Spreadsheets for Stream Processing with Unbounded Windows and Partitions

Conference on Distributed Event-Based Systems (DEBS), June 2016

Martin Hirzel, Rodric Rabbah, Philippe Suter, Olivier Tardieu, Mandana Vaziri

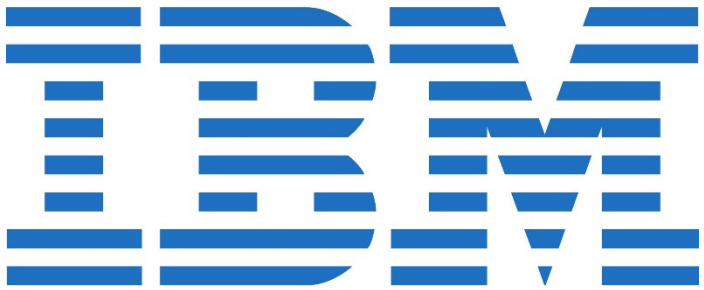

Continuous data streams arise in many different domains.

Domain	Published SPL / InfoSphere Streams customer use case
Telco 	Bouillet et al. *"Experience report: Processing 6 billion CDRs/day: From research to production"*, DEBS 2012
Medical 	Sow et al. *"Real-time analysis for short-term prognosis in intensive care"*, IBM Journal of R&D 2012
Science 	Biem et al. *"A streaming approach to radio astronomy imaging"*, ICASSP 2010
Finance 	Park et al. *"Evaluation of a high-volume, low-latency market data processing system implemented with IBM middleware"*, SP&E 2012
Energy 	
Security 	

Immediate Insights are More Valuable than Delayed Insights.

Domain

Telco 

Medical  **Notice issue before it becomes acute → save lives.**

Science 

Finance  **Grasp opportunity before it disappears → earn money.**

Energy 

Security 

There may be too much data to store to disk for offline analysis.

Domain

Telco 

Medical 

Science 

Finance 

Energy 

Security 

A grid of many antennae and many beams per antenna produces lots of data with low information content.

