An Experiment-Driven Performance Model of Stream Processing Operators in Fog Computing Environments

Hamidreza Arkian[1], Guillaume Pierre[1],
Johan Tordsson[2], Erik Elmroth[2]

[1]University of Rennes1/IRISA, France
[2]Elastisys AB, Sweden

SAC'20 - March 30-April 3, 2020 - Brno, Czech Republic

IoT-to-Cloud basic architecture

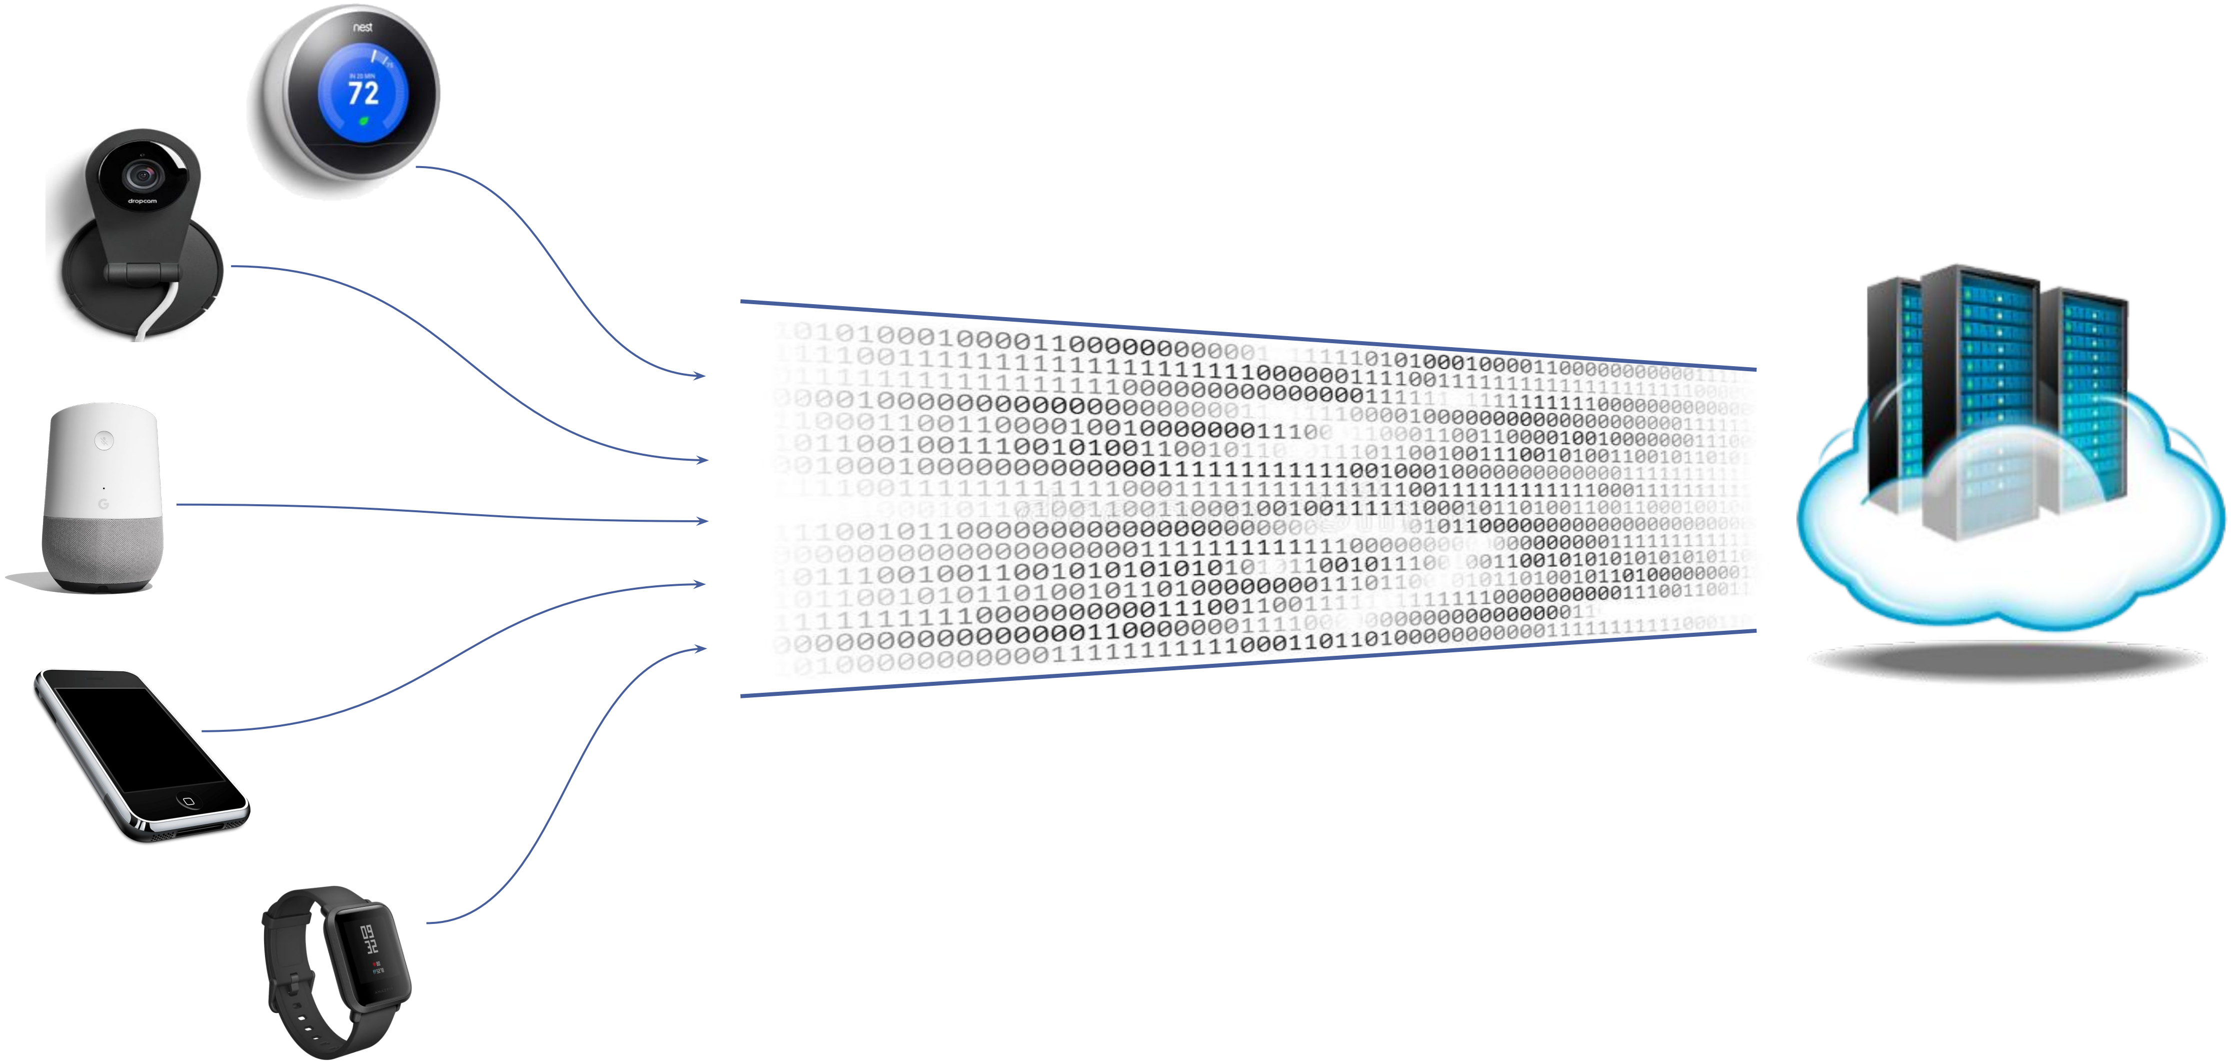

Cloud-based stream processing

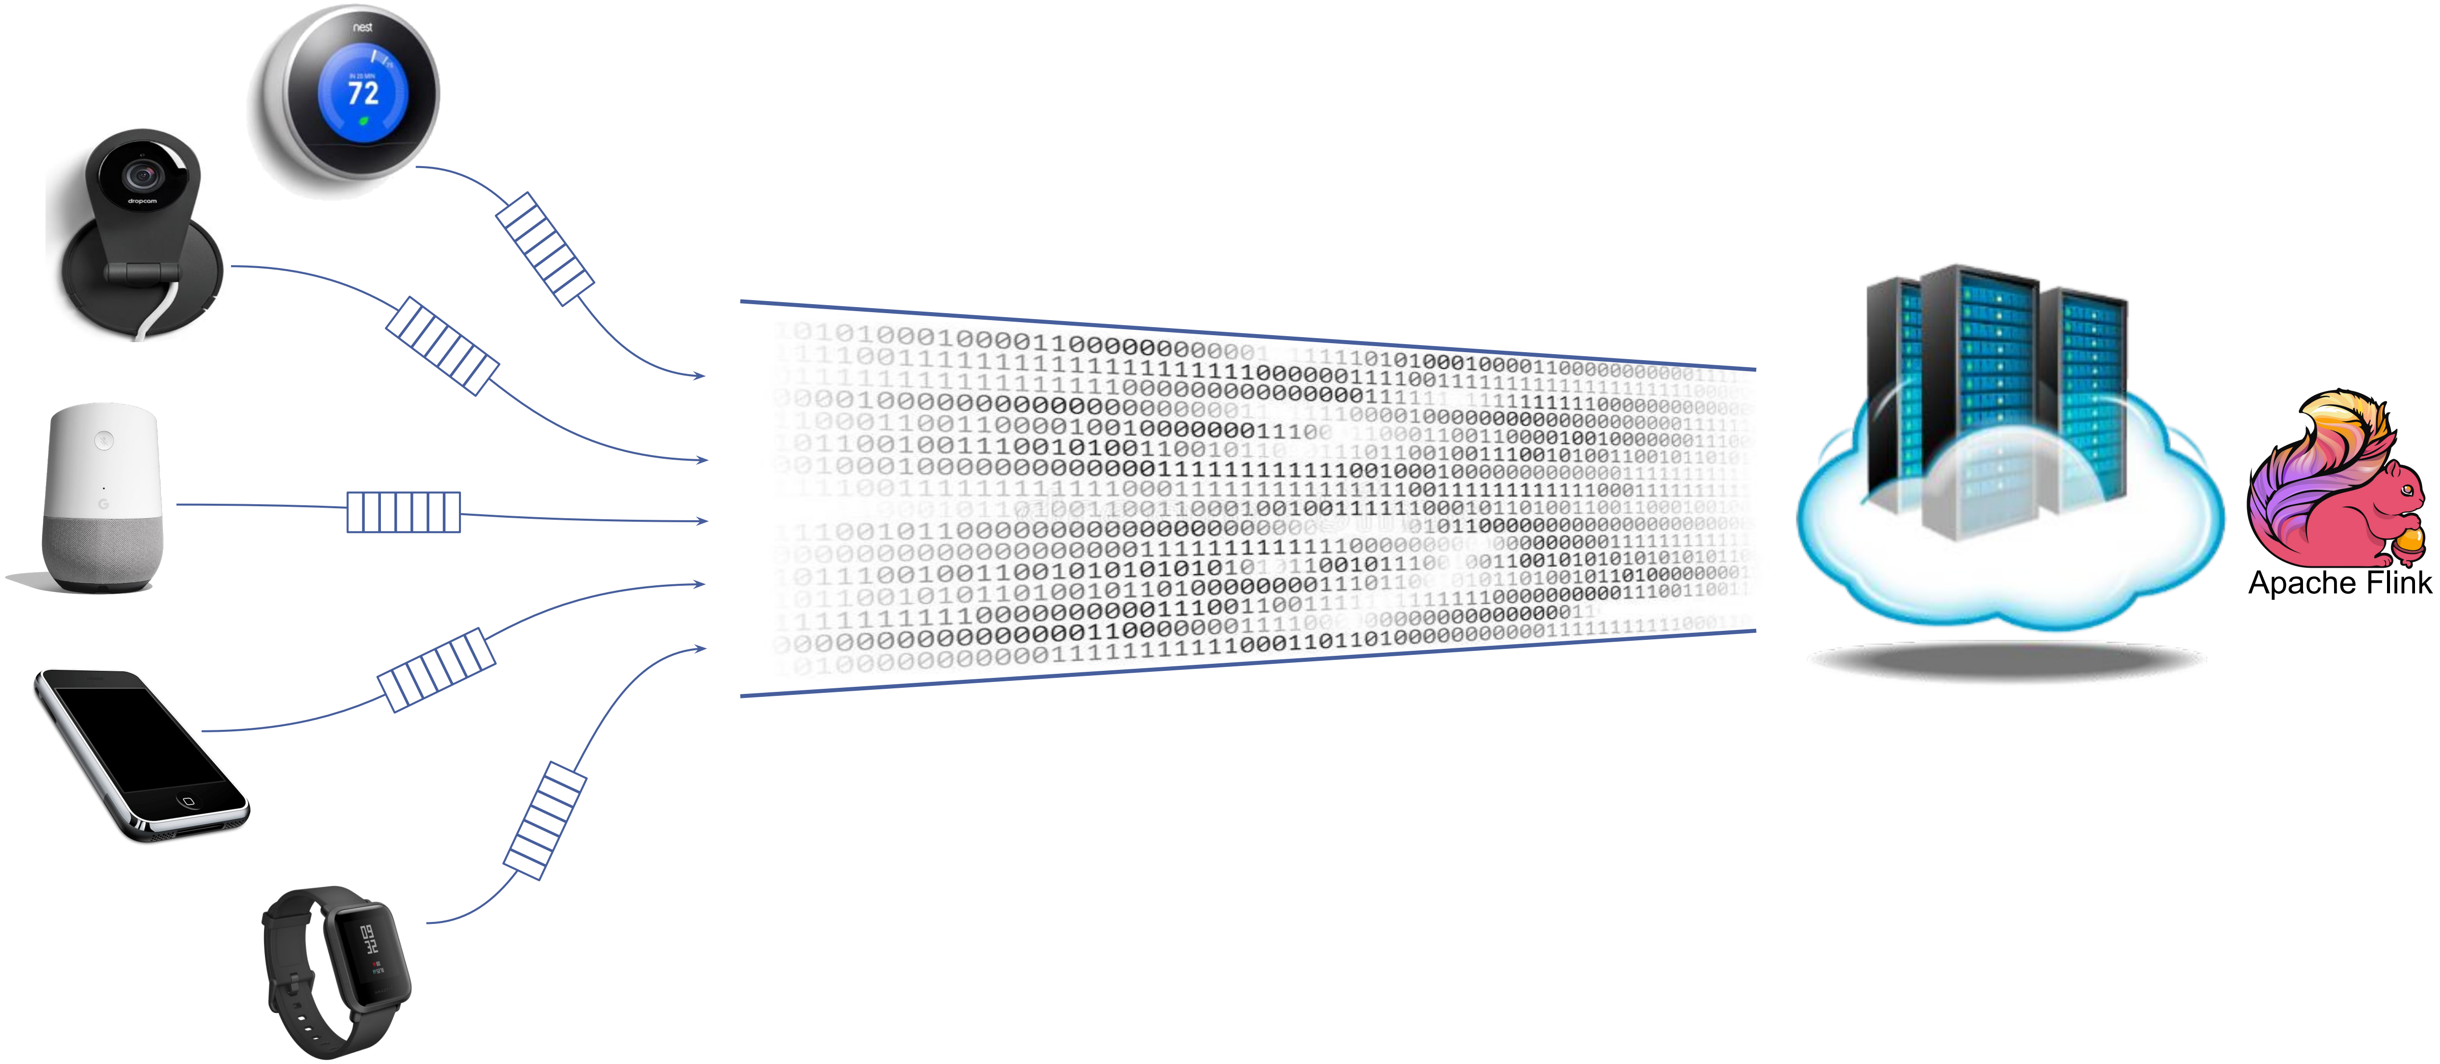

