

Private Inference Control

David Woodruff*
MIT
dpwood@mit.edu

Jessica Staddon
Palo Alto Research Center
staddon@parc.com

Abstract

Access control can be used to ensure that database queries pertaining to sensitive information are not answered. This is not enough to prevent users from learning sensitive information though, because users can combine non-sensitive information to discover something sensitive. Inference control prevents users from obtaining sensitive information via such “inference channels”, however, existing inference control techniques are not private - that is, they require the server to learn what queries the user is making in order to deny inference-enabling queries.

We propose a new primitive - private inference control (PIC) - which is a means for the server to provide inference control without learning what information is being retrieved. PIC is a generalization of private information retrieval (PIR) and symmetrically-private information retrieval (SPIR). While it is straightforward to implement access control using PIR (simply omit sensitive information from the database), it is nontrivial to implement *inference* control efficiently. We measure the efficiency of a PIC protocol in terms of its communication complexity, its round complexity, and the work the server performs per query. Under existing cryptographic assumptions, we give a PIC scheme which is simultaneously optimal, up to logarithmic factors, in the work the server performs per query, the total communication complexity, and the number of rounds of interaction. We also present a scheme requiring more communication but sufficient storage of state by the server to facilitate private user revocation. Finally, we present a generic reduction which shows that one can focus on designing PIC schemes for which the inference channels take a particularly simple threshold form.

Keywords: Private information retrieval, oblivious transfer, inference control

1 Introduction

Consider a user who wants to query a database containing sensitive information. The user should learn the response to a query if and only if it does not enable them to learn something sensitive. There are two ways to learn something sensitive: by accessing sensitive information directly or by making a sequence of non-sensitive queries whose answers, when taken together, allow the user to infer something sensitive. Preventing the former is known as *access control* and the latter as *inference control* (see, for example, [21]). A sequence of queries whose responses allow a sensitive inference to be made is known as an *inference channel*, and the goal of inference control is to prevent the user from *completing* any inference channel. As an example, consider the problem of protecting the identity of individuals in a database [18, 44, 43]. Social security numbers are clearly identifying and access control ensures that they cannot be queried without appropriate authorization. In addition, while attributes such as city of residence, year of birth, profession, marital status and political party, may not be identifying individually (and so not subject to access control), when these values are received for a single record in a database they may be identifying and thus they form an inference channel.¹ Inference control ensures that colluding users aren’t able to together make enough queries to complete such a channel.

*Most of this research was conducted while this author was an intern at PARC.

¹Note that inference control is equivalent to access control if and only if all inference channels have length one.

An equally important aspect of this setting is the potential privacy loss on the user’s side due to the exposure of their queries to the database server. To solve this problem private information retrieval (PIR) schemes (see, for example, [16, 34]) have been developed, but such schemes are not concerned with inference control. The question then arises: *How can we ensure users are not able to infer sensitive information without knowing what information they are retrieving?* To solve this problem we introduce a new primitive that we call private inference control, or PIC. In a PIC scheme, the server learns no information in the computational sense about the user’s queries, even if the user is unable to retrieve the answer to their query. Note that any PIC protocol implements oblivious transfer [42, 20, 38], so it must be interactive. Hence, the server learns the number of queries made so far simply by counting the number of interactions with the user. We emphasize that this is in fact the *only* information a computationally-bounded server learns about the user’s queries. Indeed, it is precisely the knowledge of the number of queries made so far that will enable the server to do collusion-resistant inference control.

Oblivious transfer has been studied in the context of access control [40, 3]. However, both [40] and [3] can only be applied to protect against certain collections of inference channels, rather than the more general collections that we consider here (see Section 3 for a precise description). In addition, they don’t apply to the relational databases [14] that are our main focus because they view the data as consisting of a single record. Running multiple instances of schemes in [40, 3] in parallel to accommodate multiple records causes the efficiency of their techniques, and much of the user privacy that we desire, to be lost. Indeed, we believe that new ideas are needed to solve the PIC problem.

Our enabling technique is the transfer of a small amount of extra information to the user with a query. In subsequent queries, the extra information is used to encrypt data that is communicated through the execution of a PIR or SPIR protocol. The encryption is done in such a way to ensure that users who have made a permissible sequence of queries thus far (as defined by the inference channels) are able to decrypt. The difficulty comes in determining how to do this encryption to accommodate an arbitrary collection of inference channels while keeping communication overhead and server work at reasonable levels.

As suggested by our earlier example, PIC is useful in relational databases. Once the inference channels have been identified (with a tool such as [43], for example) a PIC scheme can be designed to ensure that users are prevented from making undesired inferences while privately accessing the database, even when users collude. In addition, even when data perturbation is used to ensure privacy of the data (see, for example, [5, 48, 2]), PIC can be useful. For example, in [18] it’s shown that in a statistical database in which sums are queried, a large amount of perturbation is needed in order to preserve privacy. So much perturbation, in fact, that the data is potentially useless. If a PIC scheme were employed in such a setting it might be possible to retain the usefulness of the data by reducing the amount of perturbation while ensuring that no user can query enough data to determine the individual, unperturbed data values.² Finally, PIC may be useful in Internet content distribution, as it can allow clients to hide their items of interest from merchants (thus avoiding marketing intrusions) while providing merchants with the guarantee that only prescribed content can be retrieved.

OVERVIEW. This paper is organized as follows. Section 1.1 summarizes our results. We discuss related work in Section 1.2. Section 2 provides notation and Section 3 describes our model of PIC. In Section 4 we show that it suffices to consider inference channels of a certain “threshold” form. In Section 5 we present a simple protocol that introduces our techniques. Section 6 contains a stateful protocol with optimal server work and private user revocation based on query histories. Section 7 describes our PIC scheme that is simultaneously optimal in communication/round complexity and server work per query. We conclude in Section 8. Appendix A discusses some of the cryptographic tools we rely upon, and Appendices B and C supply proofs for Sections 4 and 6, respectively.

²However, with such an approach the user can learn the structure of the inference channels (but not the data that completes the channels) which may not occur when data is protected solely by perturbation. See Section 3 for more on this point.

1.1 Results

Besides the contribution we make by rigorously defining the security model for private inference control, we also achieve extremely efficient implementations. Note that one possible implementation is just to run a generic secure function evaluation (also referred to as a secure circuit evaluation) [30, 50]. However, the communication complexity of such schemes is at least linear in the input size. In contrast, we obtain solutions with substantially sublinear communication complexity. Our key idea to achieve low communication is to carefully divide up the keys/state of our schemes so as to run secure function evaluation on a small set of inputs while running a private information retrieval scheme on a larger set of inputs. Of course, the difficulty of doing this is that this division of inputs must respect user privacy and inference control. Our final scheme in section 7 accomplishes this division. Plugging in the best known PIR parameters [13, 36], we obtain:

Result. *For databases of size n , there exists a 1-round PIC scheme with total communication $O(\text{polylog}(n))$ and server work $O(n\text{polylog}(n))$. (Section 7)*

In addition, in Section 6 we present a scheme that incurs more overhead but has the benefit of supporting private user revocation. That is, the server can ensure that users who have accessed certain information are no longer able to query the database *without* knowing who these users are.

As a warmup to these schemes we present a scheme Section 5 achieving low communication when the user knows all of his queries in advance.

Finally, we present a reduction that shows that one can focus on designing PIC schemes for which the inference channels take a particularly simple “threshold” form and transform them into PIC schemes handling arbitrary inference channels without much cost in efficiency for a certain range of parameters.

1.2 Related Work

Private information retrieval (PIR) was introduced in a seminal paper by Chor *et al* [16]. In such schemes a server holds an n -bit binary string $x \in \{0, 1\}^n$, representing a database, and a user has some index $i \in [n]$. At the end of the protocol the user should learn x_i and the server should learn nothing about i . This notion of privacy was later relaxed [13, 34, 36] to prevent any computationally-bounded server from learning anything about i . See [4] for a survey of PIR.

A *symmetrically private information retrieval* (SPIR) protocol [23] guarantees that the user learns nothing about any index in x other than the one requested, in addition to providing user privacy. Note that neither SPIR nor PIR accounts for changes in a user’s permissions over time; these permissions must change in order to enforce inference control. However, both PIR and SPIR protocol are useful building blocks for PIC and are used in our protocols.

SPIR is essentially equivalent to the older notion of 1-out-of- n oblivious transfer (OT) [42, 9], where the emphasis in SPIR is on the communication complexity. The oblivious transfer works that bare the most similarity to PIC are [3, 40]. In [40], Naor and Pinkas present schemes enabling a user to query a database adaptively for at most a parameter τ number of times, while preserving the privacy of both the user and the database. Hence, [40] can be viewed as studying a simplified version of our problem in which the database consists of a single record and the inference channels are all sets of $\tau + 1$ attributes. It is not clear that the techniques of [40] can be efficiently extended to our setting. In particular, the natural extension in which one instance of a protocol from [40] is run for each record results in linear communication and doesn’t offer the privacy of our protocols since the server knows when each record is queried.

In [3], Aiello, Ishai and Reingold also only consider the single record setting and the natural extension of their protocols to multiple records is less efficient, and less privacy-preserving, than our protocols. However, the schemes in [3] have the advantage of allowing for more general inference channels than those in [40]. Phrased in terms of our setting, in [3] a price is associated with each attribute of a record and a user can only query attributes until their budget is exceeded. Although this allows for a wider variety of inference

channels, it is not as general as the setting we consider. We consider inference channels that are subsets of the attributes (this is made more precise in Section 3), and so an inference can be drawn when all of the attributes in a subset are received for any particular record. As an example of inference channels that cannot be realized in the “price” model of [3], consider the following two inference channels (A_i denotes the i th attribute in the database), $I_1 = \{A_1, A_2\}, I_2 = \{A_3, A_4\}$. In order to provide inference control with the protocols of [3], there must be at least one attribute in each with a price of more than $1/2$ of the budget, but then no two of these attributes can be “bought” given the budget, however no such pair of attributes forms an inference channel.

An alternate approach to inference control, and access control in general, is data perturbation [5, 48, 2]. Data perturbation has the benefit that queries are never blocked and hence the user gains no information about what’s sensitive in the database. This property is especially important when inference channels are data dependent. However, recent work [18] suggests that in some settings the amount of perturbation that must be applied to the data to guarantee privacy may be so much as to render the data useless.

A solution to the PIC problem is possible using techniques for search on encrypted data (see, [46] and the more recent work [11, 22]). The idea is that the records a user has queried are stored in encrypted form on the server and the user gives the server the capability to determine whether or not the current record of interest has been queried too much. With such a solution communication cost is essentially the cost of the SPIR protocol employed. However such an approach doesn’t offer complete privacy to the user as the server is able to determine exactly how many times a record has been queried and exactly when it was queried, even though the correct index of the record may be hidden from the server.

Tools and techniques for ensuring inference control without user privacy can be found in [43, 32, 47].

2 Preliminaries

We make use of the following tools in this paper: computational symmetrically-private information retrieval on records, homomorphic encryption, non-malleable encryption, and zero-knowledge proofs of knowledge. Background on these tools, including a specific construction of SPIR on records used in our protocols, is found in Appendix A. For simplicity, we will use the term SPIR to refer to our specific construction of SPIR on records of size k , where throughout, k is a security parameter.

For an integer m , $[m]$ denotes the set $\{1, \dots, m\}$. Further, let $2^{[m]}$ denote the set of all subsets of $[m]$.

For a vector s , s_i refers to its i th coordinate, and if s_i is itself a vector, $s_{i,j}$ denotes the j th coordinate of s_i . We repeat this notation indefinitely, so if $s_{i,j}$ is also a vector, $s_{i,j,k}$ denotes its k th coordinate. For $i \leq j$, let $s_{i,\dots,j}$ denote the $(j - i + 1)$ -tuple $(s_i, s_{i+1}, \dots, s_{j-1}, s_j)$.

For two strings or vectors s and t , let $s \circ t$ denote their concatenation. Let $|s|$ denote the length of s .

By $\eta(a, b)$ we denote an arbitrary negligible function, i.e., a function of a, b which is less than any inverse polynomial in a, b for a and b sufficiently large.

Two families of random variables U_n and V_n are computationally indistinguishable if for all probabilistic polynomial time (PPT) algorithms A , $|\Pr[A(U_n) = 1] - \Pr[A(V_n) = 1]| < \eta(n)$.

The notation \tilde{O} suppresses terms that are polylogarithmic in the number of database records n .