Stream Processing

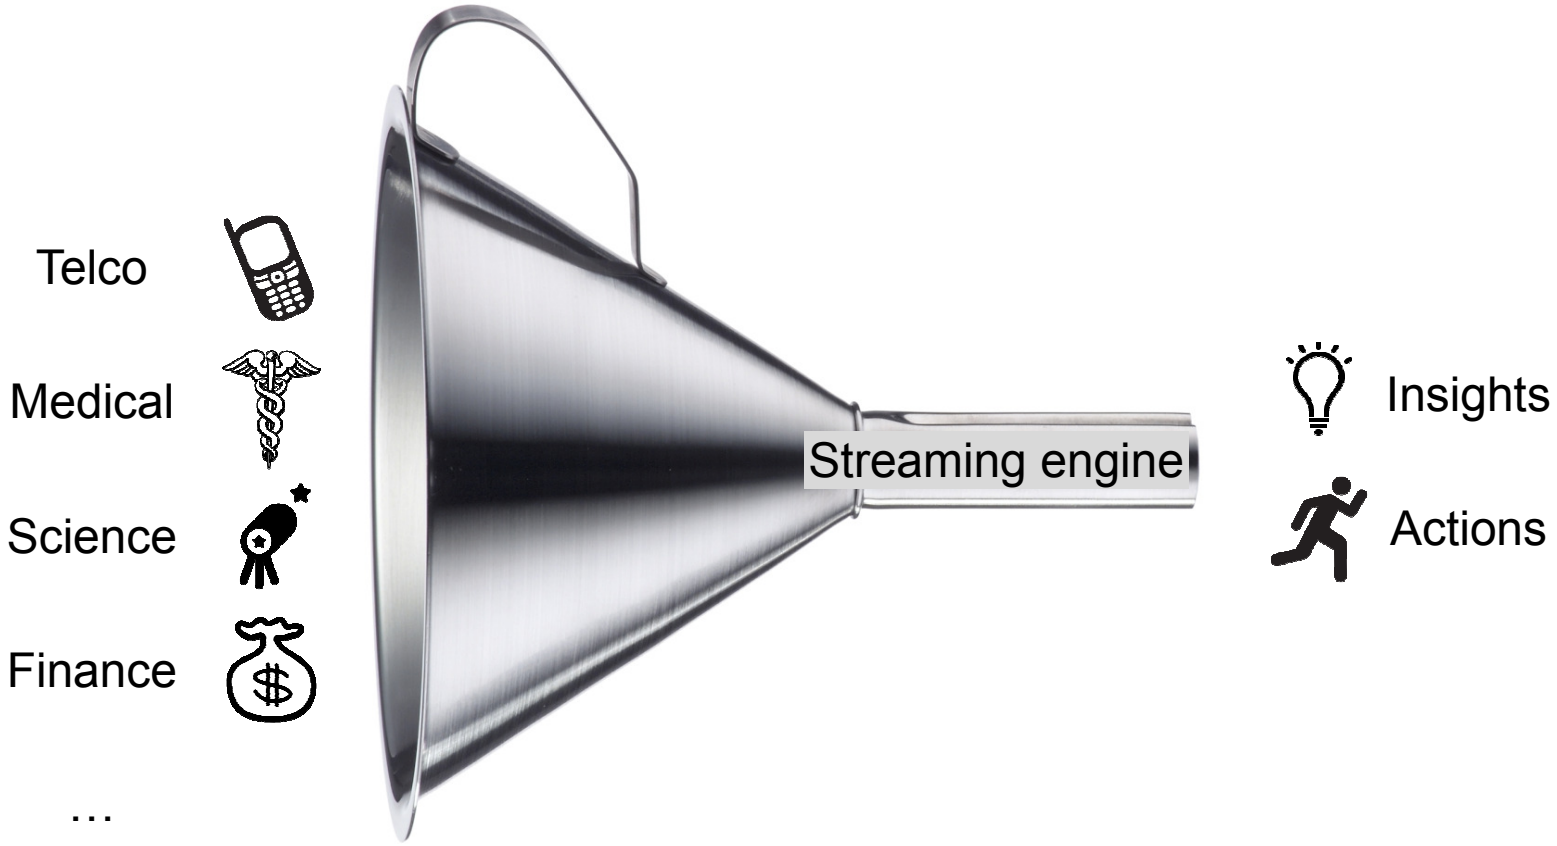

Performance Challenge

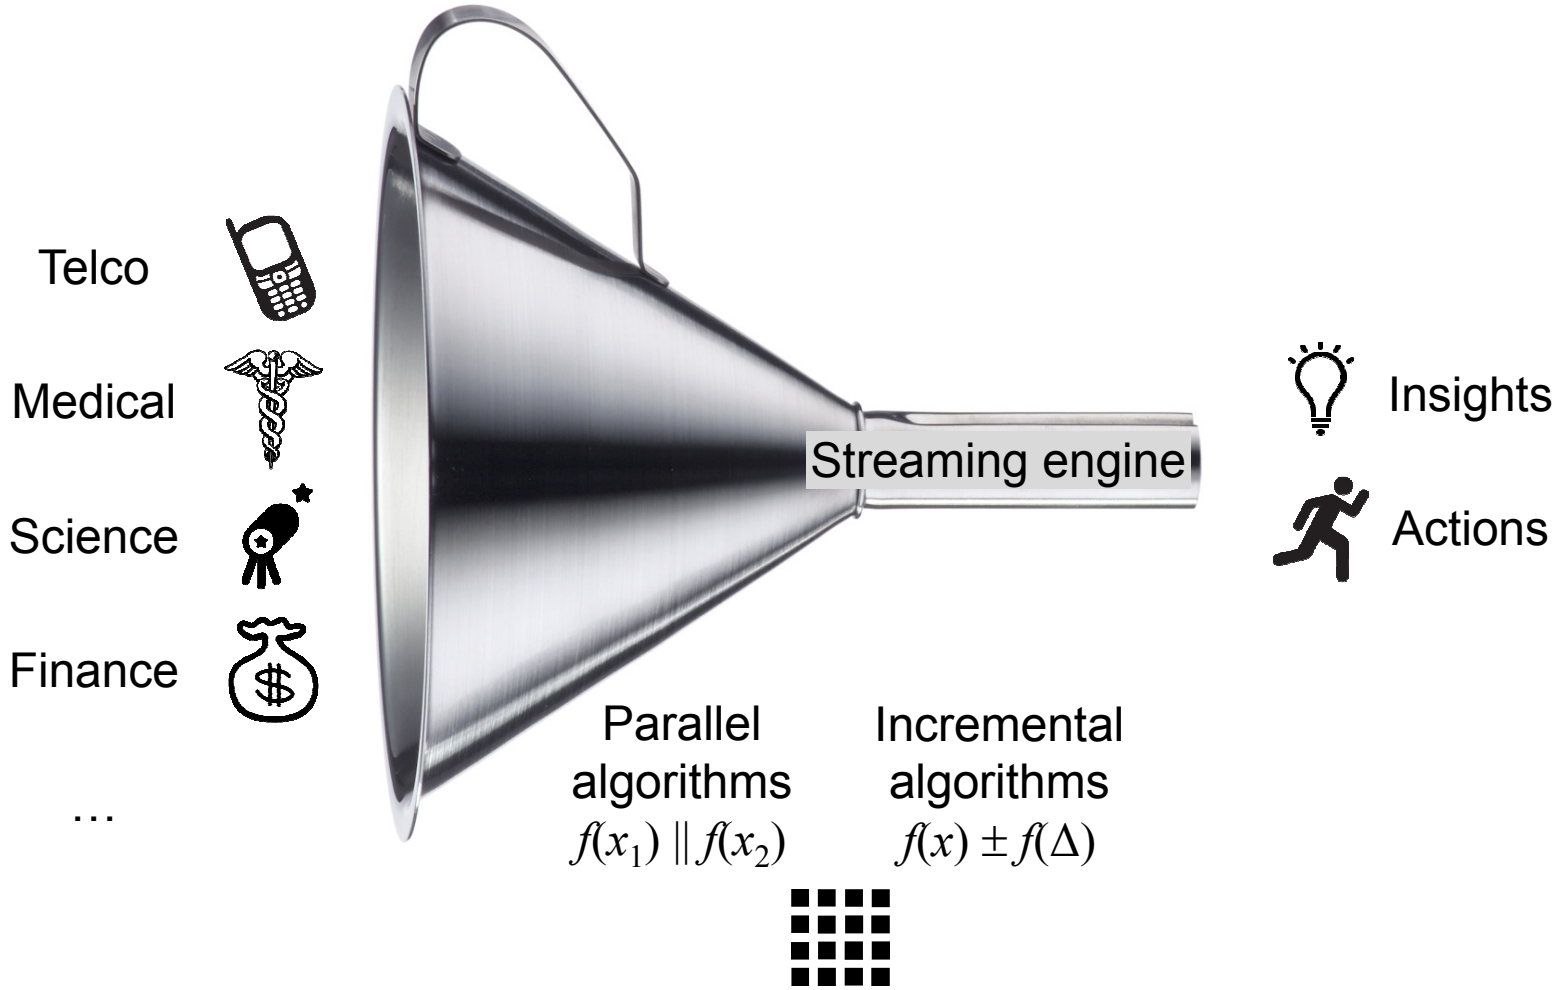

Usability Challenge

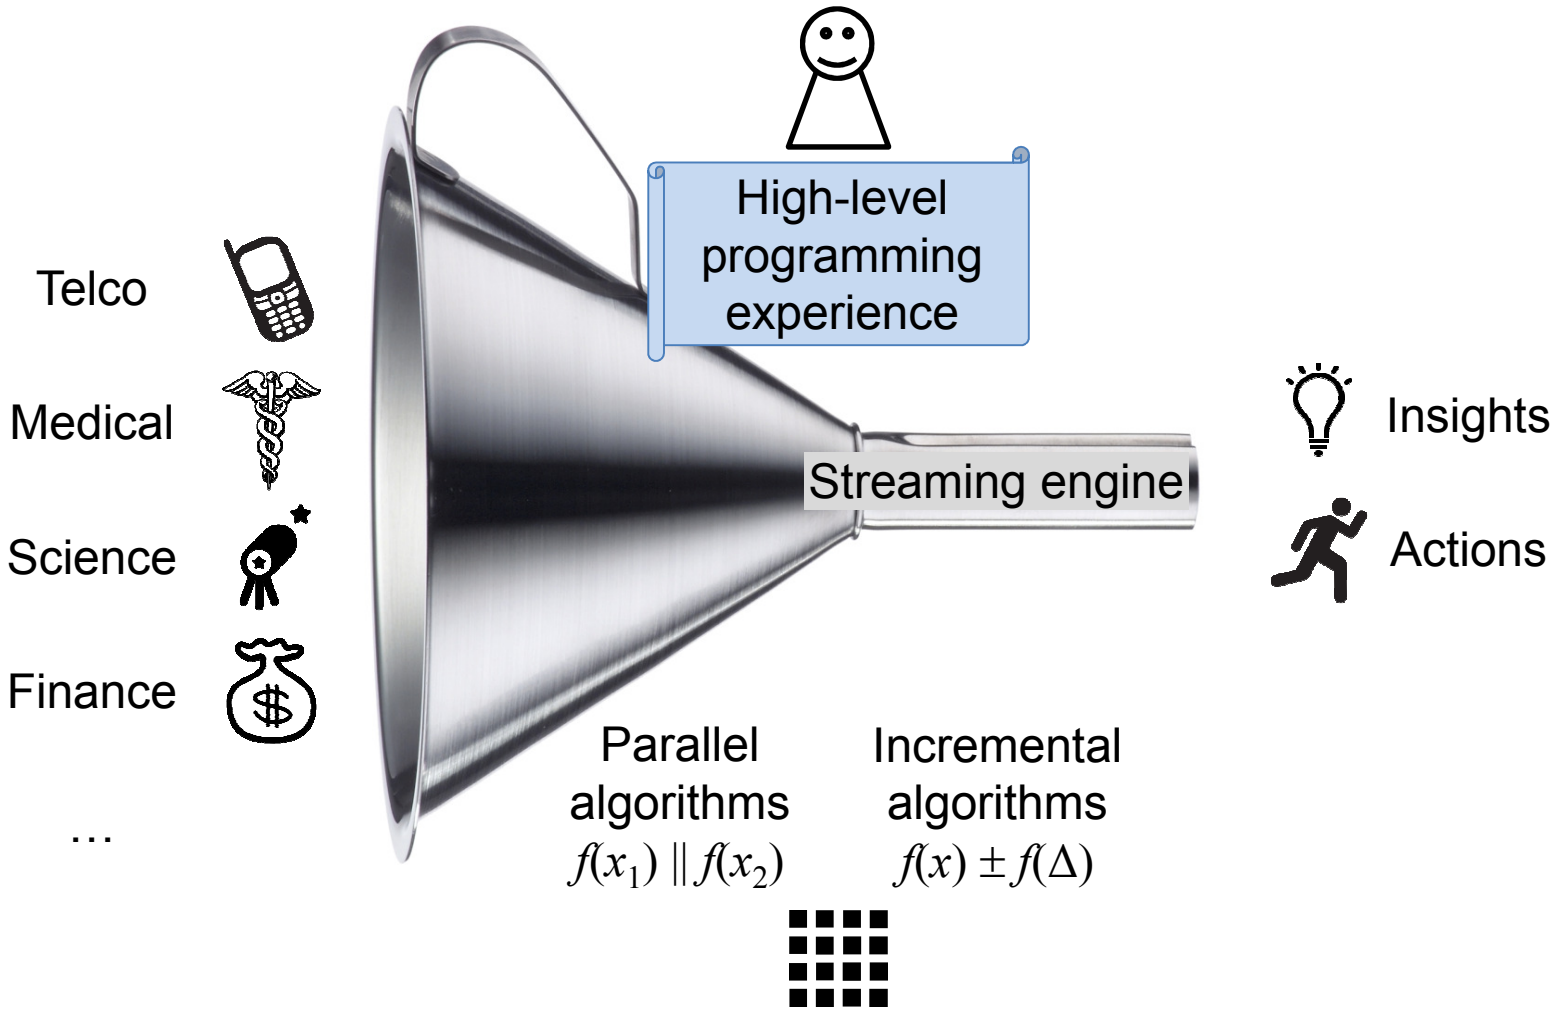

Why Spreadsheets?

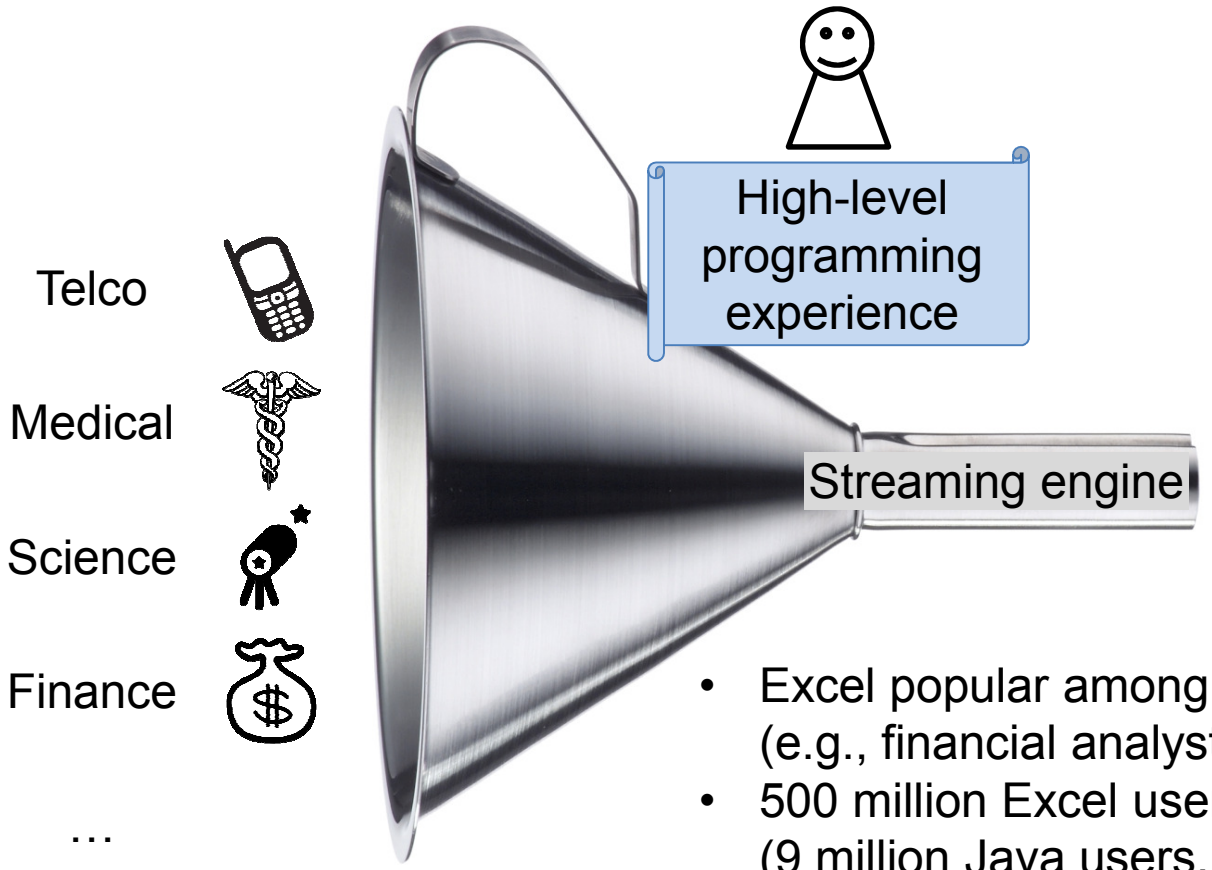

- Excel popular among non-programmers (e.g., financial analysts)
- 500 million Excel users (9 million Java users, less for other languages)
- Reactive model (spreadsheets and most streaming operators)
- 2-dimensional view (↓ tuples, → attributes)

