Database Query Processing using Distributed Shared Memory

Paul M. Aoki

CS 294-5

Fall 1994

Roadmap

- Motivation
- Methodology
- Preliminary Simulation Results
- Conclusions

Why You Should Care

- *Non-scientific* **application**

~~**HPCC**~~

- **Larger than typical benchmarks**
  - 250K lines total
  - ~3K lines in ‘‘kernel’’

Why I Care

- **Fine-grain parallel SM systems exist**
- **MPP support currently weak**
  - mainly serial servers + sockets
  - exception: Tandem MP (5 years!)

=> Modify SM support for MPPs?

Issues...

- **Feasibility**
  - Rewriting for locality
  - Speedups
- Resource Characterization
  - Communication (latency, b/w)
  - Concurrency (spinlocking)
- Semantics
  - put/get, Thekkath, segments, soft DSM, hard DSM, ...

Methodology

- **Real SM parallel DBMS server**
- **Simulated network**
- **Glue: ATOM (object modification)**
  - memory operations
  - system calls

=> No need to label SM ops

Preliminary Results: Speedups

Thekkath; 7μs ATM

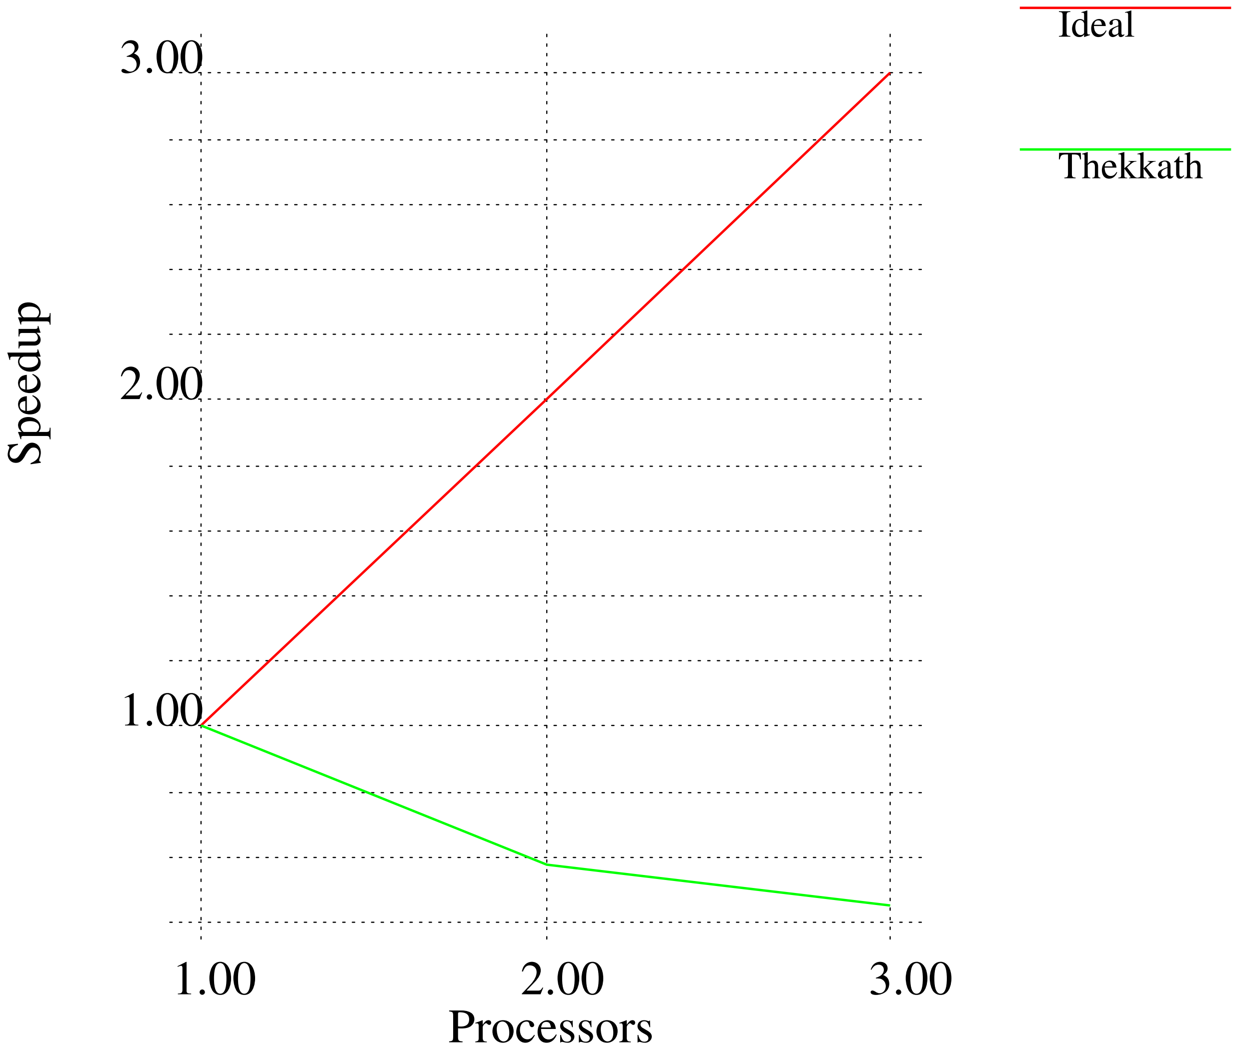

Conclusions

- **Speedups look unattractive**
  - source of performance penalty...?
- **Lots of future work!**
  - spinlocks
  - more memory types
  - parameters

