

Efficiently Performing Consistency Checks for Multi-Dimensional Preference Trade-Offs

Christoph Lofi, Wolf-Tilo Balke, Ulrich Güntzer

Abstract— Skyline Queries have recently received a lot of attention due to their intuitive query capabilities. Following the concept of Pareto optimality all ‘best’ database objects are returned to the user. However, this often results in unmanageable large result set sizes hampering the success of this innovative paradigm. As an effective remedy for this problem, trade-offs provide a natural concept for dealing with incomparable choices. Such trade-offs, however, are not reflected by the Pareto paradigm. Thus, incorporating them into the users’ preference orders and adjusting skyline results accordingly needs special algorithms beyond traditional skylining.

For the actual integration of trade-offs into skylines, the problem of ensuring the consistency of arbitrary trade-off sets poses a demanding challenge. Consistency is a crucial aspect when dealing with multi-dimensional trade-offs spanning over several attributes. If the consistency should be violated, cyclic preferences may occur in the result set. But such cyclic preferences cannot be resolved by information systems in a sensible way. Often, this problem is circumvented by restricting the trade-offs’ expressiveness, e.g. by altogether ignoring some classes of possibly inconsistent trade-offs. In this paper, we will present a new algorithm capable of efficiently verifying the consistency of any arbitrary set of trade-offs. After motivating its basic concepts and introducing the algorithm itself, we will also show that it exhibits superior average-case performance. The benefits of our approach promise to pave the way towards personalized and cooperative information systems.

Index Terms— Database Query Processing, Personalization, Skyline Queries, Preferences, Trade-Offs, Pareto Optimality, Multi-Objective Optimization

Paper Category: Technical Solution Paper

I. INTRODUCTION

PREFERENCE-based queries are important to personalize the massive amount of information residing in today’s databases and information systems. In contrast to simple SQL-style queries users are enabled to state their wishes and a suitable degree of relaxation and filtering can be applied when

queries are evaluated against a specific database instance. Since users cannot be expected to compare database items in a pairwise fashion, preferences are usually specified separately with respect to each query attribute. To answer a complex query, the attribute preferences are then combined following the Pareto semantics, i.e. all items showing equal or worse values with respect to all attributes than some other items are removed from the result set. This class of queries is referred to as *skyline queries* [8].

However, the good personalization behavior comes at the cost of rather big result set sizes, because especially for larger numbers of query attributes, many items become incomparable. This is even more problematic, since real world user preferences often have to face anti-correlation between at least some attributes in the underlying database instance. For instance, in e-commerce applications users tend to prefer less expensive offers, but also like a high degree of quality. Thus, in practical applications suitable compromises have to be found. But such compromises cannot be expressed in the framework of Pareto optimality: incomparable results (e.g. better in quality, but worse in price vs. less expensive and worse in quality) are always part of the skyline.

In [2] the concept of trade-offs was first introduced to extend the Pareto semantics and allow for compromises based on user interaction. A user simply specifies a domination or equivalence relationship between two combinations of values with respect to several attributes. The Pareto product order then can be extended by this new user-specific information. This concept is valuable for extending the expressiveness of skyline queries. But stating such trade-off information ‘amalgamates’ the attributes concerned. This complicates the query evaluation as the separability characteristic [9] of the Pareto semantics (i.e. the ability to decompose a relation on database objects back to relations on attribute values) is lost.

Moreover, to be consistent Pareto product orders just need the underlying base preferences to be acyclic, which is easy to ensure. In contrast, trade-offs might introduce inconsistencies ‘through the back door’. That means, if several trade-offs span over a set of non-disjoint attribute sets, the induced preference order might contain cyclic preference. Unfortunately, these cycles are hard to detect. A first criterion about how preference cycles can be detected even in complex product orders has been given in [2]. But when a larger number of attributes is concerned, materializing the complete product order is a costly (and mostly impractical) operation. In this paper we propose a way to check a set of trade-offs for consistency without having to materialize the product order.

Manuscript received March 18, 2008. Part of this work was supported by a grant of the German Research Foundation (DFG) within the Emmy Noether Program of Excellence.

Christoph Lofi and Wolf-Tilo Balke are with L3S Research Center, Leibniz University Hannover, Appelstr. 9a 30167 Hannover, Germany, {balke, lofi}@l3s.de.

Ulrich Güntzer is with Institute of Computer Science, University of Tübingen, Sand 13, 72076 Tübingen, Germany, ulrich.güntzer@informatik.uni-tübingen.de.

We will present a simple algorithm that (re-)uses a tree-shaped data structure which incrementally integrates one trade-off after the other. Our algorithm uses templates of possible trade-off instantiations to derive all possible inconsistencies in an iterative fashion. For each trade-off, it checks the consistency with all trade-offs that have been integrated before. If the new trade-off may lead to inconsistencies, it will immediately be rejected. Otherwise the entire set of trade-offs can be safely evaluated against the database instance. Our extensive experiments investigate typical features like e.g., runtime, memory usage, and scalability and show the practical applicability of our approach. In brief, we present a novel approach that allows to efficiently guarantee the consistency of a set of user trade-offs without having to access any database items and thus prevents costly failures at query evaluation time.

