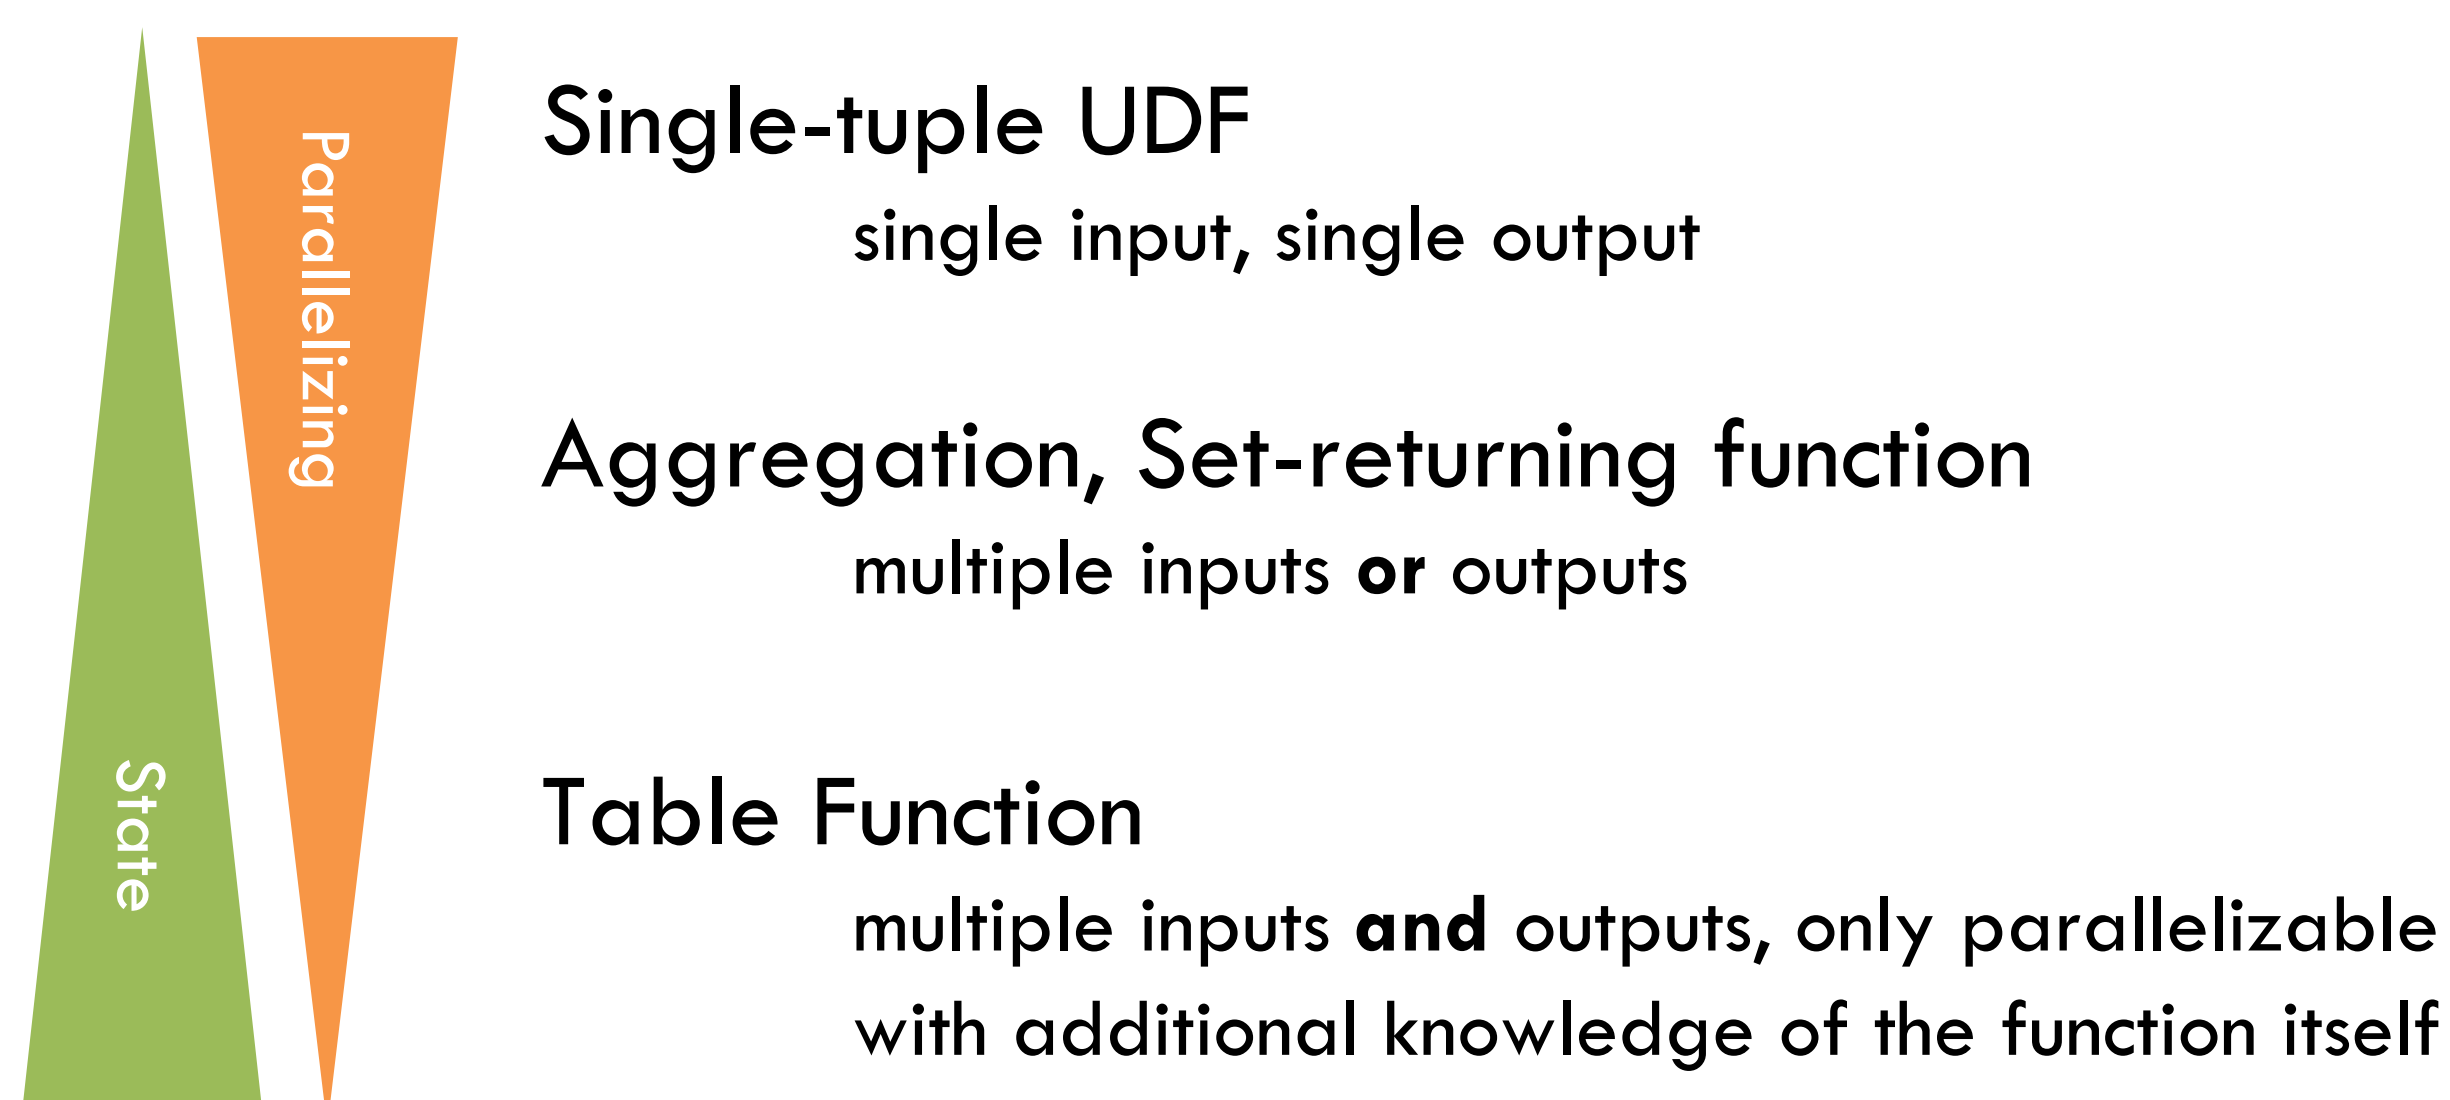


Towards Scalable UDTFs in Noria^[1]

