

Towards Speculative Offset-Based Address Translation

Chloe Alverti, Vasileios Karakostas

xalverti@cslab.ece.ntua.gr, vkarakos@cslab.ece.ntua.gr

Computing Systems Laboratory - National Technical University of Athens



1. MOTIVATION & BACKGROUND

Research Problem: Virtual Memory Overhead

Translation Lookaside Buffers miss penalty (page walks).

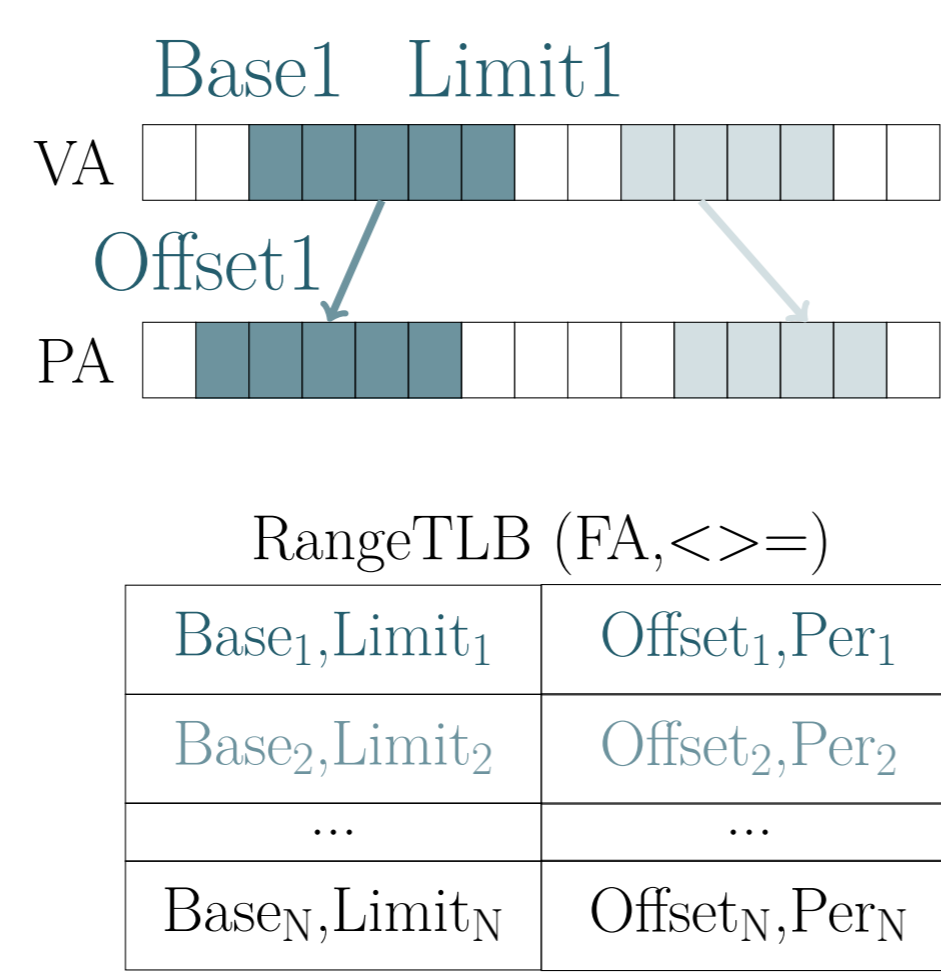
- ▶ Performance bottleneck for data intensive workloads

Prior Work: Redundant Memory Mappings [4]

Ranges→Unaligned contiguous virtual-to-physical mappings.
Eager paging→Allocation at request time to boost contiguity.
Range TLB→Accelerate translation.

Motivation: Robust Translation Performance

Eager paging→ insufficient to extract available contiguity, sensitive to fragmentation, narrows memory overcommitment.
Range TLB→Ranges can gradually break→Thrashing.



2. CONTIGUITY-AWARE PAGING

Demand paging: single random page allocation at access time.
Eager paging: high order block preallocation at request time.

Idea

Enhance system allocator with **contiguity awareness (CA)**.
Build vast offset-based ranges on a best effort basis.
▶ Boost and shield contiguity
▶ Opportunity to relax preallocation

Mechanism

Linux `struct vma` → contiguous virtual address range.
▶ new `Offset` attribute, track and store `pa - va`