3 The Model

In this paper we restrict our attention to single-server computationally-private schemes. We first consider honest users, U , and the honest server, S . We consider malicious users U^* in our definition of inference control, and honest but curious³ servers S^* in our definition of user privacy. We require all users and servers to be efficient probabilistic algorithms. For notational convenience, we assume all entries in the database are single bits, but our definitions easily extend to handle entries in $\{0, 1\}^l$ for constant l .

³Honest but curious means that, as a black box, the server behaves the same way as the honest server, but the server may keep all messages from a user and run arbitrary efficient algorithms on them.

A database is a string $x \in (\{0, 1\}^m)^n$. x_i denotes the i th *record* of the database, and $x_{i,j}$ denotes the j th *attribute* value of the i th record. In our general asymptotic analysis, we assume the number of attributes, m , is at most $O(\log \log n)$, whereas the number of records n is very large, as is the case for many relational databases. In the important case when the inference channels take a threshold form, or are otherwise easy to describe, we show how to dispense with this assumption on m . In fact in section 7 we give a best-possible PIC implementation for any m . Even in section 7, we will however assume that m is at most polynomially larger⁴ than n .

We assume that given the description of x , there is a mechanism for generating a collection \mathcal{C} of sets $F \subseteq [m]$ denoting the inference channels in x . The meaning of \mathcal{C} is that, for all $i \in [n]$ and $F \in \mathcal{C}$, the user should not learn $x_{i,j}$ for all $j \in F$. A tool such as the one in [43] can be used to determine \mathcal{C} . We take \mathcal{C} to be an input to the server. We sometimes think of \mathcal{C} and $F \in \mathcal{C}$ as sets, and other times as tuples.

As discussed in [18, 44, 43], a driving motivation is that we are concerned with information that is sensitive because it identifies individuals. We can think of an individual as a record. An individual is more likely to be identified the more of their attributes are revealed. For example, it's almost impossible to identify someone from zip code alone, but there might be a single person with a given zip code, car make and place of work, and if so, these attributes constitute an inference channel. Hence, our inference channels consist of subsets of attributes, and inference control ensures a user doesn't learn all the attributes in such a subset for any particular record. If one insists on an inference channel existing between some set S of distinct records in the database, one can merge the set S into a single record and apply our schemes with an efficiency cost.

Note that as is clear in the above example, we assume \mathcal{C} is monotone, that is, if $A \in \mathcal{C}$ and $A \subseteq B$, then $B \in \mathcal{C}$. We also assume \mathcal{C} is nonempty (else, PIR trivially solves our problem) and that \mathcal{C} is an input to the *user*. Indeed, we cannot expect to hide \mathcal{C} from the user since the user can probe the database, see which queries are blocked, and learn information about \mathcal{C} . Note that \mathcal{C} describes the *structure* of the database, and does not reveal any actual database entries.

There are two stages in a private inference control scheme: the offline stage and the online stage. In the offline stage the user and server are given their inputs, and the server is allowed to preprocess the database. In the online stage an honest user generates coordinates of a $|T|$ -tuple, $T = ((i_1, j_1), \dots, (i_{|T|}, j_{|T|}))$ one at a time, where (i_t, j_t) is allowed to depend on all previous interactions with the server. In the l th execution of the online stage an honest user should output x_{i_l, j_l} . The fact that the user generates coordinates of T one at a time reflects the fact that a user typically queries a database adaptively. However, in section 5, as an introduction to our techniques, we consider the case when T is known in advance.

For simplicity we will only require correctness when T consists of distinct pairs. Hence, $|T| \leq (m-1)n$. However, all of our protocols can be easily modified to allow for repeat queries (see Section 8). Also, we give a direct construction of a PIC scheme in section 7 which handles both repeated queries and rejected queries. That is, even if a user has a query blocked because it is inference-enabling, they can continue to make future queries and retrieve database entries so long as these queries are not inference-enabling.

We call a query sequence T of distinct pairs *permissible* if it doesn't complete any inference channels, that is, if for all $F \in \mathcal{C}$ and all $i \in [n]$, there exists an $\ell \in F$ such that $(i, \ell) \notin T$. We can think of $T = T(U, x)$ (here U denotes the code of U) as a random variable induced by the uniform distribution on ρ and γ , where ρ and γ are random strings stored by the user and server, respectively. If U is honest, T assumes a particular permissible query sequence for fixed ρ and γ .

We proceed to define the algorithms and stages that the user and server run in any PIC protocol. In every algorithm, n, m, \mathcal{C} , and 1^k are inputs, where $k = k(n)$ is a security parameter. For notational convenience we do not include these inputs in the algorithms' descriptions.

Private Inference Control: A PIC protocol consists of the following stages and algorithms.

Offline Stage: The user gets a random string ρ and the server a random string γ . $|\rho|$ and $|\gamma|$ are polynomials in n and k . The same ρ (resp. γ) will be an input to every user (resp. server) algorithm for every

⁴Notationally this means that $\tilde{O}(\text{polylog } m) = \tilde{O}(1)$

interaction with the server (resp. user). We may assume all user and server algorithms are deterministic given ρ and γ . The server is given $x \in (\{0, 1\}^m)^n$ and both parties are given $\mathcal{C} \in 2^{[m]}$. The server runs an efficient *preprocessing algorithm* $P(\cdot, \cdot)$, which outputs a string $y = P(x, \gamma)$.

Each execution of the online stage is allowed to be a multiround protocol, and we let r denote the number of rounds per execution. Without loss of generality we may assume r is the same in every execution. In the w th round of the l th execution the user computes and sends a message $m_{w,l}$ to the server and the server computes and sends an answer $a_{w,l}$ back to the user. Define:

$$M_{w,l} = \{m_{i,j} \mid 1 \leq j < l \text{ and } 1 \leq i \leq r, \text{ or } j = l \text{ and } 1 \leq i \leq w\},$$

and define:

$$A_{w,l} = \{a_{i,j} \mid 1 \leq j < l \text{ and } 1 \leq i \leq r, \text{ or } j = l \text{ and } 1 \leq i \leq w\}.$$

Note that $M_{0,l} = M_{r,l-1}$ and $A_{0,l} = A_{r,l-1}$. During the protocol the user will have some state $\sigma_U(A_{w,l}, \rho)$ and the server some state $\sigma_S(M_{w,l}, x, \gamma)$ for some w and l . For instance, $\sigma_U(A_{w,l}, \rho)$ and $\sigma_S(M_{w,l}, x, \gamma)$ could contain the entire message history.

Online Stage : Let $Q(\cdot, \cdot, \cdot, \cdot)$, $D(\cdot, \cdot, \cdot, \cdot)$ and $R(\cdot, \cdot, \cdot, \cdot)$ be efficient algorithms. Q is the *query generator*, D the *database algorithm*, and R the *reconstruction algorithm*. The following process is run:

- For $1 \leq l \leq |T|$,
 1. The user generates $T_l = (i_l, j_l)$ so that $T_{1,\dots,l}$ is permissible.
 2. For $1 \leq w \leq r$,
 - (a) The user computes $m_{w,l} = Q(i_l, j_l, \sigma_U(A_{w-1,l}, \rho), \rho)$ and sends $m_{w,l}$ to the server.
 - (b) The server computes $(a_{w,l}, \sigma_S(M_{w,l}, x, \gamma)) = D(m_{w,l}, x, y, \sigma_S(M_{w-1,l}, x, \gamma), \gamma)$, and sends $a_{w,l}$ to the user.
 - (c) The user computes $\sigma_U(A_{w,l}, \rho) = R(a_{w,l}, i_l, j_l, \sigma_U(A_{w-1,l}, \rho), \rho)$.
If $w = r$, $R(a_{w,l}, i_l, j_l, \sigma_U(A_{w-1,l}, \rho), \rho)$ also outputs out_l .

We define the *view*, $V_{U^*}(x, \rho, \gamma)$, of an arbitrary user, U^* , with respect to an arbitrary honest but curious server to consist of ρ , and $A_{r,|T|}$. We define the *view*, $V_{S^*}(U^*, x, \rho, \gamma)$, of an arbitrary server S^* with respect to an arbitrary user U^* to consist of x, γ , and $M_{r,|T|}$. As with the algorithm descriptions, n, m, \mathcal{C} , and 1^k are also part of user and server views.

Correctness and Security Definitions: The algorithms P, Q, D, R constitute a PIC protocol if the following hold for all n, m, k , and monotone $\mathcal{C} \in 2^{[m]}$:

1. **Correctness:** For all $x \in (\{0, 1\}^m)^n$ and all honest users U ,

$$\Pr_{\rho, \gamma}[\forall \alpha \in [|T(U, x)|], \text{out}_\alpha = x_{i_\alpha j_\alpha}] = 1.$$

2. **User Privacy:** For all $x \in (\{0, 1\}^m)^n$, for all honest users U and any two sequences T_1, T_2 with $|T_1| = |T_2|$, for all γ and for all honest but curious servers S^* , S^* can distinguish user query sequence T_1 from sequence T_2 with only negligible probability. That is, for all efficient (in $n, 1^k$) algorithms A ,

$$|\Pr_{\rho} [A(V_{S^*}(U, x, \rho, \gamma)) = 1 \mid T(U, x) = T_1] - \Pr_{\rho} [A(V_{S^*}(U, x, \rho, \gamma)) = 1 \mid T(U, x) = T_2]| < \eta(n, k),$$

where the probability is also taken over the coin tosses of the adversary's algorithm A , where A is additionally given⁵ n and 1^k .

⁵Note that since $m = O(1)$, A does not have to be given \mathcal{C} and m since they can be hardwired into A , even if A is uniform.

3. Inference Control: We make a comparison with the ideal model in which a trusted third party, Charlie, when given $x \in (\{0, 1\}^m)^n$ and a *permissible* T , gives the user $x_{i,j}$ for all pairs $(i, j) \in T$. More precisely, we require that for every U^* and every ρ , there exists an efficient U' , given the code of U^* and ρ , that plays U^* 's role in the ideal model, such that for every $x \in (\{0, 1\}^m)^n$, U' can find a permissible T such that the output of U^* interacting with the honest server and the output of U' interacting with Charlie given x and T are computationally indistinguishable. Note that U' is also allowed to make queries adaptively, so T may depend on x . In general, the T found by U' is a random variable induced by the uniform distribution on γ . Observe that the inference channels, \mathcal{C} , are known to both U^* and U' and are part of their views.

We phrase our definition in terms of a single user for ease of exposition. Clearly, for collusion-resistance it does not suffice to protect against inferences by a single user. It is straightforward to modify our protocols to protect against collusion – we simply permit each user to access less of each channel. The same technique is used in [47] to ensure collusion-resistant inference control without privacy, however it is easily adopted in our setting as it only depends on server knowledge of the number of queries a user has made *not* the content of the queries. Adjusting the definition is similarly straightforward, one would instead consider malicious users, U_1^*, \dots, U_c^* , where c is the desired level of collusion resistance. We revisit collusion resistance in Section 8.

Since our goal is to ensure that a malicious user doesn't learn any unintended information, the above definition is similar to what is termed “sender's security” in [40] and “vendor's security” in [3]. However, we differ from both notions in that inference channels are a more general concept than the query counts in [40] and the query prices in [3], that determine access controls in those works. We could also define inference control and user privacy to hold with respect to non-uniform adversaries (and in fact our schemes can be made secure against such adversaries), but for ease of exposition, we do not take that up here.

Efficiency: For fixed n, m , and k , we measure the efficiency of a PIC protocol as follows. All of our efficiency measures are *per query*, that is, for a single execution of the online stage. The *communication complexity* $C(n, m, k)$ is defined to be: $\max_{\mathcal{C}, x \in (\{0,1\}^m)^n, U, \gamma, \rho, l} \sum_{w=0}^r (|m_{w,l}| + |a_{w,l}|)$. The *server work* of a PIC protocol $W(n, m, k)$ is the max, over all $\mathcal{C}, x, U, \gamma, \rho$, and l , of the sum of the running times of $D(m_{w,l}, x, y, \sigma_S(M_{w-1,l}, x, \gamma), \gamma)$ over w . The *round complexity* of a PIC protocol $R(n, m, k)$ is r .

Parameter Summary:

Inputs: A database x , random user input ρ , random server input γ , collection of inference channels \mathcal{C} , number of records n , number of attributes m , and security parameter k .

