

SPLASH-4: A MODERN BENCHMARK SUITE WITH LOCK-FREE CONSTRUCTS

Eduardo José Gómez-Hernández¹ Juan M. Cebrian¹
Stefanos Kaxiras² Alberto Ros¹

¹Computer Engineering Department
University of Murcia, Spain

²Department of Information Technology
Uppsala University, Sweden

- ▶ The **cornerstone for the performance evaluation** is the benchmark suite
- ▶ Benchmark suites can **misrepresent** the performance characteristics
- ▶ **Keeping up** with architectural changes while maintaining the same workloads is a real challenge

- ▶ The **cornerstone for the performance evaluation** is the benchmark suite
- ▶ Benchmark suites can **misrepresent** the performance characteristics
- ▶ **Keeping up** with architectural changes while maintaining the same workloads is a real challenge
- ▶ We introduce **Splash-4**, a **revised version** of Splash-3 (an update on Splash-2), that introduces modern programming techniques to improve scalability on contemporary hardware

MOTIVATION



¹Woo, Steven Cameron, et al, "The SPLASH-2 programs: Characterization and methodological considerations." ACM SIGARCH computer architecture news 23, 1995

²Venetis, Ioannis E., et al, "The Modified SPLASH-2", <https://www.capsl.udel.edu//splash/> 2007

³Sakalis, Christos, et al, "Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research", ISPASS 2016

⁴Gómez-Hernández, Eduardo José et al, "Splash-4: Improving Scalability with Lock-Free Constructs", ISPASS 2021

MOTIVATION



The first major parallel benchmark suite. Still in use (+5k cites)¹

¹Woo, Steven Cameron, et al, "The SPLASH-2 programs: Characterization and methodological considerations." ACM SIGARCH computer architecture news 23, 1995

²Venetis, Ioannis E., et al, "The Modified SPLASH-2", <https://www.capsl.udel.edu//splash/> 2007

³Sakalis, Christos, et al, "Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research", ISPASS 2016

⁴Gómez-Hernández, Eduardo José et al, "Splash-4: Improving Scalability with Lock-Free Constructs", ISPASS 2021

MOTIVATION



The first major parallel benchmark suite. Still in use (+5k cites)¹

A small update that fixes bugs and updates the programming style²

¹Woo, Steven Cameron, et al, "The SPLASH-2 programs: Characterization and methodological considerations." ACM SIGARCH computer architecture news 23, 1995

²Venetis, Ioannis E., et al, "The Modified SPLASH-2", <https://www.capsl.udel.edu//splash/> 2007

³Sakalis, Christos, et al, "Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research", ISPASS 2016

⁴Gómez-Hernández, Eduardo José et al, "Splash-4: Improving Scalability with Lock-Free Constructs", ISPASS 2021

MOTIVATION



The first major parallel benchmark suite. Still in use (+5k cites)¹

A small update that fixes bugs and updates the programming style²

First major update that fixes data races and performance bugs³

¹Woo, Steven Cameron, et al, "The SPLASH-2 programs: Characterization and methodological considerations." ACM SIGARCH computer architecture news 23, 1995

²Venetis, Ioannis E., et al, "The Modified SPLASH-2", <https://www.capsl.udel.edu//splash/> 2007

³Sakalis, Christos, et al, "Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research", ISPASS 2016

⁴Gómez-Hernández, Eduardo José et al, "Splash-4: Improving Scalability with Lock-Free Constructs", ISPASS 2021

MOTIVATION



The first major parallel benchmark suite. Still in use (+5k cites)¹

A small update that fixes bugs and updates the programming style²

First major update that fixes data races and performance bugs³

Current update, exploiting lockfree and atomic operations⁴

¹Woo, Steven Cameron, et al, "The SPLASH-2 programs: Characterization and methodological considerations." ACM SIGARCH computer architecture news 23, 1995


²Venetis, Ioannis E., et al, "The Modified SPLASH-2", <https://www.capsl.udel.edu//splash/> 2007

³Sakalis, Christos, et al, "Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research", ISPASS 2016

⁴Gómez-Hernández, Eduardo José et al, "Splash-4: Improving Scalability with Lock-Free Constructs", ISPASS 2021

MOTIVATION

- ▶ Splash-2 and Splash-3 are crafted using outdated programming techniques
- ▶ Previous works noticed that the **default input sizes** limit the scalability of some applications.
 - ▶ The computation between synchronization points is not substantially longer than the synchronization
 - ▶ Using larger datasets increases the execution time, and that is a problem when using simulation infrastructures



- ▶ Splash-3 exposes **data races** and **performance bugs**, and fixed them
- ▶ In the analysis done using GEMS (in a 64-core **in-order** multicore), benchmarks reach an speedup between $16\times$ to $47\times$
- ▶ In our simulated environment (gem5-20 Intel's IceLake-like **64 out of order** cores): the scalability stops between 16 and 32 cores, with an average speedup of $2.3\times$
- ▶ In our real hardware (64-Core AMD EPYC 7702P), scalability stops between 4 and 16 cores, with an average speedup of $4.7\times$

