

An analysis of the SIGMOD 2014 Programming Contest: Complex queries on the LDBC social network graph

Márton Elekes* János Benjamin Antal Gábor Szárnyas†

This report contains an analysis of the queries defined in the SIGMOD 2014 Programming Contest. We first describe the data set, then present the queries, providing graphical illustrations for them and pointing out their caveats. Our intention is to document our lessons learnt and simplify the work of those who will attempt to create a solution to this contest. We also demonstrate the influence of this contest by listing followup works which used these queries as inspiration to design better algorithms or to define interesting graph queries.

1. Introduction

The SIGMOD conference hosts an annual Programming Contest where teams of graduate students are required to solve database-related programming tasks. Teams have approximately 3 months to implement their solution and compete on a previously agreed metric (e.g. lowest execution time). The Programming Contest of SIGMOD 2014¹ focused on graph processing problems and consisted of 4 complex graph queries on the LDBC Social Network Benchmark's schema [2]. The queries contained a mix of relational operators (such as filtering and aggregation) along with graph analytical computations (such as breadth-first search and connected components). This document contains an overview and analysis of the data set and queries of the 2014 Programming Contest.

2. Graph schema and data sets

2.1. Graph schema

The social network graph instances used in the contest can be represented as a *property graph* [1] that conforms to the schema of the LDBC Social Network Benchmark. The

*Budapest University of Technology and Economics, Department of Measurement and Information Systems, Hungary

†CWI Amsterdam, The Netherlands. Corresponding author. Email: szarnyasg@gmail.com

¹<http://www.cs.albany.edu/~sigmod14contest/>

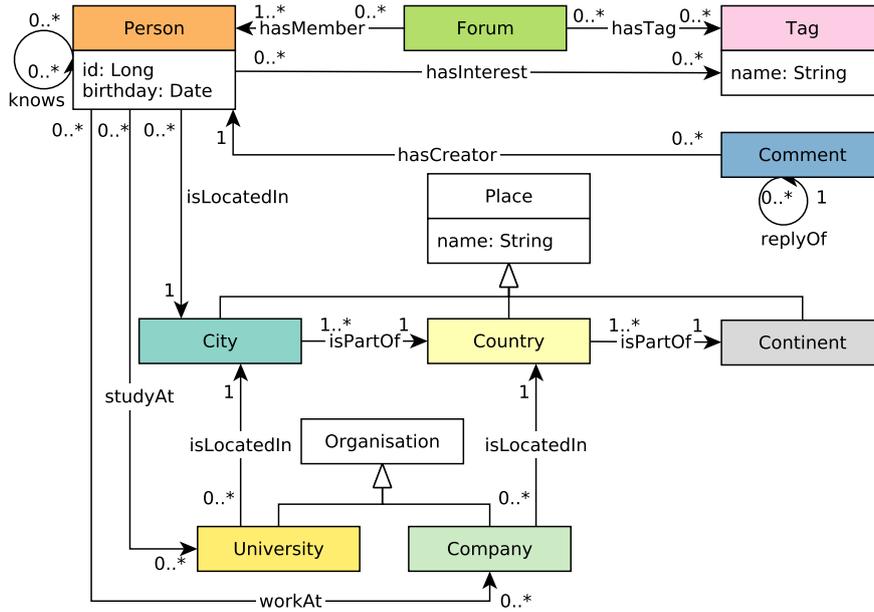


Figure 1.: Schema of the social network graphs. Only the relevant node/edge types and property names (along with their types) shown.

relevant part of the schema is shown in Figure 1 (the full schema can be found in [2]). The edges in the graph are directed with the exception of the **knows** edges which are treated as undirected.

2.2. Data sets

The data sets containing the graph instances are produced by the LDBC Datagen [20], a Hadoop-based graph generator. Datagen generates a realistic Facebook-like power-law degree distribution for the **Person**–**knows**–**Person** graph and also enforces certain correlations, e.g. people who studied together are more likely to become friends.

Small data sets with 1k and 10k **Persons** are available on the contest’s website:

- **1k**: <http://www.cs.albany.edu/~sigmod14contest/files/outputDir-1k.zip>
- **10k**: <http://www.cs.albany.edu/~sigmod14contest/files/outputDir-10k.zip>

To generate larger data sets, use the LDBC Datagen tagged as **sigmod2014contest**² with the configurations provided in Appendix A.

Caveat \triangle In the generated data, **Places** at different hierarchy levels, i.e. **Continents**, **Countries**, and **Cities** can have the same name. For example, “Australia” is both a **Continent**

²https://github.com/ldbc/ldbc_snb_datagen/releases/tag/sigmod2014contest

and a Country, while “Indonesia” is both a Country and a City.³

```

$ grep Australia place.csv
1463|Australia|\protect\vrule width0pt\protect\href{http://dbpedia.org/resource/Australia|Continent
    }{http://dbpedia.org/resource/Australia|Continent}
4|Australia|\protect\vrule width0pt\protect\href{http://dbpedia.org/resource/Australia|Country}{
    http://dbpedia.org/resource/Australia|Country}

$ grep Indonesia place.csv
39|Indonesia|\protect\vrule width0pt\protect\href{http://dbpedia.org/resource/Indonesia|Country}{
    http://dbpedia.org/resource/Indonesia|Country}
661|Indonesia|\protect\vrule width0pt\protect\href{http://dbpedia.org/resource/Indonesia|City}{http
    ://dbpedia.org/resource/Indonesia|City}

```

3. Queries

The contest defines 4 queries. In the contest, solutions were expected to execute a mix of these queries (see Section 4.4). In the following, we present the queries including their textual specification, their input parameters, output attributes and the illustration of their graph pattern.

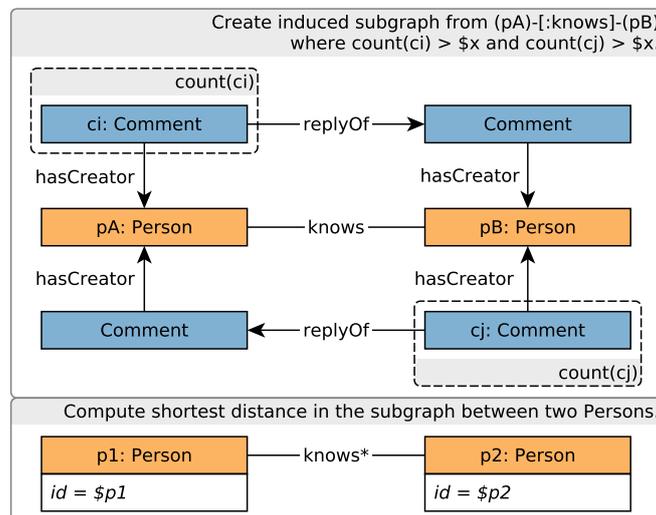


Figure 2.: Graph pattern of Q1.

Q1. Shortest Distance over Frequent Communication Paths (Figure 2). Given two integer Person identifiers $p1$ and $p2$, and another integer x , find the minimum number of hops (i.e. the shortest distance) between $p1$ and $p2$ in the graph induced by Persons who

³The Datagen version used in the contest generates multiple erroneous City names such as “India” and “Indonesia”. These have been removed in later versions of Datagen. However, even in the latest version of the data generator, Place names are not completely disjoint at different hierarchy levels as “Australia” is (correctly) both generated as a Country and a Continent.

- know each other and
- communicate frequently with each other, i.e. both have made more than x Comments in reply to the other one's Comments.

Caveat \triangle Some remarks regarding Q1:

- The frequent communication has to happen *both ways* between Person pairs.
- When determining the shortest distance, only the edges of the induced subgraph can be used.
- Notice that for $x = -1$, the original Person–knows–Person is equivalent to the induced subgraph, so it is not necessary to compute the number of interactions.

API query1(p1, p2, x)

Output One integer (the hop count) per line.

Samples 1k-sample-queries1.txt and 1k-sample-answers1.txt (Appendix B)

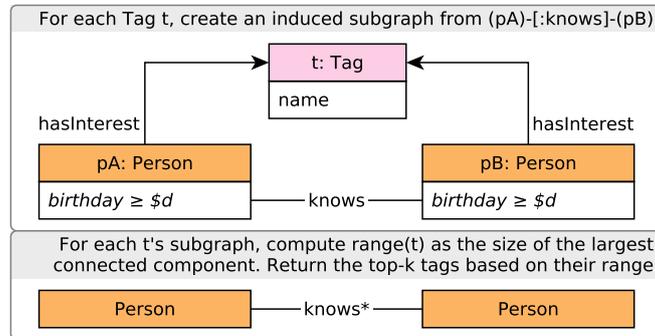


Figure 3.: Graph pattern of Q2.

Q2. Interests with Large Communities (Figure 3). Given an integer k and a birthday d , find the top- k Tags. A Tag is characterized with its *range*, i.e. the size of the largest connected component in the graph induced by Persons who

- know each other,
- are interested in the Tag, and
- were born on day d or later.

Caveat \triangle When determining the connected components, only the **knows** edges in the induced subgraph should be used.

API query2(k, d)

Output Exactly k strings (separated by a space) per line. These k strings represent Tag names of interest, ordered by range from largest to smallest, with ties broken by lexicographical ordering, ascending (e.g. “A” precedes “B” in the results).

Samples 1k-sample-queries2.txt and 1k-sample-answers2.txt (Appendix B)

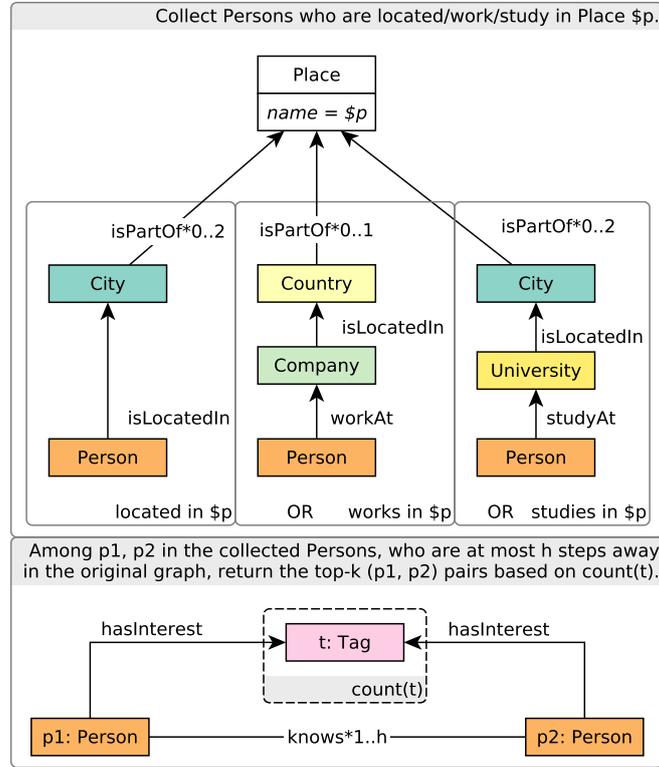


Figure 4.: Graph pattern of Q3.

Q3. Socialization Suggestion (Figure 4). Given an integer k , an integer maximum hop count h , and a Place name p , find the top- k similar pairs of Persons based on the number of common interest Tags. For each of the k pairs mentioned above, the two Persons must be located in p or study/work at Organisations in p . Furthermore, these two Persons must be no more than h hops away from each other in the original Person-knows-Person graph (i.e. their shortest distance is at most h).

Caveat \triangle When determining the Person pairs with at most h steps away, only the Person nodes in the selected Place are considered *for selecting the pairs* but all Person nodes and knows edges should be considered *when determining the shortest distance of a given pair*.