Challenges

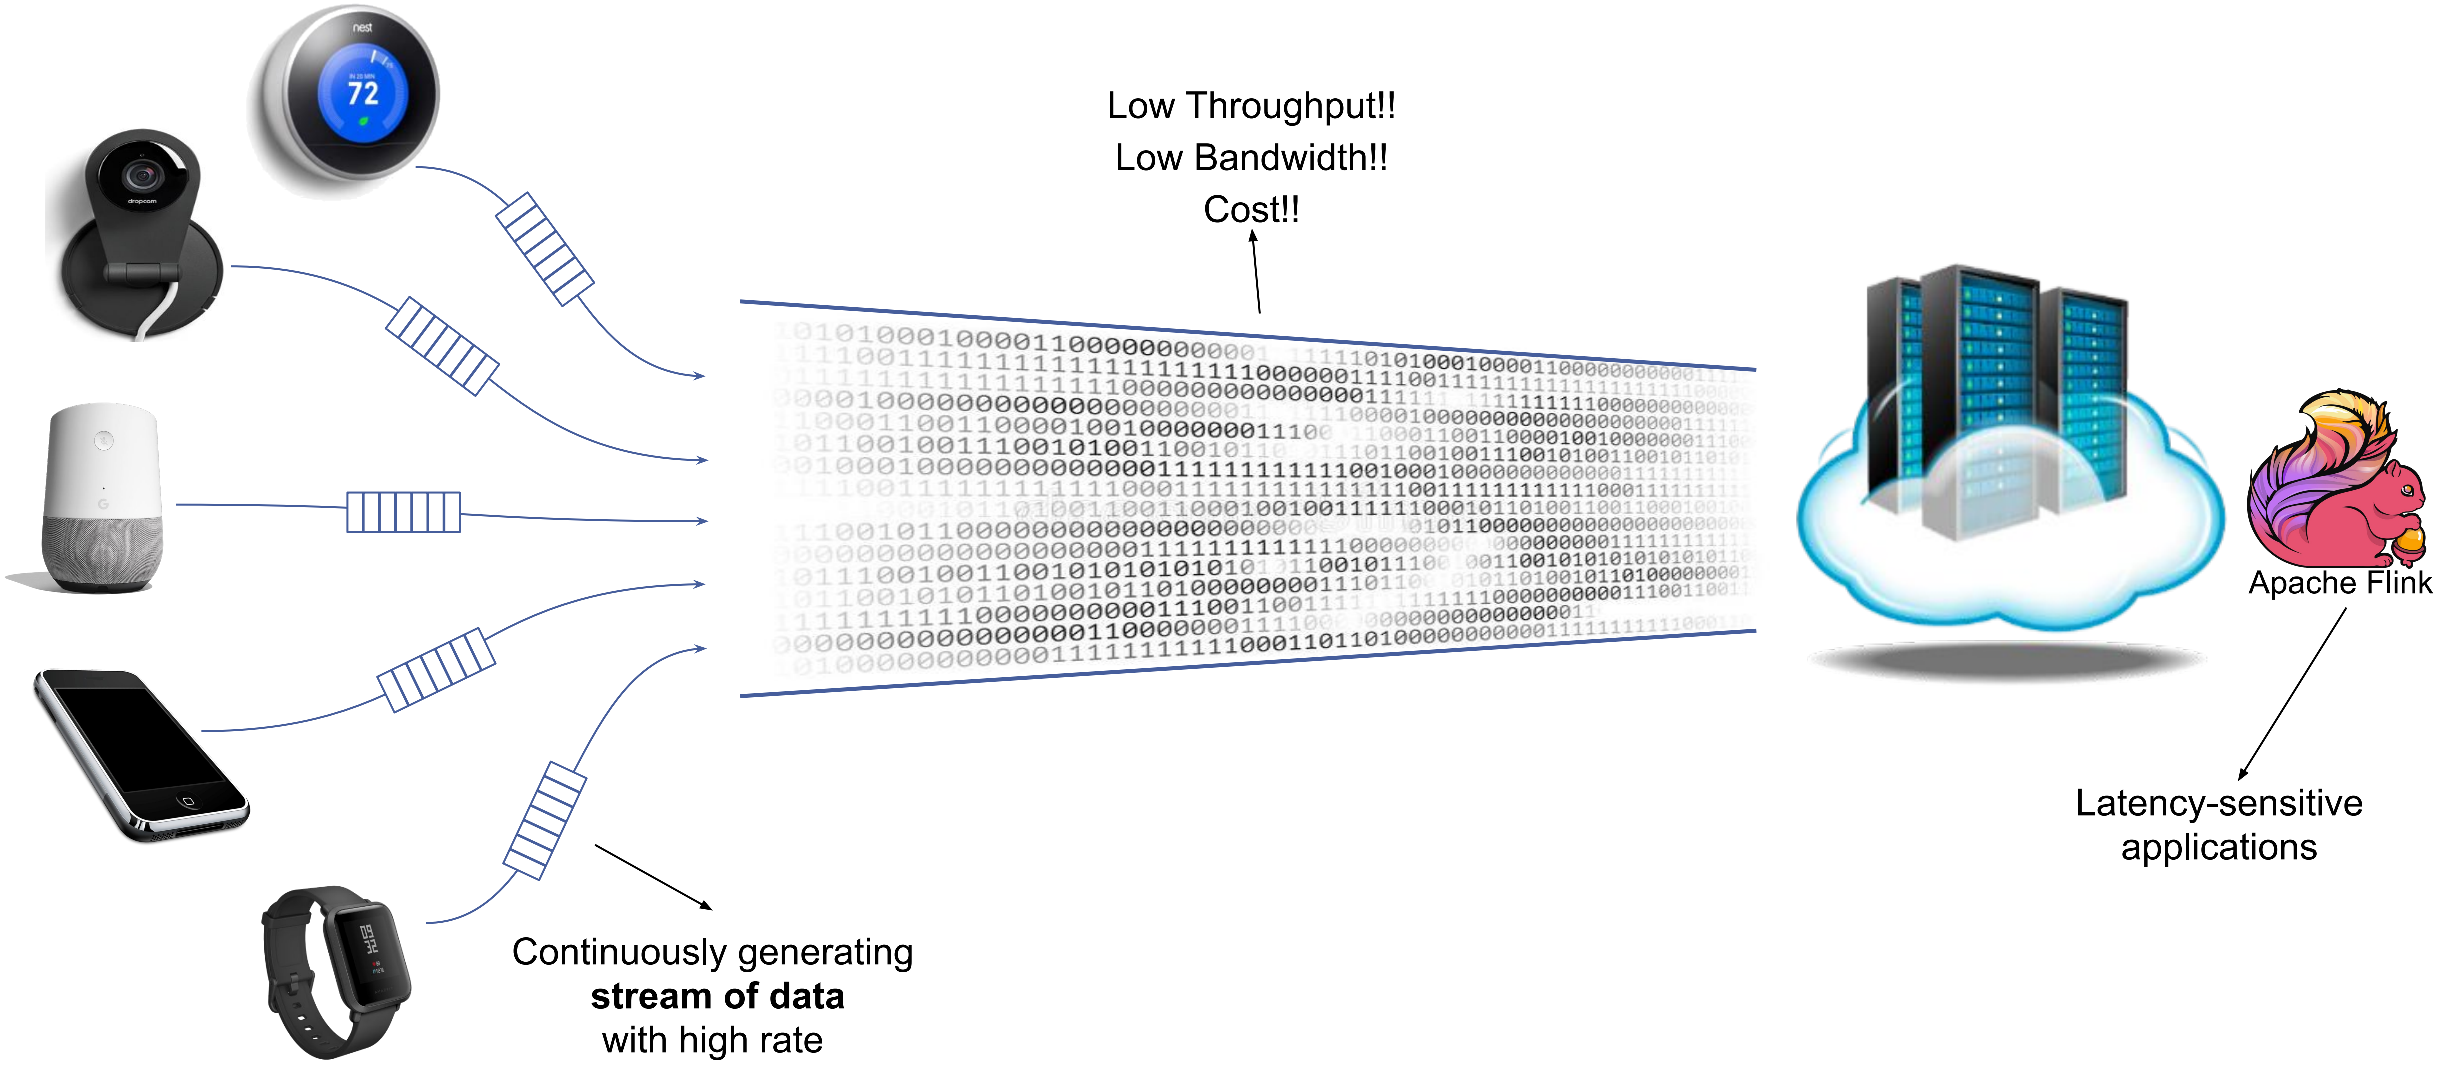

Low Throughput!!
Low Bandwidth!!
Cost!!
Apache Flink
Latency-sensitive applications
Continuously generating **stream of data** with high rate

