Generating optimal DAG-structured query evaluation plans

Thomas Neumann · Guido Moerkotte

Received: 10 October 2008 / Accepted: 2 March 2009 / Published online: 31 March 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract In many database queries relations are accessed multiple times during query processing. In these cases query processing can be accelerated by sharing scan operators and possibly other operators based upon the common relations. The standard approach to achieve sharing works as follows. In a first phase, a non-shared tree-shaped plan is generated via a traditional plan generator. In a second phase, common instances of a scan are detected and shared. After that, other possible operators are shared. The result is an operator DAG (directed acyclic graph).

The limitation of this approach is obvious. As sharing influences plan costs, a separation of the optimization into two phases comprises the danger of missing the optimal plan, since the first optimization phase does not know about sharing.

We remedy this situation by (1) introducing a general framework for reasoning about sharing and (2) sketching how this framework can be integrated into a plan generator, which then constructs optimal DAG-structured query evaluation plans.

1 Introduction

Standard query evaluation relies on tree-structured algebraic expressions which are generated by the plan generator and

then evaluated by the query execution engine [17]. Conceptually, the algebra consists of operators working on sets or bags. On the implementation side, they take one or more tuple (object) streams as input and produce a single output stream. The tree-structure thereby guarantees that every operator – except for the root – has exactly one consumer of its output. This flexible concept allows a nearly arbitrary combination of operators and highly efficient implementations.

However, this model has several limitations. Consider, e.g., the following SQL query:

```
select ckey
from   customer, order
where  ckey=ocustomer
group by ckey
having sum(price) = (select max(total)
                    from (select ckey, sum(price) as total
                          from   customer, order
                          where  ckey=ocustomer
                          group by ckey))
```

This query leads to a plan like the one at the left of Fig. 1. We observe that (1) both relations are accessed twice, (2) the join and (3) the grouping are calculated twice. To (partially) remedy this situation, Cao et al. proposed to share scans of the same relation [2]. The plan resulting from their approach is shown in the middle of Fig. 1. Still, not all sharing possibilities are exploited. Obviously, only the plan at the right exploits sharing to its full potential.

Another disadvantage of the approach by Cao et al. is that optimization is separated into two phases. In a first phase, a traditional plan generator is used to generate tree-structured plans like the one on the left of Fig. 1. In a second step, this plan is transformed into the one at the middle of Fig. 1. This approach is very nice in the sense that it does not necessitate any modification to existing plan generators: just

---

T. Neumann (✉)
Max-Planck Institute for Informatics,
Campus E1 4,
66111 Saarbrücken, Germany
e-mail: neumann@mpi-inf.mpg.de

G. Moerkotte
University of Mannheim,
Mannheim, Germany

Fig. 1 Example plans

an additional phase needs to be implemented. However, as always when more than a single optimization phase is used, there is the danger of coming up with a suboptimal plan. In our case, this is due to the fact that adding sharing substantially alters the costs of a plan. As the plan generator is not aware of this cost change, it can come up with (from its perspective) best plan, which exhibits (after sharing) higher costs than the optimal plan.

In this paper, we remedy both disadvantages of the approach by Cao et al. First, we present a general framework that allows us to reason about share equivalences. This will allow us to exploit as much sharing as possible, if this leads to the best plan. Second, we sketch a plan generator that needs a single optimization phase to generate plans with sharing. Using a single optimization phase avoids the generation of suboptimal plans. The downside is that the plan generator has to be adopted to include our framework for reasoning about share equivalence. However, we are strongly convinced that this effort is worth it.

The rest of the paper is organized as follows. Section 2 discusses related work. Section 3 precisely defines the problem. Section 4 describes the theoretical foundation for reasoning about share equivalence. Section 5 sketches the plan generator. The detailed pseudocode and its discussion is given in [22]. Section 6 explains the cost computation. Section 7 contains the evaluation. Section 8 concludes the paper. This paper is an extended and revised version of an earlier work presented at the BTW 2009 [23].

2 Related work

Let us start the discussion on related work with a general categorization. Papers discussing the generation of DAG-structured query execution plans fall into two broad categories. In the first category, a single optimal tree-structured plan is generated, which is then turned into a DAG by exploiting sharing. This approach is in danger of missing the optimal plan since the tree-structured plan is generated with costs which neglect sharing opportunities. We call this *post plan generation share detection* (PPGSD). This approach is the most prevailing one in multi-query optimiza-

tion, e.g. [28]. In the second category, common subexpressions are detected in a first phase before the actual plan generation takes place. The shared subplans are generated independently and then replaced by an artificial single operator. This modified plan is then given to the plan generator. If several sharing alternatives exist, several calls to the plan generator will be made. Although this is a very expensive endeavor due to the (in the worst case exponentially many) calls to the plan generator. Since the partial plans below and above the materialization (temp) operator are generated separately, there is a slight chance that the optimal plan is missed. We term this *loose coupling* between the share detection component and the plan generator. In stark contrast, we present a tightly integrated approach that allows to detect sharing opportunities incrementally during plan generation.

A Starburst paper mentions that DAG-structured query graphs would be nice, but too complex [13]. A later paper about the DB2 query optimizer [7] explains that DAG-structured query plans are created when considering views, but this solution materializes results in a temporary relation. Besides, DB2 optimizes the parts above and below the temp operator independently, which can lead to suboptimal plans. Similar techniques are mentioned in [3, 5].

The Volcano query optimizer [8] can generate DAGs by partitioning data and executing an operator in parallel on the different data sets, merging the result afterwards. Similar techniques are described in [9], where algorithms like select, sort, and join are executed in parallel. However, these are very limited forms of DAGs, as they always use data partitioning (i.e., in fact, one tuple is always read by one operator) and sharing is only done within one logical operator.

Another approach using loose coupling is described in [25]. A later publication by the same author [26] applies loose coupling to multi-query optimization. Another interesting approach is [4]. It also considers cost-based DAG construction for multi-query optimization. However, its focus is quite different. It concentrates on scheduling problems and uses greedy heuristics instead of constructing the optimal plan. Another loose coupling approach is described in [31]. The authors propose running the optimizer repeatedly and then using view matching mechanisms to construct DAGs by based upon solutions from the previous runs. Fi-

nally, there exist a number of papers that consider special cases of DAGs, e.g. [1, 4]. While they propose using DAGs, they either produce heuristical solutions or do not support DAGs in the generality of the approach presented here.

3 Problem definition

Before going into detail, we provide a brief formal overview of the optimization problem we are going to solve in this paper. This section is intended as an illustration to understand the problem and the algorithm. Therefore, we ignore some details like the problem of operator selection here (i.e. the set of operators does not change during query optimization).

We first consider the classical tree optimization problem and then extend it to DAG optimization. Then, we distinguish this from similar DAG-related problems in the literature. Finally, we discuss further DAG-related problems that are not covered in this paper.

3.1 Optimizing trees

It is the query optimizer's task to find the cheapest query execution plan that is equivalent to the given query. Usually this is done by algebraic optimization, which means the query optimizer tries to find the cheapest algebraic expression (e.g. in relational algebra) that is equivalent to the original query. For simplicity we ignore the distinction between physical and logical algebra in this section. Further, we assume that the query is already given as an algebraic expression. As a consequence, we can safely assume that the query optimizer transforms one algebraic expression into another.

