

EasyFlinkCEP: Big Event Data Analytics for Everyone

Nikos Giatrakos

Athena Research & Innovation Center,
Technical University of Crete

ngiatrakos@athenarc.gr

ngiatrakos@softnet.tuc.gr



INFORE

Interactive Extreme-Scale
Analytics and Forecasting

Outline

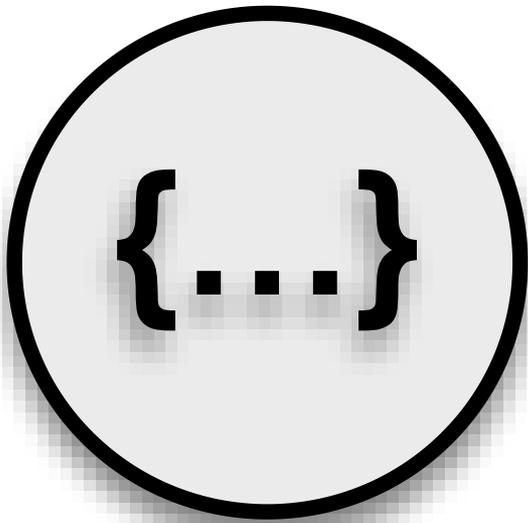
- Introduction (Streams & Apache Flink)
- CEP vs Traditional Streams
- FlinkCEP
 - Basics
 - Adoption barriers
- EasyFlinkCEP
 - EasyFlinkCEP Operator
 - EasyFlinkCEP Optimizer
- What we achieved
- Conclusions & Future Work

Outline

- Introduction (Streams & Apache Flink)
- CEP vs Traditional Streams
- FlinkCEP
 - Basics
 - Adoption barriers
- EasyFlinkCEP
 - EasyFlinkCEP Operator
 - EasyFlinkCEP Optimizer
- What we achieved
- Conclusions & Future Work

Streaming Model

- Tuple



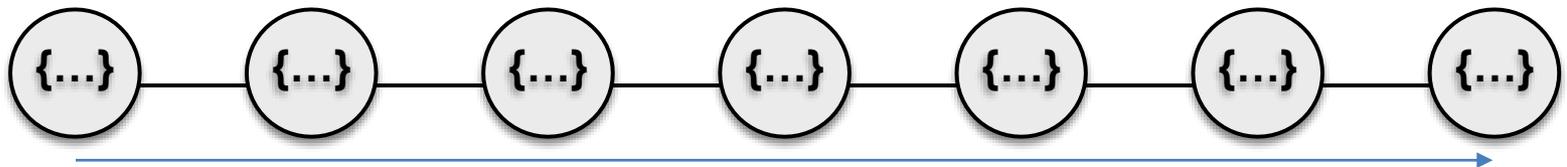
- Core Unit of Data
- Immutable Set of Key/Value Pairs
- Stock Market Example:

```
StockID Date           Time           Price  Volume
EURTRY, 01/15/2019,09:57:41, 29.21, 152
```

Streaming Model

- Stream

Unbounded Sequence of Tuples



Time

StockID	Date	Time	Price	Volume
---------	------	------	-------	--------

EURTRY,	01/15/2019,	09:57:41,	29.21,	152
---------	-------------	-----------	--------	-----

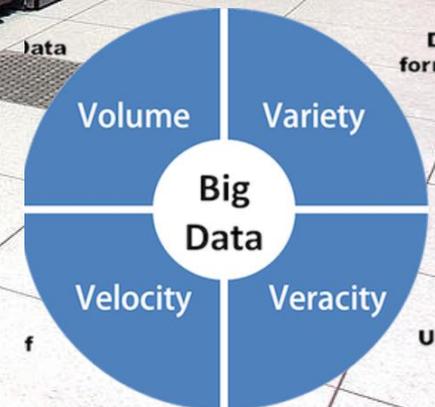
EURTRY,	01/15/2019,	09:58:45,	29.2,	190
---------	-------------	-----------	-------	-----

EURTRY,	01/15/2019,	09:58:45,	29.2,	177
---------	-------------	-----------	-------	-----

...

...

Scalability – Real Time Analytics



https://en.wikipedia.org/wiki/Blue_Gene/P



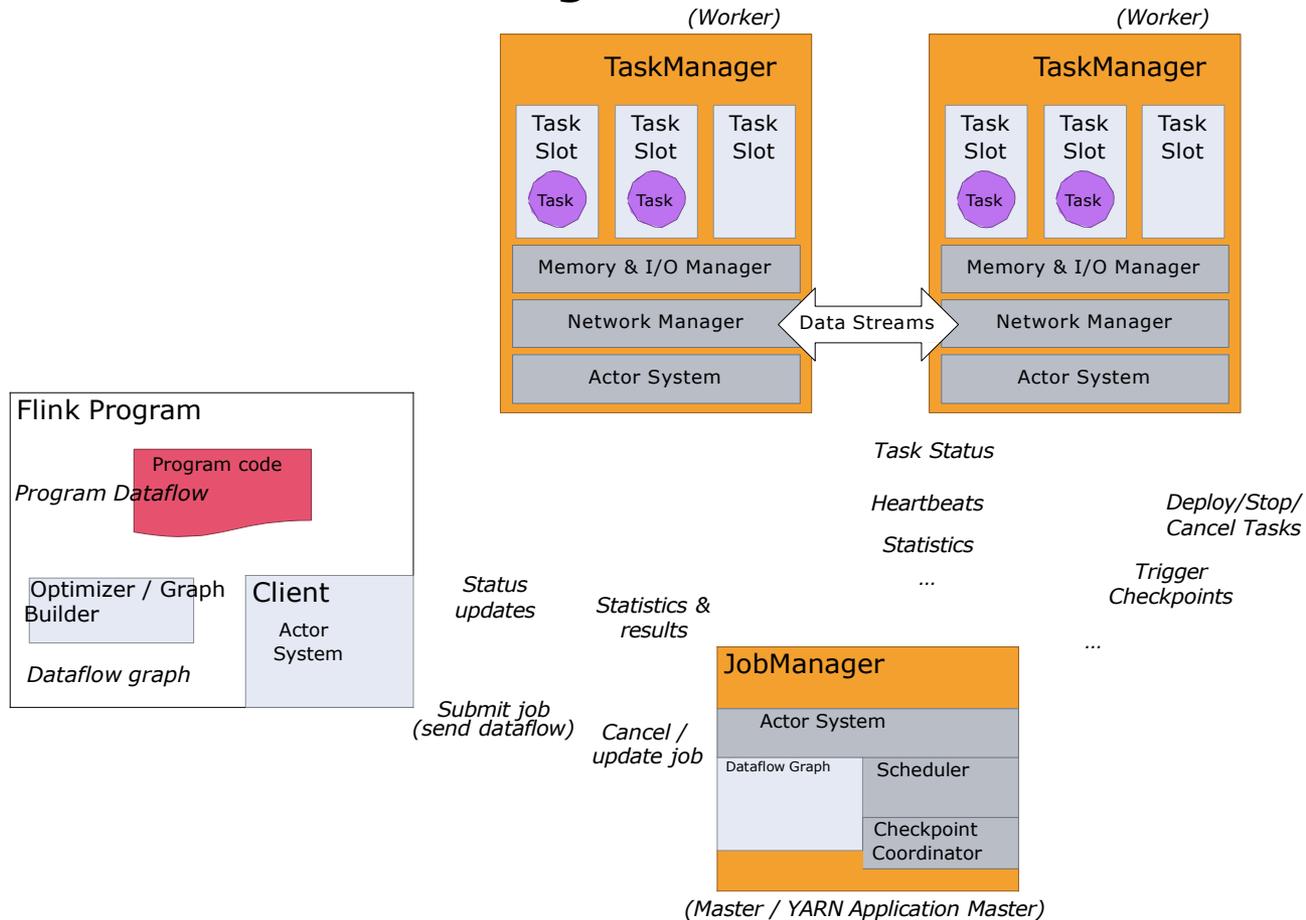
INFORE

Interactive Extreme-Scale
Analytics and Forecasting

EasyFlinkCEP

Apache Flink – Basics

- Master/worker concept can be integrated into YARN
- The client (Flink Program) is an external process
- **Notable new! Task chaining**



Programming Model

■ A DAG of streams applies transformations

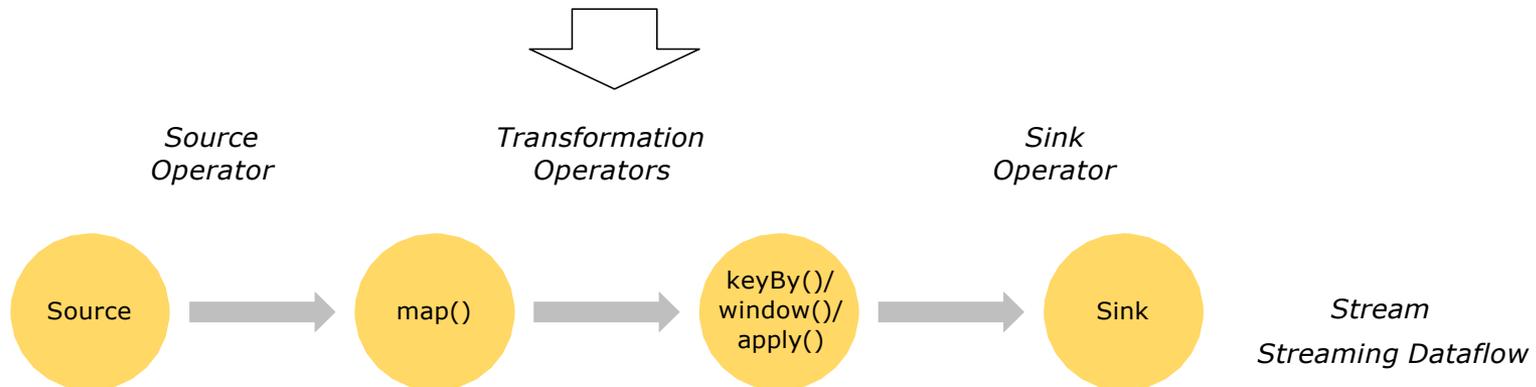
```
DataStream<String> lines = env.addSource(  
    new FlinkKafkaConsumer<> (...));  
  
DataStream<Event> events = lines.map((line) -> parse(line));  
  
DataStream<Statistics> stats = events  
    .keyBy("id")  
    .timeWindow(Time.seconds(10))  
    .apply(new MyWindowAggregationFunction());  
  
stats.addSink(new RollingSink(path));
```

Source

Transformation

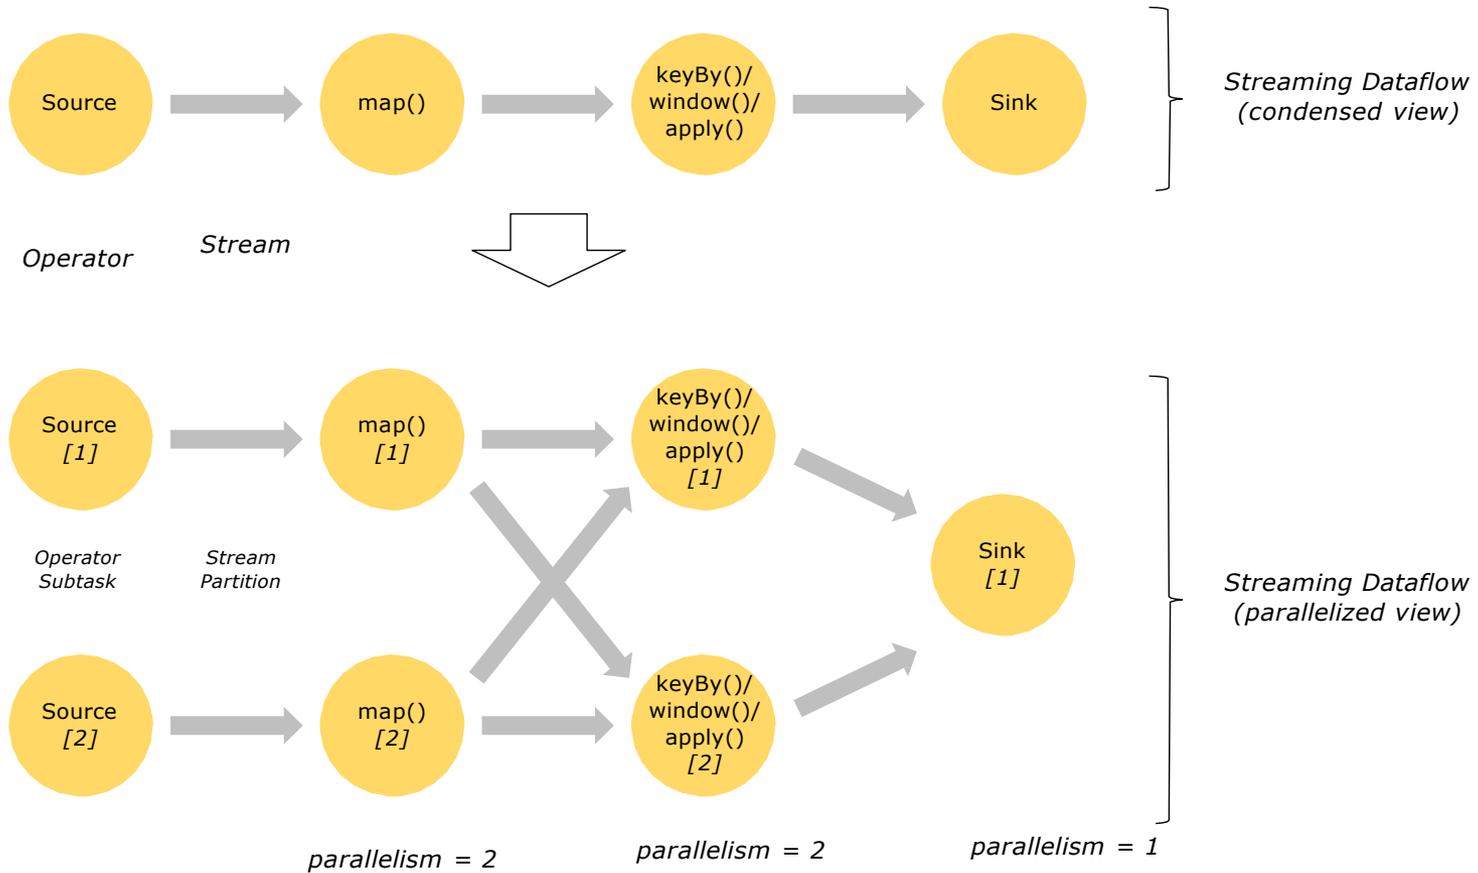
Transformation

Sink



Parallelization

- Parallelization via stream partitions and operator subtasks
- One-to-one streams preserve the order, redistribution changes them



