Prairie: A Rule Specification Framework for Query Optimizers*

Technical Report TR 94–16

Dinesh Das Don Batory
Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712–1188
{ddas,batory}@cs.utexas.edu

Abstract

From our experience, current rule-based query optimizers do not provide a very intuitive and well-defined framework to define rules and actions. To remedy this situation, we propose an extensible and structured algebraic framework called Prairie for specifying rules. Prairie facilitates rule-writing by enabling a user to write rules and actions more quickly, correctly and in an easy-to-understand and easy-to-debug manner.

Query optimizers consist of three major parts: a search space, a cost model and a search strategy. The approach we take is only to develop the algebra which defines the search space and the cost model; we do not propose a search engine (i.e., search strategy) to drive the rules. We have chosen the Volcano optimizer generator as our search engine, because it is publicly available, and also because it has an efficient branch-and-bound search strategy. Using Prairie as a front-end, we translate Prairie rules to Volcano to validate our claim that Prairie makes it easier to write rules.

We describe our algebra and present experimental results which show that using a high-level framework like Prairie to design large-scale optimizers does not sacrifice efficiency.

1 Introduction

Query optimization [9, 12, 20] is a fundamental part of database systems. It is the process of generating an efficient access plan for a database query. Informally, an access plan is an execution strategy for a query; it is the sequence of low-level database retrieval operations that, when executed, produce the database records that satisfy the query. There are three basic aspects that define and influence query optimization: the search space, the cost model, and the search strategy.

The *search space* is the set of logically equivalent access plans that can be used to evaluate a query. All plans in a query's search space return the same result; however, some plans are more efficient than others. The *cost model* assigns a cost to each plan in the search space. The cost of a plan is an estimate of the resources used when the plan is executed; the lower the cost, the better the

*This research was supported in part by grants from The University of Texas Applied Research Laboratories, Schlumberger, and Digital Equipment Corporation.

plan. The *search strategy* is a specification of which plans in the search space are to be examined. If the search space is small, a viable strategy is to enumerate and evaluate every plan. However, most search spaces, even for simple queries, are enormous, and thus query optimizers often need heuristics to control the number of plans to be examined.

Traditionally, query optimizers have been built as monolithic subsystems of a DBMS. This simply reflects the fact that traditional database systems are themselves monolithic: the algorithms used for storing and retrieving data are hard-wired and are rather difficult to change. The need to have extensible database systems, and in turn extensible optimizers, has long been recognized in systems like Genesis [1], EXODUS [3], Starburst [15] and Postgres [18]. Rule-based query optimizers are among the major conceptual advances that have been proposed to deal with query optimizer extensibility [6–8, 10, 11, 13]. The extensibility translates into the ability to incorporate new operators, algorithms, cost models, or search strategies without changing the optimization algorithm.

In this paper, we describe an algebraic framework called *Prairie* for specifying rules in a query optimizer. Prairie is similar to other rule specification languages like Starburst [13] and Volcano [8], and indeed, we have based our work on Volcano to capture most of the advantages of rule-based optimizers. However, Prairie attempts to provide some key features that, we have found, simplify the effort in writing rules:

1. A framework in which users can define a query optimizer concisely in terms of a well-defined set of operators and algorithms. *All* operators and algorithms are considered first-class objects, i.e., *any* of them can occur in any rule, and *only* these operators and algorithms can appear in rules. This scheme eliminates the need for special classes of operators and algorithms, such as enforcers in Volcano and glue in Starburst, that significantly complicate rule specification.

2. A framework in which users can define a list of properties to characterize the expressions generated in the optimization process. Again, the goal here is to allow the user to treat *all* properties as having equal status. This is different from Volcano where the user must classify properties as logical, physical, or operator/algorithm arguments.

3. A framework in which users can specify mapping functions between properties concomitantly with the corresponding rules. This contrasts with existing approaches in which mappings between properties are fragmented into multiple functions and at logically different places than the corresponding rules. Research into rule-based optimizers has revealed that property-mapping functions are a major source of user effort, so this is an important goal.

4. The format (Prairie) in which users can cleanly specify rules is not necessarily the same format needed for generating efficient optimizers. Thus, there is a need for a pre-processor (written by us) that translates between these competing representations.

Prairie strives for uniformity in dealing with issues that have been a source of most user effort and potential user errors. In the following sections, we present the Prairie framework. We explain how our P2V pre-processor maps Prairie rule specifications into Volcano rule specifications that can be processed efficiently. Experimental results to support this claim are given in Section 4, where we compare implementations of the Texas Instruments Open OODB query optimizer using both Prairie and Volcano. We conclude with a summary and related research.

2 Prairie: A language for rule specification

The basic concepts and definitions that underly the Prairie model are presented in this section. The goal is to lay a foundation for reasoning about query optimizers algebraically; this is necessary for our subsequent discussion about translating Prairie specifications to those of Volcano.

2.1 Notation and assumptions

Stored Files and Streams. A file is *stored* if its tuples reside on disk. In the case of relational databases, stored files are sometimes called *base relations*; we will denote them by R or R_i. In object-oriented schemas, stored files are *classes*; we will denote them by C or C_i. Henceforth, whenever we refer to a stored file, we mean a relation or a class; when the distinction is unimportant, we will use F or F_i. A *stream* is a sequence of tuples and is the result of a computation on one or more streams or stored files; tuples of streams are returned one at a time, typically on demand. Streams can be *named*, denoted by S_i, or *unnamed*.

Database Operations. An *operation* is a computation on one or more streams or stored files. There are two types of database operations in Prairie: abstract (or implementation-unspecified) operators and concrete algorithms. Each is detailed below.

Operators. Abstract (or conceptual) *operators* specify computations on streams or stored files; they are denoted by all capital letters (e.g., JOIN). Operators have two types of parameters: essential and additional. *Essential parameters* are the stream or file inputs to an operator; these are the primary inputs to be processed by an operator. *Additional parameters* are "fine-grain" qualifications of an operator; their purpose is to describe an operator in more detail than essential parameters. As examples, some operators are described below; for each we explicitly indicate their essential parameters and parenthetically note their additional parameters.

- SORT(S_1) sorts stream S_1. The sorting attribute is an additional parameter of SORT.
- RET(F) retrieves tuples of stored file F. Additional parameters to RET include the selection predicate, the projected attributes list, and the output tuple order.
- JOIN(S_1, S_2) joins streams S_1 and S_2. (S_1 denotes the *outer stream* and S_2 denotes the *inner stream*). Additional parameters to JOIN include the join predicate and output stream tuple order.

Other operators are defined as they are needed.

Algorithms. *Algorithms* are concrete implementations of conceptual operators; they will be represented in lower case with the first letter capitalized. Algorithms have at least the same essential and additional parameters as the conceptual operators that they implement.[1] Furthermore, there can be, and usually are, several algorithms for a particular operator. For example, File_scan, Btree_scan and Index_scan are all valid algorithms that implement the operator RET, and Merge_join and Nested_loops are algorithms that implement the JOIN operator. Different algorithms offer different execution efficiencies.

[1] Algorithms may have *tuning parameters* which are not parameters in the operators they implement.

Operator	Description	Additional Parameters	Algorithm
JOIN(S_1, S_2)	Join streams S_1, S_2	tuple_order join_predicate	Nested_loops(S_1, S_2)
			Merge_join(S_1, S_2)
RET(F)	Retrieve file F	tuple_order selection_predicate projected_attributes	File_scan(F)
			Index_scan(F)
SORT(S_1)	Sort stream S_1	tuple_order	Merge_sort(S_1)
			Null(S_1)

Table 1: Operators and algorithms in a centralized query optimizer and their additional parameters

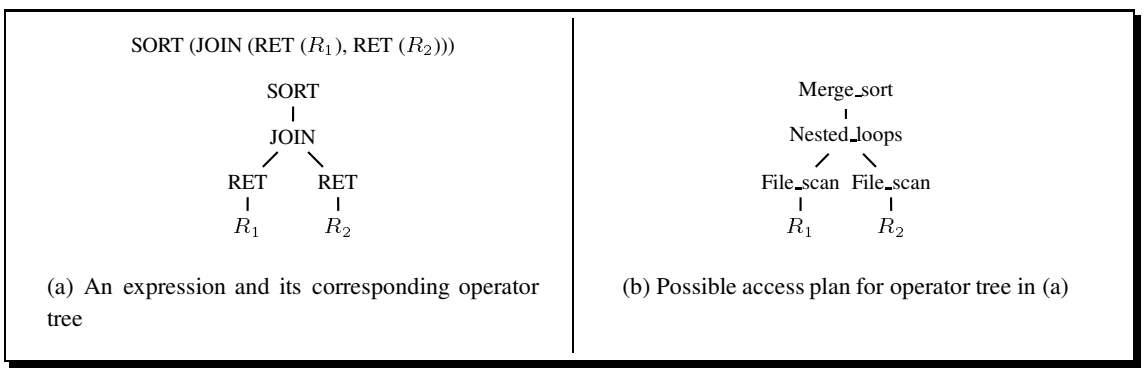

(a) An expression and its corresponding operator tree

(b) Possible access plan for operator tree in (a)

Figure 1: Example of an operator tree and access plan

Table 1 lists some operators and algorithms implementing them together with their additional parameters.