API query3(k , h , p)

Output Exactly k pairs of Person ids per line. These pairs are separated by a space and Person ids are separated by the pipe character “|”. For any Person id p , $p|p$ must be excluded. For any pairs $p_1|p_2$ and $p_2|p_1$, the second pair in lexicographical order must be excluded. These k pairs must be ordered by similarity from highest to lowest, with ties broken by lexicographical ordering based on ascending numerical order of the ids.⁴

Samples `1k-sample-queries3.txt` and `1k-sample-answers3.txt` (Appendix B)

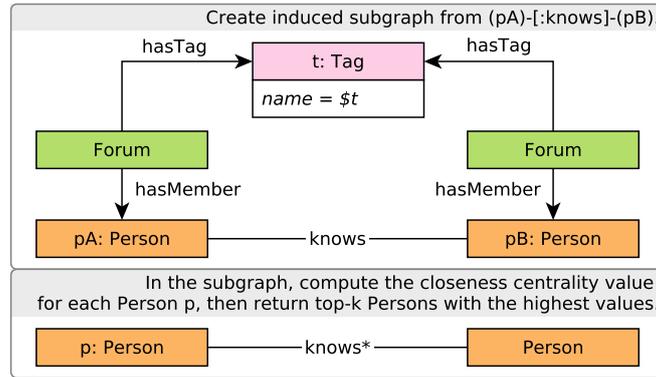


Figure 5.: Graph pattern of Q4.

Q4. Most Central People (Figure 5). Given an integer k and a Tag name t , find the top- k Persons based on the *closeness centrality value* (CCV) in the graph induced by Persons who

- know each other and
- are members of Forums which have Tag t .

For each Person p in the induced subgraph, compute

$$CCV(p) = \frac{(|C(p)| - 1)^2}{(n - 1) \cdot s(p)},$$

where

- $C(p)$ denotes the nodes in the connected component of node p ,
- n is the number of nodes in the induced subgraph, and
- $s(p)$ is the sum of shortest distances to all other reachable Persons from p , i.e.

$$s(p) = \sum_{r \in C(p)} \text{dist}(p, r)$$

For isolated nodes, the divisor in the CCV formula is 0 due to $s(p) = 0$. For these, the centrality value is defined as 0.

⁴Numerical ordering means the regular arithmetic ordering, i.e. $9 < 10$. (A common issue is comparing numbers as strings, which would give the opposite order. This should be avoided here.)

Caveat \triangle When determining the *CCV* value, only the **knows** edges in the induced subgraph should be used.

API query4(*k*, *t*)

Output Exactly *k* Person ids (separated by a space) per line. These Person ids are ordered by centrality from highest to lowest, with ties broken by Person id, in ascending order (i.e. id 1 precedes id 2 in the results).

Samples 1k-sample-queries4.txt and 1k-sample-answers4.txt (Appendix B)

4. Analysis

4.1. Data set

Distribution The power-law distribution of the Person–knows–Person graph implies that this graph has a small diameter. This means that techniques such as *direction-optimizing traversal* [3] (also known as *push/pull* and *top-down/bottom-up*) can be used to improve the performance of graph traversals.

Node id relabelling The nodes in the generated input graphs have *sparse ids*, i.e. their identifiers can take any 64-bit unsigned integer value. These identifiers cannot be used directly by systems relying on compressed matrix representations such as CSR (Compressed Sparse Row) [21]. This necessitates the use of *dense ids* with consecutive values in the $[0, |V|)$ range. To provide such identifiers, systems need to perform *node relabelling*, also known as *dense vertex relabelling* [31], *vertex permutation* [4], and *mapping from sparse to dense keys* [18]. In the context of this contest, implementers have to consider the tradeoffs of this approach. On the one hand, performing such a relabelling is an expensive step. On the other hand, it can improve the locality of neighbourhood lookup operations (compared to e.g. hash-based lookups) which form a significant portion of the operations in the queries.

4.2. Queries

We present a brief discussion on the complexity of implementing the queries of the contest efficiently. The key algorithms (kernels) required by each query are listed in Table 1.

Q1. Shortest Distance over Frequent Communication Paths

This query computes the (unweighted) shortest distance between two fixed nodes in an induced subgraph. The second part of the query, which determines the shortest distance, is rather straightforward: it can be implemented using a bidirectional search on the subgraph.

The difficulty of this query lies in amalgamating its first part, which computes the induced subgraph, with its second part. For the majority of input parameters

Algorithm	Variant	Q1	Q2	Q3	Q4
Unweighted shortest distance	Single-source, single-target	⊗	○	○	○
	Multi-source	○	○	⊗	○
	All-pairs	○	○	○	⊗
Weakly connected components	–	○	⊗	○	⊘

Table 1.: Algorithms to be used for efficient evaluation of the queries. *Notation* – ⊗: required, ○: not required, ⊘: applicable but not required.

$x \geq 0$, i.e. the induced subgraph is not equivalent to the full **Person**–**knows**–**Person** graph. In these cases, there are two key approaches for implementing the subgraph computation:

Approach (1): Precompute the induced subgraph and run the search on that graph.

Approach (2): For each edge, check the number of its interactions on-the-fly during the search.

Approach (1) is simpler to implement but often leads to suboptimal performance. This is especially the case if the path exists and is found quickly without having to traverse the entire subgraph, leaving most of the induced subgraph unused.

Therefore, approach (2) can often avoid a lot of unnecessary computation. However, implementing this approach efficiently is complicated by the setup of the original contest, where solutions were expected to evaluate multiple queries at the same time. In particular, when evaluating multiple queries, computing the number of interactions for each edge on-the-fly might result in redundant computations as different traversals will likely reach the same edges. To prevent redundant computations, implementations should use some (thread-safe) caching mechanism to store the number of interactions for each edge.

Q2. Interests with Large Communities

Unlike the other queries which use a single induced subgraph, this query defines *multiple subgraphs*, one for each **Tag**. There are 1457 **Tags** in the 1k **Person** data set and more than 10 000 **Tags** in the 1M **Person** data set. Due to this and due to the fact that the computations are completely disjoint, this query is easy to parallelize by processing different **Tags** on different threads.

To allow efficient evaluation of the selection on the birthday attributes (on the inequality condition `birthday ≥ d`), it is recommended to define an index on the `Person.birthday` attribute.

Once an induced subgraph has been created, solutions need to run a connected components algorithm to determine the size of its largest connected component. There are three key approaches to compute these:

- A naïve solution is to run repeated BFS traversals, each discovering a connected component and then restarting the BFS from a yet unvisited node, until all nodes are visited. While this approach often results in small BFS operations which often discover components of just a few nodes, it can provide a quick solution with acceptable performance in the context of this query.
- Solution authors often opt to use *Tarjan’s strongly connected components algorithm* [28]. While a sequential variant of this algorithm is simple to implement, it has a complexity of $\mathcal{O}(|V| + |E|)$ and it is difficult to parallelize.
- Authors can exploit that this query only needs *weakly connected components*. This can be computed with algorithms such as the *Shiloach-Vishkin algorithm* [23] with a complexity of $\mathcal{O}(|E| \cdot \log |V|)$, and the recent *Afforest algorithm* [26], which has a worst-case complexity of $\mathcal{O}(|V| + |E|)$ but performs well for most practical data sets and is easy to parallelize.

