


# Rethinking Address Translation in the Age of Containers, Serverless Computing, and Non-Volatile Memory

**Josep Torrellas, Dimitrios Skarlatos, Apostolos Kokolis, Tianyin Xu**

Department of Computer Science  
University of Illinois at Urbana-Champaign  
<http://iacoma.cs.uiuc.edu>


ARM Research Summit  
September 2020

# New Era in Cloud Computing: Containers



docker

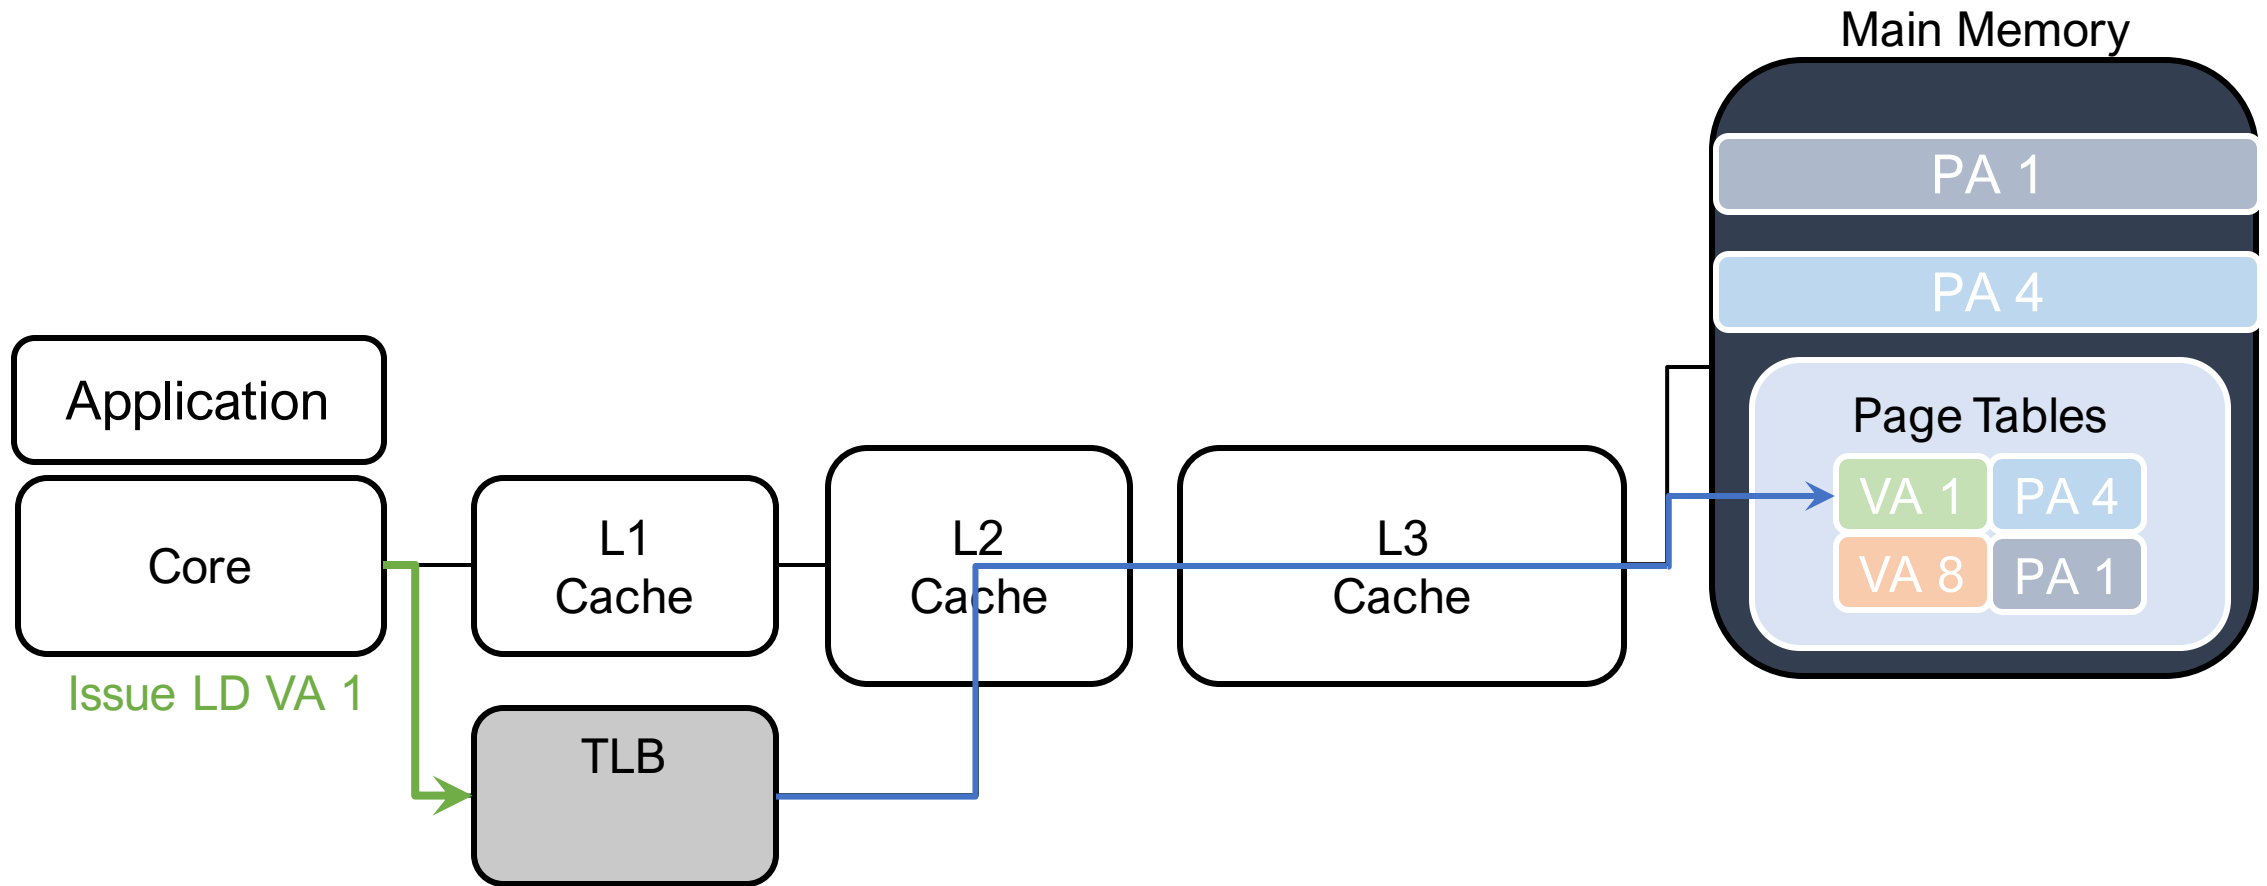
1. Lightweight
2. Faster bring-up
3. Higher consolidation