Operator Trees. An *operator tree* is a rooted tree whose non-leaf, or *interior*, nodes are database operations (operators or algorithms) and whose leaf nodes are stored files. The children of an interior node in an operator tree are the essential parameters (i.e., the stream or file parameters) of the node. Additional parameters are implicitly attached to each node. Algebraically, operator trees are compositions of database operations; thus, we will also call operator trees *expressions*; both terms will be used interchangeably.

EXAMPLE 1. A simple expression and its operator tree representation are shown in Figure 1(a). Relations R_1 and R_2 are first RETrieved, and then JOINed, and finally SORTed resulting in a stream sorted on a specific attribute. The figure shows only the essential parameters of the various operators, not the additional parameters. □

Descriptors. A *property* of a node is a (user-defined) variable that contains information used by an optimizer. An *annotation* is a ⟨*property, value*⟩ pair that is assigned to a node. A *descriptor* is a list of annotations that describes a node of an operator tree; every node has its own descriptor. As an example, Table 2 lists some typical properties that might be used in a descriptor. It must be noted

Property	Description
join_predicate	join predicate for JOIN operator
selection_predicate	selection predicate for RET operator
tuple_order	tuple order of resulting stream, DONT_CARE if none
num_records	number of tuples of resulting stream
tuple_size	size of individual tuple in stream
projected_attributes	list of projected attributes for RET operator
attributes	list of attributes
cost	estimated cost of algorithm

Table 2: Properties of nodes in an operator tree

here that streams and stored files may have different descriptor structures. The following notations will be useful in our subsequent discussions. If S_i is a stream, then $\mathbf{D_i}$ is its descriptor. Annotations of S_i are accessed by a structure member relationship, e.g., $\mathbf{D_i}$.num_records. Also, let E be an expression and let $\mathbf{D}$ be its descriptor. We will write this as $E : \mathbf{D}$.

EXAMPLE 2. The expression,

$$\text{SORT}(\text{JOIN}(\text{RET}(R_1) : \mathbf{D_3}, \text{RET}(R_2) : \mathbf{D_4}) : \mathbf{D_5}) : \mathbf{D_6}$$

corresponds to the operator tree in Figure 1(a), and represents the join of two relations R_1 and R_2. The two relations are first RETrieved, then JOINed and finally SORTed. $\mathbf{D_3}$ and $\mathbf{D_4}$ are the descriptors of the two RETs respectively, $\mathbf{D_5}$ is the descriptor of the JOIN, and $\mathbf{D_6}$ is the descriptor of the SORT. Assuming that the descriptor fields for this expression are those given in Table 2, the selection predicate for the first RET is $\mathbf{D_3}$.selection_predicate, and that for the second RET is given by $\mathbf{D_4}$.selection_predicate. The join predicate of the JOIN node is given by $\mathbf{D_5}$.join_predicate, and the attributes that are output are given by $\mathbf{D_5}$.attributes. The order in which SORT returns the output stream is given by $\mathbf{D_6}$.tuple_order. □

A notational simplification can be made here. Additional parameters of operators can be treated the same way as other properties of a node; essential parameters, however, are *expressions*. Thus, the term descriptor in the remainder of this paper will refer to a set of properties, including additional parameters, as shown in Table 2.

Currently, descriptor properties are defined entirely by the user; however, we envision providing a hierarchy of pre-defined descriptor types to aid this process.

Access Plans. An *access plan* is an operator tree in which all interior nodes are algorithms.

EXAMPLE 3. An access plan for the operator tree in Figure 1(a) is shown in Figure 1(b). Relations R_1 and R_2 are each retrieved using the File_scan algorithm, joined using Nested_loops, and finally sorted using Merge_sort. □

$E(x_1, \ldots, x_n) : \mathbf{D_1} \Longrightarrow E'(x_1, \ldots, x_n) : \mathbf{D_2}$ (1)

{{

pre-test statements

}}

test

{{

post-test statements

}}

Figure 2: General form of a T-rule

2.2 Prairie optimization paradigm

Query optimizers map operator trees to access plans.[2] Prairie admits two rather different means of optimization: top-down and bottom-up. A top-down query optimizer optimizes the parents of a node prior to optimizing the node itself. A bottom-up optimizer optimizes the children of a node prior to optimizing the node. The earliest optimizers (System R [17] and R* [16]) employed the bottom-up approach.

Our research concentrates on a top-down optimization of operator trees. We have chosen this approach because we intend to translate Prairie rules into the format required by the Volcano query optimizer generator [8]. Volcano is based on a top-down strategy. Given an appropriate search engine, Prairie can potentially also be used with a bottom-up optimization strategy; however, we will not discuss this approach in this paper.

In query optimization, there are certain annotations (such as additional parameters) that are known before any optimization is begun. These annotations can be computed at the time that the operator tree is initialized, and will not change with application of rules. Our following discussions assume operator trees are initialized.

2.3 Transformation rules

Transformation rules, or T-rules for short, define equivalences among pairs of expressions; they define mappings from one operator tree to another. Let E and E' be expressions that involve only abstract operators. Equation (1) (shown in Figure 2) shows the general form of a T-rule. The actions of a T-rule define the equivalences between the descriptors of nodes of the original operator tree E with the nodes of the output tree E'; these actions consist of a series of (C or C++) assignment[3] statements. The left-hand sides of these statements refer to descriptors of expressions on the right-hand side of the T-rule; the right-hand sides of the statements can refer to any descriptor in the T-rule. Function (called *helper* functions) calls can also appear on the right side of the assignment statements. Thus, descriptors on the *left-hand side* of a T-rule are *never* changed in the rule's actions.

[2] Actually, query optimizers operate on the output of a query compiler which translates a high-level query, e.g., one expressed in SQL, into an intermediate structure called a *query graph* and then into an operator tree. The query compilation process is well-known and we assume it is carried out before query optimization begins.

[3] The actions can be non-assignment statements (like function calls), but in this case, the P2V pre-processor (described in Section 3) needs some hints about the properties that are changed by the statement in order to correctly categorize each property. For simplicity, in this paper, we assume all actions consist of assignment statements.

A *test* is needed to determine if the transformations of the T-rule are in fact applicable.

Purely as an optimization, it is usually the case that not all statements in a T-rule's actions need to be executed prior to a T-rule's test. For this reason, the actions of a T-rule are split into two groups; those that need to be executed prior to the T-rule's test, and those that can be executed after a successful test. These groups of statements comprise, respectively, the *pre-test* and *post-test* statements of the T-rule.[4]

We now define the actions and tests of a T-rule more precisely. Let O_i be an abstract operator of E', and let $\mathbf{O_i}$ be its descriptor. Similarly, let I_i be an abstract operator of E and let $\mathbf{I_i}$ be its descriptor. (I_i is an operator that is input to the rule and O_i is an operator that is output by the rule). Let M_i denote the ith descriptor property. Thus, $\mathbf{O_i}.M_j$ is the value of the jth property of descriptor $\mathbf{O_i}$. The left-hand side of an assignment refers to an output descriptor ($\mathbf{O_i}$) or a member of an output descriptor ($\mathbf{O_i}.M_j$). The right-hand side is an expression or a helper function call that only references input descriptors and/or their members. Here are a few examples:

$$
\begin{array}{rcll}
\mathbf{O_i} & = & \mathbf{I_k}\ ; & // \text{ copy descriptor } \mathbf{I_k} \text{ to } \mathbf{O_i} \\
\mathbf{O_i}.M_j & = & \mathbf{I_k}.M_j + 4\ ; & // \text{ expression defining } \mathbf{O_i}.M_j \\
\mathbf{O_3}.M_5 & = & \text{helper } (\mathbf{I_1}.M_5, \mathbf{I_2}.M_5)\ ; & // \text{ helper function that computes } \mathbf{O_3}.M_5 \\
 & & & // \text{ from inputs } \mathbf{I_1}.M_5 \text{ and } \mathbf{I_2}.M_5.
\end{array}
$$

The test for a T-rule's applicability is a boolean expression and normally involves checks on the values of output descriptors (e.g., $\mathbf{O_3}.M_5 > 6$); occasionally, helper functions may be needed.

Again, it is important to remember that the pre-test actions are carried out prior to the test; the post-test actions are performed only if a T-rule's test evaluates to TRUE, and all post-test actions are performed immediately, with no intermediate optimization of any descendant nodes of the root of E.

Note that there are no actions that are carried out *after* the essential parameters of the root of E are optimized. This is because a T-rule only logically transforms a conceptual tree into another conceptual tree.

EXAMPLE 4. The associativity of JOINs is expressed by T-rule (2) in Figure 3(a). It rewrites a two-way join into an equivalent operator tree. The (single) pre-test statement computes the list of attributes of the new JOIN node on the right side. The test of the T-rule consists of a call to the helper function "is_associative", which returns TRUE or FALSE depending on whether the T-rule is applicable. If it is not, then the rule is rejected (e.g., because it generates a cross-product), otherwise the post-test statements are executed. The post-test statements compute various other annotations of the new nodes that are generated by applying the T-rule. Note the use of helper functions "cardinality" and "union" to compute descriptor properties.

