectivity Graph, Page-Access Sequence

When the join relationship between two relations is described at the page level, we get a page-connectivity graph. A *Page-Connectivity Graph* (PCG) [26] $B_G = (V_1, V_2, E)$ is a bipartite graph where vertex set

V_1 represents the pages from the first relation, and vertex set V_2 represents the pages from the second relation. The set of edges is constructed as follows: an edge is added between page (node) v_1^i in V_1 and page (node) v_2^j in V_2 , iff there is at least one pair of objects (r_i, s_j) in the join index such that $r_i \in v_1^i$ and $s_j \in v_2^j$. Figure 2 shows a page-connectivity graph for the join index from Figure 1(b). Nodes (a, b) represent the pages of relation R , and nodes (A, B, C) represent the pages of relation S . A *min-cut* node partition [13, 22] of graph $B_G = (V_1, V_2, E)$ partitions the nodes in V into disjoint subsets while minimizing the number of edges whose incident nodes are in two different partitions. The cut-set of a min-cut partition is the set of edges whose incident nodes are in two different partitions. Fast and efficient heuristic algorithms [20, 17] for this problem have become available in recent years. They can be used to cluster pages in PCG.

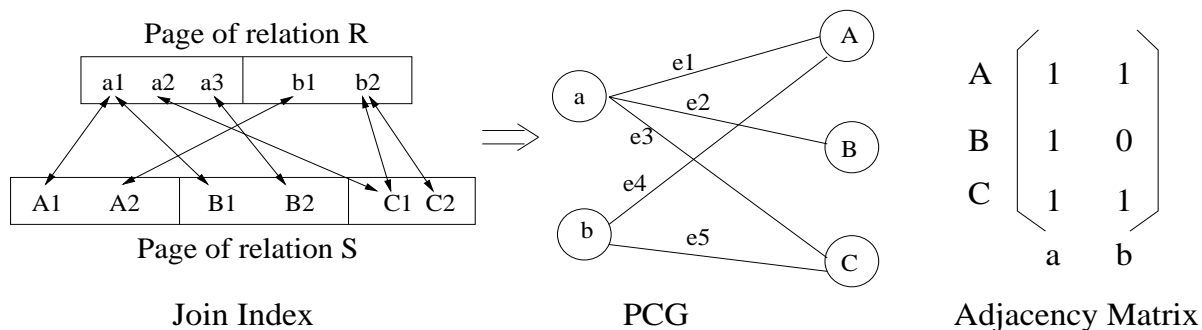


Figure 2: Construction of a Page-Connectivity Graph(PCG) from a Join Index.

A join index helps speed up the join processing, as it keeps track of all of the pairs of tuples which satisfy the join predicate. Given a join index JI , one can use the derived PCG to schedule an efficient page-access sequence to fetch the data pages. The CPU cost is fixed, as there is a fixed cost associated with joining each pair of tuples, and the number of tuples to be joined is fixed. I/O cost, on the other hand, depends on the sequence of pages accessed. When there is limited buffer space in the memory, some of the pages may have to be read multiple times from the disk. The *page-access sequence* (and in turn the join-index clustering and the clustering of the base relation) determines the I/O cost.

Example: We illustrate the dependency between the I/O cost of a join and the order in which the data pages are accessed with the help of an example, using the page-connectivity graph shown in Figure 2. Assume that the buffer space is limited to allow at most two pages of the relations in memory, after caching the whole page-connectivity graph in memory. Consider the two-page access sequences: (i) (a, A, b, B, a, C, b) and (ii) (a, A, b, C, a, B) . Each sequence allows the computation of join results using a limited buffer of two pages. However, in the first case, there are a total of seven page accesses, and in the second case there are a total of six page accesses. Note that the lower bound on the number of page accesses is five, as there are five distinct pages in the PCG. However, with two buffer spaces, there is no page-access sequence which will result in five page accesses. This is because the cycle (a, A, b, C, a) requires that at least three pages be in memory to avoid redundant page accesses. With three buffer spaces, (a, B, A, C, b) is a page-access sequence which results in five page accesses.

1.3 Problem Definition, Scope, Outline

Given that the I/O cost depends on the page-access sequence, the following optimization problem characterizes the problem of designing efficient algorithms for processing joins, given a join index and a fixed buffer size. This problem is called the Optimal Page-Access Sequence with a Fixed Buffer (OPAS-FB) problem [26], which is formally defined as follows:

OPAS-FB Problem

Given: A page-connectivity graph $PCG = (V, E)$, representing the join index, and a buffer of size $B \leq |V|$.

Find: A page-access sequence.

Objective: To minimize the number of page accesses.

Constraint: Such that the number of pages in the buffer is never more than B .

For example, the optimal page-access sequence for the PCG in Figure 2 for $B = 2$ is (a, A, b, C, a, B), which results in six page accesses.

The OPAS-FB problem is known to be NP-hard [26, 29], and heuristic solutions have been proposed in the literature for solving this problem. The heuristics in the literature can be broadly divided into two groups, namely asymmetric single-table clustering and symmetric two-table on-line clustering. We describe the relevant literature and our contribution in sections 2 and 3.

Scope: In this paper, we focus on the OPAS-FB problem. We do not address the update problems associated with managing join indices. Also, the base-relation clustering and tuple-level join-index optimization are out of the scope of this paper.

Outline: The rest of the paper is organized as follows. Section 2 discusses asymmetric single-table clustering methods and proposes our first approach, Asymmetric Clustering (**AC**), evaluating its performance with the Sorting heuristic. Section 3 describes two-table on-line clustering from the literature. In Section 4, we propose Symmetric Clustering (**SC**). In Section 5, we compare the proposed methods, AC and SC, with traditional algorithms for the OPAS-FB problem. Section 6 includes a summary and future research directions.

2 Asymmetric Methods

The main approach within **asymmetric single-table clustering** is based on sorting the join-index on one of the join keys. In the following discussion, let R and S be the two relations, with JI being the join index. The *sorting-based* asymmetric heuristic presented in [34] reads as much as possible of the join index (JI) and one relevant relation (R semi-join JI) into memory. Here JI is assumed to be sorted on relation R.ID. To reduce redundant accesses to S, access to S is clustered by sorting the list of all of the surrogates from S that are related to the subset of the join-index in memory. This heuristic ensures that no redundant accesses are performed on relation R, but it may incur redundant accesses to the second relation. The Sorting-based heuristic is most suited to the applications that have totally ordered join-keys. Rigorously speaking, the sorting-based heuristic sorts the surrogates (e.g. system-defined identifiers for pages) rather than the join-key attributes. If tables are sorted by the respective

join-keys, then surrogates for the pages in a table may be ordered by the lowest key-value for any tuple in the page. This reduces redundant page I/O in computing joins using a join-index, for join-keys with totally ordered domains. Since multi-dimensional domains, e.g. spatial data types, do not have natural total-order, sorting surrogates may not be as effective for computing spatial-joins using join-indices. We propose Asymmetric Clustering to address this problem. Asymmetric clustering uses the entries in the join-index for grouping pages of one relation, say R , based on their interaction with the pages in the other relation, say S . If the join-index (e.g. Figure 1(b)) represents the summary of a spatial join, then the pages of R are clustered using their spatial relationship with the pages of S , and the proposed method is called Asymmetric Spatial Clustering.

2.1 Basic Idea Behind Asymmetric Clustering(AC)

The example in Figure 3 highlights the different approaches of AC and the sorting heuristic. Figure 3(a) shows a 49-node PCG, with numbers 1-24 and letters A-Y corresponding to the surrogates of the two page-level relations. The figure can represent a spatial-join computation between two layers of geographic data consisting of small polygons. The pages from each relation may overlap pages from the other. Consider a memory with seven pages for buffering the pages of R and S relations. There may be additional buffering pages for managing the result, index, etc. The AC heuristic used a different clustering




Figure 3: Example of the AC and Sorting heuristic

of pages of the first relation, relative to the sorting-based approach. Figure 3(b) and 3(c) show the clusters of pages of the first relation that are used by AC and the sorting-based method. The pages of R are numbered 1-24, and within a cluster are annotated by a common symbol. For example, nodes 1,2,3,4,5,6 in Figure 3(b) are all circled, denoting that these are loaded together by sorting. Similarly, nodes 7,8,9,10,11,12 are annotated with a hexagon, and so on. Visually, one can verify that the clusters used by AC are spatially cohesive. The number of pages of the second relation which have edges to multiple clusters of pages in the first relation yield redundant I/O. One may consider using other space-filling curves(e.g. Z-order, Hilbert) to improve the performance of the sorting-based method. However, min-cut graph partitioning, the method used by AC, outperforms space-filling curves in the clustering of non-uniformly distributed spatial data, as shown in our previous work [31, 32].

The sorting heuristic clusters the nodes of the first page-level relation and then loads the pages in the sorted order. The loading sequence for the example in Figure 3 is shown in Table 2. With a

buffer size of seven, the sorting heuristic will load pages $\{1, 2, 3, 4, 5, 6\}$ from the first relation in the first six buffers and then will load one page at a time from the set of pages, $\{A, B, C, D, E, F, G, H, I\}$, in the 7th buffer. For the next round, it loads $\{7, 8, 9, 10, 11, 12\}$ in the first six buffers, and then $\{E, F, G, H, I, J, K, L, M, N, O, P\}$ one at a time in the 7th buffer, and so on, as shown in Table 2. The sorting heuristic results in 17 redundant I/Os, with a total of 66 I/Os.

	Sorting Heuristic		AC	
Round	First Six Buffers	The 7th buffer	First Six Buffers	The 7th buffer
1	1,2,3,4,5,6	A,B,C,D,E,F,G,H,I	1,2,4,5,9,10	A,B,C,D,F,G,H,L,M,N
2	7,8,9,10,11,12	E,F,G,H,I,J,K,L,M,N,O,P	3,7,8,13,14,19	B,E,F,J,K,L,Q,R,V
3	13,14,15,16,17,18	J,K,L,M,N,O,P,Q,R,S,T,U	6,11,12,17,18,22	D,H,I,N,O,P,T,U,X
4	19,20,21,22,23,24	Q,R,S,T,U,V,W,X,Y	15,16,20,21,23,24	L,M,N,R,S,T,V,W,X,Y
Total I/O	66		62	

Table 2: Loading Sequence of Sorting and APG for Figure 4

The AC clusters the nodes of the first page-level relation according to their connections with the second relation. The clustering, as shown in Figure 3(c), with a buffer size of seven, provides four clusters. Note that these clusters are different from the page-clusters used in sorting. The AC heuristic will load pages $\{1, 2, 4, 5, 9, 10\}$ of the first relation in the first six buffers, then, one by one, will load $\{A, B, C, D, F, G, H, L, M, N\}$ in the 7th buffer. In the next round, it loads $\{3, 7, 8, 13, 14, 19\}$ together into the first six buffers, and then $\{B, E, F, J, K, L, Q, R, V\}$ one at a time into the 7th buffer. Table 2 shows the total loading sequence. The AC results in 13 redundant I/Os, with a total of 62 I/Os. The difference of four I/Os out of 66 in this example may not look large. However, the relative difference in I/Os using the sorting and clustering methods will increase with an increase in data-set size. This linear characteristic of sorting yields poor clustering and limits the savings in redundant I/Os.

2.2 Description of Asymmetric Clustering Method

The goal of asymmetric clustering methods is to cluster the pages of one relation, given the join-index or its PCG. This can be formalized as a min-cut hypergraph-partitioning problem. The pages of a relation will form the nodes of the hypergraph. Each page p of the other relation will form a hyperedge, covering all pages of the first relation connected to p in PCG. Partitioning the nodes in this hypergraph will form a group of pages of the first relation that can be loaded together. Minimizing cut hyperedges during partitioning reduces the number of pages of the second relation that will need to be loaded into memory multiple times.