Outline

- Introduction (Streams & Apache Flink)
- CEP vs Traditional Streams
- FlinkCEP
 - Basics
 - Adoption barriers
- EasyFlinkCEP
 - EasyFlinkCEP Operator
 - EasyFlinkCEP Optimizer
- What we achieved
- Conclusions & Future Work

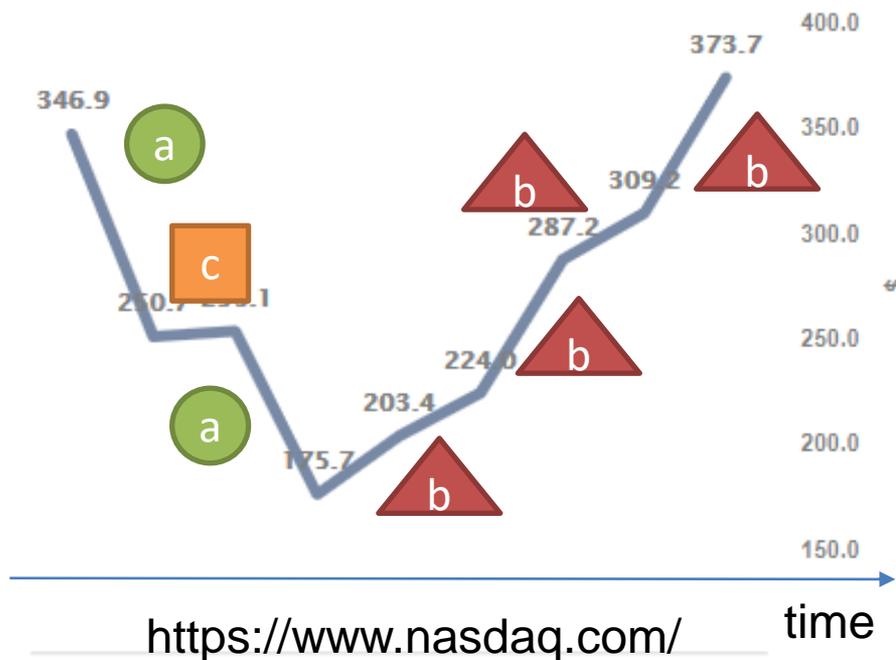
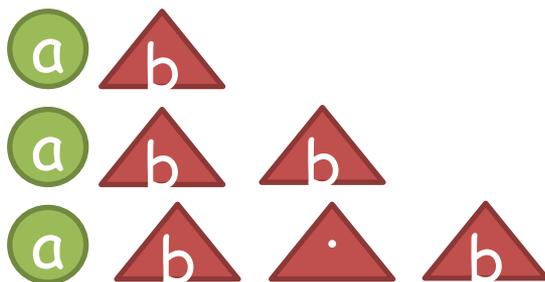
Motivating Example – CEP in Stock Market



Motivating Example – CEP in Stock Market

- a** : decreasing trend
- b** : increasing trend
- c** : steady trend
- d** : undetermined trend

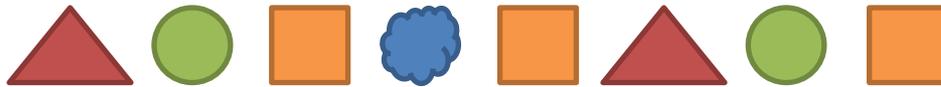
Increasing
after Decreasing
Price Trend Patterns:



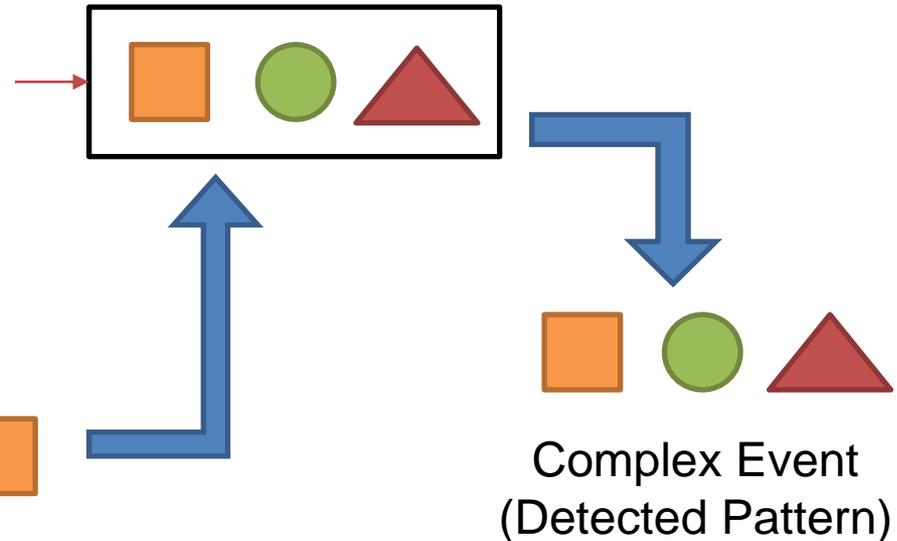
Complex Event Processing (CEP)

Query in CEP Languages
SEQ, OR, looping, use of quantifiers
and more....

Input Stream of "Simple" Events
(can be raw stream tuples)

