#### 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







2

## Large Pages for Performance

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





UPPSALA

**UNIVERSITET** 

[1] Margaritov et al., MICRO '19

## 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

UNIVERSITE

- E.g., kernel allocations, I/O buffers, etc.
- After Linux kernel build, 50% of 2MB regions had immovable pages





**Physical Addresses** 

Fragmented

Compacted

## **Overheads: Memory Bloating**

• Sparse access/use in large pages  $\rightarrow$  wasted physical space



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



UNIVERSITET





Large pages for **Performance** 

• Better TLB Coverage  $\rightarrow$  Better Performance

Large pages are **Difficult** to make

- Require contiguous regions
- Require homogeneous permissions

Large pages come with **Overheads** 

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



UNIVERSITE





Large

Difficu

Large pages for **Performance** 

#### Better TLB Coverage → Better Performance

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

Large pages come with **Overheads**  • Costly compacting to create contiguous regions

ons

- 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

- 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

UNIVERSITE

- Conserve untouched pages via holes
- 0% bloating, 17% TLB MPKI reduction (2MB: 45% TLB MPKI reduction)





11

### How it works







#### How it works: Basics



#### How it works: Basics



## How it works: Summary



UPPSALA UNIVERSITET

### 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
  - SPECCPU
  - Biobench

UPPSAL

UNIVERSITE

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



#### Microbenchmark

UNIVERSITET

#### Breakdown of entries in L2 TLB



#### Benchmarks





19

## 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







### 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



UNIVERSITE

• 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

