

Efficient Adaptive Detection of Complex Event Patterns

Ilya Kolchinsky
Technion, Israel Institute of Technology
Haifa 32000 Israel
ikolchin@cs.technion.ac.il

Assaf Schuster
Technion, Israel Institute of Technology
Haifa 32000 Israel
assaf@cs.technion.ac.il

ABSTRACT

Complex event processing (CEP) is widely employed to detect occurrences of predefined combinations (patterns) of events in massive data streams. As new events are accepted, they are matched using some type of evaluation structure, commonly optimized according to the statistical properties of the data items in the input stream. However, in many real-life scenarios the data characteristics are never known in advance or are subject to frequent on-the-fly changes. To modify the evaluation structure as a reaction to such changes, adaptation mechanisms are employed. These mechanisms typically function by monitoring a set of properties and applying a new evaluation plan when significant deviation from the initial values is observed. This strategy often leads to missing important input changes or it may incur substantial computational overhead by over-adapting.

In this paper, we present an efficient and precise method for dynamically deciding whether and how the evaluation structure should be reoptimized. This method is based on a small set of constraints to be satisfied by the monitored values, defined such that a better evaluation plan is guaranteed if any of the constraints is violated. To the best of our knowledge, our proposed mechanism is the first to provably avoid false positives on reoptimization decisions. We formally prove this claim and demonstrate how our method can be applied on known algorithms for evaluation plan generation. Our extensive experimental evaluation on real-world datasets confirms the superiority of our strategy over existing methods in terms of performance and accuracy.

PVLDB Reference Format:

Ilya Kolchinsky and Assaf Schuster. Efficient Adaptive Detection of Complex Event Patterns. *PVLDB*, 11 (11): 1346-1359, 2018. DOI: <https://doi.org/10.14778/3236187.3236190>

Categories and Subject Descriptors

H.2.4 [Database Management]: Systems

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were invited to present their results at The 44th International Conference on Very Large Data Bases, August 2018, Rio de Janeiro, Brazil.

Proceedings of the VLDB Endowment, Vol. 11, No. 11
Copyright 2018 VLDB Endowment 2150-8097/18/07... \$ 10.00.
DOI: <https://doi.org/10.14778/3236187.3236190>

General Terms

Algorithms, Design, Performance

Keywords

Stream Processing, Complex Event Processing, Adaptivity

1. INTRODUCTION

Real-time detection of complex data patterns is one of the fundamental tasks in stream processing. Many modern applications present a requirement for tracking data items arriving from multiple input streams and extracting occurrences of their predefined combinations. Complex event processing (CEP) is a prominent technology for providing this functionality, broadly employed in a wide range of domains, including sensor networks, security monitoring and financial services. CEP engines represent data items as *events* arriving from event sources. As new events are accepted, they are combined into higher-level *complex events* matching the specified patterns, which are then reported to end users.

One of the core elements of a CEP system is the *evaluation mechanism*. Popular evaluation mechanisms include non-deterministic finite automata (NFAs) [50], evaluation trees [43], graphs [8] and event processing networks (EPNs) [30]. A CEP engine uses an evaluation mechanism to create an internal representation for each pattern P to be monitored. This representation is constructed according to the *evaluation plan*, which reflects the structure of P . The evaluation plan defines how primitive events are combined into partial matches. Typically, a separate instance of the internal representation is created at runtime for every potential pattern match (i.e., a combination of events forming a valid subset of a full match).

As an example, consider the following scenario.

Example 1. A system for managing an array of smart security cameras is installed in a building. All cameras are equipped with face recognition software, and periodical readings from each camera are sent in real time to the main server. We are interested in identifying a scenario in which an intruder accesses the restricted area via the main gate of the building rather than from the dedicated entrance. This pattern can be represented as a sequence of three primitive events: 1) camera A (installed near the main gate) detects a person; 2) later, camera B (located inside the building's lobby) detects the same person; 3) finally, camera C detects the same person in the restricted area.

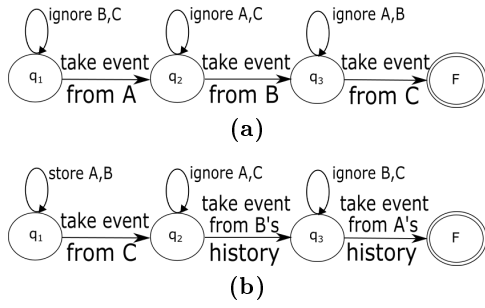


Figure 1: Evaluation structures for a sequence of events from streams A, B, C : (a) NFA without reordering; (b) Lazy NFA with reordering.

Figure 1(a) demonstrates an example of an evaluation mechanism (a non-deterministic finite automaton) for detecting this simple pattern by a CEP engine. This NFA is created according to the following simple evaluation plan. First, a stream of events arriving from camera A is inspected. For each accepted event, the stream of B is probed for subsequently received events specifying the same person. If found, we wait for a corresponding event to arrive from camera C.

Pattern detection performance can often be dramatically improved if the statistical characteristics of the monitored data are taken into account. In the example above, it can be assumed that fewer people access the restricted area than pass through the main building entrance. Consequently, the expected number of face recognition notifications arriving from camera C is significantly smaller than the expected number of similar events from cameras A and B. Thus, instead of detecting the pattern in the order of the requested occurrence of the primitive events (i.e., $A \rightarrow B \rightarrow C$), it would be beneficial to employ the “lazy evaluation” principle [37] and process the events in a different order, first monitoring the stream of events from C, and then examining the local history for previous readings of B and A. This way, fewer partial matches would be created. Figure 1(b) depicts the NFA constructed according to the improved plan.

Numerous authors proposed methods for defining evaluation plans based on the statistical properties of the data, such as event arrival rates [8, 37, 43, 46]. It was shown that systems tuned according to the a priori knowledge of these statistics can boost performance by up to several orders of magnitude, especially for highly skewed data.

Unfortunately, in real-life scenarios this a priori knowledge is rarely obtained in advance. Moreover, the data characteristics can change rapidly over time, which may render an initial evaluation plan extremely inefficient. In Example 1, the number of people near the main entrance might drop dramatically in late evening hours, making the event stream from camera A the first in the plan, as opposed to the event stream from C.

To overcome this problem, a CEP engine must continuously estimate the current values of the target parameters and, if and whenever necessary, adapt itself to the changed data characteristics. We will denote systems possessing such capabilities as *Adaptive CEP (ACEP)* systems.

A common structure of an ACEP system is depicted in Figure 2. The evaluation mechanism starts processing incoming events using some initial plan. A dedicated compo-

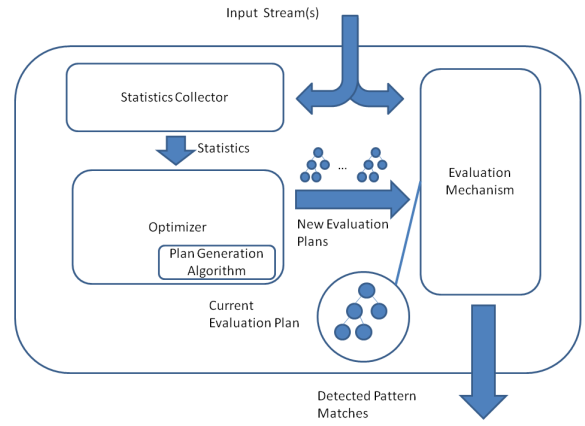


Figure 2: General structure of an adaptive CEP system.

nent calculates up-to-date estimates of the statistics (e.g., event arrival rates in Example 1) and transfers them to the optimizer. The optimizer then uses these values to decide whether the evaluation plan should be updated. If the answer is positive, a plan generation algorithm is invoked to produce a new plan (e.g., a new NFA), which is then delivered to the evaluation mechanism to replace the previously employed structure. In Example 1, this algorithm simply sorts the event types in the ascending order of their arrival rates and returns a chain-structured NFA conforming to that order.

Correct decisions by the optimizer are crucial for the successful operation of an adaptation mechanism. As the process of creating and deploying a new evaluation plan is very expensive, we would like to avoid “*false positives*,” that is, launching reoptimizations that do not improve the currently employed plan. “*False negatives*,” occurring when an important shift in estimated data properties is missed, are equally undesirable. A flawed decision policy may severely diminish or even completely eliminate the gain achieved by an adaptation mechanism.

The problem of designing efficient and reliable algorithms for reoptimization decision making has been well studied in areas such as traditional query optimization [29]. However, it has received only limited attention in the CEP domain ([37, 43]). In [37], the authors present a structure which reorganizes itself according to the currently observed arrival rates of the primitive events. Similarly to Eddies [13], this system does not adopt a single plan to maintain, but rather generates a new plan for each newly observed set of events regardless of the performance of the current one. The main strength of this method is that it is guaranteed to produce the optimal evaluation plan for any given set of events. However, it can create substantial bottlenecks due to the computational overhead of the plan generation algorithm. This is especially evident for stable event streams with little to no data variance, for which this technique would be outperformed by a non-adaptive solution using a static plan.

The second approach, introduced in [43], defines a constant threshold t for all monitored statistics. When any statistic deviates from its initially observed value by more than t , plan reconstruction is activated. This solution is much cheaper computationally than the previous one. However, some reoptimization opportunities may be missed.

