CWI

Martin Kersten

Lefteris Sidirourgos

A Database with Amnesia

monetdb

CWI

The Business Data Scientists

monetdb

CWI

Their View of Our World

monetdb

CWI

Their View on our Happiness

monetdb

CWI

Their View on our work

monetdb

CWI

there will be too much data to handle, process or even “looked at” within any given budget and time constraints

monetdb

CWI

Lets' start forgetting data

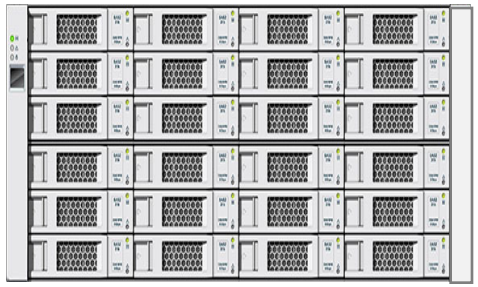

DBMS Architects

monetdb

CWI

Fundamental Design Principal

A Database should Retain Information

No Data Loss is Allowed

We can Forget Data and still Retain Information

monetdb

CWI

Data rotting

The DBMS may selectively forget data on its own initiative for the sake of storage management and responsiveness.

monetdb

CWI

The food lifecycle

rotting ⟷ refine

monetdb

CWI

Lesson 1:

Don't collect more Data than you can eat.

monetdb

CWI

Lesson 2:

Purify and refine *Data* makes consumption pleasant and lasting

monetdb

CWI

Lesson 3:

Data rotting is an evitable natural phenomena, learn to deal with Amnesia

monetdb

CWI

Amnesia

- Strategies to Forget Data while retaining information
- Depends on
  - Data Semantics
  - Data Distribution
  - Query Semantics
  - Update/Query Frequency

monetdb

CWI

Data loss & Information Retention

- How much data loss could there be?
- How much information is lost?
- What is the error and confidence level?
- How bad can it get?

monetdb

CWI

An DB Amnesia Simulator

- columnar DBMS in C
- Integers R=0..DOMAIN
- Data distributions
  - serial
  - uniform
  - normal
  - skewed (Pareto)
- A tuple may be active or forgotten
- DB has always constant size DBSIZE

monetdb

CWI

Query Workload

- SELECT-PROJECT-JOIN
- Range Selections
  - selectivity S=1.0 .. 0.01
- Aggregations
  - avg
  - sum
- Data Distribution and Query format lead to different forms of amnesia

monetdb

CWI

Information Retention Metrics

After F new tuples and forgetting F other tuples:

- $R_F(Q)$ number of tuples in query Q
- $M_F(Q)$ number of tuples missed in query result Q
- $P_F(Q)$ query precision
  - $P_F(Q)=R_F(Q)/(R_F(Q)+M_F(Q))$
- E is the error margin
  - $E=avg(R_F(Q))/avg(R_F(Q)+M_F(Q))$ computed over an entire batch of queries Q

monetdb

CWI

Temporal Biased Amnesia

Query Based Amnesia

Spatial Based Amnesia

. . . . . .

monetdb

CWI

Temporal Biased Amnesia

- FIFO-amnesia
- Reservoir based / uniform at random
- Retrograde
- Anterograde
  - prioritize historical data

monetdb

CWI

Query Based Amnesia

- Tuples that appear to often in query results are important and should not be forgotten
  - but new tuples are equally important
- Opposite, too frequent tuples are noise
  - consume or be forgotten

monetdb

CWI

Spatial Based Amnesia

- Infect areas with rot
  - randomly choose to spread rot or create new rot
    - 1..k spread
    - k+1 new
- HD Data Degradation is correlated with data interest