Bargain Finder Example

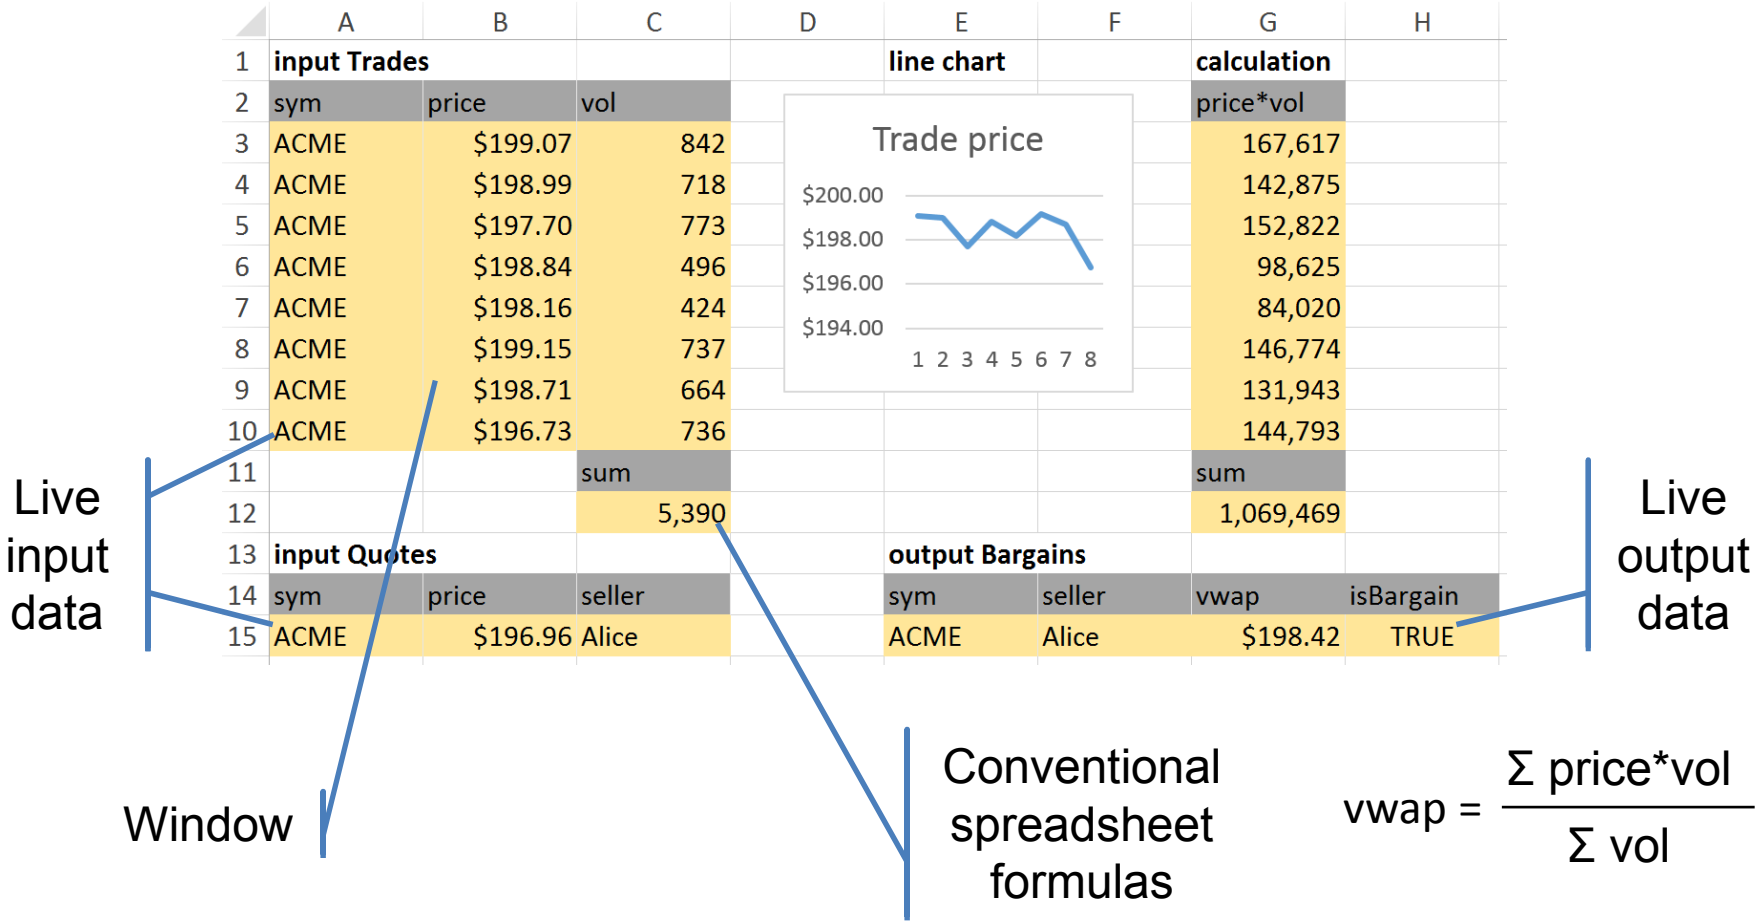

Vaziri et al. *"Stream Processing with a Spreadsheet"*, ECOOP 2014 (Distinguished Paper Award)

Limitations of Two Dimensions

Columns

Rows

	A	B	C	D	E	F	G	H
1	**input Trades**				**line chart**		**calculation**	
2	sym	price	vol				price*vol	
3	ACME	$199.07	842				167,617	
4	ACME	$198.99	718				142,875	
5	ACME	$197.70	773				152,822	
6	ACME	$198.84	496				98,625	
7	ACME	$198.16	424				84,020	
8	ACME	$199.15	737				146,774	
9	ACME	$198.71	664				131,943	
10	ACME	$196.73	736				144,793	
11			sum				sum	
12			5,390				1,069,469	
13	**input Quotes**				**output Bargains**			
14	sym	price	seller		sym	seller	vwap	isBargain
15	ACME	$196.96	Alice		ACME	Alice	$198.42	TRUE

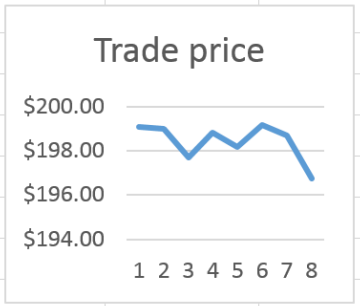

What about other stock symbols?

What about time-based windows?