Our paper is organized as follows: First, we provide a brief overview on related approaches for skyline evaluation and result refinement in chapter II. Then, we introduce the observations and definitions necessary to develop an algorithm for checking trade-off consistency in chapter III. The resulting algorithm is topic of chapter IV. In chapter 0, we evaluate the algorithm and will outline its superior performance. We close with a short summary and outlook.

II. RELATED WORK

The skyline paradigm gained a lot of momentum in recent years. The easy querying and the generally intuitive result sets comprising all optimal items are appealing characteristics. Since deriving the actual Pareto-optimal sets is computationally expensive, much research work has focused on the efficient evaluation of skyline queries using different styles of algorithms and further optimizations like indexes, see e.g. [8], [14]. However, not only the consumed computation time for query evaluation poses restrictions for the practical applicability of the Skyline paradigm. For larger numbers of query attributes, also result set sizes often become unmanageable, especially when facing a high degree of anti-correlation between different attributes in the data set. In fact, already for only about 5 to 10 attributes, skylines tend to become unmanageably large and can easily contain 30% or more of the entire database instance (see experiments in e.g., [7], [8], [11], [12]). In order to keep result sets manageable, two major approaches have been devised: On one hand, there are techniques using user independent structural properties and heuristics to reduce the skyline set without any further interaction (often called user-oblivious approaches). Here, techniques range from extending the Pareto semantics [6], over providing representative samples of the skyline [7], or giving users an overview by (approximately dominating) representatives [13], [16], to a structural result set ranking based on subspace-skyline occurrences of database items [17].

On the other hand, a more user-centered reduction and focusing the skyline set can be archived by incorporating additional information directly provided by the user. Most prominent among these approaches are techniques that allow users to interactively modify and extend their preferences in a cooperative fashion, see e.g. [10], [15]. In particular, a user might consider some attributes in the query to be more important than others. This fact is used in [15] where the user is enabled

to provide a total order on attributes. Then, the evaluation first focuses on those subspace-skylines containing only the more interesting attributes.

Trade-Offs [5] represent another approach for integrating user-provided information. Trade-offs, as part of our everyday's decision making, are indeed a very natural concept. They are able to introduce information about compromises a user is willing to accept. Such information cannot be expressed by traditional Pareto skylines. A summary of basic concepts for integrating trade-offs and interfaces for the necessary user interactions can be found in [5] and [3].

The importance of avoiding cyclic preferences for deterministic query processing is easy to see. First criteria and consistency checks were already established by [2] which, however, still relied on the materialization of the entire product order and were therefore computationally prohibitive. An obviously sufficient criterion to ensure the consistency of trade-offs is to restrict trade-offs to disjoint attribute sets only. In [4], a slightly less restrictive and easy to ensure criterion was introduced. However, this so-called characteristic of being 'cartesian' was only proven to be sufficient, but not necessary. Hence, none of the former approaches were able to provide a complete and manageable algorithm for checking a set of arbitrary trade-offs for consistency.

III. CONCEPTS FOR INTEGRATING PREFERENCE TRADE-OFFS

In this section we introduce the basic concepts needed to specify base preferences, aggregate them to a product order, and enhance the resulting order consistently with user-specific trade-off information. We will always assume that a query containing a set of attributes is given and for each attribute, a base preference is defined.

Definition 1: (base preferences and Pareto aggregation)

A base preference P on an attribute A with domain D is a strict partial order on values in D . Several base preferences P_1, \dots, P_n can be aggregated to a product order on $(D_1 \times \dots \times D_n)^2$ under the Pareto semantics such that $\forall x, y \in (D_1 \times \dots \times D_n)$:

$$(y \triangleright x) \Leftrightarrow \forall i: y_i \geq x_i \wedge \exists i: y_i > x_i, (1 \leq i \leq n)$$

The resulting product order can be enhanced by amalgamating subsets of attributes and specifying trade-offs between specific (previously incomparable) values.

Definition 2: (trade-offs)

A trade-off $t := (y_\mu \triangleright x_\mu)$ on a (sub-)set $\mu \subseteq \{1, \dots, n\}$ of attribute indices with $x_\mu, y_\mu \in \prod_{i \in \mu} D_i$ defines a preference relationship between two sets of (previously incomparable) domain values. It can be defined on all attributes or just on a subset of attributes. In the latter case it automatically enhances the preference order by all possible object extensions following the ceteris paribus semantics. That means preference relationships are added between all objects featuring domain values y_μ , resp. x_μ regarding attributes with index in μ and equivalent values regarding attributes with index not in μ . The value y_μ is referred to as the trade-off's *precondition*, whereas the value x_μ is called the *postcondition*.