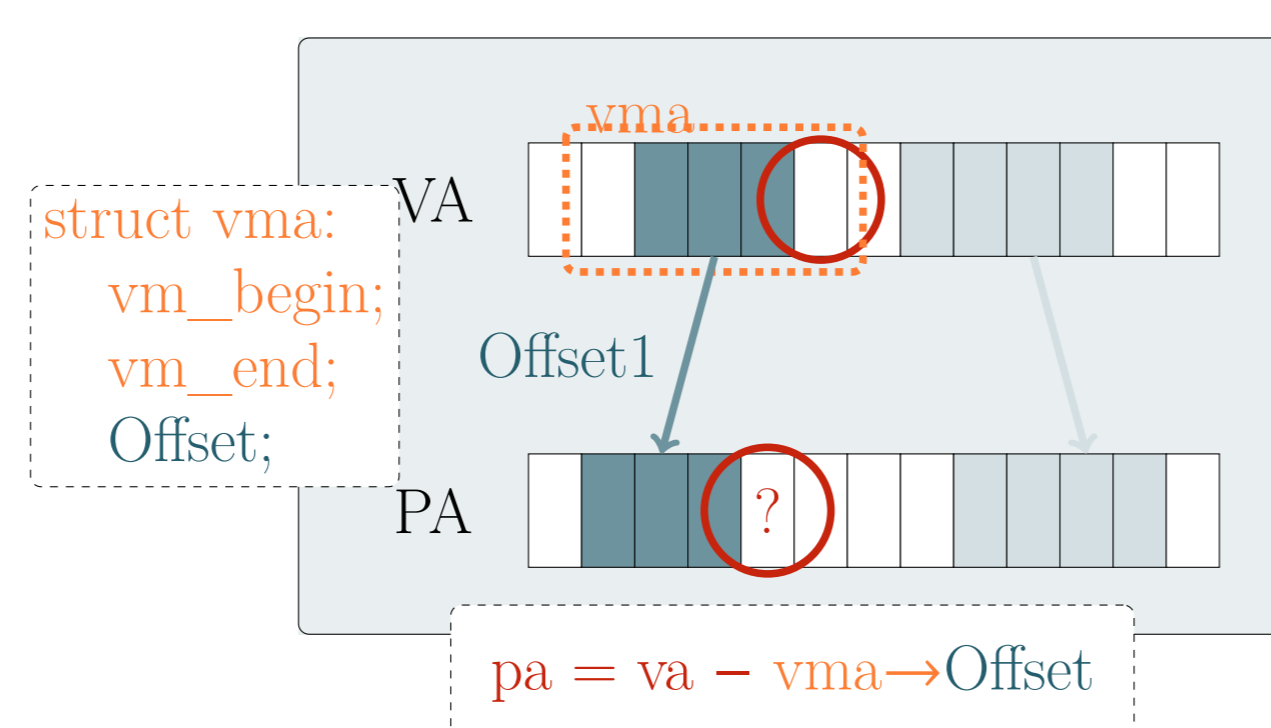
How CA-paging works on allocation requests:

- Identify physical block with respect to `vma`→`Offset`
- Available? check → `struct page* mem_map array`
- If free, reserve it from buddy allocator
Fallback → default, pick a random block
- Establish mapping → Range extended

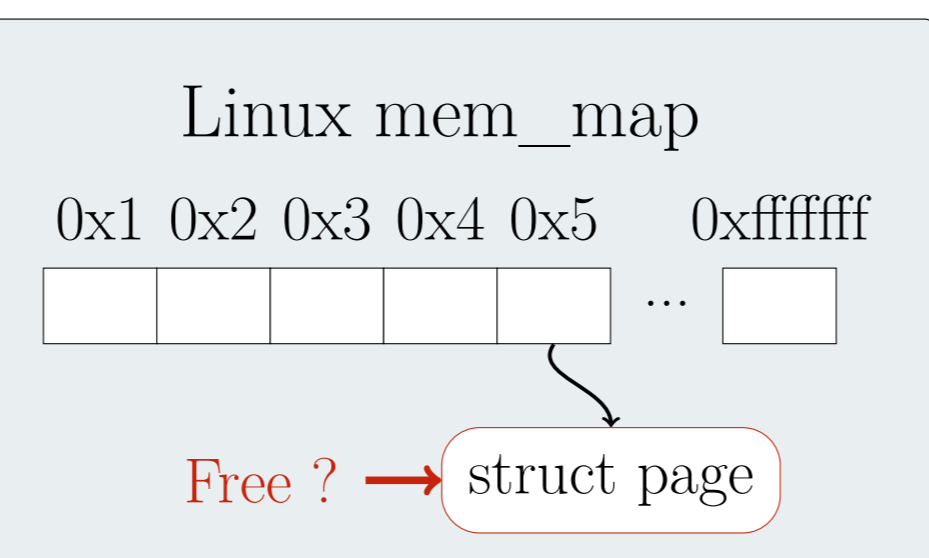
User-Level Application

```
array = malloc(2^k); ←eager (order=k)
for (i=0; i<2^order; i++)
    array[i]=i; ←demand (order=0)
```

