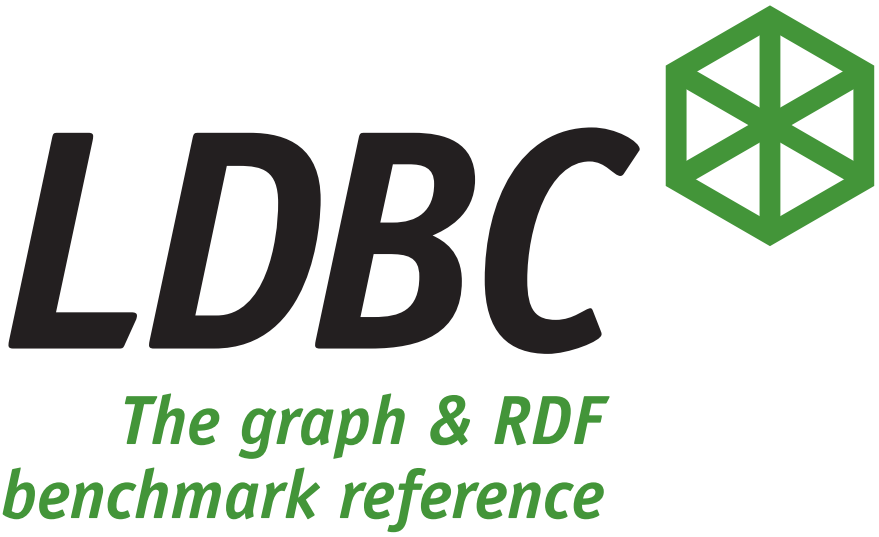

Full Disclosure Report of the LDBC Social Network Benchmark

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

January 28, 2023

General Terms

Executive Summary

This document describes an implementation of the LDBC Social Network Benchmark's Interactive workload in TuGraph (formerly known as LightGraph), a graph database management system developed by the Ant Group.

TuGraph passed two successful LDBC SNB Interactive audits in the past. The first audit was in July 2020 while the product was developed by FMA before Ant Group acquired FMA, while the second audit was in August 2022. Both of these audits were conducted in the Amazon Web Services (AWS) cloud. The current audited was performed in the Alibaba Cloud (a.k.a. Aliyun) using TuGraph's most recent version, v3.3.4, which is available as open-source software under the Apache Software License v2.0.

The benchmark implementation uses stored procedures written in C++14, which are compiled and loaded into the database as plugins. The 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. Compared to the previous audits, the benchmark uses further optimizations in the implementations of the benchmark queries.

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:
BD0CE383811840E...
1/28/2023

Dr. Márton Búr
(Auditor)

Date

DocuSigned by:
Gabor Szarnyas
FA33C2C8CB944A5...
1/28/2023

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

Date

DocuSigned by:
Chuntao Hong
3FE773CDAEA8496...
1/28/2023

Dr. Chuntao Hong
(Test Sponsor Representative)

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 Alibaba Cloud dashboard (Instance Details page).

Table 1.1: Machine Type and Location

Cloud provider	Alibaba Cloud (also known as Aliyun)
Machine region	China (Shenzhen), Shenzhen Zone F
Common name of the item	ecs.g8y.16xlarge
Operating system	Alibaba Cloud Linux 3 (Soaring Falcon)

This benchmark used two `ecs.g8y.16xlarge` instances, one for the driver and one for the system under test (SUT). Both machines were assigned to the same VPC with the same subnetwork. This is shown in Figure 1.1.

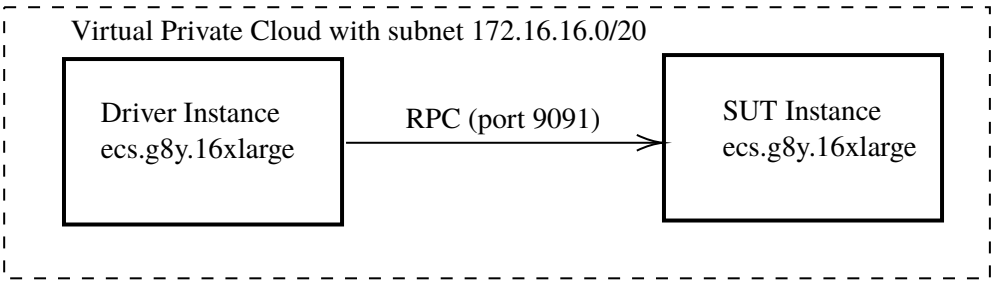

Figure 1.1: Overview of benchmark setup

1.1.2 CPU details

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

Table 1.2: CPU details summary

Type	ARM-based YiTian 710
Total number	1
Cores per CPU	64
Threads per CPU core	1
CPU clock frequency	2.75 GHz
Total cache size per CPU	L1 cache: 64KiB Data + 64KiB Instruction L2 cache: 1MiB L3 cache: 64MiB

1.1.3 Memory details

The total size of the memory installed is 256GiB and the type of memory is DDR4. 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, as the `lshw -c storage -c disk` command returned no important details. The instance has multiple disks attached. The NVMe SSD device (`/dev/nvme1n1`) is the one used during the benchmarking, i.e., the database as well as the update streams and query parameters were stored here.