Consider the example spatial-join problem depicted in Figure 4(a) with two point data-sets, (a,b,c,d) and (A,B,C,D) . Assume a blocking factor of 1 to simplify the example. The PCG of the join-index for $Distance(i, j) < \frac{L}{\sqrt{2}}$ is shown in Figure 4(c) using the overlay and distance buffer information. The nodes of the hypergraph shown in Figure 4(d) consist of the nodes of relation R, i.e. (a,b,c,d) . The hyperedges represent nodes (A,B,C,D) of S. The hyperedge corresponding to A connects a and c , since (A,a) and (A,c) satisfy the join predicate. The partition $((a,c),(b,d))$ has no cut hyperedges, and computing the join using it will result in no redundant I/O with loading sequence (a,c,A,C,b,d,B,D) ,

if 3 buffers are available to hold the pages of two relations. In contrast, the partition $((a,b),(c,d))$ cuts all four hyperedges, and computing this join will yield four redundant I/Os with loading sequence $(a,b,A,C,B,D,c,d,A,C,B,D)$, if only 3 buffers are available to hold the pages of two relations.

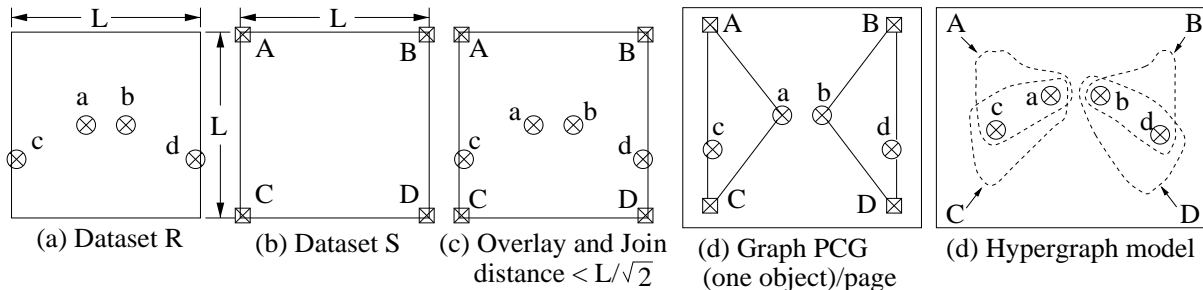


Figure 4: Construction of a one-sided hypergraph from the data set

We formally describe AC now, via the following pseudo-code.

AC Algorithm

Input: $G = (V_r, V_s, E)$ is a page connectivity graph

Output: $S = \langle P_1, P_2, \dots, P_r \rangle$ is a page access sequence with $r \geq |V_r| + |V_s|$. (P_i s need not be distinct)

```

assert(|V_r| < |V_s|);
assert(B ≥ 2); /* number of buffers */
HG_r(V_r, HE_r) = DeriveHypergraph(G); /* HG_r is a hypergraph, |HE_r| = |V_s| */
/* For each node in |V_s|, build a hyperedge to encompass all of its corresponding nodes in V_r */
PSet_r = hMetis-Partition(HG_r, B - 1) /* PSet_r is the set of partitions */
i=0;
while ((P_{i_r} = SelectUnprocessedPartition(PSet_r)) != NULL) /* Select the un-processed partition */
{
  AddPageSequence(S, P_{i_r}); /* Add all the nodes in P_{i_r} into the loading sequence */
  P_{i_s} = Sort-Eliminate-Dup(G, P_{i_r});
  /* Sort and eliminate the duplicated nodes in V_s of G which connect to nodes in P_{i_r} */
  AddPageSequence(S, P_{i_s}); /* Add all the nodes in P_{i_s} into the loading sequence */
  P_{i_r}.flag = "processed"; /* Mark this partition as "processed" */
  i++;
}

```

The first step of the AC algorithm, i.e. $\text{DeriveHypergraph}(G)$, creates a hypergraph from a given page connectivity graph G . Nodes of the first relation R form the nodes of the hypergraph. For each node v of the second relation, it builds a hyperedge to encompass a set of nodes on the first relation(R) that are connected to v in G . Next, AC partitions this hyper-graph using the min-cut hyper-graph partitioning algorithm, hMetis [17, 18, 19], which is a multi-level hypergraph-partitioning algorithm that has been shown to produce high quality bi-sections on a wide range of problems that arise in scientific and VLSI

applications. hMetis minimizes the (weighted) hyper-cut and thus tends to create partitions in which connectivity among the vertices in each partition is high, resulting in good clusters. Finally, AC loads each partition in the primary relation and its connected nodes in the second relation, one by one, to compute the join. The I/O cost of AC can be characterized via the following lemma:

Lemma 1 *Given a partition $\{V_{r_1}, V_{r_2}, \dots, V_{r_p}\}$ of V_r , i.e. pages of relation R , from the page-connectivity graph $PCG = (V_r, V_s, E)$, there is a page-access sequence of length $K = |V_r| + \sum_{v \in V_s} f(v)$ to process the spatial join, where $f(v)$ denotes the number of partitions of V_r that have an edge to node v in V_s .*

Proof: A node v in V_s is connected to $f(v)$ partitions of V_r . Therefore, the node v in V_s has to be loaded $f(v)$ times into the buffer to compute the spatial join. The total number of redundant I/Os is $\sum_{v \in V_s} (f(v) - 1)$. The total I/O cost = $|V_r| + |V_s| + \sum_{v \in V_s} (f(v) - 1) = |V_r| + |V_s| + \sum_{v \in V_s} f(v) - |V_s| = |V_r| + \sum_{v \in V_s} f(v)$ ■

We note that the min-cut hypergraph-partitioning algorithm, e.g. hMetis, minimizes the number of hyperedges connecting nodes across clusters. This does not directly minimize $\sum_{v \in V_s} f(v)$, as it does not distinguish between a hyperedge spanning four clusters and another spanning two clusters. While AC outperforms the sorting-based heuristic already, the performance of AC will improve when better algorithms for hypergraph partitioning are available which minimize the total number of cuts on cut-hyperedges. We plan to explore this in future work.

2.3 Comparison of Asymmetric Approaches: Sorting vs. AC

2.3.1 Experiment Design

We now compare the performance of the Sorting heuristic and AC, using a join index derived from the Sequoia 2000 [33] dataset. The *Point* table contains 62,584 California place names with their associated locations (Longitude and Latitude), extracted from the US Geological Survey’s Geographic Names Information System (GNIS). The *Polygon* table contains 4,388 records, representing Cropland and Pasture land use in California. Throughout Section 2.3 and 5, the *Point* and *Polygon* tables will be referenced as R and S , respectively. We plot the point and polygon data set of California records as in Figure 5.

Now, consider the following queries:

Q.A. "For each place in the *Point* table, find \mathbf{N} nearest croplands from the *Polygon* table".

Q.B. "For each place in the *Point* table, find all croplands which are within a distance \mathbf{D} ".

The spatial join of these two queries produces sets of join indices and such join indices are of interest in spatial data mining for neighborhood indexing [6]. The value of \mathbf{N} and \mathbf{D} can be increased/decreased for adjusting the edge ratio [5]. Give a join graph $G = (V_R, V_S, E)$, the edge ratio of G , denoted by Θ , is defined as the ratio of the total number of edges in G to the maximum possible number of edges in G if it is a fully connected graph; i.e., $\Theta = \frac{|E|}{|V_R||V_S|}$. The edge ratio provides a measure of the page-level join selectivity.

The variable parameters are the Buffer size, Page size and Edge ratio. The metric of evaluation is the number of page accesses required by each algorithm to implement the join. The edge ratio is controlled

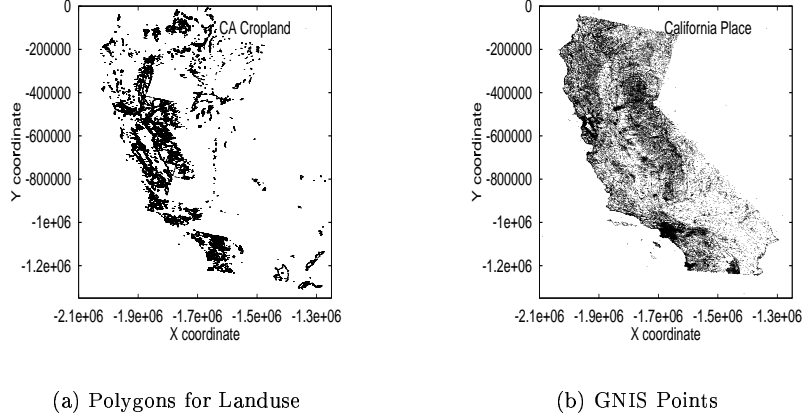


Figure 5: Example of the Sequoia 2000 data set

by **N** and **D**. The number of nearest neighbors, **N**, is varied from 1 to 5, yielding an edge-ratio of 0.002 to 0.005. The side-size of range queries, **D**, is varied from 400 to 4800 units where the extent of California is almost $1.2 * 10^6$ units (North-South) x $0.8 * 10^6$ units (East-West). This yields an edge-ratio of 0.002 to 0.003.

Page size represents the size of disk blocks and memory pages. Different values of page sizes include 2, 4, 8, 16, 32, 64 Kbytes. The size of the records in the point table is 64 bytes. The blocking factor for the Point table is the ratio of page size and record size. Point records are spatially clustered in the pages of the point table. The records in the Polygon table are of variable size. The size of a record in the Polygon table is $16 + 32 * (\# \text{ of Points in the Polygon})$ bytes. The number of points in a polygon can vary from a dozen to a few thousands, and a large polygon may span multiple pages.

The buffer size represents the ratio of available memory size as a fraction of the size of the Point table, which is the smaller of the two tables. Memory buffer size varies from 4% to 20% of the size of the smaller table.

Figure 6 shows the various process steps of the experiment design. From the base point and polygon tables, we derived families of join indices for queries **Q.A.** and **Q.B.** for different values of edge ratio. Next, a set of page-connectivity graphs(PCGs) are generated for each join index, given different values of page size. The page-connectivity graphs are input to the "Page-Access-Sequence Generator," which simulates the behavior of OPAS-FB algorithms (i.e. Sorting and AC) for a given buffer size. The page-access sequence and total page I/Os are tracked for each combination of join algorithm, page size, buffer size, and edge ratio to derive the experiment results.

2.3.2 Experiment results

Figures 7 and 8 show the comparison between the AC and Sorting heuristic for range-query join-indices and N-nearest-neighbor join-indices, respectively. The AC method is uniformly better than the Sorting heuristic.

Figures 7(a) and 8(a) show the impact of page size, which we vary from 2 kbytes to 64 kbytes. The page size and the number of page accesses are shown in logarithm scale(base two). As the page size

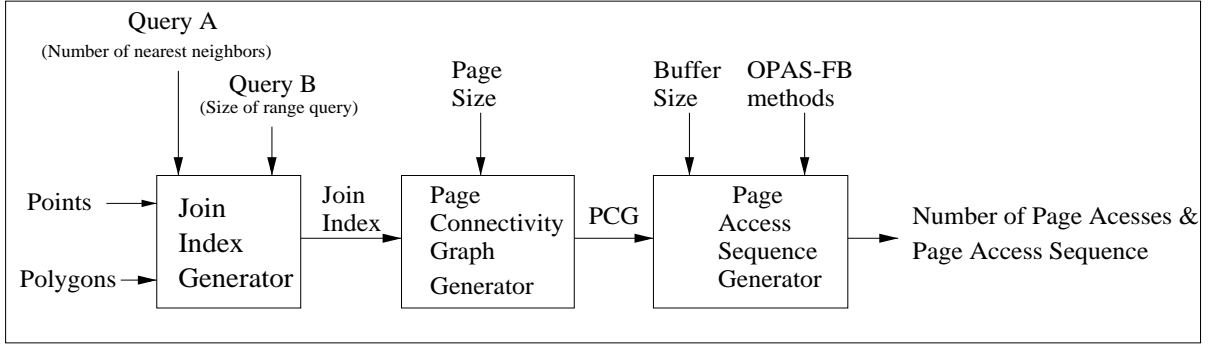


Figure 6: Experimental setup and design.

increases, the number of pages decreases, and clustering efficiency improves for all methods, reducing the performance gap between the two methods.

