

# LATR: Lazy Translation Coherence

## Lightning Talk

Mohan Kumar<sup>\*</sup>, **Steffen Maass**<sup>\*</sup>, Sanidhya Kashyap, Ján Veselý<sup>‡</sup>,  
Zi Yan<sup>‡</sup>, Taesoo Kim, Abhishek Bhattacharjee<sup>‡</sup>, Tushar Krishna


Georgia Institute of Technology    <sup>‡</sup>Rutgers University

\* Co-First Authors

March 28, 2018


# TLBs and TLB shutdowns

- TLB: Translation Lookaside Buffer
- TLB Shutdown: Process of removing TLB mappings
- Problem: TLB shutdown can bottleneck applications, example Apache:



# TLBs and TLB shutdowns

- TLB: Translation Lookaside Buffer
- TLB Shutdown: Process of removing TLB mappings
- Problem: TLB shutdown can bottleneck applications, example Apache:



- Our insight: *Synchronous* TLB shutdown is expensive:  
⇒ Up to 130  $\mu$ s on 120 cores

- LATR: **L**azy **T**ranslation Coherence
- Key insight: LATR enables *asynchronous* shutdowns
- Up to 70% reduction for `munmap()`, 60% higher throughput for Apache
- LATR scales to a 120 core system

- LATR: **L**azy **T**ranslation Coherence

Learn more in our talk:

**Memory 2 (Session 2) @ 9:20am**