

Privacy-Preserving Shortest Path Computation

(Extended Version)

David J. Wu Joe Zimmerman J  r  my Planul John C. Mitchell

Stanford University

{dwu4, jzim, mitchell}@cs.stanford.edu, jeremy.planul@ens-lyon.org

Abstract

Navigation is one of the most popular cloud computing services. But in virtually all cloud-based navigation systems, the client must reveal her location and destination to the cloud service provider in order to learn the fastest route. In this work, we present a cryptographic protocol for navigation on city streets that provides privacy for both the client’s location and the service provider’s routing data. Our key ingredient is a novel method for compressing the next-hop routing matrices in networks such as city street maps. Applying our compression method to the map of Los Angeles, for example, we achieve over tenfold reduction in the representation size. In conjunction with other cryptographic techniques, this compressed representation results in an efficient protocol suitable for fully-private real-time navigation on city streets. We demonstrate the practicality of our protocol by benchmarking it on real street map data for major cities such as San Francisco and Washington, D.C.

1 Introduction

Location privacy is a major concern among smartphone users, and there have been numerous controversies due to companies tracking users’ locations [AVD11, Che11]. Among the various applications that require location information, navigation is one of the most popular. For example, companies such as Google, Apple, and Waze have built traffic-aware navigation apps to provide users with the most up-to-date routing information. But to use these services, users must reveal their location and destination to the cloud service provider. In doing so, they may also reveal other sensitive information about their personal lives, such as their health condition, their social and political affiliations, and more.

One way to provide location privacy is for the user to download the entire map from the cloud service provider and then compute the best route locally on her own mobile device. Unfortunately, since service providers invest significant resources to maintain up-to-date routing information, they are not incentivized to publish their entire routing database in real-time. Even in the case of a paid premium service, in which the service provider does not derive compensation from learning the user’s location data, it is not obvious how to achieve fully-private navigation. The user does not trust the cloud provider with her location data, and the cloud provider does not trust the user with its up-to-date routing information, so neither party has

This is the extended version of a paper by the same name that appeared at the *Network and Distributed System Security Symposium (NDSS)* in February, 2016. Permission to freely reproduce all or part of this paper for noncommercial purposes is granted provided that copies bear this notice and the full citation on the first page. Reproduction for commercial purposes is strictly prohibited without the prior written consent of the Internet Society, the first-named author (for reproduction of an entire paper only), and the author’s employer if the paper was prepared within the scope of employment.

NDSS ’16, 21-24 February 2016, San Diego, CA, USA

Copyright 2016 Internet Society, ISBN 1-891562-41-X

<http://dx.doi.org/10.14722/ndss.2016.23052>

all of the data to perform the computation. While general-purpose cryptographic tools such as multiparty computation solve this problem in theory (see Section 7), these protocols are prohibitively expensive in practice for applications such as real-time navigation.

Our results. In this work, we present an efficient cryptographic protocol for *fully-private* navigation: the user keeps private her location and destination, and the service provider keeps private its proprietary routing information (except for the routing information associated with the specific path requested by the user and a few generic parameters pertaining to the network). We give a complete implementation of our protocol and benchmark its performance on real street map data (Section 5.3). Since our protocol is real-time (the user continues receiving directions throughout the route), we benchmark the performance “per hop”, where each hop roughly corresponds to an intersection between streets.¹ For cities such as San Francisco and Washington, D.C., each hop in our protocol requires about 1.5 seconds and less than 100 KB of bandwidth. In addition, before the protocol begins, we execute a preprocessing step that requires bandwidth in the tens of megabytes. Since this preprocessing step can be performed at any time, in practice it would likely be run via a fast Wi-Fi connection, before the mobile user needs the real-time navigation service, and thus the additional cost is very modest. To our knowledge, ours is the first fully-private navigation protocol efficient enough to be feasible in practice.

Our technical contributions. In our work, we model street-map networks as graphs, in which the nodes correspond to street intersections, and edges correspond to streets. In our model, we assume that the network topology is public (i.e., in the case of navigation on city streets, the layout of the streets is publicly known). However, only the service provider knows the up-to-date traffic conditions, and thus the shortest path information. In this case, the server’s “routing information” consists of the weights (that is, travel times) on the edges in the network.

By modeling street-maps as graphs, we can easily construct a straw-man private navigation protocol based on symmetric private information retrieval (SPIR) [GIKM00, KO97, NP05]. Given a graph \mathcal{G} with n nodes, the server first constructs a database with n^2 records, each indexed by a source-destination pair (s, t) . The record indexed (s, t) contains the shortest path from s to t . To learn the shortest path from s to t , the client engages in SPIR with the server for the record indexed (s, t) . Security of SPIR implies that the client just learns the shortest path and the server learns nothing. While this method satisfies the basic security requirements, its complexity scales quadratically in the number of nodes in the graph. Due to the computational cost of SPIR, this solution quickly becomes infeasible in the size of the graph.

Instead, we propose a novel method to compress the routing information in street-map networks. Specifically, given a graph \mathcal{G} with n nodes, we define the next-hop routing matrix $M \in \mathbb{Z}^{n \times n}$ for \mathcal{G} to be the matrix where each entry M_{st} gives the index of the first node on the shortest path from node s to node t . To apply our compression method, we first preprocess the graph (Section 3) such that each entry in the next-hop routing matrix M can be specified by two bits: $M_{st} = (M_{st}^{(\text{NE})}, M_{st}^{(\text{NW})})$ where $M^{(\text{NE})}, M^{(\text{NW})} \in \{-1, 1\}^{n \times n}$. We then compress $M^{(\text{NE})}$ by computing a *sign-preserving decomposition*: two matrices $A^{(\text{NE})}, B^{(\text{NE})} \in \mathbb{Z}^{n \times d}$ where $d \ll n$ such that $M^{(\text{NE})} = \text{sign}(A^{(\text{NE})} \cdot (B^{(\text{NE})})^T)$. We apply the same procedure to compress the other component $M^{(\text{NW})}$. The resulting compression is lossless, so there is no loss in accuracy in the shortest paths after applying our transformation. When applied to the road network for the city of Los Angeles, we obtain over 10x reduction in the size of the representation. Our compression method is highly parallelizable and by running our computation on GPUs, we can compress next-hop matrices with close to 50 million elements (for a 7000-node network) in under ten minutes.

¹In a few cases, hops in our construction occur mid-street or in instances such as traffic circles. These are rare enough that even in large cities such as Los Angeles, the total number of hops along any route is less than 200.

Moreover, our compression method enables an efficient protocol for a fully-private shortest path computation. In our protocol, the rounds of interaction correspond to the nodes in the shortest path. On each iteration of the protocol, the client learns the next hop on the shortest path to its requested destination. Abstractly, if the client is currently at a node s and navigating to a destination t , then after one round of the protocol execution, the client should learn the next hop given by $M_{st} = (M_{st}^{(\text{NE})}, M_{st}^{(\text{NW})})$. Each round of our protocol thus reduces to a two-party computation of the components $M_{st}^{(\text{NE})}$ and $M_{st}^{(\text{NW})}$. Given our compressed representation of the next-hop routing matrices, computing $M_{st}^{(\text{NE})}$ reduces to computing the sign of the inner product between the s^{th} row of $A^{(\text{NE})}$ and the t^{th} row of $B^{(\text{NE})}$, and similarly for $M_{st}^{(\text{NW})}$. In our construction, we give an efficient method for inner product evaluation based on affine encodings, and use Yao’s garbled circuits [Yao86, BHR12] to evaluate the sign function. An important component of our protocol design is a novel way of efficiently combining affine encodings and garbled circuits. Together, these methods enable us to construct an efficient, fully-private navigation protocol.

Other approaches. An alternative method for private navigation is to use generic tools for two-party computation such as Yao’s garbled circuits [Yao86, BHR12] and Oblivious RAM (ORAM) [GO96, SvDS⁺13]. While these approaches are versatile, they are often prohibitively expensive for city-scale networks (in the case of Yao circuits), or do not provide strong security guarantees against malicious clients (in the case of ORAM). For instance, the garbled-circuit approach by Carter et al. [CMTB13, CLT14] requires several minutes of computation to answer a single shortest path query in a road network with just 100 nodes. Another generic approach combining garbled circuits and ORAM [LWN⁺15] requires communication on the order of GB and run-times ranging from tens of minutes to several hours for a single query on a network with 1024 nodes. Thus, current state-of-the-art tools for general two-party computation do not give a viable solution for private navigation in city-scale networks. We survey other related methods in Section 7.

2 Preliminaries and Threat Model

We begin with some notation. For a positive integer n , let $[n]$ denote the set of integers $\{1, \dots, n\}$. For two ℓ -bit strings $x, y \in \{0, 1\}^\ell$, we write $x \oplus y$ to denote their bitwise XOR. For a prime p , we write \mathbb{F}_p to denote the finite field with p elements, and \mathbb{F}_p^* to denote its multiplicative group. Let \mathcal{D} be a probability distribution. We write $x \leftarrow \mathcal{D}$ to denote that x is drawn from \mathcal{D} . Similarly, for a finite set S we write $x \stackrel{\text{R}}{\leftarrow} S$ to denote that x is drawn uniformly at random from S . A function $f(\lambda)$ is negligible in a security parameter λ if $f = o(1/\lambda^c)$ for all $c \in \mathbb{N}$.

For two distribution ensembles $\{\mathcal{D}_1\}_\lambda, \{\mathcal{D}_2\}_\lambda$, we write $\{\mathcal{D}_1\}_\lambda \stackrel{c}{\approx} \{\mathcal{D}_2\}_\lambda$ to denote that $\{\mathcal{D}_1\}_\lambda$ and $\{\mathcal{D}_2\}_\lambda$ are computationally indistinguishable (i.e., no probabilistic polynomial-time algorithm can distinguish them, except with probability negligible in λ). We write $\{\mathcal{D}_1\}_\lambda \equiv \{\mathcal{D}_2\}_\lambda$ to denote that $\{\mathcal{D}_1\}_\lambda$ and $\{\mathcal{D}_2\}_\lambda$ are identically distributed for all λ . For a predicate $\mathcal{P}(x)$, we write $\mathbf{1}\{\mathcal{P}(x)\}$ to denote the indicator function for $\mathcal{P}(x)$, i.e., $\mathbf{1}\{\mathcal{P}(x)\} = 1$ if and only if $\mathcal{P}(x)$ is true, and otherwise, $\mathbf{1}\{\mathcal{P}(x)\} = 0$. If \mathcal{G} is a directed graph, we write (u, v) to denote the edge from node u to node v .

A function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ with key-space \mathcal{K} , domain \mathcal{X} , and range \mathcal{Y} is a PRF [GGM86] if no efficient adversary can distinguish outputs of the PRF (with key $k \stackrel{\text{R}}{\leftarrow} \mathcal{K}$, evaluated on inputs chosen adaptively by the adversary) from the corresponding outputs of a truly random function from $\mathcal{X} \rightarrow \mathcal{Y}$.

Threat model. We give a high-level survey of our desired security properties, and defer the details to Section 4.2. We operate in the two-party setting where both parties know the network topology as well as a few generic parameters about the underlying graph structure (described concretely in Section 4.2), but

only the server knows the weights (the routing information). The client holds a source-destination pair. At the end of the protocol execution, the client learns the shortest path between its requested source and destination, while the server learns nothing. The first property we require is privacy for the client’s location. Because of the sensitivity of location information, we require privacy to hold even against malicious servers, that is, servers whose behavior can deviate from the protocol specification.

The second requirement is privacy for the server’s routing information, which may contain proprietary or confidential information. The strongest notion we can impose is that at the end of the protocol execution, the client does not learn anything more about the graph other than the shortest path between its requested source and destination and some generic parameters associated with the underlying network. While this property is not difficult to achieve if the client is semi-honest (that is, the client adheres to the protocol specification), in practice there is little reason to assume that the client will behave this way. Thus, we aim to achieve security against malicious clients. In our setting, we will show that a malicious client learns only the shortest path from its requested source to its requested destination, except for failure events that occur with probability at most $\approx 2^{-30}$. For comparison, 2^{-30} is the probability that an adversary running in time $\approx 2^{50}$ is able to guess an 80-bit secret key.²

To summarize, we desire a protocol that provides privacy against a malicious server and security against a malicious client. We note that our protocol does not protect against the case of a server corrupting the map data; in practice, we assume that the map provider is trying to provide a useful service, and thus is not incentivized to provide misleading or false navigation information.

3 Graph Processing

As described in Section 1, we model street-map networks as directed graphs, where nodes correspond to intersections, and edges correspond to streets. To enable an efficient protocol for fully-private shortest path computation, we first develop an efficient method for preprocessing and compressing the routing information in the network. In this section, we first describe our preprocessing procedure, which consists of two steps: introducing dummy nodes to constrain the out-degree of the graph, and assigning a cardinal direction to each edge. Then, we describe our method for compressing the routing information in the graph; here, we exploit the geometric structure of the graph.

Bounding the out-degree. Let \mathcal{G} be the directed graph representing the road network. We assume that the nodes in \mathcal{G} have low out-degree. In a road network (see Figure 1 for an example), the nodes correspond to street intersections, and thus typically have at most four outgoing edges, one in each cardinal direction. In the first step of our preprocessing procedure, we take a weighted, directed graph \mathcal{G} and transform it into a weighted, directed graph \mathcal{G}' where each node has maximum out-degree 4. Specifically, we start by setting $\mathcal{G}' = \mathcal{G}$. Then, as long as there is a node $u \in \mathcal{G}'$ with neighbors v_1, \dots, v_ℓ and $\ell > 4$, we do the following. First, we add a new node u' to \mathcal{G}' . For $i \geq 4$, we add the edge (u', v_i) to \mathcal{G}' and remove the edge (u, v_i) from \mathcal{G}' . We also add a zero-weight edge from u to u' in \mathcal{G}' . By construction, this transformation preserves the shortest-path between nodes in \mathcal{G} and constrains the out-degree of all nodes in \mathcal{G}' to 4.

Orienting the edges. In a road network, we can associate each node by an (x, y) pair in the coordinate plane (for example, the node’s latitude and longitude). Consider a coordinate system aligned with the cardinal directions: the x -axis corresponds to the east-west axis and the y -axis corresponds to the north-south axis. Then, for each node u in the graph \mathcal{G} , we associate each of its neighbors v_i ($0 \leq i < 4$) with a direction $\text{dir}_i \in \{\text{N, E, S, W}\}$ (for north, east, south, west, respectively) relative to u . For a concrete example,

²Even in the case of these low-probability failure events, one can show that a malicious client only learns a bounded-length path emanating from its requested source, though it may not be a shortest path to any particular destination.

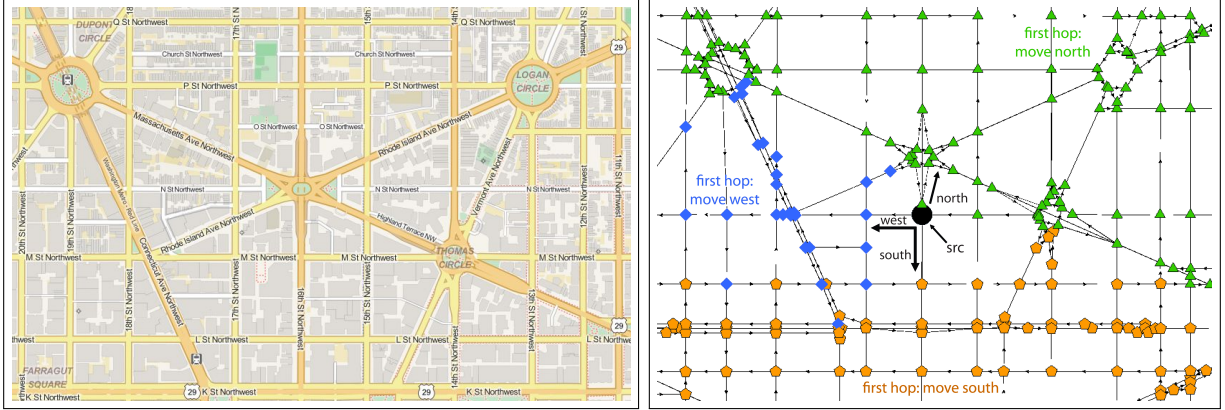


Figure 1: Subsection of map of Washington, D.C. from OpenStreetMap [Ope] (left) and visualization of the routing network after preprocessing (right). The visualization on the right shows the first hop of the shortest path from the source node (denoted by the circle) to all other nodes in the graph (denoted by a polygon). In this example, the source node has three neighbors: to the north, west, and south (as indicated in the diagram). If the first hop in the shortest path from the source to a node is to move north, then the node is represented by a green triangle. If the first hop is to move west, then the node is represented by a blue diamond, and if the first hop is to move south, then the node is represented by an orange pentagon.

refer to the visualization of the preprocessed graph in Figure 1. Here, the center node (labeled “src”) has three neighbors, each of which is associated with a cardinal direction: north, west, or south in this case. We define the orientation of an edge to be the direction associated with the edge.

To determine the orientation of the edges in \mathcal{G} , we proceed as follows. For each node $u \in \mathcal{G}$, we associate a *unique* direction $\text{dir}_i \in \{N, E, S, W\}$ with each neighbor v_i of u . In assigning the four cardinal directions to each node’s neighbors, we would like to approximate the true geographical locations of the nodes. In our setting, we formulate this assignment as a bipartite matching problem for each node u , with u ’s neighbors (at most 4) forming one partition of the graph, and the four cardinal directions $\{N, E, S, W\}$ forming the other. We define the cost of a matching between a neighbor v_i of u and a direction dir_j to be the angle formed by the vector from u to v_i and the unit vector aligned in the direction dir_j . In assigning directions to neighbors, we desire a matching that minimizes the costs of the matched neighbors. Such a matching can be computed efficiently using the Hungarian method [KY55]. In this way, we associate a cardinal direction with each edge in \mathcal{G} .

Compressing shortest paths. Next, we describe a method for compressing the next-hop routing matrix for a road network. Let \mathcal{G} be a directed graph with n nodes and maximum out-degree 4. Using the method described above, we associate a direction $\text{dir} \in \{N, E, S, W\}$ with each edge in \mathcal{G} . Since there are four possible values for dir , we can encode the direction using exactly two bits b_{NE} and b_{NW} , where $b_{NE} = 0$ if and only if $\text{dir} \in \{N, E\}$, and $b_{NW} = 0$ if and only if $\text{dir} \in \{N, W\}$. Intuitively, b_{NE} encodes the direction with respect to the northwest-southeast axis while b_{NW} encodes the direction with respect to the northeast-southwest axis. Thus, for each node $u \in \mathcal{G}$, we associate a unique two-bit index (b_{NE}, b_{NW}) with each of its outgoing edges. For notational convenience, we define a function IndexToDirection that maps an index (b_{NE}, b_{NW}) to the corresponding direction $\text{dir} \in \{N, E, S, W\}$. Specifically,

$$\begin{aligned} \text{IndexToDirection}(0, 0) &= N & \text{IndexToDirection}(0, 1) &= E \\ \text{IndexToDirection}(1, 0) &= W & \text{IndexToDirection}(1, 1) &= S. \end{aligned} \tag{1}$$

We next compute the shortest path p_{st} between all source-destination pairs (s, t) in \mathcal{G} . In our imple-

mentation, we run Dijkstra’s algorithm [Dij59] on each node in \mathcal{G} , but the precise choice of shortest-path algorithm does not matter for our compression procedure, as its cost is dominated by the other steps. After computing all-pairs shortest paths in \mathcal{G} , we define the next-hop routing matrices $M^{(\text{NE})}, M^{(\text{NW})} \in \{0, 1\}^{n \times n}$ for \mathcal{G} , where $(M_{st}^{(\text{NE})}, M_{st}^{(\text{NW})})$ encodes the direction of the first edge in the shortest path p_{st} .

Just as the geometry of road networks enables us to orient the edges, the geometry also suggests a method for compressing the next-hop routing matrices. Take for example the road network in Figure 1. From the visualization, we observe that when the destination t lies to the north of the source s , the first hop on the shortest path is usually to take the edge directed north. In our framework, this means that both $M_{st}^{(\text{NE})}$ and $M_{st}^{(\text{NW})}$ are more likely to be 0 rather than 1. Thus, by orienting the edges in the graph consistently, we find that the resulting routing matrices $M^{(\text{NE})}$ and $M^{(\text{NW})}$ have potentially compressible structure.

To compress a matrix $M \in \{0, 1\}^{n \times n}$, we first rescale the elements in M to be in $\{-1, 1\}$. Our goal is to find two matrices $A, B \in \mathbb{Z}^{n \times d}$ such that $\text{sign}(AB^T) = M$ with $d < n$.³ We can formulate the problem of computing A and B as an optimization problem with objective function $J(A, B)$:

$$J(A, B) = \sum_{j=1}^n \sum_{k=1}^n \ell \left((AB^T)_{jk}, M_{jk} \right), \quad (2)$$

where $\ell(x, t)$ is a loss function. A simple loss function is the 0-1 loss function $\ell(x, t) = \mathbf{1}\{\text{sign}(x) \neq t\}$, which assigns a uniform loss of 1 whenever the sign of the predicted value x does not match the target value t . However, from an optimization perspective, the 0-1 loss is not a good loss function since it is non-convex and neither continuous nor differentiable. Practitioners have instead used continuous convex approximations to the 0-1 loss, such as the SVM hinge loss $\ell_{\text{hinge}}(x, t) = \max(0, 1 - tx)$ [RVC⁺04] and its quadratically smoothed variant, the modified Huber hinge loss [Zha04]:

$$\ell_{\text{huber}}(x, t) = \begin{cases} \max(0, 1 - tx)^2 & tx \geq -1 \\ -4 \cdot tx & \text{otherwise.} \end{cases} \quad (3)$$

In our setting, we use the modified Huber hinge loss ℓ_{huber} . While ℓ_{huber} is convex in the input x , it is not convex in the optimization parameters A, B (due to the matrix product), and so the objective function $J(A, B)$ is not convex in A, B . Thus, standard optimization algorithms like LBFGS [BLNZ95] are not guaranteed to find the global optimum. The hope is that even a local optimum will correspond to a low-rank, sign-preserving decomposition of the matrix M , and indeed, we confirm this empirically.