Consider Example 1 again. Recall that we are interested in detecting the events by the ascending order of their arrival rates, and let the rates for events generated by cameras A, B and C be $rate_A = 100$, $rate_B = 15$, $rate_C = 10$ respectively. Obviously, events originating at A are significantly less sensitive to changes than those originating at B and C. Thus, if we monitor the statistics with a threshold $t > 6$, a growth in C to the point where it exceeds B will not be discovered, even though the reoptimization is vital in this case. Alternatively, setting a value $t < 6$ will result in detection of the above change, but will also cause the system to react to fluctuations in the arrival rate of A, leading to redundant plan recomputations.

No single threshold in the presented scenario can ensure optimal operation. However, by removing the conditions involving t and monitoring instead a pair of constraints $\{rate_A > rate_B, rate_B > rate_C\}$, plan recomputation would be guaranteed if and only if a better plan becomes available.

This paper presents a novel method for making efficient and precise on-the-fly adaptation decisions. Our method is based on defining a tightly bounded set of conditions on the monitored statistics to be periodically verified at runtime. These conditions, which we call *invariants*, are generated during the initial plan creation, and are constantly recomputed as the system adapts to changes in the input. The invariants are constructed to ensure that a violation of at least one of them guarantees that a better evaluation plan is available.

To the best of our knowledge, our proposed mechanism is the first to provably avoid false positives on reoptimization decisions. It also achieves notably low numbers of false negatives as compared to existing alternatives, as shown by our empirical study. This method can be applied to any deterministic algorithm for evaluation plan generation and used in any stream processing scenario.

The contributions and the structure of this paper can thus be summarized as follows:

- We formally define the reoptimizing decision problem for the complex event processing domain (Section 2).
- We present a novel method for detecting reoptimization opportunities in ACEP systems by verifying a set of invariants on the monitored data characteristics and formally prove that no false positives are possible when this method is used. We also extend the basic method to achieve a balance between computational efficiency and precision (Section 3).
- We demonstrate how to apply the invariant-based method on two known algorithms for evaluation structure creation, the greedy order-based algorithm [35] and ZStream algorithm [43], and discuss the generalization of these approaches to broader categories of algorithms (Section 4).
- We conduct an extensive experimental evaluation, comparing the invariant-based method to existing state-of-the-art solutions. The results of the experiments, performed on two real-world datasets, show that our proposed method achieves the highest accuracy and the lowest computational overhead (Section 5).

2. PRELIMINARIES

This section presents the notations used throughout this paper, outlines the event detection process in an ACEP system, and provides a formal definition of the reoptimizing decision problem, which will be further discussed in the subsequent sections.

2.1 Notations and Terminology

A pattern recognized by a CEP system is defined by a combination of primitive events, operators, predicates, and a time window. The patterns are formed using declarative specification languages ([24, 28, 50]).

Each event is represented by a type and a set of attributes, including the occurrence timestamp. Throughout this paper we assume that each primitive event has a well-defined type, i.e., the event either contains the type as an attribute or it can be easily inferred from the event attributes using negligible system resources. We will denote the pattern size (i.e., the number of distinct primitive events in a pattern) by n .

The predicates to be satisfied by the participating events are usually organized in a Boolean formula. Any condition can be specified on any attribute of an event, including the timestamp (e.g., for supporting multiple time windows).

The operators describe the relations between the events comprising a pattern match. Among the most commonly used operators are sequence (SEQ), conjunction (AND), disjunction (OR), negation (typically marked by '~'), requires the absence of an event from the stream) and Kleene closure (marked by '*'), accepts multiple appearances of an event in a specified position). A pattern may include an arbitrary number of operators.

To illustrate the above, consider Example 1 again. We will define three event types according to the identifiers of the cameras generating them: A, B and C. For each primitive event, we will set the attribute *person_id* to contain a unique number identifying a recognized face. Then, to detect a sequence of occurrences of the same person in three areas in a 10-minute time period, we will use the following pattern specification syntax, taken from SASE [50]:

```
PATTERN SEQ (A a, B b, C c)
WHERE ((a.person_id = b.person_id) ^
       (b.person_id = c.person_id))
WITHIN 10 minutes.
```

On system initialization, the pattern declaration is passed to the *plan generation algorithm* \mathcal{A} to create the *evaluation plan*. The evaluation plan provides a scheme for the CEP engine, according to which its internal pattern representation is created. The plan generation algorithm accepts a pattern specification P and a set of statistical data characteristic values *Stat*. It then returns the evaluation plan to be used for detection. If these values are not known in advance, a default, empty *Stat*, is passed. Multiple plan generation algorithms have been devised, efficiently supporting patterns with arbitrarily complex combinations of the aforementioned operators [35, 36, 43].

In Example 1, *Stat* contains the arrival rates of event types A, B and C, the evaluation plan is an ordering on the above types, and \mathcal{A} is a simple sorting algorithm, returning a plan following the ascending order of the arrival rates. The CEP engine then adheres to this order during pattern detection. Another popular choice for a statistic to be monitored is the set of selectivities (i.e., the probabilities of success) of the inter-event conditions defined by the pattern. Examples of plan generation algorithms requiring the knowledge of condition selectivities are presented in Section 4.

The plan generation algorithm attempts to utilize the information in *Stat* to find the best possible evaluation plan subject to some predefined set of performance metrics, which




Figure 3: Evaluation trees produced by ZStream for the sequence of events from streams A, B, C : (a) a left-deep tree; (b) a right-deep tree.

we denote as $Perf$. These metrics may include throughput, detection latency, network communication cost, power consumption, and more. For instance, one possible value for $Perf$ in Example 1 is $\{throughput, memory\}$, as processing the events according to the ascending order of their arrival rates was shown to vastly improve memory consumption and throughput of a CEP system [37].

In the general case, we consider \mathcal{A} to be a computationally expensive operation. We also assume that this algorithm is optimal; that is, it always produces the best possible solution for the given parameters. While this assumption rarely holds in practice, the employed techniques usually tend to produce empirically good solutions.

An evaluation plan is not constrained to be merely an order. Figure 3 demonstrates two possible tree-structured plans as defined by ZStream [43]. An evaluation structure following such a plan accumulates the arriving events at their corresponding leaves, and the topology of the internal nodes defines the order in which they are matched and their mutual predicates are evaluated. Matches reaching the tree root are reported to the end users. From this point on, we will denote such plans as *tree-based plans*, whereas plans similar to the one used for Example 1 will be called *order-based plans*. While the methods discussed in this paper are independent of the specific plan structure, we will use order-based and tree-based plans in our examples.

2.2 Detection-Adaptation Loop

During evaluation, an ACEP system constantly attempts to spot a change in the statistical properties of the data and to react accordingly. This process, referred to as the *detection-adaptation loop*, is depicted in Algorithm 1.

The system accepts events from the input stream and processes them using the current evaluation plan. At the same time, the values of the data statistics in $Stat$ are constantly reestimated by the dedicated component (Figure 2), often as a background task. While monitoring simple values such as the event arrival rates is trivial, more complex expressions (e.g., predicate selectivities) require advanced solutions. In this paper, we utilize existing state-of-the-art techniques from the field of data stream processing [14, 27]. These histogram-based methods allow to efficiently maintain a variety of stream statistics over sliding windows with high precision and require negligible system resources.

Opportunities for adaptation are recognized by the *reoptimizing decision function* \mathcal{D} , defined as follows:

$$\mathcal{D} : STAT \rightarrow \{true, false\},$$

Algorithm 1 Detection-adaptation loop in an ACEP system

Input: pattern specification P , plan generation algorithm \mathcal{A} , reoptimizing decision function \mathcal{D} , initial statistic values $in_stat \in STAT$

```

curr_plan ← A(P, in_stat)
while more events are available:
    process incoming events using curr_plan
    curr_stat ← estimate current statistic values
    if D(curr_stat):
        new_plan ← A(P, curr_stat)
        if new_plan is better than curr_plan:
            curr_plan ← new_plan
    apply curr_plan

```

where $STAT$ is a set of all possible collections of the measured statistic values. \mathcal{D} accepts the current estimates for the monitored statistic values and decides whether reoptimization is to be attempted. Whenever \mathcal{D} returns *true*, \mathcal{A} is invoked. The output of \mathcal{A} is a new evaluation plan, which, if found more efficient than the current plan subject to the metrics in $Perf$, is subsequently deployed.

Methods for replacing an evaluation plan on-the-fly without significantly affecting system performance or losing intermediate results are a major focus of current research [29]. Numerous advanced techniques were proposed in the field of continuous query processing in data streams [10, 38, 53]. In our work, we use the CEP-based strategy introduced in [37]. Let t_0 be the time of creation of the new plan. Then, partial matches containing at least a single event accepted before t_0 are processed according to the old plan p_{old} , whereas the newly created partial matches consisting entirely of “new” events are treated according to the new plan p_{new} . Note that since p_{old} and p_{new} operate on disjoint sets of matches, there is no duplicate processing during execution. At time $t_0 + W$ (where W is the time window of the pattern), the last “old” event expires and the system switches fully to p_{new} .

In general, we consider the deployment procedure to be a costly operation and will attempt to minimize the number of unnecessary plan replacements.

