The Chemistry of Concurrent and Distributed Programming Workshop Agadir, Morocco, May 2015

Concurrent Data Structures: Fast but Relaxed versus Strict but Slow Semantics Christoph Kirsch University of Salzburg

Joint work with A. Haas, M. Lippautz, H. Payer, and A. Sokolova

## **Concurrent Data Structures**

 We are interested in designing and implementing concurrent data structures that are <u>fast</u> and <u>scale</u> on multicore hardware.



#### 4 Processors w/ 10 Cores each w/ 2 Hyperthreads each

CPU Socket 0						
HT HT						
16 KB data						
HT HT						
L1: 32 KB instr 16 KB data						
L2: 256 KB data						
L3: Cache 24 MB						
					_ 128 GE	
					_	

L3: Cache 24 MB						
HT HT L1: 32 KB instr 16 KB data						
L2: 256 KB data						
HT HT L1: 32 KB instr 16 KB data						
L2: 256 KB data						
CPU Socket 2						

CPU Socket 1					
HT HT	HT HT	HT HT	НТ НТ	НТ НТ	
L1: 32 KB instr 16 KB data					
L2: 256 KB data					
HT HT	HT HT	НТ НТ	НТ НТ	HT HT	
L1: 32 KB instr 16 KB data					
L2: 256 KB data					

L3: Cache 24 MB

**1**emory

L3: Cache 24 MB					
HT HT	HT HT	НТ НТ	HT HT	HT HT	
L1: 32 KB instr 16 KB data					
L2: 256 KB data					
HT HT	НТ НТ	HT HT	НТ НТ	НТ НТ	
L1: 32 KB instr 16 KB data					
L2: 256 KB data					
CPU Socket 3					

## **Multicore Scalability**



number of cores





Fig. 2. Performance and scalablity of producer/consumer microbenchmarks with an increasing number of threads



- Scal is a collection of concurrent data structures designed by us (<u>underlined</u>) and others, plus a benchmarking framework:
- Treiber Stack [IBM86]

\*\*

- \* Timestamped Stack [POPL15], k-Stack (relaxed) [POPL13]
- Michael-Scott Queue [PODC96]
- LCRQ [PPoPP13], Segment Queue (relaxed) [OPODIS10]
  - Timestamped Queue [<u>POPL15</u>], Distributed Queue (relaxed) [<u>CF13</u>], k-FIFO Queue (relaxed) [<u>PaCT13</u>]
  - \* Timestamped Deque [POPL15]

## Timestamping







## **Time Sources**

- TS-atomic: fetch and increment counter
- \* TS- stutter: stuttering counter (Lamport clock)
- \* TS-hardware: hardware clock
- \* TS-interval: interval hardware clock
- \* TS-CAS: compare-and-swap interval counter





(a) High contention benchmark on the 40core machine.



(c) Low contention benchmark on the 40core machine.



(b) High contention benchmark on the 64core machine.



(d) Low contention benchmark on the 64core machine.

Figure 2.5: TS stack performance on the 40-core machine (left) and on the 64-core machine (right) in the producer-consumer benchmark.



Figure 2.9: High-contention producer-consumer benchmark using TS-interval and TS-CAS timestamping with increasing delay on the 64-core machine, exercising 32 producers and 32 consumers.

# Semantics VS. Non-Determinism VS. Performance

## Order Deviation (Queues) [RACES12]

- \* We record the times when enqueue and dequeue operations are <u>invoked</u> on a concurrent queue
- For each element E inserted by an enqueue operation invoked at time T<sub>e</sub> and removed by a dequeue operation invoked at time T<sub>d</sub> we compute the <u>order deviation</u> of E as:
  - the number of elements that were inserted by enqueue operations invoked before T<sub>e</sub> and removed by dequeue operations invoked after T<sub>d</sub>
- \* We present the <u>average order deviation</u> of all elements inserted and removed in a run



(a) High-contention producer-consumer benchmark on the 40-core machine.



(c) Low-contention producer-consumer benchmark on the 40-core machine.



(b) High-contention producer-consumer benchmark on the 64-core machine.



(d) Low-contention producer-consumer benchmark on the 64-core machine.

Figure 7.1: Order deviation in the producer-consumer benchmark on the 40-core machine with 40 producers and 40 consumers (left), and on the 64-core machine with 32 producers and 32 consumers (right).



(a) High contention multiple-producer single-consumer benchmark on the 40core machine.



(c) Low contention multiple-producer single-consumer benchmark on the 40core machine.



(b) High contention multiple-producer single-consumer benchmark on the 64core machine.



(d) Low contention multiple-producer single-consumer benchmark on the 64core machine.

Figure 7.2: Order deviation in the multiple-producer single-consumer benchmark on the 40-core machine with 40 producers (left), and on the 64-core machine with 32 producers (right).



(a) High-contention single-producer multiple consumer benchmark on the 40-core machine.



(c) Low-contention single-producer multiple consumer benchmark on the 40-core machine.



(b) High-contention single-producer multiple consumer benchmark on the 64-core machine.



(d) Low-contention single-producer multiple consumer benchmark on the 64-core machine.

Figure 7.3: Order deviation in the single-producer multiple-consumer benchmark on the 40-core machine with 40 producers (left), and on the 64-core machine with 32 producers (right).