Figures 7(b) and 8(b) show the effect of buffer size (as a fraction of the size of the smaller relation) on the I/O performance of AC and the sorting-based method. As long as the buffer is smaller than the smaller of the two relations involved in the join, both AC and the sorting-based approach use most of the buffers to load the pages of only one relation. The difference in performance comes from their difference in clustering ability: AC has a lower I/O cost than Sorting.

Figures 7(c) and 8(c) show the effect of the edge ratio. AC uniformly outperforms the Sorting-based approach. The gap between the performance of the two methods does not show any trend.

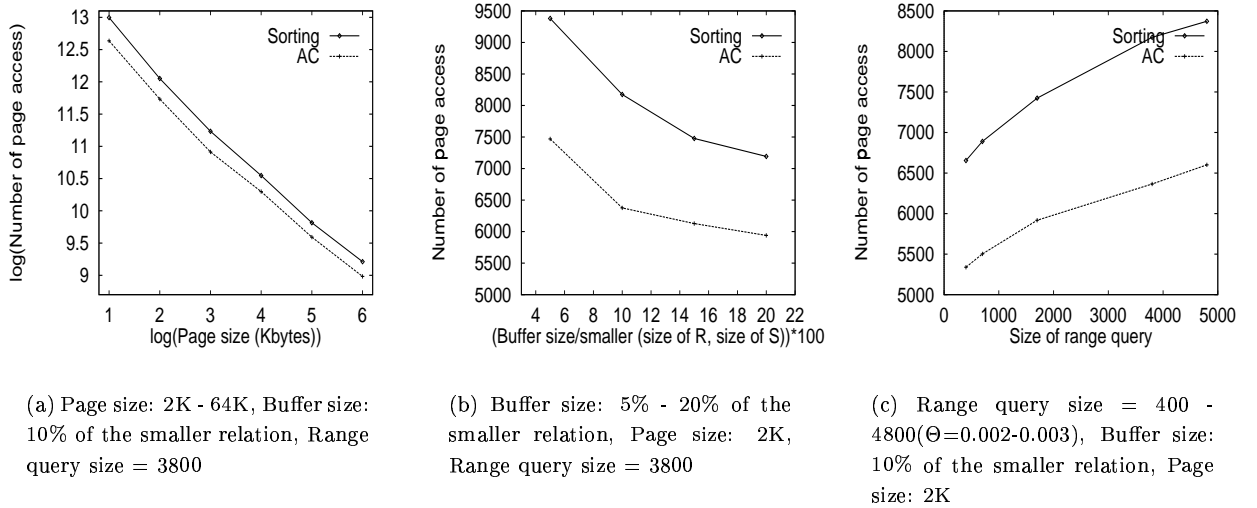


Figure 7: **Range Query Join-index:** The effect of page size, buffer size, and number of nearest neighbors on the AC and Sorting heuristic

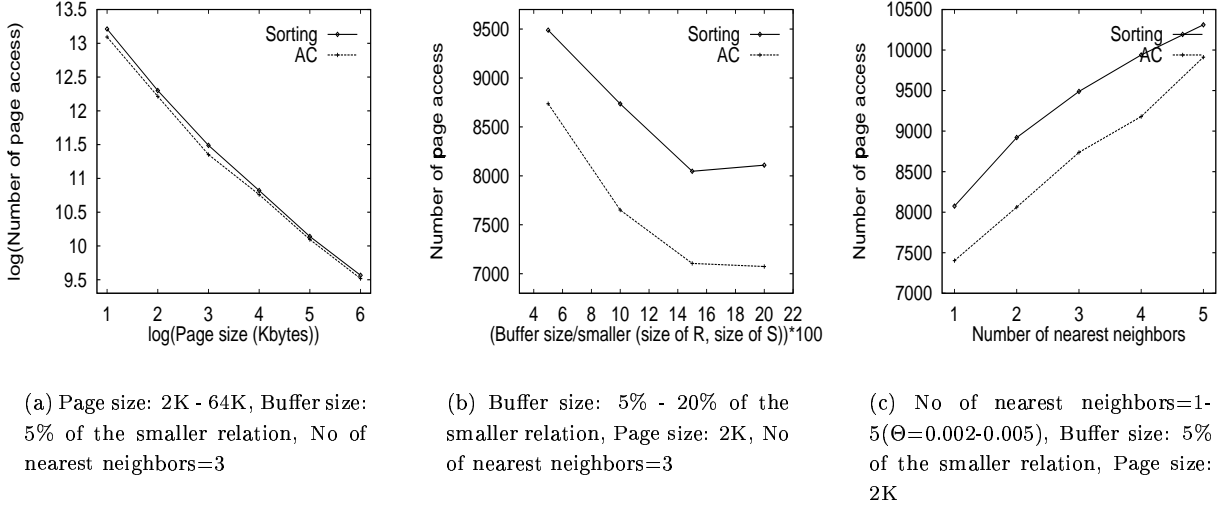


Figure 8: **Nearest Neighbor Join-index**: The effect of page size, buffer size, and number of nearest neighbors on the AC and Sorting heuristic

3 Symmetric Methods

While AC is an improvement over the Sorting-based method for spatial joins, it has a few drawbacks. For example, its buffer utilization can be poor, since it gives almost the entire buffer space to one relation. We illustrate this with the help of the spatial join problem shown in Figure 9. Figure 9(a) shows a polygon set with 6 polygons, $R_0..R_5$, and a point data set with 6 points. The adjacency matrix M_{PCG} representation of a join-index is shown in Figure 9(b), along with the page access sequence for the sorting-based algorithm with three memory buffers. Sorting requires 16 I/Os, including 4 redundant I/Os on S_1, S_2, S_3 , and S_4 , using a page-access sequence of $R_0, R_1, S_0, S_1, S_2, R_2, R_3, S_1, S_2, S_3, S_4, R_4, R_5, S_3, S_4, S_5$.

A symmetric method may alternate between the pages of the two relations, as shown in Figure 9(c), to compute the join with 12 I/O (i.e. no redundant I/O) using a page-access sequence of $R_0, S_0, S_1, R_1, S_2, R_2, S_3, R_3, S_4, R_4, S_5, R_5$. This property can be generalized to other adjacency matrices with only B-diagonal entries, where $\{M_{PCG}[i, j] = 1\} \Rightarrow \{|i - j| \leq \lfloor B/2 \rfloor\}$, and B is the number of buffers available for pages of R and S. The indices i and j refer to the row-indices and column-indices. The symmetric method can process the B-diagonal entries of an adjacency matrix with no redundant I/Os, given B buffers for R and S.

The main approaches in **symmetric two-table clustering** are based on either the Traveling Salesman Problem heuristic or on incremental clustering. A traveling salesperson(TSP)-based heuristic [10] uses a complete graph constructed by taking the nodes of one relation as the nodes of the graph. The weight on an edge between nodes a and b denotes the number of page-accesses required to fetch all of the neighbors of b , given that all of the neighbors of a are in memory. This method requires a large amount of memory, as the complete graph grows quadratically with the number of nodes in the smaller of the relations. Incremental clustering is based on selecting the next page or the next set of pages to be fetched

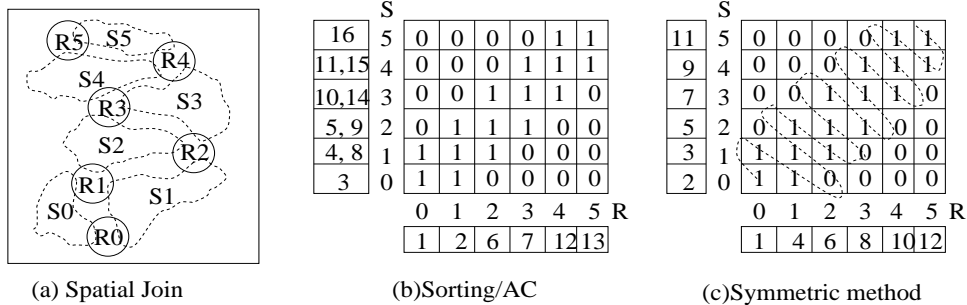


Figure 9: Comparison of symmetric and asymmetric methods

into memory, given the pages in the buffer and the remaining edges to be processed in the bi-partite page-connectivity graph. The selection is often based on the number of neighbors in the memory buffers and the number of neighbors on the disk. Details of actual heuristics follow.

Symmetric Heuristic: FP was proposed by Fotouhi and Pramanik [7]. The buffer is initialized with a node which has the smallest degree in the page-connectivity graph. The memory buffer is added with the largest resident-degree node. The resident degree of node A is the number of nodes which are connected to A and are in memory buffers. If more than one node has the largest resident degree, the algorithm chooses the one with the smallest non-resident degree. The non-resident degree of a node A is equal to $total_degree(A) - resident_degree(A)$. When the buffer is full, a node that has the smallest number of edges with the nodes on the disk can be swapped out.

Symmetric Heuristic: OM, developed by Omiecinski [28], is designed specifically for bipartite join graphs. Initially, load a pair of nodes (r_i, s_j) , where $r_i \in R$ and $s_j \in S$, from the page-connectivity graph in the memory buffers, such that (a) (r_i, s_j) is connected and (b) the sum of the degree of r_i and s_j is minimal. In each iteration, an out-of-memory least-non-residential-degree neighbor p of an in-memory lowest-non-residential-degree node q is selected to be swapped in. If the memory buffers are full, then the lowest-non-residential degree node r which is not connected to p in PCG is swapped out. Node r may come back to memory within the next few iterations.

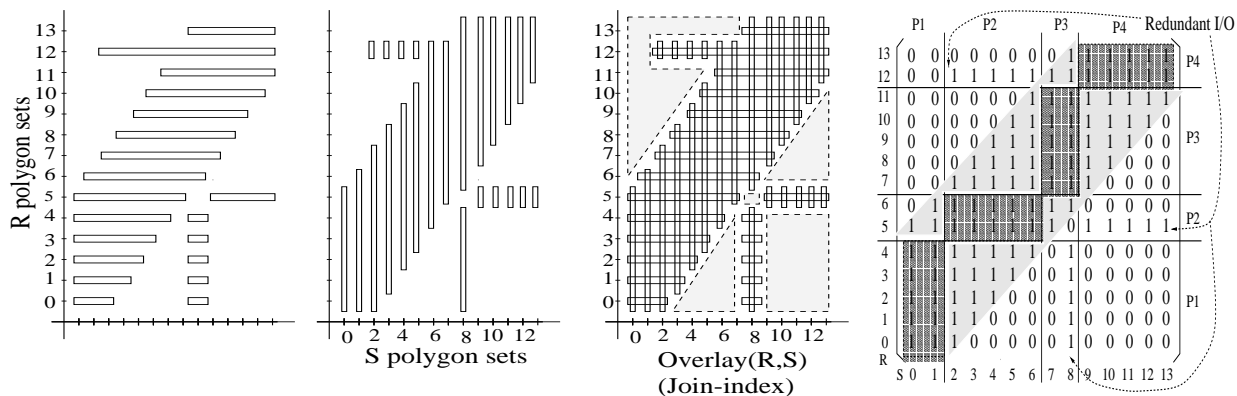
Symmetric Heuristic: Chan [5] generalized the OM heuristic by swapping in a set of nodes, i.e. segments, in each iteration. The set of nodes selected for an iteration are the unprocessed on-disk neighbors of the lowest-non-residential degree node n , either in memory or on disk. Node n can be processed and swapped out of the memory buffers at the end of the current iteration. If the memory buffers do not have enough empty space, then the page with the lowest number of a non-residential degree is swapped out.