Nearly all optimizers use a tree-structured algebra, i.e. the algebraic expression can be written as a tree of operators. The operators themselves form the nodes of the tree, the edges represent the dataflow between the operators. In order to make the distinction between trees and DAGs apparent, we give their definitions. A *tree* is a directed, cycle-free graph $G = (V, E)$, $|E| = |V| - 1$ with a distinguished root node $v_0 \in V$ such that all $v \in V \setminus \{v_0\}$ are reachable from v_0 .

Now, given a query as a tree $G = (V, E)$ and a cost function c , the query optimizer tries to find a new tree $G' = (V, E')$ such that $G \equiv G'$ (concerning the produced output) and $c(G')$ is minimal (to distinguish the tree case from the DAG case we will call this equivalence $\equiv_T$). This can be done in different ways, either transformatively by transforming G into G' using known equivalences [10–12], or constructively by building G' incrementally [16, 27]. The optimal solution is usually found by using dynamic programming or memoization. If the search space is too large then heuristics are used to find good solutions.

An interesting special case is the join ordering problem where V consists only of joins and relations. Here, the following statement holds: any tree G' that satisfies the syntax constraints (binary tree, relations are leafs) is equivalent to G . This makes constructive optimization quite simple. However, this statement does no longer hold for DAGs (see Sect. 4).

3.2 Optimizing DAGs

DAGs are directed acyclic graphs, similar to trees with overlapping (shared) subtrees. Again, the operators form the nodes, and the edges represent the dataflow. In contrast to trees, multiple operators can depend on the same input operator. We are only interested in DAGs that can be used as execution plans, which leads to the following definition. A DAG is a directed, cycle-free graph $G = (V, E)$ with a denoted root node $v_0 \in V$ such that all $v \in V \setminus \{v_0\}$ are reachable from v_0 . Note that this is the definition of trees without the condition $|E| = |V| - 1$. Hence, all trees are DAGs.

As stated above, nearly all optimizers use a tree algebra, with expressions that are equivalent to an operator tree. DAGs are no longer equivalent to such expressions. Therefore, the semantics of a DAG has to be defined. To make full use of DAGs, a DAG algebra would be required (and some techniques require such a semantics, e.g. [30]). However, the normal tree algebra can be lifted to DAGs quite easily: a DAG can be transformed into an equivalent tree by copying all vertices with multiple parents once for each parent. Of course this transformation is not really executed: it only defines the semantics. This trick allows us to lift tree operators to DAG operators, but it does not allow the lifting of tree-based equivalences (see Sect. 4).

We define the problem of optimizing DAGs as follows. Given the query as a DAG $G = (V, E)$ and a cost function c , the query optimizer has to find any DAG $G' = (V' \subseteq V, E')$ such that $G \equiv G'$ and $c(G')$ is minimal. Thereby, we defined two DAG-structured expressions to be equivalent ($\equiv_D$) if and only if they produce the same output. Note that there are two differences between tree optimization and DAG optimization: First, the result is a DAG (obviously), and second, the result DAG possibly contains fewer operators than the input DAG.

Both differences are important and both are a significant step from trees! The significance of the latter is obvious as it means that the optimizer can choose to eliminate operators by reusing other operators. This requires a kind of reasoning that current query optimizers are not prepared for. Note that this decision is made during optimization time and not beforehand, as several possibilities for operator reuse might exist. Thus, a cost-based decision is required. But also the DAG construction itself is more than just reusing operators: a real DAG algebra (e.g. [30]) is vastly more expressive and

cannot e.g. be simulated by deciding operator reuse beforehand and optimizing trees.

The algorithm described in this work solves the DAG construction problem in its full generality.

By this we mean that (1) it takes an arbitrary query DAG as input (2) constructs the optimal equivalent DAG, and (3) thereby applies equivalences, i.e. a rule-based description of the algebra. This discriminates it from the problems described below, which consider different kinds of DAG generation.

3.3 Problems not treated in depth

In this work, we concentrate on the algebraic optimization of DAG-structured query graphs. However, using DAGs instead of trees produces some new problems in addition to the optimization itself.

One problem area is the execution of DAG-structured query plans. While a tree-structured plan can be executed directly using the iterator model, this is no longer possible for DAGs. One possibility is to materialize the intermediate results used by multiple operators, but this induces additional costs that reduce the benefit of DAGs. Ideally, the reuse of intermediate results should not cause any additional costs, and, in fact, this can be achieved in most cases. As the execution problem is common for all techniques that create DAGs as well as for multi-query optimization, many techniques have been proposed. A nice overview of different techniques can be found in [14]. In addition to this generic approach, there are many special cases like e.g. application in parallel systems [8] and sharing of scans only [2]. The more general usage of DAGs is considered in [20, 25], which describe runtime systems for DAGs.

4 Algebraic optimization

In this section, we present a theoretical framework for DAG optimization. We first highlight three different aspects that differentiate DAG optimization from tree optimization. Then, we use these observations to formalize the reasoning over DAGs.

4.1 Using tree equivalences

Algebraic equivalences are fundamental to any plan generator: It uses them either directly by transforming algebraic expressions into equivalent ones, or indirectly by constructing expressions that are equivalent to the query. For tree-structured query graphs, many equivalences have been proposed (see e.g. [6, 18]). But when reusing them for DAGs, one has to be careful.

Fig. 2 Invalid transformation for DAGs

When only considering the join ordering problem, the joins are freely reorderable. This means that a join can be placed anywhere where its syntax constraints are satisfied (i.e. the join predicate can be evaluated). However, this is not true when partial results are shared. Let us demonstrate this by the example presented in Fig. 2. The query computes the same logical expression twice. In (a) the join $A \bowtie B$ is evaluated twice and can be shared as shown in (b). But the join with C may not be executed before the split, as shown in (c), which may happen when using a constructive approach to plan generation (e.g. dynamic programming or memoization) that aggressively tries to share relations and only considers syntax constraints. That is, a join can be build into a partial plan as soon as its join predicate is evaluable which in turn only requires that the referenced tables are present. This is the only check performed by a dynamic programming approach to join ordering. Intuitively, it is obvious that (c) is not a valid alternative, as it means that $\bowtie C$ is executed on both branches. But in other situations, a similar transformation is valid, e.g. selections can often be applied multiple times without changing the result. As the plan generator must not rely on intuition, we now describe a formal method to reason about DAG transformations. Note that the problem mentioned above does not occur in current query optimization systems, as they treat multiple occurrences of the same relation in a query as distinct relations. But for DAG generation, the query optimizer wants to treat them as identical relations and thus potentially avoid redundant scans.

The reason why the transformation in Fig. 2 is invalid becomes clear if we look at the variable bindings. Let us denote by $A : a$ the successive binding of variable a to members of a set A . In the relational context, a would be bound to all tuples found in relation A . As shown in Fig. 3a, the original expression consists of two different joins $A \bowtie B$ with different bindings. The join can be shared in (b) by properly applying the renaming operator (ρ) to the output. While a similar rename can be used after the join $\bowtie C$ in (c), this still means that the topmost join joins C twice, which is different from the original expression.

This brings us to a rather surprising method to use normal algebra semantics:

Fig. 3 More verbose representation of Fig. 2

A binary operator must not construct a (*logical*) DAG.

Here, logical means that the same algebra expression is executed on both sides of its input. Further:

What we do allow are *physical* DAGs, which means that we allow sharing operators to compute multiple logical expressions simultaneously.

As a consequence, we only share operators after proper renames: if an operator has more than one consumer, all but one of these must be rename operators. Thus, we use ρ to pretend that the execution plan is a tree (which it is, logically) instead of the actual DAG.

