



Eywa

# ***BLADE RUNNER*** ***2049***



Ryan Gosling / K  
Replicant



Ana de Armas / Joi  
AI

A movie about Fog Computing



# Fog Computing Sci-Fi

Blade Runner 2049's Joi lives in the Fog



## Smart Home Hologram

Joi lives on a **console** in K's home rather than the cloud. She can control all actuators in the house.



## Emanator

Joi can reside on the **portable emanator** and move around with K. A Fog Computing device?



## K's Spinner Car

Joi is connected to the car. When the spinner goes down, she loses the ability to project herself.

# The Internet of Things

Fog Computing  
An emerging technology that bridges the gap, deployed close to the source.

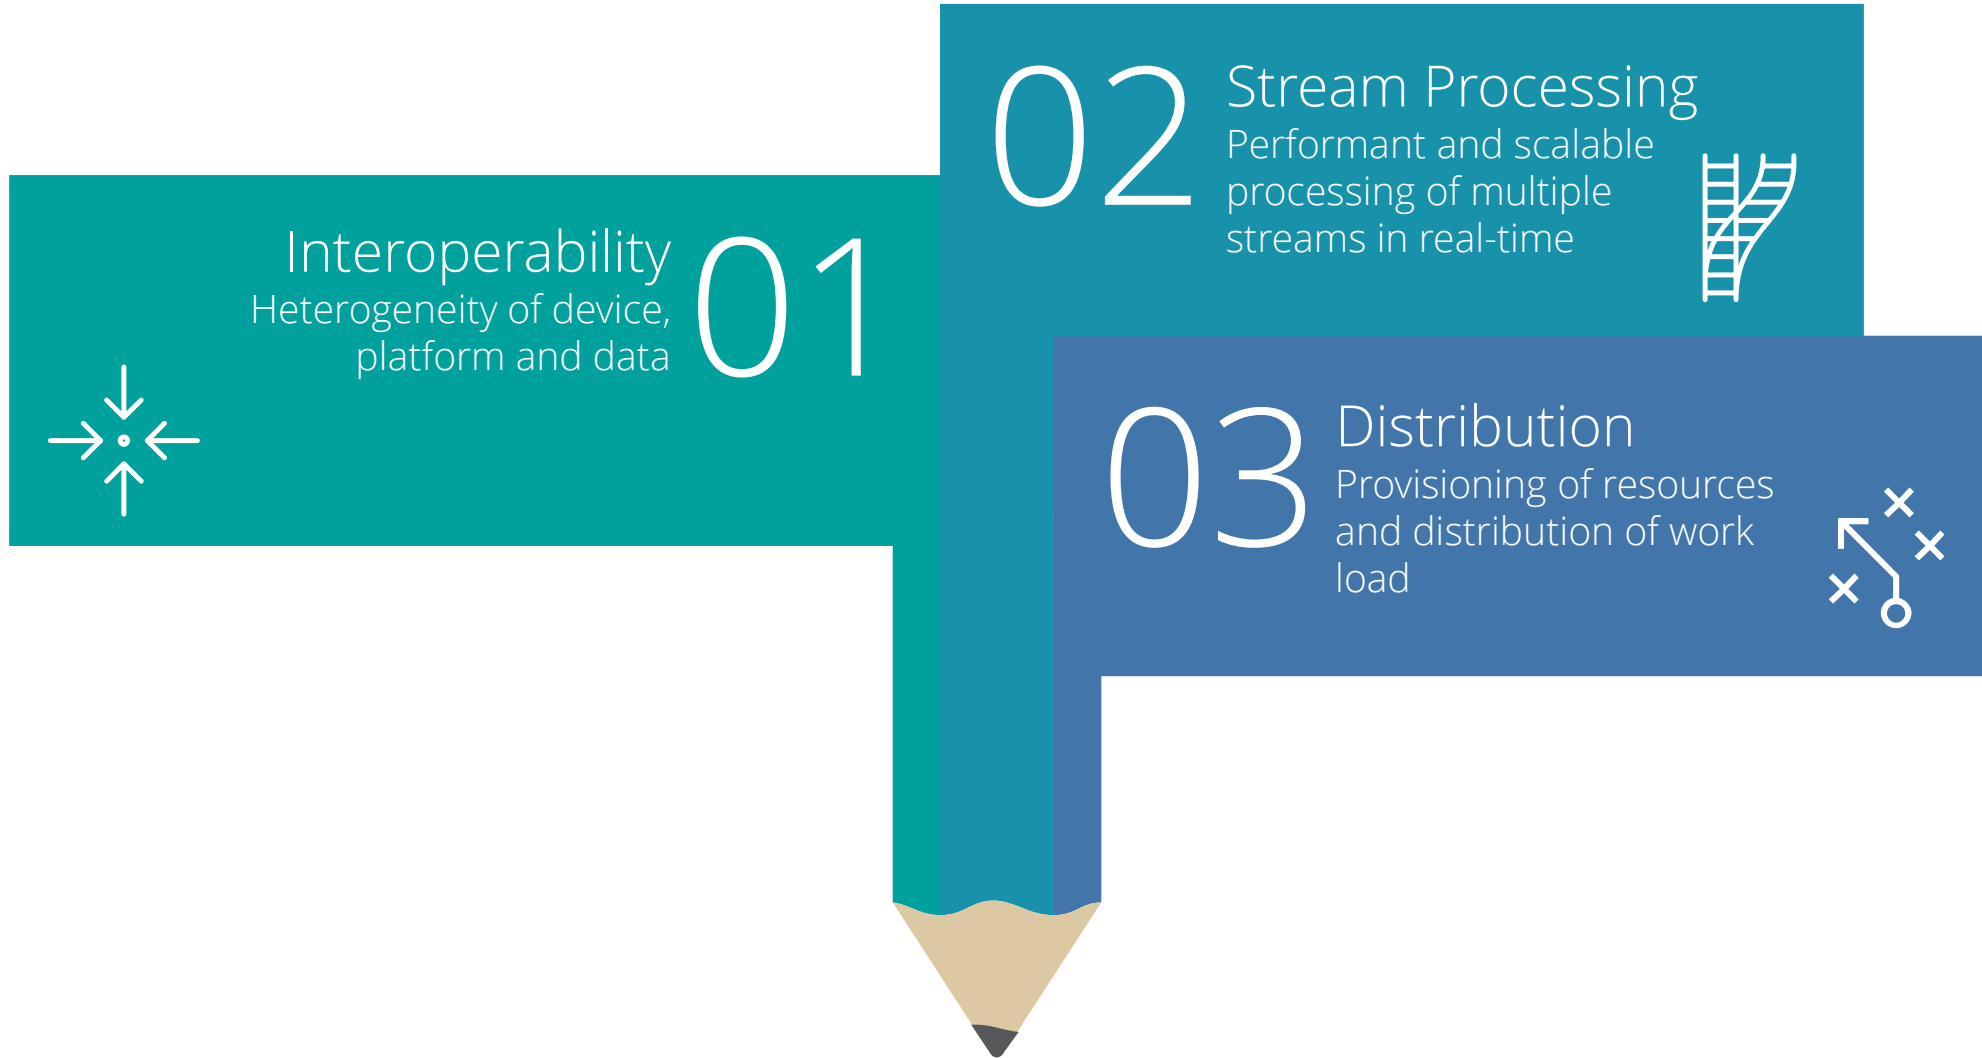


Cloud  
Dynamic provisioning of scalable resources e.g. analytics on a **huge volume** of historical data.



