Multi-Grain Coherence Directories

J. Zebchuk, B. Falsafi, and A. Moshovos

UNIVERSITY OF TORONTO
FACULTY OF APPLIED SCIENCE & ENGINEERING

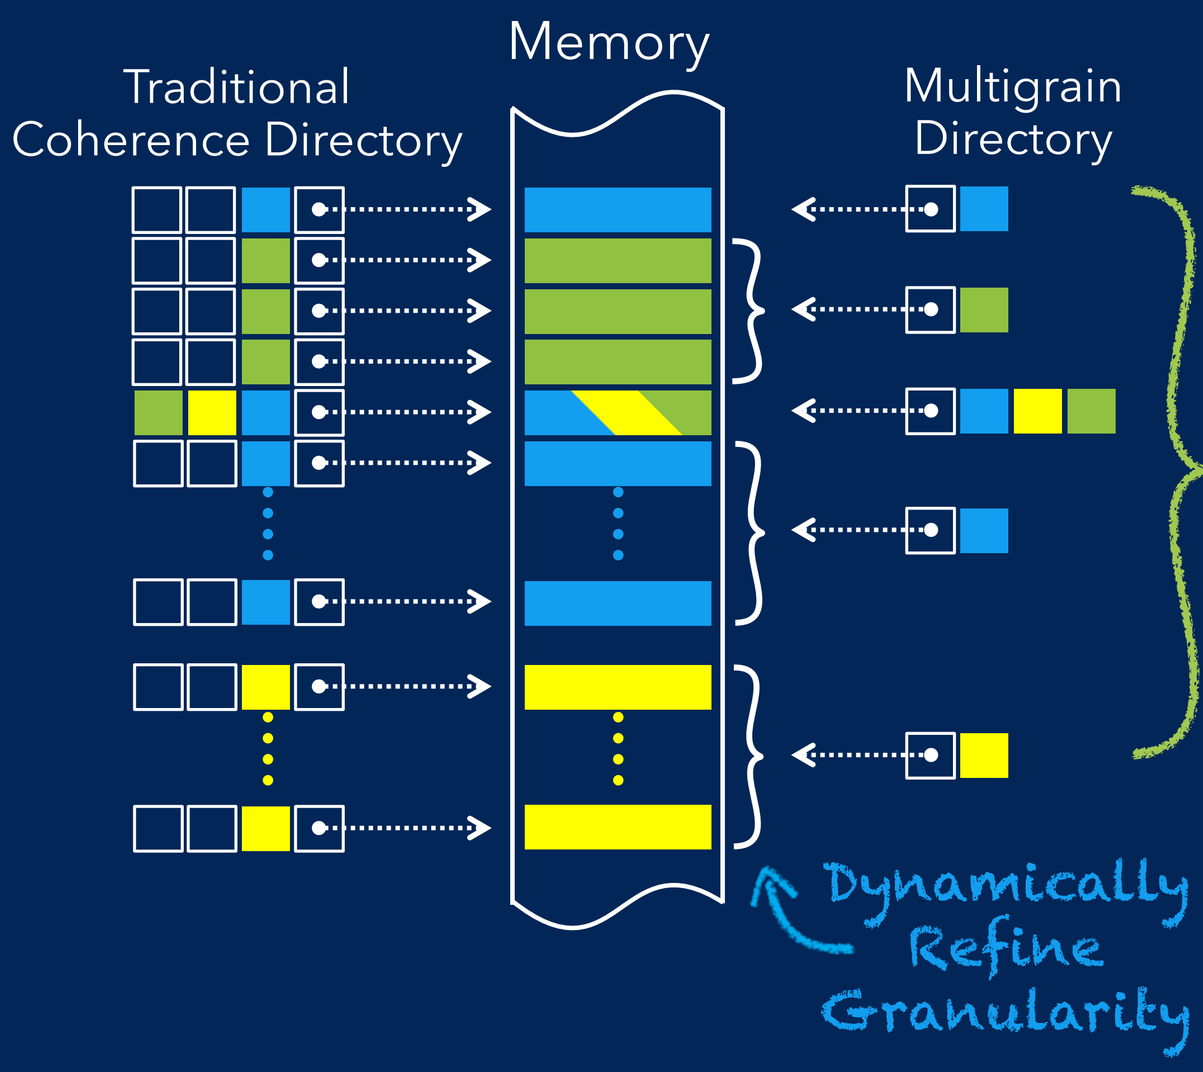

Initial thought experiment:

How many entries does the directory require if it can dynamically refine the size of each entry?

