

Perforated Page: Supporting Fragmented Memory Allocation for Large Pages

Chang Hyun Park, Sanghoon Cha, Bokyeong Kim,
Youngjin Kwon, David Black-Schaffer, and Jaehyuk Huh



UPPSALA
UNIVERSITET



Benefits and Challenges of Large Pages

Performance

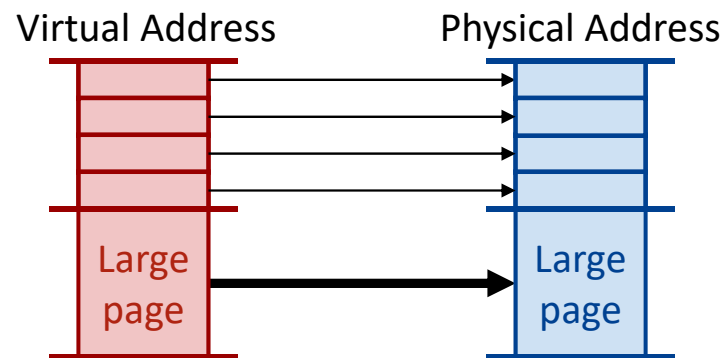
Difficult to make

Overheads



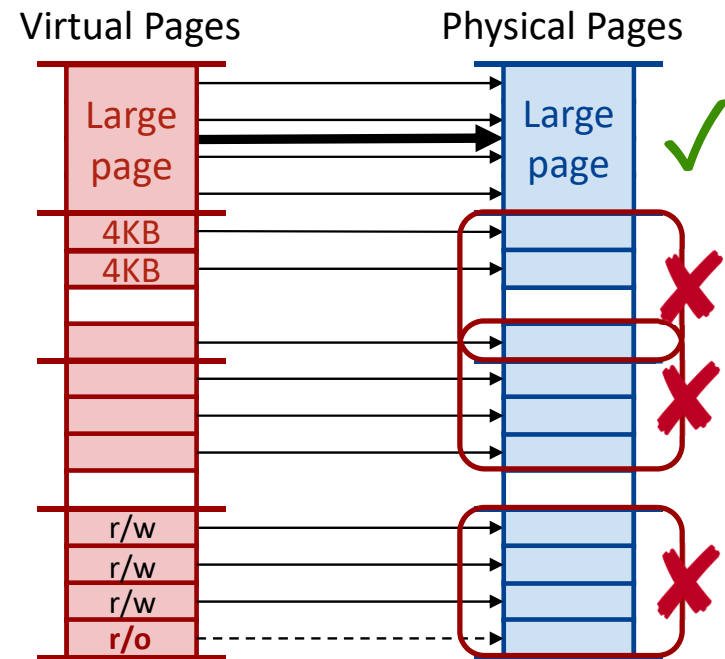
Large Pages for Performance

- 2MB large page \rightarrow 512x TLB coverage
- 2MB large page \rightarrow ~68% faster execution [1]




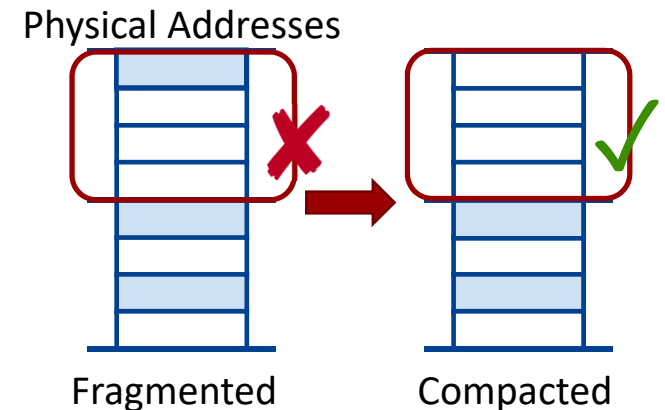
Large Pages are Difficult to Make

- Contiguous: 512 4KB pages
- Aligned: 2MB boundary
- Homogeneous: permissions




Overheads: Compaction

- Memory compaction to create large pages
 - Compaction takes up to 35% of time [1]
- Immovable pages prevent compaction
 - E.g., kernel allocations, I/O buffers, etc.
 - After Linux kernel build, 50% of 2MB regions had immovable pages