4.2 Share equivalence

Before going into more detail, we define whether two algebra expressions are *share equivalent*. This notion will express that *one expression can be computed by using the other expression and renaming the result*. Thus, given two algebra expressions A and B , we define

$$A \equiv_S B \text{ iff } \exists_{\delta_{AB}: \mathcal{A}(A) \rightarrow \mathcal{A}(B), \delta_{AB} \text{ bijective}} \rho_{\delta_{AB}}(A) \equiv_D B.$$

where we denote by $\mathcal{A}(A)$ all the attributes provided in the result of A .

As this condition is difficult to test in general, we use a constructively defined sufficient condition of share equivalence instead. First, two scans of the same relation are share equivalent, since they produce exactly the same output (with

different variable bindings):

$$\text{scan}_1(R) \equiv_S \text{scan}_2(R)$$

Note that in a constructive bottom-up approach, the mapping function $\delta_{A,B}$ is unique. Therefore, we always know how attributes are mapped.

Other operators are share equivalent if their input is share equivalent and their predicates are equivalent after applying the mapping function. The conditions for share equivalence for common operators are summarized in Fig. 4. They are much easier to check, especially when constructing plans bottom-up (as this follows the definition).

Note that share equivalence as calculated by the tests above is orthogonal to normal expression equivalence. For example, $\sigma_1(\sigma_2(R))$ and $\sigma_2(\sigma_1(R))$ are equivalent but *not* derivable as share equivalent by testing the sufficient conditions. This will not pose any problems to the plan generator, as it will consider both orderings. On the other hand, $\text{scan}_1(R)$ and $\text{scan}_2(R)$ are share equivalent, but not equivalent, as they may produce different attribute bindings.

Share equivalence is only used to detect if exactly the same operations occur twice in a plan and, therefore, cause costs only once. Logical equivalence of expressions is handled by the plan generator anyway, it is not DAG-specific.

Using this notion, the problem in Fig. 2 becomes clear: In part b), the expression $A \bowtie B$ is shared, which is ok, as $(A \bowtie B) \equiv_S (A \bowtie B)$. But in part c), the top-most join tries to also share the join with C , which is not ok, as $(A \bowtie B) \not\equiv_S ((A \bowtie B) \bowtie C)$. Note that while this might look obvious, it is not when e.g. constructing plans bottom up and assuming freely reorderable joins, as discussed in Sect. 3.1.

4.3 Optimizing DAGs

The easiest way to reuse existing equivalences is to hide the DAG structure completely: During query optimization, the query graph is represented as a tree, and only when determining the costs of a tree the share equivalent parts are determined and the costs adjusted accordingly. Only after the query optimization phase the query is converted

Fig. 4 Definition of share equivalence for common operators

$$\begin{array}{ll}
 A \cup B \equiv_S C \cup D & \text{if } A \equiv_S C \wedge B \equiv_S D \\
 A \cap B \equiv_S C \cap D & \text{if } A \equiv_S C \wedge B \equiv_S D \\
 A \cdot B \equiv_S C \cdot D & \text{if } A \equiv_S C \wedge B \equiv_S D \\
 \Pi_A(B) \equiv_S \Pi_C(D) & \text{if } B \equiv_S D \wedge \delta_{B,D}(A) = C \\
 \rho_{a \rightarrow b}(A) \equiv_S \rho_{c \rightarrow d}(B) & \text{if } A \equiv_S B \wedge \delta_{A,B}(a) = c \wedge \delta_{A,B}(b) = d \\
 \chi_{a:f}(A) \equiv_S \chi_{b:g}(B) & \text{if } A \equiv_S B \wedge \delta_{A,B}(a) = b \wedge \delta_{A,B}(f) = g \\
 \sigma_{a=b}(A) \equiv_S \sigma_{c=d}(B) & \text{if } A \equiv_S B \wedge \delta_{A,B}(a) = c \wedge \delta_{A,B}(b) = d \\
 A \times B \equiv_S C \times D & \text{if } A \equiv_S C \wedge B \equiv_S D \\
 A \times_{a=b}(B) \equiv_S C \times_{c=d}(D) & \text{if } A \equiv_S C \wedge B \equiv_S D \wedge \delta_{A,C}(a) = c \wedge \delta_{B,D}(b) = d \\
 \Gamma_{A;a:f}(B) \equiv_S \Gamma_{C;b:g}(D) & \text{if } B \equiv_S D \wedge \delta_{B,D}(A) = C \wedge \delta_{B,D}(a) = b \wedge \delta_{B,D}(f) = g
 \end{array}$$

into a DAG by merging share equivalent parts. While this reduces the changes required for DAG support to a minimum, it makes the cost function very expensive. Besides, if the query graph is already DAG-structured (e.g. for bypass plans), the corresponding tree-structured representation is much larger (e.g. exponentially for bypass plans), enlarging the search space accordingly.

A more general optimization can be done by sharing operators via rename operators. While somewhat difficult to do in a transformation-based plan generator, for a constructive plan generator it is easy to choose a share equivalent alternative and add a rename operator as needed. Logically, the resulting plans behave as if the version without renaming was executed (i.e. as if the plan was a tree instead of a DAG). Therefore, the regular algebraic equivalences can be used for optimization. This issue will come up again when we discuss the plan generator.

4.4 (No) optimal substructure

Optimization techniques like dynamic programming and memoization rely on an optimal substructure of a problem (neglecting physical properties like sortedness or groupedness for a moment). This means that Bellmann's optimality principle holds and, thus, the optimal solution can be found by combining optimal solutions for subproblems. This is true for generating optimal tree-structured query graphs, but is not necessarily true for generating optimal DAGs. To see this, consider Fig. 5, which shows two plans for $A \bowtie B \bowtie C \bowtie B \bowtie C \bowtie D$. The plan on the left-hand side was constructed bottom-up, relying on the optimal substructure. Thus, $A \bowtie B \bowtie C$ was optimized, resulting in the optimal join ordering $(A \bowtie B) \bowtie C$. Besides, the optimal solution for $B \bowtie C \bowtie D$ was constructed, resulting in $B \bowtie (C \bowtie D)$. But when these two optimal partial solutions are combined, no partial join result can be reused. When choosing the sub-optimal partial solutions $A \bowtie (B \bowtie C)$ and $(B \bowtie C) \bowtie D$, the expression $B \bowtie C$ can be shared, which might result in a better plan. Therefore, the optimal DAG cannot be constructed by just combining optimal partial solutions. Our approach avoids this problem by keeping track of sharing opportunities and considering them while pruning otherwise dominated plans.

Fig. 5 Possible non-optimal substructure for DAGs

4.5 Reasoning over DAGs

After looking at various aspects of DAG generation, we now give a formal model to reason about DAGs. More precisely, we specify when two DAGs are equivalent and when one DAG dominates another. Both operations are crucial for algebraic optimization.

As we want to lift tree equivalences to DAGs, we need some preliminaries: We name the equivalences for trees $\equiv_T$ and assume that the following conditions hold for all operators θ , that is the equivalence can be checked on a per operator basis.

$$\begin{aligned} t &\equiv_T t' \Rightarrow \theta(t) \equiv_T \theta(t') \\ t_1 &\equiv_T t'_1 \wedge t_2 \equiv_T t'_2 \Rightarrow (t_1 \theta t_2) \equiv_T (t'_1 \theta t'_2) \end{aligned}$$

These conditions are a fundamental requirement of constructive plan generation, but as seen in Sect. 4.1, they do no longer hold for DAGs in general. However, they hold for DAGs (V, E) that logically are a tree (LT), i.e. that have non-overlapping input for all operators. For an expression e let $\mathcal{A}(e)$ denote the set of attributes produced by e . We then define $LT((V, E))$ for a DAG (V, E) as follows:

$$\begin{aligned} LT((V, E)) \text{ iff } \forall v, v_1, v_2 \in V, (v, v_1) \in E, (v, v_2) \in E : \\ \mathcal{A}(v_1) \cap \mathcal{A}(v_2) = \emptyset \end{aligned}$$

Note that this definition implies that all sharing of intermediate results must be done by renaming attributes.

