

TOSTING: Investigating Total Store Ordering on ARM

ARCS'23

Lars Wrenger Dominik Töllner Daniel Lohmann

Leibniz Universität Hannover
{wrenger,toellner,lohmann}@sra.uni-hannover.de

14.06.23




- Backward compatibility with Rosetta 2 (x86 translation layer)
- Emulation is challenging due to different **Memory Ordering** models

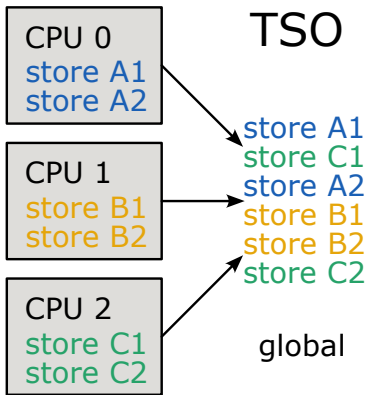
- CPUs execute instructions in parallel and out of order to hide latencies

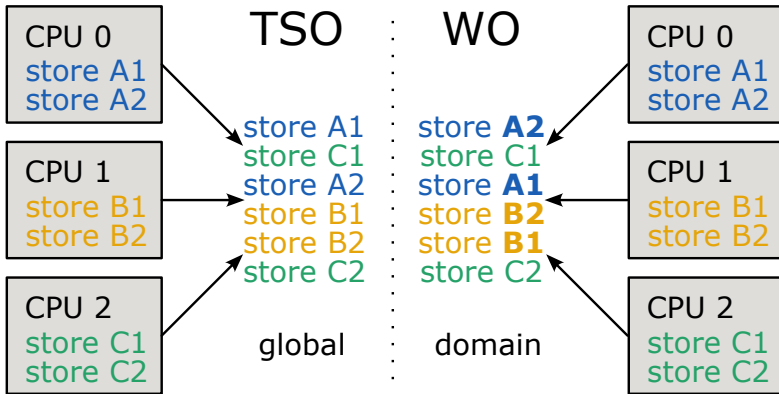
Memory Ordering – Gharachorloo et al. 1990; Higham et al. 1997; Pulte et al. 2017


Defines how accesses to shared memory can be observed by other participants.

- Stricter ordering simplifies the programming model
 - Total Store Ordering (TSO) on x86
- Weaker ordering allows more latency hiding
 - Weak Ordering (WO) on ARMv8








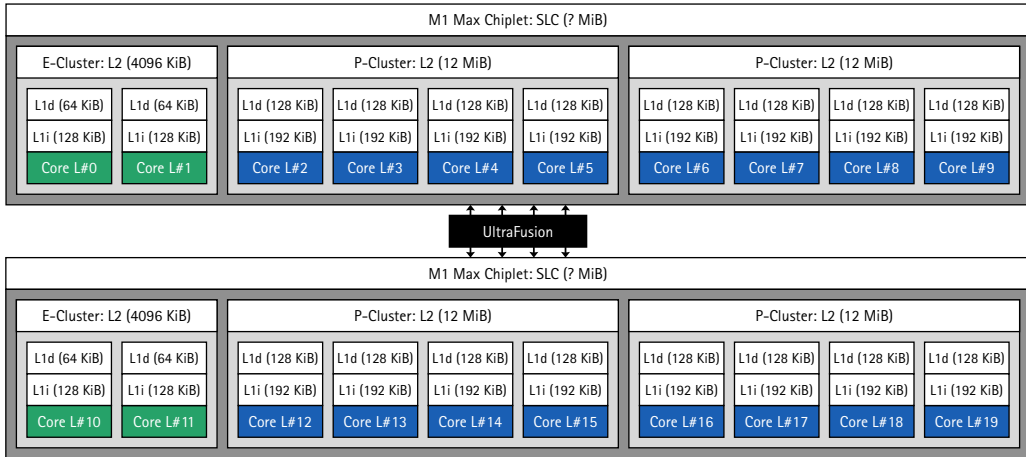


- Emulating TSO on a WO platform requires synchronizing every memory access.
- The Apple M1 implements TSO in hardware for the x86 emulation



