

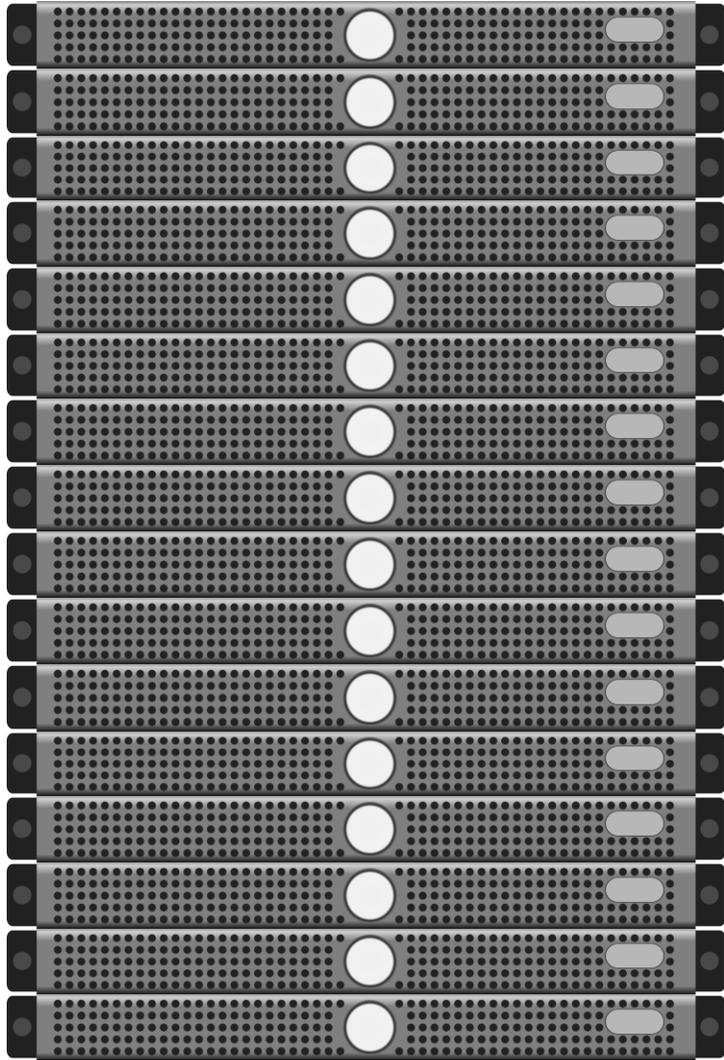
Microsoft Research
Faculty
Summit
2016



Towards a “scaling-in” approach for data processing

Jignesh M. Patel
University of Wisconsin





~2006

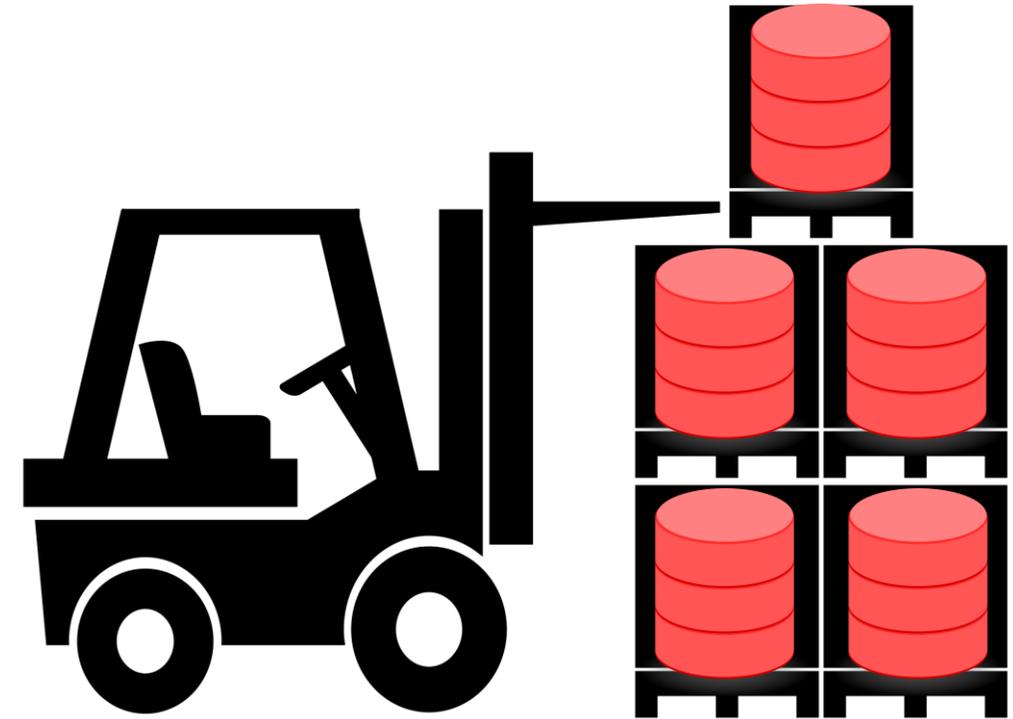
Azure G5
32 Cores
448GB RAM
\$8.69/hr



2016

Azure
D15_{v2}
20 Cores
140GB RAM
\$1.253/hr

Scope: BI
Analytics



Data Warehouse



