The graph & RDF benchmark reference

Full Disclosure Report of the LDBC Social Network Benchmark

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

July 26, 2020

General Terms

Executive Summary

TuGraph (formerly known as LightGraph) is a proprietary graph database product developed by 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 functions, which are compiled and loaded into the database as plugins. Thus, 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 data model with indices over node and edge identifiers and over properties selected by the user. TuGraph also supports precomputed properties (i.e., derived values or materialized views) and maintaining the consistency of such properties is done via specific queries at runtime. This current benchmark implementation employs two such precomputed properties to improve the operation throughput of the system. 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:
96E3EBD92BAE4C8...

Mr. Marton Bur
(Auditor)

7/28/2020
Date

DocuSigned by:
A4E0271F756C4E6...

Dr. Gabor Szarnyas
(Head of LDBC SNB Task Force)

7/30/2020
Date

DocuSigned by:
3584AB093AA0456...

Dr. Xiaowei Zhu
(Test Sponsor Representative)

7/28/2020
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	Ohio
Common name of the item	r5d.12xlarge
Operating system	18.04.1-Ubuntu

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 8175M 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: 1.5MB L2 cache: 24MB L3 cache: 33MB

1.1.3 Memory details

The total size of the memory installed is 374GB, and this information was obtained using the `cat /proc/meminfo` command issued from the virtual machine instance. The type and frequency of the memory installed in the virtual machine was not obtainable.

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 xfs filesystem. The storage size and type is from the Amazon Web Services website `https://aws.amazon.com/ec2/instance-types/r5/` (accessed: July 12, 2020).

The 4KB QD1 write performance was measured with the `fio` command and the output (Listing A.2) showed an average of 28853 IOPS.

1.1.5 Network details

The presented benchmark run only used a single machine, thus network details are not included here.

1.1.6 Machine pricing

The system pricing summary is included in the table below.

Table 1.3: Pricing summary

Item	**Price**
r5d.12xlarge reserved instance machine in AWS (standard 3-year term)	34 150 USD
Permanent TuGraph license	170 000 USD
Maintenance fee (3 years)	76 500 USD

1.1.7 System availability

The latest software version of TuGraph (version 1.10) was made available on July 16, 2020. 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	v0.3.3
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.3–A.5.

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.6) 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

```
[comment.csv]
LABEL=Comment,HEADER=0
id:INT64:ID,creationDate:INT64,locationIP:STRING,browserUsed:STRING,content:STRING,length:INT32,creator:INT64,place:INT64,replyOfPost:INT64:OPTIONAL,replyOfComment:INT64:OPTIONAL

[comment_hasCreator_person.csv]
LABEL=commentHasCreator,SRC_ID=Comment:id,DST_ID=Person:id,HEADER=0
SRC_ID,creationDate:INT64,SKIP,SKIP,SKIP,SKIP,DST_ID,SKIP,SKIP,SKIP

...
```

Data loading times are shown for each scale factor in the table below. Values were measured using the GNU Time tool with the `-v` flag, reading the *Elapsed (wall clock) 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 (the CSV conversion was executed outside the measurements and was not timed). 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* 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)
10	145.35	55.28	200.63
30	404.90	178.34	583.24
100	1 162.58	566.83	1 729.41
300	5 893.75	1 834.23	7 727.98

3 Test Driver Details

3.1 Driver implementation

A test driver adaptation for the SUT was provided by the test sponsor. The used version is available online: `https://github.com/fma-cloud/ldbc_snb_implementations/tree/ec2` (accessed: July 16, 2020). 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=0.0064 for scale factor 30,
- TCR=0.0280 for scale factor 100 and
- TCR=0.0900 for scale factor 300.

The complete configuration files for the different scale factors are shown in Listings A.7–A.9, 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.