Most symmetric methods proposed in the literature are incremental, considering local information in PCG. We propose a Symmetric Clustering(SC) method, which exploits global information across the entire PCG. SC is described in section 4.

3.1 An example

We use an example to illustrate the differences between various heuristics for computing joins, given a join-index. Readers are warned that this example is a bit detailed to bring out the differences between various methods. On the first reading, one may gloss over the details of Section 3.1, and simply look at

Table 3 to get the summary. We have tried hard to find an example which could both explain our method, i.e. SC, and explain the differences between all of the methods. We have not been able to find anything simpler. Figure 10(a) shows the polygon-clusters in R and S relations with their overlays. Some polygon-clusters have one polygon; others have two polygons, for example, $R_0, R_1, \dots, R_5, S_1, S_2, \dots, S_5$, and S_8, S_9, \dots, S_{13} . Polygon clusters are natural in geographic data as well. For example, the boundary of the United States will be represented by a collection of polygons representing the Mainland, Alaska, Hawaii, etc.



(a) Overlay of Two Relations

(b) Adjacency matrix of PCG for Join-index of overlay(R,S). Dark region=edges within a partition. Lightly shaded region=edges within B-diagonal

Figure 10: Connectivity graph of two relations

The spatial-join relationship shown uses geometric as well as topological representation. Figure 10(a) shows an overlay of the two data-sets to provide a visual representation of join relationships. Shaded areas are provided for realism. One may imagine a beautiful city with many parks and lakes(shaded areas) which break the continuity of streets (objects in R) and avenues (objects in S). Figure 10(b) provides adjacency matrix representation of the page-connectivity graph with one object per page. To simplify the example, we assume a unit blocking factor, i.e. one polygon-cluster or road per disk page.

A summary of the behavior of alternative methods is shown in Figure 11, which has five different pieces. The adjacency matrix representation of the join index is reproduced from Figure 10(b) to facilitate understanding. The IDs of the polygon clusters in R and S appear immediately to the left of the matrix. The two vectors left of and below the node-identifiers list the degree of each node in PCG. For example, the degree of R13 is 6. The nodes with the highest degree (13) are R5 and S8, followed by R12, which has a degree of 12.

The remaining two pieces of Figure 11 present the summary behavior of various algorithms in terms of page-access sequence. These tables show the rank-order of polygon-clusters (i.e. Pages assuming unit blocking factor) from R and S in their respective page-access sequences. For example, the page access sequence for OM is $R_0, S_0, S_1, S_2, S_8, R_1, S_3, R_2, S_4, R_3, S_5, R_4, R_5, R_6, \dots$ and so on. Multiple ranks for a node

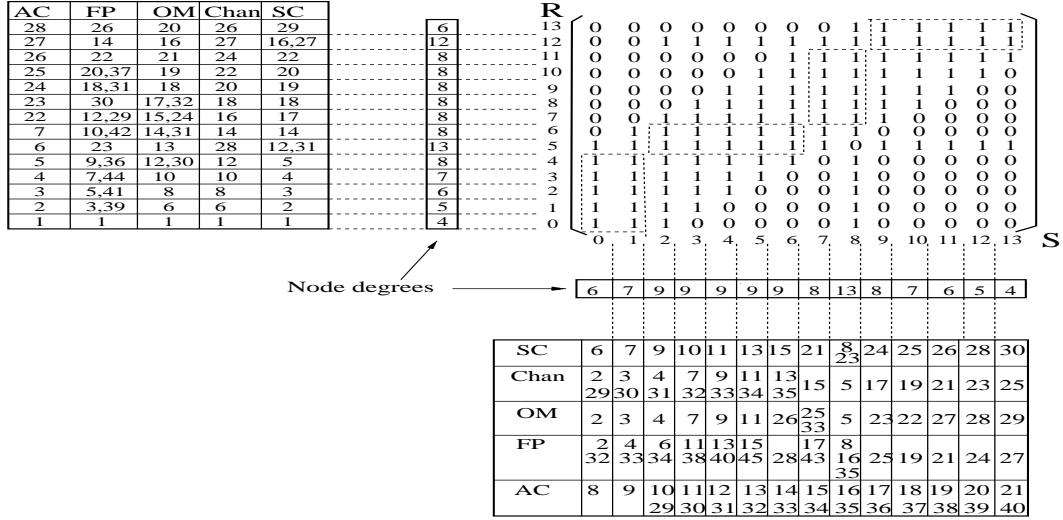


Figure 11: Comparison of different methods

indicate redundant I/O. For example, OM loads R4,R6,R7,R8, and S7 twice in memory with rankings of (12,30),(14,31),(15,24),(17,32), and (25,33). This information is further summarized in Table 3, which shows total I/O, nodes with redundant I/O, and min/max/average degree of nodes with redundant I/O. More details are available in Appendix A and B.

Method	AC	FP	OM	Chan	SC
Total I/O	40	45	33	35	31
Nodes with redundant I/O	S2-S13	R1-R4, R6, R7, R9, R10, S0-S5, S7, S8	R4,R6,R7, R8,S7	S0-S6	R12, R5, S13
Degree of nodes with redundant I/O	Avg	8	8.3	8	8.3
	Min	4	5	8	6
	Max	13	13	8	9

Table 3: Summary of page accesses for different algorithms

FP prefers to swap-in the nodes with the highest residential degree and swap out the nodes with the lowest non-residential degree. In this example, these policies tend to favor high degree nodes for buffering, leading to redundant I/O on low-to-medium degree nodes which have come memory already. Note that nodes R5 and R12 have no redundant I/O despite a very high degree, while R1, R2, and S0 have redundant I/O despite having a very low degree. Table 9(Appendix B.4) provides a detailed execution trace of FP for interested readers.

OM has a good overall performance with only 5 redundant I/Os. However, the redundant I/Os occur with nodes of medium degree, e.g. R4,R6,R7,R8, and S7(degree 8.) The nodes with a higher degree have no redundant I/O. This can be related to the swap-in and swap-out policy. Swap-in policy favors low-residential degree neighbors of low non-residential degree in-memory nodes. The policy has the

unintended effect of pulling all of the neighbors (S0,S1,S2,S8) in order of degree of the in-memory node with the lowest non-residential degree, e.g. R0. If one of these neighbors (S8) has a very high degree, it may stick around for a long time due to swap-out policy based on the lowest non-residential degree[‡]. This favored treatment of high non-residential degree nodes in memory leads to increased redundant I/O, in this example. Table 8(Appendix B.3) provides a detailed execution trace of OM for interested readers.

Chan’s method is a logical refinement of the OM method. It tries to bring in all of the neighbors of the node N with the lowest non-residential degree. This allows the complete processing of all edges incident on N . The hope is to reduce redundant I/O; however, Chan’s method leads to an unintended side effect for spatial joins. After R0 comes into the buffer, its neighbors(S0,S1,S2,S8) are brought in to process it. After R0 is processed, the next node with the lowest non-residential degree is R1, which comes in memory with its remaining neighbor S3. After R1 is processed, R2 is selected, and soon the nodes of R are favored to be processed first, since the first node to be loaded was from R. This leads to excessive redundant loading for nodes of S, in this example. Table 7(Appendix B.2) provides a detailed execution trace of Chan for interested readers.

In contrast to FP, OM, and Chan, SC uses a global clustering approach. SC uses symmetric clustering to permute the rows and columns of the adjacency matrix representation of PCG to bring as many edges as possible within the B-diagonal. All of the edges within the B-diagonal can be processed without redundant I/O. The redundant I/O for edges outside the B-diagonal are minimized by computing the vertex cover, i.e. a set of nodes covering all of the edges outside the B-diagonal. The nodes in the vertex cover are scheduled with appropriate partitions, and some of these vertices may yield multiple redundant I/Os. SC has only three redundant I/Os, resulting from a vertex cover {S8,R12,R5} of the 15 off B-diagonal edges. Incidentally, the nodes with redundant I/Os have a very high degree. Table 10(Appendix B.5) provides a detailed execution trace for interested readers.

4 Proposed Symmetric Clustering Method

In contrast to incremental approaches (e.g. FP, OM, Chan), Symmetric Clustering(SC) uses a global approach based on Band-diagonalization of the adjacency matrix representation of PCG. The number of redundant I/Os depends only on the edges outside of the B-diagonal and can often be reduced via identifying a vertex-cover.

Recall that pure B-diagonal entries for a square matrix are defined by $\{M_{PCG}[i, j] = 1\} \Rightarrow \{|i - j| \leq \lfloor B/2 \rfloor\}$, where B is the number of buffers available for pages of R and S. Band-diagonalization of a matrix rearranges the rows and columns of the matrix to bring in as many non-zero entries as possible within B-diagonals. Thus, a matrix with only B-diagonal elements is already band-diagonalized. However, a band-diagonalized matrix may have a few entries outside of the B-diagonal.

SC Algorithm

Input: $G = (V, E)$ is a page-connectivity graph; B is the number of buffers.

[‡]It is non-intuitive. See the explanation from [28] “If a node in the buffer is going to be swapped out, then in the worst case it could be brought in one additional time for each page it is connected to outside the buffer, assuming it were to be swapped out each time.” Thus swap-out kicks out the node with the lowest non-residential degree.

Output: $S = \langle P_1, P_2, \dots, P_r \rangle$ is a page-access sequence with $r \geq |V|$. (P_i s need not be distinct)

```

{Step 1} <  $G_{BD}, P_{order}$  > = Band-diagonalization( $G, B$ ); /* Get B-diagonal graph and ordered set of partitions */
{Step 2} <  $G_{OBD}$  > = Get-off-B-diagonal-entries( $G_{BD}$ ); /* Find all the off-B-diagonal edges and nodes from  $G_{BD}$  */
{Step 3} <  $VC$  > = Find-vertex-cover( $G_{OBD}$ ); /* Find the vertex cover  $VC$  for the Off-B-diagonal cut-edge  $E_{OBE}$  */
{Step 4} <  $S$  > = Access-sequence-generator( $P_{order}, G_{BD}, VC$ ); /* Generate the page access sequence */

```

First, SC derives the Band-diagonalized matrix G_{BD} by permuting the rows and columns of the adjacency matrix representation of PCG to bring in as many edges as possible within the B-diagonal. Secondly, from G_{BD} , SC gets the graph G_{OBD} for off-B-diagonal edges and their corresponding nodes. Thirdly, SC determines the vertex cover for G_{OBD} . Finally, SC generates the page-loading sequence based on the Band-diagonalized matrix, the vertex cover of the off-B-diagonal edges, and the partition ordering.

Example Revisited

Consider the join-computation problem discussed in Figures 10 and 11. The input to the SC algorithm is the page-connectivity graph shown in Figure 10(b).

In **step 1**, the nodes in the PCG of Figure 10(b) are rearranged to get as many edges with the B-diagonal as possible. The lightly shaded area in Figure 10(b) shows the B-diagonal. In this example, input graph G and output graph G_{BD} of step 1 are identical for simplicity. The nodes are partitioned in groups of 7 nodes, which is 1 less than the number of memory buffers available for R and S. Partitions $P_1 = \{R_0-R_4, S_0-S_1\}$, $P_2 = \{R_5-R_6, S_2-S_6\}$, $P_3 = \{R_7-R_{11}, S_7-S_8\}$, $P_4 = \{R_{12}-R_{13}, S_9-S_{13}\}$ are used in this example, as shown in the shaded rectangles in Figure 10(b). The partitions are loaded in the order P_1, P_2, P_3 , and then P_4 . The edges between the nodes within a common partition can be processed with no redundant I/O. The edges between the nodes that are in adjacent partitions in the loading sequence and which fall inside the B-diagonal can also be processed without any redundant I/Os, due to the availability of 1 extra buffer.

