

# Platform-Independent Robust Query Processing

Srinivas Karthik, Jayant R. Haritsa, *Fellow, IEEE*, Sreyash Kenkre, Vinayaka Pandit, and Lohit Krishnan

**Abstract**—To address the classical selectivity estimation problem for OLAP queries in relational databases, a radically different approach called `PlanBouquet` was recently proposed in [1], wherein the estimation process is completely abandoned and replaced with a calibrated discovery mechanism. The beneficial outcome of this new construction is that provable guarantees on worst-case performance, measured as Maximum Sub-Optimality (*MSO*), are obtained thereby facilitating robust query processing. The `PlanBouquet` formulation suffers, however, from a systemic drawback—the *MSO* bound is a function of not only the query, but also the optimizer’s behavioral profile over the underlying database platform. As a result, there are adverse consequences: (i) the bound value becomes highly variable, depending on the specifics of the current operating environment, and (ii) it becomes infeasible to compute the value without substantial investments in preprocessing overheads. In this paper, we first present `SpillBound`, a new query processing algorithm that retains the core strength of the `PlanBouquet` discovery process, but reduces the bound dependency to only the query. It does so by incorporating plan termination and selectivity monitoring mechanisms in the database engine. Specifically, `SpillBound` delivers a worst-case multiplicative bound, of  $D^2 + 3D$ , where  $D$  is simply the number of error-prone predicates in the user query. Consequently, the bound value becomes independent of the optimizer and the database platform, and the guarantee can be issued simply by query inspection. We go on to prove that `SpillBound` is within an  $O(D)$  factor of the *best possible* deterministic selectivity discovery algorithm in its class. We next devise techniques to bridge this quadratic-to-linear *MSO* gap by introducing the notion of *contour alignment*, a characterization of the nature of plan structures along the *boundaries* of the selectivity space. Specifically, we propose a variant of `SpillBound`, called `AlignedBound`, which exploits the alignment property and provides a guarantee in the range  $[2D + 2, D^2 + 3D]$ . Finally, a detailed empirical evaluation over the standard decision-support benchmarks indicates that: (i) `SpillBound` provides markedly superior performance w.r.t. *MSO* as compared to `PlanBouquet`, and (ii) `AlignedBound` provides additional benefits for query instances that are challenging for `SpillBound`, often coming close to the ideal of *MSO* linearity in  $D$ . From an absolute perspective, `AlignedBound` evaluates virtually all the benchmark queries considered in our study with *MSO* of around 10 or lesser. Therefore, in an overall sense, `SpillBound` and `AlignedBound` offer a substantive step forward in the long-standing quest for robust query processing.

**Index Terms**—Selectivity estimation, plan bouquets, robust query processing

## 1 INTRODUCTION

A long-standing problem plaguing database systems is that the predicate selectivity estimates used for optimizing declarative SQL queries are often significantly in error [3], [4]. This results in highly sub-optimal choices of execution plans, and corresponding blowups in query response times. The reasons for such substantial deviations are well documented [5], and include outdated statistics, coarse summaries, attribute-value independence (AVI) assumptions, complex user-defined predicates, and error propagations in the query execution tree. It is therefore of immediate practical relevance to design query processing techniques that limit the deleterious impact of these errors, and thereby provide robust query processing.

We use the notion of Maximum Sub-Optimality (*MSO*), introduced in [1], as a measure of the robustness provided

by a query processing technique to errors in selectivity estimation. Specifically, given a query, the *MSO* of the processing algorithm is the worst-case ratio, over the *entire* selectivity space, of its execution cost with respect to the optimal cost incurred by an oracular system that magically knows the correct selectivities. It has been empirically determined that *MSOs* can reach very large values on current database engines [1]—for instance, with Query 19 of the TPC-DS benchmark, it goes as high as a million!<sup>1</sup> More importantly, worryingly large sub-optimality values are *not rare*—for the same Q19, the sub-optimality values for as many as 40 percent of the locations in the selectivity space are higher than 1,000.

As explained in [1], most of the previous approaches to robust query processing (e.g., [3], [6], [7], [8]), including the influential POP and Rio frameworks, are based on heuristics that are not amenable to *bounded guarantees* on the *MSO* measure. A notable exception to this trend is the `PlanBouquet` algorithm, recently proposed in [1], which provides, for the first time, a provable *MSO* guarantee. Here, the selectivities are not estimated, but instead, systematically *discovered* at run-time through a calibrated sequence of cost-limited executions from a carefully chosen set of plans, called the “plan bouquet”. The search space for the bouquet plans is the *Parametric Optimal Set of Plans* (POSP) [9] over the selectivity

1. Assuming that estimation errors can range over the entire selectivity space.

- S. Karthik and J.R. Haritsa are with the Database Systems Lab, Indian Institute of Science, Bangalore, Karnataka 560012, India. E-mail: {srinivas, haritsa}@dsl.sec.iisc.ernet.in.
- S. Kenkre, V. Pandit, and L. Krishnan are with IBM Research, Bangalore, Karnataka 560045, India. E-mail: {srekenkr, pvinayak, lohith.namboodiri}@in.ibm.com.

Manuscript received 7 Sept. 2016; revised 24 Dec. 2016; accepted 19 Jan. 2017. Date of publication 0 . 0000; date of current version 0 . 0000. Recommended for acceptance by W. Lehner, J. Gehrke, and K. Shim. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier no. 10.1109/TKDE.2017.2664827

```

select * from lineitem, orders, part where
p_partkey = l_partkey and o_orderkey = l_orderkey
and p_retailprice < 1000

```

Fig. 1. Example query (EQ).

66 space. The PlanBouquet technique guarantees  $MSO \leq 4 * |$   
67  $PlanBouquet|$ .<sup>2</sup>

### 68 1.1 PlanBouquet

69 We describe the working of PlanBouquet with the help of  
70 the example query EQ shown in Fig. 1, which enumerates  
71 orders for cheap parts costing less than 1,000. To process  
72 this query, current database engines typically estimate three  
73 selectivities, corresponding to the two join predicates  
74 ( $part \bowtie lineitem$ ) and ( $orders \bowtie lineitem$ ), and the filter  
75 predicate ( $p.retailprice < 1,000$ ). While it is conceivable  
76 that the filter selectivity may be estimated reliably, it is often  
77 difficult to ensure similarly accurate estimates for the join  
78 predicates. We refer to such predicates as error-prone predi-  
79 cates, or epp in short (shown bold-faced in Fig. 1).

#### 80 1.1.1 Example Execution

81 Given the above query, PlanBouquet constructs a two-  
82 dimensional space, called as Error-prone Selectivity Space  
83 (ESS) corresponding to the epps, covering their entire selec-  
84 tivity range  $([0, 1] * [0, 1])$ , as shown in Fig. 2a.

85 On this selectivity space, a series of *iso-cost* contours,  $\mathcal{IC}_1$   
86 through  $\mathcal{IC}_m$ , are drawn—each iso-cost contour  $\mathcal{IC}_i$  has an  
87 associated cost  $CC_i$ , and represents the connected selectivity  
88 curve along which the cost of the optimal plan, as determined  
89 by the optimizer, is equal to  $CC_i$ . Further, the contours are  
90 selected such that the cost of the first contour  $\mathcal{IC}_1$  corresponds  
91 to the minimum query cost  $C$  at the origin of the space, and in  
92 the following intermediate contours, the cost of each contour  
93 is *double* that of the previous contour.<sup>3</sup> That is,  $CC_i = 2^{(i-1)}C$   
94 for  $1 < i < m$ . The last contour's cost,  $CC_m$ , is capped to the  
95 maximum query cost at the top-right corner of the space.

96 As a case in point, in Fig. 2a, there are five hyperbolic-  
97 shaped contours,  $\mathcal{IC}_1$  through  $\mathcal{IC}_5$ , with their costs ranging  
98 from  $C$  to  $16C$ . Each contour has a set of optimal plans cov-  
99 ering disjoint segments of the contour—for instance, con-  
100 tour  $\mathcal{IC}_2$  is covered by plans  $P_2, P_3$  and  $P_4$ .

101 The *union* of the optimal plans appearing on all the  
102 contours constitutes the “plan bouquet”—so, in Fig. 2a,  
103 plans  $P_1$  through  $P_{14}$  form the bouquet. Given this set, the  
104 PlanBouquet algorithm operates as follows: Starting  
105 with the cheapest contour  $\mathcal{IC}_1$ , the plans on each contour  
106 are sequentially executed *with a time limit equal to the con-*  
107 *tour's budget*.

108 If a plan fully completes its execution within the  
109 assigned time limit, then the results are returned to the  
110 user, and the algorithm finishes. Otherwise, as soon as  
111 the time limit of the ongoing execution expires, the plan  
112 is forcibly terminated and the partially computed results  
113 (if any) are discarded. It then moves on to the next plan  
114 in the contour and starts all over again. In the event that  
115 the entire set of plans in a contour have been tried out  
116 without any reaching completion, it *jumps* to the next con-  
117 tour and the cycle repeats.

2. A more precise bound is given later in this section.

3. A doubling factor minimizes the MSO guarantee, as proved in [1].

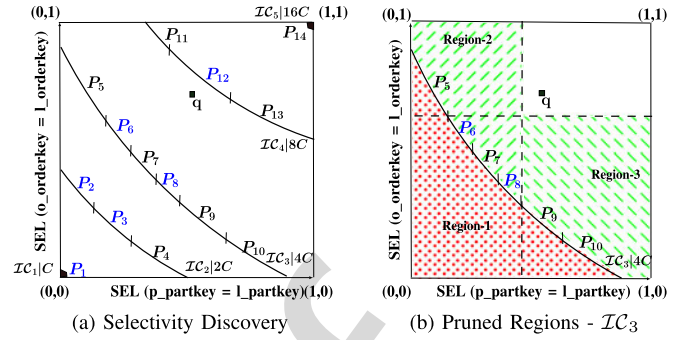


Fig. 2. PlanBouquet and SpillBound.

As a sample instance, consider the case where the query  
is located at  $q$ , in the intermediate region between contours  
 $\mathcal{IC}_3$  and  $\mathcal{IC}_4$ , as shown in Fig. 2a. To process this query,  
PlanBouquet would invoke the following budgeted execu-  
tion sequence:

$$P_1|C, P_2|2C, P_3|2C, P_4|2C, P_5|4C, \dots, P_{10}|4C, P_{11}|8C, P_{12}|8C,$$

with the execution of the final  $P_{12}$  plan completing the  
query.

#### 1.1.2 Performance Guarantees

The overheads entailed by the “trial-and-error” exercise can  
be *bounded, irrespective of the query location in the space*. In par-  
ticular,  $MSO \leq 4 * \rho$ , where  $\rho$  is the plan cardinality on the  
“maximum density” contour. The density of a contour  
refers to the *number* of plans present on it—for instance, in  
Fig. 2a, the maximum density contour is  $\mathcal{IC}_3$  which features  
six plans.

#### 1.1.3 Limitations

The PlanBouquet formulation, while breaking new ground,  
suffers from a systemic drawback—the specific value of  $\rho$ ,  
and therefore the bound, is a function of not only the query,  
but also the optimizer’s behavioral profile over the underlying  
database platform (including data contents, physical schema,  
hardware configuration, etc.). As a result, there are adverse  
consequences: (i) The bound value becomes highly variable,  
depending on the specifics of the current operating environ-  
ment—for instance, with TPC-DS Query 25, PlanBouquet’s  
MSO guarantee of 24 under PostgreSQL shot up, under an  
identical computing environment, to 36 for a commercial  
engine, due to the change in  $\rho$ ; (ii) It becomes infeasible to  
compute the value without substantial investments in prepro-  
cessing overheads; and (iii) Ensuring a bound that is small  
enough to be of practical value, is contingent on the heuristic  
of “anorexic reduction” [10] holding true.

## 1.2 SpillBound

Our objective here is to develop a robust query processing  
approach that offers an MSO bound which is *solely query-*  
*dependent*, irrespective of the underlying database platform.  
That is, we desire a “structural bound” instead of a  
“behavioral bound”. Accordingly, we present a new query  
processing algorithm, called SpillBound, that achieves  
this objective in the sense that it delivers an MSO bound  
that is only a function of  $D$ , the *number* of predicates in the  
query that are prone to selectivity estimation errors. More-  
over, the dependency is in the form of a low-order

polynomial, with MSO expressed as  $(D^2 + 3D)$ . Consequently, the bound value becomes: (i) independent of the database platform,<sup>4</sup> (ii) known upfront by merely inspecting the query, and not incurring any preprocessing overhead, (iii) indifferent to the anorexic reduction heuristic, and (iv) certifiably low in value for practical values of  $D$ .

