

# POLAR: Adaptive and Non-invasive Join Order Selection via Plans of Least Resistance

David Justen<sup>1,2</sup>, Daniel Ritter<sup>3</sup>, Campbell Fraser<sup>4</sup>, Andrew Lamb<sup>5</sup>, Nga Tran<sup>5</sup>, Allison Lee<sup>6</sup>, Thomas Bodner<sup>7,8</sup>, Mhd Yamen Haddad<sup>9,10</sup>, Steffen Zeuch<sup>1,2</sup>, Volker Markl<sup>1,2</sup>, Matthias Boehm<sup>1,2</sup>

Corresponding author's email: [david.justen@tu-berlin.de](mailto:david.justen@tu-berlin.de)

## Motivation

Wrong cardinality estimates lead to **suboptimal join orders**

**Adaptive Query Processing (AQP)** measures cardinalities to reorganize joins during execution

Despite decades of AQP research, **low adoption in practice**

- ▶ System and integration **complexity**
- ▶ Large **performance overheads**

## Design Objectives

**Non-invasive Integration, Overhead Mitigation**


**Separation** of compilation and execution phase

**Reuse** host system optimizer and operators

Allow **fallback** on original plan

**Bound** exploration overhead

## System Overview



Enhance optimized query pipelines with **additional join orders**

Extensible **multiplexer** operator with several **routing strategies** and probabilistic **regret bound**

Open-source **DuckDB** prototype


**Extension** in development

## Join Order Selection

Form **d-dimensional grid** from base table predicate and join selectivities

Discretize selectivities with **exponential decay**

**Sample grid uniformly** and invoke DPsize with sampled points



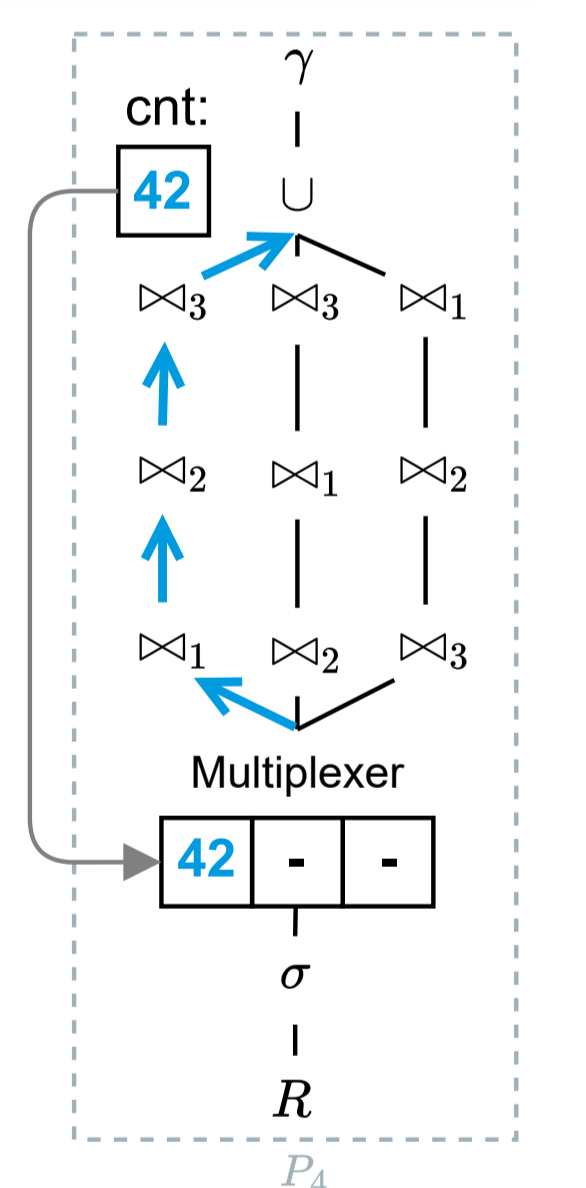
## Adaptive Pipeline Execution

Route bags of tuples consecutively through pre-selected join orders

Use **path resistance** (i.e., intermediate results per input tuple) for routing decisions


**Probabilistic Regret Bound**

Trade-off between *exploiting* well-performing join orders and *exploring* weaker paths based on past observations



## Experiments

### Individual Performance Impact




**JOB**: occasional performance improvements up to 9x

**SSB**: small impact and moderate overhead up to 7%

**SSB-skew**: improvement in almost every query

### Total Execution Times



**POLAR**: total execution time consistently  $\leq$  DuckDB

**Large benefits from reusing DuckDB components**