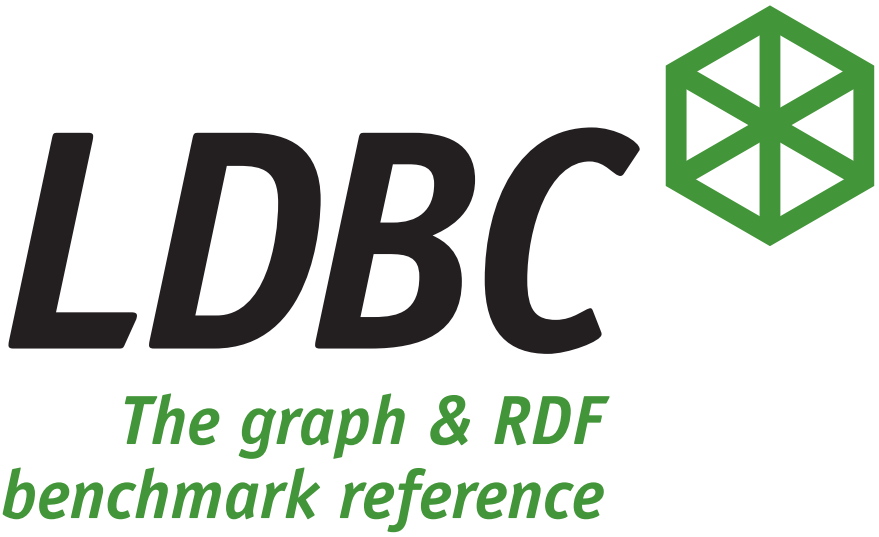

Full Disclosure Report of the LDBC Social Network Benchmark

An Implementation of the LDBC Social Network Benchmark's Interactive Workload over TuGraph

August 16, 2022

General Terms

Executive Summary

TuGraph (formerly known as LightGraph) is a proprietary graph database product developed by Ant Group. TuGraph was previously audited in July 2020 and developed by FMA before Ant Group acquired FMA. This document describes an implementation of the LDBC Social Network Benchmark's Interactive workload in TuGraph. The implementation uses stored procedures written using C++14, which are compiled and loaded into the database as plugins. The benchmark implementation uses imperative queries with manually defined query evaluation plans over the data to compute the queries specified in the workload. The data schema follows the property graph model with indices over node and edge identifiers and over properties selected by the user. TuGraph supports three precomputed properties to improve the operation throughput of the system, which are maintained during each update. The system under test and the driver communicates using remote procedure calls (RPC) over local sockets.

Declaration of Audit Success

This report contains an audited LDBC benchmark run. The results have been gathered by an independent and impartial auditor who has validated the implementation of the queries, successfully run the ACID tests associated with the claimed isolation level (serializable), and verified the overall system's configuration conformance to the description of the benchmark and its strict requirements.

DocuSigned by:
David Püroja
DD2A752E1D61426...

Mr. David Püroja
(Auditor)

8/16/2022
Date

DocuSigned by:
Gábor Szárnyas
60A78D85058140A...

Dr. Gábor Szárnyas
(Head of LDBC SNB Task Force)

8/16/2022
Date

DocuSigned by:
Chuntao Hong
3FE773CDAEA8496...

Dr. Chuntao Hong
(Test Sponsor Representative)

8/16/2022
Date

Table of Contents

1 System Description and Pricing Summary

1.1 Details of machines driving and running the workload

1.1.1 Machine overview

The details below were obtained from the Amazon Web Services console.

Table 1.1: Machine Type and Location

Cloud provider	Amazon Web Services
Machine region	N. Virginia (us-east-1)
Common name of the item	r5d.12xlarge
Operating system	18.04.6-Ubuntu

This benchmark used two `r5d.12xlarge` instances, one for the driver and one for the SUT placed in a AWS Placement group with the cluster strategy to reduce network latency. [1] Both machines were assigned to the same VPC with the same subnet. This is shown in Figure 1.1.

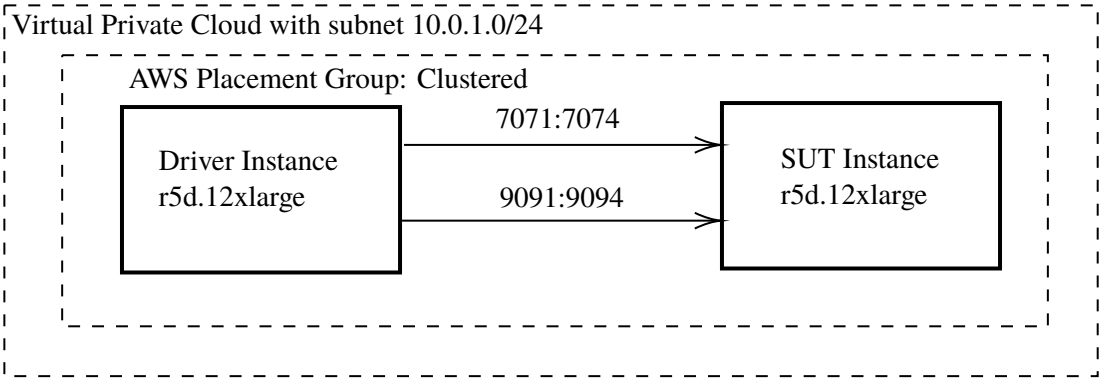

Figure 1.1: Overview of benchmark setup

1.1.2 CPU details

The details below were obtained using the command `cat /proc/cpuinfo` (Listing A.1) issued from the machine instance and the datasheet of the used CPU type.

Table 1.2: CPU details summary

Type	Intel®Intel Xeon®Platinum 8259CL CPU @ 2.5GHz
Total number	1
Cores per CPU	24
Threads per CPU	48
CPU clock frequency	2.5GHz
Total cache size per CPU	L1 cache: 1536KiB L2 cache: 24MiB L3 cache: 33MiB

[1] https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#placement-groups-cluster

1.1.3 Memory details

The total size of the memory installed is 384GiB and the type of memory is DDR4 with frequency 2666MHz. This information was obtained using the `sudo lshw -c memory` command (Listing A.2) issued from the virtual machine instance.

1.1.4 Disk and storage details

Disk controller or motherboard type was not obtainable from the virtual machine instance. The storage consists of 2 x 900GB NVMe SSD in RAID0 configuration, formatted with `ext4` filesystem. The storage size and type is from the Amazon Web Services website `https://aws.amazon.com/ec2/instance-types/r5/` (accessed: August 09, 2022).

The 4KB QD1 write performance was measured with the `fio` command and the output (Listing A.5) showed an average of 7544 IOPS.

1.1.5 Network details

The benchmark was run using two `r5d.12xlarge` instances, both deployed in the same availability zone behind a Virtual Private Cloud (VPC) configured with subnet `10.0.1.0/24`. Both instances where assigned security groups to open the following ports:

- `7071`: REST port used by TuGraph
- `9091`: RPC-port used by TuGraph

On each instance, the firewall was disabled using `sudo ufw disable` to minimize performance impact of the firewall during the benchmark.

The `r5d.12xlarge` instances use the Elastic Network Adapter provided by Amazon. This information was obtained using the `lshw -class network` command (Listing A.3). Network throughput between the two instances was measured using the `iperf` tool on port 9091 using 48 threads and the output (Listing A.4) showed an average of 12.0 Gbit/sec from client to server and 17.7 Gbit/sec from server to client.

1.1.6 Machine pricing

The system pricing summary is included in the table below. The pricing of the AWS machine instance is the price for a 3-year reserved dedicated instance machine without upfront payment.

Table 1.3: Pricing summary

Item	**Price**
r5d.12xlarge reserved instance machine in AWS (standard 3-year term)	44 676 USD
Permanent TuGraph license	170 000 USD
Maintenance fee (3 years)	76 500 USD
Total cost of ownership	291 176 USD

1.1.7 System availability

The latest software version of TuGraph (version 3.2) was made available on July 29, 2022. This version was deployed to the machine described in this section.

2 Dataset Generation

2.1 General information

The data generation settings of the LDBC Datagen are described below.

Table 2.1: Datagen settings summary

Datagen version	v0.3.5
Output format	CsvCompositeMergeForeign serializer
Scale factors	10, 30, 100, and 300

Scale factor 10 is used for query results validation only, while the other ones were used for performance measurements.

2.2 Datagen configurations

The Datagen configuration for SF10 is shown in Listing 2.1. The configurations for SF30, SF100 and SF300 are shown in Listings A.6–A.8.

Listing 2.1: Contents of `params-sf10.ini` used for scale factor 10

```
ldbc.snb.datagen.generator.scaleFactor:snb.interactive.10
ldbc.snb.datagen.serializer.numUpdatePartitions:32

ldbc.snb.datagen.serializer.dynamicActivitySerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.activity.CsvCompositeMergeForeignDynamicActivitySerializer
ldbc.snb.datagen.serializer.dynamicPersonSerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.person.CsvCompositeMergeForeignDynamicPersonSerializer
ldbc.snb.datagen.serializer.staticSerializer:ldbc.snb.datagen.serializer.snb.csv.staticserializer.CsvCompositeMergeForeignStaticSerializer

ldbc.snb.datagen.serializer.dateFormatter:ldbc.snb.datagen.util.formatter.LongDateFormatter
```

2.3 Data loading and data schema

The output produced by the Datagen is converted to a custom (i.e., vendor-specific) CSV representation which can be loaded into the database (see the attached `convert.py` script). The loading process takes a configuration file `import.conf` (see Listing 2.2 and Listing A.9) which defines the files to process along with what data types are represented in the values in each column.

