

SWAP: Effective Fine-Grain Management of Shared Last-Level Caches with Minimum Hardware Support

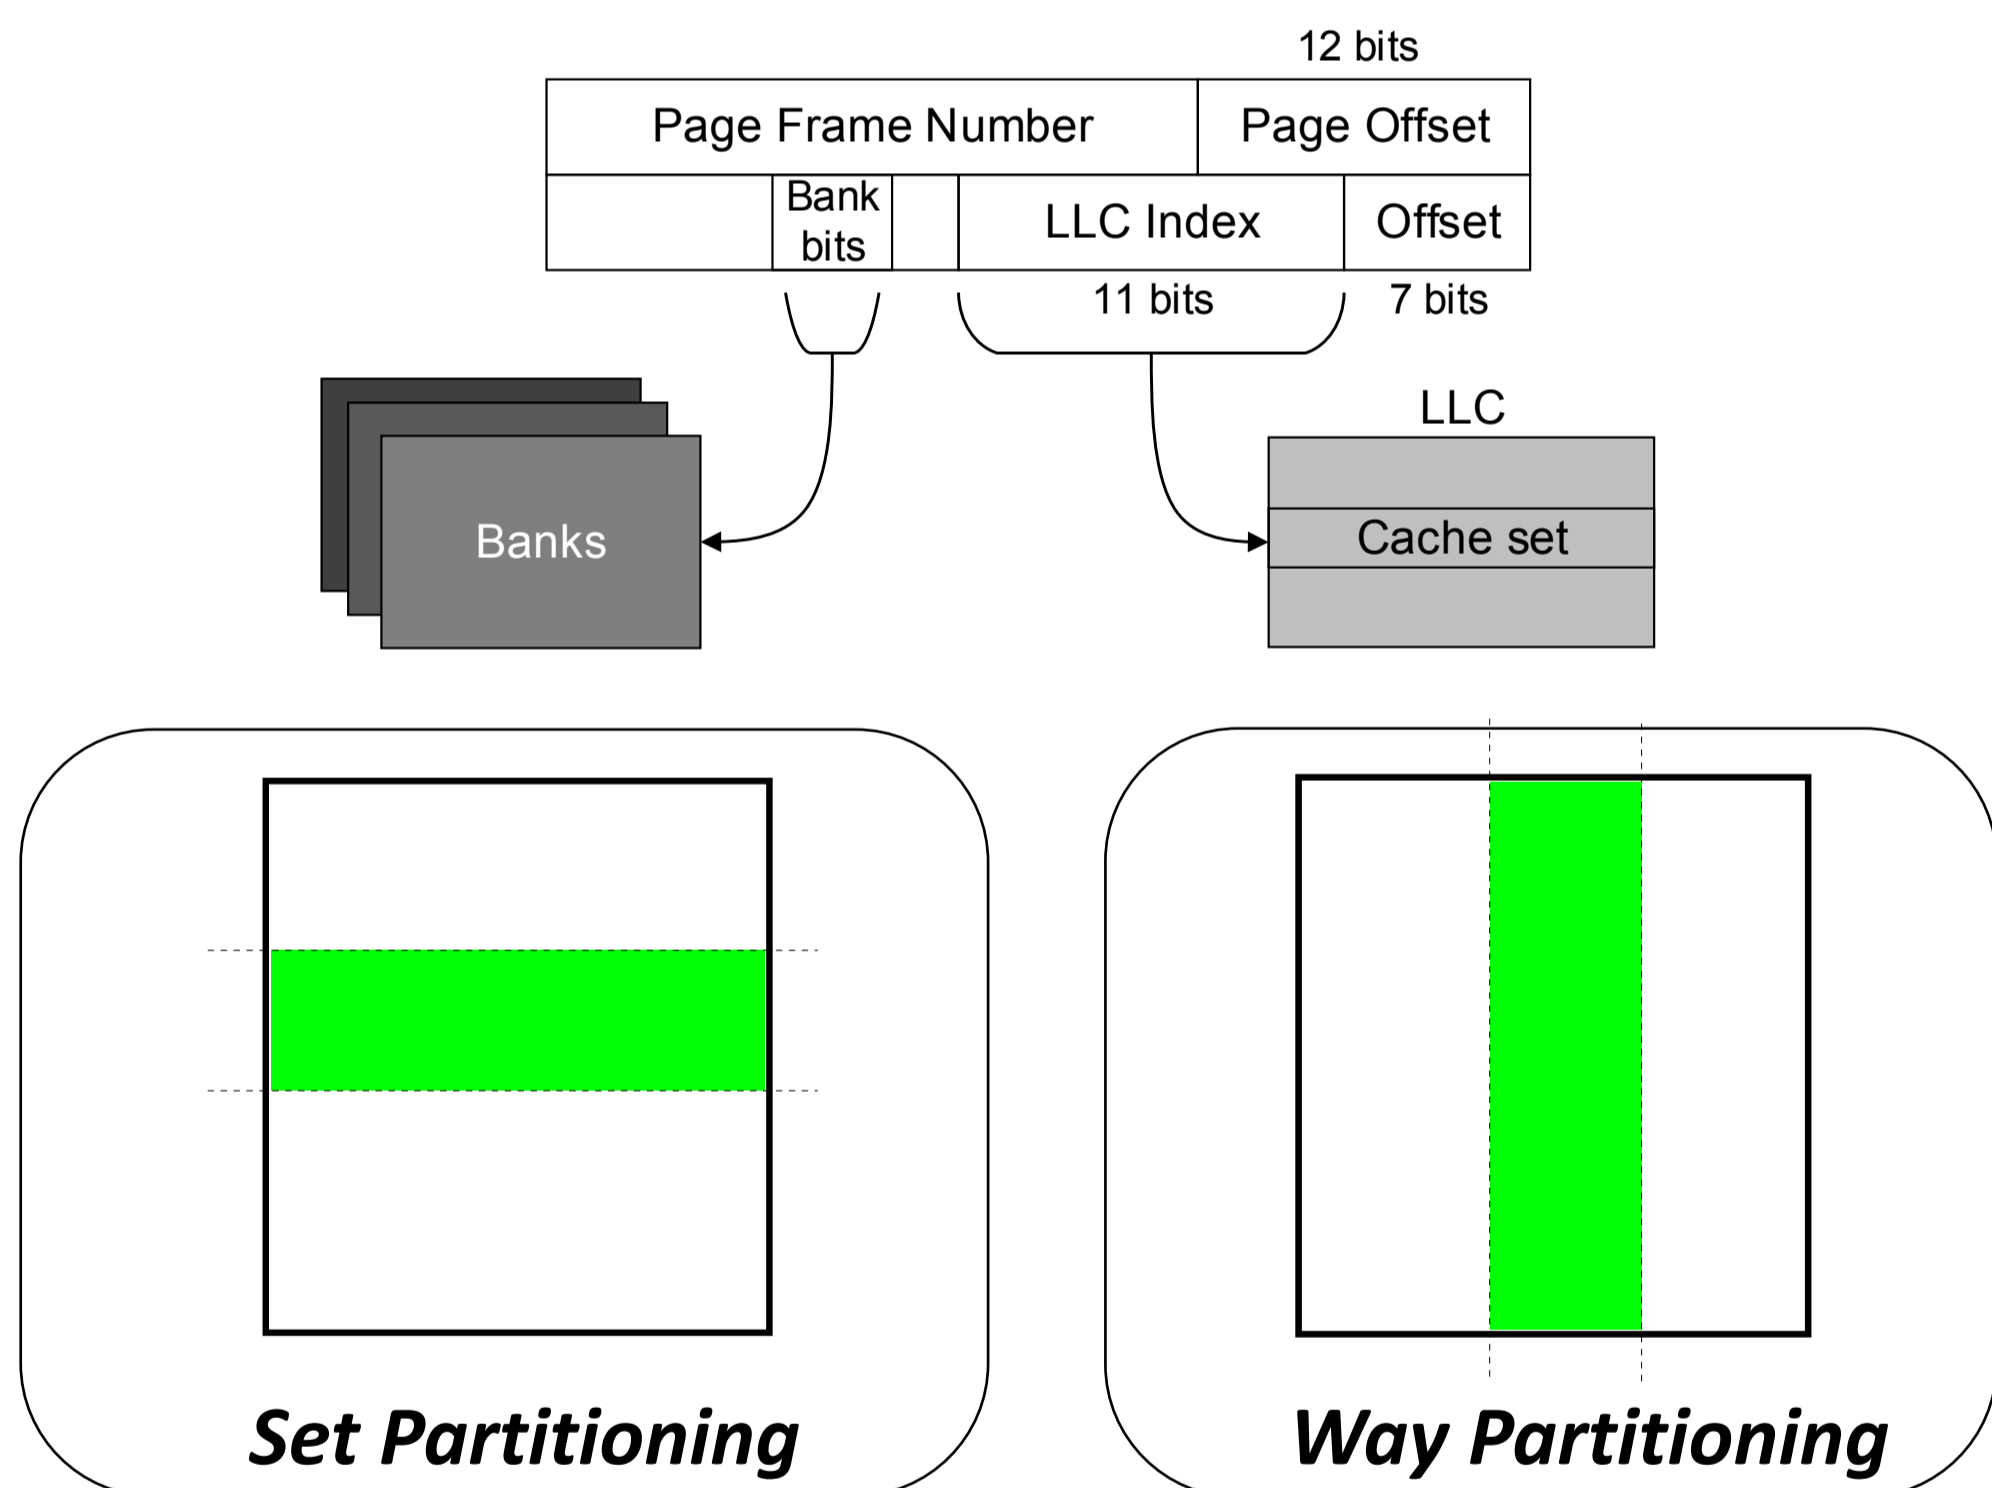
Introduction

Performance isolation is an important goal in server-class environments. Partitioning the last-level cache of a chip multiprocessor (CMP) across co-running applications has proven useful in this regard. Two popular approaches are (a) hardware support for way partitioning, or (b) operating system support for set partitioning through page coloring. Unfortunately, neither approach by itself is scalable beyond a handful of cores without incurring in significant performance overheads.

We propose SWAP, a scalable and fine-grained cache management technique that seamlessly combines set and way partitioning. By cooperatively managing cache ways and sets, SWAP (“Set and WAY Partitioning”) can successfully provide hundreds of fine-grained cache partitions for the manycore era.

