


Base-Delta-Immediate Compression: Practical Data Compression for On-Chip Caches



Gennady
Pekhimenko



Vivek Seshadri



Onur Mutlu




Todd C. Mowry

Cache compression is a promising technique to increase on-chip cache capacity and to decrease on-chip and off-chip bandwidth usage. Unfortunately, directly applying well-known compression algorithms (usually implemented in software) leads to high hardware complexity and unacceptable decompression/compression latencies, which in turn can negatively affect performance. Hence, there is a need for a simple yet efficient compression technique that can effectively compress common in-cache data patterns, and has minimal effect on cache access latency. In this paper, we introduce a new compression algorithm called Base-Delta-Immediate (BDI) compression, a practical technique for compressing data in on-chip caches. The key idea is that, for many cache lines, the values within the cache line have a low dynamic range – i.e., the differences between values stored within the cache line are small. As a result, a cache line can be represented using a base value and an array of differences whose combined size is much smaller than the original cache line (we call this the base+delta encoding, see Figure 1). Moreover, many cache lines intersperse such base+delta values with small values – our BDI technique efficiently incorporates such immediate values into its encoding.

Compared to prior cache compression approaches, our studies show that BDI strikes a sweet-spot in the tradeoff between compression ratio, decompression/compression latencies, and hardware complexity. Our results show that BDI compression improves performance for both single-core (8.1% improvement) and multi-core workloads (9.5% / 11.2% improvement for two/four cores). For many applications, BDI provides the performance benefit of doubling the cache size of the baseline system, effectively increasing average cache capacity by 1.53X.

Compressed Cache Line



Uncompressed Cache Line

Fig. 1: BDI decompressor design.




Fig 2: Single-core IPC improvement.