AMD

LARGE PAGES AND LIGHTWEIGHT MEMORY MANAGEMENT IN VIRTUALIZED ENVIRONMENTS: CAN YOU HAVE IT BOTH WAYS?

Binh Pham[§], Jan Vesely[§], Gabriel H. Loh[‡], Abhishek Bhattacharjee[§]

[§]Rutgers University

[‡]AMD Research

Large Pages Advantages

- Large pages are often used to mitigate address translation overhead
  - Increase TLB reach
  - Reduce page walk latency

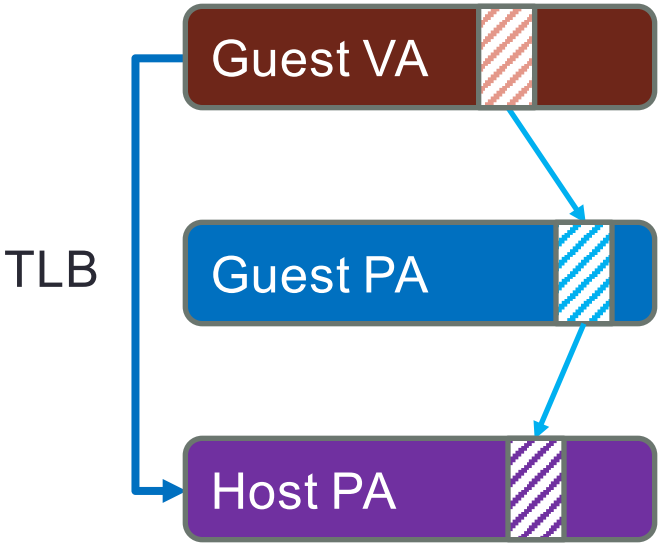

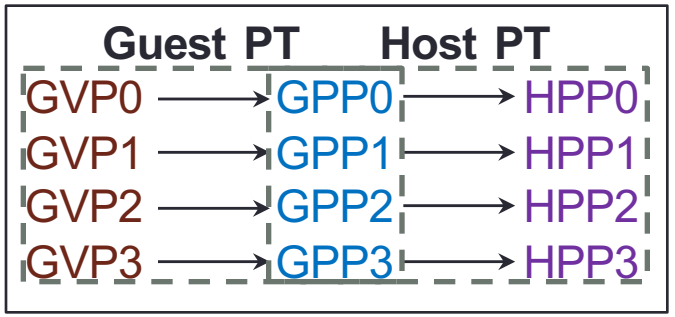

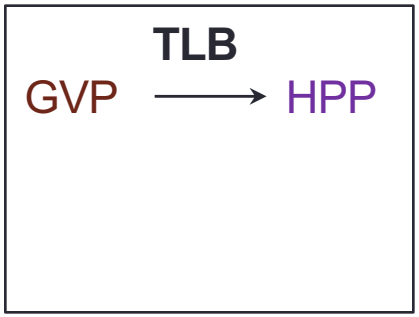

AMD

Large Pages Disadvantages

- Larger working sets [1]
- Coarse protection granularity [1]
- Hurts performance on NUMA systems due to node imbalance and poor locality [2]

[1] Talluri, Kong, Hill, Patterson. **Tradeoffs in Supporting Two Page Sizes**. ISCA 1992 .

[2] Gaud, Lepers, Decouchant, Funston, Fedorova, Quema. **Large Pages May Be Harmful on NUMA Systems**. USENIX ATC 2014.

Our work: Large Page Benefits vs Light Weight Memory Management

- Memory management techniques:
  - Page sharing
  - Memory sampling
  - Memory compression
  - Virtual machine migration
- Hypervisor splinters large pages into small pages

