

---

## **Full Disclosure Report of the LDBC Social Network Benchmark**

---

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

December 25, 2023

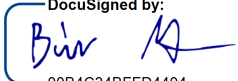
## GENERAL TERMS

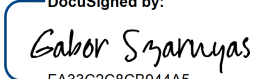
### Executive Summary

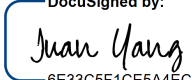
The present document describes an audited run of the Interactive Workload of LDBC's Social Network Benchmark using AtlasGraph version 3.0.0. AtlasGraph is a graph database co-developed by StarGraph (Beijing Haizhi Stargraph Technology Co., Ltd.) and MadSys Group (Tsinghua University). The benchmark was executed while ensuring serializable isolation level. The audited runs assessing query performance were executed on three scale factors: SF30, SF100, and SF300. While AtlasGraph has the option to run declarative Cypher queries, the benchmark queries were implemented in the Rust imperative programming language. The graph schema used in the benchmark runs did not contain any precomputed attributes, and the used indices over attributes are explicitly marked in the data schema definition.

### Declaration of Audit Success

This report contains an audited LDBC benchmark execution. 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:  
  
.....00B4C34BFED4404.....  
Dr. Márton Búr  
(Auditor) 12/25/2023  
Date

DocuSigned by:  
  
.....FA33C2C8CB944A5.....  
Dr. Gábor Szárnyas  
(Head of LDBC SNB Task Force) 12/26/2023  
Date

DocuSigned by:  
  
.....8E33C9F1CE5A4EG.....  
Juan Yang  
(Test Sponsor Representative) 12/26/2023  
Date



## TABLE OF CONTENTS

1	SYSTEM DESCRIPTION AND PRICING SUMMARY	4
1.1	Details of machines driving and running the workload . . . . .	4
1.1.1	Machine overview . . . . .	4
1.1.2	CPU details . . . . .	4
1.1.3	Memory details . . . . .	4
1.1.4	Disk and storage details . . . . .	4
1.1.5	Network details . . . . .	5
1.1.6	Machine pricing . . . . .	5
1.1.7	System availability . . . . .	5
2	DATASET GENERATION	6
2.1	General information . . . . .	6
2.2	Datagen configurations . . . . .	6
2.3	Data loading and data schema . . . . .	6
3	TEST DRIVER DETAILS	9
3.1	Driver implementation . . . . .	9
3.2	Benchmark configuration of driver . . . . .	9
4	PERFORMANCE RESULTS	10
5	VALIDATION OF THE RESULTS	14
6	ACID COMPLIANCE	15
6.1	Transaction isolation level . . . . .	15
6.2	SNB ACID test results . . . . .	15
6.3	Recovery and durability . . . . .	15
6.3.1	Recovery . . . . .	15
6.3.2	Durability . . . . .	15
6.3.3	Consistency after recovery . . . . .	17
7	SUPPLEMENTARY MATERIALS	18
A	APPENDIX	20
A.1	CPU details . . . . .	20
A.2	Memory details . . . . .	21
A.3	Storage details and IO performance . . . . .	22
A.4	Network performance . . . . .	24
A.5	Datagen configuration . . . . .	25
A.6	Import configuration . . . . .	25
A.7	Benchmark configuration . . . . .	42
A.8	Validation configuration . . . . .	45

## System Description and Pricing Summary

---

# 1 SYSTEM DESCRIPTION AND PRICING SUMMARY

## 1.1 Details of machines driving and running the workload

### 1.1.1 Machine overview

This benchmark used two r5d.12xlarge instances, one for the driver and one for the SUT. The instance details below were obtained from the Amazon Web Services console. The operating system name and version was obtained from running `cat /etc/os-release` command.

Table 1.1: Machine Type and Location

Cloud provider	Amazon Web Services
Machine region	China (Beijing)
Common name of the item	r5d.12xlarge
Operating system	Ubuntu 20.04.6

### 1.1.2 CPU details

The details below were obtained using the command `cat /proc/cpuinfo` (Listing A.1) and `lscpu` (Listing A.2) issued from the machine instance, and the datasheet of the used CPU type.

Table 1.2: CPU details summary

Type	Intel® Intel Xeon® Platinum 8259CL CPU @ 2.50GHz CPU
Total number	1
Cores per CPU	24
Threads per CPU	48
CPU clock frequency	2.5 GHz
Total cache size per CPU	L1d cache: 768KiB L1i cache: 768KiB L2 cache: 24MiB L3 cache: 35.8MiB

### 1.1.3 Memory details

The total size of the memory installed is 384GB, and this information was obtained using the `cat /proc/meminfo` command issued from the virtual machine instance. The details section of the `sudo lshw -c memory` returns DIMM DDR4 Static column Pseudo-static Synchronous Window DRAM 3200 MHz (0.3 ns). For the latter command, see full output in Listing A.3.

### 1.1.4 Disk and storage details

The disk controller and motherboard type were not obtainable from the virtual machine instance. The storage consists of two individual 900GB NVMe SSD, formatted with the xfs filesystem. The storage size and type is from the Amazon Web Services website <https://aws.amazon.com/ec2/instance-types/r5/> (accessed: December 13, 2023). Actual disk parameters for the SUT instance were obtained by `sudo lshw -c storage -c disk` and are shown in Listing A.4.

The 4KB QD1 write performance was measured with the `fio` command and the output (Listing A.5) showed an average of 16 898 IOPS.



### 1.1.5 Network details

The presented benchmark ran in a client-server architecture. The network throughput performance was measured using the `iperf` tool. The key parts of the command output are presented in Listing A.6.

### 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)	232 461 RMB
Software license	1 200 000 RMB
Maintenance fee (3 years)	600 000 RMB

### 1.1.7 System availability

The latest software version of AtlasGraph (version 3.0.0) was made available on November 28, 2023. 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	v1.0.0
Output format	CsvComposite serializer
Scale factors	10, 30, 100, 300

**Note:** 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.7–A.9.

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

```

1 ldbc.snb.datagen.generator.scaleFactor:snb.interactive.10
2
3 ldbc.snb.datagen.serializer.dynamicActivitySerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.
  activity.CsvCompositeDynamicActivitySerializer
4 ldbc.snb.datagen.serializer.dynamicPersonSerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.person.
  CsvCompositeDynamicPersonSerializer
5 ldbc.snb.datagen.serializer.staticSerializer:ldbc.snb.datagen.serializer.snb.csv.staticserializer.
  CsvCompositeStaticSerializer

```

### 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 `prepare_snb.sh` script). The loading process takes a schema descriptor file `schema.json` (see Listing 2.2 and Listing A.10) which defines the data types represented in the graph, as well as a mapping configuration `mapping_sf{10,30,100,300}.json` (see Listing 2.3 and Listing A.11) that specifies which CSV columns are used to populate the respective edges and nodes in the graph.

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

```

1 {
2   "graph_name": "ldbc",
3   "shard_num": 1,
4   "replica_factor": 1,
5   "vertices": {
6     "city": {
7       "properties": [
8         {
9           "name": "id",
10          "type": "int64",
11          "is_pk": true,
12          "not_null": true
13        },
14        {

```

```

15         "name": "name",
16         "type": "string"
17     },
18     {
19         "name": "url",
20         "type": "string"
21     }
22 ],
23 "max_vertex_num": 1343
24 },
25 ...
26 },
27 "edges": {
28     "likes": {
29         "properties": [
30             {
31                 "name": "creationDate",
32                 "type": "datetime"
33             }
34         ],
35         "relations": [
36             {
37                 "from_key_label_name": "person",
38                 "to_key_label_name": "post",
39                 "out_edge_strategy": "none"
40             },
41             {
42                 "from_key_label_name": "person",
43                 "to_key_label_name": "comment",
44                 "out_edge_strategy": "none"
45             }
46         ]
47     },
48     ...
49 }
50 }

```

Listing 2.3: Excerpt from mapping\_sf10.json, describing the data mapping

```

1 {
2     "graph_name": "ldbc",
3     "vertices": {
4         "city": [
5             {
6                 "has_header": true,
7                 "csv_row_fields_order": [
8                     "id",
9                     "name",
10                    "url"
11                ],
12                "csv_path": "/disk1/SF10/social_network/static/city_0_0.csv",
13                "type": "csv",
14                "delimiter": "|"
15            }
16        ],
17        ...
18    },
19    "edges": {

```



```

20     "likes": [
21         {
22             "has_header": true,
23             "csv_row_fields_order": [
24                 "from_key/person",
25                 "to_key/comment",
26                 "creationDate"
27             ],
28             "csv_path": "/disk1/SF10/social_network/dynamic/person_likes_comment_0_0.csv",
29             "type": "csv",
30             "delimiter": "|"
31         },
32         {
33             "has_header": true,
34             "csv_row_fields_order": [
35                 "from_key/person",
36                 "to_key/post",
37                 "creationDate"
38             ],
39             "csv_path": "/disk1/SF10/social_network/dynamic/person_likes_post_0_0.csv",
40             "type": "csv",
41             "delimiter": "|"
42         }
43     ],
44     ...
45 }
46 }

```

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 an index over properties marked `is_pk` (primary keys) and `is_index` in the schema. These times, however, do not include CSV conversion times (the CSV conversion was executed outside of the measurements and was not timed). The NA values in the **Data preprocessing time** column are indicating that there were no materialized views or other auxiliary data structures precomputed prior to executing the benchmark, and index creation times are included in the CSV loading times. 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	1 680	NA	1 680
100	6 791	NA	6 791
300	17 291	NA	17 291



### 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	<a href="https://github.com/ldbc/ldbc_snb_interactive_driver/releases/tag/v1.2.0">https://github.com/ldbc/ldbc_snb_interactive_driver/releases/tag/v1.2.0</a>
Implementations version	v1.0.0	<a href="https://github.com/ldbc/ldbc_snb_interactive_impls/releases/tag/1.0.0">https://github.com/ldbc/ldbc_snb_interactive_impls/releases/tag/1.0.0</a>
LDBC SNB specification version	v0.3.6	<a href="https://arxiv.org/pdf/2001.02299v3.pdf">https://arxiv.org/pdf/2001.02299v3.pdf</a>
Driver read threads	64	
Driver write threads	64	

#### 3.1 Driver implementation

A test driver adaptation for the SUT was provided by the test sponsor. 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 `io.atlasgraph.interactive.AtlasGraphInteractiveDb` extends the class `com.ldbc.driver.Db` provided in the LDBC SNB Interactive driver package. Internally, the `AtlasGraphInteractiveDb` relies on Google's packages `grpc` and `protobuf` for remote procedure calls (RPC) to communicate with the SUT.