2.3 Reoptimizing Decision Problem

The reoptimizing decision problem is the problem of finding a function \mathcal{D} that maximizes the performance of a CEP system subject to $Perf$. It can be formally defined as follows: *given the pattern specification P , the plan generation algorithm \mathcal{A} , the set of monitored statistics $Stat$, and the set of performance metrics $Perf$, find a reoptimizing decision function \mathcal{D} that achieves the best performance of the ACEP detection-adaptation loop (Algorithm 1) subject to $Perf$.*

In practice, the quality of \mathcal{D} is determined by two factors. The first factor is the correctness of the answers returned by \mathcal{D} . Wrong decisions can either fall into the category of *false positives* (returning *true* when the currently used plan is still the best possible) or *false negatives* (returning *false* when a more efficient plan is available). Both cases cause the system to use a sub-optimal evaluation plan. The second factor is the time and space complexity of \mathcal{D} . As we will see in Section 5, an accurate yet resource-consuming implementation of \mathcal{D} may severely degrade system performance regardless of its output.

We can now analyze the solutions to the reoptimizing decision problem implemented by the adaptive frameworks which we discussed in Section 1. The tree-based NFA [37] defines a trivial decision function \mathcal{D} , unconditionally returning *true*. In ZStream [43] this functions loops over all values in the input parameter *curr_stat* and returns *true* if and only if a deviation of at least t is detected.

3. INVARIANT-BASED METHOD FOR THE REOPTIMIZING DECISION PROBLEM

As illustrated above, the main drawback of the previously proposed decision functions is their coarse granularity, as the same condition is verified for every monitored data property. We propose a different approach, based on constructing a set of fine-grained *invariants* that reflect the existing connections between individual data characteristics. The reoptimizing decision function \mathcal{D} will then be defined as a conjunction of these invariants.

In this section, we present the invariant-based decision method and discuss its correctness guarantees, time and space complexity, and possible optimizations.

3.1 Invariant Creation

A *decision invariant* (or simply invariant) will be defined as an inequality of the following form:

$$f_1(stat_1) < f_2(stat_2),$$

where $stat_1, stat_2 \in STAT$ are sets of the monitored statistic values and $f_1, f_2 : STAT \rightarrow \mathbb{R}$ are arbitrary functions.

We are interested in finding a single invariant for each *building block* of the evaluation plan in current use. A building block is defined as the most primitive, indivisible part of a plan. An evaluation plan can then be seen as a collection of building blocks. For instance, the plan for detecting a sequence of three events of types A, B and C, which we discussed in Example 1, is formed by the following blocks:

1. “Accept an event of type C”;
2. “Scan the history for events of type B matching the accepted C”;
3. “Scan the history for events of type A matching the accepted C and B”.

In general, in an order-based plan each step in the selected order will be considered a block, whereas for tree-based plans a block is equivalent to an internal node.

We know that the specific plan from the above example was chosen because the plan generation algorithm \mathcal{A} sorts the event types according to their arrival rates. If, for instance, the rate of B exceeded that of A, the second block would have been “Scan the history for events of type A matching the accepted C” and the third would also have changed accordingly. In other words, the second block of the plan is so defined because, during the run of \mathcal{A} , the condition $rate_B < rate_A$ was at some point checked, and the result of this check was positive. Following the terminology defined above, in this example $STAT$ consists of all valid arrival rate values and f_1, f_2 are trivial functions, i.e., $f_1(x) = f_2(x) = x$.

We will denote any condition (over the measured statistic values) whose verification has led the algorithm to include some building block in the final plan as a *deciding condition*.

Obviously, no generic method exists to distinguish between a deciding condition and a regular one. This process is to be applied separately on any particular algorithm \mathcal{A} based on its semantics. In our example, assume that the arrival rates are sorted using a simple min-sort algorithm, selecting the smallest remaining one at each iteration. Then, any comparison between two arrival rates will be considered a deciding condition, as opposed to any other condition which may or may not be a part of the implementation of this particular algorithm.

When \mathcal{A} is invoked on a given input, locations can be marked in the algorithm’s execution flow where the deciding conditions are verified. We will call any actual verification of a deciding condition a *block-building comparison (BBC)*. For instance, assume that we start executing our min-sort algorithm and a deciding condition $rate_C < rate_A$ is verified. Further assume that $rate_C$ is smaller than $rate_A$. Then, this verification is a BBC associated with the building block “Accept an event of type C first”, because, unless this deciding condition holds, the block will not be included in the final plan. This will also be the case if $rate_C < rate_B$ is subsequently verified and $rate_C$ is smaller. If $rate_B$ is smaller, the opposite condition, $rate_B < rate_C$, becomes a BBC associated with a block “Accept an event of type B first”. Overall, $(n - 1)$ BBCs take place during the first min-sort iteration, $(n - 2)$ during the second iteration, and so forth.

In general, for each building block b of any evaluation plan, we can determine a *deciding condition set (DCS)*. A DCS of b consists of all deciding conditions that were actually checked and satisfied by BBCs belonging to b as explained above. Note that, by definition, the intersection of two DCSs is always empty. In our example, assuming that the blocks listed above are denoted as b_1, b_2, b_3 , the deciding condition sets are as follows:

$$\begin{aligned} DCS_1 &= \{rate_C < rate_B, rate_C < rate_A\}, \\ DCS_2 &= \{rate_B < rate_A\}, \\ DCS_3 &= \emptyset. \end{aligned}$$

As long as the above conditions hold, no other evaluation plan can be returned by \mathcal{A} . On the other hand, if any of the conditions is violated, the outcome of \mathcal{A} will result in generating a different plan. If we define the decision function \mathcal{D} as a conjunction of the deciding condition sets, we will recognize situations in which the current plan becomes sub-optimal with high precision and confidence.

However, verifying all deciding conditions for all building blocks is very inefficient. In our simple example, the total number of such conditions is quadratic in the number of event types participating in the pattern. For more complicated plans and generation algorithms, this dependency may grow to a high-degree polynomial or even become exponential. Since the adaptation decision is made during every iteration of Algorithm 1, the overhead may negatively affect the system throughput and the response time.

To overcome this problem, we will constrain the number of conditions to be verified by \mathcal{D} to one per building block. For each deciding condition set DCS_i , we will determine the tightest condition, that is, the one that was closest to being violated during plan generation. This tightest condition will be selected as an *invariant* of the building block b_i . In other words, we may alternatively define an invariant as a deciding condition selected for actual verification

by \mathcal{D} out of a DCS. More formally, given a set $DCS_i = \{c_1, \dots, c_m\}$ such that $c_k = (f_{k,1}(stat_{k,1}) < f_{k,2}(stat_{k,2}))$, we will select a condition that minimizes the expression $(f_{k,2}(stat_{k,2}) - f_{k,1}(stat_{k,1}))$ as an invariant of the building block b_i .

In the example above, the invariant for DCS_i is $rate_C < rate_B$, since we know that $rate_B < rate_A$, and therefore $rate_B - rate_C < rate_A - rate_C$. It is clear that $rate_B$ is a tighter bound for the value of $rate_C$ than $rate_A$.

To summarize, the process of invariant creation proceeds as follows. During the run of \mathcal{A} on the current set of statistics $Stat$, we closely monitor its execution. Whenever a block-building comparison is detected for some block b , we add the corresponding deciding condition to the DCS of b . After the completion of \mathcal{A} , the tightest condition of each DCS is extracted and added to the invariant list.

Figure 4 demonstrates the invariant creation process applied on the pattern from Example 1 and the rate-sorting algorithm \mathcal{A} discussed above. Each subfigure depicts a different stage in the plan generation and presents the DCSs and the BBCs involved at this stage.

As discussed above, this generic method has to be adapted to any specific implementation of \mathcal{A} . This is trivially done for any \mathcal{A} which constructs the solution plan in a step-by-step manner, selecting and appending one building block at a time. However, for algorithms incorporating other approaches, such as dynamic programming, it is more challenging to attribute a block-building comparison to a single block of the plan. In Section 4, we will exemplify this process on two algorithms taken from the previous work in the field and discuss its applicability on broader algorithm categories.

3.2 Invariant Verification and Adaptation

During the execution of the detection-adaptation loop (Algorithm 1), \mathcal{D} traverses the list of invariants built as described above. It returns *true* if a violated invariant was found (according to the current statistic estimates) and *false* otherwise. This list is sorted according to the order of the respective building blocks in the evaluation plan. In Example 1, first the invariant $rate_C < rate_B$ will be verified, followed by $rate_B < rate_A$. The reason is that an invariant implicitly assumes the correctness of the preceding invariants (e.g., $rate_B < rate_A$ assumes that $rate_C < rate_B$ holds; otherwise, it should have been changed to $rate_C < rate_A$). For tree-based plans, the verification proceeds in a bottom-up order. For example, for the tree plan displayed in Figure 3(a), the order is $(A, B) \rightarrow (A, B, C)$.

If a violation of an invariant is detected, \mathcal{A} is invoked to create a new evaluation plan. In this case, the currently used invariants are invalidated and a new list is created following the process described above. Subsequent verifications performed by \mathcal{D} are then based on the new invariants.

Assuming that any invariant can be verified in constant time and memory, the complexity of \mathcal{D} using the invariant-based method is $O(B)$, where B is the number of the building blocks in an evaluation plan. This number is bounded by the pattern size (the number of event types participating in a pattern) for both order-based and tree-based plans. To guarantee this result, an application of the invariant-based method on a specific implementation of \mathcal{A} has to ensure that the verification of a single invariant is a constant-time operation, as we exemplify in Section 4.

