

Meet the Walkers

Accelerating Index Traversals for In-Memory Databases

Onur Kocberber¹, Boris Grot², Javier Picorel¹, Babak Falsafi¹, Kevin Lim³, Parthasarathy Ranganathan⁴

¹EcoCloud, EPFL

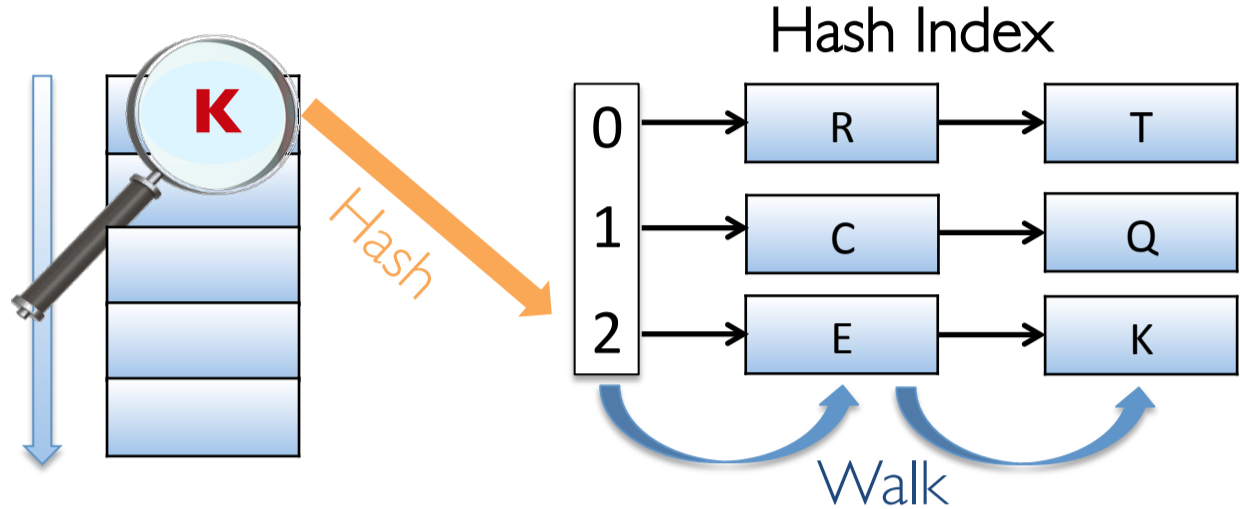
²University of Edinburgh

³HP Labs

⁴Google, Inc.

Our World is Data-Driven!


- Data resides in huge databases
- Indexes are essential for all database operators
- Hash index: a fundamental index structure



- Frequent use case: join via hash index

How Much Time is Spent Indexing?

Measurement on Xeon 5670 CPU with 100GB Dataset on MonetDB



Indexing is the biggest contributor to execution time

Dissecting Index Lookups

Hash: Avg. 30% time of each lookup

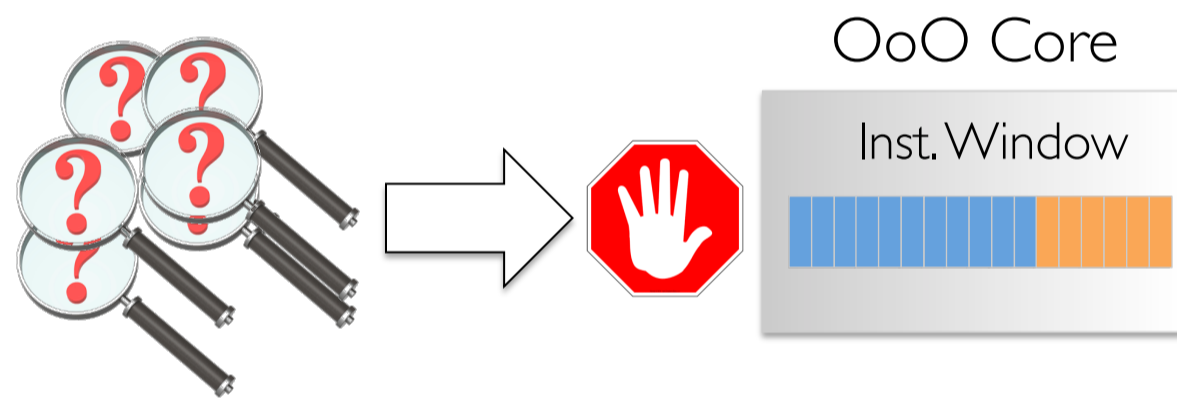
- Computationally intensive, high cache locality