The file system type was `ext4`. The 4KB QD1 write performance was measured with the `fio` command and the output (Listing A.5) showed an average of 3431 IOPS.

1.1.5 Network details

The benchmark was run using two `ecs.g8y.16xlarge` instances, both deployed in the same availability zone behind a Virtual Private Cloud (VPC) configured with subnetwork `172.16.16.0/20`. Both instances had the following ports open:

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

The `ecs.g8y.16xlarge` instances use a Virtio network device. 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 33.4 Gbit/sec from client to server and 31.6 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). The TuGraph database software is freely available under the Apache Software License v2.0, hence its permanent license cost is 0.00 RMB.

Table 1.3: Pricing summary

Item	**Price**
ecs.g8y.16xlarge reserved instance machine in Alibaba Cloud (standard 3-year term)	220 541.76 RMB
Permanent TuGraph license	0.00 RMB
Maintenance fee (3 years)	57 000.00 RMB
Total cost of ownership	277 541.76 RMB

1.1.7 System availability

The latest software version of TuGraph (version 3.3.4) was made available on January 20, 2023. This version was deployed to the machine described in this section. TuGraph is an open-source software released under the Apache Software License 2.0, the used release is available online on GitHub[1].

[1] `https://github.com/TuGraph-db/tugraph-db/releases/tag/v3.3.4`

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 it does 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	379.473	268.533	648.006
100	1 333.226	990.350	2 323.576
300	3 791.617	7 957.374	11 748.991

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.tugraph.interactive.TuGraphInteractiveDb` extends the class `com.ldbc.driver.Db` provided in the LDBC SNB Interactive driver package. Internally, the `TuGraphInteractiveDb` 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=0.00220 for scale factor 30,
- TCR=0.00781 for scale factor 100 and
- TCR=0.03400 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. In each case, the query on-time compliance is higher than the minimum required 95% [1]. The performance summary tables below highlight key performance characteristics.

Table 4.1: Summary of results for scale factor 30

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 01m 29.605s	117 603 754	16 133.08 $\frac{\text{operations}}{\text{second}}$	99.98%

Table 4.2: Summary of results for scale factor 100

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 00m 26.624s	122 608 813	16 966.26 $\frac{\text{operations}}{\text{second}}$	95.63%

Table 4.3: Summary of results for scale factor 300

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 02m 51.486s	99 755 499	13 532.62 $\frac{\text{operations}}{\text{second}}$	96.31%

During the benchmark run, the query execution times 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.

[1] The total number of late operations for each run in the results in the attachment is referred to as `excessive_delay_count`.

Table 4.4: Detailed performance benchmark results for scale factor 30 in microseconds