3.3 Correctness Guarantees and the K-invariant Method

We will now formally prove that the invariant-based method presented above guarantees that no false positive detections will occur during the detection-adaptation loop.

THEOREM 1. *Let \mathcal{D} be a reoptimizing decision function implemented according to the invariant-based method. Let \mathcal{A} be a deterministic plan generation algorithm in use and let p be the currently employed plan. Then, if at some point during execution \mathcal{D} returns true, the subsequent invocation of \mathcal{A} will return a plan p' , such that $p' \neq p$.*

By definition, if \mathcal{D} returns *true*, then there is at least one invariant whose verification failed, i.e., its deciding condition does not hold anymore. Let c be the first such condition, and let b_i be the building block such that $c \in DCS_i$ (recall that there is only one such b_i). Then, by determinism of \mathcal{A} and by the ordering defined on the invariants, the new run of \mathcal{A} will be identical to the one that produced p until the block-building comparison that checks c . At that point, by definition of the block-building comparison, the negative result of validating c will cause \mathcal{A} to reject b_i as the current building block and select a different one, thus producing a plan p' , which is different from p . ■

Since we assume \mathcal{A} to always produce the optimal solution, the above result can be extended.

COROLLARY 1. *Let \mathcal{D} be an invariant-based reoptimizing decision function and let \mathcal{A} be a deterministic plan generation algorithm in use. Then, if at some point during execution \mathcal{D} returns true, the subsequent invocation of \mathcal{A} will return a plan that is more efficient than the currently employed one.*

Note that the opposite direction of Theorem 1 does not hold. It is still possible that a more efficient evaluation plan can be deployed, yet this opportunity will not be detected by \mathcal{D} because we only pick a single condition from each deciding condition set (see Section 4.2 for an example). If we were to include the whole union of the above sets in the invariant set, even stronger guarantees could be achieved, as stated in the following theorem.

THEOREM 2. *Let \mathcal{D} be a reoptimizing decision function implemented according to the invariant-based method, **with all conditions from all DCSs included in the invariant set**. Let \mathcal{A} be a deterministic plan generation algorithm in use and let p be the currently employed plan. Then, if **and only if** at some point during the execution \mathcal{D} returns true, the subsequent invocation of \mathcal{A} will return a plan p' , such that $p' \neq p$.*

The first direction follows immediately from Theorem 1. For the second direction, let $p' \neq p$ and let $b_i \in p, b'_i \in p'$ be the first building blocks that differ in p and p' . By \mathcal{A} 's determinism, there exist $f_1, f_2, stat_1, stat_2$ s. t.

$$\begin{aligned} (f_1(stat_1) < f_2(stat_2)) &\in DCS_i \\ (f_2(stat_2) < f_1(stat_1)) &\in DCS'_i, \end{aligned}$$

as otherwise there would be no way for \mathcal{A} to deterministically choose between b_i and b'_i . Since p' was created by \mathcal{A} using the currently estimated statistic values, we can deduce that $f_2(stat_2) < f_1(stat_1)$ holds. Consequently,




Figure 4: Invariant creation for pattern $SEQ(A, B, C)$ from Example 1: (a) selecting the first event type in the detection order; (b) C is set as the first event type, and selection of the second event type is in process; (c) B is set as the second type, and only a single event type remains for the third position; (d) the evaluation plan and the invariant set are finalized.

$f_1(stat_1) < f_2(stat_2)$ does not hold. By the assumption that all deciding conditions are included in the invariant set, \mathcal{D} will necessarily detect this violation, which completes the proof. ■

The above result shows that greater precision can be gained if we do not limit the number of monitored invariants per building block. However, as discussed above, validating all deciding conditions may drastically increase the adaptation overhead.

The tradeoff between performance and precision can be controlled by introducing a new parameter K , defined as the maximal number of conditions from a deciding set to select as invariants. We will refer to the method using a specific value of K as the K -invariant method, as opposed to the *basic invariant method* discussed above. Note that the 1-invariant method is equivalent to the basic one. The K -invariant method becomes more accurate and more time-consuming for higher values of K . The total number of the invariants in this case is at most $K \cdot (B - 1)$.

3.4 Distance-Based Invariants

By Corollary 1, it is guaranteed that a new, better evaluation plan will be produced following an invariant violation. However, the magnitude of its improvement over the old plan is not known. Consider a scenario in which two event types in a pattern have very close arrival rates. Further assume that there are slight oscillations in the rates, causing the event types to swap positions periodically when ordered according to this statistic. If an invariant is defined comparing the arrival rates of these two types, then \mathcal{D} will discover these minor changes and two evaluation plans with little to no difference in performance will be repeatedly produced and deployed. Although not a “false positive” by definition, the overhead implied by this situation may exceed any benefit of using an adaptive platform.

To overcome this problem, we will introduce the notion of the minimal distance d , defined as the smallest relative difference between the two sides of the inequality required for an invariant to be considered as violated. That is, given a deciding condition $f_{k,1}(stat_{k,1}) < f_{k,2}(stat_{k,2})$, we will construct the invariant to be verified by \mathcal{D} as follows:

$$(1 + d) \cdot f_{k,1}(stat_{k,1}) < f_{k,2}(stat_{k,2}).$$

The experimental study in Section 5 demonstrates that a correctly chosen d leads to a significant performance improvement over the basic technique. However, finding a sufficiently good d is a difficult task, as it depends on the data, the type of statistics, the invariant expression, and the frequency and magnitude of the runtime changes. We identify the following directions for solving this problem:

1)Parameter scanning: empirically checking a range of candidate values to find the one resulting in the best performance. This method is the simplest, but often infeasible in real-life scenarios.

2)Data analysis methods: deriving d from the currently available statistics can provide a good estimate in some cases. For instance, it can be calculated as the average relative difference between the sides of a deciding condition obtained during the initial plan generation, or, more formally:

$$d = AVG \left(\frac{|f_{k,2}(stat_{k,2}) - f_{k,1}(stat_{k,1})|}{\min(f_{k,1}(stat_{k,1}), f_{k,2}(stat_{k,2}))} \right).$$

The effectiveness of this approach depends on the distribution and the runtime behavior of the statistical values. Specifically, false positives may be produced when the values are very close and the changes are frequent. Still, we expect it to perform reasonably well in the common case. This technique can also be utilized to produce a starting point for parameter scanning.

3)Meta-adaptive methods: dynamically tuning d on-the-fly to adapt it to the current stream statistics. This might be the most accurate and reliable solution. We start with some initial value, possibly obtained using the above techniques. Then, as invariants are violated and new plans are computed, we modify d to prevent repeated reoptimization attempts when the observed gain in plan quality is low. An even higher precision can be achieved by additionally utilizing fine-grained per-invariant distances. This advanced research direction is a subject for our future work.

We implement and experimentally evaluate the first two approaches in Section 5.

4. APPLICATIONS OF THE INVARIANT-BASED METHOD

In Section 3, we presented a generic method for defining a reoptimizing decision function \mathcal{D} as a list of invariants. As we have seen, additional steps are required in order to apply this method to a specific choice of the evaluation plan structure and the plan generation algorithm. Namely, the following should be strictly defined: 1)what is considered a building block in a plan; 2)what is considered a block-building comparison in \mathcal{A} ; 3)how we associate a BBC with a building block. Additionally, efficient verification of the invariants must be ensured. In this section, we will exemplify this process on two plan-algorithm combinations taken from previous works in the field. The experimental study in Section 5 will also be conducted on these adapted algorithms. We also discuss how the presented techniques can be generalized to several classes of algorithms.

Algorithm 2 Greedy algorithm for order-based plans

Input: event types e_1, \dots, e_n , arrival rates r_1, \dots, r_n , inter-event predicate selectivities $sel_{1,1}, \dots, sel_{n,n}$
Output: order-based evaluation plan $E = e_{p_1}, e_{p_2}, \dots, e_{p_n}$

```
 $E \leftarrow \emptyset; p_1 = \operatorname{argmin}_j \{r_j \cdot sel_{j,j}\}$   
add  $e_{p_1}$  to  $E$   
for  $i$  from 2 to  $n$ :  
   $p_i = \operatorname{argmin}_{j \notin E} \{r_j \cdot sel_{j,j} \cdot \prod_{k < i} sel_{p_k,j}\}$   
  add  $e_{p_i}$  to  $E$   
return  $E$ 
```

4.1 Greedy Algorithm for Order-Based Plans

The greedy heuristic algorithm based on cardinalities and predicate selectivities was first described in [48] for creating left-deep tree plans for join queries. It was adapted to the CEP domain in [35]. The algorithm supports all operators described in Section 2.1 and their arbitrary composition. Its basic form, which we describe shortly, only targets conjunction and sequence patterns of arbitrary complexity. Support for other operators and their composition is implemented by either activating transformation rules on the input pattern or applying post-processing steps on the generated plan (e.g., to augment it with negated events). As these additional operations do not affect the application of the invariant-based method, we do not describe them here. The reader is referred to [35] for more details.