When we perform the optimization using LBFGS, the matrices A, B are real-valued. To obtain matrices over the integers, we scale the entries in A, B by a constant factor and round. The scaling factor is empirically chosen so as to preserve the relation $\text{sign}(AB^T) = M$. We describe this in greater detail in Section 5.3.

4 Private Navigation Protocol

In this section, we describe our protocol for privately computing shortest paths. First, we describe the cryptographic building blocks we employ in our construction.

³This is not the same as computing a low-rank approximation of M . Our goal is to find low-rank matrices whose product preserves the *signs* of the entries of M . In practice, the matrix M is full-rank, and not well-approximated by a low-rank product.

Private information retrieval. A computational private information retrieval (PIR) [CMS99, CGKS95, KO97, Cha04, GR05, Lip05, OI07] protocol is a two-party protocol between a sender who holds a database $\mathcal{D} = \{r_1, \dots, r_n\}$ and a receiver who holds an index $i \in [n]$. At the conclusion of the PIR protocol, the receiver learns r_i while the sender learns nothing. A PIR protocol only ensures privacy for the receiver’s index (and not for the remaining records in the sender’s database).

Oblivious transfer. Similar to PIR, an 1-out-of- n oblivious transfer (OT) protocol [NP99, NP01, NP05, Rab05] is a two-party protocol that allows the receiver to privately retrieve a record r_i from the sender who holds a database $\{r_1, \dots, r_n\}$. In contrast with PIR, an OT protocol also provides privacy for the sender: the receiver only learns its requested record r_i , and nothing else about the other records. Closely related is the notion of symmetric PIR (SPIR) [KO97, GIKM00, NP05], which is functionally equivalent to OT.

Garbled circuits. Yao’s garbled circuits [Yao86, LP09, BHR12] were initially developed for secure two-party computation. The core of Yao’s construction is an efficient transformation that takes a Boolean circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and produces a garbled circuit \tilde{C} along with n pairs of encodings $\{k_i^0, k_i^1\}_{i \in [n]}$. Then, for any input $x \in \{0, 1\}^n$, the combination of the garbled circuit \tilde{C} and the encodings $S_x = \{k_i^{x_i}\}_{i \in [n]}$ (where x_i denotes the i^{th} bit of x) enable one to compute $C(x)$, and yet reveal nothing else about x .

4.1 Protocol Design Overview

We first give an intuitive overview of our fully-private navigation protocol. As described in Section 3, we first preprocess the network \mathcal{G} to have maximum out-degree $d = 4$ and then associate a cardinal direction with each of the edges in \mathcal{G} . As in Section 3, let $(M^{(\text{NE})}, M^{(\text{NW})})$ be the precomputed next-hop routing matrices for \mathcal{G} , and let $(A^{(\text{NE})}, B^{(\text{NE})}), (A^{(\text{NW})}, B^{(\text{NW})})$ be the compressed representation of $M^{(\text{NE})}, M^{(\text{NW})}$, respectively.

Our private shortest paths protocol is an iterative protocol that reveals the shortest path from a source s to a destination t one hop at a time. When the client engages in the protocol with input (s, t) , it learns which neighbor v of s is the next node on the shortest path from s to t . Then, on the next round of the protocol, the client issues a query (v, t) to learn the next node in the path, and so on, until it arrives at the destination node t . With this iterative approach, each round of our protocol can be viewed as a two-party computation of the entry $(M_{st}^{(\text{NE})}, M_{st}^{(\text{NW})})$ from the next-hop routing matrices. We give the full description of our private navigation protocol in Figure 3, and sketch out the important principles here. To simplify the presentation, we first present the core building blocks that suffice for semi-honest security. We then describe additional consistency checks that we introduce to obtain security against a malicious client and privacy against a malicious server.

4.1.1 Semi-honest Secure Construction

Abstractly, we can view each round of our protocol as computing the following two-party functionality twice (once for $M^{(\text{NE})}$ and once for $M^{(\text{NW})}$). The server has two matrices $A, B \in \mathbb{Z}^{n \times d}$, which we will refer to as the source and destination matrices, respectively, and the client has two indices $s, t \in [n]$. At the end of the protocol, the client should learn $\text{sign}(\langle A_s, B_t \rangle)$, where A_s and B_t are the s^{th} and t^{th} rows of A and B , respectively. The client should learn nothing else about A and B , while the server should not learn anything. Our protocol can thus be decomposed into two components:

1. Evaluation of the inner product $\langle A_s, B_t \rangle$ between the *source vector* A_s and the *destination vector* B_t .
2. Determining the sign of $\langle A_s, B_t \rangle$.

In the following, we will work over a finite field \mathbb{F}_p large enough to contain the entries in A, B . In particular, we view A, B as $n \times d$ matrices over \mathbb{F}_p .

Evaluating the inner product. The first step in our protocol is evaluating the inner product between the source vector A_s and the destination vector B_t . Directly revealing the value of $\langle A_s, B_t \rangle$ to the client, however, leaks information about the entries in the compressed routing matrices A, B . To protect against this leakage, we instead reveal a blinded version of the inner product. Specifically, on each round of the protocol, the server chooses blinding factors $\alpha \xleftarrow{R} \mathbb{F}_p^*$ and $\beta \xleftarrow{R} \mathbb{F}_p$. We then construct the protocol such that at the end of the first step, the client learns the blinded value $\alpha \langle A_s, B_t \rangle + \beta$ instead of $\langle A_s, B_t \rangle$.

One candidate approach for computing the blinded inner product is to use a garbled circuit. However, while Yao’s garbled circuits suffice for private evaluation of any two-party functionality, when the underlying operations are more naturally expressed as addition and multiplication over \mathbb{F}_p , it is more convenient to express the functionality in terms of an arithmetic circuit. In an arithmetic circuit (over \mathbb{F}_p), the “gates” correspond to field operations (addition and multiplication), and the values on the wires correspond to field elements.

In recent work, Applebaum et al. [AIK14] construct the analog of Yao’s garbling procedure for arithmetic circuits. In particular, evaluating a function of the form $f(x, y) = \langle x, y \rangle + \sum_{i \in [d]} z_i$, where $x, y \in \mathbb{F}_p^d$ and each $z_i \in \mathbb{F}_p$ is a constant can be done efficiently using the affinization gadgets from [AIK14, §5]. Specifically, for each x_i, y_i , we define the following affine encoding functions $L_{x_i}^{\text{affine}}(x_i), L_{y_i}^{\text{affine}}(y_i)$:

$$\begin{aligned} L_{x_i}^{\text{affine}}(x_i) &= \left(x_i - r_i^{(1)}, x_i r_i^{(2)} + z_i + r_i^{(3)} \right) \\ L_{y_i}^{\text{affine}}(y_i) &= \left(y_i - r_i^{(2)}, y_i r_i^{(1)} - r_i^{(1)} r_i^{(2)} - r_i^{(3)} \right), \end{aligned} \quad (4)$$

where $r_i^{(1)}, r_i^{(2)}, r_i^{(3)}$ are chosen uniformly from \mathbb{F}_p . We will also write $L_{x_i}^{\text{affine}}(x_i; r_i), L_{y_i}^{\text{affine}}(y_i; r_i)$ to denote affine encodings of x_i and y_i using randomness $r_i \in \mathbb{F}_p^3$. Given $L_{x_i}^{\text{affine}}(x_i)$ and $L_{y_i}^{\text{affine}}(y_i)$ for all $i \in [n]$, evaluating $f(x, y)$ corresponds to evaluating the expression

$$\sum_{i \in [n]} \left[L_{x_i}^{\text{affine}}(x_i) \right]_1 \cdot \left[L_{y_i}^{\text{affine}}(y_i) \right]_1 + \left[L_{x_i}^{\text{affine}}(x_i) \right]_2 + \left[L_{y_i}^{\text{affine}}(y_i) \right]_2, \quad (5)$$

where we write $[\cdot]_i$ to denote the i^{th} component of a tuple. For notational convenience, we also define $L_x^{\text{affine}}(x)$ and $L_y^{\text{affine}}(y)$ as

$$\begin{aligned} L_x^{\text{affine}}(x) &= \left(L_{x_1}^{\text{affine}}(x_1), \dots, L_{x_d}^{\text{affine}}(x_d) \right) \\ L_y^{\text{affine}}(y) &= \left(L_{y_1}^{\text{affine}}(y_1), \dots, L_{y_d}^{\text{affine}}(y_d) \right). \end{aligned} \quad (6)$$

Similarly, we write $L_x^{\text{affine}}(x; r), L_y^{\text{affine}}(y; r)$ to denote the affine encoding of vectors $x, y \in \mathbb{F}_p^d$ using randomness $r \in \mathbb{F}_p^{3d}$. The affine encodings $L_x^{\text{affine}}(x), L_y^{\text{affine}}(y)$ provides statistical privacy for the input vectors x, y [AIK14, Lemma 5.1].

Next, we describe how these affine encodings can be used to compute the blinded inner product in the first step of the protocol. At the beginning of each round, the server chooses blinding factors $\alpha \xleftarrow{R} \mathbb{F}_p^*$ and $\beta \xleftarrow{R} \mathbb{F}_p$. Then, it constructs the affine encoding functions $L_x^{\text{affine}}, L_y^{\text{affine}}$ for the function $f_{\alpha, \beta}(x, y) = \langle \alpha x, y \rangle + \beta$ according to Eq. (4). Next, the server prepares two encoding databases \mathcal{D}_{src} and \mathcal{D}_{dst} where the s^{th} record in \mathcal{D}_{src} consists of the affine encodings $L_x^{\text{affine}}(A_s)$ of each source vector, and the t^{th} record in \mathcal{D}_{dst} consists of $L_y^{\text{affine}}(B_t)$ of each destination vector. To evaluate the blinded inner product, the client performs two SPIR queries: one for the s^{th} record in \mathcal{D}_{src} to obtain the encodings of A_s and one for the t^{th} record

in \mathcal{D}_{dst} to obtain the encodings of B_t .⁴ The client then evaluates the arithmetic circuit using Eq. (5) to obtain $z = f_{\alpha,\beta}(A_s, B_t)$. To a malicious client, without knowledge of α or β , the value $f_{\alpha,\beta}(A_s, B_t)$ appears uniform over \mathbb{F}_p and independent of A_s, B_t .

Determining the sign. To complete the description, it remains to describe a way for the client to learn the sign of the inner product $\langle A_s, B_t \rangle$. The client has the value $z = \alpha \langle A_s, B_t \rangle + \beta$ from the output of the arithmetic circuit while the server knows the blinding factors α, β . Since computing the sign function is equivalent to performing a comparison, arithmetic circuits are unsuitable for the task. Instead, we construct a separate Yao circuit to unblind the inner product and compare it against zero. More specifically, let $g(x, \gamma, \delta) = \mathbf{1}\{[\gamma x + \delta]_p > 0\}$, where $[\cdot]_p$ denotes reduction modulo p , with output in the interval $(-p/2, p/2)$. Then,

$$g(z, \alpha^{-1}, -\alpha^{-1}\beta) = \text{sign}(A_s, B_t).$$

To conclude the protocol, the server garbles a Boolean circuit C^{unblind} for the unblinding function g to obtain a garbled circuit $\tilde{C}^{\text{unblind}}$ along with a set of encodings L^{unblind} . It sends the garbled circuit to the client, along with encodings of the unblinding coefficients $\gamma = \alpha^{-1}, \delta = \alpha^{-1}\beta$ to the client. The client engages in 1-out-of-2 OTs to obtain the input encodings of z , and evaluates the garbled circuit $\tilde{C}^{\text{unblind}}$ to learn $\text{sign}(\langle A_s, B_t \rangle)$.

4.1.2 Enforcing Consistency for Stronger Security

As described, the protocol reveals just a single edge in the shortest path. Repeated iteration of the protocol allows the client to learn the full shortest path. Moreover, since the server’s view of the protocol execution consists only of its view in the PIR and OT protocols, privacy of these underlying primitives ensures privacy of the client’s location, even against a malicious server.⁵

Security for the server only holds if the client follows the protocol and makes consistent queries on each round. However, a malicious client can request the shortest path for a different source and/or destination on each round, thereby allowing it to learn edges along arbitrary shortest paths of its choosing. To protect against a malicious client, we bind the client to making *consistent* queries across consecutive rounds of the protocol. We say that a sequence of source-destination queries $(s_1, t_1), \dots, (s_\ell, t_\ell)$ is *consistent* if for all $i \in [\ell]$, $t_1 = t_i$, and $s_{i+1} = v_i$ where v_i is the first node on the shortest path from s_i to t_i .

Consistency for the destinations. To bind the client to a single destination, we do the following. At the beginning of the protocol, for each row $i \in [n]$ in \mathcal{D}_{dst} , the server chooses a symmetric encryption key $k_{\text{dst},i}$. Then, on each round of the protocol, it encrypts the i^{th} record in \mathcal{D}_{dst} with the key $k_{\text{dst},i}$. Next, at the beginning of the protocol, the client OTs for the key $k_{\text{dst},t}$ corresponding to its destination t . Since this step is performed only once at the beginning of the protocol, the only record in \mathcal{D}_{dst} that the client can decrypt is the one corresponding to its original destination. Because each record in \mathcal{D}_{dst} is encrypted under a different key, the client can use a PIR protocol instead of an SPIR protocol when requesting the record from \mathcal{D}_{dst} .

Consistency for the sources. Maintaining consistency between the source queries is more challenging because the source changes each round. We use the fact that the preprocessed graph has out-degree at

⁴The databases \mathcal{D}_{src} and \mathcal{D}_{dst} are each databases over n records (as opposed to n^2 in the straw-man protocol from Section 1).

⁵While a malicious server can send the client malformed circuits or induce selective failure attacks, the server does not receive any output during the protocol execution nor does the client abort the protocol when malformed input is received. Thus, we achieve privacy against a malicious server.

Inputs: Tuples $(z_{\text{NE}}, \gamma_{\text{NE}}, \delta_{\text{NE}}), (z_{\text{NW}}, \gamma_{\text{NW}}, \delta_{\text{NW}}) \in \mathbb{F}_p^3$, PRF keys $k_{\text{NE}}^0, k_{\text{NE}}^1, k_{\text{NW}}^0, k_{\text{NW}}^1 \in \{0, 1\}^\rho$, and the source and destination nodes $s, t \in [n]$. The bit-length τ is public and fixed (hard-wired into g).

Operation of g :

- If $s = t$, then output \perp .
- If $[\gamma_{\text{NE}} z_{\text{NE}} + \delta_{\text{NE}}]_p \notin [-2^\tau, 2^\tau]$ or $[\gamma_{\text{NW}} z_{\text{NW}} + \delta_{\text{NW}}]_p \notin [-2^\tau, 2^\tau]$, output \perp .
- Let $b_{\text{NE}} = \mathbf{1}\{[\gamma_{\text{NE}} z_{\text{NE}} + \delta_{\text{NE}}]_p > 0\}$, and let $b_{\text{NW}} = \mathbf{1}\{[\gamma_{\text{NW}} z_{\text{NW}} + \delta_{\text{NW}}]_p > 0\}$. Output $(b_{\text{NE}}, b_{\text{NW}}, k_{\text{NE}}^{b_{\text{NE}}}, k_{\text{NW}}^{b_{\text{NW}}})$.

Figure 2: Neighbor-computation function g for the private routing protocol.

most four. Thus, on each round, there are at most four possible sources that can appear in a consistent query in the next round.

Our construction uses a semantically-secure symmetric encryption scheme (Enc, Dec) with key-space $\{0, 1\}^\ell$, and a PRF F with domain $\{\text{N}, \text{E}, \text{S}, \text{W}\}$ and range $\{0, 1\}^\ell$. On each round of the protocol, the server generates a new set of source keys $k_{\text{src},1}, \dots, k_{\text{src},n} \in \{0, 1\}^\ell$ for encrypting the records in \mathcal{D}_{src} in the *next* round of the protocol. The server also chooses four PRF keys $k_{\text{NE}}^0, k_{\text{NE}}^1, k_{\text{NW}}^0, k_{\text{NW}}^1$, which are used to derive directional keys $k_{\text{N}}, k_{\text{E}}, k_{\text{S}}, k_{\text{W}}$. Next, for each node $v \in [n]$ in \mathcal{D}_{src} , let v_{dir} be the neighbor of v in direction $\text{dir} \in \{\text{N}, \text{E}, \text{S}, \text{W}\}$ (if there is one). The server augments the v^{th} record in \mathcal{D}_{src} with an encryption of the source key $k_{\text{src},v_{\text{dir}}}$ under the directional key k_{dir} .

When the client requests record v from \mathcal{D}_{src} , it also obtains encryptions of the keys of the neighbors of v for the *next* round of the protocol. By ensuring the client only learns one of the directional keys, it will only be able to learn the encryption key for a single source node on the next round of the protocol. We achieve this by including the PRF keys $k_{\text{NE}}^0, k_{\text{NE}}^1, k_{\text{NW}}^0, k_{\text{NW}}^1$ used to derive the directional keys as input to the garbled circuit. Then, in addition to outputting the direction, the garbled circuit also outputs the subset of PRF keys needed to derive exactly one of the directional keys $k_{\text{N}}, k_{\text{E}}, k_{\text{S}}, k_{\text{W}}$. This ensures that the client has at most one source key in the next round of the protocol.

Consistency within a round. In addition to ensuring consistency between consecutive rounds of the protocol, we also require that the client’s input to the garbled circuit is consistent with the output it obtained from evaluating the affine encodings. To enforce this, we use the fact that the entries of the routing matrices A, B are bounded: there exists τ such that $\langle A_s, B_t \rangle \in [-2^\tau, 2^\tau]$ for all $s, t \in V$. Then, in our construction, we choose the size of the finite field \mathbb{F}_p to be much larger than the size of the interval $[-2^\tau, 2^\tau]$. Recall that the arithmetic circuit computes a blinded inner product $z \leftarrow \alpha \langle A_s, B_t \rangle + \beta$ where α, β are uniform in \mathbb{F}_p^* and \mathbb{F}_p , respectively. To unblind the inner product, the server constructs a garbled circuit that first evaluates the function $g_{\gamma,\delta}(z) = \gamma z + \delta$ with $\gamma = \alpha^{-1}$ and $\delta = -\alpha^{-1}\beta$. By construction, γ is uniform over \mathbb{F}_p^* and δ is uniform over \mathbb{F}_p . Thus, using the fact that $\{g_{\gamma,\delta}(z) \mid \gamma \in \mathbb{F}_p^*, \delta \in \mathbb{F}_p\}$ is a pairwise independent family of functions, we conclude that the probability that $g_{\gamma,\delta}(z') \in [-2^\tau, 2^\tau]$ is precisely $2^{\tau+1}/p$ for all $z' \in \mathbb{F}_p$. By choosing $p \gg 2^{\tau+1}$, we can ensure that the adversary cannot successfully cheat except with very small probability.

Lastly, we remark that when the client issues a query (s, t) where $s = t$, the protocol should not reveal the key for any other node in the graph. To address this, we also introduce an equality test into the garbled circuit such that on input $s = t$, the output is \perp . We give a complete specification of the neighbor-computation function that incorporates these additional consistency checks in Figure 2.

Fix a security parameter λ and a statistical security parameter μ . Let $\mathcal{G} = (V, E)$ be a weighted directed graph with n vertices, such that the out-degree of every vertex is at most 4. The client's input to the protocol consists of two nodes, $s, t \in V$, representing the source and destination of the shortest path the client is requesting. The server's inputs are the compressed routing matrices $A^{(\text{NE})}, B^{(\text{NE})}, A^{(\text{NW})}, B^{(\text{NW})} \in \mathbb{Z}^{n \times d}$ (as defined in Section 3).

We assume the following quantities are public and known to both the client and the server:

- The structure of the graph \mathcal{G} , but not the edge weights.
- The number of columns d in the compressed routing matrices.
- A bound on the bit-length τ of the values in the products $A^{(\text{NE})} \cdot (B^{(\text{NE})})^T$ and $A^{(\text{NW})} \cdot (B^{(\text{NW})})^T$.
- The total number of rounds R .

In the following description, let (Enc, Dec) be a CPA-secure symmetric encryption scheme with key space $\{0, 1\}^\ell$, and let $F : \{0, 1\}^\rho \times \{N, E, S, W\} \rightarrow \{0, 1\}^\ell$ be a PRF (where $\ell, \rho = \text{poly}(\lambda)$). Fix a prime-order finite field \mathbb{F}_p such that $p > 2^{\tau + \mu + 1}$.

Setup:

1. For each $i \in [n]$, the server chooses independent symmetric encryption keys $k_{\text{src},i}^{(1)}, k_{\text{dst},i} \xleftarrow{R} \{0, 1\}^\ell$.
2. The client and the server engage in two 1-out-of- n OT protocols with the client playing the role of the receiver:
 - The client requests the s^{th} record from the server's database $(k_{\text{src},1}^{(1)}, \dots, k_{\text{src},n}^{(1)})$, receiving a value $\hat{k}_{\text{src}}^{(1)}$.
 - The client requests the t^{th} record from the server's database $(k_{\text{dst},1}, \dots, k_{\text{dst},n})$, receiving a value \hat{k}_{dst} .