Consider three relations R_1, R_2 and R_3, and let a_i, b_i and c_i be their respective sets of attributes. Figures 3(b) and 3(c) show, respectively, examples of the applicability and non-applicability of the join associativity T-rule. □

[4] We suspect it is possible to use data-flow analysis to partition the assignment statements automatically, but for now, we let the rule-writer do the partitioning.

$$\text{JOIN}(\text{JOIN}(S_1, S_2) : \mathbf{D_4}, S_3) : \mathbf{D_5} \Longrightarrow \text{JOIN}(S_1, \text{JOIN}(S_2, S_3) : \mathbf{D_6}) : \mathbf{D_7} \quad (2)$$

```
{{
        D6.attributes = union (D2.attributes, D3.attributes) ;
}}
is_associative (D6.join_predicate, D6.attributes, D5.join_predicate)
{{
        D7 = D5 ;
        D7.join_predicate = D4.join_predicate ;
        D6.tuple_size = D2.tuple_size + D3.tuple_size ;
        D6.num_records = cardinality (D2, D3) ;
}}
```

(a) Join Associativity T-rule

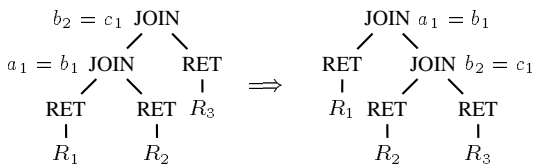

(b) Example of the associativity rule applied to an operator tree

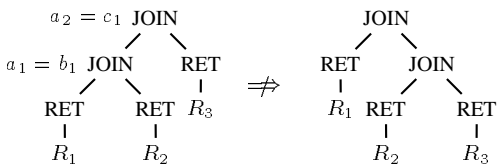

(c) Example of an operator tree where the associativity rule does not apply

Figure 3: Join associativity

$$E(x_1, \ldots, x_n) : \mathbf{D_1} \Longrightarrow A(x_1, \ldots, x_n) : \mathbf{D_2} \quad (3)$$

```
test
{{
        pre-opt statements
}}
{{
        post-opt statements
}}
```

Figure 4: General form of an I-rule

2.4 Implementation rules

Implementation rules, or I-rules for short, define equivalences between expressions and their implementing algorithms. Let E be an expression and A be an algorithm that implements E. The general form of an I-rule is given by Equation (3) (shown in Figure 4).

The actions associated with an I-rule are defined in three parts. The first part, or *test*, is a boolean expression whose value determines whether or not the rule can be applied.

The second part, or *pre-opt statements*, is a set of descriptor assignment statements that are executed only if the test is true and *before* any of the inputs x_i of E are optimized. Additional parameters of nodes are usually assigned in the pre-opt section. This is necessary before any of the

$$\text{SORT}(S_1) : \mathbf{D_2} \Longrightarrow \text{Merge_sort}(S_1) : \mathbf{D_3} \quad (4)$$
$$(\mathbf{D_2}.\text{tuple_order} \,!= \text{DONT_CARE})$$
$$\{\{$$
$$\quad \mathbf{D_3} = \mathbf{D_2} \;;$$
$$\}\}$$
$$\{\{$$
$$\quad \mathbf{D_3}.\text{cost} = \mathbf{D_1}.\text{cost} + (\mathbf{D_3}.\text{num_records}) * \log(\mathbf{D_3}.\text{num_records}) \;;$$
$$\}\}$$

Figure 5: Merge-sort sort algorithm

nodes on the right side can be optimized.

The third part, or *post-opt statements*, is a set of descriptor assignment statements that are executed *after* all x_i are optimized. Normally, the post-opt statements compute cost properties that can only be determined once the inputs to the algorithm are completely optimized and their costs known. This *does not*, however, imply a bottom-up optimization strategy. It simply means that although I-rules are applied to parents before their children are optimized, the *cost* (and other properties in the post-opt section) of the parent cannot be computed until the children have been optimized.

EXAMPLE 5. Equation (4) (in Figure 5) shows the I-rule that implements the SORT operator by Merge_sort. I-rule (4) rewrites a stream such that it is sorted using the Merge_sort algorithm. The test for this I-rule is that the tuple order of the sorted stream must not be a DONT_CARE order. The pre-opt section consists of the default statement that copies the descriptor from the left side to the expression on the right. The post-opt section consists of a cost-assigning statement to the Merge_sort node. □

EXAMPLE 6. Equation (5) (shown in Figure 6) is the I-rule that selects the Nested_loops algorithm to implement the JOIN operator. The test for this rule is TRUE since Nested_loops can be applied regardless of any property values. The pre-opt section consists of three assignment statements. The first statement sets the descriptor of Nested_loops to that of the JOIN. The next two statements express the fact that the tuple order of Nested_loops is the same as the tuple order of its left (outer) input; all other properties remain the same. The third statement in the pre-opt section ensures that this requirement is met by setting the tuple_order of S_1 on the right side.[5]

[5]Actually, it is not enough to simply set the desired tuple order of S_1; it is also necessary to ensure that *after* optimization, S_1 does indeed have the required property. One way to satisfy this is to insert a SORT node in front of S_1 that can meet the sortedness requirement of S_1. Thus, in this case, we would need a T-rule (which introduces a new operator JOPR),

$$\text{JOIN}(S_1, S_2) : \mathbf{D_3} \Longrightarrow \text{JOPR}(\text{SORT}(S_1) : \mathbf{D_4}, \text{SORT}(S_2) : \mathbf{D_5}) : \mathbf{D_6},$$

and an I-rule,

$$\text{JOPR}(S_1, S_2) : \mathbf{D_3} \Longrightarrow \text{Nested_loops}(S_1 : \mathbf{D_4}, S_2) : \mathbf{D_5}.$$

In our discussions, this additional level of detail will be ignored for the sake of simplicity.

$$\text{JOIN}(S_1, S_2) : \mathbf{D_3} \Longrightarrow \text{Nested_loops}(S_1 : \mathbf{D_4}, S_2) : \mathbf{D_5} \quad (5)$$

TRUE

{{

$\mathbf{D_5} = \mathbf{D_3}$;

$\mathbf{D_4} = \mathbf{D_1}$;

$\mathbf{D_4}$.tuple_order = $\mathbf{D_3}$.tuple_order ;

}}

{{

$\mathbf{D_5}$.cost = $\mathbf{D_4}$.cost + ($\mathbf{D_4}$.num_records) $*$ $\mathbf{D_2}$.cost ;

}}

Figure 6: Nested loops join algorithm

The post-opt section is executed after S_1 and S_2 are optimized; it consists of a single statement that assigns the cost of the Nested_loops node. The cost is indicative of the fact that in this algorithm, each tuple of the stream S_1 involves scanning the entire stream S_2; S_1 is scanned only once. □

2.5 Null algorithm

Recall that, in Section 1, we mentioned that Prairie allows users to treat all operators and algorithms as first-class objects, i.e., all operators and algorithms are explicit, in contrast to enforcers in Volcano or glue in Starburst. This requires that Prairie provide a mechanism where users can also "delete" one or more of the explicit operators from expressions. This is done by having a special class of I-rules that have the form given by Equation (6) in Figure 7(a). The left side of the rule is a single abstract operator O with one stream input S_1. The right side of the rule is an algorithm called "Null" with the same stream input but with a different descriptor. As the name suggests, the Null algorithm is supposed to pass its input unchanged to algorithms above it in an operator tree. This is accomplished in the I-rule as follows.

The test for this I-rule is always TRUE, i.e., any node in an operator tree with O as its operator can be implemented by the Null algorithm. The actions associated with this rule have a specific pattern. The pre-opt section consists of three statements. The first statement copies the descriptor of the operator O to the algorithm Null. The second statement sets the descriptor of the stream S_1 on the right side to the descriptor of the stream S_1 on the left side. Why is it necessary to do this? The key lies in the third statement. This statement copies the property "property" of the operator O node on the left side to the "property" of the input stream S_1 on the right side. Since left-hand side descriptors cannot be changed in an I-rule, a new descriptor $\mathbf{D_3}$ is necessary for S_1 to convey the property propagation information.

The post-opt section in the I-rule has only a cost-assignment statement; this simply sets the cost of the Null node to the cost of its optimized input stream.

The Null algorithm, therefore, serves to effectively transform a single operator to a no-op.

