

Robust Plans through Plan Diagram Reduction

Harish D. Pooja N. Darera Jayant R. Haritsa

Technical Report
TR-2007-02

Database Systems Lab
Supercomputer Education and Research Centre
Indian Institute of Science
Bangalore 560012, India

<http://dsl.serc.iisc.ernet.in>

Abstract

Estimates of predicate selectivities by database query optimizers often differ significantly from those actually encountered during query execution, leading to poor plan choices and inflated response times. In this paper, we investigate mitigating this problem by replacing selectivity error-sensitive plan choices with alternative plans that provide robust performance. Our approach is based on the recent observation that even the complex and dense “plan diagrams” associated with industrial-strength optimizers can be efficiently reduced to “anorexic” equivalents featuring only a few plans, without materially impacting query processing quality.

Extensive experimentation with a rich set of TPC-H and TPC-DS-based query templates in a variety of database environments indicate that plan diagram reduction typically retains plans that are substantially resistant to selectivity errors on the base relations. However, it can sometimes also be severely counter-productive, with the replacements performing much worse. We address this problem through a generalized mathematical characterization of plan cost behavior over the parameter space, which lends itself to efficient criteria of when it is safe to reduce. Our strategies are fully non-invasive and have been implemented in the Picasso optimizer visualization tool.

1 Introduction

The query execution plan choices made by database engines often turn out to be poor in practice because the optimizer’s selectivity estimates are significantly in error with respect to the actual values encountered during query execution. Such errors, which can even be in orders of magnitude in real database environments [18], arise due to a variety of reasons [23], including outdated statistics, attribute-value independence assumptions and coarse summaries.

Robust Plans. To address this problem, an obvious approach is to improve the quality of the statistical meta-data, for which several techniques have been presented in the literature ranging from improved summary structures [1] to feedback-based adjustments [23] to on-the-fly reoptimization of queries [15, 18, 3]. A complementary and conceptually different approach, which we consider in this paper, is to identify *robust plans* that are relatively less sensitive to such selectivity errors. In a nutshell, to “aim for resistance, rather than cure”, by identifying plans that provide comparatively good performance over large regions of the selectivity space. Such plan choices are especially important for industrial workloads where global stability is as much a concern as local optimality [17].

Over the last decade, a variety of strategies have been proposed to identify robust plans, including the Least Expected Cost [6, 8], Robust Cardinality Estimation [2] and Rio [3, 4] approaches. These techniques provide novel and elegant formulations (summarized in Section 6), but have to contend with the following issues:

Firstly, they are *intrusive* requiring, to varying degrees, modifications to the optimizer engine. Secondly, they require *specialized* information about the workload and/or the system which may not always be easy to obtain or model. Thirdly, their query capabilities may be *limited* compared to the original optimizer – e.g., only SPJ queries with key-based joins were considered in [2, 3]. Further, [3] has been implemented and evaluated on a non-commercial optimizer. Finally and most importantly, as explained in Section 6, none of them provide, on an individual query basis, quantitative *guarantees* on the quality of their final plan choice relative to the original (unmodified) optimizer’s selection. That is, they “cater to the crowd, not individuals”.

The SEER Algorithm. In this paper, we present **SEER** (Selectivity-Estimate-Error-Resistance), a new strategy for identifying robust plans that can be directly used on operational database environments. More concretely, it

- Treats the optimizer as a black-box and therefore is inherently (a) completely non-intrusive, and (b) capable of handling whatever SQL is supported by the system. Further, it does not expect to have any additional information beyond that provided by the engine interface.
- Provides plan performance guarantees that are *individually* applicable to queries in the selectivity space.
- Considers only the *parametric optimal set of plans* (POSP) [12] as replacement candidates and therefore delivers, for errors that lie within the replacement plan’s optimality region, robustness “for free”. In contrast, the previously proposed algorithms in the literature may evaluate plans that are not optimal anywhere in the space.
- Is validated on *commercial* optimizers on both the classical TPC-H [25] and the recent TPC-DS [26] benchmarks.

We hasten to add that SEER, due to its non-intrusive design objective, only attempts to address selectivity errors that occur on the *base relations*, similar to [1]. However, since these base errors are often the source of poor plan choices due to the multiplier effect as they progress up the plan-tree [14], minimizing their impact could be of significant value in practical environments. Further, since SEER is a purely compile-time approach, it can be used in conjunction with run-time techniques such as adaptive query processing [9] for addressing selectivity errors in the higher nodes of the plan tree.

Anorexic Reduction of Plan Diagrams. SEER is based on the *anorexic reduction of plan diagrams*, a notion that was recently presented and analyzed in [11]. Specifically, a “plan diagram” [21] is a color-coded pictorial enumeration of the plan choices of the optimizer for a parametrized query template over the relational selectivity space. That is, it visually captures the POSP geometry.

For example, consider QT8, the parametrized 2D query template shown in Figure 1, based on Query 8 of TPC-H. Selectivity variations on the SUPPLIER and LINEITEM relations are specified through the `s_acctbal :varies` and `l_extendedprice :varies` predicates, respectively. The associated plan diagram for QT8 is shown in Figure 2(a), produced with the Picasso optimizer visualization tool [20] on a popular commercial database engine.

As evident from Figure 2(a)¹, plan diagrams can be extremely complex and dense, with a large number of plans covering the space – several such instances spanning a representative set of benchmark-based query templates on industrial-strength optimizers are available at [20]. However, these dense diagrams can typically be “reduced” to much simpler pictures featuring significantly fewer plans, *without materially degrading the processing quality of any individual query*. For example in Figure 2(a), if users are willing to tolerate a minor cost increase (λ) of at most 10% for any query point in the diagram, relative to its original cost, the picture could be reduced to Figure 2(b), where only 7 plans remain – that is, most of the original plans have been “completely swallowed” by their siblings, leading to a highly reduced plan cardinality.

¹The figures in this paper should ideally be viewed from a color copy, as the grayscale version may not clearly register the features.

```

select o_year, sum(case when nation = 'BRAZIL' then volume else 0 end) / sum(volume)
from (select YEAR(o_orderdate) as o_year, l_extendedprice * (1 - l_discount) as volume,
n2.n_name as nation

from part, supplier, lineitem, orders, customer,
nation n1, nation n2, region

where p_partkey = l_partkey and s_suppkey = l_suppkey and l_orderkey = o_orderkey
and o_custkey = c_custkey and c_nationkey = n1.n_nationkey and n1.n_regionkey =
r_regionkey and s_nationkey = n2.n_nationkey and r_name = 'AMERICA' and p_type
= 'ECONOMY ANODIZED STEEL' and
s_acctbal :varies and l_extendedprice :varies

) as all_nations

group by o_year
order by o_year

```

Figure 1: Example Query Template: QT8

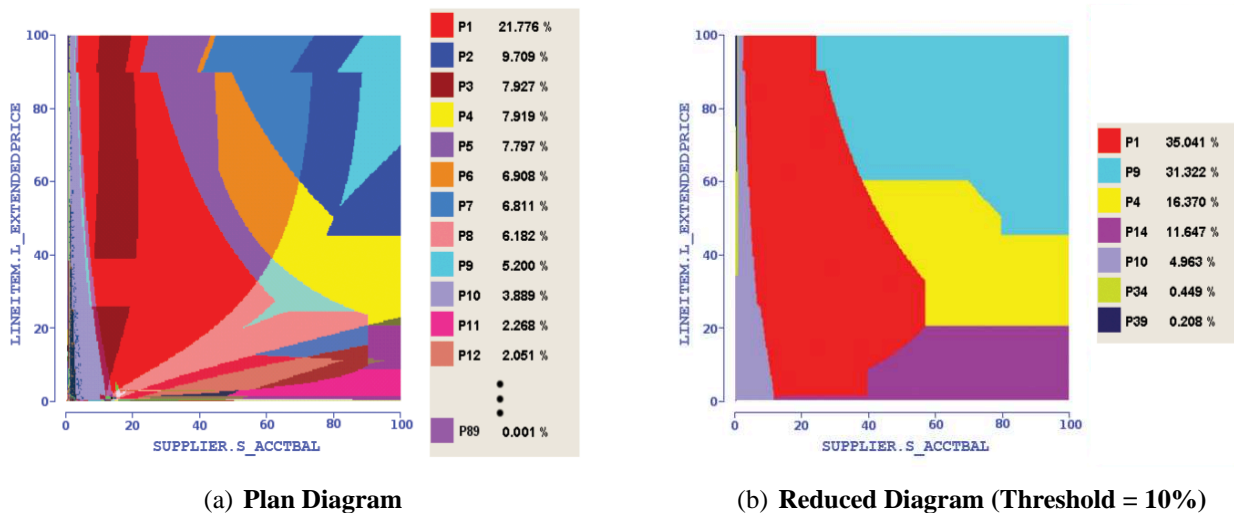


Figure 2: Sample Plan Diagram and Reduced Plan Diagram (QT8)

A detailed study of the plan diagram reduction problem was presented in [11], and it was shown that a cost increase threshold of *only 20 percent* is usually amply sufficient to bring down the absolute number of plans in the final reduced picture to *within or around ten*. In short, that complex plan diagrams can be made “anorexic” while retaining acceptable query processing performance.

Example. We now show an example of how anorexic reduction helps to identify selectivity-error-resistant plans: In Figure 2(a), estimated selectivities of around (14%,1%) lead to a choice of plan P70. However, if the actual selectivities at runtime turn out to be significantly different, say (50%,40%), executing with P70, whose cost increases steeply with selectivity, would be disastrous. In contrast, this error would have had no impact with the reduced plan diagram of Figure 2(b), since P1, the replacement plan choice at (14%,1%), remains the preferred plan for a large range of higher values, including (50%,40%). Quantitatively, at the run-time location, plan P1 has a cost of 135, while P70’s cost of 402 is about *three times* more expensive.

It is easy to see, as in the above example, that the replacement plan will, *by definition*, be a robust choice for errors that lie within its optimality region, i.e. its “*endo-optimal*” region. This is the advantage, mentioned earlier, of considering replacements only from the POSP set of plans. The obvious question then is whether the sizes of these regions are typically large enough to materially improve the system performance.

A second, and even more important question, is: What if the errors are such that the run-time locations are “*exo-optimal*” w.r.t. the replacement plan? For example, if the run-time location happens to be at (80%,90%), which is outside the optimality region of P1? In this situation, nothing can be said upfront – the replacement could be much better, similar or much worse than the original plan. Therefore, ideally speaking, we would like to have a mechanism through which one could assess whether a replacement is *globally safe* over the entire parameter space.

Contributions. In this paper, we address the above issues from both theoretical and empirical perspectives. We have conducted extensive experimentation on a leading commercial optimizer with a rich suite of multi-dimensional TPC-H and TPC-DS based query templates operating on a variety of logical and physical database designs. Our results demonstrate that *plan diagram reduction typically produces plan choices that substantially curtail the adverse effects of selectivity estimation errors*. Therefore, it clearly has potential to improve performance in general, for both the endo-optimal and exo-optimal regions.

However, we have also encountered occasional situations where a replacement plan performs much worse in its exo-optimal region than the original optimizer choice, highlighting the need to establish an efficient criterion of when a specific swallowing is globally safe. To achieve this objective, we present a generalized mathematical model of the behavior of plan cost functions over the selectivity space. The model, although simple, is sufficient to capture the cost behavior of all plans that have arisen from our query templates. Using this model, we then prove that checks on only the *perimeter* of the selectivity space are sufficient to decide the safety of reduction over the entire space. These checks involve the costing of “*foreign plans*”, that is, of costing plans in their exo-optimal regions, a feature that has become available in the current versions of several industrial-strength optimizers, including DB2[27] (Optimization Profile), SQL Server[28] (XML Plan) and Sybase[29] (Abstract Plan).

Apart from providing reduction safety, foreign-plan costing is additionally leveraged to both (a) enhance the reduction levels of the plan diagram, and (b) improve the complexity characteristics of

the reduction process, as compared to our earlier CostGreedy reduction algorithm [11]. Note that the increased diagram reduction automatically implies *larger within- λ -of-optimal regions* for the retained plans, upfront guaranteeing more robustness.

In summary, we present in this paper SEER, an efficient, effective and safe mechanism for identifying robust plans that are resistant, as compared to the optimizer’s original choices, to errors in the base relation selectivity estimates. Through a detailed study of benchmark-based query templates on commercial optimizers, we empirically demonstrate that SEER provides robust good performance for industrial-strength database environments. We also present LiteSEER, a (complexity-wise) optimally efficient heuristic algorithm, which delivers comparable robustness to that offered by SEER. Both SEER and LiteSEER have been implemented in the freely-available Picasso optimizer visualization tool [20].

Organization. The remainder of this paper is organized as follows: In Section 2, we present the overall problem background, framework and motivation. The plan cost models and the checks for replacement safety are discussed in Section 3. The design of the SEER reduction algorithm and its analysis are presented in Section 4. Our experimental framework and performance results are highlighted in Section 5. Related work is overviewed in Section 6. Finally, in Section 7, we summarize our conclusions and outline future research avenues.

2 Problem Framework

For ease of exposition, we assume in the following discussion that the SQL query template is 2-dimensional in its selectivity variations – the extension to higher dimensions is straightforward.

2.1 Plan and Reduced Plan Diagrams

From a query template Q , a plan diagram P is produced on a 2-dimensional $[0, 1]$ selectivity space S by making repeated calls to the optimizer. The selectivity space is represented by a grid of points where each point $q(x, y)$ corresponds to a unique query with selectivities x, y in the X and Y dimensions, respectively. Each q is associated with an optimal (as determined by the optimizer) plan P_i , and a cost $c_i(q)$ representing the estimated effort to execute q with plan P_i . Corresponding to each plan P_i is a unique color L_i , which is used to color all the query points that are assigned to P_i . As mentioned earlier, the plan diagram is essentially a visual characterization of the parametric optimal set of plans (POSP) [12]. We use P and S interchangeably in the remainder of the paper based on the context.

Plan Diagram Reduction Problem. This problem is defined as follows [11]: Given an input plan diagram P , and a maximum-cost-increase threshold λ ($\lambda \geq 0$), find a reduced plan diagram R with *minimum cardinality* such that for every plan P_i in P ,

1. Either $P_i \in R$, or

2. $\forall q \in P_i$, the assigned replacement plan $P_j \in R$ guarantees $\frac{c_j(q)}{c_i(q)} \leq (1 + \lambda)$

That is, find the maximum possible subset of the plans in P that can be completely “swallowed” by their sibling plans in the POSP set. A point worth reemphasizing here is that the threshold constraint

applies on an *individual query* basis. For example, setting $\lambda = 10\%$ stipulates that the cost of *each* query point in the reduced diagram is within 1.1 times its original value.

It was proved in [11] that the above problem is NP-Hard. Therefore, an efficient heuristic-based online algorithm, called **CostGreedy**, was proposed and shown to deliver near-optimal “anorexic” levels of reduction, wherein the plan cardinality of the reduced diagram usually came down to around 10 or less for a λ -threshold of only 20%. In a nutshell, complex plan diagrams can be easily made very simple without materially affecting the query processing quality.

2.2 Selectivity Estimation Errors

Consider a specific query point q_e , whose optimizer-estimated location in \mathbf{S} is (x_e, y_e) . Denote the optimizer’s optimal plan choice at point q_e by P_{oe} . Due to errors in the selectivity estimates, the *actual* location of q_e could be different at execution-time – denote this location by $q_a(x_a, y_a)$, and the optimizer’s optimal plan choice at q_a by P_{oa} . Assume that P_{oe} has been swallowed by a sibling plan during the reduction process and denote the replacement plan assigned to q_e in \mathbf{R} by P_{re} . Finally, extend the definition of query cost (which applied to the optimal plan) to have $c_i(t)$ denote the cost of an arbitrary POSP plan P_i at an arbitrary query point t in \mathbf{S} .

With respect to \mathbf{R} , the actual query point q_a will be located in one of the following disjoint regions of P_{re} that together cover \mathbf{S} :

Endo-optimal region of P_{re} : Here, q_a is located in the optimality region of the replacement plan P_{re} , which also implies that $P_{re} \equiv P_{oa}$. Since $c_{re}(q_a) \equiv c_{oa}(q_a)$, it follows that the cost of P_{re} at q_a , $c_{re}(q_a) < c_{oe}(q_a)$ (by definition of a cost-based optimizer). Therefore, improved resistance to selectivity errors is always *guaranteed* in this region.

Swallow-region of P_{re} : Here, q_a is located in the region “swallowed” by P_{re} during the reduction process. Due to the λ -threshold constraint, we are assured that $c_{re}(q_a) \leq (1 + \lambda)c_{oa}(q_a)$, and by implication that $c_{re}(q_a) \leq (1 + \lambda)c_{oe}(q_a)$. Now, there are two possibilities: If $c_{re}(q_a) < c_{oe}(q_a)$, then the replacement plan is again guaranteed to improve the resistance to selectivity errors. On the other hand, if $c_{oe}(q_a) \leq c_{re}(q_a) \leq (1 + \lambda)c_{oe}(q_a)$, the replacement is guaranteed to not cause any real harm, given the small values of λ that we consider in this paper.

Exo-optimal region of P_{re} : Here, q_a is located outside both the endo-optimal and swallow-regions of P_{re} . At such locations, we cannot apriori predict P_{re} ’s behavior, and therefore the replacement may not always be a good choice – in principle, it could be *arbitrarily worse*. Therefore, we would like to ensure that even if the replacement does not provide any improvement, it is at least guaranteed to not do any harm. That is, the *exo-optimal region should have the same performance guarantees as the swallow-region*. We show in Section 3 how this objective can be efficiently achieved through simple but powerful checks to decide when replacement is advisable.

2.3 Motivational Scenarios

Given the above framework, we now present example scenarios to motivate (a) the error-resistance utility of plan diagram reduction, and (b) the need for safety in this process.

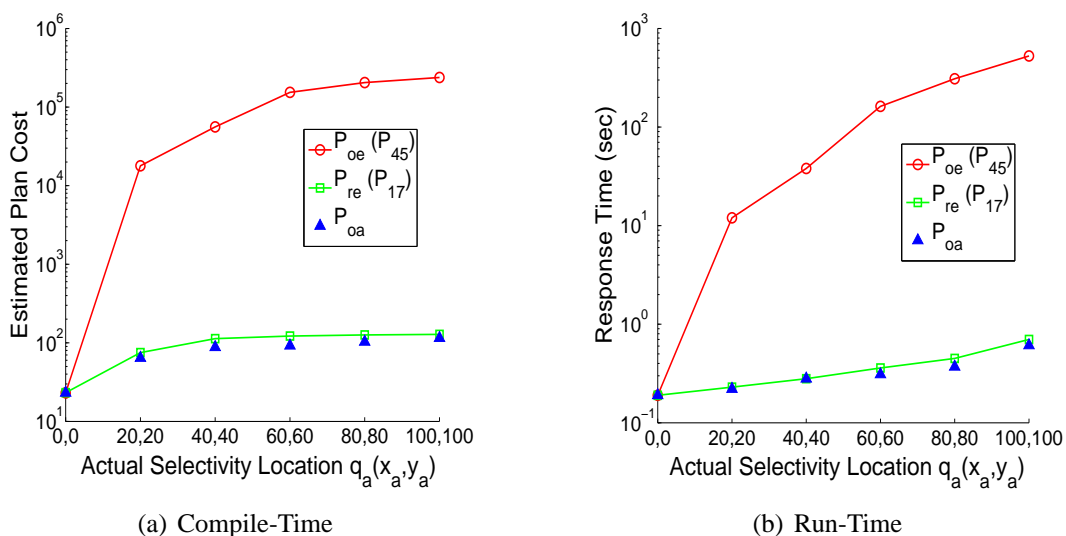


Figure 3: Beneficial Impact of Plan Replacement

Our first scenario, typical of that seen in most of our experiments, demonstrates how the replacement plan P_{re} can provide extremely substantial improvements *throughout the selectivity space*. Specifically, on a vanilla PC with a popular commercial optimizer, we generated a plan diagram for a query template based on TPC-H Q5, with selectivity variations on the CUSTOMER and SUPPLIER relations, and carried out reduction with $\lambda = 10\%$. On this diagram, consider $q_e = (0.36, 0.05)$ and a sample set of actual locations (q_a) – for instance, along the principal diagonal of \mathbf{S} . For this scenario, the costs of P_{oe} (P45), P_{re} (P17) and P_{oa} (the optimal plan at each q_a location) are shown in Figure 3(a) – note that the costs are measured on a *log scale*.

It is clear from Figure 3(a) that the replacement plan P_{re} provides *orders-of-magnitude* benefit w.r.t. P_{oe} . In fact, the error-resistance is to the extent that it virtually provides “immunity” to the error since the performance of P_{re} is close to that of the *locally optimal plan* P_{oa} throughout the space, although the endo-optimal region of P_{re} constitutes only a very small fraction of this space.

To demonstrate that the benefits anticipated from the compile-time analysis do translate to corresponding improvements *at runtime*, we show in Figure 3(b) the query response times (again measured on a *log scale*) of P_{oe} (P45), P_{re} (P17) and P_{oa} at the same q_a locations. It is vividly clear in this picture that huge savings in processing time are obtained by using the replacement plan instead of the optimizer’s original choice, and that the replacement’s performance is virtually indistinguishable from the optimal choices.

While performance improvements are usually the order of the day, there are occasional situations wherein P_{re} performs worse than P_{oe} at q_a . A particularly egregious example, arising from the *same* plan diagram described above, is shown in Figure 4(a) for $q_e = (0.03, 0.14)$ – we see here that it is now the replacement plan P_{re} (P34), which is *orders-of-magnitude* worse than P_{oe} (P26) in the presence of selectivity errors. This compile-time assessment is corroborated in Figure 4(b) which shows the corresponding query response times.

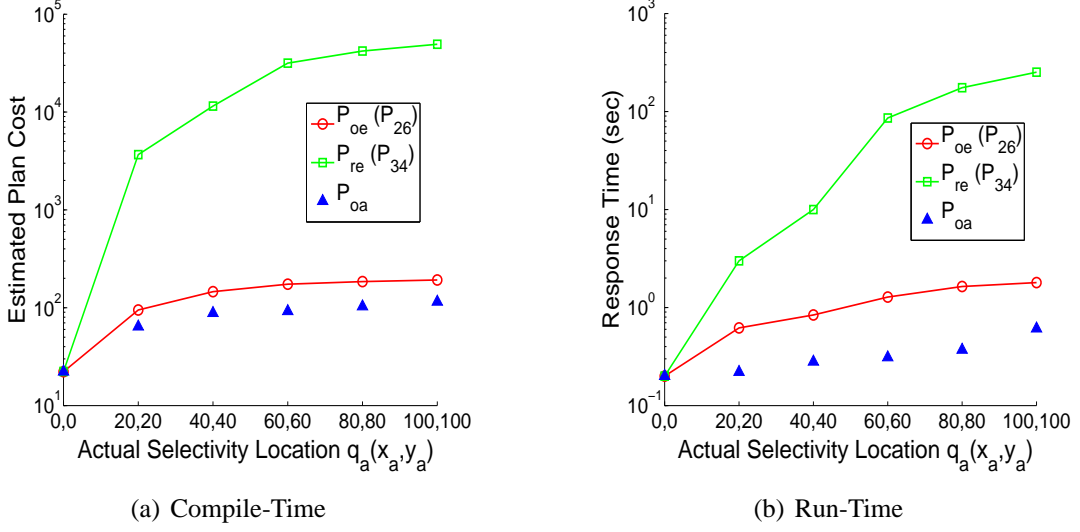


Figure 4: Adverse Impact of Plan Replacement

2.4 Robust Reduction

From the above discussion, it is clear that we need to ensure that only safe replacements are permitted. This means that replacement should be permitted only if the λ threshold criterion is satisfied not just at the estimated point, but *at all locations* in the selectivity space. At the same time, it is important to ensure that the safety check is not unnecessarily conservative, preventing most plan replacements, and in the process losing all the error-resistance benefits. Therefore, the overall goal is to maximize plan diagram reduction without violating safety considerations. More formally, our problem formulation is:

Robust Reduction Problem. Given an input plan diagram P , and a maximum cost-increase-threshold λ ($\lambda \geq 0$), find a reduced plan diagram R with *minimum plan cardinality* such that for every plan P_i in P ,

1. $P_i \in R$, or
2. $\forall q \in P_i$, the assigned replacement plan $P_j \in R$ guarantees \forall query points $q' \in P$, $\frac{c_j(q')}{c_i(q')} \leq (1 + \lambda)$

That is, find the minimum-sized error-resistant “cover” of plans that reduces the plan diagram P without increasing the cost of any reassigned query point by more than the cost increase threshold, *irrespective of the actual location of the query at run-time*.

It is easy to see that the Robust Reduction problem is NP-Hard, just like the standard Plan Diagram Reduction problem, and therefore we present a heuristic-based algorithm later in Section 4. But, prior to that, we show in the following section how replacement safety can be checked efficiently.

3 Ensuring Robust Reduction

To find an error-resistant cover of the plan diagram, we need to evaluate the behavior of each replacement plan P_{re} , w.r.t. its swallowing target P_{oe} , at *all points* in \mathbf{S} . This requires, in principle, finding the costs of P_{oe} and all potential P_{re} at every point in the diagram. Of course, P_{oe} and P_{re} need not be costed in their respective *endo-optimal* regions, since these values are already known through the plan diagram production process. The remaining *exo-optimal* costs can be obtained using the *Foreign-Plan-Costing* feature, hereafter referred to as **FPC**, that is now supported in several industrial-strength optimizers, as mentioned in the Introduction.

While the above solution is conceptually feasible, it is practically unviable due to its enormous computational overheads. Plan-costing is certainly cheaper than the optimizer’s standard optimal-plan-searching process [13], but the overall overhead is still $O(nm)$ where n and m are the number of plans and the number of points, respectively, in \mathbf{P} . Typical values of n range from the several tens to several hundreds, while m is at least in thousands, making an exhaustive approach impractical.

The above situation motivates us to study whether it is possible, based on using FPC at only a *few select locations*, to *infer* the behavior in the rest of the space. In the remainder of this section, we describe our strategy for making such an inference. We begin by designing a parametrized mathematical model for characterizing plan cost behavior. Our model is grossly simplified in comparison to those used in real optimizers, which are much more complex [19, 18]. However, what we have found in practice (with several hundred distinct plans arising out of TPC-H and TPC-DS-based query templates on industrial optimizers) is that with appropriate settings of the parameters, our simple model is quite accurate, both behaviorally and quantitatively. The reasons are that (a) in our problem space all parameters, barring the selectivities, are *constant*, resulting in complex models degenerating to comparatively simple equivalents; (b) we are *fitting* the model to the observed cost behaviors, rather than trying to *predict* them; and (c) our modeling is at the level of entire plans, aggregating the effects of several individual operators, thereby reducing the variability. Moreover, the quantitative accuracy is a bonus – it is not really required since only *behavioral* accuracy is necessary for our scheme to work.

3.1 Modeling Plan Cost Functions

For ease of presentation, we will initially assume that our objective is to model the cost behavior of plans with respect to a 2-D selectivity space (e.g. Figure 2(a)) corresponding to distinct relations R_x and R_y . The extension to higher dimensions is straightforward and is provided later in this section.

In current optimizers, the operators in the execution plan are all typically either *unary* or *binary* with regard to their inputs. Therefore, given a specific plan operator tree, like the sample one shown in Figure 5 (obtained on a $R_x = \text{LINEITEM}$, $R_y = \text{CUSTOMER}$ selectivity space), we can define the following types of nodes:

Selectivity Nodes: These are the unary nodes that implement the selection operations on relations R_x and R_y . In Figure 5, they are colored orange, corresponding to Index Scans on the LINEITEM and CUSTOMER relations, respectively.

Dependent Nodes: These are the nodes in the tree that have at least one Selectivity Node in the subtree below them. They are colored blue in Figure 5.

Node Type	Input Cardinalities		Cost Model
Selectivity Node ($\sigma = x$)	n		$a_1nx + a_2$
Dependant Unary Nodes	n_1x		$a_1n_1x + a_2$
			$a_1n_1x \log n_1x + a_2$
	n_1xy		$a_1n_1xy + a_2$
			$a_1n_1xy \log n_1xy + a_2$
Dependant Binary Nodes	n_1x	n_2	$a_1n_1x + a_2n_2 + a_3n_1n_2x + a_4$
	n_1xy	n_2	$a_1n_1xy + a_2n_2 + a_3n_1n_2xy + a_4$
	n_1x	n_2y	$a_1n_1x + a_2n_2y + a_3n_1n_2xy + a_4$

Table 1: Cost Models for Various Node Types

logarithmic model would apply to operators such as *Sort* and *Group By* that require multiple passes over the data.

Dependent Binary Nodes: These are the nodes that represent binary set operators such as *Join*, *Union*, *Minus*, etc. The different types of input possibilities and the associated cost models are shown in Table 1.

Note that we deliberately do not consider the case where *both* the inputs to the binary node are functions of x (or y or xy). This is because it is easy to prove that such a situation is not possible unless operators have *binary outputs* – we have not encountered any such operators in our study.

Lemma 1 *There cannot be a binary node that has both inputs to be functions of x (or functions of y , or functions of xy).*

Proof: If there exists a binary node N with input cardinalities n_1x and n_2x , then there should exist some node in its subtree that has a *binary output*. However, we know that all nodes in the plan tree have unary outputs (since there is no cycle in the tree). A similar argument holds for the y and xy cases. ■

3.2.1 Cost Model of a Complete Plan

The cost function of the entire plan is the aggregate sum of the costs of the individual nodes. Considering all possible cost models a node can have, we can conclude that the overall cost model of a plan for a 2D selectivity space is of the form

$$Cost(x, y) = a_1x + a_2y + a_3xy + a_4x \log x + a_5y \log y + a_6xy \log xy + a_7 \quad (1)$$

where $a_1, a_2, a_3, a_4, a_5, a_6, a_7$ are coefficients, and x, y represent the selectivities of R_x and R_y , respectively.

Modeling a specific plan requires suitably choosing the seven coefficients, and this is achieved through standard surface-fitting techniques, described in Section 5.

3.3 Extension to d-dimensional spaces

Generalizing the arguments used in the 2D case, we obtain the following cost model for a d -dimensional selectivity space.

$$\begin{aligned}
Cost(x_1, \dots, x_d) = & \sum_{i_1} (a_{i_1} x_{i_1} + b_{i_1} x_{i_1} \log x_{i_1}) + \\
& \sum_{i_1 < i_2} (a_{i_1 i_2} x_{i_1} x_{i_2} + b_{i_1 i_2} x_{i_1} x_{i_2} \log x_{i_1} x_{i_2}) \\
& + \dots + a_{12..d}(x_1 x_2 x_3 \dots x_d) \\
& + b_{12..d}(x_1 x_2 x_3 \dots x_d) \log(x_1 x_2 x_3 \dots x_d) \\
& + a_0
\end{aligned} \tag{2}$$

where the a 's and b 's are the $(2^{d+1} - 1)$ coefficients and the $x_i, i = 1 \dots d$ represent the d relational selectivities.

3.4 Replacement Safety Conditions

For the 2D scenario, using the above 7-coefficient cost model, our goal now is to come up with an efficient mechanism to assess, given an optimal plan P_{oe} , candidate replacement plan P_{re} and a cost-increase threshold λ , whether it would be safe from a *global* perspective to have P_{re} swallow P_{oe} .

Let the cost functions for P_{re} and P_{oe} be

$$f_{re}(x, y) = a_1 x + a_2 y + a_3 xy + a_4 x \log x + a_5 y \log y + a_6 xy \log xy + a_7 \tag{3}$$

and

$$f_{oe}(x, y) = b_1 x + b_2 y + b_3 xy + b_4 x \log x + b_5 y \log y + b_6 xy \log xy + b_7 \tag{4}$$

respectively. Now consider the **“safety function”**

$$f(x, y) = f_{re} - (1 + \lambda) f_{oe} \tag{5}$$

which captures the differences between the costs of P_{re} and a λ -inflated version of P_{oe} in the selectivity space. All points where $f(x, y) \leq 0$ are referred to as *SafePoints* whereas points that have $f(x, y) > 0$ are called *ViolatingPoints*. For a replacement to be globally safe, there should be no *ViolatingPoint* anywhere in the selectivity space.

In the following, we will use LR-Boundaries to collectively denote the left and right boundaries of the selectivity space, and TB-Boundaries to collectively denote the top and bottom boundaries of the space.

For a specific value of y , the safety function $f(x, y)$ can be rewritten as

$$f_y(x) = g_1 * x + g_2 * x \log x + g_3$$

for appropriate coefficients g_1, g_2, g_3 . Similarly, we can define $f_x(y)$. With this terminology, the following theorem provides us with conditions for checking whether the selectivity space is safe for the plan-pair (P_{oe}, P_{re}) with regard to replacement.

Theorem 1 For a plan-pair (P_{oe}, P_{re}) and a selectivity space \mathbf{S} with corners $[(x_1, y_1), (x_1, y_2), (x_2, y_2), (x_2, y_1)]$, the replacement is safe (i.e., within λ -threshold) in \mathbf{S} if any one of the conditions, SC1 through SC6, given in Table 2 is satisfied.

	Left Boundary	Right Boundary	Top Boundary	Bottom Boundary
SC1	Safe	Safe	$f''_{y_2}(x) \geq 0$	$f''_{y_1}(x) \geq 0$
SC2	$f'_y(x_1) \leq 0$ & Safe	Safe	$f''_{y_2}(x) < 0$	$f''_{y_1}(x) < 0$
SC3	Safe	$f'_y(x_2) \geq 0$ & Safe	$f''_{y_2}(x) < 0$	$f''_{y_1}(x) < 0$
SC4	$f''_{x_1}(y) \geq 0$	$f''_{x_2}(y) \geq 0$	Safe	Safe
SC5	$f''_{x_1}(y) < 0$	$f''_{x_2}(y) < 0$	$f'_x(y_2) \geq 0$ & Safe	Safe
SC6	$f''_{x_1}(y) < 0$	$f''_{x_2}(y) < 0$	Safe	$f'_x(y_1) \leq 0$ & Safe

Table 2: Safety Satisfaction Conditions

In order to prove the above theorem, we will start with deriving two lemmas – the first provides us with a condition that is sufficient to ensure safety of all points on the straight line segment joining a pair of safe points, while the second describes the behaviour of the slope of the safety function.

Lemma 2 (Line Safety) Given a fixed $y = y_o$, and a pair of safe points (x_1, y_o) and (x_2, y_o) with $x_2 > x_1$, the straight line joining the two points is safe if the slope $f'_y(x)$ is either

- (i) monotonically non-decreasing, OR
- (ii) monotonically decreasing with $f'_{y_o}(x_1) \leq 0$ or $f'_{y_o}(x_2) \geq 0$

A similar result holds when x is fixed.

Proof: The various possible behaviors of $f_y(x)$ are shown in Figure 6 as Curves (a) through (e). When the slope $f'_y(x)$ is monotonically non-decreasing (i.e. Condition (i) is satisfied), the safety function curve that connects the two safe points is guaranteed to lie *below* the straight line joining the two points – Curve (a) in Figure 6 shows an example of this situation. This ensures that the safety function along the given line segment is always negative and hence safe.




Figure 6: Behavior of the safety function $f_y(x)$

If, on the other hand, $f'_{y_o}(x)$ is monotonically decreasing, then the possible behaviors of the safety function $f_{y_o}(x)$ are shown in Curves (b) through (e) in Figure 6. Curves (b) and (c) denote the behaviour of the safety function when Condition (ii) is satisfied, and clearly the value of the safety function is again negative in the given range. ■

In Figure 6, Curve (d) also corresponds to a safe scenario – however, it is not possible to differentiate between Curve (d) and the unsafe case, namely Curve (e), without explicitly computing the safety function at every point on the given line-segment. Hence, we *conservatively* categorize both cases as unsafe. We have also observed that the case corresponding to Curve (e) occurs rarely in practice.