#### 3.2 Benchmark configuration of driver

The driver applied the following time compression ratio values:

- TCR=0.001 for scale factor 30,
- TCR=0.0026 for scale factor 100,
- TCR=0.0092 for scale factor 300,

The complete configuration files for the different scale factors are shown in Listings A.12–A.14, and are also included in the attached supplementary materials.

## 4 PERFORMANCE RESULTS

The performance results reported here show benchmark executions with scale factors 30, 100, and 300. The performance summary tables below highlight key performance characteristics. In each case, the query on-time compliance is higher than the minimum required 95% <sup>1</sup>. The performance summary tables below highlight key performance characteristics.

Table 4.1: Summary of results for scale factor 30

<b>Benchmark duration</b>	<b>Benchmark operations</b>	<b>Throughput</b>	<b>Query on-time compliance</b>
02h 00m 27.559s	271 982 072	37 631.25 $\frac{\text{operations}}{\text{second}}$	99.86%

Table 4.2: Summary of results for scale factor 100

<b>Benchmark duration</b>	<b>Benchmark operations</b>	<b>Throughput</b>	<b>Query on-time compliance</b>
02h 01m 35.871s	355 776 423	48 764.08 $\frac{\text{operations}}{\text{second}}$	96.22%

Table 4.3: Summary of results for scale factor 300

<b>Benchmark duration</b>	<b>Benchmark operations</b>	<b>Throughput</b>	<b>Query on-time compliance</b>
02h 04m 58.344s	362 257 201	48 311.63 $\frac{\text{operations}}{\text{second}}$	95.24%

A remark on the throughput on scale factor 30: the system under test achieved the highest possible throughput allowed by SNB Interactive. This is due to an inherent limitation of the benchmark framework. The benchmark specification prescribes a 2.5h minimum execution time (0.5h warmup, 2h measurement window), while there is also a finite amount of update operations available from the update stream. On SF30, the available update stream is exhausted before the benchmark's runtime could reach the prescribed minimum execution time, as having a higher throughput would result in an invalid benchmark run.

<sup>1</sup>The total number of late operations for each run in the results in the attachment is referred to as `excessive_delay_count`.

## Performance Results

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

Query	Total count	Min.	Max.	Mean	P <sub>50</sub>	P <sub>90</sub>	P <sub>95</sub>	P <sub>99</sub>
Complex 1	1 780 974	546	51 796	1 559.98	1 126	1 689	2 043	19 029
Complex 2	1 251 496	184	20 825	454.21	381	688	916	1 513
Complex 3	436 843	4 431	40 572	8 277.18	8 385	10 437	11 169	13 288
Complex 4	1 286 259	255	26 869	910.53	803	1 434	1 694	2 318
Complex 5	643 130	36 796	597 440	184 008.47	165 072	308 928	337 584	393 904
Complex 6	146 535	183	27 476	971.68	970	1 832	2 072	2 720
Complex 7	964 695	131	28 433	365.94	292	583	816	1 414
Complex 8	5 145 038	741	26 508	2 577.20	2 115	4 476	5 307	6 838
Complex 9	120 587	6 462	55 402	15 704.03	15 161	21 660	23 748	28 281
Complex 10	1 251 496	4 679	53 610	13 170.75	12 811	17 902	19 449	22 900
Complex 11	2 315 267	171	26 912	403.14	331	615	852	1 462
Complex 12	1 052 394	2 760	49 046	10 137.40	9 669	14 536	16 209	19 818
Complex 13	2 437 123	238	26 972	914.52	830	1 336	1 570	2 180
Complex 14	945 007	600	317 216	11 490.48	5 016	10 761	73 556	146 592
Short 1	24 823 319	94	22 948	239.02	171	421	653	1 169
Short 2	24 823 319	100	25 072	302.66	236	499	707	1 214
Short 3	24 823 319	102	26 307	408.72	292	819	1 117	1 656
Short 4	24 823 705	97	22 011	240.16	173	417	648	1 167
Short 5	24 823 705	94	20 254	232.48	167	410	624	1 138
Short 6	24 823 705	95	22 168	232.91	169	403	616	1 128
Short 7	24 823 705	107	30 206	325.97	264	517	726	1 236
Update 1	14 143	185	13 916	434.41	341	645	900	1 680
Update 2	11 307 701	145	28 617	342.84	243	609	865	1 548
Update 3	12 121 088	145	25 678	324.50	239	539	789	1 449
Update 4	253 334	159	15 925	344.27	255	564	811	1 484
Update 5	40 999 247	143	30 390	351.68	244	642	900	1 593
Update 6	3 278 856	156	23 386	349.77	262	568	810	1 455
Update 7	9 457 383	157	27 378	355.54	263	587	841	1 518
Update 8	1 008 699	148	26 946	326.45	241	543	792	1 453

## Performance Results

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

Query	Total count	Min.	Max.	Mean	P <sub>50</sub>	P <sub>90</sub>	P <sub>95</sub>	P <sub>99</sub>
Complex 1	2 157 011	178	88 324	4 305.33	2 148	2 851	31 856	42 838
Complex 2	1 515 737	136	15 307	652.76	592	997	1 154	1 593
Complex 3	455 954	12 397	68 928	25 343.93	25 281	29 965	32 654	36 132
Complex 4	1 557 841	221	27 897	1 247.22	1 158	1 877	2 128	2 700
Complex 5	719 004	1 999	510 736	202 482.60	202 200	274 560	294 064	344 208
Complex 6	129 222	239	27 371	1 862.82	1 390	3 615	3 946	4 780
Complex 7	1 475 850	153	30 777	530.29	469	845	989	1 372
Complex 8	11 216 455	125	29 074	466.88	404	783	930	1 315
Complex 9	106 418	232	67 584	19 830.86	19 346	25 749	27 825	33 014
Complex 10	1 402 057	212	63 944	18 888.03	18 637	23 558	25 514	29 771
Complex 11	2 549 194	211	28 464	617.01	555	934	1 081	1 469
Complex 12	1 274 597	144	54 010	14 390.21	13 923	19 635	21 777	26 740
Complex 13	2 951 699	498	34 720	2 238.60	2 184	2 806	3 006	3 525
Complex 14	1 144 536	1 808	333 888	36 984.07	8 019	120 024	136 808	169 800
Short 1	35 968 848	108	24 173	404.10	343	716	856	1 205
Short 2	35 968 848	115	24 377	483.59	429	795	933	1 278
Short 3	35 968 848	118	34 378	636.52	526	1 124	1 508	2 181
Short 4	35 967 338	107	78 912	399.78	337	706	846	1 199
Short 5	35 967 338	103	85 204	399.18	341	706	844	1 189
Short 6	35 967 338	106	78 616	399.66	341	705	842	1 185
Short 7	35 967 338	136	72 708	532.82	476	841	979	1 325
Update 1	14 210	222	13 256	628.66	548	988	1 161	1 815
Update 2	10 530 977	161	32 269	521.76	444	860	1 034	1 610
Update 3	15 798 757	163	74 776	525.88	449	863	1 036	1 612
Update 4	247 705	181	18 122	531.65	450	874	1 050	1 633
Update 5	33 125 847	160	36 338	517.82	438	858	1 032	1 609
Update 6	3 401 900	182	64 566	556.60	474	910	1 090	1 705
Update 7	11 086 081	179	265 296	563.09	485	900	1 074	1 684
Update 8	1 139 475	170	18 990	513.01	431	852	1 027	1 593

## Performance Results

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

Query	Total count	Min.	Max.	Mean	P <sub>50</sub>	P <sub>90</sub>	P <sub>95</sub>	P <sub>99</sub>
Complex 1	1 801 094	195	154 696	4 442.16	3 681	4 326	4 568	51 194
Complex 2	1 265 633	148	31 367	702.83	655	1 007	1 148	1 518
Complex 3	329 777	33 636	184 624	70 788.79	69 992	83 376	90 416	100 880
Complex 4	1 300 790	215	30 618	1 509.63	1 397	2 266	2 508	3 109
Complex 5	557 481	341	623 936	235 668.24	240 512	328 000	351 504	412 768
Complex 6	80 739	324	21 866	3 613.31	1 508	7 447	8 091	9 948
Complex 7	1 463 389	170	27 828	587.43	542	867	996	1 335
Complex 8	15 609 481	123	23 077	502.40	457	774	901	1 235
Complex 9	66 423	216	79 832	22 866.89	22 198	30 406	32 914	39 256
Complex 10	1 064 283	207	86 376	23 253.75	22 924	29 734	32 293	37 980
Complex 11	1 951 185	263	26 581	755.87	710	1 032	1 164	1 514
Complex 12	1 064 283	159	76 208	16 308.88	15 821	23 620	26 322	32 452
Complex 13	2 464 655	244	38 564	4 525.52	4 468	5 548	5 876	6 714
Complex 14	955 682	342	571 808	61 551.84	14 364	147 528	166 584	205 144
Short 1	37 621 276	103	29 430	449.92	406	722	845	1 160
Short 2	37 621 276	125	29 959	537.93	496	809	933	1 246
Short 3	37 621 276	119	39 540	732.23	616	1 196	1 677	2 524
Short 4	37 620 597	103	32 038	441.89	397	709	833	1 150
Short 5	37 620 597	103	23 383	444.41	402	711	833	1 145
Short 6	37 620 597	102	23 536	445.27	403	711	833	1 144
Short 7	37 620 597	148	34 610	637.40	596	906	1 030	1 348
Update 1	10 616	247	10 506	685.56	613	1 012	1 198	1 899
Update 2	9 517 761	168	36 534	558.21	484	852	1 018	1 706
Update 3	18 116 042	162	35 672	556.57	483	850	1 015	1 709
Update 4	181 915	183	19 457	584.62	510	882	1 051	1 740
Update 5	25 167 935	165	27 731	563.10	489	857	1 021	1 691
Update 6	2 787 820	189	25 973	619.16	534	949	1 145	1 945
Update 7	12 213 388	179	34 036	572.15	495	876	1 048	1 828
Update 8	940 613	174	23 101	564.46	491	858	1 022	1 704

## 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 downloaded from the LDBC repository<sup>1</sup> and it was originally created with the SNB Interactive reference implementation over Neo4j. The system with the driver configuration shown in Listing A.15 successfully returned the expected result sets for the queries of the benchmark.