For each round $r = 1, \dots, R$ of the protocol:

1. The server chooses blinding factors $\alpha_{\text{NE}}, \alpha_{\text{NW}} \xleftarrow{R} \mathbb{F}_p^*$ and $\beta_{\text{NE}}, \beta_{\text{NW}} \xleftarrow{R} \mathbb{F}_p$. Next, let $\gamma_{\text{NE}} = \alpha_{\text{NE}}^{-1}$ and $\delta_{\text{NE}} = -\alpha_{\text{NE}}^{-1}\beta_{\text{NE}} \in \mathbb{F}_p$. Define γ_{NW} and δ_{NW} analogously.
2. Let $f_{\text{NE}}, f_{\text{NW}} : \mathbb{F}_p^d \times \mathbb{F}_p^d \rightarrow \mathbb{F}_p$ where $f_{\text{NE}}(x, y) = \langle \alpha_{\text{NE}}x, y \rangle + \beta_{\text{NE}}$ and $f_{\text{NW}}(x, y) = \langle \alpha_{\text{NW}}x, y \rangle + \beta_{\text{NW}}$. The server then does the following:
 - Apply the affine encoding algorithm (Eq. 4) to f_{NE} to obtain encoding functions $L_{\text{NE},x}^{\text{affine}}, L_{\text{NW},y}^{\text{affine}}$, for the inputs x and y , respectively.
 - Apply the affine encoding algorithm to f_{NW} to obtain encoding functions $L_{\text{NW},x}^{\text{affine}}, L_{\text{NW},y}^{\text{affine}}$.
3. Let C^{unblind} be a Boolean circuit for computing the neighbor-computation function in Figure 2. The server runs Yao's garbling algorithm on C^{unblind} to obtain a garbled circuit $\tilde{C}^{\text{unblind}}$ along with encoding functions L_x^{unblind} , for each of the inputs x to the neighbor-computation function in Figure 2.
4. The server chooses symmetric encryption keys $k_{\text{src},1}^{(r+1)}, \dots, k_{\text{src},n}^{(r+1)} \xleftarrow{R} \{0, 1\}^\ell$. These are used to encrypt the contents of the source database on the next round of the protocol.
5. The server chooses four PRF keys $k_{\text{NE}}^0, k_{\text{NE}}^1, k_{\text{NW}}^0, k_{\text{NW}}^1 \xleftarrow{R} \{0, 1\}^\rho$, two for each axis. Then, the server defines the encryption keys for each direction as follows:⁶

$$\begin{aligned} k_N &= F(k_{\text{NE}}^0, N) \oplus F(k_{\text{NW}}^0, N) & k_E &= F(k_{\text{NE}}^0, E) \oplus F(k_{\text{NW}}^1, E) \\ k_S &= F(k_{\text{NE}}^1, S) \oplus F(k_{\text{NW}}^1, S) & k_W &= F(k_{\text{NE}}^1, W) \oplus F(k_{\text{NW}}^0, W). \end{aligned}$$

Figure 3: The fully-private routing protocol, as outlined in Section 4. The protocol description continues on the next page.

⁶An alternative approach is for the server to choose the encryption keys k_N, k_E, k_S, k_W uniformly at random from $\{0, 1\}^\ell$ instead of using the key-derivation procedure. In this case, the neighbor-computation function (Figure 2) would be modified to take as input the encryption keys k_N, k_E, k_S, k_W rather than the PRF keys $k_{\text{NE}}^0, k_{\text{NE}}^1, k_{\text{NW}}^0, k_{\text{NW}}^1$, and would output a single encryption key. While this approach is conceptually simpler, the resulting neighbor-computation circuits are slightly larger. In our implementation, we use the key-derivation procedure shown in the figure.

6. The server prepares the source database \mathcal{D}_{src} as follows. For each node $u \in [n]$, the u^{th} record in \mathcal{D}_{src} is an encryption under $k_{\text{src},u}^{(r)}$ of the following:

- The arithmetic circuit encodings $L_{\text{NE},x}^{\text{affine}}(A_u^{(\text{NE})})$, $L_{\text{NW},x}^{\text{affine}}(A_u^{(\text{NW})})$ of the source vectors $A_u^{(\text{NE})}$ and $A_u^{(\text{NW})}$.
- The garbled circuit encodings $L_s^{\text{unblind}}(u)$ of the source node u .
- Encryptions of the source keys for the neighbors of u in the next round of the protocol under the direction keys:

$$\kappa_{\text{N}} = \text{Enc}(k_{\text{N}}, k_{\text{src},v_{\text{N}}}^{(r+1)}), \quad \kappa_{\text{E}} = \text{Enc}(k_{\text{E}}, k_{\text{src},v_{\text{E}}}^{(r+1)}), \quad \kappa_{\text{S}} = \text{Enc}(k_{\text{S}}, k_{\text{src},v_{\text{S}}}^{(r+1)}), \quad \kappa_{\text{W}} = \text{Enc}(k_{\text{W}}, k_{\text{src},v_{\text{W}}}^{(r+1)}),$$

where $v_{\text{N}}, v_{\text{E}}, v_{\text{S}}, v_{\text{W}}$ is the neighbor of u to the north, east, south, or west, respectively. If u does not have a neighbor in a given direction $\text{dir} \in \{\text{N}, \text{E}, \text{S}, \text{W}\}$, then define $k_{\text{src},v_{\text{dir}}}^{(r+1)}$ to be the all-zeroes string 0^ℓ .

7. The server prepares the destination database \mathcal{D}_{dst} as follows. For each node $u \in [n]$, the u^{th} record in \mathcal{D}_{dst} is an encryption under $k_{\text{dst},u}$ of the following:

- The arithmetic circuit encodings $L_{\text{NE},y}^{\text{affine}}(B_u^{(\text{NE})})$, $L_{\text{NW},y}^{\text{affine}}(B_u^{(\text{NW})})$ of the destination vectors $B_u^{(\text{NE})}$ and $B_u^{(\text{NW})}$.
- The garbled circuit encodings $L_t^{\text{unblind}}(u)$ of the destination node u .

8. The client and server engage in two PIR protocols with the client playing role of receiver:

- The client requests record s from the server's database \mathcal{D}_{src} and obtains a record \hat{c}_{src} .
- The client requests record t from the server's database \mathcal{D}_{dst} and obtains a record \hat{c}_{dst} .

9. The client decrypts the records: $\hat{r}_{\text{src}} \leftarrow \text{Dec}(\hat{k}_{\text{src}}^{(r)}, \hat{c}_{\text{src}})$ and $\hat{r}_{\text{dst}} \leftarrow \text{Dec}(\hat{k}_{\text{dst}}, \hat{c}_{\text{dst}})$:

- It parses \hat{r}_{src} into two sets of arithmetic circuit encodings $\hat{L}_{\text{NE},x}^{\text{affine}}$ and $\hat{L}_{\text{NW},x}^{\text{affine}}$, a set of garbled circuit encodings $\hat{L}_s^{\text{unblind}}$, and four encryptions $\hat{\kappa}_{\text{N}}, \hat{\kappa}_{\text{E}}, \hat{\kappa}_{\text{S}}, \hat{\kappa}_{\text{W}}$ of source keys for the next round.
- It parses \hat{r}_{dst} into two sets of arithmetic circuit encodings for $\hat{L}_{\text{NE},y}^{\text{affine}}$ and $\hat{L}_{\text{NW},y}^{\text{affine}}$, and a set of garbled circuit encodings $\hat{L}_t^{\text{unblind}}$.

Using the encodings $\hat{L}_{\text{NE},x}^{\text{affine}}$ and $\hat{L}_{\text{NE},y}^{\text{affine}}$, the client evaluates the arithmetic circuit (Eq. 5) to learn \hat{z}_{NE} . Similarly, using the encodings $\hat{L}_{\text{NW},x}^{\text{affine}}$ and $\hat{L}_{\text{NW},y}^{\text{affine}}$, the server evaluates to learn \hat{z}_{NW} . If the parsing of \hat{r}_{src} or \hat{r}_{dst} fails or the arithmetic circuit encodings are malformed, the client sets $\hat{z}_{\text{NE}}, \hat{z}_{\text{NW}} \xleftarrow{\text{R}} \mathbb{F}_p$.

10. The client engages in a series of 1-out-of-2 OTs with the server to obtain the garbled circuit encodings $L_{z_{\text{NE}}}^{\text{unblind}}(\hat{z}_{\text{NE}})$ and $L_{z_{\text{NW}}}^{\text{unblind}}(\hat{z}_{\text{NW}})$ of \hat{z}_{NE} and \hat{z}_{NW} , respectively. Let $\hat{L}_{z_{\text{NE}}}^{\text{unblind}}$ and $\hat{L}_{z_{\text{NW}}}^{\text{unblind}}$ denote the encodings the client receives.

11. The server sends to the client the garbled circuit $\tilde{C}^{\text{unblind}}$ and encodings of the unblinding coefficients

$$L_{\gamma_{\text{NE}}}^{\text{unblind}}(\gamma_{\text{NE}}), L_{\gamma_{\text{NW}}}^{\text{unblind}}(\gamma_{\text{NW}}), L_{\delta_{\text{NE}}}^{\text{unblind}}(\delta_{\text{NE}}), L_{\delta_{\text{NW}}}^{\text{unblind}}(\delta_{\text{NW}}),$$

as well as encodings of the PRF keys

$$L_{k_{\text{NE}}^0}^{\text{unblind}}(k_{\text{NE}}^0), L_{k_{\text{NE}}^1}^{\text{unblind}}(k_{\text{NE}}^1), L_{k_{\text{NW}}^0}^{\text{unblind}}(k_{\text{NW}}^0), L_{k_{\text{NW}}^1}^{\text{unblind}}(k_{\text{NW}}^1).$$

Figure 3 (Continued): The fully-private routing protocol, as outlined in Section 4. The protocol description continues on the next page.

12. The client evaluates the garbled circuit $\tilde{C}^{\text{unblind}}$. If the garbled circuit evaluation is successful and the client obtain outputs $(\hat{b}_{\text{NE}}, \hat{b}_{\text{NW}}, \hat{k}_{\text{NE}}, \hat{k}_{\text{NW}})$, then the client computes a direction $\text{dir} = \text{IndexToDirection}(\hat{b}_{\text{NE}}, \hat{b}_{\text{NW}}) \in \{\text{N}, \text{E}, \text{S}, \text{W}\}$ (Eq. (1)).

(a) The client computes the direction key $\hat{k}_{\text{dir}} = F(\hat{k}_{\text{NE}}, \text{dir}) \oplus F(\hat{k}_{\text{NW}}, \text{dir})$. Next, the client decrypts the encrypted source key \hat{k}_{dir} to obtain the source key $\hat{k}_{\text{src}}^{(r+1)} = \text{Dec}(\hat{k}_{\text{dir}}, \hat{k}_{\text{dir}})$ for the next round of the protocol.

(b) Let v_{dir} be the neighbor of s in the direction given by dir (define v_{dir} to be \perp if s does not have a neighbor in the direction dir). If $v_{\text{dir}} \neq \perp$, the client outputs v_{dir} and updates $s = v_{\text{dir}}$. Otherwise, if $v_{\text{dir}} = \perp$, the client outputs \perp and leaves s unchanged.

If the OT for the input wires to the garbled circuit fails, the garbled circuit evaluation fails, or the output of the garbled circuit is \perp , then the client outputs \perp , but continues with the protocol: it leaves s unchanged and sets $\hat{k}_{\text{src}}^{(r+1)} \xleftarrow{\text{R}} \{0, 1\}^\ell$.

Figure 3 (Continued): The fully-private routing protocol, as outlined in Section 4.

4.2 Security Model

In this section, we formally specify our security model. To define and argue the security of our protocol, we compare the protocol execution in the real-world (where the parties interact according to the specification given in Figure 3) to an execution in an ideal world where the parties have access to a trusted party that computes the shortest path. Following the conventions in [Can06], we view the protocol execution as occurring in the presence of an adversary \mathcal{A} and coordinated by an environment $\mathcal{E} = \{\mathcal{E}\}_\lambda$ (modeled as a family of polynomial size circuits parameterized by a security parameter λ). The environment \mathcal{E} is responsible for choosing the inputs to the protocol execution and plays the role of distinguisher between the real and ideal experiments.

As specified in Figure 3, we assume that the following quantities are public to the protocol execution: the topology of the network $\mathcal{G} = (V, E)$, the number of columns d in the compressed routing matrices, a bound on the bit-length τ of the values in the matrix products $A^{(\text{NE})} \cdot (B^{(\text{NE})})^T$ and $A^{(\text{NW})} \cdot (B^{(\text{NW})})^T$, and the total number of rounds R (i.e., the number of hops in the longest possible shortest path). We now define the real and ideal models of execution.

Definition 4.1 (Real Model of Execution). Let π be a private navigation protocol. In the real world, the parties interact according to the protocol specification π . Let \mathcal{E} be the environment and let \mathcal{A} be an adversary that corrupts either the client or the server. The protocol execution in the real world proceeds as follows:

1. **Inputs:** The environment \mathcal{E} chooses a source-destination pair $s, t \in V$ for the client and compressed next-hop routing matrices $A^{(\text{NE})}, B^{(\text{NE})}, A^{(\text{NW})}, B^{(\text{NW})} \in \mathbb{Z}^{n \times d}$ for the server. The bit-length of all entries in the matrix products $A^{(\text{NE})} \cdot (B^{(\text{NE})})^T$ and $A^{(\text{NW})} \cdot (B^{(\text{NW})})^T$ must be at most τ . Finally, the environment gives the input of the corrupted party to the adversary.
2. **Protocol Execution:** The parties begin executing the protocol. All honest parties behave according to the protocol specification. The adversary \mathcal{A} has full control over the behavior of the corrupted party and sees all messages received by the corrupted party.
3. **Output:** The honest party computes and gives its output to the environment \mathcal{E} . The adversary computes a function of its view of the protocol execution and gives it to \mathcal{E} .

At the conclusion of the protocol execution, the environment \mathcal{E} outputs a bit $b \in \{0, 1\}$. Let $\text{REAL}_{\pi, \mathcal{A}, \mathcal{E}}(\lambda)$ be the random variable corresponding to the value of this bit.

Definition 4.2 (Ideal Model of Execution). In the ideal world, the client and server have access to a trusted party \mathcal{T} that computes the shortest paths functionality f .

1. **Inputs:** Same as in the real model of execution.
2. **Submission to Trusted Party:** If a party is honest, it gives its input to the trusted party. If a party is corrupt, then it can send any input of its choosing to \mathcal{T} , as directed by the adversary \mathcal{A} .
3. **Trusted Computation:** From the next-hop routing matrices, the trusted party computes the first R hops on the shortest path from s to t : $s = v_0, v_1, \dots, v_R$. If $v_i = t$ for some $i < R$, then the trusted party sets v_{i+1}, \dots, v_R to \perp . If the next hop in the path at v_i for some i refers to a node not in \mathcal{G} , then the trusted party sets v_{i+1}, \dots, v_R to \perp . The trusted party sends the path v_0, \dots, v_R to the client. The server receives no output.
4. **Output:** An honest party gives the sequence of messages (possibly empty) it received from \mathcal{T} to \mathcal{E} . The adversary computes a function of its view of the protocol execution and gives it to \mathcal{E} .

At the conclusion of the protocol execution, the environment \mathcal{E} outputs a bit $b \in \{0, 1\}$. Let $\text{IDEAL}_{f, \mathcal{A}, \mathcal{E}}(\lambda)$ be the random variable corresponding to the value of this bit.

To state our security theorems, we now define the environment's distinguishing advantage. Informally, we will say that a protocol is secure if no polynomial-size environment is able to distinguish the real execution from the ideal execution with non-negligible probability.

Definition 4.3 (Distinguishing Advantage — Security). Let π be a private navigation protocol, and let f be the shortest path functionality. Fix an adversary \mathcal{A} , simulator \mathcal{S} , and an environment \mathcal{E} . Then, the distinguishing advantage $\text{Adv}_{\pi, f, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{(\text{sec})}$ of \mathcal{E} in the security game is given by

$$\text{Adv}_{\pi, f, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{(\text{sec})}(\lambda) = |\Pr[\text{REAL}_{\pi, \mathcal{A}, \mathcal{E}}(\lambda) = 0] - \Pr[\text{IDEAL}_{f, \mathcal{A}, \mathcal{E}}(\lambda) = 0]|.$$

We will also work with a weaker notion of *privacy* against a malicious adversary. Informally, we say that the protocol is private if an adversary is unable to learn anything about the inputs of the other party beyond what is explicitly leaked by the inputs and outputs of the computation. To formalize this notion, we use the conventions in [IKK⁺11] and define the distinguishing advantage in the privacy game.

Definition 4.4 (Distinguishing Advantage — Privacy). Let π be a private navigation protocol, and let f be the shortest path functionality. Fix an adversary \mathcal{A} , simulator \mathcal{S} , and an environment \mathcal{E} . Define $\text{REAL}'_{\pi, \mathcal{A}, \mathcal{E}}(\lambda)$ exactly as $\text{REAL}_{\pi, \mathcal{A}, \mathcal{E}}(\lambda)$ (Definition 4.1), except in the final step of the protocol execution, the environment only receives the adversary's output (and *not* the honest party's output). Define $\text{IDEAL}'_{f, \mathcal{S}, \mathcal{E}}(\lambda)$ analogously. Then, the distinguishing advantage $\text{Adv}_{\pi, f, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{(\text{priv})}$ of \mathcal{E} in the privacy game is given by

$$\text{Adv}_{\pi, f, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{(\text{priv})}(\lambda) = |\Pr[\text{REAL}'_{\pi, \mathcal{A}, \mathcal{E}}(\lambda) = 0] - \Pr[\text{IDEAL}'_{f, \mathcal{S}, \mathcal{E}}(\lambda) = 0]|.$$

4.3 Security Theorems

The first requirement is that our protocol provides security against a malicious client. This captures the notion that a malicious client does not learn anything more about the server’s routing information beyond the shortest path between its requested endpoints and the publicly available information. In our setting, we allow a privacy-performance trade-off where the client has a small probability ($R \cdot 2^{-\mu}$, where μ is the statistical security parameter) of learning additional information about the routing information. Since the order p of the finite field must satisfy $p > 2^{\tau+\mu+1}$, using larger finite fields will decrease the failure probability, but at the expense of performance. In our experiments, $R \cdot 2^{-\mu} \approx 2^{-30}$. We now state the formal security guarantee, but defer its formal proof to Appendix A.2.

Theorem 4.5 (Security Against a Malicious Client). *Let π be the protocol in Figure 3 instantiated with a CPA-secure encryption scheme (Enc, Dec), a secure PRF F , and an OT scheme secure against a malicious client. Let λ, μ be the security parameter and statistical security parameter, respectively. Let f be the ideal shortest-paths functionality. Then, for all PPT adversaries \mathcal{A} , there exists a PPT adversary \mathcal{S} such that for every polynomial-size circuit family $\mathcal{E} = \{\mathcal{E}\}_\lambda$,*

$$\text{Adv}_{\pi, f, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{(sec)}(\lambda) \leq \text{negl}(\lambda) + R \cdot 2^{-\mu},$$

where $\text{negl}(\lambda)$ denotes a negligible function in λ .

In addition to security against a malicious client, we require our protocol to provide privacy against a malicious server. In other words, while a malicious server might be able to cause the client to receive an invalid path, it still cannot learn any information about the client’s source or destination. We formalize this in the following theorem. We defer the formal proof to Appendix A.1.

Theorem 4.6 (Privacy Against a Malicious Server). *Let π be the protocol in Figure 3 instantiated with PIR and OT primitives that provide privacy against a malicious server. Let λ be a security parameter and let f be the ideal shortest-paths functionality. Then, for all PPT adversaries \mathcal{A} , there exists a PPT adversary \mathcal{S} such that for every polynomial-size circuit family $\mathcal{E} = \{\mathcal{E}\}_\lambda$,*

$$\text{Adv}_{\pi, f, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{(priv)}(\lambda) \leq \text{negl}(\lambda),$$

where $\text{negl}(\lambda)$ denotes a negligible function in λ .

5 Experiments

In this section, we describe our implementation of the private routing protocol from Figure 3. Then, we describe our procedure for preprocessing and compressing actual road networks for major cities taken from OpenStreetMap [Ope]. Finally, we give concrete performance benchmarks for our preprocessing and compression pipeline as well as our private routing protocol on actual road networks.

5.1 Protocol Implementation

To evaluate the performance of the protocol in Figure 3, we implemented the complete protocol in C++. In this section, we describe the building blocks of our implementation. For each primitive, we choose the parameters to guarantee a minimum of 80 bits of security. The complete protocol implementation contains approximately 4000 lines of code.

PIR. We implemented the (recursive) PIR protocol based on additive homomorphic encryption from [KO97, OI07]. We instantiate the additive homomorphic encryption scheme with Paillier’s cryptosystem [Pai99], and use NTL [Sho] over GMP [Gt12] to implement the necessary modular arithmetic. We use a 1024-bit RSA modulus for the plaintext space in the Paillier cryptosystem, which provides 80 bits of security. We use two levels of recursion in the PIR protocol, so the communication scales as $O(\sqrt[3]{n})$ for an n -record database.

OT. We instantiate the OT protocol with the protocol from [HL10, §7.3] which provides security against malicious clients and privacy against malicious servers. This protocol is a direct generalization of the Naor-Pinkas OT protocol [NP01] based on the decisional Diffie-Hellman (DDH) assumption. Security against a malicious client is enforced by having the client include a zero-knowledge proof of knowledge (specifically, a Schnorr proof [Sch89]) with its OT request. To decrease the number of rounds of communication, we apply the Fiat-Shamir heuristic [FS86] to transform the interactive proof of knowledge into a non-interactive one by working in the random oracle model. We instantiate the random oracle with the hash function SHA-256. For improved performance, we implement the Naor-Pinkas OT protocol over the 256-bit elliptic curve group `numsp256d1` from [BCLN14]. We use the MSR-ECC [BCLN14] library for the implementation of the underlying elliptic curve operations. The 256-bit curve provides 128 bits of security.

Arithmetic and Yao’s circuits. We implement our arithmetic circuits over the finite field \mathbb{F}_p where $p = 2^{61} - 1$ is a Mersenne prime. Then, reductions modulo p can be performed using just two p -bit additions. We use NTL [Sho] over GMP [Gt12] for the finite field arithmetic.

For the garbled circuit implementation, we use JustGarble [BHKR13] with the “free XOR” [KS08] and row-reduction optimizations [PSSW09]. We set the parameters of the garbling framework to obtain 80-bits of security. We use the optimized addition, comparison, and multiplexer circuits from [KSS09] to implement the neighbor-computation function shown in Figure 2. For multiplication, we implement the basic “school method.”

Record encryption and PRF. We instantiate the CPA-secure encryption scheme in Figure 3 with AES in counter mode. We also instantiate the PRF used for deriving the neighbor keys (Step 5 in Figure 3) with AES. We use the implementation of AES from OpenSSL [The03].