High number of containers

Very frequent context switches

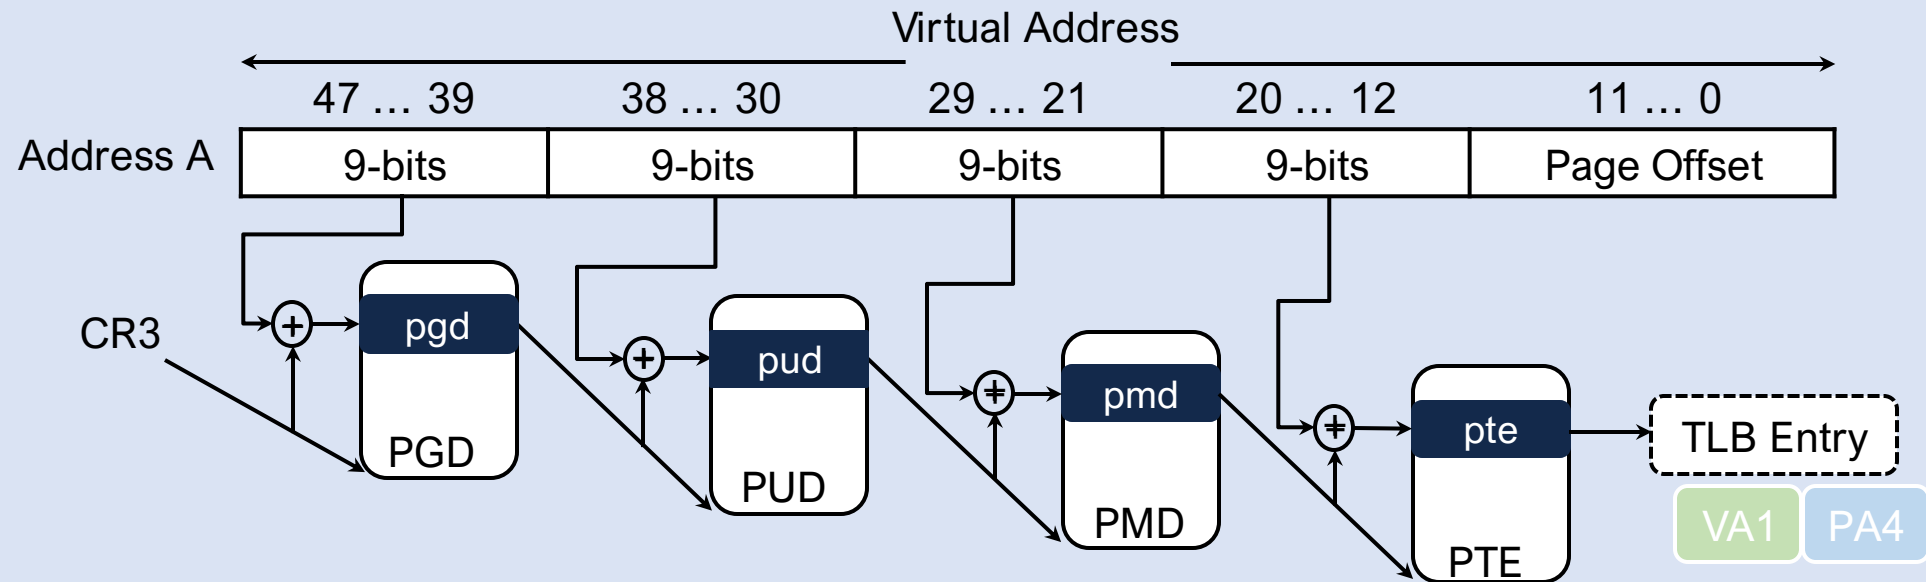
# Address Translation System



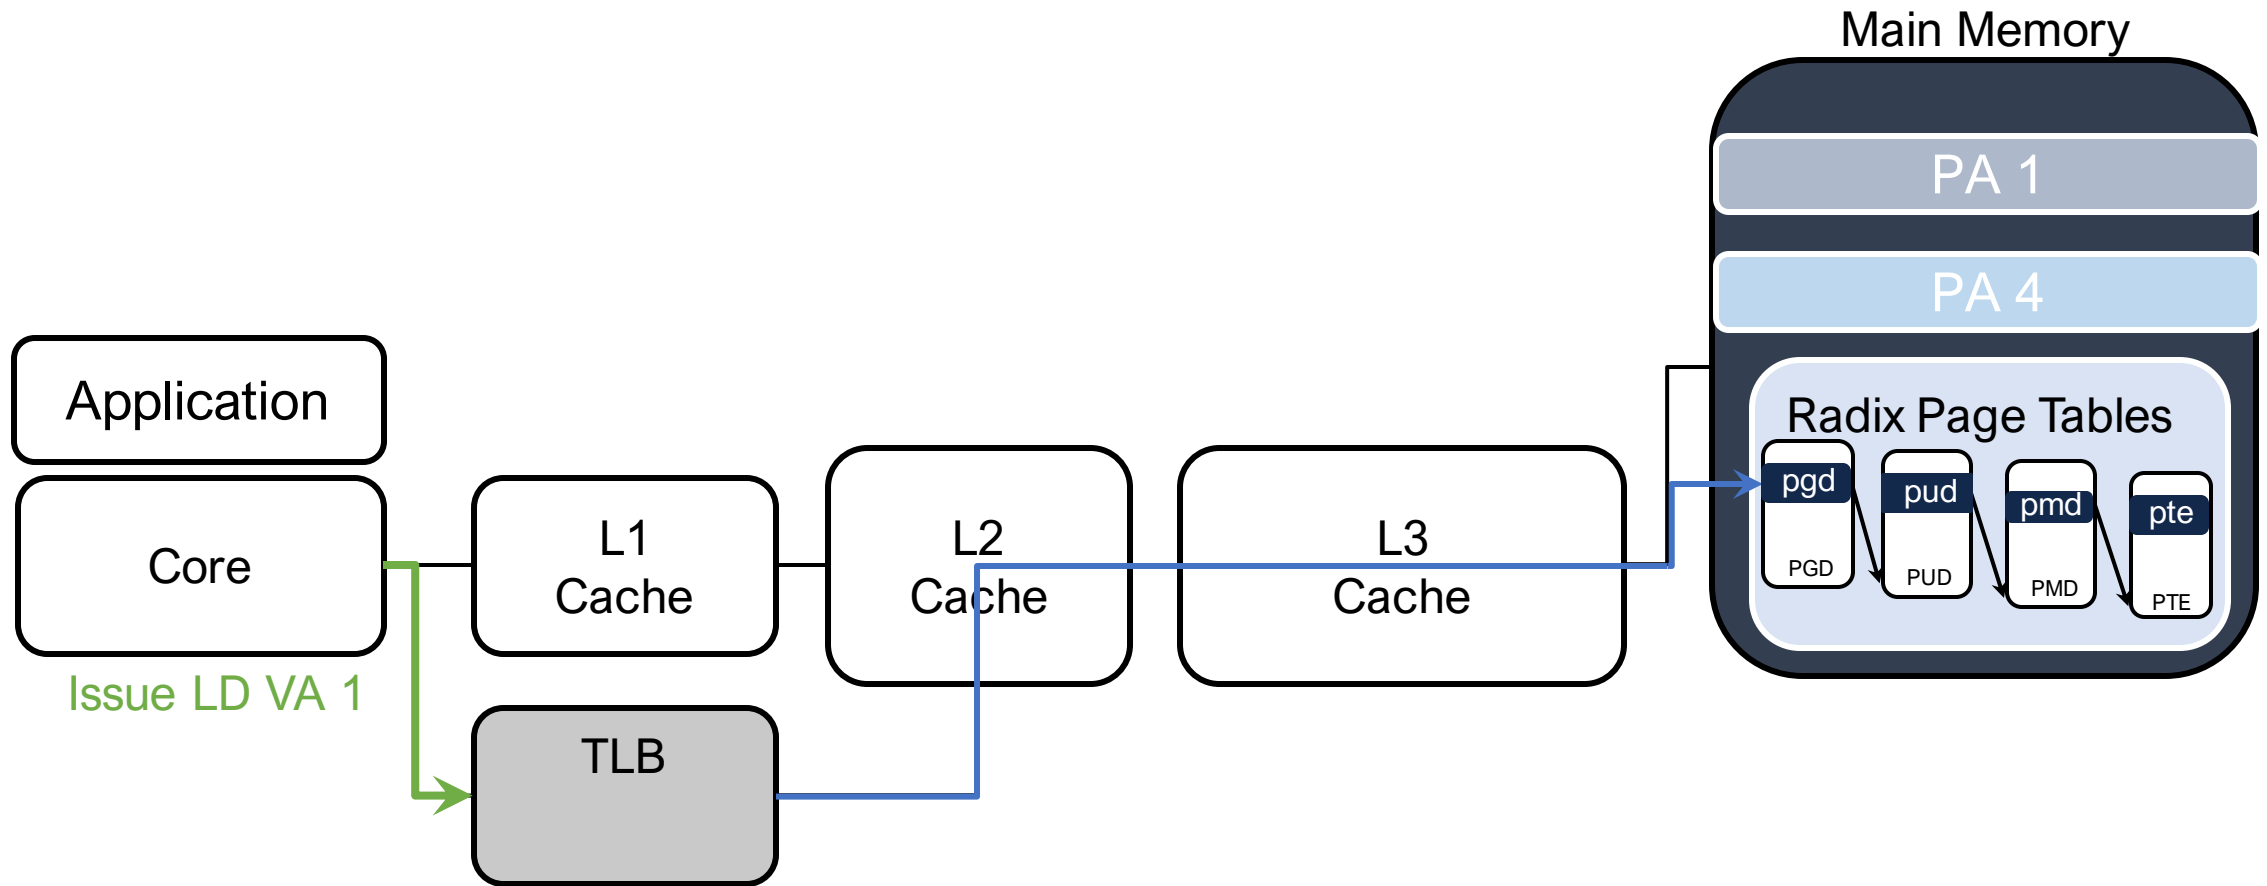
TLB Miss → “Page Walk” = Fetch entry from page table