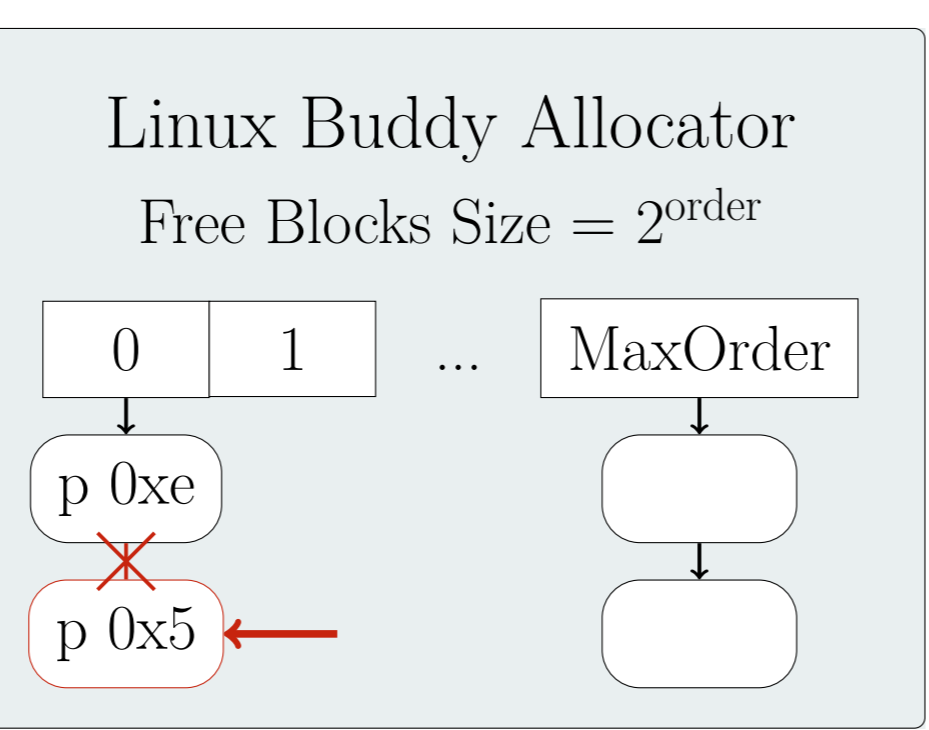
a. Request to map a VA



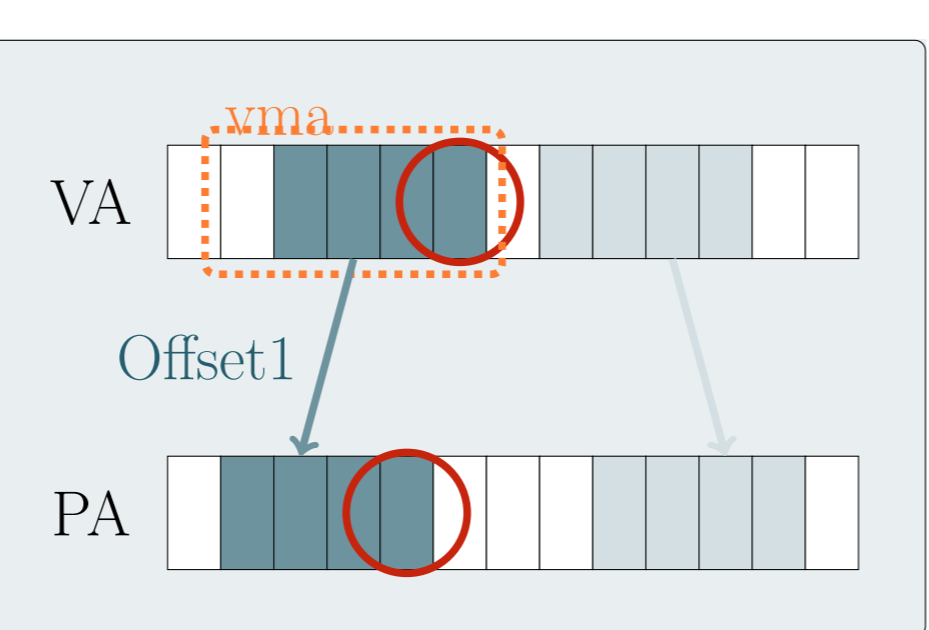
b. Preferred PA available?



c. Allocate PA



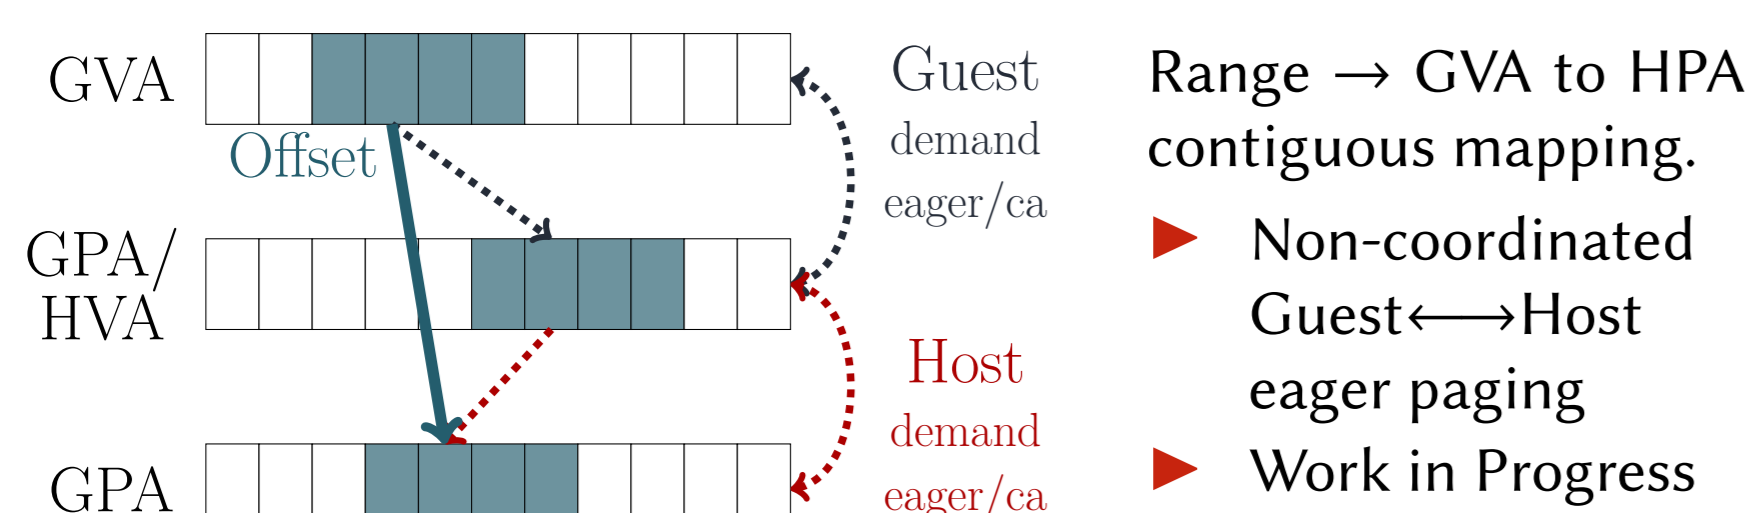
d. Establish Mapping



Advantages

- ▶ Range construction across requests
- ▶ Resilience to external fragmentation
- ▶ Independence of alignment restrictions
- ▶ Orthogonal to eager or demand paging