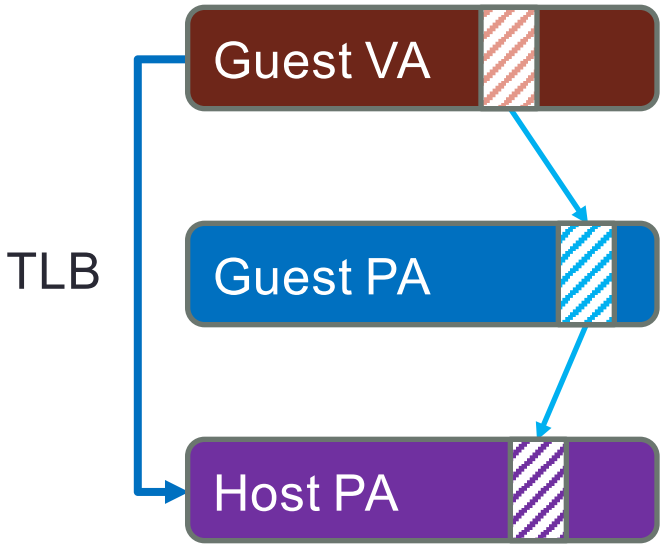

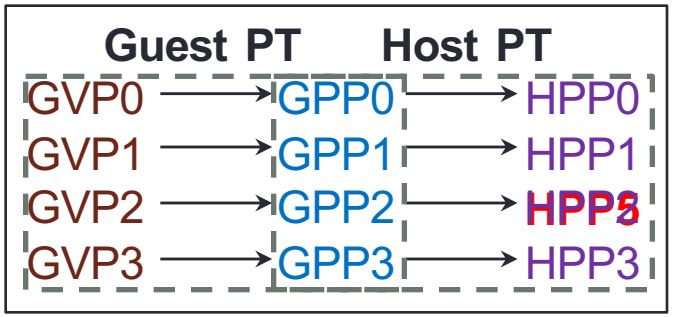

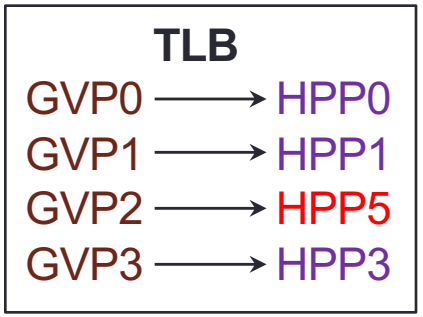

Prevalence of Page Splintering and Performance Impact