Immovable pages Not-compactable



Overheads: Memory Bloating

- Sparse access/use in large pages → wasted physical space



- Redis with large pages (4M keys, 16KB values)
 - **20% more** memory consumption (78GB → 93GB)
 - **45% fewer** TLB misses (1.8 MPKI → 1.1 MPKI)



Summary

Large pages for
Performance

- Better TLB Coverage → Better Performance

Large pages are
Difficult to make

Can we get the **benefits** of large pages
without the **difficulties** and **overheads**?

require homogeneous permissions


Large pages come
with **Overheads**

- **Costly compacting** to create contiguous regions
- **Immovable** pages common
- **Memory bloating** due to sparse access



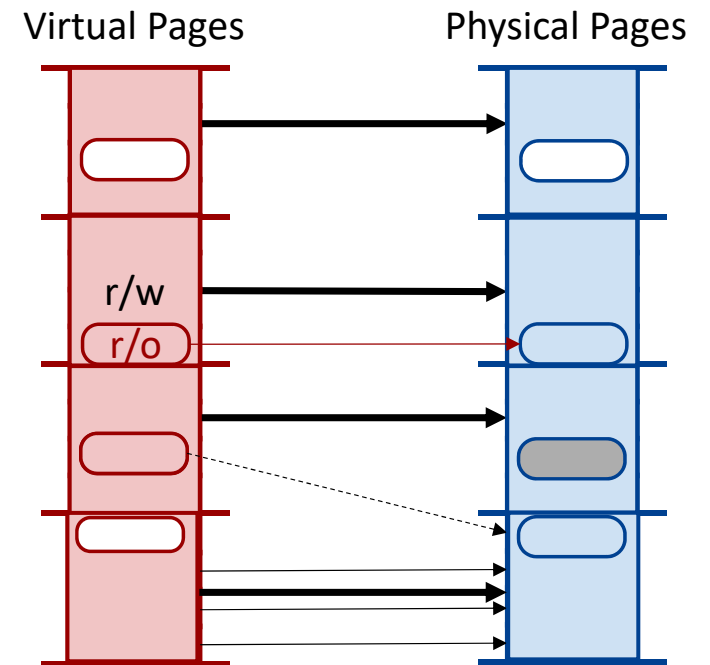
Perforated Pages

- *Perforated pages* for majority of 2MB region
- *Hole* pages for flexible fine-grained mappings