Components: Preprocessing algorithm P , query generator Q , database algorithm D , and reconstruction algorithm R .

Intermediate Variables: Query sequence T for honest users U , messages sent from user to server $M_{w,l}, m_{w,l}$, messages sent from server to user $A_{w,l}, a_{w,l}$, outputs of user out_l , state of user $\sigma_U(\cdot, \cdot)$, state of server $\sigma_S(\cdot, \cdot, \cdot)$, preprocessing output string y and round complexity r .

There is some redundancy in the notation because the intermediate variables are determined given the inputs and components. However, at times we will use the intermediate variables for simplicity.

4 A Generic Reduction

We show that it suffices to focus on developing protocols that meet the definition of PIC when the inference channels, \mathcal{C} , take a particularly simple form. We note that for m larger than $O(\log \log n)$, this reduction is not particularly efficient. Thus, in section 7 we directly construct a PIC scheme with an exponential savings

over that achieved via this reduction.

Definition 1 *A protocol is a **threshold private inference control scheme**, abbreviated TPIC, if it satisfies the requirements of a private inference control scheme under the assumption that $\mathcal{C} = \{[m]\}$.*

Note that a TPIC need not satisfy the requirements of a PIC for $\mathcal{C} \neq \{[m]\}$. However, the following theorem shows that given any TPIC, we can construct a PIC with related efficiency (for m not too large):

Theorem 1 *Let (P', Q', D', R') be a TPIC with communication $C'(n, m, k)$, server work $W'(n, m, k)$, and round complexity $R'(n, m, k)$. Then there exists a PIC scheme (P, Q, D, R) with communication complexity $C(n, m, k) = m2^m C'((m+1)2^m n, m, k)$, server work $W(n, m, k) = m2^m W'((m+1)2^m n, m, k)$, and round complexity $R(n, m, k) = m2^m R'((m+1)2^m n, m, k)$.*

Proof Sketch: The idea behind the proof is to make a new record for each channel in an original record. The new records have variable size between 1 and m , so we pad them with random bits so that they all have the same length. Bits in the pad are XORed to values in the record in order to enable the record to be queried exactly $m - 1$ times. Because of the expansion of each record, an attribute value may now occur in multiple records. We force the user to read all records containing this element by using secret sharing to “split” the element across each of these records. The only way to recover the element is to query all records containing it. The formal proof, together with a small example, can be found in Appendix B.

Again, we stress that this is a general theorem which constructs a PIC out of any TPIC, with complexity exponential in m . Note that $|\mathcal{C}|$ itself may be exponential in m , so this reduction is fairly efficient for general $|\mathcal{C}|$. In section 7 we provide a protocol with much better complexity though, which will be useful in the multi-user setting when m may be large.

5 Warm-up: PIC with Static Queries

To introduce some of the techniques behind our PIC protocols we first consider the case in which the user has a fixed set of queries T , or a query *script*, to execute. When the queries are known to the user a priori it’s possible to require the user to execute the queries in a certain order. This greatly simplifies the access control problem because it reduces the number of permissible query sequences, and thus makes it simpler to encrypt the database contents so that a user who is attempting to make an inference channel-completing query will not be able to decrypt the requested content. Our warm-up protocol is a TPIC protocol, $\mathcal{C} = \{[m]\}$.

The protocol works as follows. Upon receiving the s th query from an honest user U , the server forms a table with entries consisting of each data item $x_{i,j}$ concatenated with a key. When $s \geq m$ these entries are encrypted using a symmetric encryption algorithm $E(\cdot, \cdot)$, where the first input is the encryption key (we denote $E(k, \cdot)$ by $E_k(\cdot)$). The encryptions are generated under a set of keys that are only held by the user if they’ve made a permissible set of queries thus far *and* are following the protocol by executing entries of T in increasing order. By the latter we mean that records (i.e. the n rows of the matrix, $x \in (\{0, 1\}^m)^n$) are queried in increasing order and for a given record i , if (i, j_1) is queried before (i, j_2) , then $j_1 < j_2$. Hence, it is permissible for a user to query record i on their s th query if either a record $j < i$ was queried on their $(s - 1)$ th query *or* their $(s - 1)$ th query was also to record i but they queried a record $j' < i$ on their $(s - m + 1)$ th query. Because the variation amongst the permissible query sequences is thus reduced, the communication overhead is on the same order as the cost of a SPIR invocation on a database of size $O(kn^2)$.

Protocol 1 *Consider the following static set of queries $T = ((i_1, j_1), \dots, (i_t, j_t))$, where $1 \leq t \leq (m - 1)n$. We describe how the protocol proceeds on the s th query, (i_s, j_s) . U ’s key set, K_U , is initially empty. Let $H : \{0, 1\}^* \rightarrow \{0, 1\}^k$ be a cryptographic hash function, where k denotes the security parameter.*

1. *Query Generator, Q:* For the s th query, the user, U , and the server engage in SPIR protocol (see Appendix A) on inputs (i_s, j_s) , to generate a query, m_s , that is sent to the server.

2. *Database Algorithm, D:* To process the s th query from U , the server first selects n values at random from $\{0, 1\}^k$: $k_{1,s}, \dots, k_{n,s}$. For $s \geq m$, let $K_{i,j}^s = H(k_{i,s-1}, k_{j,s-m+1})$ for $1 \leq j \leq i-1$. The server forms the following sequence of concatenated values for $1 \leq i \leq n$: $y_{i,1} = (x_{i,1} || k_{i,s})$, $y_{i,2} = (x_{i,2} || k_{i,s})$, \dots , $y_{i,m} = (x_{i,m} || k_{i,s})$.
 - (a) If $s \leq m-1$, the server forms the following table, \mathcal{T}_s , with n rows and m columns, $(y_{i,j})_{1 \leq i \leq n, 1 \leq j \leq m}$.
 - (b) If $t \geq s > m-1$, the server forms a table \mathcal{T}_s with n rows and $2(n-1)m$ columns ($2n(n-1)m$ entries in total), where the first $2(i-1)m$ entries in row i are:
$$\{E_{k_{1,s-1}}(y_{i,j}), E_{k_{2,s-1}}(y_{i,j}), \dots, E_{k_{i-1,s-1}}(y_{i,j}), E_{K_{i,1}^s}(y_{i,j}), \dots, E_{K_{i,i-1}^s}(y_{i,j})\}_{j=1, \dots, m}$$
 - (c) The server executes the SPIR protocol on inputs \mathcal{T}_s and m_s and sends its response, a_s , to the user.
3. *Reconstruction R:* The user recovers the desired data value x_{i_s, j_s} and adds $k_{i_s, s}$ to K_U . This recovery step requires decryption using keys in K_U in steps $s > m-1$.

We do not give a formal proof of the security of this protocol since it is merely presented as an example of our approach. Note that correctness easily holds when queries are executed in order because if the user is not in danger of completing an inference channel with their s th query, they will have a key for another row from their previous query, or a key for the same row for their previous query *and* a key for an earlier row from the query they made $m-1$ queries ago. If the user doesn't have such keys, they will not be able to decrypt, which is the basis for inference control. This protocol doesn't completely preserve user privacy since there are always query sequences that the server knows could not have been executed, however within these constraints, the queries inherit privacy from the SPIR protocol.

6 Stateful PIC with Linear Work

One of the main difficulties in constructing PIC protocols is that the large number of permissible query sequences puts upward pressure on the size of the SPIR tables employed, and thus, on the communication complexity. The warm-up scheme side-steps this by dramatically reducing the number of permissible query sequences at the cost of non-adaptive queries. The scheme of this section manages communication cost by allowing the server to store encrypted information about each user's queries. This has several significant benefits. First, it allows for *adaptive* queries and second it keeps the server's work at essentially an optimal level, $O(k^2n)$, where k is a security parameter taken to be $\tilde{O}(1)$ in the following.⁶ Note that PIC necessarily implements PIR, and the best known single-server PIR scheme, even with preprocessing, uses $\Omega(n)$ work per query.

Another benefit of this approach is that the storage of encrypted query histories can facilitate user revocation as an adjustment to changes in the inference channel. If, for example, it is determined that certain information in the database is sufficiently sensitive that an user who has accessed it should no longer be permitted to query the database this can easily be privately "checked" by the database and all future queries by such a user rebuffed. User privacy is preserved since the server doesn't know whether the user passed or failed the check. Further, the presence of the encrypted query histories makes it easy to adapt to changes in the inference channels (likely with a dynamic database).

The protocol makes use of a homomorphic encryption function, $E^{hom}(\cdot)$ (instantiated with Paillier's scheme of Appendix A, for example). With $E^{hom}(\cdot)$ the user can privately send query information to the

⁶That is, we are ignoring logarithmic factors and using $\tilde{O}(1)$ for $k = O(\log^r(n))$, as is done for security against PPT (in n) algorithms for the best known PIR/SPIR schemes [13, 36, 39].

server. Then, using the homomorphic property of $E^{hom}(\cdot)$, the server can encrypt a secret, \mathcal{S} , in such a way that the user can only recover \mathcal{S} if not in danger of making an undesired inference with the user's current query. Finally, the user and server engage in SPIR⁷ (see Appendix A) on a table encrypted under \mathcal{S} and the user's encrypted query information. Hence, recovery of \mathcal{S} effectively authorizes the user to receive the answer to their query.

Protocol 2 Let U 's t^{th} query be $T_t = (i_t, j_t)$, $1 \leq i_t \leq n$, $1 \leq j_t \leq m$ and let $\mathcal{C} = \{[m]\}$. We describe how this TPIC protocol proceeds on the t^{th} query.

1. If $t = 1$, U generates public and private keys for $E^{hom}(\cdot)$ and sends the public key to the server.
2. Sending query information: U sends the server $E^{hom}(i_t)$ and $E^{hom}(j_t)$ and gives a zero-knowledge proof of knowledge that the ciphertexts are well-formed.
3. Generating authorization:
 - (a) If $t < m$ then the server sets $\mathcal{S}_t = 0$ and proceeds to step 4. Else, the server generates a secret \mathcal{S}_t and randomly generated shares, y_1, \dots, y_{t-1} , forming a $(t - m + 1)$ -out-of- $(t - 1)$ sharing of \mathcal{S}_t [45] and sends the user $E^{hom}((i_1 - i_t)y_1), \dots, E^{hom}((i_{t-1} - i_t)y_{t-1})$ using the homomorphic property of $E^{hom}(\cdot)$.
 - (b) U decrypts and recovers at least $t - m + 1$ of the values in $\{y_1, \dots, y_{t-1}\}$ if U has made a permissible sequence of queries, and thus is able to reconstruct the secret \mathcal{S}_t .
4. Query Processing:
 - (a) The server generates random values, $v_{i,j}^{(1)}, v_{i,j}^{(2)}$, for every $1 \leq i \leq n$, $1 \leq j \leq m$, and forms the table with nm entries $\mathcal{T} = (E^{hom}(v_{i,j}^{(1)}(j - j_t) + v_{i,j}^{(2)}(i - i_t) + \mathcal{S}_t + x_{ij}))_{i,j}$.
 - (b) U and the server engage in the SPIR protocol on \mathcal{T} . In order to be able to recover $x_{i,j}$, U must know \mathcal{S}_t and the following equalities must hold: $i = i_t$ and $j = j_t$. Hence, U can at most recover x_{i_t, j_t} , and can only do so if U has made a permissible sequence of queries.

We discuss the key points pertaining to the protocol's security here and refer the reader to Appendix C for the details. Inference control intuitively holds because step 2 ensures that a user who is in danger of completing an inference channel won't be able to recover \mathcal{S}_t because in step 2(a), for at least $m - 1$ values of j , $E^{hom}((i_j - i_t)y_j) = E^{hom}(0)$. The use of encryption in the SPIR phase⁸ ensures that the user will only be able to decrypt entries in the record the user is authorized to query. Note that unlike the warm-up protocol, a user who is blocked from making an inference channel-completing query *will* be able to continue to receive information from the database (provided subsequent queries are permissible given past ones). User privacy relies upon the SPIR protocol and the strength of the encryption schemes.

The per query communication cost is $O(kt)$ plus the cost of a SPIR invocation. As mentioned in Section 1.2, the $O(kt)$ part of the cost can be removed with a decrease in user privacy, by using techniques for search on encrypted data [46, 11, 22] to enable the server to determine whether or not a server is authorized by searching the user's encrypted query history, and thus removing step 2. Server work is essentially optimal in this protocol at $O(k^2n)$ with the bulk of the work occurring in step 3.

