Lecture #10 CS695

* Resource management with VMs
  - resource virtualization technique
    - VTx, shadow PT, direct mapping, PV ...
  - resource mgmt. primitives. — VM
  - managment policies

* A2 is active due 10th Feb.
* Quiz 1 paper-pen
  - 12th Feb. Wed.

[memmgmt] Memory Resource Mgmt w/ VMware ESX OSDI 2002

- Keywords: overcommitment (server consolidation) — IaaS goal!
  - maximize
  - performance vs. utilization tradeoff
  - cost
- exploit/employ statistical multiplexing (sharing resources temporally or spatially)
- ~ overprovisioning situation
  - Σ ppa ranges > mpa range } memory pressure
  - Over provisioning ratio — #VMs/PM — Σ resources promised / Σ real/mc resources

~ merging / sharing memory

~ swapping / demand paging ~ eviction / page replacement

~ ballooning

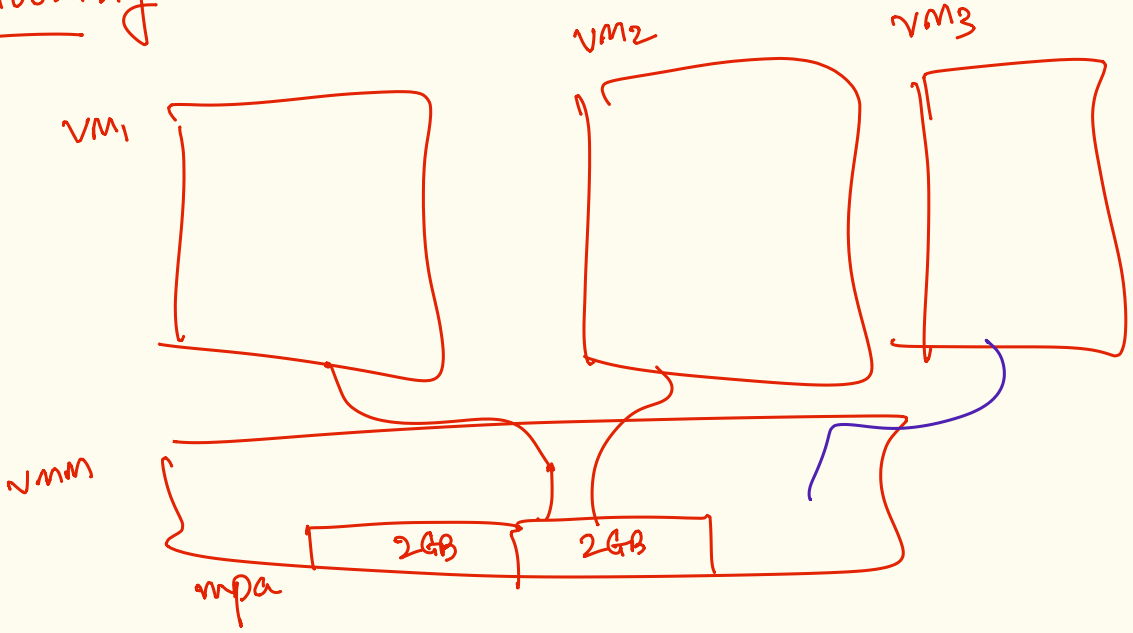

employ page-replacement / eviction by the VMM.

① p2m mappings are VMM controlled.

removing p2m updates page tables of VMs & get free machine pages.

swap to disk.

- possible but costly & needs to breach

Semantic gap — VMM needs to know per-VM utility of pages.