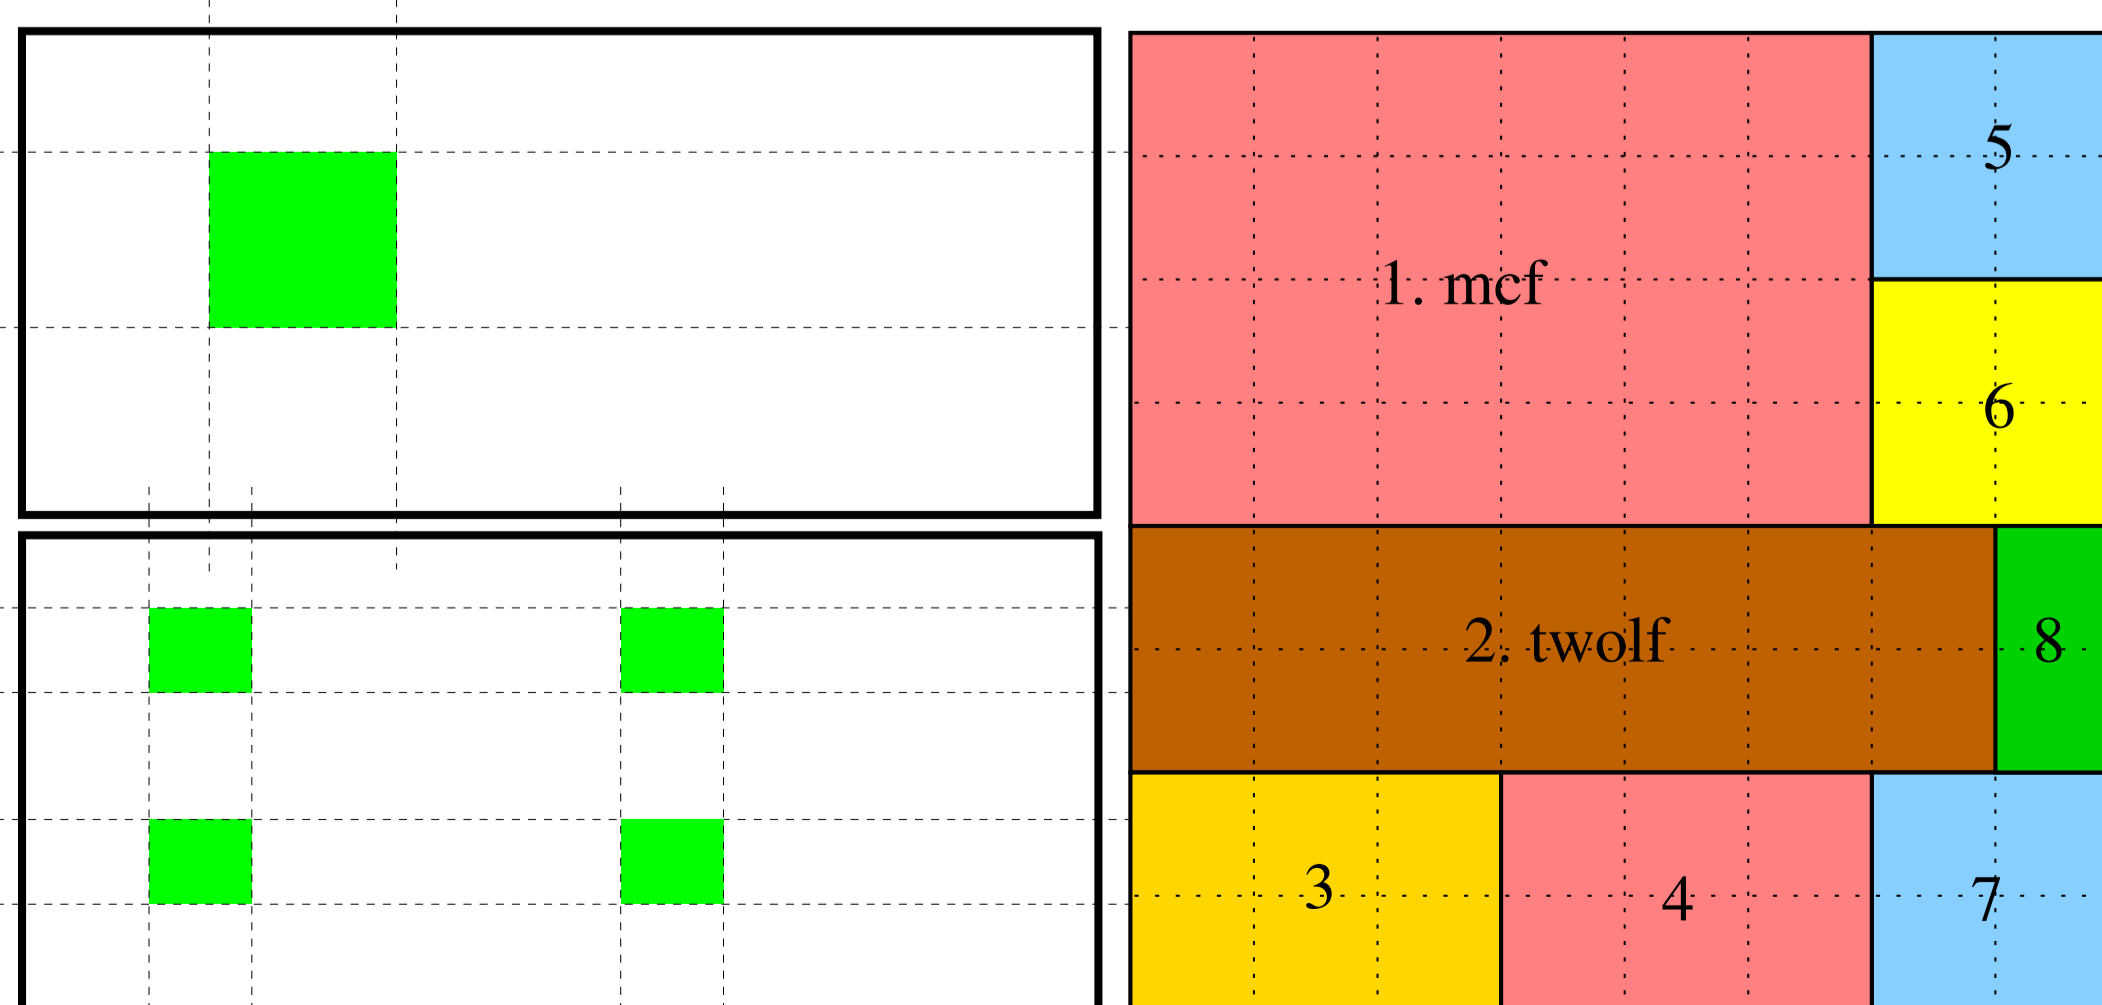
SWAP requires no additional hardware beyond way partitioning. In fact, SWAP can be readily implemented in existing commercial servers whose processors do support hardware way partitioning. In this paper, we prototype SWAP on a 48-core Cavium ThunderX platform running Linux, and show average speedups over no cache partitioning are twice as large as those attained with hardware way partitioning alone.