Q3. Socialization Suggestion

The first part of this query requires efficient lookups for Persons in a given Place. This means that they are located there and/or study at a University there and/or work at a Company there. These lookups can be assisted by representing the Places using *nested intervals* [34].

The task in the second part of this query is to determine which pairs of Persons are reachable from each other using at most h hops. This can be implemented as a *multi-source BFS* starting from each Person, then advancing the frontier of the traversal for h steps and maintaining the visited nodes for each source. Once we completed the traversal, for each source Person, the other Persons found in its visited nodes are the ones reachable within h hops.

A more efficient approach can exploit the symmetry of the undirected knows edges by running a *multi-source bidirectional BFS*, which only advances the frontier for $\lceil \frac{h}{2} \rceil$ steps, then looks for intersections between the sets of visited nodes for each pair of traversals. (Odd h values need to be treated as a special case by intersecting the set of visited nodes produced after $\lceil \frac{h}{2} \rceil - 1$ steps and $\lceil \frac{h}{2} \rceil$ steps.)

Q4. Most Central People

The first part of this query, i.e. finding the Persons who are members of a Forum which has a given Tag is relatively simple. The difficulty lies in the second part, particularly in computing the $s(p)$ value for all Person nodes. This can be formulated as an *all-pairs shortest distances* (also known as *all-pairs unweighted shortest paths*) problem. While this problem is simpler than the *all-pairs weighted shortest paths* problem (which can be tackled by e.g. the Floyd-Warshall algorithm in $\mathcal{O}(n^3)$ time), it is still computationally intensive.

For a node p , its $s(p)$ value can be computed simply by running a BFS traversal: at each step l , we increment $s(p)$ by $l \times$ the number of nodes found at level l . The challenge in this query is the sheer number of traversals to execute: on the 1M

Person data set, for a popular **Tag** with many **Forums**, there might be more than 100 000 member **Persons** whose closeness centrality values need to be computed. Therefore, to be competitive on this query, solutions need to employ some sort of optimization such as compression and efficient search space pruning.

In the following, we list a few potential optimization ideas. Some of these are mutually incompatible (e.g. it is not possible to use integer matrices and bitwise compression) but others can be combined to get an efficient implementation. We also advise the reader to consult the presentations and posters of the top-ranking teams for more ideas.

Multi-source BFS The large number of BFS traversals, all of which perform the same operation during traversal (incrementing the $s(p)$ value until a fixed-point) lend themselves to bulk processing. Multi-source BFS operations can process multiple nodes in bulk and improve locality during the computation [32, 11].

Compression Top-ranking solutions in the contest employed bitwise operations to optimize the performance of multi-source BFS traversals [32, 16]. Using bitmaps encoded as **UINT64** values has multiple advantages: it saves memory and allows 64 traversals to be handled at once. Additionally, multiple **UINT64** values can be batch processed using SIMD operations.

Integer matrix multiplication The problem can be elegantly expressed as matrix multiplications on integer matrices as demonstrated in the algorithm of [22], which requires $\mathcal{O}(\log \text{diameter}(G))$ dense matrix multiplications. Using non-naïve matrix multiplication algorithms, this can result in a lower complexity than that of the multi-source BFS-based algorithms which require $\mathcal{O}(\text{diameter}(G))$ traversal steps. This approach can be optimized further: (1) In a distributed setup, reducing communication costs between parts of the matrix multiplication can result in a significant improvement of performance [24]. (2) Using Tiskin’s algorithm [33, 25], one can formulate the problem using multiplications on a dense and a sparse matrix, making it possible to exploit the sparsity of the adjacency matrix.

Splitting the graph to connected components The definition of $s(p)$ states that for each node, only the other reachable **Person** nodes are considered. Therefore, solutions might opt to first run a (weakly) connected components algorithm on the subgraph, then compute the CCV values on the components separately. This allows establishing tighter bounds for search space pruning (see below). The limitation of this approach is that computing the connected components is fairly expensive and for popular **Tags**, the graph often consists of only a few components so the benefits of this computation are limited.

Heuristics-based search space pruning When determining the $s(p)$ values, many Person nodes p can be discarded early if we can prove that their final $s(p)$ values will be larger than the top- k $s(p)$ values in the same component (implying that their CCV value will be lower). We can exploit that the computation of $s(p)$ for nodes with a high CCV value finishes early as (by definition) these belong to a central node from which the other reachable nodes are just a few hops away. Therefore, we can use get their final $s(p)$ values early and use them to prune traversals whose $s(p)$ values will be excessively high.

Team *blxlrsmb* suggested a simple lower bound for $s(p)$ values. According to this, after step l , the following lower bound holds:

$$s(p) \geq \sum_{r \in C(p) \wedge (\text{dist}(p,r) \leq l)} \text{dist}(p,r) + \sum_{r \in C(p) \wedge (\text{dist}(p,r) > l)} (l+1),$$

where the first term is equivalent to the value of $s(p)$ after l steps and the second term can be computed with a simple multiplication:

$$|\{r \in C(p) \text{ and } r \text{ is unvisited}\}| \cdot (l+1)$$

Other algorithms Paper [6] presents a family of sophisticated algorithms for the *unweighted undirected all-pairs shortest paths problem* guaranteeing a complexity of $\mathcal{O}(|V| \cdot |E|)$.

Top-k closeness centrality Paper [19] (co-authored by Jeong-Hyon Hwang, one of the contest’s organizers) presents a top- k closeness centrality value algorithm.

4.3. Choke points

The LDBC Social Network Benchmark [17] uses *choke points* [5, 10] to characterize the difficulty of its queries. These are challenging aspects of query optimization and evaluation, which have a significant impact on a system’s performance when processing the queries. Choke points are divided into categories such as aggregation performance, join performance, data access locality, etc.