---

<sup>1</sup>[https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/interactive-v1/validation\\_params-interactive-v1.0.0-sf0.1-to-sf10.tar.zst](https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/interactive-v1/validation_params-interactive-v1.0.0-sf0.1-to-sf10.tar.zst)

## 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, and therefore complies with the requirements of the benchmark.

### 6.2 SNB 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 shutdown -r` after executing 260 million operations.

#### 6.3.2 Durability

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

```

1 $ grep LdbcUpdate1 LDBC-SNB-results_log.csv | tail -n 1
2 LdbcUpdate1AddPerson|1701303000059|1701303000059|284|0|1354402963802
3 $ grep LdbcUpdate2 LDBC-SNB-results_log.csv | tail -n 1
4 LdbcUpdate2AddPostLike|1701303000083|1701303000133|16484|0|1354402987638
5 $ grep LdbcUpdate3 LDBC-SNB-results_log.csv | tail -n 1
6 LdbcUpdate3AddCommentLike|1701303000080|1701303000133|30553|0|1354402984097
7 $ grep LdbcUpdate4 LDBC-SNB-results_log.csv | tail -n 1
8 LdbcUpdate4AddForum|1701303000076|1701303000076|361|0|1354402980867
9 $ grep LdbcUpdate5 LDBC-SNB-results_log.csv | tail -n 1
10 LdbcUpdate5AddForumMembership|1701303000082|1701303000133|30406|0|1354402986025
11 $ grep LdbcUpdate6 LDBC-SNB-results_log.csv | tail -n 1
12 LdbcUpdate6AddPost|1701303000084|1701303000132|9176|0|1354402989259
13 $ grep LdbcUpdate7 LDBC-SNB-results_log.csv | tail -n 1
14 LdbcUpdate7AddComment|1701303000083|1701303000133|16501|0|1354402988423
15 $ grep LdbcUpdate8 LDBC-SNB-results_log.csv | tail -n 1
16 LdbcUpdate8AddFriendship|1701303000090|1701303000133|12583|0|1354402994405

```

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 TCR 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 as shown below.

```

1 $ grep -rnw /disk1/SF30/social_network/ -e '1354402963802|. *|1|. *'
2 /disk1/SF30/social_network/updateStream_0_26_person.csv:430:1354402963802|0|1|37383395469397|Sanjay|Sharma|male
   |501724800000|1354402963802|1.38.110.24|Firefox|121|hi;en|Sanjay37383395469397@gmail.com;
   |Sanjay37383395469397@yahoo.com|579;1444;2062;7596;9095||544,2012;550,2004;543,2004
3 $ grep -rnw /disk1/SF30/social_network/ -e '1354402987638|. *|2|. *'
4 /disk1/SF30/social_network/updateStream_0_31_forum.csv
   :2108108:1354402987638|1347788965705|2|35184372230086|10995153908946|1354402987638

```



```

5 $ grep -rnw /disk1/SF30/social_network/ -e '1354402984097|.|*|3|.*'
6 /disk1/SF30/social_network/updateStream_0_15_forum.csv
   :2038230:1354402984097|1318413545428|3|21990232704196|37383546552852|1354402984097
7 $ grep -rnw /disk1/SF30/social_network/ -e '1354402980867|.|*|4|.*'
8 /disk1/SF30/social_network/updateStream_0_28_forum.csv:2064142:1354402980867|1301312694788|4|37383404883722|Album
   5 of K. Khan|1354402980867|15393162805969|2820
9 $ grep -rnw /disk1/SF30/social_network/ -e '1354402986025|.|*|5|.*'
10 /disk1/SF30/social_network/updateStream_0_17_forum.csv
   :1980907:1354402986025|1352796797617|5|10995116320792|37383395356068|1354402986025
11 $ grep -rnw /disk1/SF30/social_network/ -e '1354402989259|.|*|6|.*'
12 /disk1/SF30/social_network/updateStream_0_7_forum.csv
   :1995997:1354402989259|1269076577759|6|37383400272466||1354402989259|119.15.85.3|Firefox|tk|About Julius
   Caesar, government, he began extensive reforms of Roman society and government. He cent
   |101|2199023318662|26388283274501|67|5191
13 $ grep -rnw /disk1/SF30/social_network/ -e '1354402988423|.|*|7|.*'
14 /disk1/SF30/social_network/updateStream_0_5_forum.csv
   :2053658:1354402988423|1307118727048|7|37383508609169|1354402988423|196.223.2.51|Chrome|good
   |4|17592186095418|110|-1|37383508609160|
15 $ grep -rnw /disk1/SF30/social_network/ -e '1354402994405|.|*|8|.*'
16 /disk1/SF30/social_network/updateStream_0_18_forum.csv
   :1944413:1354402994405|1352311816936|8|8796093024637|35184372181281|1354402994405

```

To check whether the graph entities in the driver log entries were persisted in the database, read queries targeting the last completed update operations were executed after database restart. The queries returned the data that was committed according to the logs, so the system passed this check. These queries and their output is shown below using the Cypher REPL of the system (output is slightly formatted for better presentation in this document).

```

1 atlas[ldbc]> MATCH (v:person{id:37383395469397}) RETURN v;
2 (10-2200001e855 :person {"firstName":"Sanjay","lastName":"Sharma","locationIP":"1.38.110.24","gender":"male",
   "creationDate":"2012-12-01 23:02:43.802 UTC","browserUsed":"Firefox","birthday":"1985-11-25","id
   ":37383395469397,"language":"hi,en","email":"Sanjay37383395469397@gmail.com;Sanjay37383395469397@yahoo.com"}
   )
3 1 lines returned.
4 Time: total 3.00 ms, compile 0.32 ms, optimization 0.23 ms, execution 2.37 ms
5 Client time: 5.59
6 atlas[ldbc]> MATCH (:post{id:10995153908946})<-[e:likes]-(:person{id:35184372230086}) RETURN e;
7 (10-200000227c6)-[1c-2000005278a55:10-200000227c6:e-a00023e34d2 :likes {"creationDate":"2012-12-01 23:03:07.638
   UTC"}]->(e-a00023e34d2)
8 1 lines returned.
9 Time: total 13.90 ms, compile 0.36 ms, optimization 0.43 ms, execution 13.00 ms
10 Client time: 16.35
11 atlas[ldbc]> MATCH (:comment{id:37383546552852})<-[e:likes]-(:person{id:21990232704196}) RETURN e;
12 (10-1400000244c4)-[1c-20000052a9e14:10-1400000244c4:a-220009034214 :likes {"creationDate":"2012-12-01
   23:03:04.097 UTC"}]->(a-220009034214)
13 1 lines returned.
14 Time: total 13.55 ms, compile 0.33 ms, optimization 0.43 ms, execution 12.64 ms
15 Client time: 16.35
16 atlas[ldbc]> MATCH (v:forum{id:37383404883722}) RETURN v;
17 (8-220000918f0a :forum {"id":"37383404883722","creationDate":"2012-12-01 23:03:00.867 UTC","title":"Album 5 of K.
   Khan"} )
18 1 lines returned.
19 Time: total 2.85 ms, compile 0.21 ms, optimization 0.22 ms, execution 2.34 ms
20 Client time: 5.00
21 atlas[ldbc]> MATCH (:person{id:37383395356068})<-[e:hasMember]-(:forum{id:10995116320792}) RETURN e;
22 (8-a000000a818)-[15-20000050f77a5:8-a000000a818:10-22000002da4 :hasMember {"joinDate":"2012-12-01 23:03:06.025
   UTC"}]->(10-22000002da4)
23 1 lines returned.

```



```

24 Time: total 12.64 ms, compile 0.26 ms, optimization 0.35 ms, execution 11.96 ms
25 Client time: 15.16
26 atlas[ldbc]> MATCH (v:post{id:37383400272466}) RETURN v;
27 (e-2200004b3252 :post {"content":"About Julius Caesar, government, he began extensive reforms of Roman society
    and government. He cent","imageFile":"","language":"tk","creationDate":"2012-12-01 23:03:09.259 UTC","length
    ":101,"id":37383400272466,"locationIP":"119.15.85.3","browserUsed":"Firefox"} )
28 1 lines returned.
29 Time: total 2.66 ms, compile 0.29 ms, optimization 0.23 ms, execution 2.07 ms
30 Client time: 5.19
31 atlas[ldbc]> MATCH (v:'comment{id:37383508609169}) RETURN v;
32 (a-220006c04891 :comment {"browserUsed":"Chrome","length":4,"id":37383508609169,"locationIP":"196.223.2.51","
    creationDate":"2012-12-01 23:03:08.423 UTC","content":"good"} )
33 1 lines returned.
34 Time: total 3.12 ms, compile 0.23 ms, optimization 0.25 ms, execution 2.57 ms
35 Client time: 5.34
36 atlas[ldbc]> MATCH (:person{id:35184372181281})<-[e:knows]-(:person{id:8796093024637}) RETURN e;
37 (10-8000000097d)-[12-2000005f915a4:10-8000000097d:10-200000016921 :knows {"creationDate":"2012-12-01 23:03:14.405
    UTC"}]->(10-200000016921)
38 1 lines returned.
39 Time: total 13.92 ms, compile 0.35 ms, optimization 0.37 ms, execution 13.09 ms
40 Client time: 16.58

```

### 6.3.3 Consistency after recovery

The benchmark relied on no precomputed auxiliary data structures, thus no additional steps were required to check consistency. After the crash, the database was successfully started using the usual startup procedure.

## 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

<b>File</b>	<b>Purpose</b>
README.md	Guide with exact steps for tool installation and benchmarking
results-sf{30,100,300}.tar.gz	Driver output files for the selected scale factors
acid-test	Rust project with ACID test suite implementation
atlasgraph-3.0.0.deb	AtlasGraph binary package
query-impl	Folder containing the implementations of benchmark queries
data-generation	Folder with datagen parameters for the used scale factors
data-loading	Graph schema and CSV data mapping
data-preprocessing	CSV converter scripts and utilities
deployment	Resources for deployment of AtlasGraph using Ansible
driver	LDBC SNB driver resources
driver/atlasgraph{/src}	AtlasGraph-specific LDBC SNB driver resources and sources
build.sh	LDBC SNB driver build script
run.sh	AtlasGraph-specific LDBC SNB driver run script
*.properties	Driver parameters for SF30, SF100, SF300, and validation (SF10)

## Supplementary Materials

---

The attachment folder directory structure is as follows:

```
attachments
├── README.md
├── results-sf30.tar.gz
├── results-sf100.tar.gz
├── results-sf300.tar.gz
├── acid-test
├── atlasgraph-3.0.0.deb
├── query-impl
│   └── «query implementation file».rs
├── data-generation
│   ├── params-sf10.ini
│   ├── params-sf100.ini
│   ├── params-sf30.ini
│   └── params-sf300.ini
├── data-loading
│   ├── mapping_sf10.json
│   ├── mapping_sf100.json
│   ├── mapping_sf30.json
│   ├── mapping_sf300.json
│   └── schema.json
├── data-preprocessing
│   ├── add_entity.py
│   ├── prepare_snb.sh
│   ├── split_entity.py
│   └── split_rel.py
├── deployment
│   ├── atlas.yml
│   └── templates
├── driver
│   ├── scripts
│   │   └── build.sh
│   ├── atlasgraph
│   │   ├── benchmark-sf100.properties
│   │   ├── benchmark-sf30.properties
│   │   ├── benchmark-sf300.properties
│   │   ├── run.sh
│   │   ├── src
│   │   └── validate.properties
```

## Appendix

# A APPENDIX

## A.1 CPU details

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

```

1 processor : 0
2 vendor_id : GenuineIntel
3 cpu family : 6
4 model    : 85
5 model name : Intel(R) Xeon(R) Platinum 8259CL CPU @ 2.50GHz
6 stepping : 7
7 microcode : 0x5003604
8 cpu MHz   : 2499.998
9 cache size : 36608 KB
10 physical id : 0
11 siblings : 48
12 core id   : 23
13 cpu cores : 24
14 apicid    : 47
15 initial apicid : 47
16 fpu       : yes
17 fpu_exception : yes
18 cpuid level : 13
19 wp        : yes
20 flags     : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse sse2 ss ht
      syscall nx pdpe1gb rdtscp lm constant_tsc arch_perfmon rep_good nopl xtopology nonstop_tsc cpuid aperfmperf
      tsc_known_freq pni pclmulqdq monitor ssse3 fma cx16 pcid sse4_1 sse4_2 x2apic movbe popcnt tsc_deadline_timer
      aes xsave avx f16c rdrand hypervisor lahf_lm abm 3dnowprefetch invpcid_single pti fsgsbase tsc_adjust bmi1
      avx2 smep bmi2 erms invpcid mpx avx512f avx512dq rdseed adx smap clflushopt clwb avx512cd avx512bw avx512vl
      xsaveopt xsavec xgetbv1 xsaves ida arat pku ospke
21 bugs      : cpu_meltdown spectre_v1 spectre_v2 spec_store_bypass l1tf mds swapgs itlb_multihit mmio_stale_data
      retbleed gds
22 bogomips : 4999.99
23 clflush size : 64
24 cache_alignment : 64
25 address sizes : 46 bits physical, 48 bits virtual
26 power management:

```

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

```

1 Architecture:                x86_64
2 CPU op-mode(s):              32-bit, 64-bit
3 Byte Order:                   Little Endian
4 Address sizes:                46 bits physical, 48 bits virtual
5 CPU(s):                       48
6 On-line CPU(s) list:         0-47
7 Thread(s) per core:          2
8 Core(s) per socket:          24
9 Socket(s):                     1
10 NUMA node(s):                1
11 Vendor ID:                    GenuineIntel
12 CPU family:                   6
13 Model:                        85
14 Model name:                   Intel(R) Xeon(R) Platinum 8259CL CPU @ 2.50GHz
15 Stepping:                     7
16 CPU MHz:                      2499.998

```



```

17 BogoMIPS: 4999.99
18 Hypervisor vendor: KVM
19 Virtualization type: full
20 L1d cache: 768 KiB
21 L1i cache: 768 KiB
22 L2 cache: 24 MiB
23 L3 cache: 35.8 MiB
24 NUMA node0 CPU(s): 0-47
25 Vulnerability Gather data sampling: Unknown: Dependent on hypervisor status
26 Vulnerability Itlb multihit: KVM: Mitigation: VMX unsupported
27 Vulnerability L1tf: Mitigation; PTE Inversion
28 Vulnerability Mds: Vulnerable: Clear CPU buffers attempted, no microcode; SMT Host state unknown
29 Vulnerability Meltdown: Mitigation; PTI
30 Vulnerability Mmio stale data: Vulnerable: Clear CPU buffers attempted, no microcode; SMT Host state unknown
31 Vulnerability Retbleed: Vulnerable
32 Vulnerability Spec rstack overflow: Not affected
33 Vulnerability Spec store bypass: Vulnerable
34 Vulnerability Spectre v1: Mitigation; usercopy/swapgs barriers and __user pointer sanitization
35 Vulnerability Spectre v2: Mitigation; Retpolines, STIBP disabled, RSB filling, PBRSB-eIBRS Not affected
36 Vulnerability Srbds: Not affected
37 Vulnerability Tsx async abort: Not affected
38 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 xtopo
39    logy 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 hypervis
40    or lahf_lm abm 3dnowprefetch invpcid_single pti fsgsbase tsc_adjust bmi1 avx2
    smep bmi2 erms invpcid mpx avx512f avx512dq rdseed adx smap clflushopt clwb avx512cd avx512bw avx512
41    vl xsaveopt xsavec xgetbv1 xsaves ida arat pku ospke

```

## A.2 Memory details

Listing A.3: Output of the `lshw -c memory` command

```

1  *-firmware
2     description: BIOS
3     vendor: Amazon EC2
4     physical id: 0
5     version: 1.0
6     date: 10/16/2017
7     size: 64KiB
8     capacity: 64KiB
9     capabilities: pci edd acpi virtualmachine
10 *-cache:0
11    description: L1 cache
12    physical id: 5
13    slot: L1-Cache
14    size: 1536KiB
15    capacity: 1536KiB
16    capabilities: synchronous internal write-back instruction
17    configuration: level=1
18 *-cache:1
19    description: L2 cache
20    physical id: 6
21    slot: L2-Cache
22    size: 24MiB
23    capacity: 24MiB

```

```

24     capabilities: synchronous internal varies unified
25     configuration: level=2
26 *--cache:2
27     description: L3 cache
28     physical id: 7
29     slot: L3-Cache
30     size: 35MiB
31     capacity: 35MiB
32     capabilities: synchronous internal varies unified
33     configuration: level=3
34 *--memory
35     description: System Memory
36     physical id: 8
37     slot: System board or motherboard
38     size: 384GiB
39 *--bank
40     description: DIMM DDR4 Static column Pseudo-static Synchronous Window DRAM 3200 MHz (0.3 ns)
41     physical id: 0
42     size: 384GiB
43     width: 64 bits
44     clock: 3200MHz (0.3ns)

```

## A.3 Storage details and IO performance

Listing A.4: Output of the `lshw -c storage -c disk` command

```

1 *--storage:0
2     description: Non-Volatile memory controller
3     product: Amazon.com, Inc.
4     vendor: Amazon.com, Inc.
5     physical id: 4
6     bus info: pci@0000:00:04.0
7     version: 00
8     width: 32 bits
9     clock: 33MHz
10    capabilities: storage pciexpress msix nvme_express bus_master cap_list
11    configuration: driver=nvme latency=0
12    resources: irq:11 memory:feb0000-febf3fff
13 *--nvme0
14     description: NVMe device
15     product: Amazon Elastic Block Store
16     physical id: 0
17     logical name: /dev/nvme0
18     version: 1.0
19     serial: vol0e0c715add7e6fcc5
20     configuration: nqn=nqn.2014.08.org.nvmeexpress:1d0f1d0fvol0e0c715add7e6fcc5Amazon Elastic Block Store
state=live
21 *--namespace
22     description: NVMe namespace
23     physical id: 1
24     logical name: /dev/nvme0n1
25     size: 8GiB (8589MB)
26     capabilities: gpt-1.00 partitioned partitioned:gpt
27     configuration: guid=97b80dc3-88a5-43b8-95b7-7849c7006911 logicalsectorsize=512 sectorsize=512
28 *--storage:1
29     description: Non-Volatile memory controller

```

```

30     product: NVMe SSD Controller
31     vendor: Amazon.com, Inc.
32     physical id: 1e
33     bus info: pci@0000:00:1e.0
34     version: 00
35     width: 32 bits
36     clock: 33MHz
37     capabilities: storage pciexpress msix nvm_express bus_master cap_list
38     configuration: driver=nvme latency=0
39     resources: irq:0 memory:febfb000-febfbfff memory:fe900000-fe901fff
40     *-nvme1
41         description: NVMe device
42         product: Amazon EC2 NVMe Instance Storage
43         physical id: 0
44         logical name: /dev/nvme1
45         version: 0
46         serial: AWS22C659696258E8449
47         configuration: nqn=nqn.2014.08.org.nvmeexpress:1d0f0000AWS22C659696258E8449Amazon EC2 NVMe Instance
Storage state=live
48     *-namespace
49         description: NVMe namespace
50         physical id: 1
51         logical name: /dev/nvme1n1
52         logical name: /disk1
53         size: 838GiB (900GB)
54         configuration: logicalsectorsize=512 mount.fstype=xfs mount.options=rw,relatime,attr2,inode64,
logbufs=8,logbsize=32k,noquota sectorsize=512 state=mounted
55     *-storage:2
56         description: Non-Volatile memory controller
57         product: NVMe SSD Controller
58         vendor: Amazon.com, Inc.
59         physical id: 1f
60         bus info: pci@0000:00:1f.0
61         version: 00
62         width: 32 bits
63         clock: 33MHz
64         capabilities: storage pciexpress msix nvm_express bus_master cap_list
65         configuration: driver=nvme latency=0
66         resources: irq:0 memory:febfc000-febfffff memory:fe902000-fe903fff
67     *-nvme2
68         description: NVMe device
69         product: Amazon EC2 NVMe Instance Storage
70         physical id: 0
71         logical name: /dev/nvme2
72         version: 0
73         serial: AWS1B062B7153D68A2C7
74         configuration: nqn=nqn.2014.08.org.nvmeexpress:1d0f0000AWS1B062B7153D68A2C7Amazon EC2 NVMe Instance
Storage state=live
75     *-namespace
76         description: NVMe namespace
77         physical id: 1
78         logical name: /dev/nvme2n1
79         logical name: /disk2
80         size: 838GiB (900GB)
81         configuration: logicalsectorsize=512 mount.fstype=xfs mount.options=rw,relatime,attr2,inode64,
logbufs=8,logbsize=32k,noquota sectorsize=512 state=mounted

