TECHILA

TECHILA TECHNOLOGIES SCHEDULER BENCHMARK STUDY

TDCE VS IBM SPECTRUM SYMPHONY

24 August 2021

Executive Summary

Techila Distributed Computing Engine (TDCE) is a highly efficient middleware system providing excellent APIs and a world class scheduler that ensures maximal throughput in even the most demanding high-performance computing (HPC) situations.

This document presents new TDCE scheduler performance statistics and compares them with IBM Spectrum Symphony performance tests previously conducted and published by Microsoft [1]. As can be seen from test results, **TDCE is able to ensure a system utilization of over 95% in all tests and is able to provide higher, more consistent throughput than IBM Symphony.**

Introduction

The original study [1] showcased IBM Spectrum Symphony performance statistics using a computational capacity of 9920 CPU cores when processing different workloads. The smallest workload consisted of 100 x 10 second tasks and the largest workload consisted of 992,000 x 10 second tasks.

To compare the performance of the TDCE scheduler with IBM Symphony, Techila Technologies used Techila Distributed Computing Engine to process similar workloads using a matching compute capacity of 9920 CPU cores.

The results of these performance tests will be discussed in the following chapter. Performance metrics can be found in tabular format in the Appendix.

Comparing earlier study results and new TDCE performance data

Figure 1 shown below shows the benchmark results of the new TDCE scheduler tests superimposed over the original IBM Spectrum Symphony performance data reported in [1]. The figure also includes a line indicating the theoretical maximum throughput of the system for each task count. This theoretical maximum value represents how many 10 second tasks the system is theoretically able to process in 1 hour, assuming zero overheads.