### 1.2.1 Example Execution

SpillBound shares the core contour-wise discovery approach of PlanBouquet, but its execution strategy differs markedly. Specifically, it achieves a significant reduction in the cost of the sequence of budgeted executions employed during the selectivity discovery process. For instance, in the example scenario of Fig. 2a, the sequence of budgeted executions correspond to the plans highlighted in blue

$$P_1|C, P_2|2C, P_3|2C, P_6|4C, P_8|4C, P_{12}|8C,$$

with  $P_{12}$  again completing the query. Note that the reduced executions result in cost savings of more than 50 percent over PlanBouquet.

The advantages offered by SpillBound are achieved by the following key properties—Half-space Pruning and Contour Density Independent (CDI) execution—of the algorithm.

### 1.2.2 Half-Space Pruning

With each contour whose plans do not complete within the assigned budget, PlanBouquet is able to prune the corresponding *hypograph*—that is, the search region *below* the contour curve.

A pictorial view is shown in Fig. 2b, which focuses on contour  $\mathcal{IC}_3$ —here, the hypograph of  $\mathcal{IC}_3$  is the Region-1 marked with red dots.

However, with SpillBound, a much stronger *half-space*-based pruning comes into play. This is vividly highlighted in Fig. 2b, where the half-space corresponding to Region-2 is pruned by the (budget-limited) execution of  $P_8$ , while the half-space corresponding to Region-3 is pruned by the (budget-limited) execution of  $P_6$ . Note that Region-2 and Region-3 together subsume the entire Region-1 that is covered by PlanBouquet when it crosses  $\mathcal{IC}_3$ . Our half-space pruning property is achieved by leveraging the notion of “*spilling*”, whereby operator pipelines in the execution plan tree are *prematurely terminated* at chosen locations, in conjunction with *run-time monitoring* of operator selectivities.

### 1.2.3 Contour Density Independent Execution

Let us define a “quantum progress” to be a step in which the algorithm either (a) jumps to the next contour, or (b) fully learns the selectivity of some epp (thus reducing the effective number of epps). Then, in the example scenario, while advancing through the various contours in the discovery process, SpillBound makes quantum progress by executing at most *two plans* on each contour. In general, when there are  $D$  error-prone predicates in the user query, SpillBound is guaranteed to make quantum progress based on cost-budgeted execution of at most  $D$  carefully chosen plans on the contour.

Specifically, in each contour, for each dimension, one plan is chosen for spill-mode execution. The plan chosen for

4. Under the assumption that  $D$  remains constant across the platforms.

spill-mode execution is the one that provides the *maximal* guaranteed learning of the selectivity along that dimension. In our example,  $P_8$  and  $P_6$  are the plans chosen for the contour  $\mathcal{IC}_3$  along the  $X$  and  $Y$  dimensions, respectively.

## 1.3 Bridging the MSO Gap

At this juncture, a natural question to ask is whether some alternative selectivity discovery algorithm, based on half-space pruning, can provide better MSO bounds than SpillBound. In this regard, we prove that *no* deterministic technique in this class can provide an MSO bound less than  $D$ . Therefore, the SpillBound guarantee is no worse than a factor  $O(D)$  as compared to the best possible algorithm in its class.

### 1.3.1 Contour Alignment

Given this quadratic-to-linear gap on the MSO guarantee, we seek to characterize exploration scenarios in which SpillBound’s MSO approaches the lower bound. For this purpose, we introduce a new concept called *contour alignment*—a contour is aligned if the contour plan that is incident on the boundary of the ESS, has its selectivity learning dimension (during spill-mode execution) matching with the incident dimension. For instance, in the example of Fig. 2, contour  $\mathcal{IC}_3$  would be aligned if plan  $P_5$ , rather than  $P_6$ , happened to be the plan providing the maximal guaranteed learning along the  $Y$  dimension. Leveraging this notion, we show that the MSO bound can be reduced to  $O(D)$  if the contour alignment property is satisfied at *every contour* encountered during its execution.

Unfortunately, in practice, we may not always find the alignment property satisfied at all contours. Therefore, we design the AlignedBound algorithm which extracts the benefit of alignment wherever available, either natively or through an explicit induction. Specifically, AlignedBound delivers an MSO that is guaranteed to be in the platform-independent range  $[2D + 2, D^2 + 3D]$ .

## 1.4 Empirical Results

The bounds delivered by PlanBouquet and SpillBound are, in principle, *uncomparable*, due to the inherently different nature of their parametric dependencies. However, in order to assess whether the platform-independent feature of SpillBound is procured through a deterioration of the numerical bound, we have carried out a detailed experimental evaluation of both the approaches on standard benchmark queries, operating on the PostgreSQL engine. Moreover, we have empirically evaluated the MSO obtained for each query through an exhaustive enumeration of the selectivity space.

Our experiments indicate that for the most part, SpillBound provides similar guarantees to PlanBouquet, and occasionally, much tighter bounds. As a case in point, for TPC-DS Query 91 with six error-prone predicates, the MSO bound is 96 with PlanBouquet, but comes down to 54 with SpillBound. More pertinently, the *empirical* MSO of SpillBound is significantly better than that of PlanBouquet for *all* the queries. For instance, the empirical MSO for Q91 decreases from PlanBouquet’s 49 to 19 for SpillBound.

Turning our attention to AlignedBound, its performance is typically closer to the *lower end* of its guarantee range, i.e.,  $2D + 2$ , and often provides substantial benefits for query instances that are challenging for SpillBound. For instance,

279 AlignedBound brings the MSO of the above-mentioned  
280 Q91 test case down to **10.4**. Moreover, AlignedBound is  
281 able to complete virtually all the benchmark queries evaluat-  
282 ed in our study with a MSO of around 10 or lower.

283 In a nutshell, AlignedBound consistently collapses  
284 the enormous MSOs incurred with contemporary indus-  
285 trial-strength query optimizers, down to a single order of  
286 magnitude.

#### 287 1.4.1 Caveats

288 We hasten to add that our proposed algorithms are *not* a  
289 substitute for a conventional query optimizer. Instead, they  
290 are intended to complementarily *co-exist* with the traditional  
291 setup, leaving to the user's discretion, the specific approach  
292 to employ for a query instance. When small estimation  
293 errors are expected, the native optimizer could be sufficient,  
294 but if larger errors are anticipated, our algorithms are likely  
295 to be the preferred choice.

#### 296 1.4.2 Organization

297 The remainder of this paper is organized as follows:  
298 In Section 2, a precise description of the robust execution  
299 problem is provided, along with the associated notations.  
300 The building blocks of our algorithms are presented in  
301 Section 3.1. The SpillBound algorithm and the proof of its  
302 MSO bound are presented in Section 4, followed by the  
303 AlignedBound algorithm and its analysis in Section 5.  
304 The experimental framework and performance results are  
305 enumerated in Section 6, while pragmatic deployment  
306 aspects are discussed in Section 7. The related literature is  
307 reviewed in Section 8, and our conclusions are summarized  
308 in Section 9.

## 309 2 PROBLEM FRAMEWORK

310 In this section, we present the key concepts, notations, and  
311 the formal problem definition. For ease of presentation, we  
312 assume that the error-prone selectivity predicates (epps) for  
313 a given user query are known apriori, and defer the issue of  
314 identifying these epps to Section 7.

### 315 2.1 Error-Prone Selectivity Space

316 Consider a query with  $D$  epps. The set of all epps is denoted  
317 by  $EPP = \{e_1, \dots, e_D\}$  where  $e_j$  denotes the  $j$ th epp. The  
318 selectivities of the  $D$  epps are mapped to a  $D$ -dimensional  
319 space, with the selectivity of  $e_j$  corresponding to the  $j$ th  
320 dimension. Since the selectivity of each predicate ranges  
321 over  $[0, 1]$ , a  $D$ -dimensional hypercube  $[0, 1]^D$  results, hence-  
322 forth referred to as the *error-prone selectivity space*, or ESS. In  
323 practice, an appropriately discretized grid version of  $[0, 1]^D$   
324 is considered as the ESS. Note that each location  $q \in [0, 1]^D$   
325 in the ESS represents a specific instance where the epps of  
326 the user query happen to have selectivities corresponding  
327 to  $q$ . Accordingly, the selectivity value on the  $j$ th dimension  
328 is denoted by  $q.j$ . We call the location at which the selectiv-  
329 ity value in each dimension is 1, i.e.  $q.j = 1, \forall j$ , as the  
330 *terminus*.

331 The notion of a location  $q_1$  *dominating* a location  $q_2$  in the  
332 ESS plays a central role in our framework. Formally, given  
333 two distinct locations  $q_1, q_2 \in ESS$ ,  $q_1$  dominates  $q_2$ , denoted  
334 by  $q_1 \succeq q_2$ , if  $q_1.j \geq q_2.j$  for all  $j \in 1, \dots, D$ . In an analogous  
335 fashion, other relations, such as  $\not\succeq$ ,  $\preceq$ , and  $\not\preceq$  can be defined  
336 to capture relative positions of pairs of locations.

### 2.2 Search Space for Robust Query Processing

337 We assume that the query optimizer can identify the *optimal*  
338 query execution plan if the selectivities of all the epps  
339 are correctly known.<sup>5</sup> Therefore, given an input query and  
340 its epps, the optimal plans for *all* locations in the ESS grid  
341 can be identified through repeated invocations of the opti-  
342 mizer with different selectivity values. The optimal plan for  
343 a generic selectivity location  $q \in ESS$  is denoted by  $P_q$ , and  
344 the set of such optimal plans over the complete ESS consti-  
345 tutes the Parametric Optimal Set of Plans [9].<sup>6</sup> 346

347 We denote the cost of executing an *arbitrary* plan  $P$  at a  
348 selectivity location  $q \in ESS$  by  $Cost(P, q)$ . Thus,  $Cost(P_q, q)$   
349 represents the *optimal* execution cost for the selectivity  
350 instance located at  $q$ . In this framework, our search space  
351 for robust query processing is simply the set of tuples  
352  $\langle q, P_q, Cost(P_q, q) \rangle$  corresponding to all locations  $q \in ESS$ .

353 Throughout the paper, we adopt the convention of using  
354  $q_a$  to denote the actual selectivities of the user query epps—  
355 note that this location is unknown at compile-time, and  
356 needs to be explicitly discovered. For traditional optimizers,  
357 we use  $q_e$  to denote the *estimated* selectivity location based on  
358 which the execution plan  $P_{q_e}$  is chosen to execute the query.  
359 However, this characterization is not applicable to plan  
360 switching approaches like PlanBouquet and SpillBound  
361 because they explore a *sequence* of locations during their dis-  
362 covery process. So, we denote the deterministic sequence  
363 pursued for a query instance corresponding to  $q_a$  by  $Seq_{q_a}$ .

### 2.3 Maximum Sub-Optimality [1]

364 We now present the performance metrics proposed in [1] to  
365 quantify the robustness of query processing. 366

367 A traditional query optimizer will first estimate  $q_e$ , and  
368 then use  $P_{q_e}$  to execute a query which may actually be  
369 located at  $q_a$ . The sub-optimality of this plan choice, relative  
370 to an oracle that magically knows the correct location, and  
371 therefore uses the ideal plan  $P_{q_a}$ , is defined as

$$372 SubOpt(q_e, q_a) = \frac{Cost(P_{q_e}, q_a)}{Cost(P_{q_a}, q_a)}. \quad (1) \quad 373$$

374 The quantity  $SubOpt(q_e, q_a)$  ranges over  $[1, \infty)$ .

375 With this characterization of a specific  $(q_e, q_a)$  combina-  
376 tion, the *maximum* sub-optimality that can potentially arise  
377 over the entire ESS is given by

$$378 MSO = \max_{(q_e, q_a) \in ESS} (SubOpt(q_e, q_a)). \quad (2) \quad 379$$

380 The above definition for a traditional optimizer can be gen-  
381 eralized to selectivity discovery algorithms like PlanBou-  
382 quet and SpillBound. Specifically, suppose the discovery  
383 algorithm is currently exploring a location  $q \in Seq_{q_a}$ —it will  
384 choose  $P_q$  as the plan and  $Cost(P_q, q)$  as the associated budget.  
385 Extending this to the whole sequence, the analogue of Equa-  
386 tion (1) is defined as follows: 387

$$388 SubOpt(Seq_{q_a}, q_a) = \frac{\sum_{q \in Seq_{q_a}} Cost(P_q, q)}{Cost(P_{q_a}, q_a)}, \quad (3) \quad 389$$

5. For example, through the classical DP-based search of the plan space [11].

6. Letter subscripts for plans denote locations, whereas numeric subscripts denote identifiers.




Fig. 3. 3D cost surface on ESS.