In the following we will always refer to the aggregation of all base preferences following the Pareto semantics as the ‘*product order*’, whereas we will refer to a product order into which one or several trade-offs have been integrated as ‘*enhanced product order*’. For ease of use, in the following we will notate trade-offs in the form of generalized object pairs in square brackets with the values for attributes not in μ replaced by wildcards ‘ \star ’.

Example: Let us assume we have three attributes over binary domains $\{0, 1\}$ in a query and the base preferences are given by the natural order $1 > 0$. A user might specify a trade-off spanning over the first two attributes to make a specific combination comparable, e.g. $[1, 0, \star] \triangleright [0, 1, \star]$. Following the ceteris paribus semantics this trade-off would introduce two new preference relationships into the enhanced product order, namely $(1, 0, 0) \triangleright (0, 1, 0)$ and $(1, 0, 1) \triangleright (0, 1, 1)$.

Such trade-offs extend the expressiveness of the preference order beyond the Pareto semantics. However, their consistent integration is not easy. In contrast to the simple product order, preference cycles in the enhanced product order can occur as more trade-offs are added, even if each individual trade-off is consistent.

Example: Consider a set of three trade-offs:

$t_1: [1, 0, \star] \triangleright [0, 1, \star]$

$t_2: [\star, 1, 0] \triangleright [\star, 0, 1]$

$t_3: [0, \star, 1] \triangleright [1, \star, 0]$

Now, let’s consider some database object $(1, 0, 0)$. Applying trade-off t_1 we can deduce that $(1, 0, 0) \triangleright (0, 1, 0)$. Now trade-off t_2 applies, resulting in $(0, 1, 0) \triangleright (0, 0, 1)$ and using trade-off t_3 finally results in $(0, 0, 1) \triangleright (1, 0, 0)$. We thus get the complete trade-off sequence

$(1, 0, 0) \triangleright (0, 1, 0) \triangleright (0, 0, 1) \triangleright (1, 0, 0)$

which is obviously a cyclic preference, i.e. trade-offs t_1, t_2, t_3 and the base preferences are not consistent.

Thus, trade-offs on overlapping attribute sets can easily become inconsistent. But the inconsistency of several trade-offs is often very hard to see for the user. This is especially true if the trade-offs have been elicited implicitly, e.g. by comparing sample database items like in example critiquing frameworks [18]. To protect the information system from running into cyclic preferences, the user input therefore has to be checked for inconsistent information. In case that an inconsistent trade-off is detected, the system has to reject at least the last specified trade-off.

We will now investigate the nature of the problem and show a way to detect inconsistencies between trade-offs. Before the first trade-off is added, the Pareto product order is obviously acyclic, if and only if all base preferences are acyclic. This is a direct result of the separability characteristics of the Pareto semantics. However, when adding trade-offs, the separability is always lost and inconsistencies may occur by concatenations of base preferences with trade-offs.

In any case the question of consistency of a set of trade-offs with a set of base preferences should always be independent of

the particular database instance. Hence, we have to make sure that no combination of preferences and trade-offs for an arbitrary set of objects (i.e. arbitrary attribute value combinations) can cause a cyclic preference. The following observation states that such preference cycles in the enhanced product order are caused by sequences of trade-offs, where the post condition of the previous trade-off always matches or dominates the pre-condition of the following trade-off.

Observation 1: A cyclic preference in the enhanced product order causes that some value tuple (a_1, \dots, a_n) is preferred over itself. This of course violates the irreflexivity of strict preference orders, in this case the enhanced product order. The cycle may stretch over several value combinations using some (finite) sequence of (base) preferences interleaved with trade-offs. This sequence can be applied on (a_1, \dots, a_n) and also ends with (a_1, \dots, a_n) . Moreover, for any adjacent value combinations during the sequence (b_1, \dots, b_n) and (c_1, \dots, c_n) holds either $(b_1, \dots, b_n) \geq (c_1, \dots, c_n)$, or there exists some trade-off $t_i := (y_\mu \triangleright x_\mu)$ on a subset of attributes μ such that restricted to μ holds: $(b_1, \dots, b_n)|_\mu \geq y_\mu$ and $x_\mu \geq (c_1, \dots, c_n)|_\mu$.

We now know how trade-offs are integrated into the enhanced product order and how they can connect to each other. The sequences of connections are also the key to verify the consistency of a set of trade-offs and base preferences.

Observation 2: Since base preferences are required to be acyclic, using only preferences from the product order can never result in cyclic preferences in the enhanced product order. Thus, the consistency check has only to focus on the trade-offs. In case that a preference cycle exists, any sequence of trade-offs $(t_{i_1}, \dots, t_{i_l})$ instantiating the cycle can be repeated in the same way and therefore forms a typical pattern of two adjacent duplicate applications of trade-offs $(t_{i_1}, \dots, t_{i_l}, t_{i_1}, \dots, t_{i_l})$. Hence, as cyclic preferences always create a typical pattern (of any length), we can deduce that if and only if such adjacent duplicate patterns exist with respect to a set of trade-offs, a cyclic preference exists in the enhanced product order.