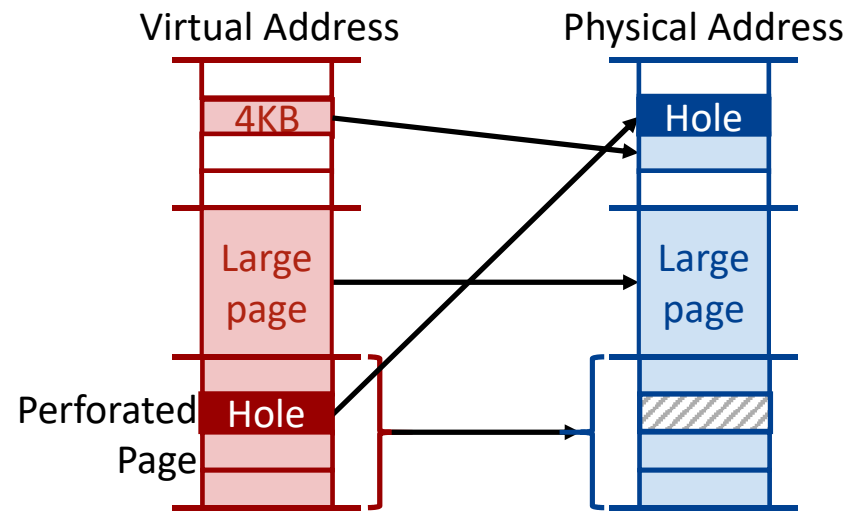


Perforated Pages

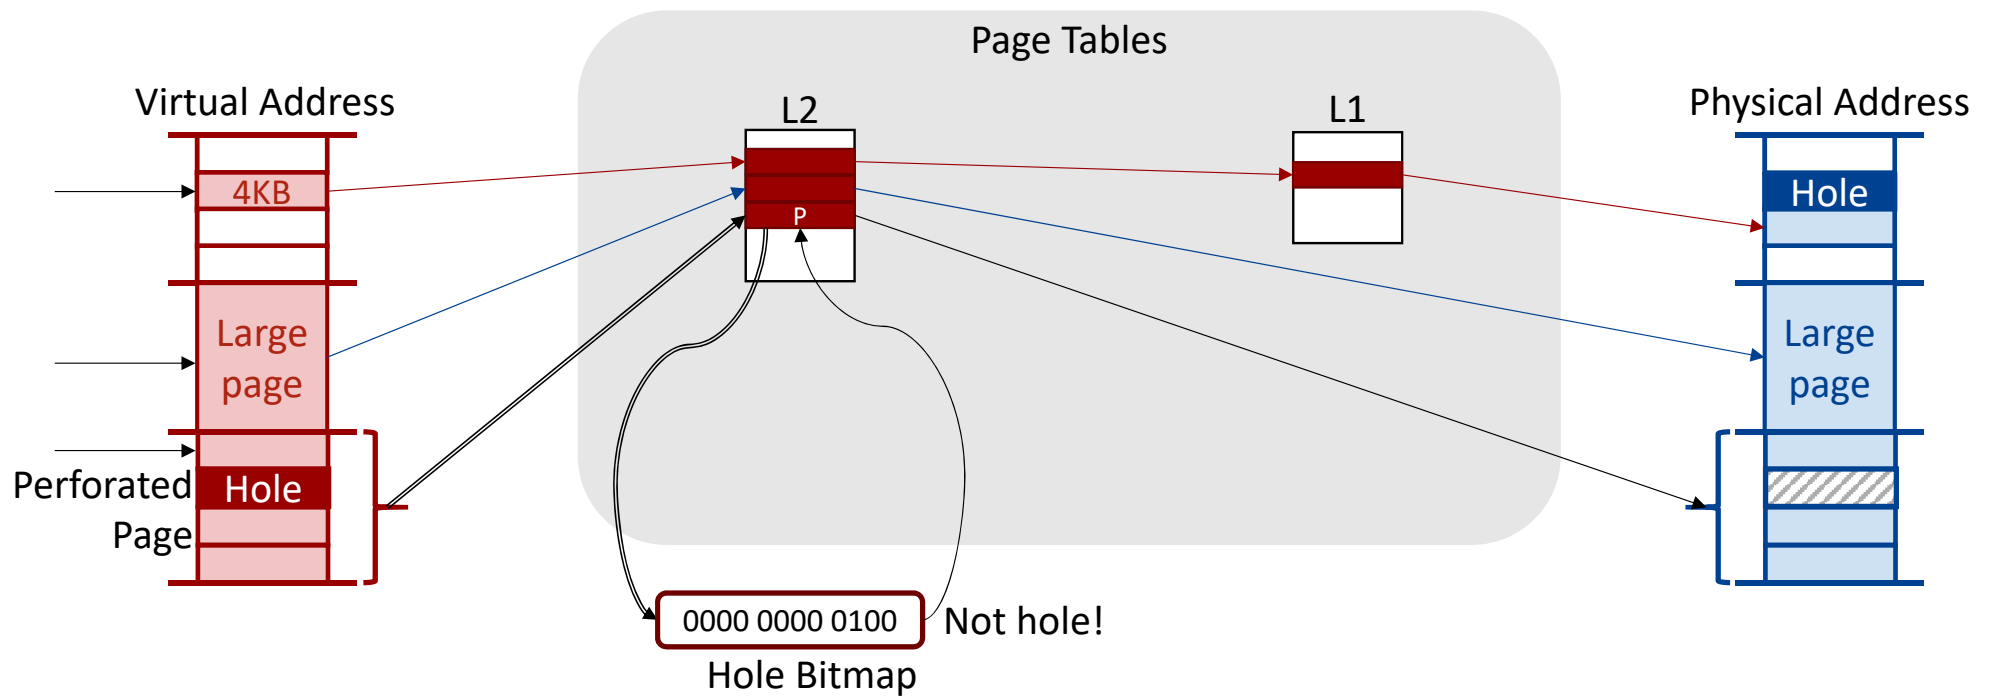
- **Large pages with holes**
 - Efficient translation *for most of the page*
 - Relaxed contiguity constraint *where needed*
 - Relaxed permission constraint *where needed*
- Tolerate immovable pages
 - Overcome by re-mapping holes
 - 50% of 2MB + 50% of perforated pages
- Avoid bloating
 - Conserve untouched pages via holes
 - 0% bloating, 17% TLB MPKI reduction (2MB: 45% TLB MPKI reduction)



