Address Translation in Virtual Machine

Peizhe Liu

System Virtualization

- You are already using a virtual machine provided by EngrIT.
- System virtual machine is powered by a software called VMM (Virtual Machine Monitor).
- We also call VM “Guest”, and VMM “Hypervisor.”

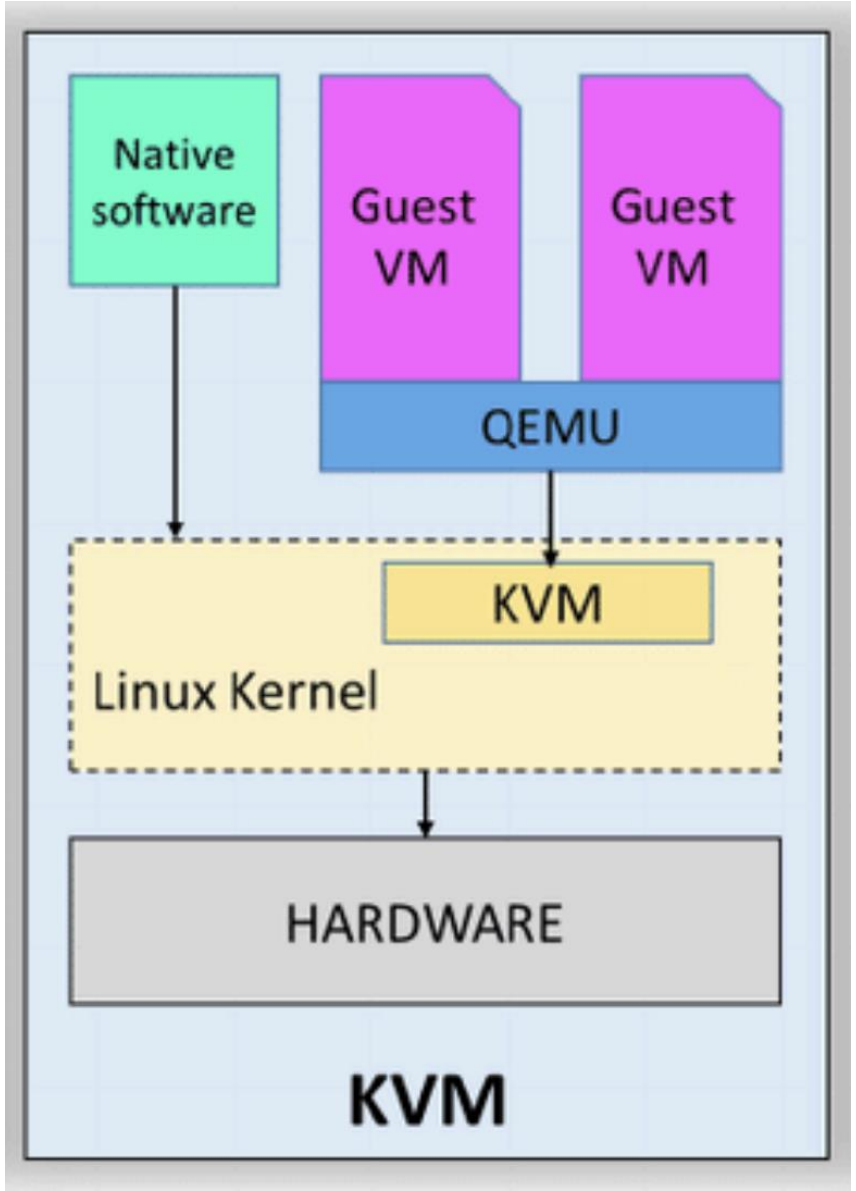

Types of Hypervisor

- Type-1 Hypervisor
  - Running on bare hardware
- Type-2 Hypervisor
  - Running on top of current OS
- Tradeoff?

Microsoft
Hyper-V

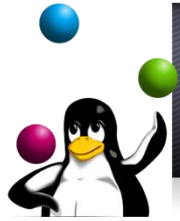

KVM

Virtualization in Cloud Computing

- Fundamental technique in today's cloud computing
- Any advantage over bare metal? Disadvantage?

Virtualization in Cloud Computing

- Efficiency (multi-tenant)
- Security (isolation)
- Flexibility (elastic scheduling)
- Manageability (snapshots)