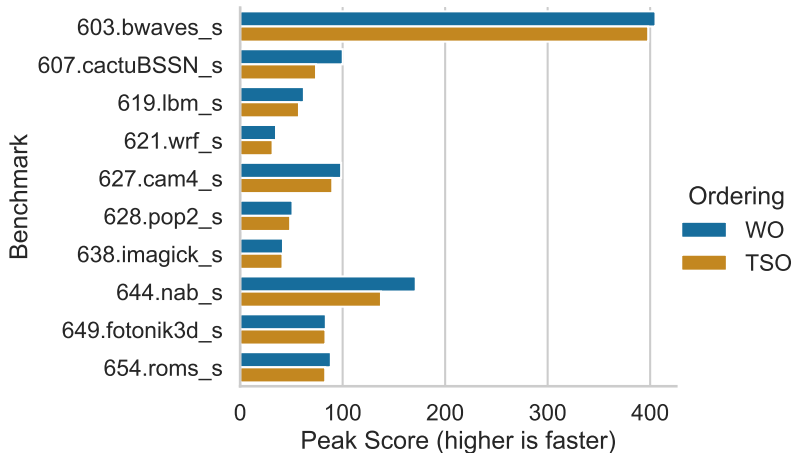
- Emulating TSO on a WO platform requires synchronizing every memory access.
- The Apple M1 implements TSO in hardware for the x86 emulation


What impact do different memory ordering models have on performance?




- ARMv8.3, Supports WO and TSO (by setting first bit of the ACTLR_EL1), 128 GiB

- Standardized benchmarking suite with realistic workloads
 - CPU- and memory-heavy applications (3D simulation, physics, image manipulation)
- We focus on multicore benchmarks
- Approach: Execute the same benchmark binaries with and without TSO







⇒ TSO is 8.94 % slower!




- Reader and writer iterate sequentially through the buffer
- Counting the number of completed iterations





- Reader and writer iterate sequentially through the buffer
- Counting the number of completed iterations

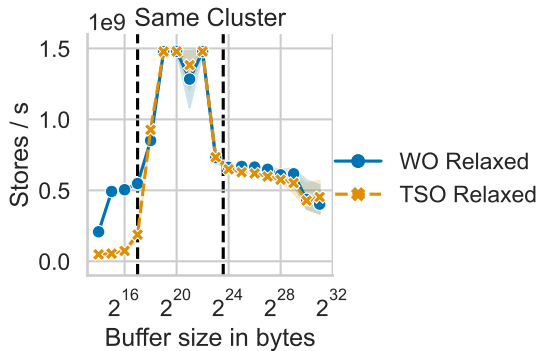



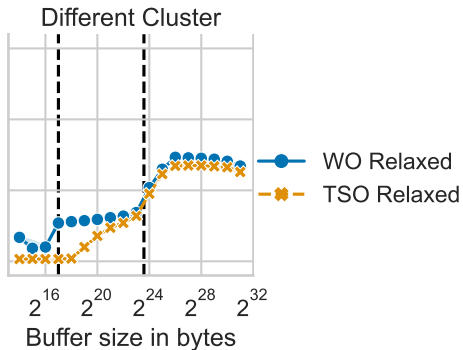
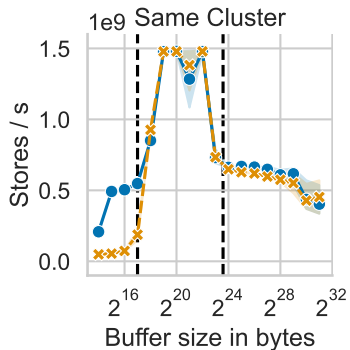


- Reader and writer iterate sequentially through the buffer
- Counting the number of completed iterations

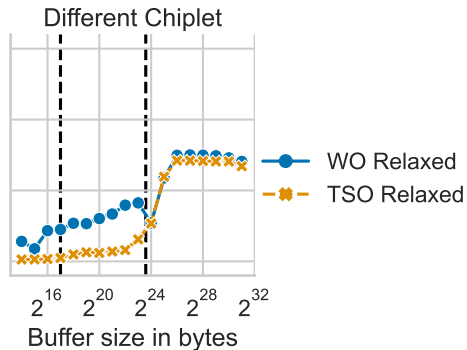
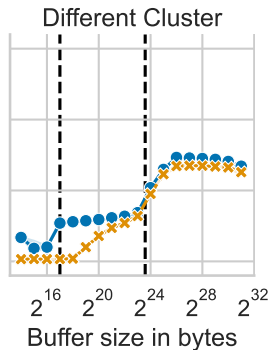