Table 4.1: Summary of results for scale factor 30

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 17m 56.784s	44 994 820	5 436.47 $\frac{\text{operations}}{\text{second}}$	100.00%

Table 4.2: Summary of results for scale factor 100

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 03m 03.076s	36 994 870	5 010.77 $\frac{\text{operations}}{\text{second}}$	100.00%

Table 4.3: Summary of results for scale factor 300

Benchmark duration	Benchmark operations	Throughput	Query on-time compliance
02h 07m 00.853s	37 003 185	4 855.52 $\frac{\text{operations}}{\text{second}}$	100.00%

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.

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}
Complex1	318 810	169	765 216	60 698.08	46 434	90 372	129 784	419 968
Complex2	224 028	192	136 888	5 186.19	4 371	7 168	9 466	26 124
Complex3	78 198	13 438	256 344	34 948.39	33 064	42 696	45 654	82 184
Complex4	230 251	196	131 656	8 001.55	6 585	13 502	15 837	31 275
Complex5	115 126	102 192	1 913 472	1 014 350.71	1 016 384	1 310 464	1 382 400	1 500 800
Complex6	26 231	442	190 152	11 946.40	14 399	23 006	24 978	37 660
Complex7	172 688	160	95 000	1 559.40	722	2 762	4 707	18 986
Complex8	921 005	5 799	224 616	15 320.57	11 899	26 989	32 669	46 132
Complex9	21 586	314	1 015 712	468 535.15	455 904	622 176	665 056	735 232
Complex10	224 028	177	544 288	103 189.22	100 820	129 032	138 136	188 280
Complex11	414 452	204	275 392	33 378.46	32 266	40 038	42 460	75 168
Complex12	188 387	203	477 488	76 743.98	74 200	102 728	112 412	148 048
Complex13	436 266	164	107 472	2 306.08	1 368	4 306	5 923	20 061
Complex14	169 164	236	533 536	15 076.53	5 453	12 044	113 580	182 504
Short1	4 444 416	115	47 348	549.07	302	928	1 639	4 575
Short2	4 444 416	121	51 618	711.75	484	1 165	1 722	4 492
Short3	4 444 416	138	59 426	1 214.22	773	2 641	4 199	6 549
Short4	4 443 714	105	47 346	565.11	310	956	1 696	4 755
Short5	4 443 714	116	45 590	503.28	288	832	1 366	4 081
Short6	4 443 714	117	46 342	506.94	292	835	1 370	4 096
Short7	4 443 714	124	45 578	597.05	386	955	1 499	4 337
Update1	2 569	560	103 980	3 405.42	1 957	5 910	10 632	30 756
Update2	1 481 994	351	124 120	1 844.40	833	4 073	5 931	14 627
Update3	2 259 829	343	115 448	1 753.36	814	3 851	5 650	13 898
Update4	43 746	450	91 668	2 354.77	917	6 289	9 655	16 791
Update5	4 449 860	349	123 624	2 178.63	1 124	4 512	6 507	15 709
Update6	534 249	473	114 664	2 076.45	980	4 929	8 075	14 494
Update7	1 410 100	473	141 008	2 343.26	1 033	5 935	9 392	16 387
Update8	164 197	405	196 080	5 923.03	3 365	13 301	20 869	39 246