- **But… overheads!**

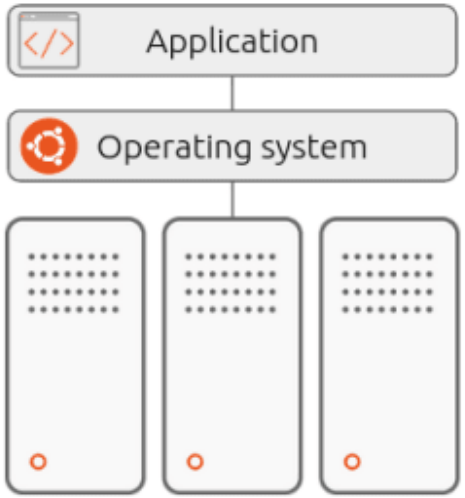

Traditional environment

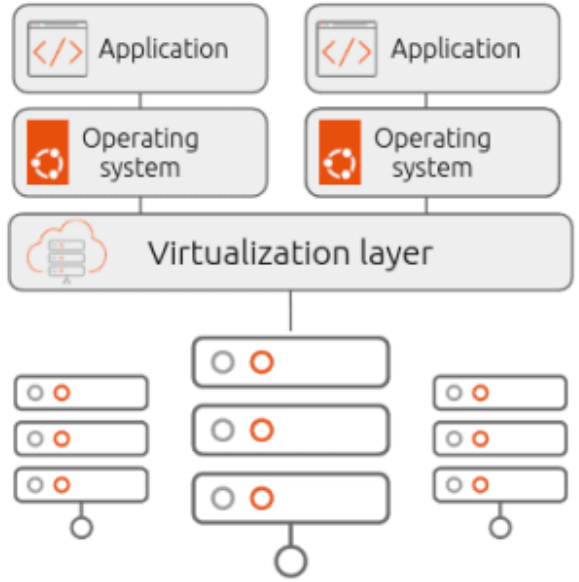

Virtualized environment

Virtual Memory

- Abstraction
- Isolation
- Efficient sharing
- Overcommit

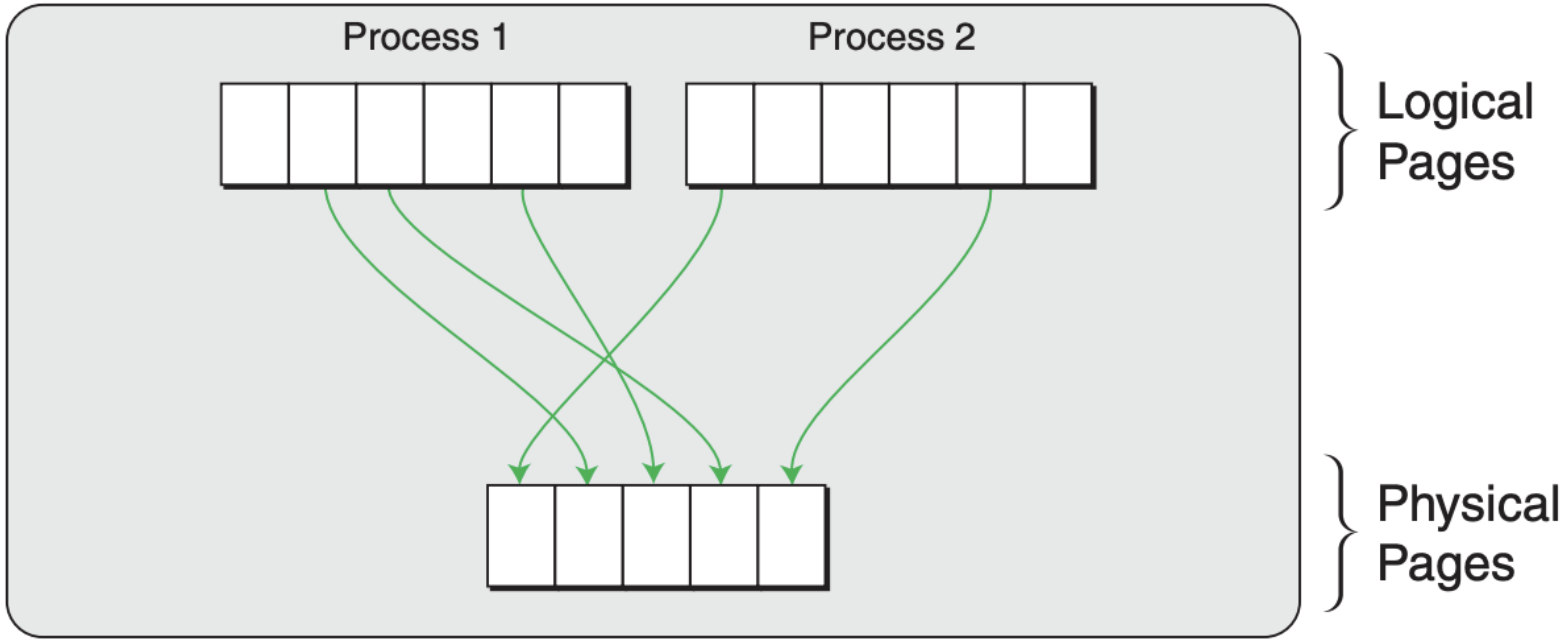

VMware, Inc. Performance Evaluation of Intel EPT Hardware Assist. https://www.vmware.com/docs/perf_esx_intel-ept-eval.pdf. 2009.

Address Translation in Bare Metal

- Handled by a hardware called MMU
- TLB hit? Done. Otherwise…

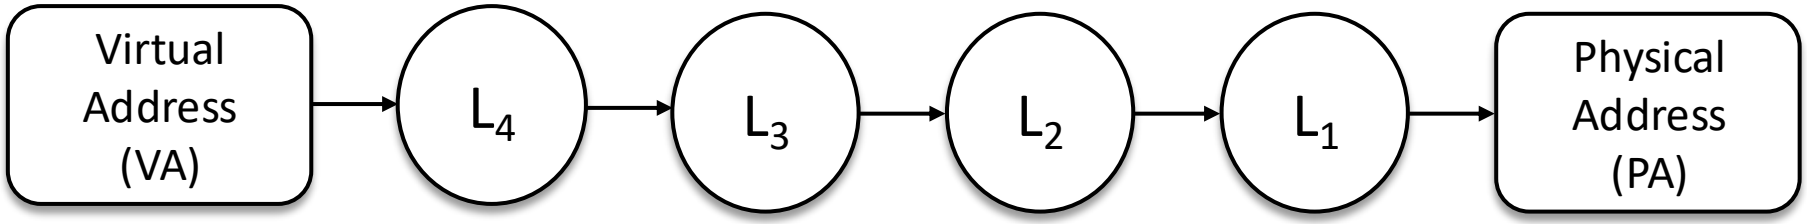

Address Translation in Virtual Machine

- Guest OS also have their own page tables.
- Guest physical address is NOT the host physical/machine address.