Efficiently



Everywh
e

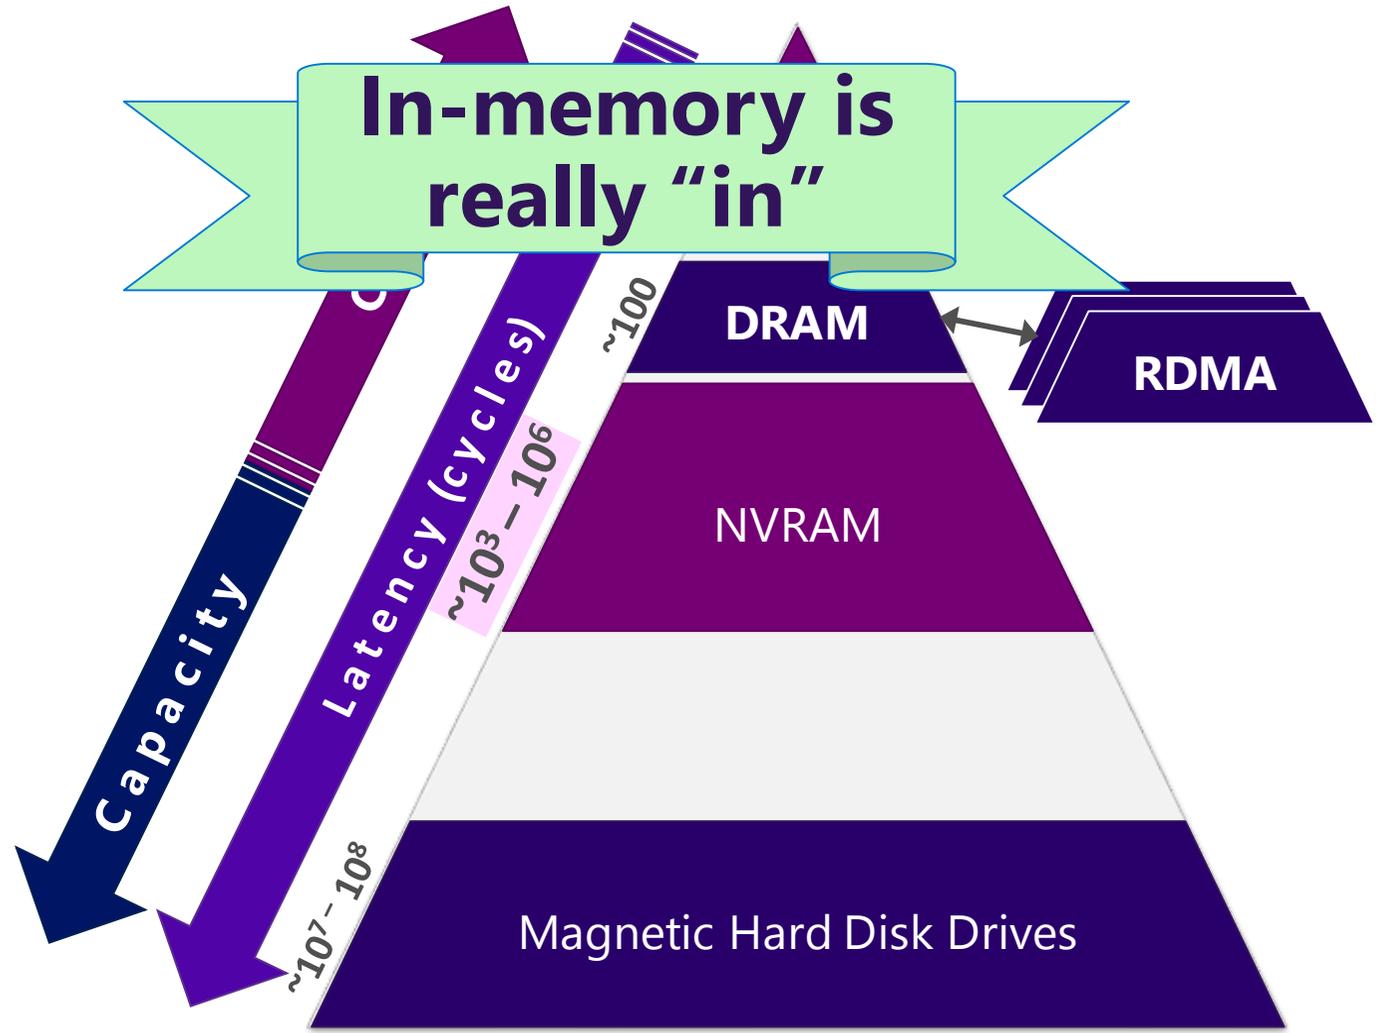


Effortlessly

Efficiently

Everywhere

Effortlessly



2016 and beyond

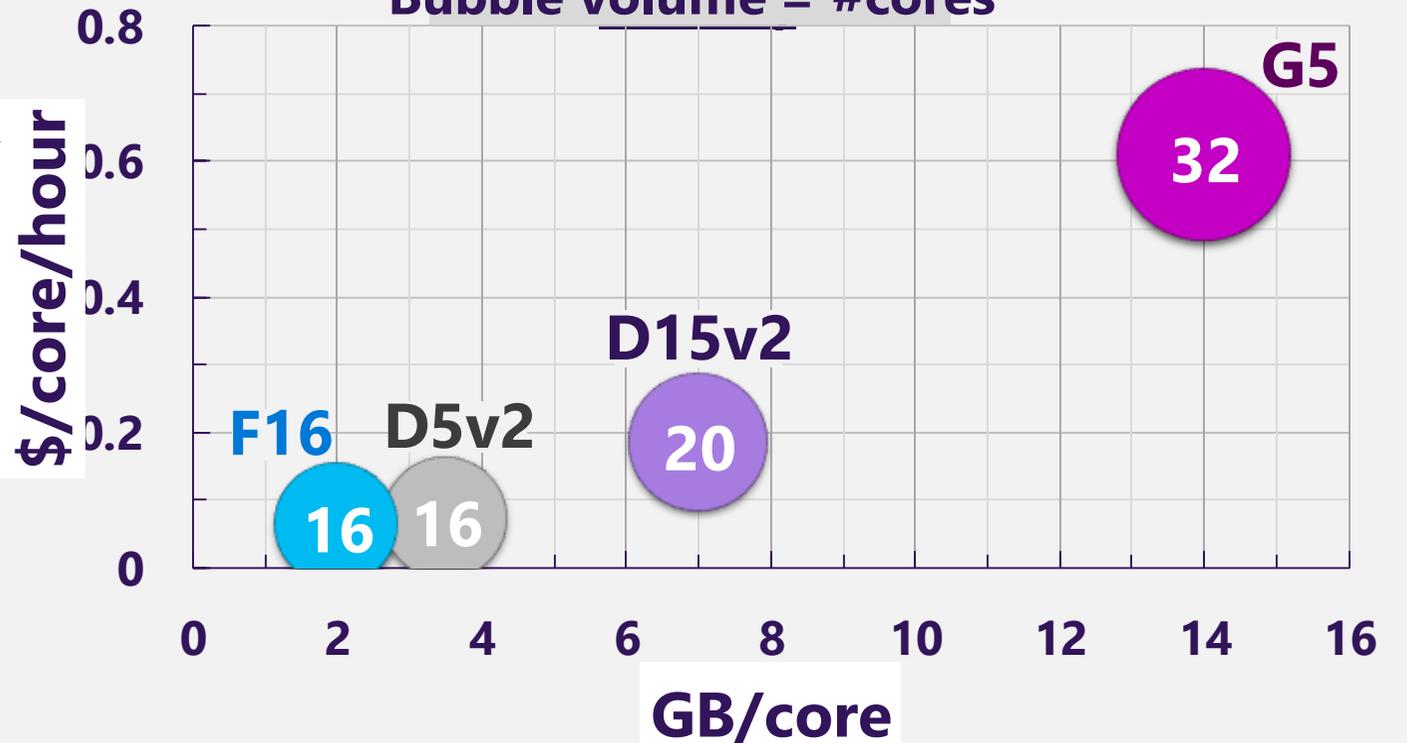
Efficiently

On modern hardware

many cores,
large
memory
configurations

Everywhere

Azure: Pricing vs Configuration (Series: D, F)
Bubble volume = #cores