# Address Translation System

## x86-64 Radix Page Tables

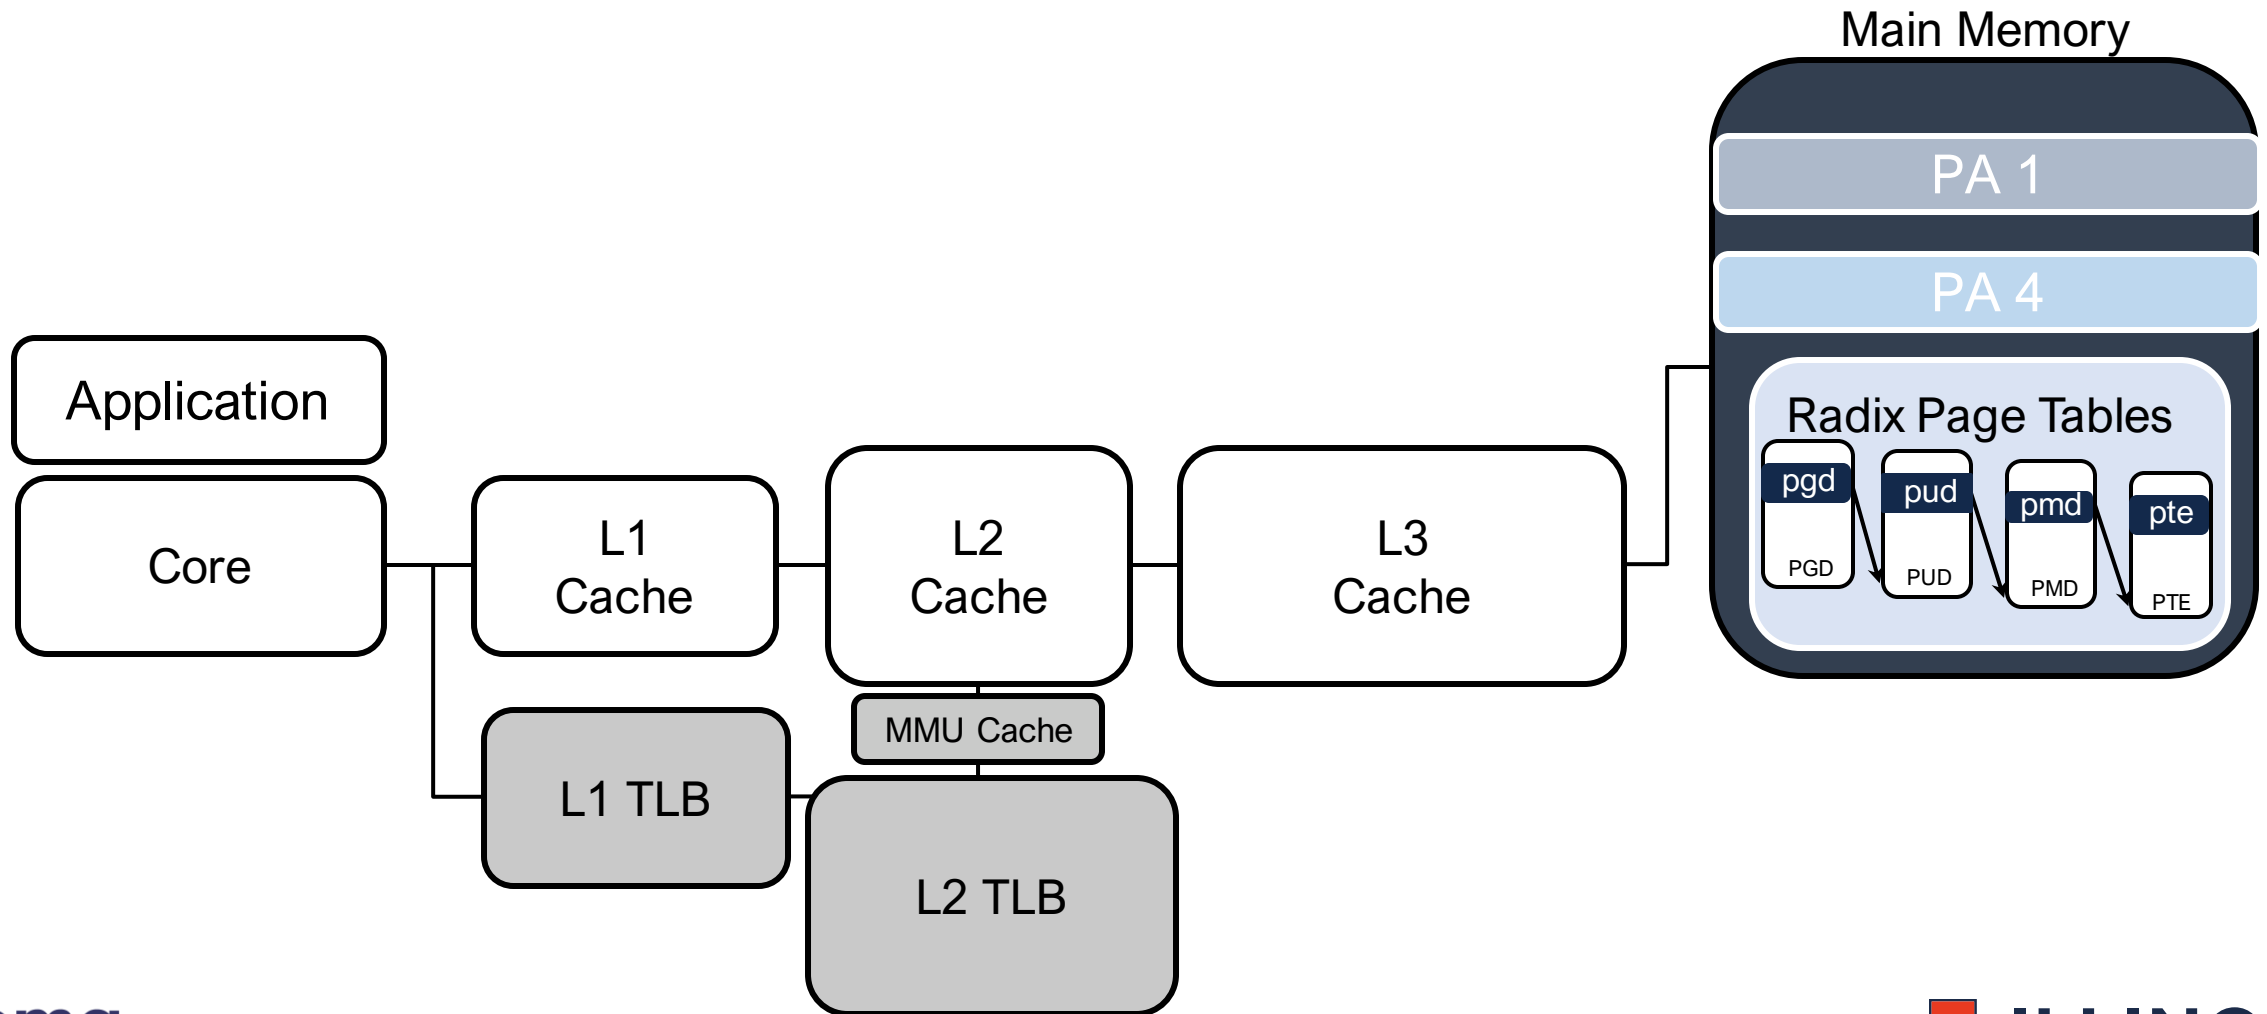


# Address Translation System



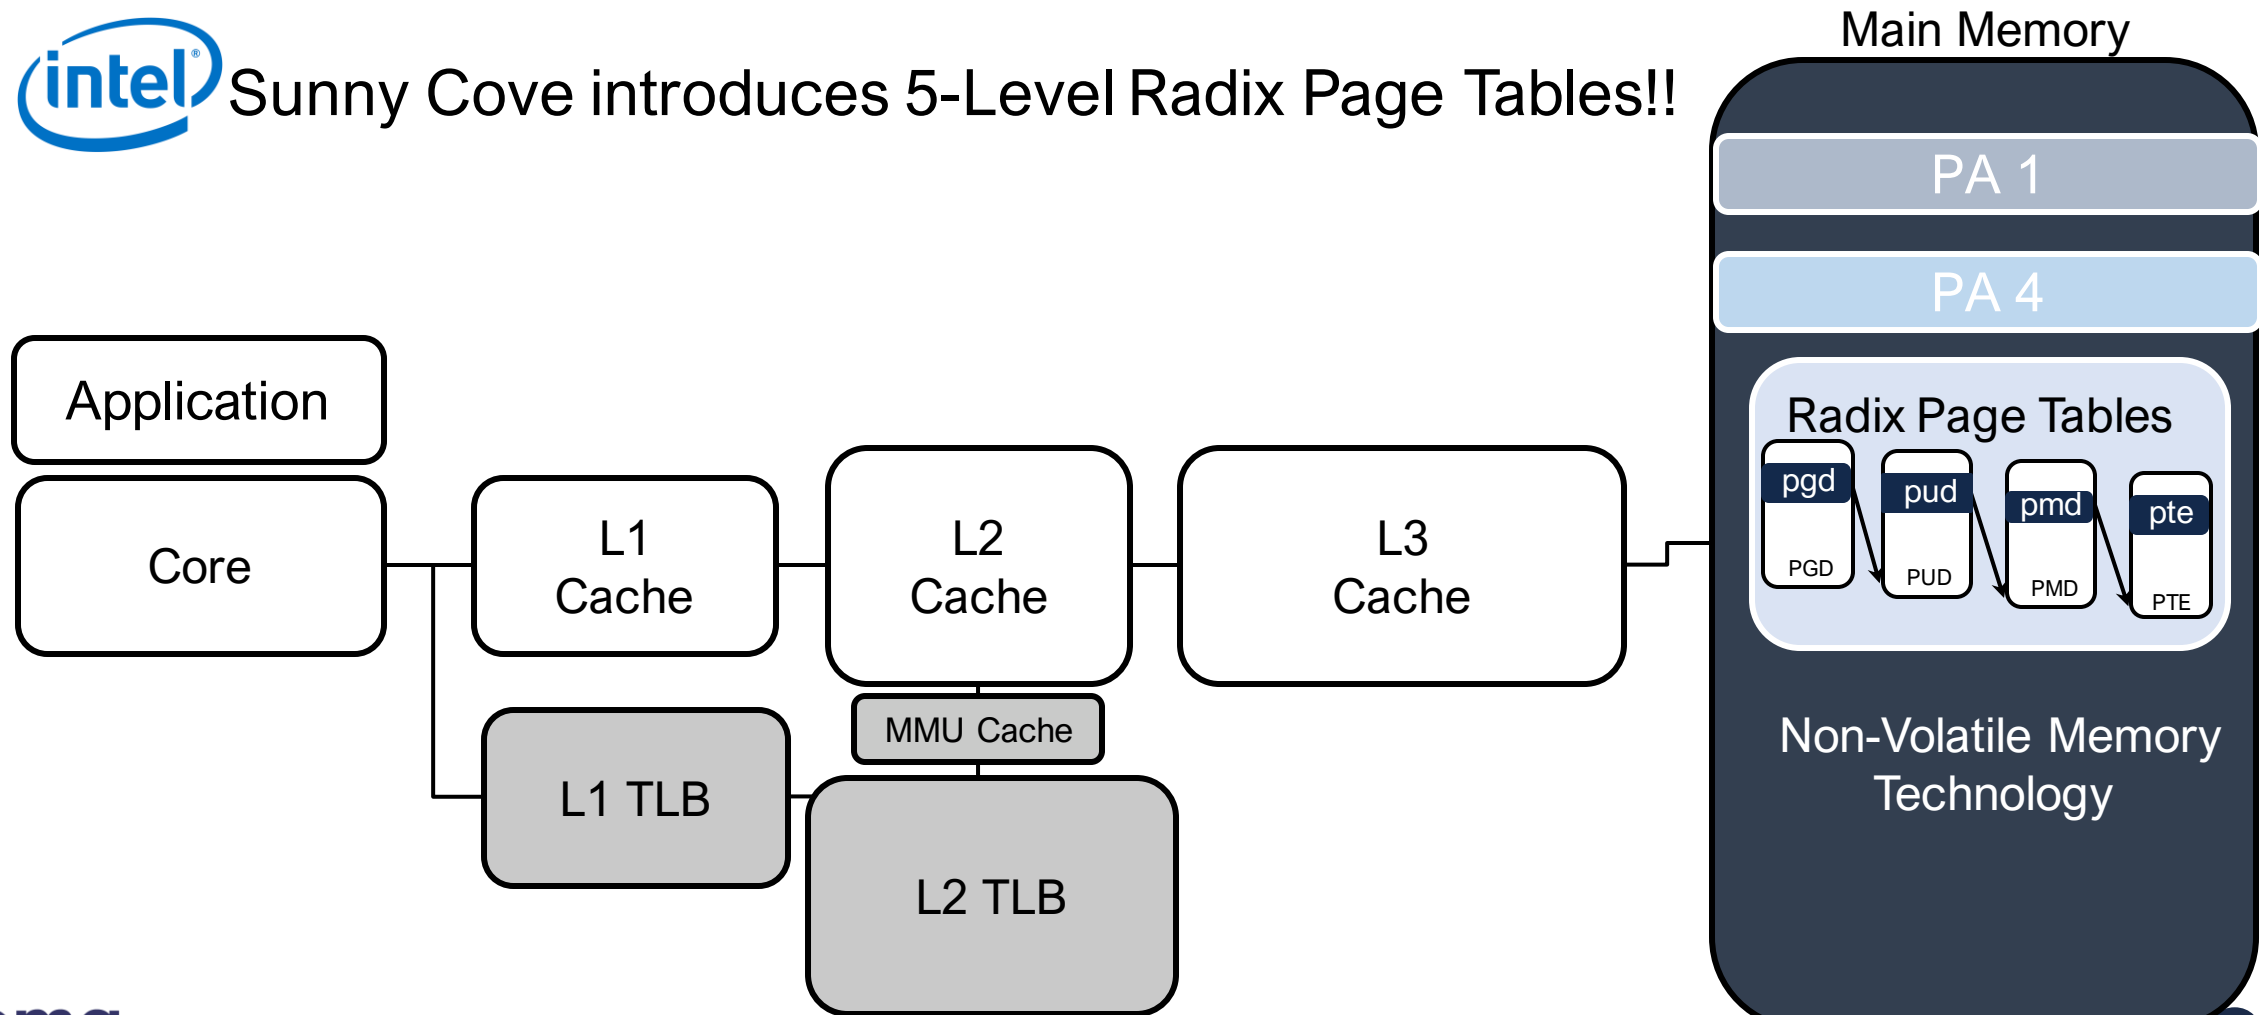
TLB Miss → “Page Walk” = Fetch entry from radix page table