Using this definition, we can now lift tree equivalences $\equiv_T$ to DAG equivalences $\equiv_D$ for DAGs d and d' : We reuse the tree equivalences directly, but thereby must make sure that the input indeed behaves like a tree, as the equivalences were originally only defined on trees:

$$d \equiv_T d' \wedge LT(d) \wedge LT(d') \Rightarrow d \equiv_D d'$$

Note that the condition $LT(d) \Rightarrow LT(d' \subseteq d)$ holds, therefore partial DAGs that violate the logical tree constraint should be discarded immediately. As a consequence, tests for $LT(d)$ are required only when adding binary (or n -ary) operators, as unary operators cannot produce a new violation.

While lifting the tree equivalences to DAGs is important, they are not enough for DAG optimization, as they only create trees. DAGs can be created by using the notion of share equivalence defined above: If two DAGs d and d' are share equivalent, the two DAGs become equivalent by renaming the result suitably:

$$d \equiv_S d' \Rightarrow d \equiv_D \rho_{\mathcal{A}(d') \rightarrow \mathcal{A}(d)} d'$$

While these two implications are not exhaustive, lifting the tree equivalences to DAG equivalences and reusing intermediate results already allows for a wide range of DAG plans. Additional equivalences can be derived e.g. from [19, 30].

In addition to checking if two plans are equivalent, the query optimizer has to decide which of two equivalent plans is better. Better usually means cheaper, according to a cost function. But sometimes two plans are incomparable, e.g. because one plan satisfies other ordering properties than the other. This is true for both, trees and DAGs. Here, we only look at the costs and DAG-specific limitations to keep the definitions short.

As shown in Sect. 4.4, DAGs cannot be compared by just looking at the costs, as one DAG might allow for more sharing of intermediate results than the other. To identify these plans, we require a labeling function that marks sharing opportunities. Here, we specify only the characteristics, see Sect. 5.3 for an explicit definition.

The DAGs are labeled using a function S . Its codomain is required to be partially ordered. We require that S assigns the same label to share equivalent DAGs. Further the partial ordering between labels must express the fact that one DAG provides more sharing opportunities than another. Thus, for two DAGs $d_1 = (V_1, E_1)$ and $d_2 = (V_2, E_2)$ we require the following formal properties for S :

$$\begin{aligned} S(d_1) = S(d_2) &\text{ iff } d_1 \equiv_S d_2 \\ S(d_1) < S(d_2) &\text{ iff } \exists d'_2 \subset_D d_2 : d_1 \equiv_S d'_2 \end{aligned}$$

Note that that $d'_2 = (V'_2, E'_2) \subset_D d_2$ iff $V'_2 \subset V_2 \wedge E'_2 = E_2|_{V'_2} \wedge d_2$ is a DAG).

Now a plan dominates another equivalent plan if it is cheaper and offers at least the same sharing opportunities:

$$\begin{aligned} d_1 \equiv_D d_2 \wedge \text{costs}(d_1) < \text{costs}(d_2) \wedge S(d_2) \leq S(d_1) \\ \Rightarrow d_1 \text{ dominates } d_2. \end{aligned}$$

Note that the characterization of S given above is overly conservative: one plan might offer more sharing opportunities than another, but these could be irrelevant for the current query. This is similar to order optimization, where the query optimizer only considers differences in interesting orderings [27, 29]. The algorithm presented in Sect. 5.3 improves the labeling by checking which of the operators could produce interesting shared intermediate results for the given query. As only whole subgraphs can be shared, the labeling function stops assigning new labels if any operator in the partial DAG cannot be reused. This greatly improves plan pruning, as more plans become comparable.

5 Plan generator skeleton

5.1 Overview

The main motivation for generating DAG structured execution plans is the ability to share intermediate results. Obviously, the fewer sharing opportunities are missed, the better.

Hence, it is important to allow for a very aggressive sharing during plan generation.

In order to share intermediate results, the plan generator has to detect that partial plans produce the same (or rather equivalent) output. This is a problem for a rule-based plan generator, as the set of operators cannot be hard-coded and tests for equivalence are expensive. We overcome this problem by using a very abstract description of logical plan properties. Instead of using a complex property vector like, e.g., [16], the plan generator uses a set of *logical properties*. The semantics of these logical properties is only known to the rules. They guarantee that two plans are equivalent if they have the same logical properties. A suitable representation is e.g. a bit vector with one bit for each property.

We assume that the logical properties are set properly by the optimization rules. This is all the plan generator has to know about the logical properties! However, to make things more concrete for the reader, we discuss some possible logical properties here. The most important ones are “relation/attribute available” and “operator applied”. In fact, these properties are already sufficient when optimizing only selections and joins as both are freely reorderable. We discuss more complex properties in Sect. 5.2.

As the logical properties describe equivalent plans, they can be used to partition the search space. This allows for a very generic formulation of a plan generator:

```

PLANGEN(goal)
1  bestPlan ← NIL
2  for each optimization rule  $r$ 
3    do if  $r$  can produce a logical property in  $goal$ 
4      then  $rem \leftarrow goal \setminus \{p \mid p \text{ is produced by } r\}$ 
5           $part \leftarrow \text{PLANGEN}(rem)$ 
6           $p \leftarrow \text{BUILDPLAN}(r, part)$ 
7          if  $bestPlan = \text{NIL}$  or  $p$  is cheaper
8            then  $bestPlan \leftarrow p$ 
9  return  $bestPlan$ 

```

This algorithm performs a top-down exploration of the search space. Starting with the complete problem (i.e. finding a plan that satisfies the query), it asks the optimization rules for partial solutions and solves the remaining (sub-) problem(s) recursively. For SPJ queries, this implies splitting the big join into smaller join problems (and selections) and solving them recursively. This algorithm is highly simplified, it only supports unary operators, does not perform memoization etc. Adding support for binary operators is easy but lengthy to describe, see [22] for a detailed discussion. Other missing and essential details will be discussed in Sect. 5.4. Let us summarize the main points of the above approach:

1. The search space is represented by abstract logical properties,

2. the plan generator recursively solves subproblems, which are fully specified by property combinations, and
3. property combinations are split into smaller problems by the optimization rules.

This approach has several advantages. First, it is very extensible, as the plan generator only reasons about abstract logical properties: the actual semantics is hidden in the rules. Second, it constructs DAGs naturally: if the same subproblem (i.e. the same logical property combination) occurs twice in a plan, the plan generator produces only one plan and thus creates a DAG. In reality, this is somewhat more complex, as queries usually do not contain exactly the same subproblem twice (with the same variable bindings etc.), but as we will see in Sect. 5.4, the notion of share equivalence can be used to identify sharable subplans.

5.2 Search space organization

An unusual feature of our plan generator is that it operates on an abstract search space, as the semantics of the logical properties are only known to concrete optimization rules. To give an intuition why this is a plausible concept, we now discuss constructing suitable properties for commonly used queries.