Fog-based stream processing

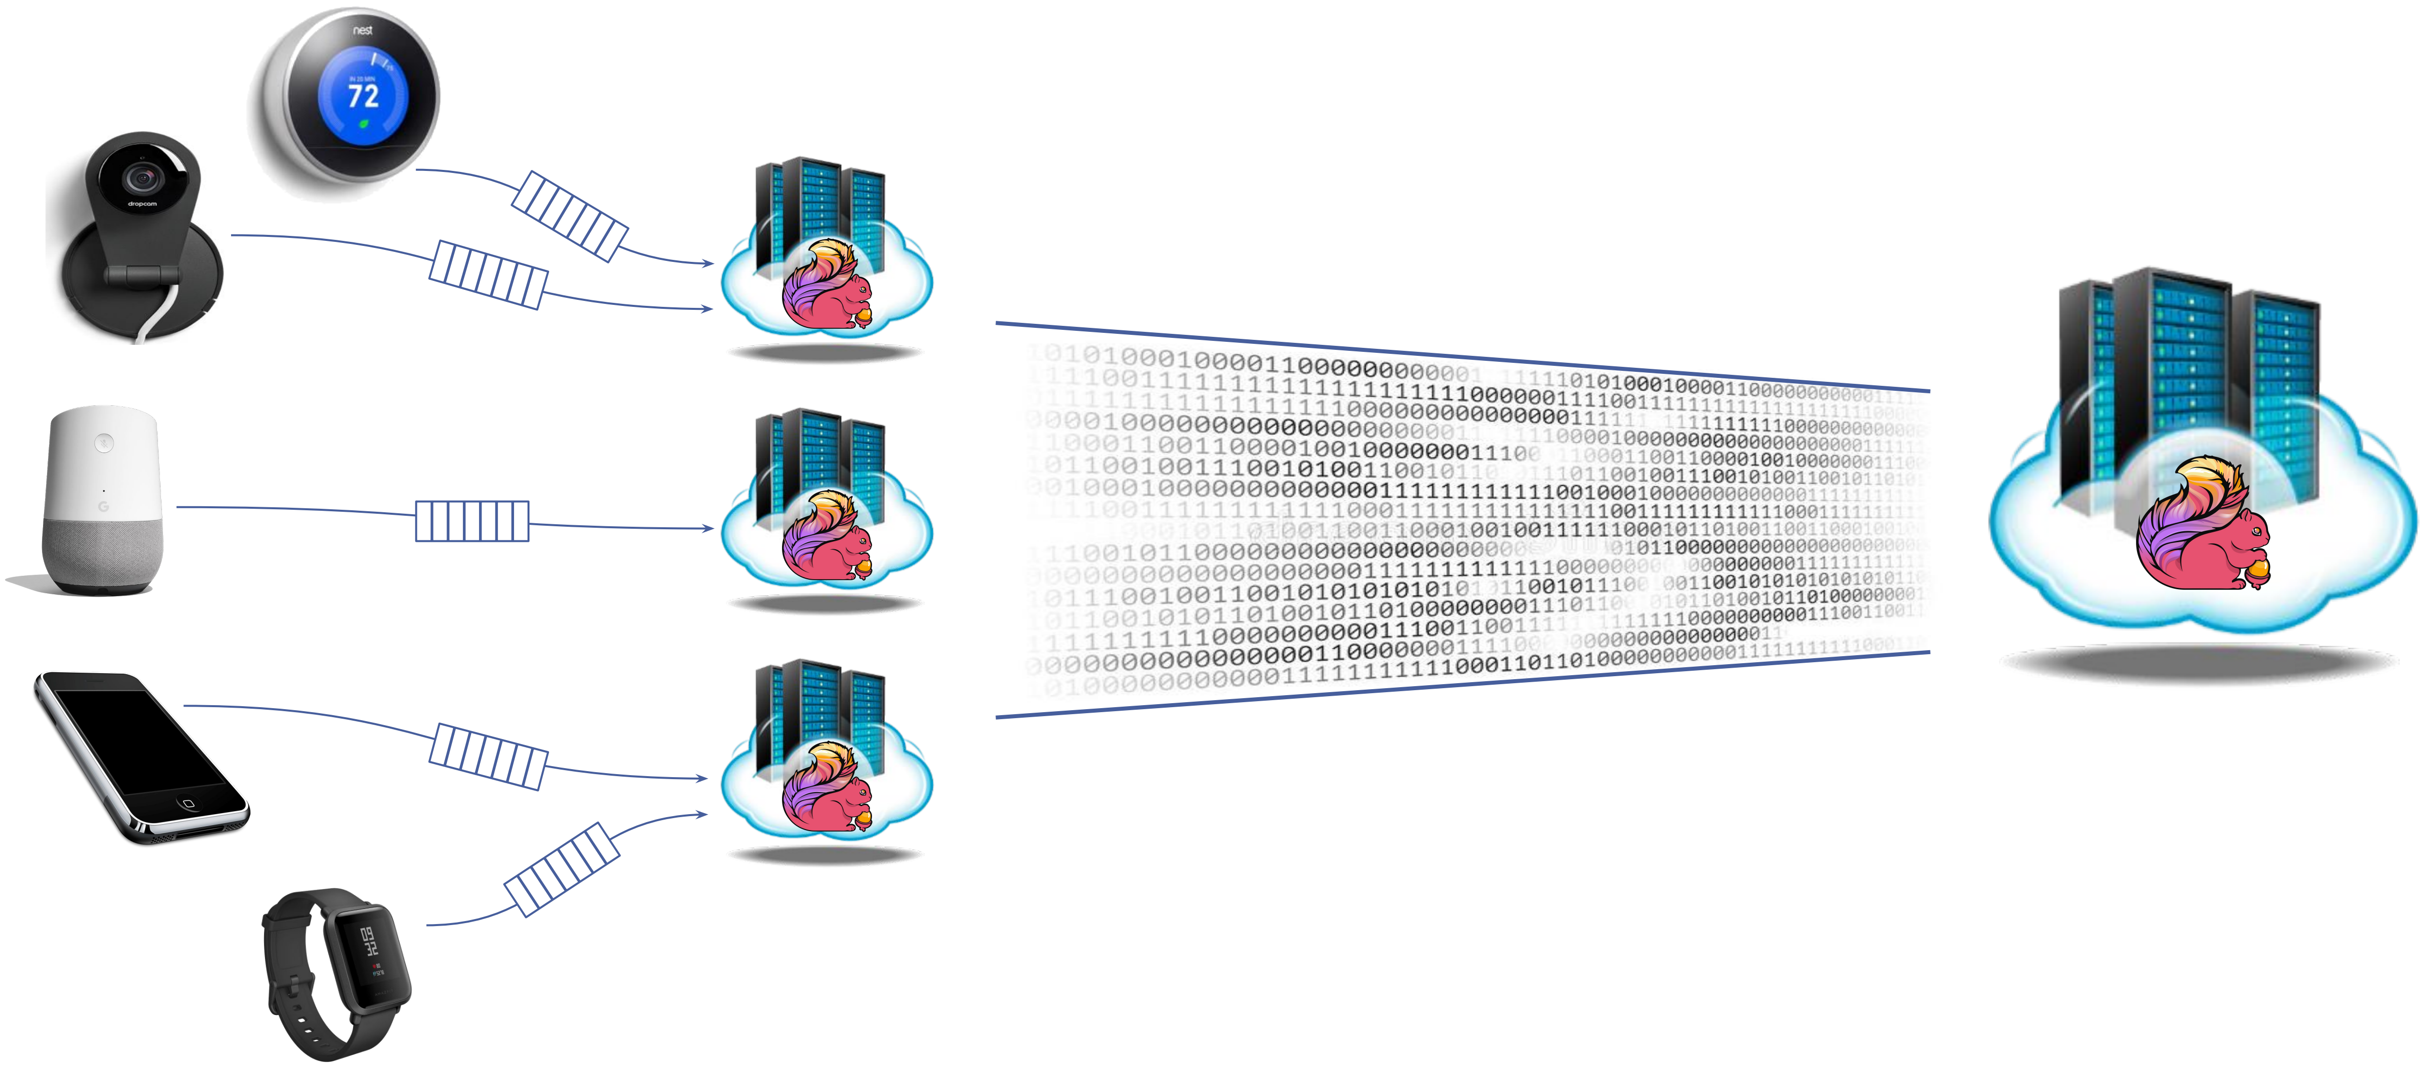

Stream processing in Fog environment

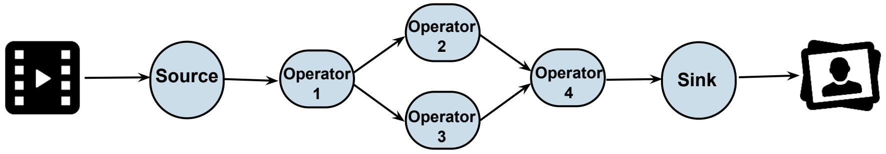

Logical graph of DSP

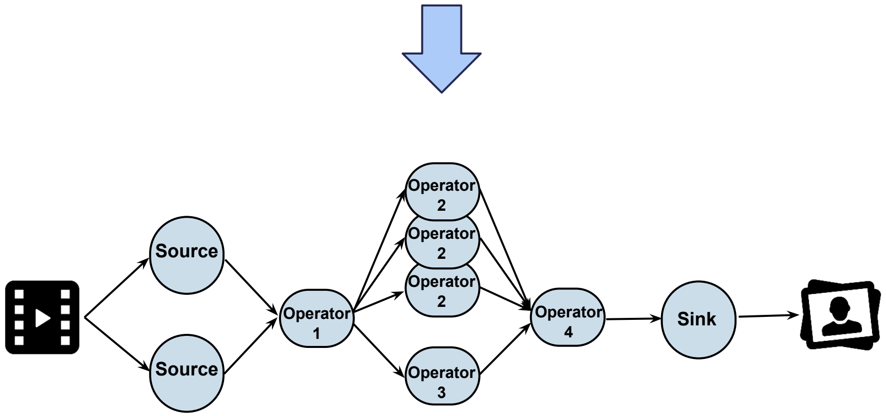

Workflow execution model

Stream processing in geo-distributed environments

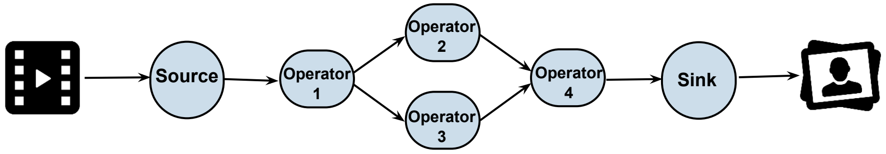

Logical graph of DSP

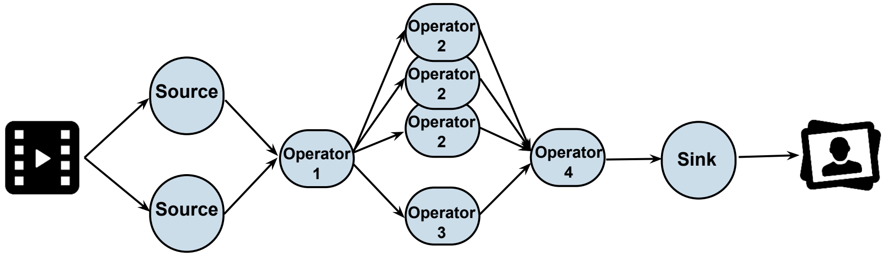

Workflow execution model

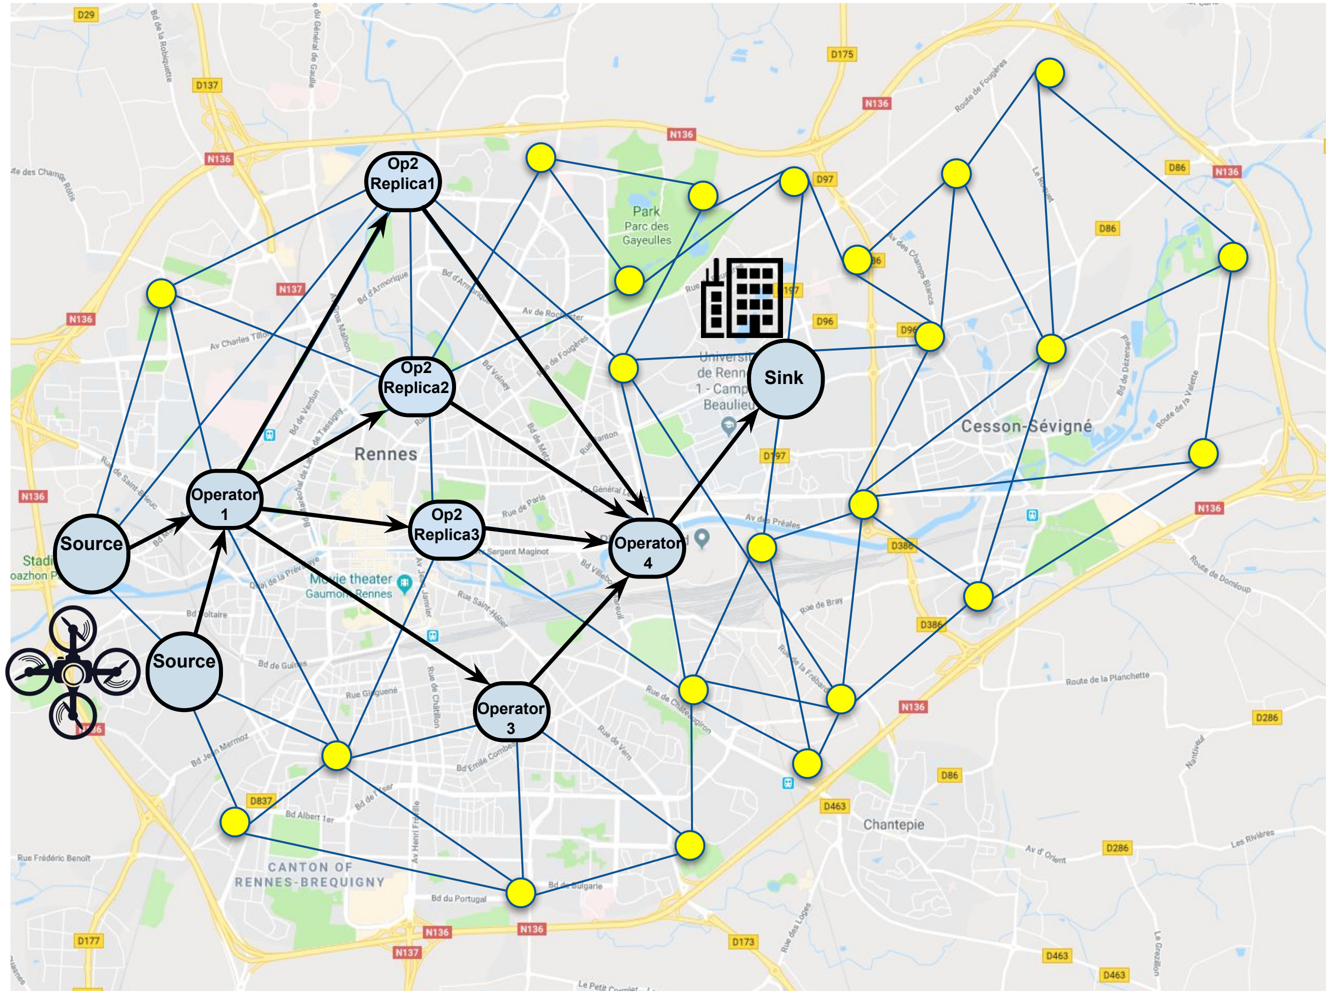