Virtualized Environments

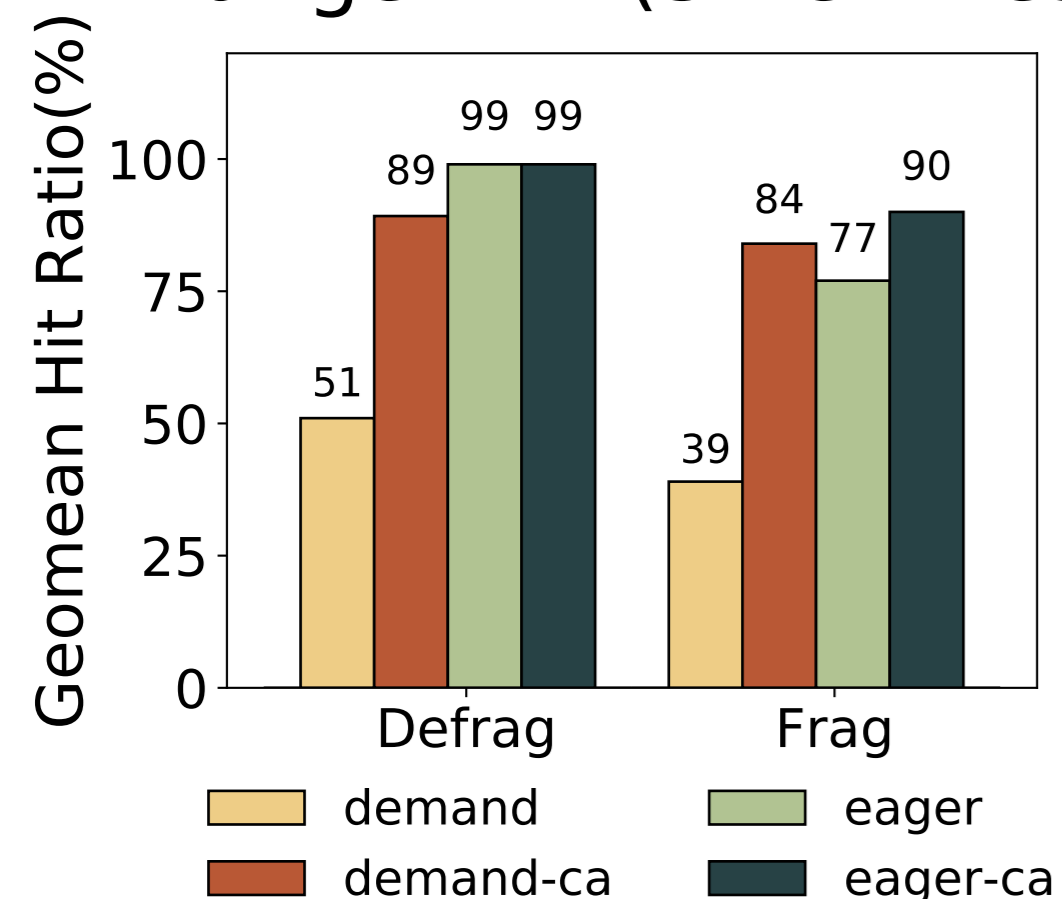


Methodology

We implement CA-paging in Linux 3.15.5.
Badger Trap [2] → RangeTLB performance native execution.
In-house VM Introspection tool → GVA to HVA mappings.
Benchmarks: data-intensive SPEC2006/Parsec workloads.

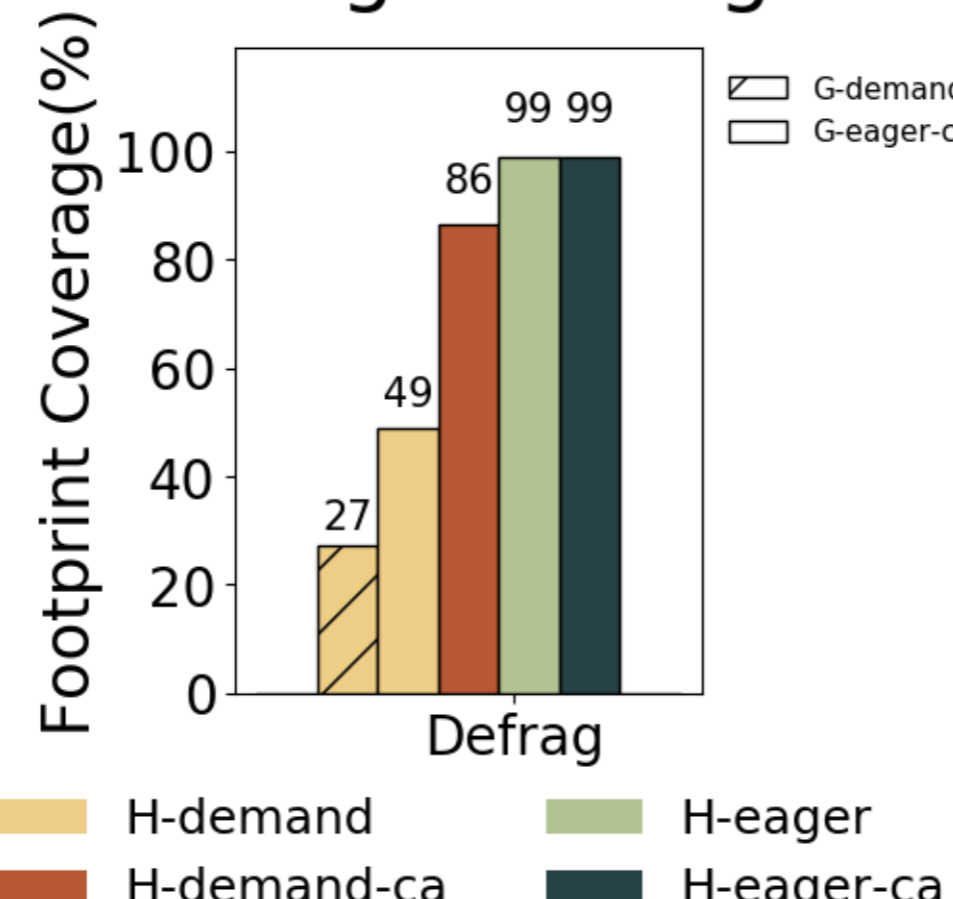
Native Execution

Range TLB (32 entries)