```

Listing A.5: Output of the fio command

```

1 $ fio --rw=write --ioengine=sync --fdatasync=1 --direct=1 --directory=io-test-data --size=2g --bs=4k --name=
  iotest
2 iotest: (g=0): rw=write, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B, ioengine=sync, iodepth=1
3 fio-3.16
4 Starting 1 process
5 iotest: Laying out IO file (1 file / 2048MiB)
6 Jobs: 1 (f=1): [W(1)][100.0%][w=66.4MiB/s][w=17.0k IOPS][eta 00m:00s]
7 iotest: (groupid=0, jobs=1): err= 0: pid=88874: Tue Dec 12 23:22:49 2023
8   write: IOPS=16.9k, BW=66.0MiB/s (69.2MB/s)(2048MiB/31025msec); 0 zone resets
9     clat (usec): min=23, max=144, avg=26.32, stdev= 2.92
10    lat (usec): min=23, max=144, avg=26.41, stdev= 2.93
11    clat percentiles (nsec):
12      | 1.00th=[24448],  5.00th=[24704], 10.00th=[24960], 20.00th=[25216],
13      | 30.00th=[25216], 40.00th=[25472], 50.00th=[25472], 60.00th=[25728],
14      | 70.00th=[25984], 80.00th=[26240], 90.00th=[27776], 95.00th=[31616],
15      | 99.00th=[40704], 99.50th=[43264], 99.90th=[51968], 99.95th=[55552],
16      | 99.99th=[65280]
17    bw (  KiB/s): min=62608, max=68864, per=99.99%, avg=67590.10, stdev=1035.48, samples=62
18    iops        : min=15652, max=17216, avg=16897.53, stdev=258.89, samples=62
19    lat (usec)  : 50=99.86%, 100=0.14%, 250=0.01%
20    fsync/fdatasync/sync_file_range:
21      sync (usec): min=27, max=3651, avg=32.06, stdev=13.04
22      sync percentiles (usec):
23        | 1.00th=[ 30],  5.00th=[ 30], 10.00th=[ 30], 20.00th=[ 31],
24        | 30.00th=[ 31], 40.00th=[ 31], 50.00th=[ 31], 60.00th=[ 32],
25        | 70.00th=[ 32], 80.00th=[ 33], 90.00th=[ 35], 95.00th=[ 40],
26        | 99.00th=[ 50], 99.50th=[ 54], 99.90th=[ 63], 99.95th=[ 70],
27        | 99.99th=[ 408]
28    cpu         : usr=3.96%, sys=17.77%, ctx=1572860, majf=0, minf=13
29    IO depths   : 1=200.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
30    submit     : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
31    complete   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
32    issued rwts: total=0,524288,0,0 short=524287,0,0,0 dropped=0,0,0,0
33    latency    : target=0, window=0, percentile=100.00%, depth=1
34
35 Run status group 0 (all jobs):
36   WRITE: bw=66.0MiB/s (69.2MB/s), 66.0MiB/s-66.0MiB/s (69.2MB/s-69.2MB/s), io=2048MiB (2147MB), run=31025-31025
      msec
37
38 Disk stats (read/write):
39   nvme1n1: ios=0/1044049, merge=0/4, ticks=0/21540, in_queue=21540, util=99.75%

```

## A.4 Network performance

Listing A.6: Output of the iperf command

```

1 $ iperf -c 172.31.30.75 -r --parallel 48 -i 1 -t 2 -p 21021
2 -----
3 Client connecting to 172.31.30.75, TCP port 21021
4 TCP window size: 325 KByte (default)
5 -----
6 [ ID] Interval      Transfer    Bandwidth
7 [ 52] local 172.31.24.90 port 43420 connected with 172.31.30.75 port 21021
8 [ 52] 0.0- 0.0 sec  323 KBytes  11.2 Gbits/sec
9 [ 45] local 172.31.24.90 port 43356 connected with 172.31.30.75 port 21021

```





```

10 [ 45]  0.0- 0.0 sec   350 KBytes  12.5 Gbits/sec
11 ..
12 [SUM]  0.0- 0.0 sec  15.5 MBytes  14.1 Gbits/sec
13 ..
14 [ 52]  0.0- 2.0 sec   53.5 MBytes   221 Mbits/sec
15 [ 51]  0.0- 2.0 sec   56.5 MBytes   234 Mbits/sec
16 ..
17 [SUM]  0.0- 2.0 sec  4.20 GBytes  17.7 Gbits/sec

```

## A.5 Datagen configuration

Listing A.7: Contents of `params-sf30.ini` used for scale factor 30

```

1 ldbc.snb.datagen.generator.scaleFactor:snb.interactive.30
2 ldbc.snb.datagen.serializer.numUpdatePartitions:32
3
4 ldbc.snb.datagen.serializer.dynamicActivitySerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.
   activity.CsvCompositeDynamicActivitySerializer
5 ldbc.snb.datagen.serializer.dynamicPersonSerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.person.
   CsvCompositeDynamicPersonSerializer
6 ldbc.snb.datagen.serializer.staticSerializer:ldbc.snb.datagen.serializer.snb.csv.staticserializer.
   CsvCompositeStaticSerializer

```

Listing A.8: Contents of `params-sf100.ini` used for scale factor 100

```

1 ldbc.snb.datagen.generator.scaleFactor:snb.interactive.100
2 ldbc.snb.datagen.serializer.numUpdatePartitions:32
3
4 ldbc.snb.datagen.serializer.dynamicActivitySerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.
   activity.CsvCompositeDynamicActivitySerializer
5 ldbc.snb.datagen.serializer.dynamicPersonSerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.person.
   CsvCompositeDynamicPersonSerializer
6 ldbc.snb.datagen.serializer.staticSerializer:ldbc.snb.datagen.serializer.snb.csv.staticserializer.
   CsvCompositeStaticSerializer

```

Listing A.9: Contents of `params-sf300.ini` used for scale factor 300

```

1 ldbc.snb.datagen.generator.scaleFactor:snb.interactive.300
2 ldbc.snb.datagen.serializer.numUpdatePartitions:32
3
4 ldbc.snb.datagen.serializer.dynamicActivitySerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.
   activity.CsvCompositeDynamicActivitySerializer
5 ldbc.snb.datagen.serializer.dynamicPersonSerializer:ldbc.snb.datagen.serializer.snb.csv.dynamicserializer.person.
   CsvCompositeDynamicPersonSerializer
6 ldbc.snb.datagen.serializer.staticSerializer:ldbc.snb.datagen.serializer.snb.csv.staticserializer.
   CsvCompositeStaticSerializer

```

## A.6 Import configuration

Listing A.10: Content of `schema.json` describing the data schema

```

1 {
2   "graph_name": "ldbc",
3   "shard_num": 1,

```

```
4   "replica_factor": 1,
5   "vertices": {
6     "tagclass": {
7       "properties": [
8         {
9           "name": "id",
10          "type": "int64",
11          "is_pk": true,
12          "not_null": true
13        },
14        {
15          "name": "name",
16          "type": "string"
17        },
18        {
19          "name": "url",
20          "type": "string"
21        }
22      ],
23      "max_vertex_num": 71
24    },
25    "tag": {
26      "properties": [
27        {
28          "name": "id",
29          "type": "int64",
30          "is_pk": true,
31          "not_null": true
32        },
33        {
34          "name": "name",
35          "type": "string"
36        },
37        {
38          "name": "url",
39          "type": "string"
40        }
41      ],
42      "max_vertex_num": 16080
43    },
44    "person": {
45      "properties": [
46        {
47          "name": "id",
48          "type": "int64",
49          "is_pk": true,
50          "not_null": true
51        },
52        {
53          "name": "firstName",
54          "type": "string"
55        },
56        {
57          "name": "lastName",
58          "type": "string"
59        },
60        {
61          "name": "gender",
```

```
62         "type": "string"
63     },
64     {
65         "name": "birthday",
66         "type": "date"
67     },
68     {
69         "name": "email",
70         "type": "string"
71     },
72     {
73         "name": "browserUsed",
74         "type": "string"
75     },
76     {
77         "name": "locationIP",
78         "type": "string"
79     },
80     {
81         "name": "language",
82         "type": "string"
83     },
84     {
85         "name": "creationDate",
86         "type": "datetime"
87     }
88 ],
89     "max_vertex_num": 1254000
90 },
91 "continent": {
92     "properties": [
93         {
94             "name": "id",
95             "type": "int64",
96             "is_pk": true,
97             "not_null": true
98         },
99         {
100            "name": "name",
101            "type": "string"
102        },
103        {
104            "name": "url",
105            "type": "string"
106        }
107    ],
108    "max_vertex_num": 6
109 },
110 "city": {
111     "properties": [
112         {
113             "name": "id",
114             "type": "int64",
115             "is_pk": true,
116             "not_null": true
117         },
118         {
119             "name": "name",
```

```
120         "type": "string"
121     },
122     {
123         "name": "url",
124         "type": "string"
125     }
126 ],
127     "max_vertex_num": 1343
128 },
129     "country": {
130         "properties": [
131             {
132                 "name": "id",
133                 "type": "int64",
134                 "is_pk": true,
135                 "not_null": true
136             },
137             {
138                 "name": "name",
139                 "type": "string"
140             },
141             {
142                 "name": "url",
143                 "type": "string"
144             }
145         ],
146         "max_vertex_num": 111
147     },
148     "university": {
149         "properties": [
150             {
151                 "name": "id",
152                 "type": "int64",
153                 "is_pk": true,
154                 "not_null": true
155             },
156             {
157                 "name": "name",
158                 "type": "string"
159             },
160             {
161                 "name": "url",
162                 "type": "string"
163             }
164         ],
165         "max_vertex_num": 6380
166     },
167     "company": {
168         "properties": [
169             {
170                 "name": "id",
171                 "type": "int64",
172                 "is_pk": true,
173                 "not_null": true
174             },
175             {
176                 "name": "name",
177                 "type": "string"
```

```
178         },
179         {
180             "name": "url",
181             "type": "string"
182         }
183     ],
184     "max_vertex_num": 1575
185 },
186 "post": {
187     "properties": [
188         {
189             "name": "id",
190             "type": "int64",
191             "is_pk": true,
192             "not_null": true
193         },
194         {
195             "name": "content",
196             "type": "string"
197         },
198         {
199             "name": "length",
200             "type": "int32"
201         },
202         {
203             "name": "browserUsed",
204             "type": "string"
205         },
206         {
207             "name": "locationIP",
208             "type": "string"
209         },
210         {
211             "name": "creationDate",
212             "type": "datetime"
213         },
214         {
215             "name": "language",
216             "type": "string"
217         },
218         {
219             "name": "imageFile",
220             "type": "string"
221         }
222     ],
223     "max_vertex_num": 187893047
224 },
225 "comment": {
226     "properties": [
227         {
228             "name": "id",
229             "type": "int64",
230             "is_pk": true,
231             "not_null": true
232         },
233         {
234             "name": "content",
235             "type": "string"
```

```

236         },
237         {
238             "name": "length",
239             "type": "int32"
240         },
241         {
242             "name": "browserUsed",
243             "type": "string"
244         },
245         {
246             "name": "locationIP",
247             "type": "string"
248         },
249         {
250             "name": "creationDate",
251             "type": "datetime"
252         }
253     ],
254     "max_vertex_num": 762949871
255 },
256 "forum": {
257     "properties": [
258         {
259             "name": "id",
260             "type": "int64",
261             "is_pk": true,
262             "not_null": true
263         },
264         {
265             "name": "title",
266             "type": "string"
267         },
268         {
269             "name": "creationDate",
270             "type": "datetime"
271         }
272     ],
273     "max_vertex_num": 12549398
274 }
275 },
276 "edges": {
277     "isSubclassOf": {
278         "properties": [],
279         "relations": [
280             {
281                 "from_key_label_name": "tagclass",
282                 "to_key_label_name": "tagclass",
283                 "out_edge_strategy": "none"
284             }
285         ]
286     },
287     "hasType": {
288         "properties": [],
289         "relations": [
290             {
291                 "from_key_label_name": "tag",
292                 "to_key_label_name": "tagclass",
293                 "out_edge_strategy": "none"