The most essential part of choosing properties is to express the query semantics properly. In particular, two plans must only have the same logical properties when they are equivalent. For the most common and (simplest) type of queries, the selection-projection-join (SPJ) queries using inner joins, it is sufficient to keep track of the available relations (and attributes) and the applied operators. For this kind of queries, two plans are equivalent if they operate on the same set of relations and they have applied the same operators. Therefore using the available attributes/relations and the applied operators as properties is sufficient for this kind of queries. Note that the properties are not only used by the plan generator but also by the optimization rules, e.g. to check if all attributes required by a selection are available.

For more general queries involving e.g. outer joins, the same properties are sufficient, but the operator dependencies are more complex: Using the extended eligibility list concept from [24], each operator specifies which other operators have to be applied before it becomes applicable. These operator dependencies allow for handling complex queries with relatively simple properties, and can be extended to express dependent subqueries etc. Note that we do not insist on a “dense” encoding of the search space that only allows for valid plans (which is desirable, but difficult to achieve in the presence of complex operator). Instead, we allow for a slightly sparse encoding and then guarantee that we only explore the valid search space.

Note that these properties are just examples, which are suitable for complex queries but not mandatory otherwise.

The plan generator makes only two assumptions about the logical properties: Two plans with the same properties are equivalent and properties can be split to form subproblems. Further, not all information about plans have to be encoded into the logical properties. In our implementation, ordering/grouping properties are handled separately, as the number of orderings/groupings can be very large. We use the data structure described in [21] to represent them. This implies that multiple plans with the same logical properties might be relevant during the search, as they could differ in some additional aspects. We will see another example for additional properties in the next section.

5.3 Sharing properties

Apart from the logical properties used to span the search space, each plan contains a *sharing* bit set to indicate potentially shared operators. It can be considered as the materialization of the labeling function S used in Sect. 4.5: The partial ordering on S is defined by the subset relation, that is, one plan can only dominate another plan if it offers at least the same sharing opportunities. We now give a constructive approach for computing S .

When considering a set of share equivalent plans, it is sufficient to keep one representative, as the other plans can be constructed by using the representative and adding a rename (note that all share equivalent plans have the same costs). Analogously, the plan generator determines all operators that are share equivalent (more precisely: could produce share equivalent plans if their subproblems had share equivalent solutions) and places them in equivalence classes. As a consequence, two plans can only be share equivalent if their top-most operators are in the same equivalence class, which makes it easier to detect share equivalence. The equivalence classes that contain only a single operator are discarded, as they do not affect plan sharing. For the remaining equivalence classes, one representative is selected and one bit in the *sharing* bit set is assigned to it.

For example, the query in Fig. 5 consists of 11 operators: A , B_1 , C_1 , B_2 , C_2 , D , $\bowtie_1$ (between A and B_1), $\bowtie_2$ (between B_1 and C_1), $\bowtie_3$ (between B_2 and C_2), $\bowtie_4$ (between C_2 and D) and $\bowtie_5$ (between A and D). Then three equivalence classes with more than one element can be constructed: $B_1 \equiv_S B_2$, $C_1 \equiv_S C_2$ and $\bowtie_2 \equiv_S \bowtie_3$. We assume that the operator with the smallest subscript was chosen as representative for each equivalence class. Then the plan generator would set the *sharing* bit B_1 for the plan B_1 , but not for the plan B_2 . The plan $A \bowtie_1 (B_1 \bowtie_2 C_1)$ would set sharing bits for B_1 , C_1 and $\bowtie_2$, as the subplan can be shared, while the plan $(A \bowtie_1 B_1) \bowtie_2 C_1$ would only set the bits B_1 and C_1 , as the join $\bowtie_2$ cannot be shared here (only whole subgraphs can be shared). The sharing bit set allows the plan generator to detect that the first plan is not dominated by the second, as

the first plan allows for more sharing. This solves the problem discussed in Sect. 4.4.

The equivalence classes are also used for another purpose: When an optimization rule requests a plan with the logical properties produced by an operator, the plan generator first checks if a share equivalent equivalence class representative exists. For example, if a rule requests a plan with B_2 , C_2 and $\bowtie_3$, the plan generator first tries to build a plan with B_1 , C_1 and $\bowtie_2$, as these are the representatives. If this rewrite is possible (i.e. a plan could be constructed), the plan constructed this way is also considered a possible solution.

In general, the plan generator uses sharing bits to explicitly mark sharing opportunities: whenever a partial plan is built using an equivalence class representative, the corresponding bit is set. In more colloquial words: the plan *offers* to share this operator. Note that it is sufficient to identify the selected representative, as all other operators in the equivalence class can be built by just using the representative and renaming the output. As sharing is only possible for whole subplans, the bit must only be set if the input is also sharable. Given, for example, three selections σ_1 , σ_2 and σ_3 , with $\sigma_1(R) \equiv_S \sigma_2(R)$. The two operator rules for σ_1 and σ_2 are in the same equivalence class, we assume that σ_1 was selected as representative. Now the plan $\sigma_1(R)$ is marked as “shares σ_1 ”, as it can be used instead of $\sigma_2(R)$. The same is done for $\sigma_3(\sigma_1(R))$, as it can be used instead of $\sigma_3(\sigma_2(R))$. But for the plan $\sigma_1(\sigma_3(R))$ the *sharing* attribute is empty, as σ_1 cannot be shared (since σ_3 cannot be shared). The plans containing σ_2 do not set the *sharing* property, as σ_1 was selected as representative and, therefore, σ_2 is never shared.

Note that the sharing bits are only set when the whole plan can be shared, but the already existing sharing bits are still propagated even if it is no longer possible to share the whole plan. Otherwise a slightly more expensive plan might be pruned early even though parts of it could be reused for other parts of the query.

Explicitly marking the sharing opportunities serves two purposes. First, it is required to guarantee that the plan generator generates the optimal plan, as one plan only dominates another if it is cheaper and offers at least the same sharing opportunities. Second, sharing information is required by the cost model, as it has to identify the places where a DAG is formed (i.e. the input overlaps). This can now be done by checking for overlapping *sharing* properties. It is not sufficient to check if the normal logical properties overlap, as the plans pretend to perform different operations (which they do, logically), but share physical operators.

5.4 Search Space Exploration under Sharing

After describing the general approach for plan generation, we now present the specific steps required to handle shared

plans. Unfortunately, the plan generator is not a single, concise algorithm but a set of algorithms that are slightly interwoven with the optimization rules. This is unavoidable, as the plan generator has to be as generic as possible and, therefore, does not understand the semantics of the operators. However, there is still a clear functional separation between the different modules: The *plan generator* itself maintains the partial plans, manages the memoization and organizes the search. The *optimization rules* describe the semantics of the operators and guide the search with their requirements. For the specific query, several optimization rules are *instantiated*, i.e. annotated with the query specific information like selectivities, operator dependencies etc. Typically each operator in the original query will cause a rule instantiation. Note that we only present a simplified plan generator here to illustrate the approach, in particular the search space navigation. The detailed algorithms are discussed in [22].

Within the search space, the plan generator tries to find the cheapest plan satisfying all logical properties required for the final solution. Note that in the following discussion, we ignore performance optimization to make the conceptual structure clearer. In practice, the plan generator prunes plans against known partial solutions, uses heuristics like KBZ [15] to get upper bounds for costs etc. However, as these are standard techniques, we do not elaborate on them here.

The core of the plan generator itself is surprisingly small and only consists of a single function that finds the cheapest plans with a given set of logical properties. Conceptually this is similar to the top-down optimization strategy known from Volcano [11, 12]. The search phase is started by requesting the cheapest plan that provides the goal properties.