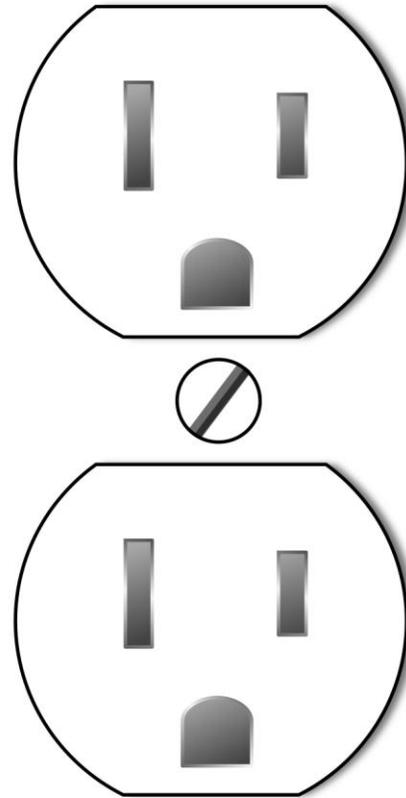


Effortlessly

Efficiently

Everywhere

Effortlessly



many cores,
large
memory
configurations

infinite and
instant
elasticity

Efficiently

On modern hardware

many cores,
large memory configurations

Everywhere

Containers of any size

infinite and instant elasticity

Effortlessly

Auto-tuning

a true utility model

