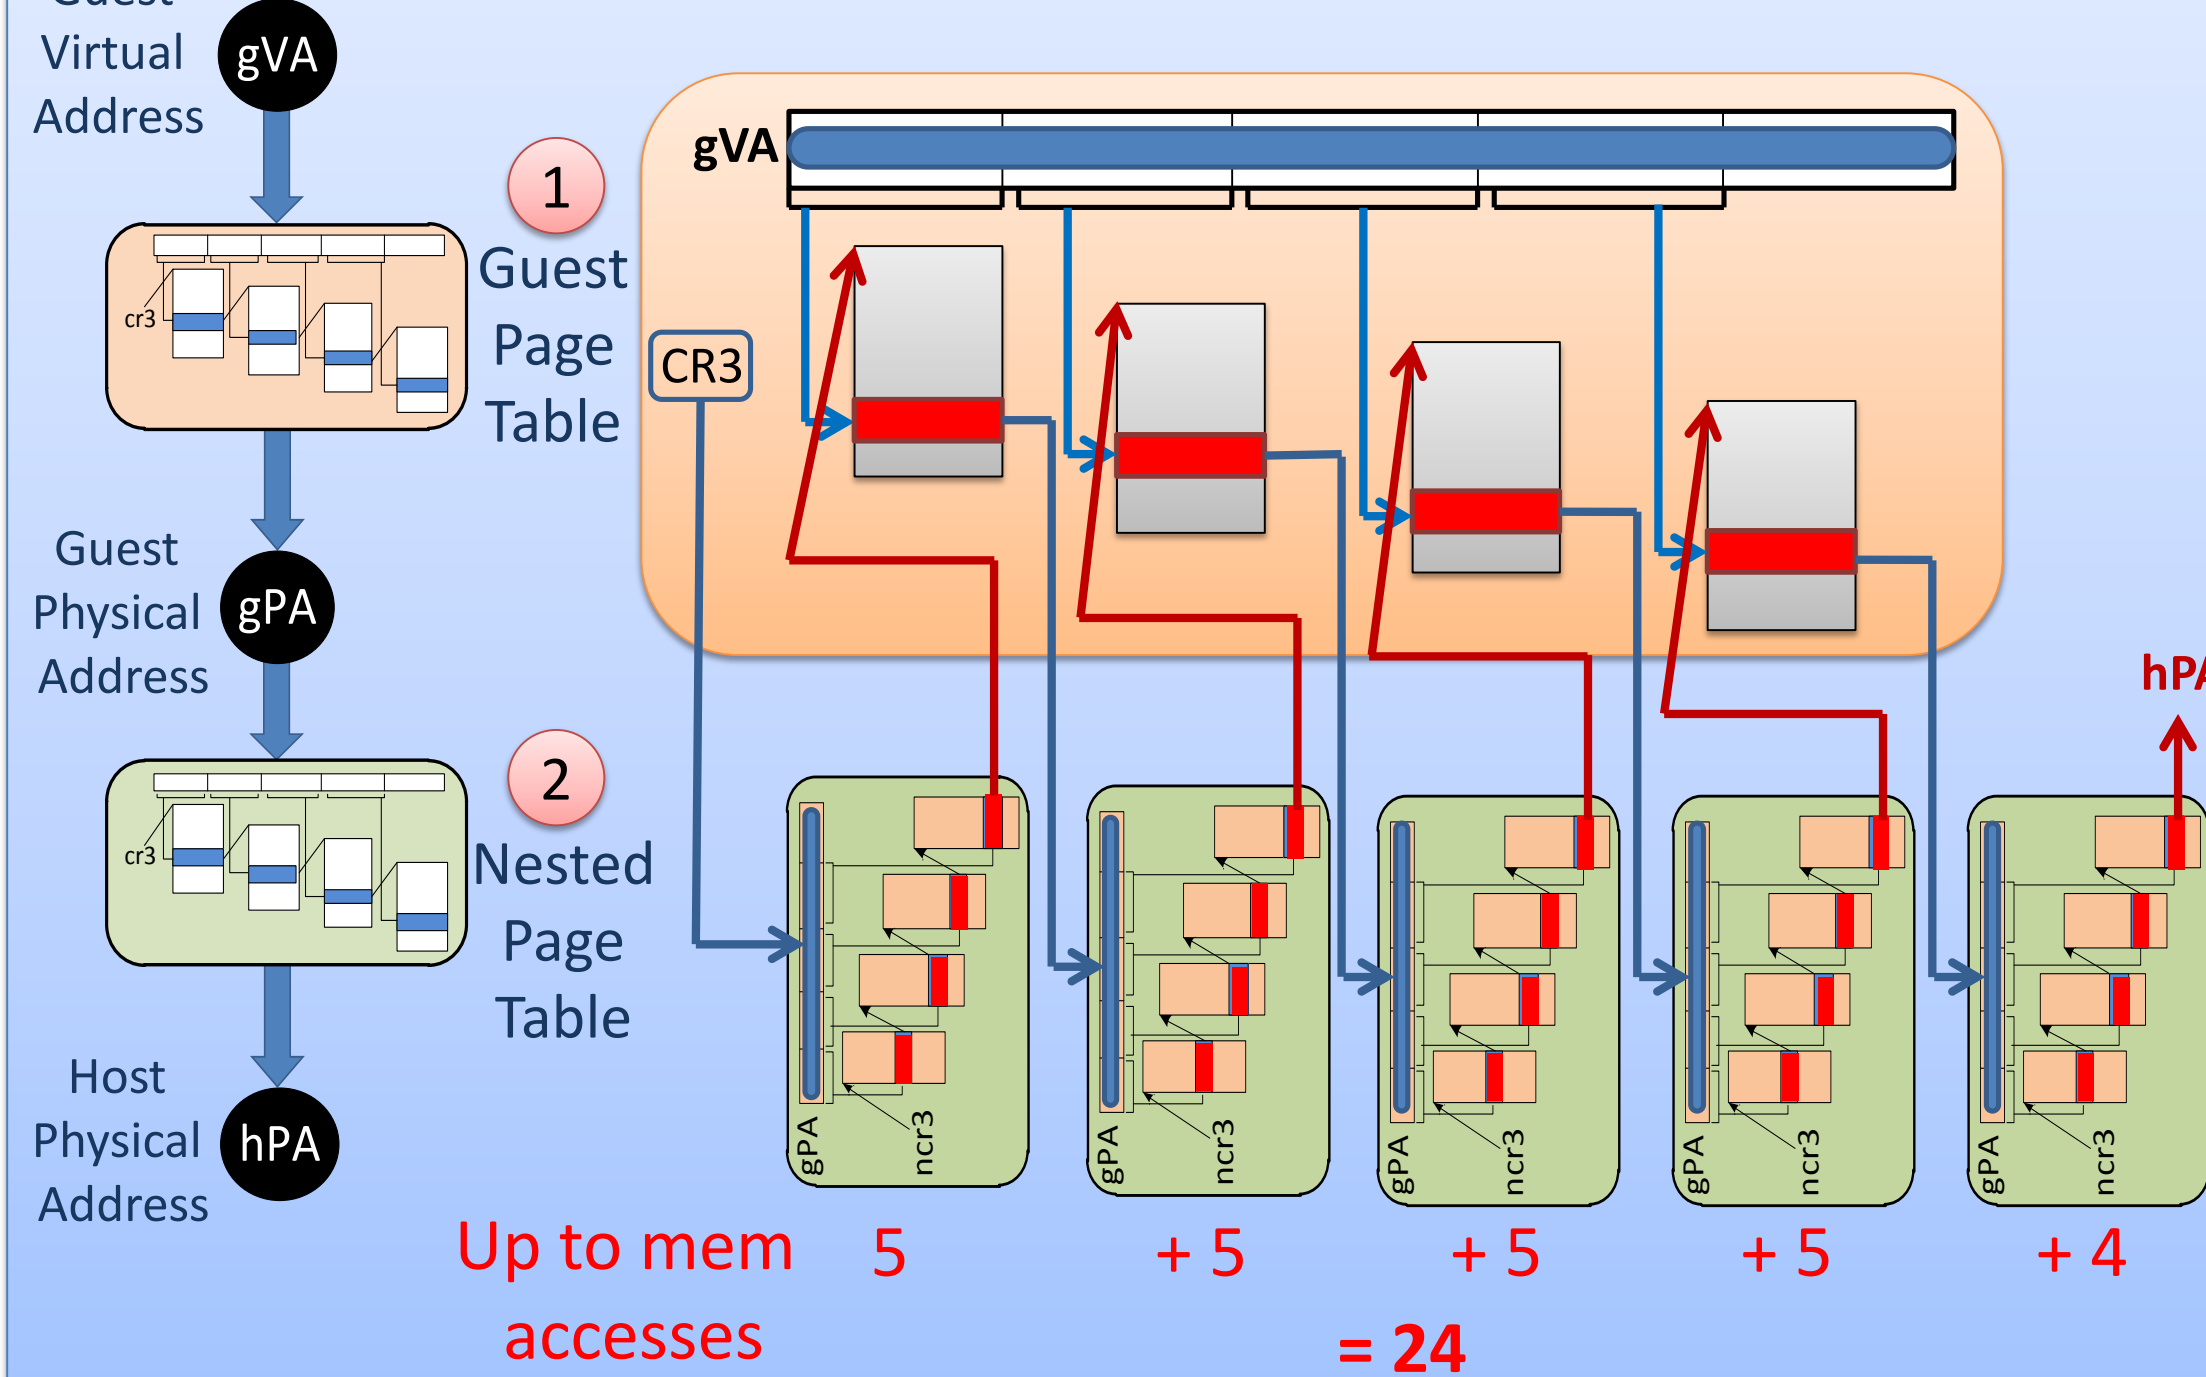


## 1. Problem

Hardware Virtualized MMUs have high overheads  
We show that that the increase in translation lookaside buffer (TLB) miss-handling costs due to the hardware-assisted memory management unit (MMU) is the largest contributor to the performance gap between native and virtual servers.