Applying Theorem 1 turns this TPIC scheme into a PIC scheme with $\tilde{O}(n)$ server work per query and $\tilde{O}(t)$ communication per query, which may be $\tilde{O}(n)$, depending on t .

Finally we note that since the server stores the encrypted query histories it is simple to accommodate user revocation based on the user's history. As an extreme example, if it is determined that access to $x_{i,j}$ is

⁷PIR may well be sufficient in order to achieve inference control here, however we opt for the stronger SPIR approach in order to demonstrate security more easily (see Appendix C).

⁸PIR may suffice instead of SPIR, however we use the stonger primitive SPIR to facilitate the security proof (Appendix C).

sufficiently valuable that a user with that information should no longer be allowed to access the database, then the server can easily “check” for $E(i)$ and $E(j)$ using the same secret sharing techniques of step 3. This does not immediately compromise the user’s privacy as the server does not know that the user was unable to recover the secret, although the user should maintain a consistent query pattern in order to ensure that the server does not learn they were revoked.

7 One-round PIC with Polylogarithmic Communication and Linear Work

In this section we present, up to logarithmic factors, a PIC protocol that is simultaneously optimal in its communication complexity, round complexity, and server work per query. Again, as in Section 6, we assume the security parameter $k = \tilde{O}(1)$.⁹ Our protocol achieves communication complexity $\tilde{O}(1)$, uses 1 round, and has $\tilde{O}(n)$ work per query. This is optimal in all three aspects with respect to known PIR schemes, even with preprocessing [7].

Our scheme enjoys a number of additional features:¹⁰

1. It directly allows for repeat queries.
2. It allows the user U to continue querying the database even if he tried to make an inference-enabling query at some point in the past. Hence the collection of inference channels \mathcal{C} need not be a user input.
3. The only information about U ’s query history the server S needs to store is the total number of queries made thus far.
4. It directly handles arbitrary inference channels instead of going through Theorem 1.
5. The work and communication depend only linearly on m and the description of \mathcal{C} , and hence m may be $O(\log \log n)$ without asymptotically affecting the work or communication. Note that this is best possible, since $|\mathcal{C}|$ may be as large as 2^m in any inference control scheme.
6. In the special case of threshold inference channels, the communication is $O(\text{polylog}(mn))$ and the work is $\tilde{O}(mn)$, so m may even be polynomial in n without asymptotically affecting the communication. Observe that $\Omega(mn)$ work is necessary for any value of m with respect to known PIR schemes.
7. The space complexity of S is dominated by the size of the database itself.
8. The server just needs $\tilde{O}(1)$ bits of randomness.

Without loss of generality, we assume n is a power of 2. We use a data structure B consisting of a balanced binary tree with n leaves, where in addition, we connect m children to each leaf of the binary tree. The idea is that the leaves of B denote entries $x_{i,j}$ of the database, and the parents of the leaves denote records x_i . U will obtain keys $K(w, z)$ associated with each leaf w indicating whether or not the value at the leaf has been accessed. Here $z \in \{0, 1\}$. Internal nodes w also have associated keys $K(w, z)$, where here z is an integer indicating the total number of times leaves in the subtree rooted at w have been accessed. These keys will be used to do inference control. When a user U retrieves a database entry they use their keys to go up the tree. If U tries to use “older” keys indicating that nodes have been accessed less than they actually have, their keys will be inconsistent with S ’s knowledge of the total number of queries made thus far, and U will not be able to recover the desired database entry.

⁹Such is the case for security against PPT (in n) algorithms for the best known PIR/SPIR schemes (e.g., transforming the PIR protocol of [13] to a SPIR protocol using [39].)

¹⁰Features 1 and 2 also hold for the scheme of Section 6. We view feature 3 as a significant advantage from a security standpoint, however as discussed in Section 6 storing state can facilitate useful access management features such as user revocation.

Before going into detail, let us fix some notation. Let α denote the root of B . We say $w \in B$ is at height d if it is d levels above the leaves. The leaves are at height 0. Each node in B of height 1 is denoted by i for some $i \in [n]$, and each of the m children of i are denoted by (i, j) for some $j \in [m]$ (we sometimes use w to refer to an arbitrary node in B , including those at height 0 and 1). For a non-root node w in B , let $\text{sib}(w)$ denote w 's siblings (there is either 1 or $m - 1$). For a non-leaf node w , let $\text{children}(w)$ denote w 's children. For a leaf node w , let $\text{anc}(w)$ denote the set of $\log n + 1$ ancestors along the path from w to α , inclusive. We say a node w is *accessed* whenever $x_{i,j}$ is successfully retrieved by the user for which $w \in \text{anc}(i, j)$. Finally, for leaves w define the set of $2 \log n + m - 1$ nodes that is the set of ancestors (as defined above) together with the siblings of the ancestors:

$$\text{sibanc}(w) = \text{anc}(w) \cup \{u \mid \exists v \in \text{anc}(w) \text{ s.t. } u = \text{sib}(v)\}$$

7.1 Protocol Description

When the honest user queries $x_{i,j}$, he will use the set of keys $\pi = \{K(w, f_w) \mid w \in \text{sibanc}(i, j)\}$, where f_w is the number of times w has been accessed. If the user is dishonest, for some $w \in \text{sibanc}(i, j)$, he may substitute $K(w, z)$ in place of $K(w, f_w)$ for some integer $z \neq f_w$. However, the invariant we maintain is that with all but negligible probability, any user U^* cannot determine $K(w, z)$ for any $z > f_w$, and so if $K(w, z)$ is substituted for $K(w, f_w)$, it holds that $z < f_w$. If the user is given $x_{i,j}$, he will also obtain the *updated* set of keys $\{K(w, f_w + 1) \mid w \in \text{sibanc}(i, j)\}$.

Our scheme enforces inference control by what we call *sum-consistency*. For any non-leaf node w and its children, $\text{children}(w)$, we say the keys $K(w, i), \{K(u, j_u) \mid u \in \text{children}(w)\}$ are *sum-consistent* if $i = \sum_{u \in \text{children}(w)} j_u$. Suppose the honest user U wants to retrieve $x_{i,j}$ on the $(t + 1)$ st query. We claim the set of keys π gives a *proof* that U is not in danger of completing an inference channel. Indeed, if U is honest, π has the following three properties:

1. For each non-leaf node w in $\text{anc}(i, j)$, $K(w, f_w)$ and $\{K(u, f_u) \mid u \in \text{children}(w)\}$ are sum-consistent.
2. $f_\alpha = t$.
3. If U is not in danger of completing an inference channel by learning $x_{i,j}$, then for all inference channels $F \in \mathcal{C}$, there is some $j' \in F$ such that $j' \neq j$ for which $K((i, j'), 0) \in \pi$.

A dishonest user U^* will not be able to furnish a proof π to obtain $x_{i,j}$ when learning $x_{i,j}$ completes an inference channel F . Indeed, if U^* does not substitute $K(w, z)$ for $K(w, f_w)$ for some $z \neq f_w$ and $w \in \text{sibanc}(i, j)$, the third property above cannot hold. On the other hand, by the invariant mentioned above, if U^* substitutes $K(w, z)$ for $K(w, f_w)$ for some $z \neq f_w$ for some w , then z is necessarily less than f_w so properties (1) and (2) cannot hold simultaneously.

Of course, for user privacy, U cannot simply give π to S . Instead, U will prove knowledge of π via a secure function evaluation (SFE) [30, 50] with S . The idea is that U will input π to the SFE, which will give the user a certain secret if and only if π is a valid proof. Now the communication complexity of an SFE for a function is linear in the circuit size computing the function, and therefore, if the keys $K(w, i)$ are completely independent of each other, the communication complexity of a generic SFE for π will be $\Omega(mn)$, which is prohibitive. Indeed, since the server cannot know which keys the user will use because of user privacy, and since the keys are independent random values, the server must input all possible user keys¹¹ into the circuit so the circuit can perform comparisons on them.

To get by with low communication, we observe that the keys associated with different nodes need not

¹¹Of course, the server can use its knowledge of the number of queries made thus far to exclude some keys from the circuit input. This does not asymptotically help though, once the user makes enough queries.

be totally independent of one another, as long as they are not predictable from each other. Let E be a non-malleable [19] (see Appendix A) symmetric-key encryption scheme, and let K be a secret key of E known only to S . Then for each node w and integer i , we define the key $K(w, i)$ to be the encryption $E_K(w, i)$ under key K . Since E is non-malleable, $E_K(w, i)$ is unpredictable from $E_K(v, j)$ for $(w, i) \neq (v, j)$. The important point here is that the server just needs to input the decryption key K of E to the secure circuit in order to perform the necessary comparisons on the user’s keys, rather than $\Omega(mn)$ truly random keys.

We would also like to handle repeated queries and rejected queries. By the former we mean that a user should be able to query record/attribute pairs already queried and by the latter we mean that even if a user’s query is rejected at some point, he should still be allowed future queries as long as they are not inference-enabling. To this end the user will have keys $E_K(\text{“reject”}, z_1)$ and $E_K(\text{“repeat”}, z_2)$ for integers z_1 and z_2 . The idea is that z_1 denotes the total number of rejected queries made so far, and z_2 denotes the total number of queries that were repeats. Since E is non-malleable and K is unknown to U , U cannot construct/modify such encryptions. These encryptions will be updated by the server “obliviously” in the secure function evaluation we now describe.

The user’s input to the circuit, C , is a set of keys π' , an $i \in [n]$, a $j \in [m]$, and two k -bit numbers P, Q , which if U is honest, denote $E_K(\text{“reject”}, z_1)$ and $E_K(\text{“repeat”}, z_2)$ for some integers z_1, z_2 . The server input is the secret key K , the total number of queries made so far t , the collection of inference channels \mathcal{C} , and a seed s to a pseudorandom function h .

The circuit functionality C first checks if $E_K((i, j), 1) \in \pi'$, which it can do efficiently since it knows K . If this is the case, C checks if Q is of the form $E_K(\text{“repeat”}, z_2)$ for some integer z_2 , and if so, outputs $E_K(\text{“repeat”}, z_2 + 1)$ to the user by decrypting and re-encrypting. If Q is not of this form, C outputs $E_K(\text{“repeat”}, 1)$. In either case C additionally gives $h(s, i, j)$ to U . Otherwise, if $E_K((i, j), 1) \notin \pi'$, C checks that the keys in π' satisfy the properties of a valid proof π , that is, they form a sum-consistent path from (i, j) to α . Note that, we replace the property that $f_\alpha = t$ with the property that $f_\alpha = t - z_1 - z_2$, since the user could not have updated his keys in the $z_1 + z_2$ queries which were rejects or repeats. Note that C can do this efficiently by decrypting the keys in π' using K along with simple arithmetic operations.

Let $g(\pi')$ denote the set of keys $\{K(w, f_w + 1) \mid K(w, f_w) \in \pi'\}$. If π' passes the above test, the user output of C is $h(s, i, j)$ and $g(\pi')$. If π' does not pass the test, C checks if P is of the form $E_K(\text{“reject”}, z_1)$ for an integer z_1 , and if so, outputs $E_K(\text{“reject”}, z_1 + 1)$ to the user. Otherwise, C outputs $E_K(\text{“reject”}, 1)$ to the user. In either case, the server has no output. Since the input size to C is $\tilde{O}(|\mathcal{C}| + m)$ (for $|K|, |s| = \tilde{O}(1)$), the total communication and server work required for this step is $\text{poly}(|\mathcal{C}|, m)$ [30, 50] since C is an efficient computation. Note that, by outputting $E_K(\text{“reject”}, 1)$ and $E_K(\text{“repeat”}, 1)$ we don’t need the user to already know $E_K(\text{“reject”}, 0)$ and $E_K(\text{“repeat”}, 0)$ at the beginning of the protocol.

Finally, the last step is for the user to use $h(s, i, j)$ to recover $x_{i,j}$. To this end, for each $i \in [n]$ and $j \in [m]$, the server prepares a table T with i, j th entry set to $x_{i,j} \oplus h(s, i, j)$. U and S then engage in SPIR¹² on T , and U should obtain $T_{i,j}$, from which he can recover $x_{i,j}$ if and only if he previously received $h(s, i, j)$. Using the best known SPIR schemes, server work in this step is $\tilde{O}(mn)$ and the communication is $\tilde{O}(1)$.