Variable-Sized Windows

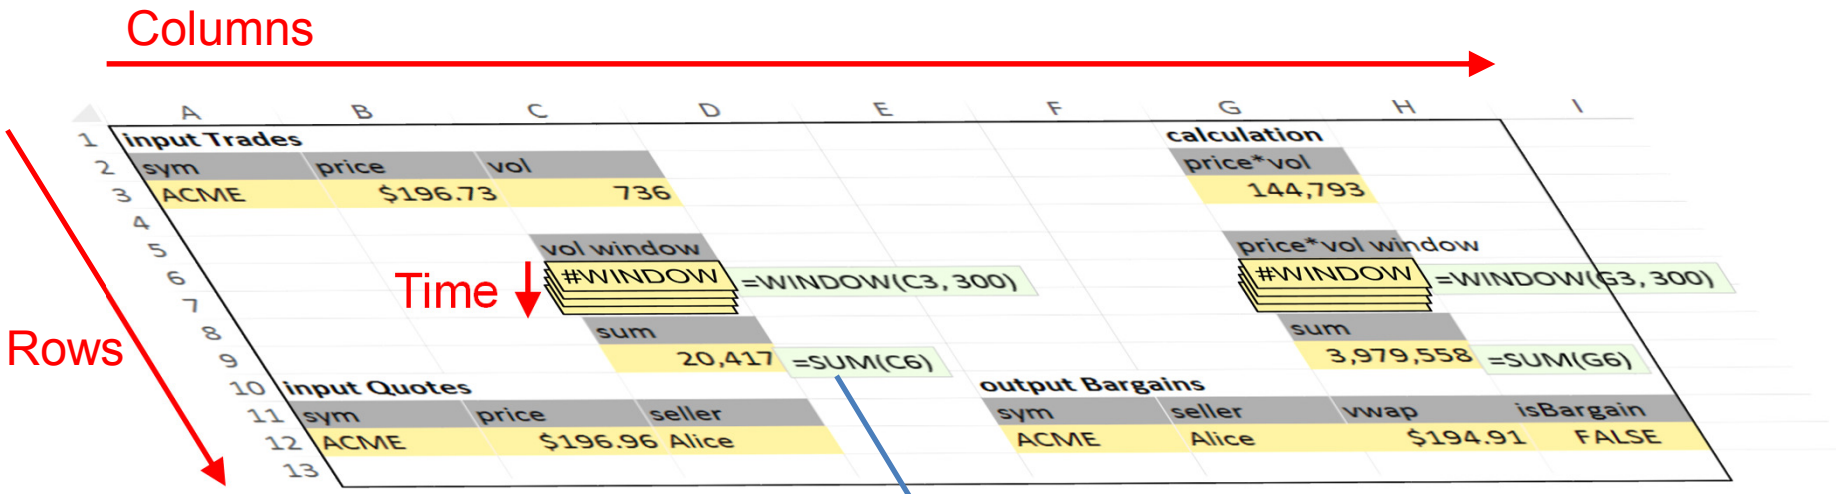

Variable-sized windows:
occupy one cell,
aggregate all elements

Shared Window Aggregation

Is the most recent value in the window unique?

COUNTIF(w, INDEX(w, COUNT(w))) = 1

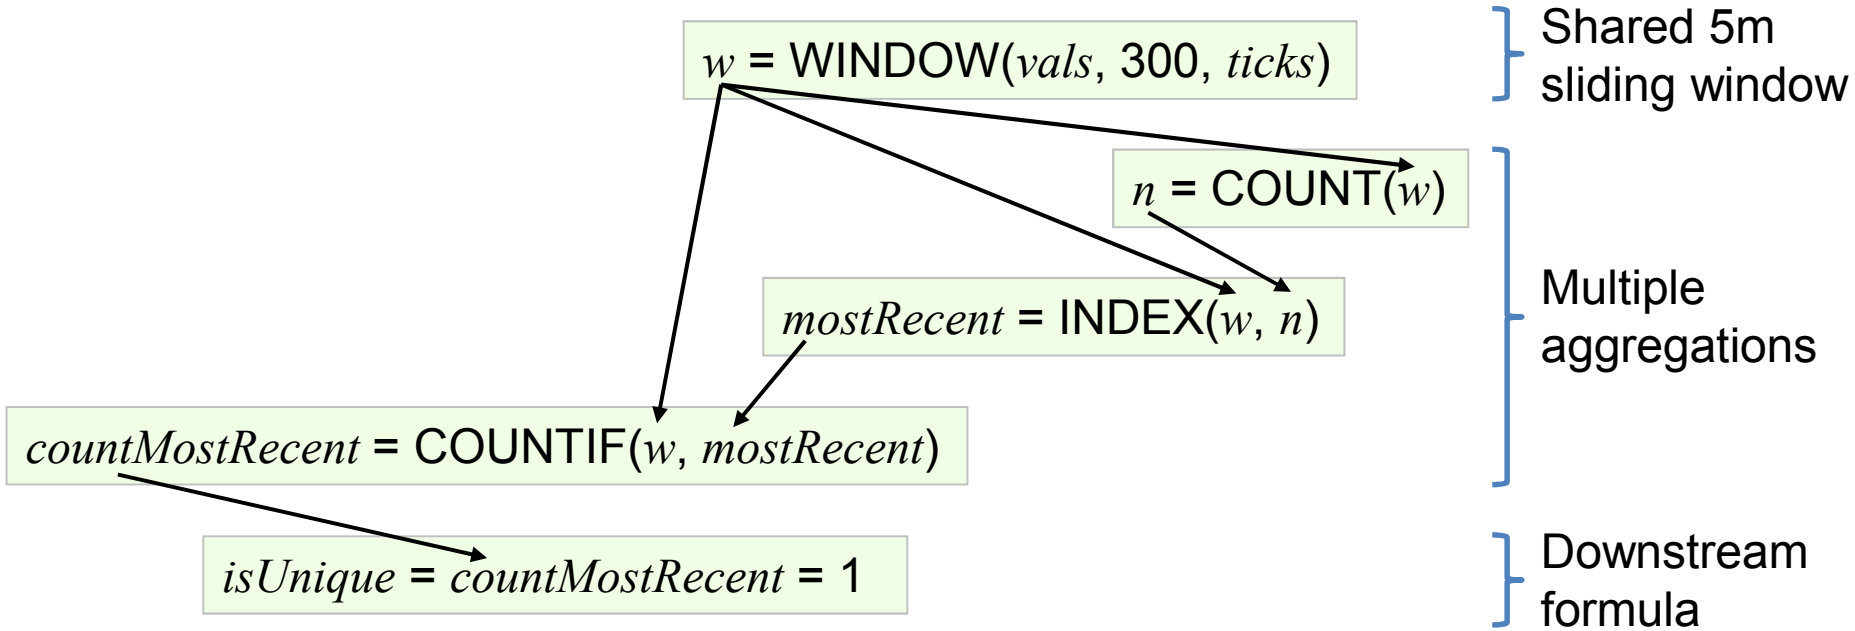

Incremental Window Aggregation

Function	Description	Data structure	Time, Space
SUM(w)	Total of the numbers in w.	Float	$O(1)$, $O(1)$
AVERAGE(w)	Arithmetic mean of the numbers in w.	Two floats	$O(1)$, $O(1)$
COUNT(w)	Number of elements in w with numbers.	Integer	$O(1)$, $O(1)$
COUNTIF(w, v)	Number of elements in w that equal v.	Hash multi-set	$O(1)$, $O(\|w\|)$
INDEX(w, i)	Element of w at index i, where 1 is the oldest.	Resizable circular buffer	$O(1)$, $O(\|w\|)$
MATCH(v, w, m)	Index of element equal to v in w if m is 0 (exact match).	Tree multi-map, integer	$O(\log \|w\|)$, $O(\|w\|)$
LARGE(w, k)	Number in w that is the kth largest, where 1 is the max.	Order statistics tree	$O(\log \|w\|)$, $O(\|w\|)$