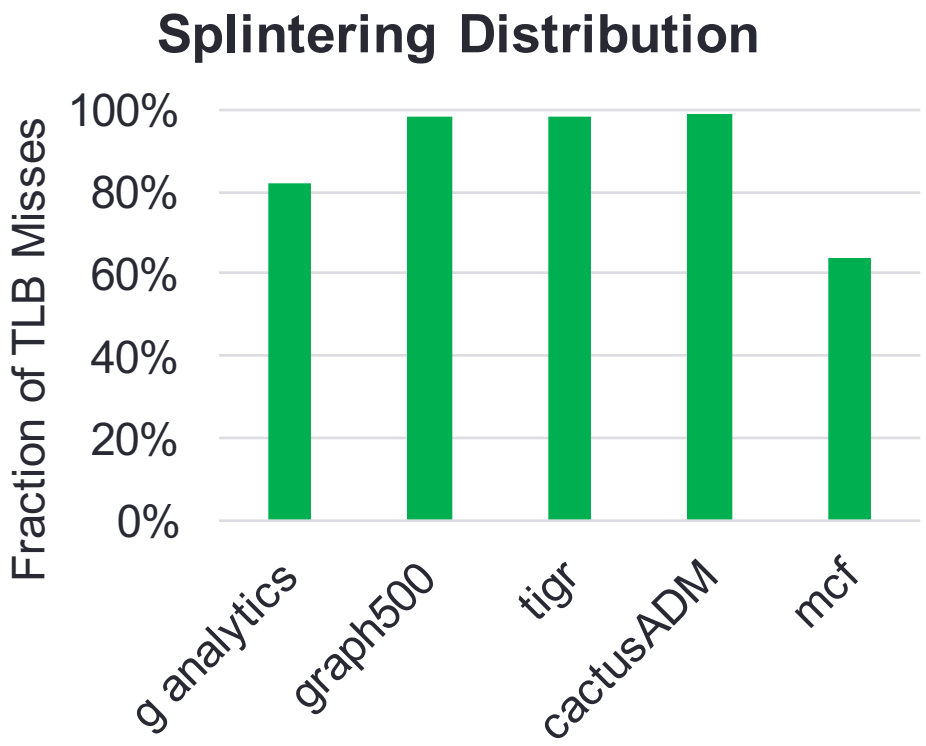

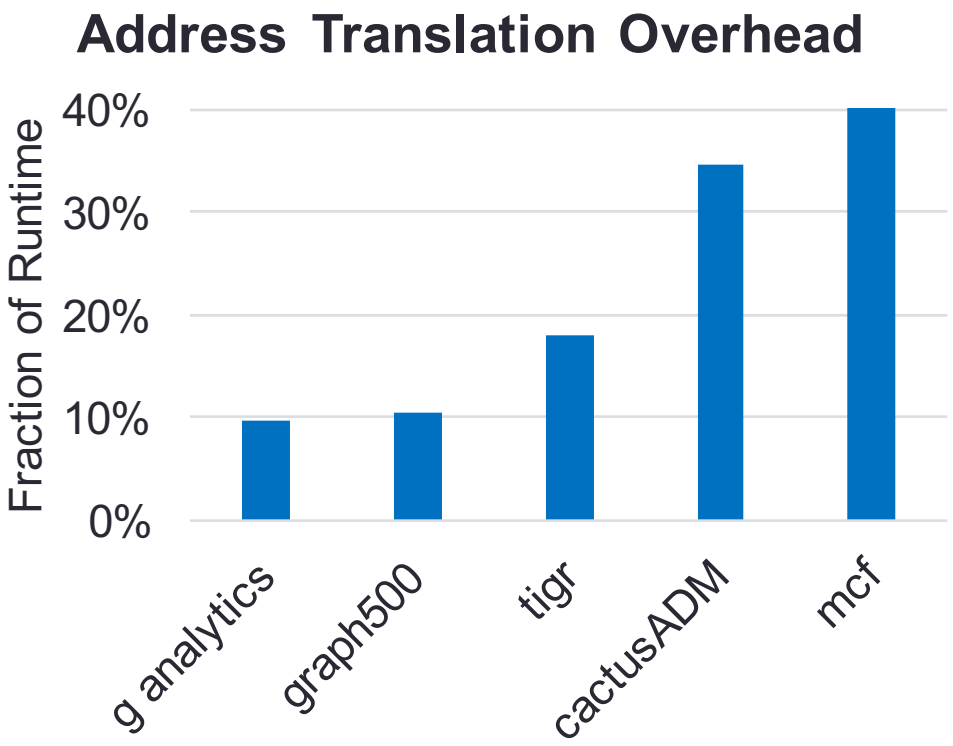

Patterns in Splintered Pages

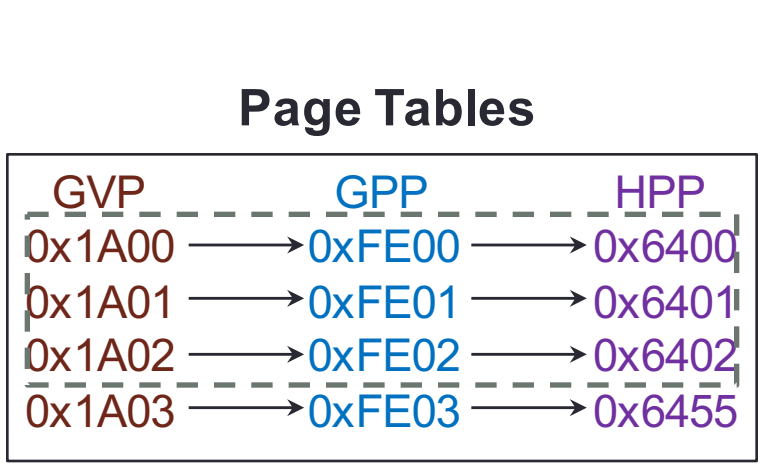

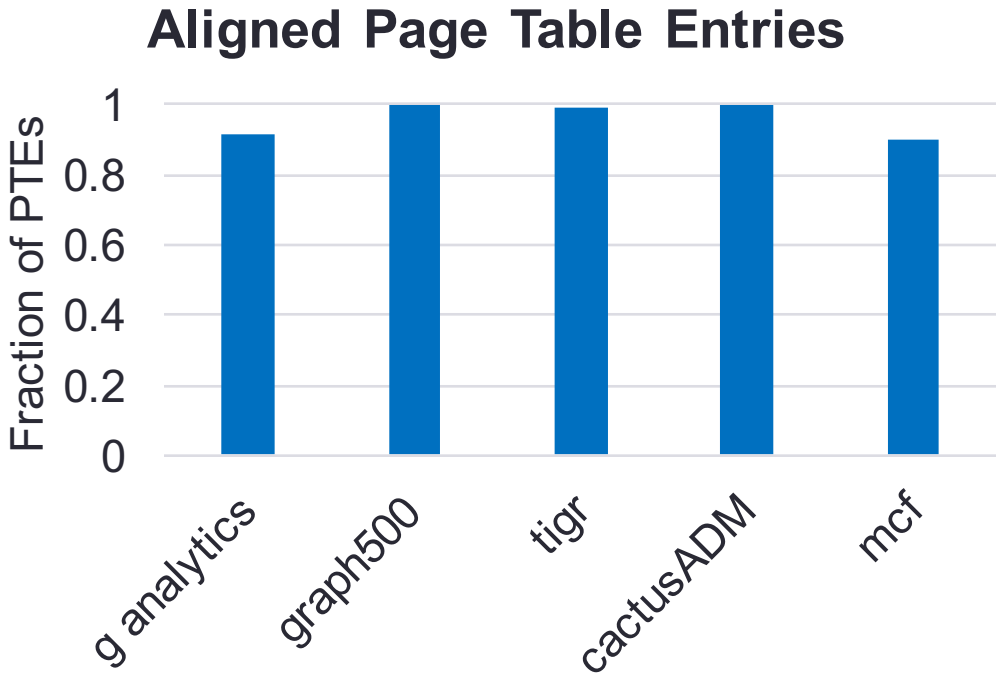

GLUE Design: Address Interpolation