The following summarizes our protocol:

Protocol 3 *In the offline stage the server S randomly chooses a seed s to a pseudorandom function h and a key K to a non-malleable encryption scheme E . In the t -th execution of the online stage, the user U first generates a query (i_t, j_t) . Then U and S perform the following SFE and SPIR reading in parallel. These*

¹²PIR may well be sufficient in order to achieve inference control here, however we opt for the stronger SPIR approach in order to use the simulator of [39].

two steps encompass the query generation algorithm Q , database algorithm D , and reconstruction algorithm R discussed in section 3:

1. Let C be a secure circuit implementing the functionality described above. U constructs the set of keys $\pi = \{K(w, f_w) \mid w \in \text{sibanc}(i_t, j_t)\}$, and feeds these along with keys $E_K(\text{"reject"}, z_1)$ and $E_K(\text{"repeat"}, z_2)$ into C , where z_1, z_2 denote the number of rejected/repeated queries made thus far. If no such queries have been made in one of these two cases (i.e., z_1 or z_2 is 0), then U substitutes a random value in the range of E in its place. S then feeds s, K , and the inference channels \mathcal{C} into C . S gets no output from C , while U 's output is divided into cases (recall the protocol description refers to honest U and S):
 - (a) If learning x_{i_t, j_t} is inference-enabling, U 's output is $E(\text{"reject"}, z_1 + 1)$.
 - (b) If x_{i_t, j_t} was previously queried, U 's output is $E(\text{"repeat"}, z_2 + 1)$ together with $h(s, i_t, j_t)$.
 - (c) Otherwise U 's output is $h(s, i_t, j_t)$ and the updated keys $\{K(w, f_w + 1) \mid K(w, f_w) \in \pi\}$.
2. S prepares a table T with i, j th entry set to $x_{i, j} \oplus h(s, i, j)$ and U and S engage in SPIR on records on T . U should obtain T_{i_t, j_t} .

If learning x_{i_t, j_t} is not inference-enabling, U reconstructs x_{i_t, j_t} from $h(s, i_t, j_t)$ and $T_{i_t, j_t} = x_{i_t, j_t} \oplus h(s, i_t, j_t)$.

7.2 Efficiency Wrap-up

Since server work and communication are dominated by the SPIR invocation, we see that the overall server work of our scheme is $\tilde{O}(mn)$ and the communication is $\tilde{O}(\text{poly}(m, |\mathcal{C}|))$. Finally we note that both the secure circuit evaluation and the SPIR invocation each take one round (with say, using the circuit evaluation of [12]), and further, they can be parallelized. Hence the whole scheme can be done in a single round.

Observe that the only state the server keeps of the user's history is the total number of queries made thus far, which of course, is unavoidable. Also, the secure circuit evaluation directly detects arbitrary inference channels, instead of passing through theorem 1.

In the special case where \mathcal{C} is of the form $\{S \subset [m] \mid |S| \geq \tau\}$, i.e., a threshold access structure, $|\mathcal{C}| = \log m$ as we just need to describe τ . In this case we can modify our scheme so that the communication is $O(\text{polylog}(mn))$ and the work remains at $\tilde{O}(mn)$. To do this, instead of having m keys for each attribute for a given record i , contributing a factor of m to the communication of the SFE, we instead have a single key for each record i encoding the total number of attributes in i which have been queried. Note that by definition of \mathcal{C} , the actual attributes queried are immaterial. This can be done as long as the key size is larger than $\log m$ bits, and hence the communication drops to $O(\text{polylog}(mn))$. We can still handle repeat queries by distributing a special key with each attribute of i , which can be an additional input to the secure circuit evaluation and checked to see if a repeat is being made. Again, this is just one additional key per secure circuit evaluation, so we get an exponential (in m) savings in communication. This allows m to be polynomial in n without asymptotically affecting the communication.

One final point is that by using a pseudorandom function h and key-dependencies through an encryption scheme E , the only randomness the server needs to keep around is the seed s and the key K . This consumes $\tilde{O}(1)$ space. When the server runs SPIR the space complexity is just $\tilde{O}(mn)$, the space necessary to store the database.

7.3 Security proofs

We give proofs that the above protocol meets the requirements of a PIC protocol:

Correctness: It is straightforward from the protocol description to check that for an honest user U , U can recover exactly those queries that are not inference-enabling. We note that the scheme handles repeat queries and rejected queries, as claimed at the beginning of this section.

User Privacy: Observe that an honest but curious server S^* 's view is that of a polynomial (in n) number of SFE and SPIR interactions with an honest U . Note that there are efficient simulators S_1 for the SFE [30, 50] and S_2 for the SPIR (to construct S_2 , simply fix an arbitrary input for U , and simulate the SPIR interaction of U with S^* for that input). Since S^* is honest but curious, its inputs (importantly x) are fixed at the beginning of the protocol, and since S^* gets no output from either the SFE or the SPIR, its view is just the concatenation of the outputs of S_1 and S_2 evaluated on its inputs at the beginning of the protocol. This defines a simulator for the view of S^* , which implies user privacy. For more background see [26].

Inference Control: Consider an arbitrary user U^* with random tape p . As with user privacy, there is an efficient simulator S_1 for the SFE simulating U^* 's view with the honest server S . Further, using the PIR to SPIR transformation of [39], there exists a simulator S_2 which given U^* 's code and p , can extract the indices requested (explicitly or implicitly) by U^* in any of the polynomial number of SPIR interactions of U^* with S .

Let U' play U^* in the ideal model. U' first chooses a random key K for E and a seed s for the pseudorandom function h . Suppose after t queries (including repeated/rejected queries), the view generated by U' is computationally indistinguishable from that generated by U^* interacting with S . By induction, we may assume that with overwhelming probability U^* has only a negligible advantage of predicting $E_K(w, j)$ for any $w \in B$ and $j > f_w$. This holds in the base case since U^* doesn't know K . Again by induction, we may assume that U^* can only generate $E_K(\text{"reject"}, j)$ for j less than or equal to the total number of rejected queries thus far. Similarly, U^* can only generate $E_K(\text{"repeat"}, j)$ for j at most the total number of repeated queries thus far. We describe U' 's simulation on the $(t + 1)$ st query, and show the inductive steps for the above.

By definition of the simulator S_1 , U' can use the code of U^* to extract the inputs π', i, j, P, Q for which U^* 's interaction with S in the SFE computation is computationally indistinguishable from the view of a user in the ideal model who just hands π', i, j, P, Q to a trusted third party. Given π', i, j and the previously requested indices I of U' to the trusted third party Charlie, U' first checks if knowledge of $x_{i,j}$ together with database values indexed by I complete an inference channel. If this is the case, U' checks if P is a valid encryption of the form $E_K(\text{"reject"}, z)$ for an integer z . If so, U' decrypts P , increments z by 1, and gives $E_K(\text{"reject"}, z + 1)$ to U^* as output. Otherwise, U' gives $E_K(\text{"reject"}, 1)$ to U^* as output. U' then checks if (i, j) already appears in I . If so, U' checks if $E_K((i, j), 1)$ appears in π' . If not, U' outputs $E_K(\text{"reject"}, z + 1)$ or $E_K(\text{"reject"}, 1)$ depending on whether or not P is a valid encryption. If $E_K((i, j), 1)$ does appear in π' and Q is a valid encryption of the form $E_K(\text{"repeat"}, z)$ for some integer z , U' outputs $E_K(\text{"repeat"}, z + 1)$ and $x_{i,j} \oplus h(s, i, j)$ to U^* , otherwise if $E_K((i, j), 1)$ appears in π' but Q is not a valid encryption, U' outputs $E_K(\text{"repeat"}, 1)$ and $x_{i,j} \oplus h(s, i, j)$ to U^* .

Finally, if (i, j) is neither inference-enabling nor already in I , U' checks if π' has the three properties of π given in the protocol description. Here, again, we replace the second property " $f_\alpha = t$ " with the property " $f_\alpha = t - z_1 - z_2$ ", where $P = E_K(\text{"reject"}, z_1)$ and $Q = E_K(\text{"repeat"}, z_2)$. If π' has these three properties, U' updates I to include (i, j) , asks the trusted third party Charlie for $x_{i,j}$, and gives $x_{i,j} \oplus h(s, i, j)$ as output to U^* , as well as the updated list $g(\pi')$ defined in the protocol description. Finally, U' then runs the simulator S_2 for SPIR to extract the requested indices (i', j') from S and gives $h(s, i', j')$ to U^* as output of the SPIR interaction. We give the user $x_{i,j} \oplus h(s, i, j)$ as the output of the circuit simulation instead of $h(s, i, j)$ so that the server doesn't have to talk to Charlie when simulating the SPIR interaction, and hence, I , as defined, is exactly the set of database values the server has already queried from Charlie.

We claim the view generated by U' is indistinguishable from that of U^* in an actual execution. By the inductive hypothesis, it holds that π' cannot satisfy the three properties of π given in the protocol description if (i, j) is inference-enabling, and so the encryptions obtained by U^* are the same as those obtained in a real

execution. Note that since h is pseudorandom, $x_{i,j} \oplus h(s, i, j)$ is indistinguishable from $h(s, i, j)$, and so both the output of the SFE and the SPIR are indistinguishable from that of a real execution.

It remains to show the inductive steps for the claims in the second paragraph. This follows from two observations. First, by construction, $E_K(\text{"reject"}, z)$ and $E_K(\text{"repeat"}, z)$ are updated for U^* precisely when a query is rejected or repeated, respectively, and $g(\pi')$ is given to U^* when in fact f_w increases by one for each w in $\text{sibanc}(i, j)$ for query (i, j) . The second observation is that E is non-malleable, so U^* cannot generate any other encryptions from these. This completes the proof.

8 Extensions and Open Problems

We've described our protocols in the context of a single user, however they naturally support many users. In the multi-user setting the database server only needs the ability to link queries made by the same user and so it's possible for users to maintain some anonymity through the use of pseudonyms. In addition, collusion resistance can easily be added to any of our protocols. To ensure c -collusion resistance for an m -channel, meaning no c users can collude to complete the channel, we modify the protocols to only allow a user to query a record $\frac{(m-1)}{c}$ times. With this modification, c users can together query at most $m - 1$ attributes in an inference channel. Although, this may restrict information access quite a bit if the database isn't queried much, it does ensure collusion resistance in a fair manner.

As mentioned in Section 3, our protocols can be modified to allow for repeat queries. Note though, that this is already directly handled by the protocol in Section 7. This is achieved by distributing a "master key", $K_{i,j}$, with each $x_{i,j}$ and requiring the server to maintain a table containing each $x_{i,j}$ encrypted under $K_{i,j}$. A user can engage with the server in SPIR protocols on this table at will. If the user wants to protect against the server learning that a repeat query is made, they can perform the repeat table lookup on every query, with at most a constant factor slowdown in efficiency, together with a query on the actual database.

Finally, as this is a new primitive there are a number of open questions. Perhaps there are applications of PIC, similar to those in [3], that lie outside of the realm of protecting sensitive subsets of information. Another question is whether our protocols can be modified to enable users to update their state more independently, e.g., without going through a secure function evaluation. Although our protocols are already one-round and communication-efficient, there still may be other PIC implementations where a user can directly provide himself with the capability to access the information they are authorized to receive, effectively bypassing the secure function evaluation.

Acknowledgements

The authors thank Benny Pinkas for helpful discussions and Diana Smetters for suggesting this problem. Much of this work was supported by DARPA contract F30602-03-C-0037. Towards the end of this work, the first author was supported by an NDSEG fellowship.

References

- [1] R. Agrawal, A. Evfimievski and R. Srikant. *Information sharing across private databases*. In SIGMOD 2003.
- [2] R. Agrawal and R. Srikant. *Privacy-preserving data mining*. In ACM SIGMOD Conference on Management of Data, pp. 439-450, 2000.
- [3] W. Aiello, Y. Ishai and O. Reingold. *Prived oblivious transfer: how to sell digital goods*. In Advances in Cryptology - Eurocrypt '01.
- [4] D. Asonov. *Private Information Retrieval. An overview and current trends*. In proceedings of ECDPvA Workshop, Informatik 2001, Vienna, Austria, September 2001.