Listing 2.2: Excerpt from `import.conf`, describing the data schema

```
{
    "schema": [
    {
        "label" : "Comment",
            "type" : "VERTEX",
            "properties" : [
            { "name" : "id", "type":"INT64"},
            { "name" : "creationDate", "type":"INT64"},
            { "name" : "locationIP", "type":"STRING"},
            { "name" : "browserUsed", "type":"STRING"},
            { "name" : "content", "type":"STRING"},
            { "name" : "length", "type":"INT32"},
            { "name" : "creator", "type":"INT64"},
            { "name" : "place", "type":"INT64"},
            { "name" : "replyOfPost", "type":"INT64", "optional":true},
            { "name" : "replyOfComment", "type":"INT64", "optional":true}
        ],
            "primary" : "id"
    },
```

Data loading times are shown for each scale factor in the table below. Values were measured using the GNU Time tool, reading the *Elapsed real time* from the output. The column **CSV loading time** shows how long it took to create a graph from the input CSV files and also to build and index over *id* properties, but they do not include CSV conversion times (CSV conversion time is outside the scope of the loading times). The column **Data preprocessing time** shows how much time it took to create initial precomputed properties (similar to materialized views) *hasMember.numPosts* and *knows.weight*, the conversion of foreign keys to vertices and to build an index over *name* properties. The column **Total time** contains the sum of the CSV loading and data preprocessing times.

Table 2.2: Data loading times

Scale factor	CSV loading time (s)	Data preprocessing time (s)	Total time (s)
30	445.796	169.415	615.211
100	1 320.981	567.084	1 888.065
300	3 890.442	1 848.022	5 738.464

3 Test Driver Details

The driver and implementations version used are described below as well as the amount of read and write threads used by the driver.

Table 3.1: Summary of test artifacts and main configuration parameters

Driver version	v1.2.0	`https://github.com/ldbc/ldbc_snb_interactive_driver/releases/tag/v1.2.0`
Implementations version	v1.0.0	`https://github.com/ldbc/ldbc_snb_interactive_impls/releases/tag/1.0.0`
LDBC SNB specification version	v0.3.6	`https://arxiv.org/pdf/2001.02299v3.pdf`
Driver Read threads	48	
Driver Write threads	32	

3.1 Driver implementation

A test driver adaptation for the SUT was provided by the test sponsor and is available as part of the attachment package. The archive created from the version of the driver used for the audited run is included in the attachments of this report.

The SUT-specific test driver class `ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb` extends the class `com.ldbc.driver.Db` provided in the LDBC SNB Interactive driver package. Internally, the `LightGraphInteractiveDb` relies on remote procedure calls (RPC) using local sockets to communicate with the SUT.

3.2 Benchmark configuration of driver

The driver applied time compression ratio values of

- TCR=0.0028 for scale factor 30,
- TCR=0.0104 for scale factor 100 and
- TCR=0.036 for scale factor 300.

The complete configuration files for the different scale factors are shown in Listings A.10–A.12, and are also included in the attached supplementary materials.

4 Performance Metrics

The performance metrics reported here show benchmark runs with scale factors 30, 100 and 300. The performance summary tables below highlight key performance characteristics. Performance metrics with 100% query on-time compliance are shown in Appendix A.1.

Table 4.1: Summary of results for scale factor 30

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 05m 08.225s	91 994 491	12 252.50 $\frac{\text{operations}}{\text{second}}$	95.85%

Table 4.2: Summary of results for scale factor 100

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 04m 59.254s	96 999 938	12 934.61 $\frac{\text{operations}}{\text{second}}$	97.02%

Table 4.3: Summary of results for scale factor 300

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 04m 26.203	94 979 387	12 721.24 $\frac{\text{operations}}{\text{second}}$	95.47%

During the benchmark run, the query executions shown in the tables below were observed using the different scale factors. Columns (except for Query and Total count) are showing duration values with microsecond (μs) precision. The notation $\mathbf{P_i}$ is used for the i^{th} percentile among all observed execution run times of a given query type. After each benchmark result table a table with the percentage of late operations per query type is shown.

Table 4.4: Detailed performance benchmark results for scale factor 30 with 95.85% on time in microseconds

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	640 475	1 121	518 320	26 220.23	17 706	28 258	34 646	361 408
Query2	450 064	351	82 784	6 016.88	5 292	8 975	11 686	19 838
Query3	157 098	13 312	128 120	37 838.88	36 964	51 580	57 198	70 216
Query4	462 565	366	94 352	9 311.54	7 817	16 264	19 488	29 053
Query5	231 282	7 844	402 016	67 564.74	64 868	103 456	119 076	158 240
Query6	52 697	647	96 348	12 176.09	12 461	25 333	30 342	40 670
Query7	346 924	276	97 268	1 891.34	1 168	3 849	5 718	11 557
Query8	1 850 260	1 252	407 440	12 020.88	7 035	27 006	37 726	70 708
Query9	43 365	37 986	917 280	531 468.22	529 472	670 688	703 904	755 520
Query10	450 064	1 728	789 824	51 673.80	42 132	90 396	116 036	187 888
Query11	832 617	463	102 528	3 177.53	2 424	5 389	7 469	13 904
Query12	378 462	590	668 768	51 096.79	41 702	91 588	115 224	181 296
Query13	876 439	276	120 276	2 309.38	1 581	4 504	6 422	12 410
Query14	339 844	344	415 792	13 580.72	4 812	11 936	102 780	170 768
Short1	8 926 009	236	53 990	1 519.34	822	3 359	5 163	10 111
Short2	8 926 009	258	44 480	1 576.58	997	3 133	4 608	9 231
Short3	8 926 009	304	76 568	2 822.06	1 580	6 636	9 780	16 372
Short4	8 926 890	242	53 904	1 607.77	830	3 717	5 670	10 650
Short5	8 926 890	234	40 314	1 369.87	794	2 877	4 334	8 801
Short6	8 926 890	239	53 816	1 350.67	794	2 798	4 216	8 665
Short7	8 926 890	251	45 700	1 419.69	865	2 876	4 303	8 796
Update1	5 112	432	66 516	3 723.00	2 133	8 384	11 903	23 581
Update2	3 177 604	321	133 192	3 847.50	2 370	9 213	12 882	20 201
Update3	4 491 147	327	106 160	3 893.92	2 419	9 297	12 965	20 256
Update4	88 004	376	87 180	3 057.83	1 675	7 168	10 343	18 268
Update5	9 775 146	341	109 436	4 158.86	2 647	9 744	13 477	20 895
Update6	1 150 120	387	64 612	2 970.33	1 663	6 814	9 753	17 462
Update7	3 370 923	402	100 276	3 501.93	2 080	8 028	11 256	19 488
Update8	338 692	360	189 288	9 467.03	5 363	19 743	28 648	80 116

Table 4.5: Detailed query on time results for scale factor 30

Query	Total count	Total late count	Late count Percentage
Query1	640 475	338 664	52.88
Query2	450 064	237 981	52.88
Query3	157 098	83 068	52.88
Query4	462 565	244 582	52.88
Query5	231 282	122 293	52.88
Query6	52 697	27 865	52.88
Query7	346 924	183 440	52.88
Query8	1 850 260	978 371	52.88
Query9	43 365	22 929	52.87
Query10	450 064	237 977	52.88
Query11	832 617	440 265	52.88
Query12	378 462	200 120	52.88
Query13	876 439	463 431	52.88
Query14	339 844	179 697	52.88
Short1	8 926 009	0	0.00
Short2	8 926 009	0	0.00
Short3	8 926 009	0	0.00
Short4	8 926 890	0	0.00
Short5	8 926 890	0	0.00
Short6	8 926 890	0	0.00
Short7	8 926 890	0	0.00
Update1	5 112	0	0.00
Update2	3 177 604	7 974	0.25
Update3	4 491 147	15 110	0.34
Update4	88 004	135	0.15
Update5	9 775 146	15 571	0.16
Update6	1 150 120	2 654	0.23
Update7	3 370 923	15 638	0.46
Update8	338 692	548	0.16
Total	91 994 491	3 818 313	4.15

Table 4.6: Detailed performance benchmark results for scale factor 100 with 97.02% on time in microseconds

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	576 459	356	1 046 208	67 165.43	22 961	38 602	652 064	778 144
Query2	405 079	348	109 556	6 723.48	6 037	9 566	12 519	20 264
Query3	121 853	32 349	281 344	103 401.92	105 104	127 028	136 136	164 008
Query4	416 332	383	87 576	10 129.20	8 482	17 749	20 729	30 842
Query5	192 153	774	272 416	66 256.41	64 142	94 148	104 368	127 464
Query6	34 535	601	133 952	30 012.09	13 233	67 064	73 536	87 016
Query7	394 420	355	87 040	1 845.73	1 189	3 550	5 530	11 545
Query8	2 997 589	382	127 768	2 757.16	1 337	6 947	10 134	16 682
Query9	28 440	415	1 211 328	728 634.74	732 096	889 824	929 760	999 616
Query10	374 699	454	859 264	62 877.67	52 460	103 480	132 232	210 528
Query11	681 270	361	80 668	3 546.20	2 852	5 482	7 847	14 597
Query12	340 636	482	1 050 432	69 274.21	54 474	128 068	163 120	262 088
Query13	788 839	361	99 320	3 131.61	2 577	5 345	7 557	14 130
Query14	305 876	417	454 144	32 178.78	5 912	101 892	119 016	164 496
Short1	9 612 581	308	49 392	1 504.47	893	2 999	4 959	10 425
Short2	9 612 581	324	50 104	1 606.57	1 094	2 840	4 411	9 600
Short3	9 612 581	371	86 796	2 897.48	1 605	6 769	9 913	16 447
Short4	9 612 966	307	48 850	1 600.02	901	3 427	5 570	11 047
Short5	9 612 966	306	48 188	1 369.62	875	2 519	4 079	9 127
Short6	9 612 966	310	48 188	1 357.20	875	2 466	3 996	9 018
Short7	9 612 966	324	50 656	1 439.95	958	2 564	4 114	9 196
Update1	3 878	587	59 262	4 344.93	2 503	10 314	13 415	22 891
Update2	3 075 601	406	101 772	3 438.39	1 974	8 278	11 849	18 958
Update3	5 688 162	398	102 416	3 421.74	1 970	8 230	11 780	18 822
Update4	65 173	452	61 918	3 586.68	1 907	8 674	11 452	19 414
Update5	7 786 408	430	109 464	3 904.07	2 367	9 123	12 732	19 948
Update6	986 412	467	80 412	4 141.84	2 309	9 629	12 317	20 149
Update7	4 149 530	472	87 428	4 522.63	2 519	10 252	13 493	21 854
Update8	296 987	461	189 040	8 605.50	5 143	18 911	27 390	51 448