Storage Manager Design

A database = A collection of tables

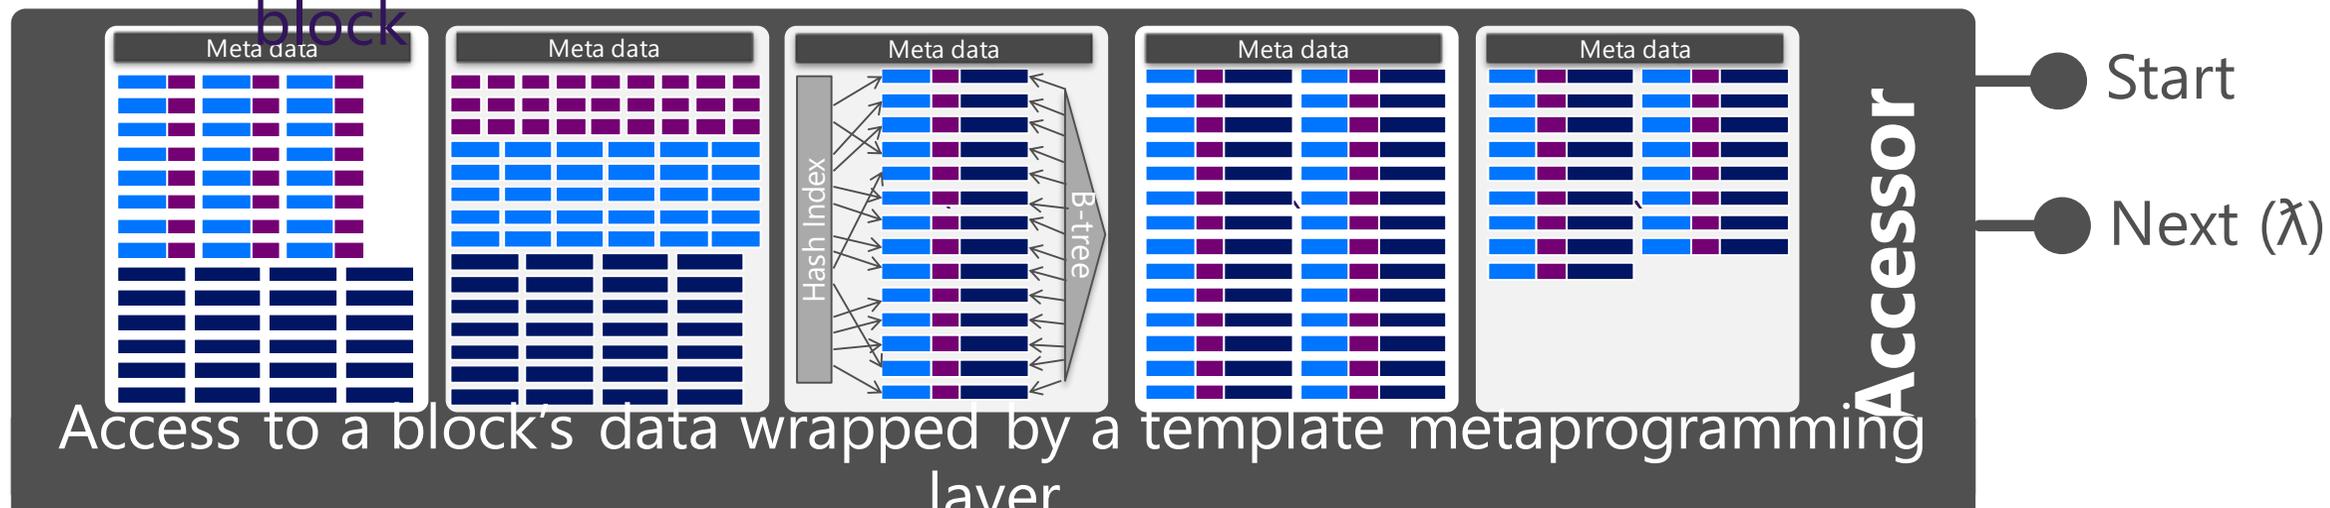
A table = set of blocks

A block = bag of tuples

A block = a mini self-contained database. Only tuples from one table per

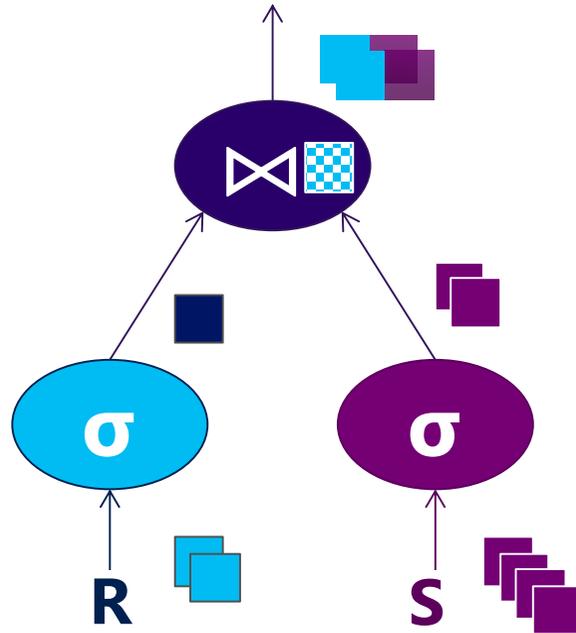
A sub-block = a storage organization within a block
e.g. row sub-block, column sub-block, or index sub-