EXAMPLE 7. Equation (7) (in Figure 7(b)) shows the I-rule that rewrites the SORT operator to use a Null algorithm. The third pre-opt statement sets the tuple order of S_1 on the right side to be

$O(S_1) : \mathbf{D_2} \Longrightarrow \text{Null}(S_1 : \mathbf{D_3}) : \mathbf{D_4}$ (6)

TRUE

{{

$\mathbf{D_4} = \mathbf{D_2}$;

$\mathbf{D_3} = \mathbf{D_1}$;

$\mathbf{D_3}$.property = $\mathbf{D_2}$.property ;

}}

{{

$\mathbf{D_4}$.cost = $\mathbf{D_3}$.cost ;

}}

(a) General form of a "Null" I-rule

SORT$(S_1) : \mathbf{D_2} \Longrightarrow \text{Null}(S_1 : \mathbf{D_3}) : \mathbf{D_4}$ (7)

TRUE

{{

$\mathbf{D_4} = \mathbf{D_2}$;

$\mathbf{D_3} = \mathbf{D_1}$;

$\mathbf{D_3}$.tuple_order = $\mathbf{D_2}$.tuple_order ;

}}

{{

$\mathbf{D_4}$.cost = $\mathbf{D_3}$.cost ;

}}

(b) Null sort algorithm

Figure 7: The "Null" algorithm concept

the tuple order of the SORT node, thus ensuring that when S_1 is optimized on the right side, it will have the same tuple order as the SORT node. □

3 The P2V pre-processor

In Section 1, we enumerated the four primary goals of Prairie, viz., uniformity in operator and algorithms; uniformity in properties; uniformity in property-transformations; and efficient generation of Prairie optimizers. The first three goals are driven by the need for conceptual simplicity; however, they alone do not necessarily generate efficient optimizers. The P2V pre-processor ensures that efficient optimizers can be realized from Prairie specifications, by translating them to the Volcano framework and then generating an optimizer by compiling with the Volcano search engine. This Prairie optimizer-generator paradigm is shown schematically in Figure 8. The pre-processor itself is 4500 lines of `flex` and `bison` code. In this section, we describe the pre-processor steps and explain why the Prairie-to-Volcano transformation is non-trivial.

3.1 Correspondence of elements in Prairie and Volcano

Table 3 shows the relationship between the basic elements of a Prairie and Volcano specification; these correspondences are preserved by the P2V pre-processor. Operators and algorithms in Prairie correspond directly to operators and algorithms, respectively, in Volcano.

The concept of enforcer-operators and enforcer-algorithms needs some explanation. In a Prairie specification, one can have a set of I-rules of the form:

$$O(S_1) : \mathbf{D_2} \Longrightarrow A_1(S_1) : \mathbf{D_3}$$

$$\ldots$$

$$O(S_1) : \mathbf{D_2} \Longrightarrow A_n(S_1) : \mathbf{D_3}$$

$$O(S_1) : \mathbf{D_2} \Longrightarrow \text{Null}(S_1 : \mathbf{D_3}) : \mathbf{D_4}$$

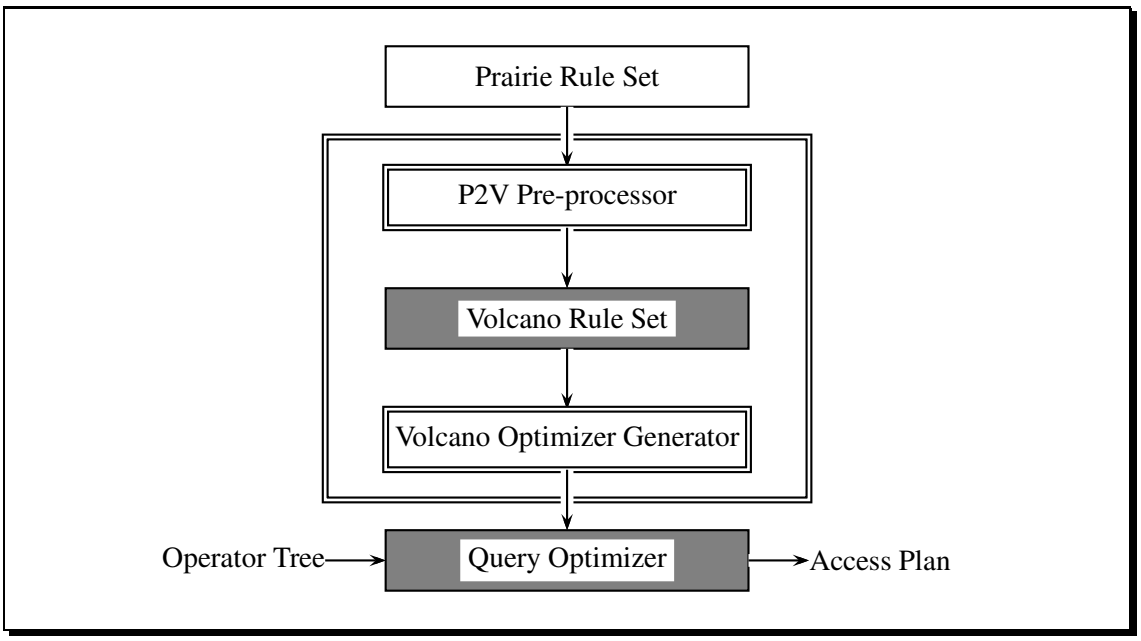

Figure 8: The Prairie optimizer-generator paradigm. Double-boxed modules represent software generators, shaded boxes represent generated programs. The outermost double-boxed portion denotes the Prairie optimizer generator.

Prairie	**Volcano**
Operator	Operator
Algorithm	Algorithm
Enforcer-operator	—
Enforcer-algorithm	Enforcer
"Null" Algorithm	—
Operator Tree	Logical Expression
Access Plan	Physical Expression
Descriptor	Operator/Algorithm Argument
	Physical property
	Cost
—	Logical Property
—	System Property

Table 3: Correspondence of elements in Prairie and Volcano

i.e., an operator O has algorithms A_1 through A_n, and Null, as implementations. The pre-processor classifies O as an *enforcer-operator*, and algorithms A_1 through A_n as *enforcer-algorithms*. An example of an enforcer-operator is the SORT operator, and an enforcer-algorithm is the Merge_sort algorithm (shown in Table 1). Enforcer-algorithms in the Prairie model are translated into enforcers in the Volcano model; enforcer-operators disappear in Volcano when the P2V pre-processor combines several I-rules to generate a Volcano rule (this is described in Section 3.3).

Another major responsibility of the P2V pre-processor is classifying different properties in the descriptor. Volcano forces users to classify properties (as logical, physical, or operator/algorithm arguments) according to their use. The classification of properties helps optimize the performance of Volcano. Unfortunately, property classification is actually rule-dependent. That is, adding another rule to a Volcano rule set may cause a property that was previously considered "logical" to become "physical", or vice versa. Migrating properties between classifications entails a considerable amount of reprogramming on the user's part, and consequently makes Volcano rule sets rather brittle. Prairie, in contrast, does not ask users to make these distinctions; the P2V pre-processor determines the classification of properties automatically and thus significantly facilitates the extensibility of rule sets. This is done by transforming a single descriptor structure into three parts (see Table 3): an operator/algorithm argument, physical properties and cost. Briefly, physical properties are properties that are requested by the user, cost represents an estimate of an algorithm's cost, and operator/algorithm arguments are all remaining properties (including additional properties). The pre-processor achieves the classification of properties by examining the actions of all rules: a property with a type "COST" is classified as a cost property, properties changed in pre-opt sections of I-rules (e.g., tuple_order in I-rule (5)) are physical properties, and all other properties are operator/algorithm arguments.

Operator trees and access plans in Prairie correspond to logical and physical expressions, respectively, in Volcano; the notations are very similar.

3.2 Correspondence of rules

3.2.1 T-rules

T-rules in Prairie are translated to trans_rules (transformation rules) in Volcano; expressions occurring on either side of a T-rule are transformed into Volcano logical expressions. Table 4(a) shows the correspondence between Prairie T-rules and Volcano trans_rules. The join associativity trans_rule (corresponding to the T-rule in Figure 3) in Volcano is as follows[6]:

(JOIN ?op_arg5 ((JOIN ?op_arg4 (?1 ?2)) ?3)) -> (JOIN ?op_arg7 (?1 (JOIN ?op_arg6 (?2 ?3))))

The pre-test and test portions of a T-rule are mapped to the cond_code (condition code) of a Volcano trans_rule, and the post-test portion is mapped to the appl_code (application code). However, expressions in Prairie T-rules can contain enforcer-operators. According to Table 3, enforcer-operators are deleted by the P2V pre-processor, so T-rules containing enforcer-operators are modified before being translated into Volcano. Sometimes, this can result in rules that can be further combined into one as an optimization measure; this is discussed in Section 3.3.

3.2.2 I-rules

I-rules in Prairie correspond to impl_rules (implementation rules) in Volcano; the complete relationship is shown in Table 4(b). However, the transformation is complicated by several factors. First, Prairie adopts a per-rule approach of specifying property transformations, i.e., properties of algorithms are specified in an I-rule, whereas Volcano adopts a per-algorithm approach where properties

[6] There are conditions and actions associated with Volcano rules that are not shown here.

Prairie	Volcano
T-rule	trans_rule
Operator	Operator
Enforcer-operator	—
Descriptor	Operator Argument
Pre-test code	Cond_code
Test	Cond_code
Post-test code	Appl_code

(a) Translation of T-rules

Prairie	Volcano
I-rule	impl_rule
Operator	Operator
Algorithm	Algorithm
Operator Descriptor	Operator Argument
Algorithm Descriptor	Algorithm Argument
Test	Cond_code
Pre-opt code	"do_any_good"
Post-opt code	"derive_phy_prop"
—	"cost"
—	"get_input_pv"

(b) Translation of I-rules. Quoted strings denote helper functions.

