

TAPIR: A Two-Server Authenticated PIR Scheme with Preprocessing

Francesca Falzon
ETH Zurich
Zurich, Switzerland
ffalzon@ethz.ch

Laura Hetz
ETH Zurich
Zurich, Switzerland
lahetz@ethz.ch

Annamira O’Toole
ETH Zurich
Zurich, Switzerland
aotoole@ethz.ch

Abstract

Authenticated Private Information Retrieval (APIR) enables a client to retrieve a record from a public database and verify that the record is “authentic” without revealing any information about which record was requested. In this work, we propose TAPIR: the first two-server APIR scheme to achieve both sublinear communication and computation complexity for queries, while also supporting updates. Our scheme builds upon the unauthenticated two-server PIR scheme SinglePass (Lazaretti and Papamanthou, USENIX’24). Due to its modular design, TAPIR provides different trade-offs depending on the underlying vector commitment scheme used.

Moreover, TAPIR is the first APIR scheme with preprocessing to support appends and edits in time linear in the database partition size. This makes it an ideal candidate for transparency applications that require support for integrity, database appends, and private lookups. We provide a formal security analysis and a prototype implementation that demonstrates our scheme’s efficiency. TAPIR incurs as little as 0.11 % online bandwidth overhead for databases of size 2^{22} , compared to the unauthenticated SinglePass. For databases of size $\geq 2^{20}$, our scheme, when instantiated with Merkle trees, outperforms all prior multi-server APIR schemes with respect to online runtime.

Keywords

Private Information Retrieval, Authenticated PIR

1 Introduction

Private Information Retrieval (PIR) enables a client to retrieve a record from a public database without revealing any information about which record was requested to the server(s) holding the database. With growing concerns around digital privacy, PIR has gained a lot of attention in both academia and industry, and has been the focus of extensive research, e.g., [KC21, CHK22, HHC⁺23, LP23, LP24]. Notably, PIR has many real-world applications such as transparency systems [HHC⁺23, CNC⁺23], private web search [HDCZ23, ABG⁺24], private messaging [CSM⁺20, AS16, LTW24], and private lookup [KC21, Res24]. PIR has even been deployed at large scale by Apple to enable private caller ID [Res24].

While PIR has been touted as a solution for various use cases, there is a gap between the functionality provided by existing PIR schemes and that required by real-world systems. Practical systems demand more than just query privacy: they require low-latency, sub-linear query performance, robust integrity guarantees against adaptive adversaries, and provably secure support for dynamic databases. For example, data structures used in transparency applications must support appends *and* provide integrity (see e.g., [MBB⁺15,

MKKS⁺23, CDGM19, TBP⁺19, HHK⁺21]), and may require additional guarantees such as append-only properties. Furthermore, these systems must operate securely even in the presence of malicious servers, and where PIR is only one part of a larger system. However, existing PIR protocols do not satisfy all of these requirements simultaneously.

A particularly subtle threat in the presence of malicious adversaries are *selective failure attacks*. If a malicious server answers a client’s query arbitrarily, the client may detect the error and choose to abort. If the server can observe whether or not the client aborts, the client’s behavior may leak information about the queried index, thereby violating PIR’s core goal of query privacy. This attack becomes especially problematic when PIR is used as a building block in a larger system—a reasonable assumption for real-world PIR use cases, such as Key Transparency (KT) or secure messaging. Simply verifying the correctness of the returned record may be insufficient to ensure that a PIR scheme achieves both *integrity* and *privacy with abort* [KO97].

To address the problem of selective failure when providing integrity in PIR, Colombo et al. [CNC⁺23] recently introduced *Authenticated PIR (APIR)*. APIR enables a client to succinctly retrieve information from a public database while guaranteeing both privacy and integrity of the retrieved record. A server’s attempt to mount a selective failure attack should result in the client aborting with equal probability, regardless of the record queried. Recent works have proposed APIR schemes for both the single- [CNC⁺23, DT24, FMS24, dCL24] and multi- [CNC⁺23, WLZ⁺23] server settings. Although these schemes ensure both privacy and integrity, they leave open the question of whether such strong guarantees can be achieved alongside the sub-linear query performance and update support required by real-world applications. In particular, prior APIR work has computational server complexity linear in the database size for answering a query. In the single-server setting, this overhead is required to ensure that the server does not learn any information about the requested record. In the multi-server setting, this overhead stems from a lack of database preprocessing, a technique commonly used in state-of-the-art PIR schemes [KC21, LP23, LP24] to reduce query-time costs.

Given the gaps in the literature, we ask the following question:

“Can we design a multi-server PIR scheme that (1) achieves query computation and communication costs that are sublinear in the database size, (2) provides integrity while mitigating selective failure attacks, and (3) supports updates?”

1.1 Contributions

In this work, we answer the above question affirmatively by proposing TAPIR: the first two-server APIR protocol with preprocessing

that achieves sublinear bandwidth and computational query complexity, *and* supports efficient updates, i.e., additions and edits.

New Definitions (§2). The combination of functionalities we aim to support requires a new interface compared to prior work. Prior APIR schemes were either multi-server without preprocessing [AB25, CNC⁺23], single-server [CNC⁺23, AR25, DT24], or did not include updates in the client-preprocessing model [FMS25]. To this end, we introduce a new syntax to formalize *updateable APIR with preprocessing*, along with formal definitions of correctness, integrity, and privacy with abort. Our formalization captures both the single-server ($k = 1$) and multi-server ($k \geq 2$) settings. Our new game-based definitions for security capture a powerful adversary who can choose which server to corrupt (for $k \geq 2$), select the database, and adaptively issue lookup and update queries. For $k \geq 2$ servers, our model assumes a malicious-but-non-colluding setting, in which all but one servers may be malicious. Privacy of the retrieved records and integrity of the database are preserved as long as one server behaves honestly.

A New Scheme (§5). We present the first two-server APIR scheme with query complexities sublinear in the database size. Our construction, TAPIR, builds on SinglePass [LP24], a state-of-the-art two-server PIR protocol with efficient client preprocessing but no authentication guarantees. To provide integrity, we combine SinglePass with a vector commitment (VC) scheme, carefully ensuring protection against selective failure attacks. This protection is not trivial, as a client should abort with equal probability when receiving incorrect information regardless of the queried record. Our scheme treats the VC as a black-box primitive, only requiring that the commitment be deterministic. This design offers implementation simplicity for practitioners while ensuring direct benefits from future advances in VC constructions. Additionally, our approach offers efficiency trade-offs between runtime and bandwidth depending on the choice of commitment scheme used. We demonstrate these trade-offs experimentally in §6. For example, instantiating TAPIR with Merkle trees yields smaller server runtime than all prior multi-server APIR schemes, whereas instantiation with Pointproofs [GRWZ20] provides online bandwidth comparable to the unauthenticated SinglePass. This adaptability allows TAPIR to be tuned to the requirements of diverse applications.

Support for Updates (§5.2). We extend our scheme to support updates, making TAPIR the first APIR scheme with preprocessing that handles dynamic databases in an efficient and provably secure manner. This extension significantly enhances TAPIR’s utility for applications such as transparency logs, where records are continually appended or modified. We provide new definitions for updatable APIR with preprocessing under adaptive adversaries in a game-based model to ensure that strong security guarantees back our treatment of updates. While SinglePass [LP24] adopts a simulation-based privacy definition that only establishes security for their static scheme under non-adaptive queries, we prove that correctness, integrity, and privacy hold even against malicious adversaries who can adaptively and arbitrarily issue both lookup and update queries.

Evaluation (§6). We implement and evaluate our scheme to showcase its query efficiency compared to prior work. TAPIR outperforms both multi-server APIR schemes of [CNC⁺23] on larger databases with respect to online runtime. Moreover, our scheme

performs similarly to non-authenticated state-of-the-art schemes of the same kind in the online phase, while incurring small overhead due to authentication. Concretely, TAPIR requires as little as 0.11 % more bandwidth than SinglePass for databases of size $N = 2^{20}$ when instantiated with a state-of-the-art algebraic VC scheme such as Pointproofs [GRWZ20] and 13.11× greater runtime compared to SinglePass when instantiated with Merkle trees. This additional cost is a small price for provable, significantly stronger security guarantees. To assess TAPIR’s practicality for KT, we conducted our experiments for 32 B records—a typical key size in KT. The results show that our scheme scales well for larger databases ($N \geq 2^{24}$) while maintaining efficient runtimes and bandwidth for the online phase and for updates. To ensure efficiency for even larger databases, our scheme can be extended in a black-box manner by leveraging database partitioning and future improvements to VCs.

1.2 Prior Work

This section provides an overview of the related works on APIR and multi-server PIR.

Multi-Server PIR. “Traditional” PIR schemes with multiple servers [CGKS95, GI14, BGI15, BGI16] incur server computation linear in the database size. The client-preprocessing model [CK20] shifts much of this cost to an *offline* preprocessing phase, in which the client and server(s) jointly process the database, allowing the client to query the database more efficiently in the *online* phase. This approach has since led to two-server schemes with sublinear online complexities [KC21, LP23, LP24]. SinglePass [LP24] improves over prior client-preprocessing schemes [CK20, KC21, LP23] with an offline phase that only requires a single pass over the database, while also supporting updates.

Authenticated PIR. The aforementioned schemes only guarantee privacy and integrity against honest servers, whereas APIR schemes extend these guarantees to malicious servers and protect against selective failure. Colombo et al. [CNC⁺23] first formalized APIR and proposed both single- and multi-server schemes. Their single-server schemes require a trusted setup to generate an honest digest. This setting is also used in VeriSimplePIR [dCL24]. The scheme of Dietz and Tessaro [DT24] avoids a trusted setup, but only supports 1-bit records without a concrete cost analysis. Balanced-PIR [AR25], concurrent work to ours, proposed a single-server stateful verifiable PIR scheme that avoids the trusted setup by having the client initially download the database and verify the digest generated by the server. It achieves sublinear amortized computation with small client storage. However, while this scheme also supports updates, its security definitions only capture the privacy of look-up queries.

Multi-server schemes eliminate the need for a trusted digest as long as one server is honest. Colombo et al. [CNC⁺23] authenticate a linear PIR scheme using Merkle trees. They also propose a second two-server APIR scheme that uses a message authentication code (MAC) [DPSZ12] to verify server responses. We refer to these constructions as APIR-Matrix-MT and APIR-DPF, both of which incur online computation linear in the database size.

Crust [WLZ⁺23] extends [CK20] to provide verifiability, but its security relies on the strong assumption that the dedicated preprocessing server is honest. This assumption is much stronger

than those of related multi-server APIR schemes [CNC⁺23], which assume at least one of the servers to be honest without specifying which one. Alon and Beimel [AB25] modify the compiler of [EKN22] to generically transform a semi-honest PIR scheme into one with security-with-abort. Falk et al. [FMS25] present a generic compiler that converts a PIR scheme into a malicious-secure PIR scheme in both single- and multi-server settings.

2 Preliminaries

We denote the set $\{1, \dots, n\}$ by $[n]$ and vectors in boldface, e.g., $\mathbf{v} = (v_1, \dots, v_n)$. For some subset of indices $S \subseteq [n]$, we use $\mathbf{v}[S] = (v_i : i \in S)$ to denote the subvector indexed by S .

We consider a database DB with $N \in \mathbb{N}$ records, each of size ℓ_r , i.e., $\text{DB} \in \{0, 1\}^{N \times \ell_r}$. Our scheme divides the database into Q partitions of size M where $Q = \lceil N/M \rceil$. We denote the q -th partition by DB_q . Our scheme also builds a second database $\widehat{\text{DB}}$ that contains opening proofs, each of size ℓ_p . We let $\text{DB}_q[m]$ denote the m -th record in DB_q and $\widehat{\text{DB}}_q[m]$ denote the corresponding proof for $m \in [M]$. The databases are padded to their maximum capacity of $M \cdot Q$ with 0-records if $N < M \cdot Q$.

To allow database updates, initial fixed values are denoted by N_{init} for the database size, Q_{init} for the number of partitions, and DB_{init} for the sequence of records at the start of the protocol. During protocol execution, these values may change; hence, we use N , Q , and DB , respectively.

For clarity and convenience, we also use the dot notation from object-oriented programming to refer to a particular element of a tuple. For example, if the server state is $\text{st}_S = (\text{DB}, \widehat{\text{DB}})$, then we write $\text{st}_S.\text{DB}$ to access the database directly. We also use this notation to append values to a tuple, e.g., $\text{st}_C.\text{pp} \leftarrow \text{pp}$ denotes adding public parameters to the client's state. In text, we might refer to common values directly instead of using the above notation, e.g., N instead of $\text{pp}.N$.

The symbol \perp either denotes an empty value or that verification has been unsuccessful and that the client has aborted.

2.1 Vector Commitments

Vector commitments (VCs) enable a party to commit to a sequence of elements and later open the commitment and prove the value at a particular index. Merkle trees are one way to implement such a primitive. However, vector commitments often additionally require that the proof be succinct, i.e., that the size of the proof is independent of the length of the vector. Examples of such succinct constructions include the original work on VCs [CF13] and Pointproofs [GRWZ20].

Definition 1 ([CF13]). An **updateable vector commitment** VC is a tuple of five algorithms with the following syntax:

- $\text{pp}_{\text{VC}} \leftarrow \text{ParamGen}(1^\lambda, n)$ takes a security parameter λ and the size of vectors to be committed $n = \text{poly}(\lambda)$. It outputs public parameters pp_{VC} . We assume the public parameters to be an implicit input to the remaining algorithms.
- $V \leftarrow \text{Commit}(\mathbf{v})$ takes a vector \mathbf{v} and outputs commitment V .
- $\pi \leftarrow \text{Open}(i, \mathbf{v})$ takes as input index $i \in [n]$ and a vector \mathbf{v} and it outputs a proof π .

- $b \leftarrow \text{Verify}(V, S, \mathbf{v}[S], \pi)$ takes as input a commitment V , index subset $S \subseteq [n]$, values $\mathbf{v}[S]$, and a proof π . It outputs bit $b = 1$ if V is consistent with $\mathbf{v}[S]$ and $b = 0$ otherwise.
- $V' \leftarrow \text{Update}(V, S, \mathbf{v}[S], \mathbf{v}'[S])$ takes as input a commitment V and updates the values at the indices in S from $\mathbf{v}[S]$ to values $\mathbf{v}'[S]$. It outputs a commitment V' to the updated vector.

For security, we require that a commitment cannot be opened to different values at the same index. More formally, a vector commitment scheme is **binding** if for every n and every adversary running in time $\text{poly}(\lambda)$, the probability of finding a commitment V and tuples (S, \mathbf{v}, π) , (S', \mathbf{v}', π') such that $\mathbf{v}[S \cap S'] \neq \mathbf{v}'[S \cap S']$ and

$$\text{Verify}(V, S, \mathbf{v}, \pi) = \text{Verify}(V, S', \mathbf{v}', \pi') = 1$$

is negligible in λ .

In the context of our APIR protocol, the **hiding** property is not required, since we assume that the database is public. For correctness, our protocol requires Commit to be **deterministic**. This is the case for Pointproofs and Merkle trees, the two VC schemes considered in this work.

2.2 Permutations

Generating permutations. Random permutations over “small” domains can be sampled using the Fisher-Yates shuffle, also known as the Fisher-Yates-Durstenfeld-Knuth shuffle [Dur64, Knu97, FY53]. At a high level, the algorithm takes as input a list of elements (from the set we wish to permute) and traverses the list, each time swapping the current element with a random element from the remaining list. This algorithm can sample an unbiased permutation of the set $[N]$ in $O(N)$ time. We restate the following lemma.

Lemma 2.2.1 ([Dur64, Knu97, FY53]). For any $N \in \mathbb{N}$, there exists an algorithm $\text{Permute}(N)$ that can output a permutation σ_N of the set $[N]$ sampled uniformly from the set of all permutations of $[N]$, in $O(N)$ time.

Both SinglePass and this work rely on this lemma for efficient and secure permutation generation.

Cycle notation. In our examples, we use *cycle notation* to compactly represent permutations. Cycle notation is used to describe a permutation $\sigma : S \rightarrow S$ as a list of disjoint cycles. Each cycle follows an element in S by repeatedly applying σ until the starting element is reached. As a concrete example, consider the permutation $\sigma = (4 \ 3 \ 1 \ 2)$ defined by $4 \mapsto 3$, $3 \mapsto 1$, $1 \mapsto 2$, $2 \mapsto 4$.

If all elements in S are contained in the cycle, then we are done, otherwise we pick an element in S not in the previous cycle(s) and start a new cycle:

$$\sigma : (x \ \sigma(x) \ \sigma(\sigma(x)) \ \dots)(y \ \sigma(y) \ \sigma(\sigma(y)) \ \dots) \dots$$

where $x, y \in S$. As a concrete example, consider the permutation σ' defined by $1 \mapsto 3$, $2 \mapsto 1$, $3 \mapsto 2$, $4 \mapsto 4$. which can be expressed in cycle notation as $(1 \ \sigma'(1) \ \sigma'(\sigma'(1)))(4) = (1 \ 3 \ 2)(4)$.