```

```
294     }
295   ]
296 },
297 "knows": {
298   "properties": [
299     {
300       "name": "creationDate",
301       "type": "datetime"
302     }
303   ],
304   "relations": [
305     {
306       "from_key_label_name": "person",
307       "to_key_label_name": "person"
308     }
309   ]
310 },
311 "hasTag": {
312   "properties": [],
313   "relations": [
314     {
315       "from_key_label_name": "post",
316       "to_key_label_name": "tag"
317     },
318     {
319       "from_key_label_name": "comment",
320       "to_key_label_name": "tag",
321       "out_edge_strategy": "none"
322     },
323     {
324       "from_key_label_name": "forum",
325       "to_key_label_name": "tag",
326       "out_edge_strategy": "none"
327     }
328   ]
329 },
330 "hasModerator": {
331   "properties": [],
332   "relations": [
333     {
334       "from_key_label_name": "forum",
335       "to_key_label_name": "person",
336       "in_edge_strategy": "none",
337       "out_edge_strategy": "single"
338     }
339   ]
340 },
341 "hasMember": {
342   "properties": [
343     {
344       "name": "joinDate",
345       "type": "datetime",
346       "is_index": true
347     }
348   ],
349   "relations": [
350     {
351       "from_key_label_name": "forum",
```

```
352         "to_key_label_name": "person",
353         "out_edge_strategy": "none"
354     }
355 ]
356 },
357 "hasInterest": {
358     "properties": [],
359     "relations": [
360         {
361             "from_key_label_name": "person",
362             "to_key_label_name": "tag",
363             "in_edge_strategy": "none"
364         }
365     ]
366 },
367 "containerOf": {
368     "properties": [],
369     "relations": [
370         {
371             "from_key_label_name": "forum",
372             "to_key_label_name": "post",
373             "in_edge_strategy": "single",
374             "out_edge_strategy": "none"
375         }
376     ]
377 },
378 "replyOf": {
379     "properties": [],
380     "relations": [
381         {
382             "from_key_label_name": "comment",
383             "to_key_label_name": "post",
384             "out_edge_strategy": "single"
385         },
386         {
387             "from_key_label_name": "comment",
388             "to_key_label_name": "comment",
389             "out_edge_strategy": "single"
390         }
391     ]
392 },
393 "hasCreator": {
394     "properties": [
395         {
396             "name": "creationDate",
397             "type": "datetime",
398             "is_index": true
399         }
400     ],
401     "relations": [
402         {
403             "from_key_label_name": "post",
404             "to_key_label_name": "person",
405             "out_edge_strategy": "single"
406         },
407         {
408             "from_key_label_name": "comment",
409             "to_key_label_name": "person",
```



```
410         "out_edge_strategy": "single"
411     }
412 ]
413 },
414 "likes": {
415     "properties": [
416         {
417             "name": "creationDate",
418             "type": "datetime"
419         }
420     ],
421     "relations": [
422         {
423             "from_key_label_name": "person",
424             "to_key_label_name": "post",
425             "out_edge_strategy": "none"
426         },
427         {
428             "from_key_label_name": "person",
429             "to_key_label_name": "comment",
430             "out_edge_strategy": "none"
431         }
432     ]
433 },
434 "workAt": {
435     "properties": [
436         {
437             "name": "workFrom",
438             "type": "int32"
439         }
440     ],
441     "relations": [
442         {
443             "from_key_label_name": "person",
444             "to_key_label_name": "company"
445         }
446     ]
447 },
448 "studyAt": {
449     "properties": [
450         {
451             "name": "classYear",
452             "type": "int32"
453         }
454     ],
455     "relations": [
456         {
457             "from_key_label_name": "person",
458             "to_key_label_name": "university",
459             "in_edge_strategy": "none"
460         }
461     ]
462 },
463 "isLocatedIn": {
464     "properties": [],
465     "relations": [
466         {
467             "from_key_label_name": "person",
```

```

468         "to_key_label_name": "city",
469         "in_edge_strategy": "none",
470         "out_edge_strategy": "single"
471     },
472     {
473         "from_key_label_name": "post",
474         "to_key_label_name": "country",
475         "out_edge_strategy": "none"
476     },
477     {
478         "from_key_label_name": "comment",
479         "to_key_label_name": "country",
480         "out_edge_strategy": "none"
481     },
482     {
483         "from_key_label_name": "company",
484         "to_key_label_name": "country",
485         "out_edge_strategy": "single"
486     },
487     {
488         "from_key_label_name": "university",
489         "to_key_label_name": "city",
490         "in_edge_strategy": "none",
491         "out_edge_strategy": "single"
492     }
493 ]
494 },
495 "isPartOf": {
496     "properties": [],
497     "relations": [
498         {
499             "from_key_label_name": "city",
500             "to_key_label_name": "country",
501             "out_edge_strategy": "none"
502         },
503         {
504             "from_key_label_name": "country",
505             "to_key_label_name": "continent",
506             "in_edge_strategy": "none"
507         }
508     ]
509 }
510 }
511 }

```

Listing A.11: Content of mapping\_sf{10,30,100,300}.json describing the data mapping

```

1 {
2   "edges": {
3     "likes": [
4       {
5         "has_header": true,
6         "csv_row_fields_order": [
7           "from_key/person",
8           "to_key/comment",
9           "creationDate"
10        ],
11        "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/person_likes_comment_0_0.csv",

```

```

12         "type": "csv",
13         "delimiter": "|"
14     },
15     {
16         "has_header": true,
17         "csv_row_fields_order": [
18             "from_key/person",
19             "to_key/post",
20             "creationDate"
21         ],
22         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/person_likes_post_0_0.csv",
23         "type": "csv",
24         "delimiter": "|"
25     }
26 ],
27 "knows": [
28     {
29         "has_header": true,
30         "csv_row_fields_order": [
31             "from_key/person",
32             "to_key/person",
33             "creationDate"
34         ],
35         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/person_knows_person_0_0.csv",
36         "type": "csv",
37         "delimiter": "|"
38     }
39 ],
40 "hasType": [
41     {
42         "has_header": true,
43         "csv_row_fields_order": [
44             "from_key/tag",
45             "to_key/tagclass"
46         ],
47         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/tag_hasType_tagclass_0_0.csv",
48         "type": "csv",
49         "delimiter": "|"
50     }
51 ],
52 "hasMember": [
53     {
54         "has_header": true,
55         "csv_row_fields_order": [
56             "from_key/forum",
57             "to_key/person",
58             "joinDate"
59         ],
60         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/forum_hasMember_person_0_0.csv",
61         "type": "csv",
62         "delimiter": "|"
63     }
64 ],
65 "isLocatedIn": [
66     {
67         "has_header": true,
68         "csv_row_fields_order": [
69             "from_key/comment",

```

```

70         "to_key/country"
71     ],
72     "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/comment_isLocatedIn_country_0_0.csv",
73     ",
74     "type": "csv",
75     "delimiter": "|"
76 },
77 {
78     "has_header": true,
79     "csv_row_fields_order": [
80         "from_key/person",
81         "to_key/city"
82     ],
83     "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/person_isLocatedIn_city_0_0.csv",
84     "type": "csv",
85     "delimiter": "|"
86 },
87 {
88     "has_header": true,
89     "csv_row_fields_order": [
90         "from_key/post",
91         "to_key/country"
92     ],
93     "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/post_isLocatedIn_country_0_0.csv",
94     "type": "csv",
95     "delimiter": "|"
96 },
97 {
98     "has_header": true,
99     "csv_row_fields_order": [
100         "from_key/company",
101         "to_key/country"
102     ],
103     "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/company_isLocatedIn_country_0_0.csv",
104     "type": "csv",
105     "delimiter": "|"
106 },
107 {
108     "has_header": true,
109     "csv_row_fields_order": [
110         "from_key/university",
111         "to_key/city"
112     ],
113     "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/university_isLocatedIn_city_0_0.csv",
114     "type": "csv",
115     "delimiter": "|"
116 }
117 ],
118 "hasModerator": [
119     {
120         "has_header": true,
121         "csv_row_fields_order": [
122             "from_key/forum",
123             "to_key/person"
124         ],
125         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/forum_hasModerator_person_0_0.csv",
126         "type": "csv",
127         "delimiter": "|"

```

```

127     }
128   ],
129   "replyOf": [
130     {
131       "has_header": true,
132       "csv_row_fields_order": [
133         "from_key/comment",
134         "to_key/comment"
135       ],
136       "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/comment_replyOf_comment_0_0.csv",
137       "type": "csv",
138       "delimiter": "|"
139     },
140     {
141       "has_header": true,
142       "csv_row_fields_order": [
143         "from_key/comment",
144         "to_key/post"
145       ],
146       "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/comment_replyOf_post_0_0.csv",
147       "type": "csv",
148       "delimiter": "|"
149     }
150   ],
151   "hasCreator": [
152     {
153       "has_header": true,
154       "csv_row_fields_order": [
155         "from_key/comment",
156         "to_key/person",
157       "creationDate"
158     ],
159     "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/comment_hasCreator_person_0_0.csv",
160     "type": "csv",
161     "delimiter": "|"
162   },
163   {
164     "has_header": true,
165     "csv_row_fields_order": [
166       "from_key/post",
167       "to_key/person",
168     "creationDate"
169   ],
170   "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/post_hasCreator_person_0_0.csv",
171   "type": "csv",
172   "delimiter": "|"
173   }
174 ],
175 "isSubclassOf": [
176   {
177     "has_header": true,
178     "csv_row_fields_order": [
179       "from_key/tagclass",
180       "to_key/tagclass"
181     ],
182     "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/tagclass_isSubclassOf_tagclass_0_0.
csv",
183     "type": "csv",