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	816 897	4 242	492 016	25 878.38	18 402	23 647	27 958	348 656
Query2	574 035	775	53 498	3 479.38	2 959	4 727	6 983	13 252
Query3	200 371	9 717	87 380	25 522.53	26 483	30 807	34 380	44 508
Query4	589 981	362	66 712	5 621.41	4 724	9 476	11 261	16 554
Query5	294 991	9 188	169 240	43 090.97	41 962	55 096	60 338	72 452
Query6	67 213	630	64 988	10 867.80	13 472	22 151	24 261	31 013
Query7	442 485	178	32 544	1 327.87	744	2 480	4 984	11 247
Query8	2 359 924	1 315	174 456	5 169.40	3 156	11 706	15 648	25 751
Query9	55 311	58 660	468 112	308 347.39	308 624	357 856	369 248	394 880
Query10	574 035	4 203	271 824	22 256.22	19 885	32 073	38 294	56 314
Query11	1 061 966	465	60 054	2 097.65	1 521	3 180	5 857	12 172
Query12	482 712	4 761	680 224	45 048.35	36 504	79 212	100 236	161 584
Query13	1 117 858	161	52 130	1 573.15	963	2 719	5 342	11 603
Query14	433 455	204	345 360	8 875.57	2 932	8 619	69 868	112 300
Short1	11 388 053	113	70 124	847.26	457	1 138	2 520	10 087
Short2	11 388 053	152	70 232	850.35	599	1 127	1 682	6 989
Short3	11 388 053	188	75 224	1 809.55	797	3 832	8 857	15 066
Short4	11 386 934	110	69 996	919.06	463	1 265	3 405	10 779
Short5	11 386 934	116	61 518	692.77	452	936	1 474	6 475
Short6	11 386 934	120	70 052	692.58	455	928	1 446	6 390
Short7	11 386 934	130	70 124	753.31	513	994	1 526	6 772
Update1	6 539	709	55 862	4 117.17	2 435	9 955	13 655	22 010
Update2	4 103 742	511	128 828	4 170.18	2 478	9 942	14 678	23 686
Update3	5 501 947	468	128 832	4 136.21	2 474	9 731	14 560	23 573
Update4	112 802	570	116 008	3 132.92	1 740	7 210	11 966	19 101
Update5	13 087 660	543	131 856	4 435.31	2 659	10 594	15 216	24 458
Update6	1 455 003	546	122 296	2 905.29	1 698	5 873	10 961	17 989
Update7	4 114 646	576	122 532	3 214.69	1 837	7 163	12 022	19 186
Update8	438 286	691	174 448	9 180.28	5 212	19 475	28 087	67 132

Table 4.5: Detailed performance benchmark results for scale factor 100 in microseconds

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	741 365	196	1 005 632	66 841.15	23 947	31 610	660 128	759 136
Query2	520 960	205	58 508	3 818.53	3 267	5 276	7 909	14 195
Query3	156 712	23 464	198 184	72 397.64	76 676	84 948	90 168	127 752
Query4	535 430	241	56 494	6 002.96	5 089	10 072	11 760	17 133
Query5	247 122	1 154	193 672	49 135.22	48 434	61 216	66 292	78 588
Query6	44 414	353	111 180	28 082.58	11 786	62 382	66 568	74 408
Query7	507 250	163	50 182	1 245.96	679	2 229	4 894	11 349
Query8	3 855 097	211	79 036	2 200.09	728	7 227	10 998	16 473
Query9	36 576	202	632 736	419 591.25	421 664	488 976	502 464	525 632
Query10	481 887	327	364 256	30 721.87	27 493	44 564	52 686	80 024
Query11	876 158	292	55 974	2 261.88	1 700	3 212	6 092	12 404
Query12	438 079	318	987 104	65 569.03	49 504	124 572	160 728	269 120
Query13	1 014 499	165	52 940	2 080.79	1 579	3 296	6 058	12 260
Query14	393 377	245	278 224	20 280.45	3 356	64 448	74 508	101 200
Short1	12 364 119	116	55 248	745.53	431	899	1 641	8 884
Short2	12 364 119	132	56 242	805.06	593	1 027	1 351	6 060
Short3	12 364 119	186	72 404	1 809.48	804	3 630	9 037	15 411
Short4	12 363 896	110	55 824	819.24	438	959	2 510	10 169
Short5	12 363 896	114	54 940	626.96	427	804	1 110	5 635
Short6	12 363 896	116	55 840	628.44	431	802	1 094	5 588
Short7	12 363 896	132	55 156	696.20	497	877	1 179	5 833
Update1	4 931	742	70 140	4 360.34	2 601	9 982	13 822	21 481
Update2	3 691 749	490	156 560	4 516.92	2 794	10 361	14 819	22 652
Update3	6 494 613	503	156 576	4 521.95	2 865	10 171	14 705	22 360
Update4	83 735	607	70 800	3 412.61	1 888	8 029	12 374	18 882
Update5	10 138 236	550	157 144	4 702.54	2 843	11 107	15 408	23 574
Update6	1 151 564	575	148 552	3 444.04	1 965	7 641	12 092	18 576
Update7	4 265 300	579	149 152	3 834.93	2 567	7 982	12 258	18 757
Update8	381 818	636	204 888	9 013.25	5 870	19 678	26 582	43 098

Table 4.6: Detailed performance benchmark results for scale factor 300 in microseconds