TABLE 1  
Notations

Notation	Meaning
epp (EPP)	Error-prone predicate (its collection)
ESS	Error-prone selectivity space
$D$	Number of dimensions of ESS
$e_1, \dots, e_D$	The $D$ epps in the query
$q \in [0, 1]^D$	A location in the ESS space
$q.j$	Selectivity of $q$ in the $j$ th dimension of ESS
$P_q$	Optimal Plan at $q \in$ ESS
$q_a$	Actual run-time selectivity
$Cost(P, q)$	Cost of plan $P$ at location $q$
$\mathcal{IC}_i$	Isocost Contour $i$
$CC_i$	Cost of an isocost contour $\mathcal{IC}_i$
$PL_i$	Set of plans on contour $\mathcal{IC}_i$

390 leading to

$$392 \quad MSO = \max_{q_a \in \text{ESS}} \text{SubOpt}(\text{Seq}_{q_a}, q_a). \quad (4)$$

## 394 2.4 Problem Definition

395 With the above framework, the problem of robust query  
396 processing is defined as follows:

397 For a given input query  $Q$  with its EPP, and the search space  
398 consisting of tuples  $\langle q, P_q, \text{Cost}(P_q, q) \rangle$  for all  $q \in$  ESS, develop  
399 a query processing approach that minimizes the MSO guarantee.

400 As in [1], the primary assumptions made in this paper  
401 that allow for systematic construction and exploration of  
402 the ESS are those of *plan cost monotonicity* (PCM) and *selectiv-*  
403 *ity independence* (SI). PCM may be stated as: For any two  
404 locations  $q_b, q_c \in$  ESS, and for any plan  $P$ ,

$$405 \quad q_b \succ q_c \Rightarrow \text{Cost}(P, q_b) > \text{Cost}(P, q_c). \quad (5)$$

406 That is, it encodes the intuitive notion that when more data  
407 is processed by a query, signified by the larger selectivities  
408 for the predicates, the cost of the query processing also  
409 increases. On the other hand, SI assumes that the selectiv-  
410 ities of the epps are all independent—while this is a com-  
411 mon assumption in much of the query optimization  
412 literature, it often does not hold in practice. In our future  
413 work, we intend to extend *SpillBound* to handle the more  
414 general case of dependent selectivities.

## 416 2.5 Geometric View and Notations

417 We now present a geometric view of the discovery space  
418 and some important notations. Consider the special case of  
419 a query with two epps, resulting in an ESS with  $X$  and  $Y$   
420 dimensions. Now, incorporate a third  $Z$  dimension to cap-  
421 ture the *cost* of the optimal plan on the ESS, i.e, for  $q \in$  ESS,  
422 the value of the  $Z$ -axis is  $\text{Cost}(P_q, q)$ . This 3D surface, which  
423 captures the cost of the optimal plan on the ESS, is called the  
424 *Optimal Cost Surface* (OCS). Associated with each point on  
425 the OCS is the POSP plan for the underlying location in the  
426 ESS. A sample OCS corresponding to the example query  $EQ$   
427 in the Introduction is shown in Fig. 3, which provides a per-  
428 spective view of this surface. In this figure, the optimality  
429 region of each POSP plan is denoted by a unique color. So,  
430 for example, the region with blue points corresponds to  
431 those locations where the “blue plan” is the optimal plan.<sup>7</sup>

7. Since Fig. 3 is only a perspective view of the OCS, it does not capture all the POSP plans.

432 *Discretization of OCS*: Let  $C_{min}$  and  $C_{max}$  denote the mini- 432  
433 mum and maximum costs on the OCS, corresponding to the  
434 origin and the terminus of the 3D space, respectively  
435 (an outcome of the PCM assumption). We define  
436  $m = \lceil \log_2(\frac{C_{max}}{C_{min}}) \rceil + 1$  hyperplanes that are parallel to the  $XY$  436  
437 plane as follows. The first hyperplane is drawn at  $C_{min}$ . For  
438  $i = 2, \dots, m - 1$ , the  $i$ th hyperplane is drawn at  $C_{min} \cdot 2^{i-1}$ .  
439 The last hyperplane is drawn at  $C_{max}$ . These hyperplanes cor- 439  
440 respond to the  $m$  isocost contours  $\mathcal{IC}_1, \dots, \mathcal{IC}_m$ . The isocost  
441 contour  $\mathcal{IC}_i$  is essentially the 2D curve obtained by intersect- 441  
442 ing the OCS with the  $i$ th hyperplane. We denote the cost of  
443  $\mathcal{IC}_i$  by  $CC_i$ . The set of plans that are on the 2D curve of  $\mathcal{IC}_i$  443  
444 are referred to as  $PL_i$ . For example, in Fig. 3,  $PL_4$  includes the pur- 444  
445 ple and maroon plans (in addition to plans that are not visible  
446 in this perspective). The *hypograph* of an isocost contour  $\mathcal{IC}_i$  is 446  
447 the set of all locations  $q \in$  ESS such that  $\text{Cost}(P_q, q) \leq CC_i$ .

448 The above geometric intuition and the formal notations  
449 readily extend to the general case of  $D$  epps, and these nota- 449  
450 tions are summarized in Table 1 for easy reference.

## 451 3 BUILDING BLOCKS OF OUR ALGORITHMS

452 The platform-independent nature of the MSO bound of the  
453 *SpillBound* is enabled by the key properties of half-space  
454 pruning and contour density independent execution. The  
455 *AlignedBound* algorithm that provides an  $O(D)$  MSO under  
456 certain special scenarios is based on the concept of contour  
457 alignment. In this section, we present these building blocks of  
458 the *SpillBound* and *AlignedBound* algorithms.

### 459 3.1 Half-Space Pruning

460 Half-space pruning is the ability to prune half-spaces from  
461 the search space based on a single cost-budgeted execution  
462 of a contour plan. We now present how half-space pruning  
463 is achieved by using *spilling* during execution of query  
464 plans. While the use of *spilling* to accelerate selectivity dis-  
465 covery had been mooted in [1], they did not consider its  
466 exploitation for obtaining guaranteed search properties.

467 We use *spilling* as the mechanism for modifying the exe-  
468 cution of a selected plan—the objective here is to utilize the  
469 assigned execution budget to extract increased selectivity  
470 information of a specific epp. Since *spilling* requires modifi-  
471 cation of plan executions, we shall first describe the query  
472 execution model.

#### 473 3.1.1 Execution Model

474 We assume the demand driven iterator model, commonly  
475 seen in database engines, for the execution of operators in the

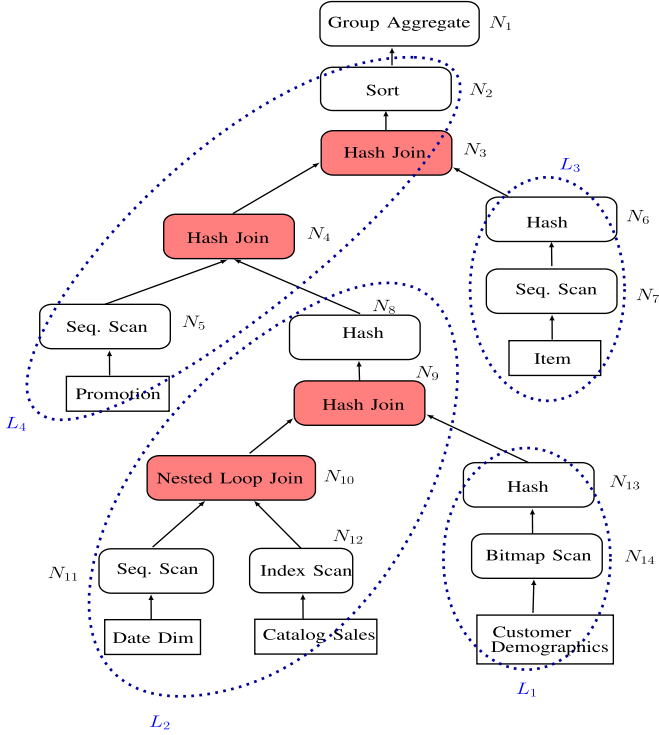


Fig. 4. Execution plan tree of TPC-DS Query 26.

476 plan tree [12]. Specifically, the execution takes place in a bot-  
 477 tom up fashion with the base relations at the leaves of the tree.

478 In conventional database query processing, the execution  
 479 of a query plan can be partitioned into a sequence of *pipe-*  
 480 *lines* [13]. Intuitively, a pipeline can be defined as the maxi-  
 481 mal concurrently executing subtree of the execution plan.  
 482 The entire execution plan can therefore be viewed as an  
 483 ordering on its constituent pipelines. We assume that only  
 484 one pipeline is executed at a time in the database system, i.  
 485 e, there is no inter-pipeline concurrency—this appears to be  
 486 the case in current engines. To make these notions concrete,  
 487 consider the plan tree shown in Fig. 4—here, the constituent  
 488 pipelines are highlighted with ovals, and are executed in  
 489 the sequence  $\{L_1, L_2, L_3, L_4\}$ .

490 Finally, we assume a standard plan costing model that  
 491 estimates the individual costs of the internal nodes, and  
 492 then aggregates the costs of all internal nodes to represent  
 493 the estimated cost of the complete plan tree.

### 494 3.1.2 Spill-Mode of Execution

495 We now discuss how to execute plans in spill-mode. For  
 496 expository convenience, given an internal node of the plan  
 497 tree, we refer to the set of nodes that are in the subtree  
 498 rooted at the node as its *upstream* nodes, and the set of nodes  
 499 on its path to the root of the complete plan tree as its *down-*  
 500 *stream* nodes.

501 Suppose we are interested in learning about the selectiv-  
 502 ity of an epp  $e_j$ . Let the internal node corresponding to  $e_j$  in  
 503 plan  $P$  be  $N_j$ . The key observation here is that the execution  
 504 cost incurred on  $N_j$ 's downstream nodes in  $P$  is *not useful*  
 505 for learning about  $N_j$ 's selectivity. So, discarding the output  
 506 of  $N_j$  without forwarding to its downstream nodes, and  
 507 devoting the entire budget to the subtree rooted at  $N_j$ , helps  
 508 to use the budget effectively to learn  $e_j$ 's selectivity. Specifi-  
 509 cally, given plan  $P$  with cost budget  $B$ , and epp  $e_j$  chosen  
 510 for spilling, the spill-mode execution of  $P$  is simply the

511 following: Create a modified plan comprised of only the  
 512 subtree of  $P$  rooted at  $N_j$ , and execute it with cost budget  $B$ .

513 Since a plan could consist of multiple epps (red coloured  
 514 nodes in Fig. 4), the sequence of spill node choices should be  
 515 made carefully to ensure guaranteed learning on the selectiv-  
 516 ity of the chosen node—this procedure is described next.

### 517 3.1.3 Spill Node Identification

518 Given a plan and an ordering of the pipelines in the plan,  
 519 we consider an ordering of epps based on the following two  
 520 rules:

521 *Inter-Pipeline Ordering*: Order the epps as per the execution  
 522 order of their respective pipelines; in Fig. 4, since  $L_4$  is  
 523 ordered after  $L_2$ , the epp nodes  $N_3$  and  $N_4$  are ordered  
 524 after  $N_9$  and  $N_{10}$ .

525 *Intra-Pipeline Ordering*: Order the epps by their upstream-  
 526 downstream relationship, i.e., if an epp node  $N_a$  is  
 527 downstream of another epp node  $N_b$  within the same  
 528 pipeline, then  $N_a$  is ordered after  $N_b$ ; in the example,  $N_3$   
 529 is ordered after  $N_4$ .

530 It is easy to see that the above rules produce a total-order-  
 531 ing on the epps in a plan—in Fig. 4, it is  $N_{10}, N_9, N_4, N_3$ .  
 532 Given this ordering, we always choose to spill on the node  
 533 corresponding to the *first* epp in the total-order. The selectiv-  
 534 ity of a spilled epp node is fully learnt when the correspond-  
 535 ing execution goes to completion within its assigned  
 536 budget. When this happens, we remove the epp from EPP  
 537 and it is no longer considered as a candidate for spilling in  
 538 the rest of the discovery process.

