

Errata for “Teseo and the Analysis of Structural Dynamic Graphs” (PVLDB 14(6): 1053 - 1066)

Dean De Leo
CWI
dleo@cwi.nl

Per Fuchs
TUM
per.fuchs@cs.tum.edu

Peter Boncz
CWI
boncz@cwi.nl

PVLDB Reference Format:

Dean De Leo, Per Fuchs, and Peter Boncz. Errata for
“Teseo and the Analysis of Structural Dynamic Graphs”
(PVLDB 14(6): 1053 - 1066). PVLDB, 14(11): 2271 - 2272, 2021.
doi:10.14778/3476249.3476278

In our paper [4], we experimentally evaluated our work, Teseo, together with five other systems under the LDBC Graphalytics benchmark [6]. We developed and publicly released [2] an ad-hoc driver for the purpose. Since the time the paper was published, a bug [1] in the driver has been found. Due to this bug, we discovered that the completion times for Graphalytics have been incorrectly measured by 1 second or slightly more than their actual values. This issue involves the results of Table 2, Figure 8 and Table 3 reported in our paper [4]. Still, because the bug equally affected all systems evaluated and it is only related to the measurements, most of the comparisons and the general conclusions in the paper still hold.

We also want to take advantage of this errata to further integrate some additional feedback, related to the driver, that we received after the publication of the paper and that we believe it can provide substantial value to the presented results. In detail:

- (1) We correct the error [1] in the measurements of our Graphalytics driver [2]. As effect, all the absolute results of Table 3 in [4] become lower of about 1 second.
- (2) We replace the output of the experiments from a hash table to a vector. The output of all algorithms in Graphalytics is a score associated to each vertex in the graph. With a hash table, all systems took several seconds to create it in Scale Factor SF 26. In some cases, this cost could be higher than the actual execution of a kernel. By replacing the hash table with a vector, the output is always created in the order of milliseconds for the CSR. This change allows to report a completion time of less than a second for the BFS execution of the CSR with SF 26.
- (3) We stabilise the results for the CSR in the CDLP kernel for SF 22 and SF 24 of Table 3 in [4], by repeating the experiments a higher number of times. The previous results were caused by “bad runs”, where the algorithm converged more slowly due to non deterministic choices made by the OpenMP scheduler.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 14, No. 11 ISSN 2150-8097.
doi:10.14778/3476249.3476278

We note that the changes 1-2 only involve the driver [3] and equally affect all systems evaluated. We request to amend Table 2, Figure 8 and Table 3 of [4] with the new measurements, attached to this errata. It follows:

- Table 2 shows the speed-up of our CSR implementation over GraphMat and GraphBlas. Compared to the results in paper [4], due to the improvements to our driver, the CSR becomes always faster than the other systems, based on their native external drivers, and the speed-up differences are larger.
- Figure 8 shows the speed-up of the kernel implementations used in our experiments over those shipped by the library authors. Compared to the same Figure 8 in [4], the speed-up differences are more pronounced, particularly for the BFS, as the completion times for the kernels are lower.
- Table 3 shows the speed-up of the CSR implementation over the systems evaluated. Compared to Table 3 in [4], the completion times of the CSR are generally lower. As consequence, the speed-up differences are more pronounced for the SF 22, SF 24 and the BFS kernel. There are only limited differences for SF 26, as the overall completion times were already dominated by the kernel executions. The speed-up of the CSR over GraphOne can now be significant for the faster kernels, due to the time spent by this system to set up a *static view*. The speed-up of the CSR over LiveGraph for the BFS kernel can also be particularly high, due to the preprocessing time required to compute the degrees of the vertices¹. These overheads were previously hidden in the measurements of [4]. The results for the CDLP kernel for the CSR are always better or comparable to the other systems.