PLANGEN(goal)

```

1  plans ← memoizationTable[goal]
2  if plans already computed
3    then return plans
4  plans ← create a new, empty PlanSet
5  shared ← goal rewritten using representatives
6  if shared ∩ goal = ∅
7    then plans ← PLANGEN(shared)
8  for each r in instantiated rules
9    do filter ← r.produces ∪ r.required
10   if filter ⊆ goal
11     then sub ← PLANGEN(goal \ r.produced)
12     plans ← plans ∪ {r(p) | p ∈ sub}
13  memoizationTable[goal] ← plans
14  return plans
```

What is happening here is that the plan generator is asked to produce plans with a given set of logical properties. First, it checks the memoization data structure (e.g. a hash table, logical properties → plan set) to see if this was already done

before. If not, it creates a new set (initially empty) and stores it in the memoization structure. Then, it checks if the goal can be rewritten to use only representatives from equivalence classes as described in Sect. 5.3 (if the operators o_1 and o_2 are in the same equivalence class, the logical properties produced by o_2 can be replaced by those produced by o_1). If the rewrite is complete, i.e., the new goal is disjoint from the original goal ($shared \cap goal = \emptyset$), the current problem can be formulated as a new one using only share equivalent operators. This is an application of the DAG equivalence given in Sect. 4.5. Thus, the plan generator tries to solve the new problem and adds the results to the set of usable plans. Afterwards, it looks at all rule instances, checks if the corresponding filter is a subset of the current goal (i.e. the rule is relevant) and generates new plans using this rule. Note that the lines 8–12 are very simplified and assume unary operators. In practice, the optimizer delegates the search space navigation (here a simple $goal \setminus r.produced$) to the rules.

6 Cost calculation

Query optimization always requires to decide whether one plan is better than another plan. This requires calculating the costs of a plan, where the costs can have different meanings. A reasonable choice would be expected query execution time, but the exact meaning is not relevant for the calculation itself. This cost calculation is relatively straight-forward for trees, but becomes much more complex for DAGs. In this section, we first look at the simple tree case, then explain why DAGs are more difficult, and finally present algorithms to calculate the costs for DAGs. We do not assume a concrete cost model here, but base the calculation on two assumptions. First, each operator produces two kinds of costs: the first amounts for producing its result and the other for reading the (already produced) output again. For example, reading the output of a sort operator twice is not twice as expensive as producing and reading it once. Second, we assume that multiple parallel reads (due to the DAG structure) do not cause additional costs. This can be achieved by using a suitable runtime system [20].

During plan generation, the costs are calculated incrementally. This means that the partial plans are annotated with the costs they cause, and when a rule creates a new partial plan, it takes the costs caused by its input and adds the costs for the newly added operator. The costs for the operator itself are the same for trees and DAGs. However, the costs caused by the input can be different (due to sharing, as we will see below). Therefore, we only discuss how to calculate the costs caused by reading the input here. Note that this is only a problem for binary operators, as unary operators cannot construct DAGs (the input can be a DAG, but it

can be treated as a scan with given costs). Therefore, the cost calculation requires a function

$INPUTCOSTS(left, leftReads, right, rightReads) : cost\_t$

which calculates the costs of reading $left$ $leftReads$ times and $right$ $rightReads$ times. This function can then be used to calculate the costs for arbitrary binary operators.

6.1 Calculation for trees

If it is clear that the operators actually form a tree (i.e. the two subgraphs are disjoint; they themselves may form a DAG) the costs can be computed easily: Just multiply the costs with the number of reads and take into account that the first read might be more expensive.

```

INPUTCOSTSTREE( $l, lReads, r, rReads$ )
1   $lCosts \leftarrow lReads * l.costs.further$ 
2   $rCosts \leftarrow rReads * r.costs.further$ 
3   $lDelta \leftarrow l.costs.first - l.costs.further$ 
4   $rDelta \leftarrow r.costs.first - r.costs.further$ 
5   $result.further \leftarrow lCosts + rCosts$ 
6   $result.first \leftarrow lCosts + rCosts + lDelta + rDelta$ 
7  return  $result$ 

```

Here, *first* denotes the first construction of a result and *further* any further access to it.

6.2 Problems when using DAGs

This simple approach does not work for DAGs. Consider the DAG-structured query plan in Fig. 6 (Γ is a group-by operator). Here, the cost of the final blockwise nested loop join (BNL) cannot be calculated in a straightforward top-down calculation. In particular, it cannot be determined by combining the cost of the two input operators. When the topmost join is treated like a normal join in an operator tree, the costs are overestimated, as the shared group-by operator is not taken into account. Since this operator serves two joins simultaneously, its costs should only be counted once. What complicates the calculation even more is that the shared op-

Fig. 6 DAG-structured query plan

erator appears twice on the right-hand side of a nested loop and, therefore, is read multiple times. The actual number of reads can only be determined by looking at all operators occurring in the plan. This makes calculating the costs complex and expensive. Ideally, the cost calculation should touch each operator at most once (resulting in linear runtime) and require only constant additional memory in the plans, as memory is usually tight during plan generation. Achieving both at the same time is difficult, thus, we present two different cost algorithms in the next section that choose different time/space tradeoffs.

6.3 Calculation for DAGs

For DAGs, the calculation is much more complex than for trees. Shared subgraphs that are read by multiple operators have to be taken into account. For the cost calculation, the actual number of passes (i.e. number of plan executions) has to be calculated, which can be lower than the number of reads in the case of sharing.

The main idea is to use a *passes* entry embedded in the plan state to keep track of the number of reads which is reset to zero after each cost calculation (it could also be stored in a side table just used for each calculation). The two input plans are traversed top-down. When visiting an operator for the first time, the costs are calculated as normal and the number of passes is stored. Further traversals can determine that the operator was already used and now only need to check if they require additional passes and calculate the costs accordingly. Some care is needed to accurately compute the number of passes, especially concerning (potential chains of) nested loops and materializing operators like temp.

INPUTCOSTSDAG(*l*, *lReads*, *r*, *rReads*)

```

1  rc  $\leftarrow$  r.rule.DAGCOSTS(r, rReads)
2  lc  $\leftarrow$  l.rule.DAGCOSTS(l, lReads)
3  reset passes to zero
4  result.first  $\leftarrow$  lc.first + rc.first
5  result.further  $\leftarrow$  lc.further + rc.further
6  return result

```

Now the operator rules have to describe how the costs propagate through the operator tree. For the basic scan operations, this is trivial, it just needs to examine *passes* to detect shared subgraphs and to check if additional reads are required:

SCAN::DAGCOSTS(*plan*, *reads*)

```

1  if plan.passes = 0
2  then C  $\leftarrow$  I/O costs for reads passes
3      plan.passes  $\leftarrow$  reads
4  return C
5  if reads > plan.passes

```

```

6  then inc  $\leftarrow$  reads - plan.passes
7      C.first  $\leftarrow$  plan.cost.further * inc
8      C.further  $\leftarrow$  plan.cost.further * inc
9      plan.passes = reads
10     return C
11 return zero costs

```

Simple unary operators like a selection basically behave the same way, they pass the number of passes down to their children. Binary operators basically behave like unary operators. However, they have to visit both input operators. As will be discussed in the next paragraph, the order of traversal is actually important.