539 As a result of this procedure, note that the selectivities of  
 540 all predicates located *upstream* of the currently spilling epp  
 541 will be known *exactly*—either because they were never epps,  
 542 or because they have already been fully learnt in the ongo-  
 543 ing discovery process. Therefore, their cost estimates are  
 544 accurate, leading to the following “half-space pruning”  
 545 lemma. The proof of the lemma can be seen in [2].

546 **Lemma 3.1.** Consider a plan  $P$  for which the spill node identi-  
 547 fication mechanism identifies the predicate  $e_j$  for spilling. Fur-  
 548 ther, consider a location  $q \in \text{ESS}$ . When the plan  $P$  is executed  
 549 with a budget  $\text{Cost}(P, q)$  in spill-mode, then we either learn (a)  
 550 the exact selectivity of  $e_j$ , or (b) that  $q_{a.j} > q.j$ .

### 551 3.2 Contour Density Independent Execution

552 We now show how the half-space pruning property can be  
 553 exploited to achieve the contour density independent execu-  
 554 tion property of the SpillBound algorithm. For this pur-  
 555 pose, we employ the term “quantum progress” to refer to a  
 556 step in which the algorithm either jumps to the next con-  
 557 tour, or fully discovers the selectivity of some epp. Informally,  
 558 the CDI property ensures that each quantum  
 559 progress in the discovery process is achieved by expending  
 560 no more than  $|\text{EPP}|$  number of plan executions.

561 For ease of understanding, we present here the technique  
 562 for the special case of two epps referred to by  $X$  and  $Y$ ,  
 563 deferring the generalization for  $D$  epps to the next section.

564 Consider the 2D ESS shown in Fig. 5, and assume that we  
 565 are currently exploring contour  $\mathcal{IC}_3$ . The two plans for spill-  
 566 mode execution in this contour are identified as follows: We  
 567 first identify the subset of plans on the contour that spill on  
 568  $X$  using the spill node identification algorithm—these plans  
 569 are identified as  $P_5^x, P_7^x, P_8^x$  in Fig. 5. The next step is to




Fig. 5. Choice of contour crossing plans.




Fig. 6. Contour alignment.

570 enumerate the subset of locations on the contour where  
 571 these  $X$ -spilling plans are optimal. From this subset, we  
 572 identify the location with the *maximum*  $X$  coordinate,  
 573 referred to as  $q_{max}^x$ , and its corresponding contour plan,  
 574 which is denoted as  $P_{max}^x$ . The  $P_{max}^x$  plan is the one chosen  
 575 to learn the selectivity of  $X$ —in Fig. 5, this choice is  $P_8^x$ .

576 By repeating the same process for the  $Y$  dimension, we  
 577 identify the location  $q_{max}^y$ , and plan  $P_{max}^y$ , for learning the  
 578 selectivity of  $Y$ —in Fig. 5, the plan choice is  $P_6^y$ . Note that  
 579 the location  $(q_{max}^x, q_{max}^y)$  is guaranteed to be either on or  
 580 beyond the  $IC_3$  contour.

581 The following lemma shows that the above plan identifi-  
 582 cation procedure satisfies the CDI property.

583 **Lemma 3.2.** *In contour  $IC_i$ , if plans  $P_{max}^x$  and  $P_{max}^y$  are executed*  
 584 *in spill-mode, and both do not reach completion, then  $Cost(P_{q_a},$*   
 585  *$q_a) > CC_i$ , triggering a jump to the next contour  $IC_{i+1}$ .*

586 **Proof.** Since the executions of both  $P_{max}^x$  and  $P_{max}^y$  do not  
 587 reach completion, we infer that  $q_{max}^x \cdot x < q_a \cdot x$  and  
 588  $q_{max}^y \cdot y < q_a \cdot y$ . Therefore,  $q_a$  strictly dominates the loca-  
 589 tion  $(q_{max}^x, q_{max}^y)$  whose cost, by PCM, is greater than  
 590  $CC_i$ . Thus  $Cost(P_{q_a}, q_a) > CC_i$ .  $\square$

591 Consider the general case of  $IC_i$  when there are more  
 592 than two epps. Corresponding to an epp  $e_j$ , the location  
 593  $q_{max}^j$  and plan  $P_{max}^j$  are defined similar to the way  $q_{max}^x$  and  
 594  $P_{max}^x$  are defined (i.e, by replacing the  $X$  coordinate with  
 595 the  $j$ th coordinate corresponding to  $e_j$ ).

### 596 3.3 Contour Alignment

597 We now introduce a key concept that helps characterize  
 598 search scenarios in which the MSO of the SpillBound  
 599 algorithm matches the lower bound. Again, for ease of  
 600 understanding, we consider the special case of a 2D ESS  
 601 with predicates  $X$  and  $Y$ .

602 Consider a contour, say  $IC_i$ , and a dimension  $j \in \{X, Y\}$ .  
 603 A location  $q_{ext}^j \in IC_i$  is said to be an *extreme location along*  
 604 *dimension  $j$*  if the location has the maximum coordinate  
 605 value for dimension  $j$  among the contour locations belong-  
 606 ing to  $IC_i$ , i.e,  $q_{ext}^j \cdot j \geq q \cdot j, \forall q \in IC_i$ . In Fig. 6, these extreme  
 607 locations are highlighted by (bold) dots.

608 A contour  $IC_i$  is said to satisfy the property of contour  
 609 alignment along a dimension  $j$  if it so happens that  
 610  $q_{max}^j = q_{ext}^j$ , i.e., the optimal plan at  $q_{ext}^j$  spills on predicate  
 611  $e_j$ . For ease of exposition, if a contour satisfies the contour  
 612 alignment property along at least one of its dimensions,  
 613 then we refer to it as an *aligned contour*. In Fig. 6, contours  
 614  $IC_2$  and  $IC_4$  are aligned along the  $X$  and  $Y$  dimensions,  
 615 respectively, and are therefore aligned contours—however,  
 616 contour  $IC_3$  is not so because it is not aligned along  
 617 either dimension.

Given a contour  $IC_i$ , Lemma 3.2 showed the sufficiency 618  
 of *two* plan executions to guarantee a quantum progress in 619  
 the discovery process. Leveraging the alignment notion, the 620  
 following lemma describes when the same progress can be 621  
 achieved with exactly *one* execution. 622

623 **Lemma 3.3.** *If a contour  $IC_i$  is aligned, then the execution of*  
 624 *exactly one plan in spill-mode with budget  $CC_i$ , is sufficient to*  
 625 *make quantum progress in the discovery process.*

626 **Proof.** Without loss of generality, let us assume that the  
 627 contour  $IC_i$  satisfies contour alignment along dimension  
 628  $j$ , i.e, the optimal plan  $P$  at the location  $q_{ext}^j$  spills on  
 629 dimension  $j$ . By Lemma 3.1, the spill-mode execution of  
 630  $P$  with budget  $CC_i$  ensures that we either learn the exact  
 631 selectivity of  $e_j$  or learn that  $q_a \cdot j > q_{ext}^j \cdot j$ . Suppose we  
 632 learn that  $q_a \cdot j > q_{ext}^j \cdot j$ , then it implies that  $q_a$  lies beyond  
 633  $IC_i$ . Thus, just the execution of  $P$  in spill-mode yields  
 634 quantum progress.  $\square$

Note that in the general ESS case of more than two 635  
 epps, there may be a *multiplicity* of  $q_{max}^j$  or  $q_{ext}^j$  locations, 636  
 but Lemma 3.3 can be easily generalized such that quan- 637  
 tum progress is achieved with a single execution in these 638  
 scenarios also. 639

## 640 4 THE SPILLBOUND ALGORITHM

In this section, we present our new robust query processing 641  
 algorithm, SpillBound, which leverages the properties of 642  
 half-space pruning and CDI execution. We begin by introduc- 643  
 ing an important notation: Our search for the actual query 644  
 location,  $q_a$ , begins at the origin, and with each spill-mode 645  
 execution of a contour plan, we monotonically move closer 646  
 towards the actual location. The running selectivity location, 647  
 as progressively learnt by SpillBound, is denoted by  $q_{run}$ . 648

649 During the entire discovery process of SpillBound,  
 only POSP plans on the isocost contour are considered for  
 spill-mode executions. Moreover, when we mention the  
 spill-mode execution of a particular plan on a contour, it  
 implicitly means that the budget assigned is equal to the  
 cost of the contour. For ease of exposition, if the epp chosen  
 to spill on is  $e_j$  for a plan  $P$ , we shall hereafter highlight this  
 information with the notation  $P^j$ . 656

657 For ease of exposition, we first present a version, called  
 2D-SpillBound, for the special case of two epps, and then  
 extend the algorithm to the general case of several epps. 659

### 660 4.1 2D-SpillBound

To provide a geometric insight into the working of 2D- 661  
 SpillBound, we will refer to the two epps,  $e_1$  and  $e_2$ , as  $X$  662  
 and  $Y$ , respectively. 2D-SpillBound explores the doubling 663

664 isocost contours  $\mathcal{IC}_1, \dots, \mathcal{IC}_m$ , starting with the minimum  
 665 cost contour  $\mathcal{IC}_1$ . During the exploration of a contour, two  
 666 plans  $P_{max}^x$  and  $P_{max}^y$  are identified, as described in Section  
 667 3.2, and executed in spill-mode. The order of execution  
 668 between these two plans can be chosen arbitrarily, and the  
 669 selectivity information learnt through their execution is used  
 670 to update the running location  $q_{run}$ . This process continues  
 671 until one of the spill-mode executions reaches completion,  
 672 which implies that the selectivity of the corresponding epp  
 673 has been completely learnt.

674 Without loss of generality, assume that the learnt selectivity  
 675 is  $X$ . At this stage, we know that  $q_a$  lies on the line  
 676  $\bar{X} = q_a.x$ . Further, the discovery problem is reduced to the  
 677 1D case, which has a unique characteristic—each isocost con-  
 678 tour of the new ESS (i.e., line  $X = q_a.x$ ) contains only *one*  
 679 plan, and this plan alone needs to be executed to cross the  
 680 contour, until eventually some plan finishes its execution  
 681 within the assigned budget. In this special 1D scenario, there  
 682 is no operational difference between PlanBouquet and  
 683 2D-SpillBound, so we simply invoke the standard Plan-  
 684 Bouquet with only the  $Y$  epp, starting from the contour cur-  
 685 rently being explored. Note that plans are *not* executed in  
 686 spill-mode in this terminal 1D phase because spilling in the  
 687 1D case weakens the bound, as explained in [14].

#### 688 4.1.1 Execution Trace

689 An illustration of the execution of 2D-SpillBound on  
 690 TPC-DS Query 91 with two epps is shown in Fig. 7. In this  
 691 example, the join predicate *Catalog Sales*  $\bowtie$  *Date Dim*,  
 692 denoted by  $X$ , and the join predicate *Customer*  $\bowtie$  *Customer*  
 693 *Address*, denoted by  $Y$ , are the two epps (both selectivities  
 694 are shown on a log scale).

695 We observe here that there are six doubling isocost con-  
 696 tours  $\mathcal{IC}_1, \dots, \mathcal{IC}_6$ . The execution trace of 2D-SpillBound  
 697 (blue line) corresponds to the selectivity scenario where the  
 698 user's query is located at  $q_a = (0.04, 0.1)$ .

699 On each contour, the plans executed by 2D-SpillBound  
 700 in spill-mode are marked in blue—for example, on  $\mathcal{IC}_2$ , plan  
 701  $P_4$  is executed in spill-mode for the epp  $Y$ . Further, upon each  
 702 execution of a plan, an axis-parallel line is drawn from the pre-  
 703 vious  $q_{run}$  to the newly discovered  $q_{run}$ , leading to the Manhat-  
 704 tan profile shown in Fig. 7. For example, when plan  $P_6$  is  
 705 executed in spill-mode for  $X$ , the  $q_{run}$  moves from  $(2E-4, 6E-4)$   
 706 to  $(8E-4, 6E-4)$ . To make the execution sequence unambigu-  
 707 ously clear, the trace joining successive  $q_{run}$ s is also annotat-  
 708 ed with the plan execution responsible for the move—to high-  
 709 light the spill-mode execution, we use  $p_i$  to denote the spilled  
 710 execution of  $P_i$ . So, for instance, the move from  $(2E-4, 6E-4)$  to  
 711  $(8E-4, 6E-4)$  is annotated with  $p_6$ .

