MMU virtualization in Intel VT-x

Deepayan Bhattacharjee

VT-x : Motivation

- To solve the problem that the x86 instructions architecture cannot be virtualized.
- Simplify VMM software by closing virtualization holes by design of Ring Compression.
- Eliminate need for software virtualization such as paravirtualization.

VPID: Virtual Processor Identifier(1)

- Hypervisors must virtualize physical memory, so that each virtual machine has the illusion of managing its own contiguous region of physical memory.
- First generation VT-x forces TLB flush on each VMX transition.
- Performance loss on all VM exits.
- Performance loss on most VM entries.
  - Guest page tables not modified always.

- Better VMM software control of TLB flushes is beneficial.

VPID: Virtual Processor Identifier(2)

- 16-bit virtual-processor-ID field.
- Cached linear translations tagged with VPID value.
- No flush of TLBs on VM entry or VM exit if VPID active.
- TLB entries of different virtual machines can all co-exist in the TLB.

Abstractions of memory

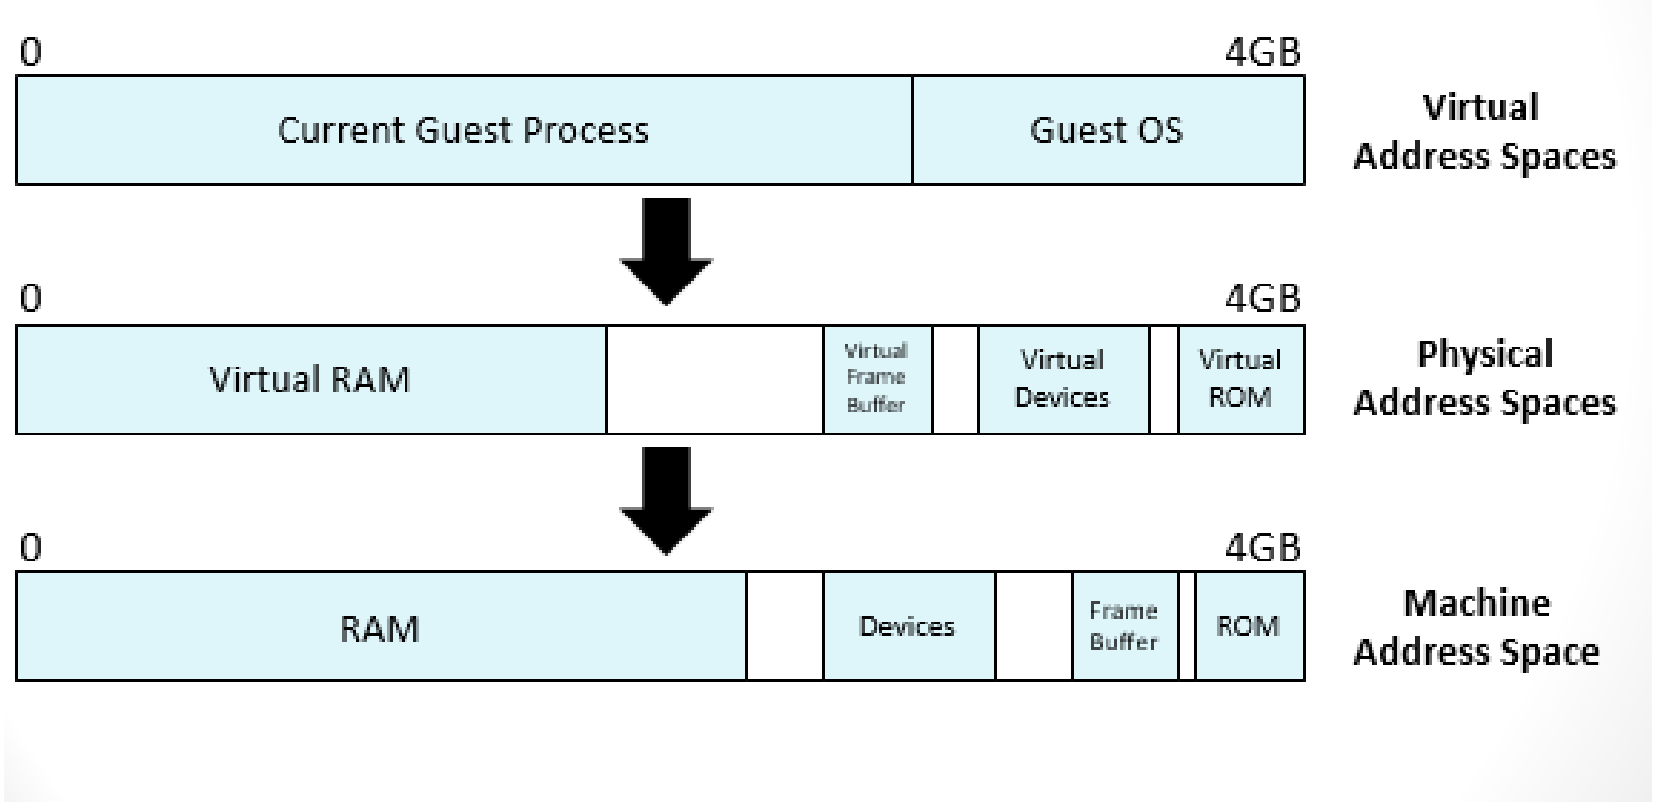

Shadow Page Tables

- To deal with these three abstractions:
  - Shadow page tables are created to map guest-virtual pages directly to machine pages.
  - Guest modifications to V to P tables synced to VMM V to M shadow page tables.
    - Guest OS page tables marked as read-only.
    - Modifications of page tables by guest OS : trapped to VMM.
    - Shadow page tables synced to the guest OS tables