2.3 Private Information Retrieval

Private Information Retrieval (PIR) allows a client to retrieve an element of a database without the server holding that database being able to determine which element was retrieved. While this problem admits a trivial solution that achieves perfect privacy—the client

downloads the entire database—the communication complexity is linear in the database size. This overhead is problematic for many real-world use cases where databases are large and client storage is limited. PIR therefore focuses on succinctness to achieve query communication complexity sublinear in the database size.

State-of-the-art two-server PIR schemes achieve both sublinear query communication and computational cost under computational security guarantees [CK20, KC21, LP23, LP24]. In these schemes, the servers and client jointly preprocess the database to compute a hint that allows the client to efficiently query the database.

Standard PIR threat models typically assume non-colluding, semi-honest servers and only guarantee privacy under these assumptions. However, in practice—especially when PIR is used as a subprotocol within a larger system—a malicious server may choose to serve a wrong answer and observe whether a client aborts. This side-channel enables an attack known as a *selective failure attack*, which can leak information about the queried index and violate the query privacy guarantees. To provide privacy in the presence of this side-channel, new security definitions and constructions for APIR have been proposed e.g., [CNC⁺23, FMS25, AB25, DT24, dCL24, AR25]. In the next section, we extend the prior work and formally introduce two-server APIR with client-preprocessing and support for updates.

3 APIR with Preprocessing

In this section, we formalize APIR with preprocessing (Def. 2) and support for updates (Def. 3). An updatable APIR scheme allows the servers to update the database by editing and deleting existing values and adding new values.

We further state the threat model (§3.1) and define the required properties of correctness (§3.2.1), integrity (§3.2.2), and privacy with abort (§3.2.3). Our formalization captures APIR schemes with $k \geq 1$ servers, thus it also captures related work [AR25]. We note that the security definitions for the single- and multi-server settings differ in their threat model. While the multi-server setting assumes at least one honest server, this is not the case in the single-server setting. We address this in §3.1; the additional steps for $k \geq 2$ servers in the security games and oracles are **highlighted in teal**.

We distinguish between initial fixed-value protocol parameters and those that might change during protocol execution via, for example, updates (see §2). Additionally, we assume that the database is stored as part of the server state.

Definition 2. A k -server APIR scheme with preprocessing, for an input vector $\text{DB}_{\text{init}} \in \{0, 1\}^{N_{\text{init}} \times \ell_r}$ where $N_{\text{init}}, \ell_r \in \mathbb{N}$, and setup parameters sp , comprises of the following seven algorithms which all take the security parameter λ as an implicit input:

- $(\text{st}_{S_b}, \text{pp}_b) \leftarrow \text{Setup}(\text{sp}, \text{DB}_{\text{init}})$ is executed by each server $b \in [k]$ and takes as input setup parameters sp and a vector DB_{init} . It outputs server state st_{S_b} and public parameters pp_b . The server state includes the database DB_b (initialized using DB_{init} and possibly modified).
- $(\text{st}_C, \text{hq}) \leftarrow \text{RequestHint}(\text{pp})$ takes the servers' public parameters $\text{pp} = (\text{pp}_b)_{b \in [k]}$ as input and outputs the client state st_C and hint queries $\text{hq} = (\text{hq}_b)_{b \in [k]}$.
- $(\text{st}'_{S_b}, \text{resp}_b) \leftarrow \text{AnsHintReq}(\text{st}_{S_b}, \text{hq}_b)$ is executed by each server $b \in [k]$ and takes as input the server state st_{S_b} and hint query hq_b . It outputs server state st'_{S_b} and response resp_b .

- $(\text{st}'_C, \text{hint}) \leftarrow \text{VerSetup}(\text{st}_C, \text{resp})$ takes as input client state st_C , and server responses $\text{resp} = (\text{resp}_b)_{b \in [k]}$. It outputs updated client state st'_C , and hint hint .
- $(\text{st}'_C, \text{qry}) \leftarrow \text{Query}(\text{st}_C, \text{idx})$ takes as input client state st_C and index $\text{idx} \in [N]$. It returns updated client state st'_C and queries $\text{qry} = (\text{qry}_b)_{b \in [k]}$.
- $(\text{st}'_{S_b}, \text{ans}_b) \leftarrow \text{Answer}(\text{st}_{S_b}, \text{qry}_b)$ is executed by each server $b \in [k]$ and takes as input the server state st_{S_b} and query qry_b . It returns the server state st'_{S_b} and answer ans_b .
- $(\text{st}'_C, x, \text{hint}') \leftarrow \text{Recon}(\text{st}_C, \text{hint}, \text{ans})$ takes as input client state st_C , hint hint , and answers $\text{ans} = (\text{ans}_b)_{b \in [k]}$. It returns updated client state st'_C , a record $x \in \{0, 1\}^{\ell_r}$, and hint hint' .

We include setup parameters sp as part of the input to Setup to account for any values used to parameterize the scheme, including parameters generated during a trusted setup phase (e.g., the public parameters of vector commitments).

Definition 3. An **updatable k -server APIR scheme with preprocessing**, is a two-server APIR scheme with preprocessing, with two additional algorithms:

- $(\text{st}'_{S_b}, \text{pp}'_b, U_b) \leftarrow \text{UpdateDB}(\text{st}_{S_b}, U)$ takes as input server state st_{S_b} , and update information U . It outputs the updated state st'_{S_b} , updated public parameters pp'_b , and update information U_b .
- $(\text{st}'_C, \text{hint}') \leftarrow \text{UpdateHint}(\text{st}_C, \text{hint}, \text{pp}, U)$ takes as input the client state st_C , hint hint , and, from each server $b \in [k]$, the parameters $\text{pp} \leftarrow (\text{pp}_b)_{b \in [k]}$ and update information $U \leftarrow (U_b)_{b \in [k]}$. It outputs the updated client state st'_C , and hint hint' .

An APIR scheme with preprocessing in both the static (Def. 2) and the dynamic (Def. 3) setting must satisfy *correctness*, *integrity*, and *privacy with abort*. We define these notions in §3.2 according to the threat model specified below.

3.1 Threat Model

Standard PIR schemes do not provide integrity of the responses and, as such, cannot guarantee correctness if one/the server is malicious. In this work, we consider a stronger threat model in which privacy and integrity are guaranteed even if one or multiple servers act malicious. Malicious servers may deviate from the protocol at any phase of the protocol and may serve incorrect responses or “garbage” to the client.

We consider different threat models for the single- and multi-server settings: in the single-server setting, privacy, and integrity should still be guaranteed, even if the single server acts malicious. In the multi-server setting, we assume that at least one server is semi-honest and does not collude with the other (possibly malicious) servers. In particular, for our two-server APIR scheme TAPIR, this assumption ensures that one of the servers behaves semi-honestly.

We note that if all servers act maliciously, they could (by coincidence or via additional information) modify the database in the same manner, hence fooling the client into accepting incorrect information. Without making strong assumptions about the database, we cannot bound this probability and, therefore, require at least one semi-honest server.

Game $G_{\text{APIR}}^{\text{Corr-APIR}}(\lambda, \text{sp})$	Games $G_{\text{APIR}}^{\text{Int-APIR}}(\lambda, \text{sp})$ and $G_{\text{APIR}}^{\text{Priv-APIR}}(\lambda, \text{sp})$
<pre> 1 : global win, hint, st_C, (st_{S_b})_{b∈[k]} 2 : win ← false, resp ← ⊥, pp ← ⊥ 3 : // Initialization 4 : (st_A, DB_{init}) ← $\mathcal{A}_0(\text{sp})$ 5 : // Setup 6 : for b ∈ [k] do 7 : (st_{S_b}, pp, pp_b) ← Setup(sp, DB_{init}) 8 : // Preprocessing 9 : (st_C, hq) ← RequestHint(pp) 10 : for b ∈ [k] do 11 : (st_{S_b}, resp, resp_b) ← AnsHintReq(st_{S_b}, hq[b]) 12 : st_{S_b} ← st'_{S_b} 13 : (st'_C, hint) ← VerSetup(st_C, resp) 14 : st_C ← st'_C 15 : // Oracle phase 16 : $\mathcal{A}_1^{O_{\text{QueryCorr}}(\cdot), O_{\text{UpdateCorr}}(\cdot)}(\text{st}_{\mathcal{A}})$ 17 : return win </pre>	<pre> 1 : global win, {b_{chall}} good, inpt, hint, st_C, st_S, t_q, t_u 2 : win ← false 3 : inpt ← ⊥, t_q, t_u ← 0 4 : // Initialization 5 : (st_A, good, (pp_b)_{b∈[k] \ good}, DB_{init}) ← $\mathcal{A}_0(\text{sp})$ 6 : // Setup 7 : (st_S, pp_{good}) ← Setup(sp, DB_{init}) 8 : b_{chall} ← {0, 1} // Pick challenge bit 9 : // Preprocessing 10 : (st_C, hq) ← RequestHint((pp_b)_{b∈[k]}) 11 : (st'_S, resp_{good}) ← AnsHintReq(st_S, hq_{good}) 12 : (st'_A, (resp_b)_{b∈[k] \ good}) ← $\mathcal{A}_1(\text{st}_{\mathcal{A}}, (\text{hq}_b)_{b∈[k] \setminus \text{good}})$ 13 : (st'_C, hint) ← VerSetup(st_C, (resp_b)_{b∈[k]}) 14 : st_C ← st'_C, st_S ← st'_S 15 : // Oracle phase 16 : $\mathcal{A}_2^{O_{\text{QueryInt}}(\cdot), O_{\text{ReconInt}}(\cdot), O_{\text{UpdateInt}}(\cdot)}(\text{st}'_{\mathcal{A}})$ 17 : b' ← $\mathcal{A}_2^{O_{\text{QueryPriv}}(\cdot), O_{\text{ReconPriv}}(\cdot), O_{\text{UpdatePriv}}(\cdot)}(\text{st}'_{\mathcal{A}})$ 18 : return win, return b_{chall} = b' </pre>

Figure 1: The games for APIR correctness (left), integrity (right, solid box), and privacy (left, dashed box). The corresponding oracles are defined in Fig. 2. Highlighted code only applies to the multi-server setting ($k \geq 2$).

3.2 Security Properties

In this section, we define the security properties of an updatable APIR scheme with preprocessing (Def. 3). We note that the definitions for APIR with a static database are the same, except the adversary has no access to any Update oracles. These static definitions are excluded for succinctness.

Before we go into the details of these definitions, we provide some context for our game-based security definitions and the capabilities of the adversary. We define the security games in Fig. 1 and the corresponding oracles in Fig. 2. The code and descriptions that apply only to the multi-server setting—where we assume a “good” server—are highlighted in teal.

Let APIR be a (possibly updateable) two-server APIR scheme with preprocessing. The games take as input security parameter λ and setup parameters sp , and they track global variables throughout oracle calls. These variables include the state of the client and of the honest server.

We assume that all adversary outputs are well-formed and within the correct domain. The adversary has the capability to choose the initial set of records DB_{init} for the protocol execution in all games. In the privacy and integrity games, the adversary also chooses an index $\text{good} \in [k]$ that defines which server is honest. Additionally, the adversary is adaptive and has access to a set of oracles. These oracles can be executed in any order and depend on the game. We assume an efficient adversary that makes at most T_q calls to a query oracle and at most T_u calls to an update oracle, such that $T_q + T_u \leq \text{poly}(\lambda)$.

The games ensure that the algorithms of the offline phase are executed in the required order, i.e., Setup, RequestHint, AnsHintReq, VerSetup. The games, oracles, and protocol description could be adapted to move this requirement into the protocols themselves. We exclude this here for simplicity.

3.2.1 Correctness. An APIR scheme is correct if, when a client requests the record at index idx , it receives the correct value stored at that index in the database, i.e., $\text{DB}[\text{idx}]$. Correctness is guaranteed to hold only when both the client and the servers behave honestly.

Definition 4. Let APIR be an updatable k -server APIR scheme with preprocessing, let $\lambda \in \mathbb{N}$ be the security parameter, and let sp be a well-formed and honestly generated setup parameter. Let $G_{\text{APIR}}^{\text{Corr-APIR}}$ be the correctness game for APIR defined in Fig. 1 and Fig. 2. The advantage of an adversary \mathcal{A} playing this game is defined as

$$\text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{Corr}}(\lambda) := \Pr [G_{\text{APIR}}^{\text{Corr-APIR}}(\lambda, \text{sp})].$$

We say that APIR is **correct** if $\text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{Corr}}(\lambda) = 0$ for all PPT \mathcal{A} .

3.2.2 Integrity. We define integrity against adaptive queries (and updates) for APIR schemes. Informally, an APIR scheme satisfies integrity if the client outputs either the correct value or aborts with \perp , except with negligible probability $\text{negl}(\lambda)$.

Definition 5. Let APIR be an updatable k -server APIR scheme with preprocessing, let $\lambda \in \mathbb{N}$ be the security parameter, and let sp be a well-formed and honestly generated setup parameter. Let $G_{\text{APIR}}^{\text{Int-APIR}}$ be the adaptive integrity game for APIR defined in Fig. 1 and Fig. 2.

Oracle $O_{\text{QueryInt}}(\text{idx})$	Oracle $O_{\text{ReconInt}}((\text{ans}_b)_{b \in [k] \setminus \text{good}})$	Oracle $O_{\text{UpdateInt}}((\text{pp}_b)_{b \in [k] \setminus \text{good}}, (U_b)_{b \in [k] \setminus \text{good}}, U)$
1: $(\text{st}'_C, \text{qry}) \leftarrow \text{Query}(\text{st}_C, \text{idx})$ 2: $(\text{st}'_S, \text{ans}_{\text{good}}) \leftarrow \text{Answer}(\text{st}_S, \text{qry}[\text{good}])$ 3: $\text{inpt} \leftarrow (\text{ans}_{\text{good}}, \text{idx})$ 4: $\text{st}_S \leftarrow \text{st}'_S, \text{st}_C \leftarrow \text{st}'_C$ 5: return $\text{qry}[[k] \setminus \text{good}]$	1: $(\text{ans}_{\text{good}}, \text{idx}) \leftarrow \text{inpt}$ 2: $(\text{st}'_C, x, \text{hint}') \leftarrow \text{Recon}(\text{st}_C, \text{hint}, (\text{ans}_b)_{b \in [k]})$ 3: if $x \neq \perp \wedge x \neq \text{DB}[\text{idx}]$ then 4: $\text{win} \leftarrow \text{true}$ 5: $\text{st}_C \leftarrow \text{st}'_C, \text{hint} \leftarrow \text{hint}', \text{inpt} \leftarrow (\perp, \perp)$	1: $(\text{st}'_S, \text{pp}_{\text{good}}, U_{\text{good}}) \leftarrow \text{UpdateDB}(\text{st}_S, U)$ 2: $(\text{st}'_C, \text{hint}') \leftarrow \text{UpdateHint}(\text{st}_C, \text{hint}, (\text{pp}_b)_{b \in [k]}, (U_b)_{b \in [k]})$ 3: $\text{st}_C \leftarrow \text{st}'_C, \text{hint} \leftarrow \text{hint}', \text{st}_S \leftarrow \text{st}'_S$
Oracle $O_{\text{QueryPriv}}(\text{idx}_0, \text{idx}_1)$	Oracle $O_{\text{ReconPriv}}((\text{ans}_b)_{b \in [k] \setminus \text{good}})$	Oracle $O_{\text{UpdatePriv}}(\text{pp}_b)_{b \in [k] \setminus \text{good}}, (U_b)_{b \in [k] \setminus \text{good}}, U)$
1: if $t_q < T_q$ then 2: $(\text{st}'_C, \text{qry}) \leftarrow \text{Query}(\text{st}_C, \text{idx}_{b_{\text{chall}}})$ 3: $(\text{st}'_S, \text{ans}_{\text{good}}) \leftarrow \text{Answer}(\text{st}_S, \text{qry}[\text{good}])$ 4: $\text{inpt} \leftarrow \text{ans}_{\text{good}}$ 5: $\text{st}_S \leftarrow \text{st}'_S, \text{st}_C \leftarrow \text{st}'_C, t_q \leftarrow t_q + 1$ 6: return $\text{qry}[[k] \setminus \text{good}]$	1: $\text{ans}_{\text{good}} \leftarrow \text{inpt}$ 2: $(\text{st}'_C, x, \text{hint}') \leftarrow \text{Recon}(\text{st}_C, \text{hint}, (\text{ans}_b)_{b \in [k]})$ 3: if $x = \perp$ then 4: $\text{abort-bit} \leftarrow \text{true}$ 5: else $\text{abort-bit} \leftarrow \text{false}$ 6: $\text{st}_C \leftarrow \text{st}'_C, \text{hint} \leftarrow \text{hint}', \text{inpt} \leftarrow \perp$ 7: return abort-bit	1: if $t_u < T_u$ then 2: $(\text{st}'_S, \text{pp}_{\text{good}}, U_{\text{good}}) \leftarrow \text{UpdateDB}(\text{st}_S, U)$ 3: $(\text{st}'_C, \text{hint}') \leftarrow \text{UpdateHint}(\text{st}_C, \text{hint}, (\text{pp}_b)_{b \in [k]}, (U_b)_{b \in [k]})$ 4: $t_u \leftarrow t_u + 1$ 5: $\text{st}_S \leftarrow \text{st}'_S, \text{st}_C \leftarrow \text{st}'_C, \text{hint} \leftarrow \text{hint}'$
Oracle $O_{\text{QueryCorr}}(\text{idx})$	Oracle $O_{\text{UpdateCorr}}(U)$	
1: $(\text{st}'_C, \text{qry}) \leftarrow \text{Query}(\text{st}_C, \text{idx})$ 2: $\text{ans} \leftarrow \perp$ 3: for $b \in [k]$ do 4: $(\text{st}'_{S_b}, \text{ans}_{\text{ans}_b}) \leftarrow \text{Answer}(\text{st}_{S_b}, \text{qry}[b])$ 5: $\text{st}_{S_b} \leftarrow \text{st}'_{S_b}$ 6: $(\text{st}'_C, x, \text{hint}') \leftarrow \text{Recon}(\text{st}'_C, \text{hint}, \text{ans})$ 7: if $x \neq \perp \wedge x \neq \text{DB}[\text{idx}]$ then 8: $\text{win} \leftarrow \text{true}$ 9: $\text{st}_C \leftarrow \text{st}'_C, \text{hint} \leftarrow \text{hint}'$	1: $(U, \text{pp}) \leftarrow \perp^2$ 2: for $b \in [k]$ do 3: $(\text{st}'_{S_b}, \text{pp}_{\text{pp}_b}, U, U_b) \leftarrow \text{UpdateDB}(\text{st}_{S_b}, U)$ 4: $\text{st}_{S_b} \leftarrow \text{st}'_{S_b}$ 5: $(\text{st}'_C, \text{hint}') \leftarrow \text{UpdateHint}(\text{st}_C, \text{hint}, \text{pp}, U)$ 6: $\text{st}_C \leftarrow \text{st}'_C, \text{hint} \leftarrow \text{hint}'$	