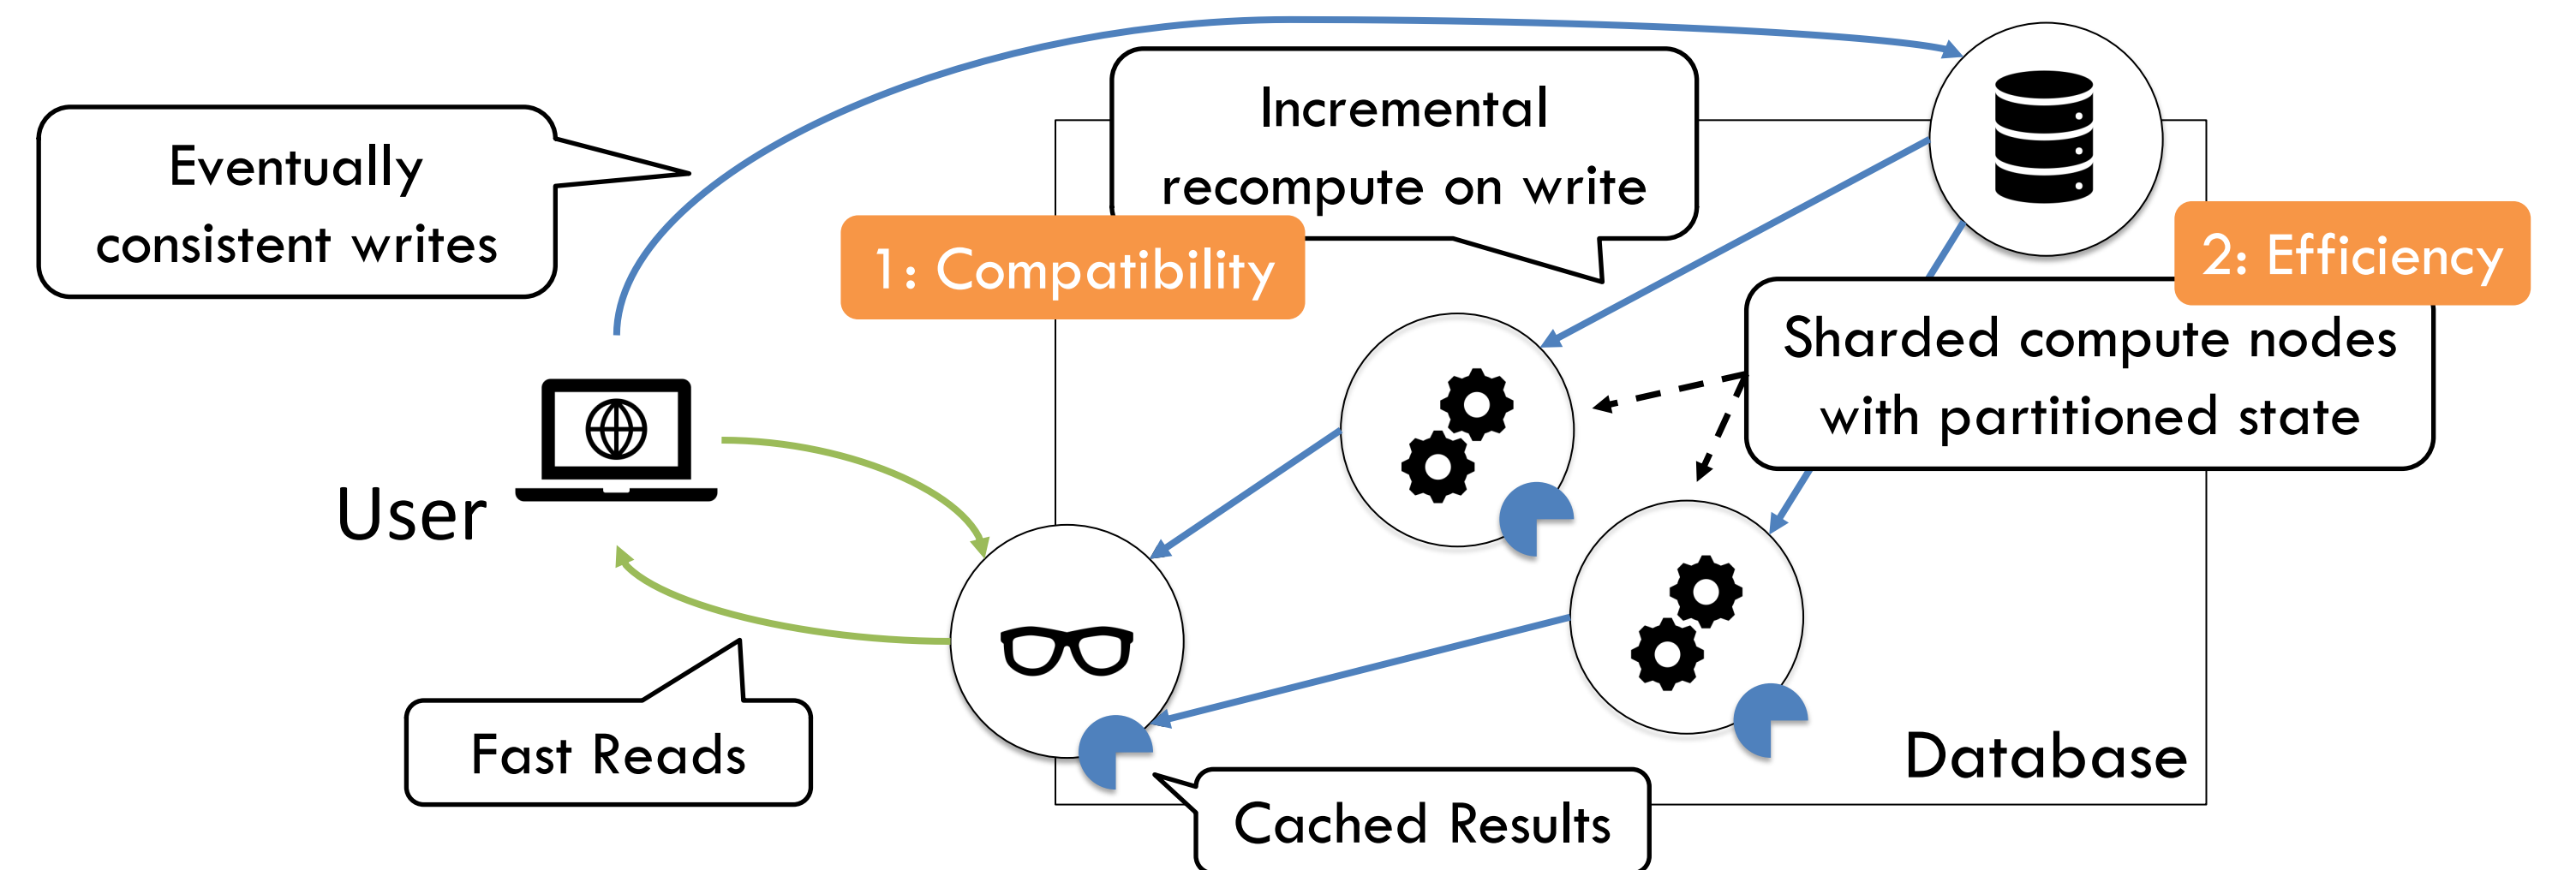
Contact: Justus Adam, Justus.Adam@tu-dresden.de

