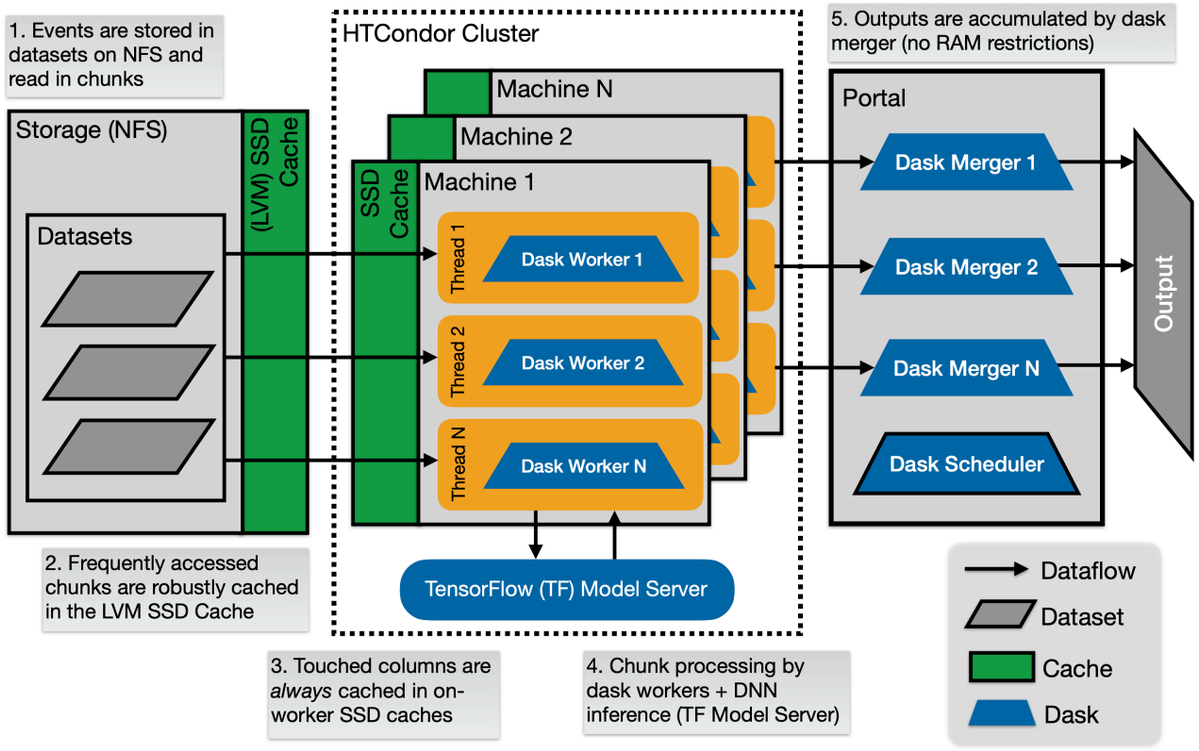

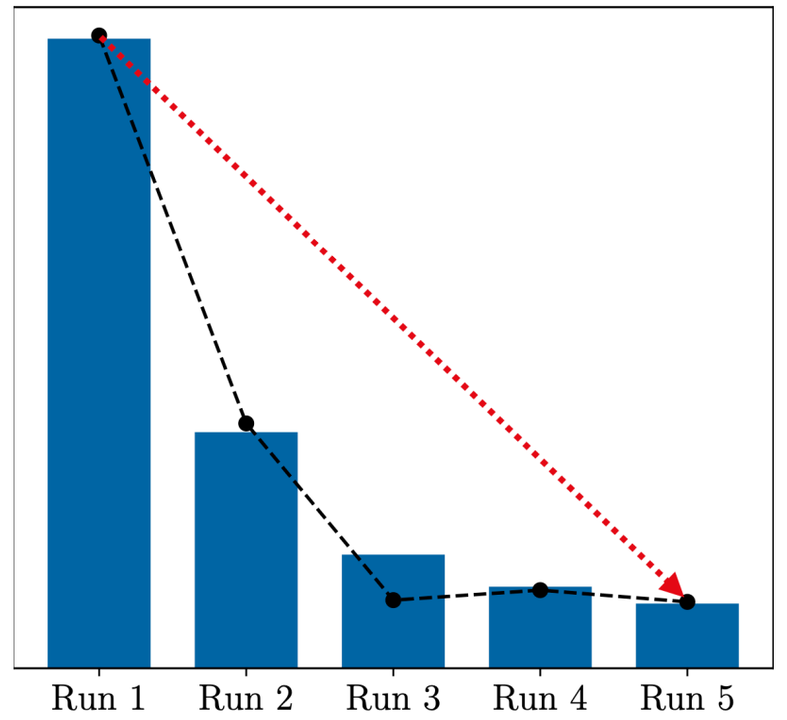

Going fast on a small-size computing cluster

Martin Erdmann, Peter Fackeldey,
Benjamin Fischer, Dennis Noll

FIDIUM - Kickoff

16.12.21

RWTH AACHEN UNIVERSITY

Starting Point

RWTH AACHEN UNIVERSITY
Peter Fackeldey - 16.12.21

Typical LHC Analysis:

- Input: Datasets $\mathcal{O}(10)$ TB