The algorithm proceeds iteratively, selecting at each step the event type which is expected to minimize the overall number of partial matches (subsets of valid pattern matches) to be kept in memory. At the beginning, the event type with the lowest arrival rate (multiplied by the selectivities of any predicates possibly defined solely on this event type) is chosen. At each subsequent step $i; i > 1$, the event type to be selected is the one that minimizes the expression $(\prod_{j=1}^i r_{p_j} \cdot \prod_{j,k \leq i} sel_{p_j,p_k})$, where r_x stands for the arrival rate of the x^{th} event type in a pattern, $sel_{x,y}$ is the selectivity of the predicate defined between the x^{th} and the y^{th} event types (equals to 1 if no predicate is defined), p_1, \dots, p_{i-1} are the event types selected during previous steps, and p_i is the candidate event type for the current step. Since a large part of this expression is constant when selecting p_i , it is sufficient to find an event type, out of those still not included in the plan, minimizing $(r_{p_i} \cdot sel_{p_i,p_i} \cdot \prod_{k < i} sel_{p_k,p_i})$.

Algorithm 2 depicts the plan generation process. When all selectivities satisfy $sel_{x,y} = 1$, i.e., no predicates are defined for the pattern, this algorithm simply sorts the events in an ascending order of their arrival rates.

We will define a building block for order-based evaluation plans produced by Algorithm 2 as a single directive of processing an event type in a specific position of a plan. That is, a building block is an expression of the form “Process the event type e_j at i^{th} position in a plan”. Obviously, a full plan output by the algorithm contains exactly n blocks, and a total of $O(n^2)$ blocks is considered during the run. Deciding conditions created for such a block are defined as:

$$r_j \cdot sel_{j,j} \cdot \prod_{k < i} sel_{p_k,j} < r_{j'} \cdot sel_{j',j'} \cdot \prod_{k < i} sel_{p_k,j'}$$

Here, $e_{j'}, j' \neq j$ is an event type which was considered to occupy i^{th} position at some point but eventually e_j was se-

Algorithm 3 ZStream algorithm for tree-based plans

Input: event types e_1, \dots, e_n , arrival rates r_1, \dots, r_n , inter-event predicate selectivities $sel_{1,1}, \dots, sel_{n,n}$
Output: tree-based evaluation plan T

```
 $subtrees \leftarrow$  new two-dimensional matrix of size  $n \times n$   
for  $i$  from 1 to  $n$ :  
   $subtrees[i][1].cardinality = subtrees[i][1].cost = r_i$   
for  $i$  from 2 to  $n$ :  
  for  $j$  from 1 to  $n - i + 1$ :  
    for  $k$  from  $j + 1$  to  $j + i$ :  
       $new\_cardinality = Card($   
         $subtrees[k - j][j].cardinality,$   
         $subtrees[i - (k - j)][k].cardinality)$   
       $new\_cost = subtrees[k - j][j].cost +$   
         $+ subtrees[i - (k - j)][k].cost + new\_cardinality$   
      if  $new\_cost < subtrees[i][j].cost$  :  
         $subtrees[i][j].tree = new\_tree($   
           $subtrees[k - j][j], subtrees[i - (k - j)][k])$   
         $subtrees[i][j].cardinality = new\_cardinality$   
         $subtrees[i][j].cost = new\_cost$   
return  $subtrees[n][1].tree$ 
```

lected. Note that, while in the worst case the products may contain up to $n - 1$ multiplicands, in most cases the number of the predicates defined over the events in a pattern is significantly lower than n^2 . Therefore, invariant verification will be executed in near-constant time.

4.2 Dynamic Programming Algorithm for Tree-Based Plans

The authors of ZStream [43] introduced an efficient algorithm for producing tree-based plans based on dynamic programming (Algorithm 3). The algorithm consists of $n - 1$ steps, where during the i^{th} step the tree-based plans for all subsets of the pattern of size $i + 1$ are calculated (for the trees of size 1, the only possible tree containing the lone leaf is assumed). During this calculation, previously memoized results for the two subtrees of each tree are used. To calculate the cost of a tree T with the subtrees L and R , the following formula is used:

$$Cost(T) = \begin{cases} r_i & T \text{ is a leaf} \\ Cost(L) + Cost(R) + Card(L, R) & \text{otherwise,} \end{cases}$$

where $Card(L, R)$ is the cardinality (the expected number of partial matches reaching the root) of T , whose calculation depends on the operator applied by the root. For example, the cardinality of a conjunction node is defined as the product of the cardinalities of its operands multiplied by the total selectivity of the conditions between the events in L and the events in R . That is,

$$Card(T) = Card(L) \times Card(R) \times SEL(L, R),$$

where $SEL(L, R)$ is a product of all predicate selectivities $sel_{i,j} : i \in L, j \in R$. Leaf cardinalities are defined as the arrival rates of the respective event types. The reader is referred to [43] for more details.

To apply the invariant-based method, we will define each internal node of a tree-based plan as a building block. This way, up to $O(n^3)$ blocks will be formed during the run of Algorithm 3, with only $O(n)$ included in the resulting plan.

A comparison between the costs of two trees will be considered a block-building comparison for the root of the less expensive tree. The deciding conditions for this algorithm will be thus defined simply as $Cost(T_1) < Cost(T_2)$, where T_1, T_2 are the two compared trees. These comparisons are invoked at each step during the search for the cheapest tree over a given subset of events. For k events, the number of candidate trees is $C_{k-1} = \frac{(2k-2)!}{(k-1)!k!}$, where C_m is the m^{th} Catalan number. Therefore, picking only one comparison as an invariant and dismissing the rest of the candidates may create a problem of false negatives, and K-invariant method is recommended instead (see discussion in Section 3.3).

The obvious problem with the above definition is that tree cost calculation is a recursive function, which contradicts our constant-time invariant verification assumption. We will eliminate this recursion by utilizing the following observation. In Algorithm 3, all block-building comparisons are performed on pairs of trees defined over the same set of event types. By invariant definition, one of these trees is always a subtree of a plan currently being in use. Recall that invariants on tree-based plans are always verified in the direction from leaves to the root. Hence, if any change was detected in one of the statistics affecting the subtrees of the two compared trees, it would be noticed during verification of earlier invariants. Thus, it is safe to represent the cost of a subtree in an invariant as a constant whose value is initialized to the cost of that subtree during invariant creation (i.e., plan construction).

4.3 General Applicability of the Invariant-Based Method

The approaches described in Sections 4.1 and 4.2 only cover two special cases. Here, we generalize the presented methodologies to apply the invariant-based method to any greedy or dynamic programming algorithm. We also discuss the applicability of our method to other algorithm categories.

A generalized variation of the technique illustrated in Section 4.1 can be utilized for any greedy plan generation algorithm. To that end, a part of a plan constructed during a single greedy iteration should be defined as a building block. Additionally, a conjunction of all conditions evaluated to select a specific block is to be defined as a block-building comparison associated with this block. Since most greedy algorithms require constant time and space for a single step, the complexity requirements for the invariant verification will be satisfied.

Using similar observations, we can generalize the approach described in Section 4.2 to any dynamic programming algorithm. A subplan memoized by the algorithm will correspond to a building block. A comparison between two subplans will serve as a BBC for the block that was selected during the initial run.

In general, the invariant-based method can be similarly adapted to any algorithm that constructs a plan in a deterministic, bottom-up manner, or otherwise includes a notion of a “building block”. To the best of our knowledge, the majority of the proposed solutions share this property.

In contrast, algorithms based on local search (adapted to CEP in [35]) cannot be used in conjunction with the invariant-based method. Rather than building a plan step-by-step, these algorithms start with a complete initial solution and modify it to create an improved version [3].

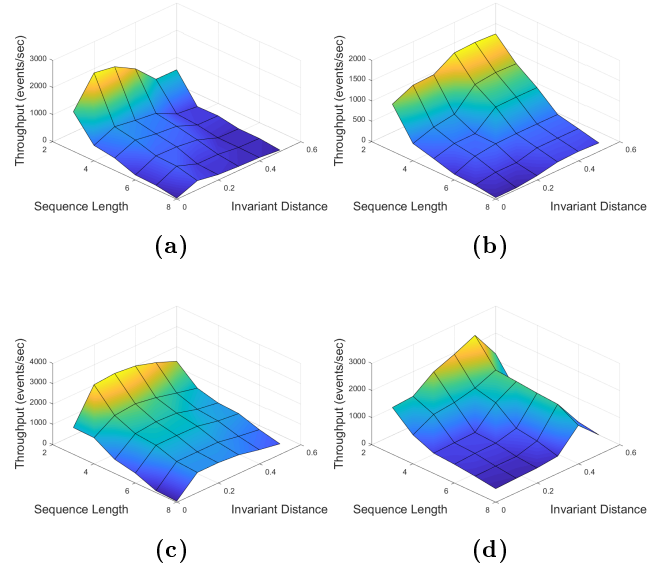


Figure 5: Throughput of the invariant-based method for different dataset-algorithm pairs as a function of the pattern size and the invariant distance d : (a) traffic dataset / greedy algorithm; (b) traffic dataset / ZStream algorithm; (c) stocks dataset / greedy algorithm; (d) stocks dataset / ZStream algorithm.

5. EXPERIMENTAL EVALUATION

In this section, the results of our experimental evaluation are presented. The objectives of this empirical study were twofold. First, we wanted to assess the overall system performance achieved by our approach and the computational overhead implied by its adaptation process as compared to the existing strategies for ACEP systems, outlined in Section 1. Our second goal was to explore how changes in the parameters of our method and of the data characteristics impact the above metrics.