5.2 Preprocessing and Map Compression.

We extract the street maps for four major cities (San Francisco, Washington, D.C., Dallas, and Los Angeles) from OpenStreetMap [Ope]. For each city, we take its most important roadways based on annotations in OpenStreetMap, and construct the resulting graph \mathcal{G} . Specifically, we introduce a node for each street intersection in the city and an edge for each roadway. We assign edge weights based on the estimated time needed to traverse the associated road segment (computed by taking the length of the segment and dividing by the approximated speed limit along the segment). Using the procedure described in Section 3, we preprocess the graph to have out-degree at most 4. We then associate each edge of \mathcal{G} with a cardinal direction by solving the assignment problem from Section 3. We use Stachniss’ implementation [Sta04] of the Hungarian method [KY55] to solve this assignment problem.

Given the graph \mathcal{G} corresponding to the road network for each city, we run Dijkstra’s algorithm [Dij59] on each node s in \mathcal{G} to compute the shortest path between all pairs of nodes. Then, using the all-pairs shortest paths information, we construct the next-hop routing matrices ($M^{(NE)}, M^{(NW)}$) for \mathcal{G} . We remark that we can substitute any all-pairs shortest path algorithm for Dijkstra’s in this step. The underlying principle we exploit in the construction of our protocol is the fact that next-hop routing matrices for road networks have a simple compressible structure amenable to cryptography.

City	n	Preprocessing Time (s)	Compression Time (s)
San Francisco	1830	0.625	97.500
Washington, D.C.	2490	1.138	142.431
Dallas	4993	4.419	278.296
Los Angeles	7010	9.188	503.007

Table 1: Average time to preprocess and compress the next-hop routing matrices for different networks. The second column gives the number of nodes n in each city’s road network. The preprocessing time column gives the average time needed to orient the edges, compute all-pairs shortest paths, and construct the next-hop routing matrix for the network. The compression time column gives the average time needed to compress the NE or NW component of the next-hop routing matrices.

Finally, we implement the optimization-based compression approach described in Section 3 to compress the next-hop routing matrices $M^{(\text{NE})}$ and $M^{(\text{NW})}$. We minimize the objective function from Eq. (2) with the loss function set to the modified Huber hinge loss from Eq. (3). Because of the highly parallelizable nature of the objective function, we write specialized CUDA kernels to evaluate the objective function and its derivative on the GPU. In our experiments, we use the LBFGS optimization algorithm [BLNZ95] from the Python scientific computation libraries NumPy and SciPy [ADH⁺01] to solve the optimization problem.

5.3 Experiments

Graph preprocessing and compression. We first measure the time needed to preprocess and compress the next-hop routing matrices for several road networks. The preprocessing time includes the time needed to orient the edges, compute all-pairs shortest paths, and construct the next-hop routing matrix for the network (as described in Section 5.2).

We also measure the time needed to compress the resulting next-hop routing matrices for the different networks. Recall that our compression method takes a matrix $M \in \{-1, 1\}^{n \times n}$ and produces two matrices $A, B \in \mathbb{Z}^{n \times d}$ such that $\text{sign}(AB^T)$ is a good approximation of M . Since the modified Huber hinge loss (Eq. 3) is an upper bound on the 0-1 loss function $\ell(x, t) = \mathbf{1}\{\text{sign}(x) \neq t\}$, when the objective value $J(A, B)$ is less than 1 (where $J(A, B)$ is the objective function in Eq. 2), we have $\text{sign}(AB^T) = M$, i.e., the matrices A, B perfectly reconstruct M . The parameter d is the number of columns in the matrices A and B . Because our objective function is non-convex in the variables A and B , LBFGS is neither guaranteed to find the globally optimal solution, nor even to converge in a reasonable number of iterations. As a heuristic for deciding whether a candidate value of d admits a feasible solution that perfectly reconstructs $M^{(\text{NE})}$ and $M^{(\text{NW})}$, we run up to 5000 iterations of LBFGS and check whether the resulting solution gives a perfect reconstruction of M . To determine the most compact representation, we search over a range of possible values for d , and choose the smallest value d that yields a perfect reconstruction of M .

We apply our compression method to the routing matrices for road networks from four cities of varying size. Then, we compare the size of the original matrix M to the size of its compressed representation A, B . The number of bits needed to represent A, B is determined by two factors: the number of columns d in each matrix A, B and the precision ν (measured in number of bits) needed to represent each entry in A, B . Recall that the optimization procedure outputs two *real-valued* matrices such that $\text{sign}(AB^T) = M$. To obtain a representation over the integers (as required by the arithmetic circuits), we scale the entries of A, B by a constant factor and round each of the resulting entries to the nearest integer. The precision ν is the number of bits needed to represent each integer component of A, B after rescaling. We choose the smallest scaling factor such that the rescaled matrices perfectly reconstruct the routing matrix M .

We run the preprocessing and compression experiments on a machine running Ubuntu 14.04 with

City	n	d	ν	τ	Compression Factor
San Francisco	1830	12	10	20	7.63
Washington, D.C.	2490	14	10	19	8.89
Dallas	4993	19	12	23	10.95
Los Angeles	7010	26	12	24	11.23

Table 2: Parameters for the compressed representation of the road networks for each city: n is the number of nodes in each network, d and ν are the number of columns and the precision, respectively, in the routing matrices $A^{(NE)}$, $B^{(NE)}$, $A^{(NW)}$, $B^{(NW)}$ of the compressed representation, and τ is the maximum number of bits needed to represent an element in the products $A^{(NE)}(B^{(NE)})^T$ and $A^{(NW)}(B^{(NW)})^T$. The last column gives the compression factor attained for each network (ratio of size of uncompressed representation to size of compressed representation).

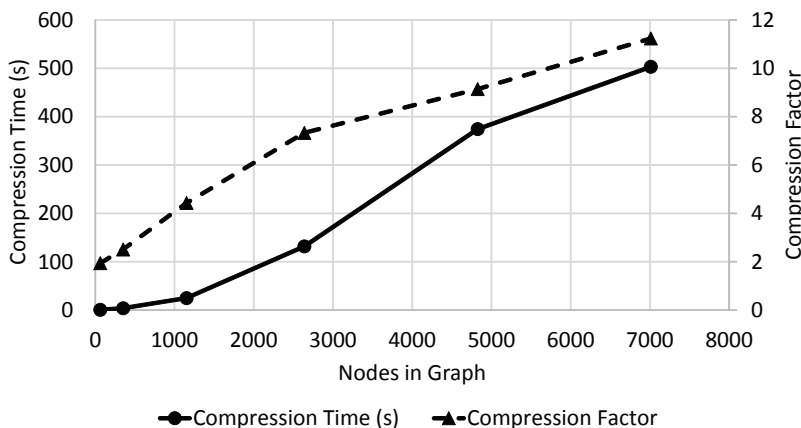


Figure 4: Average time needed to compress the next-hop routing matrix and the resulting compression factor for networks constructed from subgraphs of the road network of Los Angeles.

an 8-core 2.3 GHz Intel Core i7 CPU, 16 GB of RAM, and an Nvidia GeForce GT 750M GPU. The preprocessing and compression times for the different networks are summarized in Table 1. A description of the compressed representation of the routing matrices for each network is given in Table 2.

In Figure 4, we show the time needed to compress a single component of the next-hop routing matrix, as well as the resulting compression factor, for subgraphs of the road network for Los Angeles. The compression is quite effective, and the achievable compression factor increases with the size of the network. Moreover, even though the sizes of the next-hop routing matrices increase quadratically in the number of nodes in the graph, the optimization time remains modest. For graphs with 7000 nodes (and 350,000 optimization variables), finding a compact representation that perfectly reconstructs the next-hop matrix completes in under 10 minutes. Since we compress both the NE and NW components of the routing matrix, the total time to both preprocess and compress the shortest path information for the full city of Los Angeles is just over 15 minutes. Lastly, we note that the preprocessing time for each network is small: orienting the edges and computing all-pairs shortest paths via Dijkstra’s algorithm completes in under 10 seconds.

Performance on road networks. Next, we measure the run-time and bandwidth requirements of our private routing protocol from Figure 3. Table 2 gives the number of columns d , and the precision ν of the compressed representation of the networks for the different cities. In addition, we also compute the maximum number of bits τ needed to encode an element in the products $A^{(NE)}(B^{(NE)})^T$ and $A^{(NW)}(B^{(NW)})^T$. From Theorem 4.5, a malicious client can successfully cheat with probability at most $R \cdot 2^{-\mu}$, where μ is the

City	Total Time (s) (Single Round)	Client Computation (s)				Server Computation (s)			Bandwidth (KB)	
		Total	PIR	OT	GC	Total	PIR	OT	Upload	Download
San Francisco	1.44 ± 0.16	0.35	0.31	0.02	0.02	0.88	0.80	0.08	51.74	36.50
Washington, D.C.	1.64 ± 0.13	0.38	0.34	0.02	0.02	1.07	1.00	0.08	52.49	37.51
Dallas	2.91 ± 0.19	0.45	0.41	0.02	0.02	2.19	2.11	0.08	55.50	39.52
Los Angeles	4.75 ± 0.22	0.55	0.51	0.02	0.02	3.70	3.62	0.08	57.01	43.53

Table 3: Performance benchmarks (averaged over at least 90 iterations) for a single round of the private routing protocol described in Figure 3 on road networks for different cities. The “Total Time” column gives the average runtime and standard deviation for a single round of the protocol (including network communication times between a client and a server on Amazon EC2). The PIR, OT, and GC columns in the table refer to the time to perform the PIR for the affine encodings, the time to perform the OT for the garbled circuit encodings, and the time needed to evaluate the garbled circuit, respectively. The bandwidth measurements are taken with respect to the client (“upload” refers to communication from the client to the server)

City	R	Offline Setup		Online Setup		Total Online	Total Online
		Time (s)	Band. (MB)	Time (s)	Band. (MB)	Time (s)	Bandwidth (MB)
San Francisco	97	0.135	49.08	0.73	0.021	140.39	8.38
Washington, D.C.	120	0.170	60.72	0.76	0.023	197.48	10.57
Dallas	126	0.174	63.76	0.92	0.027	371.44	11.72
Los Angeles	165	0.223	83.49	1.00	0.028	784.34	16.23

Table 4: End-to-end performance benchmarks for the private routing protocol in Figure 3 on road networks for different cities. For each network, the number of iterations R is set to the maximum length of the shortest path between two nodes in the network. The offline computation refers to the server preparation and garbling of the R circuits for evaluating the neighbor-computation function from Figure 2. The offline computation time just includes the computational cost and does *not* include the garbled circuit download time. The online setup measurements correspond to computation and communication in the “Setup” phase of the protocol in Figure 3. The “Total Online Time” and “Total Online Bandwidth” columns give the total end-to-end time (including network communication) and total communication between the client and server in the *online* phase (navigation component) of the protocol.

statistical security parameter, and R is the total number of rounds in the protocol. For each network in our experiments, we set the number of rounds R to be the maximum length over all shortest paths between any source-destination pair in the network. This ensures both correctness (at the end of the protocol execution, the client obtains the complete shortest path from its source to its destination) as well as hides the length of the requested shortest path from the server (since the number of rounds in the protocol is independent of the client’s input). Next, recall the relation between μ and the order p of the finite field for the affine encodings: $p > 2^{\tau+\mu+1}$. In our experiments, we fix $p = 2^{61} - 1$, and R is at most $2^8 = 256$. This choice of parameters translates to μ ranging from 36 to 41, or analogously, a failure probability of 2^{-33} for the smaller networks to 2^{-28} for larger networks. Using larger fields will reduce this probability, but at the expense of performance.

To reduce the communication in each round of the protocol in Figure 3, we note that it is not necessary for the server to prepare and send a garbled circuit to the client on each round of the routing protocol. Since the neighbor-computation circuit is independent of the state of the protocol execution, the circuits can be generated and stored long before the protocol execution begins. Thus, in an *offline* phase, the server can prepare and transmit to the client a large number of garbled circuits. During the online protocol execution, on the r^{th} round, the server just sends the encodings corresponding to its input to the client; it does *not* send the description of the garbled circuit. This significantly reduces the communication cost of each round

of the online protocol. We note that even if the routing matrices $A^{(\text{NE})}, B^{(\text{NE})}, A^{(\text{NW})}, B^{(\text{NW})}$ changed (for instance, due to updates in traffic or weather conditions in the network) during the protocol execution, as long as the bound τ on the bit-length of entries in the products $A^{(\text{NE})}(B^{(\text{NE})})^T$ and $A^{(\text{NW})}(B^{(\text{NW})})^T$ remain fixed, the client and server do *not* have to redo this offline setup phase. We describe our extension for supporting live updates to the routing information in greater detail in Section 6.

We run the server on a compute-optimized Amazon EC2 instance (running Ubuntu 14.04 with a 32-core 2.7 GHz Intel Xeon E5-2680v2 processor and 60 GB of RAM) to model the computing resources of a cloud-based map provider. The throughput of our protocol is bounded by the PIR computation on the server’s side. We use up to 60 threads on the server for the PIR computation. All other parts of our system are single-threaded. For the client, we use a laptop running Ubuntu 14.04 with a 2.3 GHz Intel Core i7 CPU and 16 GB of RAM. The connection speed on the client is around 50 Mbps. Both client and server support the AES-NI instruction set, which we leverage in our implementation.

First, we measure the cost of one round of the private navigation protocol. We assume that the client has already downloaded the garbled circuits prior to the start of the protocol. Table 3 gives the total computation time and bandwidth per round of the routing protocol. When measuring the total time, we measure the end-to-end time on the client, which includes the time for the network round trips. Table 3 also gives a breakdown of the computation in terms of each component of the protocol: PIR for the arithmetic circuit encodings, OT for the garbled circuit encodings, and garbled-circuit evaluation for computing the next-hop.

We also measure the total end-to-end costs for a single shortest path query. As noted earlier, we set the number of rounds R for each network to be the maximum length of any shortest path in the network. Irrespective of the client’s source or destination, the client and server always engage in exactly R rounds of the private navigation protocol. Table 4 shows the total computation time and bandwidth required to complete a shortest-path query in the different networks. In the end-to-end benchmarks, we also measure the offline costs of the protocol, that is, the time needed for the server to garble R neighbor-computation circuits and the amount of communication needed for the client to download the circuits. In addition, we measure the computation and bandwidth needed in the online setup phase of the routing protocol (Figure 3).

In our protocol, the online setup phase of the protocol consists of three rounds of interaction. First, the server sends the client the public description of the map. Then the client OTs for the source and destination keys for the first round of the protocol, which requires two rounds of communication. As shown in Table 4, the online setup procedure completes in at most a second and requires under 30 KB of communication in our example networks.

Next, we consider the performance of each round of the protocol. From Table 3, the most computationally intensive component of our protocol is computing the responses to the PIR queries. In our implementation, we use a Paillier-based PIR, so the server must perform $O(n)$ modular exponentiations on each round of the protocol. While it is possible to use a less computationally-intensive PIR such as [MBFK14], the bandwidth required is much higher in practice. Nonetheless, our results demonstrate that the performance of our protocol is within the realm of practicality for real-time navigation in cities like San Francisco or Washington, D.C.

Lastly, we note that the offline costs are dominated essentially by communication. With hardware support for AES, garbling 100 neighbor-computation circuits on the server completes in just a quarter of a second. While garbling is fast, the size of each garbled circuit is 518.2 KB. For city networks, we typically require 100-150 circuits for each shortest-path query; this corresponds to 50-100 MB of offline download prior to the start of the navigation protocol. The experimental results, however, indicate that the number of garbled circuits required for an end-to-end execution grows sublinearly in the size of the graph. For example, the total number of rounds (and correspondingly, the number of required garbled circuits) for a graph with 1800 nodes is just under 100, while for a graph with almost four times more nodes, the number

of rounds only increases by a factor of 1.7. We also note that each neighbor-computation circuit consists of just under 50,000 non-XOR gates. In contrast, generic protocols for private navigation that construct a garbled circuit for Dijkstra’s algorithm yield circuits that contain hundreds of millions to tens of billions of non-XOR gates [CMTB13, CLT14, LWN⁺15].

Finally, we see that despite needing to pad the number of rounds to a worst-case setting, the total cost of the protocol remains modest. For the city of Los Angeles, which contains over 7000 nodes, a shortest-path query still completes in under 15 minutes and requires just over 16 MB of total bandwidth. Moreover, since the path is revealed edge-by-edge rather than only at the end of the computation, the overall protocol is an efficient solution for fully-private navigation.

Comparison to other approaches for private navigation. Many protocols [DK05, LLLZ09, XSS14] have been proposed for private navigation, but most of them rely on heuristics and do not provide strong security guarantees [DK05, LLLZ09], or guarantee privacy only for the client’s location, and not the server’s routing information [XSS14]. A different approach to fully-private navigation is to leverage generic multiparty computation techniques [Yao86, GMW87]. For instance, a generic protocol for private navigation is to construct a garbled circuit for a shortest-path algorithm and apply Yao’s protocol. This approach is quite expensive since the entire graph structure must be embedded in the circuit. For instance, Liu et al. [LWN⁺15] demonstrate that a garbled circuit for evaluating Dijkstra’s algorithm on a graph with just 1024 nodes requires over 10 *billion* AND gates. The bandwidth needed to transmit a circuit of this magnitude quickly grows to the order of GB. In contrast, even for a larger graph with 1800 nodes, the total online and offline communication required by our protocol is under 60 MB (and the online communication is under 10 MB). Carter et al. [CMTB13, CLT14] describe methods for reducing the computational and communicational cost of Yao’s protocol by introducing a third (non-colluding) party that facilitates the computation. Even with this improvement, evaluating a single shortest path on a graph of 100 nodes still requires over 10 minutes of computation. As a point of comparison, our protocols complete in around 2-3 minutes for graphs that are 15-20 times larger. Evidently, while the generic tools are powerful, they do not currently yield a practical private navigation protocol. We survey additional related work and techniques in Section 7.

6 Extensions

In this section, we describe several extensions to our protocol: supporting navigation between cities, handling updates to the routing information, and updating the source node during the protocol execution (for instance, to accommodate detours and wrong turns).

Navigating between cities. The most direct method for supporting navigation across a multi-city region is to construct a network that spans the entire region and run the protocol directly. However, since the server’s computation in the PIR protocols grows as $O(nd \log p)$, where n is the number of nodes in the graph, d is the number of columns in the compressed representation, and p is the order of the finite field used for the affine encodings, this can quickly become computationally infeasible for the server.

An alternative method that provides a performance-privacy trade-off is to introduce a series of publicly-known waypoints for each city. For example, suppose a user is navigating from somewhere in Los Angeles to somewhere in San Diego. In this case, the user would first make a private routing request to learn the fastest route from her current location to a waypoint in Los Angeles. Once the user arrives at the waypoint in Los Angeles, she requests the fastest route to a waypoint in San Diego. This second query is performed entirely in the clear, so the user reveals to the server that she is traveling from Los Angeles to San Diego. Once the user arrives at a waypoint in San Diego, she makes a final private routing request

to learn the fastest route to her destination. In this solution, the server only obtains a macro-view of the user’s location: it learns only the user’s source and destination cities, and no information about the user’s particular location within the city. As we have demonstrated, the protocol in Figure 3 is able to handle real-time navigation for a single city; thus, using this method of waypoints, we can also apply our protocol to navigation between cities with limited privacy loss.

Live updates to routing information. Routing information in road networks is dynamic, and is influenced by current traffic conditions, weather conditions, and other external factors. Ideally, the edges revealed in an iterative shortest-path protocol should always correspond to the shortest path to the destination given the current network conditions. It is fairly straightforward to allow for updates to the routing information in our protocol. Specifically, we observe that the compressed routing matrices $A^{(NE)}, A^{(NW)}, B^{(NE)}, B^{(NW)}$ need not be fixed for the duration of the protocol. As long as the total number of columns d , the bound on the bit-length τ of the values in the matrix products $A^{(NE)} \cdot (B^{(NE)})^T$ and $A^{(NW)} \cdot (B^{(NW)})^T$, and the total number of rounds R in the protocol remain fixed, the server can use a different set of routing matrices on each round of the protocol. Therefore, we can accommodate live updates to the routing information during the protocol execution by simply setting a conservative upper bound on the parameters d, τ, R . Note that we can always pad a routing matrix with fewer than d columns to one with exactly d columns by adding columns where all entries are 0. Since computing the shortest path information for a city-wide network and compressing the resulting routing matrices completes in just a few minutes, it is possible to ensure accurate and up-to-date routing information in practice.

Updating sources and destinations. Typically, in navigation, the user might take a detour or a wrong turn. While the protocol is designed to constrain the client to learn a single contiguous route through the network, it is possible to provide a functionality-privacy trade-off to accommodate deviations from the actual shortest path. One method is to introduce an additional parameter K , such that after every K iterations of the protocol, the server chooses fresh source keys for the next round of the protocol. After every K rounds, the client would also OT for a new source key. Effectively, we are resetting the protocol every K rounds and allowing the client to choose a new source from which to navigate. Correspondingly, we would need to increase the total number of rounds R in order to support the potential for detours and wrong turns. Though we cannot directly bound the number of rounds R , we can use a conservative estimate. Of course, a dishonest client can now learn multiple sub-paths to its chosen destination, namely, one sub-path each time it is allowed to choose a different source. In a similar manner, we can support updates to the destination.

7 Related Work

Numerous approaches have been proposed for private shortest path computation [DK05, LLLZ09, MY12, Mou13, XSS14, CMTB13, CLT14, BSA13, WNL⁺14, KS14, LWN⁺15]. Early works such as [DK05, LLLZ09] propose hiding the client’s location by either providing approximate locations to the server [DK05] or by having the client submit dummy sources and destinations with each shortest path query [LLLZ09]. However, these approaches only provide limited privacy for the client’s location. Later works [Mou13, MY12, XSS14] describe PIR-based solutions for hiding the client’s location. In [Mou13, MY12], the client first privately retrieves subregions of the graph that are relevant to its query [Mou13, MY12] and then locally computes the shortest path over the subgraph. In [XSS14], the client privately requests for columns of the next-hop routing matrix to learn the next hop in the shortest path. While these methods provide privacy for the client’s location, they do not hide the server’s routing information.