- Output: Histograms $\mathcal{O}(1)$ GB
- Complex:
  - Many physics objects
  - Deep learning discriminators

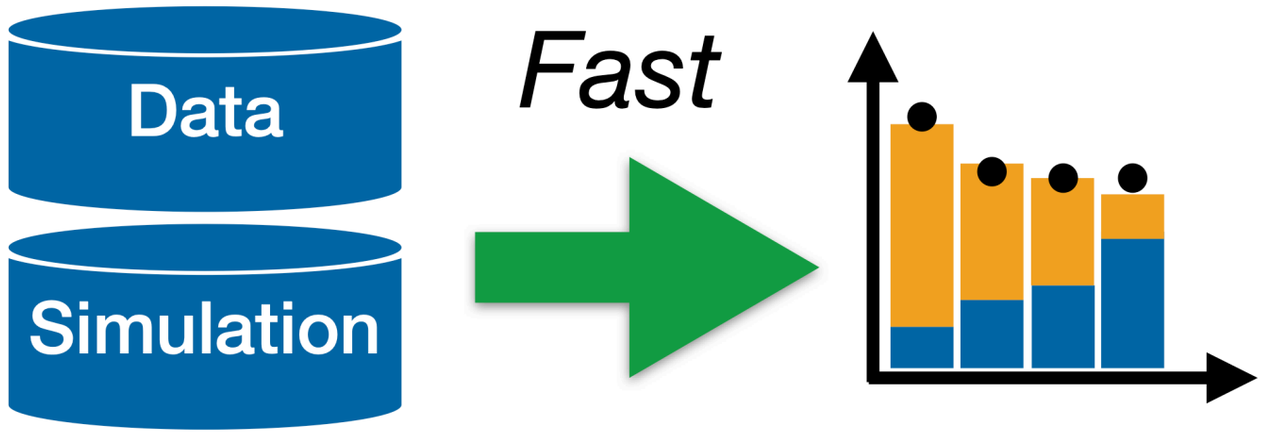

Team:

- 1 Professor
- 5 PhD Students
- 1 Master Student

Software:

- NumPy + Scikit-HEP
- Dask
- HTCondor

Hardware:

- 245 Threads
- 22TB SSD Caches
- 72TB Disk Storage

Computing Strategy:

MapReduce on *columnar* data
accelerated by caching

Dataflow & Cluster sketch

RWTH AACHEN UNIVERSITY
Peter Fackeldey - 16.12.21

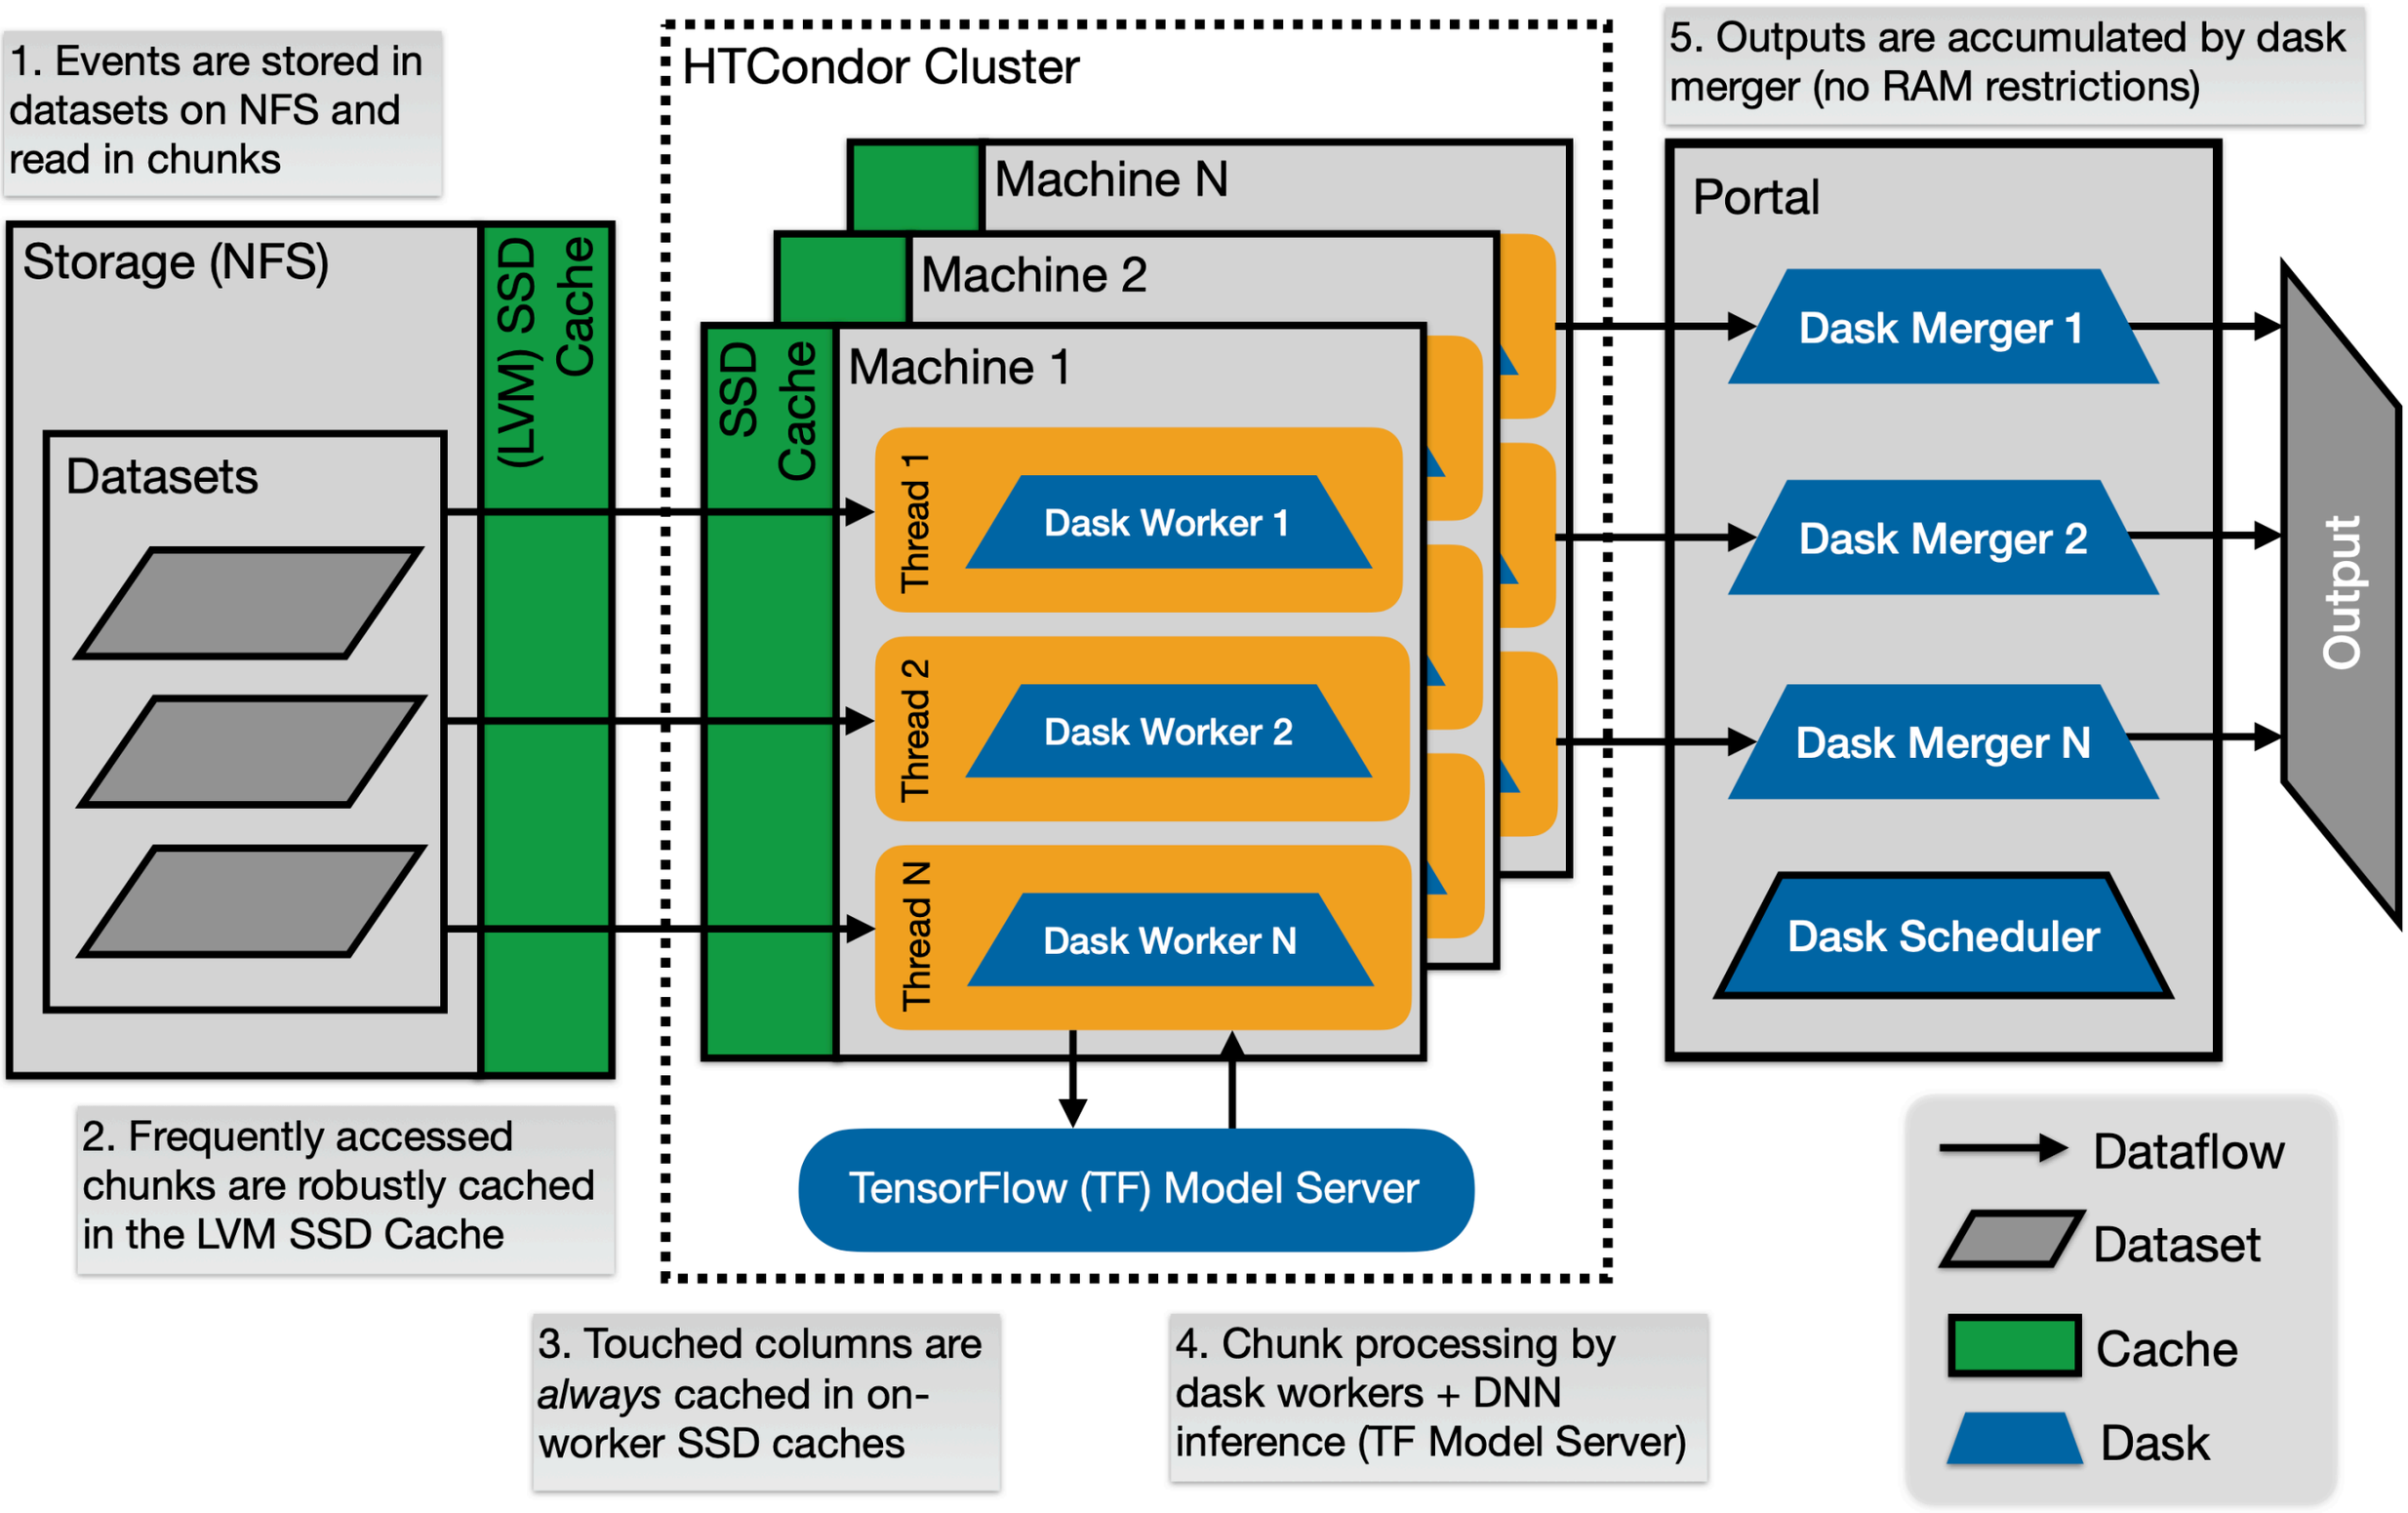

Results and Conclusion

RWTH AACHEN UNIVERSITY
Peter Fackeldey - 16.12.21

Benchmark Results:

- $4.18 \cdot 10^9$ Events (386 GB)
- Speedup of **x1.5** through SSD caching
- Bottleneck here: decompression
- Representative compared to our