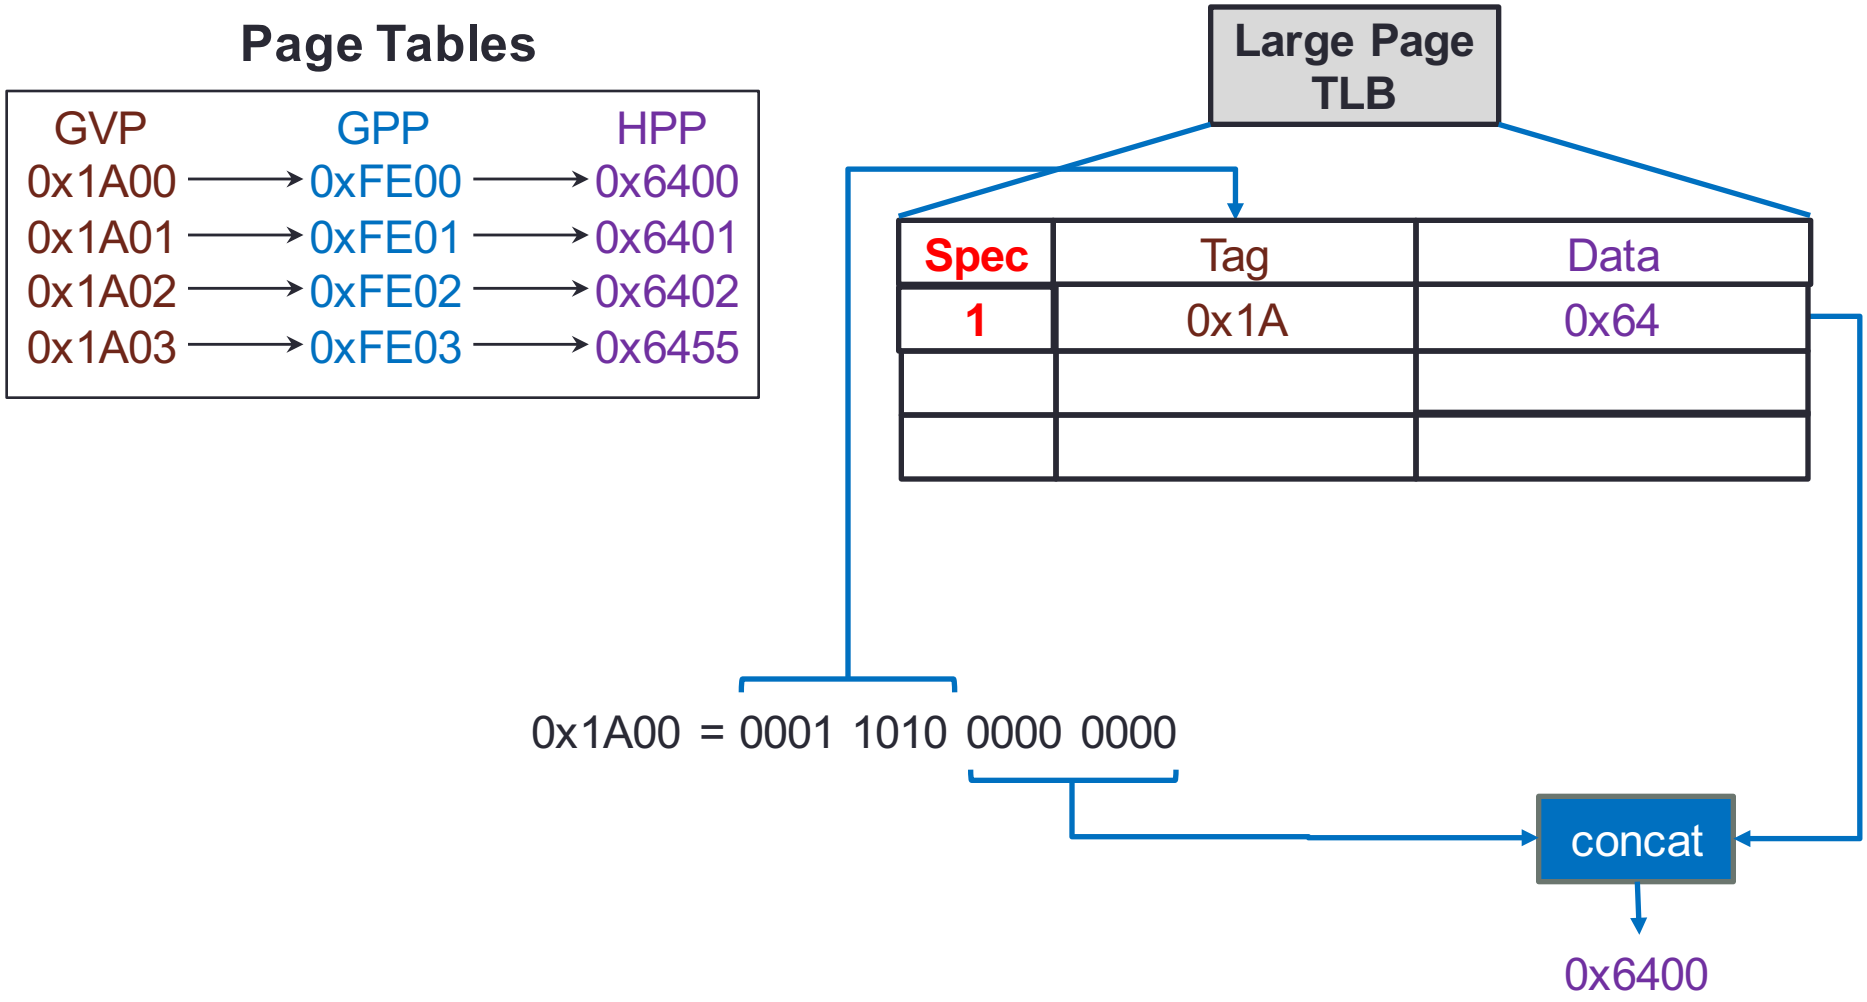

GLUE Speculative Entry Insertion

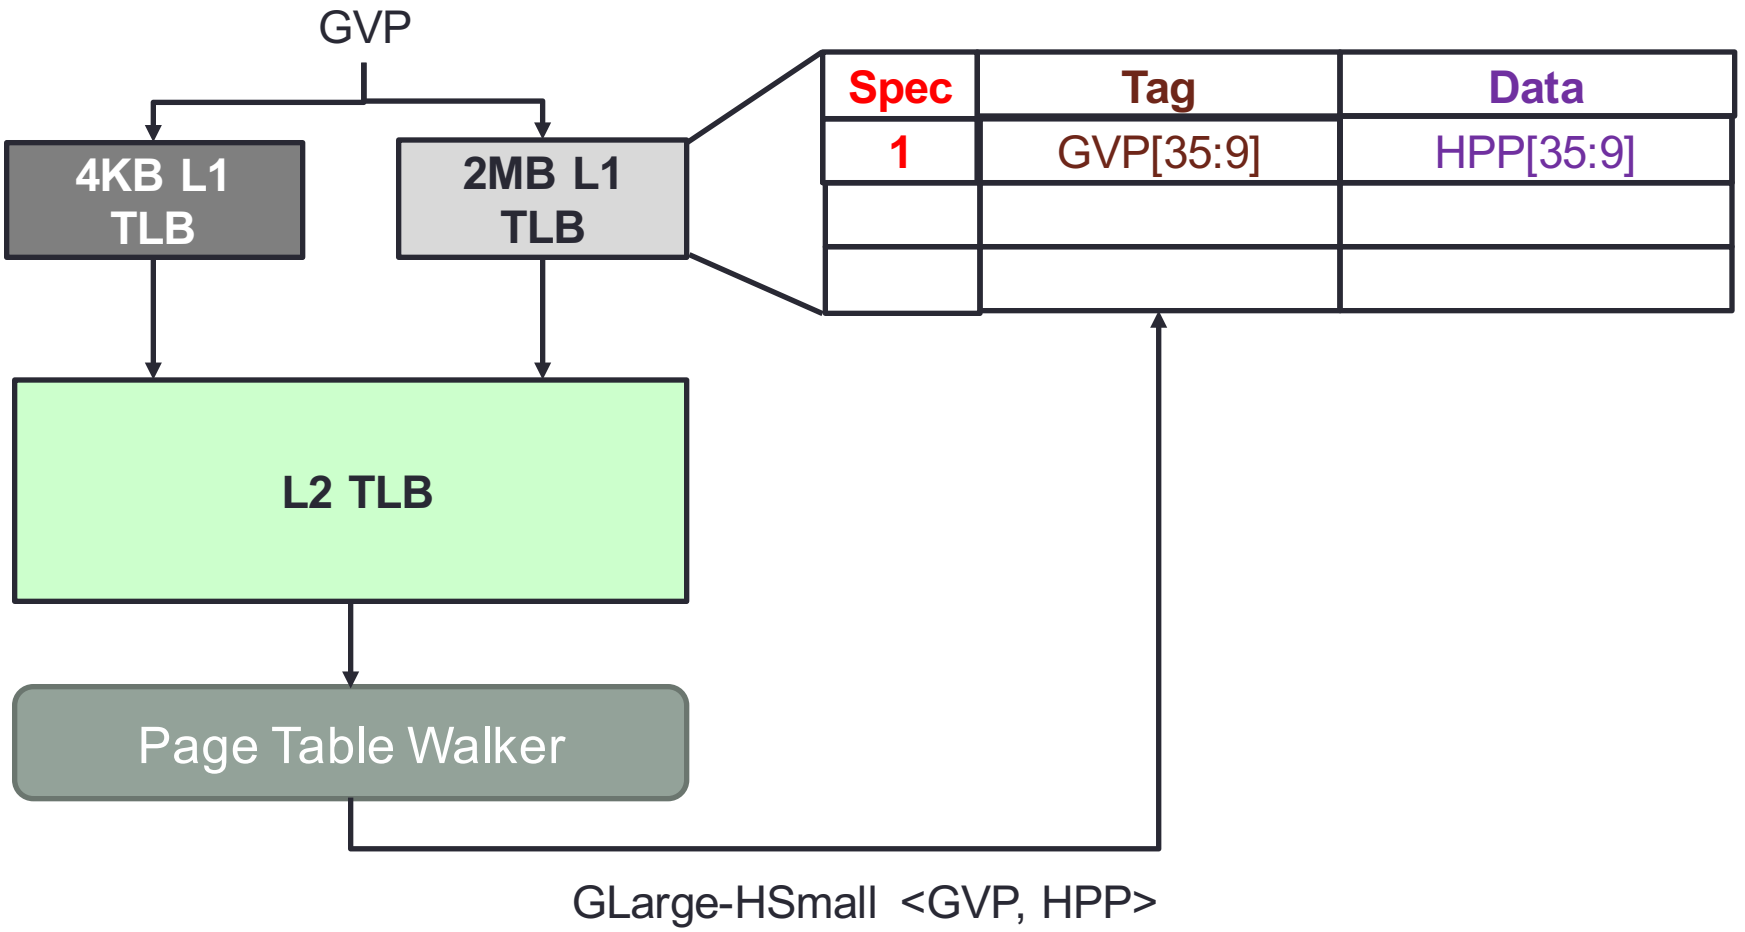

GLUE Speculation Timeline

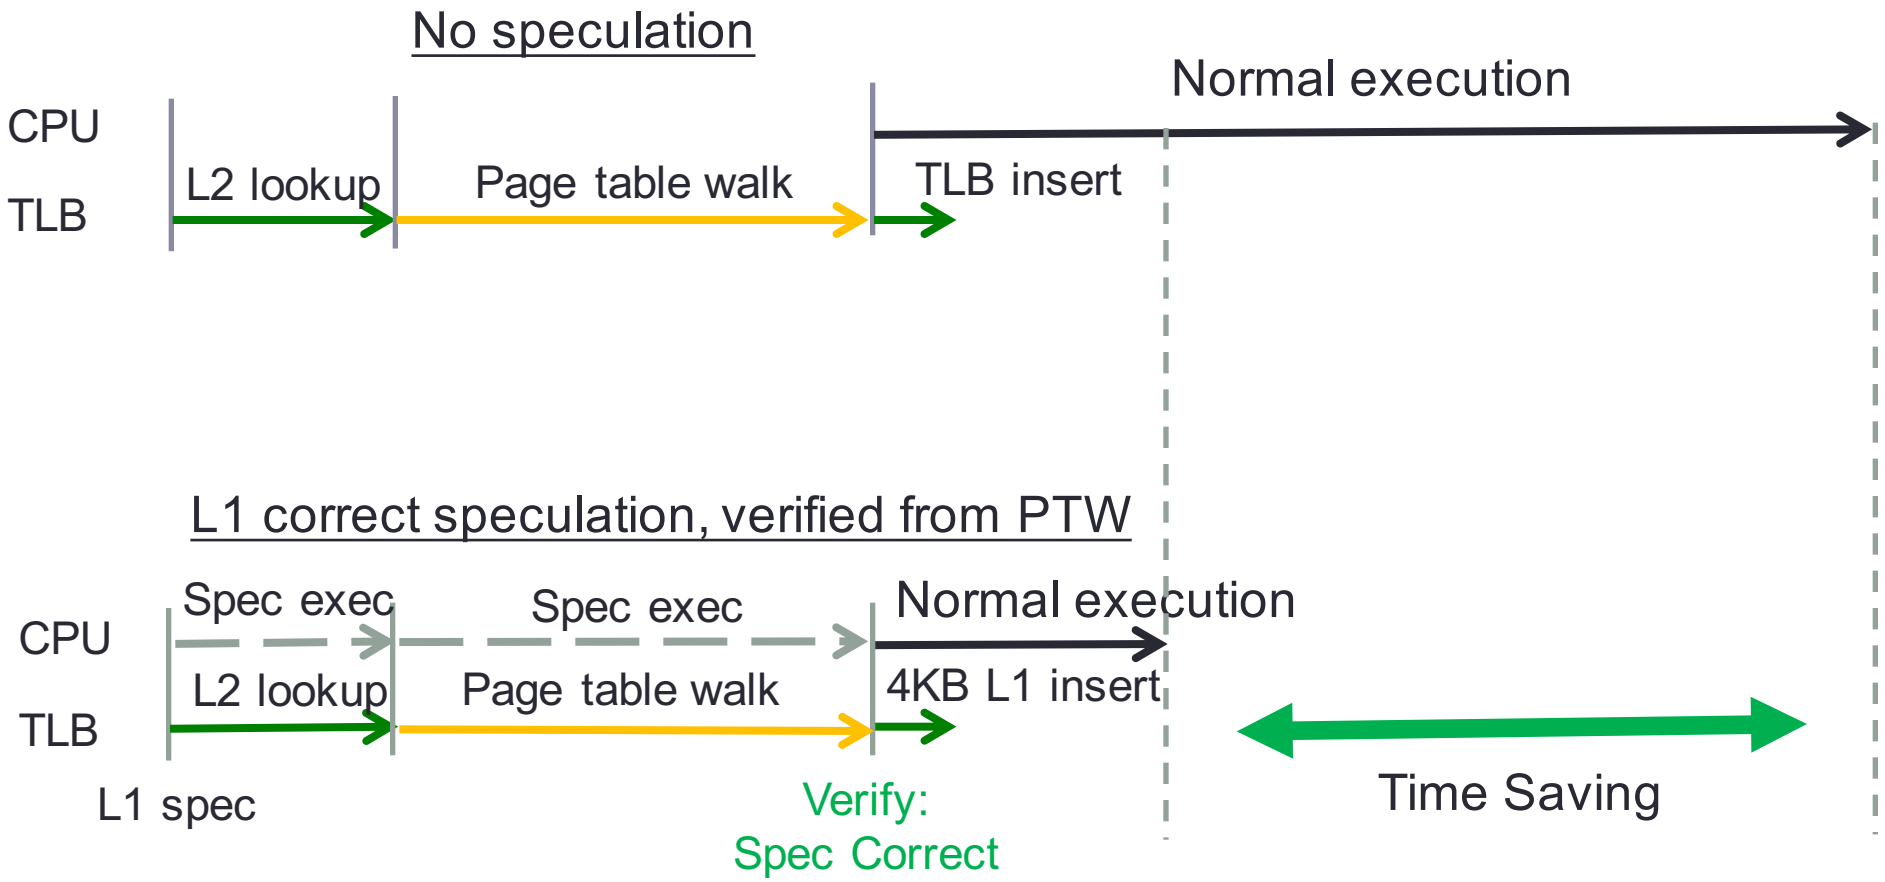

AMD

Methodology (See paper for details)

- Workloads:
  - Spec CPU2006
  - Biobench
  - Cloudsuite
- System:
  - 8 3.4 GHz cores, 24 GB RAM
  - Hypervisors: ESX, KVM
  - 8 VMs, 4 GB RAM each
- Trace-driven simulator:
  - Collect guest memory traces using Pin
  - Collect hypervisor memory traces using VMware scripts/customized KVM
  - TLB+cache simulator

How often/accurately do we speculate?