Query	Total count	Min.	Max.	Mean	P_{50}	P_{90}	P_{95}	P_{99}
Query1	486 730	175	2 127 104	45 491.16	27 872	34 672	37 110	1 296 512
Query2	342 027	184	82 180	3 548.68	3 319	4 442	5 624	11 229
Query3	89 120	63 646	1 009 728	208 519.33	220 760	260 688	329 360	419 712
Query4	351 527	216	68 312	7 125.23	6 201	12 138	13 382	17 301
Query5	150 654	452	225 808	61 464.55	61 092	74 712	79 196	92 592
Query6	21 819	272	360 816	86 821.98	10 003	200 768	216 912	241 952
Query7	395 468	170	57 912	1 037.78	667	1 630	2 996	8 476
Query8	4 218 326	232	71 716	1 224.51	649	1 551	4 493	12 603
Query9	17 950	268	1 104 704	542 760.02	539 776	659 744	690 528	771 392
Query10	287 613	331	263 040	30 868.03	29 184	42 378	48 602	66 200
Query11	527 291	321	58 716	2 230.85	1 914	2 682	4 173	9 755
Query12	287 613	352	2 485 888	89 473.15	64 724	175 552	234 536	405 968
Query13	666 052	192	58 432	2 614.54	2 615	3 898	4 903	10 414
Query14	258 265	234	395 136	26 666.23	5 463	63 464	73 820	103 128
Short1	10 166 682	114	59 146	566.76	409	772	1 023	4 222
Short2	10 166 682	145	60 572	742.80	608	1 032	1 271	3 984
Short3	10 166 682	179	80 864	1 218.97	786	1 945	2 877	11 556
Short4	10 167 146	119	81 532	637.81	448	878	1 176	4 761
Short5	10 167 146	124	60 192	570.50	438	825	1 038	3 589
Short6	10 167 146	123	60 532	621.42	480	919	1 155	3 991
Short7	10 167 146	146	62 498	643.05	510	902	1 120	3 971
Update1	2 850	778	203 104	5 759.32	4 504	10 772	14 635	23 121
Update2	2 608 328	514	309 200	5 259.24	3 923	10 233	13 530	22 431
Update3	4 256 436	500	309 184	5 309.97	4 084	10 259	13 515	22 394
Update4	48 749	595	223 904	3 663.19	2 245	7 541	10 120	18 133
Update5	6 586 437	535	309 184	5 130.22	3 509	10 146	13 450	22 467
Update6	1 190 651	599	307 840	5 985.52	4 667	12 968	15 715	23 214
Update7	5 536 489	574	307 856	4 426.29	3 553	8 348	10 891	19 163
Update8	256 474	637	370 224	14 778.02	10 027	33 482	43 490	66 056

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 150 038 operations was created with the SNB Interactive reference implementation over Neo4j, running the Community Edition of version 5.2.0. 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, the test execution returned 100% success: no atomicity or isolation tests failed with serializable isolation level transaction settings. In particular, the following isolation levels were tested successfully:

- Dirty Writes (G0)
- Aborted Reads (G1A)
- Intermediate Reads (G1B)
- Circular Information Flow (G1C)
- Item-Many-Preceders (IMP)
- Predicate-Many-Preceders (PMP)
- Observed Transaction Vanishes (OTV)
- Fractured Reads (FR)
- Lost Updates (LU)
- Write Skews (WS)

6.3 Recovery and durability

6.3.1 Recovery

Durability tests were using the regular benchmark workload with scale factor 30 and started a few minutes after 06:00AM server time. Both server machines where shutdown using the command `sudo shutdown -rf 8:00`, forced rebooting (ungracefully) both machines 2 hours after start of the benchmark, with 44 213 738 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|1674218700014|1674218700014|1811|0|1350822622755

$ grep LdbcUpdate2 LDBC-SNB-results_log.csv tail -n 1
LdbcUpdate2AddPostLike|1674218699945|1674218700033|1093|0|1350822589973

$ grep LdbcUpdate3 LDBC-SNB-results_log.csv tail -n 1
LdbcUpdate3AddCommentLike|1674218700033|1674218700033|1148|0|1350822631944

$ grep LdbcUpdate4 LDBC-SNB-results_log.csv tail -n 1
LdbcUpdate4AddForum|1674218699968|1674218699968|1567|0|1350822600741

$ grep LdbcUpdate5 LDBC-SNB-results_log.csv tail -n 1
LdbcUpdate5AddForumMembership|1674218699888|1674218700031|2489|0|1350822562835

$ grep LdbcUpdate6 LDBC-SNB-results_log.csv tail -n 1
LdbcUpdate6AddPost|1674218700023|1674218700025|1483|0|1350822626741

$ grep LdbcUpdate7 LDBC-SNB-results_log.csv tail -n 1
LdbcUpdate7AddComment|1674218699944|1674218700032|1297|0|1350822589535