We have analyzed the queries of the contest in terms of the choke points provided in the latest SNB specification, version 0.4.0-SNAPSHOT [17]. The choke points covered by the queries are shown in Table 2. The analysis confirms that the queries in the contest are complex: on average, a query has $35/4 = 8.75$ relevant choke points, making them quantifiably more complex than the queries of the Interactive workload, which have an average of 4.36 choke points/query, and the queries of the BI workload, which have 6.45.

4.4. Contest setup

The contest defined a measurement scenario where the solutions-under-benchmark were given a set of queries and their input parameters. Solutions competed on the total

Choke point	Q1	Q2	Q3	Q4
CP-1.3 Top-k pushdown	○	⊗	⊗	⊗
CP-2.3 Join type selection	○	○	⊗	○
CP-2.4 Sparse foreign key joins	○	⊗	⊗	⊗
CP-3.2 Dimensional clustering	○	⊗	○	○
CP-3.3 Scattered index access patterns	⊗	⊗	⊗	⊗
CP-5.3 Intra-query result reuse	⊗	○	⊗	○
CP-7.2 Cardinality estimation of transitive paths	⊗	○	⊗	⊗
CP-7.3 Execution of a transitive step	⊗	○	⊗	⊗
CP-7.5 Unweighted shortest paths	⊗	○	⊗	⊗
CP-7.7 Composition of graph queries	⊗	⊗	⊗	⊗
CP-8.3 Ranking-style queries	○	⊗	○	○
CP-8.4 Query composition	⊗	⊗	⊗	⊗
CP-8.6 Handling paths	⊗	○	⊗	⊗

Table 2.: Relevance of the SNB choke points for the queries in the contest. Choke points not covered by the queries were omitted. *Notation* – ⊗: relevant, ○: not relevant.

execution time, including the time of loading the data and evaluating the queries. Deciding which preprocessing computations to perform (if any), the order of queries to evaluate, and how to parallelize the evaluation was left for the solution to decide.

The contest had an uneven mix of queries, with more than 95% of the queries selected from Q1. The number of query instances per type for each data set is shown in Figure 6.

4.5. Parameter selection

Due to the highly skewed distribution of the `Person-knows-Person` graph and the correlated nature of the data, using uniform random sampling will result in unpredictable query performance as demonstrated in the results of [11]. If this is undesirable, we recommend using the *parameter curation* technique defined in [13] and selecting representative parameters categories for each query.⁵ For example, for Q1, one might find that the representative categories are the Cartesian product of the following aspects:

- **Subgraph computation:** no induced subgraph, induced subgraph with a low x value, induced subgraph with a high x value.
- **Reachability:** Person pairs that are unreachable from each other, reachable through few hops, reachable through many hops.

⁵The work on parameter curation was done as part of the LDBC project and it was published after the SIGMOD 2014 Contest. Therefore, parameter curation was presumably not used in the contest.

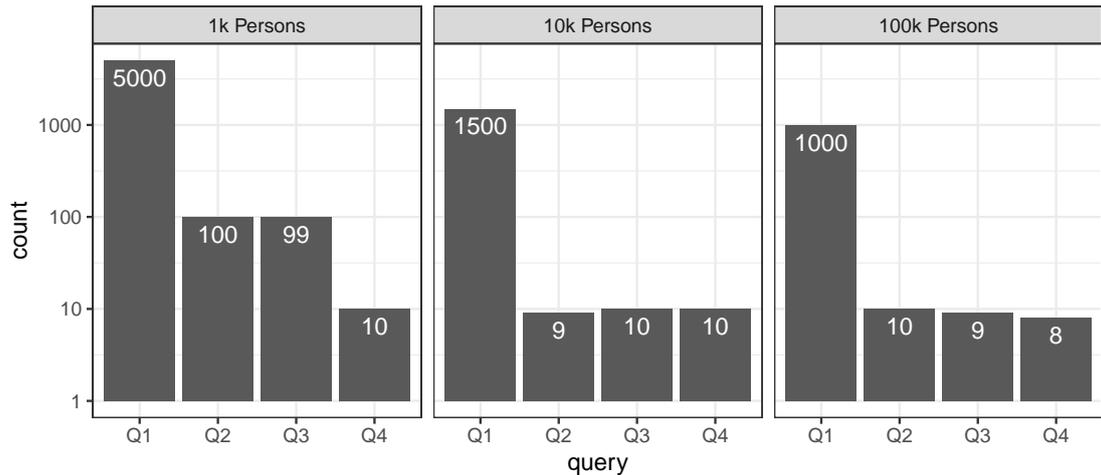


Figure 6.: Number of queries used in the original contest for the 1k, 10k, and 100k Person data sets. Note that the y axis is logarithmic.

5. Influence

Contest Overall, 33 teams have participated in the SIGMOD 2014 Programming Contest. The 5 top-ranking teams created posters and presentations detailing their solutions. These were presented at SIGMOD 2014 in Snowbird, Utah (USA), where the awards ceremony also took place.

Followup works In the last 6 years, the contest has influenced a number of works in the graph processing space. We provide a (non-exhaustive) list of papers and events that use this contest as their case study:

- Team *GenericPeople* from the Saint-Petersburg University (Russia) reported on their experience and discussed potential optimization ideas in paper [7] and technical report [8].
- The creators of two top-ranking solutions, *AWFY* (TUM, Germany) and *VIDA* (NYU, USA) co-authored a paper on efficient multi-source BFS [32]. A continuation of this work from researchers at TUM is paper [16] which presents a multi-threaded variant of the multi-source BFS algorithm.
- Members of the *AWFY* team have co-authored papers on how to evaluate graph queries efficiently [15] and on how to compute centrality values [29, 30].
- In a loosely related line of work, members of the *AWFY* team have co-authored a paper discussing the challenges of *graph loading*, including node id relabelling, and compared potential approaches [31].