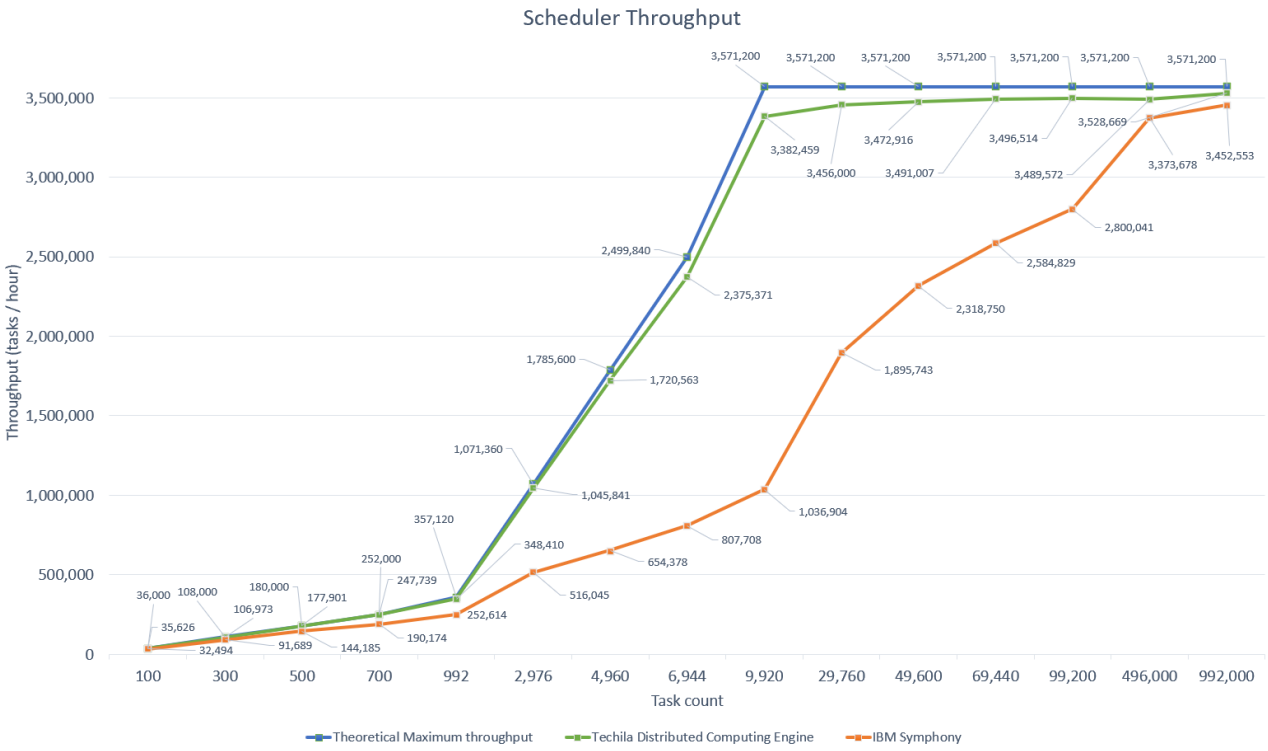

Figure 1. System throughput – TDCE closely follows the theoretical maximum throughput with all workloads while IBM Spectrum Symphony falls behind, providing only under 30% of TDCEs respective throughput in the most impacted workloads.

As can be seen from the figure above, TDCE throughput closely follows the theoretical maximum throughput line, meaning the system is able to provide a reliable and consistent turnaround time for individual tasks in all test cases. As soon as the amount of computational workload is sufficient to fully utilize all available cores, i.e. the task count is over 9920, TDCE reaches a throughput of over 3.3 million tasks per hour, **providing over 3 times the throughput of IBM Spectrum Symphony** for this specific workload. As the workload size is increased, TDCE throughput stabilizes at approximately **3.5 million tasks per hour**, providing consistently better throughput than IBM Spectrum Symphony for all workloads.

The consistent performance of the TDCE scheduler ensures that the system can provide excellent, consistent performance in all situations, which can be seen in the system utilization rates in Figure 2.

These system utilization rates have been calculated using the following formula.

$$System\ utilization = \frac{ideal\ execution\ time}{actual\ execution\ time} \cdot 100$$

More details about the performance statistics, including the ideal and actual execution times can be found in Table 1 in the Appendix.