Things  
Connected sensors and actuators producing **streams of time-series data**.

# Challenges for Fog Computing



**Eywa** is like a huge biological internet; the trees are fog computing nodes that store and process information and sensors are connected flora and fauna



Avatar

By James Cameron

# Eywa

A Fog Computing Infrastructure

IoT Sensors  
Produce Observations



Streams of  
Time-series Data

Distribution of Query Workload

Interoperable Stream  
Processing


Fog Node  
Stores and Processes






# Semantic Interoperability in Eywa

Using a Common RDF Graph Model



IoT Domains for Things and Apps




Common Structure  
Widely-used, flexible model

Data Integration  
Stores Rich Metadata

Graph Querying  
Powerful SPARQL Graph Queries





## Deliver


Inverse-publish-  
subscribe

## Distribute

Workload Distribution  
by Projection  
Pushdown

## Process

Stream Processing by  
Query Translation



Eywa Node

S

Source Node  
Publishes Data



τ


Client Node  
Issues Queries



b

Broker Node  
Facilitate Network Formation, Forwards Data





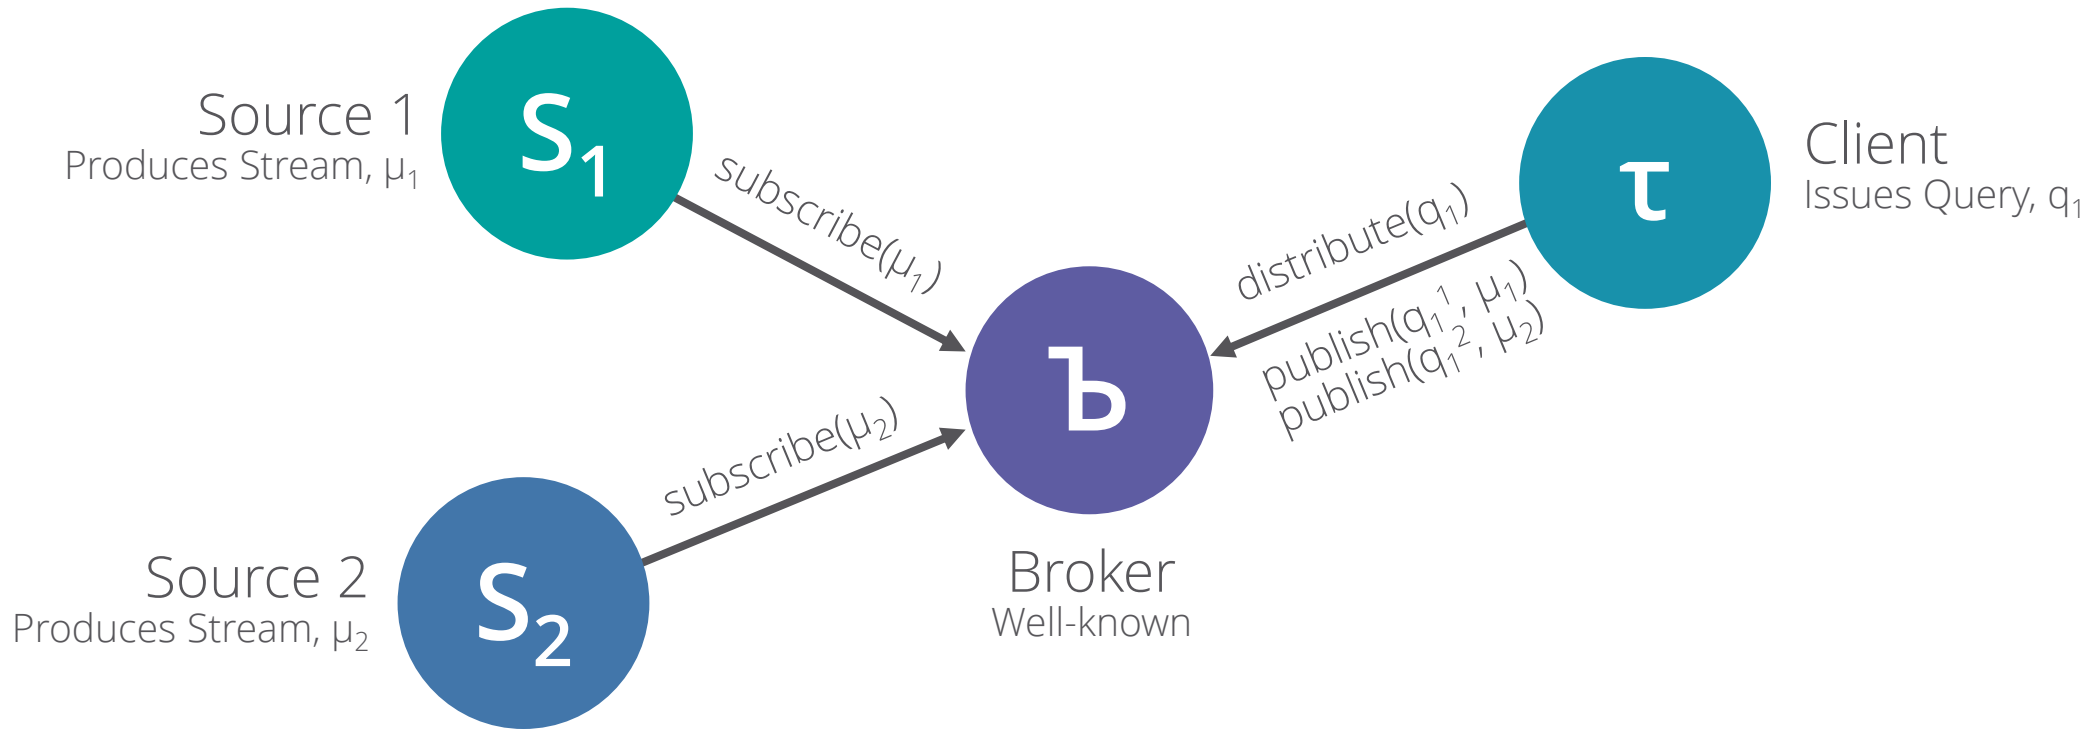
# Deliver Queries

Inverse-Publish-Subscribe

A novel system for **best-effort, co-operative query delivery** in fog computing *control-plane*


# Deliver Queries


Inverse Publish-Subscribe (1)



