

SQPR: Stream Query Planning with Reuse

Evangelia Kalyvianaki* Wolfram Wiesemann* Quang Hieu Vu+ Daniel Kuhn* Peter Pietzuch*

* Imperial College London, Dept. of Computing; + ETISALAT BT Innovation Center;

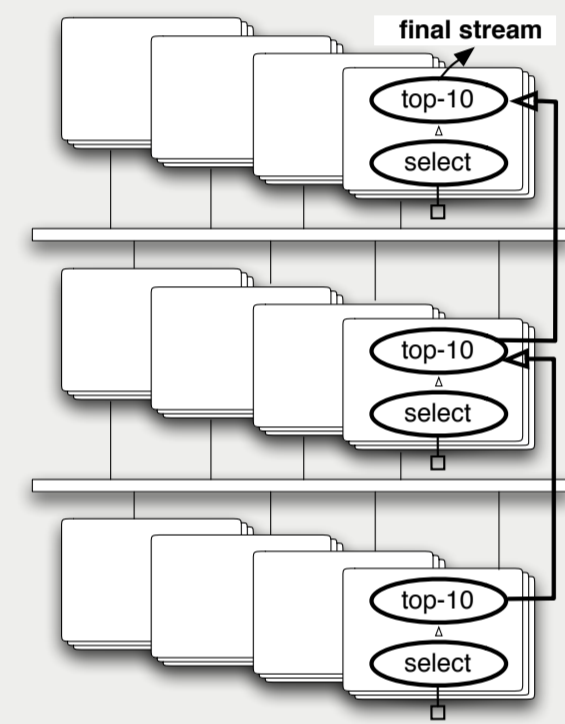
Data Stream Processing

Data streaming queries involve the processing of data from *sources* in near *real-time*. Stream processing appears in many application domains:

- ▶ environmental sensing
- ▶ healthcare systems
- ▶ financial trading
- ▶ network monitoring

Query Planning

Distributed Stream Processing System (DSPS)



Example Query:
Which are the top-10 locations, in the city center, with the highest concentration of CO?

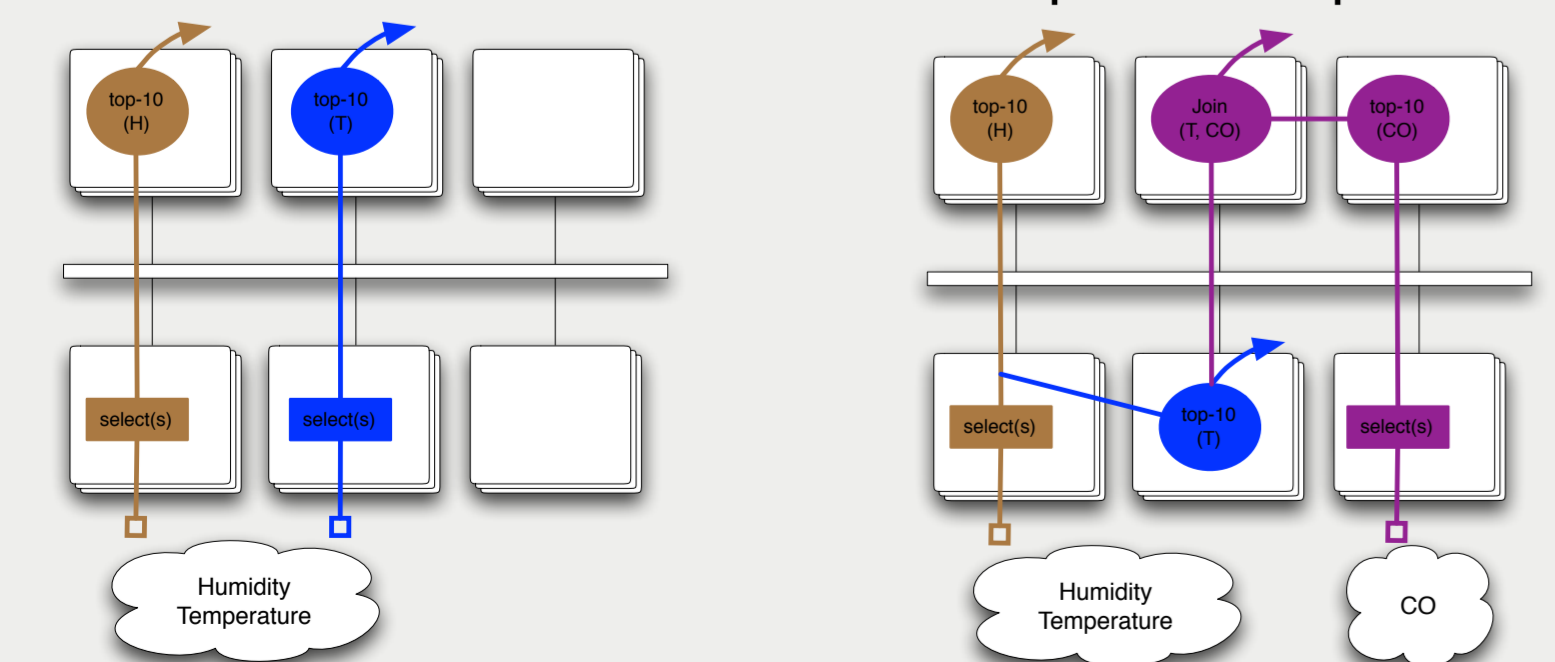
- Query plan involves:**
- ▶ operator placement
 - ▶ resource allocation: CPU, net
 - ▶ stream connections

Challenges in Query Planning

Previous approaches assume vastly provisioned DSPSs, or use heuristics. But, today, large volumes of queries soon exhaust DSPSs resources.