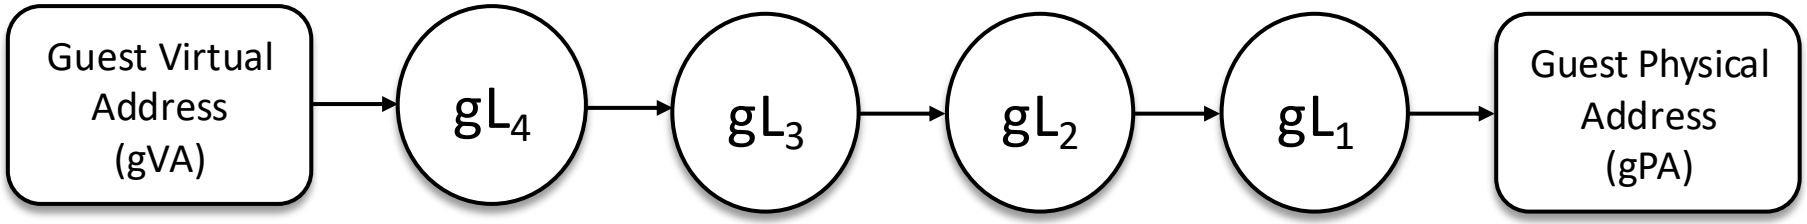

Extended Page Table (EPT)