This observation directly leads to the following lemma that will serve as a basis for the fault detection of our later algorithm:

Lemma 1: If a trade-off sequence $(t_{i_1}, \dots, t_{i_l}, t_{i_1}, \dots, t_{i_l})$ can be applied to any generic tuple of domain values, the enhanced product order contains a cyclic preference (and is thus inconsistent).

Proof: Let $\mu := \mu_{i_1} \cup \dots \cup \mu_{i_l}$ be a set of indices and v be some index $v \in \mu$. Let c_v be the value of the v -th attribute of an value tuple that is a result of applying the trade-off sequence $(t_{i_1}, \dots, t_{i_l})$.

Then every c_v can be characterized as follows:

Let j_1, \dots, j_r be the set of sub-indices with $1 \leq j_1 < \dots < j_r \leq l$ such that $v \in \mu_{i_{j_1}} \cap \mu_{i_{j_2}} \cap \dots \cap \mu_{i_{j_r}}$. Then $t_{i_{j_r}}$ is the last trade-off modifying the v -th attribute. This enforces that c_v has the value of the v -th attribute of the post-condition of $t_{i_{j_r}}$. This is still true when $(t_{i_1}, \dots, t_{i_l})$ is immediately applied for a second time.

Moreover, since this holds for every v , we can conclude that after applying $(t_{i_1}, \dots, t_{i_l})$ twice directly adjacent, the resulting value tuple after l and $2l$ trade-offs actually have the same values c_v for all attributes with $v \in \mu$. All other attributes were not affected by $(t_{i_1}, \dots, t_{i_l})$ at all (and have thus still their initial value), i.e. the value tuples after the application of l and $2l$ trade-offs are identical and therefore there exists a cyclic preference induced by the trade-offs. ■

Checking all possible trade-off sequences seems like a lot of effort. But even when applying arbitrary trade-off sequences on a generic object, the values in the postconditions of the trade-offs have to be adhered to, as stated in observation 1. For many value tuples therefore it is simply not possible to append another trade-off and the respective sequence ends. Therefore, the possible length of trade-off sequences in the case of a consistent set of trade-offs is always limited, which is important for the termination of our later algorithm.

Observation 3: When enhancing a product order by trade-offs, the set of trade-offs provided by the user is always finite. Moreover, the intermediate value tuples in each trade-off sequence do only consist of wildcards or specific values provided by the postconditions of those trade-offs already applied.

Since there is only a finite number of trade-offs, also the possible combinations of values in these intermediate tuples occurring during a sequence is limited. Thus, if there is no cyclic preference, there can only be sequences of limited length without creating two identical intermediate value tuples. If on the other hand, two identical intermediate value tuples would have been created, it is obviously possible to apply the same sub-sequence that created the second tuple again and we would be able to detect a directly adjacent duplicate pattern in the trade-off sequence.

IV. ALGORITHM

With the above observations and the fault detection condition we are now ready to design an algorithm to find all possible inconsistencies in a given set of base preferences enhanced by trade-offs.

A. Basic Algorithm for Consistency Checks

We start by constructing a tree structure containing all trade-off sequences that can possibly be build with the current set of trade-offs. The nodes of this tree always represent more or less generic objects with restriction on certain attributes (imposed by the postconditions of the trade-offs applied).

The tree is initialized as a root node which poses no restrictions on any attribute values. When considering an additional trade-off, all nodes of the existing tree are expanded, when-

ever the trade-off can be applied to them. Such an expansion will append additional nodes representing the consequences of the applied trade-off. The edges from the parent nodes to the new nodes are then labeled with the responsible trade-offs. Finally, we try to apply all already stated trade-offs to the new nodes.


Every time a node is added, we check the sequence of trade-offs from the root to the current node. If that sequence contains an adjacent duplicate pattern according to observation 2 in the previous section, the set of current trade-offs is not consistent and the algorithm terminates with a failure.

Let us reconsider the previous example and the successive integration of the three trade-offs t_1 , t_2 , and t_3 .

Example: We aim at detecting potential trade-off inconsistencies while abstracting from specific database instances. Therefore, we start with a generic root object $(*, *, *)$. Each trade-off applies to this generic object as $*$ can be replaced with any value.

First, we integrate the trade-off $t_1: [1, 0, *] \triangleright [0, 1, *]$ – this is possible since the root node $(*, *, *)$ dominates the trade-off's precondition $(1, 0, *)$. The trade-off will end at some object of the form $(0, 1, *)$ which is added as a new node to the tree.

After applying only t_1 , we obtain the following tree:




Now, we try to apply t_1 to the new node. Since $(0, 1, *)$ does not dominate the trade-off's precondition $[1, 0, *]$, it cannot be applied again. In case we would have been able to apply t_1 again, we would immediately have been able to detect the adjacent duplicate pattern (t_1, t_1) in the trade-off sequence and thus an inconsistency.