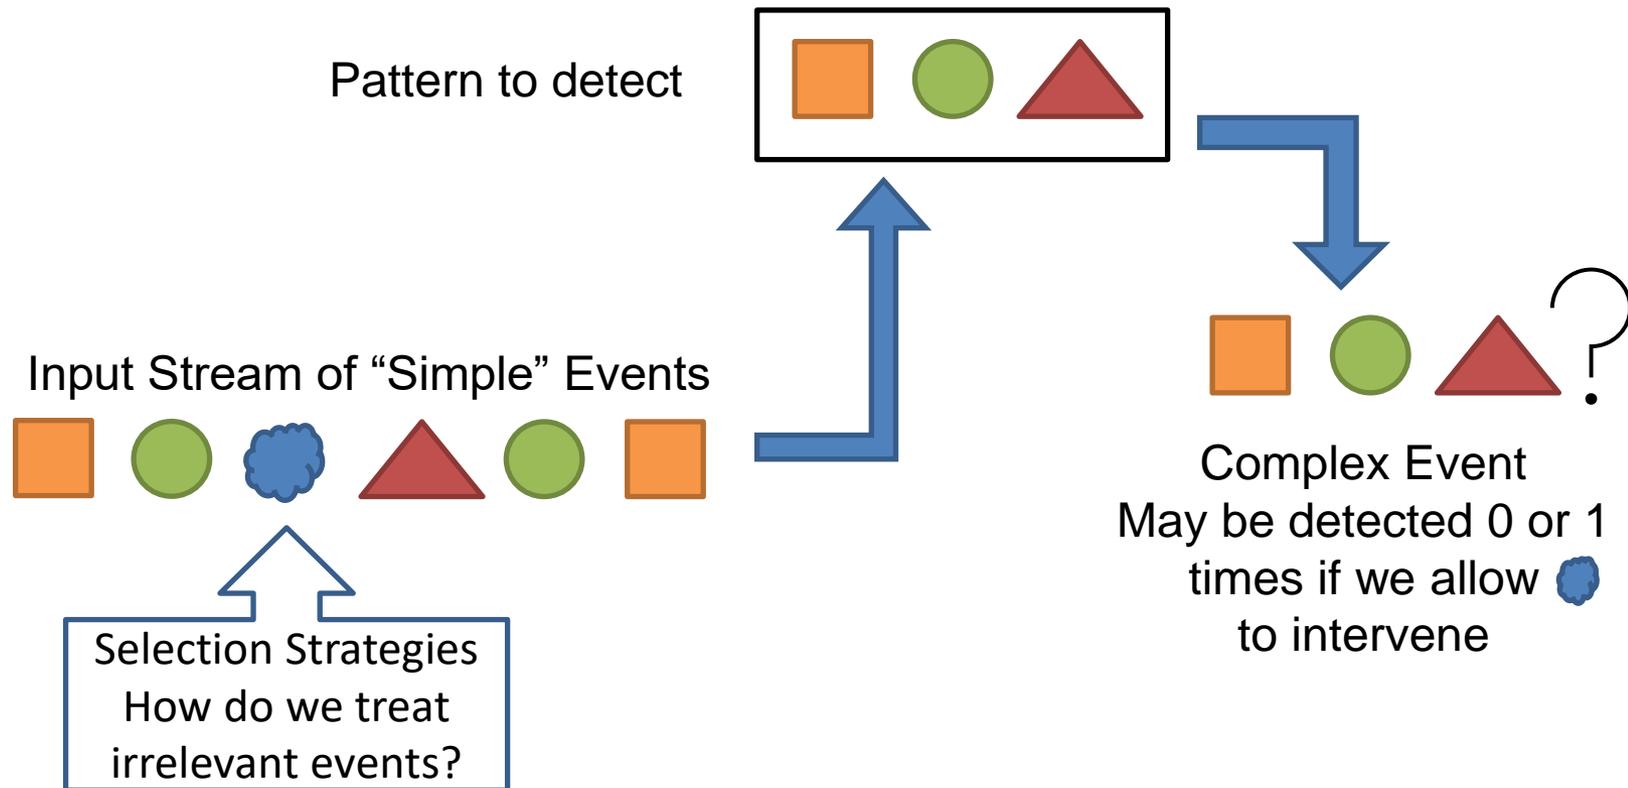


Complex Event Processing (CEP) System



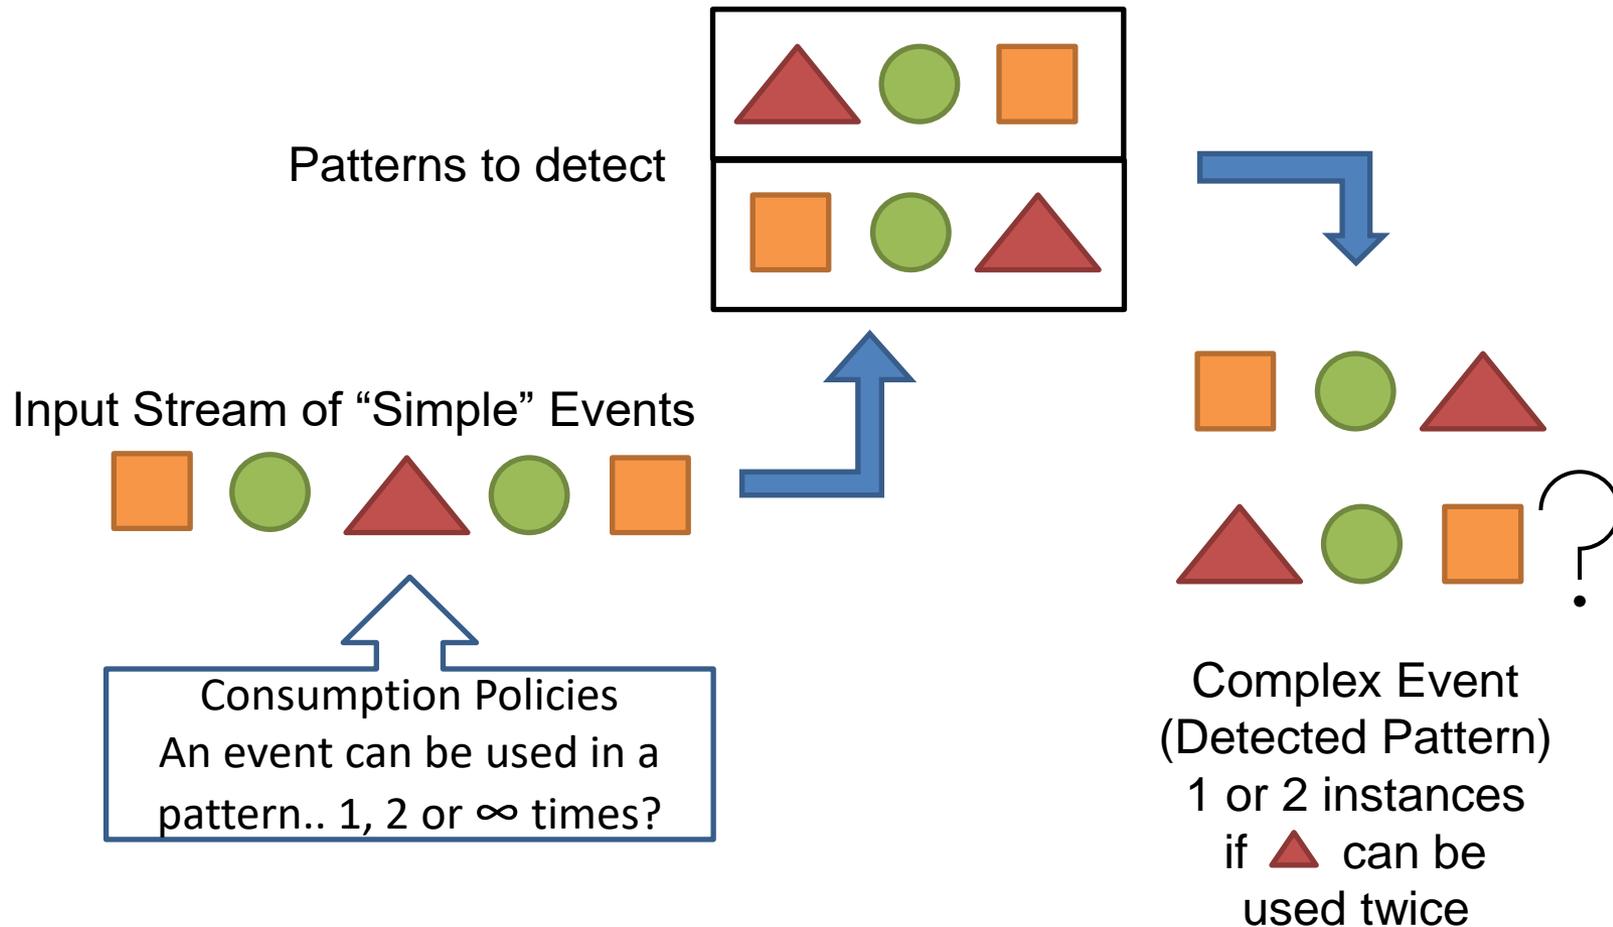
Selection Strategies

Complex Event Processing (CEP) System



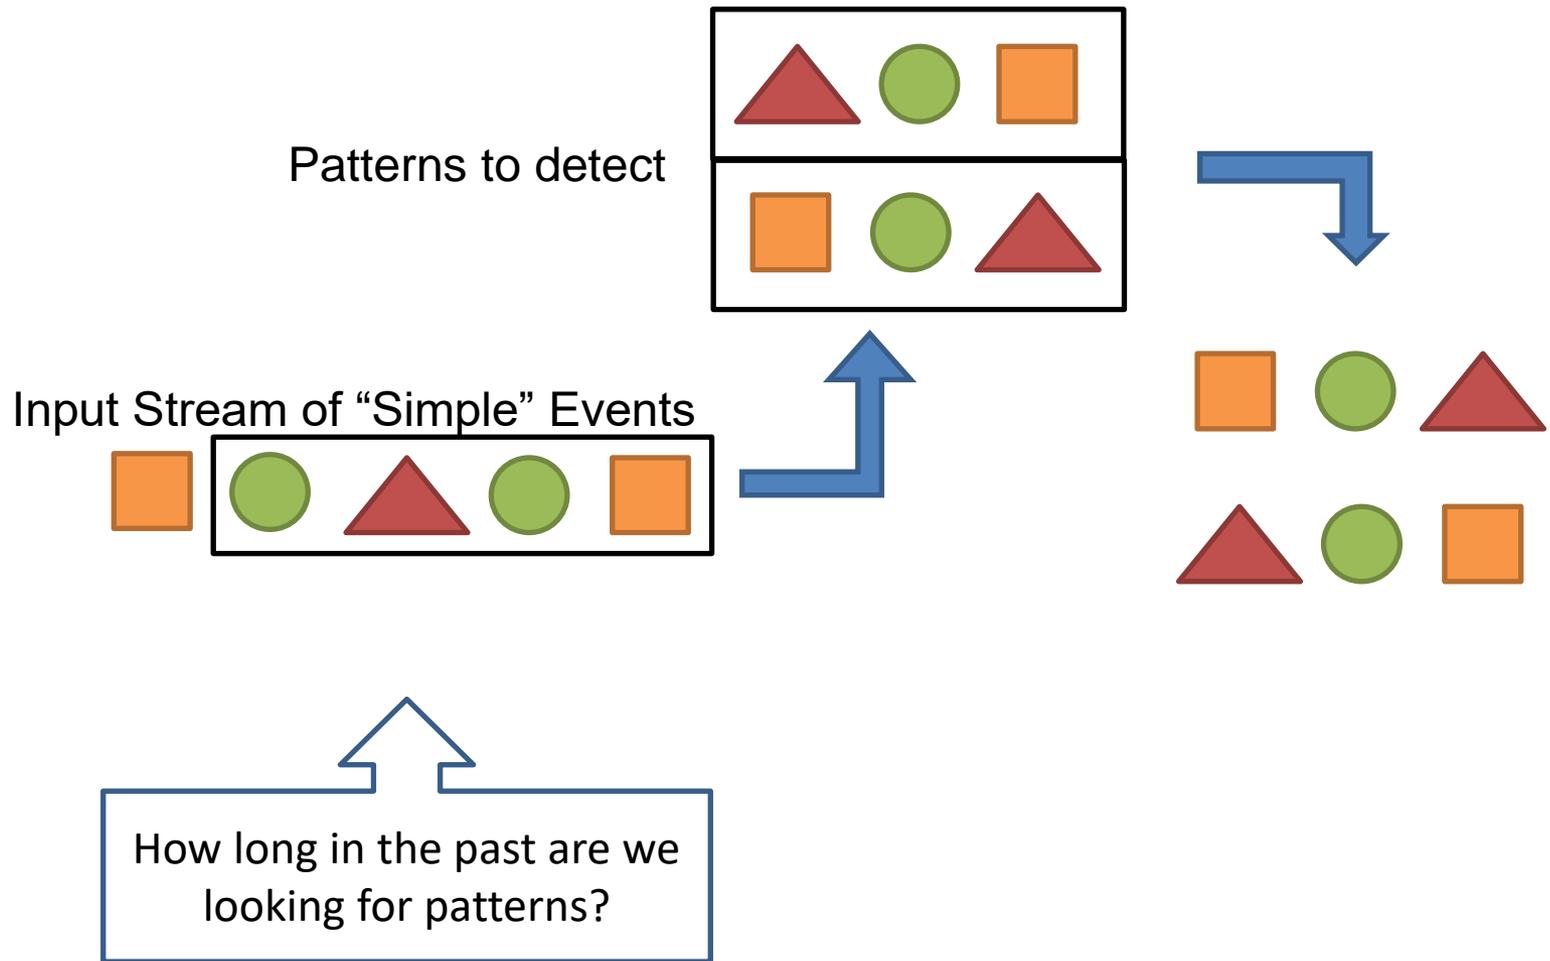
Consumption Policies

Complex Event Processing (CEP) System



Windows

Complex Event Processing (CEP) System

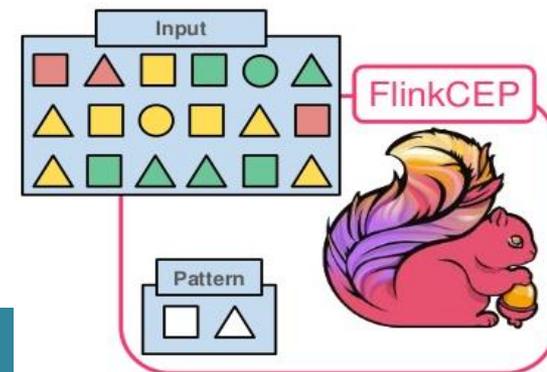


Outline

- Introduction (Streams & Apache Flink)
- CEP vs Traditional Streams
- FlinkCEP
 - Basics
 - Adoption barriers
- EasyFlinkCEP
 - EasyFlinkCEP Operator
 - EasyFlinkCEP Optimizer
- What we achieved
- Conclusions & Future Work

FlinkCEP: Native CEP on Apache Flink

- ✓ Parallel processing in a state-of-the-art Big Data platform
- ✓ Native API for defining CEP pipelines
- ✓ CEP language of high expressive power



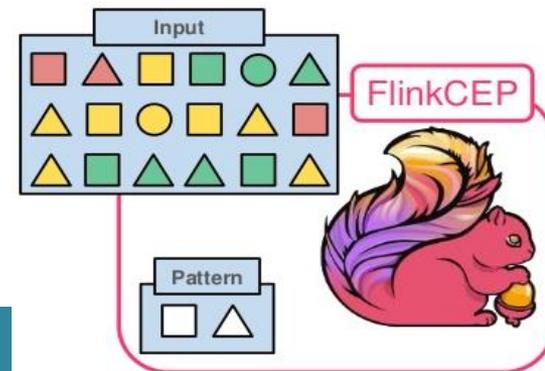
FlinkCEP: Native CEP on Apache Flink

✓ Selection Strategies (Example:  )

✓ STRICT

Input Stream:      

Matches:  



FlinkCEP: Native CEP on Apache Flink

✓ Selection Strategies (Example:  )

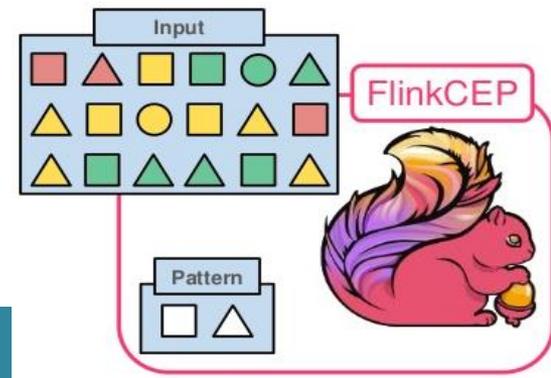
✓ Relaxed

Input Stream:      

Matches:  



FlinkCEP: Native CEP on Apache Flink

✓ Selection Strategies (Example:  )

✓ Non-deterministic

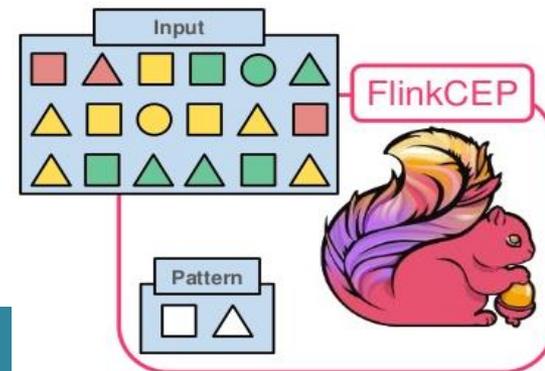
Input Stream:      

Matches:  

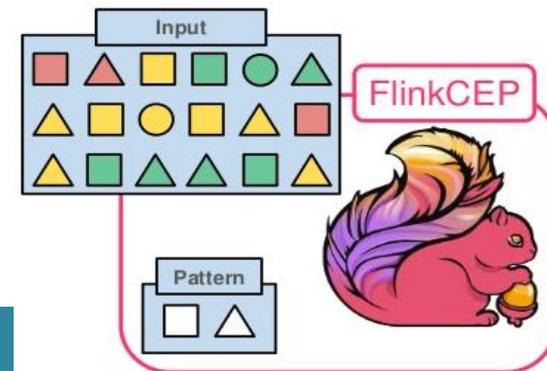
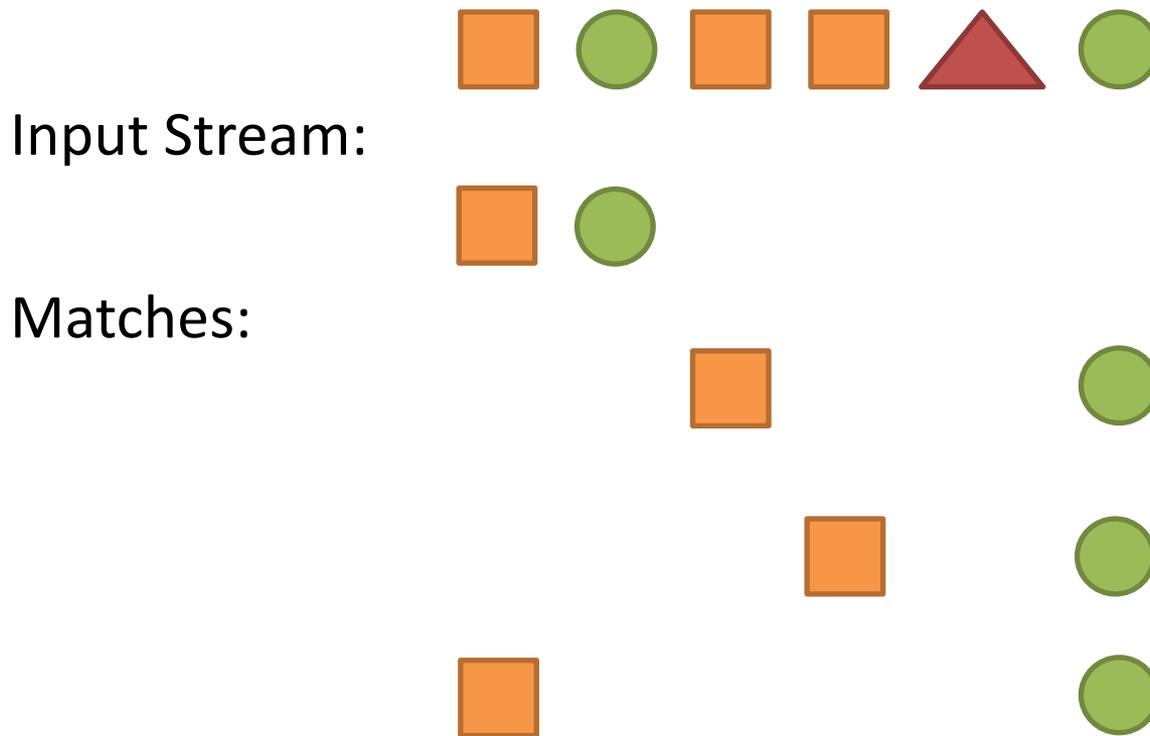
 