The redundant I/Os for the remaining edges can be reduced by computing the vertex cover via **step 2** and **3**. There are 15 edges off the B-diagonal with a vertex cover of $\{S_8, R_5, R_{12}\}$. There are 5 edges between partitions P_1 and P_3 , and they all incident on node S_8 . They can be processed with one extra I/O by bringing S_8 into the last buffer when the nodes of P_1 are in the buffer. Similarly, the 10 edges between the nodes in P_2 and P_4 can be processed in two I/Os. Since 5 of those are incident on R_{12} and R_5 , bringing R_{12} with P_2 and R_5 with P_4 will take care of all of these edges. Thus, SC results in only three redundant I/Os, resulting from a vertex cover $\{S_8, R_{12}, R_5\}$ for the 15 off B-diagonal edges. **Step 4** is to generate a page access sequence using the execution trace from previous steps.

4.1 The I/O cost of SC method

The symmetric clustering approach to minimize redundant I/O can be described in terms of the following problem statement:

Lemma 2 *Given a loading ordered partition $\{P_1, P_2, \dots, P_i, P_j, \dots, P_k\}$ of $PCG = (V, E)$, there is a page access sequence of length $k = |V| + \text{redundant-I/O}$ to compute the spatial join where the redundant-*

I/O is given by:

$$\sum_{V_i \in (\text{vertex cover of outside } B\text{-diagonal edges})} \text{Partition} - \text{degree}(V_i) \quad (1)$$

where $\text{Partition} - \text{degree}(V_i)$ is the number of distinct partitions which contain nodes V_j outside of the B-diagonals.

Proof: All of the edges within the main B-diagonals can be processed without redundant I/O, using a contour diagonalization strategy as illustrated in Figure 9(c). The redundant I/O for each node in the vertex cover is limited by the number of partitions sharing an off-B-diagonal edge.

4.2 A heuristic for Band-diagonalization

The first step of Band-diagonalization can be based on either specialized envelope-reduction algorithms [1, 9, 23] or min-cut graph partition algorithms [20, 21]. We use the latter in this paper and plan to explore the former in future work. We describe the heuristic approach that we currently use.

Band-diagonalization

Input: G is a page connectivity graph; B is the number of buffers.

Output: G_{BD} is the B-diagonal connectivity graph; P_{order} is the partition order

$PSet_m = \text{Graph-Partition}(G, B - 1); /* Using graph partition software */$

$P_{order} = \text{Order-Partition}(PSet_m); /* Order the partitions using the greedy heuristic */$

Graph-Partition: A min-cut partition algorithm, e.g. metis [20], divides the nodes of the PCG into disjoint subsets while minimizing the number of edges whose nodes are incident in two different partitions. One memory buffer is reserved for bringing in pages from the vertex cover of the off-B-diagonal entries. For example, the min-cut partitioning of PCG for the overlay(R,S) in Figure 10 may yield 4 partitions for $B=8$. The partitions shown in Figure 10(b) are $P1=\{R0-R4,S0-S1\}$, $P2=\{R5-R6,S2-S6\}$, $P3=\{R7-R11,S7-S8\}$, $P4=\{R12-R13,S9-S13\}$, resulting in 69 edges whose incident nodes are in two different partitions. The breakup of these edges by pairs of partitions of incident nodes is shown in Table 4. Formally, the min-cut graph-partitioning algorithm addresses the following problem:

Given: A connectivity graph $G = (V, E)$ with $|V| = n$, and the number of buffers, $B \geq 2$.

Find: A partition of V into p subsets, V_1, V_2, \dots, V_p such that $V_i \cap V_j = \emptyset$ for $i \neq j$ and $\bigcup_i V_i = V$.

Objective: Minimize the size of the set of edges $E_C \subseteq E$ whose incident vertices belong to different subsets.

Constraint: $|V_i| \leq (B - 1)$, and the number of partitions, $p = \lceil |V| / (B - 1) \rceil$.

Recent advances have provided scalable graph-partitioning software such as Metis [20], which can handle the large graphs relevant to databases in a relatively reasonable response time, e.g. a few seconds. We have had good results using it for database problems [24, 32].

Order-Partition chooses a partition ordering, i.e. a loading sequence, using a partition-interaction matrix M . An entry $M[P_i, P_j]$ lists the number of cut-edges between the nodes in partitions P_i and P_j . An example partition-interaction matrix for the join-index of Figure 10 is shown in Table 4. The procedure uses a simple heuristic to construct the loading sequence. It sorts the entries in $M[P_i, P_j]$ in descending order and arbitrarily breaks the tie. It initially chooses the entry with the largest value, getting a loading sequence of length 2. Then, it extends the loading on both sides in a greedy manner. For example, suppose $M[P_3, P_2]$ was selected first. Then, the loading sequence P_2 - P_3 can be extended to the right by choosing P_4 and extended to the left by choosing P_1 from among the remaining partitions. The choice is based on the highest value of $M[P_2, P_i]$ and $M[P_3, P_j]$ in the partition-interaction matrix. A better heuristic can be designed to select loading sequences that have a higher number of cut-edges between consecutive partitions. To improve the performance of the proposed SC method, we will consider these in future work.

	P_1	P_2	P_3	P_4
P_1	*	18	5	0
P_2	18	*	18	10
P_3	5	18	*	18
P_4	0	10	18	*

Table 4: # of cut-edges between partitions P_j and P_i

We use Figure 12 as an example to illustrate the steps in Band-diagonalization using the graph-partition technique. Figure 12(a) is the original PCG relation, where R and S are two relations to be joined, and where each point in the picture denotes an edge connection between the pages of two relations. We use Metis [20] to partition this PCG, and each partition has size $(B - 1)$, where B is the number of buffers available. We show the result after the partition in Figure 12(b): the pages from the R and S relations are relabeled by their partitions. Finally, we re-order these partitions to bring as many points as possible inside the B-diagonal, as shown in Figure 12(c). In Figure 12(b), 28% of the edges are outside the B-diagonal. After partition reordering, we have reduced these edges to 22% of the total edges. Edges outside of the B-diagonal can be processed by using a vertex cover, as discussed next.

4.3 Vertex Cover Computation

The redundant I/Os in the SC approach are due to the edges(i.e. non-zero matrix elements) outside the B-diagonal of a clustered adjacency-matrix representation of a join-index. These outside-B-diagonal edges are grouped via a vertex-cover algorithm to determine a small set of pages needing redundant I/Os. Determining the minimal vertex-cover for a general graph is NP-hard [8]. However, polynomial time algorithms [14] are available for determining minimal vertex covers for bi-partite graphs, e.g. PCG for join-indices.

The Find-vertex-cover procedure determines the vertex cover for all of the off-B-diagonal edges by a fast but greedy heuristic which is described below. The heuristic sorts the nodes related to the off-B-diagonal edges by their degree, i.e. the number of incident edges. The node with the highest degree is added to the vertex cover and all of the edges incident on this node are dropped. These steps are

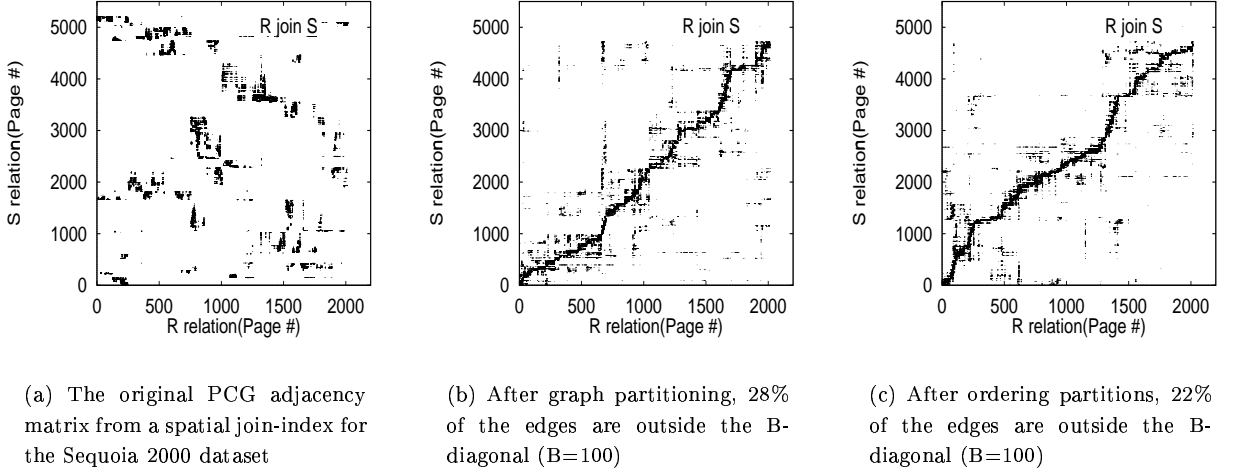


Figure 12: Using graph partitioning to derive the B-diagonal

repeated to cover all of the off-B-diagonal edges. In the future, we plan to use better algorithms which are likely to improve the performance of the proposed SC method.

Find-vertex-cover

Input: $G_{OBD} = (V_{OBD}, E_{OBD})$ is the graph for off-B-diagonal edges and corresponding nodes

Output: VC is the vertex cover for G_{OBD}

```

while( $E_{OBD}$  is not empty) {
     $V_{highest} = \text{Find\_highest\_degree\_node}(V_{OBD});$  /* Node  $V_{highest}$  has the highest degree */
     $VC = VC \cup V_{highest};$  /* Add this node to the set of vertex cover */
     $\text{Update}(G_{OBD});$  /* Update  $G_{OBD}$  by removing node  $V_{highest}$  and its corresponding edges */
}

```

In Figure 10, for example, with the loading sequence $P_1 - P_2 - P_3 - P_4$, the vertex cover for the 15 off-B-diagonal cut-edges are nodes R5, R12 and S8. Page R5 covers edges (R5,S9), (R5,S10), (R5,S11), (R5,S12) and (R5, S13). Page R12 covers edges (R12,S2), (R12,S3), (R12,S4), (R12,S5) and (R12,S6). Page S8 covers edges (R0,S8), (R1,S8), (R2,S8), (R3,S8) and (R4,S8).

4.4 Generate access sequence

The Access-sequence-generator procedure derives the page-access sequence. It loads each partition in the pre-determined order. When transferring from one partition P_i to the next scheduled partition P_{i+1} , the procedure orders the loading sequence of the nodes using the contour-diagonalization order shown in Figure 9(c). After loading a whole partition, we find all of the off-B-diagonal vertex cover nodes which connect to this partition, and load these nodes, one by one, to compute the join.

Access-sequence-generator

Input: P_{order} is the loading sequence of the partitions;
 G_{BD} is the B-diagonal connectivity graph;
 VC is the vertex cover for all the Off-B-diagonal edges.
Output: $S = \langle P_1, P_2, \dots, P_r \rangle$ is a page access sequence.

```

for( $i = 1; i \leq |P_{order}|; i++$ ){
     $P_i = \text{GetPartition}(P_{order}, i)$  /* Get the  $i$ th partition */
    if( $i==1$ ) {  $\text{AddPageSequence}(S, P_i)$ ; /* Add all the nodes within  $P_1$  into the loading sequence */ }
    else {  $\text{OrderAndAddPageSequence}(S, P_{i-1}, P_i)$ ;
        /* Order and add the nodes within  $P_i$  into the loading sequence by the following rules: */
        /* 1. Add the node within  $P_i$  which has the highest connectivity with  $P_{i-1}$  in B - diagonal */
        /* 2. Replace the node within  $P_{i-1}$  which has finished its join with the nodes in  $P_i$  */
    }
     $\text{PVC\_Set} = \text{FindConnected\_node\_from\_VC}(P_i, VC)$ ;
    /* Find if any Off-B-diagonal vertex cover which connects to this partition */
     $\text{AddPageSequence}(S, \text{PVC\_Set})$ ; /* Add these nodes into the loading sequence */
}

```

5 Comparative Evaluation of SC, AC and Competitors

The experimental setup is shown in Figure 6. Recalling that the dataset is extracted from Sequoia 2000 [33], we construct join-indices for N-nearest-neighbor(Q.A), as well as for distance-based range queries(Q.B). Variable parameters include buffer size, page size, and edge ratio. Relation R refers to the GNIS *Point* table, and relation S refers to the Landuse *Polygon* table from the Sequoia 2000 [33] dataset. For the sake of brevity, we refer readers to section 2.3.1 for details of the experiment design.