Pitfalls:

1. single resource exhaustion
2. waste of resources because of operator repetition



operator overlap stream reuse

SQPR Overview

SQPR treats query admission, **multi-resource** operator allocation and **reuse** as a single optimisation model.

Optimisation Problem

maximise:
 $\lambda_1 * (\#queries) - \lambda_2 * (CPU \text{ usage}) - \lambda_3 * (net \text{ usage}) - \lambda_4 * (balance \text{ load})$
subject to constraints:

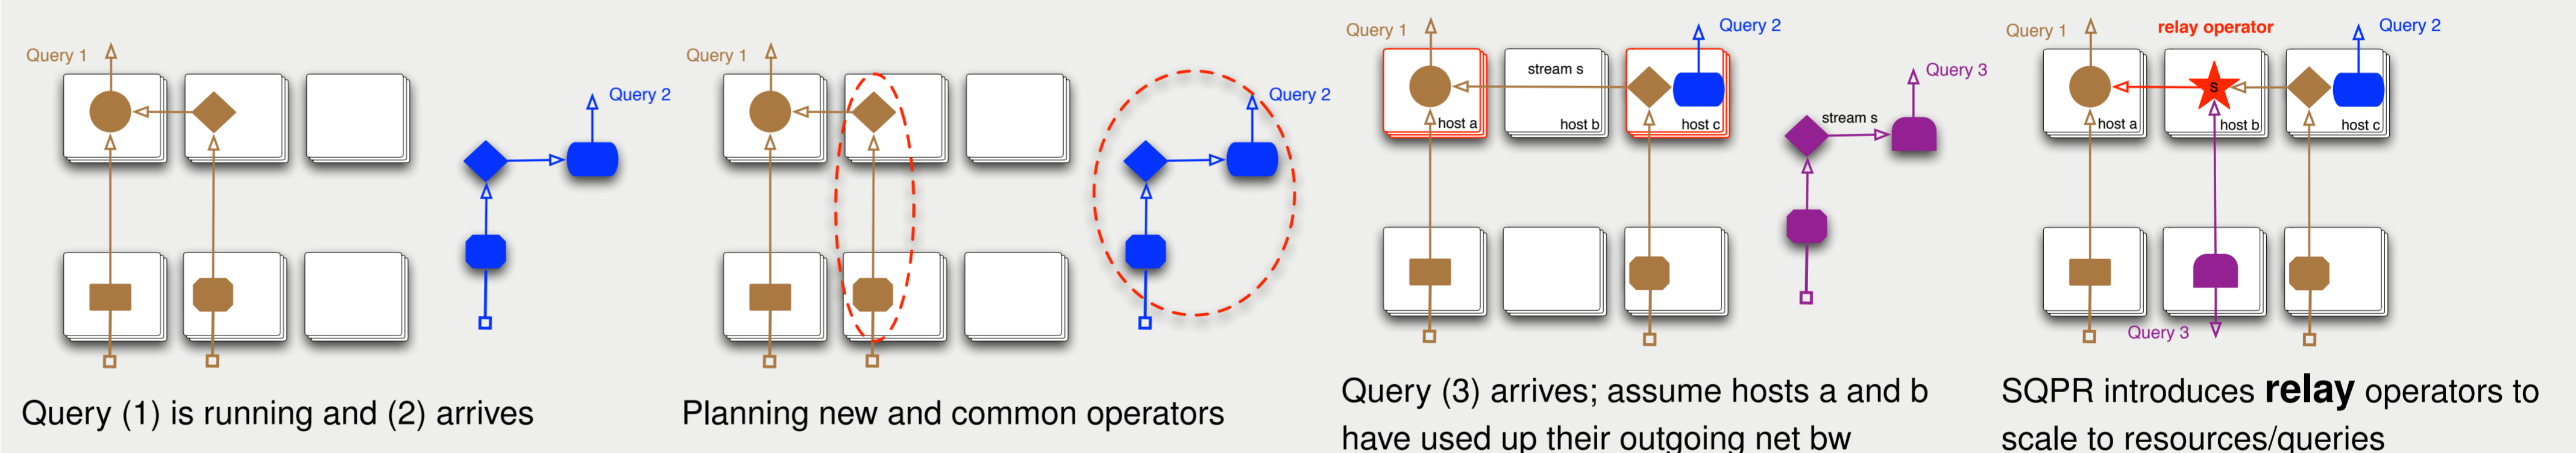
1. *availability:* streams for operators are present in hosts
2. *resources:* allocations are within the physical capacity
3. *demand:* a query is satisfied if final stream is generated
4. *acyclicity:* all streams come from a real source

→ This is an **NP-hard** problem when optimising for all existing and new queries.


SQPR Opts for *Tractable* Solutions by:

1. solving the optimisation problem for new queries
2. keeping existing queries running
3. optimising over new queries and re-planning common operators from old and new queries

SQPR Optimisation Example



Experimental Evaluation



References: [1] J. Wolf, N. Bansal, K. Hildrum, S. Parekh, D. Rajan, R. Wagle, K.L. Wu and L. Fleischer, "SODA: An Optimizing Scheduler for Large-scale Stream-based Distributed Computer Systems", in Middleware, 2008.