How it works



How it works: Basics

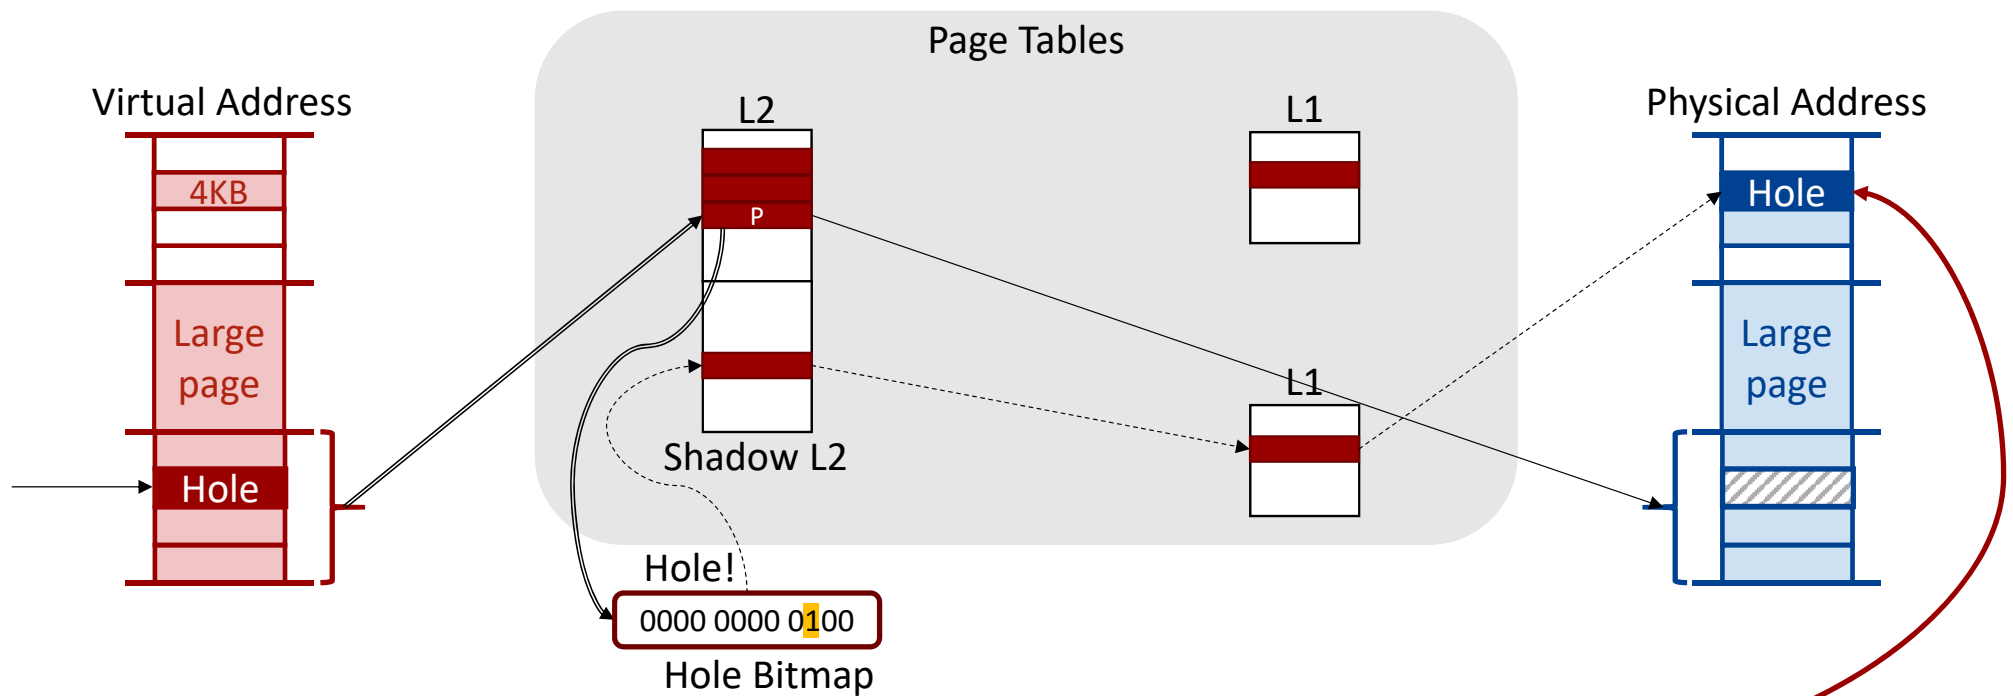


How to distinguish holes/perforated page?