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}
Complex1	211 251	176	1 317 184	137 109.84	70 220	253 840	792 768	959 936
Complex2	148 446	173	135 072	5 734.10	4 581	8 438	11 875	28 647
Complex3	44 655	32 645	478 880	93 499.36	90 280	117 796	123 468	178 904
Complex4	152 570	198	150 416	8 686.73	7 024	15 105	17 545	34 268
Complex5	70 417	207	2 441 728	1 279 996.04	1 306 432	1 717 952	1 823 296	2 000 064
Complex6	12 655	205	247 760	28 894.56	22 224	61 350	66 500	80 120
Complex7	144 541	146	101 744	1 984.95	733	4 407	7 660	21 496
Complex8	1 098 504	166	147 520	1 749.72	546	3 990	7 236	20 750
Complex9	10 423	199	1 136 640	600 302.57	586 752	788 288	845 088	956 480
Complex10	137 313	187	516 512	131 846.14	129 344	168 144	181 216	227 088
Complex11	249 660	170	295 248	41 327.69	40 212	49 926	53 254	92 360
Complex12	124 830	198	559 872	94 911.48	92 600	127 372	139 816	180 704
Complex13	289 081	163	136 032	3 855.89	2 910	6 709	10 116	24 786
Complex14	112 092	205	577 952	34 976	6 753	108 444	127 876	186 616
Short1	3 522 275	119	58 658	578.43	275	907	2 220	5 531
Short2	3 522 275	125	58 462	763.85	468	1 228	2 321	5 710
Short3	3 522 275	143	61 570	1 399.78	846	3 432	5 047	7 536
Short4	3 522 949	110	58 880	593.28	285	936	2 268	5 597
Short5	3 522 949	113	54 834	550.99	264	832	2 059	5 432
Short6	3 522 949	119	59 254	555.71	269	833	2 057	5 454
Short7	3 522 949	127	59 214	647.95	366	959	2 173	5 584
Update1	1 445	529	61 302	4 239.32	2 293	10 538	15 721	30 085
Update2	1 219 746	335	115 860	1 891.15	896	4 251	6 353	14 479
Update3	2 125 211	338	149 552	1 930.63	916	4 348	6 414	14 599
Update4	23 824	442	76 040	3 342.68	961	11 620	15 402	20 817
Update5	2 781 561	369	154 568	2 126.99	1 157	4 446	6 680	15 154
Update6	520 901	467	105 496	6 495.55	3 179	16 689	18 919	24 251
Update7	2 747 775	470	152 048	4 146.82	1 244	13 844	16 815	21 831
Update8	109 348	415	188 488	5 678.11	3 240	12 540	20 127	37 072

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}
Complex1	191 526	175	2 171 904	131 831.08	94 784	211 472	303 408	1 396 480
Complex2	134 586	170	122 876	6 168.43	5 154	8 951	12 325	28 188
Complex3	35 068	89 352	711 584	278 012.94	274 928	357 824	414 144	509 600
Complex4	138 324	202	145 552	11 349.45	9 394	20 414	22 653	37 414
Complex5	59 282	205	3 848 192	2 142 661.74	2 189 952	2 903 936	3 051 904	3 307 392
Complex6	8 586	382	345 696	78 026.94	25 151	175 696	187 344	210 792
Complex7	155 615	165	111 884	1 760.18	711	3 420	6 998	20 456
Complex8	1 659 894	174	119 300	1 582.19	578	2 999	6 548	19 894
Complex9	7 063	200	1 374 656	838 819.48	826 400	1 065 600	1 136 256	1 249 088
Complex10	113 175	196	465 376	167 285.26	165 672	221 880	239 192	278 240
Complex11	207 487	206	289 312	49 067.86	48 302	61 480	65 262	89 148
Complex12	113 174	200	401 168	112 531.55	110 288	161 016	176 848	211 256
Complex13	262 088	172	122 096	5 607.43	5 306	8 599	12 253	27 459
Complex14	101 626	179	736 704	48 026.63	11 251	112 292	131 464	194 512
Short1	4 001 189	130	60 616	470.71	299	736	1 063	3 159
Short2	4 001 189	150	54 722	696.74	522	1 089	1 489	3 134
Short3	4 001 189	154	66 152	1 542.03	1 005	2 890	5 484	8 573
Short4	4 002 336	120	53 898	491.97	306	764	1 119	3 557
Short5	4 002 336	122	63 678	436.66	289	689	957	2 470
Short6	4 002 336	118	54 638	441.24	295	693	960	2 450
Short7	4 002 336	124	55 552	550.60	401	853	1 122	2 679
Update1	1 111	618	71 628	3 567.14	2 069	6 531	12 055	30 557
Update2	805 022	377	119 880	2 023.37	846	4 101	7 199	21 137
Update3	1 252 745	385	137 000	2 005.35	832	4 095	7 190	20 866
Update4	19 337	493	68 208	2 234.62	890	5 208	9 100	21 107
Update5	2 569 147	413	142 136	2 353.41	1 158	4 538	7 694	21 822
Update6	263 653	515	142 728	2 474.88	1 002	6 086	9 900	21 324
Update7	791 004	518	124 896	2 325.95	1 037	5 231	9 151	20 635
Update8	100 761	451	174 536	6 121.67	3 160	13 968	22 561	44 454

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 was created with the SNB Interactive reference implementation over Neo4j, running the Community Edition of version 3.5.19. The system with the driver configuration showin in Listing A.10 successfully returned the expected result sets for the queries of the benchmark.