Walk: Avg. 70% time of each lookup

- Trivial computation, low cache locality

Next lookup: Inherently parallel

- But, beyond the inst. window capacity



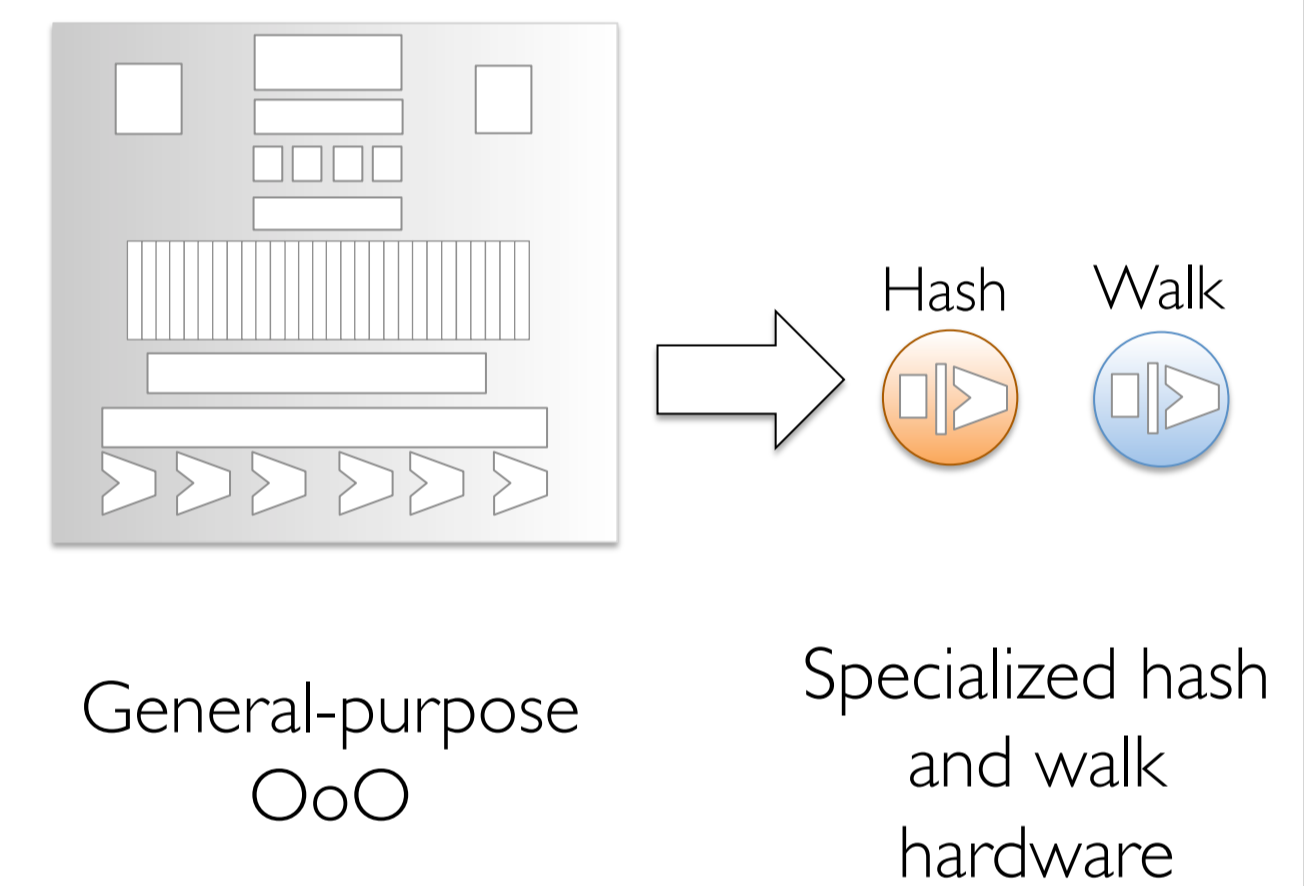
Roadmap

Roadmap for efficient and high-throughput index lookups


1. Specialize
 - Customize hardware for hash and walk
2. Parallelize
 - Perform multiple index lookups at a time
3. Generalize
 - Use a programmable building block