Per-page bitmap to mark holes
0.003% capacity overhead



How it works: Basics

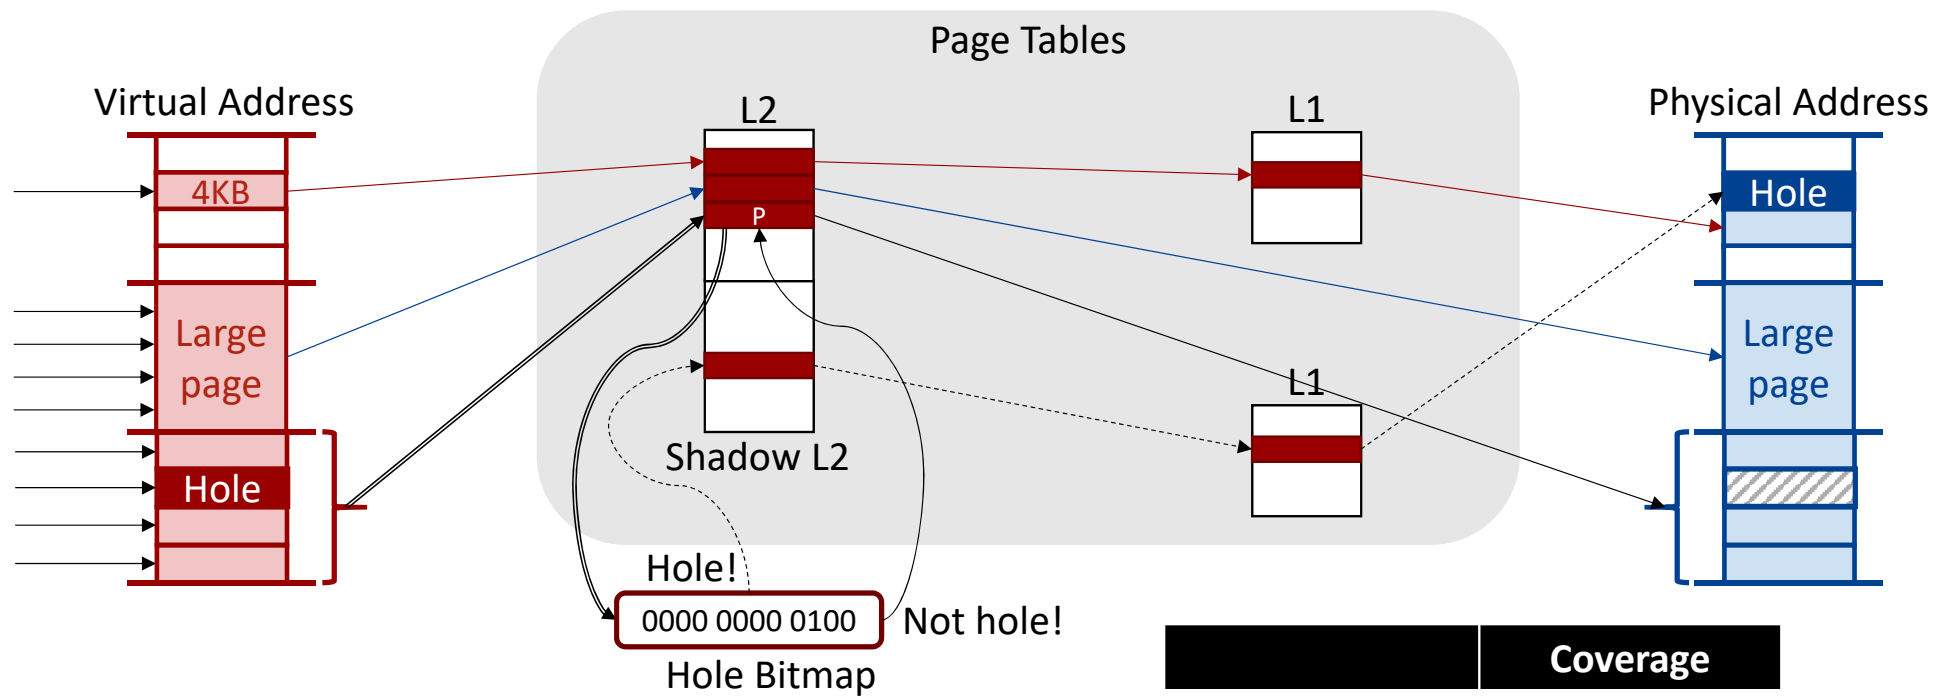


Where do we keep addresses for Hole pages?

Shadow L2 and Regular L1 Page Table!



How it works: Summary

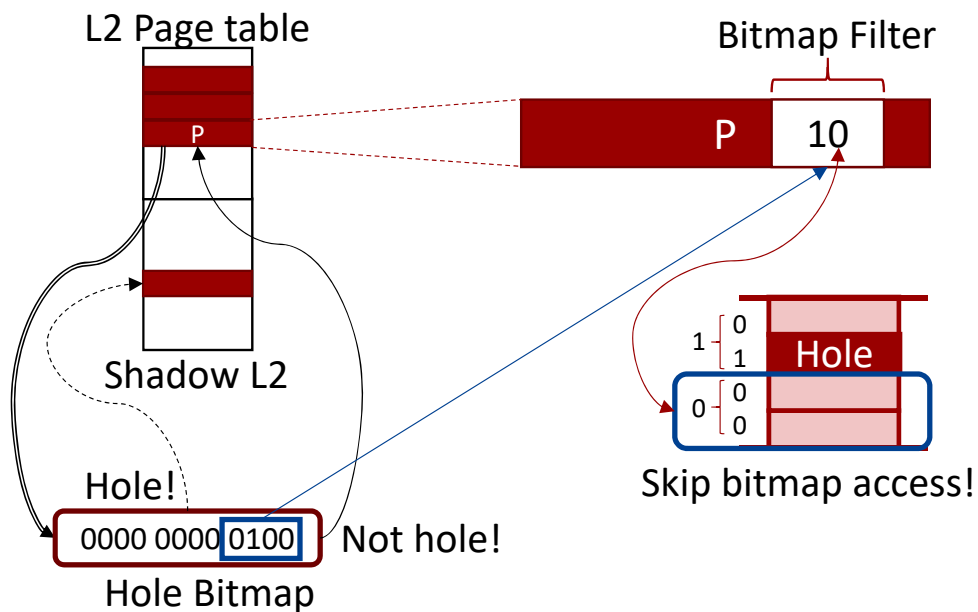


Additional accesses	Hole bitmap	Shadow L2
Perforated page walk	✓	

	Coverage
Large page	512x
Perforated page	(512 - # holes)x



How it works: Optimization



1. Coarse grain bitmap filter
 - Skip bitmap if definitely not hole
2. Hole bitmaps cached in TLB
 - 16 TLB entries per perforated page
 - Only insert accessed bitmap entries
3. Shadow L2 entries cached in Page table walker cache