Virtualized Execution

32 Largest Ranges



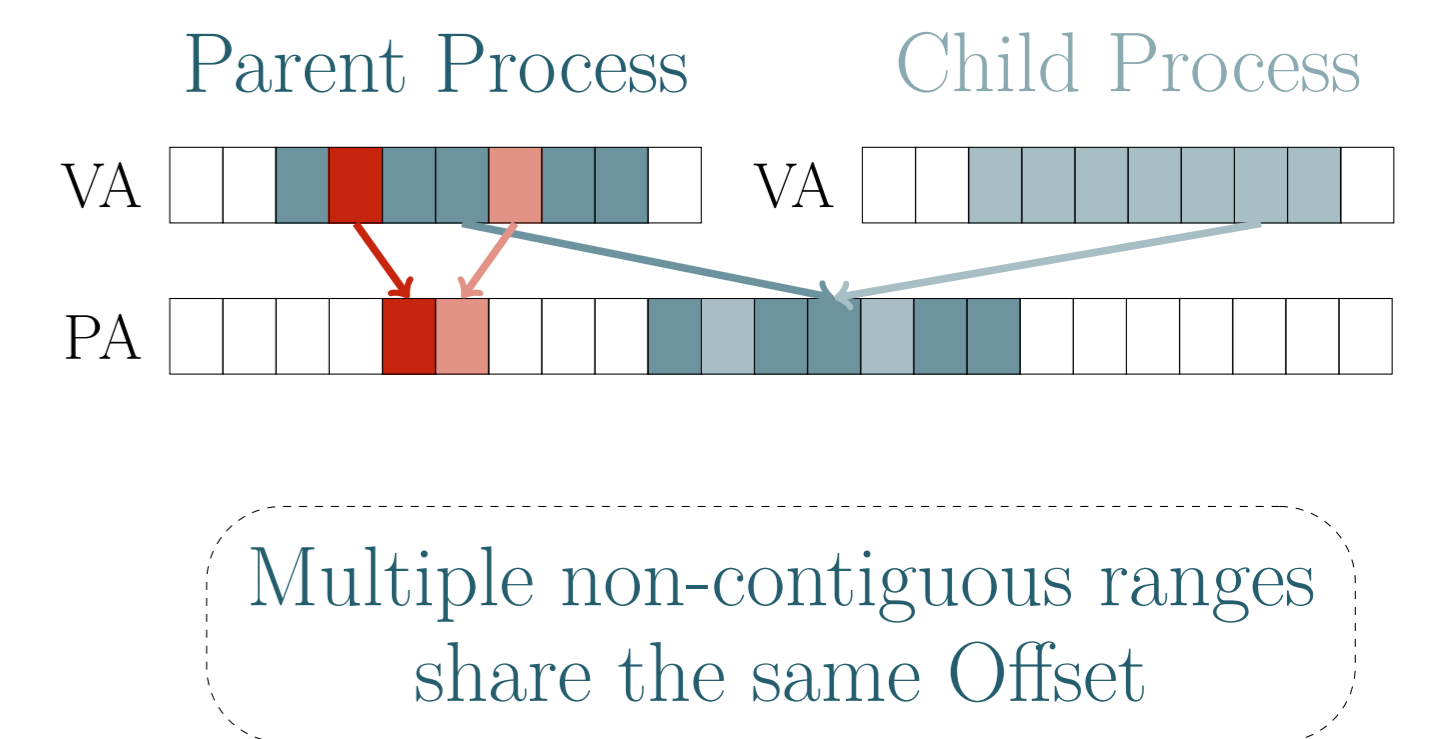
~84-89% performance with demand paging
~13% improvement fragmented environment

3. DISCONTINUITIES IN RANGES

Challenge

Discontinuities in ranges may appear due to:

- ▶ External Fragmentation
- ▶ Less Aggressive Allocator (i.e demand, eager)
- ▶ Lightweight Memory Management Techniques gradually breaking ranges
 - Copy-on-write
 - Working set sampling
 - Memory deduplication (KSM)



4. OFFSET-BASED SPECULATION

Motivation

As the number of ranges increases, scaling Range TLB alone is not a sustainable solution due to:

- ▶ fully-associative lookups
- ▶ range comparisons

Idea

Ranges Discontinuities → opportunity to **speculate** the address translation based on **Offset**.

Our Proposal

- ▶ Decouple Offsets from virtual range boundaries
- ▶ Speculate `pa = va - Offset`
- ▶ Verification page table walks on the background
- ▶ Misprediction → flush and replay

Preliminary Mechanism

Known observation → Only a few instructions are responsible for the majority of the TLB misses.
Index speculation structure by Program Counter.