monetdb

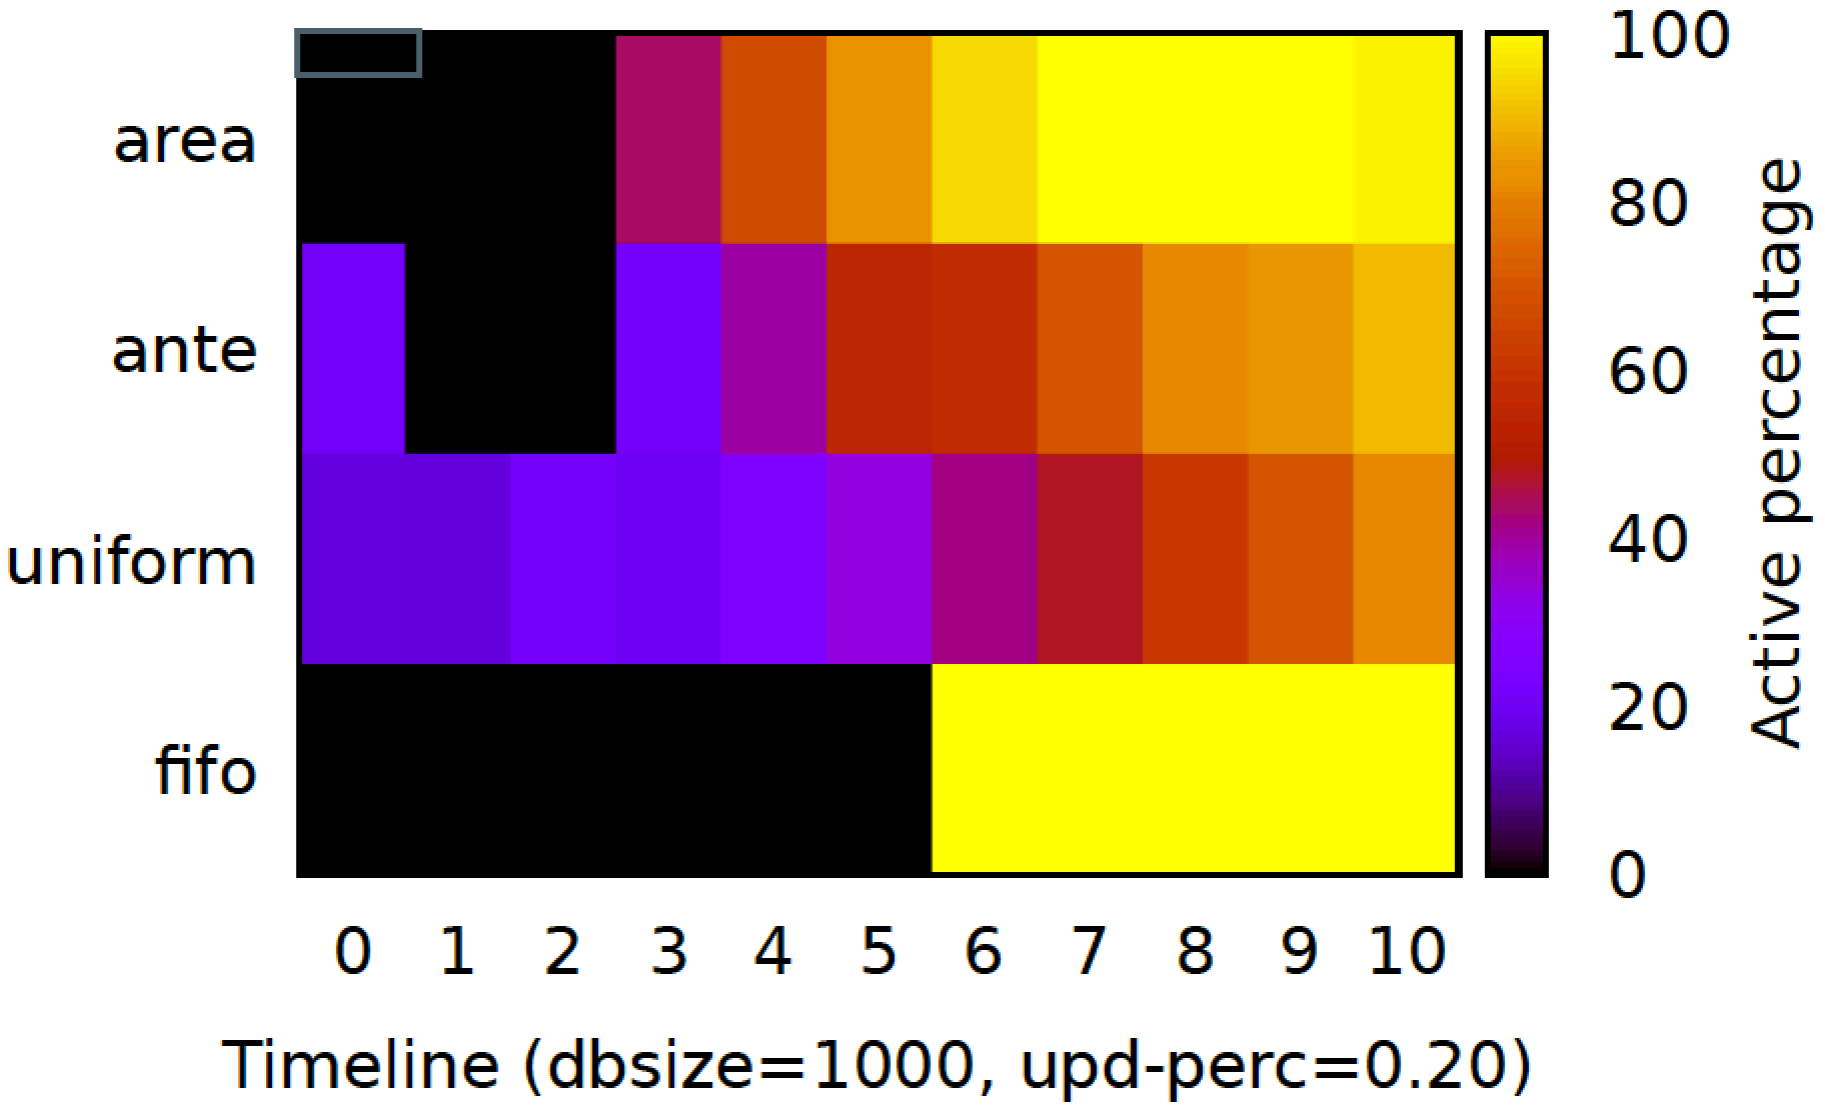

Figure 1: Database amnesia map after 10 batches of updates

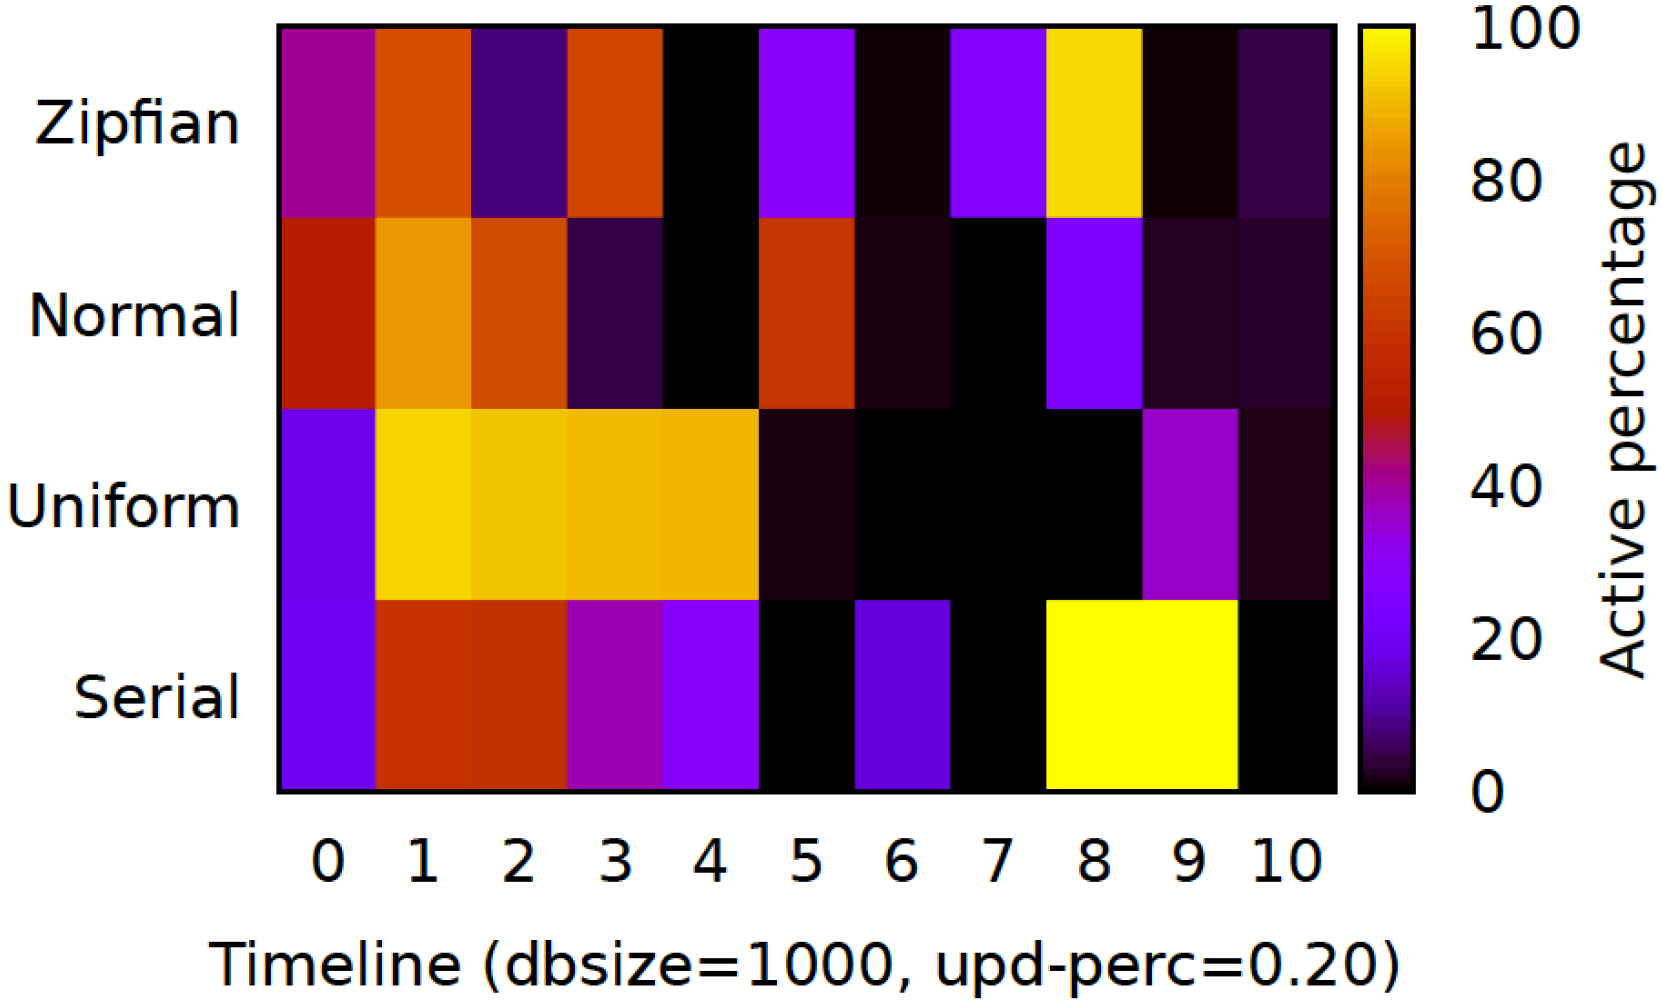

Figure 2: Database rot map after 10 batches of updates

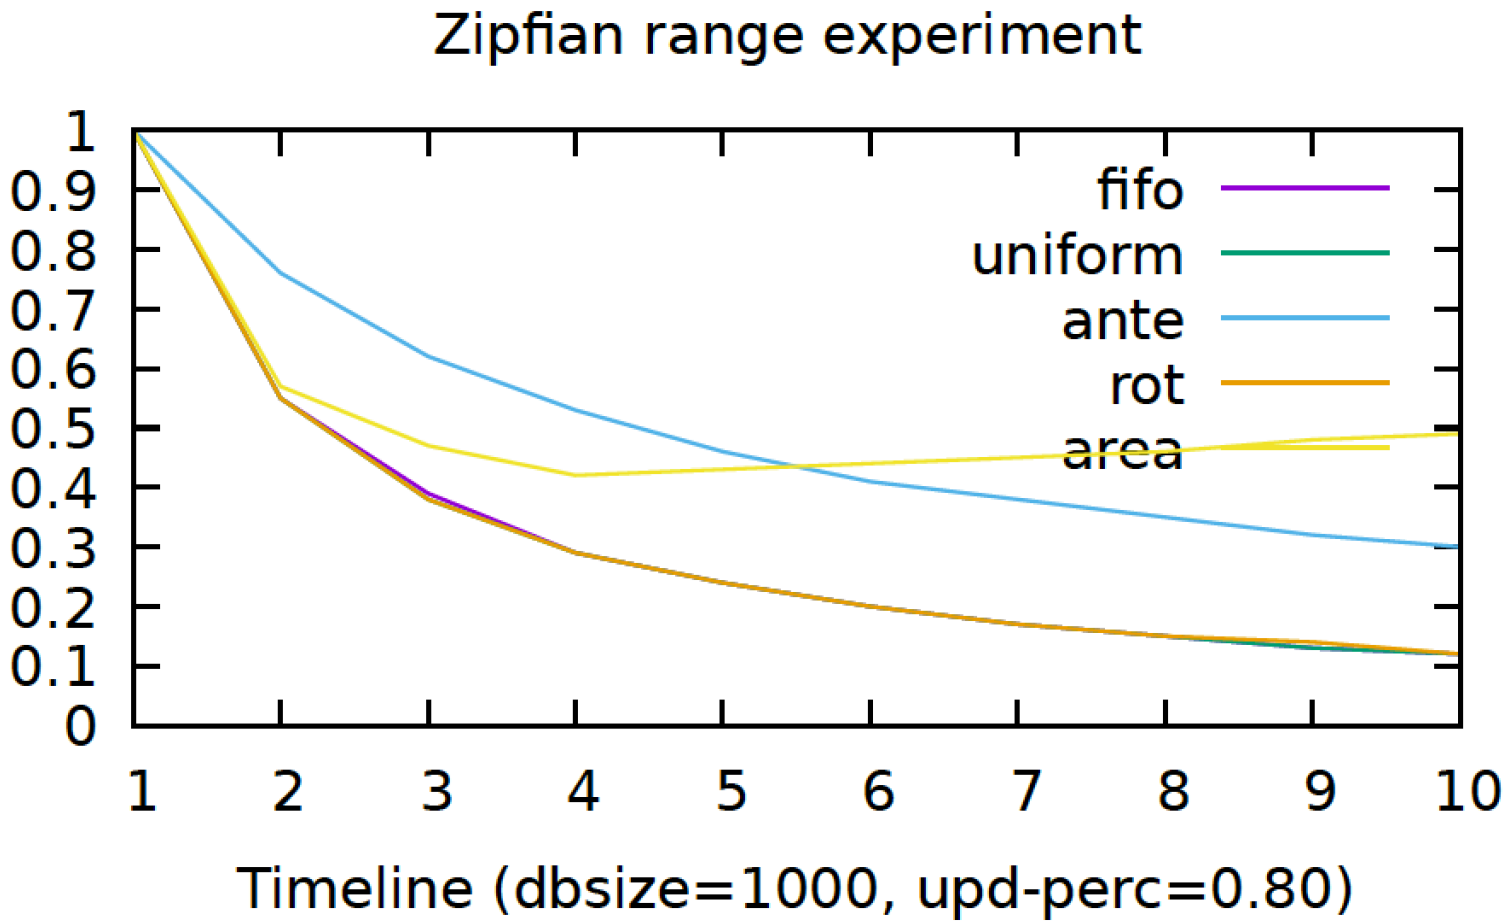

Figure 3: Range query precision ($v \in 0 .. max$)

CWI

Afterlife of Forgotten Data