Table 4: Correspondence of rules in Prairie and Volcano

of an algorithm are specified in a helper function for that algorithm. We believe that the per-rule approach is more general and intuitive. The generality arises from the fact that the property transformations can be different in different I-rules for the same algorithm; thus, the per-rule approach is a superset of the per-algorithm approach. The per-rule approach is intuitive because property mappings are made together with the rule invocation, instead of a logically disjoint helper function.

The second complicating factor in Volcano involves the reliance on several helper functions to achieve property transformations. For instance, using Prairie terminology, a Volcano implementation rule requires an additional four functions (called "do_any_good", "cost", "get_input_pv", and "derive_phy_prop"). Although the purpose of these functions is purportedly to enhance Volcano performance, they add considerable complexity to the design of rule sets. Prairie, in contrast, eliminates the need for users to specify these functions, by adopting the per-rule approach of property transformations; as shown in Table 4(b), the P2V pre-processor generates two of these functions ("do_any_good" and "derive_phy_prop") automatically from I-rule specifications. The other two helper functions are short-circuited, suggesting that the Volcano model and implementation is actually more complicated than it needs to be.

3.3 Rule merging

In the process of translating Prairie specifications to Volcano, several opportunities arise for obtaining a compact set of rules. These opportunities arise either because the user specifies a non-compact set of rules, or because of the translation process itself. An example of the latter case arises when enforcer-operators are deleted by the P2V pre-processor. Consider, for example, the following set of rules in Prairie:

$$\begin{aligned}
\text{JOIN}(S_1, S_2) : \mathbf{D_3} &\implies \text{JOPR}(\text{SORT}(S_1) : \mathbf{D_4}, \text{SORT}(S_2) : \mathbf{D_5}) : \mathbf{D_6} \\
\text{SORT}(S_1) : \mathbf{D_2} &\implies \text{Null}(S_1 : \mathbf{D_3}) : \mathbf{D_4} \\
\text{JOPR}(S_1, S_2) : \mathbf{D_3} &\implies \text{Nested_loops}(S_1 : \mathbf{D_4}, S_2) : \mathbf{D_5}
\end{aligned}$$

The first rule is a T-rule, and the next two are I-rules. Note that SORT is an enforcer-operator (because it has a Null implementation), so it is deleted in the transformation process. The first rule then becomes a mapping from one operator JOIN to another JOPR which can be viewed as an idempotence mapping; the resulting rule set can be optimized by deleting this idempotent rule and replacing all occurrences of JOPR with JOIN. Thus, the above set of Prairie rules can be combined into a single I-rule,

$$\text{JOIN}(S_1, S_2) : \mathbf{D_3} \Longrightarrow \text{Nested_loops}(S_1 : \mathbf{D_4}, S_2) : \mathbf{D_5}$$

which can then be translated into a Volcano impl_rule. In general, the number of T-rules in a Prairie rule set is equal to the number of trans_rules in a Volcano rule set plus an additional T-rule for each operator.[7] Also, the number of I-rules is the same as the number of impl_rules plus an additional I-rule for each enforcer-operator (for Null implementations, as described in Section 2.5) and an additional I-rule for each enforcer (since enforcers appear as I-rules in Prairie, instead of being implicit, as in Volcano). The larger number of rules represents a conceptual simplification in rule writing. However, the P2V pre-processor ensures that the transformation from Prairie to Volcano always produces a compact set of rules.

4 Experimental results

This section presents experimental results which demonstrate the value of Prairie in specifying rule sets of rule-based optimizers. Our experiments consist of specifying rule-based optimizers using Prairie and generating optimizers using the P2V pre-processor and the optimizer-generator paradigm of Figure 8.

In [5], we presented an implementation of a centralized relational query optimizer using Prairie. Using the P2V translator, we translated this to Volcano format and optimized several queries using the resultant optimizer. For comparison, we hand-coded the same optimizer directly in Volcano. The results presented there showed that, using Prairie (compared to directly using Volcano) resulted in approximately 50% savings in lines of code with negligible (less than 5%) increase in query optimization time. However, the optimizer was quite small in terms of the number of operators, algorithms and rules.

For a more realistic evaluation of Prairie, we needed answers to the following questions:

1. Is Prairie adequate for large-scale rule sets?
2. How is programmer productivity enhanced by the high-level abstractions of Prairie?
3. Can Prairie rule sets be translated automatically into efficient implementations?

We addressed the first question by using the Texas Instruments Open OODB query optimizer rule set, which has the largest publicly available rule set. We describe this optimizer in the next section, and then give our assessments to the last two questions in subsequent sections.

[7] This is to introduce enforcer-operators in expressions, as mentioned in footnote 5.

4.1 The Texas Instruments Open OODB query optimizer

The Texas Instruments Open Object-Oriented Database Management System [19] is an open, extensible, object-oriented database system which provides users an architectural framework that is configurable in an incremental manner. It consists of three sets of modules: a core set providing low-level primitives for creating new environments, a set of functional modules that facilitates extensibility using functional requirements, and a meta-architecture module housing the extensibility concepts of Open OODB. Examples of the core set are communication and address space management, while examples of functional modules are persistence, distribution and query processing. The meta-architecture module consists of events, sentries, and policy manager interfaces.

The query processing module provides users with a query language (OQL[C++]) based on SQL and C++. A query expressed in this high-level format is parsed and transformed into an operator tree suitable for optimization. The query optimizer generates an optimal access plan from this operator tree which is then transformed into a C++ program ready for execution.

The query optimizer in the Open OODB [2] is generated using Volcano. It is written as a set of trans_rules and impl_rules that define the algebra of an object-oriented database system. Currently, there are 17 transformation rules and 9 implementation rules together with about 13000 lines of code for support functions; this, of course, can be changed by an Open OODB user for specific needs. There are also *catalogs* which contain information about base classes that are used by the optimizer.

4.2 Programmer productivity

Programmer productivity can be measured in different ways. An admittedly simplistic metric is the number of lines of code that must be written. But there are also less tangible measures, such as the amount of conceptual effort needed to understand a particular programming task. Our experience with the Open OODB query optimizer suggests that Prairie excels on the latter, while offering modest reductions in the volume of code that needs to be written.

We converted by hand the Open OODB query optimizer's Volcano specifications to Prairie. This was a non-trivial task because of the relatively large size of the rule set and the complexity of the support functions. This was where we found Prairie helped in conceptually simplifying the rules and actions. We then used our P2V pre-processor to reconstitute these Prairie specifications as Volcano specifications. As described in Section 3, this process involved a considerable level of complexity, partly because the Prairie specification had 22 T-rules and 11 I-rules compared to 17 trans_rules and 9 impl_rules in the Volcano specification; the reconstituted Volcano specification had the same number of trans_rules and impl_rules as the original hand-coded specification.

Converting the Open OODB optimizer rule set into Prairie format actually simplified its specification as the complexities of the Volcano model were removed. The reduction in lines of code was modest — there was about a 10% savings.[8] However, as mentioned above, savings in lines of code do not adequately reflect increases in programmer productivity. We found the encapsulated specifications of Prairie — namely, the use of a single descriptor and fewer explicit support functions — made rule programming *much* easier.

[8] The original Volcano specification had 13400 lines, the Prairie specification had 12100 lines, and the P2V-generated Volcano specification had 15800 lines.

4.3 Performance results using the Open OODB optimizer

The acid test of Prairie was whether Prairie specifications could be translated into efficient optimizer implementations. Our experiments using the Open OODB consisted of optimizing 8 different queries using the two query optimizers generated, respectively, using Prairie and using Volcano directly (in the remainder of this section, we will use "Prairie" and "Volcano" to denote these two approaches). There were 4 distinct expressions that were used to generate the queries used in the experiments; these are shown in Figure 9. Each expression represents an N-way join query for varying N.

The first expression E1 is a simple retrieval and join of base classes. The second, E2, is also a join of base classes; however, after each class retrieval, an attribute has to be materialized (i.e., brought into view) before the join. The third and fourth expressions (E3 and E4) are the same as the first and second (E1 and E2) respectively, except that there is a selection of attributes (the select operator is the root of the expressions).[9]

The algebra that was used in the Prairie and Volcano optimizers for our experiments consisted of 5 relational operators SELECT, PROJECT, JOIN, RET and UNNEST (for set-valued attributes) and an object-oriented operator called MAT (for MATerialize; it is fundamentally a pointer-chasing operator for attributes of a class). There were 8 algorithms.