Deployment in Fog geo-distributed environment

Challenges

- Understanding the performance of a geo-distributed stream processing application is difficult.
- Any configuration decision can have a significant impact on performance.

We propose an **experiment-driven** predictive performance model for stream processing systems in fog computing environments.

Experimental setup

- Emulation of a real fog platform
  - 32-core server ≈ 16 fog nodes (2 cores/node)
  - Emulated network latencies
  - Apache Flink

- Test Application
  - Input stream of 100,000 *Tuple2* records
  - The operator calls the Fibonacci function Fib(24) upon every processed record

- Performance metric:
  - Processing Time (PT)

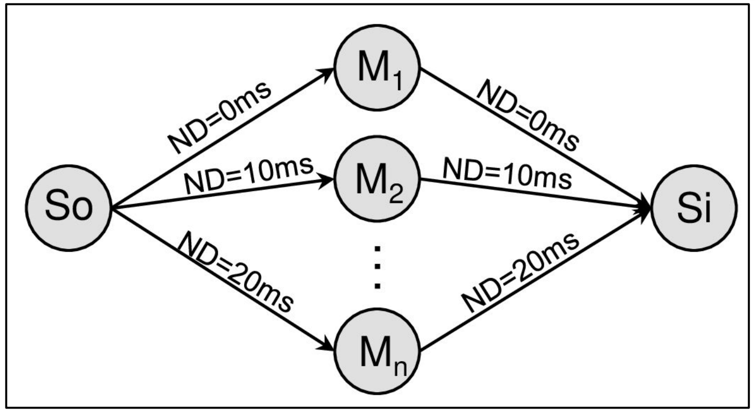

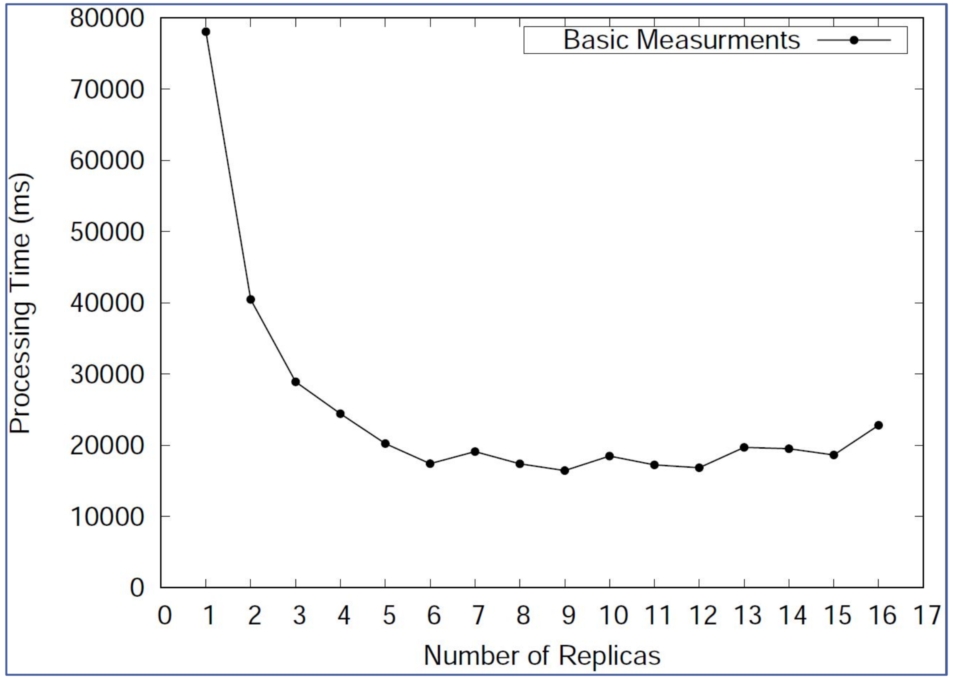

Modeling operator replication

- ***n*** operator replicas should in principle process data ***n*** times faster than a single replica