Research Challenges: User Defined Functions for distributed, incremental materialized views

- UDFs are a powerful extension point for databases
 - Third-party libraries, serialization, conversion
- Imperative source language with shared mutable state
- Table function and aggregation interfaces, are inherently stateful
- The more state is used, the harder it is to parallelize or shard



- Materialized views** offer fast read through aggressive caching
 - Incremental maintenance** makes the slow writes more efficient
- Challenge 1: Compatibility:** Support incremental computations for UDF
- Challenge 2: Efficiency:** Support sharding (parallelizing and distribution) the UDF or risk being a bottleneck to scaling

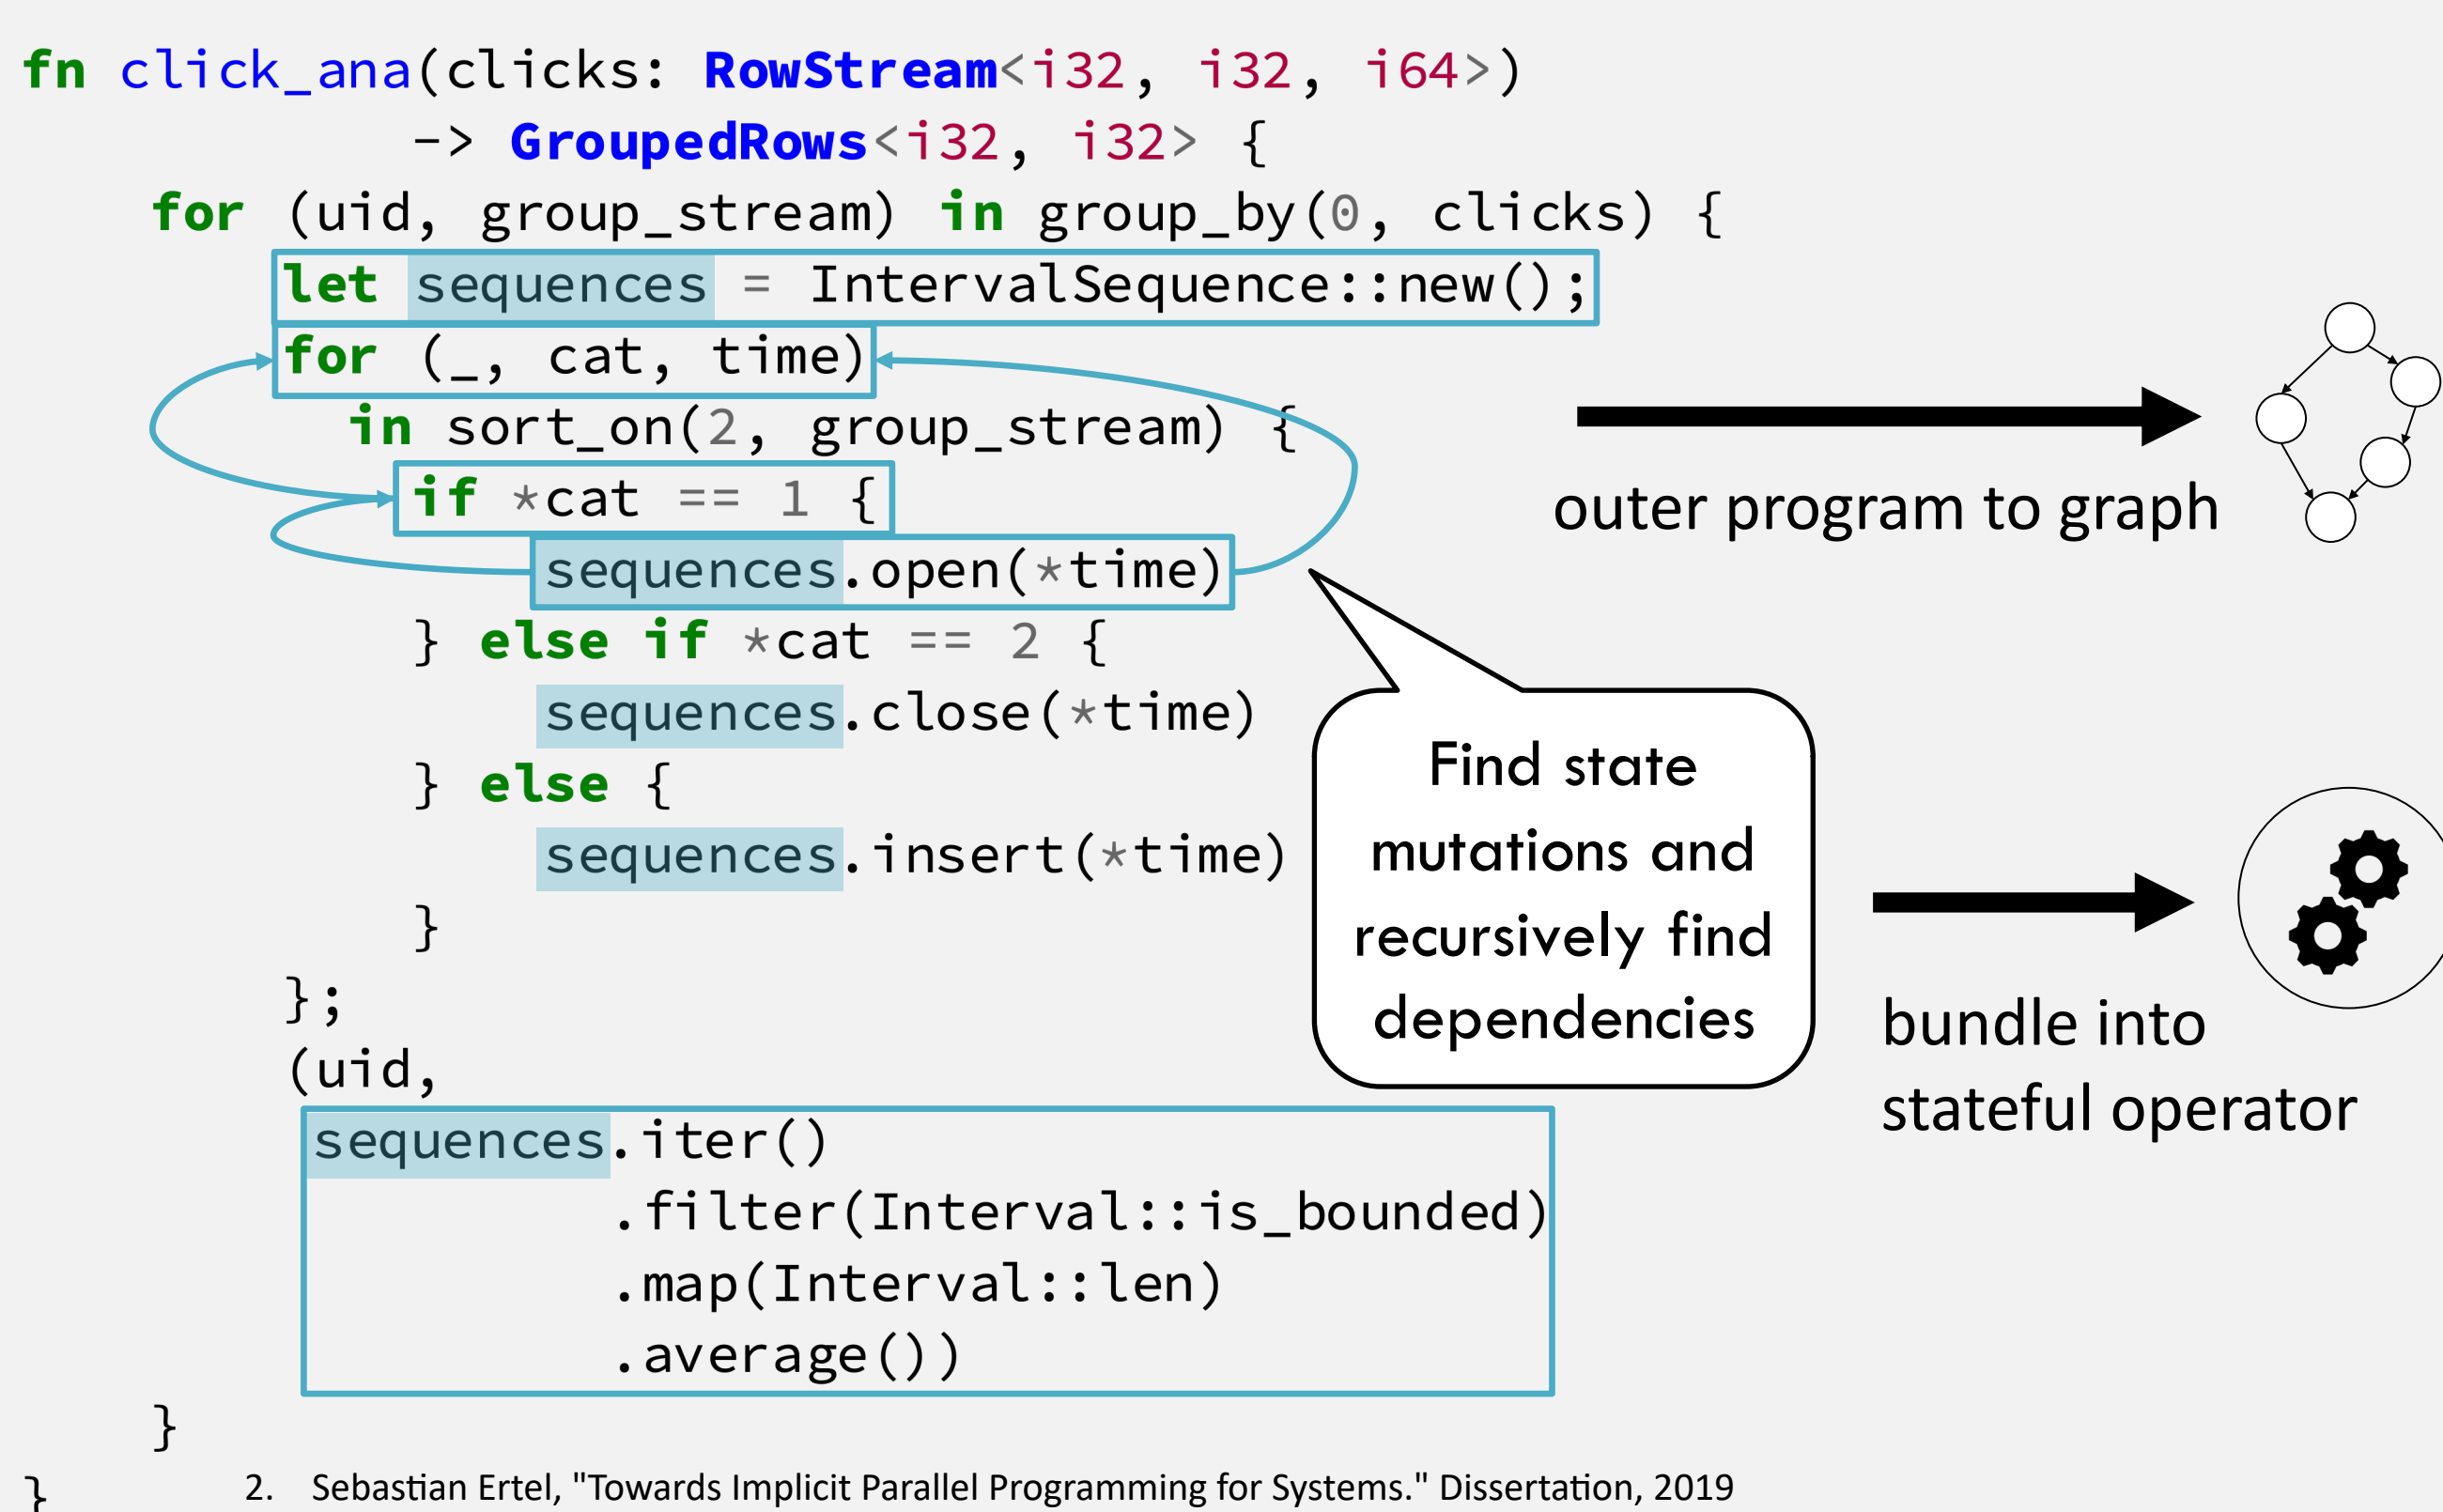


1. Jon Gjengset, Malte Schwarzkopf, Jonathan Behrens, Lara Timbó Araújo, Martin Ek, Eddie Kohler, M. Frans Kaashoek, and Robert Morris. 2018. Noria: dynamic, partially-stateful data-flow for high-performance web applications. *OSDI 18*.

Compilation Target and Strategy

- The compilation target, a Noria query, is a graph (network) of stateful operators
 - Operator state is private, not shared
 - Communication via message passing

- Use parallelizing compiler (Ohua^[2,3]) to split UDF program into
- Single "outer" program, *without shared state*, suitable for transformation into the **query graph**
 - Multiple "inner" programs, *using shared state internally*, suitable for forming the core of **stateful operators**



2. Sebastian Ertel, "Towards Implicit Parallel Programming for Systems." Dissertation, 2019

3. Sebastian Ertel, Justus Adam, Norman Rink, Andrés Goens, Jeronimo Castrillon, "STCLang: State Thread Composition as a Foundation for Monadic Dataflow Parallelism." *Haskell'19*.