block



Query Execution

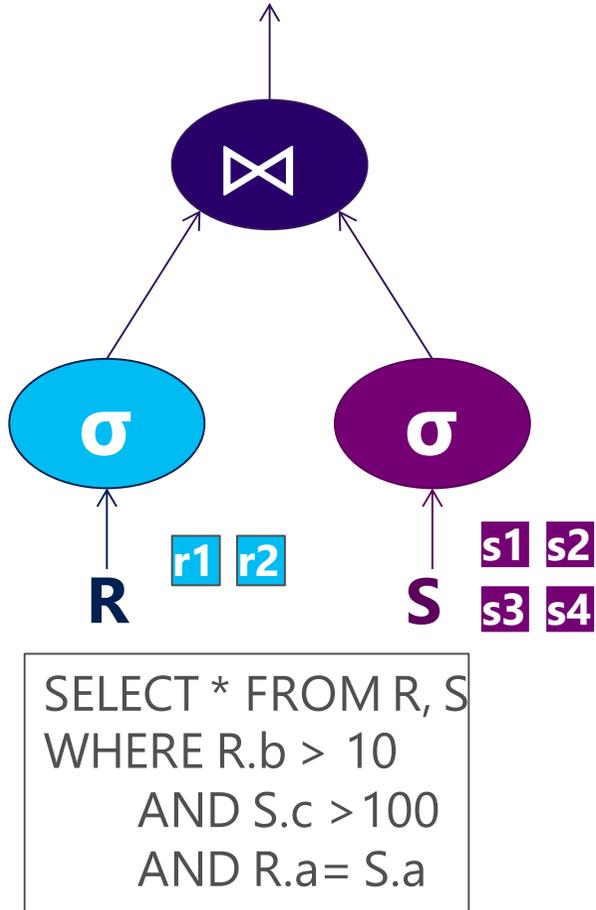
The "traditional" way



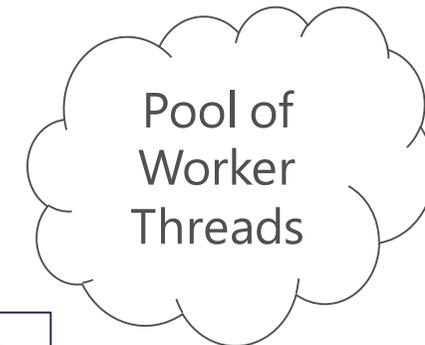
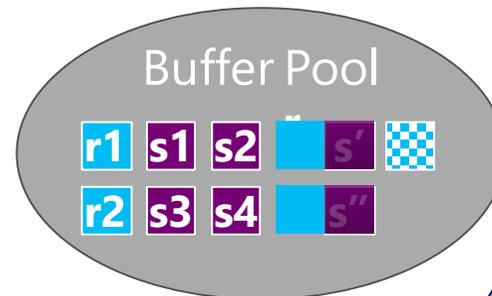
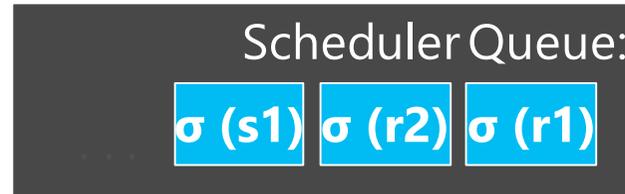
```
SELECT * FROM R, S
WHERE R.b > 10
      AND S.c > 100
      AND R.a = S.a
```