REFERENCES

- [1] 2021. Bug fix for the Graphalytics driver. https://github.com/cwida/gfe_driver/commit/7f2b79830ad55a428ae3667d4117f753ed0cb341
- [2] 2021. Graphalytics driver public repository. https://github.com/cwida/gfe_driver
- [3] 2021. Proposed changes to the Graphalytics driver. https://github.com/cwida/gfe_driver/tree/feature/vector_materialization
- [4] Dean De Leo and Peter Boncz. 2021. Teseo and the Analysis of Structural Dynamic Graphs. *Proc. VLDB Endow.* 14, 6 (Feb. 2021), 1053–1066. <https://doi.org/10.14778/3447689.3447708>
- [5] Bálint Hegyi and Gábor Szárnyas. 2020. SuiteSparse:GraphBLAS driver for LDBC Graphalytics. <https://github.com/ftsrg/ldbc-graphalytics-platforms-graphblas>
- [6] Alexandru Iosup, Tim Hegeman, Wing Lung Ngai, Stijn Heldens, Arnau Prat-Pérez, Thomas Manhardt, Hassan Chafio, Mihai Capotă, Narayanan Sundaram, Michael Anderson, Ilie Gabriel Tănase, Yinglong Xia, Lifeng Nai, and Peter Boncz. 2016. LDBC Graphalytics: A Benchmark for Large-scale Graph Analysis on Parallel

¹In LiveGraph, the API does not provide the degree of a vertex and it needs to be computed explicitly.

Graph	System	BFS	CDLP	LCC	LCC (opt)	PageRank	SSSP	WCC	Graph	System	BFS	CDLP	LCC	LCC (opt)	PageRank	SSSP	WCC
DOTA League	CSR (baseline)	< 1ms	1.48s	472s	69s	0.16s	0.37s	0.04s	Graph500 SF 22	CSR (baseline)	0.03s	4.25s	436s	18.05s	0.43s	1.26s	0.19s
	Stinger	127.21x	1.43x	3.16x	N/A	6.22x	3.32x	5.78x		Stinger	4.86x	1.30x	2.97x	N/A	4.01x	1.92x	4.26x
	LLAMA	2.31x	2.67x	DNF	N/A	1.05x	5.99x	1.30x		LLAMA	1.44x	2.51x	DNF	N/A	1.44x	4.00x	1.07x
	GraphOne	4386.22x	3.66x	1.29x	N/A	25.50x	11.72x	135.48x		GraphOne	112.18x	2.91x	1.34x	N/A	11.55x	4.91x	65.52x
	LiveGraph	84.57x	1.20x	1.43x	N/A	3.02x	2.09x	2.59x		LiveGraph	17.85x	2.16x	1.46x	N/A	4.75x	2.04x	2.51x
	Teseo, log. vtx	11.05x	1.54x	1.60x	1.49x	2.59x	2.23x	2.76x		Teseo, log. vtx	1.58x	1.76x	2.20x	3.72x	4.29x	2.98x	4.39x
	Teseo, real vtx	9.24x	0.99x	1.08x	1.27x	1.14x	1.38x	0.98x		Teseo, real vtx	1.87x	1.13x	1.18x	1.90x	1.91x	1.66x	1.56x
Uniform SF 24	CSR (baseline)	0.15s	34s	32s	4.88s	4.08s	8.26s	1.45s	Graph500 SF 24	CSR (baseline)	0.14s	21s	3195s	105s	2.63s	6.71s	1.04s
	Stinger	2.23x	1.18x	2.36x	N/A	2.49x	1.89x	1.99x		Stinger	4.70x	1.20x	DNF	N/A	3.17x	1.78x	3.49x
	LLAMA	1.63x	2.19x	DNF	N/A	2.42x	8.47x	1.51x		LLAMA	1.41x	1.98x	DNF	N/A	1.68x	5.25x	1.18x
	GraphOne	26.81x	2.03x	3.13x	N/A	3.33x	3.15x	61.32x		GraphOne	31.68x	2.00x	DNF	N/A	4.05x	2.79x	57.49x
	LiveGraph	6.33x	1.90x	1.65x	N/A	3.07x	2.99x	2.07x		LiveGraph	19.40x	1.77x	DNF	N/A	3.74x	1.94x	2.07x
	Teseo, log. vtx	1.10x	1.69x	3.18x	3.60x	3.20x	3.46x	2.95x		Teseo, log. vtx	1.52x	1.77x	DNF	6.14x	3.96x	3.23x	4.46x
	Teseo, real vtx	0.76x	1.05x	1.38x	1.89x	1.53x	2.11x	1.48x		Teseo, real vtx	1.68x	1.04x	DNF	1.93x	1.86x	1.98x	1.59x
Uniform SF 26	CSR (baseline)	0.55s	170s	137s	20s	40s	6.87s		Graph500 SF 26	CSR (baseline)	0.50s	108s	DNF	683s	15.05s	36s	5.19s
	Stinger	4.24x	1.16x	2.46x	N/A	2.15x	1.81x	1.82x		Stinger	2.55x	1.13x	DNF	N/A	2.49x	1.63x	2.91x
	GraphOne	12.24x	1.68x	3.33x	N/A	2.30x	2.76x	59.94x		GraphOne	10.26x	1.66x	DNF	N/A	2.64x	2.22x	62.01x
	LiveGraph	7.19x	1.59x	1.64x	N/A	2.75x	3.16x	1.95x		LiveGraph	21.72x	1.53x	DNF	N/A	3.14x	2.02x	1.89x
	Teseo, log. vtx	1.30x	1.62x	3.46x	4.22x	2.93x	3.48x	2.01x		Teseo, log. vtx	1.34x	1.71x	DNF	DNF	3.45x	2.97x	4.24x
	Teseo, real vtx	0.89x	1.02x	1.39x	2.03x	1.39x	2.16x	1.46x		Teseo, real vtx	1.03x	1.00x	DNF	2.08x	1.62x	1.76x	1.49x