Table 4.7: Detailed query on time results for scale factor 100

Query	Total count	Total late count	Late count Percentage
Query1	576 459	169 643	29.43
Query2	405 079	119 207	29.43
Query3	121 853	35 870	29.44
Query4	416 332	122 539	29.43
Query5	192 153	56 540	29.42
Query6	34 535	10 160	29.42
Query7	394 420	116 071	29.43
Query8	2 997 589	882 350	29.44
Query9	28 440	8 357	29.38
Query10	374 699	110 241	29.42
Query11	681 270	200 512	29.43
Query12	340 636	100 243	29.43
Query13	788 839	232 162	29.43
Query14	305 876	90 023	29.43
Short1	9 612 581	0	0.00
Short2	9 612 581	0	0.00
Short3	9 612 581	0	0.00
Short4	9 612 966	0	0.00
Short5	9 612 966	0	0.00
Short6	9 612 966	0	0.00
Short7	9 612 966	0	0.00
Update1	3 878	0	0.00
Update2	3 075 601	78 133	2.54
Update3	5 688 162	143 554	2.52
Update4	65 173	1 196	1.84
Update5	7 786 408	136 399	1.75
Update6	986 412	36 662	3.72
Update7	4 149 530	236 259	5.69
Update8	296 987	5 233	1.76
Total	96 999 938	2 891 354	2.98

Table 4.8: Detailed performance benchmark results for scale factor 300 with 95.47% on time in microseconds

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	466 919	302	1 628 736	40 436.00	24 929	34 608	41 224	1 080 960
Query2	328 105	294	103 764	6 327.71	5 733	9 320	11 453	19 506
Query3	85 492	87 124	682 144	268 495.50	266 528	340 576	381 360	473 728
Query4	337 219	283	125 740	11 395.90	9 653	20 642	23 494	33 890
Query5	144 522	545	332 816	76 863.68	74 172	110 192	122 284	148 776
Query6	20 931	403	259 712	72 132.02	15 190	165 248	178 080	199 464
Query7	379 371	276	88 060	1 602.92	1 001	3 053	4 611	10 535
Query8	4 046 633	314	126 696	2 210.67	1 027	5 314	8 603	15 559
Query9	17 220	364	1 388 224	813 791.97	805 600	1 046 688	1 109 504	1 210 560
Query10	275 907	368	729 856	59 017.41	50 378	95 988	118 236	191 872
Query11	505 829	467	110 060	3 609.58	3 038	5 299	7 295	14 176
Query12	275 907	400	1 091 136	72 113.99	54 704	139 216	179 072	292 128
Query13	638 942	286	108 204	4 001.23	3 726	6 508	8 361	15 488
Query14	247 753	325	538 912	37 848.68	9 420	90 184	105 676	150 848
Short1	9 751 564	238	42 882	1 216.59	710	2 410	3 747	8 506
Short2	9 751 564	253	41 554	1 390.84	937	2 563	3 684	8 101
Short3	9 751 564	297	78 024	2 351.06	1 307	5 192	8 155	14 814
Short4	9 750 534	236	60 394	1 261.89	707	2 566	4 094	9 086
Short5	9 750 534	240	45 360	1 142.13	701	2 220	3 319	7 627
Short6	9 750 534	240	44 010	1 142.51	706	2 213	3 293	7 548
Short7	9 750 534	253	51 624	1 228.40	796	2 315	3 406	7 735
Update1	2 733	468	59 024	4 503.15	2 159	11 485	14 739	25 123
Update2	2 423 377	335	101 220	2 903.38	1 692	6 829	10 398	17 846
Update3	3 900 359	330	119 276	2 924.45	1 723	6 858	10 420	17 858
Update4	46 692	394	68 756	4 494.35	1 854	11 206	14 684	23 783
Update5	6 314 401	345	111 692	3 080.19	1 834	7 076	10 678	18 229
Update6	1 110 425	394	97 552	6 766.27	5 749	14 056	17 895	27 146
Update7	4 907 806	411	134 344	6 219.50	4 921	13 170	17 005	26 039
Update8	246 016	383	206 280	6 890.60	3 971	15 740	22 938	43 028

Table 4.9: Detailed query on time results for scale factor 300

Query	Total count	Total late count	Late count Percentage
Query1	466 919	149 549	32.03
Query2	328 105	105 085	32.03
Query3	85 492	27 380	32.03
Query4	337 219	108 008	32.03
Query5	144 522	46 288	32.03
Query6	20 931	6 712	32.07
Query7	379 371	121 510	32.03
Query8	4 046 633	1 296 115	32.03
Query9	17 220	5 514	32.02
Query10	275 907	88 371	32.03
Query11	505 829	162 015	32.03
Query12	275 907	88 374	32.03
Query13	638 942	204 659	32.03
Query14	247 753	79 345	32.03
Short1	9 751 564	26	0.00
Short2	9 751 564	6	0.00
Short3	9 751 564	7	0.00
Short4	9 750 534	31	0.00
Short5	9 750 534	5	0.00
Short6	9 750 534	4	0.00
Short7	9 750 534	3	0.00
Update1	2 733	0	0.00
Update2	2 423 377	217 527	8.98
Update3	3 900 359	385 258	9.88
Update4	46 692	2 918	6.25
Update5	6 314 401	401 226	6.35
Update6	1 110 425	117 949	10.62
Update7	4 907 806	669 686	13.65
Update8	246 016	15 501	6.30
Total	94 979 387	4 299 072	4.53

5 Validation of the Results

The scale factor 10 data set was used for validating the correctness of the implementation over the SUT. The validation data set of size 10 000 was created with the SNB Interactive reference implementation over Neo4j, running the Community Edition of version 4.4.1. The system with the driver configuration shown in Listing A.13 successfully returned the expected result sets for the queries of the benchmark.

6 ACID Compliance

6.1 Transaction isolation level

The benchmark was executed using the *serializable* isolation level setting of the SUT, which is more strict than the *read committed* isolation level minimally required by the SNB Interactive specification.

6.2 SNB Interactive ACID test results

The ACID test implementations were reviewed to conform to the ACID test specifications, with all specified test cases implemented. Furthermore, test execution was successful, no atomicity and isolation test failed with serializable isolation level transaction settings.

6.3 Recovery and durability

6.3.1 Recovery

Durability tests were using the regular benchmark workload with scale factor 30 and started at `14:00 UTC`. Both server machines where shutdown using the command `sudo shutdown -rf 16:00`, forced rebooting (ungracefully) both machines 2 hours after start of the benchmark, with 65 669 333 completed operations. The database server process was manually started again after the crash and it was ready in 20 ms, which was not different from a regular server startup time.

6.3.2 Durability

From the driver log, the last update operations before the crash were obtained using the commands below.

```
$ grep LdbcUpdate1 LDBC-SNB-results_log.csv |tail -n 1
LdbcUpdate1AddPerson|1659542399650|1659542399650|3679|0|1350068430083
$ grep LdbcUpdate2 LDBC-SNB-results_log.csv |tail -n 1
LdbcUpdate2AddPostLike|1659542400073|1659542400079|1287|0|1350068581205
$ grep LdbcUpdate3 LDBC-SNB-results_log.csv |tail -n 1
LdbcUpdate3AddCommentLike|1659542400070|1659542400078|795|0|1350068580052
$ grep LdbcUpdate4 LDBC-SNB-results_log.csv |tail -n 1
LdbcUpdate4AddForum|1659542400080|1659542400080|751|0|1350068583811
$ grep LdbcUpdate5 LDBC-SNB-results_log.csv |tail -n 1
LdbcUpdate5AddForumMembership|1659542400069|1659542400078|1284|0|1350068580227
$ grep LdbcUpdate6 LDBC-SNB-results_log.csv |tail -n 1
LdbcUpdate6AddPost|1659542400081|1659542400081|936|0|1350068584068
$ grep LdbcUpdate7 LDBC-SNB-results_log.csv |tail -n 1
LdbcUpdate7AddComment|1659542400075|1659542400080|2629|0|1350068581845
$ grep LdbcUpdate8 LDBC-SNB-results_log.csv |tail -n 1
LdbcUpdate8AddFriendship|1659542400044|1659542400044|14085|0|1350068570839
```

From the logs, the last completed updates were retrieved for each update. The log entries include the operation name, actual and scheduled start time, the execution time, the delay between scheduled and actual start times, and the initial query start time without the scale factor multiplier (this latter one is included in the last column). Using this information, the query parameters were obtained from the initial CSV files generated by the Datagen using the initial query start time and the type number of the operation using the commands below.

```
$ grep -rnw '/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/' -e '1350068430083|.*|1|.*'
/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/updateStream_0_2_person.csv:149:1350068430083|0|1|35184372237273|
    John|Dogg|male|532569600000|1350068430083|41.223.117.50|Firefox|1449|ny;en|John35184372237273@gmail.com;
    John35184372237273@yahoo.com|1154;2939;4870;5371;12391|7952,2007|1571,2008;1572,2009;260,2008;1573,2009
$ grep -rnw '/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/' -e '1350068581205|.*|2|.*'
/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/updateStream_0_12_forum.csv
    :711616:1350068581205|1315092411454|2|19791209363615|35184478714554|1350068581205
$ grep -rnw '/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/' -e '1350068580052|.*|3|.*'
/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/updateStream_0_12_forum.csv
    :711615:1350068580052|1327168177786|3|26388279141431|35184478713887|1350068580052
$ grep -rnw '/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/' -e '1350068583811|.*|4|.*'
/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/updateStream_0_20_forum.csv
    :687146:1350068583811|1296530682491|4|35184376368866|Album 14 of Angela Chen|1350068583811|13194139714703|522
$ grep -rnw '/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/' -e '1350068580227|.*|5|.*'
/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/updateStream_0_18_forum.csv
    :659922:1350068580227|1345216491778|5|10995126813782|32985348892725|1350068580227
$ grep -rnw '/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/' -e '1350068584068|.*|6|.*'
/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/updateStream_0_17_forum.csv
    :680319:1350068584068|1312190861713|6|35184420884476|photo35184420884476.jpg|1350068584068|61.11.81.203|
    Firefox|||0|19791209334001|35184386796317|0|
$ grep -rnw '/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/' -e '1350068581845|.*|7|.*'
/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/updateStream_0_12_forum.csv
    :711617:1350068581845|1341676641489|7|35184406860398|1350068581845|196.1.98.167|Internet Explorer|About Jean-
    Jacques Rousseau, f the whole person for citAbout Edward III of En
    |77|30786325597840|96|-1|35184406860392|1909;2014;12095
$ grep -rnw '/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/' -e '1350068570839|.*|8|.*'
/data/tugraph_ldbc_snb/deps/ldbc_snb_datagen/updateStream_0_14_forum.csv
    :685409:1350068570839|1349680977315|8|24189255865173|35184372108746|1350068570839
```

To check whether the graph entities in the driver log entries were persisted in the database, custom read queries were executed after database restart. The queries returned the data that was committed according to the logs, so the system passed this check. These queries are included in the `recovery_queries.cpp` attachment.