- delete them!
- move them out of the critical path of query execution
- stop indexing them

monetdb

CWI

TAKE HOME MESSAGES

- Database amnesia techniques is a barren research landscape
- Prepare the end-users to cope with their dementing database
- Providing medicines with clear cost/effectiveness
- Re-asses all components of a DBMS to implement amnesia
  - Schema rotting
  - Query execution rotting
  - Index storage rotting
  - Record storage rotting
  - Operating system rotting,
  - Hardware rotting

monetdb

THANK YOU, ENJOY YOUR RESEARCH ADVENTURES

CWI

Hardware rotting

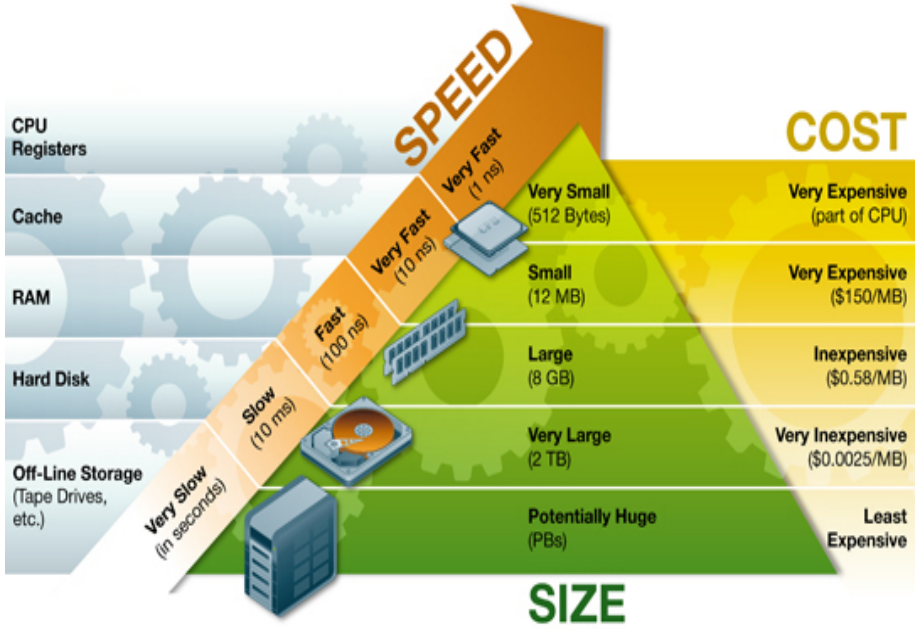

Hardware failures are a blessing, not a curse, as long as you can recognize them

monetdb

The OOM killer on Linux kicks in when it is faced with an out-of-memory condition.

Record rotting

Tuples/pages/partitions that have not been accessed over a long period become the target for automated vacuum actions.

Index rotting

Index structures, such as hashes and B-trees, are automatically capped to consume less space at the cost of re-growing its branches

Every data item contributing to a query result set is removed and no derived object can be larger then the contribution set

CWI

The medicine against rotting is to refine/purify

monetdb

Data fungi architecture