—Buell, et al. VMware Technical Journal 2013

## 2. Why is a TLB miss costlier?

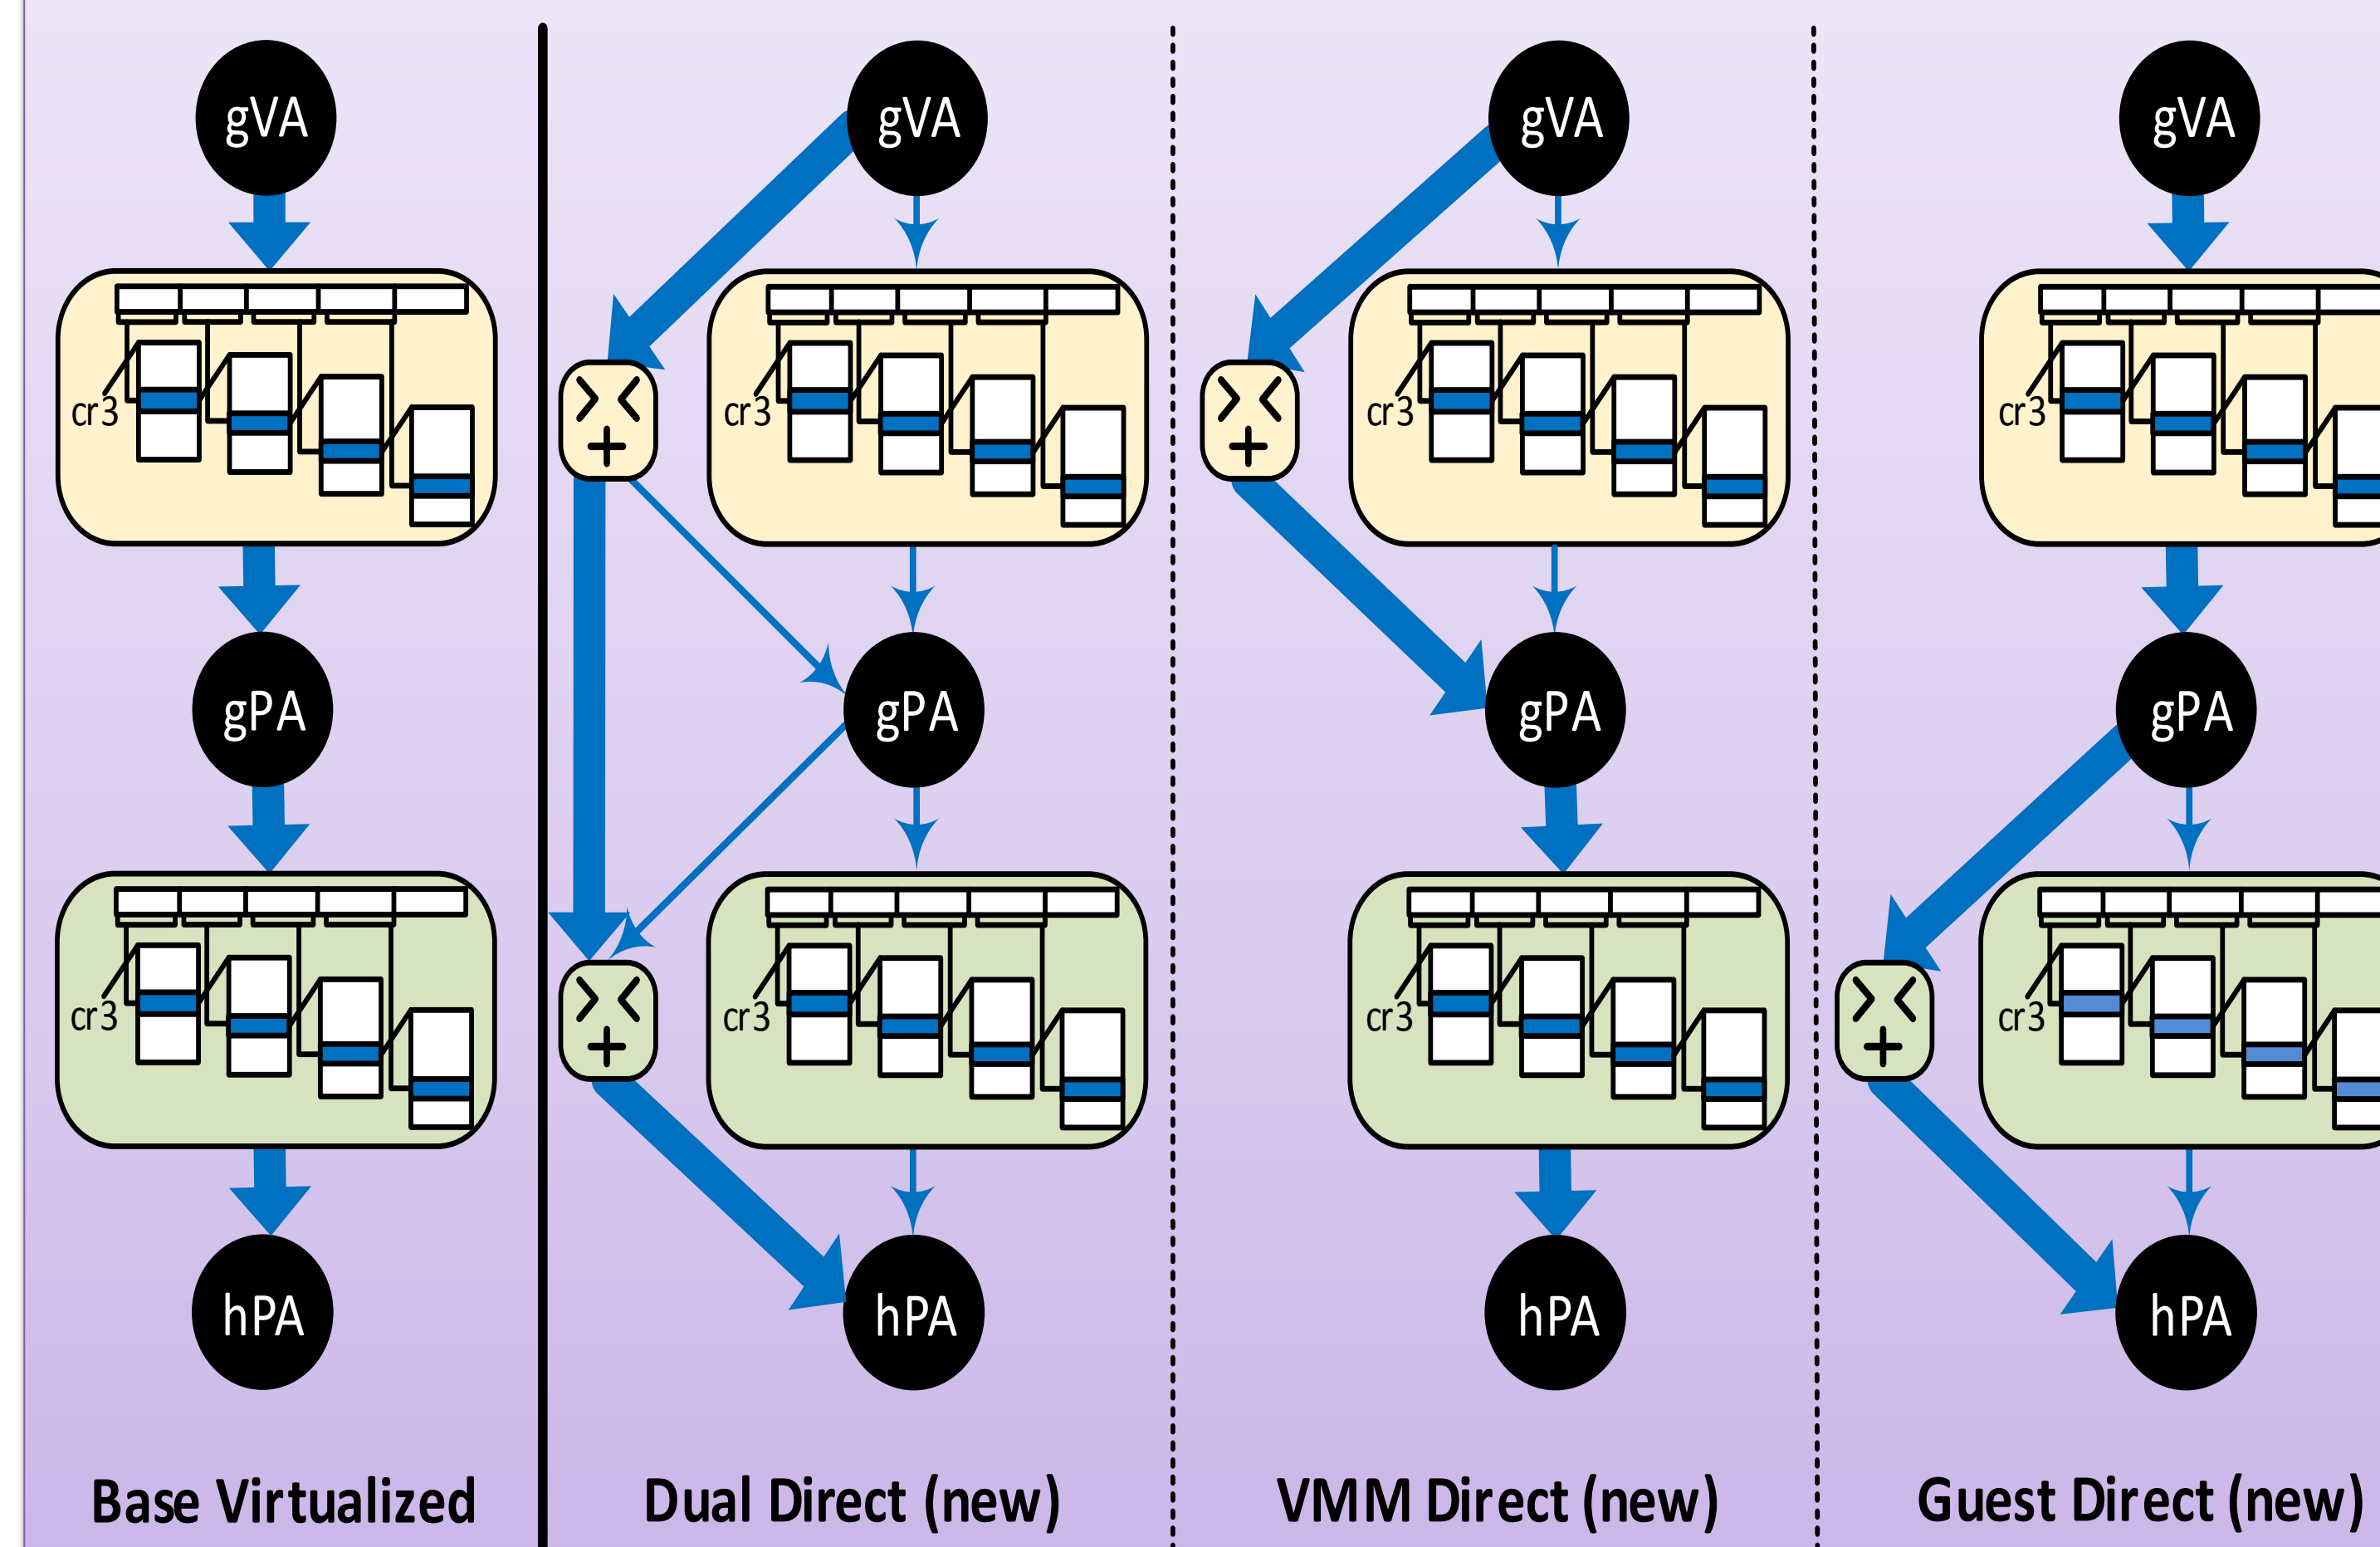


## 3. Solution

Segmentation to bypass paging

- ❖ Extend Direct Segments for virtualization
- ❖ Direct Segment at VMM, guest or both levels
- ❖ Three modes with different tradeoffs

## 4. Modes



## 5. Tradeoffs

Properties	Base Virtualized	Dual Direct	VMM Direct	Guest Direct
# of mem. accesses for most page walks	24	0	4	4
# of base-bound computations for page walks	0	1	5	1
Dimension of page walk	2D	0D	1D	1D
Guest OS modifications	✗	✓	✗	✓
VMM modifications	✗	✓	✓	minimal
Application category	Any	Big-memory	Any	Big-memory
Page sharing	✓	limited	limited	✓
Ballooning	✓	limited	✓	limited
Guest OS swapping	✓	limited	✓	limited
VMM swapping	✓	limited	limited	✓

## 6. Optimizations

- ❖ Guest physical memory fragmentation: Self-ballooning
  - Balloon-out fragmented memory and provide to VMM
  - VMM hot-adds new contiguous guest physical memory
- ❖ Host physical memory fragmentation: Compaction
  - Remap fragmented pages to create contiguous physical memory
- ❖ Permanent “hard” memory faults: Escape filter
  - Escape filter stores few pages with permanent “hard” faults
  - Escape filter checked in parallel with VMM segment register
  - If found in escape filter, get alternate translation through paging

## 7. Overheads + Results

Near- or better-than-native performance