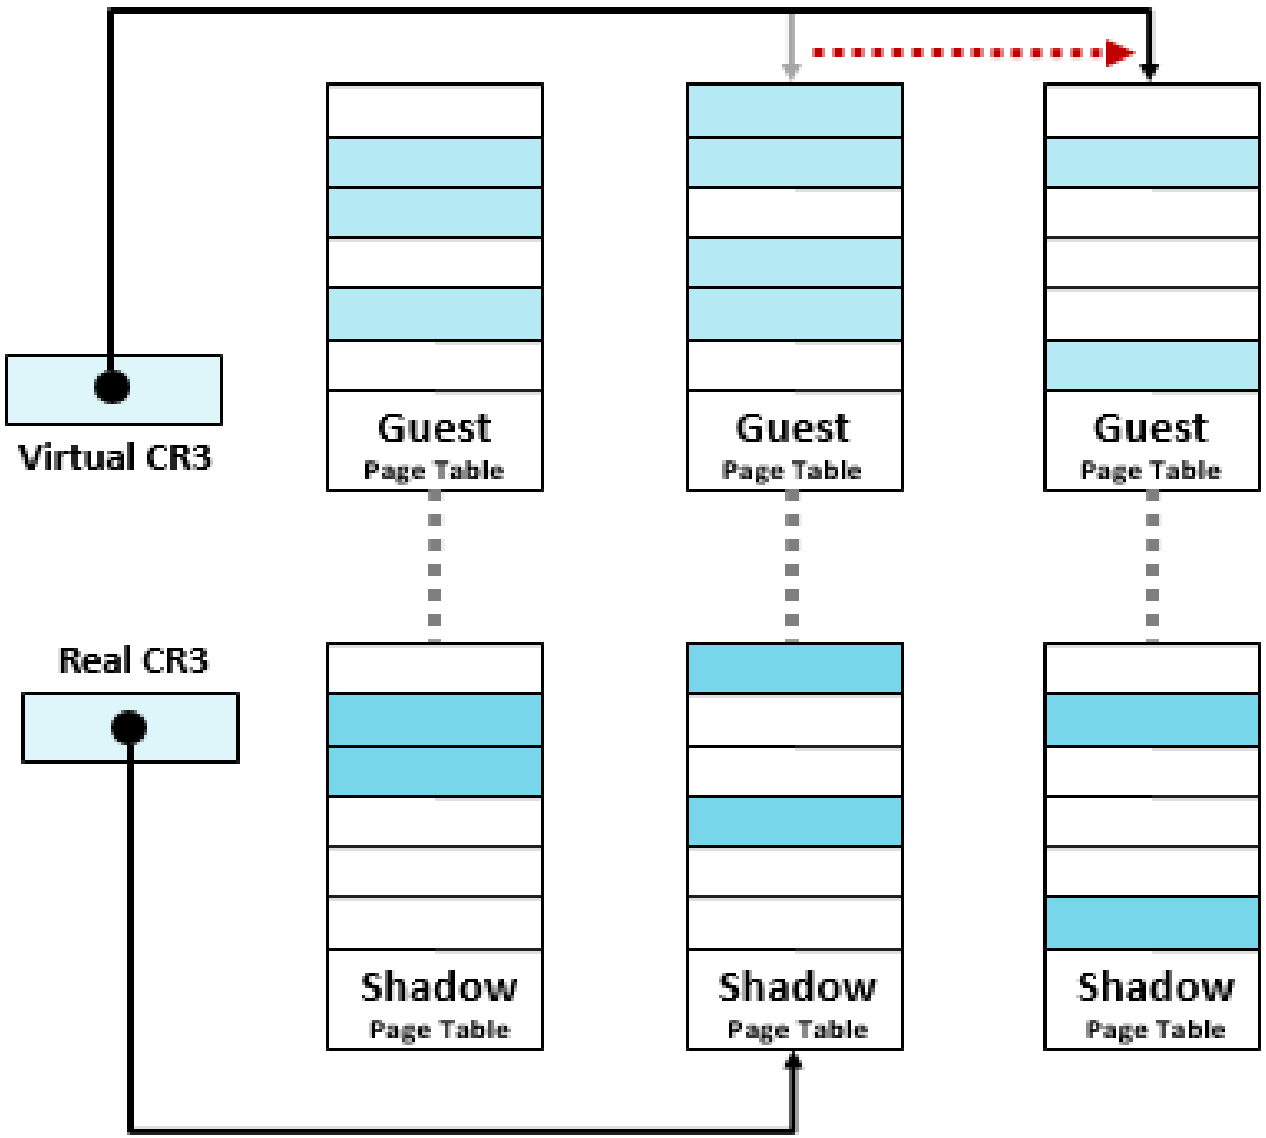

Set CR3 by guest OS (1)