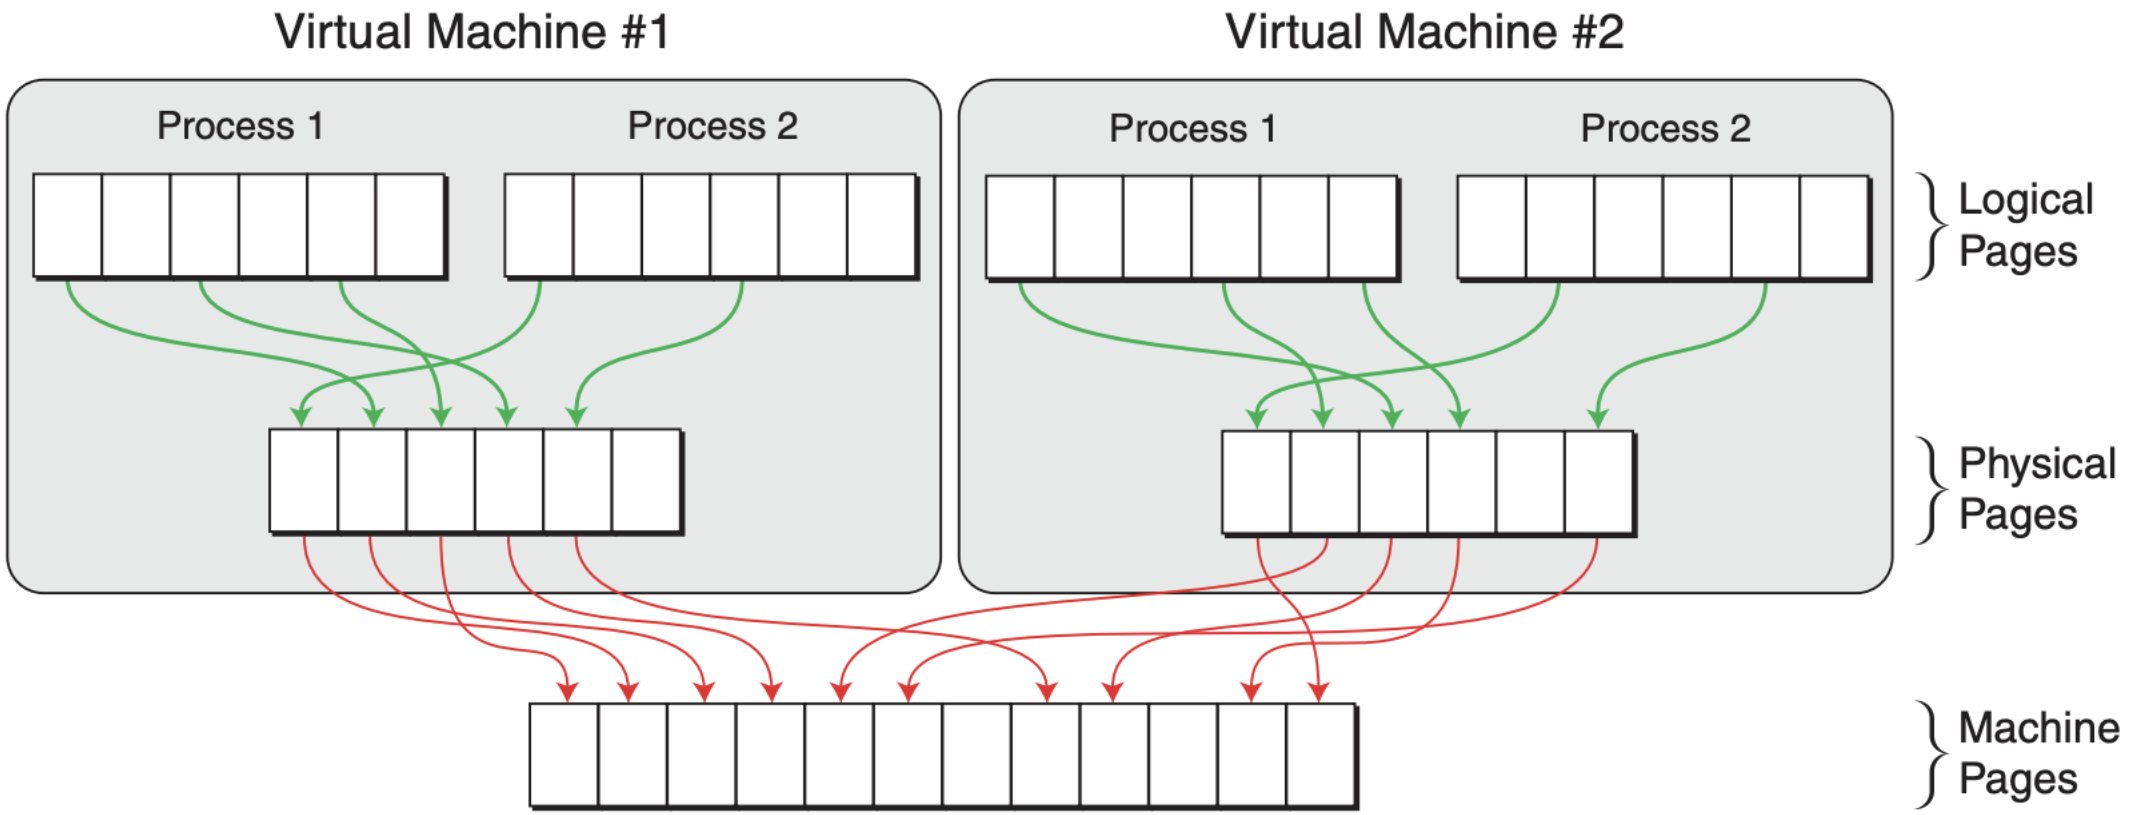

VMware, Inc. Performance Evaluation of Intel EPT Hardware Assist. https://www.vmware.com/docs/perf_esx_intel-ept-eval.pdf. 2009.

EPT Address Translation

- For every gPA and gL access, extend the translation and find the actual hPA and hL.
- This is a 2-D page walk.