# Deliver Queries

## Inverse Publish-Subscribe (2)





# Distribute Workload

Projection Pushdown

Efficient streaming of only **projected data** in the fog computing *data-plane*

# Distribute Workload


Graph Query in SPARQL

Source 1  
Processing  $q_1^1$



Project  
 $?v1, ?v2, ?v3$

Client  
Receives the Projection



Project  
Streams

No Extra Join  
Variables

temp, hum,  
congestionLevel

congestionLevel

temp, hum

$?v3$

$?v1, ?v2$

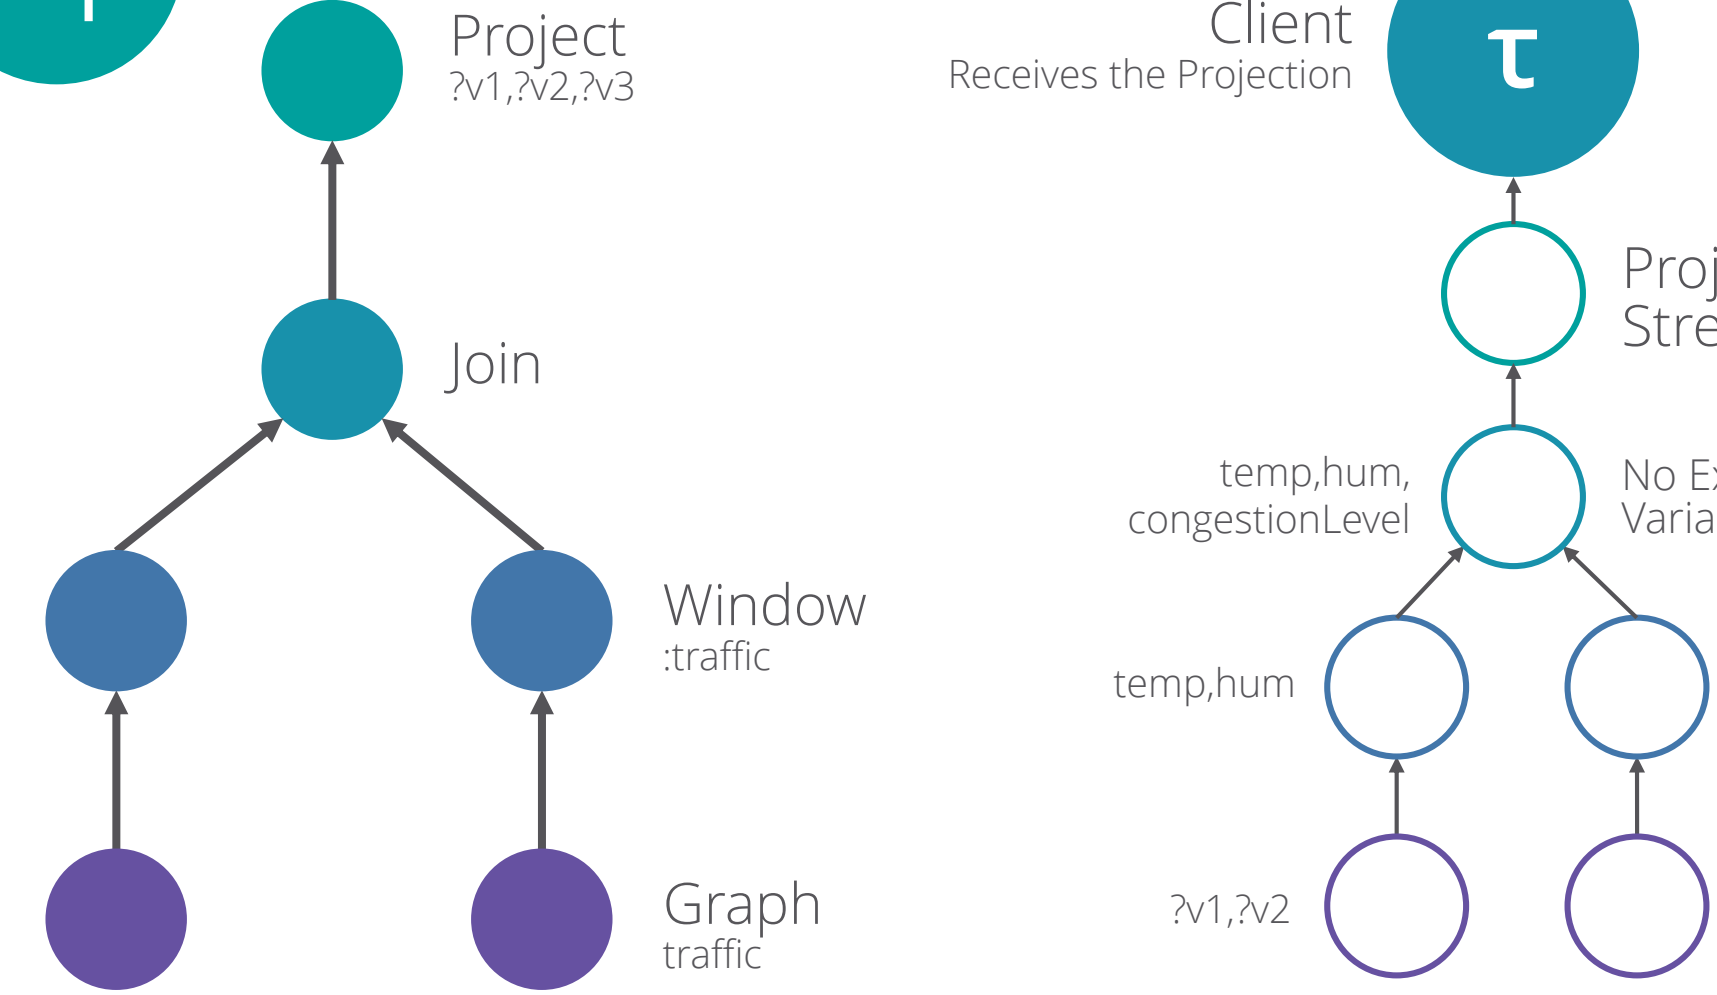
Window  
:weather

Window  
:traffic

Graph  
weather

Graph  
traffic

Join

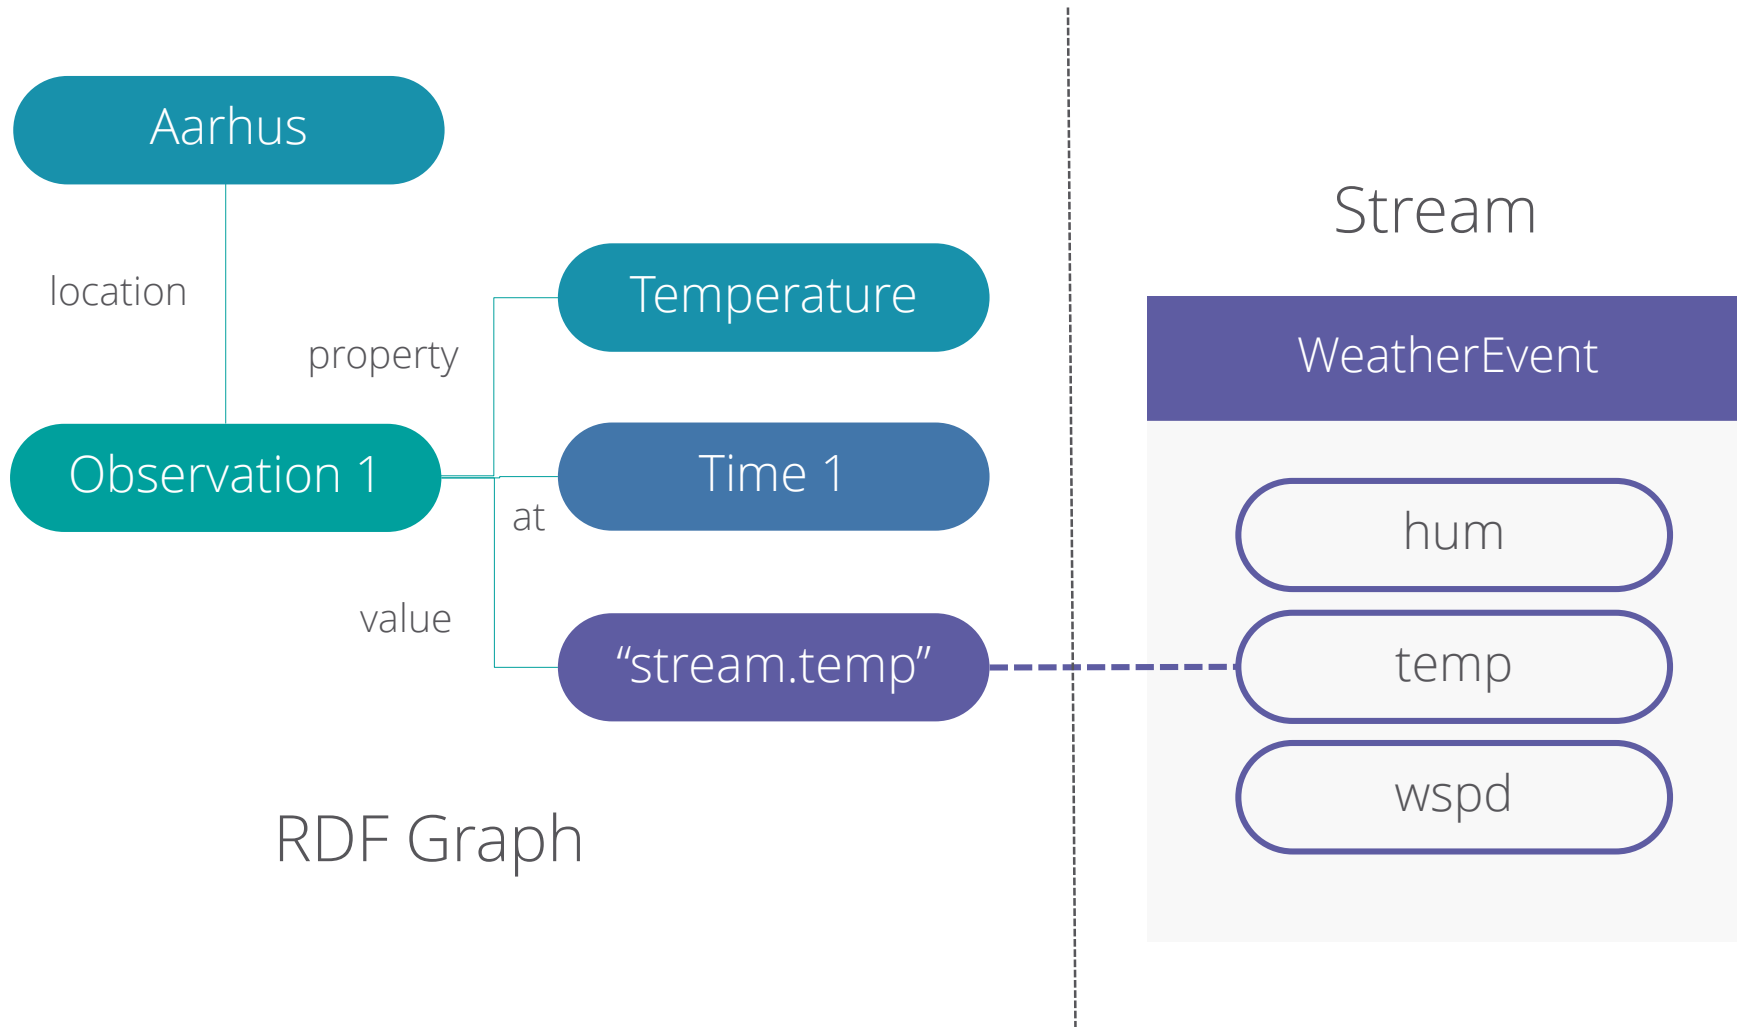


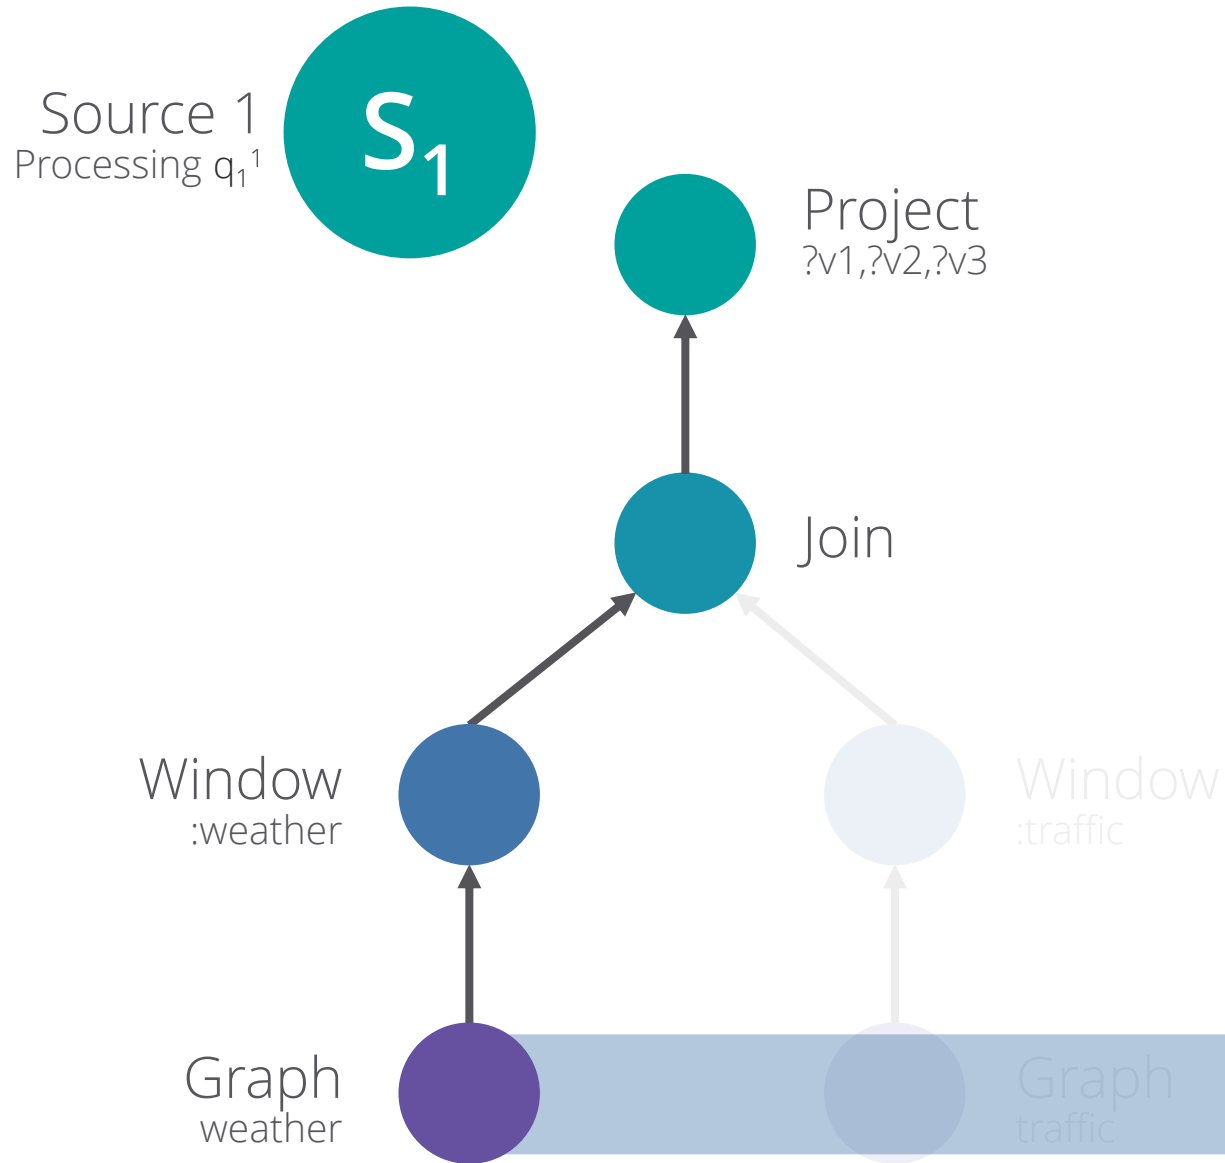
From CityBench (Smart City Streams) Query 2. Finding the traffic congestion level and weather conditions of my planned journey.