- ▶ Synchronization is done with combination of software locks, conditional variables, and barriers
- ▶ The main idea is to **replace** high-overhead synchronization operations with newer lightweight alternatives
- ▶ This translates into a performance improvement by expanding the architectural features that the benchmarks can exercise

SPLASH-4: LOCKFREE AND ATOMICS

- ▶ Modern ISAs typically provide **a basic set of atomic operations** that offer both atomicity and synchronization
- ▶ This basic set consists of atomic loads and stores, atomic read-modify-write (RMW) operations, and some atomic comparisons and exchange operations
- ▶ Typical hardware RMW atomics are only available for **integer types**

Splash-3

```
1 /* Lock */
2 ALOCK(Global->QLock, local_node);
3 work = Global->Queue[local_node
    ][0];
4 Global->Queue[local_node][0] +=
    1;
5 AULOCK(Global->QLock, local_node);
```

Splash-4

```
1 /* Lock-free */
2 work = FETCH_ADD(Global->Queue[
    local_node][0], 1);
```

ALOCK/AULOCK → mutex lock/unlock

FETCH_ADD → atomic_fetch_and_add - Sequential Consistency

SPLASH-4: LOCKFREE AND ATOMICS

WHILE&CAEXCH CONSTRUCT

- ▶ Atomic Compare-and-Swap (CAS) and Atomic Compare-and-Exchange (CAExch) are **type-agnostic**
- ▶ They can be used to implement RMW operations for more complex underlying types

```
1| CAExch(ptr, oldValue, newValue);
```

- ▶ If *oldValue* == (*ptr) then (*ptr) = *newValue*
- ▶ Else *oldValue* = (*ptr)

LOAD → Atomic Load

SPLASH-4: LOCKFREE AND ATOMICS

WHILE&CAEXCH CONSTRUCT

- ▶ Atomic Compare-and-Swap (CAS) and Atomic Compare-and-Exchange (CAExch) are **type-agnostic**
- ▶ They can be used to implement RMW operations for more complex underlying types

```
1| CAExch(ptr, oldValue, newValue);
```

- ▶ If $oldValue == (*ptr)$ then $(*ptr) = newValue$

- ▶ Else $oldValue = (*ptr)$

```
1| var oldValue = LOAD(ptr);
```

```
2| var newValue;
```

```
3| do {
```

```
4|     newValue = new;
```

```
5| } while (!CAExch(ptr, oldValue, newValue));
```

LOAD → Atomic Load

Splash-3

```
1 /* Lock */
2 LOCK(locks->error_lock)
3 if (local_err > multi->err_multi)
4     {
5     multi->err_multi = local_err;
6 }
7 UNLOCK(locks->error_lock)
```

Splash-4

```
1 /* Lock-free */
2 double expected = LOAD(multi->
3     err_multi);
4 do {
5     if (local_err <= expected)
6         break;
7 } while (!CAExch(multi->err_multi
8     , expected , local_err));
```

LOCK/UNLOCK → mutex lock/unlock

LOAD → Atomic Load - Sequential Consistency

SPLASH-4: CRITICAL SECTION SPLITTING

- ▶ Splash-4 uses lock-free constructs that manage a **single address** and replace to critical sections that modify a single address
- ▶ Splitting a larger critical section that modifies **more than one address** would enable the use of lock-free constructs in more cases
- ▶ Unfortunately, splitting large critical sections, is not possible in the general case
- ▶ Many Splash-3 critical sections, (atomic) updates of **independent variables** are **grouped together** in larger critical sections for **no apparent reason**

WATER (POTENG.C.IN 159/253)

► Atomic Double Floating Point Addition (FETCH_AND_ADD_DOUBLE)

```

1 | double oldValue = LOAD(ptr);
2 | double newValue;
3 | do {
4 |   newValue = oldValue + addition;
5 | } while (!CAExch(ptr, oldValue, newValue));

```

LOCK/UNLOCK → mutex lock/unlock

LOAD → Atomic Load - Sequential Consistency

WATER (POTENG.C.IN 159/253)

► Atomic Double Floating Point Addition (FETCH_AND_ADD_DOUBLE)

```

1 | double oldValue = LOAD(ptr);
2 | double newValue;
3 | do {
4 |   newValue = oldValue + addition;
5 | } while (!CAExch(ptr, oldValue, newValue));
  
```

Splash-3

```

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

LOCK/UNLOCK → mutex lock/unlock

LOAD → Atomic Load - Sequential Consistency

Splash-4

```

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

SPLASH-4: SENSE-REVERSING CENTRALIZED BARRIER

- ▶ The execution time between barriers is fairly short
- ▶ To minimize the overhead, we implement a sense-reversing barrier
- ▶ This construct is optimized for short waiting times, except when oversubscribing threads

LOAD → Atomic Load - Sequential Consistency
STORE → Atomic Store - Sequential Consistency