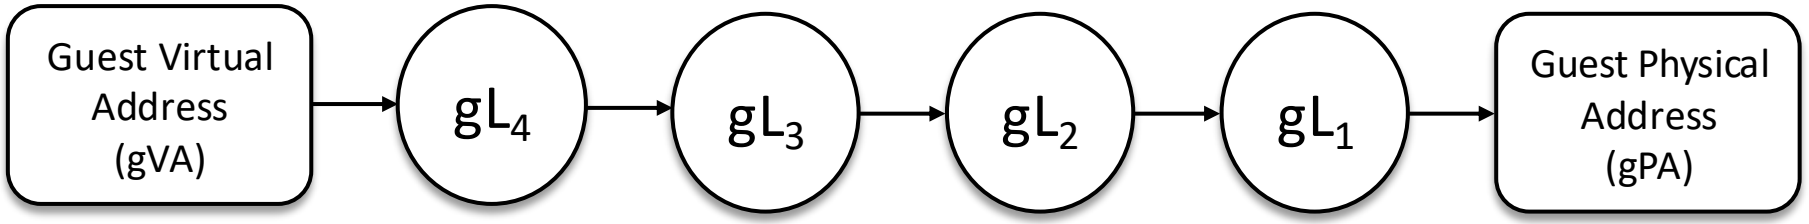

EPT Address Translation

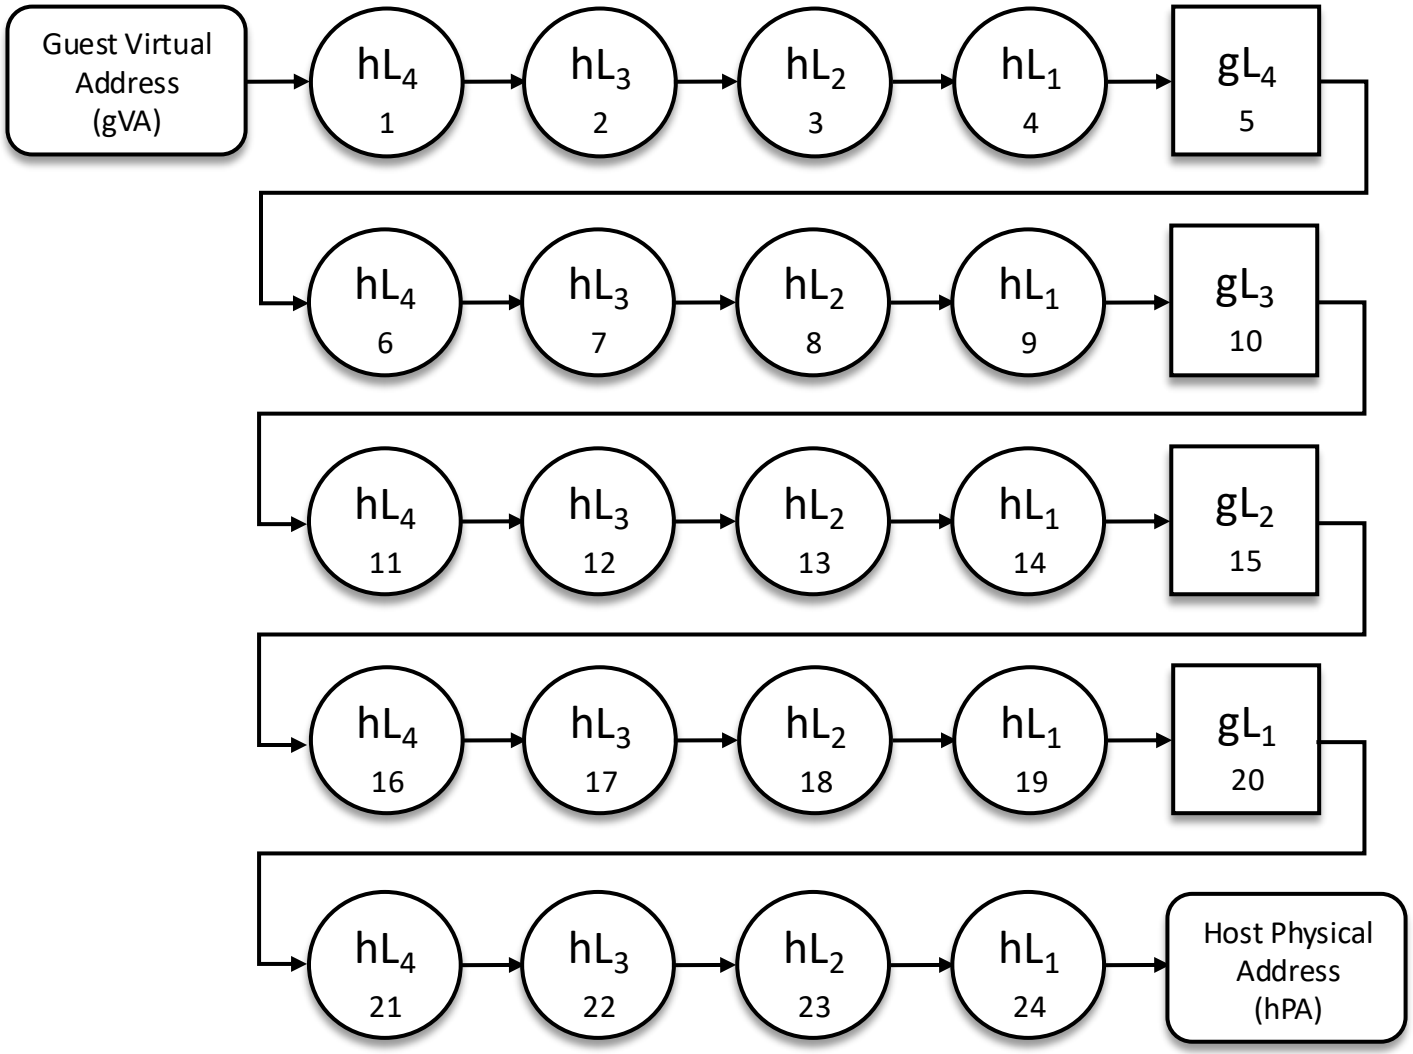

Translation Overheads

- EPT introduced translation overheads.
  - **1.46x** longer execution time versus native
  - Address translation contributed **43%** overhead

- Gets worse when scaling!
  - Higher level of page tables: **24 to 35** accesses
  - Nested virtualization: **2D to 3D** translation, **124** accesses, **4.13x** longer execution time!
  - Higher level of PT + nested: **215** accesses!

HugeGPT: Storing Guest Page Tables on Host Huge Pages to Accelerate Address Translation

PACT'23

Huge Page

- Larger page size means fewer levels
- Reduce TLB misses and PW memory accesses

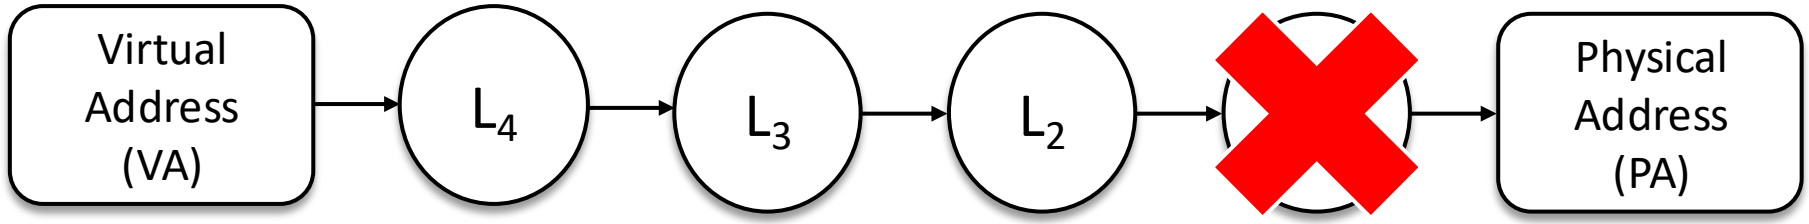

HugeGPT

- Store the **G**uest **P**age **T**able on the host **Huge** page.
- Software approach to speed up EPT address translations.