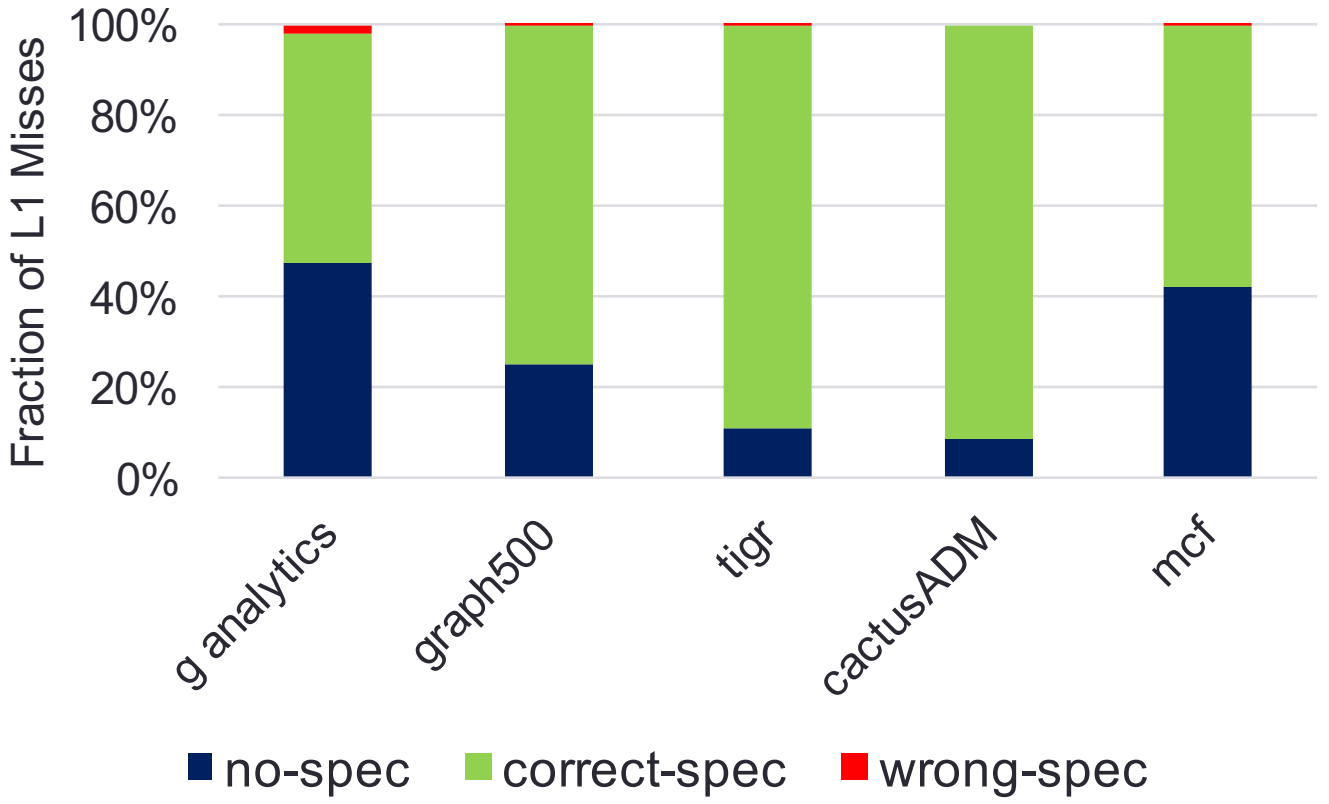

How effectively do we speculate?