# Address Translation System



# Even More Memory is Here!

 Sunny Cove introduces 5-Level Radix Page Tables!!




# Time to Rethink Virtual Memory Translation

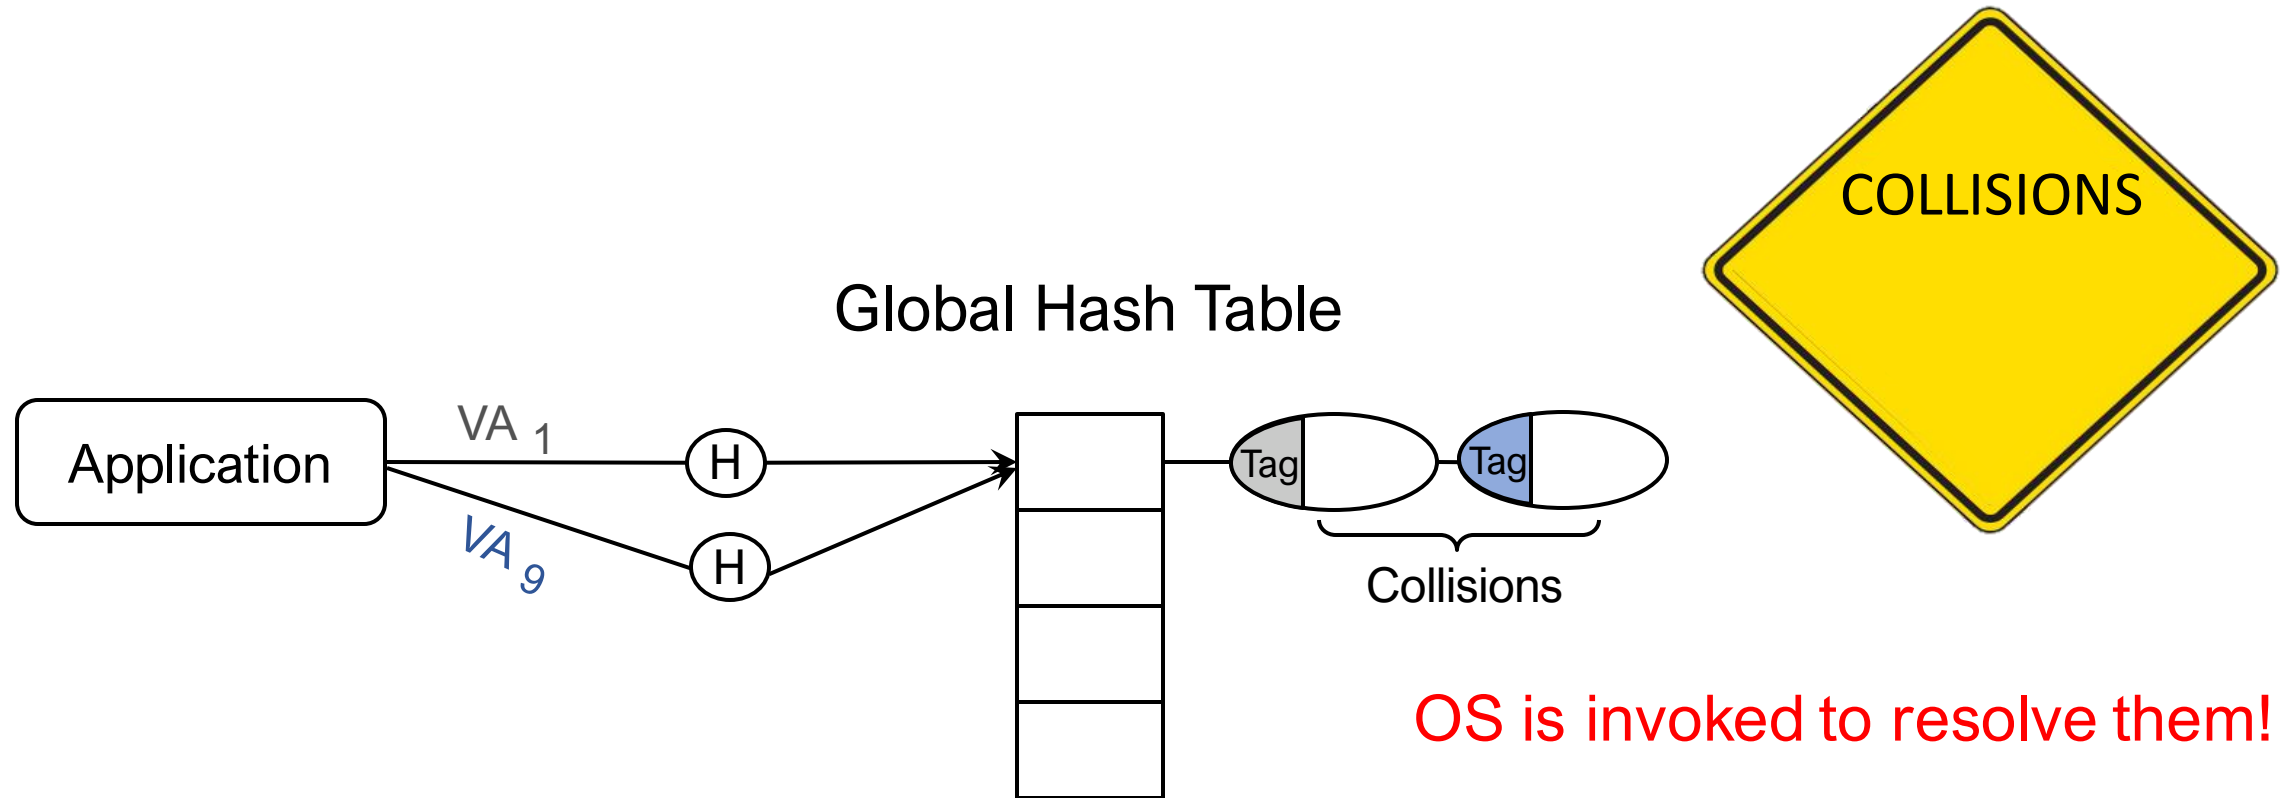
Current Radix Page  
Tables Are Not Scalable