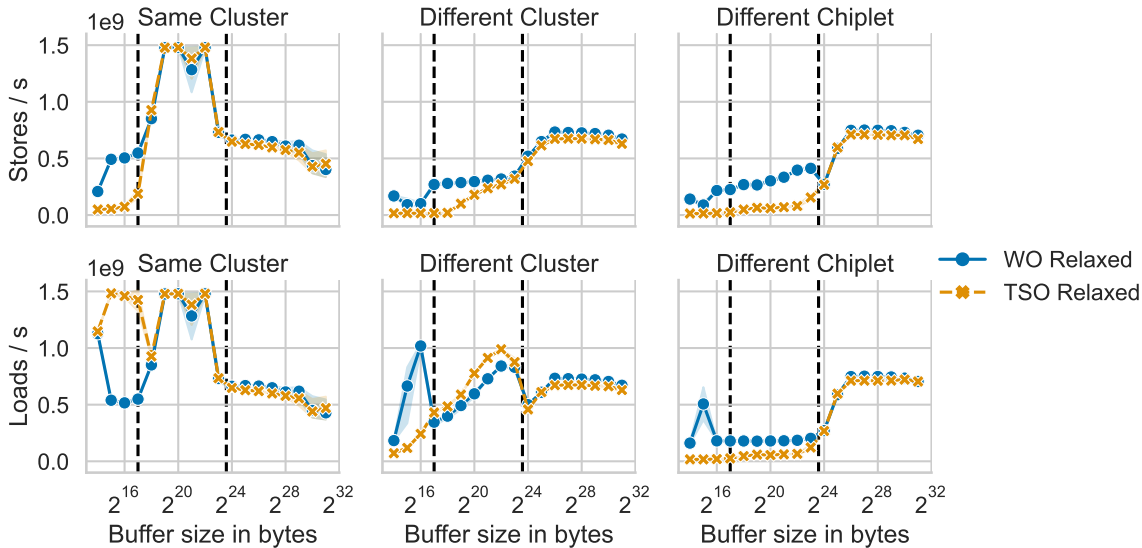
- Reader and writer iterate sequentially through the buffer
- Counting the number of completed iterations

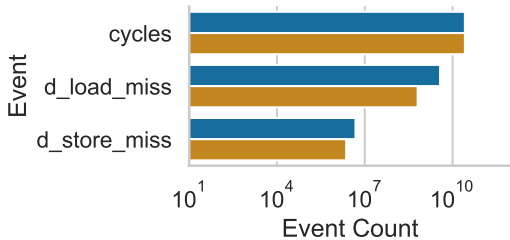
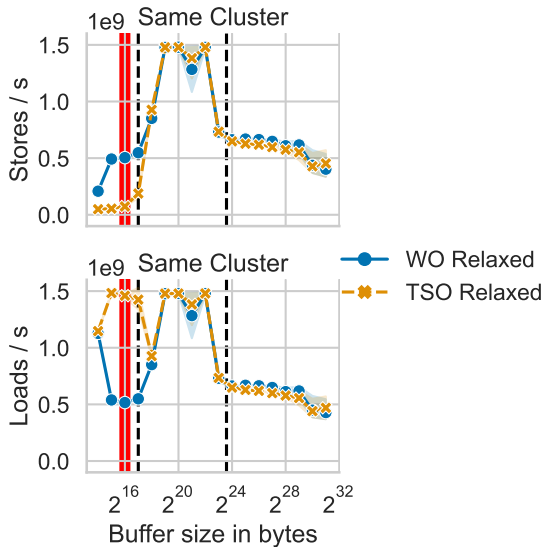