Experiment

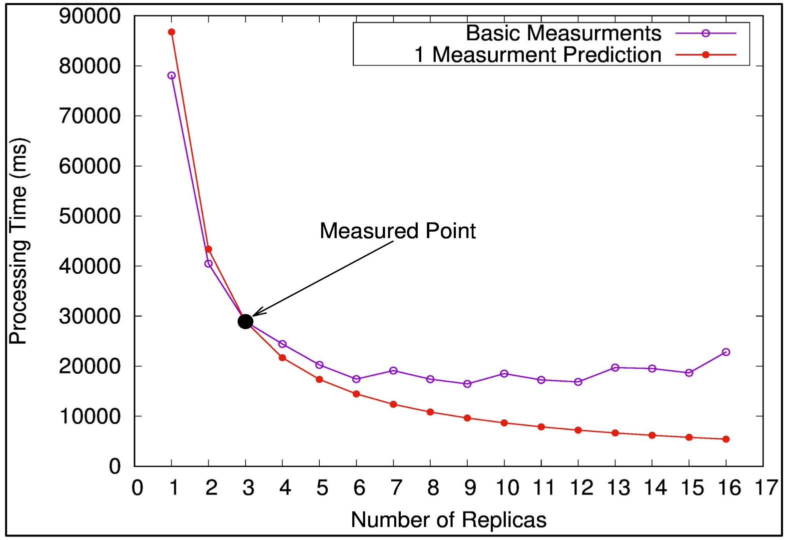

Model

$$\Pi_n = \frac{\alpha}{n}$$

- ***α*** represents the computation capacity of a single node.
- We can determine the value of *α* based on **one** measurement

Considering heterogeneous network delays

- ➢ Network delays between data sources and operator replicas slow down the whole system.
- ➢ When the network delays are heterogeneous, the dominating one is the greatest one (ND_{max}).

Experiment

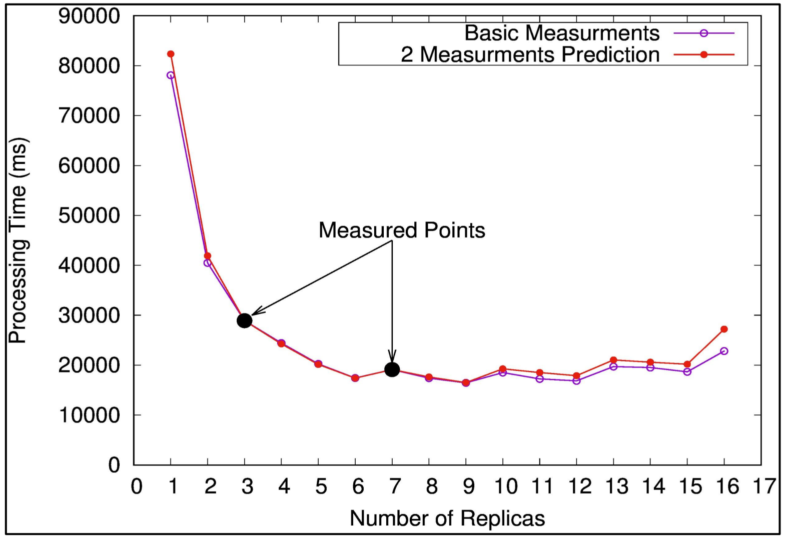

Model

$$\Pi_n = \frac{\alpha}{n} + \gamma \times ND_{max}$$

- ➢ **γ** represents the impact of network delays on overall performance.
- ➢ We can determine both α *and* γ based on **two** measurements

Improving the model's accuracy

- Operator replication incurs some amount of parallelization inefficiency
- The speedup with *n* nodes is usually a little less than *n*

Experiment

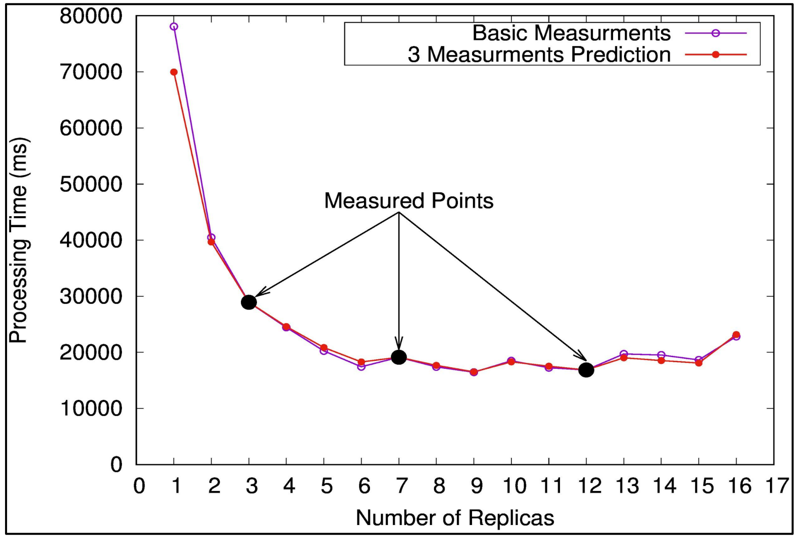

Model

$$\Pi_n = \frac{\alpha}{n^{\beta}} + \gamma \times ND_{max}$$

- β represents Flink's parallelization inefficiency
- We can determine α, β and γ based on **three or more** measurements

Prediction accuracy

Accuracy metric: $MAPE = \frac{100\%}{n} \sum_{i=1}^{n} \frac{\left|PT_i^{actual} - PT_i^{predict}\right|}{PT_i^{actual}}$