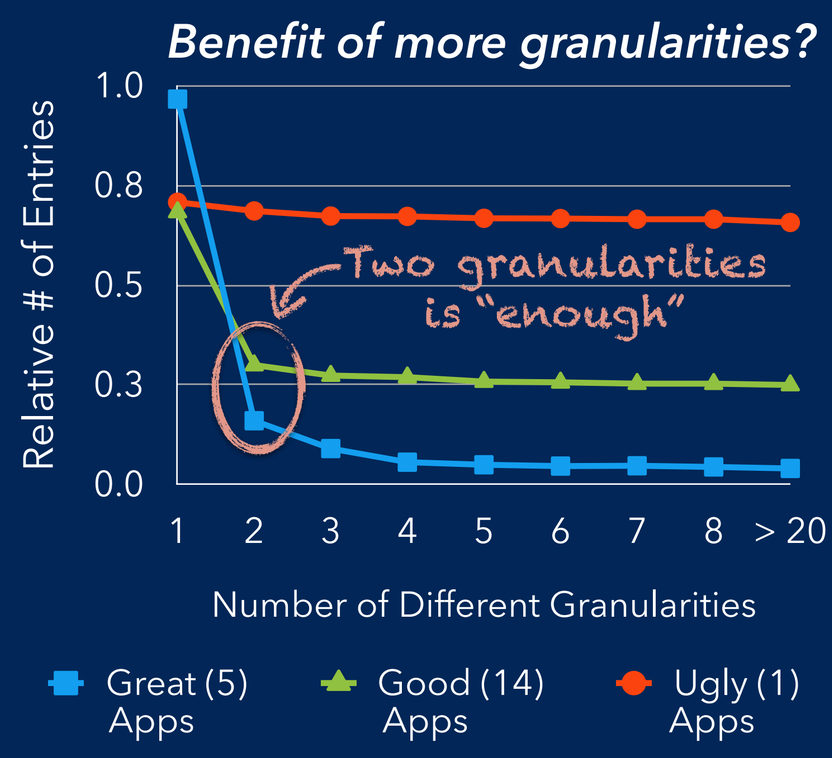

Practical dual-grain directory design

Track 64 byte blocks and 1 KB regions in a single structure.

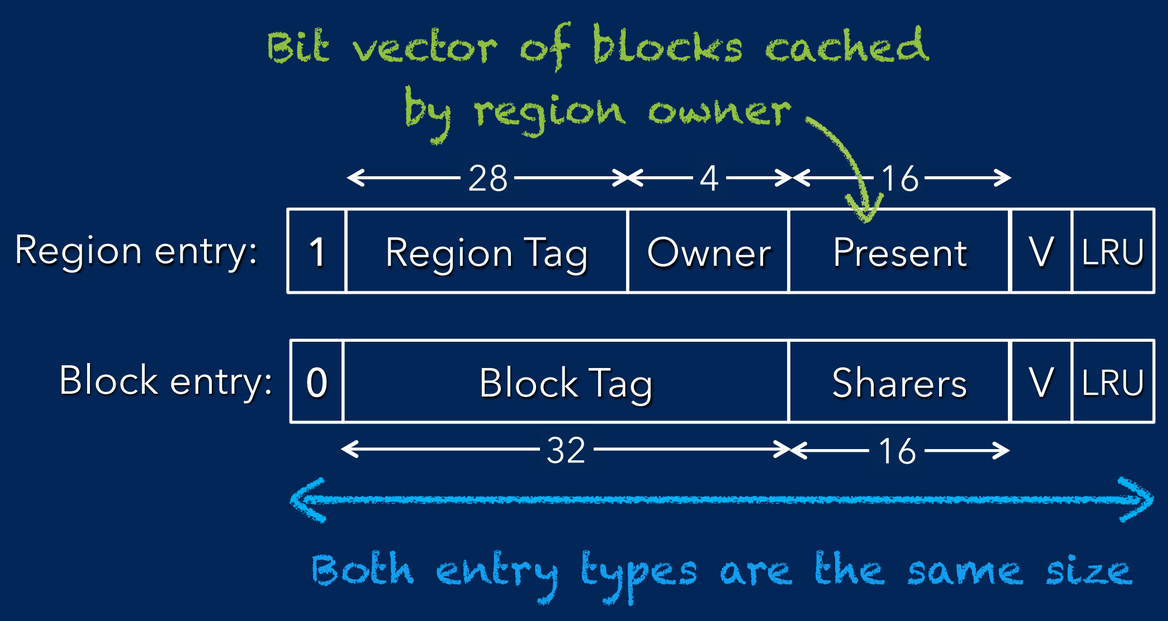

Skewed-associative structure with parallel lookups.