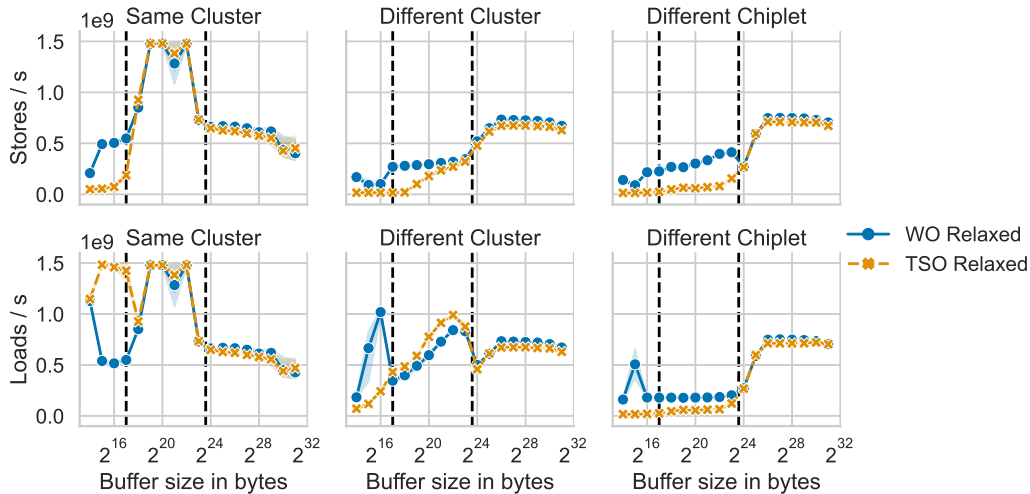


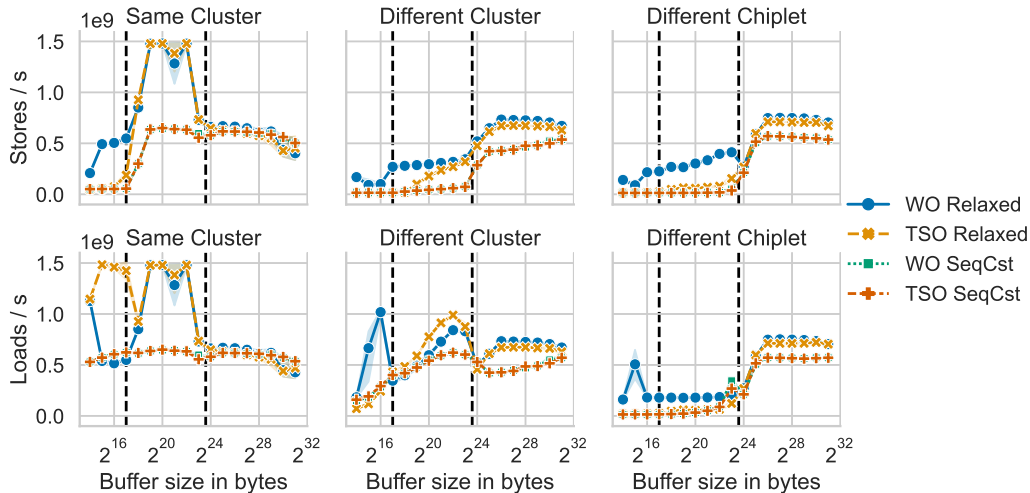


Weak stores are faster

→ More cache invalidations

→ More load misses





■ Sequentially consistent instructions (STLR, LDAR) are independent from TSO

- SPEC CPU 2017
 - TSO is 8.94 % slower
- Microbenchmarks
 - WO is faster for writes on small buffers
 - TSO is faster reads when stores are slow
- Sparse Information about the M1 Architecture



- Microarchitecture is being reversed by the community (Johnson 2023)
 - Measuring Load/Store buffer sizes
- Asahi Linux developers have documented most registers (*Asahi Linux Wiki* 2023)
- Cache hierarchy not documented
 - We have reproduced L1/L2 cache sizes
 - Cacheline size is reported as 128 B but is actually 64 B

- ▶ *Asahi Linux Wiki* (2023). <https://github.com/AsahiLinux/docs/wiki> – accessed 2023-03-23. (Visited on 03/23/2023).
- ▶ Gharachorloo, Kourosh et al. (1990). "Memory Consistency and Event Ordering in Scalable Shared-Memory Multiprocessors". In: *SIGARCH Comput. Archit. News* 18.2SI, pp. 15–26.
- ▶ Higham, Lisa et al. (1997). "Defining and comparing memory consistency models". In.
- ▶ Johnson, Dougall (2023). *Apple M1 Microarchitecture Research*.
<https://dougallj.github.io/applecpu/firestorm.html> – accessed 2023-03-23. (Visited on 03/23/2023).
- ▶ Pulte, Christopher et al. (2017). "Simplifying ARM Concurrency: Multicopy-Atomic Axiomatic and Operational Models for ARMv8". In: *Proc. ACM Program. Lang.* 2.POPL.