Background



Neither set or way partitioning is fine-grained enough.

Combine Set and Way Partitioning

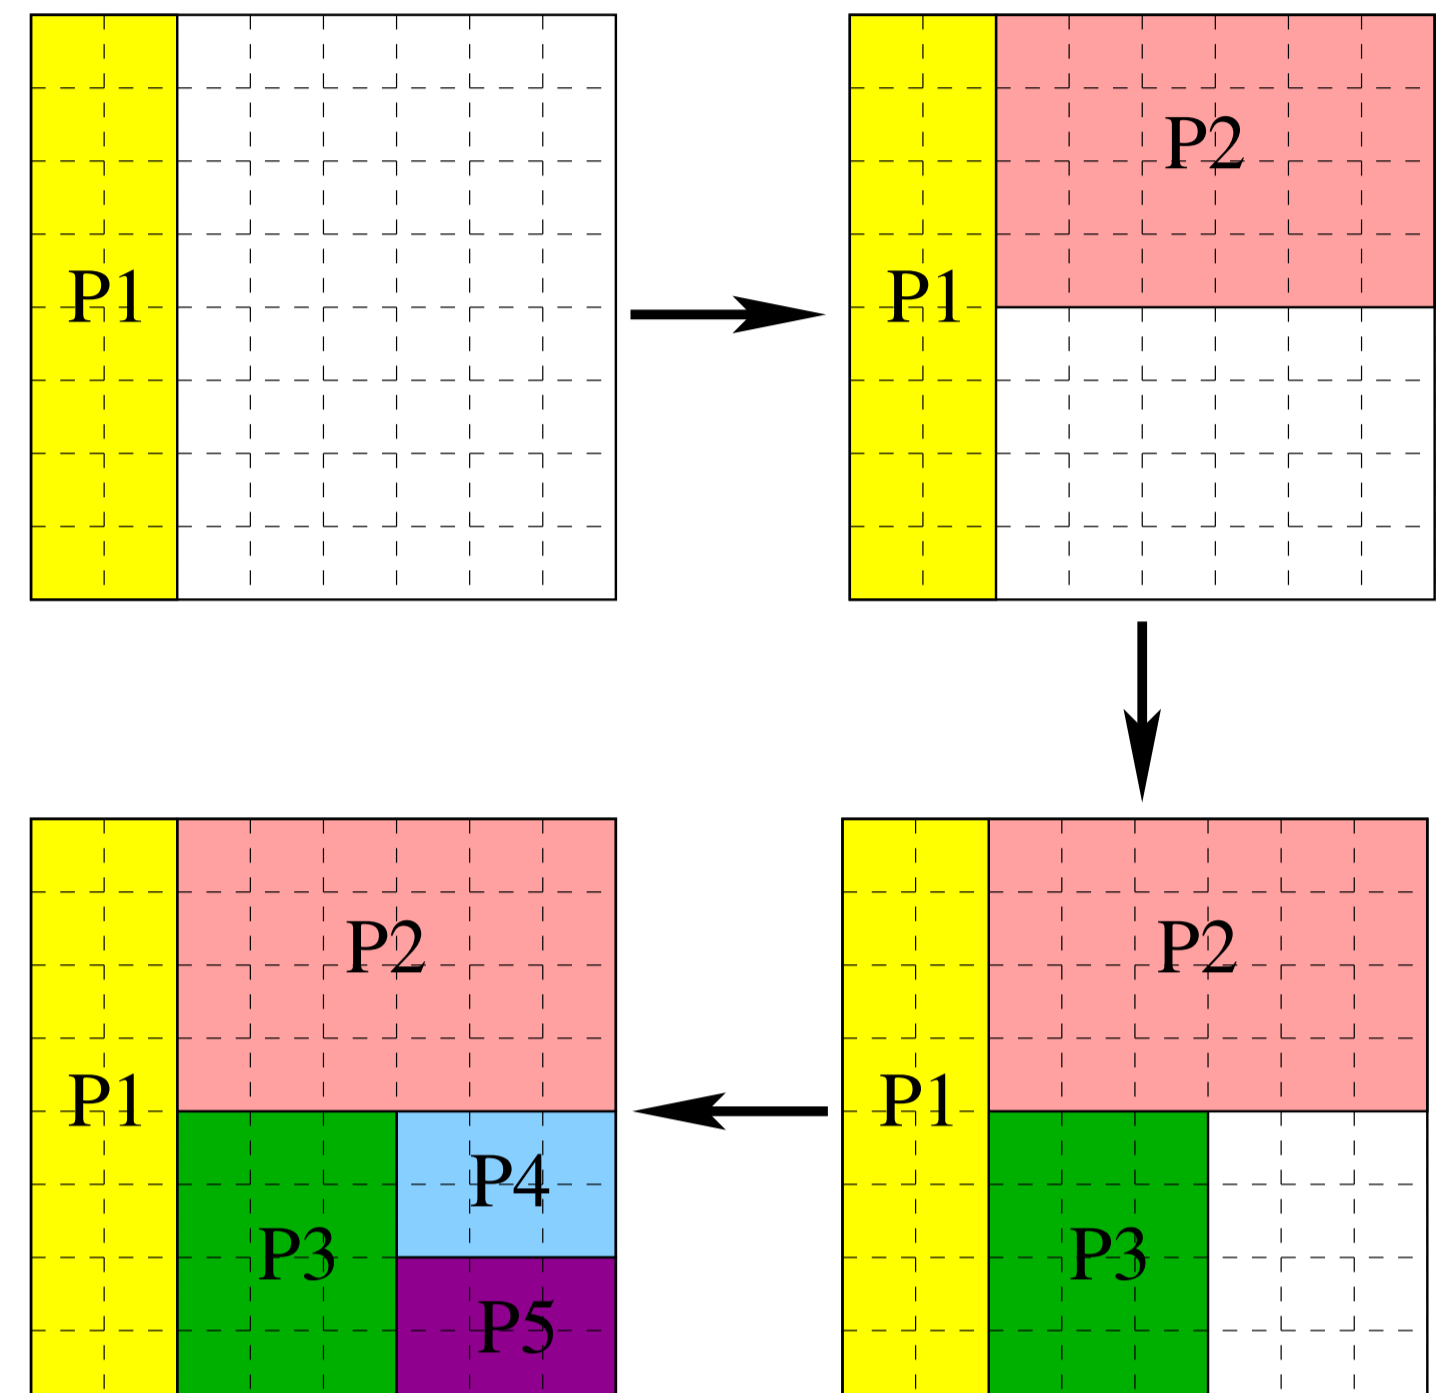


* Partition placement * Memory Pressure * Recolor overhead
* Increased conflict misses in way partitioning.