712 With the above framework, it is now easy to see that the  
 713 algorithm executes the sequence  $p_2, p_4, p_6, p_7, p_{10}, p_{11}$ , which  
 714 results in the discovery of the actual selectivity of  $Y$  epp.  
 715 After this, the 1D PlanBouquet takes over and the  
 716 selectivity of  $X$  is learnt by executing  $P_{11}$  and  $P_{19}$  in regular  
 717 (non-spill) mode.

718 This example trace of 2D-SpillBound exemplifies  
 719 how the benefits of half-space pruning and CDI execution are  
 720 realized. It is important to note that 2D-SpillBound may  
 721 execute a few plans *twice*—for example, plan  $P_{11}$ —once in  
 722 spill-mode (i.e.,  $p_{11}$ ) and once as part of the 1D PlanBouquet  
 723 exploration phase. In fact, this notion of repeating a plan ex-  
 724 ecution during the search process substantially contributes to  
 725 the MSO bound in the general case of  $D$  epps.

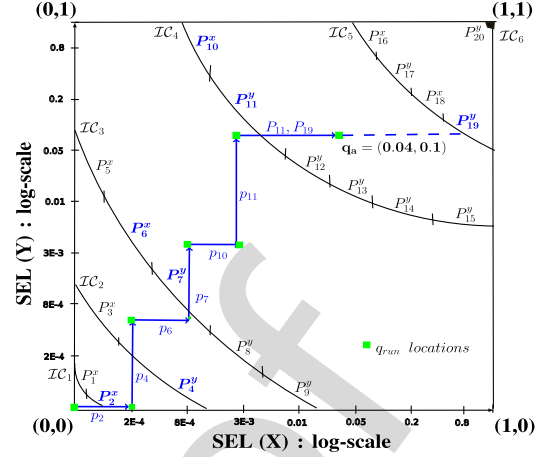


Fig. 7. Execution trace for TPC-DS Query 91.

#### 4.1.2 Performance Bounds

726 Consider the situation where  $q_a$  is located in the region  
 727 between  $\mathcal{IC}_k$  and  $\mathcal{IC}_{k+1}$ , or is directly on  $\mathcal{IC}_{k+1}$ . Then, the  
 728 2D-SpillBound algorithm explores the contours from 1 to  
 729  $k+1$  before discovering  $q_a$ . In this process,  
 730

**Lemma 4.1.** *The 2D-SpillBound algorithm ensures that at  
 731 most two plans are executed from each of the contours  
 732  $\mathcal{IC}_1, \dots, \mathcal{IC}_{k+1}$ , except for one contour in which at most three  
 733 plans are executed.*

**Proof.** Let the exact selectivity of one of the epps be learnt in  
 735 contour  $\mathcal{IC}_h$ , where  $1 \leq h \leq k+1$ . From CDI execution,  
 736 we know that 2D-SpillBound ensures that at most two  
 737 plans are executed in each of the contours  $\mathcal{IC}_1, \dots, \mathcal{IC}_h$ .  
 738 Subsequently, PlanBouquet begins operating from con-  
 739 tour  $\mathcal{IC}_h$ , resulting in three plans being executed in  $\mathcal{IC}_h$ ,  
 740 and one plan each in contours  $\mathcal{IC}_{h+1}$  through  $\mathcal{IC}_{k+1}$ .  $\square$

742 We now analyze the worst-case cost incurred by 2D-  
 743 SpillBound. For this, we assume that the contour with  
 744 three plan executions is the *costliest* contour  $\mathcal{IC}_{k+1}$ . Since the  
 745 ratio of costs between two consecutive contours is 2, the total  
 746 cost incurred by 2D-SpillBound is bounded as follows:

$$\begin{aligned}
 TotalCost &\leq 2 * CC_1 + \dots + 2 * CC_k + 3 * CC_{k+1} \\
 &= 2 * CC_1 + \dots + 2 * 2^{k-1} * CC_1 + 3 * 2^k * CC_1 \\
 &\leq 2^{k+2} * CC_1 + 2^k * CC_1 \\
 &= 5 * 2^k * CC_1.
 \end{aligned} \tag{6}$$

749 From the PCM assumption, we know that the cost for an  
 750 oracle algorithm (that apriori knows the location of  $q_a$ ) is  
 751 lower bounded by  $CC_k$ . By definition,  $CC_k = 2^{k-1} * CC_1$ .  
 752 Hence,  
 753

$$MSO \leq \frac{5 * 2^k * CC_1}{2^{k-1} * CC_1} = 10, \tag{7}$$

755 leading to the theorem:  
 756

**Theorem 4.2.** *The MSO bound of 2D-SpillBound for  
 757 queries with two error-prone predicates is bounded by 10.*

**Remark.** Note that even for a  $\rho$  value as low as 3, the MSO  
 759 bound of 2D-SpillBound is better than the bound,  
 760  $4 * 3 = 12$ , offered by PlanBouquet.  
 761



## 4.2 Extending to Higher Dimensions

We now present `SpillBound`, the generalization of the 2D-`SpillBound` algorithm to handle  $D$  error-prone predicates  $e_1, \dots, e_D$ . Before doing so, we hasten to add that the EPP set, as mentioned earlier, is constantly updated during the execution, and epps are removed from this set as and when their selectivities become fully learnt. Further, when a contour  $\mathcal{IC}_i$  is explored, the *effective search space* is the subset of locations on  $\mathcal{IC}_i$  whose selectivity along the learnt dimensions matches the learnt selectivities. From now on, in the context of exploration, references to  $\mathcal{IC}_i$  will mean its effective search space.

The primary generalization that needs to be achieved is to select, prior to exploration of a contour  $\mathcal{IC}_i$ , the best set (wrt selectivity learning) of  $|\text{EPP}|$  plans that satisfy the half-space pruning property and ensure complete coverage of the contour. To do so, we consider the (location, plan) pairs  $(q_{max}^1, P_{max}^1), \dots, (q_{max}^{|\text{EPP}|}, P_{max}^{|\text{EPP}|})$  as defined at the end of the Section 3.2. The set of  $|\text{EPP}|$  plans that satisfy the contour density independent execution property is  $\{P_{max}^1, \dots, P_{max}^{|\text{EPP}|}\}$ .

A subtle but important point to note here is that, during the exploration of  $\mathcal{IC}_i$ , the identity of  $P_{max}^j$  may change as the contour processing progresses. This is because some of the plans that were assigned to spill on other epps, may switch to spilling on  $e_j$  due to their original epps being completely learnt during the ongoing exploration. Accordingly, we term the first execution of a  $P_{max}^j$  in contour  $\mathcal{IC}_i$  as a *fresh execution*, and subsequent executions on the same epp as *repeat executions*.

---

### Algorithm 1. The `SpillBound` Algorithm

---

```

Init:  $i = 1, \text{EPP} = \{e_1, \dots, e_D\};$ 
while  $i \leq m_D$  ▷ for each contour
  if  $|\text{EPP}| = 1$  then ▷ only one epp left
    Run PlanBouquet to discover the selectivity of the
    remaining epp starting from the present contour;
    Exit;
  end if
  Run the spill node identification procedure on each plan
  in the contour  $\mathcal{IC}_i$ , i.e. plans in PLi, and use this informa-
  tion to choose plan  $P_{max}^j$  for each epp  $e_j$ ;
  exec-complete = false;
  for each epp  $e_j$  do
    exec-complete = Spill-Mode-Execution( $P_{max}^j, e_j, \text{CC}_i$ );
    Update  $q_{run.j}$  based on selectivity learnt for  $e_j$ ;
    if exec-complete then
      /* learnt the actual selectivity for  $e_j^*$  */
      Remove  $e_j$  from the set EPP;
      Break;
    end if
  end for
  if !exec-complete then
     $i = i + 1$ ; /* Jump to next contour */
  end if
  Update ESS based on learnt selectivities;
end while

```

---

Finally, it is possible that a specific epp may have *no* plan on  $\mathcal{IC}_i$  on which it can be spilled—this situation is handled by simply skipping the epp. The complete pseudocode for `SpillBound` is presented in Algorithm 1—here, `Spill-Mode-Execution`( $P_{max}^j, e_j, \text{CC}_i$ ) refers to the execution of plan  $P_{max}^j$  spilling on  $e_j$  with budget  $\text{CC}_i$ .

With the above construction, the following lemma can be proved in a manner analogous to that of Lemma 3.2:

**Lemma 4.3.** *In contour  $\mathcal{IC}_i$ , if no plan in the set  $\{P_{max}^j | e_j \in \text{EPP}\}$  reaches completion when executed in spill-mode, then  $\text{Cost}(P_{q_a}, q_a) > \text{CC}_i$ , triggering a jump to the next contour  $\mathcal{IC}_{i+1}$ .*

### 4.2.1 Performance Bounds

We now present an overview of how the MSO bound is obtained for `SpillBound`—the full proof is available in [14].

In the worst-case analysis of 2D-`SpillBound`, the exploration cost of every intermediate contour is bounded by twice the cost of the contour. Whereas the exploration cost of the last contour (i.e.,  $\mathcal{IC}_{k+1}$ ) is bounded by three times the contour cost because of the possible execution of a third plan during the `PlanBouquet` phase. We now present how this effect is accounted for in the general case.

*Repeat Executions:* As explained before, the identity of plan  $P_{max}^j$  may dynamically change during the exploration of a contour  $\mathcal{IC}_i$ , resulting in repeat executions. If this phenomenon occurs, the new  $P_{max}^j$  plan would have to be executed to ensure compliance with Lemma 4.3. We observe that each repeat execution of an epp is preceded by an event of fully learning the selectivity of some other epp, leading to the following lemma (proof in [14]):

**Lemma 4.4.** *The `SpillBound` algorithm executes at most  $D$  fresh executions in each contour, and the total number of repeat executions across contours is bounded by  $\frac{D(D-1)}{2}$ .*

Suppose that the actual selectivity location  $q_a$  is located in the region between  $\mathcal{IC}_k$  and  $\mathcal{IC}_{k+1}$ , or is directly on  $\mathcal{IC}_{k+1}$ . Then, the total cost incurred by the `SpillBound` algorithm in discovering  $q_a$  is the sum of costs from fresh and repeat executions in each of the contours  $\mathcal{IC}_1$  through  $\mathcal{IC}_{k+1}$ . Further, the worst-case cost is incurred when all the repeat executions happen at the costliest contour, namely  $\mathcal{IC}_{k+1}$ . Hence, the total cost of `SpillBound` is given by

$$\sum_{i=1}^{k+1} (\# \text{fresh executions}(\mathcal{IC}_i)) * \text{CC}_i + \frac{D(D-1)}{2} * \text{CC}_{k+1}. \quad (856)$$

Since the number of fresh executions on any contour is bounded by  $D$ , we obtain the following theorem (proof on similar lines to the 2D scenario):

**Theorem 4.5.** *The MSO bound of the `SpillBound` algorithm for any query with  $D$  error-prone predicates is bounded by  $D^2 + 3D$ .*

**Remark.** For ease of exposition of `SpillBound`, and to facilitate comparison with `PlanBouquet`, we have chosen a cost ratio of 2 between successive contours. However, it is interesting to note that cost doubling is *not* the ideal choice for `SpillBound`, unlike `PlanBouquet`, as explained in [14]—for instance, a factor of 1.8 improves `SpillBound`'s MSO guarantee from 10 to 9.9 in the 2D case. Only marginal improvements are obtained with these ideal factors for the ESS dimensionalities considered in our study.

## 4.3 Lower Bound

We now present a lower bound on MSO for a class of deterministic half-space pruning algorithms denoted by  $\mathcal{E}$ , that includes `SpillBound` in its ambit. We prove the following theorem.

**Theorem 4.6.** *For any algorithm  $A \in \mathcal{E}$  and  $D \geq 2$ , there exists a  $D$ -dimensional ESS where MSO of  $A$  is at least  $D$ .*

The proof of the above theorem is omitted due to space considerations and can be found in Section 5 of [14].

