

Splash-4

Improving Scalability with Lock-Free Constructs

Eduardo José Gómez-Hernández¹, Ruixiang Shao¹, Christos Sakalis², Stefanos Kaxiras² and Alberto Ros¹

¹Computer Engineering Department
University of Murcia, Spain

²Department of Information Technology
Uppsala University Sweden

Splash-2^[1]

- First major parallel benchmark suite
- Many works based on its behavior
- Still relevant and useful
- Quite old with outdated programming techniques and bugs

Splash-3^[2]

- Fixes many bugs of the previous version
- Focused in synchronization
- Not focused in performance only correctness

Splash-4

- Focused on atomic operations
- Better scalability in current hardware




[1] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The SPLASH-2 programs: Characterization and methodological considerations," in 22nd Int'l Symp. on Computer Architecture (ISCA), Jun. 1995, pp. 24–36.

[2] C. Sakalis, C. Leonardsson, S. Kaxiras, and A. Ros, "Splash-3: A properly synchronized benchmark suite for contemporary research," in Int'l Symp. on Performance Analysis of Systems and Software (ISPASS), Apr. 2016, pp. 101–111

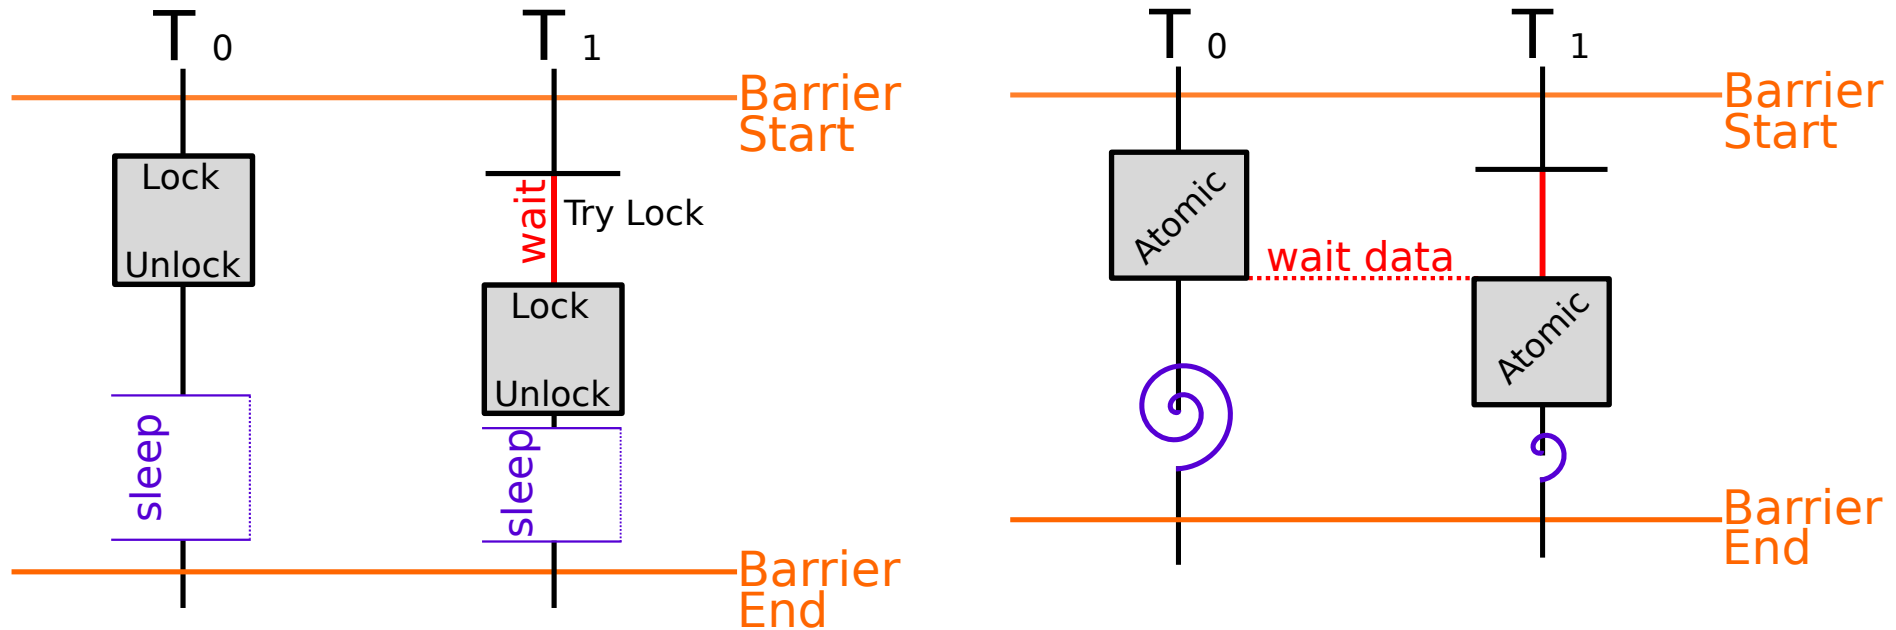
Locks (Mutexes)

Critical sections guarded by the same lock mutex, even if there is no data conflict, cannot be run in parallel, unless they are converted to atomics.




Barrier


Barriers are often implemented using mutexes and a thread sleep.
 When the time spent between barriers is high, this overhead is irrelevant.
 For contended barriers, a spinlock allow for a faster wakeup to continue the execution.



Results



Splash-3 vs Splash-4 Scalability on an **actual processor**



Splash-3 vs Splash-4 Scalability in **simulation**

Splash-4

Improving Scalability with Lock-Free Constructs

Eduardo José Gómez-Hernández¹, Ruixiang Shao¹, Christos Sakalis², Stefanos Kaxiras² and Alberto Ros¹

eduardojose.gomez@um.es

Thank you for your attention!




erc
European Research Council
Established by the European Commission

ECHO, ERC Consolidator Grant (No 819134)

This presentation and recording belong to the authors. No distribution is allowed without the authors' permission.

Barrier Groups

We define a barrier group as the set of critical sections that can be executed between two consecutive barriers (for all threads).



Barrier group 0

- CS 0
- CS 1
- CS 2
- CS 3

Critical sections 0,1,2,3 could be executed in parallel between two barriers.

This establish a relation that limits the parallelism

While&CAS

```
1 var oldValue;  
2 var newValue;  
3 do {  
4     oldValue = *(ptr);  
5     newValue = new;  
6 } while (!CAS(ptr, oldValue, newValue));
```

While&CAS structure

```
1 double oldValue;  
2 double newValue;  
3 do {  
4     oldValue = *ptr;  
5     newValue = oldValue + addition;  
6 } while (!CAS(ptr, oldValue, newValue));
```

FETCH_AND_ADD_DOUBLE operation

Atomic operations in modern processors are limited.
But using the while&cas structure is possible to craft custom “atomic constructs”[1].

In this example we propose the `FETCH_AND_ADD_DOUBLE` atomic, that allows to add 64-bits floating point numbers atomically.

[1] H. Gao and W. Hesselink, “A general lock-free algorithm using compare-and-swap,” *Information and Computation*, vol. 205, no. 2, pp. 225–241, 2007.

Sense Reversing Barrier


The barrier we used is called “sense reversing barrier”[1].

```

1 local_sense = !local_sense;
2 if (atomic_fetch_sub(&(count), 1) == 1) {
3     count = cores;
4     sense = local_sense;
5 } else {
6     do {} while (sense != local_sense);
7 }

```

Sense-reversing barrier



Ratio of time spent waiting on barriers, Original vs Sense-Reversing Centralized

[1] J. M. Mellor-Crummey and M. L. Scott, “Algorithms for scalable syn-chronization on shared-memory multiprocessors,”ACM Trans. Comput.Syst., vol. 9, no. 1, pp. 21–65, Feb. 1991.

Lock Split

In certain situations is possible to break the critical section into multiple ones without changing the result (breaking the group atomicity).

These examples are from the “water-nsquare” benchmark.

Group atomicity is not needed and neither is assumed anywhere in the code.

We surmise that in the original Splash-2 such clustering with the purpose of amortizing the high cost of the lock and unlock over many operations.

```

1  /* Lock */
2  LOCK(gl->PotengSumLock);
3  * POTA = * POTA + LPOTA;
4  * POTR = * POTR + LPOTR;
5  * PTRF = * PTRF + LPTRF;
6  UNLOCK(gl->PotengSumLock);

```

```

1  /* Lock-free */
2  FETCH_AND_ADD_DOUBLE(POTA, LPOTA);
3  FETCH_AND_ADD_DOUBLE(POTR, LPOTR);
4  FETCH_AND_ADD_DOUBLE(PTRF, LPTRF);

```

poteng.c.in 159 & poteng.c.in 253

```

1  /* Lock */
2  ALOCK(gl->MolLock, mol % MAXLCKS);
3  for ( dir = XDIR; dir <= ZDIR; dir++) {
4      temp_p = VAR[mol].F[DEST][dir];
5      temp_p[H1] += PFORCES[ProcID][mol][dir][H1];
6      temp_p[O] += PFORCES[ProcID][mol][dir][O];
7      temp_p[H2] += PFORCES[ProcID][mol][dir][H2];
8  }
9  AUNLOCK(gl->MolLock, mol % MAXLCKS);

```


```

1  /* Lock-free */
2  for ( dir = XDIR; dir <= ZDIR; dir++) {
3      FETCH_AND_ADD_DOUBLE(&(VAR[mol].F[DEST][dir][H1]), PFORCES[ProcID][mol][dir][H1]);
4      FETCH_AND_ADD_DOUBLE(&(VAR[mol].F[DEST][dir][O]), PFORCES[ProcID][mol][dir][O]);
5      FETCH_AND_ADD_DOUBLE(&(VAR[mol].F[DEST][dir][H2]), PFORCES[ProcID][mol][dir][H2]);
6  }

```

interf.c.in 156 & interf.c.in 167 & interf.c.in 179

Lock effects



Percent of time spent in critical sections out of total execution. The three bars per core count represent the original version, the straightforward C11 atomics, and the lock-free version respectively. The critical section time (CS) corresponds to the time spent in the critical section for the lock-unlock case, while for the C11 atomics and the lock-free constructs the original critical-section work is subsumed by the operations/constructs themselves.