Lemma 3 (Slope Behavior) *If the slope of the safety function, $f'_y(x)$, is non-decreasing (resp. decreasing) along the line-segments $y = y_1$ and $y = y_2$, then it is non-decreasing (resp. decreasing) for all line segments in the interval (y_1, y_2) . A similar result holds for $f'_x(y)$.*

Proof: Consider the slope of the safety function

$$f'_{y_o}(x) = \frac{df_{y_o}(x)}{dx} = g_1 + g_2(1 + \log x) \quad (6)$$

For $x \in (0, 1)$, this slope is monotonic and its behavior depends on the sign of g_2 . From Equations 3 and 4, we know that g_2 can be written as the following function of y

$$\begin{aligned} g_2(y) &= (a_4 - (1 + \lambda)b_4) + (a_6 - (1 + \lambda)b_6)y \\ &= (k_1 + k_2y) \end{aligned} \quad (7)$$

where k_1 and k_2 are constants.

Since $g_2(y)$ is a linear function of y , the Lemma immediately follows. ■

We now prove Theorem 1 using the LineSafety and SlopeBehavior lemmas:

Proof: Consider the SC1 condition in Table 2: Since $f''_y(x) \geq 0$ (i.e. slope $f'_y(x)$ is non-decreasing) at the TB-boundaries, then from Lemma 3, we know that the slope $f'_y(x)$ is non-decreasing throughout the range (y_1, y_2) .

Moving on to the SC2 and SC3 conditions: Since $f''_y(x) < 0$ (i.e. slope $f'_y(x)$ is decreasing) at the TB-boundaries, then from Lemma 3, we know that the slope $f'_y(x)$ is decreasing throughout the range (y_1, y_2) . Further, we know that for a given $y = y_o \in (y_1, y_2)$, either $f'_{y_o}(x_1) \leq 0$ (SC2) or $f'_{y_o}(x_2) \geq 0$ (SC3).

Thus, when SC1, SC2 or SC3 is satisfied, then for all lines between points (x_1, y) and (x_2, y) , $y \in (y_1, y_2)$, the end-points are safe (because the LR-boundaries are safe), and the slope conditions given in Lemma 2 are satisfied. Hence, all such line-segments are safe, the union of which is the given region.

Similar arguments can be used to show safety of the region when conditions SC4, SC5 or SC6 are satisfied. Hence the theorem. ■

The test criteria of Theorem 1 are utilized for determining reduction safety in the SafetyCheck algorithm, described next. A related point to note here is that these checks are *conservative* in that it is possible to have global safety even if none of the conditions are met – i.e. the test is sufficient, but not necessary.

4 The SEER Algorithm

In this section, we first describe the safety checking procedure, which given a plan-pair (P_{oe}, P_{re}) , responds whether the replacement of P_{oe} by P_{re} is globally safe throughout the selectivity space \mathbf{S} . We then present and analyze the SEER algorithm which uses this procedure to perform error-resistant plan diagram reduction.

In the following, we will assume that the selectivity space \mathbf{S} is represented by a grid \mathbf{G} , with $m = r \times r$ points, i.e. the grid resolution in each dimension is r .

4.1 Safety Checking

To implement safe reduction in a 2-D plan diagram, we need to be able to check for the satisfaction of any of the conditions (SC1 through SC6) stipulated in Theorem 1. A straightforward way to achieve this is the following *Perimeter Test* procedure:

Perimeter Test. First compute the safety function at all points on the *perimeter* of \mathbf{G} – this is obtained through the foreign-plan-costing (FPC) feature. Then, compute the slope behavior (non-decreasing or decreasing) along all the grid lines – this is achieved by evaluating the slopes at the matching end-points on the perimeter and comparing the values. The slope at a perimeter point is approximated by computing the value of the safety function at its immediate internal neighbor – i.e., along the “inner perimeter”, and evaluating the slope of the line segment joining these two points. Finally, use these results to verify whether any of the 6 safety conditions are satisfied.

In the Perimeter test, the number of FPC operations is $2 * 4(r - 1)$ for the perimeter (the 2 is due to having to compute both f_{re} and f_{oe}), while the computation of the slopes takes an additional $2 * 4(r - 3)$ costings of the inner perimeter, leading to a total of approximately $16r$. Note that this is much less than the $2r^2$ FPC operations required by a brute-force approach of costing both plans at all points in the diagram. For example, with $r = 100$, the overhead is brought down by over an order of magnitude. The red and blue points shown in Figure 7 are to be costed in this test.

An obvious minor improvement that could be carried out on the $16r$ overhead is to perform the inner perimeter costings only when conditions SC1 and SC4 are violated. In this case, only one of SC2 or SC3 (resp. SC5 or SC6) can be valid. Hence, we need to perform FPC operations only at *two* boundaries of the inner perimeter, one along each dimension. This reduces the FPC overhead to $12r$.

Wedge Test. We now present a powerful optimization, called *Wedge Test*, that allows conditions SC1 and SC4 to be checked with a *constant* number of FPC, specifically 24, *irrespective of the resolution*. This is based on the observation that the slope of the safety function is a monotonic function (Equation 6). Thus, by comparing the slopes at the corners of the space, we can infer the slope behaviour of the safety function along its boundaries. Applying Lemma 2, the safety of the boundaries can also be inferred. Hence, it is sufficient to perform FPC only at each corner of the space and its two adjacent points on the perimeter boundaries – that is, at the “corner wedges”. Only the red points shown in Figure 7 are to be costed in this test.

Based on the above observations, we employ a two-stage process of safety-checking – in the first stage, use the extremely cheap Wedge Test check, and only if it fails, use the more expensive Perimeter Test to verify replacement safety.




Figure 7: Perimeter and Wedge Test

Note that once a plan is costed at a given location, we store this cost in a cache for reuse later, ensuring no redundant computations.

4.2 Plan Diagram Reduction

We now show how the above safety checks are integrated into the SEER procedure for plan diagram reduction. Note that SEER's design is completely different from that of CostGreedy [11] because now reduction is permitted only if it satisfies a safety criterion that is applicable over \mathbf{S} , whereas CostGreedy's attention is limited to only P_{oe} 's endo-optimal region.

The complete SEER algorithm is shown in Figure 8. Here, a Set-Cover instance is first created from the input plan diagram \mathbf{P} . Then the two-stage global safety checking procedure of the Wedge Test, followed by the Perimeter Test, is implemented to evaluate replacement possibilities across each pair of plans in \mathbf{P} , and the Set-Cover instance is updated accordingly. Finally, the resulting instance is solved using the standard greedy techniques [22, 10] to obtain the reduced plan diagram \mathbf{R} .

Analysis. As discussed earlier, each replacement assessment of a plan-pair (P_{oe}, P_{re}) requires at most $O(r)$ FPCs to be performed. There are $O(n^2)$ such comparisons performed by the algorithm. However, since we cache the already obtained costs, the amortized number of FPC to be performed per plan is $O(r)$. Thus, for grid \mathbf{G} with $m = r \times r$ points, the comparison of all plan pairs requires only $O(n\sqrt{m} + n^2)$ time. Solving the Set-Cover problem using the Greedy Set-Cover algorithm [22, 10] requires $O(n^2)$ time. This results in an $O(n\sqrt{m} + n^2)$ reduction algorithm. Further, since the set cover instance created has $|U| = n$, the approximation factor of this reduction algorithm is $O(\log n)$.

The above bounds and approximation factors for SEER compare very favorably with those of the CostGreedy reduction algorithm [11], which has time complexity $O(nm)$ and approximation factor of $O(\log m)$, since typically $n \ll m$.

SEER (Plan Diagram P , Threshold λ)

1. Create a Set-Cover Instance $I = (U, S)$, where $S = \{S_1, S_2, \dots, S_n\}$, $U = \{1, 2, \dots, n\}$, corresponding to the n plans in the original plan diagram P .
2. Set each $S_i = \{i\}, \forall i = 1 \dots n$
3. For each pair of plans (P_i, P_j) do
 - if (WEDGE_TEST $(P_i, P_j, \lambda) == \text{Safe}$) then
 - $S_i = S_i \cup \{j\}$
 - else if (PERIMETER_TEST $(P_i, P_j, \lambda) == \text{Safe}$) then
 - $S_i = S_i \cup \{j\}$
4. Solve the Set-Cover instance I using the Greedy Set-Cover algorithm to identify the plans retained in R .

Figure 8: The SEER Reduction Algorithm

4.2.1 Extension to Higher Dimensions

The SafetyCheck algorithm used to verify the safety of the replacement of P_{oe} by P_{re} in a d -dimensional selectivity space is given in Figure 9.

SafetyCheck (Plan Diagram P , Threshold λ , Plan P_{oe} , Plan P_{re} , Dimension d)

1. if($d == 2$)
 - (a) if (WEDGE_TEST $(P, P_i, P_j, \lambda) == \text{Safe}$) then
 - return true.**
 - else if (PERIMETER_TEST $(P, P_i, P_j, \lambda) == \text{Safe}$) then
 - return true.**
 - (b) **return false.**
2. else
 - (a) $safety = true;$
 - (b) for each $(d - 1)$ -dimension slice P' of P
 - $safety = safety \wedge \text{SafetyCheck}(P', \lambda, P_{oe}, P_{re}, d - 1)$
3. **return safety.**

Figure 9: n-Dimensional SafetyCheck Algorithm

The above algorithm recursively finds the safe area of the $(d - 1)$ -dimension “slices” of the input d -dimension selectivity space. When $d = 2$, the WEDGE_TEST and PERIMETER_TEST methods are used to check for safety. The SEER algorithm incorporating this checking mechanism is shown in Figure 10.

SEER (Plan Diagram \mathbf{P} , Threshold λ)

1. Create a Set-Cover Instance $I = (U, S)$, where $S = \{S_1, S_2, \dots, S_n\}$, $U = \{1, 2, \dots, n\}$, corresponding to the n plans in the original plan diagram \mathbf{P} .
2. Set each $S_i = \{i\}$, $\forall i = 1 \dots n$
3. For each pair of plans (P_i, P_j) do
 - if (SafetyCheck ($\mathbf{P}, \lambda, P_i, P_j, d$) == true) then

$$S_i = S_i \cup \{j\}$$
4. Solve the set-cover instance I using the Greedy Setcover algorithm to identify the plans retained in \mathbf{R} .

Figure 10: n-Dimensional SEER Reduction Algorithm

Each d -dimensional plan diagram is composed of r $(d - 1)$ -dimensional plan diagrams. The time complexity of the SafetyCheck algorithm for the base case when $d = 2$ is $O(r)$. Thus the SafetyCheck algorithm runs in $O(r^{d-1})$ time. Given a grid with $m = r^d$ points, FPC is performed $O(m^{\frac{d-1}{d}})$ times for each plan pair. Thus, the time complexity of the SEER reduction algorithm for a d -dimensional plan diagram is $O(m^{\frac{d-1}{d}}n + n^2)$.

4.3 Variants of SEER

4.3.1 LiteSEER: A Fast Variant

The SEER design makes conscious efforts, as described above, to minimize the computational overheads, but these overheads do grow with increasing dimensionality of the query template. Therefore, we have also designed and evaluated LiteSEER, a light-weight heuristic-based algorithm that trades SEER's safety guarantee for providing rapid running-times. In LiteSEER, a replacement is simply assumed to be safe if *all the corner points of the selectivity space are safe*. The intuition behind this observation is that when two points are safe, then the straight line joining them is also usually safe. This is corroborated by our experimental results (Section 5) which indicate that the heuristic provides almost the same safety as that obtained through the strict-checking criteria of SEER.

Given a d -dimensional plan diagram \mathbf{P} with n plans, the LiteSEER algorithm only computes the safety function at the 2^d corners of the associated selectivity space. It immediately follows that its overall complexity is $O(2^d n + n^2)$. Since, in most practical scenarios of interest, $2^d \ll n$ (e.g. in the 2-D case, $2^d = 4$, while n is typically in the several tens, if not more), the effective complexity turns out to be $O(n^2)$. Note that, in principle, in the absence of any apriori information, this is the *minimum work* required to be executed by *any* reduction algorithm. Therefore, LiteSEER is an optimal algorithm (complexity-wise) w.r.t. efficiency.