- [5] L. Beck. *A security mechanism for statistical databases*. In ACM TODS, pp. 316-338, 1980.
- [6] A. Beimel, Y. Ishai, E. Kushilevitz and T. Malkin. *One-way functions are essential for single-server private information retrieval*. In proceedings of STOC '99.
- [7] A. Beimel, Y. Ishai and T. Malkin. *Reducing the Servers Computation in Private Information Retrieval: PIR with Preprocessing*. In Advances in Cryptology – Crypto 2000.
- [8] M. Bellare, O. Goldreich. *On Defining Proofs of Knowledge*. In Advances in Cryptology – Crypto 1992.
- [9] M. Blum. *Three applications of the oblivious transfer: Part i: Coin flipping by telephone; part ii: How to exchange secrets; part iii: How to send electronic mail*. Department of EECS, University of California, Berkeley, CA, 1981.
- [10] D. Boneh. *The decision Diffie-Hellman problem*. In Proceedings of the Third Algorithmic Number Theory Symposium, Lecture Notes in Computer Science, Vol. 1423, Springer-Verlag, pp. 48–63, 1998.
- [11] D. Boneh, G. Di Crescenzo, R. Ostrovsky and G. Persiano. *Searchable public key encryption*. To appear in Advances in Cryptology – Eurocrypt 2004. <http://eprint.iacr.org/2003/195/>
- [12] C. Cachin, J. Camenisch, J. Kilian and J. Müller. *One-round secure computation and secure autonomous mobile agents*. In Ugo Montanari, Jos P. Rolim, and Emo Welzl, editors, Proc. 27th International Colloquium on Automata, Languages and Programming (ICALP), volume 1853 of Lecture Notes in Computer Science, pages 512-523. Springer, 2000.
- [13] C. Cachin, S. Micali and M. Stadler. *Computationally private information retrieval with polylogarithmic communication*. In Advances in Cryptology – Eurocrypt '99.
- [14] S. Castano, M. Fugini, G. Martella and P. Samarati. *Database Security*. ACM Press, 1995.
- [15] B. Chor, N. Gilboa and M. Naor, *Private Information Retrieval by Keywords*, TR CS0917, Department of Computer Science, Technion, 1997.
- [16] B. Chor, O. Goldreich, E. Kushilevitz and M. Sudan. *Private information retrieval*. In proceedings of FOCS '95.
- [17] G. Di Crescenzo, T. Malkin and R. Ostrovsky. *Single database private information retrieval implies oblivious transfer*. In Advances in Cryptology – Eurocrypt '00.
- [18] I. Dinur and K. Nissim. *Revealing information while preserving privacy*. In proceedings of PODS 2003.
- [19] D. Dolev, C. Dwork, and M. Naor. *Non-malleable cryptography*, SIAM. J. Computing, Vol. 30, No. 2, 2000, pp. 391–437.
- [20] S. Even, O. Goldreich and A. Lempel. *A randomized protocol for signing contracts*. In Communications of the ACM, 1985.
- [21] C. Farkas and S. Jajodia. *The inference problem: a survey*. In SIGKDD Explorations, December 2002.
- [22] E. Goh. *How to search efficiently on encrypted compressed data*. In the Cryptology ePrint Archive, Report 2003/216, October 7, 2003. <http://eprint.iacr.org/2003/216/>
- [23] Y. Gertner, Y. Ishai, E. Kushilevitz and T. Malkin. *Protecting data privacy in private information retrieval schemes*. In proceedings of STOC '98.
- [24] G. Di Crescenzo, T. Malkin and R. Ostrovsky. *Single database private information retrieval implies oblivious transfer*. In Advances in Cryptology – Eurocrypt 2000.

- [25] O. Goldreich. *Foundations of Cryptography Teaching Notes*. Available at: [url: www.wisdom.weizmann.ac.il/ oded/PSBookFrag/teach.ps](http://www.wisdom.weizmann.ac.il/oded/PSBookFrag/teach.ps)
- [26] O. Goldreich. *Secure Multi-Party Computation*, 1998. Available at <http://philby.ucsd.edu/>
- [27] O. Goldreich. *Zero-knowledge twenty years after its invention*. Electronic Colloquium on Computational Complexity, 2002.
- [28] S. Goldwasser and S. Micali. *Probabilistic encryption*. JCSS, pp.270-299, 1984.
- [29] S. Goldwasser, S. Micali and C. Rackoff. *The knowledge complexity of interactive proof systems*. SIAM J. Comp., **18**(1), pp.186-208, 1989.
- [30] O. Goldreich, S. Micali, and A. Wigderson. *How to Play Any Mental Game*. In proceedings of 19th STOC, pp. 218-229, 1987.
- [31] O. Goldreich, S. Micali and A. Wigderson. *Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems*. J of the ACM, **38**(1), pp.691-729, 1991.
- [32] T. Keefe, M. Thuraisingham, W. Tsai. *Secure query-processing strategies*. IEEE Computer, Vol. 22, No. 3, 1989.
- [33] R. Kumar, S. Rajagopalan, A. Sahai. *Coding constructions for blacklisting problems without computational assumptions*. In Advances in Cryptology – Crypto ‘99.
- [34] E. Kushilevitz and R. Ostrovsky. *Replication is not needed: single database, computationally-private information retrieval*. In Proceedings of FOCS ‘97.
- [35] E. Kushilevitz and R. Ostrovsky. *One-way trapdoor permutations are sufficient for non-trivial single-server private information retrieval*. In Advances in Cryptology – Eurocrypt ‘00.
- [36] H. Lipmaa. *An Oblivious Transfer Protocol with Log-Squared Total Communication*, <http://eprint.iacr.org/2004/063/>
- [37] D. Naccache and J. Stern. *A new public key cryptosystem*. In Advances in Cryptology – Eurocrypt 1997, pp.27-36.
- [38] M. Naor and B. Pinkas. *Efficient oblivious transfer*. In proceedings of SODA 2001.
- [39] M. Naor and B. Pinkas. *Oblivious transfer and polynomial evaluation*. In proceedings of STOC 1999.
- [40] M. Naor and B. Pinkas. *Oblivious transfer with adaptive queries*. In Advances in Cryptology - Crypto ‘99, pp. 573-590.
- [41] P. Paillier. *Public-key cryptosystems based on composite degree residuosity classes*. In Advances in Cryptology – Eurocrypt 1999, pp. 223-238.
- [42] M. Rabin. *How to exchange secrets by oblivious transfer*. Tech report TR 81, Aiken Computation Lab, 1981.
- [43] X. Qian, M. Stickel, P. Karp, T. Lunt, and T. Garvey. *Detection and elimination of inference channels in multilevel relational database systems*. In Proc. of the IEEE Symposium on Research in Security and Privacy, pp. 196–205, 1993.
- [44] P. Samarati and L. Sweeney. *Generalizing Data to Provide Anonymity when Disclosing Information*. In PODS 1998.
- [45] A. Shamir. *How to Share a Secret*. In Communications of the ACM, 22(11):612- 613, November 1979.

- [46] D. Song, D. Wagner and A. Perrig. *Practical Techniques for Searches on Encrypted Data*. In Proc. of the 2000 IEEE Security and Privacy Symposium, May 2000.
- [47] J. Staddon. *Dynamic Inference Control*. 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery, 2003.
- [48] J. Traub, Y. Yemini and H Wozniakowski. *The statistical security of a statistical database*. In ACM TODS, pp.672-679, 1984.
- [49] A. Yamamura and T. Saito. *Private information retrieval based on the subgroup membership problem*. In ACISP 2001, pp. 206-220.
- [50] A. C. Yao. *Protocols for secure computations*. In proc. of 23rd FOCS, 1982, pp. 160-164. 16.

A Tools

We rely on the following primitives in our protocols. In the following we describe each briefly in turn, referring the reader to appropriate references for more details.

A.1 Computational Symmetrically-Private Information Retrieval on Records

Symmetrically-private information retrieval (SPIR) is introduced in [23]. With each invocation of a SPIR protocol a user learns exactly one bit of a database while giving the database no information about which bit was learned. Hence, SPIR protocols add data privacy to private information retrieval (PIR) protocols such as [16], that only ensure user privacy. We rely on single-server SPIR schemes in our protocols. Such schemes necessarily offer computational, rather than unconditional, security. In [39], for an n bit database it's shown how to transform any PIR protocol into a SPIR protocol with essentially the same communication cost; more precisely, $\log n$ invocations of a 1-out-of-2 Oblivious Transfer¹³ protocol [20], which itself has constant communication cost. We require an additional property of the SPIR protocol in our constructions, namely that it is secure in the real/ideal model, that is, to guarantee the user does not learn more than a single index i of the database x , we require that there is an efficient simulator which can extract the index i learned by any (possible malicious) user U^* interacting with the honest server. Applying the transformation of [39] to the PIR schemes of [13, 36] gives such a construction with $O(kn)$ server work and $O(\text{polylog}(n))$ communication, where k is a security parameter. k can be set to $O(\text{polylog}(n))$ for security against polynomial size circuits assuming the Φ -hiding assumption [13].

Another issue is that in all of our schemes, we actually perform SPIR on *records* of size k rather than on a database $x \in \{0, 1\}^n$. It is, however, a simple matter to convert a binary SPIR scheme into a SPIR scheme on records by running k invocations of the binary scheme in parallel. This gives us a 1-round, $O(k \text{polylog}(n))$ communication per query, $O(k^2 n)$ server work simulatable SPIR protocol on records of size k . The dependence on k can be improved using techniques of [15], but since k is polylogarithmic in n in our applications, for simplicity we don't optimize this dependence in our protocols.

A.2 Homomorphic Encryption

An encryption scheme, $E : (G_1, +) \rightarrow (G_2, \cdot)$ is homomorphic if for all $a, b \in G_1$, $E(a + b) = E(a) \cdot E(b)$. For more background on this primitive see, for example, [28, 37].

We make use of the Paillier homomorphic encryption scheme [41] in one of our protocols and so we briefly repeat it here:

1. **Initialize:** Choose two primes, p and q and set $N = p \cdot q$. Let $\lambda = \text{lcm}(p - 1, q - 1)$. Let the public key be (N, g) where the order of $g \bmod N^2$ is a multiple of N , and let the secret key be λ .
2. **Encrypt:** Given a message $M \in \mathbb{Z}_N$, choose a random value $x \in \mathbb{Z}_N^*$. The encryption of M is, $E(M) = g^M x^N \bmod N^2$.
3. **Decrypt:** Let $L(u) = \frac{u-1}{N}$, where u is congruent to 1 modulo N . To recover M from $E(M)$ calculate, $\frac{L(E(M)^\lambda \bmod N^2)}{L(g^\lambda \bmod N^2)} \bmod N$.

In [41] it's shown that this encryption scheme's semantic security is equivalent to the Decisional Composite Residuosity Assumption. To see that the scheme is homomorphic, note that $E(M_1) \cdot E(M_2) = (g^{M_1} x_1^N \bmod N^2) \cdot (g^{M_2} x_2^N \bmod N^2) = g^{M_1+M_2} (x_1 x_2)^N \bmod N^2 = E(M_1 + M_2)$.

¹³SPIR is essentially equivalent to the older notion of 1-out-of- n oblivious transfer [42]

A.3 Zero-Knowledge Proofs

Zero-knowledge proofs (ZKPs) were introduced in [29] and it's shown in [31] that all languages in NP have zero-knowledge proofs provided one-way functions exist. For a recent survey on zero-knowledge proofs see [27].

An interactive proof system for the language, L , enables a computationally unbounded prover, P , to convince a polynomially bounded verifier, V , that an element of L is in fact in the language. Conversely, for any element that's not in L , P only succeeds in convincing V of membership with negligible probability. P and V may be modelled as probabilistic Turing machines, with their interaction captured by a *transcript*. We say an interactive proof system is *zero-knowledge* if the verifier learns nothing additional (i.e. beyond membership in L or lack of membership in L) from the prover. This means that any (possibly malicious) verifier V' can efficiently generate a transcript of the interaction that is indistinguishable from that produced by P . That is, there exists a probabilistic polynomial time machine M , called the *simulator* of the interaction of V with P , such that for all $x \in L$, $\{ \langle P, V \rangle (x) \}_{x \in L}$ is indistinguishable from $\{ M(x) \}$. For our purposes, computational indistinguishability suffices.

Intuitively, a *zero-knowledge proof of knowledge* allows a prover to convince a verifier of some fact in zero-knowledge if and only if the prover “knows something”. This is formally captured by the existence of an efficient algorithm known as a knowledge extractor - see [8] for definitions. We use this primitive in Protocol 2.