Incremental Operator State

1: Compatibility

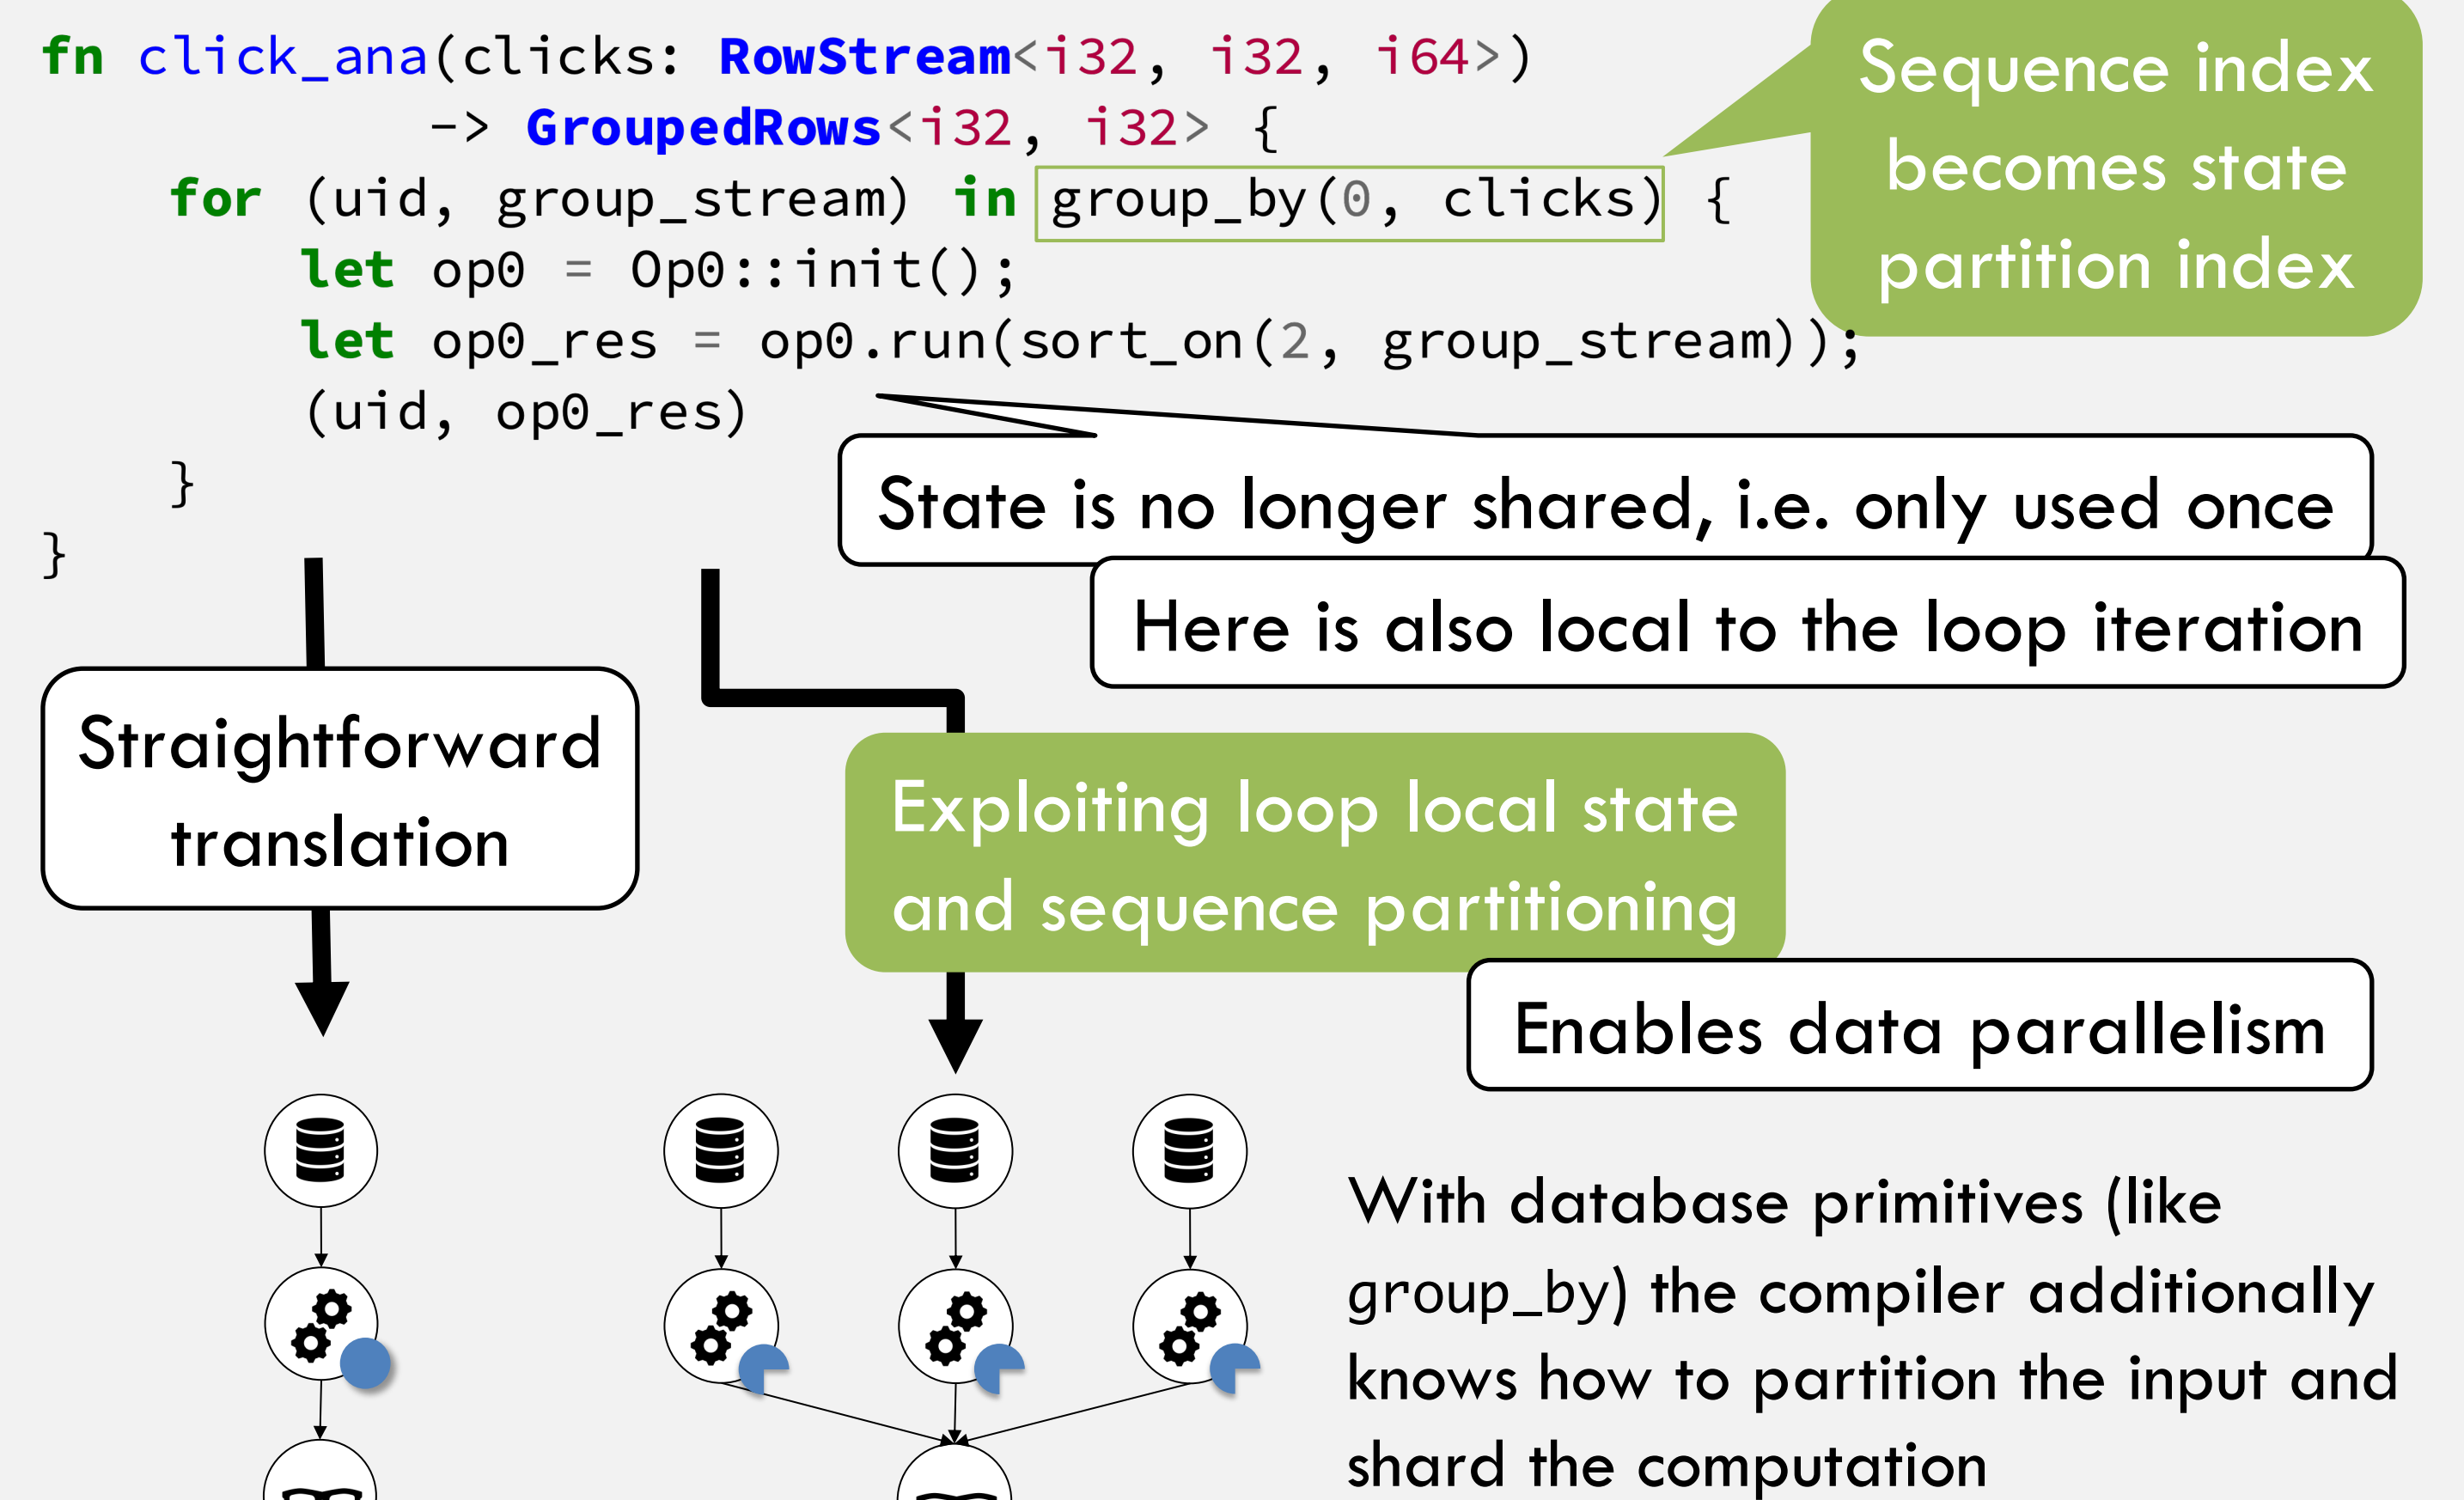
- We require custom operator state mutations to be reversible
- This allows stateful operators to be made incremental automatically

If a previously processed value is deleted, rerun operator computation but **revert** modifications instead of applying them.

Data Parallelism

2: Efficiency

Outer program after splitting




Preliminary Results

1: Compatibility Our novel technique can compile a subset of imperative, stateful Rust to incrementally maintained views in a dataflow engine

- Supports Single-tuple UDFs, aggregations, table functions and standalone queries

2: Efficiency Parallelism and sharding is implicit and effortless

- For our example query we are able to achieve near linear scaling up to 3 shards.



Diminishing returns after 3 shards are likely caused by orchestration overhead starting to dominate the small data size.