FlinkCEP: Native CEP on Apache Flink

✓ Consumption Policies (Example:  )
assuming Non-deterministic Selection

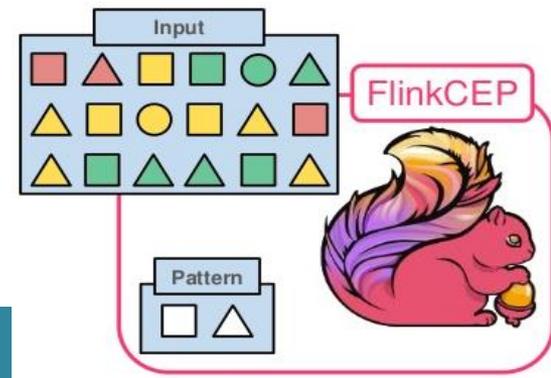
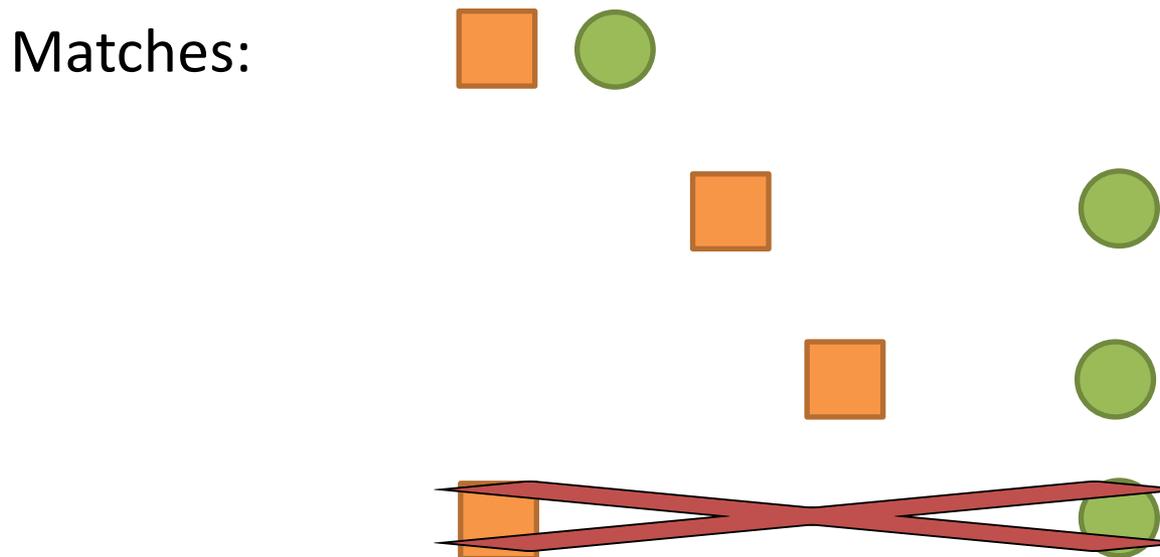
✓ NO_SKIP (Reuse everything)



FlinkCEP: Native CEP on Apache Flink

- ✓ Consumption Policies (Example:  )
 - ✓ SKIP_TO_NEXT (Uses starting event only once)

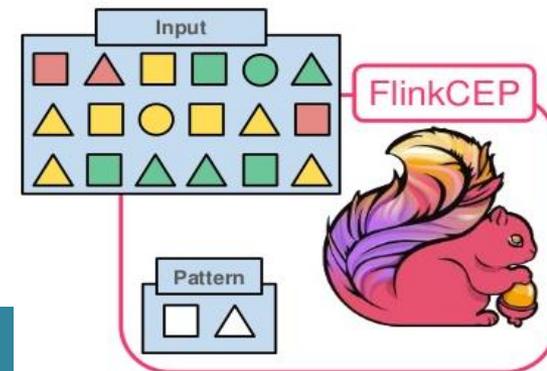
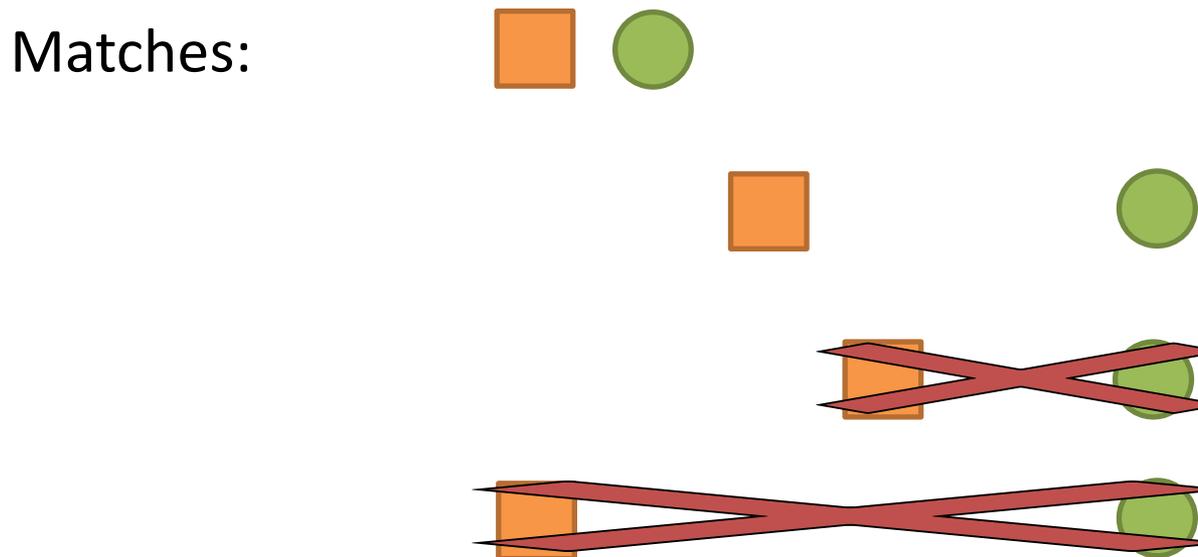
Input Stream:      



FlinkCEP: Native CEP on Apache Flink

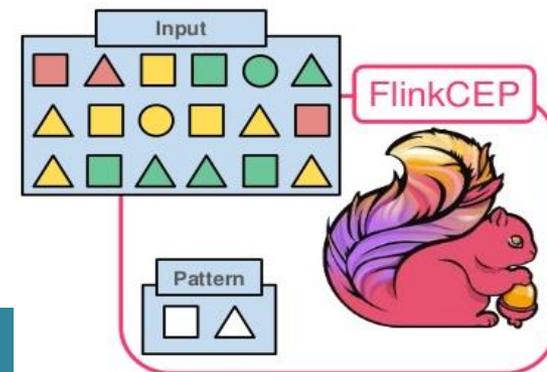
- ✓ Consumption Policies (Example:  
 - ✓ SKIP_TO_LAST () (drop partial matches in-between full match)

Input Stream:      



FlinkCEP: Native CEP on Apache Flink

- ✓ Parallel processing in a state-of-the-art Big Data platform
- ✓ Native API for defining CEP pipelines
- ✓ CEP language of high expressive power
- Difficult for non-programmer Business Analysts
 - Java/Scala Coding
- Cumbersome pattern parameterization
- Cluster administration decisions needed



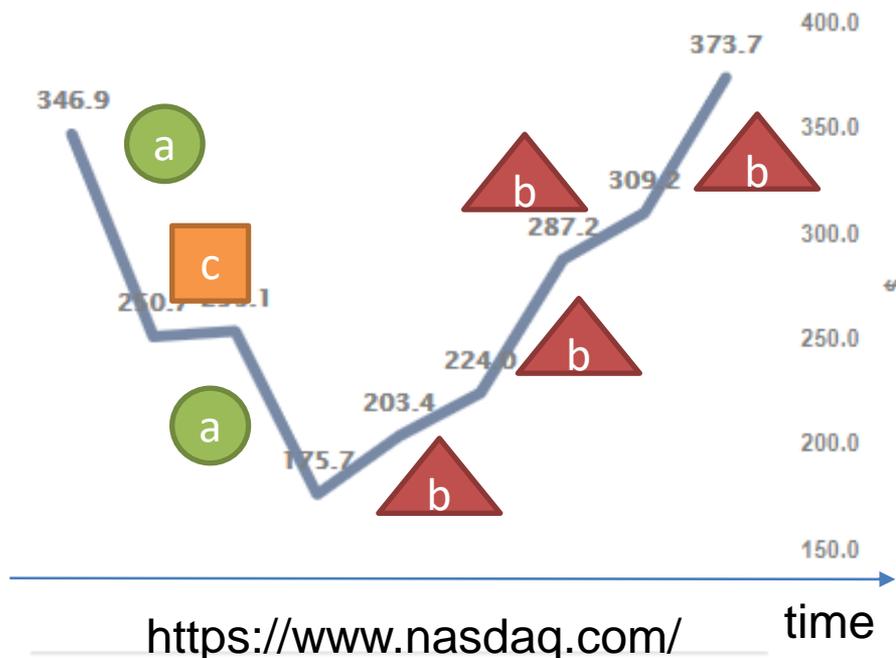
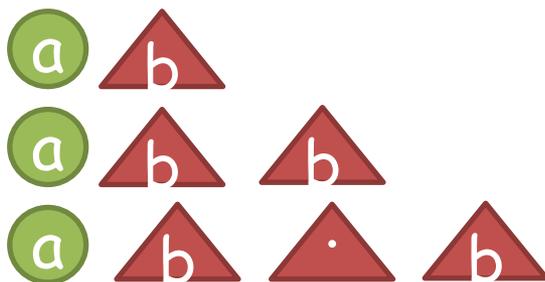
Outline

- Introduction (Streams & Apache Flink)
- CEP vs Traditional Streams
- FlinkCEP
 - Basics
 - Adoption barriers
- EasyFlinkCEP
 - EasyFlinkCEP Operator
 - EasyFlinkCEP Optimizer
- What we achieved
- Conclusions & Future Work

Motivating Example – CEP in Stock Market

-  : decreasing trend
-  : increasing trend
-  : steady trend
-  : undetermined trend

Increasing
after Decreasing
Price Trend Patterns:

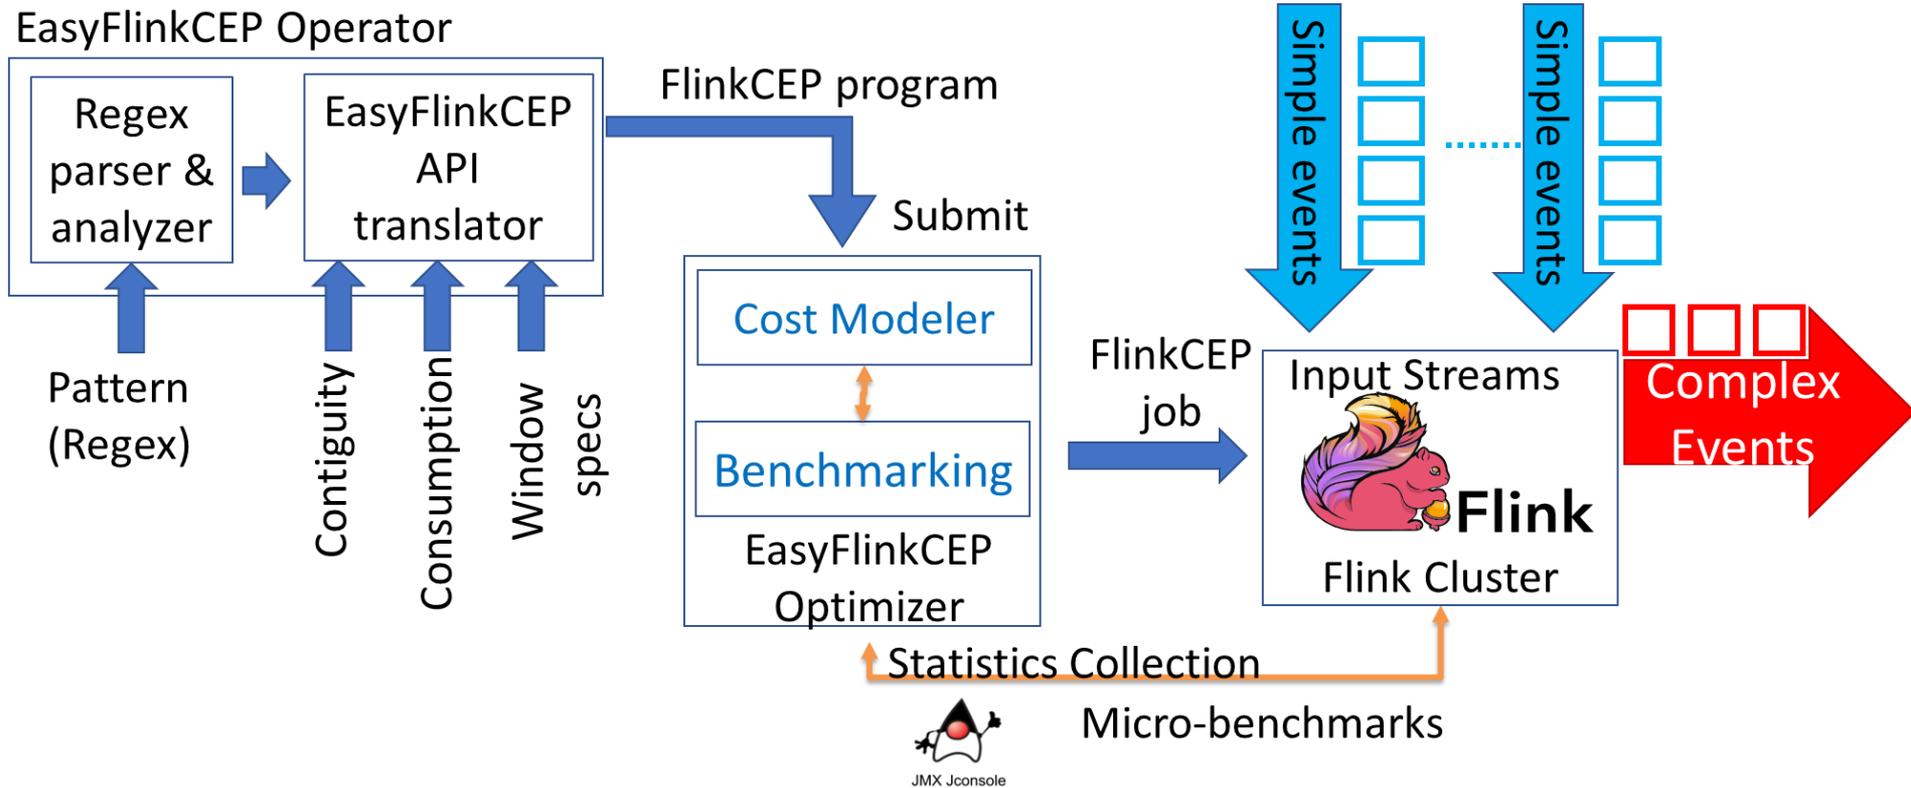


<https://www.nasdaq.com/> time

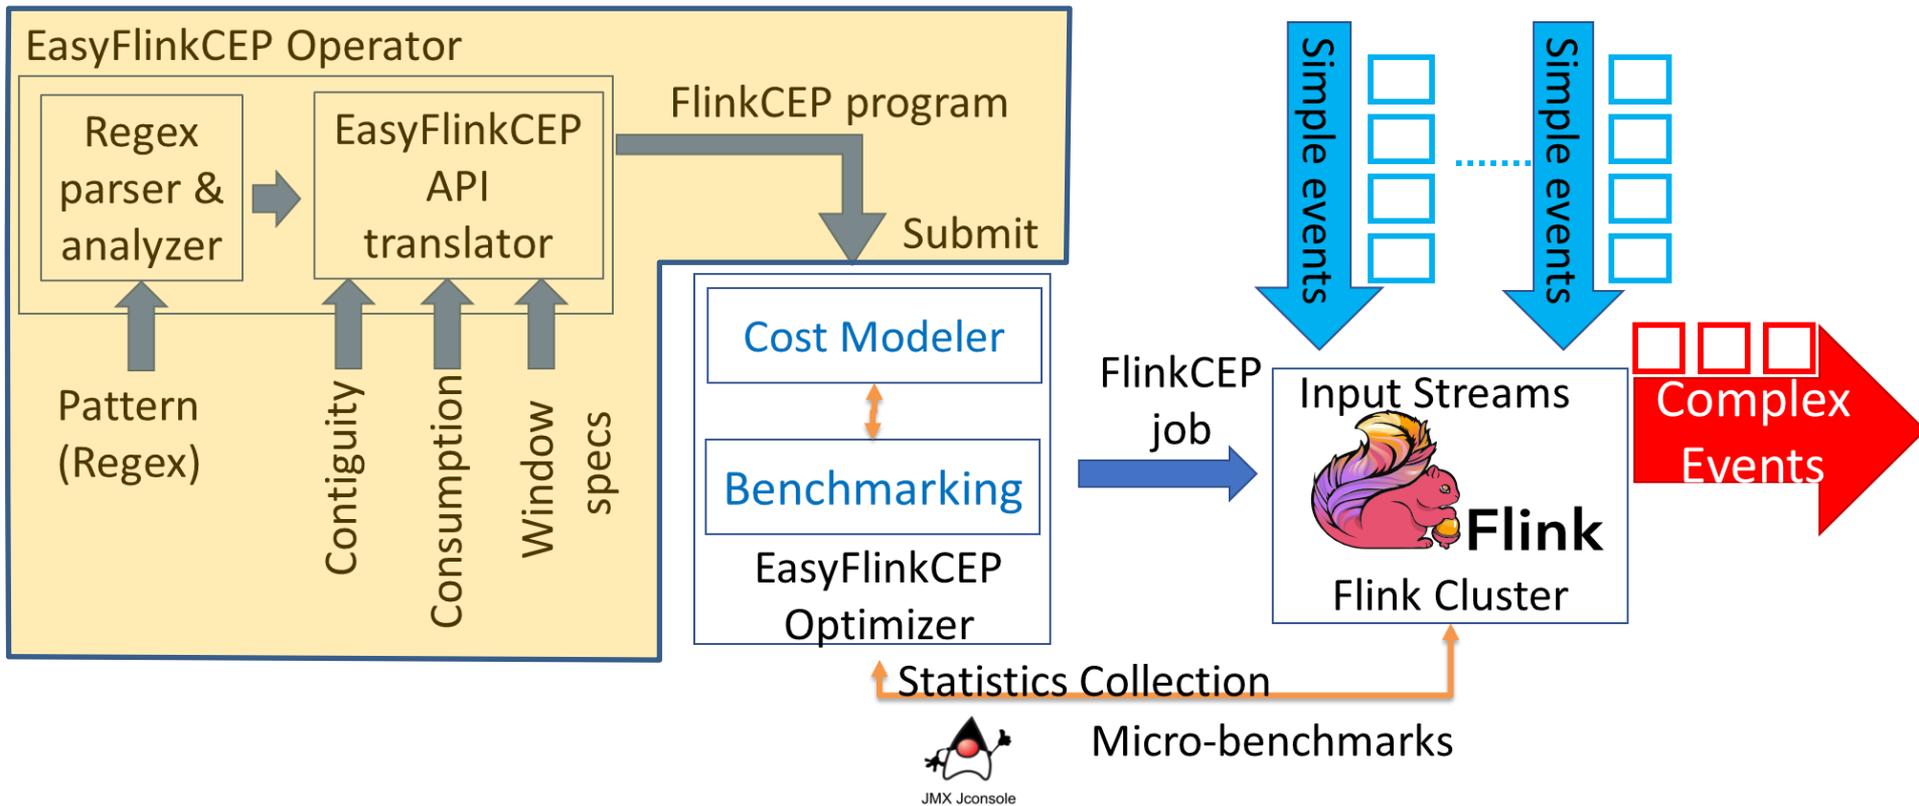
Regular Expressions Plug

- The language of the “Business Analysts”
 - Examples
 - ab
 - ab^+
 - ab^*
 - $ab\{1, 3\}c|d$
 - $a\{2, 5\}b?(c|d)$

EasyFlinkCEP Architecture & Contributions



EasyFlinkCEP Operator



Motivating Example – The FlinkCEP Program

```
// ----- Create pattern "a b+" -----
Pattern<Event, ?> pattern = Pattern.<Event>begin(" start ", skipStrategy ).where(new
SimpleCondition<>() {
    @Override
    public boolean filter (Event value) {
        return value.getName().startsWith ("a" );}
}) .followedBy("next") .where(new SimpleCondition<>() { //followedBy for Relaxed Contiguity
    @Override
    public boolean filter (Event value) {
        return value.getName().startsWith ("b" );}
}) .oneOrMore().consecutive () ; // oneOrMore().consecutive () for b+
PatternStream<Event> patternStream = CEP.pattern(input , pattern );
```

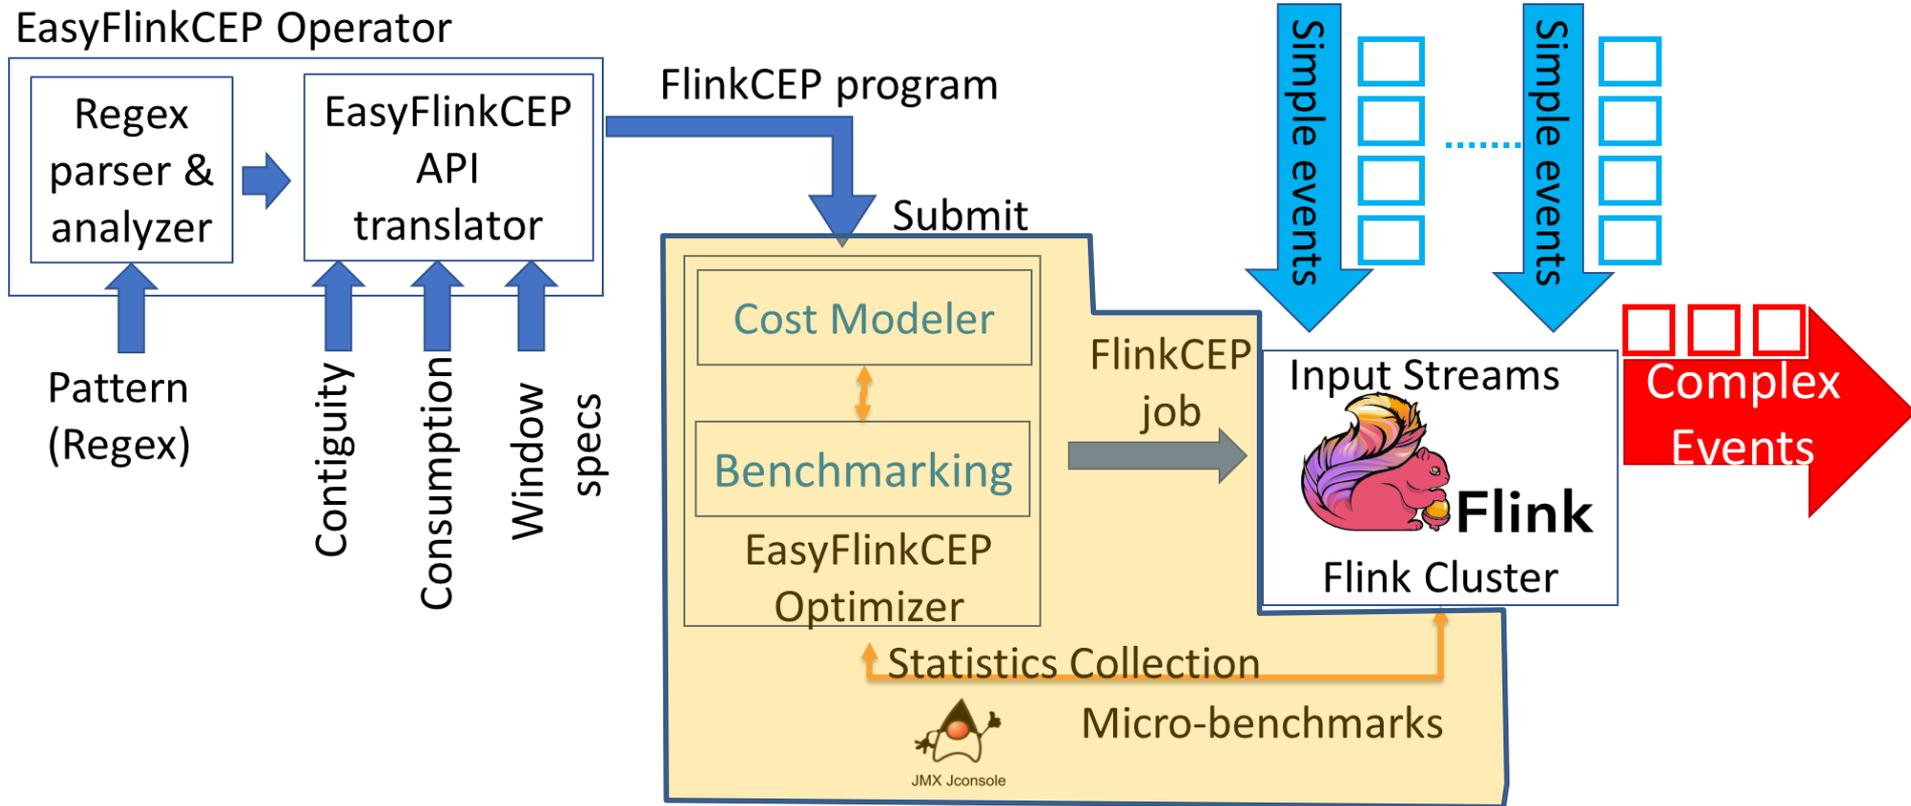
Motivating Example – EasyFlinkCEP Operator

EasyFlinkCEP Operator allows to:

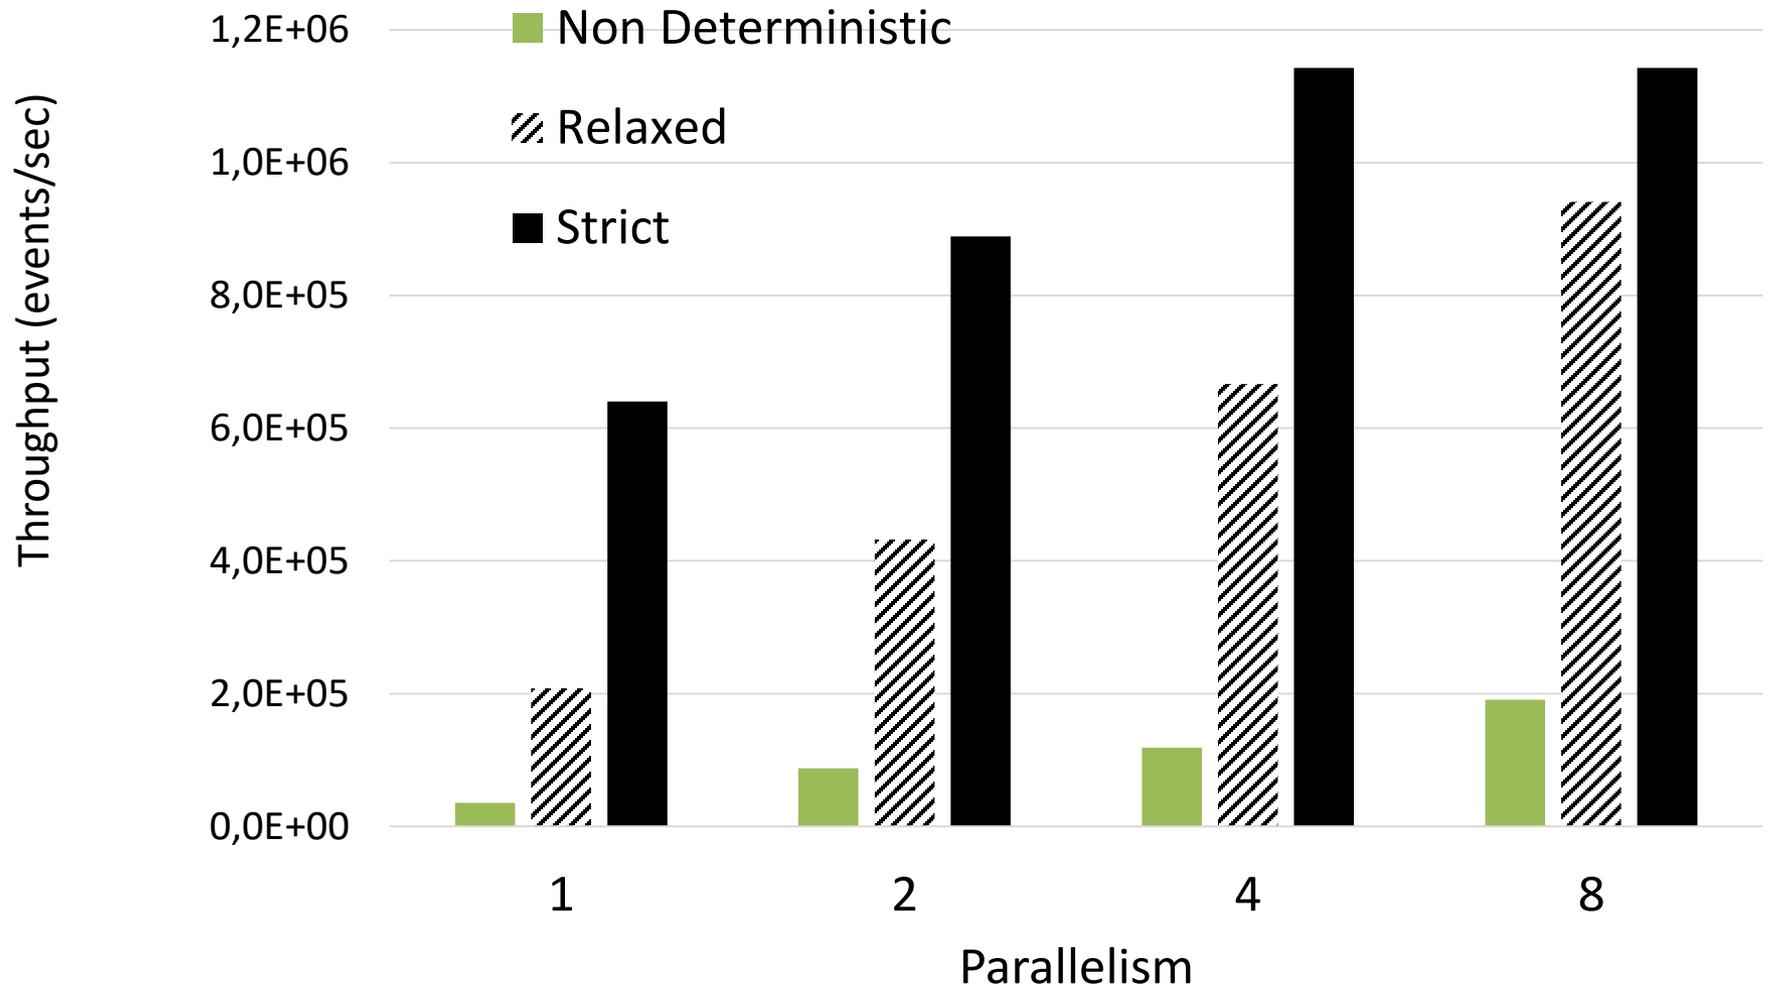
- ✓ Graphically Define
RegEx, Selection, Consumption, Window
- ✓ Internally writes the FlinkCEP program

The screenshot displays the EasyFlinkCEP interface. At the top, there is a home icon and a text input field labeled "RegEx:" containing the value "ab+". Below this, there are three main menu items: "Selection Strategies", "Consumption Policies", and "Window". The "Selection Strategies" dropdown is open, showing options: "Strict Contiguity", "Relaxed Contiguity", "Non-Deterministic ...", and "...". The "Consumption Policies" dropdown is also open, showing options: "NO_SKIP", "SKIP_TO_NEXT", "SKIP_PAST_LAST_EVENT", and "...".

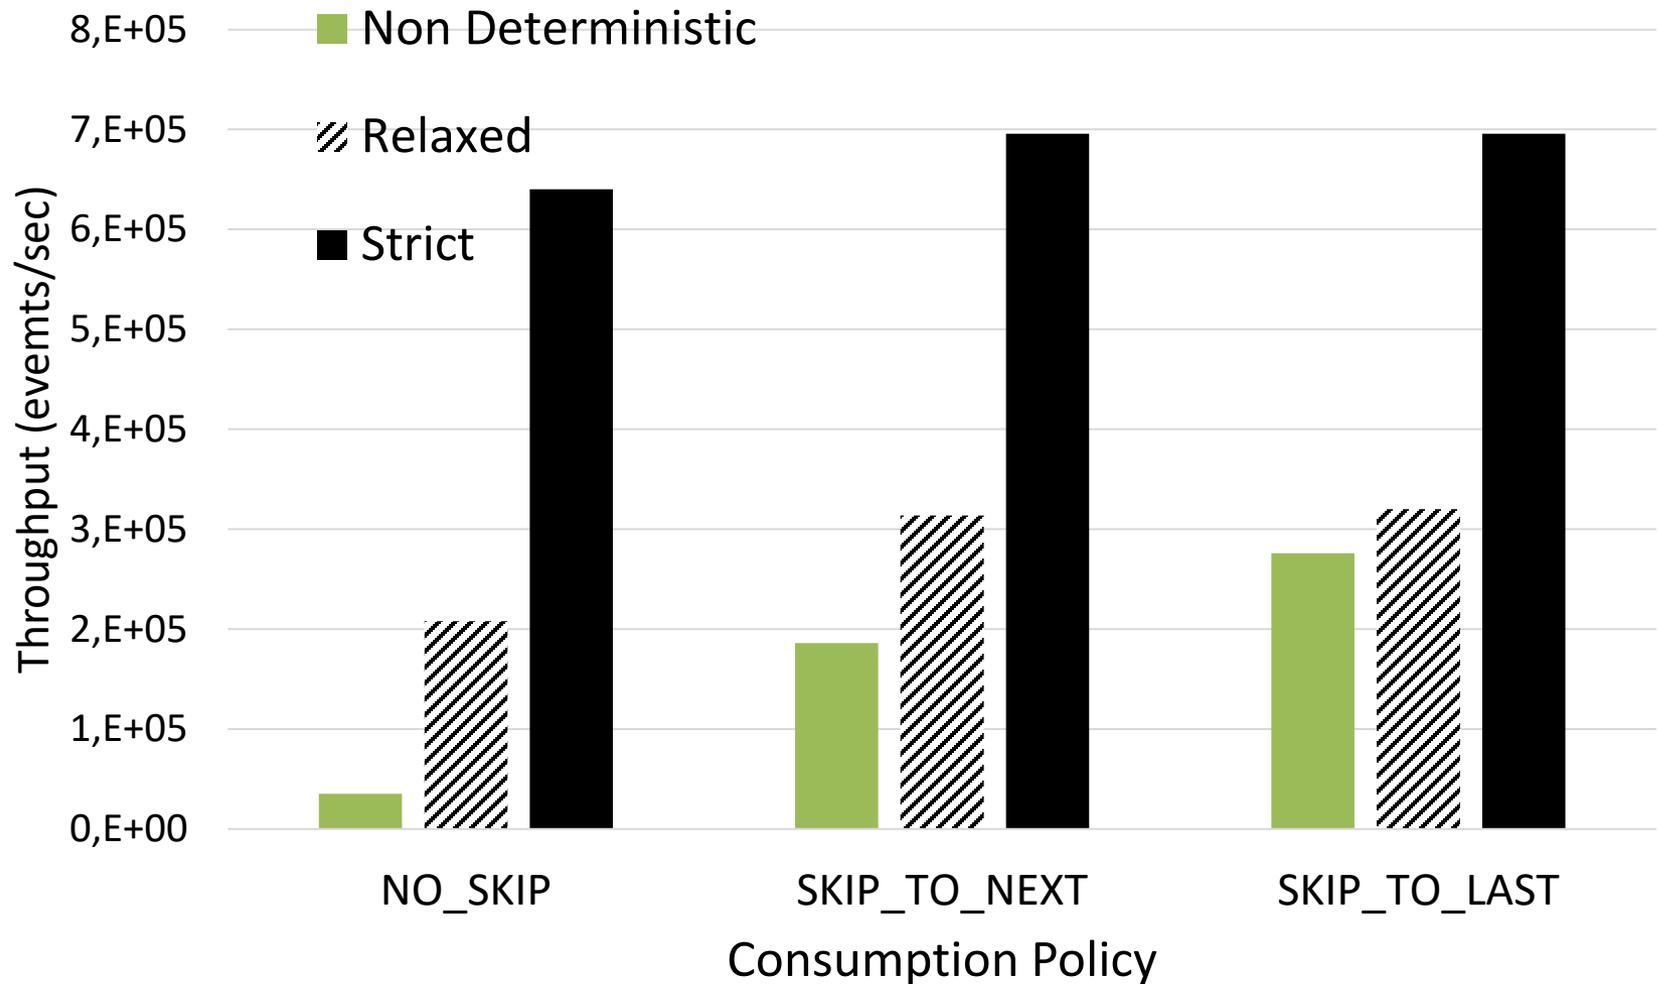
EasyFlinkCEP Optimizer



Exploring Selection Strategy Effect



Exploring Consumption Policy Effect



Motivation for Bayesian Optimization (BO)

Selection strategy:

How many partial matches to cache&check with each input?

Consumption policy:

Holds computation, check if events consumed in parallel instances

RegEx → Black-box Function

Window:

affects mem usage,
consumed/expired events

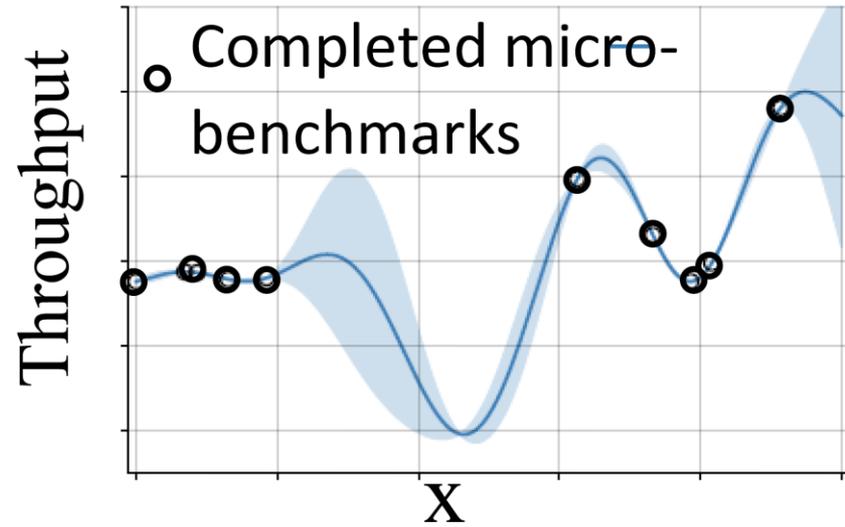
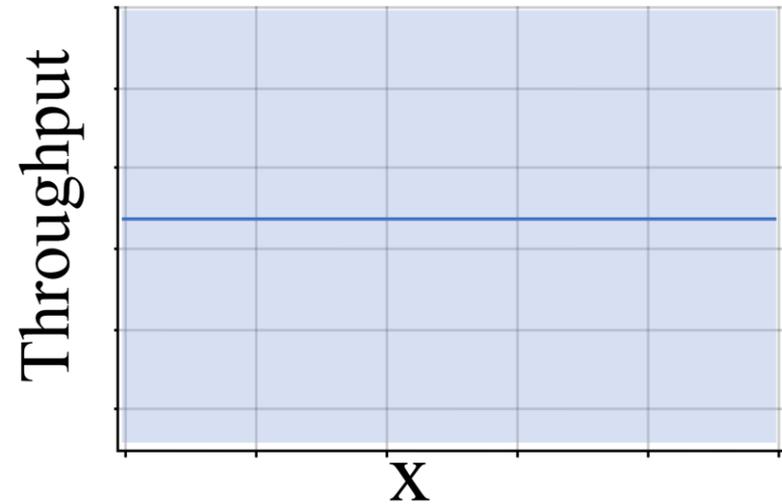
Parallelism?