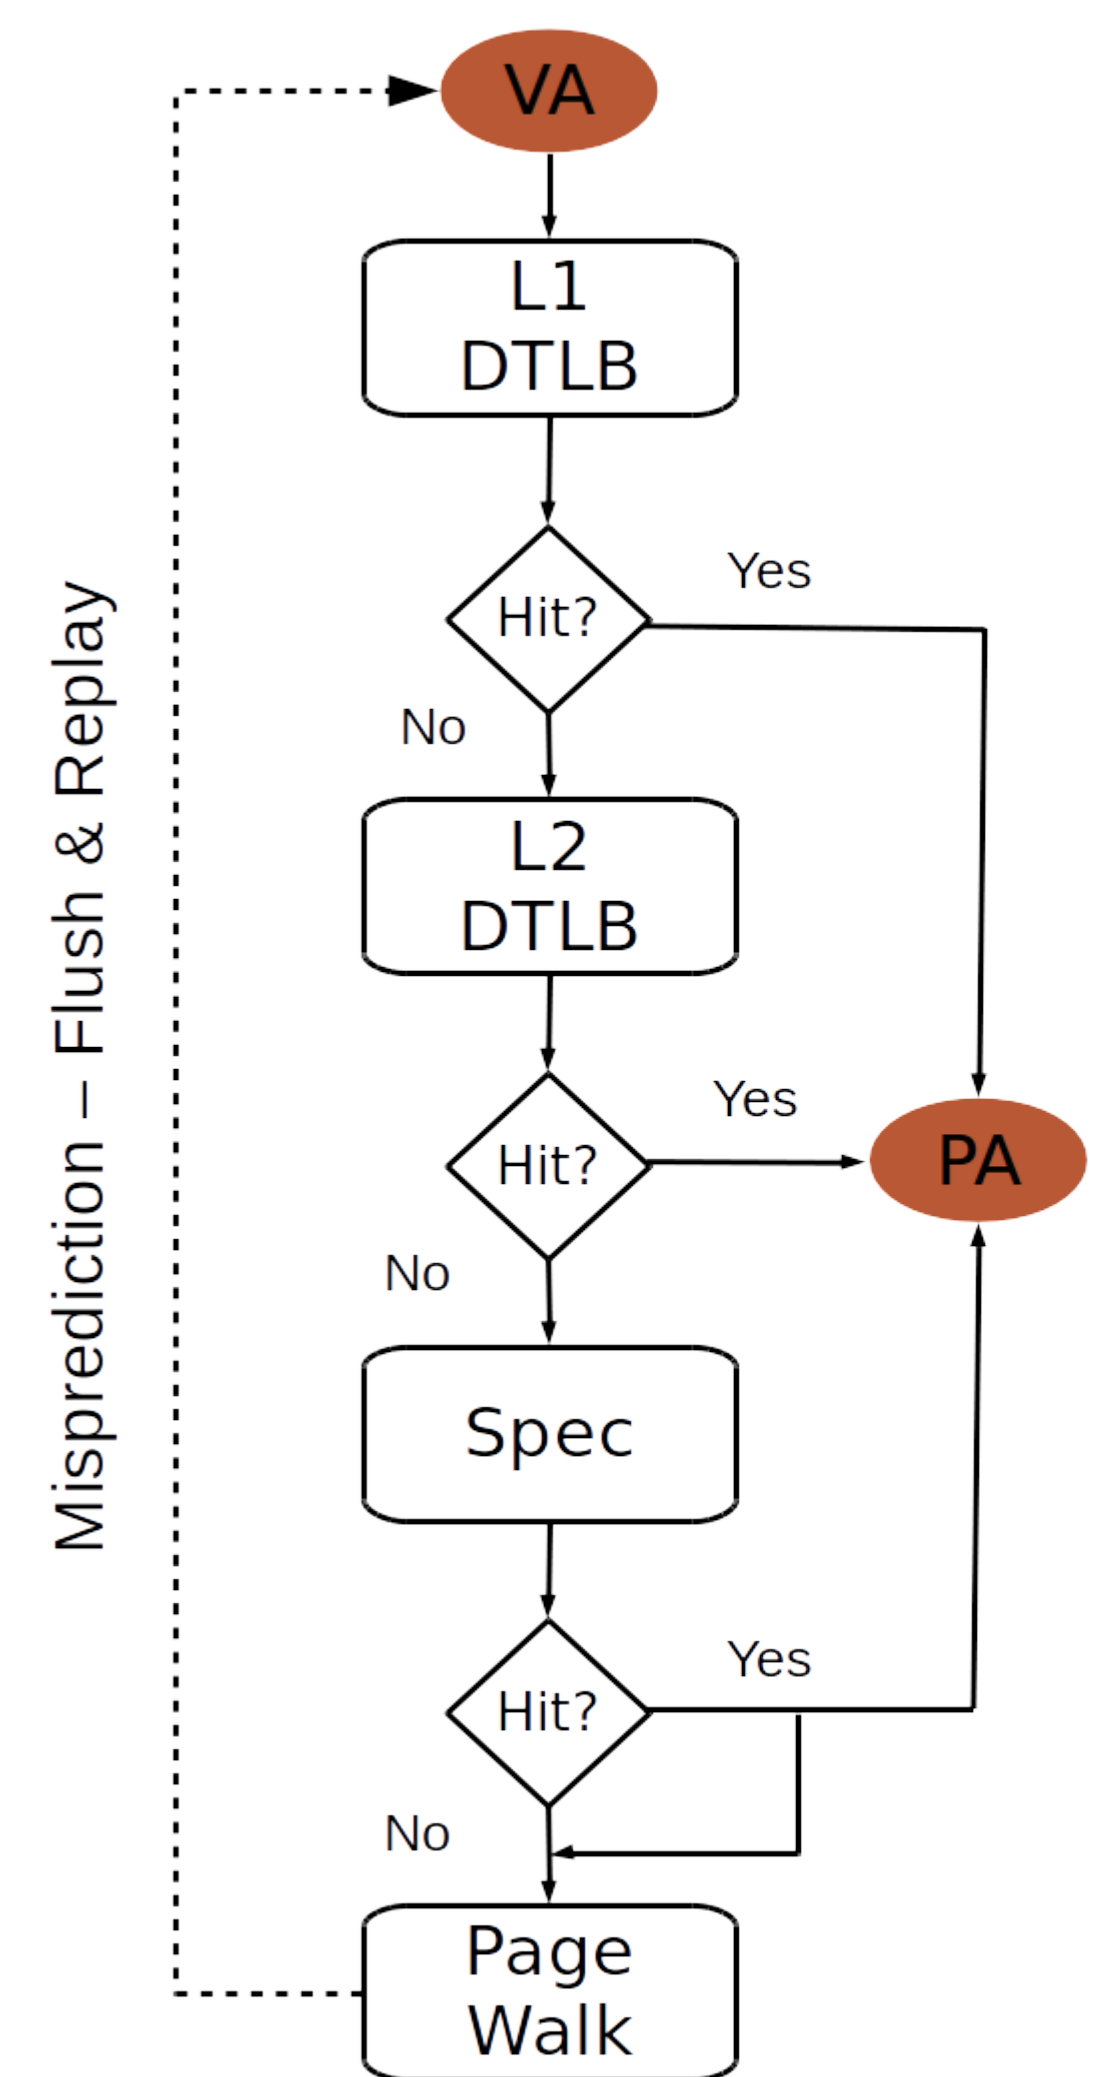
- Track the last Offset per PC
- Track 4 LRU Offsets per PC → Oracle selects the correct one → evaluate potential

