Modern locking

- Traditional goal is to make lock regions as small as possible to lower contention
  - Makes cache communication overhead dominate with any contention
  - Non deterministic performance for small regions
- Work in critical section > Transfer Latency
  - Work time should be at least several hundred ns in 1-2 socket, more on larger systems
  - Applies to IPIs, read locks, reference counts too

Lock behavior with different critical section times on 2S

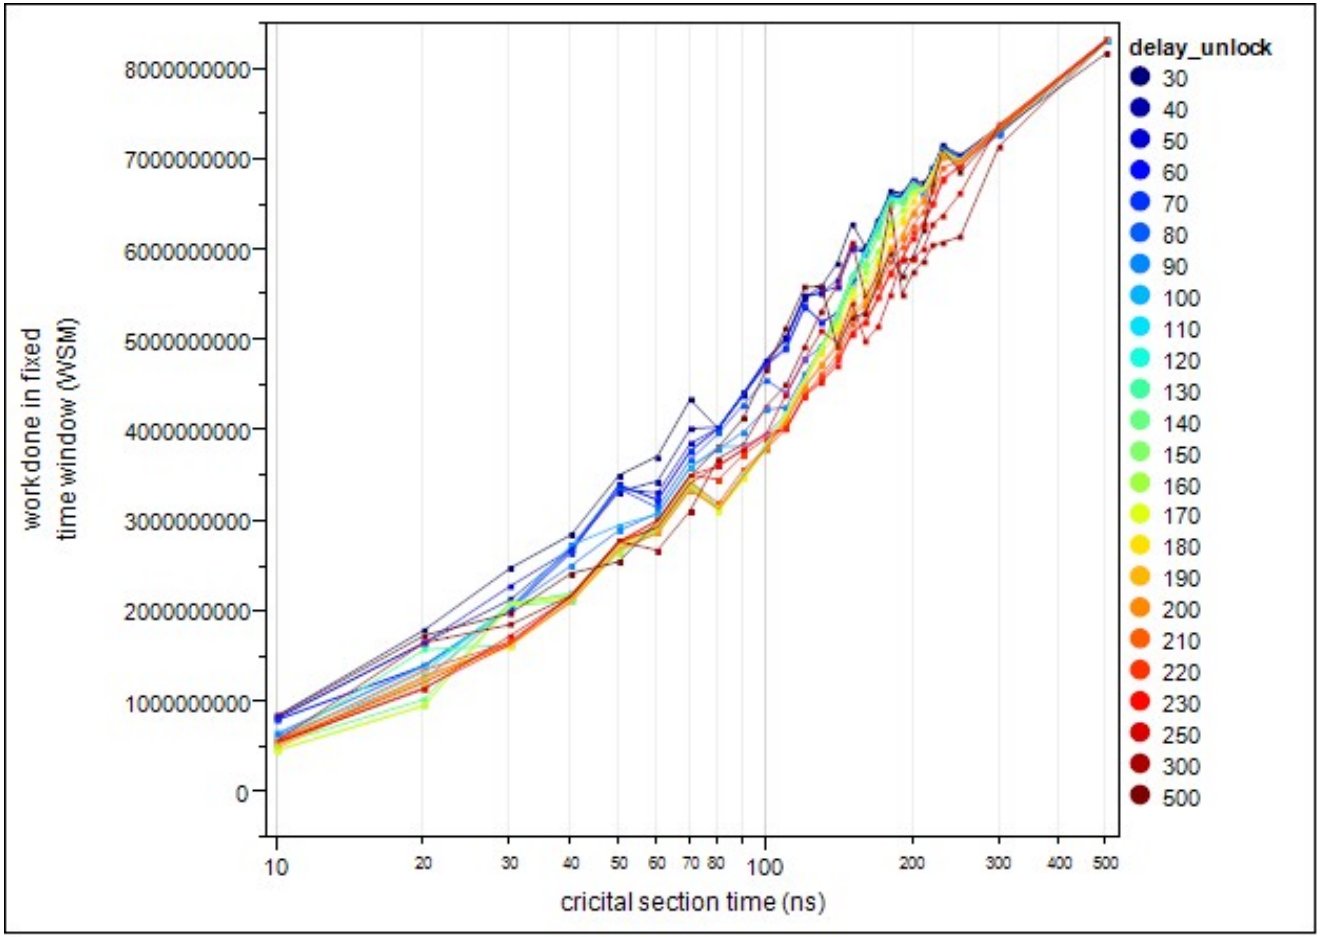

Implications

- Every time we process only a single object inside a lock this is non scalable
  - Need restructuring of many interfaces to batch
- Best batching factors may change between machines (communication latency)
  - Trade off between lock contention and batching
  - Global tunables for batch sizes?
  - http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/xeon-lock-scaling-analysis-paper.pdf