No analytic formula to
quantify workload

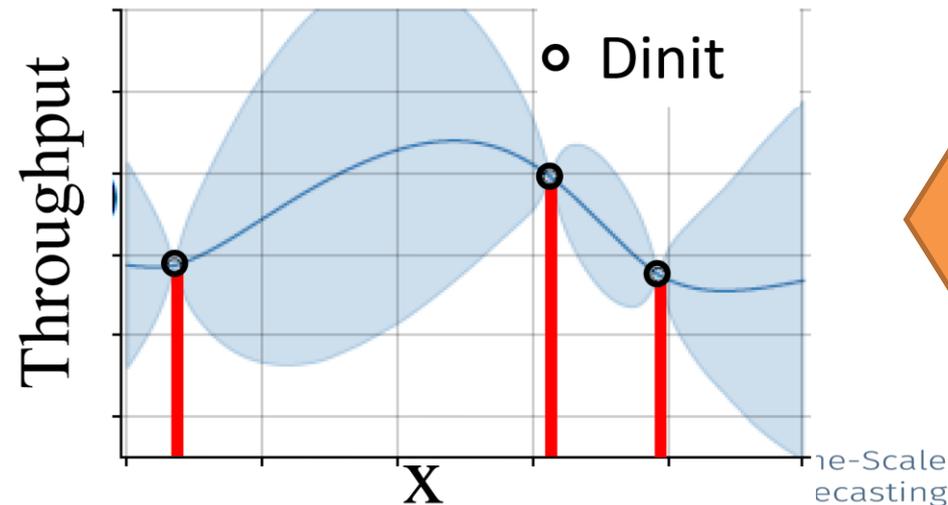
- Learn RegEx computational demands
 - Expensive function evaluations / Benchmarks
 - Time consuming



EasyFlinkCEP Benchmarking (1/2)

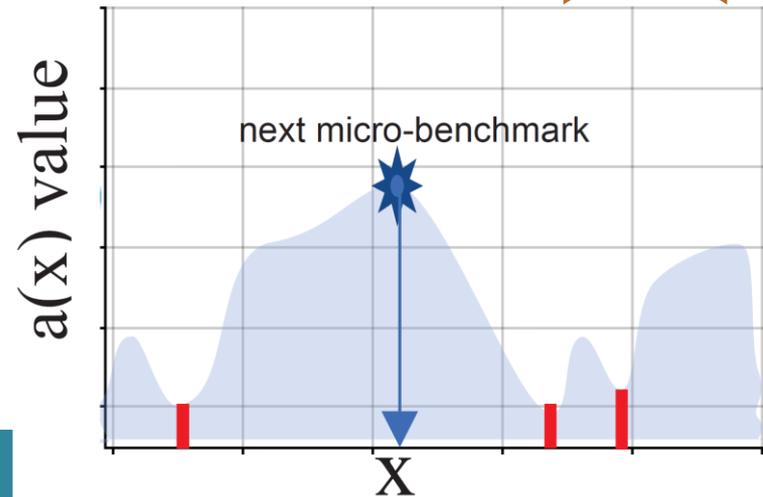


 $X=(\text{selection strategy, consumption policy, window, parallelism})$

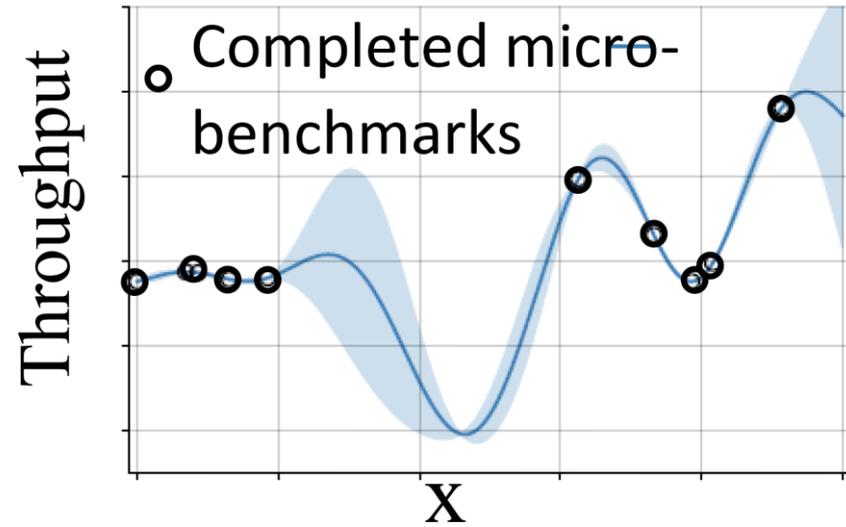
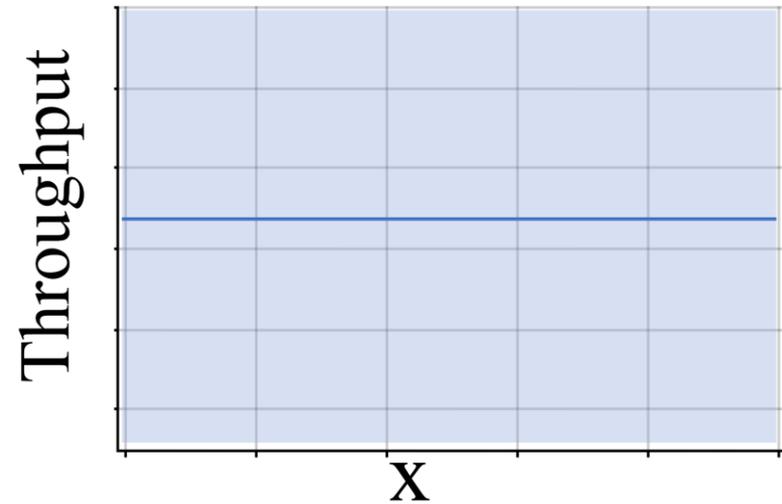


re-Scale
ecasting

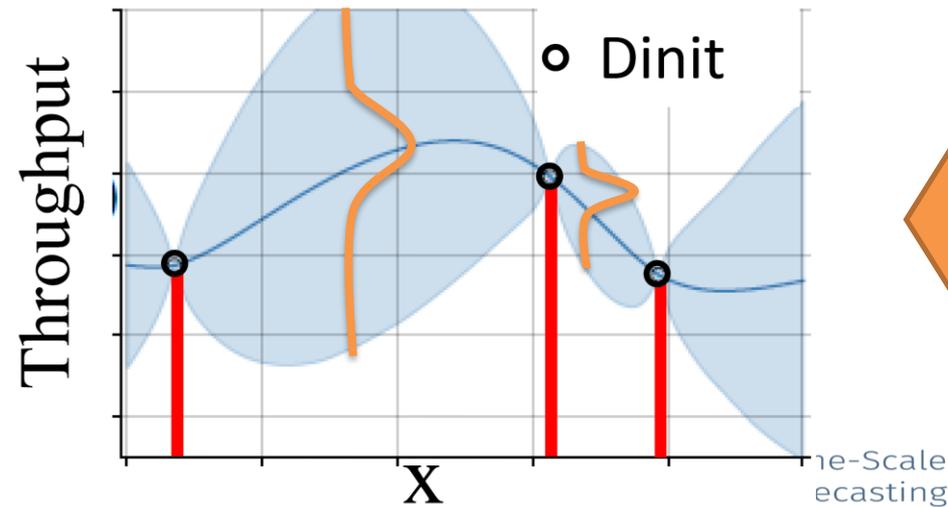
EasyFlinkCEP



EasyFlinkCEP Benchmarking (1/2)

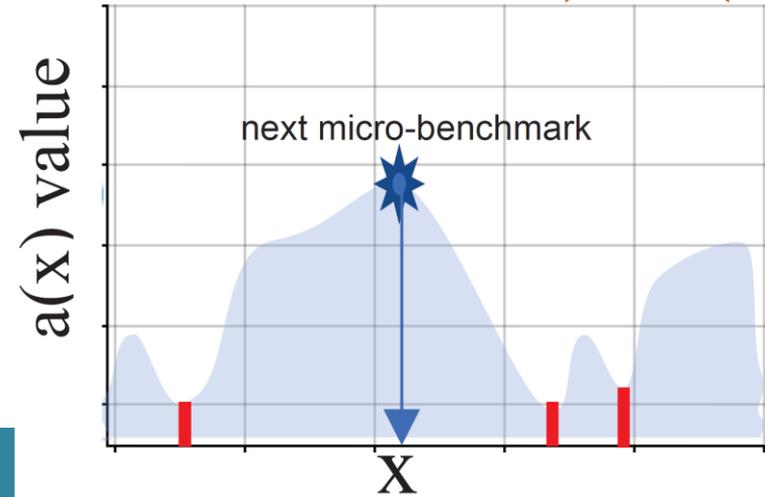


X=(selection strategy, consumption policy, window, parallelism)

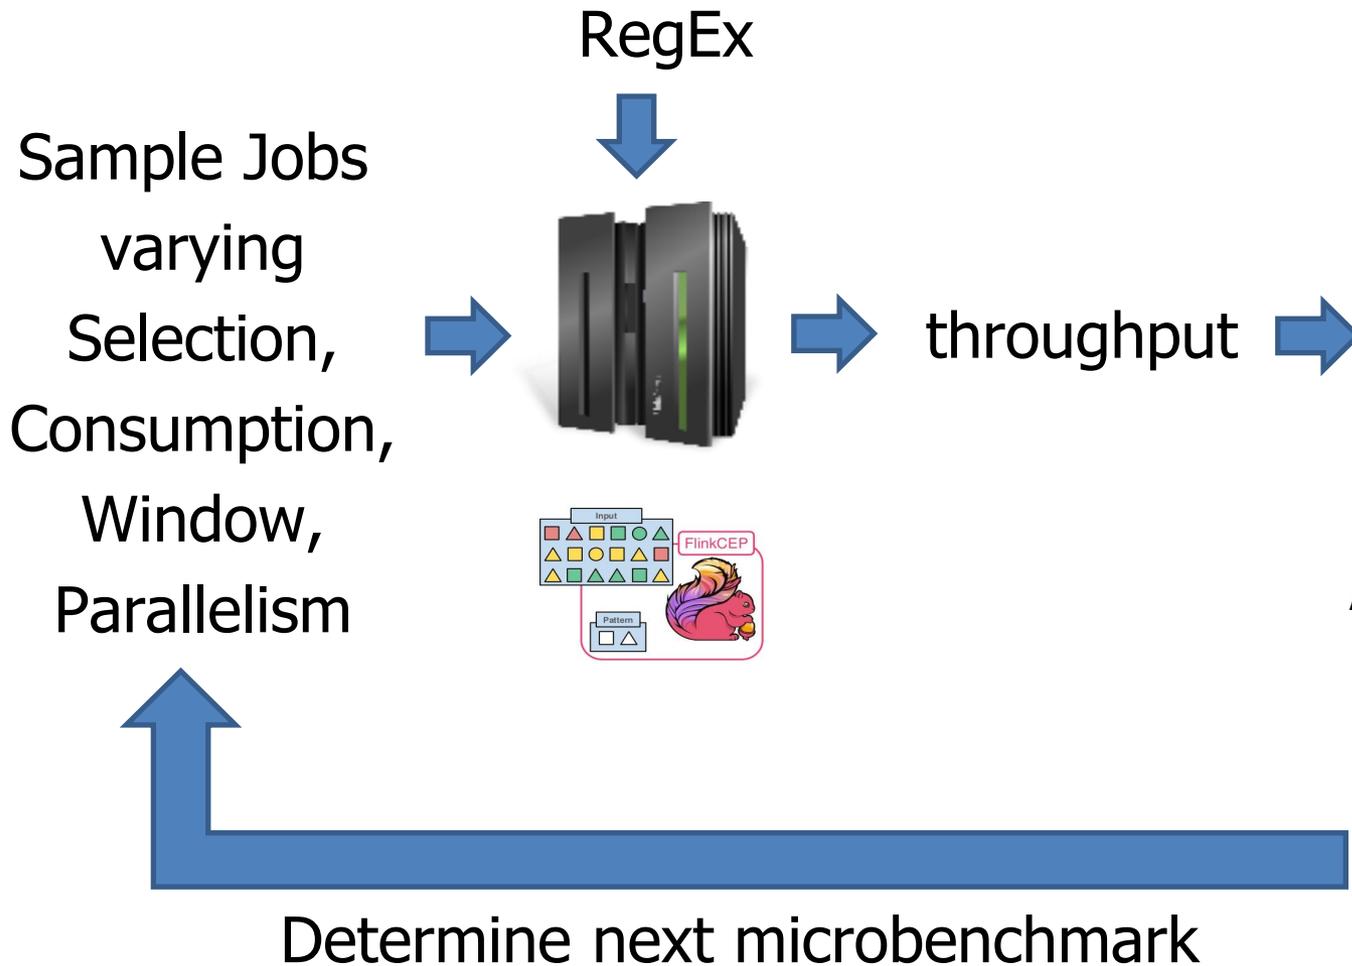


One-Step
Forecasting

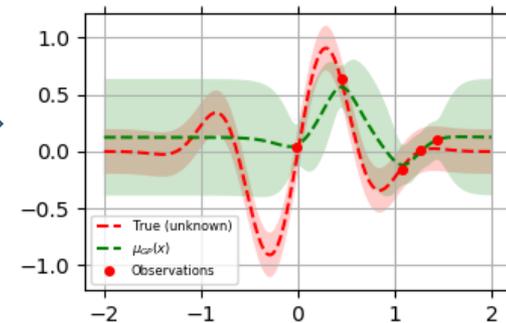
EasyFlinkCEP



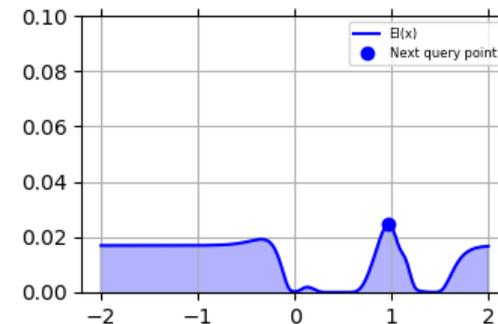
EasyFlinkCEP Benchmarking (2/2)



BO model

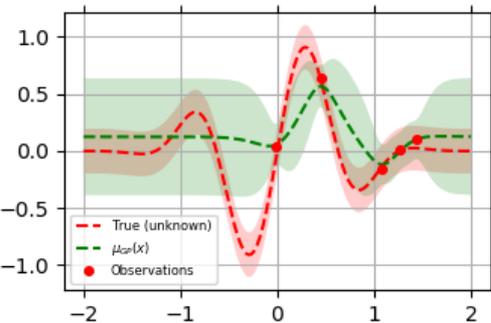


Acquisition Function



EasyFlinkCEP Optimizer

Cost
Modeler



Optimizer pick

What if?
Parallelism = π_i

Performance Oracle
Throughput (T)

RegEx,
Selection Strategy,
Consumption Policy,
Window

Performance Oracle
under parallelism = π

Loop π_i : 1..max Π

$$\begin{aligned} & \max T_{\pi_i} / T_{\pi_k}, k < i \\ & \text{s.t. } \pi_i \in \mathbb{N}, \\ & 1 \leq \pi_i \leq \max \Pi \end{aligned}$$

Good Practices

- D_{init} proportional to 10% of all (contiguity, consumption, window, parallelism) combinations
- Use Lower Confidence Bound (LCB) as the acquisition function (higher accuracy – fewer BO iterations)
- Bucketize window value ranges of interest
- Running about 30% of all possible jobs for a few minutes suffices
- Add input distribution in the X axis if streams are highly volatile