TABLE 2  
Cost of Enforcing Contour Alignment

Query	Original	$\lambda = 1.2$	$\lambda = 1.5$	$\lambda = 2.0$	Max $\lambda$
3D_Q96	18	18	27	45	130
4D_Q7	70	70	90	90	3.62
4D_Q26	20	30	40	50	66.95
4D_Q91	67	67	77	77	5.38
5D_Q29	40	70	100	-	1.35
5D_Q84	100	-	-	-	1

## 5 THE ALIGNEDBOUND ALGORITHM

Given the quadratic-to-linear gap on MSO, we now identify exploration scenarios in which the MSO of `SpillBound` matches the  $\Omega(D)$  lower bound—we do so by leveraging the contour alignment notion. Consider the scenario in which all the contours are aligned—then by Lemma 3.3, each of these contour requires only a single execution to make quantum progress. Following the lines of the analysis of `SpillBound`, and the fact that the most expensive execution sequence occurs when all the selectivities are learnt in the last contour ( $\mathcal{IC}_{k+1}$ ), the total cost incurred in the worst-case would be

$$\begin{aligned} TotalCost &= CC_1 + \dots + CC_k + D * CC_{k+1} \\ &= CC_1 + \dots + 2^{k-1}CC_1 + D * 2^kCC_1 \\ &\leq (2^{k-1}CC_1)(2D + 2), \end{aligned}$$

leading to the following theorem:

**Theorem 5.1.** *If the contour alignment property is satisfied at every step of the algorithm's execution, then the MSO bound is  $2D + 2$ .*

In practice, however, the contour alignment property may not be natively satisfied at all contours—for instance, as enumerated later in Table 2, as few as 18 percent of the contours were aligned for a 3D ESS with TPC-DS Query 96. Therefore, we propose in this section the `AlignedBound` algorithm which operates in three steps: First, it exploits the property of alignment wherever available natively. Second, it attempts to *induce* this property, by replacing the optimal plan with an aligned substitute if the substitution does not overly degrade the performance. Finally, it investigates the possibility of leveraging alignment at a finer granularity than complete contours.

To aid in description of the algorithm, we denote by  $Ext(i, j)$  the set of all extreme locations on a contour  $\mathcal{IC}_i$  along a dimension  $j$ . With this, a contour  $\mathcal{IC}_i$  is said to satisfy contour alignment along dimension  $j$  if  $q_{max}^j \in Ext(i, j)$ , i.e., at least one of the extreme locations along dimension  $j$  has an optimal plan that spills on  $e_j$ . Second, the set of all plans that spill on predicate  $e_k$  is denoted by  $\mathcal{P}^k$ .

### 5.1 Induced Contour Alignment

Given a contour  $\mathcal{IC}_i$  that does not satisfy contour alignment, we *induce* contour alignment on the contour as follows: Consider a plan  $P$  which spills on  $e_k \in EPP$ . It is a candidate replacement plan for any location  $q_{ext}^k \in Ext(i, k)$  in order to obtain alignment along dimension  $k$ —the cost of the replacement is equal to  $Cost(P, q_{ext}^k)$ . Therefore, the minimum cost of inducing contour alignment along dimension  $k$  is given by the pair  $(P^k \in \mathcal{P}^k, q_{ext}^k \in Ext(i, k))$  for which  $Cost(P^k, q_{ext}^k)$  is minimized. Next, we find the dimension  $j$  for which the cost

of the replacement pair  $(P^j, q_{ext}^j)$  is minimum across all dimensions. Finally, the optimal plan at  $q_{ext}^j$  is replaced by  $P^j$ , and the *penalty*  $\lambda$  of this replacement is the ratio of  $Cost(P^j, q_{ext}^j)$  to  $Cost(P_{q_{ext}^j}, q_{ext}^j)$ .

The usefulness of induced contour alignment depends on the penalty incurred in enforcing the property. To assess this quantitatively, we conducted an empirical study, whose results are shown in Table 2. Here, each row is a query instance. The “Original” column indicates the percentage of the contours that satisfy contour alignment without any replacements. A column with a particular  $\lambda$  value, say  $c$ , indicates the percentage of the contours satisfying contour alignment when the replacement plans are not allowed to exceed a penalty of  $c$ . The last column shows the minimum penalty that needs to be incurred for all the contours to satisfy contour alignment.

We see from the table that there are cases where full contour alignment can be induced relatively cheaply—for instance, a 50 percent penalty threshold is sufficient to make Query 5D\_Q29 completely aligned. However, there also are cases, such as 3D\_Q96, where extremely high penalty needs to be paid to achieve contour alignment. Therefore, we now develop a weaker notion of alignment, called “*predicate set alignment*” (PSA), which operates at a finer granularity than entire contours, and attempts to address these problematic scenarios.

### 5.2 Predicate Set Alignment

We say that a set  $T \subseteq EPP$  satisfies predicate set alignment with the leader dimension  $j$  if, for any location  $q \in \mathcal{IC}_i$  whose optimal plan spills on any dimension in  $T$ ,  $q \cdot j \leq q_{max}^j \cdot j$ . The set of all locations in  $\mathcal{IC}_i$  whose optimal plan spills on a dimension corresponding to a predicate in  $T$ , is denoted by  $\mathcal{IC}_i|T$ . For convenience, we assume that the predicate corresponding to the leader dimension belongs to  $T$ . Note that PSA is a weaker notion of alignment—while contour alignment with leader dimension  $j$  mandates that  $q_{max}^j \cdot j \geq q \cdot j$  for any  $q \in \mathcal{IC}_i$ , PSA only requires that  $q_{max}^j \cdot j \geq q \cdot j$  for all  $q \in \mathcal{IC}_i|T$ .

**Lemma 5.2.** *Suppose  $T_1, \dots, T_l$  are sets of *epps* satisfying predicate set alignment such that  $\cup_{k=1}^l T_k = EPP$ , then  $\cup_{k=1}^l \mathcal{IC}_i|T_k = \mathcal{IC}_i$ .*

**Proof.** Every  $q \in \mathcal{IC}_i$  spills on one of the dimensions in EPP. Therefore, it belongs to at least one  $\mathcal{IC}_i|T$ .  $\square$

**Lemma 5.3.** *Suppose  $T_1, \dots, T_l$  are sets of *epps* satisfying predicate set alignment such that  $\cup_{k=1}^l T_k = EPP$ , then spill-mode execution of  $l$  POSP plans on  $\mathcal{IC}_i$  is sufficient to make quantum progress.*

**Proof.** Let  $j_1, \dots, j_l$  be the leader dimensions for  $T_1, \dots, T_l$ , respectively. Then, the  $l$  POSP plans chosen for the execution are  $P_{q_{max}^{j_k}}^{j_k}$  for  $k = 1, \dots, l$ . After this, based on the definition of PSA, Lemma 5.2, and an argument similar to the proof of Lemma 3.3, it can be shown that spill-mode execution of the chosen  $l$  POSP plans is sufficient to make quantum progress. The complete proof is available in [14].  $\square$

#### 5.2.1 Inducing Predicate Set Alignment

Consider a contour  $\mathcal{IC}_i$ , and a candidate set  $T \subseteq EPP$  with a leader dimension  $j \in T$ . We now present a mechanism to induce predicate set alignment on  $T$  with leader dimension  $j$ .

We consider the extreme location along the dimension  $j$  among all the locations in  $\mathcal{IC}_i|T$ , i.e.  $q_T^j = \arg \max_{q \in \mathcal{IC}_i|T} q^j$  (in case of a multiplicity of such points, any one point can be picked). Consider the set  $S = \{q \in \mathcal{IC}_i \wedge q^j = q_T^j\}$ , i.e. all the locations belonging to  $\mathcal{IC}_i$  whose coordinate value on  $j$ th dimension is equal to the coordinate value on  $j$ th dimension of an extreme location in  $\mathcal{IC}_i|T$ . It is easy to see that  $T$  satisfies predicate set alignment if the optimal plan at any of the locations in  $S$  is replaced with a plan  $P$  that spills on  $e_j$ . We now find a pair  $(P \in \mathcal{P}^j, q \in S)$  such that  $Cost(P, q)$  is minimum. The predicate set alignment property is induced by replacing the optimal plan at  $q$  with the plan  $P$ . The penalty  $\lambda$  for the replacement is defined as before. We remark that the step of inducing predicate set alignment can be implemented efficiently and is discussed in [14].

### 5.2.2 Finding Minimum Cost Predicate Set Cover

Lemma 5.3 essentially says that a set of predicate sets  $T_1, \dots, T_l$  that cover EPP can be leveraged to make quantum progress. We now argue that it is sufficient to limit the search to merely the set of *partition covers* of EPP.

Consider a set  $T$  which satisfies PSA along dimension  $j$ . The *cover cost* of  $T_1, \dots, T_l$  is said to be sum of cost of enforcing PSA for each of the  $T_i$ s. We say that  $T$  satisfies *maximal PSA* with leader dimension  $j$  if no super-set of  $T$  satisfies the property with same or lesser cost. Consider  $T_1, \dots, T_l$  which cover EPP and have been enforced to satisfy maximal PSA. We now obtain a partition cover whose cover cost is at most the cover cost of  $T_1, \dots, T_l$ .

Let  $j_1, \dots, j_l$  be the leader dimensions for  $T_1, \dots, T_l$ . The maximal property of the  $T_i$ s implies that no dimension can be a leader dimension for more than one  $T_i$ . Therefore, the following sets  $\pi_1 = T_1 + \{j_1\} - \cup_{m=2}^{m=l} \{j_m\}, \pi_k = T_k + \{j_k\} - \cup_{m=1, m \neq k}^{m=l} \{j_m\} - \cup_{m=1}^{m \leq k} \pi_m$  for  $k = 2, \dots, l-1$ , and  $\pi_l = T_l - \cup_{m=1}^{m=l-1} \pi_m$  provide a partition cover with the same set of leader dimensions  $j_1, \dots, j_l$ . It follows that the cover cost of  $\pi_1, \dots, \pi_l$  is at most the cover cost of  $T_1, \dots, T_l$ . The full proof of this is presented in [14]. Therefore, we can restrict the search for EPP cover to only partition covers without incurring any increase in the penalty of the EPP cover. The benefit of this is that the number of partition covers of a set is much smaller than the number of different ways of covering a set with its subsets.

Given a partition cover  $\pi = \{\pi_1, \dots, \pi_l\}$ ,  $\pi_\lambda$  denotes the sum of the penalties incurred in enforcing PSA for each of the  $\pi_i$ s along their leader dimensions.

### 5.3 Algorithm Description

The AlignedBound algorithm is presented in Algorithm 2. The steps that are identical to the steps in SpillBound are not presented again and simply captured as comments.

The key steps of the algorithm are S1 and S2 which are executed using the partition cover and predicate set alignment techniques described in Section 5.2.

A legitimate concern at this point is whether in trying to induce alignment, the  $D^2 + 3D$  guarantee may have been lost along the way. The proof that this is not so, and that the quadratic bound is retained is available in [14]. In summary, AlignedBound delivers an MSO that is guaranteed to be in the platform-independent range  $[2D + 2, D^2 + 3D]$ .

### Algorithm 2. The AlignedBound Algorithm

```

1: Init:  $i = 1$ ,  $EPP = \{e_1, \dots, e_D\}$ ;
2: while  $i \leq m$  do ▷ for each contour
3:   /* Handle special 1-D case when it is encountered */
4:   S0:  $\Pi =$  Set of all partitions of EPP (remaining epps);
5:   S1: We pick  $\pi \in \Pi$  with minimum  $\pi_\lambda$ ;
6:   for each part  $\pi_k \in \pi$  do
7:     S2: Let  $j_k$  be the leader dimension,  $P$  the replacement
           plan along dimension  $j_k$ , and  $q$  the location whose
           optimal plan is replaced with  $P$ ;
8:     exec-complete = Spill-Mode-Execution( $P, e_{j_k}, Cost(P, q)$ );
9:     Update  $q_{run}.j_k$  based on selectivity learnt for  $e_{j_k}$ ;
10:    if exec-complete then
11:      Remove  $e_{j_k}$  from the part  $\pi_k$  and the set EPP;
12:      Break;
13:    end if
14:  end for
15:  /* Update ESS, jump contour as in SpillBound */
16: end while

```

## 6 EXPERIMENTAL EVALUATION

As mentioned earlier, the MSO guarantees delivered by PlanBouquet and SpillBound are not directly comparable, due to the inherently different nature of their dependencies on the  $\rho$  and  $D$  parameters, respectively. However, we need to assess whether the platform-independent feature of SpillBound is procured at the expense of a deterioration in the numerical bounds. Accordingly, we present in this section an evaluation of SpillBound on a representative set of complex OLAP queries, and compare its MSO performance with that of PlanBouquet. Furthermore, we also conduct an evaluation of AlignedBound over the same set of queries to appraise its performance benefits over SpillBound. The experimental framework, which is similar to that used in [1], is described first, followed by an analysis of the results.