# Alternative: Global Hashed Page Table

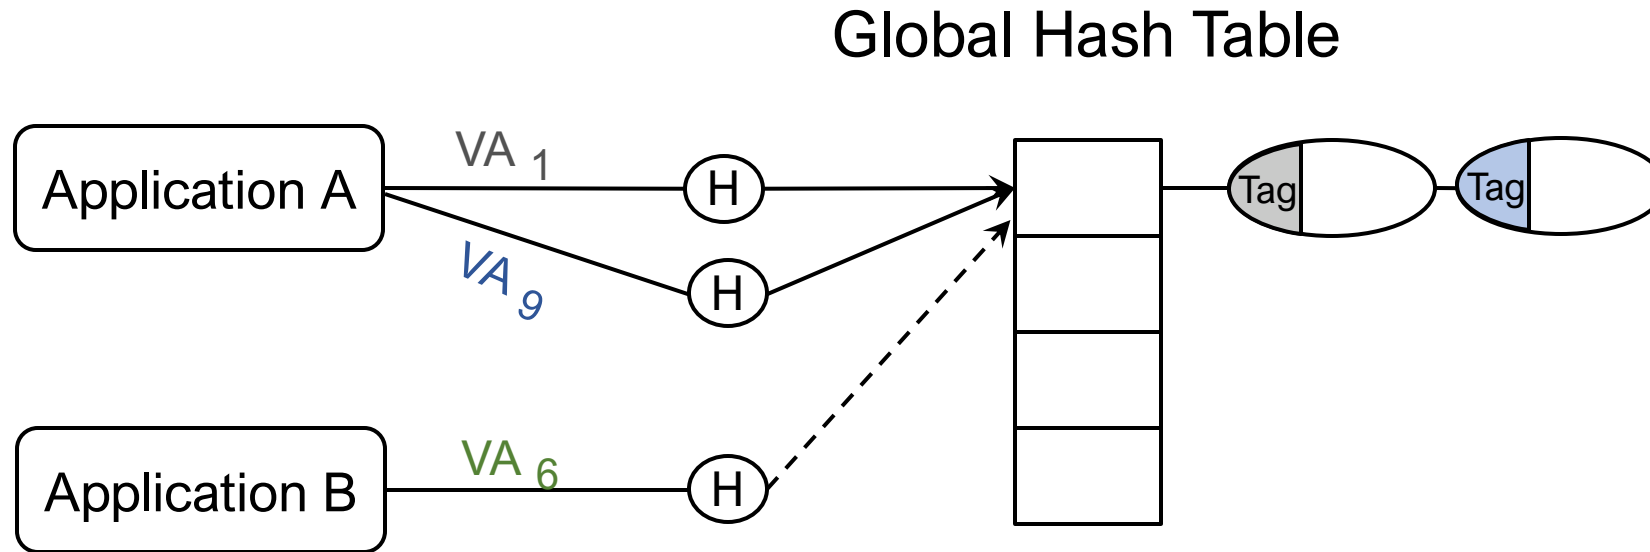


# Alternative: Global Hashed Page Table



# Alternative: Global Hashed Page Table

How to share pages?




# Alternative: Global Hashed Page Table

How to share pages?

New level of indirection!!

Global Hash Table




# Alternative: Global Hashed Page Table

How to share pages?  
Multiple page sizes?

New level of indirection!!

Global Hash Table




# Alternative: Global Hashed Page Table

How to share pages?  
Multiple page sizes?

New level of indirection!!