  HH → bbWW analysis

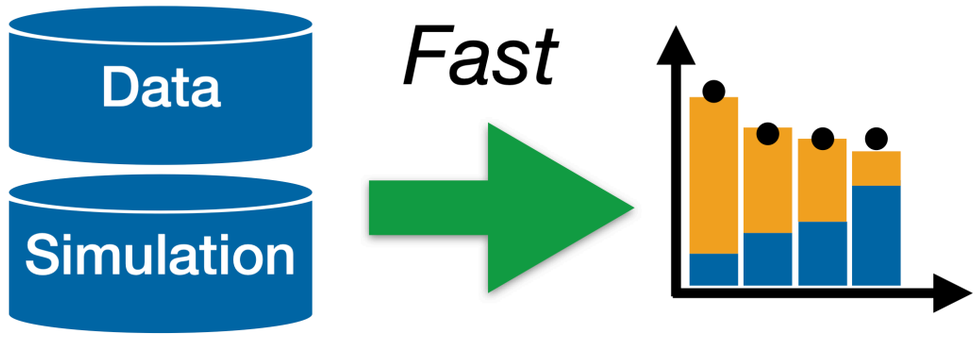

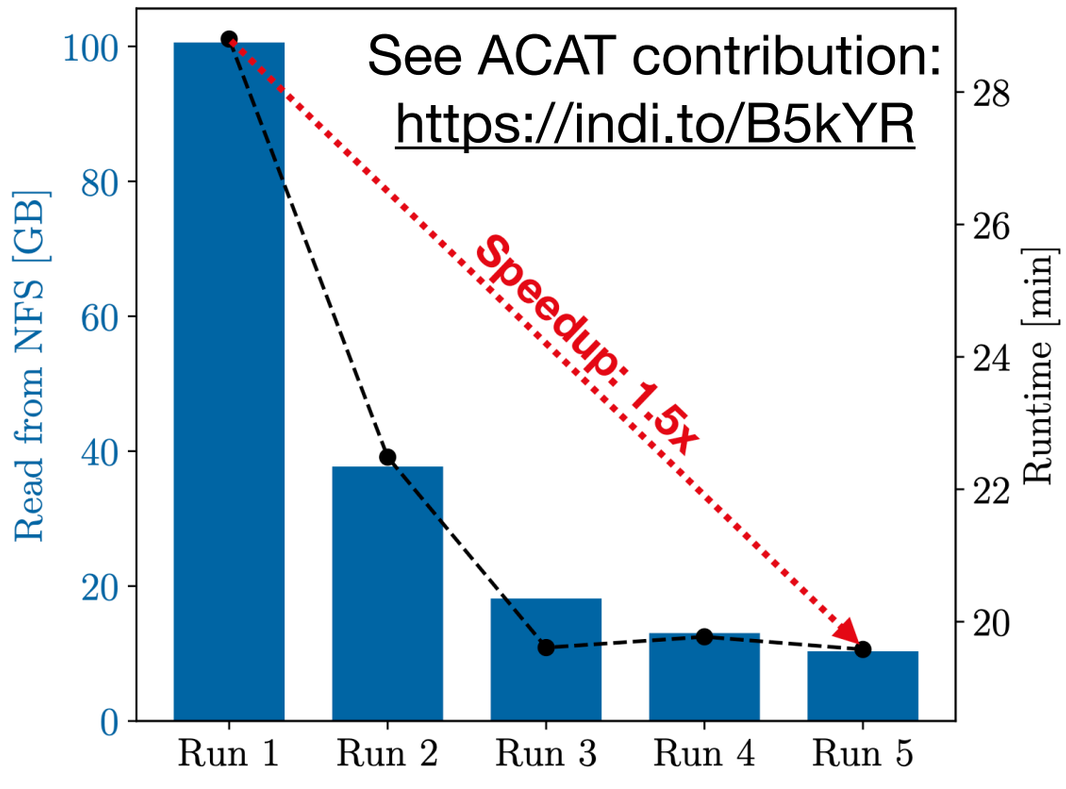

Conclusion:

- Optimise chunk processing by vectorisation and GPU-offloading
- Data caching/locality:
  - Nearline storage (NFS): reliable and low-latency
  - Deterministic use of distributed SSD caches (across workers)

We are happy to continue this work in FIDIUM and established a first collaboration with Thomas Kuhr and David Koch!