A.4 Non-malleable Encryption

This concept was formally defined and introduced in [19]. Informally, an encryption scheme is non-malleable, if in addition to being semantically secure, it is impossible to generate a different ciphertext from a given ciphertext so that the two respective plaintexts are related. See [19] for a more formal description and implementations.

B Reduction

In this appendix we prove the theorem of Section 4 and provide a small example of the database that's constructed in the offline stage. \mathbb{Z} denotes the integers and \mathbb{N} the positive integers.

Theorem 1 *Let (P', Q', D', R') be a TPIC with communication $C'(n, m, k)$, server work $W'(n, m, k)$, and round complexity $R'(n, m, k)$. Then there exists a PIC scheme (P, Q, D, R) with communication complexity $C(n, m, k) = m2^m C'((m+1)2^m n, m, k)$, server work $W(n, m, k) = m2^m W'((m+1)2^m n, m, k)$, and round complexity $R(n, m, k) = m2^m R'((m+1)2^m n, m, k)$.*

Proof: Much of the notation in this proof is defined in Sections 2 and 3. We use “primed” notation for the original TPIC scheme and unprimed notation for our PIC scheme. Thus, $x, \gamma, \rho, \mathcal{C}, n, m, k$ denote the inputs of the PIC scheme, P, Q, D , and R denote its components, and the intermediate variables are $M_{w,l}, A_{w,l}, m_{w,l}, a_{w,l}, out_l, \sigma_U(\cdot, \cdot), \sigma_S(\cdot, \cdot, \cdot), y$, and r . We construct inputs $x', \gamma', \rho', \mathcal{C}', n', m', k'$ to feed into the TPIC scheme, for which P', Q', D' , and R' are the components and the intermediate variables are $M'_{w,l}, A'_{w,l}, m'_{w,l}, a'_{w,l}, out'_l, \sigma'_U(\cdot, \cdot), \sigma'_S(\cdot, \cdot, \cdot), y'$, and r' .

Offline Stage: Fix an honest user U . The user and server are given $n, m, 1^k$, and a monotone $\mathcal{C} \in 2^{[m]}$. Define $m' = m, \mathcal{C}' = \{[m]\}, n' = (m+1)|\mathcal{C}|n$, and $k' = k$. Let $\text{poly}(n, m, k)$ be a polynomial bounding the length of ρ' and γ' in the TPIC scheme. In the PIC scheme the user is given ρ with $|\rho| = \text{poly}(m(m+1)2^m n, m, k)$ and the server a γ with $|\gamma| = \text{poly}(m(m+1)2^m n, m, k) + (m+1)m2^m n$. The server is also given an $x \in (\{0, 1\}^m)^n$. Set $\rho' = \rho$, and let γ' be the restriction of γ to its first $\text{poly}(m(m+1)2^m n, m, k)$ coordinates. Define $r = R(n, m, k)$ to be $m|\mathcal{C}|r' = m|\mathcal{C}|R(n', m', k')$.

The server constructs a database $x' \in (\{0, 1\}^m)^{n'}$, which can be thought of as a string in $(\{0, 1\}^{|\mathcal{C}|})^{(m+1)n}$. For each $i \in [n], j \in [|\mathcal{C}|]$, and $k \in [|\mathcal{C}_j|]$, let $s_{i,j,k}$ be a random bit subject to the constraints:

$$\forall \alpha \in [m], \quad \bigoplus_{j,k \text{ s.t. } \mathcal{C}_{j,k}=\alpha} s_{i,j,k} = x_{i,\alpha}.$$

For each $i \in [n], j \in [|\mathcal{C}|]$, and $|\mathcal{C}_j| < k \leq m$, let $r_{i,j,k}$ be a random bit. For each $i \in [n]$ and $j \in [|\mathcal{C}|]$,

- for $k = 1$, define $x'_{i,j,k} = s_{i,j,k} \oplus \bigoplus_{v=|\mathcal{C}_j|+1}^m r_{i,j,v}$,
- for $2 \leq k \leq |\mathcal{C}_j|$, define $x'_{i,j,k} = s_{i,j,k}$,
- and for $|\mathcal{C}_j| < k \leq m$, define $x'_{i,j,k} = r_{i,j,k}$.

The $x'_{i,j,k}$ for $|\mathcal{C}_j| < k \leq m$ are the “random pad” values. Finally, for each $i \in \{n+1, \dots, (m+1)n\}$, each $j \in [|\mathcal{C}|]$, and each $k \in [m]$, set $x'_{i,j,k}$ to be a random bit. These are “dummy” bits in the database that the user will query so that the number of rounds per online execution is the same for every execution, which we'll need when proving user privacy. Observe that there are enough random bits in γ so that every random bit needed in the construction of x' is in fact a truly random bit. This follows from the fact that there are at most $(m+1)m2^m n$ random bits needed, one for each position in x' .

A small example of a database, x' , that is output at this stage is given in Figure 1.

P: Let $y' = P'(x', \gamma')$, where P' is also given inputs $n', m', 1^{k'}, \mathcal{C}'$, and γ' . Define $y = P(x, \gamma) = x' \circ y'$.

Online Stage: We consider the online stage for a given permissible $T = ((i_1, k_1), \dots, (i_{|T|}, k_{|T|})) = T(U, x)$. For $1 \leq l \leq |T|$, let $E_l = \{j \mid k_l = \mathcal{C}_{j,1}\}$ be the set of channels whose first element is k_l , and more generally, let $F_l = \{j \mid k_l \in \mathcal{C}_j\}$ be the set of channels containing k_l . Define the sets:

$$G_l = \{((i_l - 1)|\mathcal{C}| + j, k) \mid j \in F_l, \mathcal{C}_{j,k} = k_l\}, \quad H_l = \{((i_l - 1)|\mathcal{C}| + j, k) \mid j \in E_l, |\mathcal{C}_j| < k \leq m\}.$$

Figure 1: For a database, x , with $n = 1$, $m = 4$, and $\mathcal{C} = ((1, 2), (1, 2, 3), (1, 2, 4), (1, 2, 3, 4))$, the output, x' (depicted as a $|\mathcal{C}|(m + 1)n = 20 \times 4 = m$, binary matrix), of the offline stage is shown here. Note that $s_{1,4,1} \oplus s_{1,3,1} \oplus s_{1,2,1} \oplus s_{1,1,1} = x_{1,1}$, $s_{1,4,2} \oplus s_{1,3,2} \oplus s_{1,2,2} \oplus s_{1,1,2} = x_{1,2}$, $s_{1,4,3} \oplus s_{1,2,3} = x_{1,3}$ and $s_{1,4,4} \oplus s_{1,3,3} = x_{1,4}$.

G_l denotes the $s_{i,j,k}$ values queried in the l th execution, and H_l denotes the random pad values queried in the l th execution. Order the pairs in G_l and H_l arbitrarily (so G_l and H_l become ordered tuples), and define $I_l = G_l \circ H_l$. I_l is the list of pairs the user needs to query in the l th execution, excluding dummy pairs. Let $s(l) = (l - 1)r' - \sum_{i=1}^{l-1} |I_i|$ and let $t(l) = lr' - \sum_{i=1}^l |I_i|$. Express $s(l) + |\mathcal{C}|n$ as $s_1(l)m + s_2(l)$ and $t(l) + |\mathcal{C}|n$ as $t_1(l)m + t_2(l)$ for integers $s_1(l), t_1(l), s_2(l), t_2(l)$ with $0 \leq s_2(l), t_2(l) < m$. The following is the list of dummy pairs the user queries so that the number of rounds per execution is the same in every execution:

$$D_l = ((i, j) \in \mathbb{Z} \times [m - 1] \mid ((s_1(l), s_2(l)) + (0, 1)) \leq (i, j) \leq (t_1(l), t_2(l))),$$

where for two pairs $(x_1, y_1), (x_2, y_2) \in \mathbb{Z} \times [c]$, $(x_1, y_1) \leq (x_2, y_2)$ if and only if either $x_1 < x_2$, or $x_1 = x_2$ and $y_1 \leq y_2$. Also, for a pair $(x, y) \in \mathbb{Z} \times [c]$, $(x, y) + (0, 1)$ denotes $(x, y + 1)$ if $y < c$ and $(x + 1, 1)$ if $y = c$.

Set $J_l = I_l \circ D_l$, and observe that $|J_{l_1}| = |J_{l_2}| = \frac{r}{r'}$ for any l_1 and l_2 . J_l is the list of pairs the user actually queries in the l th execution. Define $T' = J_1 \circ J_2 \circ \dots \circ J_{|T|}$. The idea is that (Q, D, R) will run (Q', D', R') on each pair in T' . There will be $|T|$ executions l , $1 \leq l \leq |T|$, of the online stage, and for each l there will be r rounds w , $1 \leq w \leq r$. For a given w and l , it will be *as if* (Q', D', R') is running on T' and x' with w' and l' , where $w' = (w - 1 \bmod r') + 1$ and $l' = (l - 1)m|\mathcal{C}| + \lceil \frac{w}{r'} \rceil$. Define the map: $\Gamma(w, l) = (w', l')$.

We can define $M_{w,l}, A_{w,l}, m_{w,l}$, and $a_{w,l}$ in terms of the simulation of (Q', D', R') on T' and x' :

$$m_{w,l} = m'_{\Gamma(w,l)}, \quad a_{w,l} = a'_{\Gamma(w,l)}, \quad \text{and hence, } M_{w,l} = M'_{\Gamma(w,l)}, \quad A_{w,l} = A'_{\Gamma(w,l)}.$$

It follows that $C(n, m, k) \leq m2^m C'((m+1)|\mathcal{C}|n, m, k)$ as claimed, since for each l , $m2^m$ invocations of the underlying (Q', D', R') protocol are run. Note also that $R(n, m, k) \leq m2^m R'((m+1)2^m n, m, k)$ as claimed.

Q: For $1 \leq l \leq |T|$ and $0 \leq w \leq r$, define¹⁴:

$$\sigma_U(A_{w,l}, \rho) = \sigma'_U(A'_{\Gamma(w,l)}, \rho') \circ s(l-1) \circ t(l-1) \circ \{out'_i \mid \Gamma(0, l)_2 \leq i \leq \Gamma(w, l)_2\}.$$

Using the definition of the online stage of a PIC protocol (and of $m_{w,l}$), this completely specifies the input/output behavior of Q . Note that we keep $s(l-1)$ and $t(l-1)$ so that the user can construct $s(l), t(l)$, and hence D_l . We keep out'_i s around so that the user will be able to compute out_l , defined below. Note that some of the lists defined above ($D_l, E_l, F_l, G_l, H_l, I_l, J_l$) could be included in the user's state to improve efficiency, but we don't attempt to optimize user-side work.

D: For $1 \leq w \leq r$, define $\sigma_S(M_{w,l}, x, \gamma) = \sigma'_S(M'_{\Gamma(w,l)}, x', \gamma')$. Using the definition of the online stage (and of $a_{w,l}$), this completely specifies D . Observe that $W(n, m, k) \leq m|\mathcal{C}|W'((m+1)|\mathcal{C}|n, m, k)$ as claimed, since we have $m|\mathcal{C}|r'$ executions of D' per query. This follows from the fact that D just simulates D' on the appropriate inputs, so the running times of D' and D are the same.

R: For all $1 \leq l \leq |T|$ define:

$$out_l = \bigoplus_{i=\Gamma(1,l)_2}^{\Gamma(1,l)_2+|I_l|-1} out'_i$$

From the definition of the online stage, this completely specifies R .

It remains to show that (P, Q, D, R) is a PIC scheme:

Correctness: Consider any honest user U and the honest server S . We need to prove that for all x and all $l \in [|T = T(U, x)|]$, $out_l = x_{T_l}$.

First, observe that we have carefully constructed T' so that if T has distinct coordinates, then so does T' . Indeed, consider any J_{l_1}, J_{l_2} for $l_1 \neq l_2$. We claim $J_{l_1} \cap J_{l_2} = \emptyset$. First note that the number of dummy values queried for a given l is less than $|\mathcal{C}|m$, and since T has distinct coordinates and \mathcal{C} is nonempty, $|T| \leq (m-1)n$. Hence, there are at most $m(m-1)n|\mathcal{C}|$ dummy queries over all l . But x' has $mn|\mathcal{C}|$ records containing dummy values, and since each such record can be queried $m-1$ times, there are $m(m-1)n|\mathcal{C}|$ dummy values that can be queried without completing an inference channel, as needed. Hence, $D_{l_1} \cap D_{l_2} = \emptyset$, so we only need to show $I_{l_1} \cap I_{l_2} = \emptyset$. But this is immediate from the definitions of E_l, F_l, G_l , and H_l .