Table 3: Results measured in the Graphalytics benchmark. For the CSR, we report the absolute completion time. For all the other systems, for ease of comparison, we show the speed-up of the CSR over the given system. The best result, outside the CSR, is marked in bold. The results refer to the median out of 15 runs.

System	Graph	BFS	CDLP	LCC	PageRank	SSSP	WCC
GraphMat	graph500-24	94.55x	40.15x	DNF	49.39x	7.23x	177.13x
	uniform-24	268x	74.12x	198x	87.06x	10.74x	179.12x
SuiteSparse	graph500-24	13.82x	3.40x	3.04x	10.44x	2.06x	4.15x
	uniform-24	7.92x	2.19x	4.62x	10.64x	2.87x	5.22x

Table 2: Speed-up achieved by our CSR baseline over GraphMat [8] and SuiteSparse:GraphBLAS [5]. The results refer to the processing time [7], that is, the completion time without the preprocessing and loading time, from the median out of 5 runs.

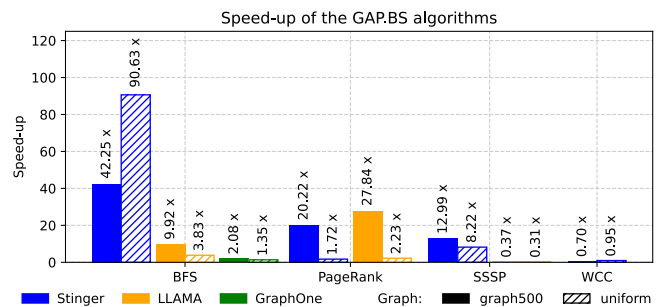


Figure 8: Speed-up of the graph algorithms from the GAP BS over the same algorithms provided by the library authors, for the graph500 and uniform graphs at SF 24.

and Distributed Platforms. *Proc. VLDB Endow.* 9, 13 (Sept. 2016), 1317–1328. <https://doi.org/10.14778/3007263.3007270>

- Alexandru Iosup, Tim Hegeman, Wing Lung Ngai, Stijn Heldens, Arnau Prat-Pérez, Thomas Manhardt, Hassan Chafiq, Mihai Capotă, Narayanan Sundaram, Michael Anderson, Ilie Gabriel Tănase, Yinglong Xia, Lifeng Nai, and Peter Boncz. 2017. LDBC Graphalytics Benchmark specification, v0.9.0. http://github.com/ldbc/ldbc_graphalytics_docs
- Wing Lung Ngai, Stijn Heldens, Mihai Capotă, Ahmed Musaafer, Tim Hegeman, Narayanan Sundaram, and Michael J. Anderson. 2017. Graphalytics GraphMat platform driver. <https://github.com/atlarge-research/graphalytics-platforms-graphmat>