5.1 Experiment results

Figures 13 and 14 compare all of the OPAS-FB heuristics for N-nearest-neighbor join-indices and range-query join-indices, respectively. For different experiments, we vary the page size, buffer size, and edge ratio.

5.1.1 Page size

Page size affects the clustering of the base relations and also the degree of the nodes in the PCGs. We study the effect of page size on the performance of the OPAS-FB methods. Figure 13(a) and 14(a) show the effect of page size. The page size is varied from 2 kbytes to 64 kbytes. The page size and the number of page accesses are shown in the logarithmic form of base two. In the N-nearest-neighbor join-indices, the SC and AC outperform all of the other methods using different page sizes. In the Range-query join-indices, SC and AC require fewer page accesses than all of the other methods. The AC outperforms SC when page size is greater than 64 kbytes. The OM's method performs well with the 4 and 8 kbyte page size.

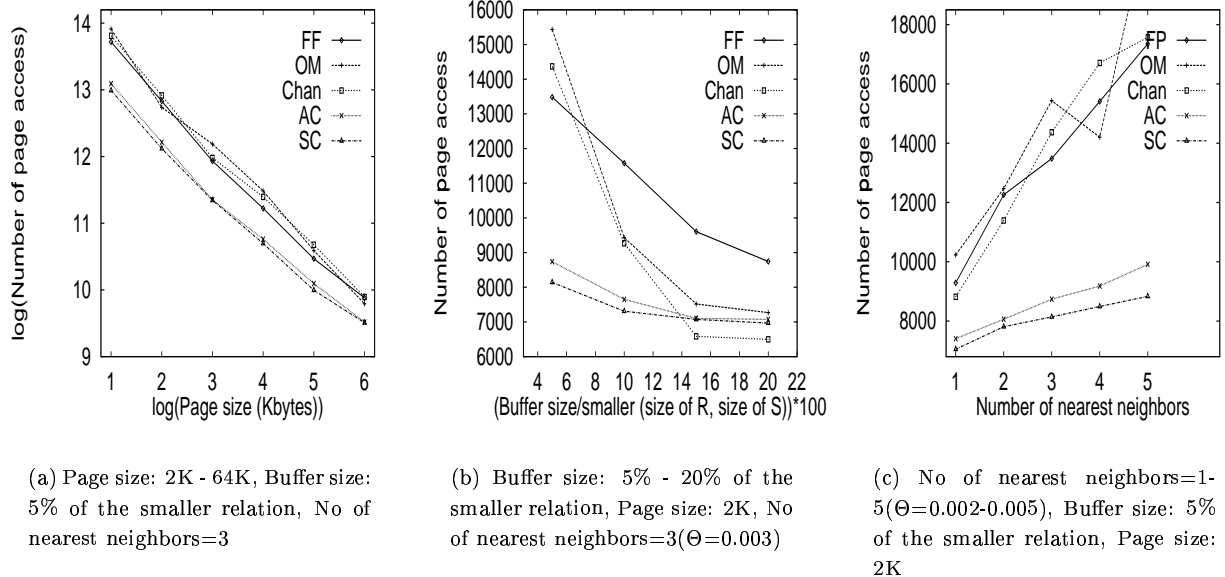


Figure 13: Nearest Neighbor Query: The effect of page size, buffer size, and number of nearest neighbors for different OPAS-FB heuristics

5.1.2 Buffer size

Figure 13(b) and 14(b) show the effect of buffer size. We varied the number of buffers as a percentage of the number of pages of the smaller relation. The percentage is changed from 5 to 20. The AC and SC perform better than all of the other methods when the buffer size is relatively low, e.g. 5 to 10 percent. The Chan and OM methods do well with large buffer sizes.

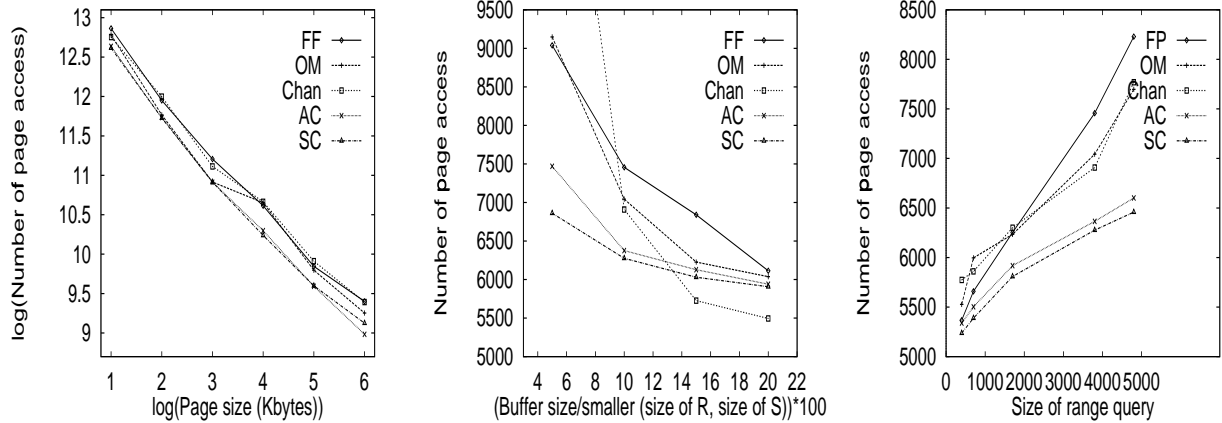
5.1.3 Edge Ratio

In this experiment, we changed the edge ratio by increasing/decreasing the number of neighbors and the size of the range query. The result of the experiment is shown in Figures 13(c) and 14(c). AC and SC uniformly outperform all of the other methods. The SC requires fewer page accesses than AC.

6 Conclusion and Future Work

In this paper, we introduced two new algorithms for spatial join computation, given a join-index and a fixed buffer size. The key idea is to use spatial clustering. The proposed AC and SC algorithms outperformed the traditional methods in our experiments with the Sequoia 2000 [33] dataset, particularly when the size of the memory buffer was small (<10%), relative to the size of the spatial relations. We also provided a formal characterization of an upper bound on the number of redundant I/Os needed by AC and SC.

In the future we would like to improve some of the heuristics chosen in the implementation of AC and SC, as discussed throughout this paper. We would also like to look at related issues regarding



(a) Page size: 2K - 64K, Buffer size: 10% of the smaller relation, Range query size = 3800

(b) Buffer size: 5% - 20% of the smaller relation, Page size: 2K, Range query size = 3800($\Theta=0.003$)

(c) Range query size = 400 - 4800($\Theta=0.002-0.003$), Buffer size: 10% of the smaller relation, Page size: 2K

Figure 14: Range Query: Effect of page size, buffer size, and number of nearest neighbors for different OPAS-FB heuristics

maintenance of join-indices in the face of updates and also the interaction of join-indices with the join-computation algorithm. Finally, we are interested in exploring the usefulness of AC and SC in data Warehouses [15, 16] (e.g. processing star-joins using the STARindex [15]) and spatial data mining(e.g. neighborhood index [6]).

Acknowledgments

This work is sponsored in part by the Army High Performance Computing Research Center under the auspices of the Department of the Army, Army Research Laboratory cooperative agreement number DAAH04-95-2-0003/contract number DAAH04-95-C-0008, the content of which does not necessarily reflect the position or the policy of the government, and no official endorsement should be inferred. This work was also supported in part by NSF grant #9631539. We would like to thank Christiane McCarthy for improving the the readability and technical accuracy of this paper. We also thank Xuan Liu, Xinhong Tan and Weili Wu for their technical comments.

References

- [1] S. T. Barnard, A. Pothan, and H. D. Simon. A Spectral Algorithm for Envelope Reduction of Sparse Matrices. *Numerical Linear Algebra with Applications*, 2:317-334, 1995.
- [2] L. Becker, K. Hinrichs, and U. Finke. A New Algorithm for Computing Joins With Grid Files. In *Proceedings of International Conference on Data Engineering*, 1993.
- [3] T. Brinkhoff, H. Kriegel, R. Schneider, and B. Seeger. Multi-Step Processing of Spatial Joins. In *Proceedings of ACM SIGMOD International Conference on Management of Data*, May 1994.
- [4] T. Brinkhoff, H. Kriegel, and B. Seeger. Efficient Processing of Spatial Joins Using R-trees. In *Proceedings of ACM SIGMOD International Conference on Management of Data*, May 1990.

- [5] Chee Yong Chan and Beng Chin Ooi. Efficient Scheduling of Page Access in Index-Based Join Processing. *IEEE Transactions on Knowledge and Data Engineering*, November/December 1997.
- [6] M. Ester and J. Sander S. Gundlach, H. Kriegel. Database Primitives for Spatial Data Mining. *Proc. Int. Conf. on Databases in Office, Engineering and Science, Freiburg, Germany*, 1999.
- [7] F. Fotouhi and S. Pramanik. Optimal Secondary Storage Access Sequence for Performing Relational Join. *IEEE Transactions on Knowledge and Data Engineering*, September 1989.
- [8] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1993.
- [9] A. George and A. Pothen. An Analysis of Spectral Envelope-Reduction via Quadratic Assignment Problems. *SIAM Journal of Matrix Analysis and its Applications*, 18(3):706–732, 1997.
- [10] P. Goyal, H.F. Li, E. Regener, and F. Sadri. Scheduling of Page Fetches in Join Operation Using Bc-Trees. In *Proceedings of International Conference on Data Engineering*, 1988.
- [11] G. Graefe. Query Evaluation Techniques for Large Databases. *Computing Surveys*, 25(2):73–170, 1993.
- [12] O. Gunther. Efficient Computation of Spatial Joins. In *Proceedings of International Conference on Data Engineering*, 1993.
- [13] L. Hagen and A. Kahng. Fast Spectral Methods for Ratio Cut Partitioning and Clustering. In *Proceedings of IEEE International Conference on Computer Aided Design*, 1991.
- [14] J.E. Hopcroft and R.M. Karp. An $n^{5/2}$ algorithm for maximum matching of graphs. *SIAM J. Comput.*, 2(4):225–231, 1973.
- [15] Informix. *White Papers*, <http://www.informix.com/informix/solutions/dw/redbrick/wpapers/star.html>.
- [16] W. H. Inmon. *Building the Data Warehouse*. John Wiley & Sons Inc, 1992.
- [17] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar. hMetis Home Page. <http://www-users.cs.umn.edu/karypis/metis/hmetis/main.html>.
- [18] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar. Multilevel hypergraph partitioning: Application in VLSI domain. *Proceedings ACM/IEEE Design Automation Conference*, 1997.
- [19] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar. Multilevel hypergraph partitioning: Application in VLSI domain. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 7(1):69–79, March 1999.
- [20] G. Karypis and V. Kumar. Metis Home Page. <http://www-users.cs.umn.edu/karypis/metis/metis/main.html>.
- [21] G. Karypis and V. Kumar. Parallel Multilevel Graph Partitioning. In *Proceedings of Supercomputing*, November 1996.
- [22] B. W. Kernighan and S. Lin. An Efficient Heuristic Procedure for Partitioning Graphs. *The Bell System Technical Journal*, 1970.
- [23] G. Kumpf and A. Pothen. Two Improved Algorithms for Envelope and Wavefront Reduction. *BIT*, 37(3):001–032, 1997.
- [24] D. R. Liu and S. Shekhar. A Similarity Graph-Based Approach to Declustering Problem and its Applications Toward Parallelizing Grid Files. In *Proceedings of the Eleventh International Conference on Data Engineering, IEEE*, pages 373–381, March 1995.
- [25] M. Lo and C. V. Ravishankar. Spatial Joins Using Seeded Trees. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data*, pages 209–220, 1994.
- [26] T. Merrett, Y. Kimbayasi, and H. Yasuura. Scheduling of Page-Fetches in Join Operations. In *Proceedings of the 7th International Conference on Very Large Databases*, 1981.
- [27] P. Mishra and M.H. Eich. Join Processing in Relational Databases. *Computing Surveys*, 24(1):63–113, 1992.
- [28] Edward R. Omiecinski. Heuristics for Join Processing Using Nonclustered Indexes. *IEEE Transactions on Software Engineering*, 15, January 1989.
- [29] S. Pramanik and D. Ittner. Use of Graph Theoretic Models for Optimal Relational Database Accesses to Perform Join. *ACM Transactions on Database Systems*, March 1985.
- [30] D. Rotem. Spatial Join Indices. In *Proceedings of International Conference on Data Engineering*, 1991.
- [31] S. Shekhar, S. Chawla, S. Ravada, A. Fetterer, X. Liu, and C. Lu. Spatial Databases: Accomplishments and Research Needs. *IEEE Transactions on Knowledge and Data Engineering*, 11(1):45–55, 1999.
- [32] S. Shekhar and D. R. Liu. CCAM: A Connectivity-Clustered Access Method for Networks and Networks Computations. *IEEE Trans. on Knowledge and Data Engineering*, 9(1), January 1997.
- [33] M. Stonebraker, J. Frew, and J. Dozier. The Sequoia 2000 Project. In *Proceedings of the Third International Symposium on Large Spatial Databases*, 1993.
- [34] P. Valduriez. Join Indices. *ACM Transactions on Database Systems*, pages 218–246, December 1987.