Methodology

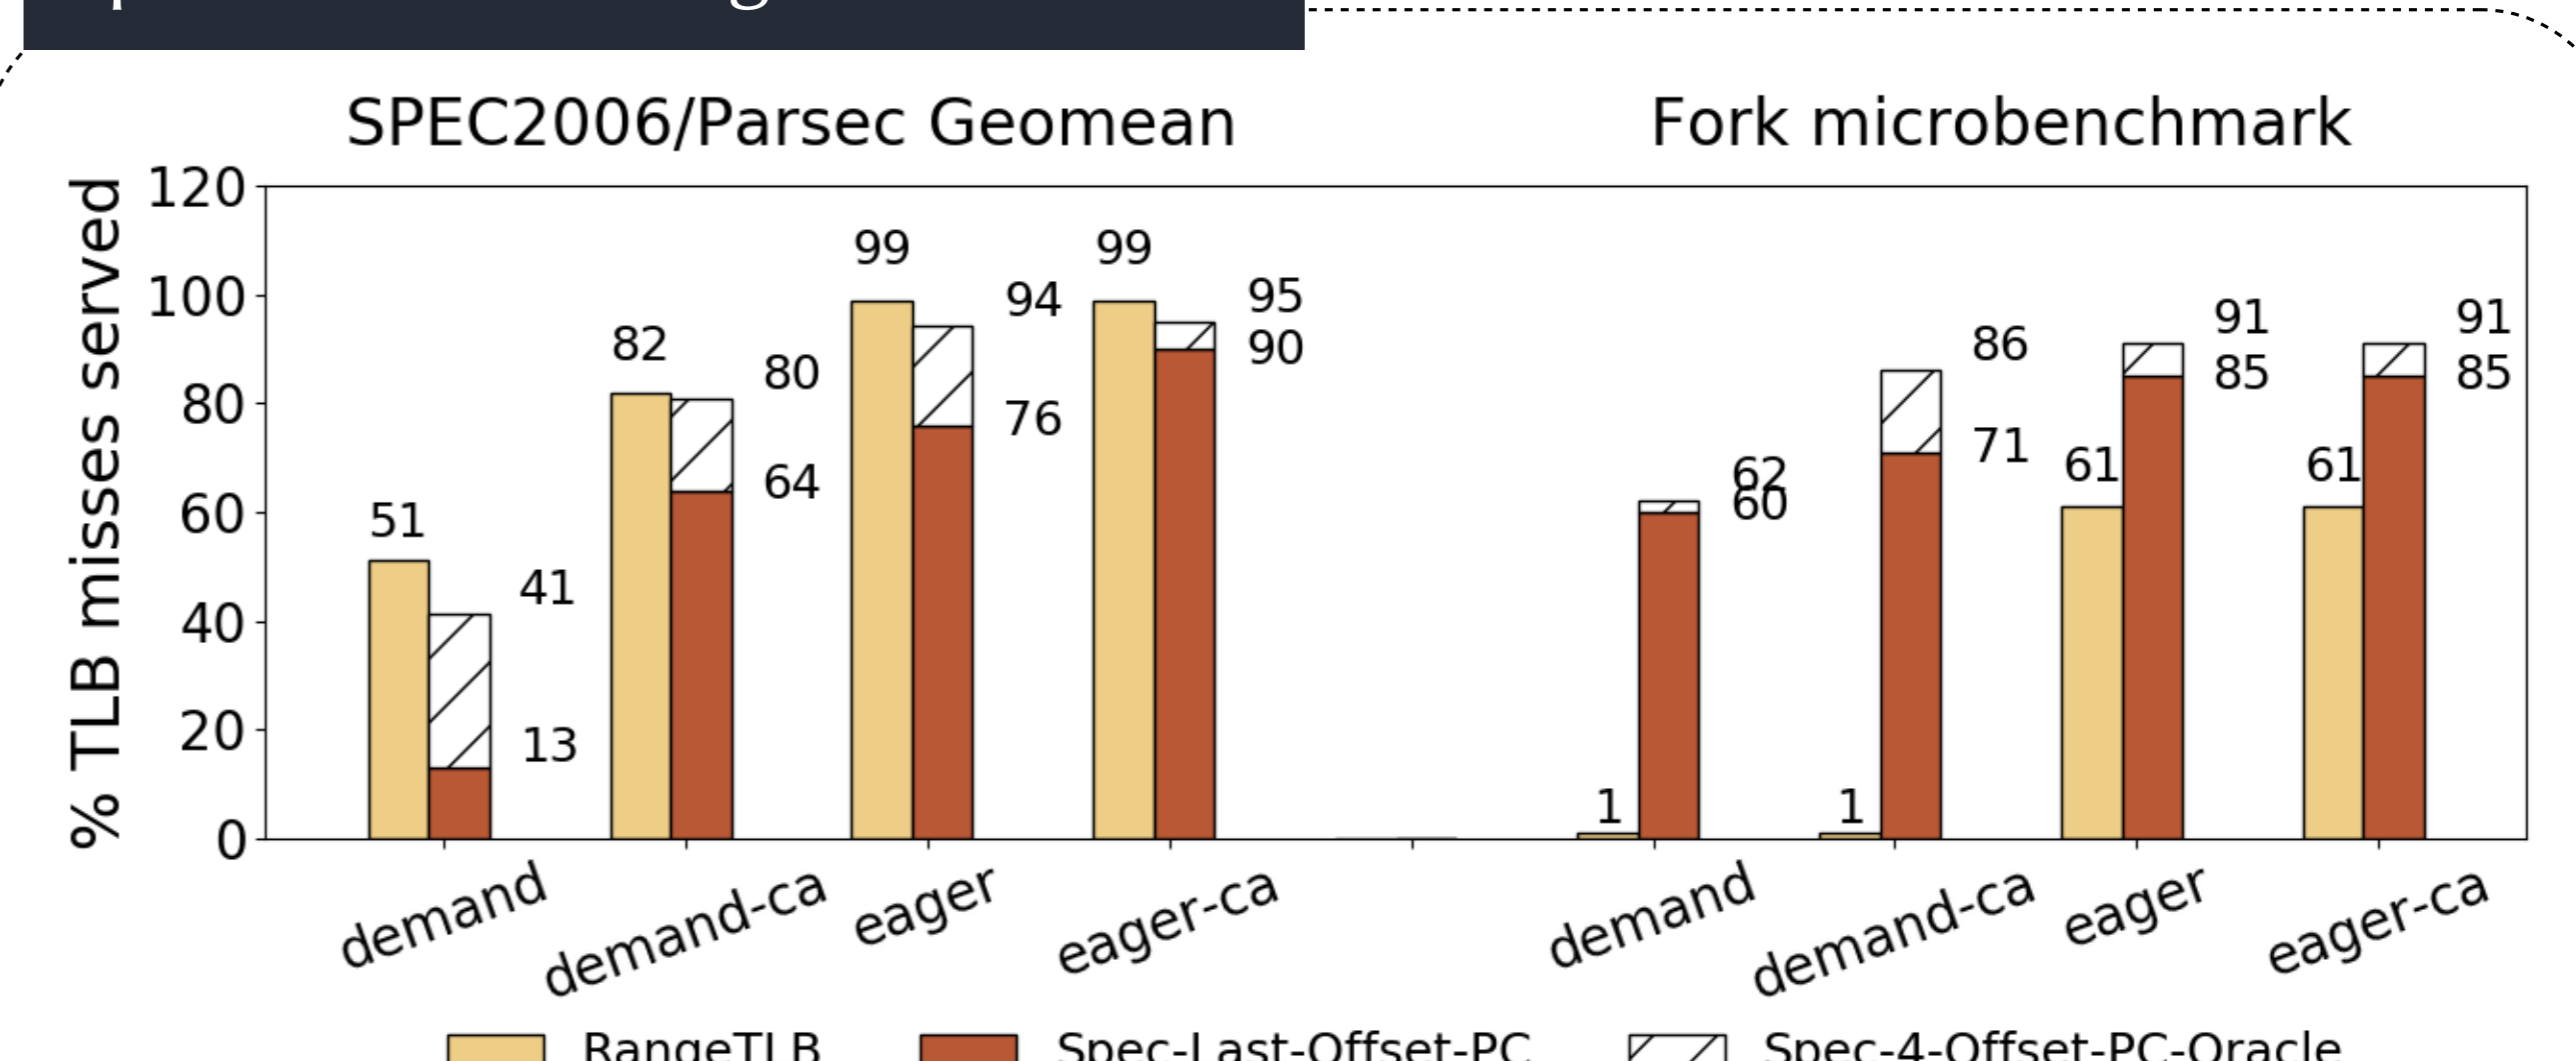
Speculation could be an auxiliary translation mechanism. For preliminary evaluation, we adopt the challenging scenario of replacing RangeTLBs.
Badger Trap [2] → RangeTLB and Speculation.