Partitioned Virtual Worksheets

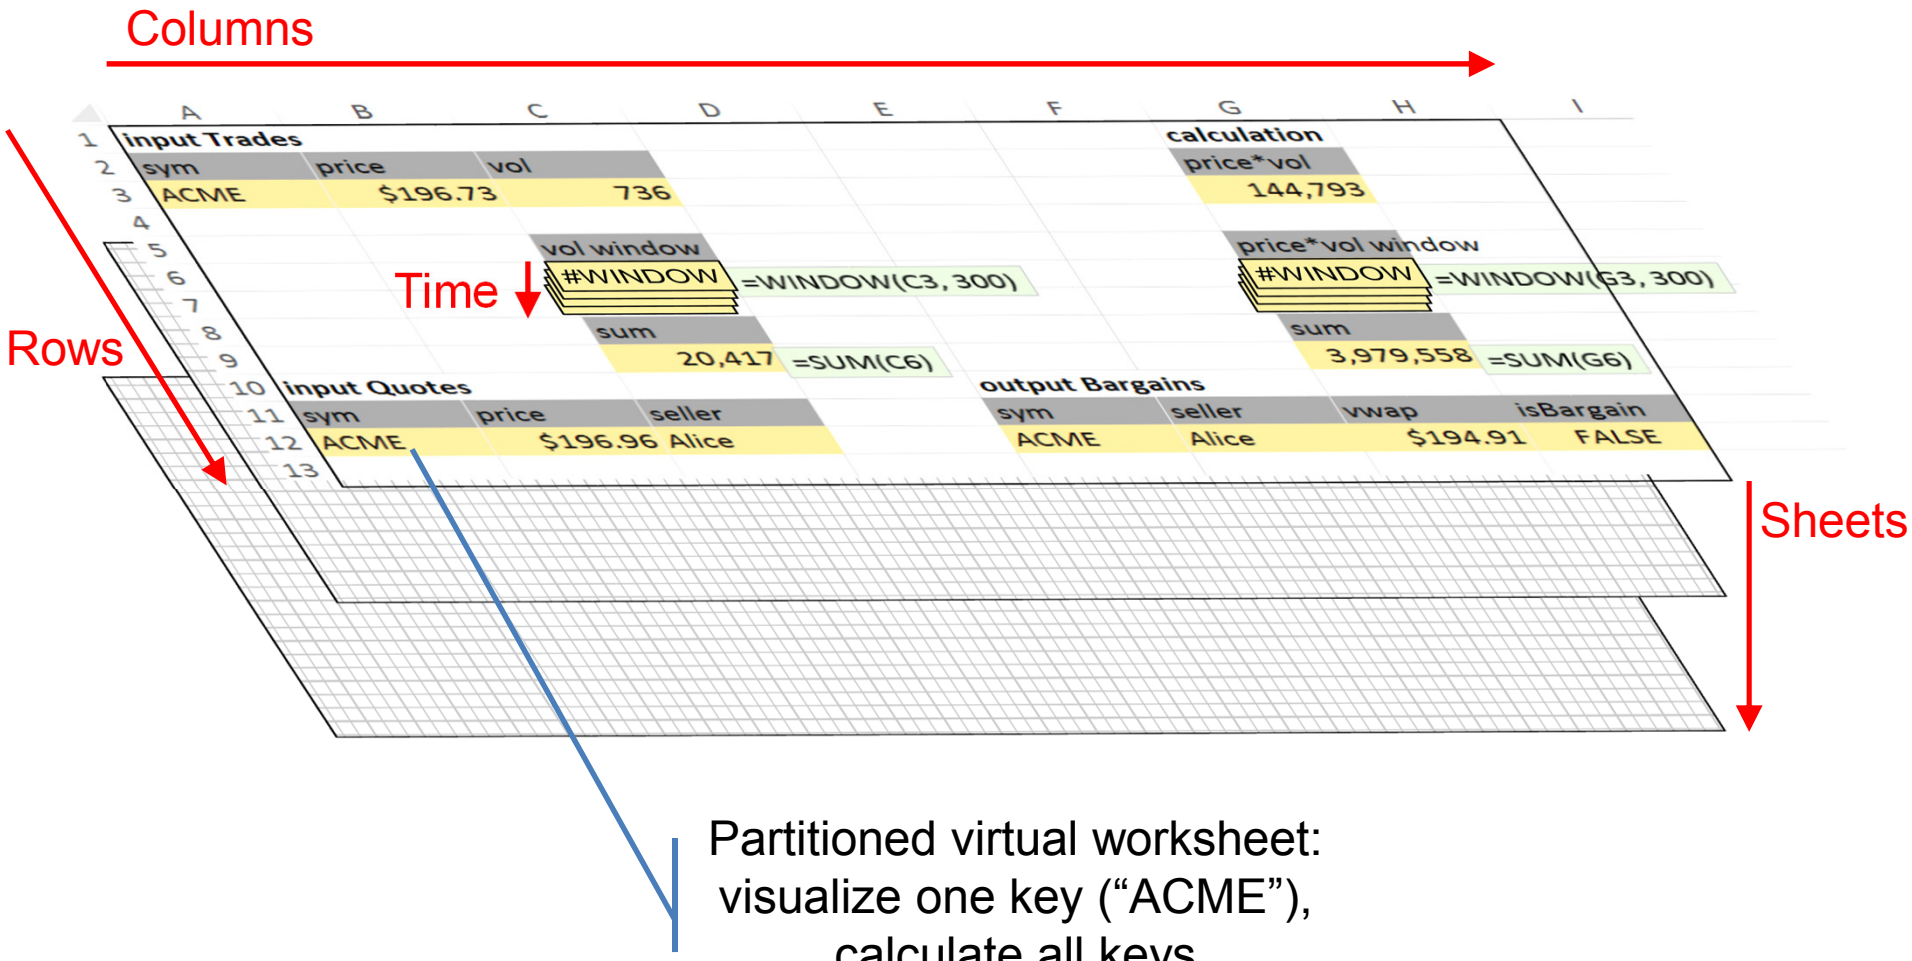

Partitioning Implementation

Spreadsheet specifies calculation for example sym=="ACME" only

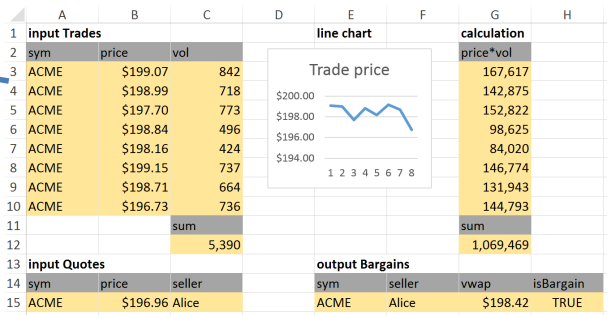

	A	B	C	D	E	F	G	H
1	**input Trades**				**line chart**		**calculation**	
2	sym	price	vol				price*vol	
3	ACME	$199.07	842				167,617	
4	ACME	$198.99	718				142,875	
5	ACME	$197.70	773				152,822	
6	ACME	$198.84	496				98,625	
7	ACME	$198.16	424				84,020	
8	ACME	$199.15	737				146,774	
9	ACME	$198.71	664				131,943	
10	ACME	$196.73	736				144,793	
11			sum				sum	
12			5,390				1,069,469	
13	**input Quotes**				**output Bargains**			
14	sym	price	seller		sym	seller	vwap	isBargain
15	ACME	$196.96	Alice		ACME	Alice	$198.42	TRUE