There are many parameters that can be varied when benchmarking a query optimizer. Since our objective was to verify that the Prairie approach did not sacrifice efficiency, our criteria for the queries was that they test a majority of the rules, with varying properties of the base classes. To this end, we tested our optimizer (and the Volcano optimizer) with 8 different queries (shown in Table 5). The eight queries Q1 through Q8 are derived from the 4 expressions in Figure 9. Each expression E1 through E4 is used to obtain two queries for a fixed number N of JOINs in the expression. The only difference between the two queries obtained from an expression is that the first one does not contain any indices on any classes, whereas the second one contains a single index on each base class occurring in the expression. In expressions where a SELECT is present (E3 and E4), the selection predicate is a conjunction of equality predicates $bc_i == const_i$, where bc_i is an attribute of class C_i, and $const_i$ is a constant (we arbitrarily set this to i, because its value doesn't affect the correctness or performance of the optimizer). In addition, for queries with a SELECT and whose base classes have indices (Q6 and Q8 in Table 5), the (single) index of each base class was chosen to be the attribute referenced in the selection predicate. For example, class C_i was chosen to have an index on attribute bc_i. The join predicates for each JOIN were chosen at random, and were always equality predicates. The choice of JOIN predicates was such that the queries corresponded to linear query graphs. In the future, we will experiment with non-linear (e.g., star) query graphs.

Table 5 also shows the number of trans_rules and impl_rules that are matched by each expression. These are the rules whose left hand sides match a sub-expression. However, not all the rules were necessarily applicable. For instance, an impl_rule with an index scan would not apply to Q3, although it might apply to Q4.

Queries Q1 through Q8 were optimized for increasing number N of JOINs. For a fixed number

[9] The most complex expression E4 consists of all operators in the algebra, except PROJECT and UNNEST. PROJECT was not considered because it appeared in only one impl_rule and no trans_rules, and thus, would not affect the size of the search space of abstract expressions. UNNEST was not considered because it appeared in exactly one trans_rule and one impl_rule; including it in our queries would have increased the number of parameters that could affect our run-times. We preferred to concentrate on simple JOIN expressions.

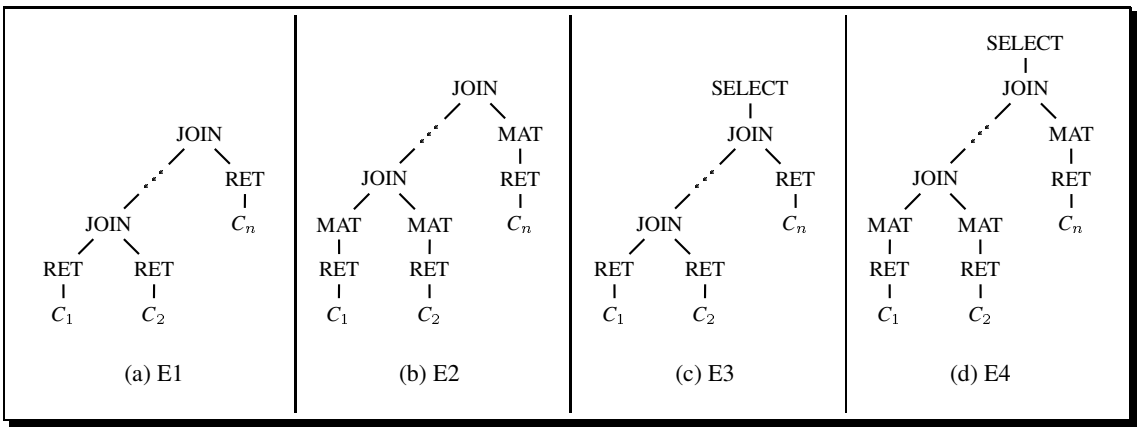

Figure 9: Expressions used in generating queries for experiments

Query	Indices?	Expression	Rules matched	
			trans_rules	impl_rules
Q1	No	E1	3	3
Q2	Yes			
Q3	No	E2	8	4
Q4	Yes			
Q5	No	E3	9	5
Q6	Yes			
Q7	No	E4	16	7
Q8	Yes			

Table 5: Queries used in experiments

of JOINs in a query, we varied the cardinalities of the base classes 5 times, each time generating a query with different class properties, and averaged the run-times over the 5 query instances to generate the per-query optimization time. Thus, each point in our graphs represents the average of 5 queries. The run-times were measured[10] using the GNU `time` command. All experiments were performed on a lightly loaded DECstation 5000/200 running Ultrix 4.2.

The optimization times for each query for both approaches (Prairie and Volcano) are shown in Figures 10 through 13. The number of joins in each set of graphs was varied to a maximum of 8, or until virtual memory was exhausted.

The first set of graphs, in Figure 10 shows the performance of a simple relational-type query. The optimization times are almost identical between Prairie and Volcano, and the notable point is that the presence of an index does not change the optimizer's behavior, i.e., the two graphs are identical. This arises because the optimizer algebra had only two join algorithms (pointer join and hash join), neither of which makes use of any indices.

The second set of graphs (Figure 11) shows the results of optimizing Q3 and Q4. Here, as in

[10] Since the run-times were too small to be measured accurately with `time`, each query instance was optimized 3000 times (in a loop) and the total time was divided by 3000 to get the per-query optimization time.

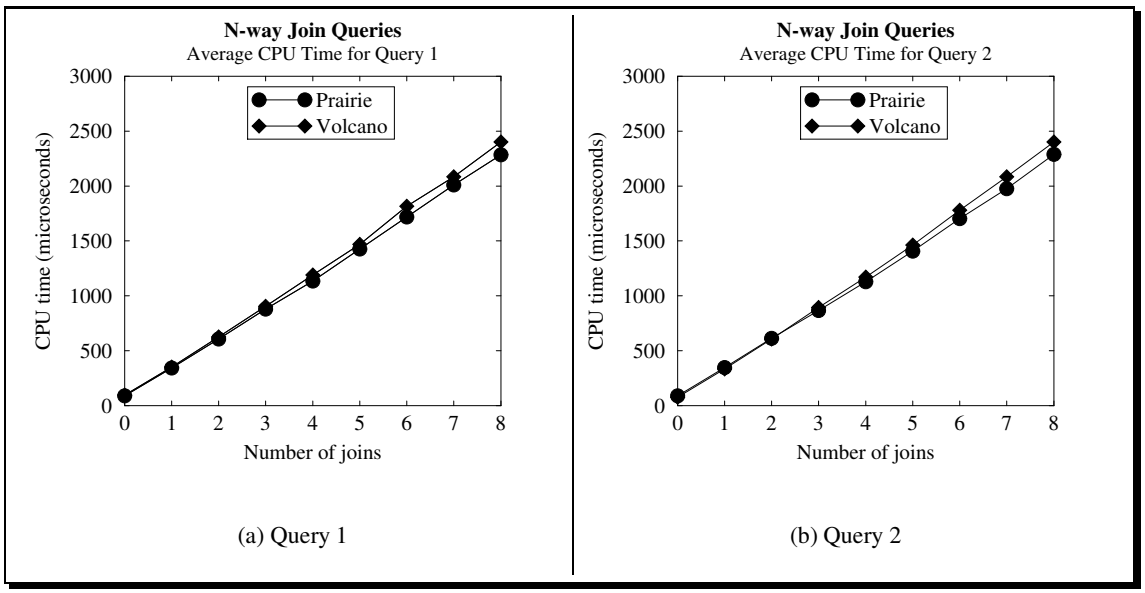

Figure 10: Query optimization times for Q1 and Q2

Figure 10, the presence (or absence) of indices makes no difference. Both the Prairie and Volcano approaches have comparable run-times. The sharp jump in the graphs from 7-way to 8-way joins is due to the fact that since all optimization is done in main memory, dynamic memory allocation (caused by `malloc` calls) results in a lot of thrashing at this point. We speculate that in systems with more virtual memory, the graphs will be smoother.

The third and fourth sets of graphs in Figures 12 and 13 are optimizations of queries with a selection predicate. In these cases, the presence of an index makes a difference if the index is referenced in the selection predicate (as we designed). Also, in these two figures, the performance of both Prairie and Volcano was almost identical, except that Prairie does slightly worse due to the larger number of `malloc` calls that the P2V translator introduces. Also, note that we could only go up to 3-way joins before virtual memory was exhausted. As the available memory decreases, there is increased thrashing (as shown by the sharp changes in slope in the plots) resulting in a much slower optimization process.

In all four sets of plots, we can see that Prairie performs with almost (less than 5% variation) the same efficiency as Volcano. In extreme cases, when memory is scarce, Prairie runs more slowly (about 15%) (e.g., Figure 12(b)), but we believe that this situation already represents a serious bottleneck for both Volcano and Prairie.