Figure 2: Oracles for $G_{\text{APIR}}^{\text{Int-APIR}}$ (top), $G_{\text{APIR}}^{\text{Priv-APIR}}$ (middle), and $G_{\text{APIR}}^{\text{Corr-APIR}}$ (bottom). Highlighted code only applies if $k \geq 2$.

The advantage of an adversary \mathcal{A} playing this game is defined as

$$\text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{int}}(\lambda) := \Pr \left[G_{\text{APIR}}^{\text{Int-APIR}}(\lambda, \text{sp}) \right].$$

We say that APIR achieves **integrity** against adaptive queries and updates if $\text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{int}}(\lambda) \leq \text{negl}(\lambda)$ for all PPT \mathcal{A} .

3.2.3 Privacy. An APIR scheme is considered private if the server(s) gain no information about which record the client queries. This guarantee must hold even if the server(s) learn whether the client outputs the error symbol \perp and aborts during reconstruction. This stronger notion of privacy was introduced in [CNC⁺23] and is designed to protect against selective-failure attacks. In this work, we formalize privacy against adaptive queries and updates via a left-or-right indistinguishability game. At the start of the game, a random bit b_{chall} is chosen. The adversary submits two indices to the $O_{\text{QueryPriv}}$ oracle, but only the b_{chall} -th index is used to execute the query. At the end of the game, the adversary must guess which bit was selected.

Definition 6. Let APIR be an updatable k -server APIR scheme with preprocessing, let $\lambda \in \mathbb{N}$ be the security parameter, and let sp be a well-formed and honestly generated setup parameter. Let $G_{\text{APIR}}^{\text{Priv-APIR}}$ be the privacy game against adaptive queries (and updates) for APIR defined in Fig. 1 and Fig. 2. The advantage of an adversary \mathcal{A}

playing this game is defined as

$$\text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{priv}}(\lambda) := \left| \Pr \left[G_{\text{APIR}}^{\text{Priv-APIR}}(\lambda, \text{sp}) \right] - \frac{1}{2} \right|.$$

We say that APIR achieves **privacy** against adaptive queries if $\text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{priv}}(\lambda) \leq \text{negl}(\lambda)$ for all PPT \mathcal{A} .

4 Upgrading SinglePass to Malicious Security

We start by providing a brief overview of the unauthenticated two-server PIR scheme with client-preprocessing, SinglePass [LP24]. We then outline the challenges and solutions to upgrading SinglePass to malicious security.

4.1 An Overview of SinglePass

To achieve privacy in SinglePass, the two servers are assigned distinct roles. Server 1 is responsible for computing the hint during the offline phase and providing updates to it during the online phase. These hint updates ensure that queries remain independent over time, thereby preventing information leakage across multiple queries. Server 2 is responsible for supplying the values required for the client to reconstruct the queried record during the online phase. Crucially, because the hint is used to generate the client's queries, the hint server must not have access to the values

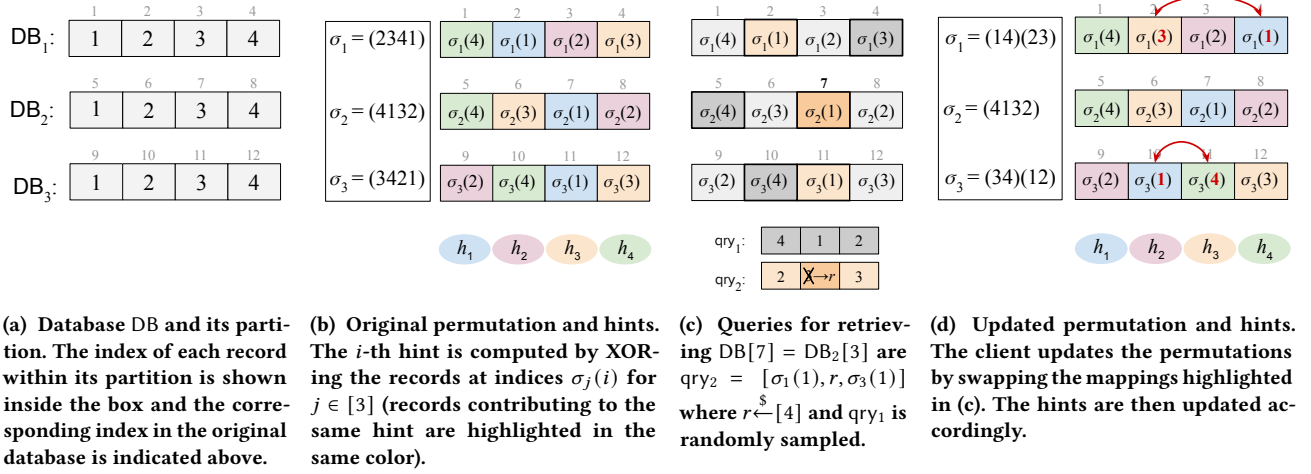


Figure 3: An example of the SinglePass scheme [LP24] for a database DB with $N = 12$, $Q = 3$, and $M = N/Q = 4$.

used to reconstruct the queried record. Otherwise, it could correlate the hint with the responses and infer the client's queried index.

We begin by assuming the servers share a database DB of size N , which is partitioned into Q disjoint sub-arrays, each of size M . The q -th partition is denoted $\text{DB}_q = \text{DB}[(q-1) \cdot M + 1 : q \cdot M]$, where $Q = \lceil N/M \rceil$. The database is padded to its maximum capacity.

We provide a running example of the scheme in Fig. 3, and show a database partition for $N = 12$ and $Q = 3$ in Fig. 3a.

Offline phase. The offline phase starts with Server 1 sampling a permutation σ_q for each partition DB_q , $q \in [Q]$. For each $m \in [M]$, the server computes the hint $h_m = \bigoplus_{q \in [Q]} \text{DB}_q[\sigma_q(m)]$ and sends hint $= (h_1, \dots, h_M)$ to the client (with the seeds for generating the permutations). In Fig. 3b, we give an example of three permutations and the hint they result in. Concretely, for h_1 , we have that

$$h_1 = \bigoplus_{q \in [3]} \text{DB}_1[\sigma_q(1)] = \text{DB}_1[2] \oplus \text{DB}_2[3] \oplus \text{DB}_3[3].$$

Online phase. To issue a query for an index $\text{idx} \in [N]$, the client must first find the partition, q^* , and index within the partition, m^* , that map to the index idx in the original database, i.e., $\text{DB}[\text{idx}] = \text{DB}_{q^*}[m^*]$. The client then computes $\text{ind} \in [M]$ such that $\sigma_{q^*}(\text{ind}) = m^*$ and an array of indices defined as $\text{qry}_2 = [\sigma_q(\text{ind}) : q \in [Q]]$. Next, the client replaces $\text{qry}_2[m^*]$ with a randomly sampled value in $[M]$ (thereby hiding the index of interest and ensuring query privacy). In Fig. 3c, we give an example for querying $\text{idx} = 7$ which corresponds to $(q^*, m^*) = (2, 3)$. The client computes $\text{ind} = 1$ since $\sigma_2(1) = 3$ and sets $\text{qry}_2 = [\sigma_1(1), r, \sigma_3(1)] = [2, r, 3]$ where $r \xleftarrow{\$} [4]$ is random.

In parallel to computing array qry_2 , the client also samples $r_1, \dots, r_Q \xleftarrow{\$} [M]^Q$, and sets $\text{qry}_1 = [\sigma_q(r_q) : q \in [Q]]$. The arrays qry_1 and qry_2 are then sent to the respective servers. The servers respond by sending an array of the records at the requested indices.

Given the answer from Server 2, $[\text{DB}_q[\text{qry}_2[q]] : q \in [Q]]$, the client can recover the desired record as follows:

$$\text{DB}[\text{idx}] = \text{DB}_{q^*}[m^*] = \left(\bigoplus_{\substack{q \in [Q] \\ q \neq q^*}} \text{DB}_q[\text{qry}_2[q]] \right) \oplus h_{\text{ind}}.$$

Finally, the client takes the answer from Server 1, $[\text{DB}_q[\text{qry}_1[q]] : q \in [Q]]$, and updates the hint, such that $q \neq q^*$ for $q \in [Q]$:

$$h_{\text{ind}} = h_{\text{ind}} \oplus \text{DB}[\text{qry}_1[q]] \oplus \text{DB}[\text{qry}_2[q]]$$

$$h_{r_q} = h_{r_q} \oplus \text{DB}[\text{qry}_1[q]] \oplus \text{DB}[\text{qry}_2[q]].$$

Figure 3d gives an example of the refreshed hint.

The privacy of SinglePass relies on the fact that the client sends a pseudorandom vector of indices—based on the permutations—to each server. If the queried index is in the set, the client replaces it with a random index from the same range. The server receiving this vector does not know the permutations and, thus, cannot identify which, or even if, an index was replaced. Lazzaretti and Papamanthou [LP24] state and prove this with their Show-and-Shuffle theorem which we restate in Theorem A.1.1 in Appendix A.1.

4.2 Upgrading to Malicious Security

We now outline the challenges we encountered in extending SinglePass to guarantee malicious security and privacy with abort.

Challenge 1: Computing the hint privately. In SinglePass, the client and one server jointly preprocess the database to compute a hint stored at the client. Upgrading this step to be secure against a malicious server poses a challenge: a malicious server cannot be trusted to compute the hint correctly. Our key insight is to make the hint generation both private and verifiable. This could be achieved using heavy cryptographic tools (e.g., distributed multi-point functions [KKEPR24, BGH⁺25] or custom MPC protocols), but we prioritize computational efficiency. Instead, we assume a streaming model, in which both servers stream the database to the client. The client checks that the records streamed from both servers match and then computes the hint locally. Streaming—a standard technique in PIR [ZPZS24, GZS24, GZSP25, AR25]—leaks no information to the servers and significantly reduces their computational cost compared to, e.g., SinglePass. Instead, the client runs the preprocessing itself.

Challenge 2: Protecting against selective failure. Note that the servers in prior multi-server APIR schemes [CNC⁺23] have

symmetric roles, i.e., both servers do the same operations (on different inputs). In existing PIR schemes with client-preprocessing, one server typically generates and updates the hint, while the other only answers queries. In contrast, to ensure privacy in the offline phase against a potentially malicious server, we require both servers to participate in the setup identically. This symmetry prevents either server from gaining an advantage in violating privacy. To achieve this, we adopt a stream-and-compare approach, as mentioned above. Other approaches may also be possible, e.g., by using zero-knowledge proof techniques, which we leave as a direction for future work.

In addition, we leverage vector commitments [CF13] to verify the correctness of queried records. Vector commitments can be viewed as a generalization of Merkle trees and Pedersen vector commitments, both of which have been used in prior work [CNC⁺23, DT24, FMS25, dCL24].

Challenge 3: Guaranteeing security of updates. Another challenge was proving that integrity and query privacy hold when adaptive lookup and update queries are made to the database. In contrast, the original paper [LP24] only proved security for queries to static databases, and their security definition only considered non-adaptive queries. We address this challenge by providing new game-based security notions for updateable APIR with preprocessing. Our games (see Fig. 1) consider an adversary with access to query, reconstruction, and update oracles. The adversary may adaptively choose to make a polynomial number of queries, and may even interleave lookup and update queries arbitrarily. Care was taken to strengthen our definitions by ensuring that trivial wins/losses were excluded.

5 The TAPIR Scheme

Our scheme proceeds in two phases: an *offline* phase, where the client preprocesses the database, followed by an *online* phase, where the client queries the server based on the preprocessed information. The preprocessing phase allows communication and computational complexities of the online phase to be sublinear in the database size. This section first provides a high-level description of TAPIR for static databases (§5.1) before extending it to updatable databases (§5.2). Lastly, §5.3 states the correctness, integrity, and security of TAPIR.

5.1 Scheme Overview

Our scheme allows the client to abort in most stages of the protocol if it does not accept the input from the servers. This is denoted by the client's internal abort state $st_C.abort$ being set to **true**. Once this state is set to **true**, the client stops all computation and only outputs \perp . On input \perp , the servers also directly output \perp . The parties do not proceed with any actual retrievals or computation; thus, the protocol is effectively aborted. SinglePass, and consequently our scheme, requires that the algorithms be executed in a strict order. In particular, the client must run Recon after Query and before invoking either Query again or UpdateHint. This ordering is enforced at the start of the client's online algorithms by checking whether $st_C.q^*$ is set (for Recon) or unset (otherwise).

5.1.1 Setup parameters. The setup parameters $sp = (N_{init}, Q_{init}, M, \ell_r, \lambda)$ specify the initial database size N_{init} , the number of partitions

Q_{init} , the size of each partition M , the size of each record ℓ_r , and the security parameter λ . Our scheme is initialized with these setup parameters sp and a vector of database records $DB_{init} \in 0, 1^{N_{init} \times \ell_r}$. The number of partitions is $Q_{init} = \lceil N_{init}/M \rceil$, which is tunable, allowing for a trade-off between communication cost and computation time. The setup parameters also include any values required to parameterize the underlying VC scheme, which we leave implicit.

5.1.2 Offline phase. The Setup algorithm—run by each server—begins by partitioning DB_{init} into Q_{init} partitions of size M and padding as necessary to obtain DB . For each $q \in [Q_{init}]$, the servers compute a vector commitment com_q to each partition DB_q such that the m -th record committed to by com_q corresponds to record $DB_q[m]$. For every record in DB , the servers also compute an opening proof, storing it in a secondary database \widehat{DB} of the same size as DB , with each proof stored at the same index as its corresponding record. Using the setup parameters, the servers initialize the public parameters pp , which include the database metadata, security parameters, and the database digest. Both servers maintain copies of DB and \widehat{DB} , along with pp , as part of their state.

When a client joins the system, it first runs RequestHint, during which it requests two things from the servers: (1) the public parameters and associated vector commitments that serve as a digest of the database, and (2) a streamed copy of the database from which it can compute a SinglePass-like hint. It first initializes the abort bit $st_C.abort \leftarrow \text{false}$. To obtain (1), the client downloads the public parameters pp_1 and pp_2 from Servers 1 and 2, respectively, and checks if $pp_1 \neq pp_2$. Since at least one server is honest, equality guarantees that the digest was generated consistently. If the public parameters differ, the client aborts. Otherwise, the client proceeds to obtain (2) by sending a hint request to both servers.