Query Execution **Clean Separation of Data Flow and Control** Flow

The "traditional" way



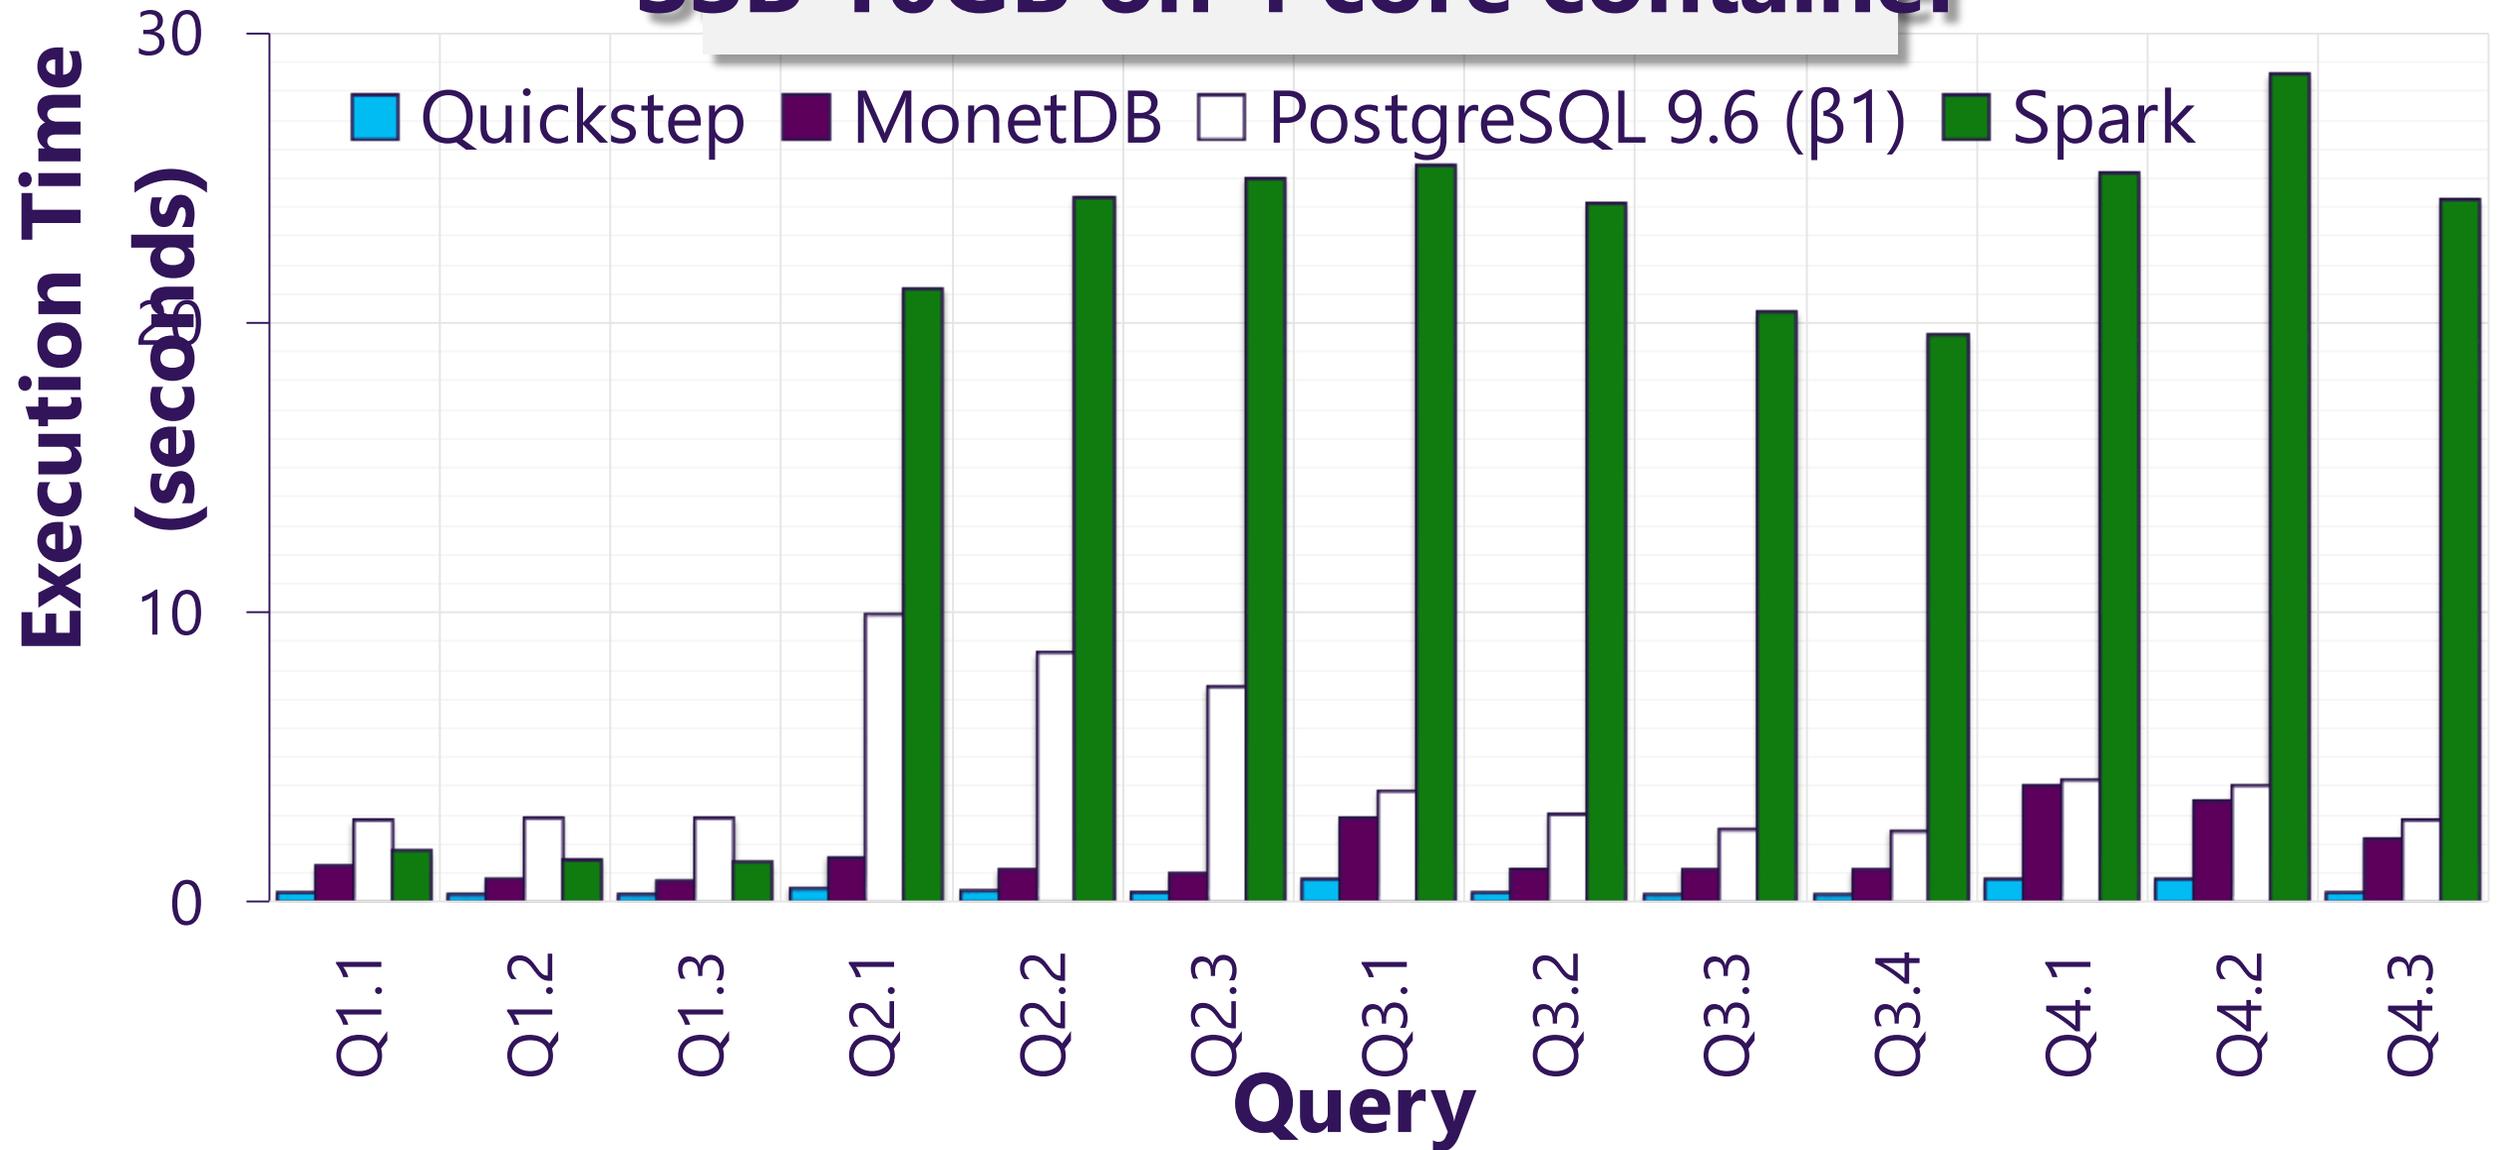
The Quickstep way



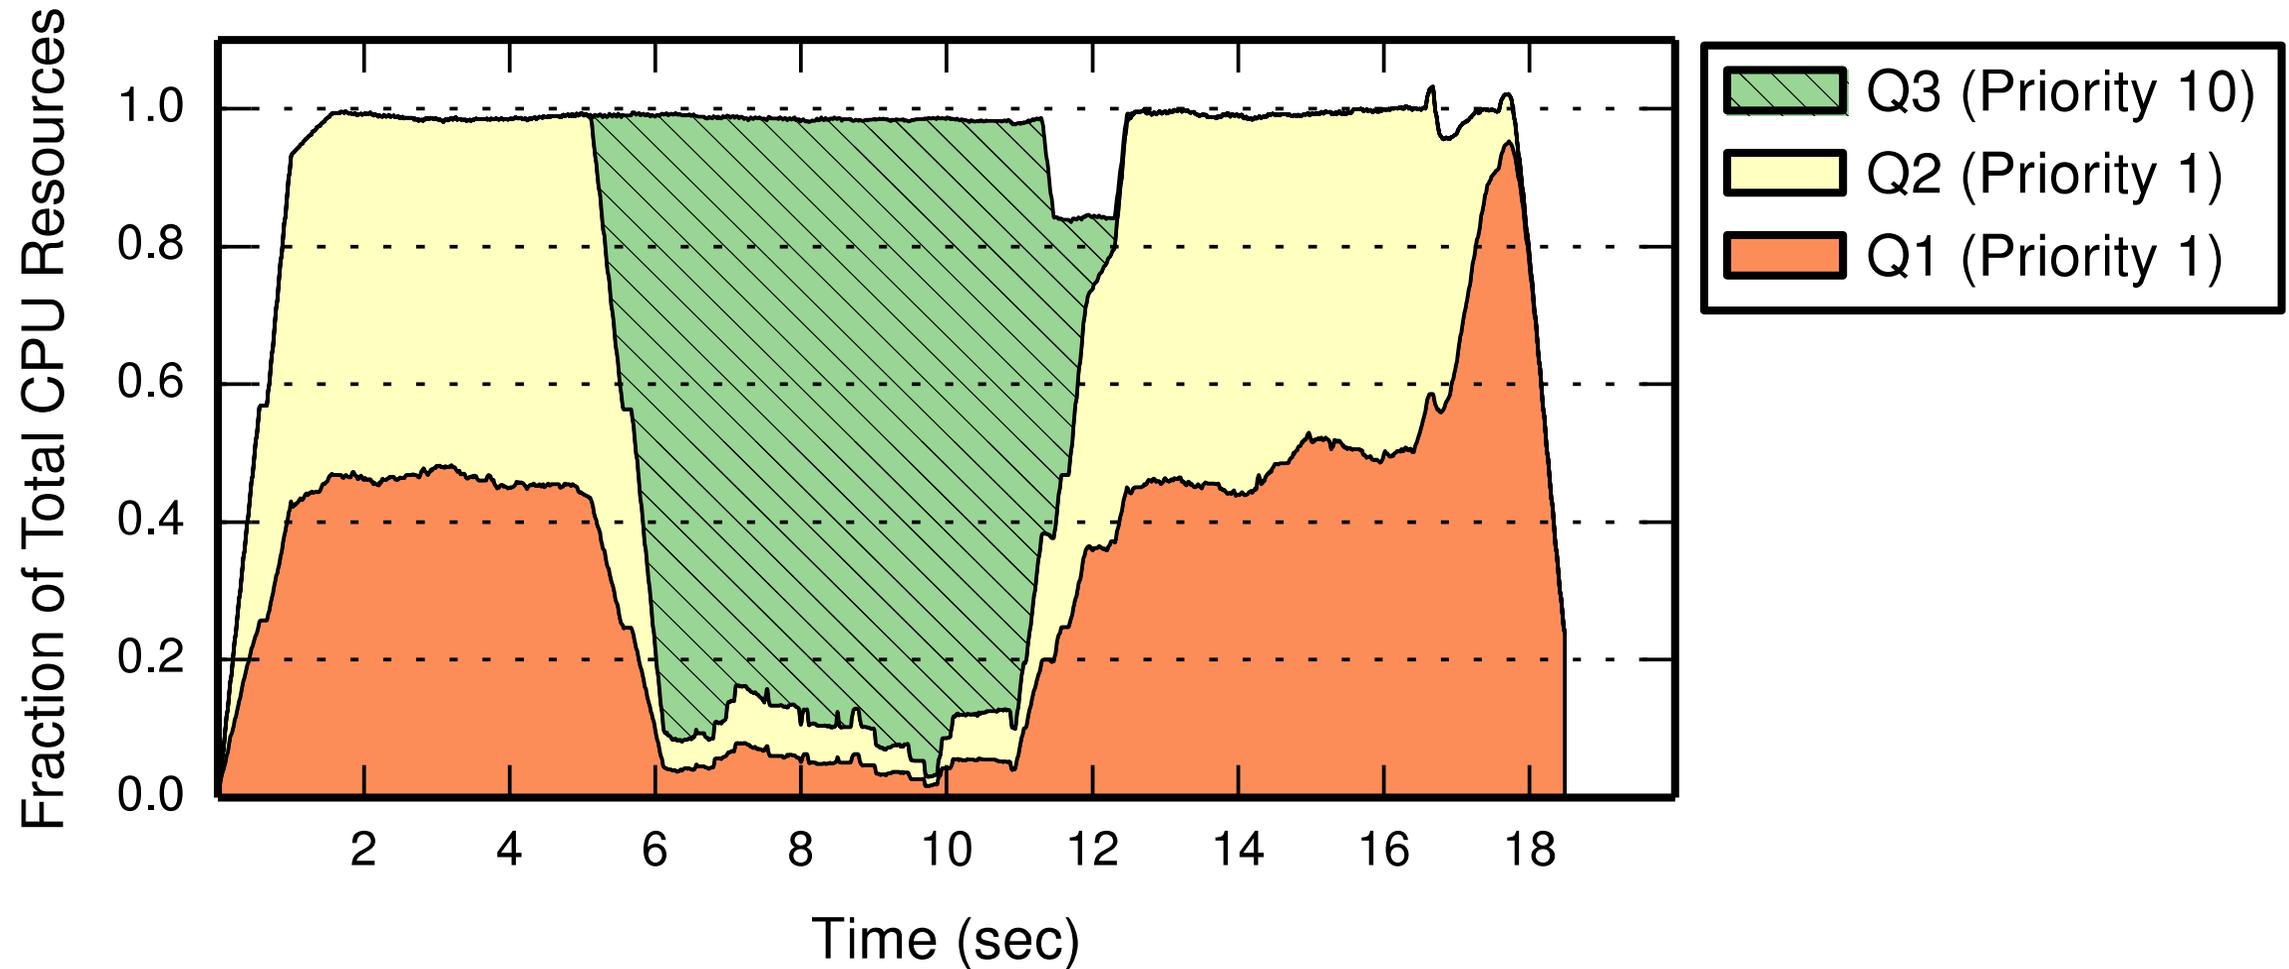
- Advantages
- + Cleaner Abstraction
 - + Dynamic Optimization
 - + Better p9X
 - + Manageability and Debug-ability



SSB 10GB on 4 core container



Priority scheduling: Example of elastic behavior



Efficiently

Bare metal performance on core data kernels, **and on any hardware**

Everywhere

elasticity for inter and intra-operator parallelism, **even for complex operators**

Effortlessly

Unified memory management, hybrid storage formats, **learning-based scheduler/optimizer**

Quickstep

An experimental data platform for the cloud-native data world

www.quickstep.io

Thanks!



**Harshad
Deshmukh**



**Rogers
Jeffrey**



**Hakan
Memisoglu**



**Navneet
Potti**



**Saket
Saurabh**



**Marc
Spehlmann**



**Subasree
Venkatsub
h-
ramaniyen**



**Zuyu
Zhang**



**Jianqiao
Zhu**