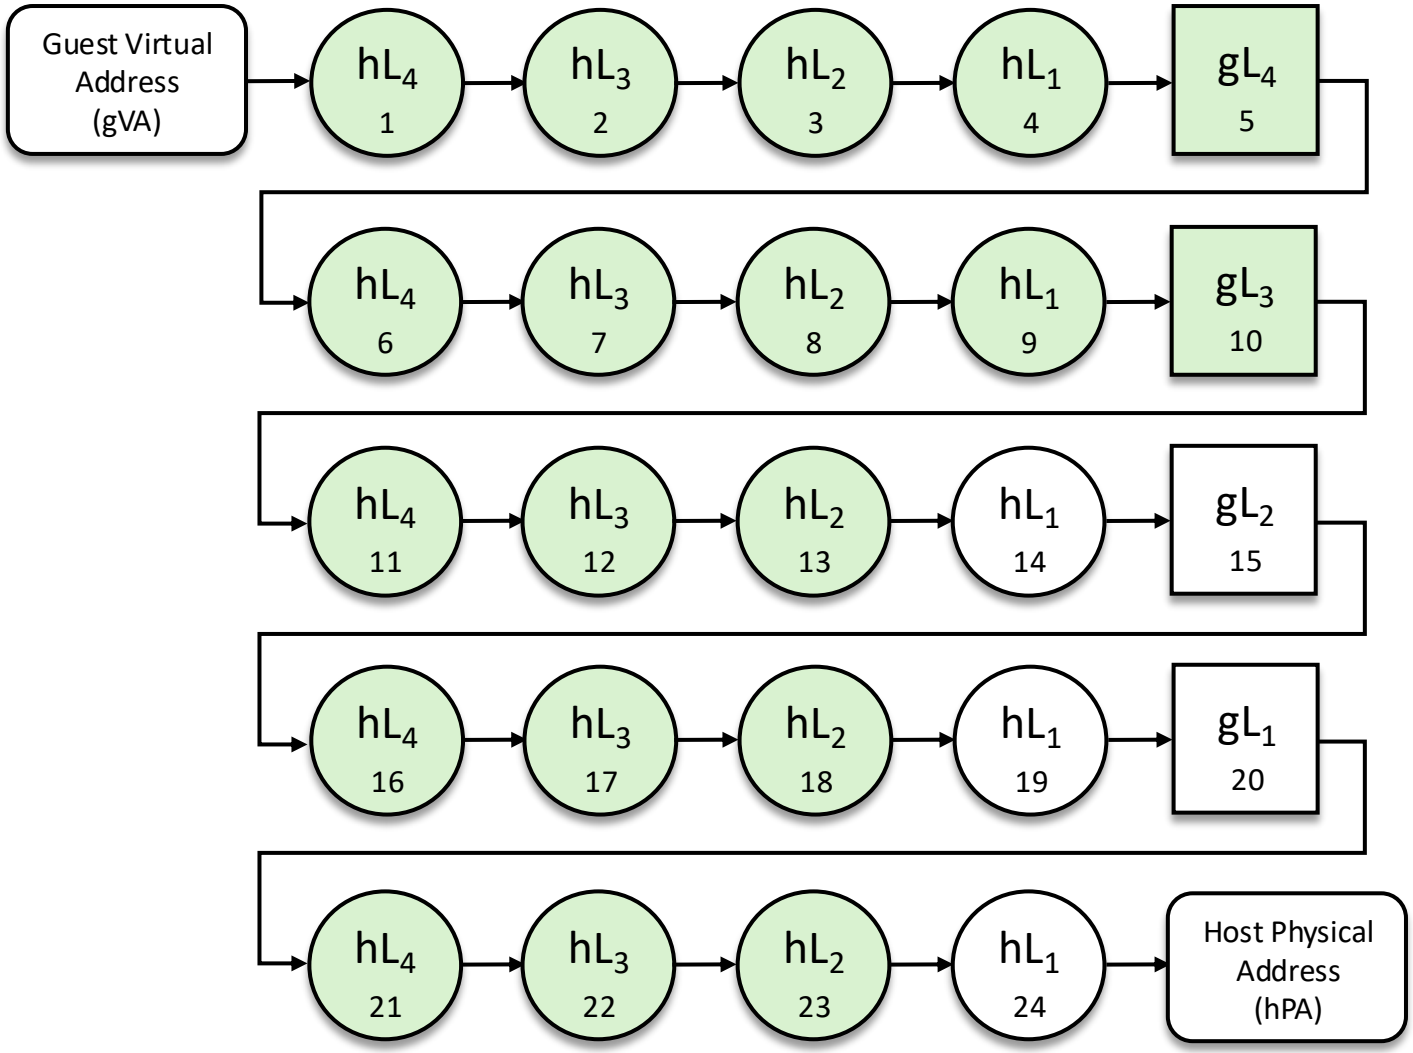

24 memory accesses in the worst case

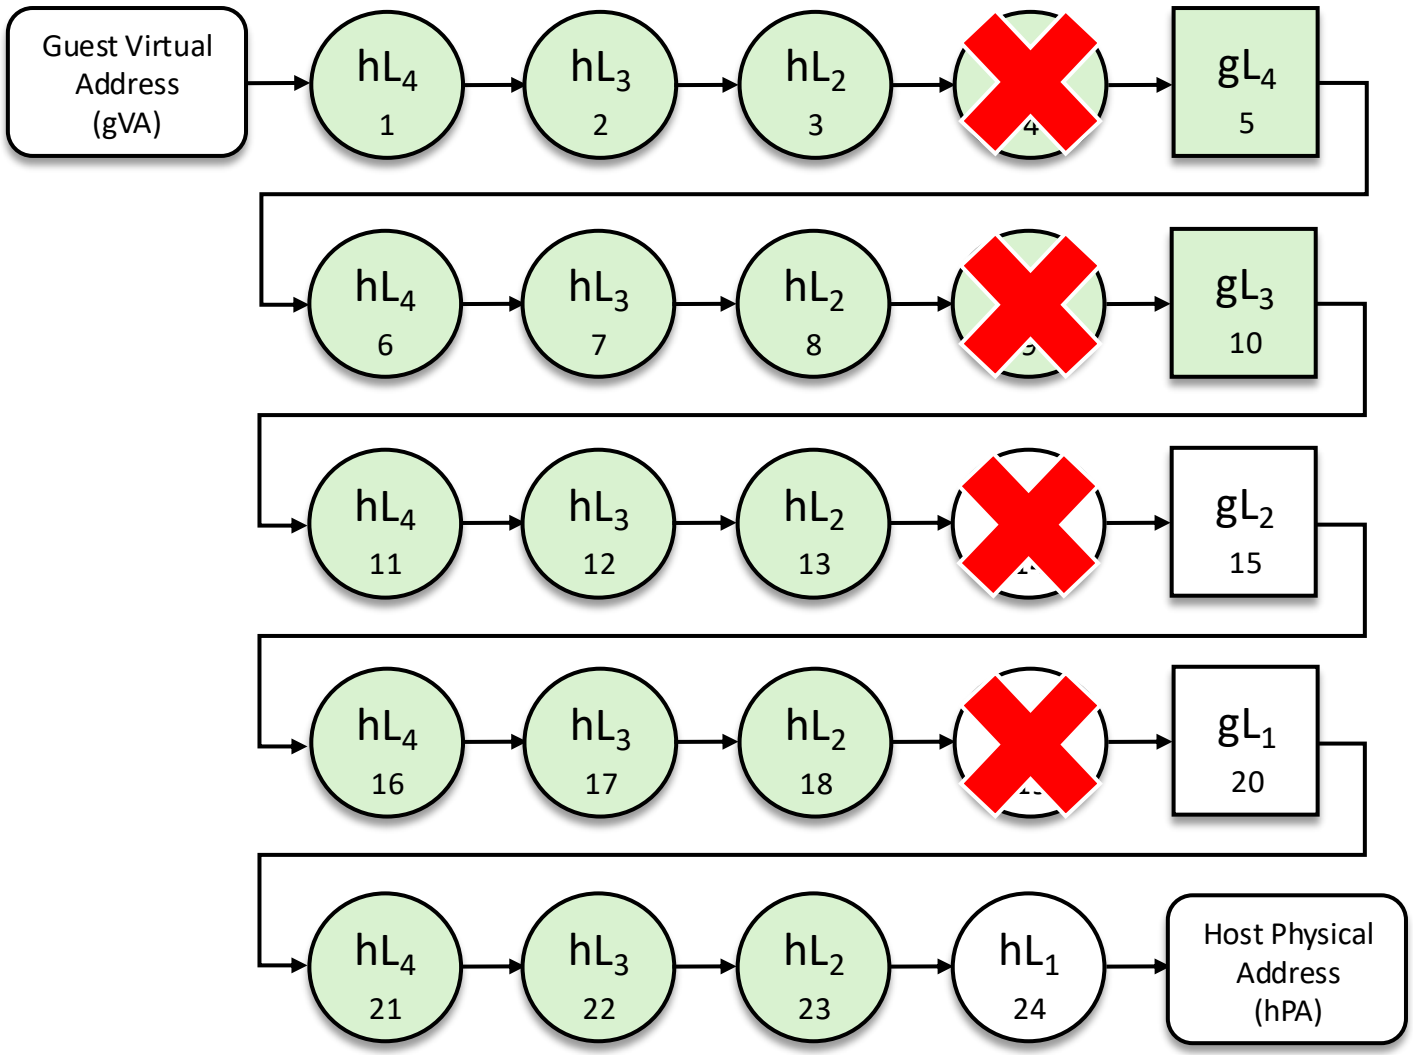

20 memory accesses in the worst case

HugeGPT Evaluation

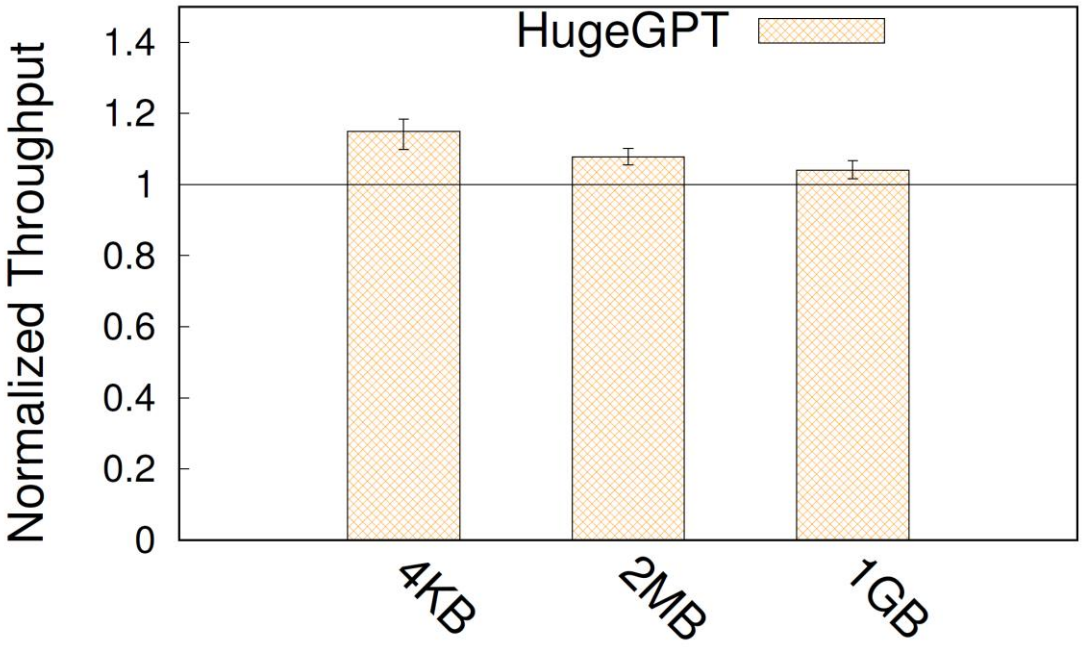