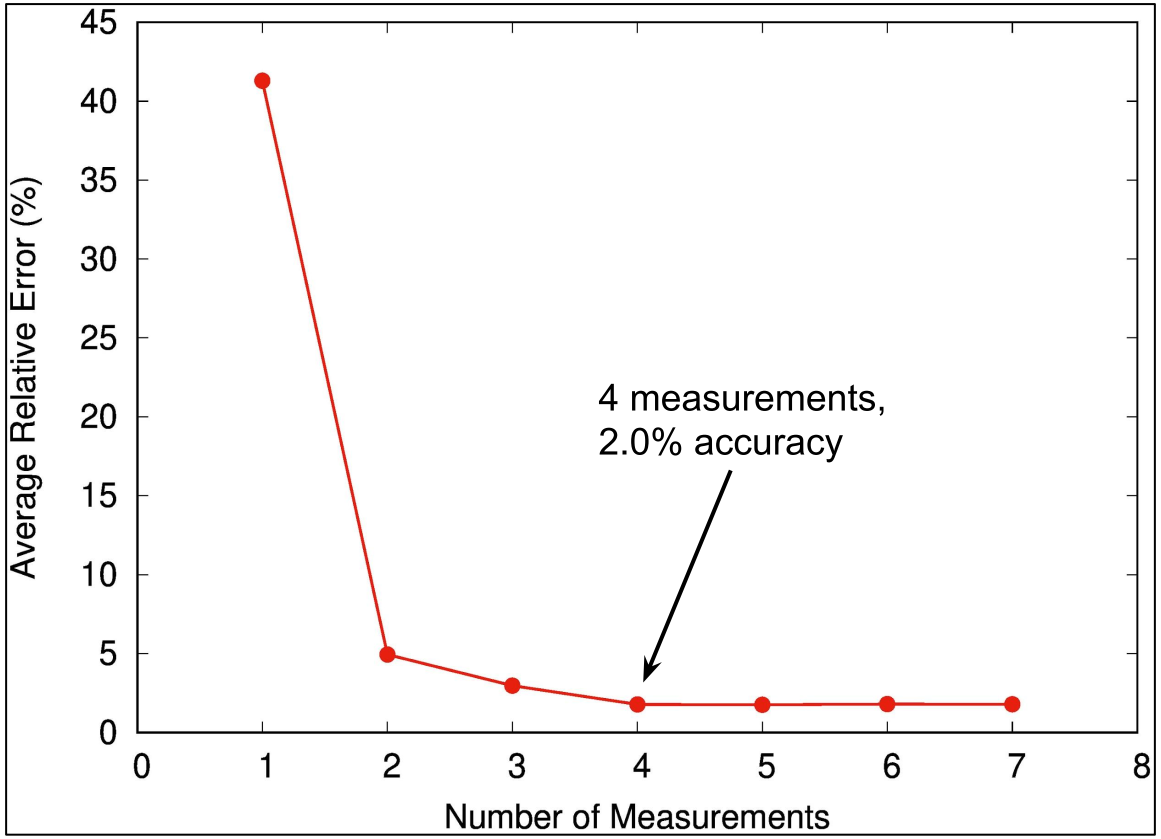

What about modeling an entire (simple) workflow?

➢ The throughput of an entire workflow is determined by the slowest operator

Workflow

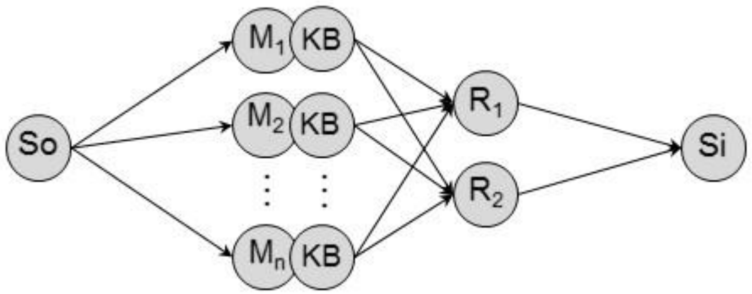

Experiment

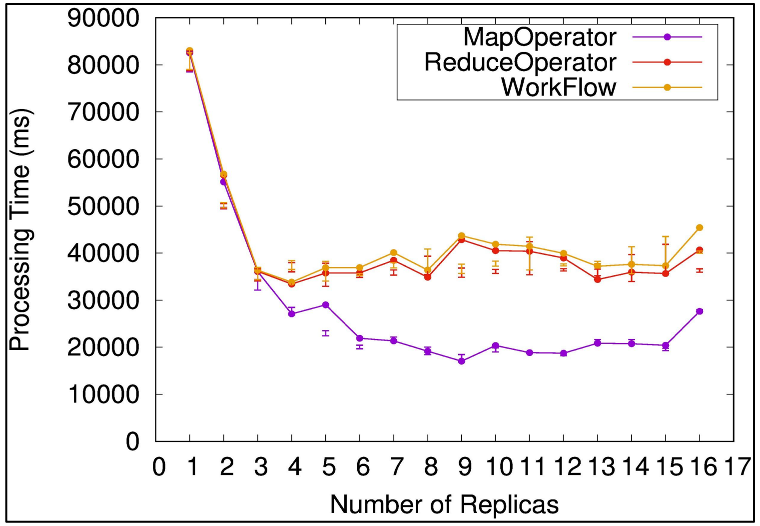

Model

$$\Pi^{Workflow} = max(\Pi^{Map+KeyBy}, \Pi^{Reduce})$$

Can we reuse the parameters instead of multiple measurements?

Calibrated model for **Operator 1**

$$\Pi_n = \frac{\alpha}{n^{\beta}} + \gamma \times ND_{max}$$

$$\left\{ \alpha_1, \beta_1, \gamma_1 \right\}$$

Uncalibrated model for **Operator 2**

$$\Pi_n = \frac{\alpha}{n}$$

$$\left\{ \boldsymbol{\alpha_2}, \beta_1, \gamma_1 \right\}$$

- ➢ α cannot be reused because it is specific to the computation complexity of one operator.
- ➢ β and γ capture properties that are independent from the nature of the computation carried out by the operator.
- ➢ β and γ values of one operator's model might be reused for other operators' models.

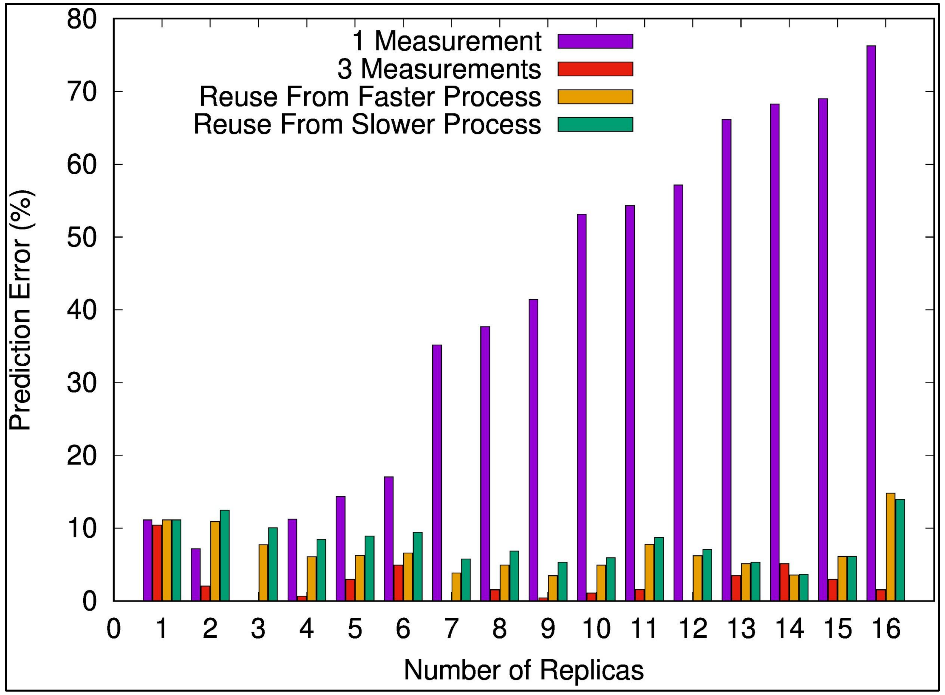

Conclusions

- Heterogeneous network characteristics make it difficult to understand the performance of stream processing engines in geo-distributed environments.
- A predictive performance model for Apache Flink operators that is backed by experimental measurements and evaluations was proposed.
- The model predictions are accurate within ±2% of the actual values.

Hamidreza Arkian

hamidreza.arkian@irisa.fr

Training the next generation of European Fog computing experts

http://www.fogguru.eu/

Acknowledgment

This work is part of a project that has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 765452. The information and views set out in this publication are those of the author(s) and do not necessarily reflect the official opinion of the European Union. Neither the European Union institutions and bodies nor any person acting on their behalf may be held responsible for the use which may be made of the information contained therein.