```

```
184         "delimiter": "|"
185     }
186 ],
187 "workAt": [
188     {
189         "has_header": true,
190         "csv_row_fields_order": [
191             "from_key/person",
192             "to_key/company",
193             "workFrom"
194         ],
195         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/person_workAt_company_0_0.csv",
196         "type": "csv",
197         "delimiter": "|"
198     }
199 ],
200 "studyAt": [
201     {
202         "has_header": true,
203         "csv_row_fields_order": [
204             "from_key/person",
205             "to_key/university",
206             "classYear"
207         ],
208         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/person_studyAt_university_0_0.csv",
209         "type": "csv",
210         "delimiter": "|"
211     }
212 ],
213 "isPartOf": [
214     {
215         "has_header": true,
216         "csv_row_fields_order": [
217             "from_key/city",
218             "to_key/country"
219         ],
220         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/city_isPartOf_country_0_0.csv",
221         "type": "csv",
222         "delimiter": "|"
223     },
224     {
225         "has_header": true,
226         "csv_row_fields_order": [
227             "from_key/country",
228             "to_key/continent"
229         ],
230         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/country_isPartOf_continent_0_0.csv",
231         "type": "csv",
232         "delimiter": "|"
233     }
234 ],
235 "hasInterest": [
236     {
237         "has_header": true,
238         "csv_row_fields_order": [
239             "from_key/person",
240             "to_key/tag"
241         ],
```

```

242         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/person_hasInterest_tag_0_0.csv",
243         "type": "csv",
244         "delimiter": "|"
245     }
246 ],
247     "hasTag": [
248         {
249             "has_header": true,
250             "csv_row_fields_order": [
251                 "from_key/comment",
252                 "to_key/tag"
253             ],
254             "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/comment_hasTag_tag_0_0.csv",
255             "type": "csv",
256             "delimiter": "|"
257         },
258         {
259             "has_header": true,
260             "csv_row_fields_order": [
261                 "from_key/forum",
262                 "to_key/tag"
263             ],
264             "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/forum_hasTag_tag_0_0.csv",
265             "type": "csv",
266             "delimiter": "|"
267         },
268         {
269             "has_header": true,
270             "csv_row_fields_order": [
271                 "from_key/post",
272                 "to_key/tag"
273             ],
274             "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/post_hasTag_tag_0_0.csv",
275             "type": "csv",
276             "delimiter": "|"
277         }
278     ],
279     "containerOf": [
280         {
281             "has_header": true,
282             "csv_row_fields_order": [
283                 "from_key/forum",
284                 "to_key/post"
285             ],
286             "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/forum_containerOf_post_0_0.csv",
287             "type": "csv",
288             "delimiter": "|"
289         }
290     ]
291 },
292 "graph_name": "ldbc",
293 "vertices": {
294     "comment": [
295         {
296             "has_header": true,
297             "csv_row_fields_order": [
298                 "id",
299                 "creationDate",

```

```
300         "locationIP",
301         "browserUsed",
302         "content",
303         "length"
304     ],
305     "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/comment_0_0.csv",
306     "type": "csv",
307     "delimiter": "|"
308 }
309 ],
310 "city": [
311     {
312         "has_header": true,
313         "csv_row_fields_order": [
314             "id",
315             "name",
316             "url"
317         ],
318         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/city_0_0.csv",
319         "type": "csv",
320         "delimiter": "|"
321     }
322 ],
323 "university": [
324     {
325         "has_header": true,
326         "csv_row_fields_order": [
327             "id",
328             "name",
329             "url"
330         ],
331         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/university_0_0.csv",
332         "type": "csv",
333         "delimiter": "|"
334     }
335 ],
336 "forum": [
337     {
338         "has_header": true,
339         "csv_row_fields_order": [
340             "id",
341             "title",
342             "creationDate"
343         ],
344         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/forum_0_0.csv",
345         "type": "csv",
346         "delimiter": "|"
347     }
348 ],
349 "country": [
350     {
351         "has_header": true,
352         "csv_row_fields_order": [
353             "id",
354             "name",
355             "url"
356         ],
357         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/country_0_0.csv",
```





```
358         "type": "csv",
359         "delimiter": "|"
360     }
361 ],
362 "company": [
363     {
364         "has_header": true,
365         "csv_row_fields_order": [
366             "id",
367             "name",
368             "url"
369         ],
370         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/company_0_0.csv",
371         "type": "csv",
372         "delimiter": "|"
373     }
374 ],
375 "person": [
376     {
377         "has_header": true,
378         "csv_row_fields_order": [
379             "id",
380             "firstName",
381             "lastName",
382             "gender",
383             "birthday",
384             "creationDate",
385             "locationIP",
386             "browserUsed",
387             "language",
388             "email"
389         ],
390         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/person_0_0.csv",
391         "type": "csv",
392         "delimiter": "|"
393     }
394 ],
395 "tagclass": [
396     {
397         "has_header": true,
398         "csv_row_fields_order": [
399             "id",
400             "name",
401             "url"
402         ],
403         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/tagclass_0_0.csv",
404         "type": "csv",
405         "delimiter": "|"
406     }
407 ],
408 "tag": [
409     {
410         "has_header": true,
411         "csv_row_fields_order": [
412             "id",
413             "name",
414             "url"
415         ],
```



```

416         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/tag_0_0.csv",
417         "type": "csv",
418         "delimiter": "|"
419     }
420 ],
421     "post": [
422     {
423         "has_header": true,
424         "csv_row_fields_order": [
425             "id",
426             "imageFile",
427             "creationDate",
428             "locationIP",
429             "browserUsed",
430             "language",
431             "content",
432             "length"
433         ],
434         "csv_path": "/disk1/SF{10,30,100,300}/social_network/dynamic/post_0_0.csv",
435         "type": "csv",
436         "delimiter": "|"
437     }
438 ],
439     "continent": [
440     {
441         "has_header": true,
442         "csv_row_fields_order": [
443             "id",
444             "name",
445             "url"
446         ],
447         "csv_path": "/disk1/SF{10,30,100,300}/social_network/static/continent_0_0.csv",
448         "type": "csv",
449         "delimiter": "|"
450     }
451 ]
452 }
453 }

```

## A.7 Benchmark configuration

Listing A.12: Contents of `interactive-benchmark-sf30.properties` used for scale factor 30

```

1 host=172.31.7.248
2 port=21021
3
4 status=1
5 thread_count=64
6 name=LDBC-SNB
7 mode=execute_benchmark
8 results_log=true
9 time_unit=MICROSECONDS
10 time_compression_ratio=0.001
11 peer_identifiers=
12 workload_statistics=false
13 spinner_wait_duration=1

```



```

14 help=false
15 ignore_scheduled_start_times=false
16
17 workload=org.ldbcouncil.snb.driver.workloads.interactive.LdbcSnbInteractiveWorkload
18 db=io.atlasgraph.interactive.AtlasGraphInteractiveDb
19
20 ldbc.snb.interactive.parameters_dir=/disk1/SF30/substitution_parameters/
21 ldbc.snb.interactive.updates_dir=/disk1/SF30/social_network/
22 ldbc.snb.interactive.short_read_dissipation=0.2
23 ldbc.snb.interactive.scale_factor=30
24
25 operation_count=272000000
26 warmup=64000000
27
28 ldbc.snb.interactive.LdbcQuery1_enable=true
29 ldbc.snb.interactive.LdbcQuery2_enable=true
30 ldbc.snb.interactive.LdbcQuery3_enable=true
31 ldbc.snb.interactive.LdbcQuery4_enable=true
32 ldbc.snb.interactive.LdbcQuery5_enable=true
33 ldbc.snb.interactive.LdbcQuery6_enable=true
34 ldbc.snb.interactive.LdbcQuery7_enable=true
35 ldbc.snb.interactive.LdbcQuery8_enable=true
36 ldbc.snb.interactive.LdbcQuery9_enable=true
37 ldbc.snb.interactive.LdbcQuery10_enable=true
38 ldbc.snb.interactive.LdbcQuery11_enable=true
39 ldbc.snb.interactive.LdbcQuery12_enable=true
40 ldbc.snb.interactive.LdbcQuery13_enable=true
41 ldbc.snb.interactive.LdbcQuery14_enable=true
42
43 ldbc.snb.interactive.LdbcShortQuery1PersonProfile_enable=true
44 ldbc.snb.interactive.LdbcShortQuery2PersonPosts_enable=true
45 ldbc.snb.interactive.LdbcShortQuery3PersonFriends_enable=true
46 ldbc.snb.interactive.LdbcShortQuery4MessageContent_enable=true
47 ldbc.snb.interactive.LdbcShortQuery5MessageCreator_enable=true
48 ldbc.snb.interactive.LdbcShortQuery6MessageForum_enable=true
49 ldbc.snb.interactive.LdbcShortQuery7MessageReplies_enable=true
50
51 ldbc.snb.interactive.LdbcUpdate1AddPerson_enable=true
52 ldbc.snb.interactive.LdbcUpdate2AddPostLike_enable=true
53 ldbc.snb.interactive.LdbcUpdate3AddCommentLike_enable=true
54 ldbc.snb.interactive.LdbcUpdate4AddForum_enable=true
55 ldbc.snb.interactive.LdbcUpdate5AddForumMembership_enable=true
56 ldbc.snb.interactive.LdbcUpdate6AddPost_enable=true
57 ldbc.snb.interactive.LdbcUpdate7AddComment_enable=true
58 ldbc.snb.interactive.LdbcUpdate8AddFriendship_enable=true

```

Listing A.13: Contents of interactive-benchmark-sf100.properties used for scale factor 100

```

1 host=172.31.7.248
2 port=21021
3
4 status=1
5 thread_count=64
6 name=LDBC-SNB
7 mode=execute_benchmark
8 results_log=true
9 time_unit=MICROSECONDS
10 time_compression_ratio=0.0026