Spec Structure

PC ₁	Offset ₁
PC ₂	Offset ₂
...	...
PC _N	Offset _N



Speculation vs RangeTLB - Native



~90% performance at optimal (eager-ca) conditions
Potentially close to RangeTLB performance
~20% better performance than RangeTLB for Fork microbenchmark

- ▶ Preliminary results from (VM Introspection tool) indicate that on a multiple VM environment with KSM, tracking offsets instead of ranges covers 10% more of an application's footprint.

[1] T. Barr, A. Cox, and S. Rixner. SpecTLB: A mechanism for speculative address translation. In *ISCA 2011*.
[2] J. Gandhi and A. Basu et al. Badgertrap: A tool to instrument x86-64 tlb misses. *SIGARCH Comput. Archit. News*, 2014.
[3] S. Haria, M. Hill, and M. Swift. Devirtualizing memory in heterogeneous systems. In *ASPLOS 2018*.
[4] V. Karakostas and J. Gandhi et al. Redundant memory mappings for fast access to large memories. In *ISCA 2015*.
[5] Binh Pham et al. Large pages and lightweight memory management in virtualized environments: Can you have it both ways? In *MICRO 2015*.
[6] T. Zheng, H. Zhu, and M. Erez. SIPT: Speculatively Indexed, Physically Tagged Caches. In *HPCA 2018*.