1. Specialize

Design a dedicated unit for hash and walk



2. Parallelize



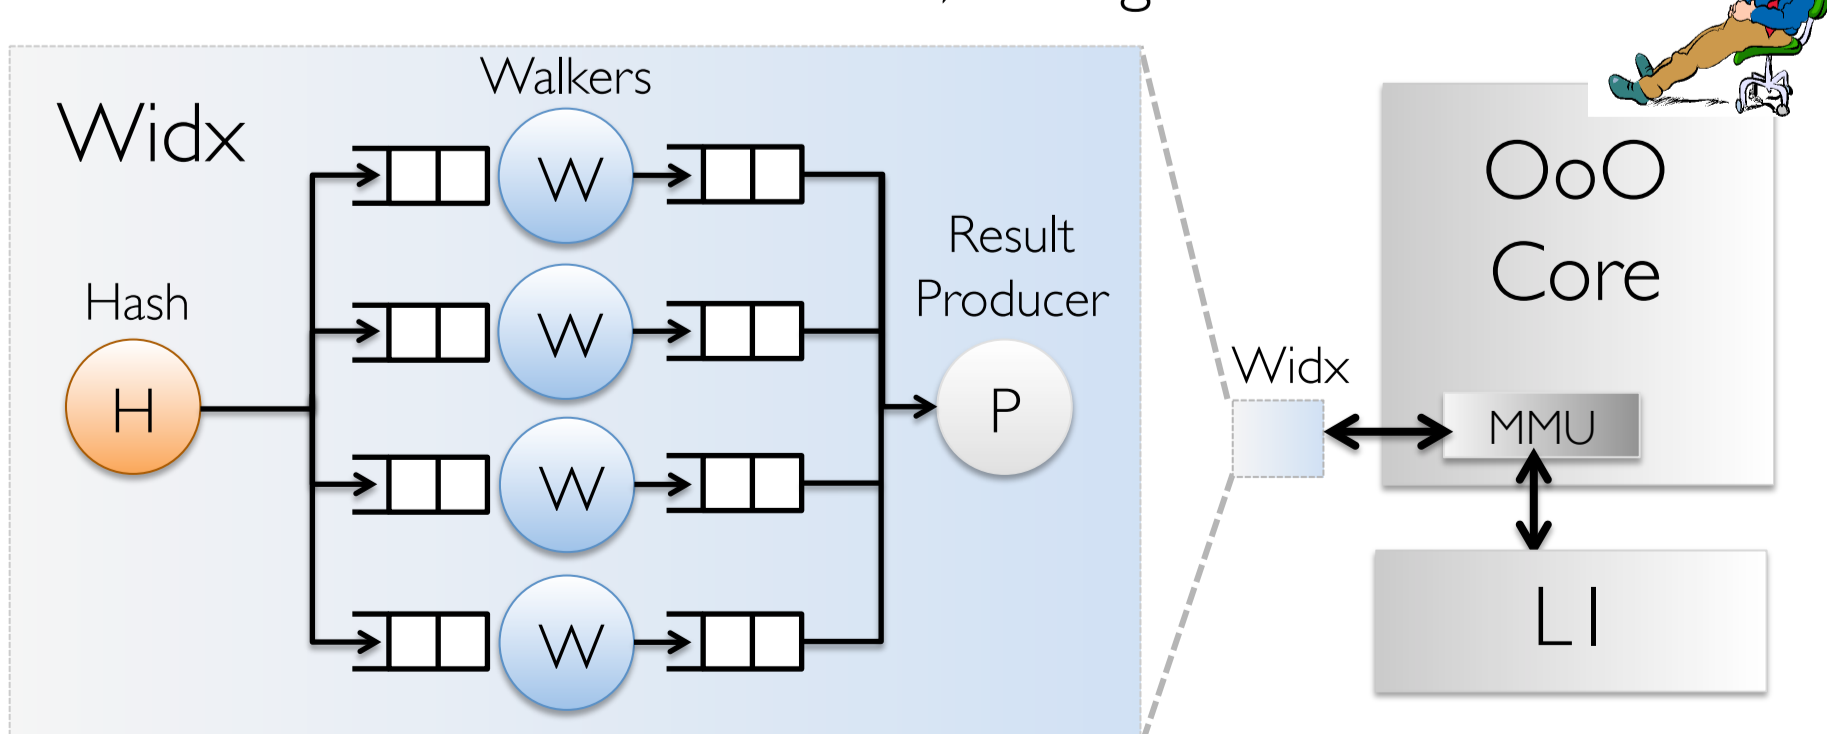
3. Generalize

- Widix unit: common building block for hash and walk
 - Two-stage RISC core
 - Custom ISA
- Widix units are programmable
 - Execute functions written in Widix ISA
 - Support limitless number of data structure layouts