As the next step, the second trade-off t_2 has to be added. The tree therefore has to be expanded accordingly for both nodes.

Example: Trade-off $t_2: [* , 1, 0] \triangleright [* , 0, 1]$ can be applied for both nodes of the current tree, since $(*, *, *)$ fulfills t_2 's precondition trivially and also the new node of the previous step $(0, 1, *)$ dominates the precondition $[*, 1, 0]$. This in turn results in two new nodes: $(0, 0, 1)$ and $(*, 0, 1)$. For each of these nodes, we again have to apply all previous trade-offs, i.e. t_1 and t_2 , recursively.

This creates following tree:



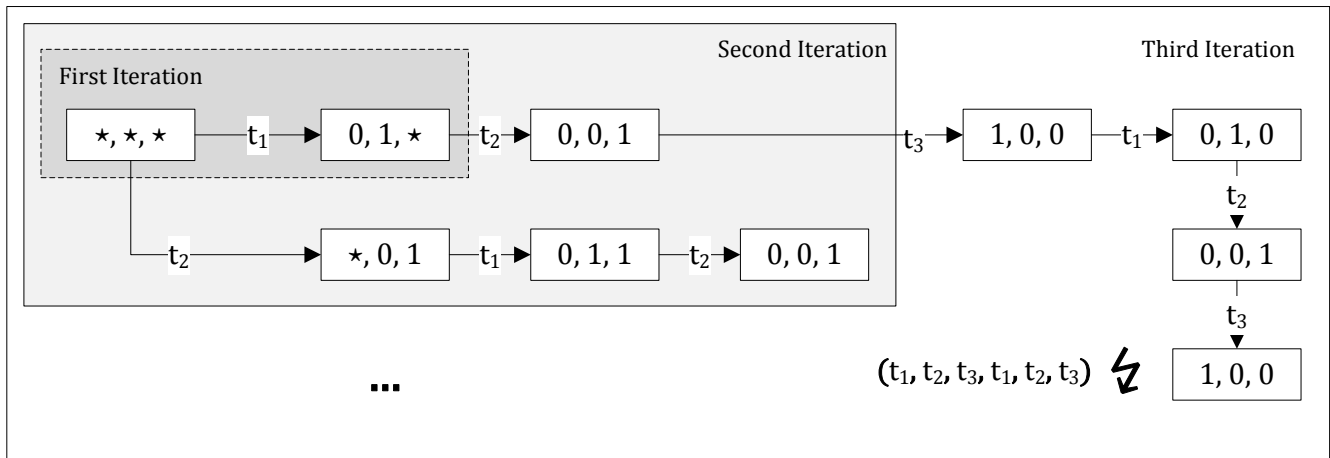


Figure 1: Search tree for three trade-offs until termination due to inconsistency

Example: The third trade-off $t_3: [0, *, 1] \triangleright [1, *, 0]$ is applied in a similar fashion. Again, t_3 extends all nodes that match its precondition $[0, *, 1]$. Figure 1 illustrates the effects of the application of t_3 : when extending node $(0,0,1)$ recursively, we soon encounter a sequence of trade-offs showing an adjacent duplicate pattern $(t_1, t_2, t_3, t_1, t_2, t_3)$ and thus proving that the set of trade-offs t_1, t_2, t_3 is indeed inconsistent and the algorithm consequentially terminates with a failure.

Please note that in Figure 1 all nodes are omitted which would have been induced by t_3 after the inconsistency was detected (this especially includes those nodes of t_3 extending from the root).

In the following, we finally present the complete algorithm. It uses following data structures:

- *Template*:
 - Contains an array of values and wildcards (denoted as ‘*’), e.g., $[\star, 12, 18, \star]$
- *Trade-Off*:
 - Contains two templates *pre* and *post* of equal length and featuring wildcards at the same positions
 - *pre* and *post* do not dominate each other, e.g., *pre*: $[\star, 12, 18, \star]$; *post*: $[\star, 7, 22, \star]$
- *Node*:
 - Each node has a set of *labeled* child nodes. The label denotes which trade-off was used to create the child.
 - Each node has a template as content.