# Distribute Workload

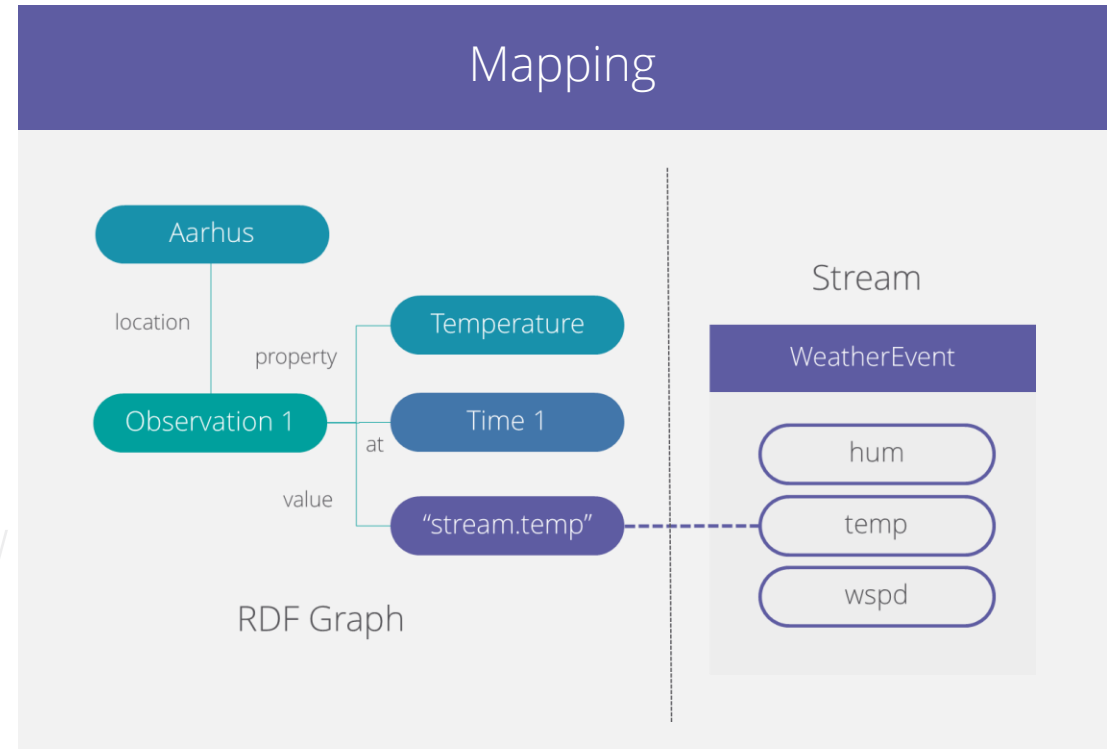
Efficient Mappings for RDF Stream Processing