There is also work on developing secure protocols for other graph-theoretic problems and under different models [BS05, FG06]. For example, Brickell and Shmatikov [BS05] consider a model where two parties hold a graph over a common set of vertices, and the goal is to compute a function over their joint graphs. Their protocols do not extend to navigation protocols where one party holds the full graph, and only the client should learn the result of the computation. In [FG06], the authors describe protocols for parties who each hold a subset of a graph to privately reconstruct the joint graph. Their methods are designed for social network analysis and do not directly apply to private navigation.

Another line of work has focused on developing data-oblivious algorithms for shortest path computation [BSA13] or combining shortest path algorithms such as Dijkstra’s with oblivious data structures or ORAM [WNL⁺14, KS14]. In these methods, the routing data is stored in an ORAM or an oblivious data structure on the server. The client then executes the shortest-path algorithm on the server to learn the path between its source and destination. Since the pattern of memory accesses is hidden from the server, these approaches provide client privacy. While these protocols can be efficient in practice, they do not provide security against a malicious client trying to learn additional details about the routing information on the server. Thus, for scenarios where the map data is proprietary (for instance, in the case of real-time traffic routing), or when the routing information itself is sensitive (for instance, when providing navigational assistance for a presidential motorcade or coordinating troop movements in a military scenario [CMTB13, CLT14]), the ORAM-based solutions do not provide sufficient security.

Also relevant are the works in secure multiparty computation (MPC) [Yao86, GMW87]. While these methods can be successfully used to build private navigation protocols [LWN⁺15, CMTB13, CLT14], they do not currently yield a practical private navigation protocol. A more comprehensive comparison of our protocol to these generic methods is provided at the end of Section 5.3.

There is also a vast literature on graph compression algorithms. For planar graphs, there are multiple methods based on computing graph separators [LT79, BBK03, BBK04]. Other methods based on coding schemes [HKL00] have also been proposed and shown to achieve information-theoretically optimal encoding. While these algorithms are often viable in practice, it is not straightforward to represent them compactly as a Boolean or an arithmetic circuit. Thus, it is unclear how to combine them with standard cryptographic primitives to construct a private shortest path protocol.

Finally, there has also been work on developing compact representations of graphs for answering *approximate* distance queries in graphs [TZ01]. These techniques have been successfully applied for privacy-preserving approximate distance computation on graphs [MKNK15]. However, these distance-oracle-based methods only provide an estimate on the *length* of the shortest path, and do not give a private navigation protocol.

8 Conclusion

In this work, we constructed an efficient protocol for privately computing shortest paths for navigation. First, we developed a method for compressing the next-hop matrices for road networks by formulating the compression problem as that of finding a sign-preserving, low-rank matrix decomposition. Not only did this method yield a significant compression, it also enabled an efficient cryptographic protocol for fully private shortest-path computation in road networks. By combining affine encodings with Yao’s circuits, we obtained a fully-private navigation protocol efficient enough to run at a city-scale.

Acknowledgments

The authors would like to thank Dan Boneh, Roy Frostig, Hristo Paskov, and Madeleine Udell for many helpful comments and discussions. While conducting this work, authors David Wu and Joe Zimmerman

were supported by NSF Graduate Research Fellowships. This work was further supported by the DARPA PROCEED research program. Opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or NSF.

References

- [ADH⁺01] David Ascher, Paul F. Dubois, Konrad Hinsen, Jim Hugunin, and Travis Oliphant. Numerical python. Technical report, Lawrence Livermore National Laboratory, 2001.
- [AIK14] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. How to garble arithmetic circuits. *SIAM J. Comput.*, 43(2):905–929, 2014.
- [AVD11] Julia Angwin and Jennifer Valentino-Devries. Apple, Google collect user data. *The Wall Street Journal*, 2011.
- [BBK03] Daniel K. Blandford, Guy E. Blelloch, and Ian A. Kash. Compact representations of separable graphs. In *SODA*, pages 679–688, 2003.
- [BBK04] Daniel K. Blandford, Guy E. Blelloch, and Ian A. Kash. An experimental analysis of a compact graph representation. In *Workshop on Analytic Algorithmics and Combinatorics*, pages 49–61, 2004.
- [BCLN14] Joppe Bos, Craig Costello, Patrick Longa, and Michael Naehrig. Specification of curve selection and supported curve parameters in MSR ECCLib. Technical Report MSR-TR-2014-92, Microsoft Research, June 2014.
- [BHKR13] Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. Efficient garbling from a fixed-key blockcipher. In *IEEE Symposium on Security and Privacy*, pages 478–492, 2013.
- [BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. Cryptology ePrint Archive, Report 2012/265, 2012. <http://eprint.iacr.org/>.
- [BLNZ95] Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM J. Comput.*, 16(5):1190–1208, 1995.
- [BS05] Justin Brickell and Vitaly Shmatikov. Privacy-preserving graph algorithms in the semi-honest model. In *ASIACRYPT*, pages 236–252, 2005.
- [BSA13] Marina Blanton, Aaron Steele, and Mehrdad Aliasgari. Data-oblivious graph algorithms for secure computation and outsourcing. In *ASIA CCS*, pages 207–218, 2013.
- [Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
- [Can06] Ran Canetti. Security and composition of cryptographic protocols: a tutorial (part I). *SIGACT News*, 37(3):67–92, 2006.
- [CGKS95] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *FOCS*, pages 41–50, 1995.
- [Cha04] Yan-Cheng Chang. Single database private information retrieval with logarithmic communication. In *ACISP*, pages 50–61, 2004.

- [Che11] Jacqui Cheng. How Apple tracks your location without consent, and why it matters. Ars Technica, 2011.
- [CLT14] Henry Carter, Charles Lever, and Patrick Traynor. Whitewash: outsourcing garbled circuit generation for mobile devices. In *ACSAC*, pages 266–275, 2014.
- [CMS99] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In *EUROCRYPT*, pages 402–414, 1999.
- [CMTB13] Henry Carter, Benjamin Mood, Patrick Traynor, and Kevin R. B. Butler. Secure outsourced garbled circuit evaluation for mobile devices. In *USENIX*, pages 289–304, 2013.
- [Dij59] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.
- [DK05] Matt Duckham and Lars Kulik. A formal model of obfuscation and negotiation for location privacy. In *PERVASIVE*, pages 152–170, 2005.
- [FG06] Keith B. Frikken and Philippe Golle. Private social network analysis: how to assemble pieces of a graph privately. In *WPES*, pages 89–98, 2006.
- [FS86] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- [GIKM00] Yael Gertner, Yuval Ishai, Eyal Kushilevitz, and Tal Malkin. Protecting data privacy in private information retrieval schemes. *J. Comput. Syst. Sci.*, 60(3):592–629, 2000.
- [GKP⁺13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 555–564, 2013.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *J. ACM*, 43(3):431–473, 1996.
- [GR05] Craig Gentry and Zulfikar Ramzan. Single-database private information retrieval with constant communication rate. In *ICALP*, pages 803–815, 2005.
- [Gt12] Torbjörn Granlund and the GMP development team. *GNU MP: The GNU Multiple Precision Arithmetic Library*, 5.0.5 edition, 2012. <http://gmplib.org/>.
- [HKL00] Xin He, Ming-Yang Kao, and Hsueh-I Lu. A fast general methodology for information-theoretically optimal encodings of graphs. *SIAM J. Comput.*, 30(3):838–846, 2000.
- [HL10] Carmit Hazay and Yehuda Lindell. *Efficient Secure Two-Party Protocols - Techniques and Constructions*. Information Security and Cryptography. Springer, 2010.

- [IKK⁺11] Yuval Ishai, Jonathan Katz, Eyal Kushilevitz, Yehuda Lindell, and Erez Petrank. On achieving the “best of both worlds” in secure multiparty computation. *SIAM J. Comput.*, 40(1):122–141, 2011.
- [Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In *STOC*, pages 20–31, 1988.
- [KO97] Eyal Kushilevitz and Rafail Ostrovsky. Replication is NOT needed: SINGLE database, computationally-private information retrieval. In *FOCS*, pages 364–373, 1997.
- [KS08] Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In *ICALP*, pages 486–498, 2008.
- [KS14] Marcel Keller and Peter Scholl. Efficient, oblivious data structures for MPC. In *ASIACRYPT*, pages 506–525, 2014.
- [KSS09] Vladimir Kolesnikov, Ahmad-Reza Sadeghi, and Thomas Schneider. Improved garbled circuit building blocks and applications to auctions and computing minima. Cryptology ePrint Archive, Report 2009/411, 2009. <http://eprint.iacr.org/>.
- [KY55] H. W. Kuhn and Bryn Yaw. The Hungarian method for the assignment problem. *Naval Res. Logist. Quart.*, pages 83–97, 1955.
- [Lip05] Helger Lipmaa. An oblivious transfer protocol with log-squared communication. In *ISC*, pages 314–328, 2005.
- [LLLZ09] Ken C. K. Lee, Wang-Chien Lee, Hong Va Leong, and Baihua Zheng. Navigational path privacy protection: navigational path privacy protection. In *CIKM*, pages 691–700, 2009.
- [LP09] Yehuda Lindell and Benny Pinkas. A proof of security of Yao’s protocol for two-party computation. *J. Cryptology*, 22(2):161–188, 2009.
- [LT79] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, (2):177–189, 1979.
- [LWN⁺15] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. OblivM: A programming framework for secure computation. In *IEEE Symposium on Security and Privacy*, 2015.
- [MBFK14] Carlos Aguilar Melchor, Joris Barrier, Laurent Fousse, and Marc-Olivier Killijian. Xpire: Private information retrieval for everyone. *IACR Cryptology ePrint Archive*, 2014:1025, 2014.
- [MKNK15] Xianrui Meng, Seny Kamara, Kobbi Nissim, and George Kollios. GRECS: graph encryption for approximate shortest distance queries. *IACR Cryptology ePrint Archive*, 2015:266, 2015.
- [Mou13] Kyriakos Mouratidis. Strong location privacy: A case study on shortest path queries. In *ICDE*, pages 136–143, 2013.
- [MY12] Kyriakos Mouratidis and Man Lung Yiu. Shortest path computation with no information leakage. *PVLDB*, 5(8):692–703, 2012.
- [NP99] Moni Naor and Benny Pinkas. Oblivious transfer and polynomial evaluation. In *STOC*, pages 245–254, 1999.

- [NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In *SODA*, pages 448–457, 2001.
- [NP05] Moni Naor and Benny Pinkas. Computationally secure oblivious transfer. *J. Cryptology*, 18(1):1–35, 2005.
- [OI07] Rafail Ostrovsky and William E. Skeith III. A survey of single-database private information retrieval: Techniques and applications. In *Public Key Cryptography*, pages 393–411, 2007.
- [Ope] OpenStreetMap Contributors. OpenStreetMap. <http://www.openstreetmap.org/>.
- [Pai99] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, pages 223–238. 1999.
- [PSSW09] Benny Pinkas, Thomas Schneider, Nigel P. Smart, and Stephen C. Williams. Secure two-party computation is practical. In *ASIACRYPT*, pages 250–267, 2009.
- [Rab05] Michael O. Rabin. How to exchange secrets with oblivious transfer. *IACR Cryptology ePrint Archive*, 2005:187, 2005.
- [RVC⁺04] Lorenzo Rosasco, Ernesto De Vito, Andrea Caponnetto, Michele Piana, and Alessandro Verri. Are loss functions all the same? *Neural Computation*, 16(5):1063–107, 2004.
- [Sch89] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In *CRYPTO*, pages 239–252, 1989.
- [Sho] Victor Shoup. NTL: A library for doing number theory. <http://www.shoup.net/ntl/>.
- [Sta04] Cyrill Stachniss. C implementation of the Hungarian method. <http://www2.informatik.uni-freiburg.de/~stachnis/misc.html>, 2004.
- [SvDS⁺13] Emil Stefanov, Marten van Dijk, Elaine Shi, Christopher W. Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. Path ORAM: an extremely simple oblivious RAM protocol. In *CCS*, pages 299–310, 2013.
- [The03] The OpenSSL Project. OpenSSL: The open source toolkit for SSL/TLS. www.openssl.org, April 2003.
- [TZ01] Mikkel Thorup and Uri Zwick. Approximate distance oracles. In *STOC*, pages 183–192, 2001.
- [WNL⁺14] Xiao Shaun Wang, Kartik Nayak, Chang Liu, T.-H. Hubert Chan, Elaine Shi, Emil Stefanov, and Yan Huang. Oblivious data structures. In *CCS*, pages 215–226, 2014.
- [XSS14] Yong Xi, Loren Schwiebert, and Weisong Shi. Privacy preserving shortest path routing with an application to navigation. *Pervasive and Mobile Computing*, 13:142–149, 2014.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167, 1986.
- [Zha04] Tong Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *ICML*, 2004.

A Security Proofs

In this section, we show that the protocol in Figure 3 securely computes the shortest paths functionality in the presence of a malicious client, and provides privacy against a malicious server. To simplify our proofs, we work in the OT-hybrid model where we assume the parties have access to an ideal 1-out-of- n OT functionality [Kil88]. Specifically, in the real protocol, we replace every OT invocation with an oracle call to a trusted party that implements the OT functionality: the sender sends the database of records (r_1, \dots, r_n) to the trusted party and the receiver sends an index $i \in [n]$ to the trusted party. The trusted party then gives the receiver the record r_i . Security in the standard model then follows by instantiating the ideal OT functionality with an OT protocol that provides security against malicious clients [HL10] and privacy against malicious servers, and then invoking the sequential composition theorem of [Can00].

A.1 Proof of Theorem 4.6

At a high level, privacy for the client’s location follows from the fact that the server’s view in the protocol execution consists only of its view in the OT and PIR protocols. By assumption, both the OT protocols and the PIR protocols provide privacy for the client’s input, so the claim follows. We now show this formally. As noted at the beginning of Appendix A, we work in the OT-hybrid model, where we replace each OT invocation with an oracle call to an ideal OT functionality. First, we state the definition of privacy as it applies to the PIR protocol.

Definition A.1 (Privacy for PIR). Fix a security parameter $\lambda \in \mathbb{N}$, and let π be a PIR protocol. Let \mathcal{A} be a non-uniform PPT server for π . Let $\text{view}_{\pi, \mathcal{A}}(1^\lambda, \mathcal{D}, i)$ denote the view of adversary \mathcal{A} in the PIR protocol on database $\mathcal{D} \in \{0, 1\}^*$ (the server’s input) and index $i \in \{0, 1\}^*$ (the client’s input). Then, π is a *private* PIR protocol if for all non-uniform PPT servers \mathcal{A} , databases $\mathcal{D} \in \{0, 1\}^*$, indices $i, i' \in \{0, 1\}^*$ where $|i| = |i'|$, we have that

$$\text{view}_{\pi, \mathcal{A}}(1^\lambda, \mathcal{D}, i) \stackrel{c}{\approx} \text{view}_{\pi, \mathcal{A}}(1^\lambda, \mathcal{D}, i').$$

Let \mathcal{A} be a malicious server for the private shortest paths protocol in Figure 3. We construct an ideal-world simulator \mathcal{S} such that the distribution of outputs of \mathcal{A} in the real protocol is computationally indistinguishable from the outputs of \mathcal{S} in the ideal world. This suffices to prove that the protocol in Figure 3 provides *privacy* against a malicious server. The simulator \mathcal{S} begins by running \mathcal{A} . Below, we describe how \mathcal{S} simulates the view for \mathcal{A} in the protocol execution.

Setup. In the setup phase of the protocol, the simulator does nothing. This is the correct behavior because in the OT-hybrid model, the adversary \mathcal{A} does not receive any messages during the setup phase.

Round. On each round of the protocol, the simulator \mathcal{S} plays the role of the client in the PIR protocol and requests for record 0 in both the source database and in the destination database. Again, since we are working in the OT-hybrid model, these are the only messages adversary \mathcal{A} obtains in the real protocol. At the end of the protocol execution, adversary \mathcal{A} will output some function of its view of the protocol execution. The simulator \mathcal{S} echoes this output to the environment.

Correctness of the simulation. To conclude the proof, it suffices to show that the view \mathcal{S} simulates for \mathcal{A} is computationally indistinguishable from the view \mathcal{A} expects in the real protocol. This condition holds vacuously in the setup phase of the protocol. Let $\text{view}_{\mathcal{A}}^{(r)}$ be the adversary’s view on the r^{th} round of the protocol. The view $\text{view}_{\mathcal{A}}^{(r)}$ may be written as $\text{view}_{\mathcal{A}}^{(r)} = \left\{ \text{view}_{\text{PIR}, \mathcal{A}}^{(r)}(1^\lambda, \mathcal{D}_{\text{src}}, i_{\text{src}}), \text{view}_{\text{PIR}, \mathcal{A}}^{(r)}(1^\lambda, \mathcal{D}_{\text{dst}}, i_{\text{dst}}) \right\}$,

where $\mathcal{D}_{\text{src}}, \mathcal{D}_{\text{dst}}$ are the encoding databases \mathcal{A} chooses in the real protocol, and $i_{\text{src}}, i_{\text{dst}}$ are the indices of the records the client chooses in the real protocol. By privacy of the PIR (Definition A.1), it follows that

$$\begin{aligned} \text{view}_{\text{PIR}, \mathcal{A}}^{(r)}(1^\lambda, \mathcal{D}_{\text{src}}, i_{\text{src}}) &\stackrel{c}{\approx} \text{view}_{\text{PIR}, \mathcal{A}}^{(r)}(1^\lambda, \mathcal{D}_{\text{src}}, 0) \\ \text{view}_{\text{PIR}, \mathcal{A}}^{(r)}(1^\lambda, \mathcal{D}_{\text{dst}}, i_{\text{dst}}) &\stackrel{c}{\approx} \text{view}_{\text{PIR}, \mathcal{A}}^{(r)}(1^\lambda, \mathcal{D}_{\text{dst}}, 0). \end{aligned}$$

Since the request for the source encodings and for the destination encodings constitute two independent instances of the PIR protocol, we conclude that the view \mathcal{S} simulates for \mathcal{A} in each round of the protocol is computationally indistinguishable from the view \mathcal{A} expects in the real protocol. Thus, the output of \mathcal{S} in the ideal-world execution is computationally indistinguishable from that of \mathcal{A} in the real world. \square

A.2 Proof of Theorem 4.5

Before we prove Theorem 4.5, we describe the simulatability requirement we require on the garbled circuit encodings used in the protocol in Figure 3. Intuitively, we require that the garbled circuit encodings can be entirely simulated given the output of the computation; that is, the garbled circuit together with one set of encodings do not reveal any information about the underlying inputs other than what is explicitly revealed by the output. Bellare et al. formalize this notion in [BHR12]. Here, we give a simplified definition adapted from [GKP⁺13] and specialized to the case of Yao’s garbling scheme [Yao86, LP09].

Definition A.2 (Yao Garbling Scheme). A Yao garbling scheme Π_{Yao} for a family of circuits $\{\mathcal{C}_n\}_{n \in \mathbb{N}}$ (where \mathcal{C}_n is a set of Boolean circuits on n input bits) consists of three algorithms (Yao.Garble, Yao.Encode, Yao.Eval) where

- Yao.Garble($1^\lambda, C$) is a randomized algorithm that takes as input a security parameter λ and a circuit $C \in \mathcal{C}_n$ for some n and outputs a garbled circuit \tilde{C} along with a secret key $\text{sk} = \{L_i^0, L_i^1\}_{i \in [n]}$ is a set containing n pairs of encodings $L_i^0, L_i^1 \in \{0, 1\}^*$.
- Yao.Encode(sk, x) is a deterministic algorithm that takes the secret key $\text{sk} = \{L_i^0, L_i^1\}_{i \in [n]}$ and an input $x = x_1 \cdots x_n \in \{0, 1\}^n$, and outputs an encoding $\tilde{x} = \{L_i^{x_i}\}_{i \in [n]}$. Specifically, the encoding \tilde{x} of x is the subset of encodings in sk associated with the bits of x .
- Yao.Eval(\tilde{C}, \tilde{x}) is a deterministic algorithm that takes a garbled circuit \tilde{C} and a set of encodings $\tilde{x} = \{L_i^{x_i}\}_{i \in [n]}$ for some $x \in \{0, 1\}^n$ and outputs a value z .

Definition A.3 (Correctness). A Yao garbling scheme $\Pi_{\text{Yao}} = (\text{Yao.Garble}, \text{Yao.Encode}, \text{Yao.Eval})$ for a family of circuits $\{\mathcal{C}_n\}_{n \in \mathbb{N}}$ is correct if for all $n = \text{poly}(\lambda)$, $C \in \mathcal{C}_n$, and $x \in \{0, 1\}^n$, the following holds. Letting $(\tilde{C}, \text{sk}) \leftarrow \text{Yao.Garble}(1^\lambda, C)$, then with overwhelming probability in λ , we have

$$\text{Yao.Eval}(\tilde{C}, \text{Yao.Encode}(\text{sk}, x)) = C(x),$$

where the probability is taken over the random coins used in Yao.Garble.

Definition A.4 (Input Privacy). A Yao garbling scheme $\Pi_{\text{Yao}} = (\text{Yao.Garble}, \text{Yao.Encode}, \text{Yao.Eval})$ for a family of circuits $\{\mathcal{C}_n\}_{n \in \mathbb{N}}$ is input-private if there exists a PPT simulator \mathcal{S}_{Yao} such that for all $n = \text{poly}(\lambda)$, $C \in \mathcal{C}_n$, $x \in \{0, 1\}^n$, the following holds:

$$\left\{ (\tilde{C}, \text{sk}) \leftarrow \text{Yao.Garble}(1^\lambda, C) ; (\tilde{C}, \text{Yao.Encode}(\text{sk}, x)) \right\} \stackrel{c}{\approx} \mathcal{S}_{\text{Yao}}(1^\lambda, C, C(x)).$$

Lemma A.5 ([Yao86, LP09]). *Assuming one-way functions exist, there exists a Yao garbling scheme (Definition A.2) that is input-private (Definition A.4).*

Next, we note that the affine encodings from Section 4, Eq. (4) provide statistical privacy.

Lemma A.6 ([AIK14, Lemma 5.1], adapted). *Fix a finite field \mathbb{F}_p of prime order p , and take $z_1, \dots, z_d \in \mathbb{F}_p$. Define the function $f : \mathbb{F}_p^d \times \mathbb{F}_p^d \rightarrow \mathbb{F}_p$ where $f(x, y) = \langle x, y \rangle + \sum_{i \in [d]} z_i$. Let L_x^{affine} and L_y^{affine} be the affine encoding functions from Eq. (6). Then, there exists a PPT simulator \mathcal{S}_{ac} such that for all $x, y \in \mathbb{F}_p^d$ and $z_1, \dots, z_d \in \mathbb{F}_p$,*

$$\mathcal{S}_{\text{ac}}(f(x, y)) \equiv \left\{ r \stackrel{\text{R}}{\leftarrow} \mathbb{F}_p^{3d} ; \left(L_x^{\text{affine}}(x; r), L_y^{\text{affine}}(y; r) \right) \right\}.$$

Proof of Theorem 4.5. To show Theorem 4.5, we first define the following hybrid experiments:

- Hybrid Hyb_0 : This is the real experiment (Definition 4.1).
- Hybrid Hyb_1 : Same as Hyb_0 , except the protocol execution aborts if the client succeeds in making an “inconsistent” query (described below).
- Hybrid Hyb_2 : This is the ideal experiment (Definition 4.2).

Informally speaking, we say that the client succeeds in making an “inconsistent” query if on some round r , it requests the garbled circuit encodings for values \hat{z}_{NE} and \hat{z}_{NW} that were not the outputs of the arithmetic circuit, and yet, the client obtains a set of garbled circuit encodings where the garbled circuit evaluation does not output \perp . We now specify this property more precisely.

Specification of Hybrid Hyb_1 . Let $s, t \in [n]$ be the source and destination nodes the client sends to the ideal OT functionality in the setup phase of the protocol in Figure 3. Let $(s = v_0, v_1, \dots, v_R)$ be the shortest path from s to t as defined by the environment’s choice of the next-hop routing matrices $A^{(\text{NE})}, B^{(\text{NE})}, A^{(\text{NW})}, B^{(\text{NW})}$ in the protocol execution (Definition 4.1). The protocol execution in Hyb_1 proceeds identically to that in Hyb_0 , except the protocol execution halts (with output \perp) if the following bad event occurs:

On a round $1 \leq r \leq R$, the client submits $\hat{z}_{\text{NE}}, \hat{z}_{\text{NW}}$ to the ideal OT functionality where either

$$\hat{z}_{\text{NE}} \neq \alpha_{\text{NE}} \langle A_{v_r}^{(\text{NE})}, B_t^{(\text{NE})} \rangle + \beta_{\text{NE}} \quad \text{or} \quad \hat{z}_{\text{NW}} \neq \alpha_{\text{NW}} \langle A_{v_r}^{(\text{NW})}, B_t^{(\text{NW})} \rangle + \beta_{\text{NW}},$$

and $C^{\text{unblind}}((\hat{z}_{\text{NE}}, \gamma_{\text{NE}}, \delta_{\text{NE}}), (\hat{z}_{\text{NW}}, \gamma_{\text{NW}}, \delta_{\text{NW}}), k_{\text{NE}}^0, k_{\text{NE}}^1, k_{\text{NW}}^0, k_{\text{NW}}^1, v_r, t) \neq \perp$, where all values other than $\hat{z}_{\text{NE}}, \hat{z}_{\text{NW}}, v_r$, and t are the round-specific values chosen by the server on round r .

Figure 5: Abort event in hybrid Hyb_1 .

Fix a security parameter λ . Let π be a private navigation protocol and let f be the ideal shortest-path functionality. For a client \mathcal{A} , a simulator \mathcal{S} and an environment \mathcal{E} , we define the following random variables:

- $\text{Hyb}_0(\lambda, \pi, \mathcal{A}, \mathcal{E})$ is the output of experiment Hyb_0 with adversary \mathcal{A} and environment \mathcal{E} . In particular, $\text{Hyb}_0(\lambda, \pi, \mathcal{A}, \mathcal{E}) = \text{REAL}_{\pi, \mathcal{A}, \mathcal{E}}(\lambda)$ (Definition 4.1).
- $\text{Hyb}_1(\lambda, \pi, \mathcal{A}, \mathcal{E})$ is the output of experiment Hyb_1 with adversary \mathcal{A} and environment \mathcal{E} .
- $\text{Hyb}_2(\lambda, f, \mathcal{S}, \mathcal{E})$ is the output of experiment Hyb_2 with simulator \mathcal{S} and environment \mathcal{E} . In particular, $\text{Hyb}_2(\lambda, f, \mathcal{S}, \mathcal{E}) = \text{IDEAL}_{f, \mathcal{S}, \mathcal{E}}(\lambda)$ (Definition 4.2).

To prove Theorem 4.5, we show the following two claims.

Claim A.7. Let λ, μ be the security parameter and statistical security parameter, respectively. Let π be the protocol in Figure 3 instantiated with secure cryptographic primitives as described in Theorem 4.5. Then, for all PPT adversaries \mathcal{A} , and every polynomial-size circuit family $\mathcal{E} = \{\mathcal{E}\}_\lambda$,

$$|\Pr[\text{Hyb}_0(\lambda, \pi, \mathcal{A}, \mathcal{E}) = 0] - \Pr[\text{Hyb}_1(\lambda, \pi, \mathcal{A}, \mathcal{E}) = 0]| \leq \text{negl}(\lambda) + R \cdot 2^{-\mu}.$$

Claim A.8. Let λ, μ be the security parameter and statistical security parameter, respectively. Let π be the protocol in Figure 3 instantiated with secure cryptographic primitives as described in Theorem 4.5. Let f be the ideal shortest-paths functionality. Then, for all PPT adversaries \mathcal{A} , there exists a PPT adversary \mathcal{S} such that for every polynomial-size circuit family $\mathcal{E} = \{\mathcal{E}\}_\lambda$,

$$|\Pr[\text{Hyb}_1(\lambda, \pi, \mathcal{A}, \mathcal{E}) = 0] - \Pr[\text{Hyb}_2(\lambda, f, \mathcal{S}, \mathcal{E}) = 0]| \leq \text{negl}(\lambda).$$

Proof of Claim A.8. We begin by showing Claim A.8. Given a real-world adversary \mathcal{A} , we construct an efficient ideal-world simulator \mathcal{S} such that the distribution of outputs of any (possibly malicious) client \mathcal{A} in Hyb_1 is computationally indistinguishable from the distribution of outputs of an ideal-world simulator \mathcal{S} in Hyb_2 . Since the server does not produce any output, this will suffice to show that hybrids Hyb_1 and Hyb_2 are computationally indistinguishable.

As stated in Section 4.2 and Figure 3, we assume that the topology of the graph $\mathcal{G} = (V, E)$, the number of columns d in the compressed routing matrices, the bound τ on the bit-length of the products of the compressed matrices, and the total number of rounds R are public and known to both parties in the protocol execution. Moreover, as described in Figure 3, we assume that $p > 2^{\tau+\mu+1}$ where $\mu \in \mathbb{N}$ is the statistical security parameter. In particular, there exists an element $\xi \in \mathbb{F}_p$ such that $\xi \notin [-2^\tau, 2^\tau]$.

Specification of the simulator. We begin by describing the simulator. We let \mathcal{T} denote the trusted party for the shortest path functionality as defined in the specification of the ideal model of execution from Section 4.2. The simulator starts running adversary \mathcal{A} . We describe the behavior of the simulator in the setup phase of the protocol as well as the behavior on each round of the routing protocol.

Setup. In the setup phase of the protocol, the simulator \mathcal{S} does the following:

1. As in the real protocol, the simulator first chooses independent symmetric encryption keys $\bar{k}_{\text{src},i}^{(1)}, \bar{k}_{\text{dst},i} \xleftarrow{R} \{0,1\}^\ell$ for all $i \in [n]$.
2. When \mathcal{A} makes an OT request for an entry s in the source key database, the simulator replies with the key $\bar{k}_{\text{src},s}^{(1)}$. Recall that in the ideal OT functionality, each party just sends its input to the ideal functionality, and the ideal functionality sends the requested element to the client.
3. When \mathcal{A} makes an OT request for an entry t in the destination key database, the simulator replies with the key $\bar{k}_{\text{dst},t}$ to \mathcal{A} .
4. Finally, the simulator sends (s, t) to the trusted party \mathcal{T} . The trusted party replies to the simulator with a path $(s = v_0, \dots, v_R)$.

Round. Next, we describe the behavior of the simulator in each round $1 \leq r \leq R$ of the protocol.

1. The simulator chooses $\bar{z}_{\text{NE}}, \bar{z}_{\text{NW}} \xleftarrow{R} \mathbb{F}_p$. Then, \mathcal{S} invokes the simulator \mathcal{S}_{ac} (Lemma A.6) on inputs \bar{z}_{NE} and \bar{z}_{NW} to obtain four sets of encodings $\bar{L}_{\text{NE},x}^{\text{affine}}, \bar{L}_{\text{NE},y}^{\text{affine}}, \bar{L}_{\text{NW},x}^{\text{affine}}, \bar{L}_{\text{NW},y}^{\text{affine}}$.

2. Let C^{unblind} be a circuit computing the neighbor-computation function in Figure 2. As in the real protocol, the simulator runs Yao's garbling algorithm on C^{unblind} to obtain a garbled circuit \bar{C}^{unblind} , along with encoding functions $\bar{L}_x^{\text{unblind}}$ for each of the inputs x to the neighbor-computation function in Figure 2.
3. As in the real protocol, the simulator chooses symmetric encryption keys $\bar{k}_{\text{src},i}^{(r+1)} \xleftarrow{R} \{0,1\}^\ell$ for all $i \in [n]$. These will be used to encrypt the elements in the source database on the next round of the protocol.
4. The simulator also chooses four PRF keys $\bar{k}_{\text{NE}}^0, \bar{k}_{\text{NE}}^1, \bar{k}_{\text{NW}}^0, \bar{k}_{\text{NW}}^1 \xleftarrow{R} \{0,1\}^\rho$, two for each axis NE, NW. As in the real scheme, the simulator defines the encryption keys for each direction as follows:

$$\begin{aligned} \bar{k}_{\text{N}} &= F(\bar{k}_{\text{NE}}^0, \text{N}) \oplus F(\bar{k}_{\text{NW}}^0, \text{N}) & \bar{k}_{\text{E}} &= F(\bar{k}_{\text{NE}}^0, \text{E}) \oplus F(\bar{k}_{\text{NW}}^1, \text{E}) \\ \bar{k}_{\text{S}} &= F(\bar{k}_{\text{NE}}^1, \text{S}) \oplus F(\bar{k}_{\text{NW}}^1, \text{S}) & \bar{k}_{\text{W}} &= F(\bar{k}_{\text{NE}}^1, \text{W}) \oplus F(\bar{k}_{\text{NW}}^0, \text{W}). \end{aligned}$$

5. The simulator prepares the source database $\overline{\mathcal{D}}_{\text{src}}$ as follows. Let $u = v_{r-1}$. If $u \neq \perp$, then the u^{th} record in $\overline{\mathcal{D}}_{\text{src}}$ is an encryption under $\bar{k}_{\text{src},u}^{(r)}$ of the following:
 - Arithmetic encodings $\bar{L}_s^{\text{affine}} = (\bar{L}_{\text{NE},x}^{\text{affine}}, \bar{L}_{\text{NW},x}^{\text{affine}})$.
 - Garbled circuit encodings $\bar{L}_s^{\text{unblind}}(u)$.
 - Encryptions of the source keys for the neighbors of u in the next round of the protocol under the direction keys:

$$\text{Enc}(\bar{k}_{\text{N}}, \bar{k}_{\text{src},v_{\text{N}}}^{(r+1)}), \quad \text{Enc}(\bar{k}_{\text{E}}, \bar{k}_{\text{src},v_{\text{E}}}^{(r+1)}), \quad \text{Enc}(\bar{k}_{\text{S}}, \bar{k}_{\text{src},v_{\text{S}}}^{(r+1)}), \quad \text{Enc}(\bar{k}_{\text{W}}, \bar{k}_{\text{src},v_{\text{W}}}^{(r+1)}),$$

where $v_{\text{N}}, v_{\text{E}}, v_{\text{S}}, v_{\text{W}}$ is the neighbor of u in \mathcal{G} to the north, east, south, or west, respectively. If u does not have a neighbor in a given direction $j \in \{\text{N}, \text{E}, \text{S}, \text{W}\}$, then define $\bar{k}_{\text{src},v_j}^{(r+1)}$ to be the all-zeroes string in $\{0,1\}^\ell$.

For all nodes $u \neq v_{r-1}$, the simulator sets the u^{th} record in $\overline{\mathcal{D}}_{\text{src}}$ to be an encryption of the all-zeroes string under $\bar{k}_{\text{src},u}^{(r+1)}$.

6. The simulator prepares the destination database $\overline{\mathcal{D}}_{\text{dst}}$ as follows. The t^{th} record in $\overline{\mathcal{D}}_{\text{dst}}$ is an encryption under $\bar{k}_{\text{dst},t}$ of the following:
 - Arithmetic encodings $\bar{L}_t^{\text{affine}} = (\bar{L}_{\text{NE},y}^{\text{affine}}, \bar{L}_{\text{NW},y}^{\text{affine}})$.
 - Garbled circuit encodings $\bar{L}_t^{\text{unblind}}(t)$.

For all nodes $u \neq t$, the simulator sets the u^{th} record in $\overline{\mathcal{D}}_{\text{dst}}$ to be an encryption of the all-zeroes string under $\bar{k}_{\text{dst},u}$.

7. When \mathcal{A} makes a PIR request for a record in the source database, the simulator plays the role of the sender in the PIR protocol using $\overline{\mathcal{D}}_{\text{src}}$ as its input database.
8. When \mathcal{A} makes a PIR request for a record in the destination database, the simulator plays the role of the sender in the PIR protocol using $\overline{\mathcal{D}}_{\text{dst}}$ as its input database.
9. When \mathcal{A} engages in OT for the garbled circuit encodings of \hat{z}_{NE} and \hat{z}_{NW} , the simulator replies with the encodings $\bar{L}_{z_{\text{NE}}}^{\text{unblind}}(\hat{z}_{\text{NE}})$ and $\bar{L}_{z_{\text{NW}}}^{\text{unblind}}(\hat{z}_{\text{NW}})$.

10. Let $\text{dir} \in \{N, E, S, W\}$ be the direction of the edge from v_{r-1} to v_r in \mathcal{G} . The simulator sets $\gamma_{NE} = 0 = \gamma_{NW}$ and δ_{NE}, δ_{NW} as follows:

- If $\hat{z}_{NE} = \bar{z}_{NE}$ and $\hat{z}_{NW} = \bar{z}_{NW}$ and $v_r \neq \perp$, then the simulator sets

$$\bar{\delta}_{NE} = \begin{cases} -1 & \text{if dir} = N \text{ or dir} = E \\ 1 & \text{if dir} = S \text{ or dir} = W \end{cases} \quad \text{and} \quad \bar{\delta}_{NW} = \begin{cases} -1 & \text{if dir} = N \text{ or dir} = W \\ 1 & \text{if dir} = S \text{ or dir} = E. \end{cases}$$

- If $\hat{z}_{NE} = \bar{z}_{NE}$ and $\hat{z}_{NW} = \bar{z}_{NW}$ and $v_r = \perp$, then the simulator sets $\bar{\delta}_{NE} = \xi = \bar{\delta}_{NW}$. Recall that $\xi \in \mathbb{F}_p$ satisfies $\xi \notin [-2^\tau, 2^\tau]$.
- If exactly one of $\hat{z}_{NE} \neq \bar{z}_{NE}$ or $\hat{z}_{NW} \neq \bar{z}_{NW}$ holds, then with probability $\varepsilon = 2^{\tau+1}/p$, the simulator aborts the simulation and outputs \perp . Otherwise, with probability $1 - \varepsilon$, the simulator sets $\bar{\delta}_{NE} = \xi = \bar{\delta}_{NW}$.
- If both $\hat{z}_{NE} \neq \bar{z}_{NE}$ and $\hat{z}_{NW} \neq \bar{z}_{NW}$, then with probability ε^2 , the simulator aborts and outputs \perp . Otherwise, with probability $1 - \varepsilon^2$, the simulator sets $\bar{\delta}_{NE} = \xi = \bar{\delta}_{NW}$.

The simulator sends to \mathcal{A} the garbled circuit \bar{C}^{unblind} , the encodings of the unblinding coefficients

$$\bar{L}_{\gamma_{NE}}^{\text{unblind}}(\bar{\gamma}_{NE}), \bar{L}_{\gamma_{NW}}^{\text{unblind}}(\bar{\gamma}_{NW}), \bar{L}_{\delta_{NE}}^{\text{unblind}}(\bar{\delta}_{NE}), \bar{L}_{\delta_{NW}}^{\text{unblind}}(\bar{\delta}_{NW}),$$

as well as encodings of the PRF keys

$$\bar{L}_{k_{NE}^0}^{\text{unblind}}(\bar{k}_{NE}^0), \bar{L}_{k_{NE}^1}^{\text{unblind}}(\bar{k}_{NE}^1), \bar{L}_{k_{NW}^0}^{\text{unblind}}(\bar{k}_{NW}^0), \bar{L}_{k_{NW}^1}^{\text{unblind}}(\bar{k}_{NW}^1).$$

At the end of the protocol execution, the adversary \mathcal{A} outputs some function of its view of the execution. If the simulator has not aborted, the simulator gives the output of \mathcal{A} to the environment. This completes the specification of the simulator \mathcal{S} . We now show that \mathcal{S} correctly simulates the view of \mathcal{A} in Hyb_1 .

Correctness of the simulation. First, we define some random variables for the view of the client in the real protocol and in the simulation. Let $\text{view}_{\text{real}}^{(0)}$ be the adversary's view during the setup phase when interacting according to the real protocol in Hyb_1 , and let $\text{view}_{\mathcal{S}}^{(0)}$ be the view that \mathcal{S} simulates for \mathcal{A} during the setup phase in the ideal world. In this case, $\text{view}_{\text{real}}^{(0)} = (\hat{k}_{\text{src}}^{(1)}, \hat{k}_{\text{dst}}^{(1)})$ where s and t correspond to the source and destination \mathcal{A} provided as input to the OT protocol. In the simulated view, $\text{view}_{\mathcal{S}}^{(0)} = (\bar{k}_{\text{src},s}^{(1)}, \bar{k}_{\text{dst},t}^{(1)})$ with s, t defined similarly.

Next, let $\text{view}_{\text{real}}^{(r)}$ be the random variable corresponding to the adversary's view during round r of the real protocol, and let $\text{view}_{\mathcal{S}}^{(r)}$ be the view that \mathcal{S} simulates for \mathcal{A} during round r in the ideal world. More explicitly, we write $\text{view}_{\text{real}}^{(r)} = (\text{PIR}_{\text{src}}, \text{PIR}_{\text{dst}}, \bar{C}^{\text{unblind}}, \bar{L}^{\text{unblind}})$, where PIR_{src} and PIR_{dst} denote the client's view in the PIR protocol on databases \mathcal{D}_{src} and \mathcal{D}_{dst} , respectively, \bar{C}^{unblind} denotes the garbled circuit, and \bar{L}^{unblind} denotes the set of garbled circuit encodings the client receives via the OT protocol (corresponding to inputs \hat{z}_{NE} and \hat{z}_{NW}) as well as the encodings of the server's inputs: the unblinding coefficients $\gamma_{NE}, \gamma_{NW}, \delta_{NE}, \delta_{NW}$, and the PRF keys $k_{NE}^0, k_{NE}^1, k_{NW}^0, k_{NW}^1$. Similarly, we define $\text{view}_{\mathcal{S}}^{(r)} = (\overline{\text{PIR}}_{\text{src}}, \overline{\text{PIR}}_{\text{dst}}, \bar{C}^{\text{unblind}}, \bar{L}^{\text{unblind}})$, where $\overline{\text{PIR}}_{\text{src}}$ and $\overline{\text{PIR}}_{\text{dst}}$ denote the client's view in the PIR protocol over databases $\overline{\mathcal{D}}_{\text{src}}$ and $\overline{\mathcal{D}}_{\text{dst}}$, respectively, and \bar{L}^{unblind} denotes the set of garbled circuit encodings the client receives from the OT protocol as well as the encodings of the server's inputs $\bar{\gamma}_{NE}, \bar{\gamma}_{NW}, \bar{\delta}_{NE}, \bar{\delta}_{NW}, \bar{k}_{NE}^0, \bar{k}_{NE}^1, \bar{k}_{NW}^0, \bar{k}_{NW}^1$. Next, define

$$\text{view}_{\text{real}}^{(0:r)} = \left\{ \text{view}_{\text{real}}^{(i)} \right\}_{i=0}^r$$

to be the joint distribution of the view of adversary \mathcal{A} in the setup and first r rounds of the protocol. We define $\text{view}_S^{(0:r)}$ analogously. We now show that

$$\text{view}_{\text{real}}^{(0:R)} \stackrel{c}{\approx} \text{view}_S^{(0:R)}.$$

We first characterize the keys in the simulation. Conceptually, we show that if a client knows at most one source key in round r , then this property also holds in $r + 1$. This effectively binds the client to a single consistent path in the course of the protocol execution.

Lemma A.9. *Let $s, t \in [n]$ be the source and destination nodes the client submits to the OT oracle in the setup phase of the simulation. Let $(s = v_0, \dots, v_R)$ be the path the simulator receives from the trusted party. Fix a round $0 < r < R$, and suppose the following conditions hold:*

- For all nodes $v \neq v_{r-1}$, the conditional distribution of $\bar{k}_{\text{src},v}^{(r)}$ given $\text{view}_S^{(0:r-1)}$ is computationally indistinguishable from the uniform distribution on $\{0, 1\}^\ell$.
- For all nodes $v \neq t$, the conditional distribution of $\bar{k}_{\text{dst},v}$ given $\text{view}_S^{(0:r-1)}$ is computationally indistinguishable from the uniform distribution on $\{0, 1\}^\ell$.

Then the corresponding conditions hold for round $r + 1$.