5.1 Experimental Setup

We implemented the two CEP models described in Section 4, the lazy NFA [37] with the greedy order-based algorithm [48] and the ZStream model with tree-based dynamic programming algorithm [43]. We also added support for three adaptation methods (i.e., implementations of \mathcal{D}): 1) the unconditional reoptimization method from [37]; 2) the constant-threshold method from [43]; 3) the invariant-based method. To accurately estimate the event arrival rates and predicate selectivities on-the-fly, we utilized the algorithm from [27] for maintaining statistics over sliding window.

Since the plan generation algorithms used during this study create plans optimized for maximal throughput, we choose throughput as a main performance metric, reflecting the effectiveness of the above algorithms in the presence of changes in the input. We believe that similar results could be obtained for algorithms targeting any other optimization goal, such as minimizing latency or communication cost.

Two real-world datasets were used in the experiments. For each of them, we created 5 sets of patterns containing different operators (Section 2.1), as follows: (1) sequences; (2) se-

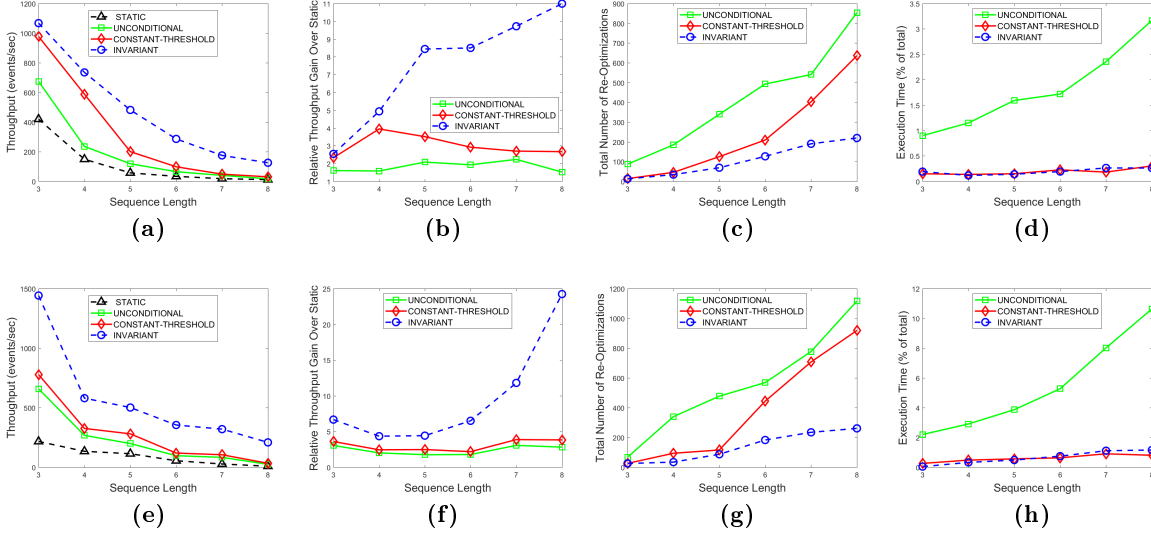


Figure 6: Comparison of the adaptation methods applied on the traffic dataset combined with the greedy algorithm ((a)-(d)) and ZStream algorithm ((e)-(h)): (a),(e) throughput (higher is better); (b),(f) relative throughput gain over the non-adaptive method (higher is better); (c),(g) total number of plan reoptimizations; (d),(h) computational overhead (lower is better).

quences with an additional event under negation; (3) conjunctions; (4) sequences with a single event under Kleene closure; (5) composite patterns, consisting of a disjunction of three sequences. Each set contained 6 patterns of sizes varying from 3 to 8. We defined pattern size as the number of events in a pattern for sets 1-4 and the number of events in each subpattern for set 5. Further details are specified below for each dataset. For lack of space, we only include the graphs averaged over all pattern sets unless otherwise stated. The specific results obtained for each set can be found in the extended version of this paper [34].

The first dataset contains the vehicle traffic sensor data, provided by City of Aarhus, Denmark [9] and collected over a period of 4 months from 449 observation points, with 13,577,132 primitive events overall. Each event represents an observation of traffic at the given point. The attributes of an event include, among others, the point ID, the average observed speed, and the total number of observed vehicles during the last 5 minutes. The arrival rates and selectivities for this dataset were highly skewed and stable, with few on-the-fly changes. However, the changes that did occur were mostly very extreme. The patterns for this dataset were motivated by normal driving behavior, where the average speed tends to decrease with the increase in the number of vehicles on the road. We requested to detect the violations of this model, i.e., combinations (sequences, conjunctions, etc., depending on the operator involved) of three or more observations with either an increase or a decline in both the number of vehicles and the average speed.

The second dataset was taken from the NASDAQ stock market historical records [1]. Each record in this dataset represents a single update to the price of a stock, spanning a 1-year period and covering over 2100 stock identifiers with prices updated on a per minute basis. Our input stream contained 80,509,033 primitive events, each consisting of a stock identifier, a timestamp, and a current price. For each stock identifier, a separate event type was defined. In ad-

dition, we preprocessed the data to include the difference between the current and the previous price. Contrary to the traffic dataset, low skew in data statistics was observed, with the initial values nearly identical for all event types. The changes were highly frequent, but mostly minor. The patterns to evaluate were then defined as combinations of different stock identifiers (types), with the predefined price differences (e.g., for a conjunction pattern $AND(A, B, C)$ we require $A.diff < B.diff < C.diff$).

All models and algorithms under examination were implemented in Java. All experiments were run on a machine with 2.20 Ghz CPU and 16.0 GB RAM.

5.2 Experimental Results

In our first experiment, we evaluated the performance of the invariant-based method for different values of the invariant distance d , obtained by parameter scanning (Section 3.4). In this experiment, only the sequence pattern sets were used. For each of the four possible dataset-algorithm combinations, the system throughput was measured as a function of the tested pattern size and of d , with its values ranging from 0 (which corresponds to the basic method) to 0.5.

The results are displayed in Figure 5. It can be observed that in each scenario there exists an optimal value d_{opt} , which depends on the data and the algorithm in use, consistently outperforming the other values for all pattern sizes. For distances higher than d_{opt} , too many changes in the statistics are undetected, while the lower values trigger unnecessary adaptations. Overall, the throughput achieved by using invariants with distance d_{opt} is 2 to 25 times higher than that of the basic method ($d = 0$).

Then, we validated the average relative difference method described in Section 3.4 by comparing its output value d_{avg} to d_{opt} (obtained via parameter scanning as described above) for each scenario. For lack of space, we only outline here the main observations of this study. Full results are available in the extended paper [34]. For the traffic dataset, the com-

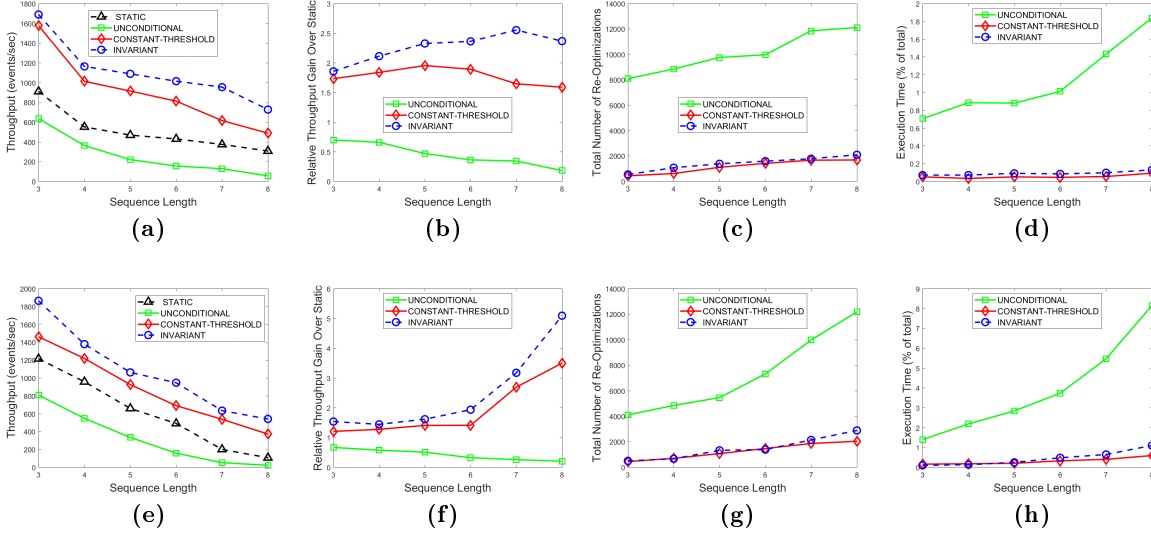


Figure 7: Comparison of the adaptation methods applied on the stocks dataset combined with the greedy algorithm ((a)-(d)) and ZStream algorithm ((e)-(h)): (a),(e) throughput (higher is better); (b),(f) relative throughput gain over the non-adaptive method (higher is better); (c),(g) total number of plan reoptimizations; (d),(h) computational overhead (lower is better).

puted values were considerably close to the optimal ones for patterns of length 6 and above, with precision reaching at least 87% (for ZStream algorithm and pattern length 7) and as high as 92% (Greedy algorithm, length 8). For the stocks dataset, the achieved accuracy was only 31-44%. This can be attributed to the low data skew, matching our earlier expectations and highlighting the need for developing better solutions, which is the goal of our future work.

Next, we performed an experimental comparison of all previously described adaptation methods. The comparison was executed separately for each dataset-algorithm combination. For the invariant-based method, the d_{opt} values obtained during the first experiment were used. For the constant-threshold method, an optimal threshold t_{opt} was empirically found for each of the above combinations using a similar series of runs.

Figures 6-7 summarize the results. Each figure consists of two rows. Every row represents the measurements for a particular dataset-algorithm combination and contains four graphs, presenting different statistics as a function of the pattern size. The first graph shows the throughput achieved using each of the adaptation methods. Here, we have also included the “static” method in our study, where no adaptation is supported and the dataset is processed using a single, predefined plan. The second graph is a different way of viewing the previous one, comparing the adaptation methods by the relative speedup they achieve over the “static plan” approach. The third graph depicts the total number of reoptimizations (actual plan replacements) recorded during each run. Finally, we report the computational overhead of each method, that is, a percentage of the total execution time spent on executions of \mathcal{D} and \mathcal{A} (i.e., checking whether a reoptimization is necessary and computing new plans).

The throughput comparison demonstrates the superiority of the invariant-based method over its alternatives for all scenarios. Its biggest performance gain is achieved in the traffic scenario, characterized by high skew and major

statistic shifts (Figure 6). This gain reaches its peak for larger patterns, with the maximal recorded performance of more than 6 times that of the second-best constant-threshold method: the greater the discrepancy between the data characteristics, the more difficult it is to find a single threshold to accurately monitor all the changes. Since this discrepancy may only increase as more statistic values are added to the monitored set, we expect the superiority of this method to keep growing with the pattern size beyond the values we experimented with. Figures 6(b) and 6(f) provide a clear illustration of the above phenomenon and of the invariant-based method scalability. Also, for larger patterns the constant-threshold method nearly converges to the unconditional one due to the increasing number of its false positives.

For the stocks dataset (Figure 7), the throughput measurements for the constant-threshold and the invariant-based methods are considerably closer. Due to the near-uniformity of the statistic values and of their variances, finding a single t_{opt} is sufficient to recognize most important changes. Hence, the precision of the constant-threshold method is very high on this input. Nevertheless, the invariant-based method achieves a performance speedup for this dataset as well (albeit only about 30-60%) without adding significant overhead. Also, for the same reason, the static plan performs reasonably well in this scenario, decidedly outperforming the unconditional method. The latter suffers from extreme over-adapting to the numerous small-scale statistic shifts.

The total number of reoptimizations performed in each scenario (Figures 6(c), 6(g), 7(c), 7(g)) backs up and augments the above results. The invariant-based method requires few plan replacements while also achieving the best throughput. The extremely high numbers produced by the unconditional strategy lead to its poor performance. For the traffic dataset, the constant-threshold method tends to approach these numbers for larger patterns. This can either be a sign of multiple false positives or over-adapting. For the stocks dataset, this method is similar to the invariant-based.

Figures 6(d), 6(h), 7(d) and 7(h) present the computational overhead of the compared approaches. Here, the same behavior is observed for all dataset-algorithm combinations. While the invariant-based and the constant-threshold methods consume negligible system resources, unconditional re-optimization results in up to 11% of the running time devoted to the adaptation process.

As evident by the experiments with stock market data, smaller number of reoptimizations and lower computational overhead do not necessarily result in better overall system performance. On this dataset, the invariant-based method achieves the highest throughput despite a slightly higher overhead as compared to the second-best constant-threshold method. This can be attributed to the false negatives of the latter, that is, cases in which it missed a reoptimization opportunity and kept using an old plan despite a better one being available.

In all experiments, the relative gain of the invariant-based method was considerably higher for ZStream algorithm than for the greedy one. There are two reasons for this result. First, the more complex structure of the tree-based plans makes it more difficult to capture the dependencies between plan components without fine-grained invariants. Second, as this algorithm is more computationally expensive, the penalty for a redundant reoptimization is higher. Following these observations, we believe that the invariant-based method is capable of achieving even larger benefit for more advanced and precise (and hence more complex) plan generation algorithms. Utilizing this method will thus encourage the adoption of such algorithms by CEP engines.

6. RELATED WORK

Complex event processing is an increasingly active research field [26]. The origins of CEP systems can be traced to older data stream managements systems (DSMSs), including Aurora/Borealis [4], Stream [12], TelegraphCQ [21], and NiagaraCQ [22]. This was followed by the emergence of a broad variety of solutions for detecting occurrences of situations of interest, as opposed to generic data, including frameworks such as SASE/SASE+ [50, 7, 52], CEDR [18], Cayuga [28], T-Rex [25] and Amit [6]. Esper [2] and IBM System S [11] are examples of widely used commercial CEP providers.

Many CEP approaches incorporate NFAs as their primary evaluation structure [50, 28, 25]. Various extensions to this model were developed, such as AFA [20] and lazy NFA [37]. ZStream [43] utilizes tree-based detection plans. Event processing networks [30] is another conceptual model, presenting a pattern as a network of simple agents.

Multiple works have addressed the broad range of CEP optimization opportunities arising when the statistical characteristics of the primitive events are taken into account. In [8] “plan-based evaluation” is described, where the arrival rates of events are exploited to reduce network communication costs. The authors of NextCEP [46] propose a framework for pattern rewriting in which operator properties are utilized to assign a cost to every candidate evaluation plan. Then, a search algorithm (either greedy or dynamic) is applied to select the lowest cost detection scheme. ZStream [43] applies a set of algebraic rule-based transformations on a given pattern, and then reorders the operators to minimize the cost of a plan.

Adaptive query processing (AQP) is the widely studied problem of adapting a query plan to the unstable data characteristics [29]. Multiple solutions consider traditional databases [5, 32, 15, 42, 33, 47]. The mid-query reoptimization mechanism [33], one of the first to possess adaptive properties, collects statistics at the predefined checkpoints and compares them to the past estimates. If severe deviation is observed, the remainder of the data is processed using a new plan. The methods described in [15] and [42] are the closest in spirit to our work. Rather than executing reoptimization on a periodic basis or upon a constant change, the authors compute an individual range for each monitored value within which the current plan is considered close-to-optimal.

The field of stream processing has developed adaptive techniques of its own. A-Greedy [16] is an algorithm for adaptive ordering of pipelined filters, providing strong theoretical guarantees. Similarly to our method, it detects violations of invariants defined on the filter drop probabilities. The authors of [40] describe “incremental reoptimization,” where the optimizer constantly attempts to locate a better plan using efficient search and pruning techniques. Eddy [13, 19, 41] presents stateless routing operators, redirecting incoming tuples to query operators according to a predefined routing policy. This system discovers execution routes on-the-fly in a per-tuple manner. Query Mesh [44] is a middle-ground approach, maintaining a set of plans and using a classifier to select a plan for each data item. Large DSMSs have also incorporated adaptive mechanisms [49, 17].

The majority of the proposed CEP techniques are derived from adaptivity considerations [31]. The two notable exceptions, ZStream [43] and tree-based NFA [37] were covered in detail above. Additional works labeled as ‘adaptive’ refer to on-the-fly switching between several detection algorithms [45, 51] or dynamic rule mining [23, 39].

7. CONCLUSIONS AND FUTURE WORK

In this paper, we discussed the problem of efficient adaptation of a CEP system to on-the-fly changes in the statistical properties of the data. A new method was presented to avoid redundant reoptimizations by periodically verifying a small set of simple conditions defined on the monitored data characteristics. We proved that validating this set of conditions will only fail if a better evaluation plan is available. We applied our method on two real-life algorithms and experimentally demonstrated the achieved performance gain.

One area of interest that was not yet addressed by the existing approaches is the multi-pattern adaptive CEP, where the system is given a set of patterns possibly containing common subexpressions. In this case, the detection process typically follows a single global plan that exploits sharing opportunities. While our method can be trivially applied to multi-pattern systems with no sharing, substantially more sophisticated optimization techniques are required for the general case. We intend to target this research direction in our future work.

8. ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Union’s Horizon 2020 Research And Innovation Programme under grant agreement no. 688380 and was partially supported by the Israel Science Foundation (grant no. 919191) and by HPI-Technion Research School.