6.3.3 Consistency after recovery

The provided `check_consistency` program was executed to verify that the precomputed values (i.e., materialized views) are still consistent after a crash. This check completed successfully.

7 Supplementary Materials

The table below shows the list of supplementary materials. These materials are made available with this full disclosure report to allow reproducibility of results.

Table 7.1: Supplementary materials

File	Purpose
`results-sf{30,100,300}.tar.gz`	Driver output files for the selected scale factors
`results-sf{30,100,300}-100-on-time.tar.gz`	Driver output files for the selected scale factors with 100% query on time compliance
`console-output-sf{30,100,300}.txt`	Driver console outputs for the selected scale factors
`console-output-sf{30,100,300}-100-ontime.txt`	Driver console outputs for the selected scale factors with 100% query on time compliance
`params-sf{10,30,100,300}.ini`	Datagen parameters for the used scale factors
`convert.py` and `convert-csvs.sh`	CSV converter scripts
`import_data.sh`	Database bulk importer tool
`import.conf`	Data schema descriptor file
`lgraph_standalone.json`	Database configuration file
`ldbc_snb_interactive_impls.zip`	SUT-specific LDBC driver implementation
`interactive-benchmark-sf{30,100,300}.properties`	Driver configurations
`interactive-validate.properties`	Results validation driver settings
`«query».cpp`	Benchmark query implementation files
`acid.cpp`	ACID tests implementation
`compile_embedded.sh`	Used to compile the ACID tests implementation
`check_consistency.cpp`	Tool to verify materialized views
`install.{sh,py}` and `compile_plugin.sh`	Scripts to install queries as stored procedures
`compile_embedded.sh`	Script to compile standalone TuGraph applications
`preprocess.cpp`	Plugin to calculate materialized views
`snb_common.h` and `snb_constants.h`	Data schema-specific headers
`generate_snb_constants.cpp`	Tool to generate data schema-specific header
`recovery_queries.cpp`	Test cases to check commited transactions for recovery test
`TuGraph-3.2.0-1.x86_64.deb`	Linux installer package of database

The attachment folder directory structure is as follows:

```
attachments
├── results-sf30.tar.gz
├── results-sf100.tar.gz
├── results-sf300.tar.gz
├── results-sf30-100-on-time.tar.gz
├── results-sf100-100-on-time.tar.gz
├── results-sf300-100-on-time.tar.gz
├── console-output-sf30.txt
├── console-output-sf100.txt
├── console-output-sf300.txt
├── console-output-sf30-100-ontime.txt
├── console-output-sf100-100-ontime.txt
├── console-output-sf300-100-ontime.txt
├── TuGraph-3.2.0-1.x86_64.deb
└── tugraph_ldbc_snb
    ├── images
    │   └── security_groups.jpg
    ├── deps
    │   └── ldbc_snb_interactive_impls.zip
    ├── README.md
    ├── NOTES.md
    ├── LEGAL.md
    ├── params-sf10.ini
    ├── params-sf30.ini
    ├── params-sf100.ini
    ├── params-sf300.ini
    ├── convert_csvs.sh
    ├── import_data.sh
    ├── lgraph_standalone.json
    ├── load-scripts
    │   ├── convert.py
    │   └── import-data
    │       └── import.conf
    └── plugins
        ├── «query».cpp
        ├── acid.cpp
        ├── check_consistency.cpp
        ├── compile_plugin.sh
        ├── compile_embedded.sh
        ├── install.py
        ├── install.sh
        ├── preprocess.cpp
        ├── snb_common.h
        ├── snb_constants.h
        ├── generate_snb_constants.cpp
        └── recovery_queries.cpp
```

A Appendix

A.1 Performance metrics 100% on-time compliance

The performance metrics reported here show benchmark runs with scale factors 30, 100 and 300 for 100% query on-time compliance which can be compared with previous audited results. The driver version used in this audit, v1.2.0, used the tolerated late operation count of 95% of all the queries, where the previous driver version, v.0.3.x, allowed only for 10 late operations per query type, see the GitHub issue #154 [1] in the driver repository for more details.

The driver applied time compression ratio values of

- TCR=0.0034 for scale factor 30,
- TCR=0.011475 for scale factor 100 and
- TCR=0.0432 for scale factor 300.

Table A.1: Summary of results for scale factor 30

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 01m 59s	75 985 520	10 380.65 $\frac{\text{operations}}{\text{second}}$	100%

Table A.2: Summary of results for scale factor 100

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 05m 33s	89 004 182	11 814.59 $\frac{\text{operations}}{\text{second}}$	100%

Table A.3: Summary of results for scale factor 300

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 13m 25s	84 008 662	10 494.29 $\frac{\text{operations}}{\text{second}}$	100%

During the benchmark run, the query executions shown in the tables below were observed using the different scale factors. Columns (except for Query and Total count) are showing duration values with microsecond (μs) precision. The notation $\mathbf{P_i}$ is used for the i^{th} percentile among all observed execution run times of a given query type. After each benchmark result table a table with the percentage of late operations per query type is shown.

[1] https://github.com/ldbc/ldbc_snb_interactive_driver/issues/154

Table A.4: Detailed performance benchmark results for scale factor 30 in microseconds for 100% query on-time compliance

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	530 774	1 050	442 160	20 974.74	14 434	17 810	19 097	313 104
Query2	372 976	186	43 570	3 861.46	3 608	5 235	5 870	8 452
Query3	130 189	11 996	94 468	28 181.26	28 170	35 010	37 156	42 546
Query4	383 337	249	44 996	6 442.49	5 586	11 065	12 890	16 241
Query5	191 668	7 100	105 732	41 629.09	41 286	53 866	57 798	66 000
Query6	43 672	487	43 352	8 824.70	10 858	18 405	19 834	22 529
Query7	287 502	170	43 624	784.44	554	1 260	1 928	4 815
Query8	1 533 348	1 119	70 684	3 129.24	2 550	5 082	6 083	8 822
Query9	35 938	32 418	723 872	435 920.02	433 872	523 056	545 344	590 368
Query10	372 976	1 485	130 052	23 516.33	22 979	29 861	32 497	39 332
Query11	690 006	295	52 006	1 776.61	1 586	2 279	2 702	5 786
Query12	313 639	485	213 464	29 355.30	25 152	48 544	58 992	82 124
Query13	726 322	163	48 006	1 103.17	810	2 119	2 654	4 971
Query14	281 635	219	315 232	9 962.63	3 384	5 445	82 564	128 964
Short1	7 396 228	136	20 264	365.92	300	544	667	1 371
Short2	7 396 228	147	22 885	521.78	448	768	960	1 647
Short3	7 396 228	192	24 143	673.40	538	1 119	1 658	2 416
Short4	7 397 240	134	23 038	370.55	301	551	682	1 464
Short5	7 397 240	133	22 939	354.89	299	528	635	1 193
Short6	7 397 240	138	20 178	360.53	304	535	643	1 216
Short7	7 397 240	145	16 102	416.84	362	609	722	1 288
Update1	4 246	319	34 290	1 328.92	1 033	2 166	2 539	5 973
Update2	2 598 827	221	49 934	631.50	464	896	1 380	3 456
Update3	3 759 560	219	56 914	634.14	466	900	1 387	3 473
Update4	72 738	268	44 706	698.40	498	1 057	1 702	3 817
Update5	7 854 131	237	102 496	849.05	651	1 359	1 917	3 835
Update6	953 316	279	44 106	763.41	558	1 186	1 910	3 971
Update7	2 793 941	295	48 466	769.05	583	1 058	1 643	4 139
Update8	277 135	250	127 904	4 055.75	2 105	8 516	15 442	28 450

Table A.5: Detailed performance benchmark results for scale factor 100 in microseconds for 100% query on-time compliance