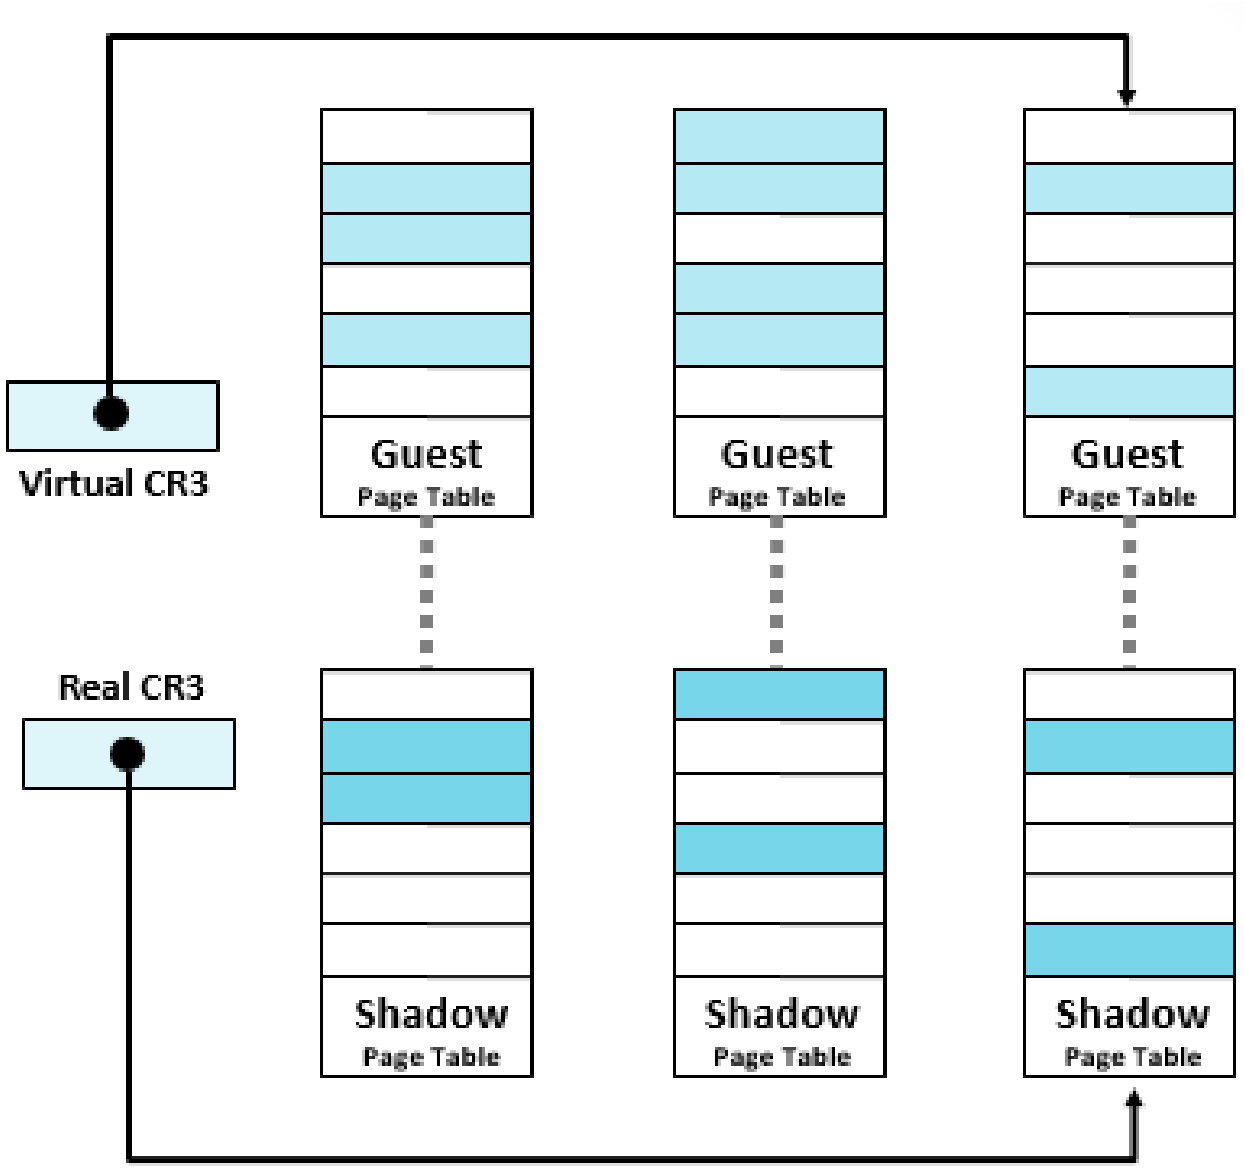

Set CR3 by guest OS (2)

Drawbacks

- Maintaining consistency between guest page tables and shadow page tables leads to an overhead: VMM traps
- Loss of performance due to TLB flush on every “world-switch”.
- Memory overhead due to shadow copying of guest page tables.

Nested / Extended Page Tables

- Extended page-table mechanism (EPT) is used to support the virtualization of physical memory.
- Translates the guest-physical addresses used in VMX non-root operation.
- Guest-physical addresses are translated by traversing a set of EPT paging structures to produce physical addresses that are used to access memory.