# Distribute Workload

## Graph Matching to Projection (1)





Window  
:weather




Window  
:traffic



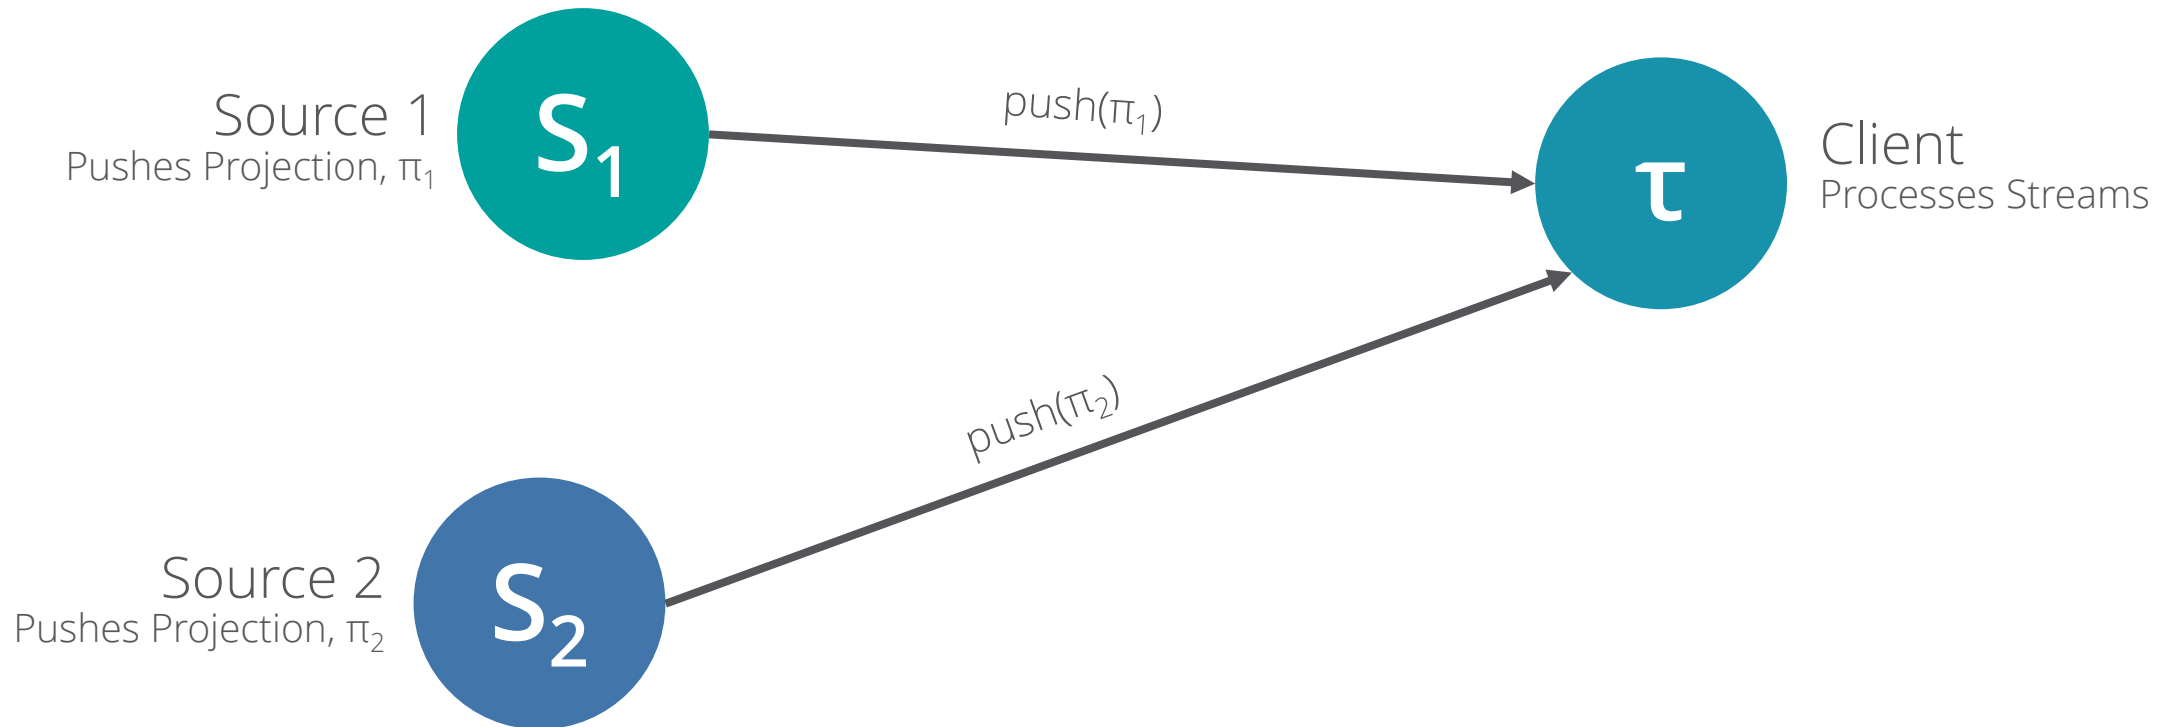
# Distribute Workload

Graph Matching to Projection (2)



# Distribute Workload

Projection Pushdown, Push Projection







# Process

Query Translation

Efficient stream processing of graphs for IoT time-series in the fog computing *data-plane*



Client  
Processes Query,  $q_1$



# Processing Streams

Query Translation for Stream Processing

## Event Processing Language Query

SELECT

temp

AS

v1

hum

AS

v2

congestionLevel


AS

v3

FROM

weather(3s)

congestionLevel(3s)



# Evaluation

CityBench Smart City Benchmark

Latency and Scalability



# Evaluation on 3 Stream Processing Engines

## Smart City RDF Streams

### CITYBENCH



Real-time streams (e.g. vehicle traffic, parking, weather, pollution)



Based on smart city applications (e.g. parking space finder, admin console)



Run on resource-constrained Raspberry Pis as Fog Nodes (~500mhz CPU, 512mb ram, SD CARD)

01

### C-SPARQL

Barbieri et al. "C-SPARQL: SPARQL for continuous querying." WWW2009.

02

### CQELS


Le-Phuoc et al. "A native and adaptive approach for unified processing of linked streams and linked data." ISWC2011

03

### Eywa


# Latency Evaluation

CityBench Query 1 (traffic congestion level on two roads)