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	526 013	174	983 968	60 077.70	19 612	30 236	610 848	706 592
Query2	369 630	175	107 320	5 239.62	4 660	7 235	8 895	18 286
Query3	111 189	30 653	237 984	87 838.36	87 696	110 260	118 544	143 056
Query4	379 898	207	110 348	7 888.52	6 585	14 263	16 772	25 065
Query5	175 338	513	215 592	50 305.85	48 628	68 764	79 016	103 024
Query6	31 513	218	126 716	25 578.69	11 627	57 028	62 334	74 860
Query7	359 903	164	110 052	1 168.46	617	1 950	3 429	11 451
Query8	2 735 265	215	113 968	1 351.37	549	2 582	5 593	14 248
Query9	25 951	217	1 049 408	646 450.21	645 152	785 024	820 096	887 680
Query10	341 908	268	586 336	39 795.71	34 376	59 470	75 960	124 268
Query11	621 651	201	111 808	2 636.52	2 148	3 420	5 239	14 003
Query12	310 826	270	815 488	46 409.40	36 512	85 260	106 664	161 016
Query13	719 806	169	113 616	2 276.70	1 727	3 705	5 026	13 470
Query14	279 109	242	424 656	26 964.97	4 760	85 988	101 624	143 632
Short1	8 772 509	139	31 902	573.22	342	891	1 454	5 171
Short2	8 772 509	156	47 756	734.52	526	1 165	1 669	4 435
Short3	8 772 509	198	57 800	1 208.73	659	2 279	3 769	10 656
Short4	8 771 451	136	36 962	607.20	346	931	1 635	6 008
Short5	8 771 451	132	31 192	527.09	338	809	1 272	4 066
Short6	8 771 451	137	35 002	527.07	344	808	1 254	3 954
Short7	8 771 451	152	42 592	599.23	419	908	1 351	4 078
Update1	3 539	338	56 000	2 300.04	1 327	4 589	7 982	17 909
Update2	2 869 080	227	91 296	1 272.86	568	2 539	4 757	12 598
Update3	5 390 198	221	107 816	1 263.14	566	2 513	4 701	12 502
Update4	59 493	269	74 764	1 415.14	604	3 352	5 589	11 836
Update5	7 083 612	242	103 916	1 607.91	813	3 187	5 651	13 774
Update6	926 299	272	76 592	2 200.83	766	6 745	8 726	12 123
Update7	4 009 664	302	122 944	1 486.51	708	3 573	6 018	11 322
Update8	270 966	274	153 664	5 195.38	2 674	12 226	19 360	36 402

Table A.6: Detailed performance benchmark results for scale factor 300 in microseconds for 100% query on-time compliance

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	419 204	165	1 581 504	35 235.78	21 444	27 797	30 333	999 552
Query2	294 576	167	92 524	4 776.41	4 299	7 074	8 174	15 145
Query3	76 756	83 692	579 392	229 354.41	233 744	288 624	333 744	404 896
Query4	302 759	183	120 768	8 935.10	7 552	15 806	19 198	24 393
Query5	129 753	335	225 216	61 108.70	60 282	82 344	88 392	104 836
Query6	18 792	223	240 608	64 623.06	86 208	142 128	152 216	170 048
Query7	340 604	168	101 272	981.62	549	1 482	2 555	9 446
Query8	3 633 106	210	127 572	966.92	501	1 214	2 703	11 104
Query9	15 460	194	1 153 920	721 141.18	722 720	899 552	946 624	1 020 864
Query10	247 712	251	345 520	38 928.17	37 072	53 056	60 786	86 728
Query11	454 138	337	102 044	2 655.16	2 201	3 308	4 399	12 454
Query12	247 712	254	602 784	46 865.57	37 232	87 112	109 084	158 512
Query13	573 649	174	112 188	3 000.16	2 877	4 907	5 466	12 856
Query14	222 435	196	480 672	30 726.89	6 924	72 964	85 324	122 744
Short1	8 758 164	139	29 925	429.69	312	636	907	2 405
Short2	8 758 164	150	43 056	621.01	494	947	1 289	2 590
Short3	8 758 164	195	47 236	911.79	634	1 646	2 414	5 144
Short4	8 758 659	134	29 934	437.50	314	644	924	2 589
Short5	8 758 659	137	28 365	414.26	312	609	819	2 185
Short6	8 758 659	140	30 825	419.69	317	615	824	2 197
Short7	8 758 659	149	28 789	490.33	388	719	935	2 305
Update1	2 456	347	75 216	2 098.10	1 197	3 058	7 177	17 956
Update2	2 043 745	226	105 456	926.62	534	1 482	2 622	8 754
Update3	3 216 673	230	119 460	927.54	527	1 501	2 664	8 822
Update4	42 083	286	50 340	1 323.97	538	2 689	6 355	13 013
Update5	5 662 928	248	98 076	1 132.61	722	1 909	2 985	8 745
Update6	894 275	297	94 600	3 590.19	901	10 021	11 171	17 532
Update7	3 639 853	303	109 596	1 758.67	667	5 159	8 342	15 157
Update8	220 865	265	100 988	4 207.02	2 056	9 600	15 567	32 854

A.2 CPU details

Listing A.1: Output of the `cat /proc/cpuinfo` command for one core

```
processor : 0
vendor_id : GenuineIntel
cpu family  : 6
model   : 85
model name  : Intel(R) Xeon(R) Platinum 8259CL CPU @ 2.50GHz
stepping  : 7
microcode : 0x500320a
cpu MHz   : 1337.803
cache size  : 36608 KB
physical id : 0
siblings  : 48
core id   : 0
cpu cores : 24
apicid    : 0
initial apicid  : 0
fpu   : yes
fpu_exception : yes
cpuid level : 13
wp    : yes
flags   : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse sse2 ss ht
      syscall nx pdpe1gb rdtscp lm constant_tsc arch_perfmon rep_good nopl xtopology nonstop_tsc cpuid aperfmperf
     tsc_known_freq pni pclmulqdq monitor ssse3 fma cx16 pcid sse4_1 sse4_2 x2apic movbe popcnt tsc_deadline_timer
      aes xsave avx f16c rdrand hypervisor lahf_lm abm 3dnowprefetch invpcid_single pti fsgsbase tsc_adjust bmi1
     avx2 smep bmi2 erms invpcid mpx avx512f avx512dq rdseed adx smap clflushopt clwb avx512cd avx512bw avx512vl
     xsaveopt xsavec xgetbv1 xsaves ida arat pku ospke
bugs    : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit mmio_stale_data
bogomips  : 4999.99
clflush size  : 64
cache_alignment : 64
address sizes : 46 bits physical, 48 bits virtual
power management:
```

A.3 Memory details

Listing A.2: Output of the `lshw -c memory` command

```
  *-memory
       description: System Memory
       physical id: 8
       slot: System board or motherboard
       size: 384GiB
     *-bank
          description: DIMM DDR4 Static column Pseudo-static Synchronous Window DRAM 2666 MHz (0.4 ns)
          physical id: 0
          size: 384GiB
          width: 64 bits
          clock: 2666MHz (0.4ns)
```

A.4 Network details

Listing A.3: Output of the `lshw -class network` command

```
        *-network
             description: Ethernet interface
             product: Elastic Network Adapter (ENA)
             vendor: Amazon.com, Inc.
             physical id: 5
             bus info: pci@0000:00:05.0
             logical name: ens5
             version: 00
             serial: 0a:e9:7b:76:39:0f
             width: 32 bits
             clock: 33MHz
             capabilities: bus_master cap_list ethernet physical
             configuration: broadcast=yes driver=ena ip=10.0.1.56 latency=0 link=yes multicast=yes
             resources: irq:0 memory:febf4000-febf7fff memory:fe800000-fe8fffff memory:febe0000-febeffff
```

A.5 Network performance

Some of the output has been omitted for briefness. First sum is from client to server, second sum from server to client.

Listing A.4: Output of the `iperf` command

```
Processing triggers for man-db (2.8.3-2ubuntu0.1) ...
ubuntu@ip-10-0-1-53:~$ iperf -c 10.0.1.56 -r --parallel 48 -i 1 -t 2 -p 9091

Server listening on TCP port 9091
TCP window size:  128 KByte (default)
------------------------------------------------------------
------------------------------------------------------------
Client connecting to 10.0.1.56, TCP port 9091
TCP window size:  715 KByte (default)
------------------------------------------------------------
[ 51] local 10.0.1.53 port 34524 connected with 10.0.1.56 port 9091
[ 48] local 10.0.1.53 port 34520 connected with 10.0.1.56 port 9091
...
[ ID] Interval       Transfer     Bandwidth
[ 51]  0.0- 1.0 sec  47.1 MBytes   395 Mbits/sec
[ 48]  0.0- 1.0 sec  47.4 MBytes   397 Mbits/sec
[ 11]  0.0- 1.0 sec  12.2 MBytes   103 Mbits/sec
...
[SUM]  0.0- 2.0 sec  2.81 GBytes  12.0 Gbits/sec
...
[  3] local 10.0.1.53 port 9091 connected with 10.0.1.56 port 58644
[  4] local 10.0.1.53 port 9091 connected with 10.0.1.56 port 58646
...
[SUM]  0.0- 2.0 sec  4.22 GBytes  17.7 Gbits/sec
```

A.6 IO performance

Listing A.5: Output of the `fio` command

```
ubuntu@ip-10-0-1-48:/data$ fio --rw=write --ioengine=sync --fdatasync=1 --direct=1 --directory=io-test-data --
    size=2g --bs=4k --name=iotest
iotest: (g=0): rw=write, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B, ioengine=sync, iodepth=1
fio-3.1
Starting 1 process
iotest: Laying out IO file (1 file / 2048MiB)
Jobs: 1 (f=1): [W(1)][100.0%][r=0KiB/s,w=29.8MiB/s][r=0,w=7630 IOPS][eta 00m:00s]
iotest: (groupid=0, jobs=1): err= 0: pid=15076: Tue Aug  9 10:05:22 2022
  write: IOPS=7544, BW=29.5MiB/s (30.9MB/s)(2048MiB/69490msec)
    clat (usec): min=28, max=998, avg=39.56, stdev=26.94
     lat (usec): min=28, max=998, avg=39.66, stdev=26.96
    clat percentiles (usec):
     |  1.00th=[   31],  5.00th=[   32], 10.00th=[   32], 20.00th=[   33],
     | 30.00th=[   34], 40.00th=[   34], 50.00th=[   35], 60.00th=[   35],
     | 70.00th=[   36], 80.00th=[   40], 90.00th=[   56], 95.00th=[   60],
     | 99.00th=[   68], 99.50th=[  227], 99.90th=[  482], 99.95th=[  529],
     | 99.99th=[  586]
   bw (  KiB/s): min=28864, max=31144, per=100.00%, avg=30182.07, stdev=461.89, samples=138
   iops        : min= 7216, max= 7786, avg=7545.51, stdev=115.47, samples=138
  lat (usec)   : 50=83.39%, 100=16.04%, 250=0.14%, 500=0.36%, 750=0.07%
  lat (usec)   : 1000=0.01%
  cpu          : usr=2.89%, sys=12.91%, ctx=1048781, majf=0, minf=12
  IO depths    : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
     complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
     issued rwt: total=0,524288,0, short=0,0,0, dropped=0,0,0
     latency   : target=0, window=0, percentile=100.00%, depth=1

Run status group 0 (all jobs):
  WRITE: bw=29.5MiB/s (30.9MB/s), 29.5MiB/s-29.5MiB/s (30.9MB/s-30.9MB/s), io=2048MiB (2147MB), run=69490-69490
    msec

Disk stats (read/write):
    md0: ios=0/2716297, merge=0/0, ticks=0/0, in_queue=0, util=0.00%, aggrios=0/836961, aggrmerge=0/523142,
     aggrticks=0/42393, aggrin_queue=0, aggrutil=99.51%
  nvme2n1: ios=0/836936, merge=0/523124, ticks=0/42683, in_queue=0, util=99.41%
  nvme1n1: ios=0/836987, merge=0/523160, ticks=0/42103, in_queue=0, util=99.51%
ubuntu@ip-10-0-1-48:/data$
```