- Queries similar to Q1 and Q3, i.e. ones that compute shortest distances between Persons pairs or Person sets, are included in the LDBC Social Network Benchmark’s *Interactive* [12, 2] and *Business Intelligence* workloads [27, 17].
- The Grand Challenge of the DEBS 2016 (Distributed Event-Based Systems) conference defined a problem based on the data sets of the SIGMOD 2014 Programming Contest.⁶ Its description states the following:

The data for the DEBS 2016 Grand Challenge is based on the dataset provided together with the LDBC Social Network Benchmark. DEBS 2016 Grand Challenge takes up the general scenario from the 2014 SIGMOD Programming Contest, however, in contrast to the SIGMOD contest, it explicitly focuses on processing streaming data and thus dynamic graphs.

- The 2018 Transformation Tool Contest, an annual contest held at the STAF (Software Technologies: Applications and Foundations) federation of conferences, presented the *Social Media* benchmark [14] as its live case study which participants were required to solve during the conference. Similarly to the DEBS 2016 Grand Challenge, this case study defined two queries over a simplified social network schema which are evaluated on a continuously changing data set (thus favouring solutions that employ *incremental view maintenance* techniques).
- In 2020, we have created a solution using the linear algebra-based concepts of the GraphBLAS C API and implemented in the SuiteSparse:GraphBLAS parallel library [9]. Our solution was published in paper [11] and is available online.⁷

6. Online resources

The source code, presentations, and posters of the top-5 teams are available on the contest’s original website.⁸ To provide strong baseline implementations for evaluating our GraphBLAS solution presented in [11], we have updated the solutions of two teams, *AWFY* and *blxlrsmb*, and made them available online.⁹ The updated solutions can be compiled with GCC version 9 and we made it possible to run each query implementation individually (compared to the contest’s setup where multiple queries were expected to run in parallel).

Acknowledgements

We would like to thank Jeong-Hyon Hwang for providing us the parameter files. We would also like to thank Attila Nagy and Dávid Sándor for thorough discussions on the

⁶<https://web.archive.org/web/20190131182518/https://debs.org/debs-2016-grand-challenge-social-networks/>

⁷<https://github.com/ldbc/sigmod2014-contest-graphblas>

⁸<http://www.cs.albany.edu/~sigmod14contest/leaders.html>

⁹<https://github.com/ftsrg/sigmod2014-pc-top-solutions>

graph queries of the contest.

References

- [1] Renzo Angles et al. “Foundations of Modern Query Languages for Graph Databases”. In: *ACM Comput. Surv.* 50.5 (2017), 68:1–68:40. DOI: [10.1145/3104031](https://doi.org/10.1145/3104031).
- [2] Renzo Angles et al. “The LDBC Social Network Benchmark”. In: *CoRR* abs/2001.02299 (2020). URL: <http://arxiv.org/abs/2001.02299>.
- [3] Scott Beamer, Krste Asanovic, and David A. Patterson. “Direction-optimizing breadth-first search”. In: *SC*. IEEE/ACM, 2012. DOI: [10.1109/SC.2012.50](https://doi.org/10.1109/SC.2012.50).
- [4] Maciej Besta and Torsten Hoefler. “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations”. In: *CoRR* abs/1806.01799 (2018). URL: <http://arxiv.org/abs/1806.01799>.
- [5] Peter A. Boncz, Thomas Neumann, and Orri Erling. “TPC-H Analyzed: Hidden Messages and Lessons Learned from an Influential Benchmark”. In: *TPCTC*. 2013, pp. 61–76. DOI: [10.1007/978-3-319-04936-6_5](https://doi.org/10.1007/978-3-319-04936-6_5).
- [6] Timothy M. Chan. “All-pairs shortest paths for unweighted undirected graphs in $o(mn)$ time”. In: *ACM Trans. Algorithms* 8.4 (2012), 34:1–34:17. DOI: [10.1145/2344422.2344424](https://doi.org/10.1145/2344422.2344424).
- [7] George A. Chernishev et al. “On Several Social Network Analysis Problems”. In: *All-Russian Scientific Conference “Digital libraries: Advanced Methods and Technologies, Digital Collections”*. 2014, pp. 234–242. URL: http://ceur-ws.org/Vol-1297/234-242_paper-34.pdf.
- [8] George A. Chernishev et al. *On Several Social Network Analysis Problems: A Report*. Tech. rep. Saint-Petersburg University, Russia, 2014. URL: <https://www.math.spbu.ru/user/chernishev/papers/sigmod2014contest-report.pdf>.
- [9] Timothy A. Davis. “Algorithm 1000: SuiteSparse:GraphBLAS: Graph Algorithms in the Language of Sparse Linear Algebra”. In: *ACM Trans. Math. Softw.* 45.4 (2019), 44:1–44:25. DOI: [10.1145/3322125](https://doi.org/10.1145/3322125).
- [10] Markus Dreseler et al. “Quantifying TPC-H Choke Points and Their Optimizations”. In: *Proc. VLDB Endow.* 13.8 (2020), pp. 1206–1220. URL: <http://www.vldb.org/pvldb/vol13/p1206-dreseler.pdf>.
- [11] Márton Elekes et al. “A GraphBLAS solution to the SIGMOD 2014 Programming Contest using multi-source BFS”. In: *High Performance Extreme Computing (HPEC)*. IEEE, 2020.
- [12] Orri Erling et al. “The LDBC Social Network Benchmark: Interactive Workload”. In: *SIGMOD*. 2015, pp. 619–630. DOI: [10.1145/2723372.2742786](https://doi.org/10.1145/2723372.2742786).
- [13] Andrey Gubichev and Peter A. Boncz. “Parameter Curation for Benchmark Queries”. In: *TPCTC*. Vol. 8904. Lecture Notes in Computer Science. Springer, 2014, pp. 113–129. DOI: [10.1007/978-3-319-15350-6_8](https://doi.org/10.1007/978-3-319-15350-6_8).