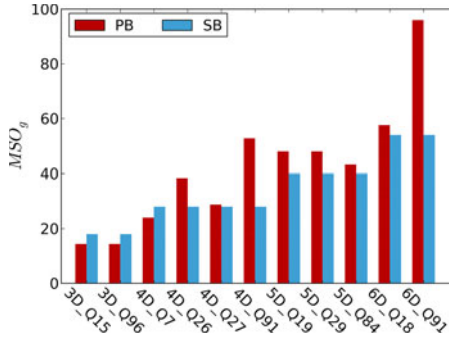

### 6.1 Database and System Framework

Our test workload is comprised of representative SPJ queries from the TPC-DS benchmark, operating at the base size of 100 GB. The number of relations in these queries range from 4 to 10, and a spectrum of join-graph geometries are modeled, including *chain*, *star*, *branch*, etc. The number of epps range from 2 to 6, all corresponding to join predicates, giving rise to challenging multi-dimensional ESS spaces.

To succinctly characterize the queries, the nomenclature  $xD\_Qz$  is employed, where  $x$  specifies the number of epps, and  $z$  the query number in the TPC-DS benchmark. For example, 3D\_Q15 indicates TPC-DS Query 15 with three of its join predicates considered to be error-prone.

The database engine used in our experiments is a modified version of PostgreSQL 8.4 [15] engine, with the primary changes being the (1) selectivity injection—to generate the ESS, (2) abstract plan execution—to instruct the execution engine to execute a particular plan, (3) time-limited execution of plans and (4) spilling—to execute plans in spill-mode. In addition, we implement a feature that obtains a least cost plan from optimizer which spills on a user-specified epp. This is primarily needed for AlignedBound algorithm to find the minimum penalty replacement pair which is mentioned in Section 5.

The remainder of this section is organized as follows. For ease of presentation, first we compare the performance of

Fig. 8. Comparison of MSO guarantees ( $MSO_g$ ).Fig. 9. Variation of  $MSO_g$  with dimensionality (Q91).

1106 PlanBouquet and SpillBound, and subsequently move  
 1107 on to comparing SpillBound and AlignedBound. We use  
 1108 the abbreviations PB, SB and AB to refer to PlanBouquet,  
 1109 SpillBound and AlignedBound, respectively. Further,  
 1110 we use  $MSO_g$  (MSO guarantee) and  $MSO_e$  (MSO empirical)  
 1111 to distinguish between the MSO guarantee and the empiri-  
 1112 cally evaluated MSO obtained on our suite of queries.

## 1113 6.2 SpillBound versus PlanBouquet

1114 The MSO guarantee for PlanBouquet on the original ESS  
 1115 typically turns out to be very high due to the large values of  
 1116  $\rho$ . Therefore, as in [1], we conduct the experiments for  
 1117 PlanBouquet only after carrying out the *anorexic reduction*  
 1118 transformation [10] at the default  $\lambda = 0.2$  replacement  
 1119 threshold—we use  $\rho_{RED}$  to refer to this reduced value.

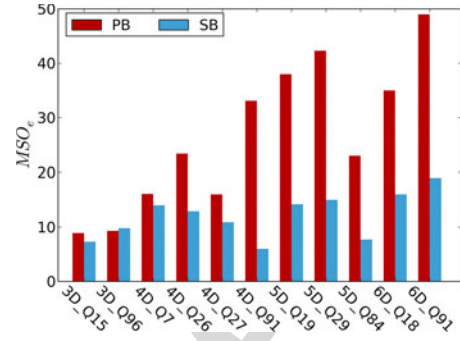
### 1120 6.2.1 Comparison of MSO Guarantees ( $MSO_g$ )

1121 A summary comparison of  $MSO_g$  for PB and SB over almost  
 1122 a dozen TPC-DS queries of varying dimensionality is shown  
 1123 in Fig. 8—for PB, they are computed as  $4(1 + \lambda)\rho_{RED}$ ,  
 1124 whereas for SB, they are computed as  $D^2 + 3D$ .

1125 We observe here that in a few instances, specifically  
 1126 4D\_Q26, 4D\_Q91 and 6D\_Q91, SB's guarantee is noticeably  
 1127 tighter than that of PB—for instance, the values are 28 and  
 1128 52.8, respectively, for 4D\_Q91. In the remaining queries, the  
 1129 bound quality is roughly similar between the two algo-  
 1130 rithms. Therefore, contrary to our fears, the MSO guarantee  
 1131 is not found to have suffered due to incorporating platform  
 1132 independence.

### 1133 6.2.2 Variation of MSO Guarantee with Dimensionality

1134 In our next experiment, we investigated the behavior of  $MSO_g$   
 1135 as a function of ESS dimensionality for a given query. We  
 1136 present results here for an example TPC-DS query, namely  
 1137 Query 91, wherein the number of epps were varied from 2  
 1138 upto 6—the corresponding performance profile is shown in

Fig. 10. Comparison of empirical MSO ( $MSO_e$ ).

1139 Fig. 9. We observe here that while SB is marginally worse at  
 1140 the lowest dimensionality of 2, it becomes appreciably better  
 1141 than PB with increasing dimensionality—in fact, at 6D, the  
 1142 values are 96 and 54 for PB and SB, respectively.

### 1143 6.2.3 Comparison of Empirical MSO ( $MSO_e$ )

1144 We now turn our attention to evaluating the empirical MSO,  
 1145  $MSO_e$ , incurred by the two algorithms. There are two rea-  
 1146 sons that it is important to carry out this exercise: First, to  
 1147 evaluate the looseness of the guarantees. Second, to evaluate  
 1148 whether PB, although having weaker bounds in theory, pro-  
 1149 vides better performance in practice, as compared to SB.

1150 The assessment was accomplished by explicitly and  
 1151 exhaustively considering each and every location in the ESS  
 1152 to be  $q_a$ , and then evaluating the sub-optimality incurred for  
 1153 this location by PB and SB. Finally, the maximum of these  
 1154 values was taken to represent the  $MSO_e$  of the algorithm.

1155 The  $MSO_e$  results are shown in Fig. 10 for the entire suite  
 1156 of test queries. Our first observation is that the empirical per-  
 1157 formance of SB is far better than the corresponding guaran-  
 1158 tees in Fig. 8. In contrast, while PB also shows improvement,  
 1159 it is not as dramatic. For instance, considering 6D\_Q18, PB  
 1160 reduces its MSO from 57.6 to 35.2, whereas SB goes down  
 1161 from 54 to just 16. A detailed analysis of the significant gap  
 1162 between SB's  $MSO_g$  and  $MSO_e$  values is provided in [2].

1163 The second observation is that the gap between SB and  
 1164 PB is *accentuated* here, with SB performing substantially bet-  
 1165 ter over a larger set of queries. For instance, consider query  
 1166 5D\_Q29, where the  $MSO_g$  values for PB and SB were 52.8  
 1167 and 40, respectively—the corresponding empirical values  
 1168 are 42.3 and 15.1 in Fig. 10.

1169 Finally, even for a query such as 4D\_Q7, where PB had a  
 1170 marginally *better* bound (24 for PB and 28 for SB in Fig. 8),  
 1171 we find that it is SB which behaves better in practice (16.1  
 1172 for PB and 13.9 for SB in Fig. 10).

### 1173 6.2.4 Average-Case Performance (ASO)

1174 A legitimate concern with our choice of MSO metric is that  
 1175 its improvements may have been purchased by degrading  
 1176 average-case behavior. To investigate this possibility, we  
 1177 have considered ASO, the average case equivalent of MSO,  
 1178 which is defined as follows under the assumption that all  
 1179  $q_a$ 's are equally likely

$$1181 \quad ASO = \frac{\sum_{q_a \in \text{ESS}} \text{SubOpt}(q_e, q_a)}{\sum_{q_a \in \text{ESS}} 1}. \quad (8)$$

1182 We evaluated the ASO of PB and SB for all the test  
 1183 queries, and these results are shown in Fig. 11. Observe  
 1184 that, contrary to our fears, SB provides much better  
 1185




Fig. 11. Comparison of ASO performance.

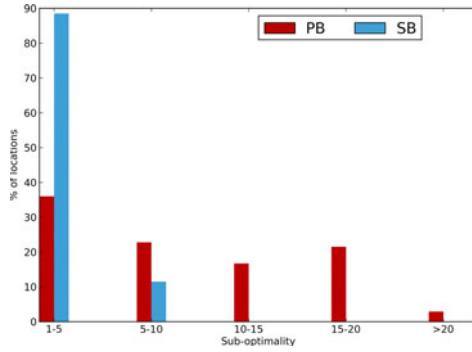


Fig. 12. Sub-optimality distribution (4D\_Q91).

performance, especially at higher dimensions, as compared to PB. For instance, with 5D\_Q19, the ASO for SB is nearly 100 percent better than PB, going down from 17 to 8.6. Thus, SB offers significant benefits over PB in terms of both worst-case and average-case behavior.

### 6.2.5 Sub-Optimality Distribution

In our final analysis, we profile the *distribution* of sub-optimality over the ESS. That is, a histogram characterization of the number of locations with regard to various sub-optimality ranges. A sample histogram, corresponding to query 4D\_Q91, is shown in Fig. 12, with sub-optimality ranges of width 5. We observe here that for over 90 percent of the ESS locations, the sub-optimality of SB is less than 5. Whereas this performance is achieved for only 35 percent of the locations using PB. Similar patterns were observed for the other queries as well, and these results indicate that from both *global and local* perspectives, SB has desirable performance characteristics as compared to PB.


### 6.3 Wall-Clock Time Experiments

All the experiments thus far were based on optimizer cost values. We have also carried out experiments wherein the actual *query response times* were explicitly measured for the native optimizer, SB and AB. As a representative example, we have chosen TPC-DS Q91 featuring 4 error-prone predicates, referred to as  $e_1, \dots, e_4$ . In this experiment, the *optimal* plan took less than a minute (44 secs) to complete the query. However, the *native optimizer* required more than 10 minutes (628 secs) to process the data, thus incurring a sub-optimality of 14.3.

In contrast, SB took only around 4 minutes (246 secs), corresponding to a sub-optimality of 5.6. Table 3 shows the drilled down information of plan executions for every contour with SB. In addition, the selectivities learnt for the

TABLE 3  
SpillBound Execution on TPC-DS Query 91

Contour no.	$e_1$ (plan)	$e_2$ (plan)	$e_3$ (plan)	$e_4$ (plan)	Time (sec.)
1	0	0	0	<b>0.08</b> ( $p_1$ )	1.3
2	<b>0.02</b> ( $p_3$ )	0	0	<b>0.3</b> ( $p_2$ )	7.5
3	<b>0.08</b> ( $p_4$ )	0	0	<b>1</b> ( $p_5$ )	21
4	<b>0.2</b> ( $p_4$ )	0	0	<b>12</b> ( $p_5$ )	51.2
5	<b>5</b> ( $p_9$ )	<b>0.8</b> ( $p_6$ )	0	12	86.3
5	<b>30</b> ( $p_9$ )	0.8	<b>5</b> ( $p_8$ )	<b>60</b> ( $p_7$ )	176.4
6	<b>80</b> ( $P_{10}$ )	0.8	5	60	246.4

Fig. 13. Comparison of empirical MSO ( $MSO_e$ ).

corresponding  $e_{pp}$  during every execution are also captured. The selectivity information learnt in each contour, shown in %, is indicated by boldfaced font in the table. Further, for each execution, the plan employed, and the overheads accumulated so far, are enumerated. A plan  $P$  executed in spill-mode is indicated with a  $p$ . As can be seen in the table, the execution sequence consists of partial executions of 13 plans spanning six consecutive contours, and culminates in the full execution of plan  $P_{10}$  which produces the query results.

Finally, we also conducted the above mentioned Q91 experiment with AB. The algorithm needed less than 3 minutes (165.1 secs) for completing the query, involving 10 partial plan executions before the culminating full execution. Thus, AB brings the sub-optimality down to just 3.8 in this example.

### 6.4 AlignedBound versus SpillBound

We now turn our attention to evaluating how the predicate set alignment property, exploited by AB, impacts its empirical performance as compared to SB. Specifically, we assess the  $MSO_e$  incurred by the two algorithms, with the comparison on other metrics, such as ASO and sub-optimality distribution, deferred to [14].

#### 6.4.1 Comparison of Empirical MSO

The  $MSO_e$  numbers for SB and AB are captured in Fig. 13. First, we highlight that the  $MSO_e$  values for AB are consistently less than around 10, for all the queries. Second, AB significantly brings down the  $MSO_e$  numbers for the several queries whose  $MSO_e$  values with SB are greater than 15. As a case in point, AB brings down the  $MSO_e$  of 6D\_Q91 from 19 to 10.4.