```



```

11 peer_identifiers=
12 workload_statistics=false
13 spinner_wait_duration=1
14 help=false
15 ignore_scheduled_start_times=false
16
17 workload=org.ldbcouncil.snb.driver.workloads.interactive.LdbcSnbInteractiveWorkload
18 db=io.atlasgraph.interactive.AtlasGraphInteractiveDb
19
20 ldbc.snb.interactive.parameters_dir=/disk1/SF100/substitution_parameters/
21 ldbc.snb.interactive.updates_dir=/disk1/SF100/social_network/
22 ldbc.snb.interactive.short_read_dissipation=0.2
23 ldbc.snb.interactive.scale_factor=100
24
25 operation_count=355800000
26 warmup=95000000
27
28 ldbc.snb.interactive.LdbcQuery1_enable=true
29 ldbc.snb.interactive.LdbcQuery2_enable=true
30 ldbc.snb.interactive.LdbcQuery3_enable=true
31 ldbc.snb.interactive.LdbcQuery4_enable=true
32 ldbc.snb.interactive.LdbcQuery5_enable=true
33 ldbc.snb.interactive.LdbcQuery6_enable=true
34 ldbc.snb.interactive.LdbcQuery7_enable=true
35 ldbc.snb.interactive.LdbcQuery8_enable=true
36 ldbc.snb.interactive.LdbcQuery9_enable=true
37 ldbc.snb.interactive.LdbcQuery10_enable=true
38 ldbc.snb.interactive.LdbcQuery11_enable=true
39 ldbc.snb.interactive.LdbcQuery12_enable=true
40 ldbc.snb.interactive.LdbcQuery13_enable=true
41 ldbc.snb.interactive.LdbcQuery14_enable=true
42
43 ldbc.snb.interactive.LdbcShortQuery1PersonProfile_enable=true
44 ldbc.snb.interactive.LdbcShortQuery2PersonPosts_enable=true
45 ldbc.snb.interactive.LdbcShortQuery3PersonFriends_enable=true
46 ldbc.snb.interactive.LdbcShortQuery4MessageContent_enable=true
47 ldbc.snb.interactive.LdbcShortQuery5MessageCreator_enable=true
48 ldbc.snb.interactive.LdbcShortQuery6MessageForum_enable=true
49 ldbc.snb.interactive.LdbcShortQuery7MessageReplies_enable=true
50
51 ldbc.snb.interactive.LdbcUpdate1AddPerson_enable=true
52 ldbc.snb.interactive.LdbcUpdate2AddPostLike_enable=true
53 ldbc.snb.interactive.LdbcUpdate3AddCommentLike_enable=true
54 ldbc.snb.interactive.LdbcUpdate4AddForum_enable=true
55 ldbc.snb.interactive.LdbcUpdate5AddForumMembership_enable=true
56 ldbc.snb.interactive.LdbcUpdate6AddPost_enable=true
57 ldbc.snb.interactive.LdbcUpdate7AddComment_enable=true
58 ldbc.snb.interactive.LdbcUpdate8AddFriendship_enable=true

```

Listing A.14: Contents of `interactive-benchmark-sf300.properties` used for scale factor 300

```

1 host=172.31.7.248
2 port=21021
3
4 status=1
5 thread_count=64
6 name=LDBC-SNB
7 mode=execute_benchmark

```



```
8 results_log=true
9 time_unit=MICROSECONDS
10 time_compression_ratio=0.0092
11 peer_identifiers=
12 workload_statistics=false
13 spinner_wait_duration=1
14 help=false
15 ignore_scheduled_start_times=false
16
17 workload=org.ldbcouncil.snb.driver.workloads.interactive.LdbcSnbInteractiveWorkload
18 db=io.atlasgraph.interactive.AtlasGraphInteractiveDb
19
20 ldbc.snb.interactive.parameters_dir=/disk1/SF300/substitution_parameters/
21 ldbc.snb.interactive.updates_dir=/disk1/SF300/social_network/
22 ldbc.snb.interactive.short_read_dissipation=0.2
23 ldbc.snb.interactive.scale_factor=300
24
25 operation_count=362300000
26 warmup=87300000
27
28 ldbc.snb.interactive.LdbcQuery1_enable=true
29 ldbc.snb.interactive.LdbcQuery2_enable=true
30 ldbc.snb.interactive.LdbcQuery3_enable=true
31 ldbc.snb.interactive.LdbcQuery4_enable=true
32 ldbc.snb.interactive.LdbcQuery5_enable=true
33 ldbc.snb.interactive.LdbcQuery6_enable=true
34 ldbc.snb.interactive.LdbcQuery7_enable=true
35 ldbc.snb.interactive.LdbcQuery8_enable=true
36 ldbc.snb.interactive.LdbcQuery9_enable=true
37 ldbc.snb.interactive.LdbcQuery10_enable=true
38 ldbc.snb.interactive.LdbcQuery11_enable=true
39 ldbc.snb.interactive.LdbcQuery12_enable=true
40 ldbc.snb.interactive.LdbcQuery13_enable=true
41 ldbc.snb.interactive.LdbcQuery14_enable=true
42
43 ldbc.snb.interactive.LdbcShortQuery1PersonProfile_enable=true
44 ldbc.snb.interactive.LdbcShortQuery2PersonPosts_enable=true
45 ldbc.snb.interactive.LdbcShortQuery3PersonFriends_enable=true
46 ldbc.snb.interactive.LdbcShortQuery4MessageContent_enable=true
47 ldbc.snb.interactive.LdbcShortQuery5MessageCreator_enable=true
48 ldbc.snb.interactive.LdbcShortQuery6MessageForum_enable=true
49 ldbc.snb.interactive.LdbcShortQuery7MessageReplies_enable=true
50
51 ldbc.snb.interactive.LdbcUpdate1AddPerson_enable=true
52 ldbc.snb.interactive.LdbcUpdate2AddPostLike_enable=true
53 ldbc.snb.interactive.LdbcUpdate3AddCommentLike_enable=true
54 ldbc.snb.interactive.LdbcUpdate4AddForum_enable=true
55 ldbc.snb.interactive.LdbcUpdate5AddForumMembership_enable=true
56 ldbc.snb.interactive.LdbcUpdate6AddPost_enable=true
57 ldbc.snb.interactive.LdbcUpdate7AddComment_enable=true
58 ldbc.snb.interactive.LdbcUpdate8AddFriendship_enable=true
```

## A.8 Validation configuration

Listing A.15: The contents of `validate.properties`



```
1 host=172.31.7.248
2 port=21021
3
4 status=1
5 thread_count=1
6 name=LDBC-SNB
7 mode=validate_database
8 results_log=true
9 time_unit=MICROSECONDS
10 time_compression_ratio=0.001
11 peer_identifiers=
12 workload_statistics=false
13 spinner_wait_duration=1
14 help=false
15 ignore_scheduled_start_times=true
16
17 workload=org.ldbcouncil.snb.driver.workloads.interactive.LdbcSnbInteractiveWorkload
18 db=io.atlasgraph.interactive.AtlasGraphInteractiveDb
19 operation_count=10000
20
21 validate_workload=true
22 validate_database=validation_params-sf10.csv
23 ldbc.snb.interactive.parameters_dir=/disk1/SF10/substitution_parameters/
24 ldbc.snb.interactive.short_read_dissipation=0.2
25 ldbc.snb.interactive.update_interleave=466
26 ldbc.snb.interactive.scale_factor=10
27
28 ## frequency of read queries (number of update queries per one read query)
29 ldbc.snb.interactive.LdbcQuery1_freq=1
30 ldbc.snb.interactive.LdbcQuery2_freq=1
31 ldbc.snb.interactive.LdbcQuery3_freq=1
32 ldbc.snb.interactive.LdbcQuery4_freq=1
33 ldbc.snb.interactive.LdbcQuery5_freq=1
34 ldbc.snb.interactive.LdbcQuery6_freq=1
35 ldbc.snb.interactive.LdbcQuery7_freq=1
36 ldbc.snb.interactive.LdbcQuery8_freq=1
37 ldbc.snb.interactive.LdbcQuery9_freq=1
38 ldbc.snb.interactive.LdbcQuery10_freq=1
39 ldbc.snb.interactive.LdbcQuery11_freq=1
40 ldbc.snb.interactive.LdbcQuery12_freq=1
41 ldbc.snb.interactive.LdbcQuery13_freq=1
42 ldbc.snb.interactive.LdbcQuery14_freq=1
43
44 # *** For debugging purposes ***
45
46 ldbc.snb.interactive.LdbcQuery1_enable=true
47 ldbc.snb.interactive.LdbcQuery2_enable=true
48 ldbc.snb.interactive.LdbcQuery3_enable=true
49 ldbc.snb.interactive.LdbcQuery4_enable=true
50 ldbc.snb.interactive.LdbcQuery5_enable=true
51 ldbc.snb.interactive.LdbcQuery6_enable=true
52 ldbc.snb.interactive.LdbcQuery7_enable=true
53 ldbc.snb.interactive.LdbcQuery8_enable=true
54 ldbc.snb.interactive.LdbcQuery9_enable=true
55 ldbc.snb.interactive.LdbcQuery10_enable=true
56 ldbc.snb.interactive.LdbcQuery11_enable=true
57 ldbc.snb.interactive.LdbcQuery12_enable=true
58 ldbc.snb.interactive.LdbcQuery13_enable=true
```

```
59 ldbc.snb.interactive.LdbcQuery14_enable=true
60
61 ldbc.snb.interactive.LdbcShortQuery1PersonProfile_enable=true
62 ldbc.snb.interactive.LdbcShortQuery2PersonPosts_enable=true
63 ldbc.snb.interactive.LdbcShortQuery3PersonFriends_enable=true
64 ldbc.snb.interactive.LdbcShortQuery4MessageContent_enable=true
65 ldbc.snb.interactive.LdbcShortQuery5MessageCreator_enable=true
66 ldbc.snb.interactive.LdbcShortQuery6MessageForum_enable=true
67 ldbc.snb.interactive.LdbcShortQuery7MessageReplies_enable=true
68
69 ldbc.snb.interactive.LdbcUpdate1AddPerson_enable=true
70 ldbc.snb.interactive.LdbcUpdate2AddPostLike_enable=true
71 ldbc.snb.interactive.LdbcUpdate3AddCommentLike_enable=true
72 ldbc.snb.interactive.LdbcUpdate4AddForum_enable=true
73 ldbc.snb.interactive.LdbcUpdate5AddForumMembership_enable=true
74 ldbc.snb.interactive.LdbcUpdate6AddPost_enable=true
75 ldbc.snb.interactive.LdbcUpdate7AddComment_enable=true
76 ldbc.snb.interactive.LdbcUpdate8AddFriendship_enable=true
```