In addition, to further verify the correctness of the SUT, validation was performed for scale factors 1 and 3, where reference results were generated with PostgreSQL and Neo4j, respectively. In both cases the results were sound and complete.

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 the server machine was shut down using the command `sudo reboot` after 41M 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|1595172523653|1595172523653|11103|0|1349009994880
$ grep LdbcUpdate2 LDBC-SNB-results_log.csv | tail -n 1
LdbcUpdate2AddPostLike|1595172524081|1595172524082|878|0|1349010061787
$ grep LdbcUpdate3 LDBC-SNB-results_log.csv | tail -n 1
LdbcUpdate3AddCommentLike|1595172524075|1595172524075|572|0|1349010060928
$ grep LdbcUpdate4 LDBC-SNB-results_log.csv | tail -n 1
LdbcUpdate4AddForum|1595172524032|1595172524032|792|0|1349010054022
$ grep LdbcUpdate5 LDBC-SNB-results_log.csv | tail -n 1
LdbcUpdate5AddForumMembership|1595172524084|1595172524084|937|0|1349010062198
$ grep LdbcUpdate6 LDBC-SNB-results_log.csv | tail -n 1
LdbcUpdate6AddPost|1595172524006|1595172524006|1060|0|1349010050019
$ grep LdbcUpdate7 LDBC-SNB-results_log.csv | tail -n 1
LdbcUpdate7AddComment|1595172524081|1595172524081|1053|0|1349010061766
$ grep LdbcUpdate8 LDBC-SNB-results_log.csv | tail -n 1
LdbcUpdate8AddFriendship|1595172523977|1595172523977|1268|0|1349010045625
```

From the logs, the last completed updates were retrieved for each update operation. 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. 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 qureies 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
`console-output-sf{30,100,300}.txt`	Driver console outputs 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_implementations-ec2.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
`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`	Simple test cases to check committed transactions
`TuGraph-1.10.0-1.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
├── console-output-sf30.txt
├── console-output-sf100.txt
├── console-output-sf300.txt
├── TuGraph-1.10.0-1.x86_64.deb
├── ldbc_snb_implementations-ec2.zip
└── tugraph_ldbc_snb
    ├── README.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
        ├── install.py
        ├── install.sh
        ├── preprocess.cpp
        ├── snb_common.h
        ├── snb_constants.h
        ├── generate_snb_constants.cpp
        └── recovery_queries.cpp
```

A Appendix

A.1 CPU details

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

```
processor : 1
vendor_id : GenuineIntel
cpu family  : 6
model   : 85
model name  : Intel(R) Xeon(R) Platinum 8175M CPU @ 2.50GHz
stepping  : 4
microcode : 0x2000069
cpu MHz   : 1580.304
cache size  : 33792 KB
physical id : 0
siblings  : 48
core id   : 23
cpu cores : 24
apicid    : 47
initial apicid  : 47
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
    hle avx2 smep bmi2 erms invpcid rtm 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 taa itlb_multihit
bogomips  : 5000.00
clflush size  : 64
cache_alignment : 64
address sizes : 46 bits physical, 48 bits virtual
power management:
```