Algorithm: Check Trade-Off Consistency

1. *Initialize Algorithm*
 - 1.1. Initialize *root*, a node with a template consisting only of wildcards and an empty list of children
 - 1.2. Initialize *allTradeOffs*, an empty list of trade-offs
2. *Check Trade-Off Consistency*
 - 2.1. Request a new trade-off t from the user

- 2.2. Add t to the list *allTradeOffs*
- 2.3. Call *expandNode(root, {t})*;
- 2.4. **If** the user wants to specify another trade-off
 - 2.4.1. **Goto** 2

Procedure *expandNode*(Node *node*, List of trade-offs *newTradeOffs*)

1. **For each** *child* in *node.children*
 - 1.1. Call *expandNode(child, newTradeOffs)*
2. **For each** t in *newTradeOffs* with $node.template \triangleright t.pre$ (i.e. t is applicable on the current node)
 - 2.1. **If** the trade-off labels of the path from the root-node to the current node followed by t contains an adjacent duplicate sequence
 - 2.1.1. **Exit** Algorithm with failure: Set of trade-offs inconsistent
 - 2.2. Create template *newTemplate* containing all values of $t.post$ where wildcards are replaced by the respective values in *node.template*
 - 2.3. Create a new node *newNode* having *newTemplate* as template
 - 2.4. Call *expandNode(newNode, allTradeOffs)*
 - 2.5. Add *newNode* to *node.children* and label the respective edge with t

B. Optimization of Matching Pre-/Postconditions

After establishing the basic algorithm, we can focus on an additional optimization. It aims at reducing the number of template comparisons required in *expandNode* step 2 ($node.template \triangleright t.pre$). While a single template comparison is a relatively cheap operation, it has to be performed for each node of the tree for every new trade-off. With growing tree sizes, template comparisons easily contribute a significant amount of the required overall effort.

However, in each set of trade-offs usually some sequences are impossible because the pre- and postconditions of trade-offs are incomparable, i.e. they will never match. Such trade-offs obviously cannot be applied adjacently in a sequence. Therefore in many cases the result of pre-/postcondition compari-

sons can be predicted by simply considering the current trade-off and the trade-off labeling the edge which leads to the current node.

Example: Consider two trade-offs $t_1: [1, 0, \star] \triangleright [0, 1, \star]$ and $t_4: [1, \star, 0] \triangleright [0, \star, 1]$. t_4 can never be applied directly after t_1 as the postcondition of t_1 does not dominate the precondition of t_4 regardless of the values of the wildcards. In the same way t_4 's postcondition will never dominate t_1 's precondition. In fact, t_1 and t_4 can never occur adjacently in any sequence of trade-offs.

Hence, a good optimization of the algorithm is to precompute a matrix containing the information, whether it possible at all (in the most general case) that a given trade-off can be applied directly after some other. The algorithm is now modified in such a way that it performs an (extremely cheap) lookup in the matrix before starting to compare the two respective templates. If the value in the matrix indicates that the current trade-off can never be applied on the actual node (by considering the trade-off which created the current node), the actual comparison can be skipped. Comparisons only need to be performed if it is generally possible to apply the trade-off. Moreover, it can be restricted to comparing only those attributes containing wildcards in any of the two trade-offs.

V. EXPERIMENTAL EVALUATION

In this section, we evaluate our algorithm using various synthetic performance tests. The usage of synthetic tests was necessary as there is no practical test set of typical user trade-offs readily available. In each evaluation, we check different sets of trade-offs for consistency and measure the resulting resource and time consumption.

A. Experimental Setup

We performed our experiments on a notebook computer featuring an Intel Pentium M 2.0 GHz CPU with 1.5 GB RAM and Sun Java 6.

For the various evaluations, randomly generated trade-off sets were used. Trade-off sets could be adjusted in a) their size (default 20 trade-offs per set if not mentioned otherwise), b) the number of total attributes (default 6), c) the number of attributes involved in a trade-off (default randomly chosen between 2 and 4 for each trade-off), and d) the size of the attribute domains (default 20).

B. Sequences of Fixed Length

In this section, random trade-off sets of the size 20 are examined. The procedure for this evaluation is as follows: 25,000 sets of 20 trade-offs each are randomly generated. For each set, the trade-offs are incrementally checked for consistency as described in chapter IV, one after the other. During these checks, the search tree of all possible trade-off chains (i.e. all possibilities to apply the current trade-offs to generic database objects and objects generated by other trade-offs) is constructed. The tree is then tested for the occurrence of adjacent duplicate trade-off patterns indicating inconsistency of the trade-off set (see observation 2 in chapter III). The algorithm

Table 1. Measured Statistics for 20 trade-offs

Measure	Time	Tree Size	Tree Depth	Inconsistent at
Median	16 ms	4,288	9	14
Quartil 1	<1 ms	975	7	10
Quartil 3	94 ms	17,556	12	17
Min	<1 ms	2	1	2
Max	36.9 sec	3,949,421	27	20
Mean	375 ms	48,049	9.46	13.23
Std.Dev.	2,040 ms	235,889	3.51	4.65

terminates when all trade-offs have been checked or an inconsistency is found.

During that course, four key measurements were taken: The time needed to check a full set of 20 trade-offs for consistency, the size of the search tree constructed to perform that check, and the depth of that tree. Finally, it was measured after how many trade-offs the sequence became inconsistent (if at all). The summarized values of the 15,000 trade-off sets are compiled in Table 1.

A sequence of 20 different trade-offs already contains a large amount of user provided information. During our evaluation, 41% of all randomly generated trade-off sequences had been inconsistent; most of them broke after stating more than 14 trade-offs – shorter trade-off chains were most of the time consistent.