On input of hint query hq_b , Server $b \in \{1, 2\}$ runs AnshintReq. If $hq_b \neq \perp$, the server proceeds to stream a copy of the database. Streaming allows the client to verify the equality of corresponding records across both servers, mitigating the risk of a malicious server providing inconsistent data. If the received information is not equal, the client aborts. Otherwise, it samples secret random permutations $\sigma_1, \dots, \sigma_Q : [M] \rightarrow [M]$ —one for each partition of size M —and computes the hint as $h_m = \bigoplus_{q=1}^Q DB_q[\sigma_q(m)]$ and sets $hint = (h_1, \dots, h_M)$. This ensures that the hint is consistent with the committed database while hiding its value from both servers. The client then stores the verified public parameters pp , the computed hint $hint$, and the client key ck (i.e., the set of secret permutations) as part of its state, completing the offline phase.

5.1.3 Online phase. To retrieve a record from the database, the client runs Query on this record's index idx corresponding to $(q^*, m^*) \in [pp.Q] \times [pp.M]$. If the client aborted earlier, it simply outputs \perp to the servers. Otherwise, the client continues with the protocol and issues a query to each server. These queries are computed as in SinglePass and consist of an array of database indices.

On receiving query qry_b , server b executes Answer. If the client did not abort, the server returns the records at the requested indices, along with the corresponding opening proofs from \widehat{DB} , to the client.

Upon receiving ans_1 and ans_2 from the servers, the client executes Recon, checking whether $st_C.abort = \text{true}$ and exiting early if yes. Otherwise, it verifies *all* the proofs output by the servers

Setup(sp, DB _{init}) → (st _{S_b} , pp _b)	RequestHint(pp) → (st _C , hq)	VerSetup(st _C , resp) → (st' _C , hint)
1: (N _{init} , Q _{init} , M, ℓ _r , λ) ← sp	1: st _C ← ∅, st _C .abort ← false	1: // Verify streamed records
2: Initialize (DB, DB̂) ← 0 ^{2·((Q_{init}·M)×ℓ_r)} , d ← ⊥	2: // Verify public parameters	2: if st _C .abort ∨ resp[1] ≠ resp[2] then
3: DB[N _{init}] ← DB _{init}	3: if pp[1] ≠ pp[2] then	3: st _C .abort ← true
4: for q ∈ [Q _{init}] do	4: st _C .abort ← true	4: return (st _C , ⊥)
5: d.com _q ← VC.Commit(DB _q)	5: return (st _C , ⊥)	5: (ck, hint) ← ⊥ ²
6: // Compute and store proofs	6: (hq ₁ , hq ₂) ← ("Stream DB, please!") ²	6: // Generate permutations
7: for m ∈ [M] do	7: st _C .pp ← pp[1]	7: for q ∈ [st _C .pp.Q] do
8: DB̂ _q [m] ← VC.Open(m, DB _q)	8: return (st _C , (hq ₁ , hq ₂))	8: ck.σ _q ← \$ Permute(st _C .pp.M)
9: pp _b ← (M, ℓ _r , λ, d), pp _b .N ← N _{init} , pp _b .Q ← Q _{init}	AnsHintReq(st _{S_b} , hq _b) → (st' _{S_b} , resp _b)	9: DB ← resp[1]
10: return ((DB, DB̂, pp _b), pp _b)	1: if hq _b = ⊥ then return (st _{S_b} , ⊥)	10: // Compute hint as in Singlepass
	2: // Stream database	11: for m ∈ [st _C .pp.M] do
	3: return (st _{S_b} , st _{S_b} .DB)	12: hint.h _m = $\bigoplus_{q=1}^{st_C.pp.Q} DB_q[ck.σ_q(m)]$
		13: st _C .ck ← ck, st _C .q* ← ⊥
		14: return (st _C , hint)

Figure 4: The offline phase of our two-server APIR scheme.

Query(st _C , idx) → (st' _C , qry)	Recon(st _C , hint, ans) → (st' _C , x, hint')
1: if st _C .abort ∨ st _C .q* ≠ ⊥ then return (st _C , ⊥, ⊥)	1: if st _C .abort ∨ st _C .q* = ⊥ then return (st _C , ⊥, ⊥)
2: pp ← st _C .pp, {σ ₁ , ..., σ _{pp.Q} } ← st _C .ck	2: pp ← st _C .pp, ind ← st _C .ind, q* ← st _C .q*
3: q* ← ⌈idx/pp.M⌉	3: (qry _{1,1} , ..., qry _{1,pp.Q}) ← st _C .qry ₁
4: m* ← ((idx - 1) mod pp.M) + 1	4: (qry _{2,1} , ..., qry _{2,pp.Q}) ← st _C .qry ₂
5: Find ind ∈ [pp.M], s.t. σ _{q*} (ind) = m*	5: ((rec _{1,1} π _{1,1}), ..., (rec _{1,pp.Q} π _{1,pp.Q})) ← ans[1]
6: (r*, r ₁ , ..., r _{pp.Q}) ← \$ [pp.M] ^{pp.Q+1}	6: ((rec _{2,1} π _{2,1}), ..., (rec _{2,pp.Q} π _{2,pp.Q})) ← ans[2]
7: (qry ₁ , qry ₂) ← ⊥ ²	7: (com ₁ , ..., com _{pp.Q}) ← pp.d
8: for q ∈ [pp.Q] do	8: (h ₁ , ..., h _{pp.M}) ← hint
9: qry ₁ .qry _{1,q} ← σ _q (r _q)	9: // Verify commitments
10: qry ₂ .qry _{2,q} ← σ _q (ind)	10: for q ∈ [pp.Q] do
11: if q ≠ q* then	11: for b ∈ {1, 2} do
12: Swap σ _q (ind) and σ _q (r _q)	12: v ← VC.Verify(com _q , qry _{b,q} , rec _{b,q} , π _{b,q})
13: else qry ₂ .qry _{2,q*} ← r*	13: if ¬v then
14: st _C .ck ← {σ ₁ , ..., σ _{pp.Q} }	14: st _C .abort ← true
15: st _C ← (q*, ind, qry ₁ , qry ₂)	15: return (st _C , ⊥, ⊥)
16: return (st _C , (qry ₁ , qry ₂))	16: // Recover record
Answer(st _{S_b} , qry _b) → (st' _{S_b} , ans _b)	17: x ← $\left(\bigoplus_{q \in [pp.Q] \setminus \{q^*\}} rec_{2,q} \right) \oplus h_{ind}$
1: if qry _b = ⊥ then return (st _{S_b} , ⊥)	18: for q ∈ [pp.Q] \ {q*} do // Update hint
2: DB ← st _{S_b} .DB, DB̂ ← st _{S_b} .DB̂	19: h _{ind} ← h _{ind} ⊕ rec _{1,q} ⊕ rec _{2,q}
3: Q ← st _{S_b} .pp.Q	20: h _{r_q} ← h _{r_q} ⊕ rec _{1,q} ⊕ rec _{2,q}
4: (qry _{b,1} , ..., qry _{b,Q}) ← qry _b	21: st _C .ck ← {σ _q q ∈ [pp.Q]}, st _C .q* ← ⊥
5: return [DB _q [qry _{b,q}] DB̂ _q [qry _{b,q}] q ∈ [Q]]	22: return (st _C , x, (h ₁ , ..., h _{pp.M}))

Figure 5: The online Phase of our two-server APIR scheme.

and, if any verification fails, aborts. Finally, if all checks pass, then the client reconstructs the requested record by combining the returned records with the locally-stored hint, exactly as in SinglePass. It updates the hint using the fresh randomness obtained from ans₁, ensuring consistency for future queries. Finally, the client updates its state with the updated permutations and outputs its state st_C, the record of interest x, and the updated hint.

Intuition for privacy. Observe that since r* is uniformly sampled from [M], the probability that the client directly queries for DB_{q*}[m*] (the record of interest) is just as likely as it querying for any other record in DB_{q*} (see Fig. 4 Query line 13). To show that the set of indices in a query reveals no information to the server, we leverage the Show-and-Shuffle theorem introduced by [LP24]

and restated in Appendix A.1. The queried record is correctly reconstructed if (1) the hint is correct up until the client runs Query and (2) the answers returned by the two servers are correct. As we show in §5.3, the hint is always correct; otherwise, the client aborts. Thus, it is sufficient to check the proofs returned in the answers. Importantly, since the requested records leak nothing about the queried index, aborting if the opening proofs in the answers are invalid, leaks nothing about the queried index. In contrast, if we were to only check the opening proof of the reconstructed record, then the servers could carry out a selective failure attack.

5.2 Supporting Updates

This section provides an overview of how TAPIR can be extended to support appends (add) and edits (edit). We give the pseudocode for our updatable APIR scheme in Fig. 6. Observe that updates require the servers to update their database entries and generate/update the digests and commitments. The client must then update its state and hint accordingly.

5.2.1 Updating the Database. On input of server state $st_{S_b} = (pp_b, DB, \widehat{DB})$ and a sequence of updates U , Server b executes UpdateDB. Let $pp' \leftarrow pp_b$, $st'_{S_b} \leftarrow st_{S_b}$, and initialize set $U_b \leftarrow \emptyset$. The procedure iterates through the updates in U and assigns appropriate indices to all new records. This ensures that all operations are applied consistently, since add operations to an existing partition $q \in [pp'.Q]$ are treated equivalently to edit operations.

The procedure then computes the updated database parameters and stores them in pp' , where the new number of partitions is $pp'.Q = \lceil pp'.N/M \rceil$ and $pp'.Q$ denotes the number of partitions before the update. Let U'_q denote the subset of updates corresponding to partition q . The algorithm iterates over each partition $q \in [pp'.Q]$ and applies the associated update operations U'_q .

For each update $(op, idx, x_{new}) \in U'_q$, where op denotes the operation, idx is the position m in the partition DB_q , x_{new} is the new value, and x_{old} is the old value, the following steps are performed:

- Update the database record: $DB_q[m] \leftarrow x_{new}$.
- Update the partition's vector commitment at index m to x_{new} if the partition is not new, i.e., if $q \leq pp'.Q$.
- Append the update information $(op, idx, x_{new} \oplus x_{old})$ to U_b . Note that U_b stores the delta between the old and new values.

After all updates to a partition have been applied, if the partition is new (i.e., $q > pp'.Q$), then a fresh vector commitment com_q is computed for the partition and added to the digest $pp'.d$. Finally, UpdateDB refreshes the opening proofs in \widehat{DB} and outputs the updated server state st'_{S_b} , the updated public parameters pp' , and the update information U_b .

5.2.2 Updating the Hint. Updates on the client side require three steps: (i) verifying the equality of the new public parameters—which contain the updated digests—and the update information returned by the servers, (ii) generating permutations for new partitions, and (iii) updating the hint using the verified update information and secret permutations.

UpdateHint takes as input the client state st_C , the current hint $hint$, the updated public parameters pp_1, pp_2 , and the update information U_1, U_2 from the two servers. Step (i) verifies consistency of the new parameters and updates. If $pp_1 \neq pp_2$, $U_1 \neq U_2$, or

UpdateDB(st_{S_b}, U) $\rightarrow (st'_{S_b}, pp'_b, U_b)$

```

1:  $pp' \leftarrow st_{S_b}.pp$ ,  $st'_{S_b} \leftarrow st_{S_b}$ ,  $Q_{old} \leftarrow pp'.Q$ 
2:  $U', U_b \leftarrow \emptyset$ 
3: for  $(op, idx, x) \in U$  do
4:   if  $op = \text{add}$  then // Add index to new records
5:      $pp'.N \leftarrow pp'.N + 1$ 
6:      $U' \leftarrow U' \cup \{(op, pp'.N, x)\}$ 
7:   else  $U' \leftarrow U' \cup \{(op, idx, x)\}$ 
8:  $pp'.Q \leftarrow \lceil pp'.N/pp'.M \rceil$  // Updated num. partitions
9: Extend  $st'_{S_b}.DB$ ,  $st'_{S_b}.\widehat{DB}$ ,  $pp'.d$  to  $pp'.Q$  partitions/values
10:  $(com'_1, \dots, com'_{pp'.Q}) \leftarrow pp'.d$ 
11: for  $q \in [pp'.Q]$  do
12:   // Get all update operations for partition  $q$ 
13:    $U'_q \leftarrow \{(op, idx, x) \in U' : pp'.M \cdot (q-1) < idx \leq pp'.M \cdot q\}$ 
14:   for  $(op, idx, x_{new}) \in U'_q$  do
15:      $m \leftarrow ((idx-1) \bmod pp'.M) + 1$ 
16:      $x_{old} \leftarrow st'_{S_b}.DB_q[m]$ 
17:      $st'_{S_b}.DB_q[m] \leftarrow x_{new}$ 
18:     if  $q \leq Q_{old}$  then // Update partition commitment
19:        $com'_q \leftarrow VC.Update(com'_q, m, x_{old}, x_{new})$ 
20:      $U'_q \leftarrow U'_q \cup \{(op, idx, x_{new} \oplus x_{old})\}$ 
21:     if  $q > Q_{old}$  then // New partition commitment
22:        $com'_q \leftarrow VC.Commit(st'_{S_b}.DB_q)$ 
23:     for  $m \in [pp'.M]$  do // Update opening proofs.
24:        $st'_{S_b}.\widehat{DB}_q[m] \leftarrow VC.Open(m, st'_{S_b}.DB_q)$ 
25:    $U_b \leftarrow U_b \cup U'_q$ 
26:  $pp'.d \leftarrow (com'_1, \dots, com'_{pp'.Q})$ ,  $st'_{S_b}.pp \leftarrow pp'$ 
27: return  $(st'_{S_b}, pp', U_b)$ 

```

UpdateHint($st_C, hint, pp, U$) $\rightarrow (st'_C, hint')$

```

1: if  $pp[1] \neq pp[2] \vee U[1] \neq U[2]$  then // Verify update info
2:    $st_C.abort \leftarrow \text{true}$ 
3: if  $st_C.abort \vee st_C.q \neq \perp$  return  $(st_C, \perp)$ 
4:  $Q' \leftarrow st_C.pp.Q$ ,  $ck' \leftarrow st_C.ck$ 
5:  $(hint_1, \dots, hint_{pp[1].M}) \leftarrow hint$ 
6: for  $q \in [pp[1].Q - st_C.pp.Q]$  // Generate new permutation(s)
7:    $Q' = Q' + 1$ 
8:    $ck'.\sigma_q \leftarrow \text{Permute}(pp[1].M)$ 
9: for  $(op, idx, x) \in U[1]$  do // Update hints
10:   $\hat{q} \leftarrow \lceil idx/pp[1].M \rceil$ 
11:   $\hat{m} \leftarrow ((idx-1) \bmod pp[1].M) + 1$ 
12:   $ind \leftarrow ck'.\sigma_{\hat{q}}^{-1}(\hat{m})$ 
13:   $hint_{ind} \leftarrow hint_{ind} \oplus x$ 
14:  $st_C.pp \leftarrow pp[1]$ ,  $st_C.ck \leftarrow ck'$ 
15: return  $(st_C, (hint_1, \dots, hint_{pp[1].M}))$ 

```

Figure 6: UpdateDB and UpdateHint algorithm descriptions for TAPIR.

$st_C.abort = \text{true}$, the client sets $st_C.abort \leftarrow \text{true}$ and returns (st_C, \perp) . Since at least one server is semi-honest, equality guarantees correctness. Step (ii) generates permutations for any new

database partitions and appends them to the client key $st_C.ck$. Step (iii) updates the hint using the verified update information. For each $(op, idx, x) \in U_1$, the client uses its secret permutations $\{\sigma_q\}_{q \in [Q]}$ to locate the position in the hint. Let (\hat{q}, \hat{m}) be the partition and offset of idx . The client computes $ind \leftarrow \sigma_{\hat{q}}^{-1}(\hat{m})$ and sets $hint_{ind} \leftarrow hint_{ind} \oplus x$. Finally, the client outputs the updated state st_C (including the new public parameters) and the updated hint.

5.3 Scheme Correctness and Security

TAPIR requires $k = 2$ servers, one is assumed to be semi-honest. Both servers hold a copy of the same public database. Depending on the VC scheme used to instantiate our APIR scheme, we may also require that the public parameters are honestly generated (e.g., through a trusted setup, as in Pointproofs [GRWZ20]).

Theorem 5.3.1. Let VC be a vector commitment scheme with binding such that VC.Commit is deterministic. The updateable two-server APIR scheme with preprocessing described in Figs. 4, 5, and 6 satisfies correctness (Def. 4).

Theorem 5.3.2. Let $\lambda \in \mathbb{N}$ be the security parameter, $N_{init}, N = \text{poly}(\lambda)$, and VC be a secure vector commitment scheme with binding such that VC.Commit is deterministic. If both servers are non-colluding, then the updateable two-server APIR scheme with preprocessing in Figs. 4, 5, and 6 satisfies integrity (Def. 5).

Theorem 5.3.3. Let $\lambda \in \mathbb{N}$ be the security parameter, $N_{init}, N = \text{poly}(\lambda)$, and VC be a secure vector commitment scheme with binding such that VC.Commit is deterministic. If both servers are non-colluding and integrity holds, then the updateable two-server APIR scheme with preprocessing in Figs. 4, 5, and 6 satisfies privacy with abort (Def. 6).

The proofs are in Appendices B.1, B.2, and B.3, respectively.