A.7 Datagen configuration

Listing A.6: Contents of `params-sf30.ini` used for scale factor 30

```
ldbc.snb.datagen.generator.scaleFactor:snb.interactive.30
ldbc.snb.datagen.serializer.numUpdatePartitions:32

ldbc.snb.datagen.serializer.dynamicActivitySerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.activity.CsvCompositeMergeForeignDynamicActivitySerializer
ldbc.snb.datagen.serializer.dynamicPersonSerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.person.CsvCompositeMergeForeignDynamicPersonSerializer
ldbc.snb.datagen.serializer.staticSerializer:ldbc.snb.datagen.serializer.snb.csv.staticserializer.CsvCompositeMergeForeignStaticSerializer

ldbc.snb.datagen.serializer.dateFormatter:ldbc.snb.datagen.util.formatter.LongDateFormatter
```

Listing A.7: Contents of `params-sf100.ini` used for scale factor 100

```
ldbc.snb.datagen.generator.scaleFactor:snb.interactive.100
ldbc.snb.datagen.serializer.numUpdatePartitions:32

ldbc.snb.datagen.serializer.dynamicActivitySerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.activity.CsvCompositeMergeForeignDynamicActivitySerializer
ldbc.snb.datagen.serializer.dynamicPersonSerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.person.CsvCompositeMergeForeignDynamicPersonSerializer
ldbc.snb.datagen.serializer.staticSerializer:ldbc.snb.datagen.serializer.snb.csv.staticserializer.CsvCompositeMergeForeignStaticSerializer

ldbc.snb.datagen.serializer.dateFormatter:ldbc.snb.datagen.util.formatter.LongDateFormatter
```

Listing A.8: Contents of `params-sf300.ini` used for scale factor 300

```
ldbc.snb.datagen.generator.scaleFactor:snb.interactive.300
ldbc.snb.datagen.serializer.numUpdatePartitions:32

ldbc.snb.datagen.serializer.dynamicActivitySerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.activity.CsvCompositeMergeForeignDynamicActivitySerializer
ldbc.snb.datagen.serializer.dynamicPersonSerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.person.CsvCompositeMergeForeignDynamicPersonSerializer
ldbc.snb.datagen.serializer.staticSerializer:ldbc.snb.datagen.serializer.snb.csv.staticserializer.CsvCompositeMergeForeignStaticSerializer

ldbc.snb.datagen.serializer.dateFormatter:ldbc.snb.datagen.util.formatter.LongDateFormatter
```