A Summary trace for different algorithms in Figure 11

Table 5 provides a summary trace of various heuristics. Each row represents a new page being fetched into main memory. Thus, the number of rows represent the total number of pages fetched. For simplicity, the clustering of I/O for multiple pages is not modeled. An entry $(+R0)$ for Sorting in iteration 1 means that page $R0$ was fetched. An entry $(+S1, = S1)$ for iteration 9 of Sorting implies that $S1$ was fetched into memory $(+S1)$ and that all edges incident on $S1$ were processed with the pages available in the buffer. The set of pages available in the buffer in this iteration are $\{R0, R1, R2, \dots, R6\}$, since we have fetched those in previous iterations. The buffer containing $S1$ can be reused in the next iteration to bring in $S2$, as shown by entry $(+S2)$ for Sorting in iteration 10. The next interesting entry is $(-S2, +S3)$ in iteration 11 for Sorting, where the buffer containing $S2$ is overwritten by incoming page $S3$ even though some edges (e.g. $R7 - S2$) incident cannot be processed right away. This is due to a buffer size fixed at eight. Note that $S2$ returns to memory in iteration 29 to process the edge $(R7 - S2)$. This leads to a redundant I/O. Because the graph is highly connected in spatial order, the AC will not re-cluster the R relation, and uses the same loading sequence as the Sorting heuristic.

Traces of other algorithms are shown in other columns using 8 buffers. Note that the number of the last rows with a '+' entry (page fetch) designates the total number of page I/Os for an algorithm. In other words, the Sorting/AC-based algorithm has 40 I/Os, Chan's heuristic has 35, OM has 33, FF has 45 and SGP has 31, as shown in the last row labeled I/O count.

Algorithms for computing Join with a Join Index					
Iteration	Sorting/ASC	Chan	OM	FP	SSC
1	+R0	+R0	+R0	+R0	+R0
2	+R1	+S0	+S0	+S0	+R1
3	+R2	+S1	+S1	+R1	+R2
4	+R3	+S2	+S2	+S1	+R3
5	+R4	+S8,=R0	+S8,=R0	+R2	+R4
6	+R5	+R1	+R1	+S2	+S0
7	+R6	+S3,=R1	+S3,=R1	+R3	+S1
8	+S0,=S0	+R2	+R2	+S8,=R0	+S8
9	+S1,=S1	+S4,=R2	+S4,=R2	+R4	-S8,+S2,=R0
10	+S2	+R3	+R3	-R1,+R6	+S3,=R1
11	-S2,+S3	+S5,=R3	+S5,=R3	-S0,+S3	+S4,=R2
12	-S3,+S4	+R4	+R4	-R2,+R7	+R5,=S0
13	-S4,+S5	-S0,+S6,=R4	-R4,+R5,=S0	-S1,+S4	+S5,=R3
14	-S5,+S6	+R6	+R6,=S1	-R3,+R12	+R6,=S1
15	-S6,+S7	-S1,+S7,=R6	+R7	-S2,+S5	+S6,=R4
16	-S7,+S8	+R7	-R6,+R12,=S2	-R4,+S8	+R12
17	-S8,+S9	-S2,+S9,=R7	+R8,=S3	-R6,+S7	-R12,+R7,=S2
18	-S9,+S10	+R8	+R9,=S4	-R7,+R9	+R8,=S3
19	-S10,+S11	-S3,+S10,=R8	+R10,=S5	-S3,+S10	+R9,=S4
20	-S11,+S12	+R9	+R13	-R8,+R10	+R10,=S5
21	-S12,+S13,=R0	-S4,+S11,=R9	-R7,+R11,=S8	-S4,+S11	+S7
22	-S13,+R7,=R1	+R10	+S10,=S10	-R9,+R11	-R5,+R11,=S6
23	+R8,=R2	-S5,+S12,=R10	+S9	-S8,+R5	+S8,=R6
24	+R9,=R3	+R11	-R8,+R7,=S9	-S5,+S12	+S9,=R7
25	+R10,=R4	-S6,+S13,=R11	+S7	-S7,+S9	+S10,=R8
26	+R11,=R5	+R13,=R13	-S7,+S6,=R7	-R10,+R13,=S10	+S11,=R9
27	+R12,=R6	+R12,=S8	+S11,=R9,=S11	+S13,-S13,=S11	+R12,=S7
28	+R13	+R5,=S7	+S12,=R10,=S12	+S6,=R11,=R12	+S12,=R10
29	+S2,=S2	+S0,=S0,=S9	+S13,=R5,=R11	+R7,=R7,=S12	+R13,=S8
30	+S3,=S3	+S1,=S1,=S10	+R4,=R4,=R12	+R8,=R8	+S13,=R11
31	+S4,=S4	+S2,=S2,=S11	+R6,=R13,=S13	+R9,=R9,=S9	+R5
32	+S5,=S5	+S3,=S3,=S12	+R8,=S6	+S0,=S0	
33	+S6,=S6	+S4,=S4,=S13	+S7	+S1,=S1	
34	+S7,=S7	+S5,=S5		+S2,=S2	
35	+S8,=S8	+S6		+S8,=R13,=S8	
36	+S9,=S9			+R4,=R4	
37	+S10,=S10			+R10,=R10	
38	+S11,=S11			+S3	
39	+S12,=S12			+R1,=R1,=S3	
40	+S13			+S4,=R5	
41				+R2,=R2,=S4	
42				+R6,=S6	
43				+S7,=R6,=S7	
44				+R3	
45				+S5	
I/O count	40	35	33	45	31

Table 5: Summary of results for all algorithms(Buffer size=8, -:Swap Out, +:Add, =: Done)

B Trace for different algorithms in Figure 11

B.1 Execution Trace for ASC/Sorting method

Table 6 shows the behavior of the Sorting method for computing a join, given the join-index of Figure 10. Table 6 has five columns. The first column shows the iteration number. The second column shows the node swapped out in the current iteration. The third column shows the node selected and brought into the memory buffers. The fourth and fifth columns show the pages of R and S in the main memory buffer. The last column shows the nodes which have been processed completely and need not come into the memory buffer again.

Iteration	Swap out	Add	Page of R in Buffer	Page of S in Buffer	Done
1		R0	R0		
2		R1	R0, R1		
3		R2	R0, R1, R2		
4		R3	R0, R1, R2, R3		
5		R4	R0, R1, R2, R3, R4		
6		R5	R0, R1, R2, R3, R4, R5		
7		R6	R0, R1, R2, R3, R4, R5, R6		
8		S0	R0, R1, R2, R3, R4, R5, R6	S0	S0
9		S1	R0, R1, R2, R3, R4, R5, R6	S1	S1
10		S2	R0, R1, R2, R3, R4, R5, R6	S2	
11	S2	S3	R0, R1, R2, R3, R4, R5, R6	S3	
12	S3	S4	R0, R1, R2, R3, R4, R5, R6	S4	
13	S4	S5	R0, R1, R2, R3, R4, R5, R6	S5	
14	S5	S6	R0, R1, R2, R3, R4, R5, R6	S6	
15	S6	S7	R0, R1, R2, R3, R4, R5, R6	S7	
16	S7	S8	R0, R1, R2, R3, R4, R5, R6	S8	
17	S8	S9	R0, R1, R2, R3, R4, R5, R6	S9	
18	S9	S10	R0, R1, R2, R3, R4, R5, R6	S10	
19	S10	S11	R0, R1, R2, R3, R4, R5, R6	S11	
20	S11	S12	R0, R1, R2, R3, R4, R5, R6	S12	
21	S12	S13	R0, R1, R2, R3, R4, R5, R6	S13	R0, R1, R2, R3, R4, R5, R6
22	S13	R7	R7		
23		R8	R7, R8		
24		R9	R7, R8, R9		
25		R10	R7, R8, R9, R10		
26		R11	R7, R8, R9, R10, R11		
27		R12	R7, R8, R9, R10, R11, R12		
28		R13	R7, R8, R9, R10, R11, R12, R13		
29		S2	R7, R8, R9, R10, R11, R12, R13	S2	S2
30		S3	R7, R8, R9, R10, R11, R12, R13	S3	S3
31		S4	R7, R8, R9, R10, R11, R12, R13	S4	S4
32		S5	R7, R8, R9, R10, R11, R12, R13	S5	S5
33		S6	R7, R8, R9, R10, R11, R12, R13	S6	S6
34		S7	R7, R8, R9, R10, R11, R12, R13	S7	S7
35		S8	R7, R8, R9, R10, R11, R12, R13	S8	S8
36		S9	R7, R8, R9, R10, R11, R12, R13	S9	S9
37		S10	R7, R8, R9, R10, R11, R12, R13	S10	S10
38		S11	R7, R8, R9, R10, R11, R12, R13	S11	S11
39		S12	R7, R8, R9, R10, R11, R12, R13	S12	S12
40		S13	R7, R8, R9, R10, R11, R12, R13	S13	R7, R8, R9, R10, R11, R12, R13, S13

Table 6: Example: the Sorting/ASC method(Buffer size=8)

B.2 Execution Trace for Chan’s method

Table 7 shows the behavior of Chan’s method for computing a join, given the join-index of Figure 10. Table 7 has five columns. The first column shows the iteration number. The second column shows the node swapped out in the current iteration. The third column shows the node selected and brought into the memory buffers. The fourth and fifth columns show the pages of R and S in the main memory buffer. The last column shows the nodes which have been processed completely and need not come into the memory buffer again.