#### 6.4.2 Rationale for AB's Performance Benefits

Recall that AB provides an MSO guarantee in the range  $[2D + 2, D^2 + 3D]$ . As can be seen in Fig. 13, the  $MSO_e$  values for AB are closer to the corresponding  $2D + 2$  bound

1254 value, shown with dotted lines in the figure. These results  
1255 suggest that the empirical performance of AB approaches  
1256 the  $\mathcal{O}(D)$  lower bound on MSO.

1257 We now shift our focus to examining the reasons for AB's  
1258 MSO<sub>e</sub> performance benefits over SB. In Table 4, the maximum  
1259 penalty over all partitions encountered during execution is  
1260 tabulated for the various queries. The important point to note  
1261 here is that these penalty values are lower than 3, even for 6D  
1262 queries. Since the highest cost investment for quantum prog-  
1263 ress in any contour is the maximum penalty times the cost of  
1264 the contour, the low value for penalty results in the observed  
1265 benefits, especially for higher dimensional queries.

## 1266 6.5 Evaluation on the JOB Benchmark

1267 All the above experiments were conducted on the TPC-DS  
1268 benchmark, an industry standard. Recently a new bench-  
1269 mark, called Join Order Benchmark (JOB), specifically  
1270 designed to provide challenging workloads for current opti-  
1271 mizers, was proposed in [4]. Given its design objective, it  
1272 appears appropriate to evaluate our query processing algo-  
1273 rithms on this platform. A difficulty, however, is that all the  
1274 queries in the JOB benchmark feature *cyclic predicates*,  
1275 directly nullifying our selectivity independence assump-  
1276 tion. Therefore, as an interim work-around, we shut off the  
1277 optimizer's automatic inclusion of implicit join predicates,  
1278 and verified that the consequent optimizer plans either  
1279 remained the same or were only marginally sub-optimal.

1280 We now present results for a representative Query 1a  
1281 from the JOB benchmark. For this query, we found that, as  
1282 expected by design, the native optimizer's performance was  
1283 substantially worse, with the MSO going well above 6,000.  
1284 In marked contrast, SB continued to retain its strong per-  
1285 formance profile with an MSO of only around 12. And AB  
1286 reduced this even further to below 9.

## 1287 7 DEPLOYMENT ASPECTS

1288 Over the preceding sections, we have conducted a theoret-  
1289 ical characterization and empirical evaluation of our pro-  
1290 posed algorithms. We now discuss some pragmatic aspects  
1291 of its usage in real-world contexts. Most of these issues have  
1292 already been previously discussed in [1], in the context of  
1293 the PlanBouquet algorithm, and we therefore only sum-  
1294 marize the salient points here for easy reference.

1295 First, our assumption of a perfect cost model. If we were  
1296 to be assured that the cost modeling errors, while non-zero,  
1297 are *bounded* within a  $\delta$  error factor, then the MSO guarantees  
1298 in this paper will carry through modulo an inflation by a  
1299 factor of  $(1 + \delta)^2$  [14]. That is, the MSO guarantee of Spill-  
1300 Bound (and AlignedBound) would be  $(D^2 + 3D)(1 + \delta)^2$ .  
1301 Moreover, the errors induced by cost model are fairly small.  
1302 For instance,  $\delta = 0.3$  is reported in [16].

1303 Second, with regard to identification of the epps that  
1304 constitute the ESS, we could leverage application domain  
1305 knowledge and query logs to make this selection, or simply  
1306 be conservative and assign all uncertain combination of  
1307 predicates to be epps.

1308 Third, the construction of the contours in the ESS is cer-  
1309 tainly a computationally intensive task since it is predi-  
1310 cated on repeated calls to the optimizer, and the overheads  
1311 increase exponentially with ESS dimensionality. However,  
1312 for canned queries, it may be feasible to carry out an offline  
1313 enumeration; alternatively, when a multiplicity of hard-  
1314 ware is available, the contour constructions can be carried

TABLE 4  
Maximum Penalty for AB

Query	max. penalty for AB
3D_Q15	2.42
3D_Q96	3
4D_Q7	2
4D_Q26	2.25
4D_Q27	2
4D_Q91	2.05
5D_Q19	2.5
5D_Q29	1.81
5D_Q84	1.1
6D_Q18	1.92
6D_Q91	1.25

1315 out in parallel since they do not have any dependence on  
1316 each other.

1317 Finally, while PlanBouquet can directly work off the  
1318 API of existing query optimizers, SpillBound and  
1319 AlignedBound are *intrusive* since they require changes in  
1320 the core engine to support plan spilling and monitoring of  
1321 operator selectivities. However, our experience with Post-  
1322 greSQL is that these facilities can be incorporated relatively  
1323 easily—the full implementation required only a few hun-  
1324 dred lines of code.

## 1325 8 RELATED WORK

1326 Our work materially extends the PlanBouquet approach  
1327 presented in [1], which is the first work to provide worst-case  
1328 guarantees for query processing performance. As already  
1329 highlighted, the primary new contribution is the provision of  
1330 a structural bound with SpillBound (and AlignedBound),  
1331 whereas PlanBouquet delivered a behavioral bound. Fur-  
1332 ther, the performance characteristics of both our algorithms  
1333 are substantively superior to those of PlanBouquet, as illus-  
1334 trated in the experimental study.

1335 A detailed comparison to the prior literature on selectiv-  
1336 ity estimation issues is provided in [1]. Since SpillBound  
1337 and AlignedBound belong to the class of plan switching  
1338 approaches, they may appear similar at first sight to influen-  
1339 tial systems such as POP [3] and Rio [6]. However, there are  
1340 key differences: First, they start with the optimizers estimate  
1341 as the initial seed and then conduct a full-scale re-optimiza-  
1342 tion if the estimate is found to be significantly in error. In  
1343 contrast, our proposed algorithms always start executing  
1344 plans from the *origin* of the selectivity space, ensuring both  
1345 repeatability of the query execution strategy as well as con-  
1346 trolled switching overheads.

1347 Second, both POP and Rio are based on heuristics and do  
1348 not provide any performance bounds. In particular, POP  
1349 may get stuck with a poor plan since its selectivity validity  
1350 ranges are defined using structure-equivalent plans only.  
1351 Similarly, Rios sampling-based heuristics for monitoring  
1352 selectivities may not work well for join-selectivities, and its  
1353 definition of plan robustness based solely on the perfor-  
1354 mance at the corners of the ESS has not been validated.

## 1355 9 CONCLUSION AND FUTURE WORK

1356 We presented SpillBound, a query processing algorithm  
1357 that deliver a worst-case performance guarantee depen-  
1358 dent solely on the dimensionality of the selectivity  
1359 space  $(D^2 + 3D)$ . This substantive improvement over

PlanBouquet is achieved through a potent pair of conceptual enhancements: half-space pruning of the ESS thanks to a spill-based execution model, and bounded number of executions for jumping from one contour to the next. The new approach facilitates porting of the bound across database platforms, easy knowledge and low magnitudes of the bound value, and indifference to the efficacy of the anorexic reduction heuristic. Further, we also showed that SpillBound is within an  $\mathcal{O}(D)$  factor of the best deterministic selectivity discovery algorithm in its class. Finally, we introduced the contour alignment and predicate set alignment properties and leveraged them to design AlignedBound with the objective of bridging the quadratic-to-linear MSO gap between SpillBound and the lower bound.

A detailed experimental evaluation on complex high-dimensional OLAP queries demonstrated that our algorithms provide competitive guarantees to their PlanBouquet counterpart, while their empirical performance is significantly superior. Moreover, AlignedBound's performance often approaches the ideal of MSO linearity in  $D$ .

In our future work, we wish to develop automated assistants for guiding users in deciding whether to use the native query optimizer or our algorithms for executing their queries. We also plan to work on extending SpillBound and AlignedBound to handle the case of dependent predicate selectivities.

## ACKNOWLEDGMENTS

The authors thank the anonymous reviewers and Anshuman Dutt for their valuable comments on this work. This paper is an extended version of [2], with the primary new contributions being the AlignedBound algorithm and its evaluation.

## REFERENCES

- [1] A. Dutt and J. Haritsa, "Plan bouquets: A fragrant approach to robust query processing," *ACM Trans. Database Syst.*, vol. 41, no. 2, pp. 11:1–11:37, 2016.
- [2] S. Karthik, J. Haritsa, S. Kenkre, and V. Pandit, "Platform-independent robust query processing," in *Proc. 32nd IEEE Int. Conf. Data Eng.*, May 2016, pp. 325–336.
- [3] V. Markl, V. Raman, D. Simmen, G. Lohman, H. Pirahesh, and M. Cilimdžić, "Robust query processing through progressive optimization," in *Proc. ACM SIGMOD 30th Int. Conf. Manage. Data*, Jun. 2004, pp. 659–670.
- [4] V. Leis, A. Gubichev, A. Mirchev, P. Boncz, A. Kemper, and T. Neumann, "How good are query optimizers, really?" in *Proc. 42nd Int. Conf. Very Large Data Bases*, Sep. 2016.
- [5] M. Stillger, G. Lohman, V. Markl, and M. Kandil, "LEO-DB2's learning optimizer," in *Proc. 27th Int. Conf. Very Large Data Bases*, Sep. 2001, pp. 19–28.
- [6] S. Babu, P. Bizarro, and D. DeWitt, "Proactive re-optimization," in *Proc. ACM SIGMOD 31st Int. Conf. Manage. Data*, Jun. 2005, pp. 107–118.
- [7] T. Neumann and C. Galindo-Legaria, "Taking the edge off cardinality estimation errors using incremental execution," in *Proc. 15th Conf. Database Syst. Business Technol. Web*, Mar. 2013, pp. 73–92.
- [8] N. Kabra and D. DeWitt, "Efficient mid-query re-optimization of sub-optimal query execution plans," in *Proc. ACM SIGMOD 24th Int. Conf. Manage. Data*, Jun. 1998, pp. 106–117.
- [9] A. Hulgeri and S. Sudarshan, "Parametric query optimization for linear and piecewise linear cost functions," in *Proc. 28th Int. Conf. Very Large Data Bases*, Aug. 2002, pp. 167–178.
- [10] D. Harish, P. Darera, and J. Haritsa, "On the production of anorexic plan diagrams," in *Proc. 33rd Int. Conf. Very Large Data Bases*, Sep. 2007, pp. 1081–1092.
- [11] P. Selinger, M. Astrahan, D. Chamberlin, R. Lorie, and T. Price, "Access path selection in a relational database management system," in *Proc. ACM SIGMOD 5th Int. Conf. Manage. Data*, Jun. 1979, pp. 23–34.

- [12] G. Graefe, "Query evaluation techniques for large databases," *ACM Comput. Surveys*, vol. 25, no. 2, pp. 73–170, 1993.
- [13] S. Chaudhuri, V. Narasayya, and R. Ramamurthy, "Estimating progress of execution for SQL queries," in *Proc. ACM SIGMOD 30th Int. Conf. Manage. Data*, Jun. 2004, pp. 803–814.
- [14] S. Karthik, J. Haritsa, S. Kenkre, V. Pandit, and L. Krishnan, "Platform-independent robust query processing," Tech. Report TR-2016-02, DSL/CDS, IISc, 2016, <http://dsl.cds.iisc.ac.in/publications/report/TR/TR-2016-02.pdf>.
- [15] PostgreSQL. [Online]. Available: <http://www.postgresql.org/docs/8.4/static/release.html>
- [16] G. Lohman, "Is query optimization a solved problem?" [Online]. Available: <http://wp.sigmod.org/?p=1075>



**Srinivas Karthik** is currently working toward the PhD degree in the Indian Institute of Science, Bangalore. His interests include robust query processing and testing.



**Jayant R. Haritsa** is a faculty of the CDS and CSA Departments, Indian Institute of Science, Bangalore, India. His interests include database engine design and testing. He is a fellow of the IEEE and the ACM.



**Sreyash Kenkre** is a research staff member with IBM Research India, and is part of the Cognitive Industry Solutions Group. His interests include graph theory and algorithms.



**Vinayaka Pandit** is a senior technical staff member at IBM Research and leads the Cognitive Industry Solutions Group, IBM India Research Lab. His interests include approximation algorithms, combinatorial optimization and randomization techniques, and their application to information management problems.



**Lohit Krishnan** is a technical staff member at IBM Research and is part of the Cognitive Industry Solutions Group, IBM India Research Lab. His interests include database systems and machine learning.

► For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).