NESTEDLOOP::DAGCOSTS(*plan*, *reads*)

```

1  if plan.passes = 0
2  then C  $\leftarrow$  r.rule.DAGCOSTS(r, reads * l.card)
3      C  $\leftarrow$  C + l.rule.DAGCOSTS(l, reads)
4      plan.passes  $\leftarrow$  reads
5  return C
6  if reads > plan.passes
7  then C  $\leftarrow$  r.rule.DAGCOSTS(r, reads * l.card)
8      C  $\leftarrow$  C + l.rule.DAGCOSTS(l, reads)
9      plan.passes  $\leftarrow$  reads
10     return C
11 return zero costs

```

Note that the nested loop code shown can require an exponential runtime, as both input sources are visited (given a chain of nested loops where *left* = *right*, this could result in a runtime of 2^n). However, only the right input source is actually read multiple times. By visiting the right input first, we make sure that the *passes* entry is set to a large value. Ideally, all further visits require a lower or equal number of passes, resulting in a linear time consumption.

While this is true most of the time, some extreme trees indeed trigger the exponential behavior. An example for this is shown in Fig. 7. Depending on the selectivities and the cardinalities of the relations, the left-hand side might actu-

Fig. 7 A DAG requiring exponential runtime

ally read the shared operators more often than the right-hand side: Assume that C consists of 100 tuples, D of 1 tuple and $C \bowtie D$ of 10 tuples. During the cost calculation, $\bowtie_5$ is asked to calculate its costs. As its left-hand side consists of 10 tuples, it asks D for the costs of reading D 10 times. Afterwards, it asks $\bowtie_6$ for its costs, which determines that D has to be read 100 times, which is larger than the previous value, requiring a revisit. The same situation can happen with B , visiting the right-hand side of $\bowtie_3$ twice and thus D four times (as the revisit increases the *passes* entries in the whole subgraph). The same happens for A , which leads to eight accesses to D , doubling with each additional nested loop pair.

In reality, this kind of DAGs do not occur, resulting in linear runtime of `INPUTCOSTSDAG`. However, as they might occur and it is unsatisfactory to have an exponential step during cost calculation, we present an algorithm with linear bounds in the next sections. However, the algorithm is more involved and slower for most DAGs. Therefore, it might be preferable to try the exponential algorithm first. Since the linear case visits each plan node at most twice, it is safe to abort the algorithm after visiting more than $2n$ operators and switch to a linear algorithm. This guarantees both good best case and worst case performance.

The only reason why the first algorithm is exponential instead of linear is that a plan node might be visited again with a *reads* > *passes*. If we can always guarantee that this does not happen, each node is visited at most twice, resulting in a linear runtime. This can be achieved by visiting the plan nodes in topological order: Each plan node passes the number of reads down to its children (iteratively, not recursively). Since the nodes are visited in topological order, the number of reads does not increase after the node is visited, resulting in linear time and only requiring linear space for the topological sort. The only disadvantage is that the topological sort has to be repeated for each cost calculation. Although this can be done in linear time, it results in larger constants than the incremental approach described in the previous section.

A more detailed evaluation of the different algorithms, and another algorithm with linear runtime is included in our technical report [22]. Regardless of the actual algorithm used, calculating the costs of a DAG is much more expensive than calculating the costs of a tree. Fortunately, one can easily check if an operator creates a DAG by inspecting if the *sharing* properties of the input plans overlap, as they mark sharable operators present in the graphs. If they do not overlap, we have the cheap tree case. Otherwise, we can compute a lower bound for the actual costs by taking the maximum costs of each input. If this lower bound is already dominated by a known alternative, the cost calculation can be canceled. Only if the plan seems interesting, the expensive calculation is performed.

7 Evaluation

In the previous sections, we discussed several aspects of optimizing DAG-structured query graphs. However, we still have to show three claims: 1) creating DAG-structured query plans is actually beneficial for common queries, 2) sharing base scans as proposed by Cao et al. [2] is inferior to full DAG support, and 3) the overhead of generating DAG-structured plans is negligible. Therefore, we present several queries for which we create tree-structured and DAG-structured query plans. Both the compile time and the runtime of the resulting plans are compared to see if the overhead for DAGs is worthwhile. For the DAG plans, we created the full DAGs in a single phase optimizer using our DAG reasoning, and the scan sharing plans by a two phase optimizer based upon the tree plans as proposed in [2]. All experiments were executed on a 2.2 GHz Athlon64 system running Windows XP. The plans were executed using a runtime system that could execute DAGs natively [20].

Each operator (join, group-by etc.) is given 1 MB buffer space. This is somewhat unfair against the DAG-structured query plans, as they need fewer operators and could therefore allocate larger buffers. But dynamic buffer sizes would affect the cost model, and the space allocation should probably be a plan generator decision. As this is beyond the scope of this work, we just use a static buffer size here.

As a comparison with the state of the art in commercial database systems, we included results for DB2 9.2.1, which factorizes common subexpressions using materialization of intermediate results. As we were unable to measure the optimization time accurately enough, we only show the total execution time (which includes optimization) for DB2.

7.1 TPC-H

The TPC-H benchmark is a standard benchmark to evaluate relational database systems. It tests ad-hoc queries which result in relatively simple plans allowing for an illustrative comparison between tree- and DAG-structured plans. We used a scale factor 1 database (1 GB).

Before looking at some exemplary queries, we would like to mention that queries without sharing opportunities are unaffected by DAG support: The plan generator produces exactly the same plans with and without DAG support, and their compile times are identical. Therefore, it is sufficient to look at queries, which potentially benefit from DAGs.

Query 11

Query 11 is a typical query that benefits from DAG-structured query plans. It determines the most important subset of suppliers' stock in a given country (Germany in the reference query). The available stock is determined by joining

partsupp, supplier and nation. As the top fraction is requested, this join is performed twice, once to get the total sum and once to compare each part with the sum. When constructing a DAG, this duplicate work can be avoided. The compile time and runtime characteristics are shown below:

	tree	Scan Sharing	Full DAG	DB2
compilation [ms]	10.5	10.5	10.6	–
execution [ms]	4793	4256	2436	3291

While the compile time is slightly higher when considering DAGs (profiling showed this is due to the checks for share equivalence), the runtime is much smaller. The corresponding plans are shown in Fig. 8: In the tree version, the relations partsupp, supplier and nation are joined twice, once to get the total sum and once to get the sum for each part. In the DAG version, this work can be shared, which nearly halves the execution time. The scan sharing approach from [2] only slightly better than the tree plan as it still has to join twice. DB2 performs between trees and full DAGs, as it reuses intermediate results but materializes to support multiple reads.

Query 2

Query 2 selects the supplier with the minimal supply costs within a given region. Structurally, this query is similar to

Query 11, as it performs a large join twice, once for the result and once to get the minimum. However, it is more complex, as the nested minimum query depends on the outer query. We assume that the rewrite step unnests the query (by grouping and using a join), but still the nested query lacks a relation used in the outer query, which prevents sharing the whole join. To allow for greater sharing, the relation (and the corresponding predicates) can be re-added by a magic set like transformation [19]. Here, we consider three alternatives: Normal tree construction, DAG construction and DAG construction with rules for magic set transformation enabled. We omit the sharing approach from [2] here, as it is nearly identical to the second alternative for this query (one additional (cheap) join). Compile time and runtime are shown below.

	tree	DAG	DAG + magic set	DB2
compilation [ms]	9.3	9.2	9.7	–
execution [ms]	11933	7480	3535	8705

The compile times for tree and DAG are about the same (the DAG is slightly faster, as it can ignore some dominated alternatives), while the magic set variant is about 5% slower due to the increased search space. The runtime behavior of the alternatives is very different. Figure 9 shows the plans. The tree variant simply calculates the outer query and the nested query independently and joins the result. The DAG

Fig. 8 Execution plans for TPC-H Query 11

Fig. 9 Execution plans for TPC-H Query 2

variant tries to reuse some intermediate results (reducing the runtime by 37%). It still performs most of the joins in both parts, as the subqueries are not identical. It would be possible to share more by applying the join with *part* later, but this does not pay off due to selectivity of the join. When using the magic set transformation, large parts of the query become equivalent, which results in much greater sharing and also reduced aggregation effort. The consequence is a runtime reduction by 70% compared to the tree variant. DB2 performs better than the tree plan (again due to materializing intermediate results), but worse than both DAG plans.