6 Evaluation

This section evaluates the performance of TAPIR and compares it to related work. We implement our scheme with both Pointproofs [GRWZ20] and Merkle trees as the vector commitment (VC) scheme (we refer to these implementations as TAPIR-PP and TAPIR-MT, respectively). We compare our approach against the two-server APIR schemes of [CNC⁺23], specifically the DPF-based APIR scheme (denoted as APIR-DPF) and the linear APIR scheme implemented with both Pointproofs and Merkle trees (which we denote as APIR-Matrix-PP and APIR-Matrix-MT, respectively). Furthermore, we compare our scheme to SinglePass [LP24] to demonstrate the overhead incurred by making it maliciously secure.

We aim to answer the following questions:

1. How much communication and computation overhead does TAPIR have over its base protocol SinglePass? (§6.2)
2. How does TAPIR compare to related multi-server APIR schemes, especially in the online phase? (§6.3)
3. What are the amortized communication and computation costs of batch updates? (§6.4)

N	(A)PIR	Offline		Online	
		BW [kiB]	RT [s] (1-Time)	RT [s] (Per-Client)	BW [kiB] RT [s]
2 ¹⁸	APIR-DPF	-	-	-	1.43 0.04
	APIR-Matrix-MT	0.10	0.58	0.00	49.68 0.01
	TAPIR-MT	18 982.05	0.33	0.08	354.00 0.01
	TAPIR-PP	19 025.02	13 162.05	0.10	34.66 0.99
	SinglePass	18.52	0.00	0.01	34.50 0.00
2 ²⁰	APIR-DPF	-	-	-	1.56 0.17
	APIR-Matrix-MT	0.10	2.74	0.00	104.36 0.04
	TAPIR-MT	75 852.05	1.45	0.39	776.98 0.02
	TAPIR-PP	75 938.00	105 304.90	0.46	69.65 2.63
	SinglePass	37.02	0.00	0.05	69.50 0.00
2 ²²	APIR-DPF	-	-	-	1.69 0.72
	APIR-Matrix-MT	0.10	11.54	0.00	218.26 0.16
	TAPIR-MT	303 256.05	5.51	1.86	1 690.98 0.05
	TAPIR-PP	303 427.92	844 441.96	1.63	139.66 8.79
	SinglePass	74.03	0.00	0.24	139.51 0.00
2 ²⁴	APIR-DPF	-	-	-	1.83 2.97
	APIR-Matrix-MT	0.10	47.73	0.00	454.87 0.73
	TAPIR-MT	1 212 720.05	22.14	7.65	3 662.97 0.12
	TAPIR-PP	148.03	0.00	0.96	279.51 0.01
	SinglePass	148.03	0.00	0.96	279.51 0.01

Table 1: Offline and online performance comparison of (A)PIR schemes for record size 32 B. Offline costs are split into one-time server costs and per-client costs. Best APIR performance for each database size N is highlighted. Results for APIR-Matrix-PP and TAPIR-PP ($N > 2^{22}$) are omitted due to high Pointproofs setup costs.

Our protocol is implemented in Go, utilizing and adapting existing libraries for Merkle trees¹, and Pointproofs². For our benchmarks, we extended the (A)PIR implementations of related work, namely, SinglePass³ and the linear matrix PIR⁴, and implemented the APIR-DPF scheme⁵.

For TAPIR-PP, we reduce per-query bandwidth and verification cost by aggregating Pointproofs across database partitions, thereby lowering both online bandwidth and runtime. A single Pointproofs proof requires two exponentiations and two pairings, whereas aggregating across ℓ vectors requires only $1 + \ell$ of each—saving one exponentiation and one pairing per vector. Our implementation performs aggregation in Answer, shifting part of the verification cost to the server and allowing it to send a single proof instead of Q . In contrast, APIR-Matrix-PP sends only one proof per query and thus does not benefit from aggregation.

Our source code is available at <https://github.com/laurahetz/TAPIR>.

6.1 Experimental Setup

All benchmarks are executed on a single machine with two Intel Xeon Gold 6258R CPU 2.7 GHz, each with 28 cores, and 384 GB of DDR4 memory.

For TAPIR and related work, we use initial database sizes of $N \in \{2^{10}, 2^{12}, 2^{14}, 2^{16}, 2^{18}, 2^{20}, 2^{22}, 2^{24}\}$ database record 32 B, and where applicable Q database partitions of size $M = \sqrt{N}$.

¹<https://github.com/dedis/apir-code>

²<https://github.com/yacovm/Pol/tree/main/pp>

³<https://github.com/SinglePass712/Submission>

⁴<https://github.com/dimakogan/checklist>

⁵<https://github.com/osu-crypto/libOTE>




Figure 7: Offline per-client performance comparison between our authenticated scheme TAPIR and the unauthenticated base scheme SinglePass. Results shown for databases with N records of size 32 B. All axes use logarithmic scaling. We denote the schemes as follows: TAPIR-MT (• - -), TAPIR-PP (• - -), and SinglePass (▲ - -).

Due to the higher one-time runtime cost of Pointproofs in the offline phase, we were unable to run APIR-Matrix-PP and TAPIR-PP for databases larger than 2^{16} and 2^{22} , respectively. All measurements of the online phase are averaged over 25 runs.

6.2 Comparison to SinglePass

Since TAPIR modifies and extends SinglePass to achieve malicious security, we expect TAPIR to have higher runtime and communication cost in all protocol phases. In the offline phase, the increase in cost stems from the servers’ digest generation and the need to stream the database from both servers to ensure the integrity of the preprocessed information. In the online phase, the increase in bandwidth stems from the retrieved records’ proofs that are sent together with the records, and the increase in runtime results from the client having to verify these proofs.

Table 1 reports the benchmarking results for the total offline and online costs for all considered (A)PIR schemes. Figure 7 highlights the per-client bandwidth and runtime of TAPIR compared to SinglePass. This excludes the digest generation, as it is a one-time setup cost and is reusable for all clients.

As expected, TAPIR requires significantly more bandwidth overhead in the offline phase due to the digest and the client’s need to authenticate the database records included in the hint. The client-dependent offline runtime of TAPIR does not depend on the selected vector commitment scheme and is thus the same for TAPIR-MT and TAPIR-PP. Compared to SinglePass, this offline runtime is up to $8\times$ higher for $N = 2^{24}$ due to the setup verification.

Figure 8 shows the online costs of both SinglePass and TAPIR as compared to three related APIR schemes. Once again, the experimental results are consistent with our expectations.

Our scheme demonstrates very modest bandwidth increases in the online phase with respect to SinglePass, with as little as 0.11% overhead for TAPIR-PP for $N = 2^{22}$ and at most $13.11\times$ overhead for TAPIR-MT for $N = 2^{24}$. We see that TAPIR-PP significantly outperforms TAPIR-MT as Pointproofs allow optimizing bandwidth and verification cost via proof aggregation. This optimization results in communication costs comparable to SinglePass. While query costs are the same across all three schemes, the additional overhead of our scheme comes from the proofs returned as part of the

answer. A Pointproofs proof is compact and only comprises a single group element, whereas each Merkle tree proof comprises an authentication path of size $O(\log_2 M)$.

Both TAPIR variants show higher online runtimes than SinglePass, with TAPIR-MT outperforming TAPIR-PP. The higher online cost of TAPIR-PP comes from generating and verifying the aggregated proof using group exponentiations and a pairing product. In contrast, proof verification of TAPIR-MT only requires $Q \log_2 M$ hash evaluations by the client, which are cheap.

The online runtime cost of TAPIR-MT (for $N = 2^{24}$) is overall $22.64\times$ higher than SinglePass, with only $1.37\times$ overhead for the server and $54.15\times$ overhead for the client. The online runtime cost of TAPIR-PP (for $N = 2^{22}$) is overall 5.8 magnitudes higher than SinglePass. The majority of the overhead comes from the Pointproofs aggregation in the Answer algorithm and verification in the Recon algorithm, as observed in Fig. 8. We discuss the trade-off between vector commitment schemes in more detail in §7.

6.3 Comparison to Related Work

This section compares the benchmarking results of TAPIR (TAPIR-MT and TAPIR-PP) to prior multi-server APIR schemes, namely APIR-DPF, APIR-Matrix-MT [CNC⁺23], and APIR-Matrix-PP.

Using a linear PIR scheme with a vector commitment scheme, as done in APIR-Matrix-MT and APIR-Matrix-PP, requires each server to generate a digest, which the client verifies as part of the offline phase. TAPIR also relies on a vector commitment scheme for integrity in the online phase and thus also incurs the cost of digest generation and communication, as well as proof generation, verification, and communication. Since TAPIR executes the vector commitment schemes over each of the Q database partitions instead of the entire database (as done in APIR-Matrix-MT and APIR-Matrix-PP), the digest generation in TAPIR is expected to be faster in comparison. In addition, in TAPIR the client and servers preprocess the database to ensure integrity of the client’s retrieved hint. We expect this hint computation and verification to cause significant overhead in the offline phase of TAPIR. As APIR-DPF does not require any preprocessing, we expect all other considered APIR schemes to have a higher offline overhead. APIR-Matrix-MT and APIR-Matrix-PP overall have very low client offline overhead as only the digests of the servers need to be checked for equality. In the online phase, we expect TAPIR-MT to have lower runtimes than all other APIR schemes for larger database sizes, facilitated by the preprocessing and fast verification of proofs.

Table 1 reports the benchmarking results, and Fig. 8 visualizes the online costs of the benchmarked schemes.

With respect to online bandwidth, we see that the schemes instantiated with Merkle trees (TAPIR-MT, APIR-Matrix-MT) are more expensive than their Pointproofs counterparts due to the logarithmically sized proofs of the Merkle trees. We note that communication in the APIR-MATRIX schemes is lower as only secret shares of one record and its proof are communicated, while Q records and proofs are required in TAPIR. APIR-DPF has the lowest bandwidth of all considered (A)PIR schemes, at the cost of server runtime linear in the size of the database.

Regarding client runtime, APIR-DPF is even faster than the unauthenticated SinglePass while APIR-Matrix-MT and APIR-Matrix-PP




Figure 8: Online performance comparison of our scheme TAPIR and prior work. Results shown for databases with N records of size 32 B. All axes use logarithmic scaling.

have similar cost, and TAPIR incurs the highest overhead. This is expected, as the Recon algorithm of TAPIR is more complex and computationally intensive compared to the cheap XOR operations in the APIR-Matrix schemes. Regarding server runtime, TAPIR-MT outperforms all APIR schemes with only $1.37\times$ overhead compared to its unauthenticated base scheme. As expected, we observe a linear trend for the server runtime in APIR-DPF and the APIR-Matrix schemes, whereas TAPIR has sublinear runtime in the database size.

For $N \geq 2^{20}$, TAPIR-MT is the fastest APIR scheme online, outperforming APIR-Matrix-MT and APIR-DPF by a factor of up to $5.83\times$ and $23.82\times$, respectively.

6.4 Database Updates

This section evaluates the practicality of database updates in TAPIR. For this, we measured the bandwidth and runtime cost of applying a batch of 500 update operations to unique database indices for databases with N 32 B records. We differentiate between batches of only additions (ADD), only edits (EDIT), and BOTH additions and edits, and report the amortized costs in Fig. 9.

We expect edits and additions within existing partitions to be equally fast, as they are handled the same way. Client runtime and bandwidth are expected to be independent of the chosen VC scheme. In contrast, we expect the server runtime to be higher for Pointproofs than for Merkle trees, due to the higher cost of generating and updating commitments and proofs. We expect bandwidth costs to be similar over all update types, with the exception that new partitions require communicating new digests. These only occur with ADD and BOTH, while BOTH has, on average, only half the additions of ADD, and thus will incur less bandwidth.

We observe that the client runtime is very fast, as updates only require lightweight operations such as generating new permutations and performing basic arithmetic (e.g., XOR, mod, division) to update the hint. For instance, the largest observed amortized client runtime is just 0.005 ms at $N = 2^{22}$ for TAPIR-MT. TAPIR-PP incurs slightly higher cost than TAPIR-MT, due to the equality check of the public parameters containing bilinear group elements as opposed to hash function outputs.

The server runtime is higher but remains practical, especially since update computations are client-independent. Amortized runtimes range from 0.005 ms ($N = 2^{10}$, TAPIR-MT) to 370.5 ms ($N = 2^{14}$, TAPIR-PP). TAPIR-PP is generally slower due to the costly

bilinear pairings required for Pointproofs commitments. EDIT operations also tend to be more expensive than ADD, as each of the 500 edits may require updating commitments for up to 500 existing partitions, whereas additions only require initializing and committing to a few new partitions.

Bandwidth costs are modest, ranging from 83.48 B ($N = 2^{10}$, TAPIR-MT) to 86.01 B ($N = 2^{22}$, TAPIR-MT). The bandwidth of TAPIR-MT and TAPIR-PP is comparable across ADD, EDIT, and BOTH operations, since the client only receives the update tuples (identical for both schemes) and the new public parameters with the updated digest. The digests differ slightly in form—TAPIR-MT uses the Merkle tree root and TAPIR-PP bilinear group element—but both are small and similar in size. Among the three operations, ADD incurs the highest bandwidth due to the transmission of new vector commitments for added partitions. EDIT incurs the least since no new partitions are created, and BOTH falls in between as half of the operations are additions on average. A noticeable bandwidth gap emerges around $N = 2^{16}$ due to partitioning; with $M = 2^8 = 256$, ADD requires approximately two new partitions per batch, EDIT requires none, and BOTH requires about one.

6.5 Application: Key Transparency

Systems for Key Transparency (KT) are widely deployed by companies such as Google [Goo24], WhatsApp [LL23], Proton [GH24], and Keybase [key22] to enable users to verify public keys for end-to-end encrypted communication, and are even undergoing standardization [Gro25]. Typically, KT schemes use an authenticated dictionary hosted by an untrusted server that maps user IDs to public keys. A digest of the dictionary is computed, allowing clients to perform verifiable lookups.

Because KT requires integrity for public key lookups, it presents a promising application for APIR. For example, KT can be instantiated by extending index-based APIR to support keyword search (see [CD24, HLP⁺25] and § 7). Alternatively, Index-APIR can be used to index public keys directly; clients can then obtain the index of a public key for lookups (e.g., via gossiping, an auditor, or a previous lookup). Standard deployed KT systems often overlook a critical aspect: query privacy for users requesting another user’s key. APIR fills this gap by providing query privacy with abort even against adversarial servers.

Both WhatsApp’s [Met23] and Proton’s [GH24] white papers state that their KT systems use Curve25519 as the Public Identity




Figure 9: Amortized update costs for TAPIR with Merkle trees (MT) and Pointproofs (PP), comparing ADD, EDIT, and BOTH operations. Results are for databases with N records of 32 B sizes and update batches of size 500. Axes use logarithmic scaling.

Key, meaning keys stored in the dictionary are 32 B long. In our experiments (see, e.g., Fig. 8), we demonstrate the performance of our schemes TAPIR-MT and TAPIR-PP on 32 B records to reflect this use case. In particular, TAPIR-MT scaled easily to $N = 2^{24}$ and can be extended to larger databases common in KT by leveraging database partitioning as described in § 7. Moreover, the efficient client and server runtimes, along with the small bandwidth required for updates of a batch size of 500 (as discussed in § 6.4), further bolster the practicality of our scheme for transparency.

7 Conclusion

Optimizations. The complexity of TAPIR grows with the database size N . To improve scalability, the database can be partitioned into p sub-databases [HSW23], with TAPIR running on each in parallel. This adds minor overhead but reduces per-database cost and enables trivial parallelization of both phases. We leave this for future work.

Support for Keywords. While (A)PIR usually assumes knowledge of a record’s index, many applications rely on keywords. Our scheme can be extended in a black-box way to support keyword queries [CD24, HLP⁺25]. This incurs extra bandwidth and computation, but allows retrieval based on labels rather than indices.

Trade-offs and Use Cases. TAPIR supports a flexible runtime and bandwidth trade-off (Fig. 5 and Tab. 1). For example, using Merkle trees yields the smallest runtimes among multi-server APIR schemes, while using Pointproofs [GRWZ20] matches the online bandwidth of SinglePass. In contrast, linear schemes like [CNC⁺23] show little variation across different VCs. Thus, bandwidth-sensitive applications with frequent queries (e.g., key transparency, oblivious message detection) benefit from TAPIR-PP, while infrequent-query settings (e.g., contact discovery, medical look-up) favor TAPIR-MT.

Acknowledgments

The authors thank Tianxin Tang, Christian Weinert, Matilda Backendal for their helpful feedback and input. The authors would also like to thank Shannon Veitch for her suggestion to adopt a game-based security definition. Francesca Falzon is supported by Armasuisse Science and Technology.

References

- [AB25] Bar Alon and Amos Beimel. On the definition of malicious private information retrieval. In *The sixth Information-Theoretic Cryptography (ITC)*, 2025.