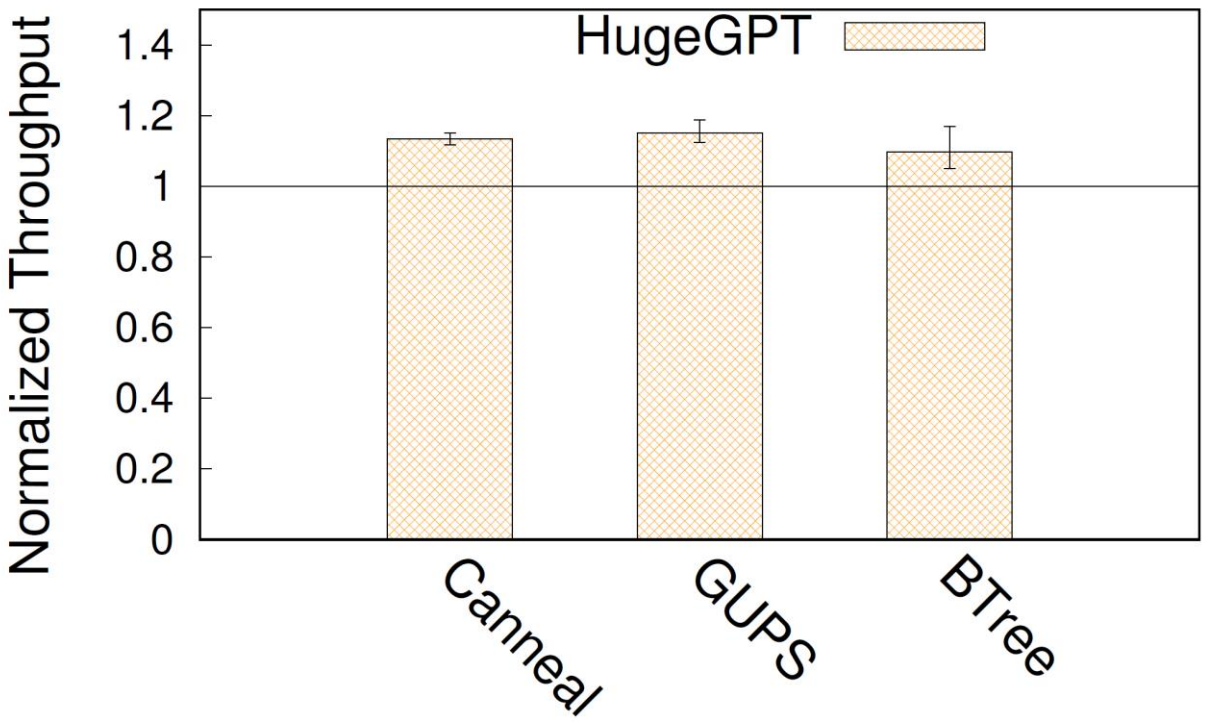

Jia, W., Zhang, J., Shan, J., Du, Y., Ding, X., and Xu, T. HugeGPT: Storing Guest Page Tables on Host Huge Pages to Accelerate Address Translation. In *Proceedings of the 32nd International Conference on Parallel Architectures and Compilation Techniques (PACT'23)* (Oct. 2023).

Artifact

- Implemented on Linux v.6.1.81
- https://github.com/xlab-uiuc/hugegpt-linux

Thank You!

Q/A