4.3.2 PartialSEER

The problem formulation for robust reduction required the replacement plan to be *globally safe*. As a generalized variant, the safety criteria can be relaxed to allow a plan P_{re} to replace plan P_{oe} if P_{re} is safe in at least a user-defined *minimum safe fraction (MSF)* of the area covered by \mathbf{S} ($MSF \leq 1$).

PartialSafetyCheck (Plan Diagram \mathbf{P} , Threshold λ , Area $allowedViolation$, Plan P_{oe} , Plan P_{re} , Dimension d)

1. if($d == 2$)
 - (a) if (WEDGE_TEST ($\mathbf{P}, P_{oe}, P_{re}, \lambda$) == Safe) **return** $allowedViolation$
 - (b) if (PERIMETER_TEST ($\mathbf{P}, P_{oe}, P_{re}, \lambda$) == Safe) **return** $allowedViolation$
 - (c) if the slope criteria of the six conditions of Theorem 1 are not satisfied, **return** -1
 - (d) if ($allowedViolation = 0$) **return** -1
 - (e) if no *two* adjacent boundaries are safe, **return** -1 .
 - (f) Let the first violating point at the top-boundary of the grid \mathbf{G} occur at $x = x_v$. Set $x = x_v$, $y = r - 1$, $NumViolatingPoints = 0$
 - (g) While $x \neq r$ and $y \neq -1$
 - i. Set $count = 0$
 - ii. While current point is violating (i.e. $f(x, y) > 0$) and $y \neq -1$
 - A. move down (i.e. $y--$)
 - B. if ($NumViolatingPoints + (r - y - 1) \times (r - x - 1)$) $> allowedViolation$, **return** -1
 - iii. While current point is safe (i.e. $f(x, y) \leq 0$) and $x \neq r$
 - A. move right (i.e. $x++$), $count++$
 - B. if ($NumViolatingPoints + count \times (r - y - 1)$) $> allowedViolation$, **return** -1
 - iv. $NumViolatingPoints + = count \times (r - y - 1)$
 - (h) $allowedViolation - = NumViolatingPoints$
 - (i) **return** $allowedViolation$
2. else
 - (a) for each $(d - 1)$ -dimension slice \mathbf{P}' of \mathbf{P}
 - i. $allowedViolation = \text{PartialSafetyCheck}(\mathbf{P}', \lambda, allowedViolation, P_{oe}, P_{re}, d - 1)$
 - ii. if ($allowedViolation < 0$) **return** $allowedViolation$;
3. **return** $allowedViolation$.

Figure 11: The PartialSafetyCheck Algorithm




Figure 12: The PartialSafety-Check Algorithm

PartialSEER (Plan Diagram P , Threshold λ , MinSafeFraction MSF)

1. Create a Set-Cover Instance $I = (U, S)$, where $S = \{S_1, S_2, \dots, S_n\}$, $U = \{1, 2, \dots, n\}$, corresponding to the n plans in the original plan diagram P .
2. Set each $S_i = \{i\}, \forall i = 1 \dots n$
3. For each pair of plans (P_i, P_j) do
 - (a) Set $allowedViolation = (1 - MSF) \times Area(P)$.
 - (b) if $(PartialSafetyCheck(P, \lambda, allowedViolation, P_i, P_j, d) \geq 0)$ then

$$S_i = S_i \cup \{j\}$$
4. Solve the set-cover instance I using the Greedy Setcover algorithm to identify the plans retained in R .

Figure 13: The PartialSEER Reduction Algorithm

In order to assess partial safety, we first perform the WEDGE_TEST and PERIMETER_TEST checks for global safety. If this fails, we verify whether the slope criteria of any of the 6 conditions given in Theorem 1 is satisfied. If true, we allow plan P_{re} to replace plan P_{oe} if

1. At least two adjacent boundaries in the perimeter of S are safe; and
2. The MSF requirement is met in S .

The reason for restricting our attention to situations where at least two adjacent boundaries are safe is that, for this case, an efficient algorithm can be set up to check satisfaction of the area requirement, as described below. Figure 11 shows the modified SafetyCheck algorithm that finds the safe area when the left and bottom boundaries of S are safe. The algorithm is similar when other boundaries are safe.

From Theorem 1 we know that the safe (and violating) points form contiguous regions in S when the slope criteria of at least one of the size conditions are satisfied. Since the left and bottom boundaries of the grid are safe, the x and y axes form a part of the boundary of the safe region. The PartialSafetyCheck algorithm traces the remainder of this boundary.

Figure 12 shows the flow of the algorithm while tracing the boundary between the safe (green) and violating (red) regions of the selectivity space for a pair of plans. In this figure, the top and right boundaries of the region violate the safety requirement.

We start from the first violating point on the top-boundary of the grid, and at each stage either move down or right in the grid. At each interior point that we move to, we perform the costing of the plans P_{oe} and P_{re} . The algorithm stops when we reach the bottom or right boundaries of the grid.

The number of *right* or *down* movements required to reach this termination situation is atmost $2r$ movements for a $r \times r$ grid. Hence, for a pair of plans, atmost $4r$ extra costings are needed to obtain the error-resistant area. Steps 1 through 3 of the algorithm require an additional $12r$ costings in the

worst case scenario – it is usually much smaller. Thus, the overall time complexity of the modified PartialSafetyCheck algorithm for a 2-dimensional selectivity space is $O(r)$.

The PartialSEER reduction algorithm, which employs the PartialSafetyCheck safety-checking technique, is shown in Figure 13.

5 Experimental Results

The testbed used in our experiments is the Picasso optimizer visualization tool [20], executing on a Sun Ultra 20 workstation equipped with an Opteron Dual Core 2.5GHz processor, 4 GB of main memory and 720 GB of hard disk, running the Windows XP Pro operating system. The experiments were conducted over plan diagrams produced from a variety of two and three-dimensional **TPC-H** and **TPC-DS**-based query templates operating on the OptCom commercial optimizer. The TPC-H database contains *uniformly* distributed data of size 1GB, while the TPC-DS database hosts *skewed* data that occupies 100GB. The cost-increase threshold used in all the plan diagram reductions is $\lambda = 20\%$.

Physical Design. Following a methodology similar to that advocated in [5], we considered three different physical design configurations in our study: **PrimaryKey (PK)**, **AllIndex (AI)**, and **TunedIndex (TI)**. PK represents the default physical design of our database engine, wherein a clustered index is created on each primary key. AI, on the other hand, represents an “index-rich” situation wherein (single-column) indices are available on all query-related schema attributes. Finally, TI represents the index environment obtained by implementing the recommendations of the database engine’s index tuning advisor (which include multi-column indices).

In the subsequent discussion, we use QT x to refer to a query template based on Query x of the TPC-H benchmark, and DSQT x to refer to a query template based on Query x of the TPC-DS benchmark, operating in the default PK configuration. We prefix AI and TI to the query template identifiers in describing our results for these specialized configurations.

Query Location Distribution. All the performance results shown initially in this section are for plan diagrams generated with *exponentially* distributed locations for the query points across the selectivity space, resulting in higher query densities near the selectivity axes and towards the origin. This choice is based on earlier observations in the literature (e.g. [12, 13, 21]) that plans tend to be densely packed in precisely these regions of the selectivity space. From a performance perspective, these diagrams represent the “tough-nut” challenging situations with respect to obtaining anorexic reduction due to their high plan densities and substantially broader range of plan cost values.

For completeness, we have also conducted all the experiments with a *uniform* distribution of query locations – these results, which are qualitatively similar to those presented here, are detailed in Section 5.6.

Performance Metrics. In the remainder of this section, we evaluate the SEER reduction algorithm with regard to the following performance parameters: (a) Diagram Reduction Quality, (b) Error-resistance obtained through Reduction, (c) Safety of Reduction, and (d) Computational Efficiency. As a precursor, we first evaluate the validity of the plan cost function model (Section 3.1).

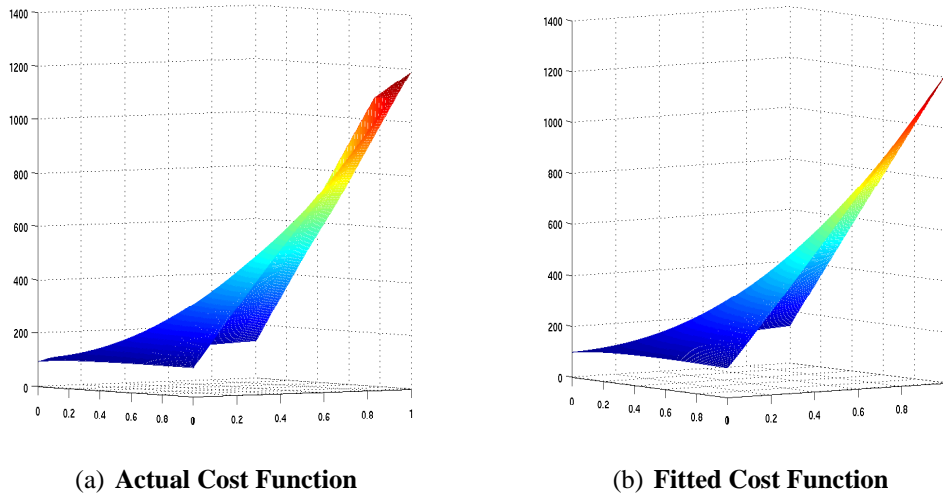


Figure 14: Plan Cost Function Modeling

5.1 Validity of Plan Cost Model

The validity of the plan cost model presented in Equation 1 was assessed by attempting to fit the costs of plans generated by OptCom². The experimental data consisted of optimizer-estimated execution costs over the selectivity space of the plans that appeared in the various plan diagrams (taken from both exponentially and uniformly distributed query templates). As mentioned earlier, the foreign-plan-costing (FPC) feature was used to evaluate plans outside of their endo-optimal regions.

The surface fitting was carried out with the classical Linear Least Squares method [16] and implemented using Matlab 7.4 [24]. An example 2-D fitted cost function is:

$$Cost(x, y) = 17.9x + 45.9y + 1046xy - 39.5x \log x + 4.5y \log y + 27.6xy \log xy + 97.3$$

For this plan, the complete plots of the actual cost surface and the fitted cost surface, as a function of the selectivities of the two base relations, are shown in Figure 14. It is visually evident that the fit is very good.


As further evidence of the accuracy of our model, Table 3 shows the quality-of-fit, measured in terms of the maximum and average *Root-Mean-Square(RMS)* errors, over a large number of plans featuring in the plan diagrams arising from our suite of multi-dimensional query templates. The consistently low RMS values suggest that the model is sufficiently accurate for our purposes.

Finally, as an additional precaution, we deliberately searched for plan cost functions with complex shapes to assess the quality-of-fit in these difficult cases. An example is shown in Figure 15, and we see that even here, the fit is of high quality (the RMS Error is only around 10%). This can be attributed to the fact that our cost model has 7 parameters which gives sufficient freedom to fit most of the plan cost functions found in practice. Our curve fitting technique does not impose any restriction on the behaviour of the cost function, and hence we see for this example, a small PCM violation in the fitted curve.

²We have also validated this plan cost model on another commercial database, and found the results to be similar

Dimension	Number of Plans	Maximum RMS Error (%)	Average RMS Error(%)
2D (TPC-H)	614	14.20	1.82
2D (TPC-DS)	168	7.31	2.87
3D (TPC-H)	28	6.98	1.92
3D (TPC-DS)	100	2.71	1.58

Table 3: RMS Errors in Fitted Cost Surfaces



(a) Actual Cost Function

(b) Fitted Cost Function

$$Cost(x, y) = 184.3x + 619.9y + 524.5xy - 1090x \log x - 1179.9y \log y - 836.2xy \log xy - 1000$$

$$RMSError = 10.96\%$$

Figure 15: Complex Plan Cost Function

5.2 Plan Diagram Reduction Quality

A potentially worrisome aspect of our quest to obtain globally robust reduction is whether it might result in losing out on the anorexic reduction levels observed in the localized reduction processes of [11]. This concern is quantitatively allayed in Table 4, which presents a comparison between SEER and CostGreedy (CG) of the number of plans in the reduced diagram for a diverse suite of multi-dimensional query templates on the TPC-H database. The PK physical design configuration was operational in these experiments.