SPLASH-4: SENSE-REVERSING CENTRALIZED BARRIER

- ▶ The execution time between barriers is fairly short
- ▶ To minimize the overhead, we implement a sense-reversing barrier
- ▶ This construct is optimized for short waiting times, except when oversubscribing threads

```

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

```

LOAD → Atomic Load - Sequential Consistency

STORE → Atomic Store - Sequential Consistency

SPLASH-4: SUMMARY

- ▶ A total of 50 critical sections (33%) has been transformed.
 - ▶ 27 critical sections (18%) were converted into C11 atomics
 - ▶ 18 critical sections (12%) were converted into CAS-construct (without splitting)
 - ▶ Splitting Critical sections allowed 5 more critical sections (3%) to be converted into 15 CAS-construct


EVALUATION: METHODOLOGY

- ▶ Measurements account for the **region of interest** (ROI)
- ▶ We used the **recommended inputs from Splash-3** for all the executions

- ▶ The **hardware** used is an AMD Epyc 7702P CPU


- ▶ The **simulated** machine (in gem5-20) is mimicking an Intel's IceLake processor
 - ▶ Measurements reset stats after a warm-up period (as suggested by the original Splash-2 developers)

EVALUATION: SIMULATOR OVERHEAD BREAKDOWN




- ▶ Most of the cases the execution time is dominated by the barriers
- ▶ Raytrace gets a huge overhead reduction due to the lock-free constructs

EVALUATION: TOP-DOWN (OCEAN-CONT)




- ▶ Top-Down groups several hardware counters to understand easier the reason of the stalling
- ▶ In Splash-4 we observe an increase on the back-end pressure

EVALUATION: SIMULATOR SCALABILITY




► In general, scalability improvement is moved from 16 to 32

EVALUATION: REAL MACHINE SCALABILITY



► In general, scalability improvement is moved from 4 to 16

EVALUATION: REAL MACHINE 64 THREADS SUMMARY



- ▶ With the lock-free constructs the execution time is reduced 11%
- ▶ With the new barriers the execution time is reduced 40%
- ▶ The combined reduction is 52%

CONCLUSIONS

- ▶ We presented **The Splash-4 Benchmark Suite**
- ▶ We performed a detailed performance analysis comparing with Splash-3

CONCLUSIONS

- ▶ We presented **The Splash-4 Benchmark Suite**
- ▶ We performed a detailed performance analysis comparing with Splash-3
- ▶ Our study shows:
 - ▶ A significant improvement on scalability
 - ▶ Execution time is reduced significantly on the simulated and real hardware

CONCLUSIONS

- ▶ We presented **The Splash-4 Benchmark Suite**
- ▶ We performed a detailed performance analysis comparing with Splash-3
- ▶ Our study shows:
 - ▶ A significant improvement on scalability
 - ▶ Execution time is reduced significantly on the simulated and real hardware
- ▶ In summary:
 - ▶ The Splash benchmark suite is still a cornerstone in computer architecture research
 - ▶ It should be updated to be able to exploit modern hardware features

CONCLUSIONS

- ▶ We presented **The Splash-4 Benchmark Suite**
- ▶ We performed a detailed performance analysis comparing with Splash-3
- ▶ Our study shows:
 - ▶ A significant improvement on scalability
 - ▶ Execution time is reduced significantly on the simulated and real hardware
- ▶ In summary:
 - ▶ The Splash benchmark suite is still a cornerstone in computer architecture research
 - ▶ It should be updated to be able to exploit modern hardware features

Splash-4 removes some of the synchronization overheads, showing the hidden reasons that prevents applications from scaling, leaving a door open for researchers to further improve their designs

SPLASH-4: A MODERN BENCHMARK SUITE WITH LOCK-FREE CONSTRUCTS

Eduardo José Gómez-Hernández¹ Juan M. Cebrian¹
Stefanos Kaxiras² Alberto Ros¹

eduardojose.gomez@um.es

Thank you for your attention!

ECHO, ERC Consolidator Grant (No 819134)
Spanish Ministerio de Economía, Industria y Competitividad – Agencia Estatal de
Investigación (ERC2018-092826)

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

INDEPENDENT VARIABLES

- ▶ During the possible window of execution time of the critical section, all variables cannot be related on any way, even not giving information from one to another

```

1 | LOCK(lock);
2 | ptr = NULL;
3 | some_random_boolean = false;
4 | UNLOCK(lock);

```

INDEPENDENT VARIABLES

- ▶ During the possible window of execution time of the critical section, all variables cannot be related on any way, even not giving information from one to another

```

1 | LOCK(lock);
2 | ptr = NULL;
3 | some_random_boolean = false;
4 | UNLOCK(lock);
  
```

```

1 | LOCK(lock);
2 | if (some_random_boolean) {
3 |     local = *ptr;
4 | }
5 | UNLOCK(lock);
  
```

OVERHEAD SUMMARY