$ grep LdbcUpdate8 LDBC-SNB-results_log.csv tail -n 1
LdbcUpdate8AddFriendship|1674218700029|1674218700029|4266|0|1350822630179
```

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 TCR multiplier (the latter 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/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/ -e '1350822622755|.*|1|.*'
/data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/updateStream_0_26_person.csv
    :201:1350822622755|0|1|35184372090839|John|Santos|male|590976000000|1350822622755|110.5.68.38|Chrome|819|en;
    tl|John35184372090839@ausi.com;John35184372090839@gmx.com
    |138;141;283;565;781;807;1175;1182;1551;1677;1993;1995;2784;2806;2835;2851;2902;2913;2990;3030;3081;5114;
    5147;5176;5381;6086;6165;6940;6980;7521;9317;9329;9633;10542;13177;13359|5530,2006|951,2007;956,2008

$ grep -rnw /data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/ -e '1350822589973|.*|2|.*'
/data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/updateStream_0_8_forum.csv
    :935854:1350822589973|1335882766026|2|28587302461301|35184381869814|1350822589973

$ grep -rnw /data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/ -e '1350822631944|.*|3|.*'
/data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/updateStream_0_28_forum.csv
    :932015:1350822631944|1348362511973|3|35184372191371|30786371900977|1350822631944

$ grep -rnw /data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/ -e '1350822600741|.*|4|.*'
/data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/updateStream_0_22_forum.csv
    :876171:1350822600741|1331078927440|4|35184379460092|Album 8 of Scott Smith|1350822600741|26388279089128|2804

$ grep -rnw /data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/ -e '1350822562835|.*|5|.*'
/data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/updateStream_0_0_forum.csv
    :928519:1350822562835|1296195715670|5|28587303423943|13194139656712|1350822562835

$ grep -rnw /data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/ -e '1350822626741|.*|6|.*'
/data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/updateStream_0_22_forum.csv
    :876192:1350822626741|1331078927440|6|35184413844659|photo35184413844659.jpg|1350822626741|24.52.147.161|
    Firefox|||0|26388279089128|35184379460092|57|

$ grep -rnw /data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/ -e '1350822589535|.*|7|.*'
/data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/updateStream_0_8_forum.csv
    :935853:1350822589535|1308990772693|7|35184516104099|1350822589535|192.160.50.65|Internet Explorer|About
    Iraqi Kurdistan, of Iraq, and establishes Arabic and Kurdish as Iraq's joi
    |80|17592186163970|49|-1|35184516104096|11714

$ grep -rnw /data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/ -e '1350822630179|.*|8|.*'
/data/audit/tugraph_ldbc_snb/deps/ldbc_snb_datagen_hadoop/social_network/updateStream_0_21_forum.csv
    :924327:1350822630179|1350264281399|8|15393162805079|35184372140886|1350822630179
```

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, confirming that the system successfully performed the recovery. These queries are included in the `recovery_queries.cpp` attachment.

6.3.3 Consistency after recovery

The provided `check_consistency.cpp` 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
`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
`run_datagen.sh`	Utility script to run LDBC datagen
`TuGraph-3.3.4-1.el7.aarch64.rpm`	Linux installer package of database

The attachment folder directory structure is as follows:

```
attachments
├── TuGraph-3.3.4-1.el7.aarch64.rpm
├── results-sf30.tar.gz
├── results-sf100.tar.gz
├── results-sf300.tar.gz
└── tugraph_ldbc_snb
    ├── README.md
    ├── NOTES.md
    ├── LEGAL.md
    ├── convert_csvs.sh
    ├── deps
    │   └── ldbc_snb_interactive_impls.zip
    ├── import_data.sh
    ├── lgraph_standalone.json
    ├── load-scripts
    │   ├── convert.py
    │   └── import-data
    │       └── import.conf
    ├── params-sf10.ini
    ├── params-sf30.ini
    ├── params-sf100.ini
    ├── params-sf300.ini
    ├── plugins
    │   ├── «query».cpp
    │   ├── acid.cpp
    │   ├── check_consistency.cpp
    │   ├── compile_embedded.sh
    │   ├── compile_plugin.sh
    │   ├── generate_snb_constants.cpp
    │   ├── install.py
    │   ├── install.sh
    │   ├── preprocess.cpp
    │   ├── recovery_queries.cpp
    │   ├── snb_common.h
    │   └── snb_constants.h
    └── run_datagen.sh
```

A Appendix

A.1 CPU details

Listing A.1: Output of the `lscpu` command for one core

```
Architecture:        aarch64
Byte Order:          Little Endian
CPU(s):              64
On-line CPU(s) list: 0-63
Thread(s) per core:  1
Core(s) per socket:  64
Socket(s):           1
NUMA node(s):        1
Vendor ID:           ARM
BIOS Vendor ID:      Alibaba Cloud
Model:               0
BIOS Model name:     virt-rhel7.6.0
Stepping:            r0p0
CPU max MHz:         2750.0000
CPU min MHz:         2750.0000
BogoMIPS:            100.00
L1d cache:           64K
L1i cache:           64K
L2 cache:            1024K
L3 cache:            65536K
NUMA node0 CPU(s):   0-63
Flags:               fp asimd evtstrm aes pmull sha1 sha2 crc32 atomics fphp asimdhp cpuid asimdrdm jscvt fcma
    lrcpc dcpop sha3 sm3 sm4 asimddp sha512 sve asimdfhm dit uscat ilrcpc flagm ssbs sb dcpodp sve2 sveaes
    svepmull svebitperm svesha3 svesm4 flagm2 frint svei8mm svebf16 i8mm bf16 dgh
```

A.2 Memory details

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

```
  *-memory
       description: System Memory
       physical id: 1000
       size: 256GiB
       capabilities: ecc
       configuration: errordetection=multi-bit-ecc
     *-bank:0
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 0
          slot: DIMM 0
          size: 16GiB
     *-bank:1
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 1
          slot: DIMM 1
          size: 16GiB
     *-bank:2
          description: DIMM RAM
```

```
          vendor: Alibaba Cloud
          physical id: 2
          slot: DIMM 2
          size: 16GiB
     *-bank:3
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 3
          slot: DIMM 3
          size: 16GiB
     *-bank:4
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 4
          slot: DIMM 4
          size: 16GiB
     *-bank:5
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 5
          slot: DIMM 5
          size: 16GiB
     *-bank:6
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 6
          slot: DIMM 6
          size: 16GiB
     *-bank:7
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 7
          slot: DIMM 7
          size: 16GiB
     *-bank:8
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 8
          slot: DIMM 8
          size: 16GiB
     *-bank:9
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: 9
          slot: DIMM 9
          size: 16GiB
     *-bank:10
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: a
          slot: DIMM 10
          size: 16GiB
     *-bank:11
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: b
          slot: DIMM 11
          size: 16GiB
```

```
     *-bank:12
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: c
          slot: DIMM 12
          size: 16GiB
     *-bank:13
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: d
          slot: DIMM 13
          size: 16GiB
     *-bank:14
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: e
          slot: DIMM 14
          size: 16GiB
     *-bank:15
          description: DIMM RAM
          vendor: Alibaba Cloud
          physical id: f
          slot: DIMM 15
          size: 16GiB
  *-firmware
       description: BIOS
       vendor: EFI Development Kit II / OVMF
       physical id: 0
       version: 0.0.0
       date: 02/06/2015
       size: 96KiB
       capabilities: uefi virtualmachine
```

A.3 Network details

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

```
    *-network
    description: Ethernet controller
    product: Virtio network device
    vendor: Red Hat, Inc.
    physical id: 0
    bus info: pci@0000:05:00.0
    version: 00
    width: 64 bits
    clock: 33MHz
    capabilities: msix bus_master cap_list
    configuration: driver=virtio-pci latency=0
    resources: iomemory:a00-9ff iomemory:a00-9ff irq:0 memory:a000804000-a000804fff memory:a000800000-a000803fff
  *-virtio1
       description: Ethernet interface
       physical id: 0
       bus info: virtio@1
       logical name: eth0
       serial: 00:16:3e:1a:46:61
       capabilities: ethernet physical
       configuration: autonegotiation=off broadcast=yes driver=virtio_net driverversion=1.0.0 ip=172.16.18.134
     link=yes multicast=yes
```

A.4 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

```
# iperf -c 172.16.18.134 -r --parallel 48 -i 1 -t 2 -p 9091
------------------------------------------------------------
Server listening on TCP port 9091
TCP window size:  128 KByte (default)
------------------------------------------------------------
[ 30] local 172.16.18.136 port 54884 connected with 172.16.18.134 port 9091
[ 11] local 172.16.18.136 port 54786 connected with 172.16.18.134 port 9091
...
[ ID] Interval       Transfer     Bandwidth
[ 43] 0.00-1.00 sec  53.5 MBytes   449 Mbits/sec
[ 30] 0.00-1.00 sec  23.6 MBytes   198 Mbits/sec
...
[SUM] 0.00-1.00 sec  3.89 GBytes  33.4 Gbits/sec
...
[ 49] local 172.16.18.136 port 9091 connected with 172.16.18.134 port 48434
[ 50] local 172.16.18.136 port 9091 connected with 172.16.18.134 port 48440
...
[SUM] 1.00-2.00 sec  3.68 GBytes  31.6 Gbits/sec
```

A.5 IO performance

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

```
$ 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-3.22
Starting 1 process
iotest: Laying out IO file (1 file / 2048MiB)
Jobs: 1 (f=1): [W(1)][100.0%][w=13.6MiB/s][w=3478 IOPS][eta 00m:00s]
iotest: (groupid=0, jobs=1): err= 0: pid=275496: Tue Jan 17 23:33:56 2023
  write: IOPS=3431, BW=13.4MiB/s (14.1MB/s)(2048MiB/152803msec); 0 zone resets
    clat (usec): min=75, max=6606, avg=90.14, stdev=31.45
     lat (usec): min=75, max=6607, avg=90.21, stdev=31.45
    clat percentiles (usec):
     |  1.00th=[   80],  5.00th=[   81], 10.00th=[   82], 20.00th=[   83],
     | 30.00th=[   84], 40.00th=[   85], 50.00th=[   86], 60.00th=[   87],
     | 70.00th=[   88], 80.00th=[   90], 90.00th=[   96], 95.00th=[  106],
     | 99.00th=[  198], 99.50th=[  262], 99.90th=[  461], 99.95th=[  578],
     | 99.99th=[  898]
   bw (  KiB/s): min=12488, max=14448, per=100.00%, avg=13740.14, stdev=309.13, samples=305
   iops        : min= 3122, max= 3612, avg=3435.04, stdev=77.28, samples=305
  lat (usec)   : 100=92.68%, 250=6.76%, 500=0.48%, 750=0.06%, 1000=0.01%
  lat (msec)   : 2=0.01%, 4=0.01%, 10=0.01%
  fsync/fdatasync/sync_file_range:
    sync (usec): min=172, max=7917, avg=200.66, stdev=44.17
    sync percentiles (usec):
     |  1.00th=[  180],  5.00th=[  184], 10.00th=[  186], 20.00th=[  188],
     | 30.00th=[  190], 40.00th=[  192], 50.00th=[  194], 60.00th=[  196],
     | 70.00th=[  200], 80.00th=[  204], 90.00th=[  217], 95.00th=[  233],
     | 99.00th=[  314], 99.50th=[  375], 99.90th=[  734], 99.95th=[  930],
     | 99.99th=[ 1500]
  cpu          : usr=0.68%, sys=3.30%, ctx=1048581, majf=0, minf=14
  IO depths    : 1=200.0%, 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 rwts: total=0,524288,0,0 short=524287,0,0,0 dropped=0,0,0,0
     latency   : target=0, window=0, percentile=100.00%, depth=1

Run status group 0 (all jobs):
  WRITE: bw=13.4MiB/s (14.1MB/s), 13.4MiB/s-13.4MiB/s (14.1MB/s-14.1MB/s), io=2048MiB (2147MB), run=152803-152803msec

Disk stats (read/write):
  nvme1n1: ios=0/1571325, merge=0/1047562, ticks=0/139246, in_queue=139247, util=100.00%
```

A.6 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.7 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.8 Benchmark configuration

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

```
host=172.16.18.136
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=0.0022
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.tugraph.interactive.TuGraphInteractiveDb
operation_count=117600000
ldbc.snb.interactive.parameters_dir=../../ldbc_snb_datagen_hadoop/substitution_parameters/
ldbc.snb.interactive.updates_dir=../../ldbc_snb_datagen_hadoop/social_network/
ldbc.snb.interactive.short_read_dissipation=0.2
ldbc.snb.interactive.update_interleave=156
ldbc.snb.interactive.scale_factor=30

warmup=29400000

## frequency of read queries (number of update queries per one read query)
ldbc.snb.interactive.LdbcQuery1_freq=26
ldbc.snb.interactive.LdbcQuery2_freq=37
ldbc.snb.interactive.LdbcQuery3_freq=106
ldbc.snb.interactive.LdbcQuery4_freq=36
ldbc.snb.interactive.LdbcQuery5_freq=72
ldbc.snb.interactive.LdbcQuery6_freq=316
ldbc.snb.interactive.LdbcQuery7_freq=48
ldbc.snb.interactive.LdbcQuery8_freq=9
ldbc.snb.interactive.LdbcQuery9_freq=384
ldbc.snb.interactive.LdbcQuery10_freq=37
ldbc.snb.interactive.LdbcQuery11_freq=20
ldbc.snb.interactive.LdbcQuery12_freq=44
ldbc.snb.interactive.LdbcQuery13_freq=19
ldbc.snb.interactive.LdbcQuery14_freq=49

# *** 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=172.16.18.136
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=0.00781
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.tugraph.interactive.TuGraphInteractiveDb
operation_count=122600000
ldbc.snb.interactive.parameters_dir=../../ldbc_snb_datagen_hadoop/substitution_parameters/
ldbc.snb.interactive.updates_dir=../../ldbc_snb_datagen_hadoop/social_network/
ldbc.snb.interactive.short_read_dissipation=0.2
ldbc.snb.interactive.update_interleave=48
ldbc.snb.interactive.scale_factor=100

```

```
warmup=30650000

## frequency of read queries (number of update queries per one read query)
ldbc.snb.interactive.LdbcQuery1_freq=26
ldbc.snb.interactive.LdbcQuery2_freq=37
ldbc.snb.interactive.LdbcQuery3_freq=123
ldbc.snb.interactive.LdbcQuery4_freq=36
ldbc.snb.interactive.LdbcQuery5_freq=78
ldbc.snb.interactive.LdbcQuery6_freq=434
ldbc.snb.interactive.LdbcQuery7_freq=38
ldbc.snb.interactive.LdbcQuery8_freq=5
ldbc.snb.interactive.LdbcQuery9_freq=527
ldbc.snb.interactive.LdbcQuery10_freq=40
ldbc.snb.interactive.LdbcQuery11_freq=22
ldbc.snb.interactive.LdbcQuery12_freq=44
ldbc.snb.interactive.LdbcQuery13_freq=19
ldbc.snb.interactive.LdbcQuery14_freq=49

# *** 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=172.16.18.136
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=0.034
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.tugraph.interactive.TuGraphInteractiveDb
operation_count=99760000
ldbc.snb.interactive.parameters_dir=../../ldbc_snb_datagen_hadoop/substitution_parameters/
ldbc.snb.interactive.updates_dir=../../ldbc_snb_datagen_hadoop/social_network/
ldbc.snb.interactive.short_read_dissipation=0.2
ldbc.snb.interactive.update_interleave=17
ldbc.snb.interactive.scale_factor=300

warmup=24940000

## frequency of read queries (number of update queries per one read query)
ldbc.snb.interactive.LdbcQuery1_freq=26
ldbc.snb.interactive.LdbcQuery2_freq=37
ldbc.snb.interactive.LdbcQuery3_freq=142
ldbc.snb.interactive.LdbcQuery4_freq=36
ldbc.snb.interactive.LdbcQuery5_freq=84
ldbc.snb.interactive.LdbcQuery6_freq=580
ldbc.snb.interactive.LdbcQuery7_freq=32
ldbc.snb.interactive.LdbcQuery8_freq=3
ldbc.snb.interactive.LdbcQuery9_freq=705
ldbc.snb.interactive.LdbcQuery10_freq=44
ldbc.snb.interactive.LdbcQuery11_freq=24
ldbc.snb.interactive.LdbcQuery12_freq=44
ldbc.snb.interactive.LdbcQuery13_freq=19
ldbc.snb.interactive.LdbcQuery14_freq=49

# *** 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.9 Validation configuration

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

```
host=172.16.18.134
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.tugraph.interactive.TuGraphInteractiveDb
operation_count=10000

validate_workload=true
validate_database=validation_params.csv
ldbc.snb.interactive.parameters_dir=../../ldbc_snb_datagen_hadoop/substitution_parameters/
ldbc.snb.interactive.short_read_dissipation=0.2
ldbc.snb.interactive.update_interleave=466
ldbc.snb.interactive.scale_factor=10

## frequency of read queries (number of update queries per one read query)
ldbc.snb.interactive.LdbcQuery1_freq=1
```

```
ldbc.snb.interactive.LdbcQuery2_freq=1
ldbc.snb.interactive.LdbcQuery3_freq=1
ldbc.snb.interactive.LdbcQuery4_freq=1
ldbc.snb.interactive.LdbcQuery5_freq=1
ldbc.snb.interactive.LdbcQuery6_freq=1
ldbc.snb.interactive.LdbcQuery7_freq=1
ldbc.snb.interactive.LdbcQuery8_freq=1
ldbc.snb.interactive.LdbcQuery9_freq=1
ldbc.snb.interactive.LdbcQuery10_freq=1
ldbc.snb.interactive.LdbcQuery11_freq=1
ldbc.snb.interactive.LdbcQuery12_freq=1
ldbc.snb.interactive.LdbcQuery13_freq=1
ldbc.snb.interactive.LdbcQuery14_freq=1

# *** 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
```