- [14] Georg Hinkel. “The TTC 2018 Social Media Case”. In: *Transformation Tool Contest, co-located with the 2018 Software Technologies: Applications and Foundations (TTC@STAF)*. Vol. 2310. CEUR Workshop Proceedings. CEUR-WS.org, 2018, pp. 39–43. URL: <http://ceur-ws.org/Vol-2310/paper5.pdf>.
- [15] Moritz Kaufmann et al. “Hochperformante Analyse von Graph-Datenbanken”. In: *Datenbanksysteme für Business, Technologie und Web (BTW)*. Vol. P-241. LNI. GI, 2015, pp. 311–330. URL: <https://dl.gi.de/20.500.12116/2414>.
- [16] Moritz Kaufmann et al. “Parallel Array-Based Single- and Multi-Source Breadth First Searches on Large Dense Graphs”. In: *EDBT*. 2017. DOI: [10.5441/002/edbt.2017.02](https://doi.org/10.5441/002/edbt.2017.02).
- [17] LDBC Social Network Benchmark task force. *The LDBC Social Network Benchmark (version 0.4.0-SNAPSHOT)*. Tech. rep. Linked Data Benchmark Council, 2020. URL: https://ldbc.github.io/ldbc_snb_docs/ldbc-snb-specification.pdf.
- [18] Tobias Mühlbauer et al. “Instant Loading for Main Memory Databases”. In: *VLDB* 6.14 (2013), pp. 1702–1713. DOI: [10.14778/2556549.2556555](https://doi.org/10.14778/2556549.2556555).
- [19] Paul W. Olsen, Alan G. Labouseur, and Jeong-Hyon Hwang. “Efficient top-k closeness centrality search”. In: *ICDE*. IEEE Computer Society, 2014, pp. 196–207. DOI: [10.1109/ICDE.2014.6816651](https://doi.org/10.1109/ICDE.2014.6816651).
- [20] Minh-Duc Pham, Peter A. Boncz, and Orri Erling. “S3G2: A Scalable Structure-Correlated Social Graph Generator”. In: *TPCTC*. Springer, 2012, pp. 156–172. DOI: [10.1007/978-3-642-36727-4_11](https://doi.org/10.1007/978-3-642-36727-4_11).
- [21] Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003. ISBN: 978-0-89871-534-7. DOI: [10.1137/1.9780898718003](https://doi.org/10.1137/1.9780898718003).
- [22] Raimund Seidel. “On the All-Pairs-Shortest-Path Problem in Unweighted Undirected Graphs”. In: *J. Comput. Syst. Sci.* 51.3 (1995), pp. 400–403. DOI: [10.1006/jcss.1995.1078](https://doi.org/10.1006/jcss.1995.1078).
- [23] Yossi Shiloach and Uzi Vishkin. “An $O(\log n)$ Parallel Connectivity Algorithm”. In: *J. Algorithms* 3.1 (1982), pp. 57–67. DOI: [10.1016/0196-6774\(82\)90008-6](https://doi.org/10.1016/0196-6774(82)90008-6).
- [24] Edgar Solomonik, Aydin Buluç, and James Demmel. “Minimizing Communication in All-Pairs Shortest Paths”. In: *IPDPS*. IEEE Computer Society, 2013, pp. 548–559. DOI: [10.1109/IPDPS.2013.111](https://doi.org/10.1109/IPDPS.2013.111).
- [25] Edgar Solomonik and Torsten Hoefler. “Sparse Tensor Algebra as a Parallel Programming Model”. In: *CoRR* abs/1512.00066 (2015). URL: <http://arxiv.org/abs/1512.00066>.
- [26] Michael Sutton, Tal Ben-Nun, and Amnon Barak. “Optimizing Parallel Graph Connectivity Computation via Subgraph Sampling”. In: *IPDPS*. IEEE Computer Society, 2018, pp. 12–21. DOI: [10.1109/IPDPS.2018.00012](https://doi.org/10.1109/IPDPS.2018.00012).
- [27] Gábor Szárnyas et al. “An early look at the LDBC Social Network Benchmark’s Business Intelligence workload”. In: *GRADES-NDA at SIGMOD/PODS*. ACM, 2018, 9:1–9:11. DOI: [10.1145/3210259.3210268](https://doi.org/10.1145/3210259.3210268).

- [28] Robert Endre Tarjan. “Depth-First Search and Linear Graph Algorithms”. In: *SIAM J. Comput.* 1.2 (1972), pp. 146–160. DOI: [10.1137/0201010](https://doi.org/10.1137/0201010).
- [29] Manuel Then et al. “Efficient Batched Distance and Centrality Computation in Unweighted and Weighted Graphs”. In: *Datenbanksysteme für Business, Technologie und Web (BTW)*. Vol. P-265. LNI. GI, 2017, pp. 247–266. URL: <https://dl.gi.de/20.500.12116/632>.
- [30] Manuel Then et al. “Efficient Batched Distance, Closeness and Betweenness Centrality Computation in Unweighted and Weighted Graphs”. In: *Datenbank-Spektrum* 17.2 (2017), pp. 169–182. DOI: [10.1007/s13222-017-0261-x](https://doi.org/10.1007/s13222-017-0261-x).
- [31] Manuel Then et al. “Evaluation of parallel graph loading techniques”. In: *GRADES at SIGMOD*. ACM, 2016. DOI: [10.1145/2960414.2960418](https://doi.org/10.1145/2960414.2960418).
- [32] Manuel Then et al. “The More the Merrier: Efficient Multi-Source Graph Traversal”. In: *PVLDB* 8.4 (2014), pp. 449–460. DOI: [10.14778/2735496.2735507](https://doi.org/10.14778/2735496.2735507).
- [33] Alexandre Tiskin. “All-Pairs Shortest Paths Computation in the BSP Model”. In: *ICALP*. Vol. 2076. Lecture Notes in Computer Science. Springer, 2001, pp. 178–189. DOI: [10.1007/3-540-48224-5_15](https://doi.org/10.1007/3-540-48224-5_15).
- [34] Vadim Tropashko. “Nested intervals tree encoding in SQL”. In: *SIGMOD Rec.* 34.2 (2005), pp. 47–52. DOI: [10.1145/1083784.1083793](https://doi.org/10.1145/1083784.1083793).

A. Configuration of the LDBC Datagen

We list the configuration files used for generating the data sets used in the contest. These configurations were passed in the form of the `params.ini` to the Datagen.

```
numtotalUser:100000  
startYear:2013  
numYears:1  
serializerType:csv
```

Listing 1.: Configuration for the 100k Person data set.

```
numtotalUser:250000  
startYear:2010  
numYears:2  
serializerType:csv
```

Listing 2.: Configuration for the 250k Person data set.

```
numtotalUser:500000  
startYear:2010  
numYears:2  
serializerType:csv
```

Listing 3.: Configuration for the 500k Person data set.

```
numtotalUser:1000000  
startYear:2010  
numYears:2  
serializerType:csv
```

Listing 4.: Configuration for the 1M Person data set.