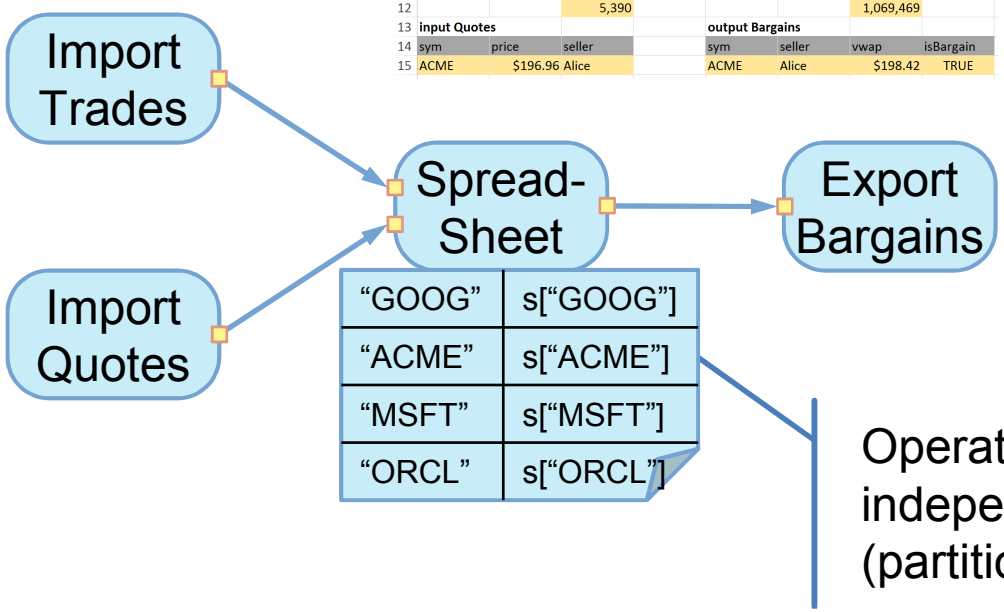

Data Parallelism

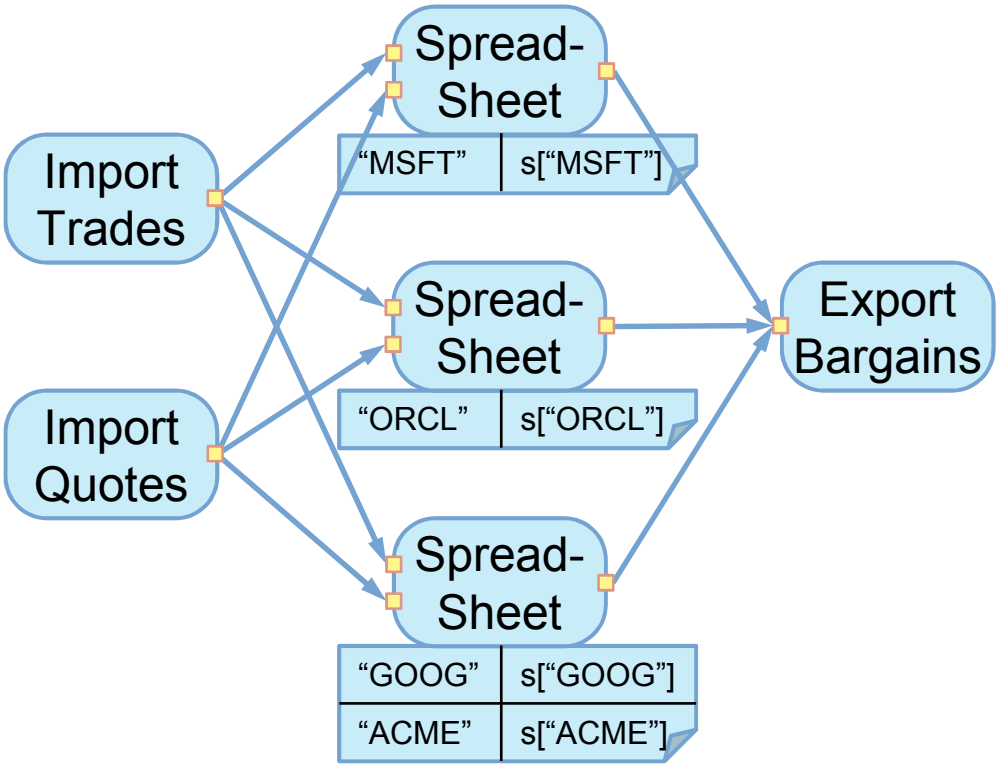


```
@parallel(width=3, partitionBy=[{port=Trades,attributes=[sym]},
                                {port=Quotes,attributes=[sym]}])
```

Operator Configuration from SPL

```
1   stream<Bargain> Bargains = SpreadSheet(Trades; Quotes) {
2    param
3      spreadsheet     : "vwap.xls";
4      inputs          : {
5        A3 =Trades.sym, B3 =Trades.price, C3 =Trades.vol,
6        A12=Quotes.sym, B12=Quotes.price, C12=Quotes.seller };
7      partitionByLHS : Trades.sym;
8      partitionByRHS : Quotes.sym;
9      timeByLHS       : Trades.ts;
10     timeByRHS       : Quotes.ts;
11   output
12     Bargains        :
13       sym =RString("E12"), seller =RString("F12"),
14       vwap=Float64("G12"), bargain=Boolean("H12");
15  }
```

Hirzel et al. *"IBM Streams Processing Language: Analyzing big data in motion"*, IBM Journal of R&D 2013