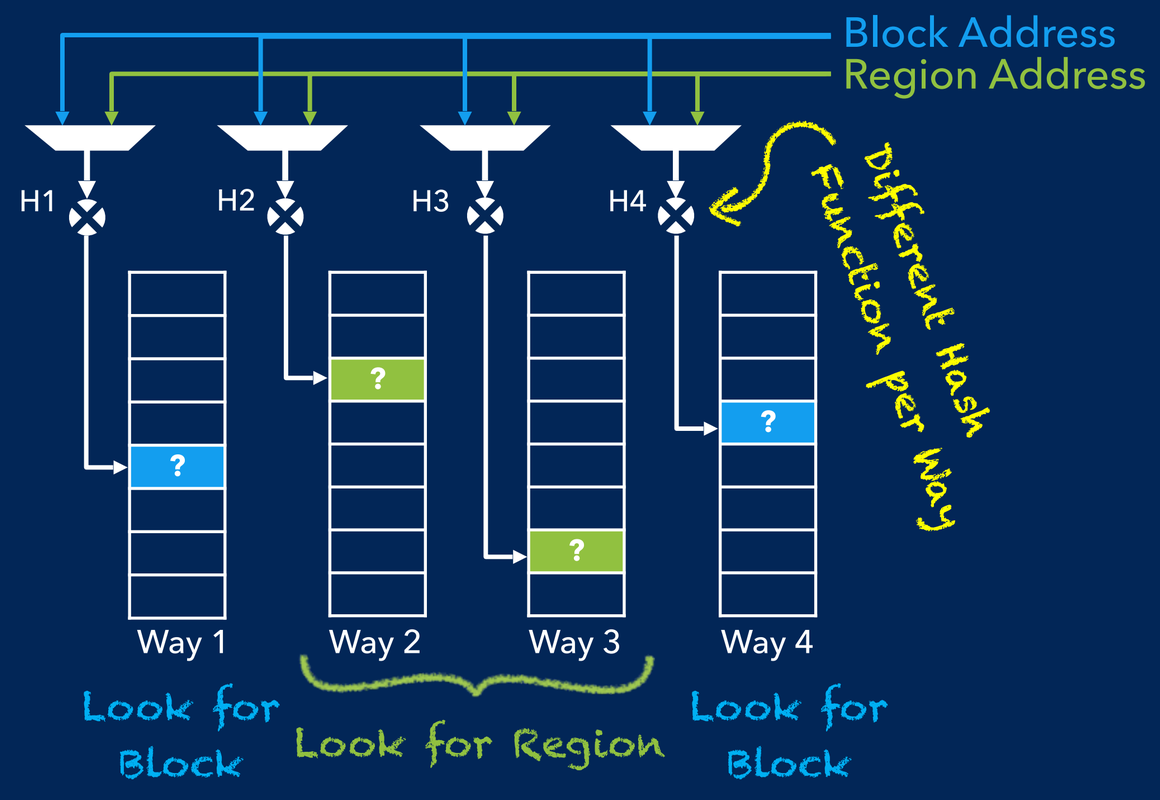

Related Works:

Two recent proposals also track coherence for individual blocks and larger regions. Neither proposal precisely tracks individual blocks within regions, and their designs potentially require multiple lookups for a single request and are susceptible to corner cases that behave poorly.

M. Alisafaee. Spatiotemporal coherence tracking. MICRO, 2012.

L. Fang, P. Liu, Q. Hu, M. C. Huang, and G. Jiang. Building expressive, area-efficient coherence directories. PACT, 2013.

What's different about MGD?

- ✓ Explore more than two granularities
- ✓ Efficient, parallel lookups
- ✓ Explicit, precise information for all blocks
- ✓ No coherence protocol changes
- ✓ No new race conditions
- ✓ Robust performance

Methodology:

- 16-core chip-multiprocessor
- *64 KB L1 Instruction and Data caches*
- *256 KB private L2 cache*
- *16 MB distributed shared L3 cache*
- 8-way set-associative directory
- Simics + Flexus full-system simulation
- *SPARCv9 ISA, Solaris OS (version 8 or 10)*
- *Detailed timing of out-of-order cores and memory hierarchy*
- Selection of 18 diverse applications
- *SpecWeb, Parsec, TPC-C, TPC-H, CloudSuite*

Final Result: