SecGDB: Graph Encryption for Exact Shortest Distance Queries with Efficient Updates

Qian Wang[†], Kui Ren[‡], Minxin Du[†], Qi Li[§], and Aziz Mohaisen[‡]

[†]School of CS, Wuhan University, Wuhan, China
`{qianwang,duminxin}@whu.edu.cn`
[‡]Department of CSE, University at Buffalo, SUNY, Buffalo, USA
`{kuiren,mohaisen}@buffalo.edu`
[§]Graduate School at Shenzhen, Tsinghua University, Shenzhen, China
`qi.li@sz.tsinghua.edu.cn`

Abstract. In the era of big data, graph databases have become increasingly important for NoSQL technologies, and many systems (*e.g.*, online social networks, world-wide web and electrical grids, etc.) can be modeled as graphs for semantic queries. Meanwhile, with the advent of cloud computing, data owners are highly motivated to outsource and store their massive potentially-sensitive graph data on remote untrusted servers in an encrypted form, expecting to retain the ability to query over the encrypted graphs.

To allow effective and private queries over encrypted data, the most well-studied class of *structured encryption* schemes are searchable symmetric encryption (SSE) designs, which encrypt search structures (*e.g.*, inverted indexes based on keyword-file pairs) for retrieving data files of interest from remote servers. So far, however, the problem of graph data encryption that supports customized queries has received limited attention in the literature. In this paper, we tackle the challenge of designing a Secure Graph DataBase encryption scheme (SecGDB) to encrypt graph structures and enforce private graph queries over the encrypted graph database. Specifically, our construction strategically makes use of efficient additively homomorphic encryption and garbled circuits to support the shortest distance queries with optimal time and storage complexities. To achieve better amortized time complexity over multiple queries, we further propose an auxiliary data structure called *query history* and store it on the remote server to act as a "caching" resource. Compared with the state-of-the-art solutions, our design returns exact shortest distance query results instead of approximate ones and allows efficient graph update queries over large-scale encrypted graphs. We prove that our construction is adaptively semantically-secure in the random oracle model and finally implement and evaluate it on various representative real-world datasets, showing that our approach is practically efficient in terms of both storage and computation.

Keywords: Graph encryption; shortest distance query; homomorphic encryption; garbled circuit.

1   Introduction

Graphs are used in a wide range of application domains, including social networks, interactive and online games, online knowledge discovery, computer networks, and the world-wide web, among others. For example, online social networks (OSN) such as Facebook and LinkedIn employ large social graphs with millions or even billions of vertices and edges in their operation. As a result, various systems have been recently proposed to handle massive graphs efficiently, where examples include Pregel [36], GraphLab [35], Horton [45] and TurboGraph [21]. These database applications allow for querying, managing and analyzing large-scale graphs in an intuitive and expressive way.

With the increased popularity of cloud computing, data users, including both individuals and enterprises, are highly motivated to outsource their (potentially huge amount of sensitive) data that may be abstracted and modeled as large graphs to remote cloud servers to reduce the local storage and management costs [28, 3, 13, 47]. However, database outsourcing also raises data confidentiality and privacy concerns due to data owners' loss of physical data control. Privacy-sensitive data therefore should be encrypted locally before outsourcing it to the untrusted cloud. Data encryption, however, hinders data utilization and computation, making it difficult to efficiently retrieve or query data of interest as opposed to the case with plaintext.

To address this challenge, the notion of *structured encryption* was first introduced by Chase and Kamara [8]. Roughly speaking, a structured encryption scheme encrypts structured data in such a way that it can be privately queried through the use of a specific token generated with knowledge of the secret key. Specifically, they presented approaches for encrypting (structured) graph data while allowing for efficient neighbor queries (*i.e.*, queries that return all vertices adjacent to a specified vertex), adjacency queries (*i.e.*, queries that would return whether two vertices are adjacent or not) and focused subgraph queries on labeled graphs (*i.e.*, queries to obtain pages ranking for a search keyword on a graph).

Despite all of these important types of queries, finding the shortest distance between two vertices, one of the most fundamental graph operations, was not supported. The shortest distance queries are not only building blocks for various more complex algorithms, but also have applications of their own. Such applications include finding the shortest path for one person to meet another in an encrypted social network graph, seeking the shortest path with the minimum delay in an encrypted networking or telecommunications abstracted graph, or performing a privacy-preserving GPS guidance in which one party holds the encrypted map while the other knows his origin and destination.

Recently, Meng *et al.* [38] addressed the graph encryption problem by pre-computing a data structure called the *distance oracle* from an original graph. They leveraged somewhat homomorphic encryption and standard private key encryption for their construction, thus answering shortest distance queries *approximately* over the encrypted distance oracle. Although their experimental results show that their schemes are practically efficient, the accuracy is sacri-

ficed for using the *distance oracle* (*i.e.*, only the approximate distance or even the negative result is returned). On the one hand, the distance oracle based methods only provide an estimate on the length of the shortest path. On the other hand, the exact path itself could also be necessary and important in many of the aforementioned application scenarios. Furthermore, both of the previous solutions only deal with static graphs [8, 38]: the outsourced encrypted graph structure cannot explicitly support efficient graph updates, since it requires to either re-encrypt the entire graph (*e.g.*, in [38], one has to pre-compute the corresponding distance oracle based on the updated graph, encrypt the graph, and then outsource it to the server for future queries), or make use of generic and expensive dynamization techniques similar to [9].

To tackle the practical limitations of the state-of-the-art, we propose a new Secure Graph DataBase encryption scheme (SecGDB) that supports both exact shortest distance queries and efficient dynamic operations. Specifically, our construction addresses four major challenges. First, to seek the best tradeoff between accuracy and efficiency, we process the graph itself instantiated by adjacency lists instead of encrypting either the *distance oracle* pre-computed from the original graph or the adjacency matrix instantiation. As a result, our scheme can be built on general graphs (*i.e.*, sparse or dense graphs) with the benefit of the adjacency list representation. Second, to compute the exact shortest path over the encrypted graph, we propose a hybrid approach that combines additively homomorphic encryption and garbled circuits to implement Dijkstra's algorithm [12] with the priority queue in a secure manner. Third, to enable dynamic updates of encrypted graphs, we carefully design an extra encrypted data structure to store the relevant information (*e.g.*, neighbor information of nodes in adjacency lists) which will be used to perform modifications homomorphically over the graph ciphertexts. Fourth, to further optimize the performance of the query phase, we introduce an auxiliary data structure called the *query history* by leveraging the previous queried results stored on the remote server as a "caching" resource; namely, the results for subsequent queries can be returned immediately without incurring further cost.

Our main contributions are summarized as follows.

- *Functionality and efficiency.* We propose SecGDB to support exact shortest distance queries with optimal time and storage complexity. We further obtain an improved amortized running time over multiple queries with the auxiliary data structure called "*query history*".
- *Dynamics.* We design an additional encrypted data structure to facilitate efficient graph updates. Compared with the state-of-the-art [8, 38], which consider only static data, SecGDB performs dynamic (*i.e.*, addition or removal of specified edges over the encrypted graph) operations with $O(1)$ time complexity.
- *Security, implementation and evaluation.* We formalize our security model using a simulation-based definition and prove the adaptive semantic security of SecGDB under the random oracle model with reasonable leakage. We im-

plement and evaluate the performance of SecGDB on various representative real-world datasets to demonstrate its efficiency and practicality.

1.1 Related Work

Song *et al.* [46] introduced the notion of searchable symmetric encryption (SSE), a technique that can be viewed as a specialization of the *structured encryption*, which received an extensive attention in the past few years [16, 19, 9, 28, 27, 7, 47, 20]. The notion of adaptive semantic security of SSE schemes was first introduced by Curtmola *et al.* [9] and further generalized to the setting of structured encryption in [8]. Since then, a great effort has been devoted to constructing efficient SSE schemes. Generally speaking, these SSE schemes usually encrypt inverted vectors or search trees for the special purpose of performing keyword based (*e.g.*, single keyword or multiple keywords) search over the encrypted document collections. Apparently, generic cryptographic primitives such as the fully homomorphic encryption [15] and oblivious RAM [17] can also be applied to realize almost all functionalities for structured encryptions. In practice, however, these generic solutions would incur prohibitively a huge amount computation and communication that could be impractical for both resource-constrained clients and the more powerful cloud servers.

Recently, other customized approaches have been developed for privacy-preserving shortest path computation. Aly *et al.* [1] investigated the problem of solving traditional graph problems, such as the shortest path problem, using multi-party computation techniques. However, their scheme has a cubic complexity in the number of vertices, which makes it impractical for large sparse graphs due to the adjacency matrix representation of the graph. Another line of work has focused on developing data oblivious algorithms for shortest path computation [3] or combining Dijkstra's algorithm with oblivious data structures to compute on sparse planer graphs [48], thus hiding partial information about the access pattern (*e.g.*, the type of operations etc.). However, the former has a quadratic complexity in number of vertices and the latter incurs bandwidth blowup and requires expensive offline computations. An oblivious secure computation framework combining garbled circuits and ORAM proposed by [34] requires communication on the order of GB and running times ranging from tens of minutes to several hours for a single query on a network graph with 1024 vertices. It is clear that the framework is not practically efficient despite providing strong security guarantees. Recently, Wu *et al.* [49] developed a privacy-preserving navigation protocol based on private information retrieval (PIR) and garbled circuits, and the routing information of the original street-map graph is compressed for the shortest path computation. We note that the navigation protocol is not applicable in our model where the graph itself stores sensitive information and must be blinded. Furthermore, interactions between the client and the server are involved in their protocol to compute each intermediate hop for every requested shortest path query.

The most related work to ours is due to Meng *et al.* [38], in which they presented structured encryption schemes for supporting *approximate* shortest

distance queries. However, their distance oracle based constructions return only the approximate distance results without generating the exact path after query, and they cannot explicitly support efficient graph updates.

2 Preliminaries and Notations

We begin by outlining some notations. Given a graph $G = (V, E)$ which consists of a set of vertices V and edges E, we denote its total number of vertices as $n = |V|$ and its number of edges as $m = |E|$. G is either *undirected* or *directed*. If G is undirected, then each edge in E is an *unordered* pair of vertices, and we use $\mathsf{len}(u, v)$ to denote the length of edge (u, v), otherwise, each edge in E is an *ordered* pair of vertices. In an undirected graph, $\mathsf{deg}(v)$ is used to denote the number of vertices adjacent to the vertex v (*i.e.*, *degree*). For a directed graph, we use $\mathsf{deg}^-(v)$ and $\mathsf{deg}^+(v)$ to denote the number of edges directed to vertex v (*indegree*) and out of vertex v (*outdegree*), respectively. A shortest distance query $q = (s, t)$ asks for the length (along with the route) of the shortest path between s and t, which we denote by $\mathsf{dist}(s, t)$ or dist_q. $[n]$ denotes the set of positive integers less than or equal to n, *i.e.*, $[n] = \{1, 2, \ldots, n\}$. We write $x \xleftarrow{\$} X$ to represent an element x being uniformly sampled at random from a set X. The output x of a probabilistic algorithm $\mathcal{A}$ is denoted by $x \leftarrow \mathcal{A}$ and that of a deterministic algorithm $\mathcal{B}$ by $x := \mathcal{B}$. Given a sequence of elements $\mathbf{v}$, we refer to the i^{th} element as $\mathbf{v}[i]$ or $\mathbf{v}_i$ and to the total number of elements in $\mathbf{v}$ by $|\mathbf{v}|$. If A is a set then $|A|$ refers to its cardinality, and if s is a string then $|s|$ refers to its bit length. We denote the concatenation of n strings $s_1, \ldots, s_n$ by $\langle s_1, \ldots, s_n \rangle$, and also denote the high-order $|s_2|$-bit of the string s_1 by $s_1^{|s_2|}$.

We also use various basic data structures including linked lists, arrays and dictionaries. Specifically, a dictionary T (also known as a map or associative array) is a data structure that stores key-value pairs (k, v). If the pair (k, v) is in T, then $\mathsf{T}[k]$ is the value v associated with k. An insertion operation of a new key-value pair (k, v) to the dictionary T is denoted by $\mathsf{T}[k] := v$. Similarly, a lookup operation takes a dictionary T and a specified key k as input, then returns the associated value v denoted by $v := \mathsf{T}[k]$.

2.1 Cryptographic Tools

Homomorphic encryption. Homomorphic encryption allows certain computations to be carried out on ciphertexts to generate an encrypted result which matches the result of operations performed on the plaintext after being decrypted. In a nutshell, a fully homomorphic encryption scheme (FHE) [44, 15, 6] allows evaluation of arbitrarily complex functions on encrypted data, and the cornerstone of FHE is the notion of a "somewhat homomorphic" encryption (SHE) scheme, which allows evaluation of functions below some complexity threshold. In this work, we only require the evaluation to efficiently support any number of additions, and there are many cryptosystems satisfying with this property. In particular, we use the Paillier cryptosystem [42] in our construction.

In the Paillier cryptosystem, the public (encryption) key is $pk_p = (n = pq, g)$, where $g \in \mathbb{Z}_{n^2}^*$, and p and q are two large prime numbers (of equivalent length) chosen randomly and independently. The private (decryption) key is $sk_p = (\varphi(n), \varphi(n)^{-1} \mod n)$. Given a message a, we write the encryption of a as $[\![a]\!]_{pk}$, or simply $[\![a]\!]$, where pk is the public key. The encryption of a message $x \in \mathbb{Z}_n$ is $[\![x]\!] = g^x \cdot r^n \mod n^2$, for some random $r \in \mathbb{Z}_n^*$. The decryption of the ciphertext is $x = L([\![x]\!]^{\varphi(n)} \mod n^2) \cdot \varphi^{-1}(n) \mod n$, where $L(u) = \frac{u-1}{n}$. The homomorphic property of the Paillier cryptosystem is given by $[\![x_1]\!] \cdot [\![x_2]\!] = (g^{x_1} \cdot r_1^n) \cdot (g^{x_2} \cdot r_2^n) = g^{x_1+x_2}(r_1 r_2)^n \mod n^2 = [\![x_1 + x_2]\!]$.

Pseudo-random functions (PRFs) and permutations (PRPs). Let $F : \{0,1\}^\lambda \times \{0,1\}^* \to \{0,1\}^*$ be a PRF, which is a polynomial-time computable function that cannot be distinguished from random functions by any probabilistic polynomial-time adversary. A PRF is said to be a PRP when it is bijective. Readers can refer to [29] for the formal definition and security proof.

Oblivious transfer. Parallel 1-out-of-2 Oblivious Transfer (OT) of m l-bit strings [39, 25], denoted as OT_l^m, is a two-party protocol run between a chooser $\mathcal{C}$ and a sender $\mathcal{S}$. For $i = 1, \ldots, m$, the sender $\mathcal{S}$ inputs a pair of l-bit strings $s_i^0, s_i^1 \in \{0,1\}^l$ and the chooser $\mathcal{C}$ inputs m choice bits $b_i \in \{0,1\}$. At the end of the protocol, $\mathcal{C}$ learns the chosen strings $s_i^{b_i}$ but nothing about the unchosen strings $s_i^{1-b_i}$, whereas $\mathcal{S}$ learns nothing about the choice b_i.

Garbled circuits. Garbled circuits were first proposed by Yao [50, 51] for secure two-party computation and later proven practical by Malkhi *et al.* [37]. At a high level, garbled circuits allow two parties holding inputs x and y, respectively, to jointly evaluate an arbitrary function $f(x, y)$ represented as a boolean circuit without leaking any information about their inputs beyond what is implied by the function output. In a garbled circuits protocol, one party (the *generator*) converts a circuit computing f into an "encrypted" version. The other party (the *evaluator*) then obliviously computes the output of the circuit without learning any intermediate values.

Several optimization techniques have been proposed in the literature to construct the standard garbled circuits. Kolensikov *et al.* [31] introduced an efficient method for creating garbled circuits which allows "free" evaluation of XOR gates, namely without incurring any communication or cryptographic operations. Pinkas *et al.* [43] proposed an approach to reduce the size of garbled gates from four to three entries, thus saving 25% of the communication overhead.

2.2 Fibonacci heap

Fibonacci heap [14] is a data structure for implementing priority queues, which consists of a collection of *trees* satisfying the *minimum-heap* property; that is, the key of a child is always greater than or equal to the key of the parent. This implies that the minimum key is always at the root of one of the trees. Generally, a heap data structure supports the following six operations.

– Make-Heap() creates and returns a new heap containing no elements.

- Insert(H, x) inserts node x, whose key field $key(x)$ has already been filled, into heap H.
- Minimum(H) returns a pointer to the node with the minimum key in the heap H.
- Extract-MIN(H) deletes the node with the minimum key from heap H, and returns a pointer to that node.
- Decrease-Key(H, x) assigns to node x in the heap H the new key value $key(x)$, which is assumed to be no greater than its current key value.
- Delete(H, x) deletes node x from heap H.

Compared with many other priority queue data structures including the *Binary heap* and *Binomial heap*, the Fibonacci heap achieves a better amortized running time [14].

3   System model and definitions

In this work, we consider the problem of designing a structured encryption scheme that supports the shortest distance queries and dynamic operations over an encrypted graph stored on remote servers efficiently.

At a high level, as shown in Figure 1, our construction contains three entities, namely the client $\mathcal{C}$, the server $\mathcal{S}$ and the proxy $\mathcal{P}$. In the initialization stage, the client $\mathcal{C}$ processes the original graph G to obtain its encrypted form Ω_G and outsources it to the cloud server $\mathcal{S}$. Meanwhile, the client $\mathcal{C}$ distributes partial secret key sk to the proxy $\mathcal{P}$. The privacy holds as long as the server $\mathcal{S}$ and the proxy $\mathcal{P}$ do not collude (*e.g.*, they respectively belong to two independent cloud service providers). This architecture of two non-colluding entities has been commonly used in the related literature [4, 13, 41]. Subsequently, to enable the shortest distance query over the encrypted graph Ω_G, the client generates a query token τ_q based on the query q and submits it to the cloud server $\mathcal{S}$. During the query phase, we adopt Yao's garbled circuits: the cloud server $\mathcal{S}$ who acts as the generator, and the proxy $\mathcal{P}$ who acts as the evaluator jointly run the secure comparison protocol. Finally, the encrypted shortest distance along with the path are returned to the client $\mathcal{C}$. In addition, the graph storage service in consideration is *dynamic*, such that the client $\mathcal{C}$ may add or remove edges to or from the encrypted graph Ω_G as well as modify the length of the specified edge. To do so, the client generates an update token τ_u corresponding to the dynamic operations. Given τ_u, the server $\mathcal{S}$ can securely update the encrypted graph Ω_G.

Formally, the core functionalities of our system are listed as below.

Definition 1. *An encrypted graph database system supporting the shortest distance query and dynamic updates consists of the following five (possibly probabilistic) polynomial-time algorithms/protocols:*

$sk \leftarrow \mathsf{Gen}(1^\lambda)$*: is a probabilistic key generation algorithm run by the client. It takes as input a security parameter λ and outputs the secret key sk.*

$\Omega_G \leftarrow \mathsf{Enc}(sk, G)$*: is a probabilistic algorithm run by the client. It takes as input*

Fig. 1: System model

Fig. 2: The secure comparison circuit.

a secret key sk and a graph G, and outputs an encrypted graph Ω_G.

$\mathsf{dist}_q \leftarrow \mathsf{Dec}(sk, c_q)$: is a deterministic algorithm run by the client. It takes as input a secret key sk and an encrypted result c_q, and outputs dist_q including the shortest distance as well as its corresponding path.

$(c_q; \sigma') \leftarrow \mathsf{DistanceQuery}(sk, q; \Omega_G, \sigma)$: is a (possibly interactive and probabilistic) protocol run between the client and the server[1]. The client takes as input a secret key sk and a shortest distance query q, while the server takes as input the encrypted graph Ω_G and the query history σ (which is empty in the beginning). During the protocol execution, a query token τ_q is generated by the client based on the query q and then sent to the server. Upon completion of the protocol, the client obtains an encrypted result c_q while the server gets a (possibly new) query history σ'.

$(\perp; \Omega'_G, \sigma) \leftarrow \mathsf{UpdateQuery}(sk, u; \Omega_G)$: is a (possibly interactive and probabilistic) protocol run between the client and the server. The client takes as input a secret key sk and an update object u (e.g., the edges to be updated), while the server takes as input the encrypted graph Ω_G. During the protocol execution, an update token τ_u is generated by the client based on the object u and then sent to the server. Upon completion of the protocol, the client gets nothing while the server obtains an updated encrypted graph Ω'_G and a new empty query history σ.

3.1 Security Definitions

Intuitively, an queryable encrypted graph database system should meet some security guarantees. First, an adversary cannot forge the shortest distance or update queries without the secret key. Second, given an encrypted graph, an adversary cannot learn any information about the underlying graph structure.

---

[1] A protocol P run between the client and the server is denoted by $(u; v) \leftarrow P(x; y)$, where x and y are the client's and the server's inputs, respectively, and u and v are the client's and the server's outputs, respectively.
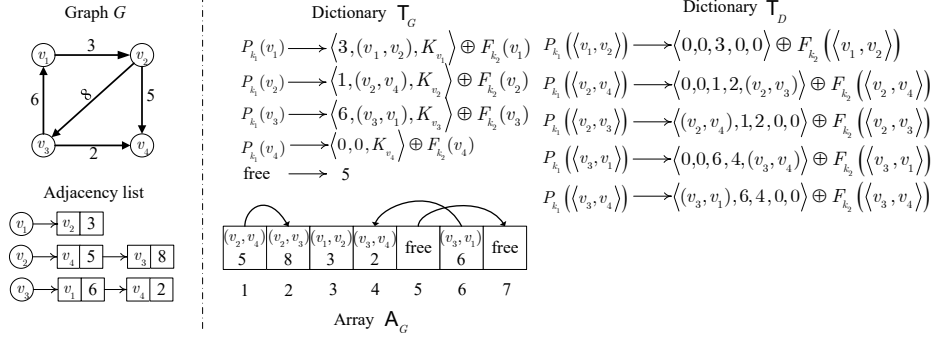
Fig. 3: An example of the encrypted graph construction.

Third, the tokens generated for a sequence of adaptive queries (including the shortest distance or the update queries) do not disclose any information about the original query objects. As in previous SSE systems [9, 28, 27, 20, 7] we also relax the security requirements appropriately by allowing some reasonable information leakage to the adversary in order to obtain higher efficiency. To capture this relaxation, we follow [9, 8, 28, 20] to parameterize the information by using a tuple of well-defined leakage functions (see Section 5).

Note that a secure comparison protocol based on the Yao's garbled circuits is constructed as a subroutine in the shortest distance query phase, and we assume that the server and the proxy are both semi-honest entities in our setting; they both run the protocol exactly as specified without any deviations, but try to learn extra information from their views of the protocol. A formal proof of security in the semi-honest model was given in [33]. Besides, various tools (*e.g.*, cut-and-choose techniques) can be applied to further extend the setting to handle the presence of malicious adversaries [37, 32, 26, 40].

In the following definition, we adapt the the notion of adaptive semantic security from [9, 8, 28] to our encrypted graph database system.

Definition 2. *(Adaptive semantic security) Let* (Gen, Enc, Dec, DistanceQuery, UpdateQuery) *be a dynamic encrypted graph database system and consider the following experiments with a stateful adversary $\mathcal{A}$, a stateful simulator $\mathcal{S}$ and three stateful leakage functions $\mathcal{L}_1$, $\mathcal{L}_2$ and $\mathcal{L}_3$:*
Real$_{\mathcal{A}}(\lambda)$: *The challenger runs* Gen(1^λ) *to generate the key sk. $\mathcal{A}$ outputs G and receives $\Omega_G \leftarrow$ Enc(sk, G) from the challenger. $\mathcal{A}$ then makes a polynomial number of adaptive shortest distance queries q or update queries u. For each q, the challenger acts as a client and runs DistanceQuery with $\mathcal{A}$ acting as a server. For each update query u, the challenger acts as a client and runs UpdateQuery with $\mathcal{A}$ acting as a server. Finally, $\mathcal{A}$ returns a bit b as the output of the experiment.*
Ideal$_{\mathcal{A},\mathcal{S}}(\lambda)$: *$\mathcal{A}$ outputs G. Given $\mathcal{L}_1(G)$, $\mathcal{S}$ generates and sends Ω_G to $\mathcal{A}$. $\mathcal{A}$ makes a polynomial number of adaptive shortest distance queries q or update*

queries u. For each q, $\mathcal{S}$ is given $\mathcal{L}_2(G, q)$, and simulates a client who runs DistanceQuery *with $\mathcal{A}$ acting as a server. For each update query u, $\mathcal{S}$ is given* $\mathcal{L}_3(G, u)$, *and simulates a client who runs* UpdateQuery *with $\mathcal{A}$ acting as a server. Finally, $\mathcal{A}$ returns a bit b as the output of the experiment.*

We say such a queryable encrypted graphs database system is adaptively $(\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3)$-*semantically secure if for all probabilistic polynomial-time (PPT) adversaries $\mathcal{A}$, there exists a probabilistic polynomial-time simulator $\mathcal{S}$ such that*

$$|\Pr[\mathbf{Real}_{\mathcal{A}}(\lambda) = 1] - \Pr[\mathbf{Ideal}_{\mathcal{A}, \mathcal{S}}(\lambda) = 1]| \leq \mathsf{negl}(\lambda),$$

where $\mathsf{negl}(\cdot)$ *is a negligible function.*

4   Our construction: SecGDB

In this section, we present our encrypted graph database system–SecGDB, which efficiently supports the shortest distance query and the update query (*i.e.*, to add, remove and modify a specified edge).

4.1   Overview

We assume that an original graph is instantiated by adjacency lists, and the length of edge (u, v) is stored in the node for v in the adjacency list for u. Namely, every node in each adjacency list contains a pair of the neighboring vertex and the length of the corresponding edge (*i.e.*, vertex and length pair).

Our construction is inspired by [28]. The key idea is as follows. During the initialization phase, we randomly shuffle all the nodes and place them in an array. That is, we place every node of each adjacency list at a random location in the array while updating the pointers so that the "logical" integrity of the lists are preserved. We then use the Paillier cryptosystem to encrypt the length of the edge in each node, and use a "standard" private-key encryption scheme [29] to blind the entire node. In the shortest distance query phase, if the query has been submitted before or was a subset of the *query history*, the encrypted result can be immediately returned to the client; otherwise, we implement the Dijkstra's algorithm (conceptually a breadth-first search) with the aid of an advanced data structure (*i.e.*, Fibonacci heap) in a secure manner to answer the shortest distance query. Specifically, we propose a hybrid approach by leveraging the additive homomorphic property and the garbled circuits to achieve our goal. Meanwhile, the query history is updated based on the query results. To support efficient dynamic operations on the encrypted graph, we generate the relevant update token, which allows the server to add or remove the specified entry to and from the array. After finishing the updates, the query history is rebuilt for future use.

4.2   Initialization Phase

We now describe the details of the preparation stage as well as the construction of the encrypted graph. Intuitively, the initialization phase consists of Gen and

---

Algorithm 1 Graph Enc algorithm

Input: $G = (V, E), sk$
Output: Ω_G

1: Set $n = |V|, m = |E|$;
2: Initialize an array A_G of size $m + z$;
3: Initialize two dictionaries $\mathsf{T}_G, \mathsf{T}_D$ of size $n+1$ and m;
4: Initialize a random permutation π over $[m + z]$;
5: Initialize a counter $\mathtt{ctr} = 1$;
6: **for** each vertex $u \in V$ **do**
7:     Generate $K_u := G_{k_3}(u)$;
8:     **for** $i = 1$ to $\deg^+(u)$ **do**
9:         Encrypt the length of the edge (u, v_i) under the Paillier cryptosystem $c_i \leftarrow [\![\mathsf{len}(u, v_i)]\!]_{pk_p}$
10:        **if** $i = 1$ and $i \neq \deg^+(u)$ **then**
11:            Set $\mathtt{N}_i := \langle P_{k_1}(v_i), F_{k_2}(v_i), c_i, \pi(\mathtt{ctr}+1)\rangle$;
12:            Set $\mathtt{D}_i := \langle \mathbf{0}, \mathbf{0}, \pi(\mathtt{ctr}), \pi(\mathtt{ctr} + 1), P_{k_1}(\langle u, v_{i+1}\rangle)\rangle$;
13:        **else if** $i \neq 1$ and $i = \deg^+(u)$ **then**
14:            Set $\mathtt{N}_i := \langle P_{k_1}(v_i), F_{k_2}(v_i), c_i, \mathtt{NULL}\rangle$;
15:            Set $\mathtt{D}_i := \langle P_{k_1}(\langle u, v_{i-1}\rangle), \pi(\mathtt{ctr} - 1), \pi(\mathtt{ctr}), \mathbf{0}, \mathbf{0}\rangle$;
16:        **else if** $i = 1$ and $i = \deg^+(u)$ **then**
17:            Set $\mathtt{N}_i := \langle P_{k_1}(v_i), F_{k_2}(v_i), c_i, \mathtt{NULL}\rangle$;
18:            Set $\mathtt{D}_i := \langle \mathbf{0}, \mathbf{0}, \pi(\mathtt{ctr}), \mathbf{0}, \mathbf{0}\rangle$;
19:        **else**
20:            Set $\mathtt{N}_i := \langle P_{k_1}(v_i), F_{k_2}(v_i), c_i, \pi(\mathtt{ctr}+1)\rangle$;
21:            Set $\mathtt{D}_i := \langle P_{k_1}(\langle u, v_{i-1}\rangle), \pi(\mathtt{ctr} - 1), \pi(\mathtt{ctr}), \pi(\mathtt{ctr}+1), P_{k_1}(\langle u, v_{i+1}\rangle)\rangle$;
22:        **end if**
23:        Sample $r_i \overset{\$}{\leftarrow} \{0,1\}^\lambda$;
24:        Store the encrypted $\mathtt{N}_i$ in the array $\mathsf{A}_G[\pi(\mathtt{ctr})] := \langle \mathtt{N}_i \oplus H(K_u, r_i), r_i\rangle$;
25:        Store the encrypted $\mathtt{D}_i$ in the dictionary $\mathsf{T}_D[P_{k_1}(\langle u, v_i\rangle)] := \mathtt{D}_i \oplus F_{k_2}(\langle u, v_i\rangle)$;
26:        Increase $\mathtt{ctr} = \mathtt{ctr} + 1$;
27:    **end for**
28:    Store a pointer to the head n-ode of the adjacency list for u in the dictionary $\mathsf{T}_G[P_{k_1}(u)] := \langle \mathtt{addr}(\mathtt{N}_1), P_{k_1}(\langle u, v_1\rangle), K_u\rangle \oplus F_{k_2}(u)$;
29: **end for**
30: **for** $i = 1$ to z **do**
31:    Set $\mathtt{F}_i := \langle \mathbf{0}, \pi(\mathtt{ctr}+1)\rangle$;
32:    **if** $i = z$ **then**
33:        Set $\mathtt{F}_i := \langle \mathbf{0}, \mathtt{NULL}\rangle$;
34:    **end if**
35:    Store the unencrypted $\mathtt{F}_i$ in the array $\mathsf{A}_G[\pi(\mathtt{ctr})] := \mathtt{F}_i$;
36:    Increase $\mathtt{ctr} = \mathtt{ctr} + 1$;
37: **end for**
38: Store a pointer to the head node of the free list in the dictionary $\mathsf{T}_G[\mathtt{free}] := \langle \mathtt{addr}(\mathtt{F}_1), \mathbf{0}\rangle$;
39: Output the encrypted graph $\Omega_G = (\mathsf{A}_G, \mathsf{T}_G, \mathsf{T}_D)$;

Enc as presented in Definition 1. The scheme uses the Paillier cryptosystem, and three pseudo-random functions P, F and G, where P is defined as $\{0,1\}^\lambda \times \{0,1\}^* \to \{0,1\}^\lambda$, F is defined as $\{0,1\}^\lambda \times \{0,1\}^* \to \{0,1\}^*$ and G is defined as $\{0,1\}^\lambda \times \{0,1\}^* \to \{0,1\}^\lambda$. We also use a random oracle H which is defined as $\{0,1\}^* \to \{0,1\}^*$.

Gen(1^λ): Given a security parameter λ, generate the following keys uniformly at random from their respective domains:

- three PRF keys $k_1, k_2, k_3 \overset{\$}{\leftarrow} \{0,1\}^\lambda$ for $P_{k_1}(\cdot)$, $F_{k_2}(\cdot)$ and $G_{k_3}(\cdot)$, respectively;
- (sk_p, pk_p) for the Paillier cryptosystem.

The output is $sk = (k_1, k_2, k_3, sk_p, pk_p)$, where sk_p is sent to the proxy through a secure channel.

As shown in Algorithm 1, the setup procedures are done in the first five steps. From line 6 to 29, the length of the edge is encrypted under the Paillier cryptosystem and the entire node $\mathtt{N}_i$ is encrypted by XORing an output of the random oracle H. Meanwhile, the neighboring information of each node $\mathtt{N}_i$ (*i.e.*, the nodes following and previous to $\mathtt{N}_i$ in the original adjacency lists, and the corresponding positions in A_G) constitutes the dual node $\mathtt{D}_i$, and the encrypted dual node will be stored in the dictionary T_D. Generally speaking, T_D stores

the pointer to each edge, and it is used to support efficient delete updates on the encrypted graph. After the aforementioned operations are done, the address of each head node will be encrypted and stored in the dictionary T_G, namely, T_G stores the pointer to the head of each adjacency list. The remaining z cells in the array construct an unencrypted `free` list, which is used in the add updates. To ensure the size of all the entries in A_G, T_G and T_D is identical, we should pad by a string of 0's (*i.e.*, $\mathbf{0}$). Finally, we output the encrypted graph Ω_G.

Figure 3 gives an illustrative example of the Enc algorithm to construct the encrypted graph from a directed graph with four vertices v_1, v_2, v_3 and v_4 as well as five edges. All the nodes contained in the original (three) adjacency lists are now stored at random locations in A_G, and the dictionaries T_G and T_D are also shown in Figure 3. Note that in a real encrypted graph, there would be padding to hide partial structural information of the original graph (as will be discussed in Section 5); we omit this padding for simplicity in this example.

4.3   Shortest Distance Query Phase

In this section, we will describe the process of performing the exact shortest distance query over the encrypted graph, as summarized in Algorithm 2.

The DistanceQuery protocol shown in Algorithm 2 works as follows. First, the client generates the query token τ_q based on a query $q = (s, t)$, and then sends it to the server. If the token has been queried before or acts as a subpath of the *query history* σ, the server returns the result c_q ($c_q \subset \sigma$) to the client immediately; otherwise, the server executes the Dijkstra's algorithm with the aid of a Fibonacci heap H in a private way. Concretely, the server first reads off the vertices that are adjacent to the source s and inserts to the heap H (line 14 to 22). Subsequently, each iteration of the loop from line 23 to 49 starts by extracting the vertex α with the minimum key. If the vertex α is the requested destination τ_2, the server updates the query history σ based on the newly-obtained path, computes the encrypted result c_q via reverse iteration and returns it to the client. Else, the server recovers the pointer to the head of the adjacency list for the vertex α, and then retrieves nodes in the adjacency list. Specifically, for the node N_i, once an update of $\xi[\alpha_i]$ occurs it indicates that a shorter path to α_i via α has been discovered, the server then updates the path. Next, the server either runs Insert(H, α_i) (if α_i is not in H) or Decrease-Key($\mathsf{H}, \alpha_i, key(\alpha_i)$). It is worth noting that both the conditional statement $\xi[\alpha] \cdot c_i < \xi[\alpha_i]$ and some specific operations on the Fibonacci heap (*e.g.*, Extract-MIN) require performing a comparison on the encrypted data. Hence we build a secure comparison protocol (see Section 4.3) based on the garbled circuits and invoke it as a subroutine.

Finally, the client runs Dec(c_q, sk) to obtain the dist_q as follows. Given c_q, the client parses it as a sequence of $\langle c_1, c_2 \rangle$ pairs, and for each pair, the client decrypts c_1 (the path) and c_2 (the distance) by using k_1 and sk_p, respectively.

Remarks. Conceptually, the history σ consists of all previous de-duplicated queried results. For a new query, the server traverses σ and checks whether the new query belongs to a record in σ. For example, let history σ consist of a shortest path from s to t (*i.e.*, $\{s, \ldots, u, \ldots, v, \ldots, t\}$), then for a new query

---

Algorithm 2 DistanceQuery protocol

Input:
    The client $\mathcal{C}$'s input is $sk, q = (s, t)$;
    The server $\mathcal{S}$'s input is Ω_G, σ;
Output:
    The client $\mathcal{C}$'s output is c_q;
    The server $\mathcal{S}$'s output is σ';
1: $\mathcal{C}$: compute $\tau_q := (P_{k_1}(s), P_{k_1}(t), F_{k_2}(s))$;
2: $\mathcal{C} \Rightarrow \mathcal{S}$: output τ_q to the server;
3: $\mathcal{S}$: parse τ_q as (τ_1, τ_2, τ_3);
4: **if** $\mathsf{T}_G[\tau_1] = \perp$ or $\mathsf{T}_G[\tau_2] = \perp$ **then**
5:      $\mathcal{S} \Rightarrow \mathcal{C}$: return $\perp$ to the client;
6: **else if** $\{\tau_1, \tau_2\} \subset \sigma$ **then**
7:      $\mathcal{S} \Rightarrow \mathcal{C}$: return c_q to the client;
8: **else**
9:      $\mathcal{S}$: initialize a Fibonacci heap $\mathsf{H} \leftarrow$ Make-Heap();
10:      $\mathcal{S}$: initialize two dictionaries ξ and path;
11:      $\mathcal{S}$: compute $\langle \mathsf{addr}_1, \mathsf{str}, K_s \rangle := \mathsf{T}_G[\tau_1] \oplus \tau_3$;
12:      $\mathcal{S}$: parse $\mathsf{A}_G[\mathsf{addr}_1]$ as $\langle \mathsf{N}'_1, r_1 \rangle$;
13:      $\mathcal{S}$: compute $\mathsf{N}_1 := \mathsf{N}'_1 \oplus H(K_s, r_1)$;
14:      **while** $\mathsf{addr}_{i+1} \neq \texttt{NULL}$ **do**
15:          $\mathcal{S}$: parse N_i as $\langle \alpha_i, \beta_i, c_i, \mathsf{addr}_{i+1} \rangle$;
16:          $\mathcal{S}$: store $\mathsf{path}[\alpha_i] := \langle \tau_1, c_i \rangle$;
17:          $\mathcal{S}$: set $\xi[\alpha_i] := c_i$ and $key(\alpha_i) := \xi[\alpha_i]$;
18:          $\mathcal{S}$: run Insert(H, α_i) with the $key(\alpha_i)$;
19:          $\mathcal{S}$: parse $\mathsf{A}_G[\mathsf{addr}_{i+1}]$ as $\langle \mathsf{N}'_{i+1}, r_{i+1} \rangle$;
20:          $\mathcal{S}$: compute $\mathsf{N}_{i+1} := \mathsf{N}'_{i+1} \oplus H(K_s, r_{i+1})$;
21:          $\mathcal{S}$: increase $i = i + 1$;
22:      **end while**
23:      **repeat**
24:          $\mathcal{S}$: parse Extract-MIN(H) as $\langle \alpha, key(\alpha) \rangle$;
25:          **if** $\alpha = \tau_2$ **then**
26:              $\mathcal{S}$: update σ' based on path;
27:              $\mathcal{S} \Rightarrow \mathcal{C}$: return c_q to the client;
28:              $\mathcal{S}$: break;
29:          **end if**
30:          $\mathcal{S}$: compute $\langle \mathsf{addr}_1, \mathsf{str}, K_u \rangle := \mathsf{T}_G[\alpha] \oplus \beta$;
31:          $\mathcal{S}$: parse $\mathsf{A}_G[\mathsf{addr}_1]$ as $\langle \mathsf{N}'_1, r_1 \rangle$;
32:          $\mathcal{S}$: compute $\mathsf{N}_1 := \mathsf{N}'_1 \oplus H(K_u, r_1)$;
33:          **while** $\mathsf{addr}_{i+1} \neq \texttt{NULL}$ **do**
34:              $\mathcal{S}$: parse N_i as $\langle \alpha_i, \beta_i, c_i, \mathsf{addr}_{i+1} \rangle$;
35:              **if** $\xi[\alpha] \cdot c_i < \xi[\alpha_i]$ **then**
36:                  $\mathcal{S}$: update $\xi[\alpha_i] := \xi[\alpha] \cdot c_i$;
37:                  $\mathcal{S}$: set $key(\alpha_i) := \xi[\alpha_i]$;
38:                  $\mathcal{S}$: store $\mathsf{path}[\alpha_i] := \langle \alpha, c_i \rangle$;
39:              **end if**
40:              **if** $\alpha_i \notin \mathsf{H}$ **then**
41:                  $\mathcal{S}$: run Insert(H, α_i) with the $key(\alpha_i)$;
42:              **else**
43:                  $\mathcal{S}$: run Decrease-Key($\mathsf{H}, \alpha_i, key(\alpha_i)$);
44:              **end if**
45:              $\mathcal{S}$: parse $\mathsf{A}_G[\mathsf{addr}_{i+1}]$ as $\langle \mathsf{N}'_{i+1}, r_{i+1} \rangle$;
46:              $\mathcal{S}$: compute $\mathsf{N}_{i+1} := \mathsf{N}'_{i+1} \oplus H(K_u, r_{i+1})$;
47:              $\mathcal{S}$: increase $i = i + 1$;
48:          **end while**
49:      **until** H is empty
50: **end if**

$q = (u, v)$, the corresponding encrypted result $c_q = \{u, \ldots, v\}$ where $c_q \subset \sigma$ can be returned immediately. Note that only lookup operations (of dictionary) are required, thus making the whole process highly efficient.

Secure Comparison Protocol We now present the secure comparison protocol which is based on the garbled circuits [24, 50] for selecting the minimum of two encrypted values. This subroutine is implemented by the circuit shown in Figure 2, and we use a CMP circuit and two SUB circuits constructed in [30] to realize the desired functionality.

At the beginning, the server has two encrypted values $[\![a_1]\!]$ and $[\![a_2]\!]$ and the proxy has the secret key sk_p. W.l.o.g., we assume that the longest shortest distance between any pair of vertices (*i.e.*, diameter [22]) lies in $[2^l]$, namely, a_1 and a_2 are two l-bit integers. Instead of sending $[\![a_1]\!]$ and $[\![a_2]\!]$ to the proxy, the server first masks them with two k-bit random numbers r_1 and r_2 (*e.g.*, $[\![a_1 + r_1]\!] = [\![a_1]\!] \cdot [\![r_1]\!]$) respectively, where k is a security parameter ($k > l$). Then the server's inputs are r_1 and r_2, and the proxy's inputs are $a_1 + r_1$ and $a_2 + r_2$. Finally, the output single bit x implies the comparison result: if $x = 1$, then $a_1 > a_2$; 0 otherwise. Note that masking here is done by performing addition

SUB circuit	CMP circuit	Total circuit
k	l	$2k + l$

Table 1: The number of non-free binary gates.

over the integers which is a form of statistical hiding. More precisely, for a l-bit integer a_i and a k-bit integer r_i, releasing $a_i + r_i$ gives statistical security of roughly 2^{l-k} for the potential value a_i. Therefore, by choosing the security parameter k properly, we can make this statistical difference arbitrarily low [24].

Since we adopt the free-XOR technique [31] in our construction, the XOR gates do not contribute much to the cost of the garbled circuits since no cryptographic operations are needed. Table 1 summarizes the number of non-XOR gates in each of our circuits.

Packing Optimization. It is worth noting that the message space of the Paillier cryptosystem is much greater than the space of the blinded values. We can therefore provide a great improvement in both computation time and bandwidth by leveraging the packing technique. The basic idea is to send one ciphertext in the form $[\![\langle(a_{i+1} + r_{i+1}), \ldots, (a_{i+p} + r_{i+p})\rangle]\!]$ instead of p ciphertexts of the form $[\![a_i + r_i]\!]$, where $p = \frac{1024}{k}$ (1024-bit modulus used in Paillier cryptosystem). More specifically, given two blinded ciphertexts $[\![a_1 + r_1]\!]$ and $[\![a_2 + r_2]\!]$, the server can aggregate them into a single ciphertext of the form $[\![\langle(a_1 + r_1), (a_2 + r_2)\rangle]\!] = [\![a_1 + r_1]\!]^{2^k} \cdot [\![a_2 + r_2]\!]$. Then the "packed" ciphertext can be obtained via aggregating p ciphertexts in the same manner.

We note that another line of work using specialized homomorphic encryption, such as Goldwasser-Micali cryptosystem [18, 5] or DGK cryptosystem [10, 11], can also solve the above millionaires problem. However, applying the homomorphic encryption based solution is not only time-consuming but also incurs prohibitively high communication overhead due to the bitwise encryption. Moreover, the entire cost of each comparison will be included in the query phase, and this is not desirable in practice. Fortunately, using garbled circuits allows us to move as many of expensive operations (*e.g.*, the computationally expensive OT and the creation of garbled circuits as well as the transfer of garbled circuits) into a pre-computation (offline) stage. In fact, the server can prepare the garbled circuits and send them together with the translation-tables to the proxy before a shortest distance query is submitted. Besides, we use the standard technique of [2] to pre-compute OTs, and the extensions of [25] can be used to reduce an arbitrary number of OTs to a constant number c, where c is a security parameter. Later in the query (online) phase, the server sends the garbled values corresponding to its input bits to the proxy, and the online part of the OT protocol is executed by the proxy to obtain its own garbled values. Using the garbled inputs, the proxy evaluates the garbled circuit for each comparison, obtains the result from the translation-table and sends it back to the server.

4.4   Supporting Encrypted Graph Dynamics

We next discuss the support of update operations over the encrypted graph. Here, we do not particularly consider the addition and removal of vertices, because the update of the vertex can be viewed as the update of a collection of related edges.

Algorithm 3 UpdateQuery protocol

Input:
  The client $\mathcal{C}$'s input is sk, u;
  The server $\mathcal{S}$'s input is Ω_G;
Output:
  The client $\mathcal{C}$'s output is $\perp$;
  The server $\mathcal{S}$'s output is Ω'_G, σ;

a) Adding new edges
At the client $\mathcal{C}$:
1) u contains information about newly-added edge (v_1, v_2) with the length $\mathsf{len}(v_1, v_2)$;
2) compute the update token $\tau_u := (P_{k_1}(v_1), F_{k_2}(v_1)^{|\langle \mathsf{addr}, \mathsf{str} \rangle|}, P_{k_1}(\langle v_1, v_2 \rangle), F_{k_2}(\langle v_1, v_2 \rangle), \mathbb{N})$, where $\mathbb{N} = \langle \langle P_{k_1}(v_2), F_{k_2}(v_2), [\![\mathsf{len}]\!], \mathbf{0} \rangle \oplus H(K_{v_1}, r), r \rangle$;
$\mathcal{C} \Rightarrow \mathcal{S}$: output τ_u to the server;
At the server $\mathcal{S}$:
1) parse τ_u as $(\tau_1, \tau_2, \tau_3, \tau_4, \tau_5)$ and return $\perp$ if τ_1 is not in T_G;
2) compute $\langle \mathsf{addr}_1, \mathbf{0} \rangle := \mathsf{T}_G[\mathsf{free}]$;
3) parse $\mathsf{A}_G[\mathsf{addr}_1]$ as $\langle \mathbf{0}, \mathsf{addr}_2 \rangle$;
4) update the pointer to the next free node $\mathsf{T}_G[\mathsf{free}] := \langle \mathsf{addr}_2, \mathbf{0} \rangle$;
5) compute $\langle \mathsf{addr}_3, \mathsf{str} \rangle := \mathsf{T}_G[\tau_1]^{|\langle \mathsf{addr}, \mathsf{str} \rangle|} \oplus \tau_2$;
6) parse τ_5 as $\langle \mathbb{N}', r \rangle$ and set $\mathsf{A}_G[\mathsf{addr}_1] := \langle \mathbb{N}' \oplus \langle \mathbf{0}, \mathsf{addr}_3 \rangle, r \rangle$;
7) update the pointer to the newly-added node $\mathsf{T}_G[\tau_1] := \mathsf{T}_G[\tau_1]^{|\langle \mathsf{addr}, \mathsf{str} \rangle|} \oplus \langle \mathsf{addr}_3, \mathsf{str} \rangle \oplus \langle \mathsf{addr}_1, \tau_3 \rangle$;

8) store $\mathsf{T}_D[\tau_3] := \langle \mathbf{0}, \mathbf{0}, \mathsf{addr}_1, \mathsf{addr}_3, \mathsf{str} \rangle \oplus \tau_4$;
9) update $\mathsf{T}_D[\mathsf{str}] := \mathsf{T}_D[\mathsf{str}]^{|\langle \mathsf{addr}, \mathsf{str} \rangle|} \oplus \langle \tau_3, \mathsf{addr}_1 \rangle$;
10) obtain an updated graph Ω'_G and rebuild σ;

b) Deleting existing edges
At the client $\mathcal{C}$:
1) u contains information about the existing edge (v_1, v_2) to be deleted;
2) compute $\tau_u := (P_{k_1}(\langle v_1, v_2 \rangle), F_{k_2}(\langle v_1, v_2 \rangle))$;
$\mathcal{C} \Rightarrow \mathcal{S}$: outputs τ_u to the server;
At the server $\mathcal{S}$:
1) parse τ_u as (τ_1, τ_2) and return $\perp$ if τ_1 is not in T_D;
2) look up in T_D and computes $\langle \mathsf{str}_1, \mathsf{addr}_1, \mathsf{addr}_2, \mathsf{addr}_3, \mathsf{str}_3 \rangle := \mathsf{T}_D[\tau_1] \oplus \tau_2$;
3) compute $\langle \mathsf{addr}_4, \mathbf{0} \rangle := \mathsf{T}_G[\mathsf{free}]$;
4) free the node and set $\mathsf{A}_G[\mathsf{addr}_2] := \langle \mathbf{0}, \mathsf{addr}_4 \rangle$;
5) update the pointer $\mathsf{T}_G[\mathsf{free}] := \langle \mathsf{addr}_2, \mathbf{0} \rangle$;
6) parse $\mathsf{A}_G[\mathsf{addr}_1]$ as $\langle \mathbb{N}'_1, r_1 \rangle$;
7) update node $\mathsf{A}_G[\mathsf{addr}_1] := \langle \mathbb{N}'_1 \oplus \mathsf{addr}_2 \oplus \mathsf{addr}_3, r_1 \rangle$;
8) update the corresponding entry $\mathsf{T}_D[\mathsf{str}_1] := \mathsf{T}_D[\mathsf{str}_1] \oplus \langle \mathsf{addr}_2, r_1 \rangle \oplus \langle \mathsf{addr}_3, \mathsf{str}_3 \rangle$;
9) update the corresponding entry $\mathsf{T}_D[\mathsf{str}_3] := \mathsf{T}_D[\mathsf{str}_3] \oplus \langle \mathsf{addr}_2, r_1 \rangle \oplus \langle \mathsf{addr}_1, \mathsf{str}_1 \rangle$;
10) obtain an updated graph Ω'_G and rebuild σ;

To add new edges, the client generates the corresponding token τ_u for an update object u and sends it to the server. After receiving τ_u, the server locates the first free node addr_1 in the array A_G, and modifies the pointer in T_G to point to the second one. Later, the server retrieves the high-order useful information (without the key K_{v_1}) of the head node $\mathbb{N}_1$, stores $\mathbb{N}$ that represents the newly edge at location addr_1 and modifies its pointer to point to the original head node $\mathbb{N}_1$ without decryption. Then, the server updates the pointer in T_G to point to the newly-added node, and finally updates the corresponding entries in the dictionary T_D. To remove the existing edges, the client generates the update token τ_u and submits it to the server. Subsequently, the server looks up in the T_D and recovers the adjacency information of the specified edge. In the following steps, the server frees the node, inserts it into the head of the free list and then homomorphically modifies the pointer of the previous node to point to the next node in A_G. Eventually, the server updates the related entries in the dictionary T_D. Note that modifying a specified edge can be easily achieved by removing the "old" edge first, and adding a "new" edge with the modified length later. After the encrypted graph has been updated, the old query history is deleted and a new empty history will be rebuilt simultaneously.

4.5   Performance Analysis

The time cost of initialization phase is dominated by encrypting all the edges using Paillier cryptosystem and processing all the vertices to obtain the encrypted dictionary, thus the time complexity of this part is $O(m+n)$, where m is the number of edges and n is the number of vertices in the original graph. The generated encrypted graph, which consists of an array and two dictionaries, has the storage complexity $O(m+n)$ at the server side. In the query phase, we use the Fibonacci heap to speed up the Dijkstra's algorithm in a private manner, where each iteration to find a vertex in the shortest path requires one Extract-MIN operation. In addition, each edge which satisfies the inequality in line 35 of Algorithm 2 requires either an Insert or a Decrease-Key operation. There are at most $(n-1)$ Extract-MIN operations, $(n-1)$ Insert operations and $(m-n+1)$ Decrease-Key operations in total. Thus, we obtain an $O(n \log n + m)$ time complexity which is optimal among other priority queue optimization techniques (*e.g.*, binary or binomial heap) [14]. By maintaining an auxiliary structure history σ at the server, we can have an even better amortization time complexity over multiple queries, *i.e.*, the query time for subsequent queries that can be looked up in the history are (almost) constant. It is obvious that the time complexity for both addition and removal operations on the encrypted graph are only $O(1)$.

During the execution of the secure comparison protocol, the computation and communication costs between the server and the proxy are directly related to the number of gates in the comparison circuit. As we discussed above, many expensive operations of the garbled comparison circuits can be pushed into a pre-computation phase, and thus most of the computation cost will be relieved from the query phase. On the other hand, the communication cost in the offline computation phase (pre-computation) is dominated by the transfer of the circuits. More concretely, the server transmits $(6k + 3l)t$ bits for each comparison circuit. Besides, about $6kt$ bits cost for transmitting garbled inputs are incurred in the online (query) phase, where t is the bit length of a garbled value for a wire. The communication cost of the OT protocol is omitted here since it only needs to be performed a constant number of times (see Section 4.3).

5   Security

We allow reasonable leakage to the server to trade it for efficiency. Now, we provide a formal description of the three leakage functions $\mathcal{L}_1$, $\mathcal{L}_2$ and $\mathcal{L}_3$ considered in our scheme as follows.

- (*Leakage function $\mathcal{L}_1$*). Given a graph G, $\mathcal{L}_1(G) = \{n, m, \#\mathsf{A}_G\}$, where n is the total number of vertices, m is the total number of edges in the graph G and $\#\mathsf{A}_G$ denotes the number of entries (*i.e.*, $m + z$) in the array A_G.
- (*Leakage function $\mathcal{L}_2$*). Given a graph G, a query q, $\mathcal{L}_2(G, q) = \{\mathsf{QP}(G, q), \mathsf{AP}(G, q)\}$, where $\mathsf{QP}(G, q)$ denotes the query pattern and $\mathsf{AP}(G, q)$ denotes the access pattern, both of which are given in the following definitions.

- (*Leakage function* $\mathcal{L}_3$). Given a graph G, an update object u, $\mathcal{L}_3(G, u) = \{\mathsf{id}_v, \mathsf{id}_{new}, \mathsf{next}\}$ is for add updates, and $\mathcal{L}_3(G, u) = \{\mathsf{id}_{del}, \mathsf{next}, \mathsf{prev}\}$ is for delete updates, where id_v denotes the identifier of the start vertex in the newly edge, id_{new} and id_{del} denote the identifiers of the edges to be added and deleted, respectively. prev and next contain the neighboring information (*i.e.*, the identifiers of the neighboring edges) of the edge to be updated. If there are no nodes in A_G before and after the edge to be updated then prev and next are set to $\perp$.

Definition 3. *(Query Pattern). For two shortest distance queries $q = (s, t), q' = (s', t')$, define $\mathsf{sim}(q, q') = (s = s', s = t', t = s', t = t')$, i.e., whether each of the vertices in q matches each of the vertices in q'. Let $\mathsf{q} = (q_1, \ldots, q_\delta)$ be a sequence of δ queries, the query pattern $\mathsf{QP}(G, q)$ induced by q is a $\delta \times \delta$ symmetric matrix such that for $1 \leq i, j \leq \delta$, the element in the i^{th} row and j^{th} column equals $\mathsf{sim}(q_i, q_j)$. Namely, the query pattern reveals whether the vertices in the query have appeared before.*

Definition 4. *(Access Pattern). Given a shortest distance query q for the graph G, the access pattern is defined as $\mathsf{AP}(G, q) = \{\mathsf{id}(c_q), \mathsf{id}(c_q)', \mathsf{id}^*(c_q)\}$, where $\mathsf{id}(c_q)$ denotes the identifiers of vertices in the encrypted result c_q, $\mathsf{id}(c_q)'$ denotes the identifiers of vertices contained in the dictionary path and it reveals the subgraph consisting of vertices reachable from the source ($\mathsf{id}(c_q) \subset \mathsf{id}(c_q)'$), and $\mathsf{id}^*(c_q)$ denotes the identifiers of the edges with one of its endpoints is the head node of retrieved adjacency lists.*

Discussion. Trading security for efficiency is a common practice in SSE designs. We follow the state-of-the-art of SSE solutions [9, 8, 28, 27, 7, 47, 20] to allow limited information to be revealed for higher efficiency. In practice, the query pattern implies whether a new query has been issued before, and the access pattern discloses the structural information such as graph connectivity associated with the query. The leakage is not revealed unless its corresponding query has been issued. This is similar to keyword-based SSE schemes, where the leakage (*i.e.*, patterns associated with a keyword query) is revealed only if the corresponding keyword is searched. Fortunately, we can guarantee some level of privacy to the original graph (*e.g.*, structural information) with slightly lower efficiency in our setting. To be specific, we can add some form of noise (*i.e.*, padding carefully designed fake entries [9, 8, 28] to each original adjacency list) during the generation of the encrypted graph. Hence, some leakage information (*i.e.*, the accurate number of edges in the original graph and the adjacency information of each vertex) can be mitigated. Moreover, in various application scenarios where the data may be abstracted and modeled as sparse graphs (see Table 2), the leakage would not be a big problem. Fully protecting the above two patterns (also forward privacy defined in [47]) without using expensive ORAM techniques remains an open challenging problem, which is our future research focus.

Dataset	Type	Vertices	Edges	Storage
Talk	directed	2,394,385	5,021,410	63.3 MB
Youtube	undirected	1,134,890	2,987,624	36.9 MB
EuAll	directed	265,214	420,045	4.76 MB
Gowalla	undirected	196,591	1,900,654	21.1 MB
Vote	directed	7,115	103,689	1.04 MB
Enron	undirected	36,692	367,662	3.86 MB

Table 2: The characteristics of datasets.

Dataset	Time (min.)	Storage (MB)			
		T_G	T_D	A_G	Total
Talk	1042.1	3.6	172.3	1460.5	1636.4
Youtube	460.6	8.93	102	874	984.93
EuAll	76.8	5.37	14.4	122	141.77
Gowalla	307.77	4.69	65.24	556.42	626.35
Vote	17.8	0.14	3.55	30.3	33.99
Enron	69.4	0.88	12.6	107	120.48

Table 3: The cost of initialization phase.

Theorem 1. *If Paillier cryptosystem is CPA-secure and P, F and G are pseudo-random, then the encrypted graph query database system is adaptively $(\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3)$-semantically secure in the random oracle model.*

Due to the space limitation, please refer to the Appendix for the proof Theorem 1.

6   Experimental Evaluation

In this section, we present experimental evaluations of our dynamic graph encryption scheme on a number of large-scale graphs. The experiments are performed on separate machines with different configurations. Concretely, the client runs on a machine with an Intel Core CPU with 4-core operating at 2.90GHz and equipped with 12GB RAM, and runs a Windows 10 operating system. Both the server and the proxy run on machines with an Intel Xeon CPU with 24-core operating at 2.10GHz and equipped with 128GB RAM and running Linux. We implemented algorithms described in Section 4 in Java, used HMAC for PRF/PRPs and instantiated the random oracle with HMAC-SHA-256 (both are contained in the default Java library). Our secure comparison protocol is built on top of FastGC [23], a Java-based open-source framework that enables developers to define arbitrary circuits. Several optimizations (*i.e.*, free-XOR technique, reduction of garbled tables and OT extension) discussed in the previous sections have been provided by the framework.

Our implementation used the following parameters: in the initialization phase, we use Paillier cryptosystem with a 1024-bit modulus. In the secure comparison protocol, the bit length allocated for the diameter l is 16 and the bit length of each random mask is 32. Besides, the FastGC framework provides a 80-bit security level; namely, it uses 80-bit wire labels for garbled circuits and security parameter $c = 80$ for the OT extension.

6.1   Datasets

We used real-world graph datasets publicly available from the Stanford SNAP website (available at https://snap.stanford.edu/data/), and selected the following six representative datasets including both directed and undirected graphs, with the scale ranging from thousands to millions of vertices and edges: *wiki-Talk*, a large network extracted from all user talk pages; *com-Youtube*, a large social network based on the Youtube web site; *email-EuAll*, an email network

Phase	Offline		Online	
Time/Bandwidth	s	KB	s	KB
Garbled circuits	1.26	3357.47	0.66	None
Oblivious transfer	0.454	21.91	1.23	1943.34

Table 4: The cost of secure comparison protocol.

generated from a European research institution; *loc-Gowalla*, a location-based social network; *wiki-Vote*, a network that contains all the Wikipedia voting data; and *email-Enron*, an email communication network. Table 2 summarizes the main characteristics of these datasets.

6.2 Experimental results

Table 3 shows the performance of the initialization phase (one-time cost) including the time to setup encrypted graphs as well as the corresponding storage cost. As can be seen, the time to encrypt a graph ranges from a few minutes to several hours which is practical. For example, it takes only 17.4 hours to obtain an encryption of the *wiki-Talk* graph including 2.4 million vertices and 5.1 million edges. Besides, we note that this phase is highly-parallelizable; namely, we bring the setup time down to just over 30 minutes by utilizing a modest cluster of 32 nodes. Furthermore, the storage cost of an encrypted graph is dominated by A_G with the total size ranging from 33.99MB for *wiki-Vote* to 1.60GB for *wiki-Talk*. We also note that our construction has less storage space requirements compared to Meng *et al.* [38] (*e.g.*, 2.07GB for *com-Youtube* in [38], whereas our scheme takes 984.93MB).

We first measured the time to query an encrypted graph without query history stored on the server. To simulate realistic queries that work in a similar manner with [20], we choose the query vertices in a random fashion weighted according to their outdegrees; that is, the probability of being selected grows with the number of outdegrees. The average time at the server (taken over 1,000 random queries) is given in Figure 4(a) for all encrypted graphs. In general, the results show that the query time ranges from 20.4s for *wiki-Vote* to 46.4 minutes for *wiki-Talk*. The computation at the proxy side in the shortest distance query phase mainly consists of the homomorphic decryptions and the comparison circuit evaluations. Additionally, we can obtain an order-of-magnitude improvement in both computation time and bandwidth by using the packing optimization presented in Section 4.3. The actual time for the client to generate the token and decrypt the encrypted result per each query is always less than 0.1s which is very fast. In addition, about 1.5KB communication overhead is required to transfer the token and the encrypted result for each query.

Next, the performance of the query phase with the help of history stored on the server is illustrated in Figure 4(b) and 4(c), and a block of 1,000 random executions results in one measurement point in both figures. In Figure 4(b), the y-axis represents the ratio of the average query time using history to that without using history. Generally, it reflects that the average query time decreases

(a) Query without histo-ry

(b) Query with history

(c) Query with history

Fig. 4: The cost of distance query phase.

with the increase of the number of queries, because subsequent queries can first be answered by leveraging the history. Furthermore, as can be seen, after 10,000 queries, it obtains about 86% reduction of the query time for *wiki-Vote* compared to that without using history, *i.e.*, it only needs roughly 2.9s to answer a subsequent query. Figure 4(c) demonstrates the increasing size of history (instantiated by HashMap in our implementation) with the increasing amount of total shortest distance queries. In particular, for example, the additional storage cost for history at the server side is about 650MB for *wiki-Talk* after 10,000 runs.

For clarity, we run the secure comparison protocol for 1,000 times to demonstrate its performance. The results are summarized in Table 4. At the server side, the time cost of generating garbled circuits for 1,000 comparisons is only 1.26s and about 3357.47KB communication overhead to transfer the circuits. For the proxy side, it only needs 0.66s to evaluate the circuits with no bandwidth cost. Since we adopt the OT extension optimization, only a constant number (*i.e.*, 80) of OTs are required in the offline stage. In particular, the time cost and communication cost in this stage are 0.454s and 21.91KB, respectively. During the online stage, the server transmits the wire labels representing its inputs to the proxy, and the online part of OT protocol is executed by the proxy to obtain its garbled inputs. Hence, the total time consumption and communication overhead of this stage are 1.23s and 1943.34KB, respectively.

Figure 5 shows the execution time (averaged over 1,000 runs) for adding and deleting an edge over all the encrypted graphs. Obviously, both addition and deletion operations are practically efficient and independent of the scale of the graphs. As shown in Figure 5(a), the time cost at the client side is dominat-
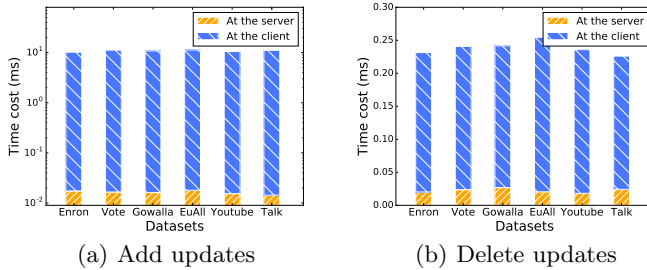

(a) Add updates

(b) Delete updates

Fig. 5: The time cost of dynamic updates.

ed by generating an encryption of the length of the edge to be updated (roughly 10ms), while the server side has a negligible running time. Similar results can be obtained in Figure 5(b) for the delete updates. It only needs about 0.25ms to delete a specified edge, and the time to generate the delete token at the client side dominates the time cost of the entire process. In addition, about 0.3KB and tens of bytes are consumed when performing adding and deleting operations, respectively.

7   Conclusion

In this paper, we designed a new graph encryption scheme–SecGDB to encrypt graph structures and enforce private graph queries. In our construction, we used additively homomorphic encryption and garbled circuits to support shortest distance queries with optimal time and storage complexities. On top of this, we further proposed an auxiliary data structure called *query history* stored on the remote server to achieve better amortized time complexity over multiple queries. Compared to the state-of-the-art, SecGDB returns the exact distance results and allows efficient graph updates over large-scale encrypted graph database. SecGDB is proven to be adaptively semantically-secure in the random oracle model. We finally evaluated SecGDB on representative real-world datasets, showing its efficiency and practicality for use in real-world applications.

Acknowledgment

References

1. A. Aly, E. Cuvelier, S. Mawet, O. Pereira, and M. Van Vyve, "Securely solving simple combinatorial graph problems," in *Proc. of FC'13*.   Springer, 2013, pp. 239–257.
2. D. Beaver, "Precomputing oblivious transfer," in *Proc. of CRYPTO'95*.   Springer, 1995, pp. 97–109.
3. M. Blanton, A. Steele, and M. Alisagari, "Data-oblivious graph algorithms for secure computation and outsourcing," in *Proc. of ASIACCS'13*.   ACM, 2013, pp. 207–218.
4. D. Boneh, C. Gentry, S. Halevi, F. Wang, and D. J. Wu, "Private database queries using somewhat homomorphic encryption," in *Proc. of ACNS'13*.   Springer, 2013, pp. 102–118.

5. R. Bost, R. A. Popa, S. Tu, and S. Goldwasser, "Machine learning classification over encrypted data," in *Proc. of NDSS'15*, 2015.
6. Z. Brakerski and V. Vaikuntanathan, "Fully homomorphic encryption from ring-lwe and security for key dependent messages," in *Proc. of CRYPTO'11*. Springer, 2011, pp. 505–524.
7. D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Roşu, and M. Steiner, "Highly-scalable searchable symmetric encryption with support for boolean queries," in *Proc. of CRYPTO'13*. Springer, 2013, pp. 353–373.
8. M. Chase and S. Kamara, "Structured encryption and controlled disclosure," in *Proc. of ASIACRYPT'10*. Springer, 2010, pp. 577–594.
9. R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky, "Searchable symmetric encryption: improved definitions and efficient constructions," in *Proc. of CCS'06*. ACM, 2006, pp. 79–88.
10. I. Damgard, M. Geisler, and M. Kroigard, "Homomorphic encryption and secure comparison," *International Journal of Applied Cryptography*, vol. 1, no. 1, pp. 22–31, 2008.
11. ——, "A correction to'efficient and secure comparison for on-line auctions'," *International Journal of Applied Cryptography*, vol. 1, no. 4, pp. 323–324, 2009.
12. E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
13. Y. Elmehdwi, B. K. Samanthula, and W. Jiang, "Secure k-nearest neighbor query over encrypted data in outsourced environments," in *Proc. of ICDE'14*. IEEE, 2014, pp. 664–675.
14. M. L. Fredman and R. E. Tarjan, "Fibonacci heaps and their uses in improved network optimization algorithms," *JACM*, vol. 34, no. 3, pp. 596–615, 1987.
15. C. Gentry, "Fully homomorphic encryption using ideal lattices." in *Proc. of S-TOC'09*. ACM, 2009, pp. 169–178.
16. E.-J. Goh *et al.*, "Secure indexes," *IACR Cryptology ePrint Archive*, vol. 2003, p. 216, 2003.
17. O. Goldreich and R. Ostrovsky, "Software protection and simulation on oblivious rams," *JACM*, vol. 43, no. 3, pp. 431–473, 1996.
18. S. Goldwasser and S. Micali, "Probabilistic encryption," *Journal of computer and system sciences*, vol. 28, no. 2, pp. 270–299, 1984.
19. P. Golle, J. Staddon, and B. Waters, "Secure conjunctive keyword search over encrypted data," in *Proc. of ACNS'04*. Springer, 2004, pp. 31–45.
20. F. Hahn and F. Kerschbaum, "Searchable encryption with secure and efficient updates," in *Proc. of CCS'14*. ACM, 2014, pp. 310–320.
21. W.-S. Han, S. Lee, K. Park, J.-H. Lee, M.-S. Kim, J. Kim, and H. Yu, "Turbograph: a fast parallel graph engine handling billion-scale graphs in a single pc," in *Proc. of SIGKDD'13*. ACM, 2013, pp. 77–85.
22. F. Harary, "Graph theory," *Westview Press*, 1969.
23. Y. Huang, D. Evans, J. Katz, and L. Malka, "Faster secure two-party computation using garbled circuits," in *Proc. of USENIX Security'11*. USENIX, 2011.
24. Y. Huang, L. Malka, D. Evans, and J. Katz, "Efficient privacy-preserving biometric identification," in *Proc. of NDSS'11*, 2011, pp. 250–267.
25. Y. Ishai, J. Kilian, K. Nissim, and E. Petrank, "Extending oblivious transfers efficiently," in *Proc. of CRYPTO'03*. Springer, 2003, pp. 145–161.
26. S. Jarecki and V. Shmatikov, "Efficient two-party secure computation on committed inputs," in *Proc. of EUROCRYPT'07*. Springer, 2007, pp. 97–114.
27. S. Kamara and C. Papamanthou, "Parallel and dynamic searchable symmetric encryption," in *Proc. of FC'13*. Springer, 2013, pp. 258–274.