9. REFERENCES

- [1] <http://www.eoddata.com>.
- [2] <http://www.espertech.com>.
- [3] E. Aarts and J. Lenstra, editors. *Local Search in Combinatorial Optimization*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1997.
- [4] D. J. Abadi, Y. Ahmad, M. Balazinska, M. Cherniack, J. Hwang, W. Lindner, A. S. Maskey, E. Rasin, E. Ryzkina, N. Tatbul, Y. Xing, and S. Zdonik. The design of the Borealis stream processing engine. In *CIDR*, pages 277–289, 2005.
- [5] M. Acosta, M. Vidal, T. Lampo, J. Castillo, and E. Ruckhaus. Anapsid: An adaptive query processing engine for sparql endpoints. In *International Semantic Web Conference (1)*, volume 7031, pages 18–34. Springer, 2011.
- [6] A. Adi and O. Etzion. Amit - the situation manager. *The VLDB Journal*, 13(2):177–203, 2004.
- [7] J. Agrawal, Y. Diao, D. Gyllstrom, and N. Immerman. Efficient pattern matching over event streams. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, pages 147–160, New York, NY, USA, 2008. ACM.
- [8] M. Akdere, U. Çetintemel, and N. Tatbul. Plan-based complex event detection across distributed sources. *PVLDB*, 1(1):66–77, 2008.
- [9] M. Ali, F. Gao, and A. Mileo. Citybench: A configurable benchmark to evaluate rsp engines using smart city datasets. In *Proceedings of ISWC 2015 - 14th International Semantic Web Conference*, pages 374–389, Bethlehem, PA, USA, 2015. W3C.
- [10] A. Aly, W. Aref, M. Ouzzani, and H. Mahmoud. JISC: adaptive stream processing using just-in-time state completion. In *Proceedings of the 17th International Conference on Extending Database Technology, Athens, Greece, March 24-28, 2014.*, pages 73–84.
- [11] L. Amini, H. Andrade, R. Bhagwan, F. Eskesen, R. King, P. Selo, Y. Park, and C. Venkatramani. Spc: A distributed, scalable platform for data mining. In *Proceedings of the 4th International Workshop on Data Mining Standards, Services and Platforms*, pages 27–37, New York, NY, USA, 2006. ACM.
- [12] A. Arasu, B. Babcock, S. Babu, J. Cieslewicz, M. Datar, K. Ito, R. Motwani, U. Srivastava, and J. Widom. *STREAM: The Stanford Data Stream Management System*, pages 317–336. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
- [13] R. Avnur and J. Hellerstein. Eddies: Continuously adaptive query processing. *SIGMOD Rec.*, 29(2):261–272, May 2000.
- [14] B. Babcock, M. Datar, R. Motwani, and L. O’Callaghan. Maintaining variance and k-medians over data stream windows. In *Proceedings of the Twenty-second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 234–243, New York, NY, USA, 2003. ACM.
- [15] S. Babu, P. Bizarro, and D. DeWitt. Proactive re-optimization. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, pages 107–118, New York, NY, USA. ACM.
- [16] S. Babu, R. Motwani, K. Munagala, I. Nishizawa, and J. Widom. Adaptive ordering of pipelined stream filters. In *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, pages 407–418, New York, NY, USA, 2004. ACM.
- [17] S. Babu and J. Widom. Streamon: An adaptive engine for stream query processing. In *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, pages 931–932, New York, NY, USA, 2004. ACM.
- [18] R. S. Barga, J. Goldstein, M. H. Ali, and M. Hong. Consistent streaming through time: A vision for event stream processing. In *CIDR*, pages 363–374, 2007.
- [19] P. Bizarro, S. Babu, D. J. DeWitt, and J. Widom. Content-based routing: Different plans for different data. In *Proceedings of the 31st International Conference on Very Large Data Bases, Trondheim, Norway, August 30 - September 2, 2005*, pages 757–768. ACM, 2005.
- [20] B. Chandramouli, J. Goldstein, and D. Maier. High-performance dynamic pattern matching over disordered streams. *PVLDB*, 3(1-2):220–231, 2010.
- [21] S. Chandrasekaran, O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. Madden, V. Raman, F. Reiss, and M. A. Shah. Telegraphicq: Continuous dataflow processing for an uncertain world. In *CIDR*, 2003.
- [22] J. Chen, D. J. DeWitt, F. Tian, and Y. Wang. Niagaraq: A scalable continuous query system for internet databases. *SIGMOD Rec.*, 29(2):379–390, 2000.
- [23] J. Coffi, C. Marsala, and N. Museux. Adaptive complex event processing for harmful situation detection. *Evolving Systems*, 3(3):167–177, Sep 2012.
- [24] G. Cugola and A. Margara. Tesla: a formally defined event specification language. In *DEBS*, pages 50–61. ACM, 2010.
- [25] G. Cugola and A. Margara. Complex event processing with T-REX. *J. Syst. Softw.*, 85(8):1709–1728, 2012.
- [26] G. Cugola and A. Margara. Processing flows of information: From data stream to complex event processing. *ACM Comput. Surv.*, 44(3):15:1–15:62, 2012.
- [27] M. Datar, A. Gionis, P. Indyk, and R. Motwani. Maintaining stream statistics over sliding windows. *SIAM J. Comput.*, 31(6):1794–1813, June 2002.
- [28] A. Demers, J. Gehrke, M. Hong, M. Riedewald, and W. White. Towards expressive publish/subscribe systems. In *Proceedings of the 10th International Conference on Advances in Database Technology*, pages 627–644. Springer-Verlag.
- [29] A. Deshpande, Z. Ives, and V. Raman. Adaptive query processing. *Foundations and Trends in Databases*, 1(1):1–140, January 2007.
- [30] O. Etzion and P. Niblett. *Event Processing in Action*. Manning Publications Co., 2010.
- [31] I. Flouris, N. Giatrakos, A. Deligiannakis, M. Garofalakis, M. Kamp, and M. Mock. Issues in complex event processing: Status and prospects in the big data era. *Journal of Systems and Software*, 127:217 – 236, 2017.
- [32] Z. Ives, A. Halevy, and D. Weld. Adapting to source properties in processing data integration queries. In *Proceedings of the 2004 ACM SIGMOD International*

- Conference on Management of Data*, pages 395–406, New York, NY, USA, 2004. ACM.
- [33] N. Kabra and D. DeWitt. Efficient mid-query re-optimization of sub-optimal query execution plans. *SIGMOD Rec.*, 27(2):106–117, June 1998.
- [34] I. Kolchinsky and A. Schuster. Efficient adaptive detection of complex event patterns. *CoRR*, abs/1801.08588, 2017.
- [35] I. Kolchinsky and A. Schuster. Join query optimization techniques for complex event processing applications. *PVLDB*, 11(11):1332–1345, 2018.
- [36] I. Kolchinsky, A. Schuster, and D. Keren. Efficient detection of complex event patterns using lazy chain automata. *CoRR*, abs/1612.05110, 2016.
- [37] I. Kolchinsky, I. Sharfman, and A. Schuster. Lazy evaluation methods for detecting complex events. In *DEBS*, pages 34–45. ACM, 2015.
- [38] J. Krämer, Y. Yang, M. Cammert, B. Seeger, and D. Papadias. Dynamic plan migration for snapshot-equivalent continuous queries in data stream systems. In *Proceedings of the 2006 International Conference on Current Trends in Database Technology*, pages 497–516, Berlin, Heidelberg, 2006. Springer-Verlag.
- [39] O. Lee, E. You, M. Hong, and J. Jung. *Adaptive Complex Event Processing Based on Collaborative Rule Mining Engine*, pages 430–439. Springer International Publishing, Cham, 2015.
- [40] M. Liu, Z. Ives, and B. Loo. Enabling incremental query re-optimization. In *Proceedings of the 2016 International Conference on Management of Data*, pages 1705–1720, New York, NY, USA. ACM.
- [41] S. Madden, M. Shah, J. Hellerstein, and V. Raman. Continuously adaptive continuous queries over streams. In *Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data*, pages 49–60, New York, NY, USA, 2002. ACM.
- [42] V. Markl, V. Raman, D. Simmen, G. Lohman, H. Pirahesh, and M. Cilimdžić. Robust query processing through progressive optimization. In *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, pages 659–670, New York, NY, USA. ACM.
- [43] Y. Mei and S. Madden. ZStream: a cost-based query processor for adaptively detecting composite events. In *SIGMOD Conference*, pages 193–206. ACM, 2009.
- [44] R. Nehme, K. Works, C. Lei, E. Rundensteiner, and E. Bertino. Multi-route query processing and optimization. *J. Comput. Syst. Sci.*, 79(3):312–329, May 2013.
- [45] M. Sadoghi and H. Jacobsen. Adaptive parallel compressed event matching. In *IEEE 30th International Conference on Data Engineering, 2014*, pages 364–375, 2014.
- [46] N. P. Schultz-Møller, M. M., and P. R. Pietzuch. Distributed complex event processing with query rewriting. In *DEBS*. ACM, 2009.
- [47] M. Stillger, G. Lohman, V. Markl, and M. Kandil. Leo - db2's learning optimizer. In *Proceedings of the 27th International Conference on Very Large Data Bases*, pages 19–28, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- [48] A. Swami. Optimization of large join queries: Combining heuristics and combinatorial techniques. *SIGMOD Rec.*, 18(2):367–376, 1989.
- [49] N. Tatbul, U. Çetintemel, S. Zdonik, M. Cherniack, and M. Stonebraker. Load shedding in a data stream manager. In *Proceedings of the 29th International Conference on Very Large Data Bases - Volume 29*, pages 309–320. VLDB Endowment, 2003.
- [50] E. Wu, Y. Diao, and S. Rizvi. High-performance complex event processing over streams. In *SIGMOD Conference*, pages 407–418. ACM, 2006.
- [51] I. Yi, J. G. Lee, and K. Y. Whang. Apam: Adaptive eager-lazy hybrid evaluation of event patterns for low latency. In *Proceedings of the 25th ACM Conference on Information and Knowledge Management*, pages 2275–2280. ACM, 2016.
- [52] H. Zhang, Y. Diao, and N. Immerman. On complexity and optimization of expensive queries in complex event processing. In *SIGMOD*, pages 217–228, 2014.
- [53] Y. Zhu, E. Rundensteiner, and G. Heineman. Dynamic plan migration for continuous queries over data streams. In *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, pages 431–442, New York, NY, USA, 2004. ACM.