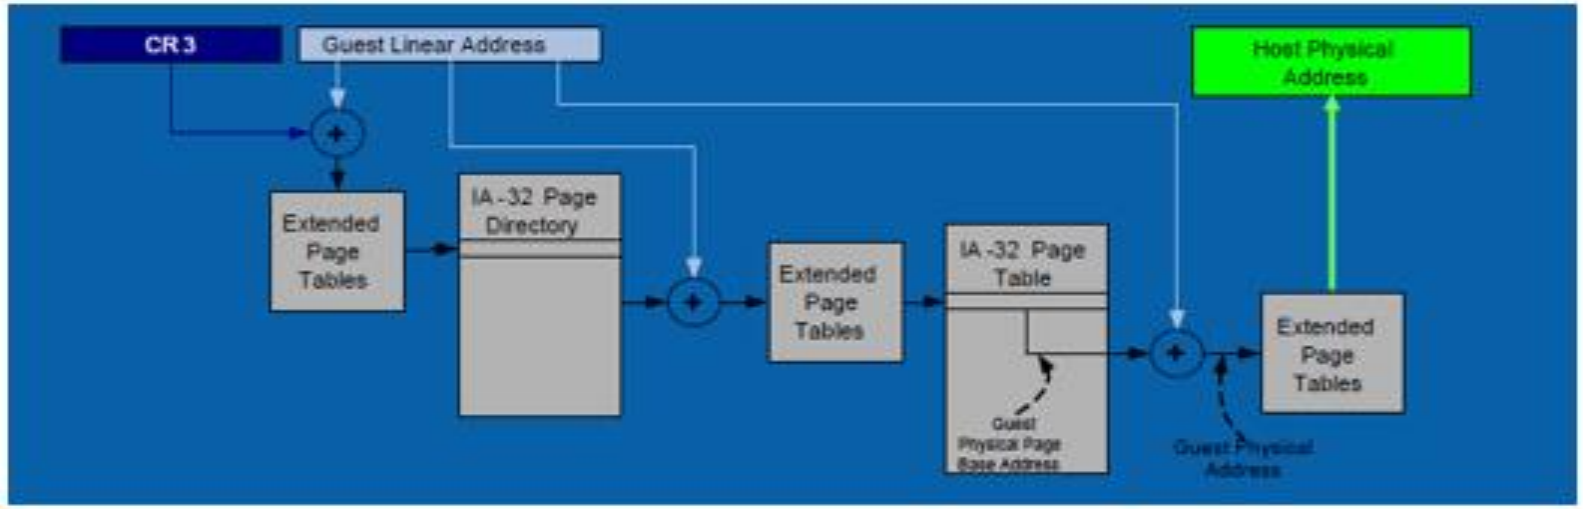

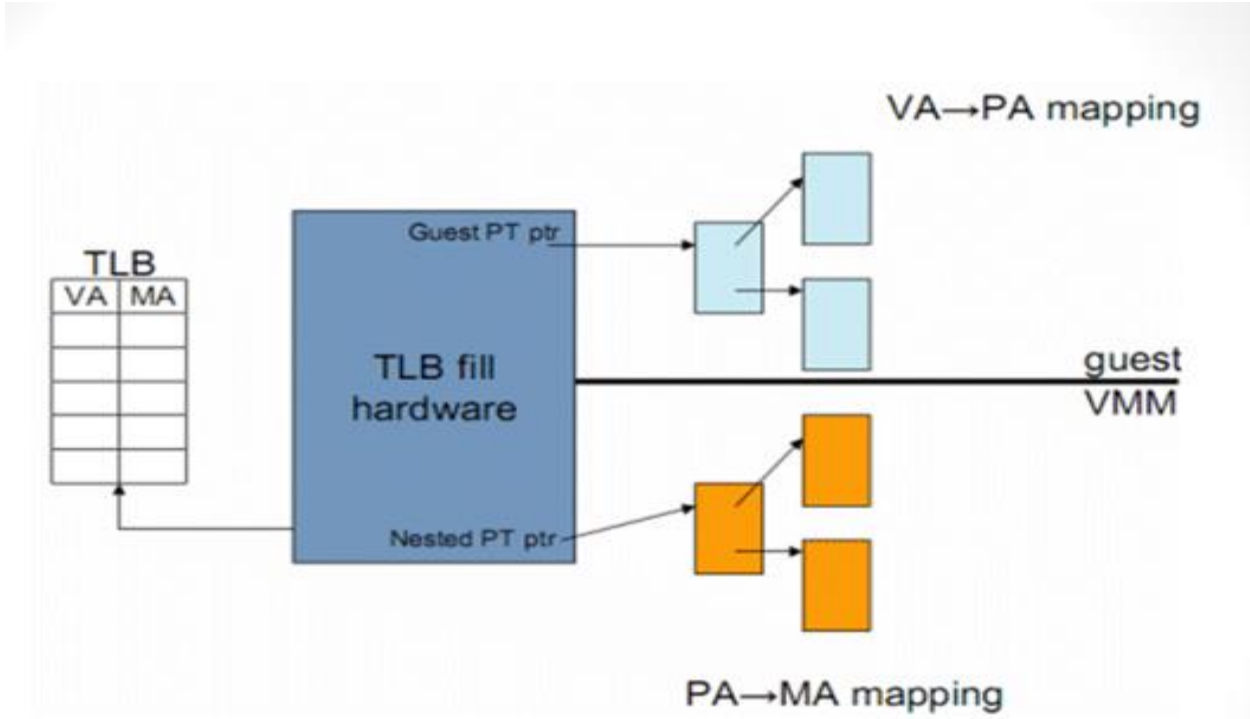

Nested / Extended Page Tables

Pros and Cons of EPT

- Pros:
  - Simplified VMM design.
  - Guest page table modifications are not to be trapped, hence VM exits are minimized.
  - Reduced memory footprint compared to shadow page table algorithms.

- Cons:
  - TLB miss is very costly since guest-physical address to machine address needs an extra EPT walk for each stage of guest-virtual address translation.

Sources:

- Materials are taken from:

  Hardware and Software Support for Virtualization, by Edouard Bugnion, Jason Nieh, Dan Tsafrir.

- Materials and diagrams are taken from:
  - Hardware-assisted Virtualization presentation by Pratik Shah and Rohan Patil, Carnegie Mellon University.

- Intel Manual: www.intel.com/content/dam/www/public/us/en/documents/white-papers/virtualization-enabling-intel-virtualization-technology-features-and-benefits-paper.pdf

Thank You!