At first glance, SEER might have been expected to perform worse than CostGreedy because its additional safety checks may prevent some plan swallowings permitted by CostGreedy— in fact, this was the source of our concern. However, in Table 4, we actually find the *converse* – while CostGreedy does provide anorexic reduction, SEER does even better. The reason for this is that CostGreedy follows a conservative cost-bounding approach to estimate the costs of plans outside their endo-optimal regions (details in [11]). SEER, on the other hand, uses the foreign-plan-costing feature to obtain the exact

Query Template	Original No. of plans	CG	CG-FPC	SEER	LiteSEER
QT2 (2D)	60	14	3	6	6
QT5 (2D)	51	7	2	2	2
QT8 (2D)	121	7	2	2	2
QT9 (2D)	137	9	3	4	4
QT10 (2D)	44	3	3	3	3
QT16 (2D)	32	11	3	3	3
QT5 (3D)	68	8	3	3	3
QT8 (3D)	191	8	3	3	3
QT10 (3D)	75	10	3	4	4

Table 4: Plan Diagram Reduction Quality (TPC-H)

Query Template	Original No. of plans	CG	CG-FPC	SEER	LiteSEER
DSQT12 (2D)	25	6	3	2	2
DSQT18 (2D)	114	13	2	2	2
DSQT19 (2D)	55	11	3	4	4
DSQT12 (3D)	33	11	2	2	2
DSQT18 (3D)	222	15	2	4	4
DSQT19 (3D)	98	15	2	4	4

Table 5: Plan Diagram Reduction Quality (TPC-DS)

costs in these regions, and therefore has superior reduction possibilities. Therefore, the FPC feature comes in handy from both quality and safety perspectives.

A question that immediately arises is how SEER would compare against a CostGreedy variant that also utilized the FPC feature. This issue is also addressed in Table 4, where the performance of this variant (CG-FPC) is presented. We see that CG-FPC does perform better or as well as SEER, as should be expected – however, the gap, if any, is always very small. A related point to note here is that the SEER reduction quality remains excellent even for the 3D query templates, in spite of the fact that the additional dimension increases the possibility of the safety conditions being violated.

Finally, we observe in Table 4 that the LiteSEER fast variant happens to provide reduction quality identical to SEER. Under the AI (and TI) configurations, however, it occasionally performs slightly better (see Section 5.2.1), as should be expected due to its being less stringent in allowing replacements.

TPC-DS Results. The above results were generated on the TPC-H database, which has uniformly distributed data. Table 5 shows a corresponding set of results for plan diagrams generated on the TPC-DS database, which features skewed data. It is immediately evident that the reduction profiles of the various reduction algorithms are very similar to those seen with TPC-H.

Query Template	Original No. of plans	CG	CG-FPC	SEER	LiteSEER
AIQT2 (2D)	87	12	2	2	2
AIQT5 (2D)	126	14	4	6	5
AIQT8 (2D)	121	7	3	3	3
AIQT9 (2D)	132	14	3	4	3
AIQT10 (2D)	37	8	4	5	5
AIQT16 (2D)	35	9	2	2	2
AIQT5 (3D)	139	14	5	7	5
AIQT8 (3D)	168	14	4	6	5
AIQT10 (3D)	77	16	7	8	8

Table 6: Plan Diagram Reduction Quality (TPCH-AI)

5.2.1 Reduction Quality with AllIndex Configuration

While the PK configuration had only 8 primary-key indices, AllIndex includes an additional 53 (non-clustered) single-column indices covering all the remaining query-related schema attributes. The reduction quality results for this index-rich configuration are shown in Table 6. We first notice that the number of plans in the original diagram usually increases, often substantially, as should be expected since the optimizer’s search space has increased due to the availability of the additional indices. For example, the number of plans for AIQT5(2D) goes up to 125 from 51, while AIQT5(3D) jumps to 139 from 68. However, when we consider the reduction quality of the various algorithms, we find that they continue to *materially adhere to anorexic levels*, although the actual cardinalities may have gone up by a couple of plans. For example, SEER on AIQT5(2D) retains 6 plans as compared to 2 under PK.

Another point to note in Table 6 is that we now see LiteSEER occasionally permitting slightly greater reduction than SEER, due to its relaxed constraint in allowing replacements.

5.2.2 Reduction Quality with TunedIndex Configuration

The reduction quality results for the *Tuned Index* (TI) configuration which implements the recommendations of the index tuning advisor shipped with OptCom is shown in Table 7. The parameters of the tuning advisor were set to their default values, and the TPC-H benchmark queries (generated with the QGen utility) formed the input workload. For this setup, the advisor recommended 20 additional indices beyond the default Primary Key configuration.

We see here that the reduction performance is very similar to that obtained with the PK and AI configurations, testifying to SEER’s consistent behavior over a wide variety of database environments.

5.3 Error-resistance and Safety

Having established the retention of diagram reduction quality, we now move on to assessing the extent to which resistance to selectivity errors is provided through SEER reduction. We begin with defining a metric that quantitatively captures this effect:

Query Template	Original No. of plans	CG	CG-FPC	SEER	LiteSEER
TIQT2 (2D)	52	10	4	5	5
TIQT8 (2D)	108	16	3	3	3
TIQT9 (2D)	101	16	6	5	5
TIQT10 (2D)	50	10	4	2	2
TIQT16 (2D)	36	14	4	6	6
TIQT5 (3D)	84	10	4	5	5
TIQT8 (3D)	181	14	4	6	5
TIQT10 (3D)	78	12	6	8	7

Table 7: Plan Diagram Reduction Quality (TPCH-TI)

Error Resistance Metric. Given an estimated query location q_e and an actual location q_a , the *Selectivity Error Resistance Factor* (**SERF**) of a replacement plan P_{re} w.r.t. the optimal plan P_{oe} is defined as,

$$SERF(q_e, q_a) = 1 - \frac{c_{re}(q_a) - c_{oa}(q_a)}{(1 + \lambda)c_{oe}(q_a) - c_{oa}(q_a)}$$

Intuitively, SERF captures the fraction of the performance gap between P_{oe} and P_{oa} that is closed by P_{re} . In principle, SERF values can range over $(-\infty, 1]$, with the following interpretations: SERF in the range $(\lambda, 1]$ indicates that the replacement is beneficial, with values close to 1 implying “immunity” to the selectivity error. For SERF in the range $[0, \lambda]$, the replacement is indifferent in that it neither helps nor hurts, while SERF values below 0 highlight a harmful replacement that materially worsens the performance.

The above formula applies to a specific instance of replacement. To capture the net impact of reduction on improving the resistance in an *entire plan diagram*, we compute the following

$$AvgSERF = \frac{\sum_{q_e \in rep(P)} \sum_{q_a \in exo_{oe}(P)} SERF(q_e, q_a)}{\sum_{q_e \in rep(P)} \sum_{q_a \in exo_{oe}(P)} 1}$$

where $rep(P)$ is the set of points in the plan diagram P that were replaced during the reduction process, and $exo_{oe}(P)$ is the set of points lying in the exo-optimal region defined with respect to P_{oe} , the optimizer’s plan choice for q_e . The normalization is with respect to the number of possible selectivity errors in the diagram. (To ensure meaningful AvgSERF values from a robustness perspective, we exclude the uninteresting scenarios wherein both c_{re} and c_{oe} have extremely low absolute values, or are both within λ -threshold of c_{oa} .)

Note that in the above formulation, we assume for simplicity that the actual location q_a is equally likely to be anywhere in P_{oe} ’s exo-optimal space, that is, that the errors are uniformly distributed over this space. However, our conceptual framework is also applicable to the more generic case where the error locations have an associated probability distribution.

Resistance Results. For CostGreedy, SEER and LiteSEER, we show in Table 8, the AvgSERF, as defined above, as well as MinSERF and MaxSERF, the minimum and maximum values of SERF over all replacement instances, for the various query templates. We first see here that for all the algorithms,

Query Template	CG			SEER			LiteSEER		
	MinSERF	AvgSERF	MaxSERF	MinSERF	AvgSERF	MaxSERF	MinSERF	AvgSERF	MaxSERF
QT2 (2D)	-58.6	0.2	1	0	0.38	0.98	0	0.38	0.98
QT5 (2D)	-15.6	0.5	1	0.14	0.51	0.99	0.14	0.51	0.99
QT8 (2D)	-2.9	0.82	1	0.3	0.93	1	0.3	0.93	1
QT9 (2D)	-46.3	0.37	1	0	0.77	1	0	0.77	1
QT10 (2D)	-23.8	0.09	1	0	0.35	1	0	0.35	1
QT16 (2D)	-132.5	0.03	0.99	0.04	0.33	0.96	0.04	0.33	0.96
QT5 (3D)	-15.4	0.6	1	0.03	0.38	1	0.03	0.38	1
QT8 (3D)	-177.7	-0.5	1	0	0.63	1	0	0.63	1
QT10 (3D)	-130.7	-2	1	0	0.45	1	0	0.45	1
AIQT2 (2D)	-202	-1.1	0.99	0	0.8	0.99	0	0.8	0.99
AIQT5 (2D)	-1336	-3.1	1	0	0.58	1	-10	0.54	1
AIQT8 (2D)	-62.4	0.35	1	0	0.54	1	0	0.54	1
AIQT9 (2D)	-9486	-3.1	1	0	0.62	1	-5	0.66	1
AIQT10 (2D)	-20.2	0.24	1	0	0.25	0.98	0	0.25	0.98
AIQT16 (2D)	-76	0.07	1	0	0.75	1	0	0.75	1
AIQT5 (3D)	-2151	0.24	1	0.05	0.62	1	-2	0.66	1
AIQT8 (3D)	-103.7	0.43	1	0	0.44	1	-6	0.44	1
AIQT10 (3D)	-4680	-4.4	1	0	0.33	1	0	0.33	1
TIQT2 (2D)	-161.3	-0.12	0.99	0.01	0.32	0.94	0.01	0.32	0.94
TIQT8 (2D)	-550.1	-3.4	1	0	0.54	1	0	0.54	1
TIQT9 (2D)	-87.5	0.01	1	0	0.44	1	0	0.44	1
TIQT10 (2D)	-671.4	-15.7	1	0	0.27	0.99	0	0.27	0.99
TIQT16 (2D)	-14.4	0.11	0.99	0.02	0.39	0.97	0.02	0.39	0.97
TIQT5 (3D)	-166.3	-1.42	1	0	0.71	1	0	0.71	1
TIQT8 (3D)	-4188.5	-4.23	1	0	0.61	1	-1.3	0.61	1
TIQT10 (3D)	-574	-4.97	1	0	0.31	1	-8.9	0.31	1
DSQT12 (2D)	-4	0.61	1	0.07	0.44	1	0.06	0.44	1
DSQT18 (2D)	-194.75	-2	1	0.07	0.81	1	0.07	0.81	1
DSQT19 (2D)	-86	0.17	1	0	0.63	1	0	0.63	1
DSQT12 (3D)	-142	-0.32	1	0.06	0.53	1	0.06	0.53	1
DSQT18 (3D)	-3104	-0.14	1	0	0.85	1	0	0.85	1
DSQT19 (3D)	-106.6	0.33	1	0.02	0.82	1	0.02	0.82	1

Table 8: Characterization of Error-Resistance through Reduction

plan diagram reduction is capable, across the board, of providing complete immunity (MaxSERF tending to 1) to selectivity errors for individual replacement instances. Secondly, and more importantly, the AvgSERF is also quite substantial for SEER. For example, in DSQT18, on average, more than three-quarters of the performance gap due to selectivity errors is bridged by the SEER reduction process.