Global Hash Table



# Elastic Cuckoo Page Tables

Rethinking virtual memory translation for parallelism


# Cuckoo Hashing



$d$ -ary Cuckoo Hash Table




# Insertions with Cuckoo Hashing




$d$ -ary Cuckoo Hash Table

# Insertions with Cuckoo Hashing




$d$ -ary Cuckoo Hash Table

# Insertions with Cuckoo Hashing




*d*-ary Cuckoo Hash Table

# Insertions with Cuckoo Hashing




$d$ -ary Cuckoo Hash Table

# Insertions with Cuckoo Hashing




$d$ -ary Cuckoo Hash Table

# Insertions with Cuckoo Hashing




$d$ -ary Cuckoo Hash Table

# Insertions with Cuckoo Hashing




*d*-ary Cuckoo Hash Table

# Private Hashed Page Tables



*d*-ary Cuckoo Hash Table





# Cannot Be Too Big → Waste Memory




App A

App B

Main Memory

Page Tables A

Page Tables B



Private page tables cannot be too big

# Need to Dynamically Resize



App A

Main Memory

Page Tables A

VA 1

PA 4


VA 8

PA 1

Private page tables cannot be too big



Need to dynamically resize



# Cannot Rehash All Entries at Once



App A


Main Memory

Page Tables A

VA 1 PA 4

VA 8 PA 1

New Page Tables A



Private page tables cannot be too big




Need to dynamically resize




GRADUAL REHASHING  