    network

    more RAM

    more processors

Current Practice

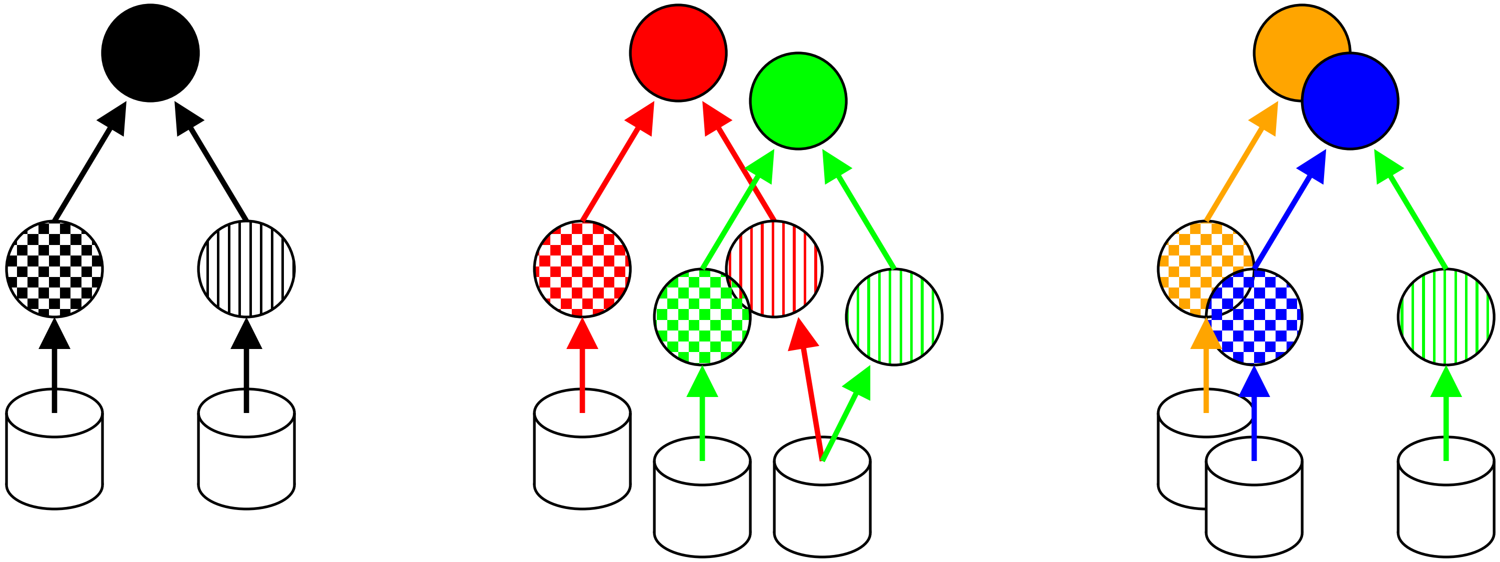

serial simple parallel general parallel

- serial servers + glue
- poor I/O balance
- no scheduling