- [ABG⁺24] Hilal Asi, Fabian Boemer, Nicholas Genise, Muhammad Haris Mughees, Tabitha Ogilvie, Rehan Rishi, Guy N. Rothblum, Kunal Talwar, Karl Tarbe, Ruiyu Zhu, and Marco Zuliani. Scalable Private Search with Wally. *CoRR*, abs/2406.06761, 2024.
- [AR25] Pranav Shriram Arunachalam and Ling Ren. Single-server stateful PIR with verifiability and balanced efficiency. *Cryptology ePrint Archive*, Paper 2025/1055, 2025.
- [AS16] Sebastian Angel and Srinath T. V. Setty. Unobservable communication over fully untrusted infrastructure. In *OSDI*, pages 551–569. USENIX Association, 2016.
- [BGH⁺25] Elette Boyle, Niv Gilboa, Matan Hamilis, Yuval Ishai, and Yaxin Tu. Improved Constructions for Distributed Multi-Point Functions. In *2025 IEEE Symposium on Security and Privacy*, pages 2414–2432. IEEE Computer Society Press, May 2025.
- [BGI15] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 337–367. Springer, Berlin, Heidelberg, April 2015.
- [BGI16] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1292–1303. ACM Press, October 2016.
- [CD24] Sofia Celi and Alex Davidson. Call me by my name: Simple, practical private information retrieval for keyword queries. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *ACM CCS 2024*, pages 4107–4121. ACM Press, October 2024.
- [CDGM19] Melissa Chase, Apoorva Deshpande, Esha Ghosh, and Harjasleen Malvai. SEEMless: Secure end-to-end encrypted messaging with less trust. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 1639–1656. ACM Press, November 2019.
- [CF13] Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 55–72. Springer, Berlin, Heidelberg, February / March 2013.
- [CGKS95] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *36th FOCS*, pages 41–50. IEEE Computer Society Press, October 1995.
- [CHK22] Henry Corrigan-Gibbs, Alexandra Henzinger, and Dmitry Kogan. Single-server private information retrieval with sublinear amortized time. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 3–33. Springer, Cham, May / June 2022.
- [CK20] Henry Corrigan-Gibbs and Dmitry Kogan. Private information retrieval with sublinear online time. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 44–75. Springer, Cham, May 2020.
- [CNC⁺23] Simone Colombo, Kirill Nikitin, Henry Corrigan-Gibbs, David J. Wu, and Bryan Ford. Authenticated private information retrieval. In Joseph A. Calandrino and Carmela Troncoso, editors, *USENIX Security 2023*, pages 3835–3851. USENIX Association, August 2023.
- [CSM⁺20] Raymond Cheng, William Scott, Elisaweta Masserova, Irene Zhang, Vipul Goyal, Thomas E. Anderson, Arvind Krishnamurthy, and Bryan Parno. Talek: Private group messaging with hidden access patterns. In *ACM CCS*, pages 84–99. ACM, 2020.
- [dCL24] Leo de Castro and Keewoo Lee. VeriSimplePIR: Verifiability in SimplePIR at no online cost for honest servers. In Davide Balzarotti and Wenyan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.

- [DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 643–662. Springer, Berlin, Heidelberg, August 2012.
- [DT24] Marian Dietz and Stefano Tessaro. Fully malicious authenticated PIR. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part IX*, volume 14928 of *LNCS*, pages 113–147. Springer, Cham, August 2024.
- [Dur64] Richard Durstenfeld. Algorithm 235: Random permutation. *Commun. ACM*, 7(7):420, July 1964.
- [EKN22] Reo Eriguchi, Kaoru Kurosawa, and Koji Nuida. On the optimal communication complexity of error-correcting multi-server PIR. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part III*, volume 13749 of *LNCS*, pages 60–88. Springer, Cham, November 2022.
- [FMS24] Brett Falk, Pratyush Mishra, and Matan Shtepele. Malicious security for PIR (almost) for free. *Cryptology ePrint Archive*, Report 2024/964, 2024.
- [FMS25] Brett Hemenway Falk, Pratyush Mishra, and Matan Shtepele. Malicious security for PIR (almost) for free. In *CRYPTO 2025, Part I*, LNCS. Springer, Cham, August 2025.
- [FY53] Ronald Aylmer Fisher and Frank Yates. *Statistical tables for biological, agricultural and medical research*. Hafner Publishing Company, 1953.
- [GH24] Thore Göbel and Daniel Huigens. Proton key transparency whitepaper. https://proton.me/files/proton_keytransparency_whitepaper.pdf, 2024.
- [GI14] Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 640–658. Springer, Berlin, Heidelberg, May 2014.
- [Goo24] Inc. Google. Key transparency. <https://github.com/google/keytransparency>, 2024. Accessed: 12/08/2025.
- [Gro25] KT Internet Engineering Task Force Working Group. Key transparency (keytrans). <https://datatracker.ietf.org/wg/keytrans/about/>, 2025.
- [GRWZ20] Sergey Gorbunov, Leonid Reyzin, Hoeteck Wee, and Zhenfei Zhang. Pointproofs: Aggregating proofs for multiple vector commitments. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 2007–2023. ACM Press, November 2020.
- [GZS24] Ashrujit Ghoshal, Mingxun Zhou, and Elaine Shi. Efficient pre-processing PIR without public-key cryptography. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 210–240. Springer, Cham, May 2024.
- [GZSP25] Ashrujit Ghoshal, Mingxun Zhou, Elaine Shi, and Bo Peng. Pseudorandom functions with weak programming privacy and applications to private information retrieval. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 284–313. Springer, Cham, May 2025.
- [HDCZ23] Alexandra Henzinger, Emma Dauterman, Henry Corrigan-Gibbs, and Nikolai Zeldovich. Private Web Search with Tiptoe. In *SOSP*, pages 396–416. ACM, 2023.
- [HHC+23] Alexandra Henzinger, Matthew M. Hong, Henry Corrigan-Gibbs, Sarah Meiklejohn, and Vinod Vaikuntanathan. One server for the price of two: Simple and fast single-server private information retrieval. In Joseph A. Calandrino and Carmela Troncoso, editors, *USENIX Security 2023*, pages 3889–3905. USENIX Association, August 2023.
- [HHK+21] Yuncong Hu, Kian Hooshmand, Harika Kalidhindi, Seung Jin Yang, and Raluca Ada Popa. Merkle²: A low-latency transparency log system. In *2021 IEEE Symposium on Security and Privacy*, pages 285–303. IEEE Computer Society Press, May 2021.
- [HLP+25] Meng Hao, Weiran Liu, Liqiang Peng, Cong Zhang, Pengfei Wu, Lei Zhang, Hongwei Li, and Robert H. Deng. Practical keyword private information retrieval from key-to-index mappings. In *USENIX Security 2025*. USENIX Association, August 2025.
- [HSW23] Laura Hetz, Thomas Schneider, and Christian Weinert. Scaling mobile private contact discovery to billions of users. In Gene Tsudik, Mauro Conti, Kaitai Liang, and Georgios Smaragdakis, editors, *ESORICS 2023, Part I*, volume 14344 of *LNCS*, pages 455–476. Springer, Cham, September 2023.
- [KC21] Dmitry Kogan and Henry Corrigan-Gibbs. Private blacklist lookups with checklist. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 875–892. USENIX Association, August 2021.
- [key22] keybase.io. Keybase chat. book.keybase.io/docs/chat, 2022. Accessed: 12/08/2025.
- [KKEPR24] Erki Külaots, Toomas Krips, Hendrik Eerikson, and Pille Pullonen-Raudvere. SLAMP-FSS: Two-party multi-point function secret sharing from simple linear algebra. *Cryptology ePrint Archive*, Report 2024/1394, 2024.
- [Knu97] Donald Ervin Knuth. *The art of computer programming*, volume 3. Pearson Education, 1997.
- [KO97] Eyal Kushilevitz and Rafail Ostrovsky. Replication is NOT needed: SINGLE database, computationally-private information retrieval. In *38th FOCS*, pages 364–373. IEEE Computer Society Press, October 1997.
- [LL23] Kevin Lewi and Sean Lawlor. Introducing auditable key transparency for end-to-end encrypted messaging. <https://engineering.fb.com/2023/04/13/security/whatsapp-key-transparency/>, April 2023. Meta Engineering Blog.
- [LP23] Arthur Lazzaretti and Charalampos Papamanthou. TreePIR: Sublinear-time and polylog-bandwidth private information retrieval from DDH. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 284–314. Springer, Cham, August 2023.
- [LP24] Arthur Lazzaretti and Charalampos Papamanthou. Single pass client-preprocessing private information retrieval. In Davide Balzarotti and Wenyan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.
- [LTW24] Zeyu Liu, Eran Tromer, and Yunhao Wang. PerfOMR: Oblivious message retrieval with reduced communication and computation. In Davide Balzarotti and Wenyan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.
- [MBB+15] Marcela S. Melara, Aaron Blankstein, Joseph Bonneau, Edward W. Felten, and Michael J. Freedman. CONIKS: Bringing key transparency to end users. In Jaeyeon Jung and Thorsten Holz, editors, *USENIX Security 2015*, pages 383–398. USENIX Association, August 2015.
- [Met23] Meta. Whatsapp key transparency overview. <https://www.whatsapp.com/security/WhatsApp-Key-Transparency-Whitepaper.pdf>, 2023.
- [MKKS+23] Harjasleen Malvai, Lefteris Kokoris-Kogias, Alberto Sonnino, Esha Ghosh, Ercan Oztürk, Kevin Lewi, and Sean F. Lawlor. Parakeet: Practical key transparency for end-to-end encrypted messaging. In *NDSS 2023*. The Internet Society, February 2023.
- [Res24] Apple Machine Learning Research. Combining machine learning and homomorphic encryption in the apple ecosystem. <https://machinelearning.apple.com/research/homomorphic-encryption>, oct 2024. Accessed: 2025-07-14.
- [TBP+19] Alin Tomescu, Vivek Bhupatiraju, Dimitrios Papadopoulos, Charalampos Papamanthou, Nikos Triandopoulos, and Srinivas Devadas. Transparency logs via append-only authenticated dictionaries. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 1299–1316. ACM Press, November 2019.
- [WLZ+23] Yinghao Wang, Xuanming Liu, Jiawen Zhang, Jian Liu, and Xiaohu Yang. Crust: Verifiable and efficient private information retrieval with sublinear online time. *Cryptology ePrint Archive*, Report 2023/1607, 2023.
- [ZPZS24] Mingxun Zhou, Andrew Park, Wenting Zheng, and Elaine Shi. Piano: Extremely simple, single-server PIR with sublinear server computation. In *2024 IEEE Symposium on Security and Privacy*, pages 4296–4314. IEEE Computer Society Press, May 2024.

A Additional Preliminaries

A.1 Show-and-Shuffle

In SinglePass [LP24], the Show-and-Shuffle game $G^{\text{Ind-SaS}}$ was introduced to capture the main steps of a single protocol round and to prove its privacy. Since TAPIR extends SinglePass to achieve malicious security, we can leverage this game to also prove privacy for our scheme. Hence, we restate the Show-and-Shuffle game in Fig. 10 and the Show-and-Shuffle Indistinguishability lemma in Theorem A.1.1, and refer to [LP24, § 3.1] for the full proof. The game is parameterized over $M, Q \in \mathbb{N}$.

Lemma A.1.1. For the Show-and-Shuffle game $G^{\text{Ind-SaS}}$ (Fig. 10), and any $M, Q \in \mathbb{N}$, the advantage of any adversary playing this game is defined as:

$$\text{Adv}_{\mathcal{A}}^{\text{ind-sas}}(M, Q) = \left| \Pr \left[G^{\text{Ind-SaS}}(M, Q) \right] - \frac{1}{2} \right| = 0.$$

B Additional Proof Content

We now give the proofs for Theorem 5.3.1 (Correctness), Theorem 5.3.2 (Integrity), and Theorem 5.3.3 (Privacy).

Game $G^{\text{Ind-SaS}}(M, Q)$

- 1 : $(\sigma_1, \dots, \sigma_Q) \leftarrow \text{Permute}(M)^Q$
- 2 : $(\text{st}_{\mathcal{A}}, x) \leftarrow \mathcal{A}_0(M, Q)$ where $x = (q^*, m^*) \in ([Q] \times [M])$
- 3 : Find $\text{ind} \in [M]$, s.t. $\sigma_{q^*}(\text{ind}) = m^*$
- 4 : $v \leftarrow \perp$
- 5 : **for** $q \in [Q]$ **do**
- 6 : **if** $q \neq q^*$ **then** $v.v_q \leftarrow \sigma_q(\text{ind})$
- 7 : **else** $v.v_q \leftarrow \text{st}_{\mathcal{A}}[q]$
- 8 : $b \leftarrow \{0, 1\}$
- 9 : $\mathcal{R}_0 = (F_1, \dots, F_Q) \leftarrow \text{Permute}(M)^Q$
- 10 : **for** $q \in [Q] \wedge q \neq q^*$ **do**
- 11 : $r_q \leftarrow \text{st}_{\mathcal{A}}[q]$
- 12 : $\sigma'_q \leftarrow \sigma_q$
- 13 : Swap $\sigma'_q(\text{ind})$ and $\sigma'_q(r_q)$
- 14 : $\mathcal{R}_1 \leftarrow (\sigma'_1, \dots, \sigma'_{q^*-1}, \sigma_{q^*}, \sigma'_{q^*+1}, \dots, \sigma'_Q)$
- 15 : $b' \leftarrow \mathcal{A}_1(\text{st}_{\mathcal{A}}, q, m, v, \mathcal{R}_b)$
- 16 : **return** $b = b'$

Figure 10: The Show-and-Shuffle game of SinglePass.

B.1 Proof of Correctness (Theorem 5.3.1)

PROOF. We note that correctness is only defined for honest servers. Let DB_{init} be the initial database records generated by \mathcal{A}_0 . Let $\sigma_q : [\text{pp}.M] \rightarrow [\text{pp}.M]$ denote the permutation for partition $q \in [\text{pp}.Q_{\text{init}}]$. By correctness of the vector commitment scheme VC and since VC.Commit is assumed to be deterministic, for each $q \in [\text{pp}.Q_{\text{init}}]$ we have

$$\text{com}_{1,q} = \text{com}_{2,q} = \text{VC.Commit}(\text{DB}_q).$$

In the offline phase, the digests satisfy

$$d_1 = (\text{com}_{1,1}, \dots, \text{com}_{1,Q_{\text{init}}}) = (\text{com}_{2,1}, \dots, \text{com}_{2,Q_{\text{init}}}) = d_2.$$

Since the remaining public parameters are derived directly from the setup parameters provided to Setup, both servers output identical values. We thus have $\text{pp}[1] = \text{pp}[2]$ (Fig. 4, line 3), and RequestHint completes successfully.

Lookup correctness. We now prove that the initial lookup query succeeds, and then proceed by induction over the sequence of T_q lookups. Since Server b sets $\text{st}_{S_b}.\text{DB}[[N_{\text{init}}]] \leftarrow \text{DB}_{\text{init}}$ (Fig. 4, line 3) and, by correctness of the streaming procedure (Fig. 4, line 3), we obtain $\text{DB}_1 = \text{DB}_2$. It follows that $\text{resp}_1 = \text{resp}_2$ (Fig. 4, line 15). Thus, the VerSetup algorithm computes a hint (Fig. 4, line 12) of the form $\text{hint} = (h_1, \dots, h_M)$ where, for each $m \in [M]$:

$$h_m = \bigoplus_{q \in [Q_{\text{init}}]} \text{DB}[(q-1) \cdot M + \sigma_q(m)] = \bigoplus_{q \in [Q_{\text{init}}]} \text{DB}_q[\sigma_q(m)].$$

Let $\text{idx} \in [N]$ be any index queried for by the client and $(\hat{q}, \hat{m}) \in [Q_{\text{init}}] \times [M]$ be such that $\text{DB}[\text{idx}] = \text{DB}_{\hat{q}}[\hat{m}]$. Let $\text{ind} \in [M]$ be such that $\sigma_{\hat{q}}(\text{ind}) = \hat{m}$.

Consider that Server 2 responds to the first query honestly. Applying the above hint equation to line 17 (Fig. 5) yields

$$\begin{aligned} x &= \bigoplus_{q \in [Q_{\text{init}}] \setminus \{\hat{q}\}} \text{rec}_{2,q} \oplus h_{\text{ind}} \\ &= \left(\bigoplus_{q \in [Q_{\text{init}}] \setminus \{\hat{q}\}} \text{DB}_q[\sigma_q(\text{ind})] \right) \oplus h_{\text{ind}} \\ &= \left(\bigoplus_{\substack{q \in [Q_{\text{init}}] \\ q \neq \hat{q}}} \text{DB}_q[\sigma_q(\text{ind})] \right) \oplus \left(\bigoplus_{q \in [Q_{\text{init}}]} \text{DB}_q[\sigma_q(\text{ind})] \right) \\ &= \text{DB}_{\hat{q}}[\sigma_{\hat{q}}(\text{ind})] = \text{DB}_{\hat{q}}[\hat{m}] = \text{DB}[\text{idx}]. \end{aligned}$$

By correctness of VC, and since the servers are assumed to be honest and return the correct record corresponding to qry_b , it follows that for every $q \in [Q_{\text{init}}]$ and every $b \in \{1, 2\}$,

$$1 = \text{VC.Verify}(\text{com}_q, \text{qry}_{b,q}, \text{rec}_{b,q}, \pi_{b,q})$$

where $\text{ans}_{b,q} = (\text{rec}_{b,q} \parallel \pi_{b,q})$. Thus, the client reconstructs and outputs $x = \text{DB}[\text{idx}]$ and the win is set to **true** with probability 0.