while the program is running

# Gradual Rehashing of Cuckoo Hash Tables

At every insert  $\rightarrow$  Rehash one element




T1      T2      T3  
Old  $d$ -ary Cuckoo Hash Table




T1'      T2'      T3'  
New  $d$ -ary Cuckoo Hash Table

# Gradual Rehashing of Cuckoo Hash Tables

At every insert  $\rightarrow$  Rehash one element




T1 T2 T3  
Old  $d$ -ary Cuckoo Hash Table


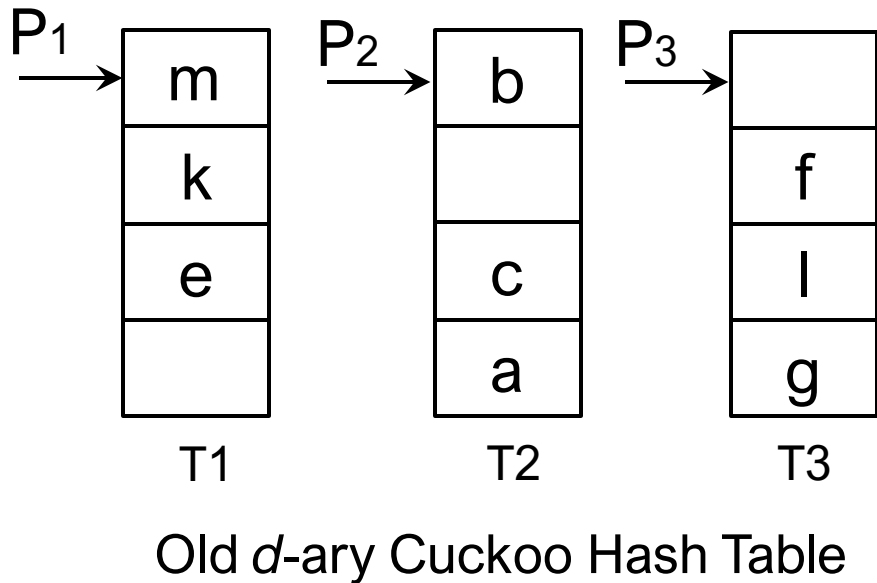


T1' T2' T3'  
New  $d$ -ary Cuckoo Hash Table


# Lookup During Gradual Rehashing



# Solution: Rehashing Pointers




# Solution: Rehashing Pointers



Old  $d$ -ary Cuckoo Hash Table


 Migrated Region



New  $d$ -ary Cuckoo Hash Table




# Solution: Rehashing Pointers




Old  $d$ -ary Cuckoo Hash Table

 Migrated Region




New  $d$ -ary Cuckoo Hash Table

# Lookup with Rehashing Pointers




Old  $d$ -ary Cuckoo Hash Table

 Migrated Region




New  $d$ -ary Cuckoo Hash Table

# Lookup with Rehashing Pointers



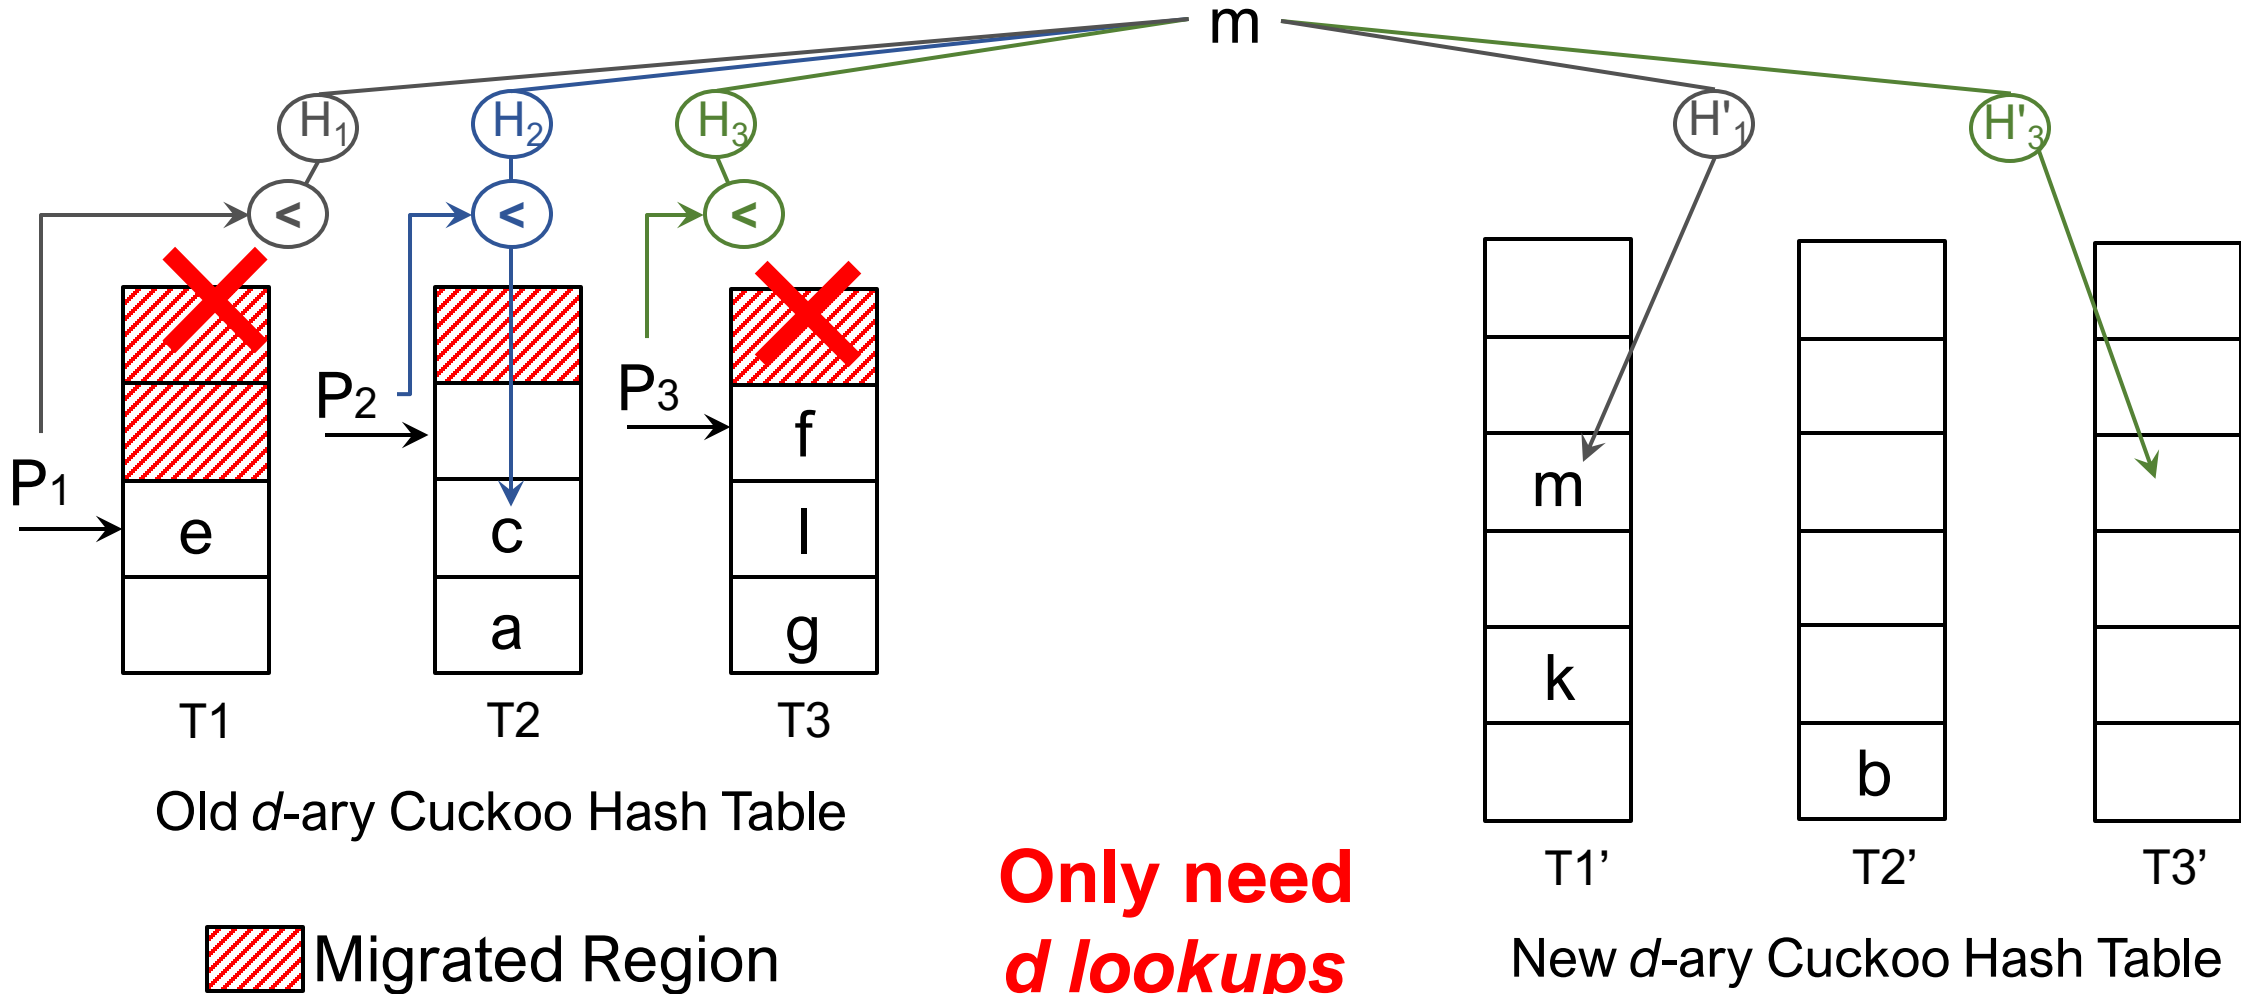
Old  $d$ -ary Cuckoo Hash Table

 Migrated Region




New  $d$ -ary Cuckoo Hash Table

# Lookup with Rehashing Pointers





**Only need  
 $d$  lookups  
during resizing**