A.2 IO performance

Listing A.2: Output of the `fio` command

```
$ fio --rw=write --ioengine=sync --fdatasync=1 --direct=1 --directory=io-test-data --size=2g --bs=4k --name=
    iotest
mytest: (g=0): rw=write, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B, ioengine=sync, iodepth=1
fio-3.1
Starting 1 process
Jobs: 1 (f=1): [W(1)][100.0%][r=0KiB/s,w=113MiB/s][r=0,w=29.0k IOPS][eta 00m:00s]
mytest: (groupid=0, jobs=1): err= 0: pid=3100: Thu Jul 16 20:15:03 2020
   write: IOPS=28.8k, BW=113MiB/s (118MB/s)(2048MiB/18176msec)
   clat (usec): min=19, max=637, avg=28.61, stdev= 8.34
       lat (usec): min=19, max=637, avg=28.69, stdev= 8.37
   clat percentiles (usec):
       |  1.00th=[   22],  5.00th=[   23], 10.00th=[   24], 20.00th=[   24],
       | 30.00th=[   25], 40.00th=[   25], 50.00th=[   26], 60.00th=[   26],
```

```
        | 70.00th=[   34], 80.00th=[   35], 90.00th=[   37], 95.00th=[   39],
        | 99.00th=[   52], 99.50th=[   83], 99.90th=[  113], 99.95th=[  121],
        | 99.99th=[  172]
    bw (  KiB/s): min=111703, max=118808, per=100.00%, avg=115413.58, stdev=1881.82, samples=36
    iops        : min=27925, max=29702, avg=28853.31, stdev=470.40, samples=36
    lat (usec)   : 20=0.01%, 50=98.89%, 100=0.85%, 250=0.25%, 500=0.01%
    lat (usec)   : 750=0.01%
    cpu          : usr=8.31%, sys=25.87%, ctx=1048573, 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=113MiB/s (118MB/s), 113MiB/s-113MiB/s (118MB/s-118MB/s), io=2048MiB (2147MB), run=18176-18176msec

Disk stats (read/write):
    md0: ios=0/1046845, merge=0/0, ticks=0/0, in_queue=0, util=0.00%, aggrios=0/262144, aggrmerge=0/0, aggrticks
     =0/5852, aggrin_queue=0, aggrutil=94.76%
    nvme0n1: ios=0/262144, merge=0/0, ticks=0/5807, in_queue=0, util=93.86%
    nvme1n1: ios=0/262144, merge=0/0, ticks=0/5898, in_queue=0, util=94.76%
```

A.3 Datagen configuration

Listing A.3: 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.4: 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.5: 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.4 Import configuration