We now show that subsequent lookup queries correctly recover the requested record and that the invariant, that the hint remains correct after every step, is preserved. In particular, we want to show that for every $m \in [M]$, the following invariant holds:

$$h_m = \bigoplus_{q \in [\text{pp}.Q]} \text{DB}[\sigma_q(m)]. \quad (1)$$

Assume that $\text{hint} = (h_1, \dots, h_M)$ is correct after responding to the k -th query and let $\text{pp}.Q$, $\text{pp}.M$, and $\text{pp}.N$ denote the current public parameters at this step.

We appeal to the argument of correctness in [LP24] and reproduce the details for completeness below. Note that for each swap between $\sigma_q(\text{ind})$ and $\sigma_q(r_q)$ the hint is modified such that

$$h_{\text{ind}} \leftarrow h_{\text{ind}} \oplus \text{DB}_q[\sigma_q(\text{ind})] \oplus \text{DB}_q[\sigma_q(r_q)].$$

This XOR removes the old element in the hint and adds the new one. Thus at the end of the $(k+1)$ -th query, $h_m = \bigoplus_{q \in [\text{pp}.Q]} \text{DB}_q[\sigma_q(m)]$ for all $m \in [\text{pp}.M]$. Since the hint is still of the correct form (Eq. 1), then, by an argument similar to the one above, the client must always recover the correct database element using the hint and the win is set to **true** with probability 0.

Update correctness. We have already shown that the hint is computed correctly at setup and after any number of T_q lookup queries. Our goal now is to prove that update queries preserve the invariant that the hint remains correct after executing UpdateDB and UpdateHint (thereby implying that lookups executed after an update are correct). Let $\text{pp}'.Q$, $\text{pp}'.M$, and $\text{pp}'.N$ denote the public parameters before the first update is made. Suppose that the hint satisfies Eq. 1 before an update is made and that for all $q \in [\text{pp}'.Q]$ and $m \in [\text{pp}.M]$, $\widehat{\text{DB}}_q[m]$ is a valid opening for $\text{DB}_q[m]$ under $\text{com}_q = \text{VC.Commit}(\text{DB}_q)$.

UpdateDB first initializes an empty set U' and the for loop on line 3 (Fig. 6) cycles through the set of updates $(\text{op}, \text{idx}, x) \in U$ and does the following: if $\text{op} = \text{add}$ then it increments $\text{pp}'.N$ and adds $(\text{add}, \text{pp}'.N, x)$ to U' , otherwise $\text{op} = \text{edit}$ and it adds $(\text{op}, \text{idx}, x)$ to U' . This step ensures that the public parameters are correct for the new database and that new values are assigned the correct index.

Using the new $pp'.N$, the server also extends the size of DB and \overline{DB} accordingly and initializes the new entries.

UpdateDB then processes updates by partition (Fig. 6, line 11). For $q \in [pp'.Q]$, let $U'_q \subseteq U'$ denote the subset of updates targeting partition q . For every $(op, idx, x) \in U'_q$ and the corresponding shift $m \leftarrow ((idx - 1) \bmod pp'.M) + 1$ the algorithm executes the following:

- If $op = \text{edit}$ (Fig. 6, line 18), the server updates $DB_q[m] \leftarrow x_{new}$ and sets $com_q \leftarrow VC.\text{Update}(com_q, m, x_{old}, x_{new})$.
- If DB_q is a new partition (Fig. 6, line 21), the server commits to this new partition, i.e., $com_q \leftarrow VC.\text{Commit}(DB_q)$.
- For every q and every $m \in [pp'.M]$ (Fig. 6, line 23), the server recomputes the opening as $\overline{DB}_q[m] \leftarrow VC.\text{Open}(m, DB_q)$.

Thus, by correctness of VC, for all q , com_q commits to DB_q and the opening proofs are all correct with respect to the new DB.

Now consider UpdateHint. With honest servers, we have $pp[1] = pp[2]$ and $U[1] = U[2]$, so no abort is triggered (i.e., Fig. 6, line 2 is skipped). UpdateHint then samples fresh permutations $\sigma_q \leftarrow \text{Permute}(M)$ for each new partition. Let U' denote the update information returned by the servers. For each $(op, idx, x) \in U'$, let (\hat{q}, \hat{m}) be the partition and offset of idx . We consider two cases:

- If $op = \text{edit}$, the servers must have set $x = x_{new} \oplus x_{old}$ where x_{new} and x_{old} are the new and old records at index idx respectively (Fig. 6, line 20). The client computes $ind = \sigma_{\hat{q}}^{-1}(\hat{m})$ and updates $h_{ind} \leftarrow h_{ind} \oplus x = h_{ind} \oplus x_{new} \oplus x_{old}$ thereby removing the old element in the hint and adding the new one.
- If $op = \text{add}$, then $x = x_{new} \oplus x_{old} = x_{new} \oplus 0^\ell = x_{new}$ since new database entries are initialized to 0. As in the previous case, the client computes $ind = \sigma_{\hat{q}}^{-1}(\hat{m})$ and updates $h_{ind} \leftarrow h_{ind} \oplus x = h_{ind} \oplus x_{new}$, thereby XOR-ing the new value x_{new} into h_{ind} .

In either case, the invariant for h_{ind} is preserved, and all other h_m are unaffected by the update and remain correct. The hint remains correct with respect to the updated database after the update is made. Thus, any later honest Answer returns the correct value, and thus win is set to **true** with probability 0. \square

B.2 Proof of Integrity (Theorem 5.3.2)

PROOF. Let $good \in \{1, 2\}$ denote the index of the honest server and, for ease of notation, let $bad = 3 - good$ denote the index of the malicious server. Let pp_{bad} and DB_{init} be the public parameters and the initial database records generated by \mathcal{A}_0 . Suppose for a contradiction that the adversary wins the integrity game, i.e., the client queries for an index $idx \in [N]$ and outputs a tuple $(st'_C, x, hint')$ such that $x \neq \perp$ and $x \neq DB[idx]$. In other words, the client reconstructs and accepts an incorrect value. We proceed with respect to the possible actions of $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2)$:

\mathcal{A}_0 generates malicious public parameters. Let d_{bad} denote the digest contained in pp_{bad} . If $pp_1 \neq pp_2$, then the client sets $st_C.abort \leftarrow \text{true}$ (Fig. 4, RequestHint, line 4). Consequently, when the client executes Recon, it always returns (st_C, \perp, \perp) (Fig. 5, Recon, line 1) and never proceeds with the remainder of the Recon algorithm. Thus, the probability that win = **true** is 0.

\mathcal{A}_1 generates a malicious response. Let DB and DB' be the databases streamed by Server 1 and Server 2, respectively. Since at least one server is honest, there exists an index $idx \in [st_C.pp.N]$

such that $DB[idx] \neq DB'[idx]$. This implies $\text{resp}[1] \neq \text{resp}[2]$. By the definition of the VerSetup algorithm, whenever $\text{resp}[1] \neq \text{resp}[2]$, the client sets $st_C.abort \leftarrow \text{true}$ (Fig. 4, VerSetup, line 3). As in the previous case, executing Recon then always outputs (st_C, \perp, \perp) (Fig. 5, Recon, line 1), so the probability that win is set to **true** is 0.

\mathcal{A}_2 generates malicious and valid update. Let pp_{bad} denote the updated public parameters and U_{bad} the update information provided by the malicious server. If $pp_1 \neq pp_2$ or $U_1 \neq U_2$, then by the definition of the UpdateHint algorithm the client sets $st_C.abort \leftarrow \text{true}$ (Fig. 6, UpdateHint, line 2). Consequently, any subsequent call to Recon will always return (st_C, \perp, \perp) (Fig. 5, Recon, line 1). Since one of the servers is honest, then the probability that the client accepts a malformed update and win is set to **true** is 0.

\mathcal{A}_2 generates a malicious answer. Parse the answer as $ans_{bad} = (ans_{bad,1}, \dots, ans_{bad,pp.Q})$ and the digest as $(com_1, \dots, com_{pp.Q}) \leftarrow pp.d$. Suppose that there exists some partition index $q \in [pp.Q]$ such that

$$VC.\text{Verify}(com_q, qry_{bad,q}, rec_{bad,q}, \pi_{bad,q}) = 1,$$

where $ans_{bad,q} = (rec_{bad,q} || \pi_{bad,q})$ and $rec_{bad,q} \neq DB_q[qry_{bad,q}]$. Because the commitment com_q is correct (due to the determinism of VC.Commit we must have equality of $pp_1 = pp_2$ for honestly generated parameters), this would violate the binding property of VC. The probability that this happens is $\text{negl}(\lambda)$. Taking a union bound over $Q = \text{poly}(\lambda)$ database partitions and $\text{poly}(\lambda)$ -many queries still yields at most negligible probability.

Hence, we reach a contradiction, completing the proof. \square

B.3 Proof of Privacy (Theorem 5.3.3)

In this section, we prove Theorem 5.3.3 with a sequence of game hops. As TAPIR extends SinglePass to achieve malicious security, this proof follows the privacy proof in [LP24] with additional consideration for integrity and updates.

For simplicity, we denote $[k] = \{1, 2\}$ with $\{0, 1\}$ in this proof. This allows us to reference the malicious server with $\neg\text{good}$.

Proof Sketch. The hint generation and online phase of TAPIR follows that of SinglePass, except for the addition of integrity checks and the client's abort status. In the privacy game, Query is run for one of the two adversarially chosen indices, and it outputs a sequence of database indices ($pp.Q$ values in $[pp.M]$) for each server. Each server answers the query by sending the database records and proofs at these indices back to the client. The client verifies the correctness of all the received records based on the corresponding proof and the database commitments stored in the client's state. The client aborts if a single verification fails. Hence, integrity (and thereby privacy with abort) is ensured by the deterministic commitment generation and binding of VC. To show that neither the set of indices in a query nor the abort status reveals any information to the server, we leverage the Show-and-Shuffle theorem introduced by [LP24] and restated in Fig. 10.

PROOF. Let \mathcal{A} be the adversary in the privacy game $G_{\text{APIR}}^{\text{Priv-APIR}}$ in Fig. 1. The adversary has the capabilities to corrupt an honest server defined by bit $good$ and to select any well-formed sequence of records DB_{init} and update information U for use in the game. We denote the maximum number of calls to the query and update

oracles by T_q and T_u . For all $N, T_q, T_u \in \text{poly}(\lambda)$ where $T_q + T_u \in \text{poly}(\lambda)$. Let $\Pr[G_x]$ denote the probability that game G_x outputs 1.

We now describe the sequence of games G_0 to G_6 for proving privacy with abort of TAPIR as stated in Theorem 5.3.3. See Figures 11 and 12 for the detailed pseudocode.

- **Game G_0** : This is equal to the privacy game $G_{\text{APIR}}^{\text{Priv-APIR}}$ (Fig. 1). The call to Recon has been replaced with the according pseudocode of this algorithm. A flag *BAD* is set to true if the adversary breaks binding of the VC scheme (G_0 , line 17).
- **Game G_1** : This game is equal to G_0 until *BAD* is set, so $\Pr[G_0] - \Pr[G_1] \leq \Pr[\text{BAD}]$ where $\Pr[\text{BAD}]$ denotes the probability of the flag *BAD* being set to **true**.

- **Bounding $\Pr[\text{BAD}]$** : Let G'_1 be identical to G_1 , except that it returns *BAD* instead of $b = b_{\text{chall}}$ such that $\Pr[\text{BAD}] = \Pr[G'_1]$. Let G''_1 be identical to G'_1 , except that it additionally keeps track of the adversary's oracle inputs and outputs to Query, Recon, and Update as tuples in the sets S_q, S_r, S_u . Hence, $\Pr[\text{BAD}] = \Pr[G'_1] = \Pr[G''_1]$.

Let \mathcal{B} be an adversary in the integrity game $G_{\text{APIR}}^{\text{Int-APIR}}$. This adversary \mathcal{B} simulates G'_1 for \mathcal{A} , where it samples a challenger bit b_{chall} , honestly executes the steps in the game, and stores the oracle inputs and outputs in the respective variables. Once \mathcal{A} halts, \mathcal{B} executes the same oracle calls in their game as \mathcal{A} does in G''_1 . For calls to oracle QueryInt, \mathcal{B} inputs idx_b from $(\text{idx}_0, \text{idx}_1, \text{qry}) \leftarrow S_q$. For calls to oracle UpdateInt, \mathcal{B} inputs $(\text{pp}, U) \leftarrow S_u$. For calls to oracle ReconInt, \mathcal{B} inputs ans from $(\text{ans}, \text{abort-bit}) \leftarrow S_r$.

If \mathcal{A} finds a valid proof $\pi'_{\neg\text{good},q}$ for a value $\text{rec}'_{\neg\text{good},q} \neq \text{rec}_{\neg\text{good},q}$ to a commitment com_q , then \mathcal{B} will win the integrity game. Thus, $\Pr[\text{BAD}] \leq \text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{int}}(\lambda)(\mathcal{B})$ and

$$\Pr[G_0] = \text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{int}}(\lambda)(\mathcal{B}) + \Pr[G_1].$$

- **Game G_2** : This is equal to the privacy game $G_{\text{APIR}}^{\text{Priv-APIR}}$ (Fig. 1). The calls to RequestHint, VerSetup, Query, and UpdateHint have been replaced with the according pseudocode of these algorithms.
- **Game $G_{(3,i)}$ for $i \in [T_q]$** : This game is equal to G_2 , except for the following change: the first i queries are constructed based on freshly sampled pseudorandom permutations $P_1, \dots, P_{\text{pp},Q}$ over $[M]$. For the last $T_q - i$ queries, the game runs exactly as G_0 using the permutations stored in st_{ck} and swapping out elements. Distinguishing between G_2 and $G_{(3,1)}$ is equivalent to breaking the Show-and-Shuffle game $G^{\text{Ind-SaS}}$ Fig. 10 as shown in [LP24]. For every $i \in [T_q - 1]$, the same argument as above holds for $G_{(3,i)}$ and $G_{(3,i+1)}$. After T_q game hops, all queries are generated using freshly sampled permutations instead of the stored ones used in G_2 . Hence, by the indistinguishability of $G^{\text{Ind-SaS}}$ (Theorem A.1.1)

$$\begin{aligned} |\Pr[G_2] - \Pr[G_{(3,1)}]| &= |\Pr[G_{(3,i)}] - \Pr[G_{(3,i+1)}]| \\ &= \text{Adv}_{\mathcal{A}}^{\text{ind-sas}}(\text{pp}, Q, \text{pp}, M) = 0 \end{aligned}$$

- **Game G_4** : This game is equal to $G_{(3,T_q)}$.
- **Game $G_{(5,i)}$ for $i \in [T_q]$** : This game is equal to the previous game, except for the following change: the first i queries, qry_b for $b \in \{0, 1\}$, are sampled pseudorandomly and uniformly from $[\text{pp}, M]$. For the last $T_q - i$ queries, the game runs exactly as G_4 . Consider the distribution for qry_b for the t -th call to the query

oracle for $b \in \{0, 1\}$ and $t \in [T_q]$. If $t \leq i$, then by definition qry_0 and qry_1 are uniformly sampled from $[M]$. For all other queries, each $P_q(\text{ind})$ is uniformly distributed over $[M]$, as the permutations are sampled independently and uniformly. For $b = 0$, it follows that all values $(\text{qry}_{0,q})_{q \in [Q]}$ are uniformly distributed over $[M]$. For $b = 1$, the same follows but for $(\text{qry}_{1,q})_{q \in [Q] \setminus q^*}$. By definition qry_{1,q^*} is also uniformly distributed over $[M]$. Hence, the entire query vector is uniformly distributed, and thus the distributions between G_4 and $G_{(5,1)}$ are equal. An analogous argument applies to the games $G_{(5,i)}$ and $G_{(5,i+1)}$ for $i \in [T_q - 1]$.

$$|\Pr[G_4] - \Pr[G_{(5,1)}]| = 0$$

- **Game G_6** : This game is equal to $G_{(5,T_q)}$.

