Two-level Page Table

Slides produced by:
Jayneel Gandhi, Arkaprava Basu,
Michael M. Swift, Mark D. Hill

Overheads of Virtualizing Memory

We show that that the increase in translation lookaside buffer (TLB) miss-handling costs due to hard-ware-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

Unvirtualized x86 translation

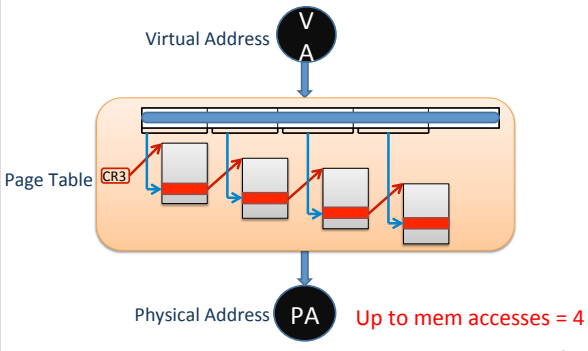

3

Two-levels of translation

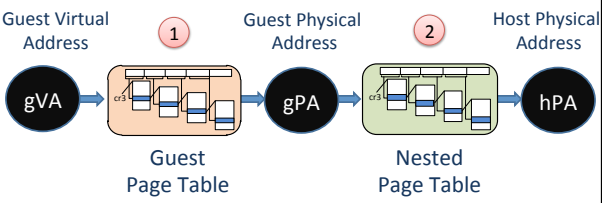

Base Virtualized

4

Two-level Translation

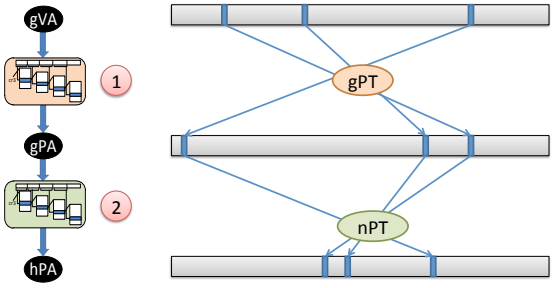

Base Virtualized

5

Support for Virtualizing Memory

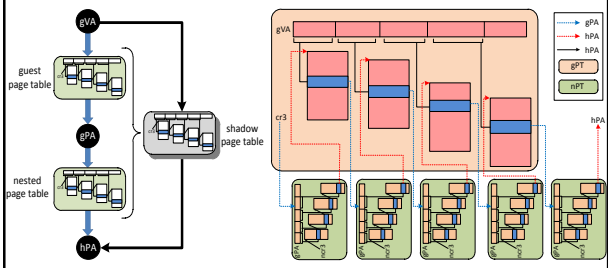

Shadow Page Table

Software only solution

2D Hardware Page Table

Hardware Support Required

6

Translation Hardware

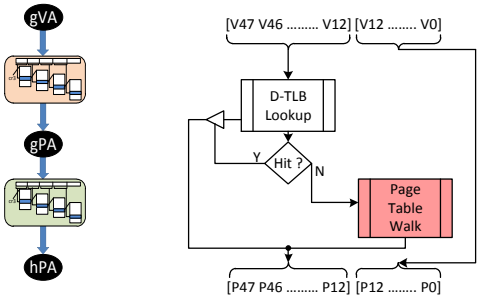

7

Support for Virtualizing Memory

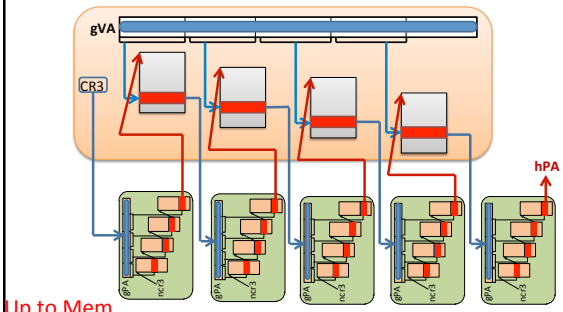

Up to Mem accesses: 5 + 5 + 5 + 5 + 4 = **24**

8

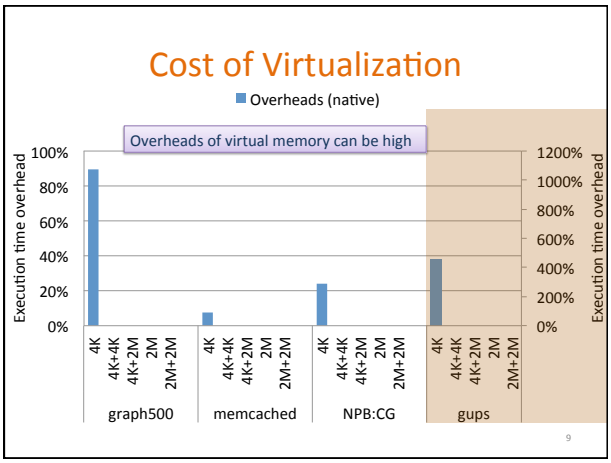
Cost of Virtualization
Overheads (native)
Overheads of virtual memory can be high
Execution time overhead
100%
80%
60%
40%
20%
0%
1200%
1000%
800%
600%
400%
200%
0%
4K
4K+4K
4K+2M
2M
2M+2M
graph500
memcached
NPB:CG
gups
9