Listing A.6: Content of `import.conf` describing the data schema

```
[comment.csv]
LABEL=Comment,HEADER=0
id:INT64:ID,creationDate:INT64,locationIP:STRING,browserUsed:STRING,content:STRING,length:INT32,creator:INT64,
    place:INT64,replyOfPost:INT64:OPTIONAL,replyOfComment:INT64:OPTIONAL

[comment_hasCreator_person.csv]
LABEL=commentHasCreator,SRC_ID=Comment:id,DST_ID=Person:id,HEADER=0
SRC_ID,creationDate:INT64,SKIP,SKIP,SKIP,SKIP,DST_ID,SKIP,SKIP,SKIP

[comment_hasTag_tag.csv]
LABEL=commentHasTag,SRC_ID=Comment:id,DST_ID=Tag:id,HEADER=0
SRC_ID,DST_ID

[comment_isLocatedIn_place.csv]
LABEL=commentIsLocatedIn,SRC_ID=Comment:id,DST_ID=Place:id,HEADER=0
SRC_ID,creationDate:INT64,SKIP,SKIP,SKIP,SKIP,SKIP,DST_ID,SKIP,SKIP

[comment_replyOf_comment.csv]
LABEL=replyOf,SRC_ID=Comment:id,DST_ID=Comment:id,HEADER=0
SRC_ID,DST_ID,creationDate:INT64

[comment_replyOf_post.csv]
LABEL=replyOf,SRC_ID=Comment:id,DST_ID=Post:id,HEADER=0
SRC_ID,DST_ID,creationDate:INT64

[forum.csv]
LABEL=Forum,HEADER=0
id:INT64:ID,title:STRING,creationDate:INT64,moderator:INT64

[forum_containerOf_post.csv]
LABEL=containerOf,SRC_ID=Forum:id,DST_ID=Post:id,HEADER=0
DST_ID,SKIP,SKIP,SKIP,SKIP,SKIP,SKIP,SKIP,SKIP,SRC_ID,SKIP

[forum_hasMember_person.csv]
LABEL=hasMember,SRC_ID=Forum:id,DST_ID=Person:id,HEADER=0
SRC_ID,DST_ID,joinDate:INT64,numPosts:INT32

[forum_hasModerator_person.csv]
LABEL=hasModerator,SRC_ID=Forum:id,DST_ID=Person:id,HEADER=0
SRC_ID,SKIP,SKIP,DST_ID

[forum_hasTag_tag.csv]
LABEL=forumHasTag,SRC_ID=Forum:id,DST_ID=Tag:id,HEADER=0
SRC_ID,DST_ID

[organisation.csv]
LABEL=Organisation,HEADER=0
id:INT64:ID,type:STRING,name:STRING,url:STRING,place:INT64
```

```
48
49 [organisation_isLocatedIn_place.csv]
50 LABEL=organisationIsLocatedIn,SRC_ID=Organisation:id,DST_ID=Place:id,HEADER=0
51 SRC_ID,SKIP,SKIP,SKIP,DST_ID
52
53 [person.csv]
54 LABEL=Person,HEADER=0
55 id:INT64:ID,firstName:STRING,lastName:STRING,gender:STRING,birthday:INT64,creationDate:INT64,locationIP:STRING,
       browserUsed:STRING,place:INT64,speaks:STRING,email:STRING
56
57 [person_hasInterest_tag.csv]
58 LABEL=hasInterest,SRC_ID=Person:id,DST_ID=Tag:id,HEADER=0
59 SRC_ID,DST_ID
60
61 [person_isLocatedIn_place.csv]
62 LABEL=personIsLocatedIn,SRC_ID=Person:id,DST_ID=Place:id,HEADER=0
63 SRC_ID,SKIP,SKIP,SKIP,SKIP,SKIP,SKIP,SKIP,DST_ID,SKIP,SKIP
64
65 [person_knows_person.csv]
66 LABEL=knows,SRC_ID=Person:id,DST_ID=Person:id,HEADER=0
67 SRC_ID,DST_ID,creationDate:INT64,weight:DOUBLE
68
69 [person_likes_comment.csv]
70 LABEL=likes,SRC_ID=Person:id,DST_ID=Comment:id,HEADER=0
71 SRC_ID,DST_ID,creationDate:INT64
72
73 [person_likes_post.csv]
74 LABEL=likes,SRC_ID=Person:id,DST_ID=Post:id,HEADER=0
75 SRC_ID,DST_ID,creationDate:INT64
76
77 [person_studyAt_organisation.csv]
78 LABEL=studyAt,SRC_ID=Person:id,DST_ID=Organisation:id,HEADER=0
79 SRC_ID,DST_ID,classYear:INT32
80
81 [person_workAt_organisation.csv]
82 LABEL=workAt,SRC_ID=Person:id,DST_ID=Organisation:id,HEADER=0
83 SRC_ID,DST_ID,workFrom:INT32:OPTIONAL
84
85 [place.csv]
86 LABEL=Place,HEADER=0
87 id:INT64:ID,name:STRING,url:STRING,type:STRING,isPartOf:INT64:OPTIONAL
88
89 [place_isPartOf_place.csv]
90 LABEL=isPartOf,SRC_ID=Place:id,DST_ID=Place:id,HEADER=0
91 SRC_ID,DST_ID
92
93 [post.csv]
94 LABEL=Post,HEADER=0
95 id:INT64:ID,imageFile:STRING:OPTIONAL,creationDate:INT64,locationIP:STRING,browserUsed:STRING,language:STRING:
       OPTIONAL,content:STRING:OPTIONAL,length:INT32,creator:INT64,container:INT64,place:INT64
96
97 [post_hasCreator_person.csv]
98 LABEL=postHasCreator,SRC_ID=Post:id,DST_ID=Person:id,HEADER=0
99 SRC_ID,SKIP,creationDate:INT64,SKIP,SKIP,SKIP,SKIP,SKIP,DST_ID,SKIP,SKIP
100
101 [post_hasTag_tag.csv]
102 LABEL=postHasTag,SRC_ID=Post:id,DST_ID=Tag:id,HEADER=0
103 SRC_ID,DST_ID
```

```
[post_isLocatedIn_place.csv]
LABEL=postIsLocatedIn,SRC_ID=Post:id,DST_ID=Place:id,HEADER=0
SRC_ID,SKIP,creationDate:INT64,SKIP,SKIP,SKIP,SKIP,SKIP,SKIP,SKIP,DST_ID

[tag.csv]
LABEL=Tag,HEADER=0
id:INT64:ID,name:STRING,url:STRING,hasType:INT64

[tag_hasType_tagclass.csv]
LABEL=hasType,SRC_ID=Tag:id,DST_ID=Tagclass:id,HEADER=0
SRC_ID,SKIP,SKIP,DST_ID

[tagclass.csv]
LABEL=Tagclass,HEADER=0
id:INT64:ID,name:STRING,url:STRING,isSubclassOf:INT64:OPTIONAL

[tagclass_isSubclassOf_tagclass.csv]
LABEL=isSubclassOf,SRC_ID=Tagclass:id,DST_ID=Tagclass:id,HEADER=0
SRC_ID,DST_ID
```