In this final game, we note that all queries are now pseudorandomly sampled and fully independent of any server input. We can conclude that through the above series of game hops, the adversary has an advantage negligible in the security parameter λ in winning the privacy game $G_{\text{APIR}}^{\text{Priv-APIR}}$ over a random guess:

$$\begin{aligned} \text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{priv}}(\lambda) &:= \text{Adv}_{\mathcal{A}}^{\text{ind-sas}}(\text{pp}, Q, \text{pp}, M) + \text{Adv}_{\mathcal{A}, \text{APIR}}^{\text{int}}(\lambda) \\ &\leq \text{negl}(\lambda) \end{aligned}$$

Both N and T are polynomial in λ , and so we have a polynomial number of game hops. \square

$G_0: \mathcal{O}_{\text{ReconPriv}}(\text{ans}_{\text{good}})$	$G_1 \quad G'_1: \mathcal{O}_{\text{ReconPriv}}(\text{ans}_{\text{bad}})$	$G_2: \mathcal{G}_{\text{APIR}}^{\text{Priv-APIR}}(\mathcal{A}, \lambda, \text{sp})$
<pre> 1 : if $\neg \text{st}_C.\text{abort} \wedge \text{st}_C.q \neq \perp$ then 2 : $\text{ans}_{\text{good}} \leftarrow \text{inpt}$ $(\text{st}'_C, x, \text{hint}') \leftarrow \text{Recon}(\text{st}_C, \text{hint}, (\text{ans}_b)_{b \in \{0,1\}})$ 3 : $\text{pp} \leftarrow \text{st}_C.\text{pp}$, $\text{ind} \leftarrow \text{st}_C.\text{ind}$, $q^* \leftarrow \text{st}_C.q^*$ 4 : $(\text{qry}_{0,1}, \dots, \text{qry}_{0,\text{pp}.Q}) \leftarrow \text{st}_C.\text{qry}_0$ 5 : $(\text{qry}_{1,1}, \dots, \text{qry}_{1,\text{pp}.Q}) \leftarrow \text{st}_C.\text{qry}_1$ 6 : $((\text{rec}_{0,1} \pi_{0,1}), \dots, (\text{rec}_{0,\text{pp}.Q} \pi_{0,\text{pp}.Q})) \leftarrow \text{ans}_0$ 7 : $((\text{rec}_{1,1} \pi_{1,1}), \dots, (\text{rec}_{1,\text{pp}.Q} \pi_{1,\text{pp}.Q})) \leftarrow \text{ans}_1$ 8 : $(\text{com}_1, \dots, \text{com}_{\text{pp}.M}) \leftarrow \text{pp.d}$ 9 : $(h_1, \dots, h_{\text{pp}.M}) \leftarrow \text{hint}$ 10 : // Verify commitments 11 : for $q \in [\text{pp}.Q]$ do 12 : for $b \in \{0,1\}$ do 13 : $v \leftarrow \text{VC.Verify}(\text{com}_q, \text{qry}_{b,q}, \text{rec}_{b,q}, \pi_{b,q})$ 14 : if $\neg v$ then 15 : $\text{st}_C.\text{abort} \leftarrow \text{true}$ 16 : elseif $\text{rec}_{b,q} \neq \text{st}_S.\text{DB}[\text{qry}_{b,q}]$ 17 : $\text{BAD} \leftarrow \text{true}$ 18 : // Recover record 19 : if $\neg \text{st}_C.\text{abort}$ 20 : $x \leftarrow \left(\bigoplus_{q \in [\text{pp}.Q] \setminus \{q^*\}} \text{rec}_{1,q} \right) \oplus h_{\text{ind}}$ 21 : for $q \in [\text{pp}.Q] \setminus \{q^*\}$ do // Update hint 22 : $h_{\text{ind}} \leftarrow h_{\text{ind}} \oplus \text{rec}_{0,q} \oplus \text{rec}_{1,q}$ 23 : $h_{r,q} \leftarrow h_{r,q} \oplus \text{rec}_{0,q} \oplus \text{rec}_{1,q}$ 24 : if $x = \perp$ then 25 : $\text{abort-bit} \leftarrow \text{true}$ 26 : else $\text{abort-bit} \leftarrow \text{false}$ 27 : $\text{st}_C.\text{ck} \leftarrow \{\sigma_q \mid q \in [\text{pp}.Q]\}$ 28 : $\text{st}_C.q^* \leftarrow \perp$, $\text{hint} \leftarrow (h_1, \dots, h_{\text{pp}.M})$ 29 : return $\text{st}_C.\text{abort}$ abort-bit </pre>	<pre> 13 : $v \leftarrow \text{VC.Verify}(\text{com}_q, \text{qry}_{b,q}, \text{rec}_{b,q}, \pi_{b,q})$ 14 : if $\neg v$ then 15 : $\text{st}_C.\text{abort} \leftarrow \text{true}$ 16 : elseif $\text{rec}_{b,q} \neq \text{st}_S.\text{DB}[\text{qry}_{b,q}]$ 17 : $\text{BAD} \leftarrow \text{true}$ 18 : $\text{st}_C.\text{abort} \leftarrow \text{true}$ 19 : if $\neg \text{st}_C.\text{abort}$ 20 : $x \leftarrow \left(\bigoplus_{q \in [\text{pp}.Q] \setminus \{q^*\}} \text{rec}_{1,q} \right) \oplus h_{\text{ind}}$ 21 : for $q \in [\text{pp}.Q] \setminus \{q^*\}$ do 22 : $h_{\text{ind}} \leftarrow h_{\text{ind}} \oplus \text{rec}_{0,q} \oplus \text{rec}_{1,q}$ 23 : $h_{r,q} \leftarrow h_{r,q} \oplus \text{rec}_{0,q} \oplus \text{rec}_{1,q}$ 24 : $\text{st}_C.\text{ck} \leftarrow \{\sigma_q \mid q \in [\text{pp}.Q]\}$ 25 : $\text{st}_C.q^* \leftarrow \perp$, $\text{hint} \leftarrow (h_1, \dots, h_{\text{pp}.M})$ 26 : return BAD return $\text{st}_C.\text{abort}$ </pre>	<pre> 8 : // Preprocessing $(\text{st}_C, \text{hq}) \leftarrow \text{RequestHint}((\text{pp})_{b \in \{0,1\}})$ 9 : $\text{st}_C.\text{abort} \leftarrow \text{true}$, $(\text{hq}_0, \text{hq}_1) \leftarrow \perp^2$ 10 : // Verify public parameters 11 : if $\text{pp}_0 = \text{pp}_1$ then 12 : $\text{st}_C.\text{abort} \leftarrow \text{false}$ 13 : $(\text{hq}_0, \text{hq}_1) \leftarrow (\text{"Stream DB, please!"})^2$ 14 : $\text{st}_C.\text{pp} \leftarrow \text{pp}_0$ 15 : $(\text{st}'_S, \text{resp}_{\text{good}}) \leftarrow \text{AnsHintReq}(\text{st}_S, \text{hq}_{\text{good}})$ 16 : $(\text{st}'_{\mathcal{A}}, \text{resp}_{\text{good}}) \leftarrow \mathcal{A}_1(\text{st}_{\mathcal{A}}, \text{hq}_{\text{good}})$ $(\text{st}'_C, \text{hint}) \leftarrow \text{VerSetup}(\text{st}_C, (\text{resp}_b)_{b \in \{0,1\}})$ $\text{st}_C \leftarrow \text{st}'_C$ 17 : if $\text{st}_C.\text{abort} \vee \text{resp}_0 \neq \text{resp}_1$ then 18 : $\text{st}_C.\text{abort} \leftarrow \text{true}$ 19 : else 20 : $(\text{ck}, \text{hint}) \leftarrow \perp^2$ 21 : // Generate permutations 22 : for $q \in [\text{st}_C.\text{pp}.Q]$ do 23 : $\text{ck}.\sigma_q \leftarrow \text{Permute}(\text{st}_C.\text{pp}.M)$ 24 : $\text{DB} \leftarrow \text{resp}_0$ 25 : for $m \in [\text{st}_C.\text{pp}.M]$ do 26 : $\text{hint}.h_m \leftarrow \bigoplus_{q=1}^{\text{st}_C.\text{pp}.Q} \text{DB}_q[\text{ck}.\sigma_q(m)]$ 27 : $\text{st}_C.\text{ck} \leftarrow \text{ck}$ </pre>

Figure 11: Games G_0 (left), G_1 and G'_1 (middle), and $G_2 \mathcal{G}_{\text{APIR}}^{\text{Priv-APIR}}$ (right) for the game hops for proving $\mathcal{G}_{\text{APIR}}^{\text{Priv-APIR}}$ for TAPIR (Theorem 5.3.3).

$G_2: \mathcal{O}_{\text{QueryPriv}}(\text{id}x_0, \text{id}x_1)$	$G_{(3,i)}: \mathcal{O}_{\text{QueryPriv}}(\text{id}x_0, \text{id}x_1)$	$G_4: \mathcal{O}_{\text{QueryPriv}}(\text{id}x_0, \text{id}x_1)$
<pre> 1: if $t_q < T_q \wedge \neg \text{st}_C.\text{abort} \wedge \text{st}_C.q^* = \perp$ then $(\text{st}'_C, (\text{qry}_b)_{b \in \{0,1\}}) \leftarrow \text{Query}(\text{st}_C, \text{id}x_{b_{\text{chall}}})$ 2: $\text{pp} \leftarrow \text{st}_C.\text{pp}, \{\sigma_q \mid q \in [\text{pp}.Q]\} \leftarrow \text{st}_C.\text{ck}$ 3: $q^* \leftarrow \lceil \text{id}x_{b_{\text{chall}}} / \text{pp}.M \rceil$ 4: $m^* \leftarrow ((\text{id}x_{b_{\text{chall}}} - 1) \bmod \text{pp}.M) + 1$ 5: Find $\text{ind} \in [\text{pp}.M]$, s.t. $\sigma_{q^*}(\text{ind}) = m^*$ 6: $(r^*, r_1, \dots, r_{\text{pp}.Q}) \leftarrow \mathcal{S}[\text{pp}.M]^{\text{pp}.Q+1}$ 7: $(\text{qry}_0, \text{qry}_1) \leftarrow \perp^2$ 8: for $q \in [\text{pp}.Q]$ do 9: $\text{qry}_0.\text{qry}_{0,q} \leftarrow \sigma_q(r_q)$ 10: $\text{qry}_1.\text{qry}_{1,q} \leftarrow \sigma_q(\text{ind})$ 11: if $q \neq q^*$ then 12: Swap $\sigma_q(\text{ind})$ and $\sigma_q(r_q)$ 13: else $\text{qry}_1.\text{qry}_{1,q^*} \leftarrow r^*$ 14: $\text{st}_C.\text{ck} \leftarrow \{\sigma_q \mid q \in [\text{pp}.Q]\}$ 15: $\text{st}_C.q^* \leftarrow q^*, \text{st}_C.\text{ind} \leftarrow \text{ind}$ 16: $\text{st}_C.\text{qry}_0 \leftarrow \text{qry}_0, \text{st}_C.\text{qry}_1 \leftarrow \text{qry}_1$ 17: $(\text{st}'_S, \text{ans}_{\text{good}}) \leftarrow \text{Answer}(\text{st}_S, \text{qry}_{\text{good}})$ </pre>	<pre> 7: $(\text{qry}_0, \text{qry}_1) \leftarrow \perp^2$ 8: if $t_q < i$ then 9: $(P_1, \dots, P_{\text{pp}.Q}) \leftarrow \mathcal{S} \text{Permute}(\text{pp}.M)^{\text{pp}.Q}$ 10: for $q \in [\text{pp}.Q]$ do 11: if $q \neq q^*$ then 12: $\text{qry}_0.\text{qry}_{0,q} \leftarrow P_q(r_q)$ 13: $\text{qry}_1.\text{qry}_{1,q} \leftarrow P_q(\text{ind})$ 14: else 15: $\text{qry}_1.\text{qry}_{1,q} \leftarrow \mathcal{S}[\text{pp}.M]$ 16: else 17: for $q \in [\text{pp}.Q]$ do 18: $\text{qry}_0.\text{qry}_{0,q} \leftarrow \sigma_q(r_q)$ 19: $\text{qry}_1.\text{qry}_{1,q} \leftarrow \sigma_q(\text{ind})$ 20: if $q \neq q^*$ then 21: Swap $\sigma_q(\text{ind})$ and $\sigma_q(r_q)$ 22: else $\text{qry}_1.\text{qry}_{1,q^*} \leftarrow r^*$ 23: $\text{st}_C.\text{ck} \leftarrow \{\sigma_q \mid q \in [\text{pp}.Q]\}$ </pre>	<pre> : : : 8: if $t_q < i$ then 9: $(P_1, \dots, P_{\text{pp}.Q}) \leftarrow \mathcal{S} \text{Permute}(\text{pp}.M)^{\text{pp}.Q}$ 10: for $q \in [\text{pp}.Q]$ do 11: if $q \neq q^*$ then 12: $\text{qry}_0.\text{qry}_{0,q} \leftarrow P_q(r_q)$ 13: $\text{qry}_1.\text{qry}_{1,q} \leftarrow P_q(\text{ind})$ 14: else 15: $\text{qry}_1.\text{qry}_{1,q} \leftarrow \mathcal{S}[\text{pp}.M]$ 16: else 17: for $q \in [\text{pp}.Q]$ do 18: $\text{qry}_0.\text{qry}_{0,q} \leftarrow \sigma_q(r_q)$ 19: $\text{qry}_1.\text{qry}_{1,q} \leftarrow \sigma_q(\text{ind})$ 20: if $q \neq q^*$ then 21: Swap $\sigma_q(\text{ind})$ and $\sigma_q(r_q)$ 22: else $\text{qry}_1.\text{qry}_{1,q^*} \leftarrow r^*$ 23: $\text{st}_C.\text{ck} \leftarrow \{\sigma_q \mid q \in [\text{pp}.Q]\}$ 15: $\text{st}_C.q^* \leftarrow q^*, \text{st}_C.\text{ind} \leftarrow \text{ind}$ </pre>
$G_{(5,i)}: \mathcal{O}_{\text{QueryPriv}}(\text{id}x_0, \text{id}x_1)$	$G_6: \mathcal{O}_{\text{QueryPriv}}(\text{id}x_0, \text{id}x_1)$	
<pre> : 2: $\text{pp} \leftarrow \text{st}_C.\text{pp}, \{\sigma_q \mid q \in [\text{pp}.Q]\} \leftarrow \text{st}_C.\text{ck}$ 3: if $t_q < i$ then 4: $\text{qry}_0, \text{qry}_1 \leftarrow [\text{pp}.M]^{2 \cdot \text{pp}.Q}$ 5: else 6: $q^* \leftarrow \lceil \text{id}x_{b_{\text{chall}}} / \text{pp}.M \rceil$ 7: $m^* \leftarrow ((\text{id}x_{b_{\text{chall}}} - 1) \bmod \text{pp}.M) + 1$ 8: : </pre>	<pre> 1: if $t_q < T_q \wedge \neg \text{st}_C.\text{abort} \wedge \text{st}_C.q^* = \perp$ then 2: $\text{pp} \leftarrow \text{st}_C.\text{pp}$ 3: if $t_q < i$ then 4: $(\text{qry}_0, \text{qry}_1) \leftarrow [\text{pp}.M]^{2 \cdot \text{pp}.Q}$ 5: else 6: $q^* \leftarrow \lceil \text{id}x_{b_{\text{chall}}} / \text{pp}.M \rceil$ 7: $m^* \leftarrow ((\text{id}x_{b_{\text{chall}}} - 1) \bmod \text{pp}.M) + 1$ 8: Find $\text{ind} \in [\text{pp}.M]$, s.t. $\sigma_{q^*}(\text{ind}) = m^*$ 9: $(r^*, r_1, \dots, r_{\text{pp}.Q}) \leftarrow \mathcal{S}[\text{pp}.M]^{\text{pp}.Q+1}$ 10: $\text{qry}_0, \text{qry}_1 \leftarrow \perp$ 11: $(P_1, \dots, P_{\text{pp}.Q}) \leftarrow \mathcal{S} \text{Permute}(\text{pp}.M)^{\text{pp}.Q}$ 12: for $q \in [\text{pp}.Q]$ do 13: $\text{qry}_0.\text{qry}_{0,q} \leftarrow P_q(r_q)$ 14: $\text{qry}_1.\text{qry}_{1,q} \leftarrow P_q(\text{ind})$ 15: $\text{st}_C.q^* \leftarrow q^*, \text{st}_C.\text{ind} \leftarrow \text{ind}$ 16: $\text{st}_C.\text{qry}_0 \leftarrow \text{qry}_0, \text{st}_C.\text{qry}_1 \leftarrow \text{qry}_1$ 17: : </pre>	

Figure 12: Games $G_2^{\text{OracleQueryPriv}}$ (top left), $G_{(3,i)}$ (top middle), G_4 (top right), $G_{(5,i)}$ (bottom left) and G_6 (bottom middle) for the game hops for proving $G_{\text{APIR}}^{\text{Priv-APIR}}$ for TAPIR (Theorem 5.3.3).