With CostGreedy, on the other hand, the AvgSERF is comparatively very poor, and occasionally even negative! The important point to note here is that these low averages are an artifact arising out of a small fraction of points (around 10-20% points occurring with probability of around 0.1) whose performance is grossly adversely affected by plan replacement. That is, plan reduction does help in the vast majority of cases but the “few very bad apples”, reflected by the hugely negative MinSERF values (which sometimes even run into the thousands), ruin the overall performance statistics. More pertinently, these results serve to quantitatively and vividly substantiate the need for safe replacement, the motivation underlying our design of the SEER algorithm.

Finally, turning our attention to LiteSEER, we see that its error-resistance profile is very similar to that of SEER – in fact, the AvgSERF and MaxSERF numbers are identical for most templates. Further, although like CostGreedy it does not guarantee safety, as testified to by the negative values in the MinSERF column, note that (a) the templates having negative values are relatively rare, (b) even in these cases, unsafe replacements occur only for about 1% of the points (with probability ≤ 0.01), and (c) most importantly, their magnitudes are small in comparison (the maximum is -10 for AIQT5(2D)).

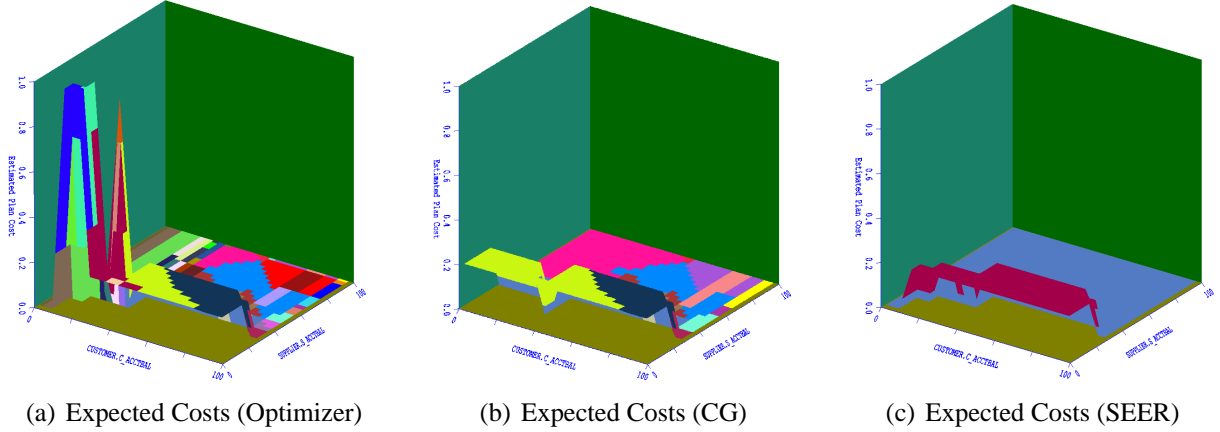


Figure 16: Safe Error-resistance with SEER

Safety Example

In the example of Figure 4, plan diagram reduction without explicitly checking for safety led to situations wherein P_{re} performed much worse than P_{oe} at q_a . The effectiveness of SEER in avoiding such unsafe replacements is visually highlighted in the sequence of pictures in Figure 16, corresponding to the same example.

Assuming that the actual location of a query at run-time q_a is uniformly distributed over S , Figure 16(a) shows the *expected cost* for each query point q_e , when executed with its optimizer-selected plan P_{oe} . Note that the peaks in the picture correspond to situations where the plan-choice is highly sensitive to selectivity errors.

Then, Figure 16(b) shows the expected cost of each query point q_e when executed with P_{re} from the reduced plan diagram obtained using CostGreedy. Note that virtually all the peaks in Figure 16(a) are substantively eliminated through the replacement choices in the reduced plan diagram – for example, the dark-blue peak at the left-top corner of Figure 16(a) is largely removed. However, on the down side, some plans suffer injurious replacements – for e.g., the earth-brown colored plan in the left-bottom corner of Figure 16(a) is now replaced by the fluorescent-green colored plan in Figure 16(b), whose expected cost is orders of magnitude greater. That is, CostGreedy in the process of eliminating existing peaks, may introduce *new peaks*.

Finally, in Figure 16(c), we show the performance of SEER reduction. We see here that (a) it removes all the peaks of Figure 16(a) like CostGreedy, and (b) it does not introduce any new peaks courtesy its safety criterion. In a nutshell, “it provides virtually all the good, and doesn’t introduce any harm”.

5.4 Efficiency of Reduction Process

We now move on to profiling the time taken to complete the reduction process by SEER as compared to CostGreedy. These results are shown in Table 9 for our query template suite.

Focussing initially on the 2D query templates, we see that SEER’s performance is quite acceptable in terms of absolute times (a few minutes per reduction), especially in comparison to the original plan

Query Template	CG (ms)	CG-FPC (min)	SEER (min)	LiteSEER (sec)
QT2 (2D)	15	53.1	3.6	14.2
QT5 (2D)	16	45.0	1.0	12
QT8 (2D)	17	108	9.6	28.8
QT9 (2D)	13	122.4	10.6	32.6
QT10 (2D)	15	38.7	3.0	10.3
QT16 (2D)	15	27.9	1.3	7.5
QT5 (3D)	25	67	19.0	32
QT8 (3D)	21	190	65.0	91
QT10 (3D)	17	74	16.5	4.5
AIQT2 (2D)	17	77.4	5.0	20.6
AIQT5 (2D)	12	112.5	3.7	30.0
AIQT8 (2D)	11	108.0	6.9	28.8
AIQT9 (2D)	18	107.9	9.1	31.4
AIQT10 (2D)	12	32.4	2.0	8.6
AIQT16 (2D)	12	30.6	2.0	8.2
AIQT5 (3D)	26	138	37.7	66.2
AIQT8 (3D)	19	167	47.3	80.2
AIQT10 (3D)	24	76	14.9	36.5
TIQT2 (2D)	18	45.9	2.9	12.2
TIQT8 (2D)	12	96.3	4.9	25.7
TIQT9 (2D)	16	90.0	7.2	24.0
TIQT10 (2D)	14	44.1	2.6	11.8
TIQT16 (2D)	12	31.5	2.0	8.4
TIQT5 (3D)	28	83	20.8	39.8
TIQT8 (3D)	24	180	67.8	86.4
TIQT10 (3D)	19	78	15.9	37.0
DSQT12 (2D)	14	21.6	2.6	5.8
DSQT18 (2D)	13	101.7	9.4	27.1
DSQT19 (2D)	14	48.6	6.4	13.0
DSQT12 (3D)	20	32.0	7.4	15.4
DSQT18 (3D)	25	221.0	89.1	106.1
DSQT19 (3D)	23	97.0	35.8	46.6

Table 9: Efficiency of Reduction Process

diagram *production time*. However, it is much slower relative to CostGreedy, which offers sub-second response times. This might seem surprising in light of our analysis in Section 4 showing that SEER is an $O(n\sqrt{m} + n^2)$ algorithm, whereas CostGreedy is $O(nm)$. The reason for the higher running time of SEER is that the basic cost-bounding computation in CostGreedy is much faster than the foreign-plan-costing operator provided by the commercial optimizers. Our discussions with the development team of OptCom have indicated that this is not due to the costing itself, but is largely an artifact of

setting up the contexts for the costing, including verifying the validity of the plan with respect to the query. Therefore, it is possible that future better implementations of the FPC feature may bring SEER’s running time closer to CostGreedy. (In fact, our own implementation of FPC in a public-domain optimizer indicates that its cost can be further brought down by an *order of magnitude* [7].)

When we consider the 3D query templates, however, the running times of SEER can be quite large. It is here that LiteSEER shows its worth since its running times are only a few minutes or even less, across the board for all the query templates. Taken in conjunction with its good safety performance (Section 5.3), it suggests that LiteSEER offers an extremely attractive compromise between the speed of CostGreedy and the robustness of SEER, making it a viable first-cut reduction technique in real-world installations.

Finally, to normalize the effect of the different costing implementations, the running time of the CG-FPC algorithm is also shown in Table 9 – we see here that CG-FPC takes in the order of *several tens or few hundreds of minutes* to complete the reduction process. In comparison, SEER’s selective usage of the FPC operator, courtesy Theorem 1 and the two-stage checking process, does succeed in substantially bringing down the overheads.

5.5 Performance of PartialSEER

Sample results obtained with the PartialSEER algorithm for $MSA = 0.8$ are shown in Table 10, 11 and 12, for the various metrics of reduction quality, error resistance and efficiency of the reduction.

Query Template	Original No. of plans	PartialSEER
QT5(2D)	51	1
QT16(2D)	32	3
DSQT18(2D)	114	2

Table 10: Plan Diagram Reduction Quality ($\lambda = 20\%$, $MSA = 0.8$)

Query	MinSERF	AvgSERF	MaxSERF
QT5(2D)	-15.6	0.37	1
QT16(2D)	-0.25	0.37	0.99
DSQT18(2D)	-0.26	0.83	1

Table 11: Error-Resistance of PartialSEER ($\lambda = 20\%$, $MSA = 0.8$)

Query	No of FPC	Time (min)
QT5(2D)	8738	8.7
QT16(2D)	5046	5.0
DSQT18(2D)	27284	27.3

Table 12: Efficiency of PartialSEER ($\lambda = 20\%$, $MSA = 0.8$)

5.6 Uniform Query Distribution

The results shown thus far were produced with an exponential distribution of query points across the selectivity space. We present here the corresponding results for plan diagrams generated with a *uniform* distribution of query points. Tables 13 and 14 show the reduction quality over our suite of query templates on the TPC-H and TPC-DS databases, respectively, operating with a Primary Key physical configuration. The performance on an All Index configuration is detailed in Table 15. Finally, the error-resistance quality and the reduction efficiency are shown in Tables 16 and 17, respectively.

These results are behaviorally similar to those obtained with the exponential distribution.

Query Template	Original No. of plans	CG	CG-FPC	SEER	LiteSEER
QT2 (2D)	25	5	3	3	3
QT5 (2D)	10	3	1	1	1
QT8 (2D)	31	4	2	2	2
QT9 (2D)	21	2	1	1	1
QT10 (2D)	13	3	2	2	2
QT16 (2D)	26	9	2	3	3
QT5 (3D)	18	1	1	1	1
QT8 (3D)	18	6	3	3	3
QT10 (3D)	18	4	2	2	2

Table 13: Plan Diagram Reduction Quality (TPC-H)

Query Template	Original No. of plans	CG	CG-FPC	SEER	LiteSEER
DSQT12 (2D)	7	4	2	2	2
DSQT18 (2D)	21	3	1	1	1
DSQT19 (2D)	28	5	2	2	2
DSQT12 (2D)	8	2	1	1	1
DSQT18 (3D)	36	2	1	1	1
DSQT19 (2D)	64	2	1	1	1

Table 14: Plan Diagram Reduction Quality (TPC-DS)

Query Template	Original No. of plans	CG	CG-FPC	SEER	LiteSEER
AIQT2 (2D)	30	8	3	3	3
AIQT5 (2D)	25	6	2	2	2
AIQT8 (2D)	25	3	2	3	3
AIQT9 (2D)	25	5	1	1	1
AIQT10 (2D)	16	4	3	3	3
AIQT16 (2D)	22	14	3	4	4
AIQT5 (3D)	37	4	2	2	2
AIQT8 (3D)	39	5	2	3	3
AIQT10 (3D)	50	9	4	3	3

Table 15: Plan Diagram Reduction Quality (TPCH-AI)

6 Related Work

Over the last decade, a variety of *compile-time* strategies have been proposed to identify robust plans. For example, in the Least Expected Cost (LEC) approach [6, 8], it is assumed that the distribution of predicate selectivities is a priori available, and then the plan that has the least-expected-cost over the distribution is chosen for execution. While the performance of this approach is likely to be good on average, it could be arbitrarily poor for a specific query as compared to the optimizer’s optimal choice for that query. Moreover, it may not always be feasible to provide the selectivity distributions.