② memory ballooning: v2p2m

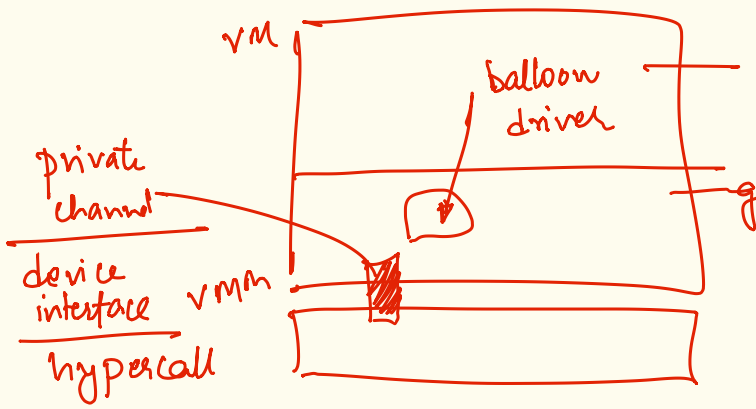

LKM which is listening on the private channel using

inflate deflate

ballooning

- inflate —
  - balloon driver allocates memory
    - kalloc → kva → pagewalk → ppa
    - occupies ppa range
    - pin pages in memory
    - conveys ppa address to VMM
    - VMM removes p2m mappings & reassigns to other demands.

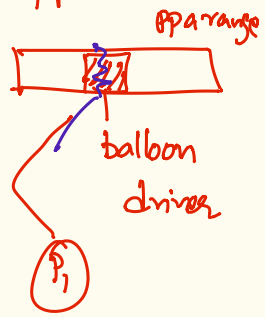

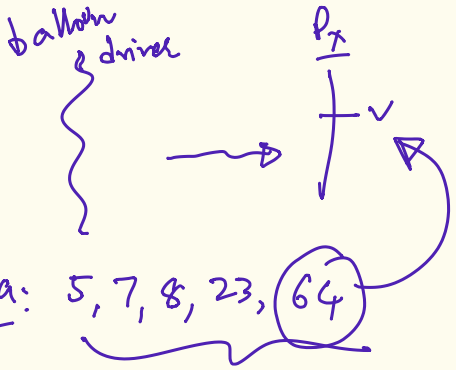

release to VMM

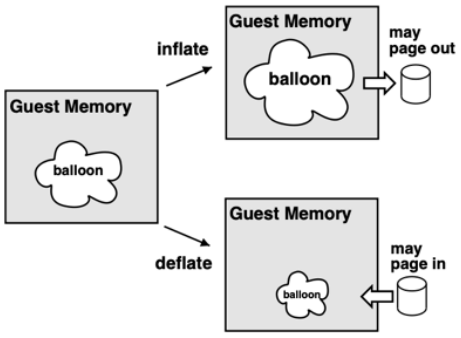

ballooning is the primitive to change memory allocation to VMs.

mgmt: OS
- which VM?
- how much to balloon?
- when to balloon?

V_b 2 P_b 2 m_b ballooning

V_u 2 P_b - guest OS

swapping

double paging problem

guest OS wants to swap P_b to disk

⇒ on access ⇒ page fault

vmm fixes v2p2m for P_b

guest reads/access P_b

& swaps page / removes mapping to disk

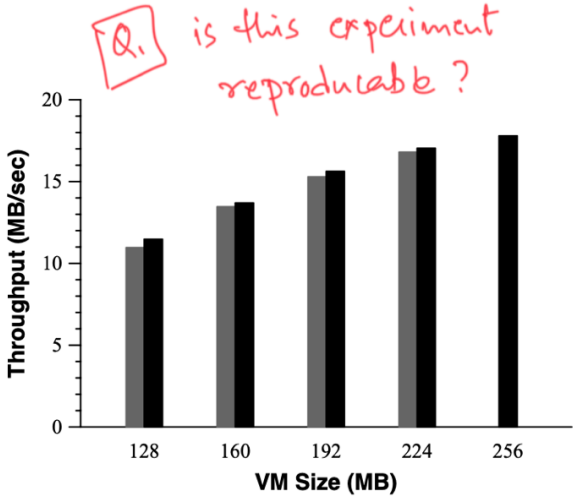

Figure 2: **Balloon Performance.** Throughput of single Linux VM running `dbench` with 40 clients. The black bars plot the performance when the VM is configured with main memory sizes ranging from 128 MB to 256 MB. The gray bars plot the performance of the same VM configured with 256 MB, ballooned down to the specified size.

or removed from a VM in order to rapidly adjust its