Hybrid Hashjoin

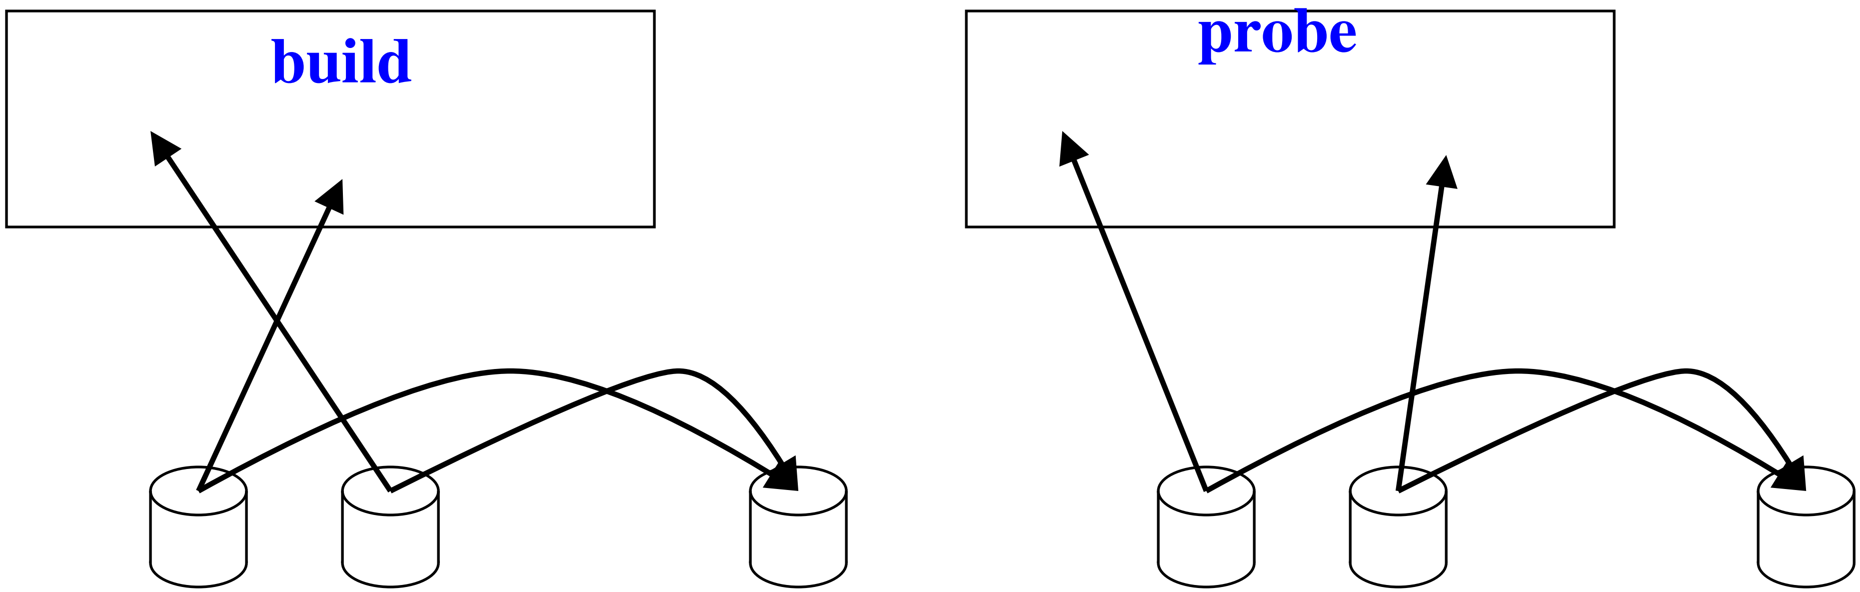

...repeat using spool files

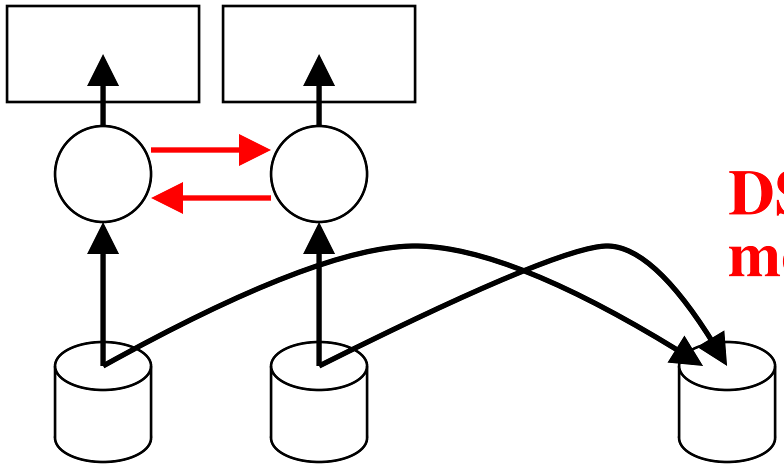

DSM version requires more partitioning

Preliminary Results: Validation