An alternative Robust Cardinality Estimation (RCE) strategy proposed in [2] is to model the selectivity dependency of the cost functions of the various competing plan choices. Then, given a user-specified “confidence threshold” T , the plan that is expected to have the *least upper bound* with regard to cost in

Query Template	CG			SEER			LiteSEER		
	MinSERF	AvgSERF	MaxSERF	MinSERF	AvgSERF	MaxSERF	MinSERF	AvgSERF	MaxSERF
QT2 (2D)	0	0.76	0.98	0.02	0.62	0.98	0.02	0.62	0.98
QT5 (2D)	0.99	0.99	0.99	0.96	0.98	1	0.96	0.98	1
QT8 (2D)	-0.47	0.73	0.99	0.27	0.44	0.99	0.27	0.44	0.99
QT9 (2D)	-55.3	-0.56	1	0.86	0.99	1	0.86	0.99	1
QT10 (2D)	-0.81	0.4	0.99	0.14	0.34	0.60	0.14	0.34	0.60
QT16 (2D)	-0.33	0.49	0.97	0.04	0.46	0.92	0.04	0.46	0.92
QT5 (3D)	0	0.01	0.02	0.78	0.95	1	0.78	0.95	1
QT8 (3D)	-22.6	0.41	1	0.15	0.36	0.99	0.15	0.36	0.99
QT10 (3D)	-5.5	0.43	1	0.19	0.69	1	0.19	0.69	1
AIQT2 (2D)	-3.95	0.16	0.97	0.15	0.64	0.94	0.15	0.64	0.94
AIQT5 (2D)	0.17	0.38	1	0.28	0.51	1	0.28	0.51	1
AIQT8 (2D)	-0.27	0.48	0.99	0.19	0.46	0.99	0.19	0.46	0.99
AIQT9 (2D)	-1999.8	-2.1	1	0.96	0.99	1	0.96	0.99	1
AIQT10 (2D)	-5.7	0.32	1	0.16	0.33	0.68	0.16	0.33	0.68
AIQT16 (2D)	-0.39	0.58	0.96	0	0.58	0.98	0	0.58	0.98
AIQT5 (3D)	0.3	0.6	1	0.05	0.9	1	0.05	0.9	1
AIQT8 (3D)	-9.25	0.32	0.99	0.03	0.75	0.99	0.03	0.75	0.99
AIQT10 (3D)	-24.6	0.44	1	0.08	0.5	1	0.08	0.5	1
DSQT12 (2D)	-1	0.16	1	0.25	0.32	0.54	0.25	0.32	0.54
DSQT18 (2D)	-12.2	0.38	1	0.59	0.99	1	0.59	0.99	1
DSQT19 (2D)	-11.8	0.08	1	0.18	0.47	1	0.18	0.47	1
DSQT12 (3D)	1	1	1	1	1	1	1	1	1
DSQT18 (3D)	0.99	0.99	0.99	0.72	0.86	1	0.72	0.86	1
DSQT19 (3D)	-0.53	0.74	1	0.77	0.87	1	0.77	0.87	1

Table 16: Characterization of Error-Resistance through Reduction

Query Template	CG (ms)	CG-FPC (min)	SEER (min)	LiteSEER (sec)
QT2 (2D)	15	21.6	2.2	5.8
QT5 (2D)	14	8.1	0.7	2.2
QT8 (2D)	14	27.0	1.9	7.2
QT9 (2D)	13	18.0	2.1	4.8
QT10 (2D)	14	10.8	0.7	2.9
QT16 (2D)	13	22.5	1.4	6.0
QT5 (3D)	25	17.0	5.1	8.2
QT8 (3D)	21	29.0	10.7	13.9
QT10 (3D)	22	23.0	7.8	11.0
AIQT2 (2D)	16	26.1	2.8	7.0
AIQT5 (2D)	16	21.6	0.7	5.8
AIQT8 (2D)	14	21.5	2.1	5.8
AIQT9 (2D)	15	21.6	1.7	5.7
AIQT10 (2D)	13	13.5	0.7	3.6
AIQT16 (2D)	13	18.9	0.4	5.0
AIQT5 (3D)	23	36.0	12.0	17.3
AIQT8 (3D)	20	38.0	14.0	18.2
AIQT10 (3D)	20	49.0	14.0	23.5
DSQT12 (2D)	19	5.4	0.07	1.4
DSQT18 (2D)	17	18.0	1.2	4.8
DSQT19 (2D)	14	24.3	1.7	6.5
DSQT12 (3D)	20	7.0	1.2	3.4
DSQT18 (3D)	30	35.0	7.2	16.8
DSQT19 (3D)	26	63.0	12.7	30.2

Table 17: Efficiency of Reduction Process

T percentile of the queries is selected as the preferred choice. The choice of T determines the level of risk that the user is willing to sustain with regard to worst-case behavior. Like the LEC approach, this too may be arbitrarily poor for a specific query as compared to the optimizer’s optimal choice.

Finally, in the (initial) optimization phase of the Rio approach [3, 4], a set of uncertainty modeling rules from [15] are used to classify selectivity errors into one of six categories (ranging from “no uncer-

tainty” to “very high uncertainty”) based on their derivation mechanisms. Then, these error categories are converted to hyper-rectangular error boxes drawn around the optimizer’s point estimate. Finally, if the plans chosen by the optimizer at the corners of the principal diagonal of the box are the same as that chosen at the point estimate, then this plan is *assumed* to be robust throughout the box. However, the conditions under which this assumption is likely to be valid are not outlined.

7 Conclusions

Errors in selectivity estimates are well-documented causes of poor plan choices by database optimizers. In this paper, we investigated whether the optimizer’s choices could be replaced by alternative plans, more resilient to these errors, from the parametric optimal set over the selectivity space. In particular, the recently proposed notion of anorexic reduction of plan diagrams was used to provide replacements that had large endo-optimal regions, making them error-resistant by definition in these areas. Further, the empirical evidence suggested that error-resistance was provided even in the exo-optimal regions in the vast majority of the cases. However, there were occasional situations where the replacement could turn out to be significantly worse. To prevent this, we developed a simple but accurate model of plan cost behavior. To our knowledge, this model is the first such characterization for industrial-strength query optimizers. Using this formulation, we devised efficient checks that operate only on the boundaries of the space to decide safety in the entire space. These checks are implemented utilizing foreign-plan costing, a recent feature addition in commercial database engines. A particularly attractive feature of our approach is that the safety guarantee applies on an *individual query basis*. As a bonus, the foreign-plan costing, in addition to providing safety, was leveraged to further improve the quality and complexity of the plan diagram reduction process.

The above techniques were integrated into the SEER algorithm and the intended benefits validated on a representative range of TPC-H and TPC-DS-based query templates on a leading commercial optimizer. We observed that typically at least one-third of the performance gap due to selectivity errors was bridged by the SEER reduction process, while in some instances, virtually *complete immunity* against selectivity errors was obtained. Our results remained consistent across different data distributions and physical design configurations.

Overall, SEER provides an effective and safe compile-time mechanism for substantially increasing resistance to selectivity errors on base relations, without requiring modifications to the optimizer or specialized knowledge of the workload/system. We also presented LiteSEER, an optimally efficient light-weight heuristic version of SEER that very cheaply provides a high degree of safety by restricting its attention to only the corners of the selectivity space. LiteSEER could be viably used in practice as a first-cut almost-safe reduction technique, especially if running time is a critical concern.

Currently, SEER operates as a post-processor after production of the plan diagram. In our future work, we intend to investigate how optimizers could internalize these ideas and be redesigned to directly produce safe reduced plan diagrams. Also, while we assumed a uniform distribution of selectivity estimation errors, it would be interesting to extend our results to incorporate user-defined probability distributions.

Acknowledgements. This work was supported in part by research grants from IBM, Microsoft and Google.

References

- [1] A. Aboulnaga and S. Chaudhuri, "Self-tuning Histograms: Building Histograms without Looking at Data", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, May 1999.
- [2] B. Babcock and S. Chaudhuri, "Towards a Robust Query Optimizer: A Principled and Practical Approach", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, June 2005.
- [3] S. Babu, P. Bizarro and D. DeWitt, "Proactive Re-Optimization", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, June 2005.
- [4] S. Babu, P. Bizarro and D. DeWitt, "Proactive Re-Optimization with Rio", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, June 2005.
- [5] N. Bruno, "A Critical Look at the TAB Benchmark for Physical Design Tools", *SIGMOD Record*, 36(4), 2007.
- [6] F. Chu, J. Halpern and P. Seshadri, "Least Expected Cost Query Optimization: An Exercise in Utility", *Proc. of ACM Symp. on Principles of Database Systems (PODS)*, May 1999.
- [7] A. Dey, S. Bhaumik, Harish D. and J. Haritsa, "Efficiently Approximating Query Optimizer Plan Diagrams", *Proc. of 34th Intl. Conf. on Very Large Data Bases (VLDB)*, August 2008.
- [8] F. Chu, J. Halpern and J. Gehrke, "Least Expected Cost Query Optimization: What Can We Expect", *Proc. of ACM Symp. on Principles of Database Systems (PODS)*, May 2002.
- [9] A. Deshpande, Z. Ives and V. Raman "Adaptive Query Processing", *Foundations and Trends in Databases*, 2007.
- [10] U. Feige, "A threshold of $\ln n$ for approximating set cover", *Journal of ACM*, 45(4), 1998.
- [11] Harish D., P. Darera and J. Haritsa, "On the Production of Anorexic Plan Diagrams", *Proc. of 33rd Intl. Conf. on Very Large Data Bases (VLDB)*, September 2007.
- [12] A. Hulgeri and S. Sudarshan, "Parametric Query Optimization for Linear and Piecewise Linear Cost Functions", *Proc. of 28th Intl. Conf. on Very Large Data Bases (VLDB)*, August 2002.
- [13] A. Hulgeri and S. Sudarshan, "AniPQO: Almost Non-intrusive Parametric Query Optimization for Nonlinear Cost Functions", *Proc. of 29th Intl. Conf. on Very Large Data Bases (VLDB)*, September 2003.
- [14] Y. Ioannidis and S. Christodoulakis, "On the Propagation of Errors in the Size of Join Results", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, May 1991.
- [15] N. Kabra and D. DeWitt, "Efficient Mid-Query Re-Optimization of Sub-Optimal Query Execution Plans", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, May 1998.
- [16] E. Kreyszig, *Advanced Engineering Mathematics*, New Age International, 5th ed, 1997.
- [17] L. Mackert and G. Lohman, "R* Optimizer Validation and Performance Evaluation for Local Queries", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, May 1986.
- [18] V. Markl, V. Raman, D. Simmen, G. Lohman, H. Pirahesh and M. Cilimdžić, "Robust Query Processing through Progressive Optimization", *Proc. of ACM SIGMOD Intl. Conf. on Management of Data*, June 2004.
- [19] J. Patel, M. Carey and M. Vernon, "Accurate Modeling of the Hybrid Hash Join Algorithm", *Proc. of ACM SIGMETRICS Intl. Conf. on Measurement and Modeling of Computer Systems*, May 1994.
- [20] Picasso Database Query Optimizer Visualizer, <http://dsl.serc.iisc.ernet.in/projects/PICASSO/picasso.html>
- [21] N. Reddy and J. Haritsa, "Analyzing Plan Diagrams of Database Query Optimizers", *Proc. of 31st Intl. Conf. on Very Large Data Bases (VLDB)*, August 2005.

- [22] P. Slavik, "A tight analysis of the greedy algorithm for set cover", *Proc. of 28th ACM Symp. on Theory of Computing*, 1996.
- [23] M. Stillger, G. Lohman, V. Markl and M. Kandil, "LEO – DB2's LEarning Optimizer", *Proc. of 27th Intl. Conf. on Very Large Data Bases (VLDB)*, September 2001.
- [24] MATLAB, <http://www.mathworks.com>
- [25] <http://www.tpc.org/tpch>
- [26] <http://www.tpc.org/tpcds>
- [27] <http://publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=/com.ibm.db2.udb.admin.doc/doc/t0024533.htm>
- [28] <http://msdn2.microsoft.com/en-us/library/ms189298.aspx>
- [29] http://infocenter.sybase.com/help/index.jsp?topic=/com.sybase.dc34982_1500/html/mig_gde/BABIFCAF.htm