Iteration	Swap out	Add	R Buffer	S Buffer	Done
1		R0	R0		
2		S0	R0	S0	
3		S1	R0	S0,S1	
4		S2	R0	S0,S1,S2	
5		S8	R0	S1,S2, S8	R0
6		R1	R1	S0,S1,S2, S8	
7		S3	R1	S0,S1,S2, S8 ,S3	R1
8		R2	R2	S0,S1,S2, S8 ,S3	
9		S4	R2	S0,S1,S2, S8 ,S3,S4	R2
10		R3	R3	S0,S1,S2, S8 ,S3,S4	
11		S5	R3	S0,S1,S2, S8 ,S3,S4,S5	R3
12		R4	R4	S0,S1,S2, S8 ,S3,S4,S5	
13	S0	S6	R4	S1,S2, S8 ,S3,S4,S5,S6	R4
14		R6	R6	S1,S2, S8 ,S3,S4,S5,S6	
15	S1	S7	R6	S2,S8,S3,S4,S5,S6,S7	R6
16		R7	R7	S2,S8,S3,S4,S5,S6,S7	
17	S2	S9	R7	S8,S3,S4,S5,S6,S7,S9	R7
18		R8	R8	S8,S3,S4,S5,S6,S7,S9	
19	S3	S10	R8	S8,S4,S5,S6,S7,S9,S10	R8
20		R9	R9	S8,S4,S5,S6,S7,S9,S10	
21	S4	S11	R9	S8,S5,S6,S7,S9,S10,S11	R9
22		R10	R10	S8,S5,S6,S7,S9,S10,S11	
23	S5	S12	R10	S8,S6,S7,S9,S10,S11,S12	R10
24		R11	R11	S8,S6,S7,S9,S10,S11,S12	
25	S6	S13	R11	S8,S7,S9,S10,S11,S12,S13	R11
26		R13	R13	S8,S7,S9,S10,S11,S12,S13	R13
27		R12	R12	S8,S7,S9,S10,S11,S12,S13	S8
28		R5	R5, R12	S7,S9,S10,S11,S12,S13	S7,S9,S10,S11,S12,S13
29		S0	R5, R12	S0	S0
30		S1	R5, R12	S1	S1
31		S2	R5, R12	S2	S2
32		S3	R5, R12	S3	S3
33		S4	R5, R12	S4	S4
34		S5	R5, R12	S5	S5
35		S6	R5, R12	S6	R12,R5,S6

Table 7: Example: Chan’s method(Buffer size=8)

B.3 Execution Trace for OM's method

Table 8 shows the behavior of OM's method for computing a join, given the join-index of Figure 10. Table 8 has five columns. The first column shows the iteration number. The second column shows the node swapped out in the current iteration. The third column shows the node selected and brought into the memory buffers. The fourth and fifth columns show the pages of R and S in the main memory buffer. The last column shows the nodes which have been processed completely and need not come into the memory buffer again.

Iteration	Swap out	Add	R Buffer	S Buffer	Done
1		R0	R0		
2		S0	R0	S0	
3		S1	R0	S0,S1	
4		S2	R0	S0,S1,S2	
5		S8	R0	S0,S1,S2, S8	R0
6		R1	R1	S0,S1,S2, S8	
7		S3	R1	S0,S1,S2, S8 ,S3	R1
8		R2	R2	S0,S1,S2, S8 ,S3	
9		S4	R2	S0,S1,S2, S8 ,S3,S4	R2
10		R3	R3	S0,S1,S2, S8 ,S3,S4	
11		S5	R3	S0,S1,S2, S8 ,S3,S4,S5	R3
12		R4	R4	S0,S1,S2, S8 ,S3,S4,S5	
13	R4	R5	R5	S0,S1,S2, S8 ,S3,S4,S5	S0
14		R6	R5,R6	S1,S2, S8 ,S3,S4,S5	S1
15		R7	R5,R6,R7	S2, S8 ,S3,S4,S5	
16	R6	R12	R5,R7, R12	S2, S8 ,S3,S4,S5	S2
17		R8	R5,R7, R12 ,R8	S8 ,S3,S4,S5	S3
18		R9	R5,R7, R12 ,R8,R9	S8 ,S4,S5	S4
19		R10	R5,R7, R12 ,R8,R9,R10	S8 ,S5	S5
20		R13	R5,R7, R12 ,R8,R9,R10,R13	S8	
21	R7	R11	R5,R12,R8,R9,R10,R13,R11	S8	S8
22		S10	R5,R12,R8,R9,R10,R13,R11	S10	S10
23		S9	R5,R12,R8,R9,R10,R13,R11	S9	
24	R8	R7	R5,R12,R9,R10,R13,R11,R7	S9	S9
25		S7	R5,R12,R9,R10,R13,R11,R7	S7	
26	S7	S6	R5,R12,R9,R10,R13,R11,R7	S6	R7
27		S11	R5,R12,R9,R10,R13,R11	S6, S11	R9,S11
28		S12	R5,R12,R10,R13,R11	S6, S12	R10,S12
29		S13	R5,R12,R13,R11	S6, S13	R5,R11,R12,R13,S13
30		R4	R4	S6	R4
31		R6	R6	S6	
32		R8	R6, R8	S6	S6
33		S7	R6, R8	S7	R6,R8,S7

Table 8: Example: OM's method(Buffer size =8)

B.4 Execution Trace for FP's method

Table 9 shows the behavior of FP's method for computing a join, given the join-index of Figure 10. Table 9 has five columns. The first column shows the iteration number. The second column shows the node swapped out in the current iteration. The third column shows the node selected and brought into the memory buffers. The fourth and fifth columns show the pages of R and S in the main memory buffer. The last column shows the nodes which have been processed completely and need not come into the memory buffer again.

Iteration	Swap out	Add	R Buffer	S Buffer	Done
1		R0	R0		
2		S0	R0	S0	
3		R1	R0,R1	S0	
4		S1	R0,R1	S0,S1	
5		R2	R0,R1,R2	S0,S1	
6		S2	R0,R1,R2	S0,S1,S2	
7		R3	R0,R1,R2,R3	S0,S1,S2	
8		S8	R0,R1,R2,R3	S0,S1,S2, S8	R0
9		R4	R1,R2,R3,R4	S0,S1,S2, S8	
10	R1	R6	R2,R3,R4, R6	S0,S1,S2, S8	
11	S0	S3	R2,R3,R4, R6	S1,S2, S8 ,S3	
12	R2	R7	R3,R4, R6 ,R7	S1,S2, S8 ,S3	
13	S1	S4	R3,R4, R6 ,R7	S2, S8 ,S3,S4	
14	R3	R12	R4, R6 ,R7,R12	S2, S8 ,S3,S4	
15	S2	S5	R4, R6 ,R7,R12	S8 ,S3,S4,S5	
16	R4	R8	R6,R7, R12 ,R8	S8 ,S3,S4,S5	
17	R6	S7	R7, R12 ,R8	S8,S3,S4,S5, S7	
18	R7	R9	R12 ,R8,R9	S8,S3,S4,S5, S7	
19	S3	S10	R12 ,R8,R9	S8,S4,S5, S7 , S10	
20	R8	R10	R12 ,R9,R10	S8,S4,S5, S7 , S10	
21	S4	S11	R12 ,R9,R10	S8,S5, S7 , S10 ,S11	
22	R9	R11	R12,R10,R11	S8,S5, S7 , S10 ,S11	
23	S8	R5	R12, R10 ,R11,R5	S5, S7 , S10 ,S11	
24	S5	S12	R12, R10 ,R11,R5	S7, S10 ,S11,S12	
25	S7	S9	R12, R10 ,R11,R5	S10,S11,S12,S9	
26	R10	R13	R12, R11 ,R5,R13	S10,S11,S12,S9	S10,S11,S12
27		S13	R12, R11 ,R5,R13	S9,S13	S13
28		S6	R12, R11 ,R5,R13	S9 ,S6	R11,R12
29		R7	R5, R13 , R7	S9 ,S6	R7
30		R8	R5, R13 , R8	S9 ,S6	R8
31		R9	R5, R13 , R9	S9 ,S6	R9,S9
32		S0	R5, R13	S6 ,S0	S0
33		S1	R5, R13	S6 ,S1	S1
34		S2	R5, R13	S6 ,S2	S2
35		S8	R5, R13	S6, S8	R13,S8
36		R4	R5,R4	S6	R4
37		R10	R5, R10	S6	R10
38		S3	R5	S6 ,S3	
39		R1	R5 ,R1	S6 ,S3	R1,S3
40		S4	R5	S6 ,S4	R5
41		R2	R2	S6 ,S4	R2,S4
42		R6	R6	S6	S6
43		S7	R6	S7	R6,S7
44		R3	R3		
45		S5	S5		R3,S5

Table 9: Example: FP's method(Buffer size =8)

B.5 Execution Trace for SC

Table 10 shows the behavior of SC's method for computing a join, given the join-index of Figure 10. Table 10 has five columns. The first column shows the iteration number. The second column shows the node swapped out in the current iteration. The third column shows the node selected and brought into the memory buffers. The fourth and fifth columns show the pages of R and S in the main memory buffer. The sixth column shows the nodes which have been processed completely and need not come into the memory buffer again. The last column shows the partition number.

Iteration	Swap out	Add	R Buffer	S Buffer	Done	Partition
1		R0	R0			1
2		R1	R0,R1			1
3		R2	R0,R1,R2			1
4		R3	R0,R1,R2,R3			1
5		R4	R0,R1,R2,R3,R4			1
6		S0	R0,R1,R2,R3,R4	S0		1
7		S1	R0,R1,R2,R3,R4	S0,S1		1
8		S8	R0,R1,R2,R3,R4	S0,S1, S8		1
9	S8	S2	R0,R1,R2,R3,R4	S0,S1,S2	R0	2
10		S3	R1,R2,R3,R4	S0,S1,S2,S3	R1	2
11		S4	R2,R3,R4	S0,S1,S2,S3,S4	R2	2
12		R5	R3,R4,R5	S0,S1,S2,S3,S4	S0	2
13		S5	R3,R4,R5	S1,S2,S3,S4,S5	R3	2
14		R6	R4,R5,R6	S1,S2,S3,S4,S5	S1	2
15		S6	R4,R5,R6	S2,S3,S4,S5,S6	R4	2
16		R12	R5,R6, R12	S2,S3,S4,S5,S6		2
17	R12	R7	R5,R6,R7	S2,S3,S4,S5,S6	S2	3
18		R8	R5,R6,R7,R8	S3,S4,S5,S6	S3	3
19		R9	R5,R6,R7,R8,R9	S4,S5,S6	S4	3
20		R10	R5,R6,R7,R8,R9,R10	S5,S6	S5	3
21		S7	R5,R6,R7,R8,R9,R10	S6,S7		3
22	R5	R11	R6,R7,R8,R9,R10,R11	S6,S7	S6	3
23		S8	R6,R7,R8,R9,R10,R11	S7,S8	R6	3
24		S9	R7,R8,R9,R10,R11	S7,S8,S9	R7	4
25		S10	R8,R9,R10,R11	S7,S8,S9,S10	R8	4
26		S11	R9,R10,R11	S7,S8,S9,S10,S11	R9	4
27		R12	R10,R11,R12	S7,S8,S9,S10,S11	S7	4
28		S12	R10,R11,R12	S8,S9,S10,S11,S12	R10	4
29		R13	R11,R12,R13	S8,S9,S10,S11,S12	S8	4
30		S13	R11,R12,R13	S9,S10,S11,S12,S13	R11,R12,R13	4
31		R5	R5	S9,S10,S11,S12,S13	R5,S9,S10,S11,S12,S13	4

Table 10: Example: SC method(Buffer size =8)