Compilation

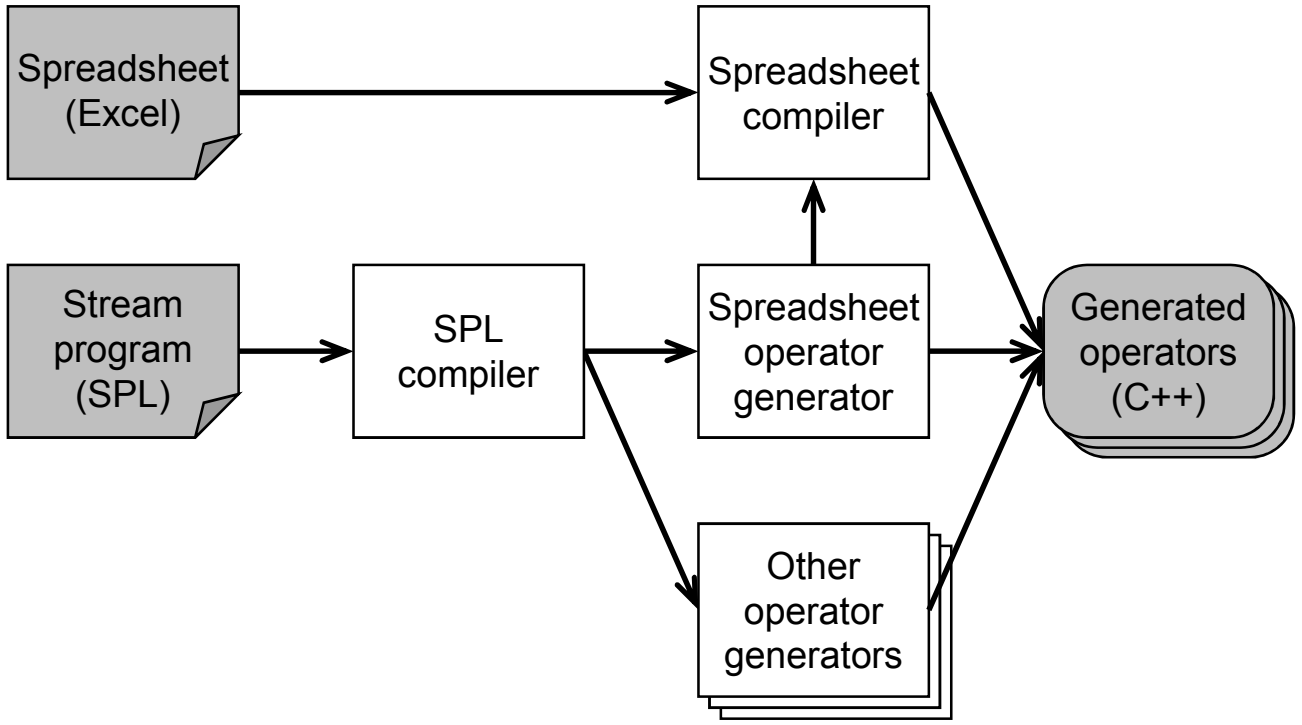

Benchmarks

Benchmark	Cells	Exprs	Windows	Partition
vwap	9	14	2x5m	ticker
twitter	22	36	3x5m	lang
linearroad	20	18	3x30s & 2x5m	segment
average	33	27	2x6	-
kalman	14	21	2x2	target id
tax	21	37	-	-
pong	35	86	-	game id
forecast	43	60	2x6	location

Throughput

Benchmark	Cells	Exprs	Windows	Partition	Ktps	vs. SPL
vwap	9	14	2x5m	ticker	960.6	2.41
twitter	22	36	3x5m	lang	45.9	1.52
linearroad	20	18	3x30s & 2x5m	segment	964.3	1.21
average	33	27	2x6	-	3,937.0	0.49
kalman	14	21	2x2	target id	3,816.8	0.47
tax	21	37	-	-	1,383.1	0.28
pong	35	86	-	game id	480.8	0.12
forecast	43	60	2x6	location	913.2	0.12

Faster

Slower

Parallel Performance

Benchmark	Cells	Exprs	Windows	Partition	Ktps	vs. SPL
vwap	9	14	2x5m	ticker	960.6	2.41
twitter	22	36	3x5m	lang	45.9	1.52
linearroad	20	18	3x30s & 2x5m	segment	964.3	1.21
average	33	27	2x6	-	3,937.0	0.49
kalman	14	21	2x2	target id	3,816.8	0.47
tax	21	37	-	-	1,383.1	0.28
pong	35	86	-	game id	480.8	0.12
forecast	43	60	2x6	location	913.2	0.12

Faster

Slower

mandelbrot

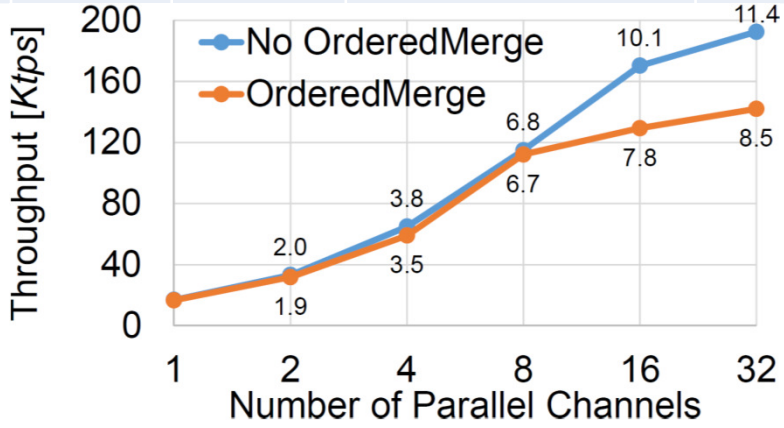

Stream Processing for the Masses

Telco

Medical

Science

Finance

…

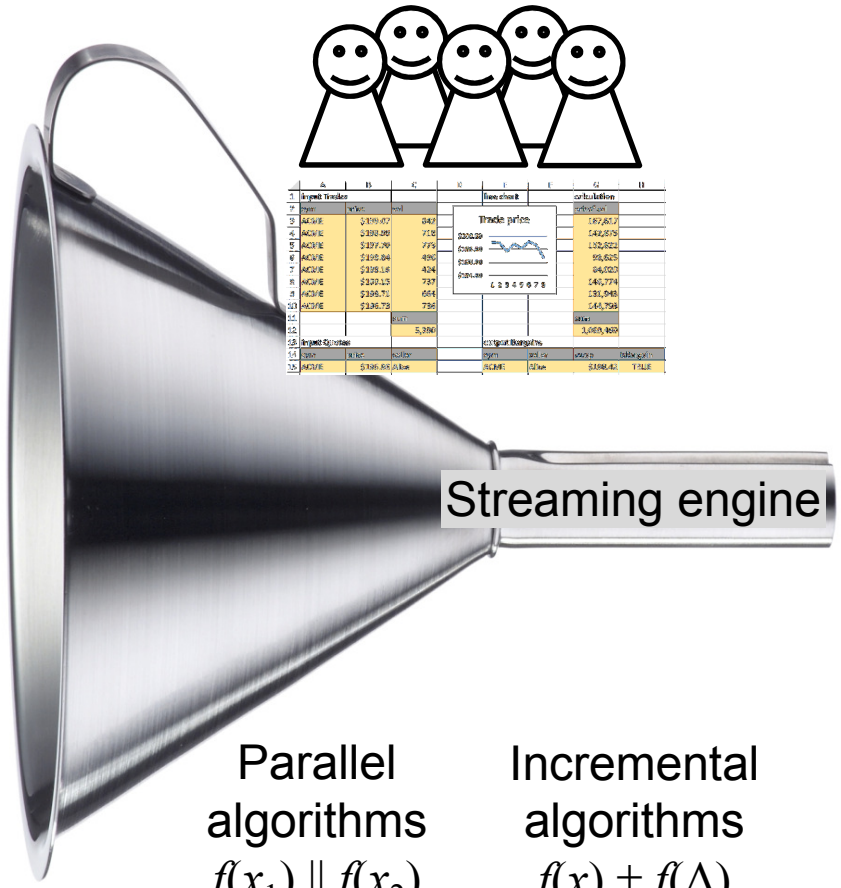

Insights

Actions

Parallel algorithms

$f(x_1) \parallel f(x_2)$

Incremental algorithms

$f(x) \pm f(\Delta)$