Proof. We first describe the client's view of the protocol execution on the r^{th} round of the protocol. For notational convenience, we set $u = v_{r-1}$. We now consider each component in $\text{view}_S^{(r)}$ separately:

- The client's view of the PIR protocol on the source database. We can express the view PIR_{src} as a (possibly randomized) function $f_1(\text{view}_S^{(0:r-1)}, \mathcal{D}_{\text{src}})$ in the client's view of the protocol execution thus far and the server's database $\overline{\mathcal{D}}_{\text{src}}$. Each record $v \neq u$ in $\overline{\mathcal{D}}_{\text{src}}$ is an encryption of the all-zeroes string. As long as $u \neq \perp$, the u^{th} record in $\overline{\mathcal{D}}_{\text{src}}$ contains the following:
 - Arithmetic circuit encodings $\bar{L}_s^{\text{affine}} = (\bar{L}_{\text{NE},x}^{\text{affine}}, \bar{L}_{\text{NW},x}^{\text{affine}})$ of the source node u .
 - Garbled circuit encodings $\bar{L}_s^{\text{unblind}}$ of the node u .
 - Encryptions $\bar{k}_N, \bar{k}_E, \bar{k}_S, \bar{k}_W$ of the source keys for the neighbors of $u = v_{r-1}$ in the next round of the protocol under the direction keys $\bar{k}_N, \bar{k}_E, \bar{k}_S, \bar{k}_W$:

$$\bar{k}_N = \text{Enc}(\bar{k}_N, \bar{k}_{\text{src},v_N}^{(r+1)}), \quad \bar{k}_E = \text{Enc}(\bar{k}_E, \bar{k}_{\text{src},v_E}^{(r+1)}), \quad \bar{k}_S = \text{Enc}(\bar{k}_S, \bar{k}_{\text{src},v_S}^{(r+1)}), \quad \bar{k}_W = \text{Enc}(\bar{k}_W, \bar{k}_{\text{src},v_W}^{(r+1)}),$$

where v_N, v_E, v_S, v_W is the neighbor of u in \mathcal{G} to the north, east, south, or west, respectively. If u does not have a neighbor in a given direction $\text{dir} \in \{N, E, S, W\}$, then $\bar{k}_{\text{src},v_{\text{dir}}}^{(r+1)} = \{0, 1\}^\ell$.

- The client's view of the PIR protocol on the destination database. Similar to the previous case, we can express the view $\overline{\text{PIR}}_{\text{dst}}$ as a (possibly randomized) function $f_2(\text{view}_S^{(0:r-1)}, \overline{\text{PIR}}_{\text{src}}, \overline{\mathcal{D}}_{\text{dst}})$. In hybrid H_2 , every record $v \neq t$ is an encryption of the all-zeroes string. The t^{th} record consists of the following:
 - Arithmetic circuit encodings $\bar{L}_t^{\text{affine}} = (\bar{L}_{\text{NE},y}^{\text{affine}}, \bar{L}_{\text{NW},y}^{\text{affine}})$ of the destination node t .
 - Garbled circuit encodings $\bar{L}_t^{\text{unblind}}$ of the destination node t .
- The garbled circuit \bar{C}^{unblind} .

- The set of garbled circuit encodings from the OT protocol. In the OT hybrid model, each OT is replaced by an oracle call to the ideal OT functionality. Thus, \bar{L}^{unblind} consists of a set of garbled circuit encodings $\bar{L}_{z_{\text{NE}}}^{\text{unblind}}, \bar{L}_{z_{\text{NW}}}^{\text{unblind}}$ from the OT protocol, encodings $\bar{L}_{\gamma_{\text{NE}}}^{\text{unblind}}, \bar{L}_{\gamma_{\text{NW}}}^{\text{unblind}}, \bar{L}_{\delta_{\text{NE}}}^{\text{unblind}}, \bar{L}_{\delta_{\text{NW}}}^{\text{unblind}}$ of the server's unblinding coefficients and encodings $\bar{L}_{k_{\text{NE}}^0}^{\text{unblind}}, \bar{L}_{k_{\text{NE}}^1}^{\text{unblind}}, \bar{L}_{k_{\text{NW}}^0}^{\text{unblind}}, \bar{L}_{k_{\text{NW}}^1}^{\text{unblind}}$ of the server's PRF keys.

We can express the adversary's view as

$$\text{view}_{\mathcal{S}}^{(r)} = \left\{ \bar{f} \left(\text{view}_{\mathcal{S}}^{(0:r-1)}, \bar{k}_{\text{N}}, \bar{k}_{\text{E}}, \bar{k}_{\text{S}}, \bar{k}_{\text{W}}, \overline{\text{affine}}, \bar{L}_{\text{s}}^{\text{unblind}}, \bar{L}_t^{\text{unblind}} \right), \bar{C}^{\text{unblind}}, \bar{L}^{\text{unblind}} \right\}, \quad (7)$$

where \bar{f} is a (possibly randomized) function, and $\overline{\text{affine}} = (\bar{L}_{\text{s}}^{\text{affine}}, \bar{L}_t^{\text{affine}})$ are the affine encodings of the source and destination vectors.

By construction, the set of encodings $\{\bar{L}_{\text{s}}^{\text{unblind}}, \bar{L}_t^{\text{unblind}}, \bar{L}^{\text{unblind}}\}$ constitute a complete set of encodings for a single unique input \bar{x} to the neighbor-computation function, and so invoking Lemma A.5, we conclude that there exists a PPT algorithm $\bar{\mathcal{S}}_1$ such that

$$\text{view}_{\mathcal{S}}^{(r)} \stackrel{c}{\approx} \bar{\mathcal{S}}_1 \left(1^\lambda, \text{view}_{\mathcal{S}}^{(0:r-1)}, \bar{k}_{\text{N}}, \bar{k}_{\text{E}}, \bar{k}_{\text{S}}, \bar{k}_{\text{W}}, \overline{\text{affine}}, C^{\text{unblind}}, C^{\text{unblind}}(\bar{x}) \right).$$

To conclude the proof, we condition on the possible outputs of $C^{\text{unblind}}(\bar{x})$. There are two cases:

- Suppose $C^{\text{unblind}}(\bar{x}) = (\hat{b}_{\text{NE}}, \hat{b}_{\text{NW}}, \hat{k}_{\text{NE}}, \hat{k}_{\text{NW}})$. By definition, this means that $\hat{k}_{\text{NE}} = \bar{k}_{\text{NE}}^{\hat{b}_{\text{NE}}}$ and $\hat{k}_{\text{NW}} = \bar{k}_{\text{NW}}^{\hat{b}_{\text{NW}}}$. Let $\text{dir} \in \{\text{N}, \text{E}, \text{S}, \text{W}\}$ be the direction of the edge from v_{r-1} to v_r in \mathcal{G} . The simulator chooses the garbled circuit encodings such that $\text{dir} = \text{IndexToDirection}(\hat{b}_{\text{NE}}, \hat{b}_{\text{NW}})$. Thus, by construction of the direction keys $\bar{k}_{\text{N}}, \bar{k}_{\text{E}}, \bar{k}_{\text{S}}, \bar{k}_{\text{W}}$, and PRF security, we conclude that the conditional distribution of $\bar{k}_{\text{dir}}^{(r)}$ given $\text{view}_{\mathcal{S}}^{(r)}$ is computationally indistinguishable from uniform for all directions $\text{dir}' \neq \text{dir}$. Invoking semantic security of (Enc, Dec) , we conclude that there exists a PPT algorithm $\bar{\mathcal{S}}_2$ such that

$$\text{view}_{\mathcal{S}}^{(r)} \stackrel{c}{\approx} \bar{\mathcal{S}}_2 \left(1^\lambda, \text{view}_{\mathcal{S}}^{(0:r-1)}, \bar{k}_{\text{dir}}, \overline{\text{affine}}, C^{\text{unblind}}, C^{\text{unblind}}(\bar{x}) \right). \quad (8)$$

By definition, v_r is the node in direction dir with respect to v_{r-1} , so

$$\bar{k}_{\text{dir}} = \text{Enc}(\bar{k}_{\text{dir}}, \bar{k}_{\text{src}, v_{\text{dir}}}^{(r+1)}) = \text{Enc}(\bar{k}_{\text{dir}}, \bar{k}_{\text{src}, v_r}^{(r+1)}). \quad (9)$$

Since for all $v \in [n]$, the simulator chooses $\bar{k}_{\text{src}, v}^{(r+1)}$ uniformly and independently (of all other quantities), we conclude from the characterization in Eq. (8) and (9) that for all nodes $v \neq v_r$, the conditional distribution of $\bar{k}_{\text{src}, v}^{(r+1)}$ given $\text{view}_{\mathcal{S}}^{(0:r)}$ is computationally indistinguishable from uniform. The first condition follows.

- Suppose $C^{\text{unblind}}(\bar{x}) = \perp$. Using an argument similar to that made for the previous case, the conditional distribution of the direction keys $\bar{k}_{\text{N}}, \bar{k}_{\text{E}}, \bar{k}_{\text{S}}, \bar{k}_{\text{W}}$ given $\text{view}_{\mathcal{S}}^{(0:r)}$ is computationally indistinguishable from uniform. By semantic security of (Enc, Dec) , we conclude that the adversary's view can be entirely simulated independently of $\bar{k}_{\text{src}, v}^{(r+1)}$ for all $v \in [n]$. Once again, the first condition holds.

To conclude the proof, we note that in the characterization from Eq. (7), all quantities in $\text{view}_{\mathcal{S}}^{(r)}$ are independent of $\bar{k}_{\text{dst}, v}$ for $v \neq t$. Thus, if the conditional distribution of $\bar{k}_{\text{dst}, v}$ given $\text{view}_{\mathcal{S}}^{(0:r-1)}$ is computationally indistinguishable from the uniform distribution, then the corresponding condition continues to hold at the end of round r . \square

With this preparation, we now argue inductively that $\text{view}_{\text{real}}^{(0:R)} \stackrel{c}{\approx} \text{view}_{\mathcal{S}}^{(0:R)}$. We begin with the base case.

Claim A.10. *Let $s, t \in [n]$ be the indices of the source and destination nodes, respectively, that the client submits to the OT oracle in the setup phase of the protocol. The following conditions hold at the end of the setup phase of the protocol:*

- *The simulator perfectly simulates the view of the \mathcal{A} : $\text{view}_{\text{real}}^{(0)} \equiv \text{view}_{\mathcal{S}}^{(0)}$.*
- *For all nodes $v \neq s$, the conditional distributions $k_{\text{src},v}^{(1)}$ given $\text{view}_{\text{real}}^{(0)}$, and $\bar{k}_{\text{src},v}^{(1)}$ given $\text{view}_{\mathcal{S}}^{(0)}$ are both uniform.*
- *For all nodes $v \neq t$, the conditional distributions $k_{\text{dst},v}$ given $\text{view}_{\text{real}}^{(0)}$ and $\bar{k}_{\text{dst},v}$ given $\text{view}_{\mathcal{S}}^{(0)}$ are both uniform.*

Proof. We consider each claim separately:

- In the real scheme, $k_{\text{src},s}^{(1)}, k_{\text{dst},t}$ are chosen uniformly and independently from $\{0, 1\}^\ell$. The same is true of the keys $\bar{k}_{\text{src},s}^{(1)}, \bar{k}_{\text{dst},t}$ in the simulation. Since $\text{view}_{\text{real}}^{(0)} = (k_{\text{src},s}^{(1)}, k_{\text{dst},t})$ and $\text{view}_{\mathcal{S}}^{(0)} = (\bar{k}_{\text{src},s}^{(1)}, \bar{k}_{\text{dst},t})$, we conclude that $\text{view}_{\text{real}}^{(0)} \equiv \text{view}_{\mathcal{S}}^{(0)}$.
- In the real protocol (resp., the simulation), for all $v \in [n]$, the encryption key $k_{\text{src},v}^{(1)}$ (resp., $\bar{k}_{\text{src},v}^{(1)}$) is chosen uniformly and independently of all other quantities in the protocol. Thus, for $v \neq s$, the distribution of $k_{\text{src},v}^{(1)}$ is independent of $\text{view}_{\text{real}}^{(0)} = (k_{\text{src},s}^{(1)}, k_{\text{dst},t})$. Similarly, the distribution of $\bar{k}_{\text{src},v}^{(1)}$ is independent of $\text{view}_{\mathcal{S}}^{(0)}$. The claim follows.
- Similar to the previous statement, for all $v \in [n]$, the encryption keys $k_{\text{dst},v}$ and $\bar{k}_{\text{dst},v}$ are chosen uniformly and independently of all other quantities in the protocol, which proves the claim. \square

Claim A.11. *Fix $0 < r < R$. Suppose the following conditions hold in round r :*

- *The view $\text{view}_{\text{real}}^{(0:r-1)}$ of the adversary interacting in the real protocol is computationally indistinguishable from the view $\text{view}_{\mathcal{S}}^{(0:r-1)}$ of the adversary interacting with the simulator.*
- *For all nodes $v \in [n]$, the conditional distribution of $k_{\text{src},v}^{(r)}$ given $\text{view}_{\text{real}}^{(0:r-1)}$ is computationally indistinguishable from the uniform distribution over $\{0, 1\}^\ell$ if and only if the conditional distribution of $\bar{k}_{\text{src},v}^{(r)}$ given $\text{view}_{\mathcal{S}}^{(0:r-1)}$ is computationally indistinguishable from the uniform distribution over $\{0, 1\}^\ell$.*
- *For all nodes $v \in [n]$, the conditional distribution of $k_{\text{dst},v}$ given $\text{view}_{\text{real}}^{(0:r-1)}$ is computationally indistinguishable from the uniform distribution over $\{0, 1\}^\ell$ if and only if the conditional distribution of $\bar{k}_{\text{dst},v}$ given $\text{view}_{\mathcal{S}}^{(0:r-1)}$ is computationally indistinguishable from the uniform distribution over $\{0, 1\}^\ell$.*

Then, the conditions also hold in round $r + 1$.

Proof. Let $s, t \in [n]$ be the source and destination nodes the client submits to the OT oracle in the setup phase of the protocol, and let $(s = v_0, \dots, v_R)$ be the path the simulator receives from the trusted party. Consider the view of \mathcal{A} in the simulation. By Claim A.10, for all $v \neq s$, the conditional distribution of the keys $\bar{k}_{\text{src},v}^{(1)}$ given $\text{view}_{\mathcal{S}}^{(0)}$ is uniform. Similarly, for all $v \neq t$, the conditional distribution of the keys $\bar{k}_{\text{dst},v}$ given $\text{view}_{\mathcal{S}}^{(0)}$ is uniform. By iteratively applying Lemma A.9, we conclude that for all nodes $v \neq v_{r-1}$, the conditional distribution of $\bar{k}_{\text{src},v}^{(r)}$ given $\text{view}_{\mathcal{S}}^{(0:r-1)}$ is computationally indistinguishable from uniform. Similarly, for all $v \neq t$, the conditional distribution of $\bar{k}_{\text{dst},t}$ given $\text{view}_{\mathcal{S}}^{(0:r-1)}$ is computationally indistinguishable from uniform.

Invoking the inductive hypothesis, we have that for all nodes $v \neq v_{r-1}$, the conditional distribution of $k_{\text{src},v}^{(r)}$ given $\text{view}_{\text{real}}^{(0:r-1)}$ is computationally indistinguishable from uniform. Similarly, for all nodes $v \neq t$, the conditional distribution of $k_{\text{dst},v}$ given $\text{view}_{\text{real}}^{(0:r-1)}$ is computationally indistinguishable from uniform. We now show that on round r , $(\text{PIR}_{\text{src}}, \text{PIR}_{\text{dst}}) \stackrel{c}{\approx} (\overline{\text{PIR}}_{\text{src}}, \overline{\text{PIR}}_{\text{dst}})$. The client's view in the PIR protocol can be regarded as a (possibly randomized) function of the client's view in the first $r-1$ rounds of the protocol and the server's input database \mathcal{D} . Since $\text{view}_{\text{real}}^{(0:r-1)} \stackrel{c}{\approx} \text{view}_{\mathcal{S}}^{(0:r-1)}$, it suffices to argue that the conditional distribution of $(\mathcal{D}_{\text{src}}, \mathcal{D}_{\text{dst}})$ given $\text{view}_{\text{real}}^{(0:r-1)}$ is computationally indistinguishable from the conditional distribution of $(\overline{\mathcal{D}}_{\text{src}}, \overline{\mathcal{D}}_{\text{dst}})$ given $\text{view}_{\mathcal{S}}^{(0:r-1)}$.

For notational convenience, let $u = v_{r-1}$. We use a hybrid argument. We define the following hybrid experiments:

- Hybrid H_0 is the real game where the server prepares \mathcal{D}_{src} and \mathcal{D}_{dst} as described in Figure 3.
- Hybrid H_1 is identical to H_0 , except the server substitutes an encryption of the all-zeroes string under $k_{\text{src},v}^{(r)}$ for all records $v \neq u$ in \mathcal{D}_{src} .
- Hybrid H_2 is identical to H_1 , except the server substitutes an encryption of the all-zeroes string under $k_{\text{dst},v}$ for all records $v \neq t$ in \mathcal{D}_{dst} .

Since for all $v \neq u$, the conditional distribution of $k_{\text{src},v}^{(r)}$ given $\text{view}_{\text{real}}^{(0:r-1)}$ is computationally indistinguishable from uniform, we can appeal to the semantic security of (Enc, Dec) to conclude that hybrid experiments H_0 and H_1 are computationally indistinguishable. Similarly, since for all $v \neq t$, the conditional distribution of $k_{\text{dst},v}$ given $\text{view}_{\text{real}}^{(0:r-1)}$ is computationally indistinguishable from uniform, we have that H_1 and H_2 are computationally indistinguishable.

We now show that the joint distribution of $(\mathcal{D}_{\text{src}}, \mathcal{D}_{\text{dst}})$ in H_2 is computationally indistinguishable from the joint distribution of $(\overline{\mathcal{D}}_{\text{src}}, \overline{\mathcal{D}}_{\text{dst}})$ in the simulation. In H_2 , every record $v \neq u$ in \mathcal{D}_{src} is an encryption of the all-zeroes string under a key $k_{\text{src},v}^{(r)}$ that is computationally indistinguishable from uniform given the adversary's view of the protocol thus far; the same is true in the simulation. Similarly, every record $v \neq t$ in \mathcal{D}_{dst} is an encryption of the all-zeroes string under a key $k_{\text{dst},v}$ that looks computationally indistinguishable from uniform to the adversary. This is the case in the simulation. Let $r_{\text{src},u}$ be the u^{th} record in \mathcal{D}_{src} and let $r_{\text{dst},t}$ be the t^{th} record in \mathcal{D}_{dst} . Similarly, let $\bar{r}_{\text{src},u}$ be the u^{th} record in $\overline{\mathcal{D}}_{\text{src}}$ and let $\bar{r}_{\text{dst},t}$ be the t^{th} record in $\overline{\mathcal{D}}_{\text{dst}}$. It suffices now to show that $(r_{\text{src},u}, r_{\text{dst},t})$ is computationally indistinguishable from $(\bar{r}_{\text{src},u}, \bar{r}_{\text{dst},t})$. In the real scheme, the record $r_{\text{src},u}$ contains the following components:

- Arithmetic circuit encodings $\tilde{L}_{\text{NE},x}^{\text{affine}}(A_u^{(\text{NE})}), \tilde{L}_{\text{NW},x}^{\text{affine}}(A_u^{(\text{NW})})$ for the source u .
- Garbled circuit encodings $\tilde{L}_s^{\text{unblind}}(u)$ for the source u .
- Encryptions $\kappa_{\text{N}}, \kappa_{\text{E}}, \kappa_{\text{S}}, \kappa_{\text{W}}$ of the source keys for the neighbors of u in the next round of the protocol under the direction keys $k_{\text{N}}, k_{\text{E}}, k_{\text{S}}, k_{\text{W}}$:

$$\kappa_{\text{N}} = \text{Enc}(k_{\text{N}}, k_{\text{src},v_{\text{N}}}^{(r+1)}), \quad \kappa_{\text{E}} = \text{Enc}(k_{\text{E}}, k_{\text{src},v_{\text{E}}}^{(r+1)}), \quad \kappa_{\text{S}} = \text{Enc}(k_{\text{S}}, k_{\text{src},v_{\text{S}}}^{(r+1)}), \quad \kappa_{\text{W}} = \text{Enc}(k_{\text{W}}, k_{\text{src},v_{\text{W}}}^{(r+1)}),$$

where $v_{\text{N}}, v_{\text{E}}, v_{\text{S}}, v_{\text{W}}$ is the neighbor of u in \mathcal{G} to the north, east, south, or west, respectively. If u does not have a neighbor in a given direction $\text{dir} \in \{\text{N}, \text{E}, \text{S}, \text{W}\}$, then $k_{\text{src},v_{\text{dir}}}^{(r+1)} = \{0, 1\}^\ell$.

The record $r_{\text{dst},t}$ contains the following components:

- Arithmetic circuit encodings $\tilde{L}_{\text{NE},y}^{\text{affine}}(B_t^{(\text{NE})}), \tilde{L}_{\text{NW},y}^{\text{affine}}(B_t^{(\text{NW})})$ for the destination t .

- Garbled circuit encodings $\tilde{L}_t^{\text{unblind}}(t)$ for the destination t .