Figure 14 shows the variation of the number of equivalence classes (they are the same in Prairie and Volcano) as a function of the number of joins in the query. The growth rate of the number of equivalence classes increases with the increase in the complexity of the expressions. In particular, for E3 and E4, the introduction of the SELECT operator results in a dramatic increase in the search space, since this operator has many interactions with the other operators of the algebra; this is also reflected in Table 5 where E3 matches three times as many trans_rules as E1 and E4 matches twice as many trans_rules as E2. The lesson to be learnt here is that extending an existing query optimizer by

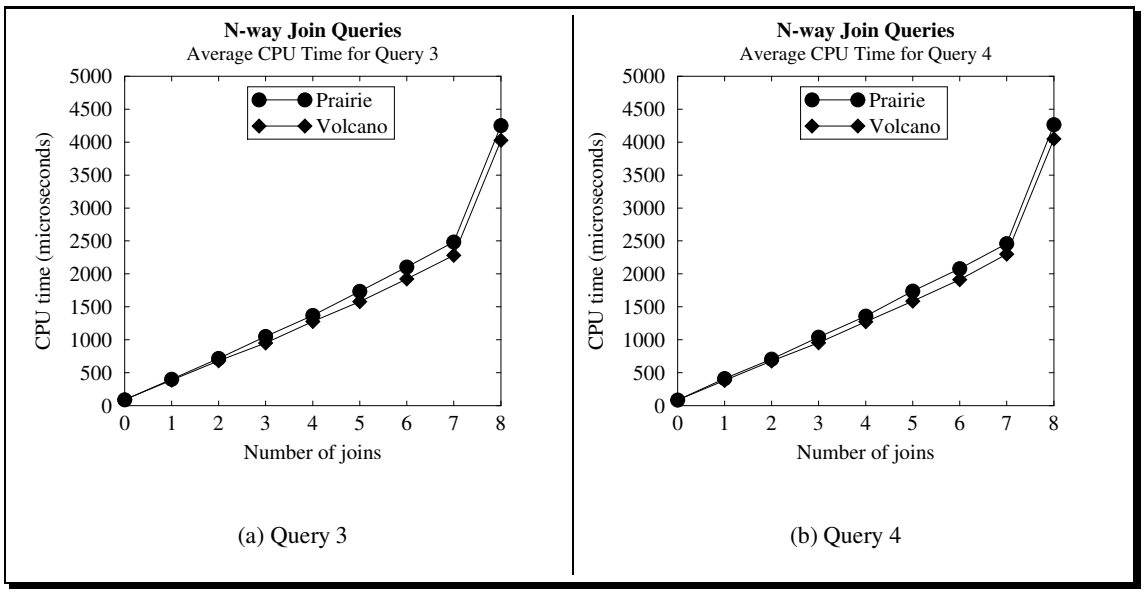

Figure 11: Query optimization times for Q3 and Q4

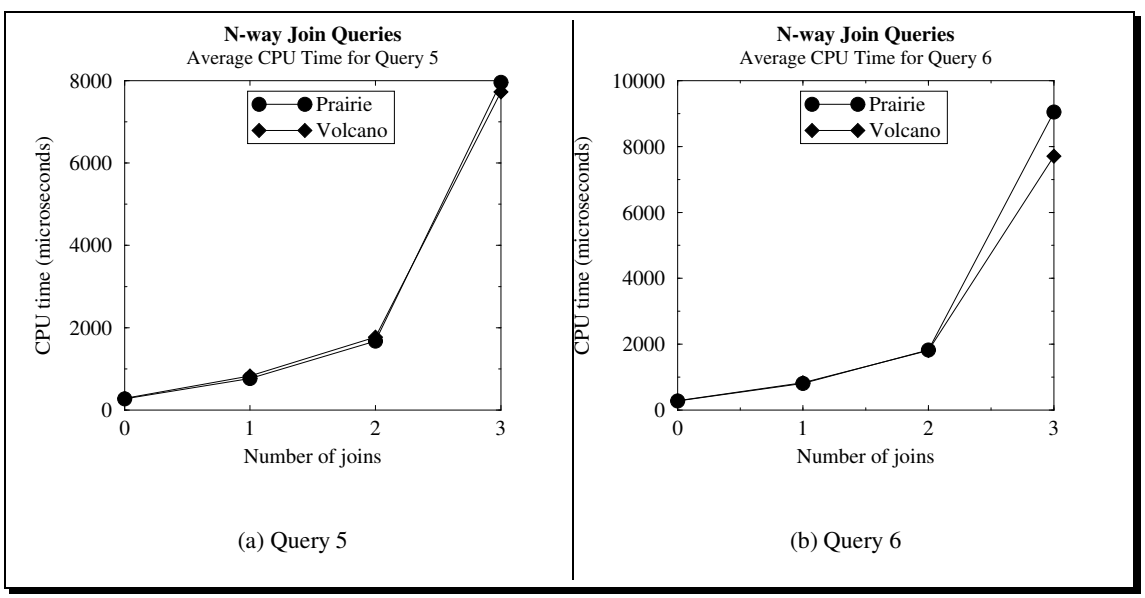

Figure 12: Query optimization times for Q5 and Q6

adding operators, algorithms or rules can result in an enormous increase in optimization complexity, especially if the additions impact a significant fraction of the equivalence classes. Extensibility, thus, must be judiciously coupled with user heuristics to avoid unpleasant surprises.

The results presented in this section show that Prairie optimizers need not sacrifice efficiency for clarity, even for large rule sets. More research and validation is necessary to verify that Prairie is an

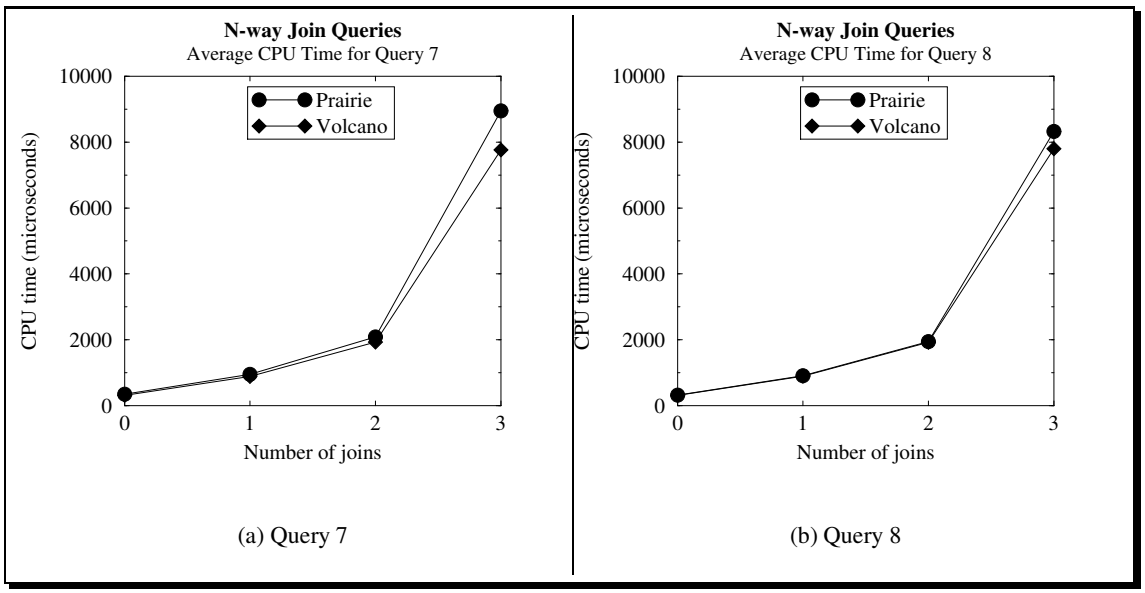

Figure 13: Query optimization times for Q7 and Q8

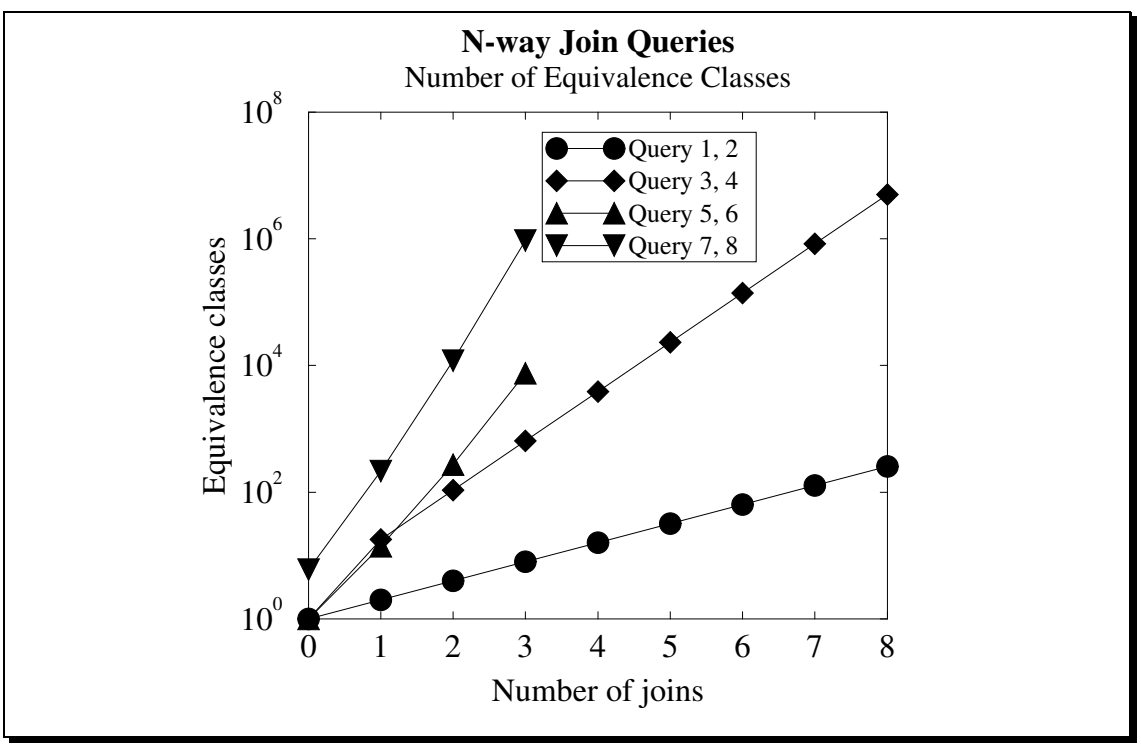

Figure 14: Number of equivalence classes vs. number of joins

efficient tool for optimizer specification.