# Exploiting Parallelism in Virtual Translation



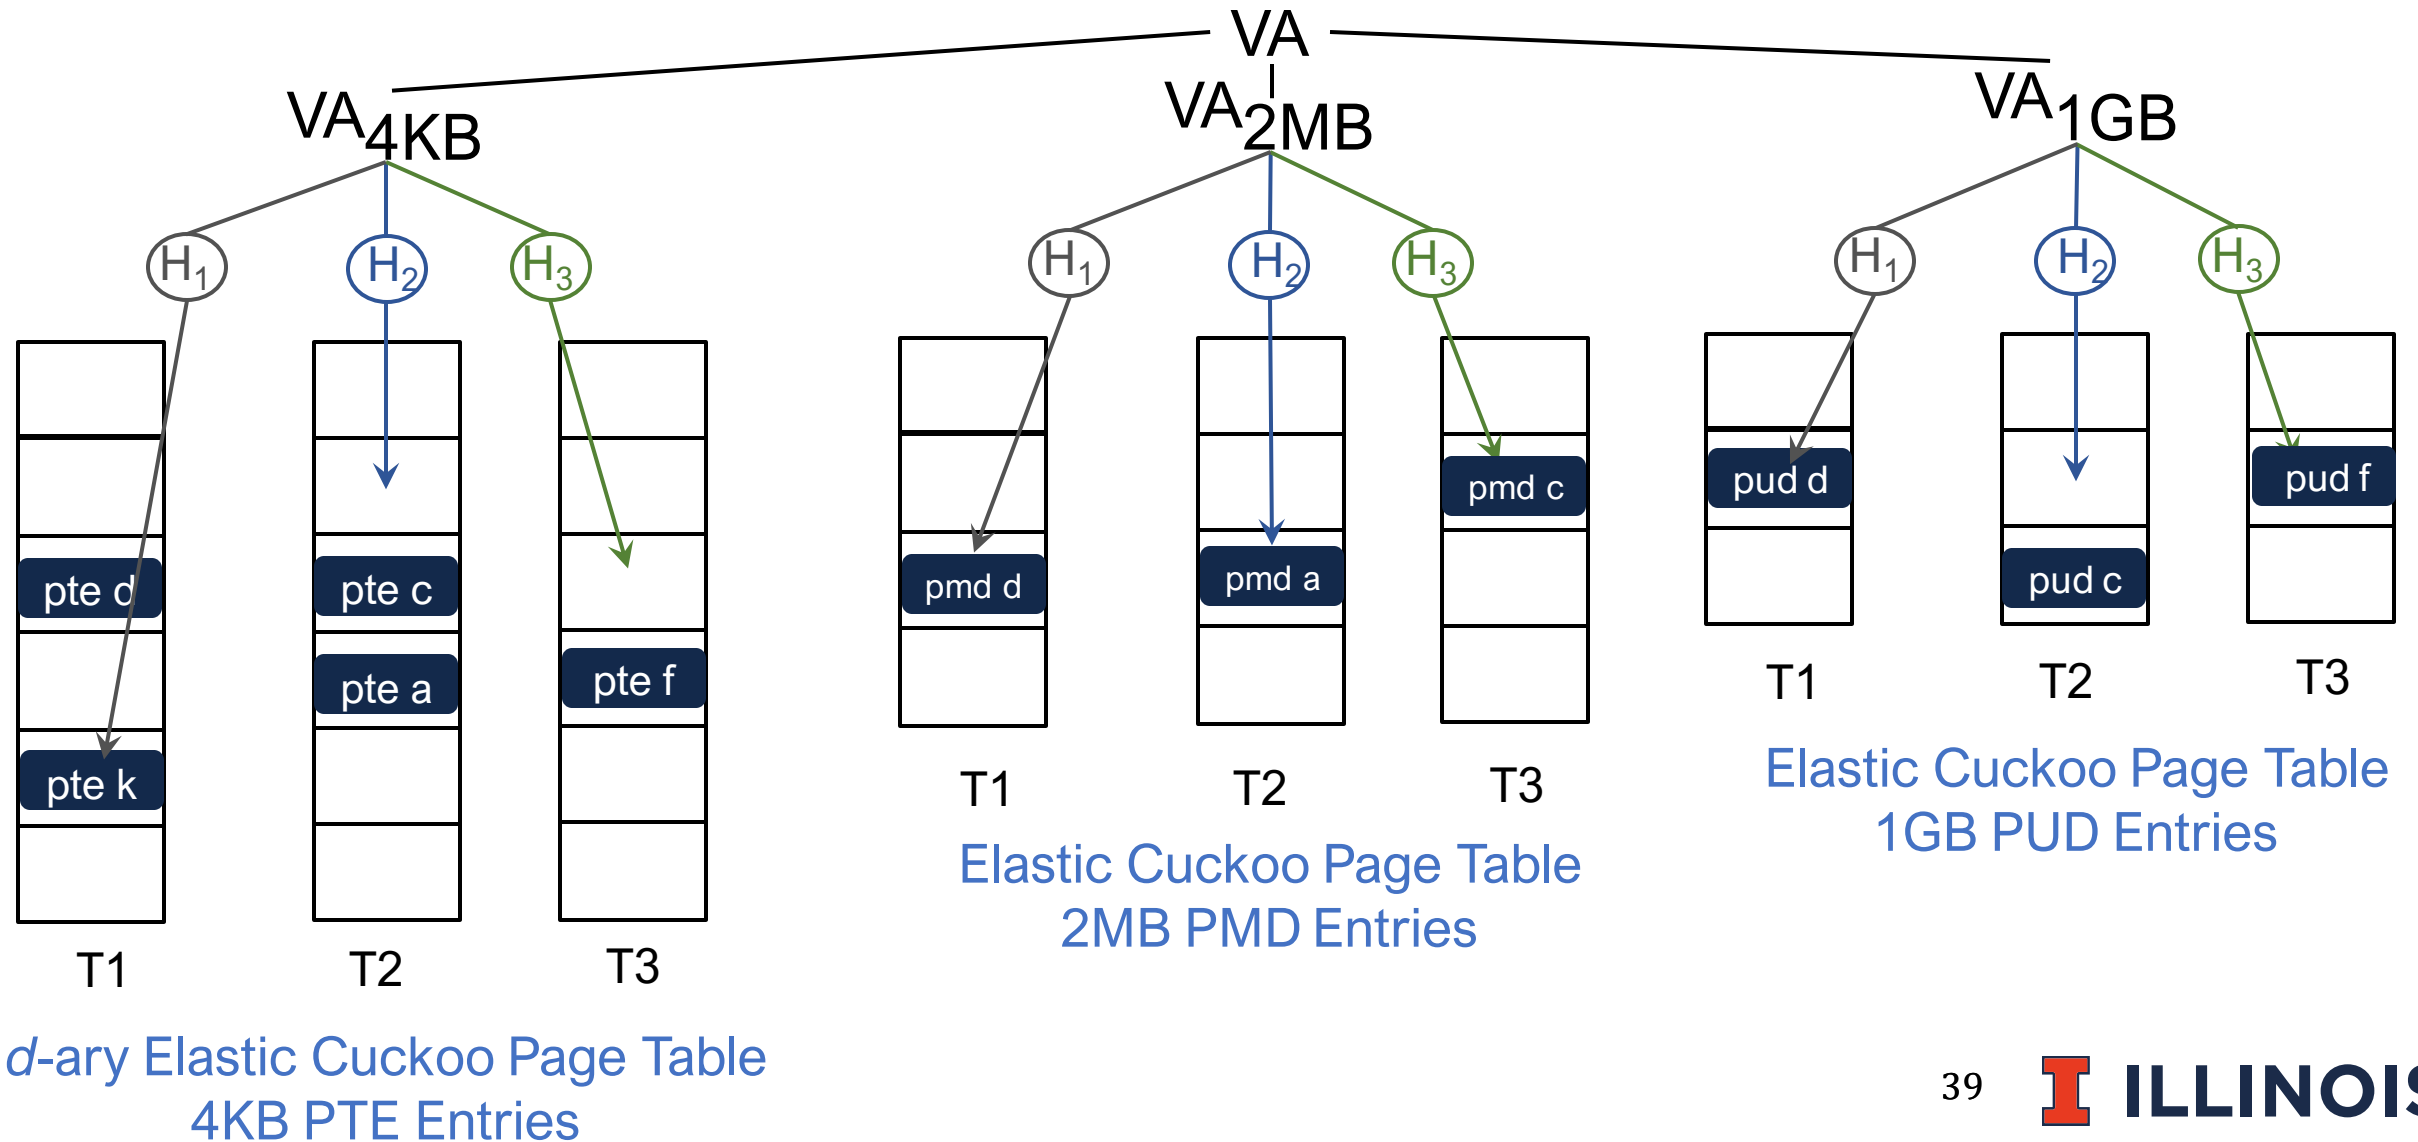
$d$ -ary Elastic Cuckoo Page Table

# Exploiting Parallelism in Virtual Translation

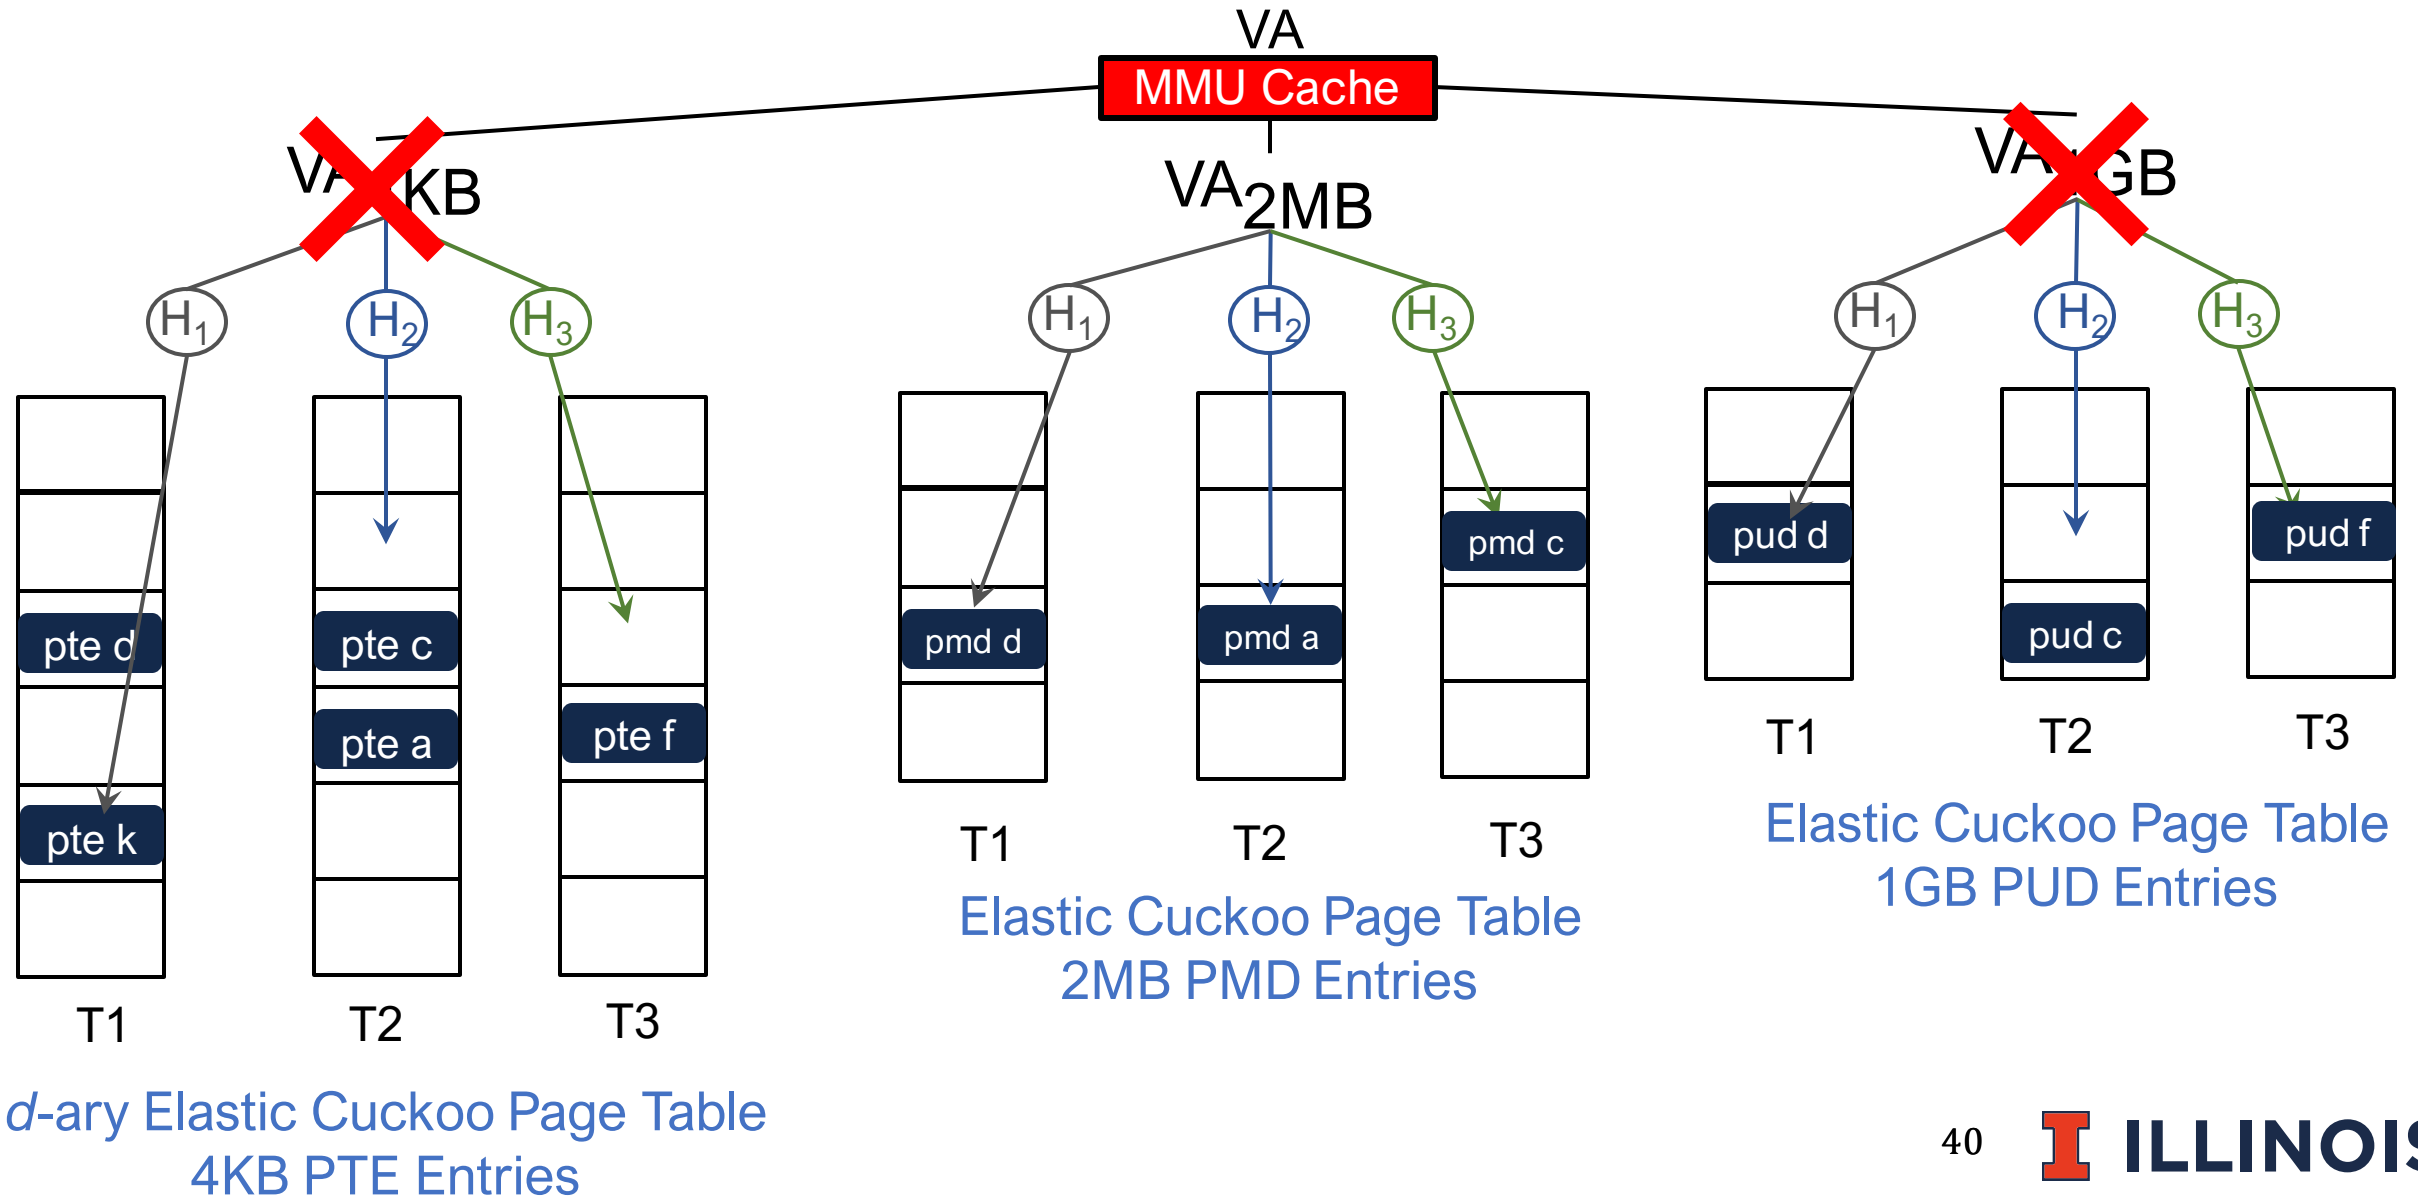


$d$ -ary Elastic Cuckoo Page Table

# Lookup Multiple Page Sizes in Parallel




# New MMU Cache to Prune Parallelism






# New MMU Cache to Prune Parallelism



# Application Speedup



# Takeaway: Elastic Cuckoo Page Tables

## Scalable alternative to existing radix page tables

- Exploits parallelism in virtual translation for the first time
- Reduces the cost of hash table resizing
- Improves application performance over radix by
  - 3-28% with only 4KB pages
  - 3-18% with additionally huge pages
- Expect even higher performance impact on virtualized environments
- Current work: graceful migration from current Radix designs to Elastic Cuckoo

# Rethinking Address Translation in the Age of Containers, Serverless Computing, and Non-Volatile Memory

**Josep Torrellas, Dimitrios Skarlatos, Apostolos Kokolis, Tianyin Xu**

Department of Computer Science  
University of Illinois at Urbana-Champaign  
<http://iacoma.cs.uiuc.edu>

ARM Research Summit  
September 2020