Outline

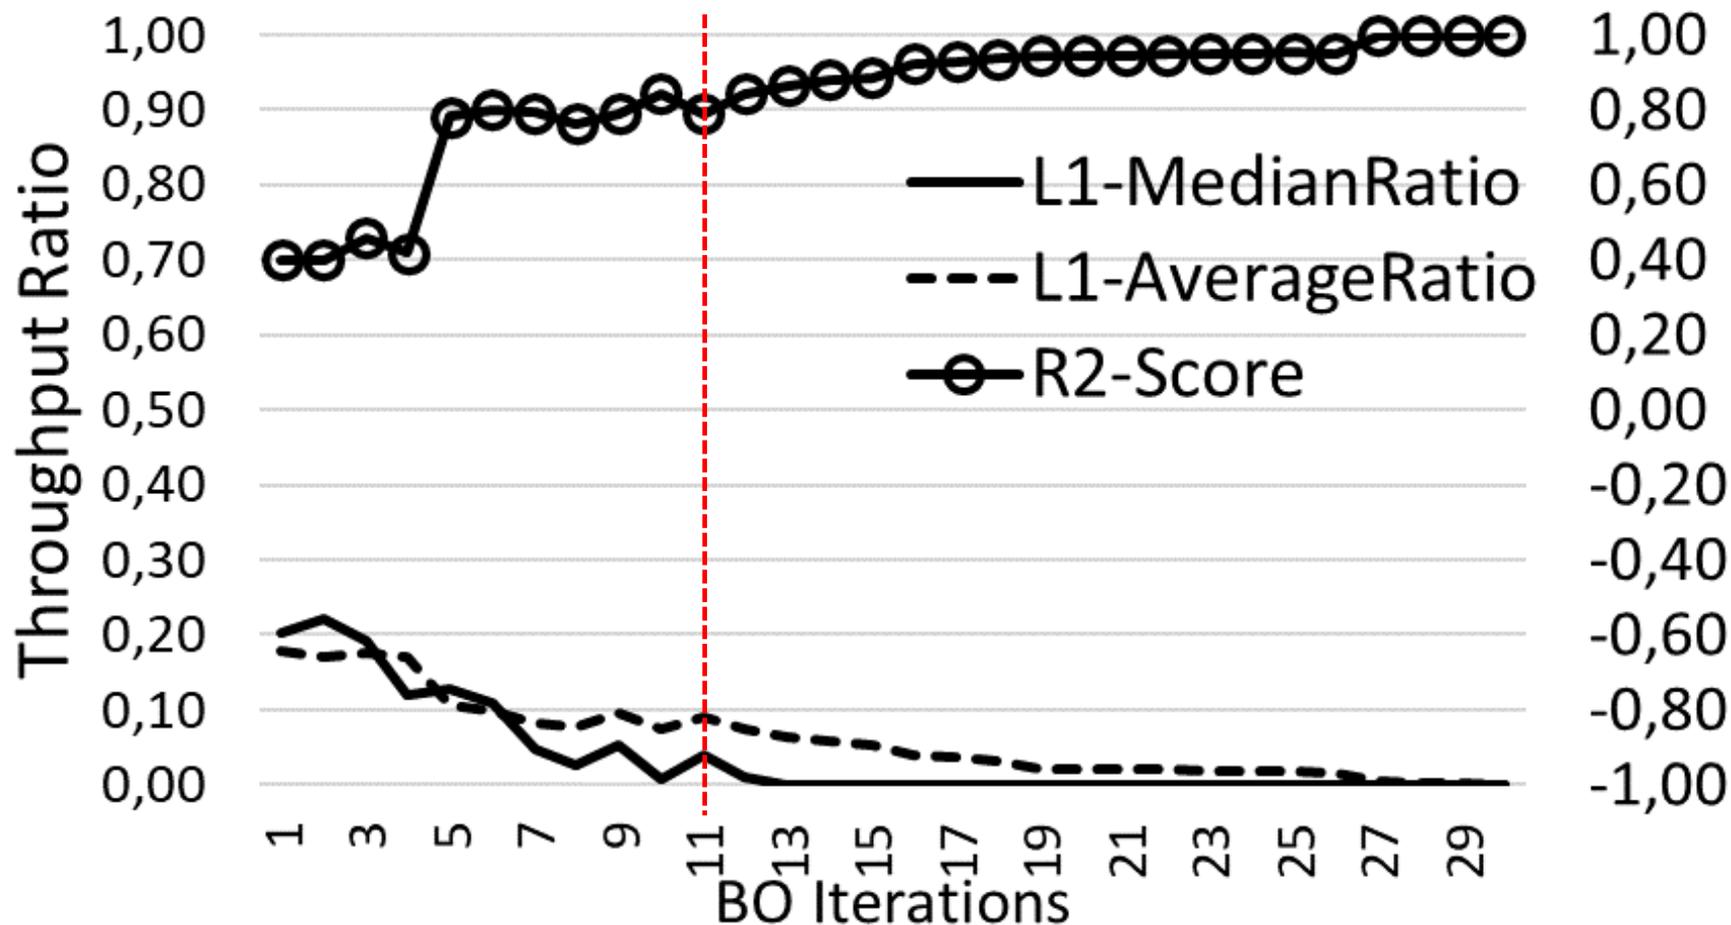
- Introduction (Streams & Apache Flink)
- CEP vs Traditional Streams
- FlinkCEP
 - Basics
 - Adoption barriers
- EasyFlinkCEP
 - EasyFlinkCEP Operator
 - EasyFlinkCEP Optimizer
- What we achieved
- Conclusions & Future Work

Experimental Data & Setup

- What we evaluate
 - Accurate Cost modeler?
 - ...with Few benchmarks required?
 - Right Optimizer Choices?
- Datasets
 - 8 stock streams w 2M events each
 - RegEx: $ab\{1, 3\}c|d$ varying selection/consumption params
 - Window of 128 events/stream
- Cluster Setup
 - VM with 16 cores and 16GB of main memory running Ubuntu 18.04

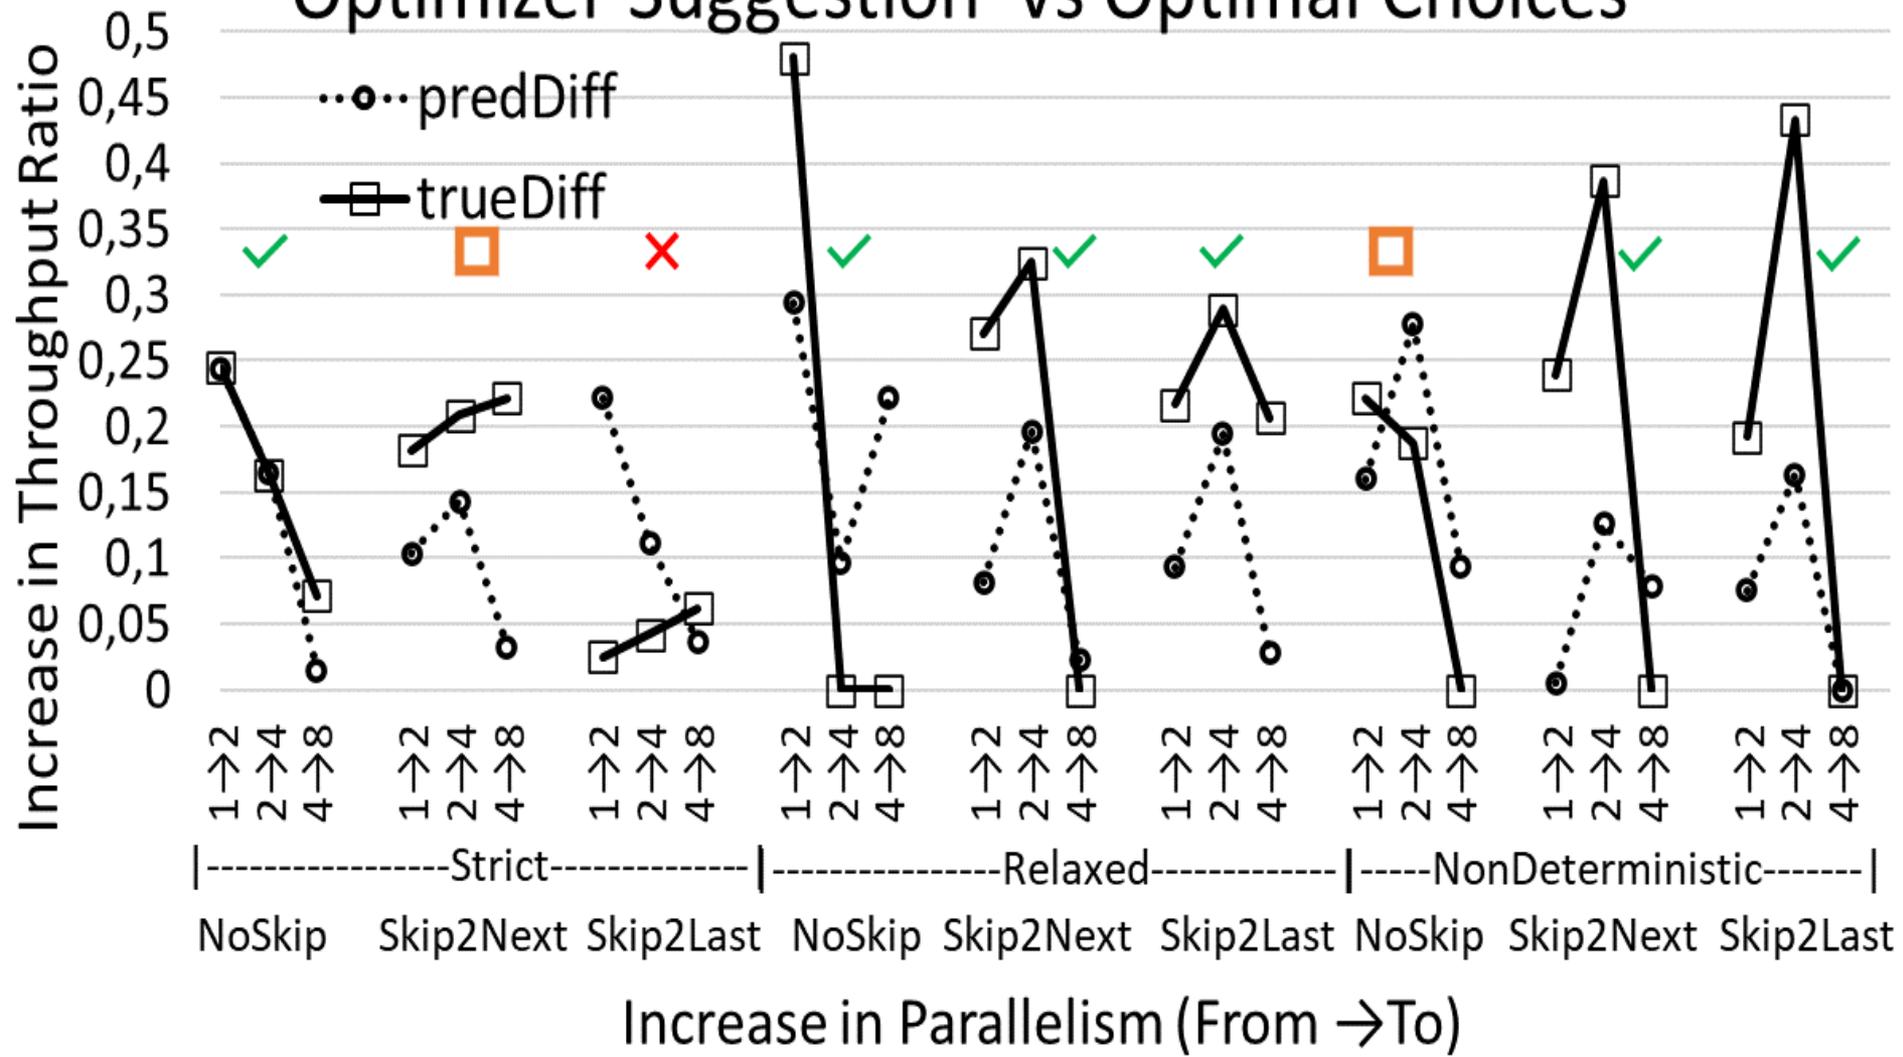
Experimental Results

Throughput Prediction Accuracy vs BO Iterations



Experimental Results

Optimizer Suggestion vs Optimal Choices



Increase in Parallelism (From → To)

Outline

- Introduction (Streams & Apache Flink)
- CEP vs Traditional Streams
- FlinkCEP
 - Basics
 - Adoption barriers
- EasyFlinkCEP
 - EasyFlinkCEP Operator
 - EasyFlinkCEP Optimizer
- What we achieved
- Conclusions & Future Work

Conclusions & Future Work

- EasyFlinkCEP
 - Putting event analysts back to business
 - No coding required
 - No cluster administration decisions
 - 1st parallel CEP Optimizer
- Future Steps:
 - Better FlinkCEP Language Support & Event Hierarchies
 - More complex selection strategy combinations
 - More complex consumption policy combinations
 - Harder to properly define windows
 - Runtime Adaptation

Few References

- EasyFlinkCEP: Big Event Data Analytics for Everyone
Nikos Giatrakos, Eleni Kougioumtzi, Antonios Kontaxakis,
Antonios Deligiannakis, Yannis Kotidis:. CIKM 2021
- Complex event recognition in the Big Data era: a survey
Nikos Giatrakos, Elias Alevizos, Alexander Artikis, Antonios
Deligiannakis, Minos Garofalakis:. VLDB J. 29(1): 313-352 (2020)

<http://infore-project.eu/>

EasyFlinkCEP:
Big Event Data Analytics for Everyone

Thank you!

ngiatrakos@athenarc.gr
ngiatrakos@softnet.tuc.gr