Backup

RWTH Aachen Group - Prof. Erdmann

Peter Fackeldey - 16.12.21

- Team:
  - 1 Professor
  - 5 PhD Students
  - 1 Master Student

- Hardware - VISPA Cluster
  - 13 Worker machines, in total:
    - 245 threads, 2GB RAM each
    - 22TB SSD (FS-Cache)
  - Storage (NFS):
    - 6x12TB HDD (striped)
    - 1TB LVM SSD Cache

- Software:
  - NumPy and python-HEP ecosystem
  - Dask (dask-jobqueue)
  - Packaged by conda

https://vispa.physik.rwth-aachen.de

Cache Affinity

On-Worker SSD Cache

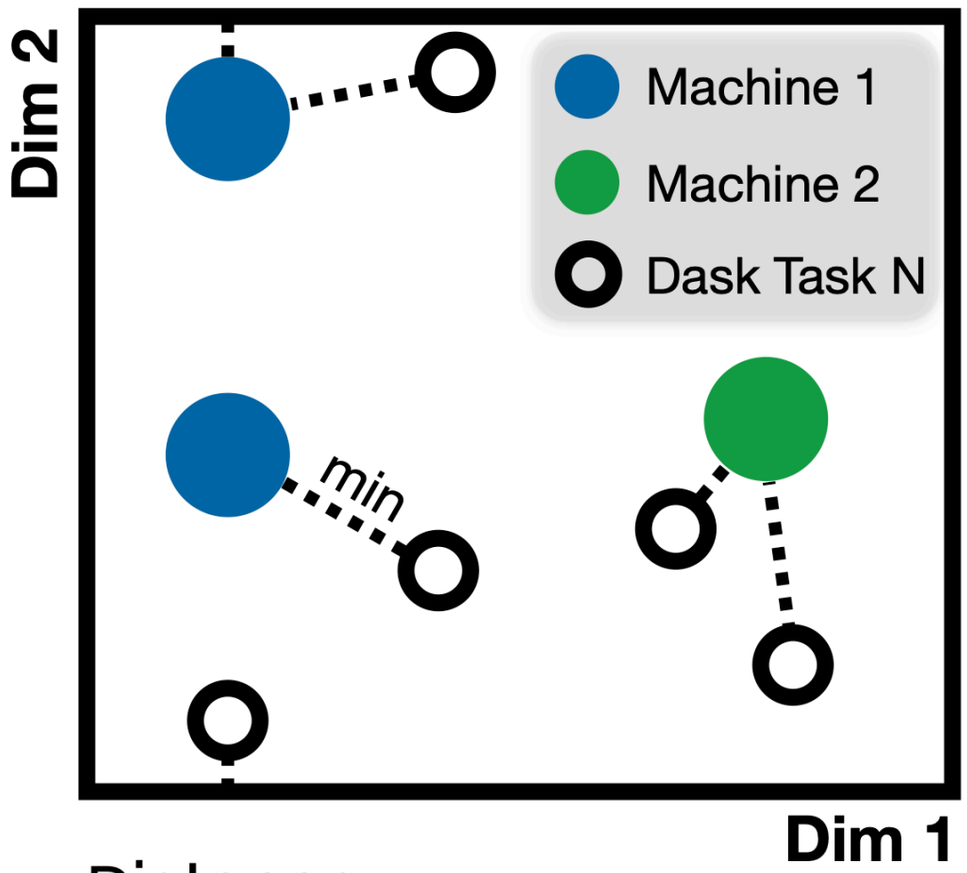

FS-Cache Properties:
- Transparent
- Shared by multiple users

Affine assignment:
- Uses hash distance
- Deterministic
- Smoothly degrading under changes of workers & tasks

Distance: ·····

$$\sum_{i}^{N} min\left(|\circ - \bullet|, 1 - |\circ - \bullet|\right)$$

with N: number of dimensions

Comparable to CRUSH of Ceph