


Quantifying Temporal and Spatial Localities

Florida State University
fox, lojpur, awang@cs.fsu.edu

Quantifying Localities

- Describe workloads
 - _ Compare different workloads
 - _ Create accurate synthetic workloads
- Understand locality transformations
 - _ How are workloads transformed by caching?
 - _ Do workloads properly stress data path components?
 - _ What opportunities remain for further data path optimizations?



Florida State University, {fox, lojpur,
awang@cs.fsu.edu}

Current Metrics

Cache hit ratios

- Measures the effectiveness of caching
- Does not cleanly separate temporal and spatial locality
- Does not show transformations
- Only applies to cache

Reference distance

- Reference distance is the number of unique blocks referenced before accessing the same block

Block distance

- Block Distance is the difference between block locations on disk


New Metrics

Affinity

- Builds on block and reference distance
- 0%-100% easy to understand
- Less sensitive to generations of hardware


Example use

Web-trace



Synthetic

Shuffled references, same inter-arrival rate, reference distribution



Future Work

- Study how each storage component transforms locality characteristics
- Verify fidelity of various benchmarks
- Improve storage components and benchmarks