The size of the search tree and the time needed to check for inconsistencies is highly correlated (Pearson coefficient >0.99). Both are influenced by the specific nature of trade-offs within a set. Some trade-offs are highly active and can often be applied after others. Many such active trade-offs in a set will result in the construction of many and long trade-off sequences and thus amount to creating a larger tree. While it is hard to grasp which properties of trade-offs allow for the construction of a larger number of trade-off sequences, it fortunately seems to be quite a rare event: 75% of all evaluated trade-off sets resulted to search trees smaller than 17,500 nodes.

Trees of that size can be processed in less than 100 milliseconds even on our relatively small test machine. However, the mean values are at 48,049 nodes and 385 ms, while the median is just 16 ms. When further examining the results, it is revealed that most trade-off sets can be checked efficiently (95% are verified in less a second) while only very few of them need considerably more effort (1% take longer than 8 seconds, with a maximum of 36 seconds). We can conclude the following: In most cases, the check for trade-off consistency can be performed fast and efficiently. However, this cannot be guaranteed – under rare, unfortunate circumstances, the computation may take considerably more time.

C. Varying the Number of Trade-Offs

The tree size (and thus the runtime) increases with every additional trade-off that needs to be checked. From a complexity theoretical point of view, our algorithm performs exponen-




Figure 2: Tree size plotted by number of trade-offs

tially within the number of trade-offs. However, we will show in the following show that this usually does not pose a problem for the practical application.

In this evaluation, we examine the impact of the size of trade-off set on the number of nodes of the search tree. We reuse the trade-offs generated during the last evaluation. Now, incrementally up to 20 times, a single trade-off is added to the present set of trade-offs to be checked for consistency. After each added trade-off, the tree-size required by the algorithm is measured.

The median and maximum tree sizes of 15,000 generated trade-off sets are plotted in Figure 3 (note the usage of a logarithmic scale on the y-axis). The exponential nature of the algorithm is clearly visible. However, note that the median size most of the times stays 2-4 orders of magnitude below the worst case measured. Also, it should be mentioned that we expect that users in a trade-off based system rarely state more than 5-10 trade-offs – our assumption of 20 trade-offs clearly beyond the expected complexity of the problem.

D. Varying the Number of Attributes

For this evaluation, the number of considered attributes per database object has been varied. The number of attributes was adjusted starting from three up to nine attributes. For each number of attributes, we again generated 15,000 sets of 20 trade-offs each. The generated trade-offs randomly span at least two up to 3/4 of the available attributes (i.e., for 8 total attributes, we had trade-offs spanning 2 to 6 attributes).

The first observation is the decrease of inconsistent trade-off sets with increasing number of attributes: While regarding three attributes per possible database object, 82% off all generated trade-off sets have shown to be inconsistent. This number decreased to 51% when considering nine attributes. Figure 3 illustrates this relationship. The reason is that the average amount of overlap of each two trade-offs in the set decreases and there are fewer options to create preference cycles.

Other significant relationships could not really be established. For any number of attributes, computation time still stayed below 1 second for around 95% of all trade-off sets. For three and four attributes, the median tree size was measured at 653 and 1358 nodes, respectively, which are lower than those for




Figure 3: Percentage of inconsistent trade-offs for increasing number of attributes

more attributes. However, this trend could not be further observed for increasing numbers of attributes. This suggests that the actual runtime of the algorithm may vary significantly and is hard to estimate.

E. Effectiveness of Optimization

In chapter IV.B, we introduced an optimization of the basic algorithm which can save template comparisons. This was facilitated by pre-computing whether two given trade-offs can be applied directly after each other in the most general case. By evaluating 15,000 sets of 20-trade-offs each, it turned out that this heuristic can save 59% of all template comparisons in average (min: 36%, max: 84%, standard deviation: 6%).

VI. SUMMARY AND OUTLOOK

Trade-Offs have shown to be a natural and powerful addition to skyline queries. By extending the Pareto semantics, they can successfully remedy the fast growing result set sizes, one of the paradigm’s major weaknesses. However, before trade-offs can be reliably integrated into the skyline computation, their consistency has to be ensured. In this paper, we presented a novel algorithm that is able to solve the essential problem of verifying arbitrary sets of multi-dimensional trade-offs. This ability helps to meet the long-term research challenge for cooperative and personalized information systems.

In contrast to former approaches, our algorithm is able to reliably detect any inconsistency in a set of trade-offs without posing restrictions on their expressiveness in a time-efficient manner. This was archived by abstracting from materializing prohibitively complex product order and without the need to access actual database instances. Instead, we focus on discovering specific patterns within possible trade-off sequences which indicate a cyclic preference in the product order. We haven proven that certain patterns containing adjacent duplicate subsequences of trade-offs will always lead to such cyclic preferences and thus the trade-off set is inconsistent.