Evaluation Methodology

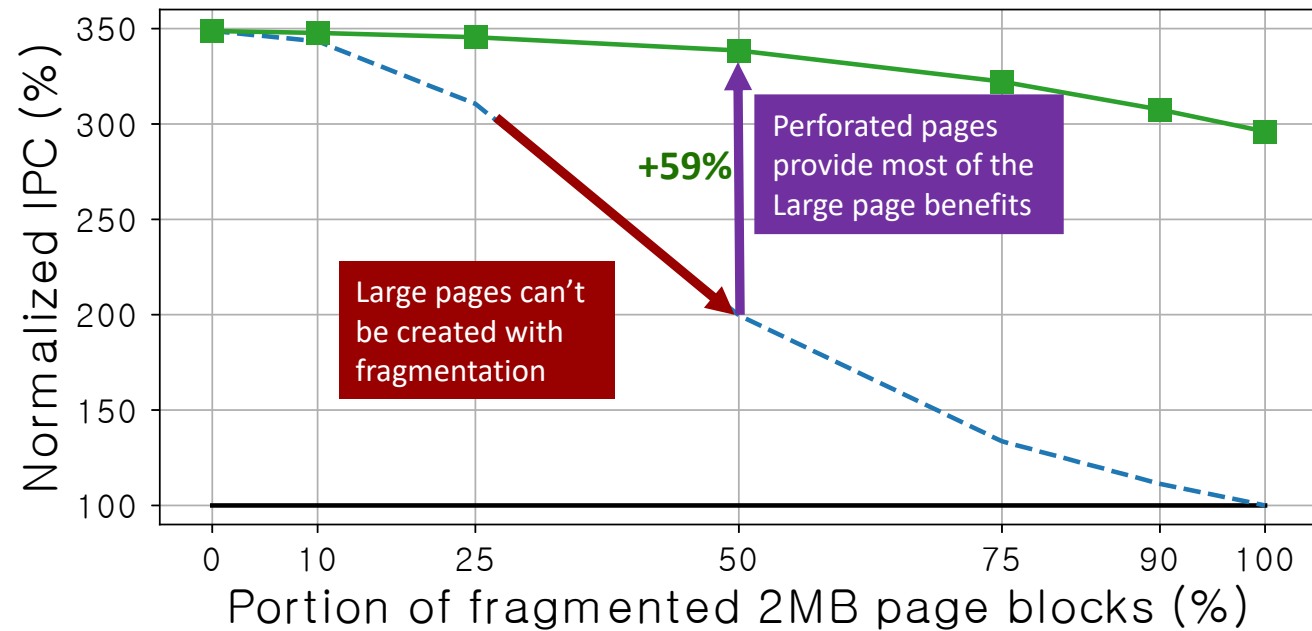
- Simulation configuration
 - Gem5
 - System-call emulation mode
- Microbenchmark
 - Random access (worst case)
- Real world benchmarks
 - SPEC CPU
 - Biobench

Component	Configuration
Processor	2GHz, OoO x86
Caches	32KB L1 I/D 2MB L2
Memory	DDR4-2400, 4 channels