A.5 Benchmark configuration

Listing A.7: Contents of `interactive-benchmark-sf30.properties` used for scale factor 30

```
host=localhost
rest_port=7071
port=9091
user=admin
pass=admin123456
is_admin=true

status=1
thread_count=48
name=LDBC-SNB
results_log=true
time_unit=MICROSECONDS
time_compression_ratio=0.0064
peer_identifiers=
workload_statistics=false
spinner_wait_duration=1
help=false
ignore_scheduled_start_times=false

workload=com.ldbc.driver.workloads.ldbc.snb.interactive.LdbcSnbInteractiveWorkload
db=ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb
operation_count=45000000
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.update_interleave=156

warmup=11000000

## 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.8: Contents of `interactive-benchmark-sf100.properties` used for scale factor 100

```
host=localhost
rest_port=7071
port=9091
user=admin
pass=admin123456
is_admin=true

status=1
```

```
thread_count=48
name=LDBC-SNB
results_log=true
time_unit=MICROSECONDS
time_compression_ratio=0.028
peer_identifiers=
workload_statistics=false
spinner_wait_duration=1
help=false
ignore_scheduled_start_times=false

workload=com.ldbc.driver.workloads.ldbc.snb.interactive.LdbcSnbInteractiveWorkload
db=ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb
operation_count=37000000
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.update_interleave=48

warmup=9000000

## 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.9: Contents of `interactive-benchmark-sf300.properties` used for scale factor 300

```
host=localhost
rest_port=7071
port=9091
user=admin
pass=admin123456
is_admin=true

status=1
thread_count=48
name=LDBC-SNB
results_log=true
time_unit=MICROSECONDS
time_compression_ratio=0.09
peer_identifiers=
workload_statistics=false
spinner_wait_duration=1
help=false
ignore_scheduled_start_times=false

workload=com.ldbc.driver.workloads.ldbc.snb.interactive.LdbcSnbInteractiveWorkload
db=ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb
operation_count=37000000
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.update_interleave=17

warmup=9000000

## 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.6 Validation configuration

Listing A.10: The contents of `interactive-validate.properties`

```
host=localhost
rest_port=7071
port=9091
user=admin
pass=admin123456
is_admin=true

status=1
thread_count=1
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=com.ldbc.driver.workloads.ldbc.snb.interactive.LdbcSnbInteractiveWorkload
db=ai.fma.impls.workloads.ldbc.snb.lightgraph.interactive.LightGraphInteractiveDb
operation_count=10000

validate_workload=true
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.update_interleave=466

## 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
```