A.8 Import configuration

Listing A.9: Content of `import.conf` describing the data schema

```
{
    "schema": [
    {
        "label" : "Comment",
            "type" : "VERTEX",
            "properties" : [
            { "name" : "id", "type":"INT64"},
            { "name" : "creationDate", "type":"INT64"},
```

```
            { "name" : "locationIP", "type":"STRING"},
            { "name" : "browserUsed", "type":"STRING"},
            { "name" : "content", "type":"STRING"},
            { "name" : "length", "type":"INT32"},
            { "name" : "creator", "type":"INT64"},
            { "name" : "place", "type":"INT64"},
            { "name" : "replyOfPost", "type":"INT64", "optional":true},
            { "name" : "replyOfComment", "type":"INT64", "optional":true}
        ],
            "primary" : "id"
    },
    {
        "label" : "Forum",
        "type" : "VERTEX",
        "properties" : [
        { "name" : "id", "type":"INT64"},
        { "name" : "title", "type":"STRING"},
        { "name" : "creationDate", "type":"INT64"},
        { "name" : "moderator", "type":"INT64"}
        ],
            "primary" : "id"
    },
    {
        "label" : "Organisation",
        "type" : "VERTEX",
        "properties" : [
        { "name" : "id", "type":"INT64"},
        { "name" : "type", "type":"STRING"},
        { "name" : "name", "type":"STRING"},
        { "name" : "url", "type":"STRING"},
        { "name" : "place", "type":"INT64"}
        ],
            "primary" : "id"
    },
    {
        "label" : "Person",
        "type" : "VERTEX",
        "properties" : [
        { "name" : "id", "type":"INT64"},
        { "name" : "firstName", "type":"STRING"},
        { "name" : "lastName", "type":"STRING"},
        { "name" : "gender", "type":"STRING"},
        { "name" : "birthday", "type":"INT64"},
        { "name" : "creationDate", "type":"INT64"},
        { "name" : "locationIP", "type":"STRING"},
        { "name" : "browserUsed", "type":"STRING"},
        { "name" : "place", "type":"INT64"},
        { "name" : "speaks", "type":"STRING"},
        { "name" : "email", "type":"STRING"}
        ],
            "primary" : "id"
    },
    {
        "label" : "Place",
        "type" : "VERTEX",
        "properties" : [
        { "name" : "id", "type":"INT64"},
        { "name" : "name", "type":"STRING"},
```

```
        { "name" : "url", "type":"STRING"},
        { "name" : "type", "type":"STRING"},
        { "name" : "isPartOf", "type":"INT64", "optional":true}
        ],
            "primary" : "id"
    },
    {
        "label" : "Post",
        "type" : "VERTEX",
        "properties" : [
        { "name" : "id", "type":"INT64"},
        { "name" : "imageFile", "type":"STRING", "optional":true},
        { "name" : "creationDate", "type":"INT64"},
        { "name" : "locationIP", "type":"STRING"},
        { "name" : "browserUsed", "type":"STRING"},
        { "name" : "language", "type":"STRING", "optional":true},
        { "name" : "content", "type":"STRING", "optional":true},
        { "name" : "length", "type":"INT32"},
        { "name" : "creator", "type":"INT64"},
        { "name" : "container", "type":"INT64"},
        { "name" : "place", "type":"INT64"}
        ],
            "primary" : "id"
    },
    {
        "label" : "Tag",
        "type" : "VERTEX",
        "properties" : [
        { "name" : "id", "type":"INT64"},
        { "name" : "name", "type":"STRING"},
        { "name" : "url", "type":"STRING"},
        { "name" : "hasType", "type":"INT64"}
        ],
            "primary" : "id"
    },
    {
        "label" : "Tagclass",
        "type" : "VERTEX",
        "properties" : [
        { "name" : "id", "type":"INT64"},
        { "name" : "name", "type":"STRING"},
        { "name" : "url", "type":"STRING"},
        { "name" : "isSubclassOf", "type":"INT64", "optional":true}
        ],
            "primary" : "id"
    },
    {
        "label" : "commentHasCreator",
        "type" : "EDGE",
        "properties" : [
        { "name" : "creationDate", "type":"INT64"}
        ],
            "constraints" : [["Comment", "Person"]]
    },
    {
        "label" : "commentHasTag",
        "type" : "EDGE",
        "properties" : [],
```

```
        "constraints" : [["Comment", "Tag"]]
    },
    {
        "label" : "commentIsLocatedIn",
        "type" : "EDGE",
        "properties" : [
        { "name" : "creationDate", "type":"INT64"}
        ],
            "constraints" : [["Comment", "Place"]]
    },
    {
        "label" : "replyOf",
        "type" : "EDGE",
        "properties" : [
        { "name" : "creationDate", "type":"INT64"}
        ],
            "constraints" : [["Comment", "Comment"], ["Comment", "Post"]]
    },
    {
        "label" : "containerOf",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Forum", "Post"]]
    },
    {
        "label" : "hasMember",
        "type" : "EDGE",
        "primary" : "joinDate",
        "properties" : [
        { "name" : "joinDate", "type":"INT64"},
        { "name" : "numPosts", "type":"INT32"},
        { "name" : "forumId", "type":"INT64"}
        ],
            "constraints" : [["Forum", "Person"]]
    },
    {
        "label" : "hasModerator",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Forum", "Person"]]
    },
    {
        "label" : "forumHasTag",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Forum", "Tag"]]
    },
    {
        "label" : "organisationIsLocatedIn",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Organisation", "Place"]]
    },
    {
        "label" : "hasInterest",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Person", "Tag"]]
```

```
    },
    {
        "label" : "personIsLocatedIn",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Person", "Place"]]
    },
    {
        "label" : "knows",
        "type" : "EDGE",
        "properties" : [
        { "name" : "creationDate", "type":"INT64"},
        { "name" : "weight", "type":"DOUBLE"}
        ],
            "constraints" : [["Person", "Person"]]
    },
    {
        "label" : "likes",
        "type" : "EDGE",
        "properties" : [
        { "name" : "creationDate", "type":"INT64"}
        ],
            "constraints" : [["Person", "Comment"], ["Person", "Post"]]
    },
    {
        "label" : "studyAt",
        "type" : "EDGE",
        "properties" : [
        { "name" : "classYear", "type":"INT32"}
        ],
            "constraints" : [["Person", "Organisation"]]
    },
    {
        "label" : "workAt",
        "type" : "EDGE",
        "properties" : [
        { "name" : "workFrom", "type":"INT32", "optional":true}
        ],
            "constraints" : [["Person", "Organisation"]]
    },
    {
        "label" : "isPartOf",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Place", "Place"]]
    },
    {
        "label" : "postHasCreator",
        "type" : "EDGE",
        "properties" : [
        { "name" : "creationDate", "type":"INT64"}
        ],
            "constraints" : [["Post", "Person"]]
    },
    {
        "label" : "postHasTag",
        "type" : "EDGE",
        "properties" : [],
```

```
        "constraints" : [["Post", "Tag"]]
    },
    {
        "label" : "postIsLocatedIn",
        "type" : "EDGE",
        "properties" : [
        { "name" : "creationDate", "type":"INT64"}
        ],
            "constraints" : [["Post", "Place"]]
    },
    {
        "label" : "hasType",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Tag", "Tagclass"]]
    },
    {
        "label" : "isSubclassOf",
        "type" : "EDGE",
        "properties" : [],
        "constraints" : [["Tagclass", "Tagclass"]]
    }
    ],
        "files" : [
        {
            "path" : "comment.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "Comment",
            "columns" : ["id","creationDate","locationIP","browserUsed","content","length","creator","place","
    replyOfPost","replyOfComment"]
        },
        {
            "path" : "forum.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "Forum",
            "columns" : ["id","title","creationDate","moderator"]
        },
        {
            "path" : "organisation.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "Organisation",
            "columns" : ["id","type","name","url","place"]
        },
        {
            "path" : "person.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "Person",
            "columns" : ["id","firstName","lastName","gender","birthday","creationDate","locationIP","browserUsed
    ","place","speaks","email"]
        },
        {
            "path" : "place.csv",
            "header" : 0,
            "format" : "CSV",
```

```
            "label" : "Place",
            "columns" : ["id","name","url","type","isPartOf"]
        },
        {
            "path" : "post.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "Post",
            "columns" : ["id","imageFile","creationDate","locationIP","browserUsed","language","content","length
    ","creator","container","place"]
        },
        {
            "path" : "tag.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "Tag",
            "columns" : ["id","name","url","hasType"]
        },
        {
            "path" : "tagclass.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "Tagclass",
            "columns" : ["id","name","url","isSubclassOf"]
        },
        {
            "path" : "comment_hasCreator_person.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "commentHasCreator",
            "SRC_ID" : "Comment",
            "DST_ID" : "Person",
            "columns" : ["SRC_ID","creationDate","SKIP","SKIP","SKIP","SKIP","DST_ID","SKIP","SKIP","SKIP"]
        },
        {
            "path" : "comment_hasTag_tag.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "commentHasTag",
            "SRC_ID" : "Comment",
            "DST_ID" : "Tag",
            "columns" : ["SRC_ID","DST_ID"]
        },
        {
            "path" : "comment_isLocatedIn_place.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "commentIsLocatedIn",
            "SRC_ID" : "Comment",
            "DST_ID" : "Place",
            "columns" : ["SRC_ID","creationDate","SKIP","SKIP","SKIP","SKIP","SKIP","DST_ID","SKIP","SKIP"]
        },
        {
            "path" : "comment_replyOf_comment.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "replyOf",
            "SRC_ID" : "Comment",
```

```
            "DST_ID" : "Comment",
            "columns" : ["SRC_ID","DST_ID","creationDate"]
        },
        {
            "path" : "comment_replyOf_post.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "replyOf",
            "SRC_ID" : "Comment",
            "DST_ID" : "Post",
            "columns" : ["SRC_ID","DST_ID","creationDate"]
        },
        {
            "path" : "forum_containerOf_post.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "containerOf",
            "SRC_ID" : "Forum",
            "DST_ID" : "Post",
            "columns" : ["DST_ID","SKIP","SKIP","SKIP","SKIP","SKIP","SKIP","SKIP","SRC_ID","SKIP"]
        },
        {
            "path" : "forum_hasMember_person.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "hasMember",
            "SRC_ID" : "Forum",
            "DST_ID" : "Person",
            "columns" : ["SRC_ID","DST_ID","joinDate","numPosts","forumId"]
        },
        {
            "path" : "forum_hasModerator_person.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "hasModerator",
            "SRC_ID" : "Forum",
            "DST_ID" : "Person",
            "columns" : ["SRC_ID","SKIP","SKIP","DST_ID"]
        },
        {
            "path" : "forum_hasTag_tag.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "forumHasTag",
            "SRC_ID" : "Forum",
            "DST_ID" : "Tag",
            "columns" : ["SRC_ID","DST_ID"]
        },
        {
            "path" : "organisation_isLocatedIn_place.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "organisationIsLocatedIn",
            "SRC_ID" : "Organisation",
            "DST_ID" : "Place",
            "columns" : ["SRC_ID","SKIP","SKIP","SKIP","DST_ID"]
        },
        {
```

```
            "path" : "person_hasInterest_tag.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "hasInterest",
            "SRC_ID" : "Person",
            "DST_ID" : "Tag",
            "columns" : ["SRC_ID","DST_ID"]
        },
        {
            "path" : "person_isLocatedIn_place.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "personIsLocatedIn",
            "SRC_ID" : "Person",
            "DST_ID" : "Place",
            "columns" : ["SRC_ID","SKIP","SKIP","SKIP","SKIP","SKIP","SKIP","SKIP","DST_ID","SKIP","SKIP"]
        },
        {
            "path" : "person_knows_person.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "knows",
            "SRC_ID" : "Person",
            "DST_ID" : "Person",
            "columns" : ["SRC_ID","DST_ID","creationDate","weight"]
        },
        {
            "path" : "person_likes_comment.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "likes",
            "SRC_ID" : "Person",
            "DST_ID" : "Comment",
            "columns" : ["SRC_ID","DST_ID","creationDate"]
        },
        {
            "path" : "person_likes_post.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "likes",
            "SRC_ID" : "Person",
            "DST_ID" : "Post",
            "columns" : ["SRC_ID","DST_ID","creationDate"]
        },
        {
            "path" : "person_studyAt_organisation.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "studyAt",
            "SRC_ID" : "Person",
            "DST_ID" : "Organisation",
            "columns" : ["SRC_ID","DST_ID","classYear"]
        },
        {
            "path" : "person_workAt_organisation.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "workAt",
```

```
            "SRC_ID" : "Person",
            "DST_ID" : "Organisation",
            "columns" : ["SRC_ID","DST_ID","workFrom"]
        },
        {
            "path" : "place_isPartOf_place.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "isPartOf",
            "SRC_ID" : "Place",
            "DST_ID" : "Place",
            "columns" : ["SRC_ID","DST_ID"]
        },
        {
            "path" : "post_hasCreator_person.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "postHasCreator",
            "SRC_ID" : "Post",
            "DST_ID" : "Person",
            "columns" : ["SRC_ID","SKIP","creationDate","SKIP","SKIP","SKIP","SKIP","SKIP","DST_ID","SKIP","SKIP
    "]
        },
        {
            "path" : "post_hasTag_tag.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "postHasTag",
            "SRC_ID" : "Post",
            "DST_ID" : "Tag",
            "columns" : ["SRC_ID","DST_ID"]
        },
        {
            "path" : "post_isLocatedIn_place.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "postIsLocatedIn",
            "SRC_ID" : "Post",
            "DST_ID" : "Place",
            "columns" : ["SRC_ID","SKIP","creationDate","SKIP","SKIP","SKIP","SKIP","SKIP","SKIP","SKIP","DST_ID
    "]
        },
        {
            "path" : "tag_hasType_tagclass.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "hasType",
            "SRC_ID" : "Tag",
            "DST_ID" : "Tagclass",
            "columns" : ["SRC_ID","SKIP","SKIP","DST_ID"]
        },
        {
            "path" : "tagclass_isSubclassOf_tagclass.csv",
            "header" : 0,
            "format" : "CSV",
            "label" : "isSubclassOf",
            "SRC_ID" : "Tagclass",
            "DST_ID" : "Tagclass",
```

```
            "columns" : ["SRC_ID","DST_ID"]
        }
    ]
}
```

A.9 Benchmark configuration

Listing A.10: Contents of `interactive-benchmark-sf30.properties` used for scale factor 30

```
host=10.0.1.47
rest_port=7071
port=9091
user=admin
pass=73@TuGraph
is_admin=true

status=1
thread_count=48
name=LDBC-SNB
mode=execute_benchmark
results_log=true
time_unit=MICROSECONDS
time_compression_ratio=0.0028
peer_identifiers=
workload_statistics=false
spinner_wait_duration=1
help=false
ignore_scheduled_start_times=false

workload=org.ldbcouncil.snb.driver.workloads.interactive.LdbcSnbInteractiveWorkload
db=ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb
operation_count=92000000
ldbc.snb.interactive.parameters_dir=../../ldbc_snb_datagen/substitution_parameters/
ldbc.snb.interactive.updates_dir=../../ldbc_snb_datagen/social_network/
ldbc.snb.interactive.short_read_dissipation=0.2
ldbc.snb.interactive.scale_factor=30

warmup=24000000

# *** For debugging purposes ***
ldbc.snb.interactive.LdbcQuery1_enable=true
ldbc.snb.interactive.LdbcQuery2_enable=true
ldbc.snb.interactive.LdbcQuery3_enable=true
ldbc.snb.interactive.LdbcQuery4_enable=true
ldbc.snb.interactive.LdbcQuery5_enable=true
ldbc.snb.interactive.LdbcQuery6_enable=true
ldbc.snb.interactive.LdbcQuery7_enable=true
ldbc.snb.interactive.LdbcQuery8_enable=true
ldbc.snb.interactive.LdbcQuery9_enable=true
ldbc.snb.interactive.LdbcQuery10_enable=true
ldbc.snb.interactive.LdbcQuery11_enable=true
ldbc.snb.interactive.LdbcQuery12_enable=true
ldbc.snb.interactive.LdbcQuery13_enable=true
ldbc.snb.interactive.LdbcQuery14_enable=true

ldbc.snb.interactive.LdbcShortQuery1PersonProfile_enable=true
ldbc.snb.interactive.LdbcShortQuery2PersonPosts_enable=true
ldbc.snb.interactive.LdbcShortQuery3PersonFriends_enable=true
ldbc.snb.interactive.LdbcShortQuery4MessageContent_enable=true
ldbc.snb.interactive.LdbcShortQuery5MessageCreator_enable=true
ldbc.snb.interactive.LdbcShortQuery6MessageForum_enable=true
ldbc.snb.interactive.LdbcShortQuery7MessageReplies_enable=true