B. Sample queries and answers

We present sample query parameters and expected outputs (answers) for the **1k** data set.

Note There is a one-to-one mapping between queries and answers. The answers files contain comments after the `%` character. These are for debugging purposes. Files produced by solutions must not contain such comments.

```
query1(576, 400, -1)  
query1(58, 402, 0)  
query1(266, 106, -1)  
query1(313, 523, -1)  
query1(858, 587, 1)  
query1(155, 355, -1)  
query1(947, 771, -1)  
query1(105, 608, 3)  
query1(128, 751, -1)  
query1(814, 641, 0)
```

Listing 5.: `1k-sample-queries1.txt`

```
3 % path 576-618-951-400 (other shortest paths may exist)  
3 % path 58-935-808-402 (other shortest paths may exist)  
3 % path 266-23-592-106 (other shortest paths may exist)
```

```

-1 % path none
4 % path 858-46-31-162-587 (other shortest paths may exist)
3 % path 155-21-0-355 (other shortest paths may exist)
2 % path 947-625-771 (other shortest paths may exist)
-1 % path none
3 % path 128-459-76-751 (other shortest paths may exist)
3 % path 814-109-557-641 (other shortest paths may exist)

```

Listing 6.: 1k-sample-answers1.txt

```

query2(3, 1980-02-01)
query2(4, 1981-03-10)
query2(3, 1982-03-29)
query2(3, 1983-05-09)
query2(5, 1984-07-02)
query2(3, 1985-05-31)
query2(3, 1986-06-14)
query2(7, 1987-06-24)
query2(3, 1988-11-10)
query2(4, 1990-01-25)

```

Listing 7.: 1k-sample-queries2.txt

```

Chiang_Kai-shek Augustine_of_Hippo Napoleon % component sizes 22 16 16
Chiang_Kai-shek Napoleon Mohandas_Karamchand_Gandhi Sukarno % component sizes 17 13 11 11
Chiang_Kai-shek Mohandas_Karamchand_Gandhi Napoleon % component sizes 13 11 10
Chiang_Kai-shek Mohandas_Karamchand_Gandhi Augustine_of_Hippo % component sizes 12 10 8
Chiang_Kai-shek Aristotle Mohandas_Karamchand_Gandhi Augustine_of_Hippo Fidel_Castro % component
  sizes 10 7 6 5 5
Chiang_Kai-shek Mohandas_Karamchand_Gandhi Joseph_Stalin % component sizes 6 6 5
Chiang_Kai-shek Mohandas_Karamchand_Gandhi Joseph_Stalin % component sizes 6 6 5
Chiang_Kai-shek Augustine_of_Hippo Genghis_Khan Haile_Selassie_I Karl_Marx Lyndon_B._Johnson
  Robert_John_\\"Mutt\"_Lange % component sizes 4 3 3 3 3 3
Aristotle Ho_Chi_Minh Karl_Marx % component sizes 2 2 2
Arthur_Conan_Doyle Ashoka Barack_Obama Benito_Mussolini % component sizes 1 1 1 1

```

Listing 8.: 1k-sample-answers2.txt

```

query3(3, 2, Asia)
query3(4, 3, Indonesia)
query3(3, 2, Egypt)
query3(3, 2, Italy)
query3(5, 4, Chengdu)
query3(3, 2, Peru)
query3(3, 2, Democratic_Republic_of_the_Congo)
query3(7, 6, Ankara)
query3(3, 2, Luoyang)
query3(4, 3, Taiwan)

```

Listing 9.: 1k-sample-queries3.txt

```

361|812 174|280 280|812 % common interest counts 4 3 3
396|398 363|367 363|368 363|372 % common interest counts 2 1 1 1
110|116 106|110 106|112 % common interest counts 1 0 0
420|825 421|424 10|414 % common interest counts 1 1 0
590|650 590|658 590|614 590|629 590|638 % common interest counts 1 1 0 0 0
65|766 65|767 65|863 % common interest counts 0 0 0
99|100 99|101 99|102 % common interest counts 0 0 0
891|898 890|891 890|895 890|898 890|902 891|895 891|902 % common interest counts 1 0 0 0 0 0 0
565|625 653|726 565|653 % common interest counts 2 1 0
795|798 797|798 567|795 567|796 % common interest counts 1 1 0 0

```

Listing 10.: 1k-sample-answers3.txt

```
query4(3, Bill_Clinton)
query4(4, Napoleon)
query4(3, Chiang_Kai-shek)
query4(3, Charles_Darwin)
query4(5, Ronald_Reagan)
query4(3, Aristotle)
query4(3, George_W._Bush)
query4(7, Tony_Blair)
query4(3, William_Shakespeare)
query4(4, Augustine_of_Hippo)
```

Listing 11.: 1k-sample-queries4.txt

```
385 492 819 % centrality values 0.5290135396518375 0.5259615384615384 0.5249520153550864
722 530 366 316 % centrality values 0.5411255411255411 0.5405405405405406 0.5387931034482758
0.5382131324004306
592 565 625 % centrality values 0.5453460620525059 0.5421115065243179 0.5408284023668639
458 305 913 % centrality values 0.5415676959619953 0.5371024734982333 0.5345838218053928
953 294 23 100 405 % centrality values 0.5446859903381642 0.5394736842105263 0.5388291517323776
0.5375446960667462 0.5343601895734598
426 819 429 % centrality values 0.5451219512195121 0.5424757281553397 0.5366146458583433
323 101 541 % centrality values 0.553596806524207 0.5433196380862576 0.5413098243818695
465 647 366 722 194 135 336 % centrality values 0.535629340423861 0.5317101013475888
0.5310624641961301 0.530416402804164 0.5297719114277313 0.5259376153257211 0.5253039555482203
424 842 23 % centrality values 0.5316276893212858 0.5296265813313688 0.5256692220686188
385 562 659 323 % centrality values 0.5506329113924051 0.54375 0.54375 0.5291970802919708
```

Listing 12.: 1k-sample-answers4.txt