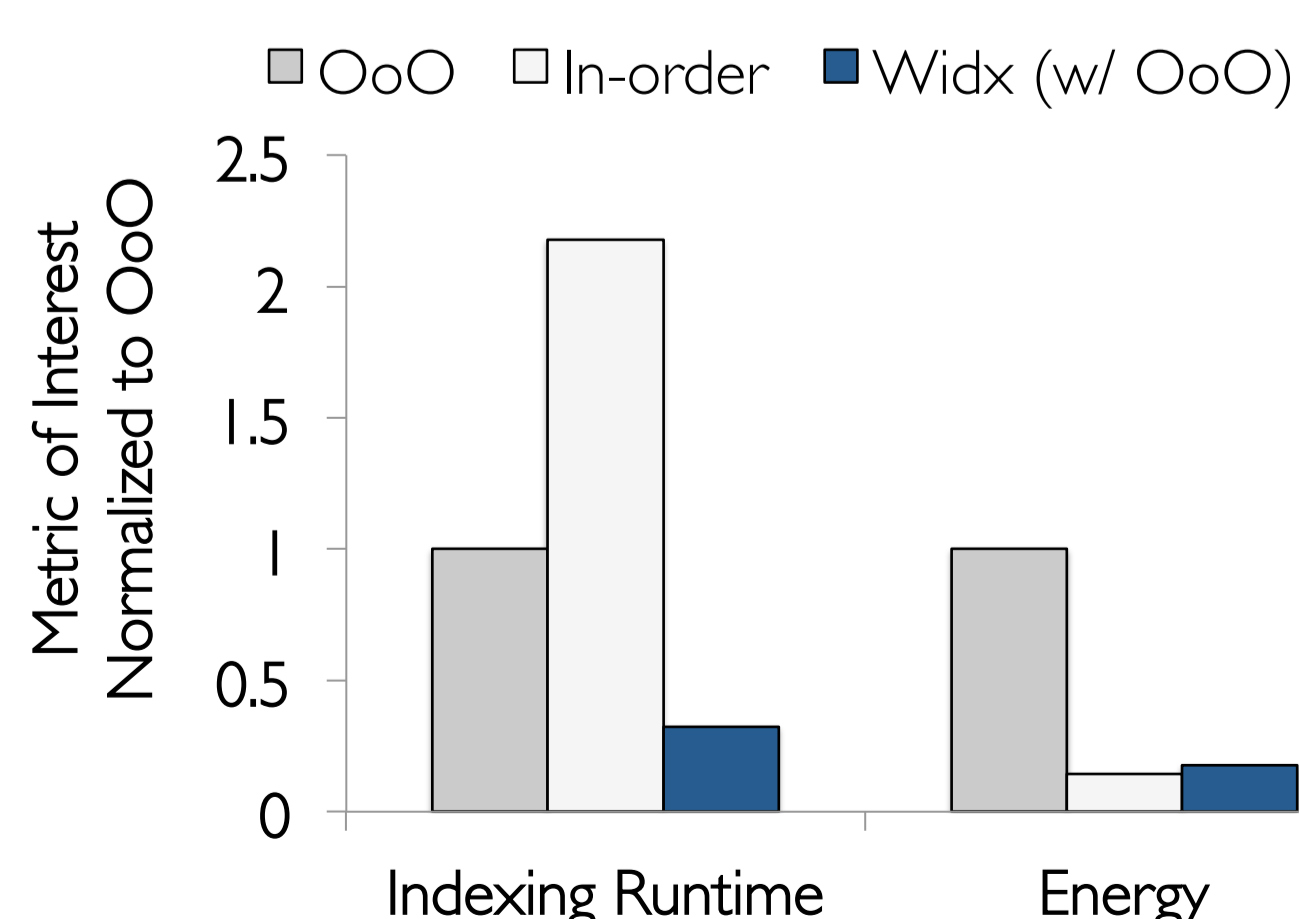
Putting it all together:Widix

When Widix runs, core goes idle



Simple, parallel hardware

Result Highlights



More details in the paper!

- Per query results
- Join kernel evaluation
- Programming model
- Bottleneck analysis
- Widix cycle breakdowns