28. S. Kamara, C. Papamanthou, and T. Roeder, "Dynamic searchable symmetric encryption," in *Proc. of CCS'12*. ACM, 2012, pp. 965–976.
29. J. Katz and Y. Lindell, *Introduction to modern cryptography*. CRC Press, 2014.
30. V. Kolesnikov, A.-R. Sadeghi, and T. Schneider, "Improved garbled circuit building blocks and applications to auctions and computing minima," in *Proc. of CNS'09*. Springer, 2009, pp. 1–20.
31. V. Kolesnikov and T. Schneider, "Improved garbled circuit: Free xor gates and applications," in *Proc. of ICALP'08*. Springer, 2008, pp. 486–498.
32. Y. Lindell and B. Pinkas, "An efficient protocol for secure two-party computation in the presence of malicious adversaries," in *Proc. of EUROCRYPT'07*. Springer, 2007, pp. 52–78.
33. ——, "A proof of security of yao's protocol for two-party computation," *Journal of Cryptology*, vol. 22, no. 2, pp. 161–188, 2009.
34. C. Liu, X. S. Wang, K. Nayak, Y. Huang, and E. Shi, "Oblivm: A programming framework for secure computation," in *Proc. of S&P'15*. IEEE, 2015, pp. 359–376.
35. Y. Low, D. Bickson, J. Gonzalez, C. Guestrin, A. Kyrola, and J. M. Hellerstein, "Distributed graphlab: a framework for machine learning and data mining in the cloud," *PVLDB*, vol. 5, no. 8, pp. 716–727, 2012.
36. G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski, "Pregel: a system for large-scale graph processing," in *Proc. of SIGMOD'10*. ACM, 2010, pp. 135–146.
37. D. Malkhi, N. Nisan, B. Pinkas, Y. Sella *et al.*, "Fairplay-secure two-party computation system." in *Proc. of USENIX Security'04*. USENIX, 2004, pp. 287–302.
38. X. Meng, S. Kamara, K. Nissim, and G. Kollios, "Grecs: graph encryption for approximate shortest distance queries," in *Proc. of CCS'15*. ACM, 2015, pp. 504–517.
39. M. Naor and B. Pinkas, "Efficient oblivious transfer protocols," in *Proc. of SODA'01*. SIAM, 2001, pp. 448–457.
40. J. B. Nielsen and C. Orlandi, "Lego for two-party secure computation," in *Proc. of TC'09*. Springer, 2009, pp. 368–386.
41. V. Nikolaenko, U. Weinsberg, S. Ioannidis, M. Joye, D. Boneh, and N. Taft, "Privacy-preserving ridge regression on hundreds of millions of records," in *Proc. of S&P'13*. IEEE, 2013, pp. 334–348.
42. P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proc. of EUROCRYPT'99*. Springer, 1999, pp. 223–238.
43. B. Pinkas, T. Schneider, N. P. Smart, and S. C. Williams, "Secure two-party computation is practical," in *Proc. of AISACRYPT'09*. Springer, 2009, pp. 250–267.
44. R. L. Rivest, L. Adleman, and M. L. Dertouzos, "On data banks and privacy homomorphisms," *Foundations of Secure Computation*, vol. 4, no. 11, pp. 169–180, 1978.
45. M. Sarwat, S. Elnikety, Y. He, and G. Kliot, "Horton: Online query execution engine for large distributed graphs," in *Proc. of ICDE'12*. IEEE, 2012, pp. 1289–1292.
46. D. X. Song, D. Wagner, and A. Perrig, "Practical techniques for searches on encrypted data," in *Proc. of S&P'00*. IEEE, 2000, pp. 44–55.
47. E. Stefanov, C. Papamanthou, and E. Shi, "Practical dynamic searchable encryption with small leakage," in *Proc. of NDSS'14*, 2014.
48. X. S. Wang, K. Nayak, C. Liu, T. Chan, E. Shi, E. Stefanov, and Y. Huang, "Oblivious data structures," in *Proc. of CCS'14*. ACM, 2014, pp. 215–226.
49. D. J. Wu, J. Zimmerman, J. Planul, and J. C. Mitchell, "Privacy-preserving shortest path computation," in *Proc. of NDSS'16*, 2016.

50. A. Yao, "Protocols for secure computations," in *Proc. of FOCS'82*.   IEEE, 1982, pp. 160–164.

51. ——, "How to generate and exchange secrets," in *Proc. of FOCS'86*.   IEEE, 1986, pp. 162–167.

Appendix

Proof Sketch: We describe a polynomial time simulator $\mathcal{S}$ such that for any PPT adversary $\mathcal{A}$, the outputs of $\textbf{Real}_{\mathcal{A}}(\lambda)$ and $\textbf{Ideal}_{\mathcal{A},\mathcal{S}}(\lambda)$ are computationally indistinguishable. Consider the simulator $\mathcal{S}$ that works as follows.

[Setup] Given $\mathcal{L}_1(G) = \{n, m, \#\mathsf{A}_G\}$, $\mathcal{S}$ constructs the simulated graph $\widetilde{\Omega}_G$ as below. To simulate T_G, it generates a dictionary $\widetilde{\mathsf{T}}_G$ of size $n + 1$ and for all $i \in [n + 1]$, stores a random $(\log \#\mathsf{A}_G + 2\lambda)$-bit string $\widetilde{\nu}_i$ in $\widetilde{\mathsf{T}}_G$ under a random λ-bit key $\widetilde{\kappa}_i$. To simulate A_G, it generates an array $\widetilde{\mathsf{A}}_G$ of size $m + z$ and fills m of these cells (chosen at random) with random strings of size $(2 \log \#\mathsf{A}_G + 4\lambda + \ell)$-bit, where ℓ denotes the bit length of ciphertext in Paillier cryptosystem, the rest cells are marked as `free`. To simulate T_D, it generates a dictionary $\widetilde{\mathsf{T}}_D$ of size m and for all $i \in [m]$, adds a random $(3 \log \#\mathsf{A}_G + 2\lambda)$-bit string $\widetilde{\nu}'_i$ associated with a random λ-bit key $\widetilde{\kappa}'_i$. Besides, let RO be an empty dictionary. Finally, it outputs the simulated graph $\widetilde{\Omega}_G = (\widetilde{\mathsf{T}}_G, \widetilde{\mathsf{A}}_G, \widetilde{\mathsf{T}}_D)$.

[Simulating DistanceQuery] Given $\mathcal{L}_2(G, q) = \{\mathsf{QP}(G, q), \mathsf{AP}(G, q)\}$ to simulate query token τ_q, $\mathcal{S}$ first checks if either of the query vertex s or t has appeared before. If s appeared previously, $\mathcal{S}$ sets $\widetilde{\tau}_1$ and $\widetilde{\tau}_3$ to the values that were previously used. Otherwise, $\mathcal{S}$ sets $\widetilde{\tau}_1 := \widetilde{\kappa}_i$ for some previously unused $\widetilde{\kappa}_i$ and $\widetilde{\tau}_3$ as follows. It chooses a previously unused cell in $\widetilde{\mathsf{A}}_G$ at random with its location `addr`, a random λ-bit string `str` (marked with the identifier in $\mathsf{id}^*(c_q)$), a random λ-bit string K, and sets $\widetilde{\tau}_3 := \widetilde{\mathsf{T}}_G[\widetilde{\kappa}_i] \oplus \langle \mathtt{addr}, \mathtt{str}, K \rangle$. It then records the association between K and the related adjacency list (*i.e.*, marks corresponding cells in $\widetilde{\mathsf{A}}_G$ with the related identifiers in $\mathsf{id}(c_q)'$). In addition, $\mathcal{S}$ does analogously for the remaining identifiers appeared in $\mathsf{id}(c_q)$. For the query vertex t, if it appeared previously, $\mathcal{S}$ sets $\widetilde{\tau}_2$ to the value that were previously used, otherwise sets $\widetilde{\tau}_2 := \widetilde{\kappa}_i$ for some previously unused $\widetilde{\kappa}_i$. Finally, $\mathcal{S}$ outputs the simulated token $\widetilde{\tau}_q = (\widetilde{\tau}_1, \widetilde{\tau}_2, \widetilde{\tau}_3)$.

[Simulating UpdateQuery] For add updates, $\mathcal{S}$ is given leakage $\mathcal{L}_3(G, u) = \{\mathsf{id}_v, \mathsf{id}_{new}, \mathsf{next}\}$ to simulate the update token τ_u. If the id_{new} has been added in the past, it just sets $(\widetilde{\tau}_1, \widetilde{\tau}_2, \widetilde{\tau}_3, \widetilde{\tau}_4)$ that were previously used. On the other hand, if id_v has appeared before, it sets $\widetilde{\tau}_1, \widetilde{\tau}_2$ that were previously used, else it first sets $\widetilde{\tau}_1 := \widetilde{\kappa}_i$ for some previously unused $\widetilde{\kappa}_i$, then it chooses a previously unused cell in $\widetilde{\mathsf{A}}_G$ at random with its location $\mathtt{addr}_1$, a random λ-bit string $\mathtt{str}_1$ (marked with the identifier in next), and sets $\widetilde{\tau}_2 := \widetilde{\mathsf{T}}_G[\widetilde{\kappa}_i]^{|\langle \mathtt{addr}_1, \mathtt{str}_1 \rangle|} \oplus \langle \mathtt{addr}_1, \mathtt{str}_1 \rangle$. In the following, $\mathcal{S}$ simulates $\widetilde{\tau}_3 := \widetilde{\kappa}'_i$, where $\widetilde{\kappa}'_i$ is a random λ-bit string marked with id_{new}, simulates $\widetilde{\tau}_4 := \widetilde{\nu}'_i$, where $\widetilde{\nu}'_i$ is a random $(3 \log \#\mathsf{A}_G + 2\lambda)$-bit string associated with $\widetilde{\kappa}'_i$. $\mathcal{S}$ finally samples a $(2 \log \#\mathsf{A}_G + 4\lambda + \ell)$-bit string at random as $\widetilde{\tau}_5$ and outputs the simulated token $\widetilde{\tau}_u = (\widetilde{\tau}_1, \widetilde{\tau}_2, \widetilde{\tau}_3, \widetilde{\tau}_4, \widetilde{\tau}_5)$. Similarly, $\mathcal{S}$

simulates the token $\widetilde{\tau}_u$ for delete updates.

[Answering H queries] Given query (K, r), $\mathcal{S}$ checks if this query was submitted before, if this is the case, it returns $\rho := \mathsf{RO}[\langle K, r \rangle]$ immediately. Otherwise, $\mathcal{S}$ checks if K has been associated with an adjacency list, if so, it finds all the related entries in $\widetilde{\mathsf{A}}_G$, and parses each entry as the form $\langle \mathsf{N}', r \rangle$, then it returns $\mathsf{N}' \oplus \langle \alpha_1, \alpha_2, \alpha_3, \alpha_4 \rangle$, where α_1 is the identifier of the entry associated with K, α_2 is the string matched with α_1, α_3 is a Paillier encryption of 0 and α_4 is an unused address in $\widetilde{\mathsf{A}}_G$ chosen at random or $\mathtt{NULL}$ (the last entry). If not, it returns a random $|\rho|$-bit string and stores it in RO under the key $\langle K, r \rangle$ to stay consistent on future queries.

In summary, the indistinguishability of $\widetilde{\Omega}_G$ from Ω_G follows from the pseudo-randomness of P, F and G and the CPA-security of Paillier cryptosystem. The indistinguishability of $\widetilde{\tau}_q$ follows from the pseudo-randomness of P and F and that of $\widetilde{\tau}_u$ from the pseudo-randomness of P, F and the CPA-security of Paillier cryptosystem.