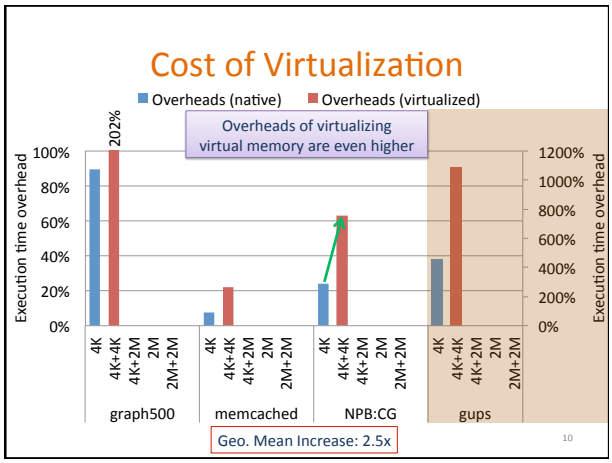
Cost of Virtualization
Overheads (native)
Overheads (virtualized)
202%
Overheads of virtualizing virtual memory are even higher
Execution time overhead
100%
80%
60%
40%
20%
0%
1200%
1000%
800%
600%
400%
200%
0%
4K
4K+4K
4K+2M
2M
2M+2M
graph500
memcached
NPB:CG
gups
Geo. Mean Increase: 2.5x
10

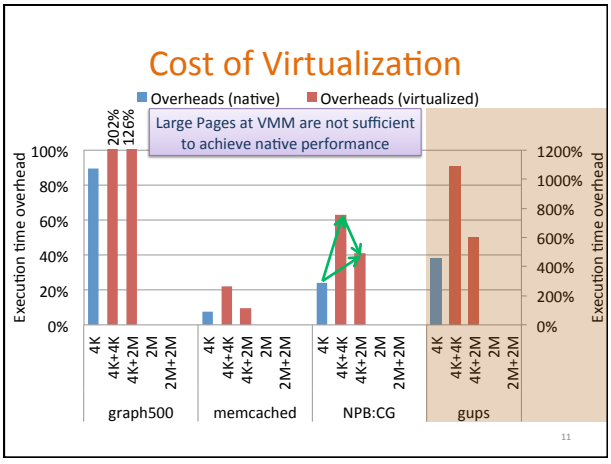
Cost of Virtualization
Overheads (native)
Overheads (virtualized)
202%
126%
Large Pages at VMM are not sufficient to achieve native performance
Execution time overhead
100%
80%
60%
40%
20%
0%
1200%
1000%
800%
600%
400%
200%
0%
4K
4K+4K
4K+2M
2M
2M+2M
graph500
memcached
NPB:CG
gups
11

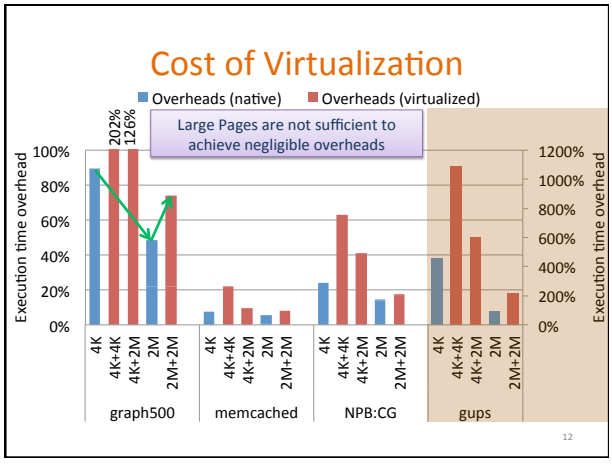
Cost of Virtualization
Overheads (native)
Overheads (virtualized)
202%
126%
Large Pages are not sufficient to achieve negligible overheads
Execution time overhead
100%
80%
60%
40%
20%
0%
1200%
1000%
800%
600%
400%
200%
0%
4K
4K+4K
4K+2M
2M
2M+2M
graph500
memcached
NPB:CG
gups
12