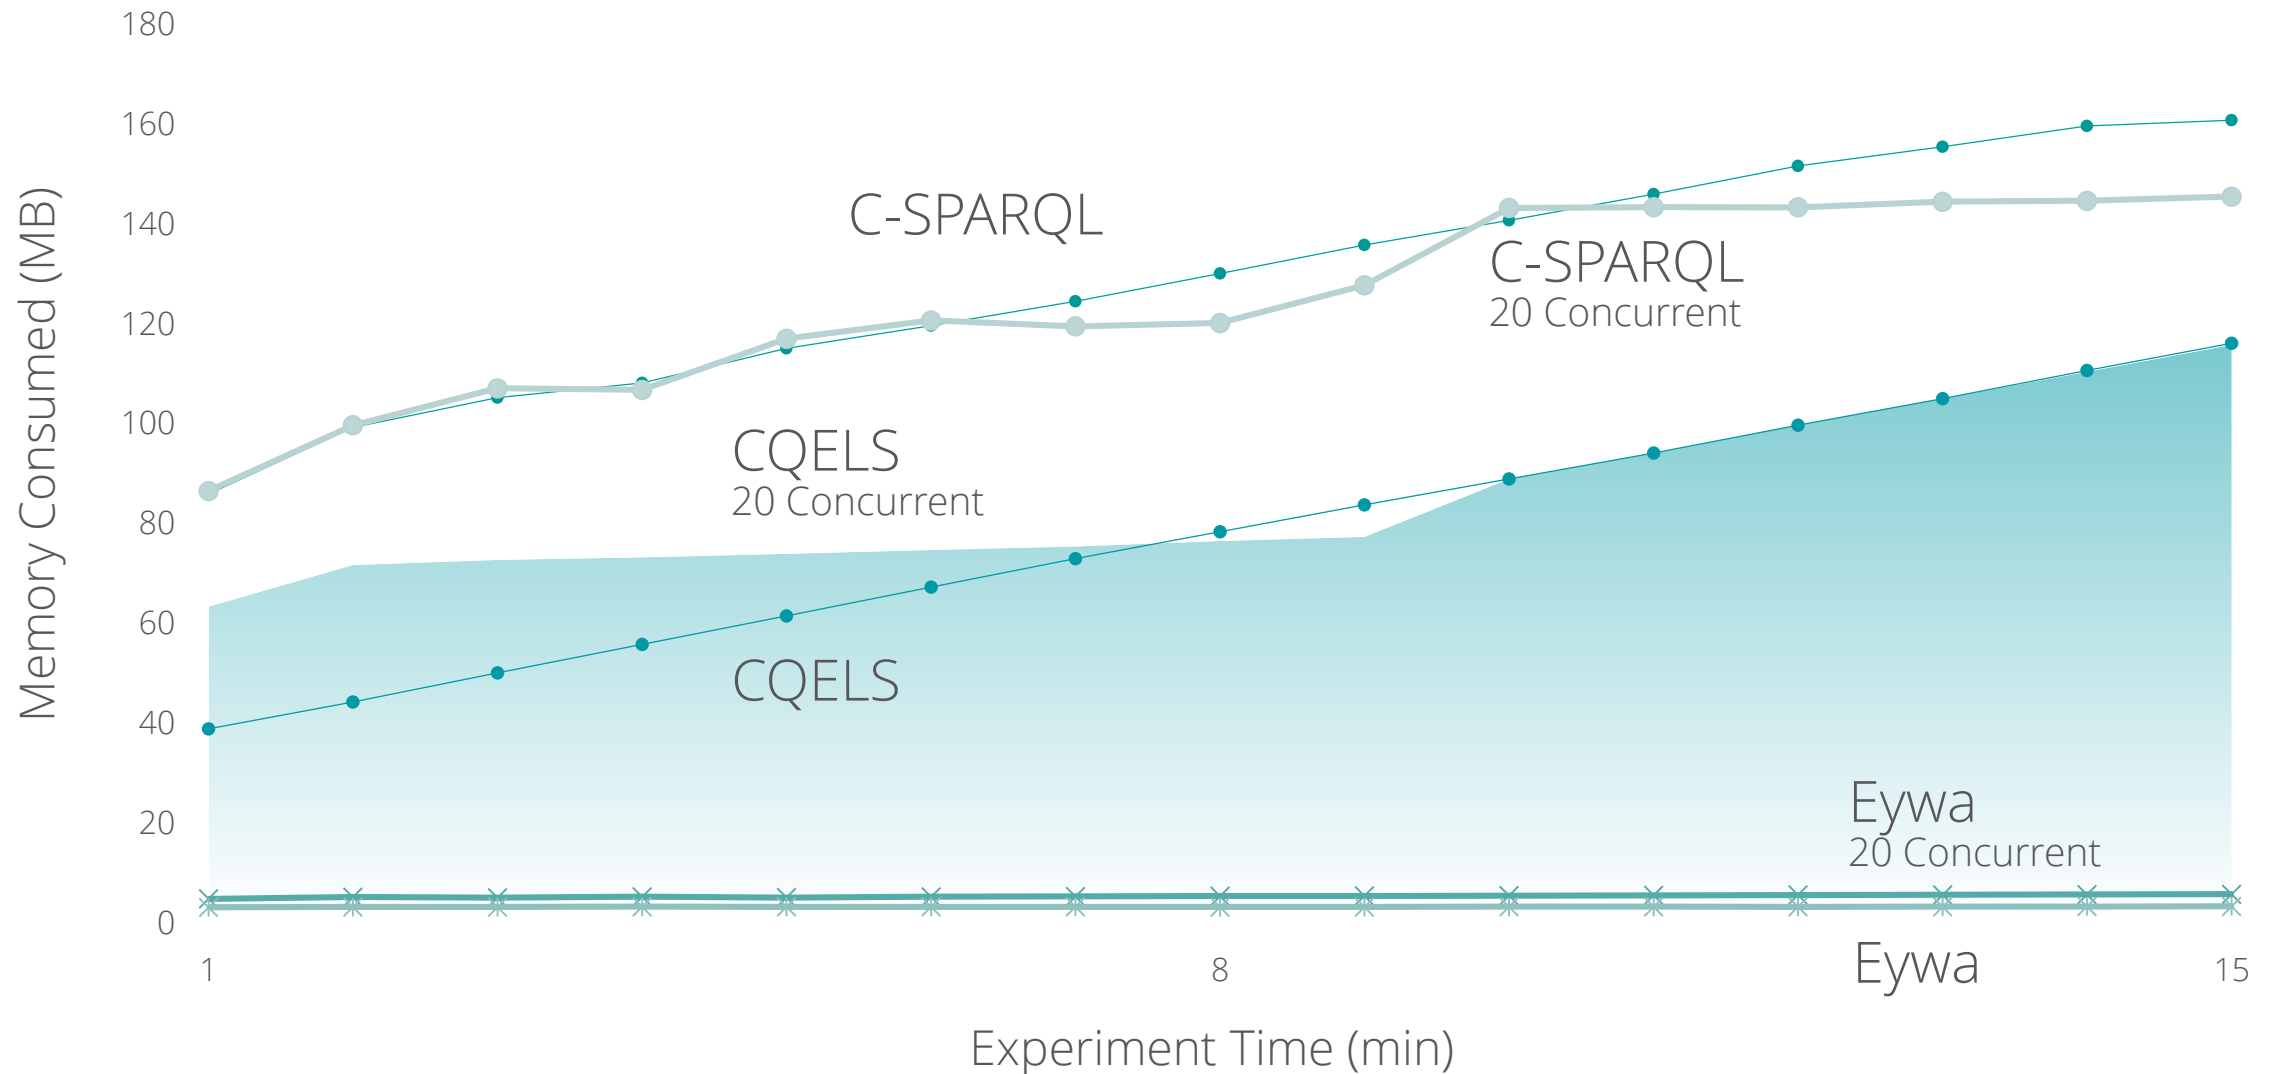
# Scalability Evaluation

CityBench Query 2 (traffic congestion level and weather)