5 Related research

The System R optimizer [17] was the most important development in query optimization research. It was a cost-based centralized relational query optimizer and introduced a variety of key concepts like "interesting" expressions, cardinality estimation using selectivity factors and dynamic programming with pruning of search space. These concepts continue to be important in query optimizer research.

The query optimizer in R^* [4, 14, 16] works in essentially the same way as that of System R, except that R^* is a distributed database system which introduces some subtle complications in its query optimizer.

The Starburst query optimizer [11, 13, 15] uses rules for all decisions that need to be taken by the query optimizer. The rules are functional in nature and transform a given operator tree into another. The rules are commonly those that reflect relational calculus facts. In Starburst, the query rewriting phase is different from the optimization phase. The rewriting phase transforms the query itself into equivalent operator trees based on relational calculus rules. The plan optimization phase selects algorithms for each operator in the operator tree that is obtained after rewriting. The disadvantage of separating the query rewrite and the optimization phases is that pruning of the search space is not possible during query rewrite, since the rewrite phase is non-cost-based.

Freytag [6] describes a rule-based query optimizer similar to Starburst. The rules are based on LISP-like representations of access plans. The rules themselves are recursively defined on smaller expressions (operator trees). Although several expressions can contain a common sub-expression, Freytag doesn't consider the possibility of sharing. Expressions are evaluated each time they are encountered. This is obviously inefficient. In addition, as in Starburst, he doesn't consider the cost transformations inherent in any query optimizer; rules are syntactic transformation rules.

EXODUS [7, 10] provides an optimizer generator which accepts a rule-based specification of the data model as input. The optimizer generator compiles these rules, together with pre-defined rules, to generate an optimizer for the particular data model and set of operators. Unlike Freytag, the optimizer generator for EXODUS allows for C code along with definitions of new rules. This allows the database implementor the freedom to associate any action with a particular rule. Operator trees in EXODUS are constructed bottom-up from previously constructed trees.

The Volcano optimizer generator project [8] evolved from the EXODUS project. It is different from all the above optimizers in one significant way: it is a top-down optimizer compared with the bottom-up strategy of the others. Operator trees are optimized starting from the root while sub-trees are not yet optimized. This leads to a constraint-driven generation of the search space. While this method results in a tight control of the search space, it is unconventional and requires careful attention on the part of the optimizer implementor to ensure that legal operator trees are not accidently left out of the search space. We have used Volcano as our back-end search engine.

6 Conclusion and future work

Current rule-based query optimizers do not provide a very intuitive and conceptually streamlined framework to define rules and actions. Our experiences with the Volcano optimizer generator suggest that its model of rules and the expression of these rules is much more complicated and too low-level than it needs to be. As a consequence, rule sets in Volcano are fragile, hard to write, and debug.

Similar problems may exist in other contemporary rule-based query optimizers.

We believe that rule-based query optimizers will be standard tools of future database systems. The pragmatic difficulties of using existing rule-based optimizers led us to develop Prairie, an extensible and structured algebraic framework for specifying rules. Prairie is similar to existing optimizers in that it supports both transformation rules and implementation rules. However, Prairie makes several improvements:

1. it offers a conceptually more streamlined model for rule specification;
2. rules are encapsulated, there are no "hidden" operators or "hidden" algorithms;
3. implementation hints (e.g., enforcers) are deduced automatically;
4. and it has efficient implementations.

We have explained how the first three points are important for simplifying rule specifications and making rule sets less brittle for extensibility. A consequence is that Prairie rules are simpler and more robust than rules of existing optimizers (e.g., Volcano). We addressed the fourth point by building a P2V pre-processor which uses sophisticated algorithms to compose and compact a Prairie rule set into a Volcano rule set. To demonstrate the scalability of our approach, we rewrote the TI Open OODB rule set as a Prairie rule set, generated its Volcano counterpart, and showed that the performance of the synthesized Volcano rule set closely matches that of the hand-designed Volcano rule set.

Our future work will concentrate on developing higher-level abstractions using Prairie, including automatically generating Prairie rule sets, and combining multiple Prairie rule sets to automatically generate efficient optimizers.

Acknowledgments

We wish to thank Texas Instruments, Inc. for making the Open OODB source code available to us. Comments by José Blakeley, Anne Ngu, Vivek Singhal and Thomas Woo greatly improved the quality of the paper.

References

[1] D. S. Batory. Building blocks of database management systems. Technical Report TR–87–23, The University of Texas at Austin, February 1988.

[2] José A. Blakeley, William J. McKenna, and Goetz Graefe. Experiences building the Open OODB query optimizer. In *Proceedings 1993 ACM SIGMOD International Conference on Management of Data*, pages 287–296, Washington, May 1993.

[3] Michael J. Carey, David J. DeWitt, Daniel Frank, Goetz Graefe, M. Muralikrishna, Joel E. Richardson, and Eugene J. Shekita. The architecture of the EXODUS extensible DBMS. In *Proceedings International Workshop on Object-Oriented Database Systems*, pages 52–65, Asilomar, September 1986.

[4] Dean Daniels, Patricia Selinger, Laura Haas, Bruce Lindsay, C. Mohan, Adrian Walker, and Paul Wilms. An introduction to distributed query compilation in R*. In *Proceedings 2nd International Conference on Distributed Databases*, pages 291–309, Berlin, September 1982.

[5] Dinesh Das and Don Batory. Prairie: An algebraic framework for rule specification in query optimizers. In *Proceedings of the Workshop on Database Query Optimizer Generators and Rule-Based Optimizers*, pages 139–154, Dallas, September 1993.

[6] Johann Christoph Freytag. A rule-based view of query optimization. In *Proceedings 1987 ACM SIGMOD International Conference on Management of Data*, pages 173–180, San Francisco, May 1987.

[7] Goetz Graefe. *Rule-Based Query Optimization in Extensible Database Systems*. PhD thesis, University of Wisconsin–Madison, 1987.

[8] Goetz Graefe. Volcano, an extensible and parallel query evaluation system. Technical Report CU–CS–481–90, University of Colorado at Boulder, July 1990.

[9] Goetz Graefe. Query evaluation techniques for large databases. *ACM Computing Surveys*, 25(2):73–170, June 1993.

[10] Goetz Graefe and David J. DeWitt. The EXODUS optimizer generator. In *Proceedings 1987 ACM SIGMOD International Conference on Management of Data*, pages 387–394, San Francisco, May 1987.

[11] L. M. Haas, J. C. Freytag, G. M. Lohman, and H. Pirahesh. Extensible query processing in Starburst. Research Report RJ 6610, IBM Almaden Research Center, December 1988.

[12] Won Kim, David S. Reiner, and Don S. Batory, editors. *Query Processing in Database Systems*. Springer-Verlag, 1985.

[13] Guy M. Lohman. Grammar-like functional rules for representing query optimization alternatives. In *Proceedings 1988 ACM SIGMOD International Conference on Management of Data*, pages 18–27, Chicago, June 1988.

[14] Guy M. Lohman, C. Mohan, Laura M. Haas, Bruce G. Lindsay, Patricia G. Selinger, Paul F. Wilms, and Dean Daniels. Query processing in R*. Research Report RJ 4272, IBM Research Laboratory, San Jose, April 1984.

[15] P. Schwarz, W. Chang, J. C. Freytag, G. Lohman, J. McPherson, C. Mohan, and H. Pirahesh. Extensibility in the Starburst database system. In *Proceedings International Workshop on Object-Oriented Database Systems*, pages 85–92, Asilomar, September 1986.

[16] P. G. Selinger and M. Adiba. Access path selection in distributed data base management systems. In Deen and Hammersly, editors, *Proceedings International Conference on Databases*, pages 204–215, University of Aberdeen, July 1980.

[17] P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price. Access path selection in a relational database management system. In *Proceedings 1979 ACM SIGMOD International Conference on Management of Data*, pages 23–34, Boston, May 1979.

[18] Michael Stonebraker and Lawrence A. Rowe. The design of Postgres. In *Proceedings 1986 ACM SIGMOD International Conference on Management of Data*, pages 340–355, Washington, May 1986.

[19] David L. Wells, José A. Blakeley, and Craig W. Thompson. Architecture of an open object-oriented database management system. *IEEE Computer*, 25(10):74–82, October 1992.

[20] C. T. Yu and C. C. Chang. Distributed query processing. *ACM Computing Surveys*, 16(4):399–433, December 1984.