Microbenchmark

— Baseline 4KB - - - Baseline 2MB+ 4KB ■ Perforated 25% holes



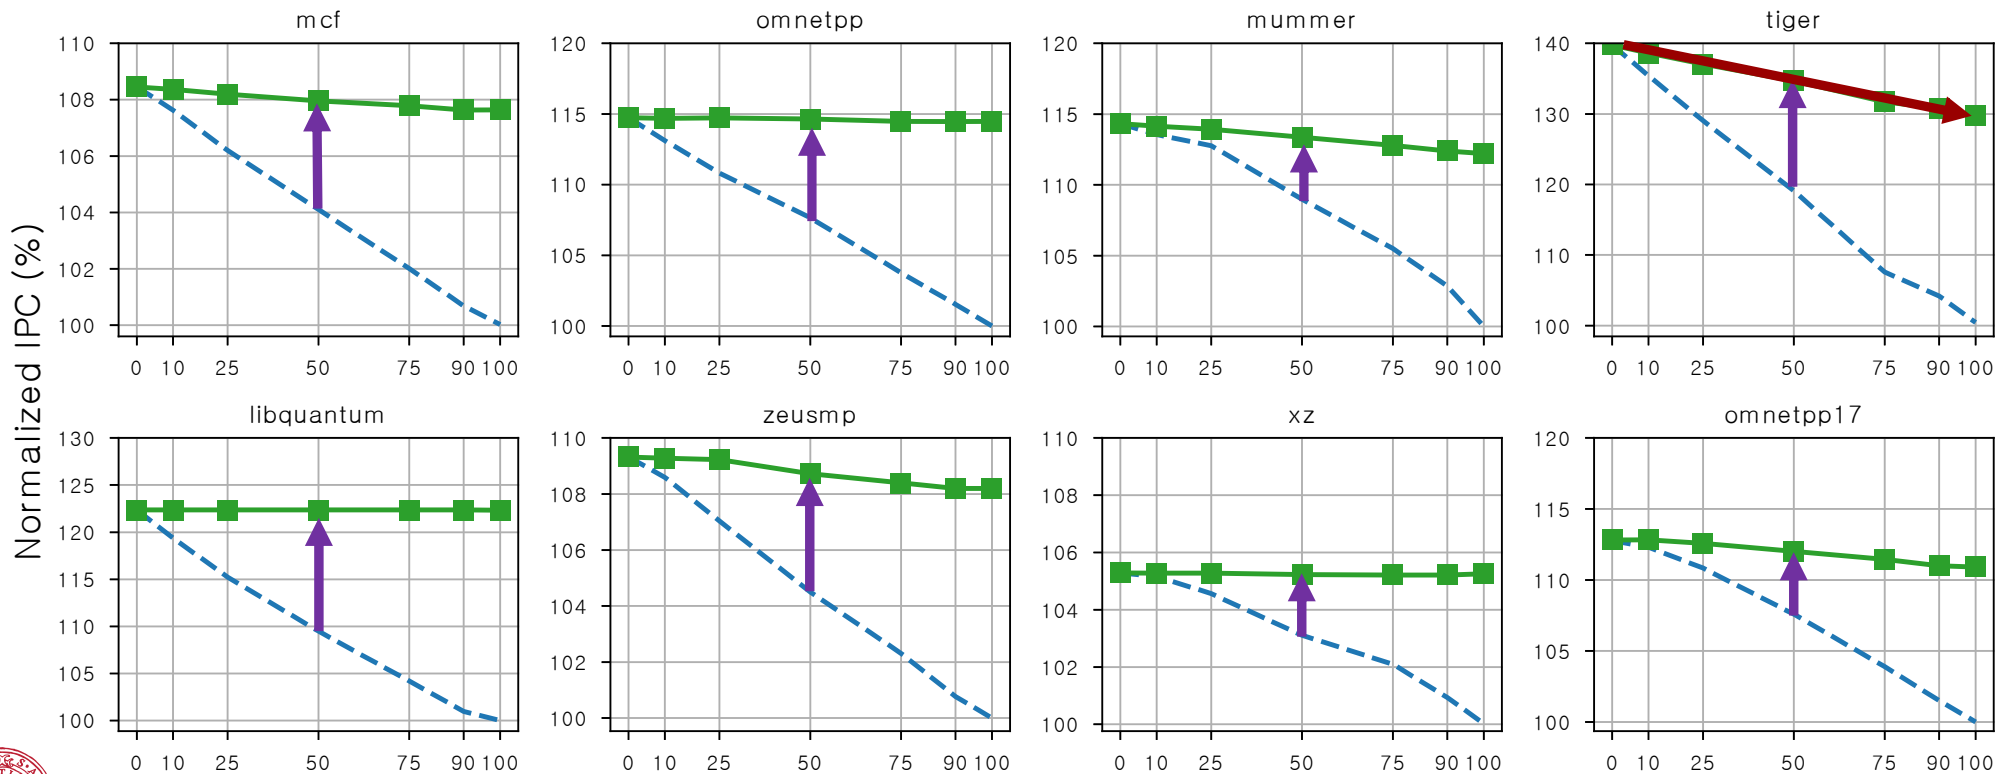
Breakdown of entries in L2 TLB

■ perforated pages ■ regular/hole 4KB
■ 2MB ■ hole bitmaps



Benchmarks

--- Baseline 2MB+ 4KB ■ Perforated 25% holes



2% - 11% performance improvement over Large page TLB

93% - 99% performance of ideal 2MB mapping (x=0)



More in the paper

- Details of TLB
- OS issues:
 - Advise for OS
 - TLB Shootdowns
- Virtualization Support
- More evaluation:
 - More hole % scenarios
 - Dispersed holes scenario
 - Virtualization
 - Comparison to prior work
- And more...



Conclusion

- Large pages deliver performance, but has challenges:
 - Contiguous, homogenous
 - Compaction, bloating
 - Immovable pages
- Perforated page provides **flexible** large page
 - Large-page translations for most of the data
 - Holes to handle pages that differ
- **Minimal changes** to existing translation HW and data structure
- Performance **similar to ideal** large page mappings
 - Retains 93-99% of performance



Perforated Page: Supporting Fragmented Memory Allocation for Large Pages

Chang Hyun Park, Sanghoon Cha, Bokyeong Kim,
Youngjin Kwon, David Black-Schaffer, and Jaehyuk Huh



UPPSALA
UNIVERSITET