```

```
ldbc.snb.interactive.LdbcUpdate1AddPerson_enable=true
ldbc.snb.interactive.LdbcUpdate2AddPostLike_enable=true
ldbc.snb.interactive.LdbcUpdate3AddCommentLike_enable=true
ldbc.snb.interactive.LdbcUpdate4AddForum_enable=true
ldbc.snb.interactive.LdbcUpdate5AddForumMembership_enable=true
ldbc.snb.interactive.LdbcUpdate6AddPost_enable=true
ldbc.snb.interactive.LdbcUpdate7AddComment_enable=true
ldbc.snb.interactive.LdbcUpdate8AddFriendship_enable=true
```

Listing A.11: Contents of `interactive-benchmark-sf100.properties` used for scale factor 100

```
host=10.0.1.58
rest_port=7071
port=9091
user=admin
pass=73@TuGraph
is_admin=true

status=1
thread_count=48
name=LDBC-SNB
mode=execute_benchmark
results_log=true
time_unit=MICROSECONDS
time_compression_ratio=0.0104
peer_identifiers=
workload_statistics=false
spinner_wait_duration=1
help=false
ignore_scheduled_start_times=false

workload=org.ldbcouncil.snb.driver.workloads.interactive.LdbcSnbInteractiveWorkload
db=ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb
operation_count=97000000
ldbc.snb.interactive.parameters_dir=../../ldbc_snb_datagen/substitution_parameters/
ldbc.snb.interactive.updates_dir=../../ldbc_snb_datagen/social_network/
ldbc.snb.interactive.short_read_dissipation=0.2
ldbc.snb.interactive.scale_factor=100

warmup=23000000

# *** For debugging purposes ***
ldbc.snb.interactive.LdbcQuery1_enable=true
ldbc.snb.interactive.LdbcQuery2_enable=true
ldbc.snb.interactive.LdbcQuery3_enable=true
ldbc.snb.interactive.LdbcQuery4_enable=true
ldbc.snb.interactive.LdbcQuery5_enable=true
ldbc.snb.interactive.LdbcQuery6_enable=true
ldbc.snb.interactive.LdbcQuery7_enable=true
ldbc.snb.interactive.LdbcQuery8_enable=true
ldbc.snb.interactive.LdbcQuery9_enable=true
ldbc.snb.interactive.LdbcQuery10_enable=true
ldbc.snb.interactive.LdbcQuery11_enable=true
ldbc.snb.interactive.LdbcQuery12_enable=true
ldbc.snb.interactive.LdbcQuery13_enable=true
ldbc.snb.interactive.LdbcQuery14_enable=true

ldbc.snb.interactive.LdbcShortQuery1PersonProfile_enable=true
```

```
ldbc.snb.interactive.LdbcShortQuery2PersonPosts_enable=true
ldbc.snb.interactive.LdbcShortQuery3PersonFriends_enable=true
ldbc.snb.interactive.LdbcShortQuery4MessageContent_enable=true
ldbc.snb.interactive.LdbcShortQuery5MessageCreator_enable=true
ldbc.snb.interactive.LdbcShortQuery6MessageForum_enable=true
ldbc.snb.interactive.LdbcShortQuery7MessageReplies_enable=true

ldbc.snb.interactive.LdbcUpdate1AddPerson_enable=true
ldbc.snb.interactive.LdbcUpdate2AddPostLike_enable=true
ldbc.snb.interactive.LdbcUpdate3AddCommentLike_enable=true
ldbc.snb.interactive.LdbcUpdate4AddForum_enable=true
ldbc.snb.interactive.LdbcUpdate5AddForumMembership_enable=true
ldbc.snb.interactive.LdbcUpdate6AddPost_enable=true
ldbc.snb.interactive.LdbcUpdate7AddComment_enable=true
ldbc.snb.interactive.LdbcUpdate8AddFriendship_enable=true
```

Listing A.12: Contents of `interactive-benchmark-sf300.properties` used for scale factor 300

```
host=10.0.1.37
rest_port=7071
port=9091
user=admin
pass=73@TuGraph
is_admin=true

status=1
thread_count=48
name=LDBC-SNB
mode=execute_benchmark
results_log=true
time_unit=MICROSECONDS
time_compression_ratio=0.036
peer_identifiers=
workload_statistics=false
spinner_wait_duration=1
help=false
ignore_scheduled_start_times=false

workload=org.ldbcouncil.snb.driver.workloads.interactive.LdbcSnbInteractiveWorkload
db=ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb
operation_count=95000000
ldbc.snb.interactive.parameters_dir=../../ldbc_snb_datagen/substitution_parameters/
ldbc.snb.interactive.updates_dir=../../ldbc_snb_datagen/social_network/
ldbc.snb.interactive.short_read_dissipation=0.2
ldbc.snb.interactive.scale_factor=300

warmup=23000000

# *** For debugging purposes ***
ldbc.snb.interactive.LdbcQuery1_enable=true
ldbc.snb.interactive.LdbcQuery2_enable=true
ldbc.snb.interactive.LdbcQuery3_enable=true
ldbc.snb.interactive.LdbcQuery4_enable=true
ldbc.snb.interactive.LdbcQuery5_enable=true
ldbc.snb.interactive.LdbcQuery6_enable=true
ldbc.snb.interactive.LdbcQuery7_enable=true
ldbc.snb.interactive.LdbcQuery8_enable=true
ldbc.snb.interactive.LdbcQuery9_enable=true
```

```
ldbc.snb.interactive.LdbcQuery10_enable=true
ldbc.snb.interactive.LdbcQuery11_enable=true
ldbc.snb.interactive.LdbcQuery12_enable=true
ldbc.snb.interactive.LdbcQuery13_enable=true
ldbc.snb.interactive.LdbcQuery14_enable=true

ldbc.snb.interactive.LdbcShortQuery1PersonProfile_enable=true
ldbc.snb.interactive.LdbcShortQuery2PersonPosts_enable=true
ldbc.snb.interactive.LdbcShortQuery3PersonFriends_enable=true
ldbc.snb.interactive.LdbcShortQuery4MessageContent_enable=true
ldbc.snb.interactive.LdbcShortQuery5MessageCreator_enable=true
ldbc.snb.interactive.LdbcShortQuery6MessageForum_enable=true
ldbc.snb.interactive.LdbcShortQuery7MessageReplies_enable=true

ldbc.snb.interactive.LdbcUpdate1AddPerson_enable=true
ldbc.snb.interactive.LdbcUpdate2AddPostLike_enable=true
ldbc.snb.interactive.LdbcUpdate3AddCommentLike_enable=true
ldbc.snb.interactive.LdbcUpdate4AddForum_enable=true
ldbc.snb.interactive.LdbcUpdate5AddForumMembership_enable=true
ldbc.snb.interactive.LdbcUpdate6AddPost_enable=true
ldbc.snb.interactive.LdbcUpdate7AddComment_enable=true
ldbc.snb.interactive.LdbcUpdate8AddFriendship_enable=true
```

A.10 Validation configuration

Listing A.13: The contents of `interactive-validate.properties`

```
host=localhost
rest_port=7071
port=9091
user=admin
pass=73@TuGraph
is_admin=true

status=1
thread_count=1
mode=validate_database
name=LDBC-SNB
results_log=true
time_unit=MICROSECONDS
time_compression_ratio=0.001
peer_identifiers=
workload_statistics=false
spinner_wait_duration=1
help=false
ignore_scheduled_start_times=true

workload=org.ldbcouncil.snb.driver.workloads.interactive.LdbcSnbInteractiveWorkload
db=ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb
operation_count=10000

validate_database=validation_params.csv
ldbc.snb.interactive.parameters_dir=../../ldbc_snb_datagen/substitution_parameters/
ldbc.snb.interactive.short_read_dissipation=0.2
ldbc.snb.interactive.scale_factor=10

```

```
# *** For debugging purposes ***
ldbc.snb.interactive.LdbcQuery1_enable=true
ldbc.snb.interactive.LdbcQuery2_enable=true
ldbc.snb.interactive.LdbcQuery3_enable=true
ldbc.snb.interactive.LdbcQuery4_enable=true
ldbc.snb.interactive.LdbcQuery5_enable=true
ldbc.snb.interactive.LdbcQuery6_enable=true
ldbc.snb.interactive.LdbcQuery7_enable=true
ldbc.snb.interactive.LdbcQuery8_enable=true
ldbc.snb.interactive.LdbcQuery9_enable=true
ldbc.snb.interactive.LdbcQuery10_enable=true
ldbc.snb.interactive.LdbcQuery11_enable=true
ldbc.snb.interactive.LdbcQuery12_enable=true
ldbc.snb.interactive.LdbcQuery13_enable=true
ldbc.snb.interactive.LdbcQuery14_enable=true

ldbc.snb.interactive.LdbcShortQuery1PersonProfile_enable=true
ldbc.snb.interactive.LdbcShortQuery2PersonPosts_enable=true
ldbc.snb.interactive.LdbcShortQuery3PersonFriends_enable=true
ldbc.snb.interactive.LdbcShortQuery4MessageContent_enable=true
ldbc.snb.interactive.LdbcShortQuery5MessageCreator_enable=true
ldbc.snb.interactive.LdbcShortQuery6MessageForum_enable=true
ldbc.snb.interactive.LdbcShortQuery7MessageReplies_enable=true

ldbc.snb.interactive.LdbcUpdate1AddPerson_enable=true
ldbc.snb.interactive.LdbcUpdate2AddPostLike_enable=true
ldbc.snb.interactive.LdbcUpdate3AddCommentLike_enable=true
ldbc.snb.interactive.LdbcUpdate4AddForum_enable=true
ldbc.snb.interactive.LdbcUpdate5AddForumMembership_enable=true
ldbc.snb.interactive.LdbcUpdate6AddPost_enable=true
ldbc.snb.interactive.LdbcUpdate7AddComment_enable=true
ldbc.snb.interactive.LdbcUpdate8AddFriendship_enable=true
```