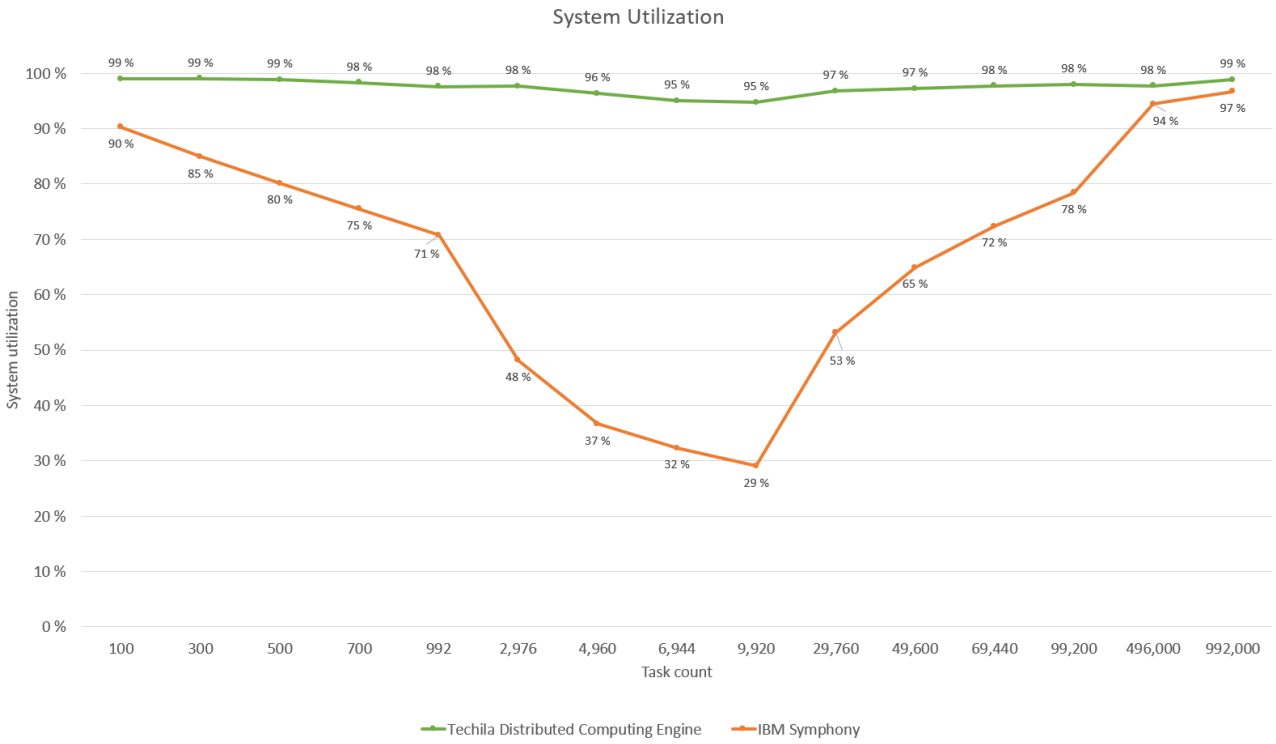

Figure 2. System utilization – TDCE ensures a consistent system utilization of 95% while the respective rates drop to as low as 29% with IBM Spectrum Symphony.

Disclaimer and parting words

TDCE benchmark was done by Techila Technologies in 2021. IBM was not involved in these tests in any way. All IBM Spectrum Symphony performance statistics were obtained from the previous study published by Microsoft [1].

If you are looking to improve the utilization and throughput of your computing environment, **please feel free to contact sales@techilatechnologies.com for additional information about the benchmark methods used or to schedule a demo.**

References and further material

[1] https://azure.microsoft.com/es-es/blog/ibm-and-microsoft-azure-support-spectrum-symphony-and-spectrum-lsf/

More information about Techila Distributed Computing Engine can be found in the TDCE product description. Additional details about the scheduler functionality can be found in the TDCE scheduler document.

Appendix

System			IBM Symphony	IBM Symphony	IBM Symphony	TDCE	TDCE	TDCE
Metric	Theoretical execution time	Theoretical maximum throughput	Throughput	Execution time	System utilization	Throughput	Execution time	System utilization
Task count	seconds	tasks / h	tasks/h	seconds	%	task/h	seconds	%
100	10	36,000	32,494	11.079	90.261 %	35,626	10.105	98.961 %
300	10	108,000	91,689	11.779	84.897 %	106,973	10.096	99.049 %
500	10	180,000	144,185	12.484	80.103 %	177,901	10.118	98.834 %
700	10	252,000	190,174	13.251	75.466 %	247,739	10.172	98.309 %
992	10	357,120	252,614	14.137	70.736 %	348,410	10.250	97.561 %
2,976	10	1,071,360	516,045	20.761	48.167 %	1,045,841	10.244	97.618 %
4,960	10	1,785,600	654,378	27.287	36.648 %	1,720,563	10.378	96.358 %
6,944	10	2,499,840	807,708	30.950	32.310 %	2,375,371	10.524	95.021 %
9,920	10	3,571,200	1,036,904	34.441	29.035 %	3,382,459	10.558	94.715 %
29,760	30	3,571,200	1,895,743	56.514	53.084 %	3,456,000	31.000	96.774 %
49,600	50	3,571,200	2,318,750	77.007	64.929 %	3,472,916	51.415	97.248 %
69,440	70	3,571,200	2,584,829	96.712	72.380 %	3,491,007	71.608	97.754 %
99,200	100	3,571,200	2,800,041	127.541	78.406 %	3,496,514	102.136	97.909 %
496,000	500	3,571,200	3,373,678	529.274	94.469 %	3,489,572	511.696	97.714 %
992,000	1000	3,571,200	3,452,553	1034.365	96.678 %	3,528,669	1012.053	98.809 %

Table 1. Performance statistics. IBM Spectrum Symphony statistics were derived from the performance statistics reported in [1].