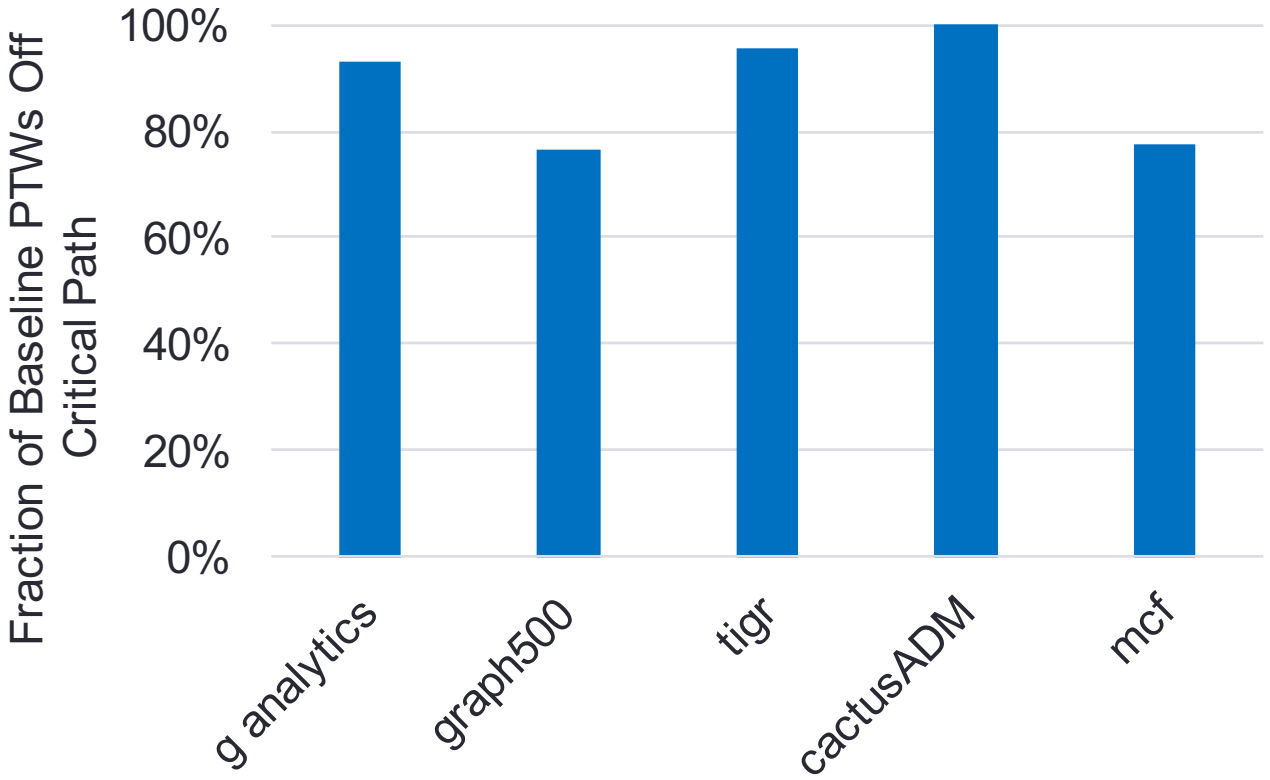

How much performance do we get?

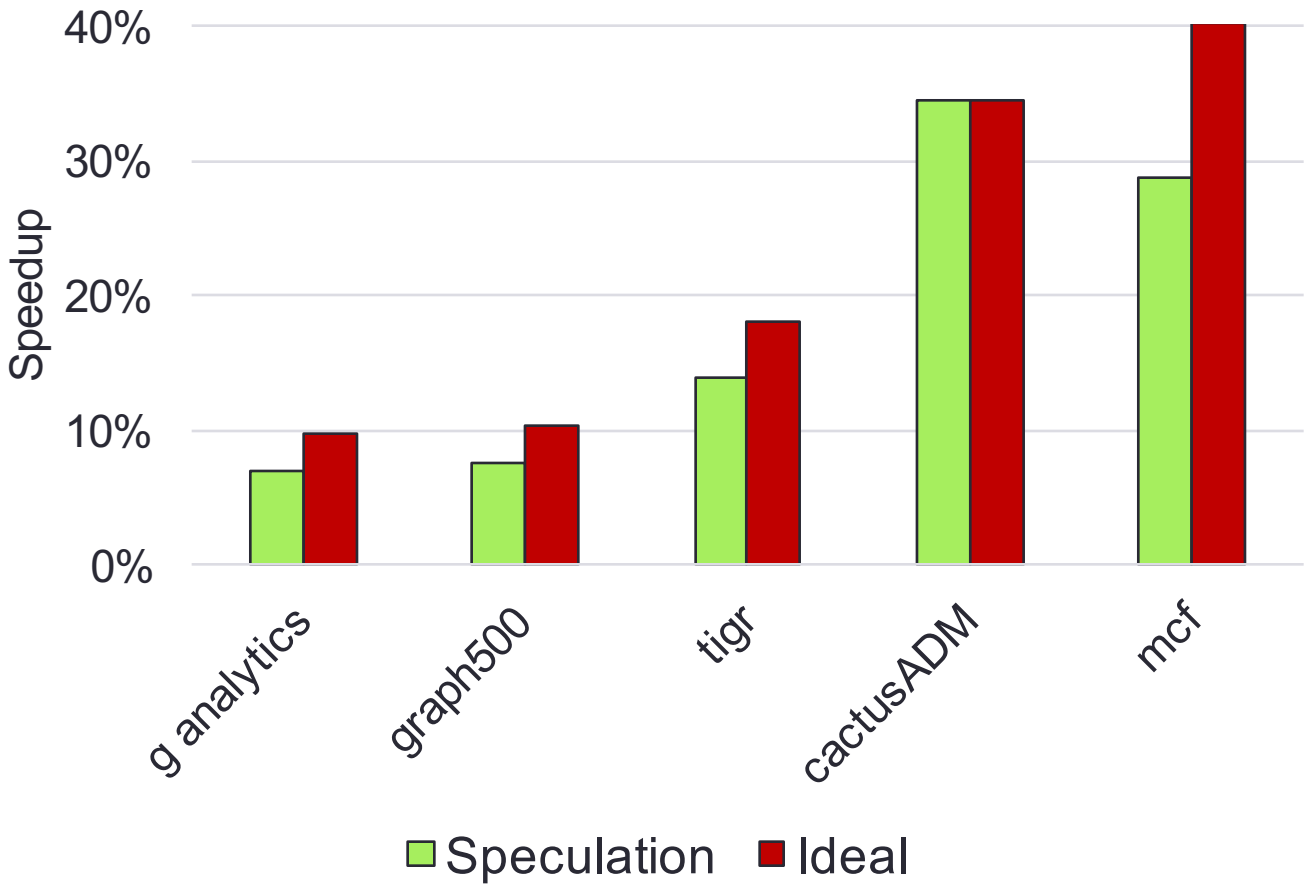

AMD

Conclusion

- Conflicting decision can be made in complex system with many layers.
- There likely still exists patterns that we can exploit.
- Our work demonstrates one example where this process can be handled effectively.

AMD

LARGE PAGES AND LIGHTWEIGHT MEMORY MANAGEMENT IN VIRTUALIZED ENVIRONMENTS: CAN YOU HAVE IT BOTH WAYS?

Binh Pham[§], Jan Vesely[§], Gabriel H. Loh[‡], Abhishek Bhattacharjee[§]

[§]Rutgers University

[‡]AMD Research