7.2 Summary

The experiments have shown that by considering sharing intermediate results, the compile time is mainly unaffected. It is somewhat affected when additionally magic sets are used (Query 2). Still, the compile time is negligible and clearly dominated by the runtime. The runtime effect of DAG support is non-negligible, as sharing can drastically reduce the runtime of many queries. Further, the real benefit comes from sharing intermediate results. Simply sharing scans as proposed in [2] improves runtimes a bit, but not nearly as much as full DAG support.

8 Conclusion

We presented the first plan generator which exploits the full potential of shared subplans by being able to explore the complete search space in a fashion that integrates reasoning about sharing and plan generation.

As we have seen, generating DAG-structured query evaluation plans is not easy. To solve the problem, we introduced the novel notion of *share equivalence* and saw that it is orthogonal to the standard notion of equivalence. Rules to reason about share equivalence made this hard to decide notion manageable for the plan generator.

In order to be able to build upon well-known algebraic equivalences, we separated regular equivalence and share equivalence. The former notion is applied only to logical plans which thus remain tree-structured. Reasoning about sharing is moved exclusively to the physical level. The central idea here was to annotate sharing possibilities explicitly and let the plan generator reason about them. To make this more efficient, plans are grouped into equivalence classes where those plans leading to the same sharing possibilities find themselves in the same class. Thereby, the reasoning step was reduced to a simple and efficient rewrite of the optimization goal. The resulting plan generator is the first one that guarantees two important features: (1) all sharing possibilities are detected and consid-

ered, and (2) the resulting plan is guaranteed to be optimal. Finally, experiments demonstrated that compared to generating tree-structured plans, the plan generation overhead is negligible, while the gains during query execution are tremendous.

There is plenty of room for future research. Let us mention just two important areas. The first is to develop and explore more optimization techniques that inherently require DAG support. With bypass plans, we have already seen an example of them. The second area where DAGs play a crucial role is stream processing.

Acknowledgement We thank Simone Seeger for her help preparing the manuscript.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Babcock B, Babu S, Datar M, Motwani R, Thomas D (2004) Operator scheduling in data stream systems. *VLDB J* 13(4): 333–353
2. Cao Y, Das G, Chan C-Y, Tanx K-L (2008) Optimizing complex queries with multiple relation instances. In: *SIGMOD*, pp 525–538
3. Chamberlin D (1998) A complete guide to DB2 universal database. Morgan Kaufmann, San Francisco, CA
4. Dalvi NN, Sanghai SK, Roy P, Sudarshan S (2001) Pipelining in multi-query optimization. In: *PODS*
5. Galindo-Legaria CA, Joshi M (2001) Orthogonal optimization of subqueries and aggregation. In: *SIGMOD*
6. Garcia-Molina H, Ullman JD, Widom J (1999) Database System Implementation. Prentice Hall, Upper Saddle River, New Jersey
7. Gassner P, Lohman GM, Schiefer KB, Wang Y (1993) Query optimization in the ibm db2 family. *IEEE Data Eng Bull* 16(4): 4–18
8. Graefe G (1990) Encapsulation of parallelism in the volcano query processing system. In: *SIGMOD*, pp 102–111
9. Graefe G (1993) Query evaluation techniques for large databases. *ACM Comput Surv* 25(2):73–170
10. Graefe G (1994) Volcano – an extensible and parallel query evaluation system. *IEEE Trans Knowl Data Eng* 6(1):120–135
11. Graefe G (1995) The cascades framework for query optimization. *IEEE Data Eng Bull* 18(3):19–29
12. Graefe G, McKenna WJ (1993) The volcano optimizer generator: Extensibility and efficient search. In: *ICDE*, pp 209–218
13. Haas LM, Freytag JC, Lohman GM, Pirahesh H (1989) Extensible query processing in starburst. In: *SIGMOD*, pp 377–388
14. Harizopoulos S, Shkapenyuk V, Ailamaki A (2005) Qpipe: A simultaneously pipelined relational query engine. In: *SIGMOD*, pp 383–394
15. Krishnamurthy R, Boral H, Zaniolo C (1986) Optimization of nonrecursive queries. In: *VLDB’86*, pp 128–137
16. Lohman GM (1988) Grammar-like functional rules for representing query optimization alternatives. In: *SIGMOD*, pp 18–27
17. Lorie RA (1974) Xrm – an extended (n-ary) relational memory. IBM Research Report, G320-2096

18. Maier D (1983) The Theory of Relational Databases. Computer Science Press, Rockville, Maryland
19. Mumick IS, Finkelstein SJ, Pirahesh H, Ramakrishnan R (1990) Magic is relevant. In: SIGMOD, pp 247–258
20. Neumann T (2005) Efficient generation and execution of dag-structured query graphs. PhD thesis, Universität Mannheim
21. Neumann T, Moerkotte G (2004) An efficient framework for order optimization. In: ICDE, pp 461–472
22. Neumann T, Moerkotte G (2008) Single phase construction of optimal DAG-structured QEPs. Research Report MPI-I-2008-5-002, Max-Planck-Institut für Informatik
23. Neumann T, Moerkotte G (2009) Single phase construction of optimal DAG-structured QEPs. In: BTW
24. Rao J, Lindsay BG, Lohman GM, Pirahesh H, Simmen DE (2001) Using eels, a practical approach to outerjoin and antijoin reordering. In: ICDE, pp 585–594
25. Roy P (1998) Optimization of dag-structured query evaluation plans. M.tech. thesis, Dept. of Computer Science and Engineering, IIT-Bombay, January 1998
26. Roy P, Seshadri S, Sudarshan S, Bhohe S (2000) Efficient and extensible algorithms for multi query optimization. In: SIGMOD
27. Selinger PG, Astrahan MM, Chamberlin DD, Lorie RA, Price TG (1979) Access path selection in a relational database management system. In: SIGMOD, pp 23–34
28. Sellis TK (1988) Multiple-query optimization. ACM Trans Database Syst 13(1):23–52
29. Simmen DE, Shekita EJ, Malkemus T (1996) Fundamental techniques for order optimization. In: SIGMOD, pp 57–67
30. Steinbrunn M, Peithner K, Moerkotte G, Kemper A (1995) Bypassing joins in disjunctive queries. In: VLDB, pp 228–238
31. Zhou J, Larson P-Å, Freytag JC, Lehner W (2007) Efficient exploitation of similar subexpressions for query processing. In: SIGMOD, pp 533–544

Thomas Neumann received his diploma and PhD degree from the University of Mannheim. In 2005 he joined the databases and information systems group at the Max-Planck Institute for Computer Science in Saarbrücken. His research is focused on various aspects of query optimization and query processing, in particular for RDF data.

Guido Moerkotte is a member of the database research group of the University of Mannheim. He is mainly interested in query optimization.