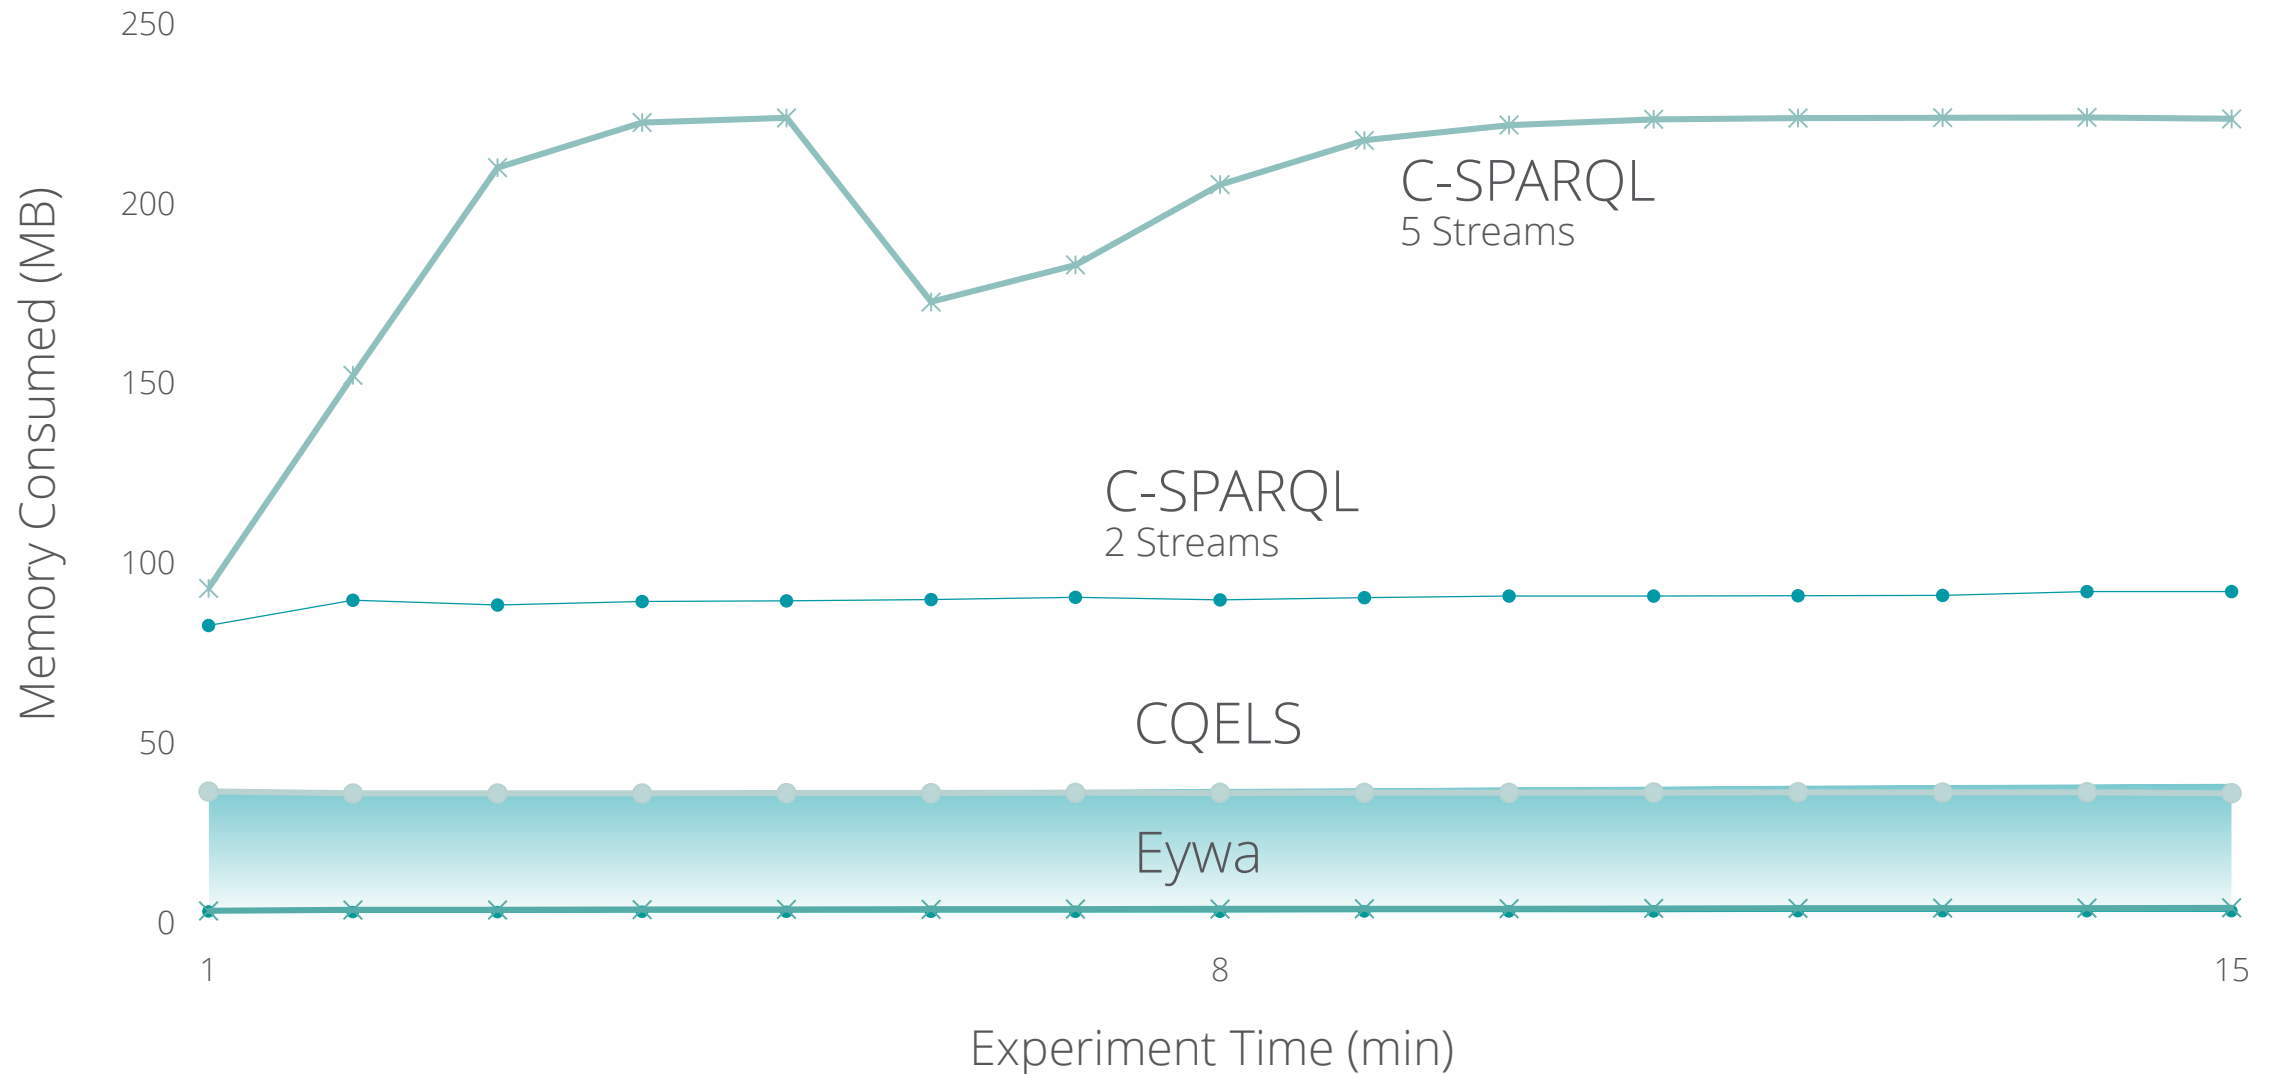
# Scalability Evaluation

CityBench Query 5 (traffic congestion where event is happening)




# Scalability Evaluation

CityBench Query 10 (most polluted area in the city in real-time)



# Fog Scalability Evaluation

CityBench Query 10 (most polluted area in the city in real-time)



# Conclusions

Latency  
Eywa is fast and performant  
at stream processing 01



Interoperability  
From the RDF graph model  
for metadata and data 03



Utility for the IoT  
Data locality, offline  
access, data ownership 05



Scalability  
Eywa Fog consumes less  
memory than other engines 02



Distributed  
Eywa is a means for  
distributed Fog Computing 04