SWAP Design

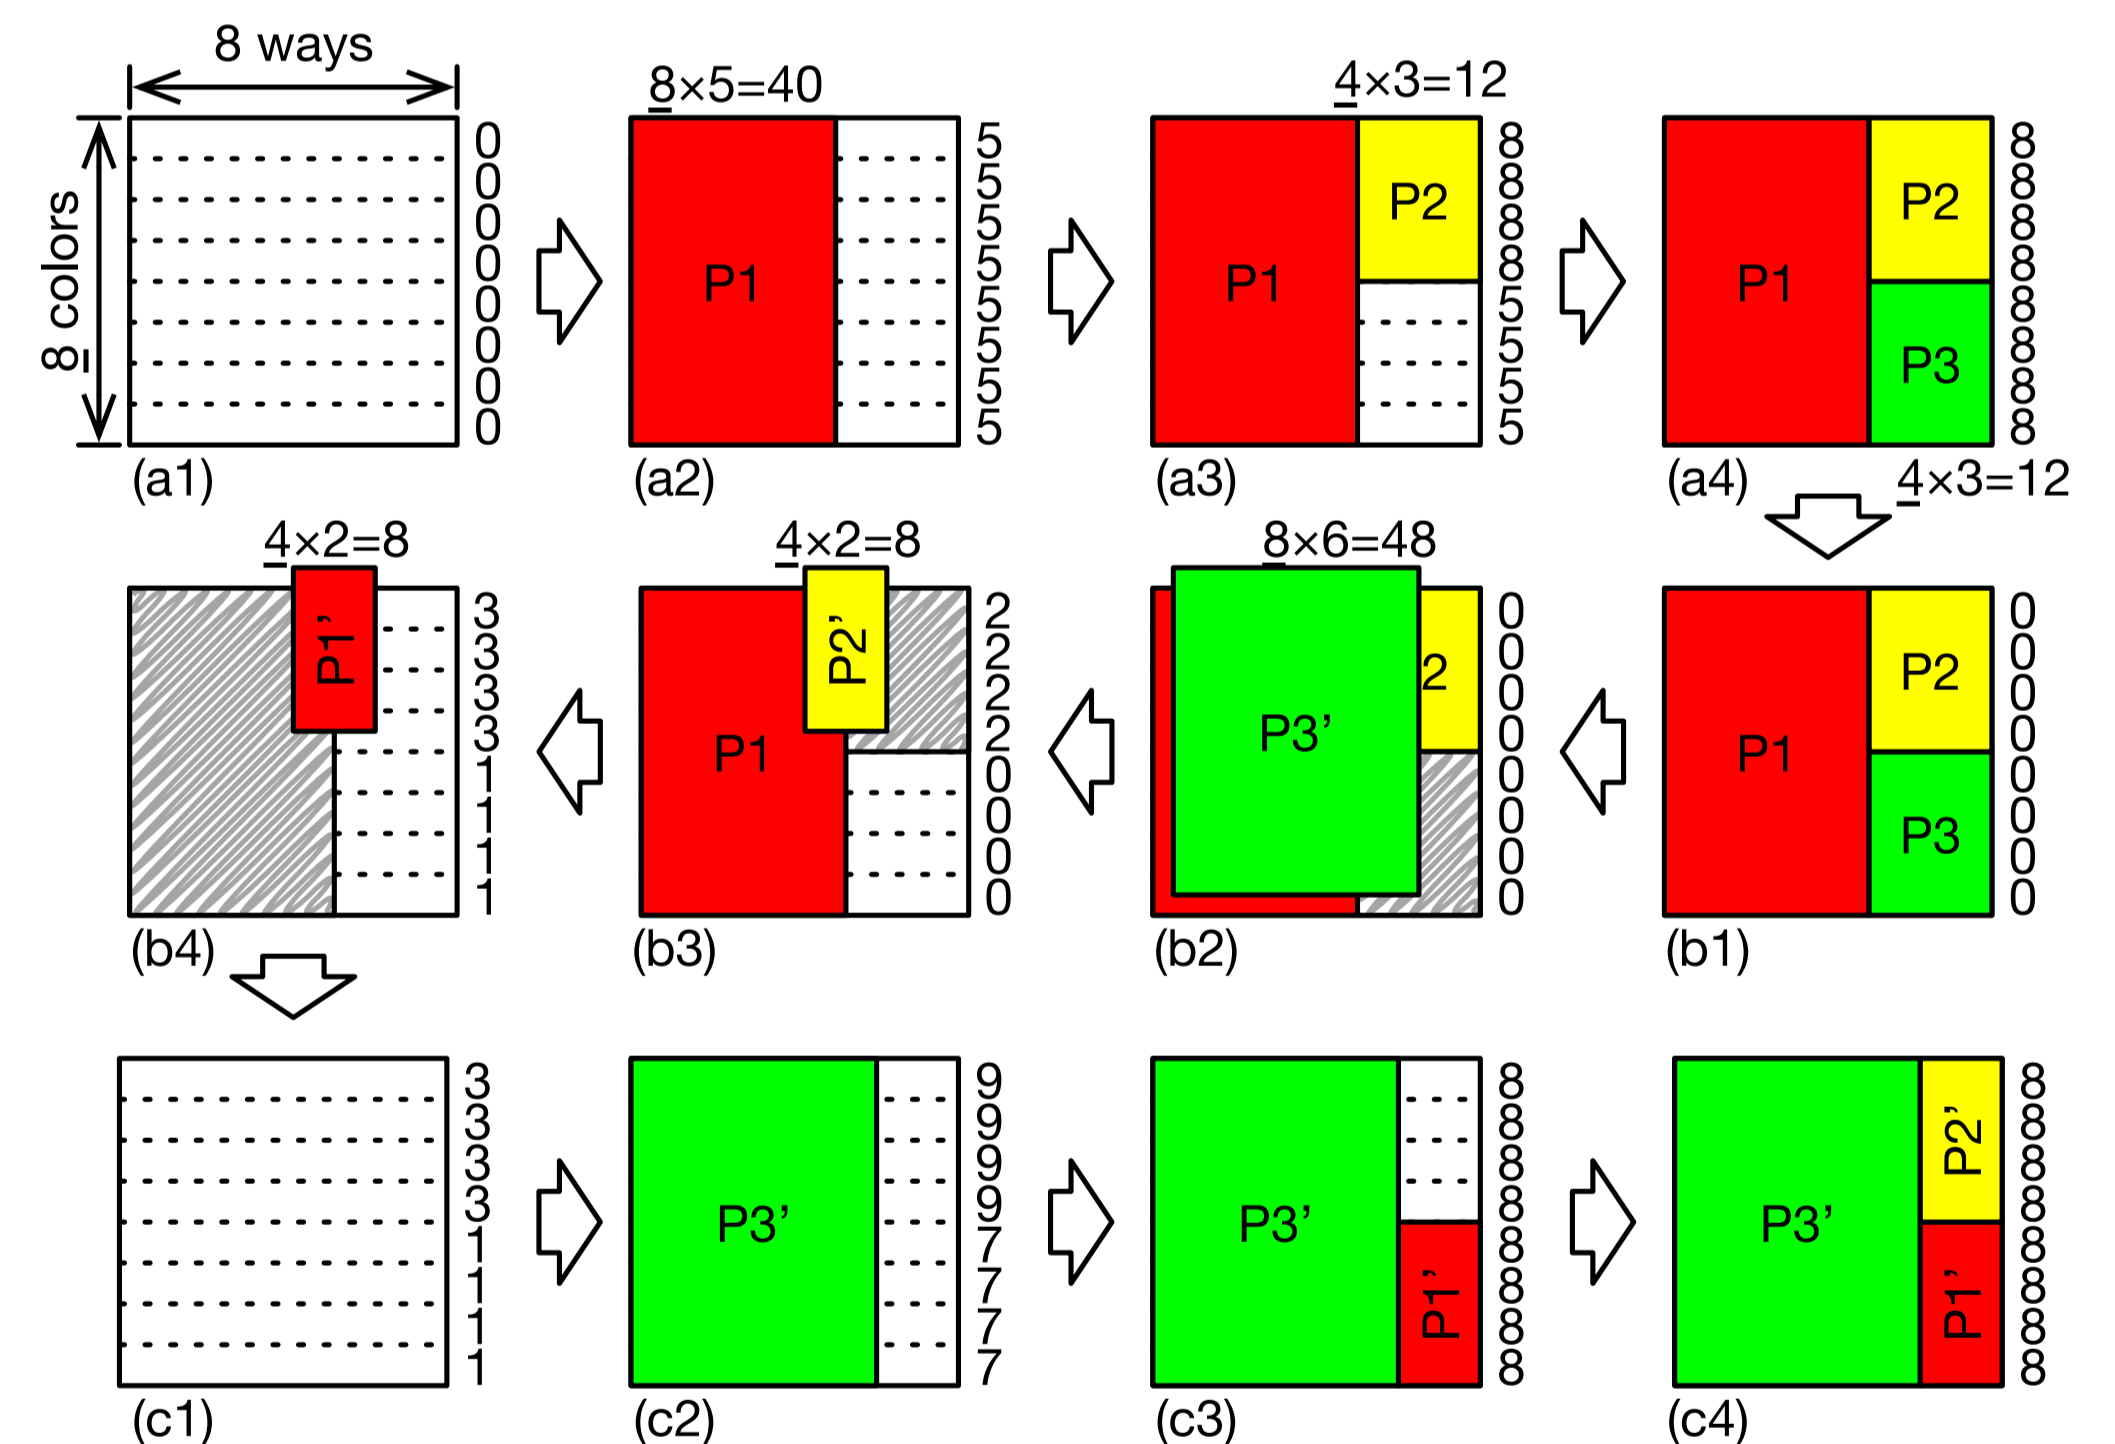
Static Partitioning:

- Offline/online profiling to get cache sensitivity curve of each app.
- Run lookahead algorithm to find allocation size.
- Sort and place applications by their classes as follows:



Dynamic Repartitioning:

- ReRun lookahead algorithm to find allocation size.
- Reset usage counters of each color.
- Try reusing previous colors as much as possible.



Evaluation

Platform: 48-core Cavium ThunderX

Results: Improve system throughput by 12-14.5%.

SWAP overhead: negligible

More potentials: guarantee QoS and improve

