

Splash-4 released, optimizing the Splash parallel benchmark suite

Eduardo José Gómez Hernández, Juan Manuel Cebrian, Alberto Ros (all University of Murcia) and Stefanos Kaxiras (Uppsala University)

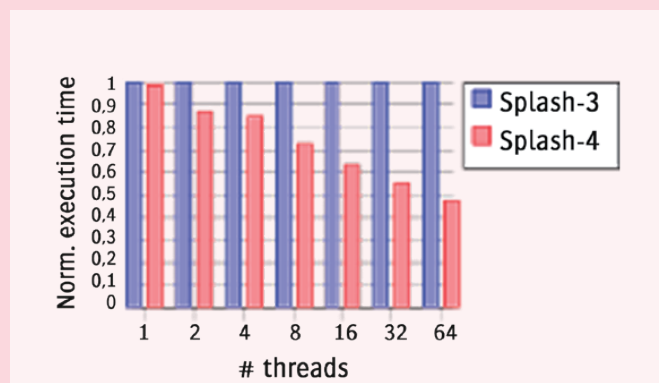


The CAPS research group at the University of Murcia, in collaboration with the UART group at Uppsala University, has released the Splash-4 benchmark suite. Splash-4 comes with an impressive performance uplift and improves scalability with respect to its predecessors Splash-3 and Splash-2.

Its first appearance was as a poster at the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 2021). A major update removing data-races was released at the 25th anniversary of the IEEE International Symposium on Workload Characterization (IISWC 2022).

The Splash-2 benchmark suite, among the many benchmark suites used, is a cornerstone in advancing knowledge for both academia and industry, with over 5,276 citations since it was published in 1995. Its successor, the Splash-3 benchmark suite, has gained over 100 citations since its appearance in 2016.

New instructions are constantly being added to hardware, and keeping up with those new features is a real challenge, especially while trying to maintain the same workloads and algorithms (for comparative purposes). In this environment, benchmark suites can misrepresent the performance characteristics of a system. They are not able to reflect the available features of the hardware, causing computer architects to overestimate or underestimate the impact of proposed techniques.



Normalized execution time of Splash-3 vs Splash-4 on AMD's EPYC 7002 for different thread counts

Splash-4 introduces modern programming techniques to improve scalability on contemporary hardware. On AMD's EPYC 7002P 64-core processor, Splash-4 reduces the normalized execution time by an average of 52% (see graph, below). The algorithms remain the same, therefore the revised applications with upgraded synchronization primitives maintain the same computational patterns as their previous versions.

This new release of the Splash benchmark suite removes a large chunk of the synchronization overhead, revealing previously hidden bottlenecks on the cores and opening the door for researchers to further improve their designs.

The Splash-4 release is the latest in a series of collaborations between the University of Murcia and Uppsala University. In addition to the Splash benchmark, previous joint research includes work on speculative execution security and cache coherence.

FURTHER INFORMATION:

Splash-4 Source Code github.com/Odnetninl/Splash-4

Eduardo José Gómez-Hernández, Rubén Titos-Gil, Juan Manuel Cebrian, Stefanos Kaxiras, Alberto Ros, 'Efficient, Distributed, and Non-Speculative Multi-Address Atomic Operations'. 54th International Symposium on Microarchitecture (MICRO), pp. 337-349, Worldwide event, October 2021

dl.acm.org/doi/10.1145/3466752.3480073

Eduardo José Gómez-Hernández, Juan Manuel Cebrian, Stefanos Kaxiras, Alberto Ros, 'Splash-4: A Modern Benchmark Suite with Lock-Free Constructs'. 2022 IEEE International Symposium on Workload Characterization (IISWC), Austin, TX (USA), November 2022

ieeexplore.ieee.org/document/9975421

S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, 'The SPLASH-2 programs: Characterization and methodological considerations'. 22nd International Symposium on Computer Architecture (ISCA), pp. 24-36, Santa Margherita Ligure (Italy), June 1995

ieeexplore.ieee.org/document/524546

Christos Sakalis, Carl Leonardsson, Stefanos Kaxiras, Alberto Ros, 'Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research'. International Symposium on Performance Analysis of Systems and Software (ISPASS), pp. 101-111, Uppsala (Sweden), April 2016 ieeexplore.ieee.org/document/7482078