In our evaluations, this approach has also shown a superior average-case performance. In fact, the consistency check for 3/4 of all evaluated trade-offs could be performed below 100 ms, 95% of all cases still stayed below 1 second.

In our future work, we intend to focus on the problem of actually incorporating trade-offs into a user's skyline after the consistency check has been performed. Also, intuitive methods for eliciting suitable trade-offs from the user and guiding them over the entire process of applying compromises to the selection problem is a promising research area. Additionally, heuristics for further improving the performance of the algorithm presented in this work might be rewarding. A closer inspection of the rare cases of longer runtimes will also help to better understand the nature of trade-offs in human decision-making.

REFERENCES

- [1] W.-T. Balke, U. Güntzer. Multi-objective Query Processing for Database Systems. *Int. Conf. on Very Large Data Bases (VLDB)*, Toronto, Canada, 2004.
- [2] W.-T. Balke, U. Güntzer, C. Lofi. Eliciting Matters – Controlling Skyline Sizes by Incremental Integration of User Preferences. *Int. Conf. on Database Systems for Advanced Applications (DASFAA)*, Bangkok, Thailand, 2007.
- [3] W.-T. Balke, U. Güntzer, C. Lofi. User Interaction Support for Incremental Refinement of Preference-Based Queries. *1st International Conference on Research Challenges in Information Science (RCIS)*, Ouarzazate, Morocco, 2007.
- [4] W.-T. Balke, U. Güntzer. Consistently Adding Amalgamations to Preference Orders. *3rd Multidisciplinary Workshop on Advances in Preference Handling (M-Pref 2007)*, Vienna, Austria, 2007.
- [5] W.-T. Balke, U. Güntzer, and C. Lofi. Incremental Trade-Off Management for Preference Based Queries. In *International Journal of Computer Science & Applications (IJCSA)*, Vol. 4(2), 2007.
- [6] W.-T. Balke, U. Güntzer, W. Siberski. Restricting Skyline Sizes using Weak Pareto Dominance. In *Informatik - Forschung und Entwicklung (IFE)*, Vol. 21(3), Springer, 2007.
- [7] W.-T. Balke, J. Zheng, U. Güntzer. Approaching the Efficient Frontier: Cooperative Database Retrieval Using High-Dimensional Skylines. In *Proc. of the Int. Conf. on Database Systems for Advanced Applications (DASFAA)*, Beijing, China, 2005.
- [8] S. Börzsönyi, K. Stocker, and D. Kossmann. The Skyline Operator. In *Proc. of the Int. IEEE Conf. on Data Engineering (ICDE)*, Heidelberg, Germany, 2001.
- [9] W. J. Bradley, J. K. Hodge, and D. M. Kilgour. Separable discrete preferences. *Mathematical Social Sciences*, 49(3):335-353, 2005.
- [10] J. Chomicki. Iterative Modification and Incremental Evaluation of Preference Queries. In *Proc. of the Int. Symp. on Found. of Inf. and Knowledge Systems (FoIKS)*, Budapest, Hungary, 2006.
- [11] P. Godfrey. Skyline cardinality for relational processing. How many vectors are maximal? In *Proc. of the Int. Symp. on Foundations of Information and Knowledge Systems (FoIKS 2004)*. Vienna, Austria, 2004.
- [12] P. Godfrey, R. Shipley, J. Gryz. Maximal Vector Computation in Large Data. In *Proc. of the Int. Conf. on Very Large Databases (VLDB)*, Trondheim, Norway, 2005.
- [13] V. Koltun, C. Papadimitriou. Approximately Dominating Representatives. In *Proc. of the Int. Conf. on Database Theory (ICDT)*, Edinburgh, UK, 2005.
- [14] D. Kossmann, F. Ramsak, S. Rost. Shooting Stars in the Sky: An Online Algorithm for Skyline Queries. In *Proc. of the Int. Conf. on Very Large Data Bases (VLDB)*, Hong Kong, China, 2002.
- [15] J. Lee, G. You, S. Hwang. Telescope: Zooming to Interesting Skylines. In *Proc. of the Int. Conf. on Database Systems for Advanced Applications (DASFAA)*, Bangkok, Thailand, 2007.
- [16] X. Lin, Y. Yuan, Q. Zhang, Y. Zhang. Selecting Stars: The k Most Representative Skyline Operator. In *Proc. of the Int. IEEE Conf. on Data Engineering (ICDE)*, Istanbul, Turkey, 2007.
- [17] J. Pei, Y. Yuan, X. Lin, W. Jin, M. Ester, Q. Liu, W. Wang, Y. Tao, J. X. Yu, Q. Zhang. Towards multidimensional subspace skyline analysis. In *ACM Trans. on Database Systems (TODS)*. Vol. 31(4), 2006.
- [18] P. Viappiani, B. Faltings, P. Pu. Preference-based Search using Example-Critiquing with Suggestions. In *Journal of Artificial Intelligence Research (JAIR)*, Vol. 27, 2006.