We also claim that if for all $i \in [n]$ and all $k \in [|\mathcal{C}|]$, $\mathcal{C}_k \not\subseteq \{j \mid \exists a \text{ s.t. } T_a = (i, j)\}$, then for all $i \in [n']$, $[m] \neq \{j \mid \exists a \text{ s.t. } T'_a = (i, j)\}$. To see this, first note that when $i \in [n'] \setminus [|\mathcal{C}|n]$, this is always true, since by construction, queries to dummy values in the database never complete a channel. Then simply note that if there is some $i \in [|\mathcal{C}|n]$ for which the set $\{j \mid \exists a \text{ s.t. } T'_a = (i, j)\} = [m]$, then every attribute of the $\lceil \frac{i}{|\mathcal{C}|} \rceil$ th record in x in the $((i \bmod n) + 1)$ st channel was queried, which means that the set $\mathcal{C}_{(i \bmod n)+1} \subseteq \{j \mid \exists a \text{ s.t. } T_a = (\lceil \frac{i}{|\mathcal{C}|} \rceil, j)\}$, which shows the claim.

¹⁴Here we take $\Gamma(0, l) = (r', (l-1)m|\mathcal{C}|)$, which is what one obtains by extending the definitions of w' and l' to $w = 0$. Also, we take $s(0) = t(0) = 0$.

Correctness of (P, Q, D, R) thus follows from correctness of (P', Q', D', R') :

$$\begin{aligned}
out_l &= \bigoplus_{i=\Gamma(1,l)_2}^{\Gamma(1,l)_2+|I_l|-1} out'_i = \bigoplus_{s \in G_l} x'_s \oplus \bigoplus_{s \in H_l} x'_s \\
&= \bigoplus_{j \in E_l} \left(s_{i_l, j, 1} + \bigoplus_{k=|\mathcal{C}_j|+1}^m r_{i_l, j, k} \right) \oplus \bigoplus_{j \in F_l \setminus E_l, k \text{ s.t. } \mathcal{C}_{j,k}=k_l} s_{i_l, j, k} \oplus \bigoplus_{j \in E_l} \bigoplus_{k=|\mathcal{C}_j|+1}^m r_{i_l, j, k} \\
&= \bigoplus_{j, k \text{ s.t. } \mathcal{C}_{j,k}=k_l} s_{i_l, j, k} = x_{i_l, k_l}.
\end{aligned}$$

User Privacy: Suppose (P, Q, D, R) does not have user privacy, that is, suppose there is an efficient algorithm A , an honest user U , an honest but curious server S_1^* , an integer m , a monotone $\mathcal{C} \in 2^{[m]}$, a polynomial q , and an infinite sequence of integers $\bar{n} = \{n(\alpha)\}_{\alpha \in \mathbb{N}}$, such that for every $\alpha \in \mathbb{N}$, there are two sequences $\tau(\alpha)_1 \neq \tau(\alpha)_2$ with $|\tau(\alpha)_1| = |\tau(\alpha)_2| < mn(\alpha)$, a random input $\gamma(\alpha)$, and a database $x(\alpha) \in (\{0, 1\}^m)^{n(\alpha)}$ such that:

$$|\Pr_{\rho} [A(V_{S_1^*}(U, x(\alpha), \rho, \gamma(\alpha))) = 1 \mid T = \tau_1(\alpha)] - \Pr_{\rho} [A(V_{S_1^*}(U, x(\alpha), \rho, \gamma(\alpha))) = 1 \mid T = \tau_2(\alpha)]| \geq \frac{1}{q(n(\alpha), k(n(\alpha)))},$$

where A is additionally given inputs $n(\alpha), m, 1^{k(n(\alpha))}$, and \mathcal{C} .

The claim is that this immediately gives a distinguisher A' for (P', Q', D', R') for the infinite sequence $\{n'(\alpha) = (m+1)2^m n(\alpha)\}_{\alpha \in \mathbb{N}}$, sequences $\tau(\alpha)'_1, \tau(\alpha)'_2$, where $\tau(\alpha)'_1, \tau(\alpha)'_2$ are the sequences fed into the simulation of (P', Q', D', R') in the (P, Q, D, R) scheme given $\tau(\alpha)_1, \tau(\alpha)_2$, respectively, databases $x(\alpha)'$ used in the simulation given $x(\alpha), \mathcal{C}' = \{[m]\}$, the randomness $\gamma'(\alpha)$ used in the simulation given $\gamma(\alpha)$, and the same m, p , and q . Note that since we forced the number of rounds to be the same in every execution of the online stage of (P, Q, D, R) , $\tau(\alpha)'_1$ and $\tau(\alpha)'_2$ have the same length for all α . Since $\tau(\alpha)_1 \neq \tau(\alpha)_2$, we have $\tau(\alpha)'_1 \neq \tau(\alpha)'_2$ by construction.

We consider the view of server S_2^* who behaves as follows. First, we hardwire m and \mathcal{C} into S_2^* . On inputs $x'(\alpha)$ and $\gamma'(\alpha)$, S_2^* simply behaves as S_1^* would in the (P, Q, D, R) scheme. First, S_2^* runs D instead of D' . Because $\sigma_S(M_{w,l}, x(\alpha), \gamma(\alpha)) = \sigma'_S(M'_{\Gamma(w,l)}, x'(\alpha), \gamma'(\alpha))$ for the honest servers S, S' , because $a_{w,l} = a_{\Gamma(w,l)}$ for all α , and because S_2^* can compute $y = x'(\alpha) \circ \gamma'(\alpha)$, S_2^* can easily simulate S_1^* 's input/output behavior. Actually, S_2^* can simulate *every* efficient algorithm S_1^* runs. This is because, given $x'(\alpha)$, S_2^* can use \mathcal{C} to recover $x(\alpha)$ and then recover $\gamma(\alpha)$. Moreover, since for all w and l we have that $m_{w,l}$ is the same as $m'_{\Gamma(w,l)}$ for the honest users in the (P, Q, D, R) and (P', Q', D', R') schemes, it follows that the views of S_2^* and S_1^* are identically distributed. Hence A' simply runs A on the views of S_2^* on the two sequences $\tau'_1(\alpha)$ and $\tau'_2(\alpha)$ together with $x(\alpha)$ and $\gamma(\alpha)$, and by the discussion above, if A has a non-negligible advantage of distinguishing $\tau_1(\alpha)$ and $\tau_2(\alpha)$, then A' has a non-negligible advantage of distinguishing $\tau'_1(\alpha)$ and $\tau'_2(\alpha)$. Note that $n'(\alpha) = \Theta(n(\alpha))$, $k'(n'(\alpha)) = k(n(\alpha))$, and given $n'(\alpha)$, A' can compute and feed $n(\alpha)$ into A . This contradicts the user privacy of (P', Q', D', R') .

Inference Control: Let S denote the honest server in the (P, Q, D, R) scheme and S' the honest server in the (P', Q', D', R') scheme. Let γ^* be γ minus the prefix γ' . To show that (P, Q, D, R) has inference control, the crucial observation is that for every x , every user U_1^* and every user randomness ρ used in the (P, Q, D, R) scheme, there exists a user U_2^* and a user randomness ρ' such that $V_{U_1^*}(x, \rho, \gamma)$ is identically distributed to $V_{U_2^*}(x'(\gamma^*), \rho', \gamma')$, where the probability space is induced by the uniform distribution on γ , and $x'(\gamma^*)$ is the database which is fixed given γ^* . The x', ρ' are just the values defined in the reduction of (P, Q, D, R) to (P', Q', D', R') , and the fact that U_2^* exists follows from the fact that S just simulates S' .

Since (P', Q', D', R') has inference control, there exists a U'_2 interacting with the trusted third party Charlie such that for every γ^* , U'_2 can efficiently output a permissible random variable¹⁵ $T_*(\gamma^*) = T_*(\gamma^*)(U'_2, x'(\gamma^*))$ for which $U'_2(\{x'(\gamma^*)_{i,j} \mid (i,j) \in T_*(\gamma^*)\})$ is computationally indistinguishable from $V_{U'_2}(x'(\gamma^*), \rho', \gamma')$. We show there is a U'_1 interacting with Charlie which can efficiently find a permissible T (again, T is a random variable) for which $U'_1(\{x_{i,j} \mid (i,j) \in T\})$ is computationally indistinguishable from $V_{U'_1}(x, \rho, \gamma)$. We show this by constructing a U'_1 for which $U'_1(\{x_{i,j} \mid (i,j) \in T_*\})$ and $U'_2(\{x'_{i,j}(\gamma^*) \mid (i,j) \in T'_*(\gamma^*)\})$ for random γ are identically distributed.

U'_1 simply simulates U'_2 . There are three types of values U'_2 may ask Charlie for: (1) a value $x'_{i,j,k}$ of the form $s_{i,j,k} \oplus \bigoplus_{k=|C_j|+1}^m r_{i,j,k}$, (2) a value of the form $s_{i,j,k}$, and (3) a value of the form $r_{i,j,k}$. U'_1 feeds random bits into the simulation of U'_2 until it is forced to obtain an $x_{i,j}$ from Charlie, at which point it requests $x_{i,j}$. Note that $x_{i,j}$ is perfectly hidden until all random pad values $r_{i,k,l}$ in channels C_k for which $C_{k,1} = j$ and $|C_k| < l \leq m$ are obtained, in addition to all of the $s_{i,k,l}$ for which $C_{k,l} = j$ are determined. Only when the last such value is requested does U'_1 actually request $x_{i,j}$ from Charlie.

It remains to show that the sequence of queries T requested by U'_1 is permissible. Since $T_*(\gamma^*)$ is permissible for any γ , for every $i \in [n]$, $j \in [|\mathcal{C}|]$, there is a $k \in [m]$ such that $x'(\gamma)_{i,j,k}$ is not in $T_*(\gamma^*)$. By construction of U'_1 , then, it is easy to see that for every $i \in [n]$ and every channel $F \in \mathcal{C}$, there is some $j \in F$ such that $x_{i,j}$ is not requested by U'_2 . This shows that (P, Q, D, R) has inference control.

This completes the proof. ■

¹⁵ $T_*(\gamma^*)$ is induced by the uniform distribution on γ' .

C Security of Stateful PIC

In this section we provide proof sketches of the security of the scheme of Section 6. Many of the details are left to the reader.

Lemma 1 *The protocol of Section 6 meets the requirements of a TPIC protocol, and hence applying theorem 1 gives a protocol meeting the requirements of a PIC protocol.*

Proof:

Correctness: Assuming the correctness of the PIR protocol, an honest user receives, $E^{hom}(\mathcal{S}_t + x_{i_t, j_t})$ on their t th query and decrypts to obtain, $\mathcal{S}_t + x_{i_t, j_t}$. If the query sequence thus far, $((i_1, j_1), \dots, (i_t, j_t))$ is permissible, and thus consists of at most $m - 1$ queries to row i_t of the database, then the user possesses the secret \mathcal{S}_t from step 2(a) and so can recover the desired bit, x_{i_t, j_t} .

User Privacy: Assuming the semantic security of E^{hom} and the zero-knowledge property of step 1, user privacy follows from the user privacy of SPIR. The proof is a standard hybrid argument and is omitted. See [26] for more background.

Inference Control: Consider an arbitrary user U^* with random tape ρ . Using the SPIR protocol of Appendix A, there exists a simulator Sim which given U^* 's code and ρ , can extract the indices (i_t, j_t) requested by U^* in step 3(b). Let U' interact with Charlie in the ideal model. U' runs the code of U^* and uses the knowledge extractor of the zero-knowledge proof of knowledge to obtain $E^{hom}(i'_t, j'_t)$ and (i'_t, j'_t) for some i'_t and j'_t . If $(i'_t, j'_t) \neq (i_t, j_t)$ (the extracted indices) or if (i_t, j_t) together with the previous values requested from Charlie complete an inference channel, in step 3(a) U' will provide the encryption $E^{hom}(r)$ for r a random value, otherwise, U' will ask Charlie for x_{i_t, j_t} and follow the protocol just as the honest server would. Note that U' carries out step 2 as the honest server would. It is not hard to see that U' can generate a computationally indistinguishable view to U^* 's, and hence a computationally indistinguishable output. ■