In the real protocol, the arithmetic circuit encodings are constructed independently of the garbled circuit for the neighbor-computation function. The neighbor keys k_{dir} for $\text{dir} \in \{N, E, S, W\}$ and source keys $k_{\text{src}, u_{\text{dir}}}^{(r+1)}$ for the subsequent round of the protocol are also generated independently of both the arithmetic circuit encodings and the garbled circuit. Thus, the joint distribution decomposes into three product distributions over the arithmetic circuit encodings, the garbled circuit encodings, and the encryptions of the source keys for the next round. We reason about each distribution separately:

- The simulator constructs the garbled circuit for the neighbor-computation function exactly as in the real scheme. Thus, the garbled circuit encodings in $r_{\text{src}, u}, \bar{r}_{\text{src}, u}$ and $r_{\text{dst}, t}, \bar{r}_{\text{dst}, t}$ are identically distributed.
- The neighbor keys \bar{k}_{dir} for $\text{dir} \in \{N, E, S, W\}$ in the simulation are generated exactly as the keys k_{dir} in the real scheme.
- In the real scheme, the affine encodings $\tilde{L}_{\text{NE}, x}^{\text{affine}}(A_u^{(\text{NE})})$ and $\tilde{L}_{\text{NE}, y}^{\text{affine}}(B_t^{(\text{NE})})$ evaluate to

$$z_{\text{NE}} = \alpha_{\text{NE}} \langle A_u^{(\text{NE})}, B_t^{(\text{NE})} \rangle + \beta_{\text{NE}},$$

where α_{NE} is uniform in \mathbb{F}_p^* and β_{NE} is uniform in \mathbb{F}_p . In particular, this means that z_{NE} is distributed uniformly over \mathbb{F}_p . This is precisely the same distribution from which the simulator samples \bar{z}_{NE} . Together with Lemma A.6, we conclude that

$$\mathcal{S}_{\text{ac}}(\bar{z}_{\text{NE}}) \equiv \mathcal{S}_{\text{ac}}(z_{\text{NE}}) \equiv (\tilde{L}_{\text{NE}, x}^{\text{affine}}(A_u^{(\text{NE})}), \tilde{L}_{\text{NE}, y}^{\text{affine}}(B_t^{(\text{NE})})).$$

An analogous argument shows that

$$\mathcal{S}_{\text{ac}}(\bar{z}_{\text{NW}}) \equiv \mathcal{S}_{\text{ac}}(z_{\text{NW}}) \equiv (\tilde{L}_{\text{NW}, x}^{\text{affine}}(A_u^{(\text{NW})}), \tilde{L}_{\text{NW}, y}^{\text{affine}}(B_t^{(\text{NW})})),$$

where $z_{\text{NW}} = \alpha_{\text{NW}} \langle A_u^{(\text{NW})}, B_t^{(\text{NW})} \rangle + \beta_{\text{NW}}$. We conclude that the arithmetic circuit encodings are identically distributed in both the real scheme and the simulation.

We conclude that $(r_{\text{src}, u}, r_{\text{dst}, t}) \equiv (\bar{r}_{\text{src}, u}, \bar{r}_{\text{dst}, t})$, and correspondingly, $(\mathcal{D}_{\text{src}}, \mathcal{D}_{\text{dst}}) \stackrel{c}{\approx} (\overline{\mathcal{D}_{\text{src}}}, \overline{\mathcal{D}_{\text{dst}}})$. In the real protocol, the client's view $\text{PIR}_{\text{src}}, \text{PIR}_{\text{dst}}$ in the PIR protocols can be expressed as an efficiently-computable and possibly randomized function f of its view $\text{view}_{\text{real}}^{(0:r-1)}$ in the first $r-1$ rounds of the protocol and the server's databases $\mathcal{D}_{\text{src}}, \mathcal{D}_{\text{dst}}$:

$$(\text{PIR}_{\text{src}}, \text{PIR}_{\text{dst}}) \equiv f(\text{view}_{\text{real}}^{(0:r-1)}, r_{\text{src}, u}, r_{\text{dst}, t}).$$

In the simulation, the simulator synthesizes databases $\overline{\mathcal{D}_{\text{src}}}, \overline{\mathcal{D}_{\text{dst}}}$, and then plays the role of the server in the PIR protocol. Thus, we have that

$$(\overline{\text{PIR}_{\text{src}}}, \overline{\text{PIR}_{\text{dst}}}) \equiv f(\text{view}_{\mathcal{S}}^{(0:r-1)}, \bar{r}_{\text{src}, u}, \bar{r}_{\text{dst}, t}).$$

Moreover, we note that the records $r_{\text{src}, u}, r_{\text{dst}, t}, \bar{r}_{\text{src}, u}, \bar{r}_{\text{dst}, t}$ are constructed independently of variables from the previous round of the protocol. Thus, using the inductive hypothesis, $\text{view}_{\text{real}}^{(0:r-1)} \stackrel{c}{\approx} \text{view}_{\mathcal{S}}^{(0:r-1)}$, and the fact that $(r_{\text{src}, u}, r_{\text{dst}, t}) \equiv (\bar{r}_{\text{src}, u}, \bar{r}_{\text{dst}, t})$, we conclude that

$$(\text{PIR}_{\text{src}}, \text{PIR}_{\text{dst}}) \equiv f(\text{view}_{\text{real}}^{(0:r-1)}, r_{\text{src}, u}, r_{\text{dst}, t}) \stackrel{c}{\approx} f(\text{view}_{\mathcal{S}}^{(0:r-1)}, \bar{r}_{\text{src}, u}, \bar{r}_{\text{dst}, t}) \equiv (\overline{\text{PIR}_{\text{src}}}, \overline{\text{PIR}_{\text{dst}}}). \quad (10)$$

Using this characterization, we can write

$$\text{view}_{\text{real}}^{(0:r)} = \left\{ \text{view}_{\text{real}}^{(0:r-1)}, f(\text{view}_{\text{real}}^{(0:r-1)}, r_{\text{src},u}, r_{\text{dst},t}), \tilde{C}^{\text{unblind}}, \tilde{L}^{\text{unblind}} \right\}, \quad (11)$$

and similarly,

$$\text{view}_{\mathcal{S}}^{(0:r)} = \left\{ \text{view}_{\mathcal{S}}^{(0:r-1)}, f(\text{view}_{\mathcal{S}}^{(0:r-1)}, \bar{r}_{\text{src},u}, \bar{r}_{\text{dst},t}), \bar{C}^{\text{unblind}}, \bar{L}^{\text{unblind}} \right\}. \quad (12)$$

From this, we express $\text{view}_{\text{real}}^{(0:r)}$ as an efficiently-computable and possibly randomized function f' in the garbled circuit components and the auxiliary components:

$$\text{view}_{\text{real}}^{(r)} \equiv f' \left(\tilde{L}_s^{\text{unblind}}(u), \tilde{L}_t^{\text{unblind}}(t), \tilde{L}^{\text{unblind}}, \tilde{C}^{\text{unblind}}, \text{aux} \right), \quad (13)$$

where aux contains the additional variables $\text{view}_{\text{real}}^{(r)}$ depends on that are independent of the garbled circuit encodings: $\text{view}_{\text{real}}^{(0:r-1)}$, the arithmetic circuit encodings, and the encryptions of the source keys for the next round of the protocol.

Since the simulator prepares the garbled circuit exactly as in the real protocol, from the characterization of $\text{view}_{\mathcal{S}}^{(0:r)}$ in Eq. (12), we can similarly write

$$\text{view}_{\mathcal{S}}^{(r)} \equiv f' \left(\bar{L}_s^{\text{unblind}}(u), \bar{L}_t^{\text{unblind}}(t), \bar{L}^{\text{unblind}}, \bar{C}^{\text{unblind}}, \overline{\text{aux}} \right), \quad (14)$$

where $\overline{\text{aux}}$ contains the same auxiliary variables as aux . From Eq. (10), we have in particular that $\text{aux} \stackrel{c}{\approx} \overline{\text{aux}}$.

By construction, the encodings $\tilde{L}_s^{\text{unblind}}(u), \tilde{L}_t^{\text{unblind}}(t), \tilde{L}^{\text{unblind}}$ in the real scheme constitutes a complete set of encodings for the garbled circuit $\tilde{C}^{\text{unblind}}$. Let \tilde{x} be the associated input to the neighbor-computation circuit C^{unblind} . In particular, $\tilde{x} = (u, t, \hat{z}_{\text{NE}}, \hat{z}_{\text{NW}}, \gamma_{\text{NE}}, \gamma_{\text{NW}}, \delta_{\text{NE}}, \delta_{\text{NW}}, k_{\text{NE}}^0, k_{\text{NE}}^1, k_{\text{NW}}^0, k_{\text{NW}}^1)$, where \hat{z}_{NE} and \hat{z}_{NW} are the encodings the client OTs for in the protocol execution. Similarly, the encodings $\bar{L}_s^{\text{unblind}}(u), \bar{L}_t^{\text{unblind}}(t), \bar{L}^{\text{unblind}}$ constitutes a complete set of encodings for the garbled circuit \bar{C}^{unblind} . Let $\bar{x} = (u, t, \bar{z}_{\text{NE}}, \bar{z}_{\text{NW}}, \bar{\gamma}_{\text{NE}}, \bar{\gamma}_{\text{NW}}, \bar{\delta}_{\text{NE}}, \bar{\delta}_{\text{NW}}, \bar{k}_{\text{NE}}^0, \bar{k}_{\text{NE}}^1, \bar{k}_{\text{NW}}^0, \bar{k}_{\text{NW}}^1)$ be the associated input to C^{unblind} in the simulation. Moreover, by the characterization of the client's view in Eq. (11) and (12), on each round of the protocol execution, we can associate two unique sets of affine encodings with the client's view, one for each axis. Let z_{NE} and z_{NW} denote the values to which these two sets of affine encodings evaluate. In the real execution, we thus have $z_{\text{NE}} = \alpha_{\text{NE}} \langle A_u^{(\text{NE})}, B_t^{(\text{NE})} \rangle + \beta_{\text{NE}}$ and $z_{\text{NW}} = \alpha_{\text{NW}} \langle A_u^{(\text{NW})}, B_t^{(\text{NW})} \rangle + \beta_{\text{NW}}$. In the simulation, we have $z_{\text{NE}} = \bar{z}_{\text{NE}}$ and $z_{\text{NW}} = \bar{z}_{\text{NW}}$. We now consider three cases:

- Suppose that the client OTs for the encodings of inputs consistent with the outputs of the arithmetic circuit. In other words, $\hat{z}_{\text{NE}} = z_{\text{NE}}$ and $\hat{z}_{\text{NW}} = z_{\text{NW}}$. By definition, if $u = t$, then $C^{\text{unblind}}(\tilde{x}) = \perp$. Otherwise, by correctness of the affine encodings, $C^{\text{unblind}}(\tilde{x}) = (\hat{b}_{\text{NE}}, \hat{b}_{\text{NW}}, k_{\text{NE}}^{\hat{b}_{\text{NE}}}, k_{\text{NW}}^{\hat{b}_{\text{NW}}})$, where $\text{dir} = \text{IndexToDirection}(\hat{b}_{\text{NE}}, \hat{b}_{\text{NW}})$ is the direction of travel from u to t , as determined by the next-hop routing matrices $A^{(\text{NE})}, B^{(\text{NE})}, A^{(\text{NW})}, B^{(\text{NW})}$. In particular, dir is the direction of the edge from $u = v_{r-1}$ to v_r . In the simulation, when $\hat{z}_{\text{NE}} = \bar{z}_{\text{NE}}$ and $\hat{z}_{\text{NW}} = \bar{z}_{\text{NW}}$, the simulator chooses unblinding factors $\bar{\gamma}_{\text{NE}}, \bar{\gamma}_{\text{NW}}, \bar{\delta}_{\text{NE}}, \bar{\delta}_{\text{NW}}$ such that $C^{\text{unblind}}(\bar{x}) = C^{\text{unblind}}(\tilde{x})$. Now, invoking the input-privacy of the garbling scheme (Definition A.4), we conclude that

$$\begin{aligned} \left\{ \tilde{L}_s^{\text{unblind}}(u), \tilde{L}_t^{\text{unblind}}(t), \tilde{L}^{\text{unblind}}, \tilde{C}^{\text{unblind}} \right\} &\stackrel{c}{\approx} \mathcal{S}_{\text{Yao}}(1^\lambda, C^{\text{unblind}}, C^{\text{unblind}}(\tilde{x})) \\ &\stackrel{c}{\approx} \left\{ \bar{L}_s^{\text{unblind}}(u), \bar{L}_t^{\text{unblind}}(t), \bar{L}^{\text{unblind}}, \bar{C}^{\text{unblind}} \right\}. \end{aligned}$$

Since the garbled circuit components of $\text{view}_{\text{real}}^{(r)}$ are computationally indistinguishable from $\text{view}_S^{(r)}$, we conclude that $\text{view}_{\text{real}}^{(r)} \stackrel{c}{\approx} \text{view}_S^{(r)}$.

To see the second condition of Claim A.11 holds, we appeal to input privacy of the garbling scheme to argue that the conditional distribution of the keys $k_{\text{NE}}^{1-\hat{b}_{\text{NE}}}$ and $k_{\text{NW}}^{1-\hat{b}_{\text{NW}}}$ is computationally indistinguishable from uniform given $\text{view}_{\text{real}}^{(0:r)}$. As in the proof of Lemma A.9, the conditional distribution of k'_{dir} for all directions $\text{dir}' \neq \text{dir} \in \{\text{N, E, S, W}\}$ is computationally indistinguishable from uniform given $\text{view}_{\text{real}}^{(0:r)}$. Finally, invoking semantic security of (Enc, Dec) , we conclude that for all $v \neq v_r$, the conditional distribution of $k_{\text{src},v}^{(r+1)}$ is computationally indistinguishable from uniform. As shown in the proof of Lemma A.9, this is precisely the case in the simulation.

The third condition of Claim A.11 holds since the components in the client's view in round r of the real scheme (Eq. 11) are independent of the destination keys $k_{\text{dst},v}$ for all $v \neq t$. Since the conditional distribution of $k_{\text{dst},v}$ for $v \neq t$ given $\text{view}_{\text{real}}^{(0:r-1)}$ is computationally indistinguishable from uniform, the conditional distribution remains uniform conditioned on $\text{view}_{\text{real}}^{(0:r)}$. This precisely matches the distribution in the simulation (Lemma A.9).

- Suppose that $u = t$. Then, $C^{\text{unblind}}(\tilde{x}) = \perp = C^{\text{unblind}}(\bar{x})$. As in the previous case, input privacy of the garbling scheme yields $\text{view}_{\text{real}}^{(r)} \stackrel{c}{\approx} \text{view}_S^{(r)}$.

It is not difficult to see that the second condition of Claim A.11 holds. Since $C^{\text{unblind}}(\tilde{x}) = \perp$, the conditional distribution of the keys $k_{\text{NE}}^0, k_{\text{NE}}^1, k_{\text{NW}}^0, k_{\text{NW}}^1$ is computationally indistinguishable from uniform given $\text{view}_{\text{real}}^{(0:r)}$. By semantic security of (Enc, Dec) , we conclude that the conditional distribution of $k_{\text{src},v}^{(r+1)}$ given $\text{view}_{\text{real}}^{(0:r)}$ is computationally indistinguishable from uniform for all $v \in [n]$. By the case analysis in the proof of Lemma A.9, this is also the case in the simulation.

The third condition follows as in the previous case.

- Suppose that $u \neq t$ and moreover, the client OTs for encodings of inputs that are inconsistent with the outputs of the arithmetic circuit. We consider two possibilities.
 - Suppose that exactly one of $\hat{z}_{\text{NE}} \neq z_{\text{NE}}$ and $\hat{z}_{\text{NW}} \neq z_{\text{NW}}$ hold. Without loss of generality, suppose that $\hat{z}_{\text{NE}} \neq z_{\text{NE}}$. In the real scheme, this means that $\hat{z}_{\text{NE}} \neq \alpha_{\text{NE}} \langle A_u^{(\text{NE})}, B_t^{(\text{NE})} \rangle + \beta_{\text{NE}}$. Next, since $\gamma_{\text{NE}} = \alpha_{\text{NE}}^{-1}$ and $\delta_{\text{NE}} = \alpha_{\text{NE}}^{-1} \beta_{\text{NE}}$, where α_{NE} is uniform in \mathbb{F}_p^* and β_{NE} is uniform in \mathbb{F}_p . Thus, γ_{NE} and δ_{NE} are also distributed uniformly over \mathbb{F}_p^* and \mathbb{F}_p , respectively. Since the family of functions $\{h_{\gamma,\delta}(z) = \gamma z + \delta \pmod{p} \mid \gamma \in \mathbb{F}_p^*, \delta \in \mathbb{F}_p\}$ is pairwise independent, it follows that for any distinct $z_{\text{NE}}, z_{\text{NE}}^* \in \mathbb{F}_p$, and all $a, b \in \mathbb{F}_p$,

$$\Pr [h_{\gamma_{\text{NE}}, \delta_{\text{NE}}}(z_{\text{NE}}) = a \wedge h_{\gamma_{\text{NE}}, \delta_{\text{NE}}}(z_{\text{NE}}^*) = b] = \frac{1}{p^2},$$

where the probability is taken over the randomness in γ_{NE} and δ_{NE} . The client's choice of \hat{z}_{NE} and \hat{z}_{NW} depends only on its view $\text{view}_{\text{real}}^{(0:r-1)}$ of the protocol execution in the previous rounds of the protocol, as well as its view PIR_{src} and PIR_{dst} in the PIR protocol. The quantities γ_{NE} and δ_{NE} are sampled independently of $\text{view}_{\text{real}}^{(0:r-1)}$. By the above characterization of PIR_{src} and PIR_{dst} , the joint distribution of $(\text{PIR}_{\text{src}}, \text{PIR}_{\text{dst}})$ is entirely simulatable given only z_{NE} and variables that are independent of γ_{NE} and δ_{NE} (by invoking the simulator for the affine encodings). Thus, by

pairwise independence, we conclude that

$$\Pr [h_{\gamma_{\text{NE}}, \delta_{\text{NE}}}(\hat{z}_{\text{NE}}) \in [-2^\tau, 2^\tau]] = \frac{2^{\tau+1}}{p} = \varepsilon.$$

If $\hat{z}_{\text{NE}} \neq z_{\text{NE}}$, then with probability $1 - \varepsilon$, $C^{\text{unblind}}(\tilde{x}) = \perp$. With probability ε , $C^{\text{unblind}}(\tilde{x}) \neq \perp$, but this precisely corresponds to the first abort condition in experiment Hyb_1 . Thus, in Hyb_1 , with probability $1 - \varepsilon$, $C^{\text{unblind}}(\tilde{x}) = \perp$ and with probability ε , the protocol execution aborts.

In the simulation, the simulator chooses the unblinding factors $\bar{\gamma}_{\text{NE}}, \bar{\gamma}_{\text{NW}}, \bar{\delta}_{\text{NE}}, \bar{\delta}_{\text{NW}}$, such that $C^{\text{unblind}}(\tilde{x}) = \perp$ with probability $1 - \varepsilon$. With probability ε , the simulator aborts. We conclude by input privacy of the garbling scheme that the simulation is correct. The analysis for the case where $\hat{z}_{\text{NW}} \neq z_{\text{NW}}$, but $\hat{z}_{\text{NE}} = z_{\text{NE}}$ is entirely analogous.

- Suppose that both $\hat{z}_{\text{NE}} \neq z_{\text{NE}}$ and $\hat{z}_{\text{NW}} \neq z_{\text{NW}}$. By the same analysis as in the first case, we have that

$$\Pr [h_{\gamma_{\text{NE}}, \delta_{\text{NE}}}(\hat{z}_{\text{NE}}) \in [-2^\tau, 2^\tau]] = \varepsilon = \Pr [h_{\gamma_{\text{NW}}, \delta_{\text{NW}}}(\hat{z}_{\text{NW}}) \in [-2^\tau, 2^\tau]].$$

Since the two events are independent, $C^{\text{unblind}}(\tilde{x}) \neq \perp$ with probability ε^2 , and the experiment aborts in Hyb_1 . With probability $1 - \varepsilon^2$, $C^{\text{unblind}}(\tilde{x}) = \perp$. In the simulation, the simulator chooses the unblinding factors such that $C^{\text{unblind}}(\tilde{x}) = \perp$ with probability $1 - \varepsilon^2$, and aborts with probability ε^2 . Correctness of the simulation follows by input privacy of the garbling scheme.

Since $C^{\text{unblind}}(\tilde{x})$ is either equal to \perp or the experiment aborts, the proof of the second and third statements of Claim A.11 follows exactly as in the previous case.

Combining Claim A.10 and Claim A.11, we conclude by induction on r that $\text{view}_{\text{real}}^{(0:R)} \stackrel{c}{\approx} \text{view}_{\mathcal{S}}^{(0:R)}$. Thus, the view of the protocol execution simulated by \mathcal{S} for \mathcal{A} is computationally indistinguishable from the view of \mathcal{A} interacting with the server in Hyb_1 . Correctness of the simulation follows. \square

Proof of Claim A.7. To conclude the proof of Theorem 4.5, we show Claim A.7, or equivalently, that no efficient adversary is able to distinguish Hyb_0 from Hyb_1 except with advantage $\text{negl}(\lambda) + R \cdot 2^{-\mu}$. Let \mathcal{A} be a distinguisher between Hyb_0 and Hyb_1 with distinguishing advantage Adv :

$$\text{Adv} = |\Pr [\text{Hyb}_0(\lambda, \pi, \mathcal{A}, \mathcal{E})] - \Pr [\text{Hyb}_1(\lambda, \pi, \mathcal{A}, \mathcal{E})]|.$$

By construction, Hyb_0 and Hyb_1 are identical experiments, except experiment Hyb_1 terminates if the abort event in Figure 5 occurs. Thus, it must be the case that \mathcal{A} is able to cause the bad event to occur with probability Adv in the real experiment. But by Claim A.8, the real protocol execution experiment with the abort event is computationally indistinguishable from the ideal-world execution with the simulator \mathcal{S} described in the proof of Claim A.8. Thus, if \mathcal{A} is able to trigger the abort event in the real protocol execution with probability Adv , it is able to trigger the abort event when interacting with the simulator \mathcal{S} with probability that is negligibly close to Adv . It suffices now to bound the probability that the simulator \mathcal{S} aborts the protocol execution. On each round r , the simulator \mathcal{S} aborts with probability at most $\varepsilon = 2^{\tau+1}/p \leq 2^{-\mu}$ since $p > 2^{\mu+\tau+1}$ (irrespective of the computational power of the adversary). By a union bound over the total number of rounds R , we conclude that \mathcal{S} aborts with probability at most $R \cdot 2^{-\mu}$ in the protocol execution. Thus, the probability that \mathcal{A} can trigger the abort event in the real protocol must be negligibly close to $R \cdot 2^{-\mu}$. We conclude that $\text{Adv} \leq \text{negl}(\lambda) + R \cdot 2^{-\mu}$, which proves the claim. \square

Claims A.7 and A.8 show that no no efficient environment can distinguish the real-world execution (Hyb_0) from the ideal-world execution (Hyb_2), except with advantage negligibly close to $R \cdot 2^{-\mu}$. This proves Theorem 4.5. \square