

Private Information Retrieval and Systems

From Theory to Practice with Bernoulli Types

Alexander Towell
`atowell@siue.edu`

January 24, 2026

Abstract

Private Information Retrieval (PIR) enables a client to retrieve data from a server without revealing which item was accessed. Despite decades of theoretical development, practical PIR systems have remained elusive due to fundamental lower bounds and efficiency challenges. This paper bridges theory and practice by applying the Bernoulli type framework to PIR, enabling efficient approximate retrieval with provable privacy guarantees.

Our key contribution is *tuple encoding*: instead of sending separate hash probes $h(x)$ and $h(y)$ for a conjunctive query $x \wedge y$, we send a single probe $h((x, y))$ that hides the correlation between terms. Combined with frequency-proportional sizing ($|\text{Valid}((x, y))| \propto 1/\text{Freq}(x, y)$), this achieves uniform output distribution regardless of query patterns. We extend this to multi-level query obfuscation that combines real queries with noise injection and distribution shifting.

We survey oblivious data structures—the building blocks of PIR systems—and their composition properties. Case studies of deployed systems (Signal’s contact discovery, CryptDB, Tor) demonstrate that these techniques work at scale, serving billions of users. The paper provides both the theoretical foundations for understanding PIR and practical guidance for building privacy-preserving systems.

1 Introduction

1.1 The Private Information Retrieval Problem

A user wishes to query a database without revealing their query. This fundamental problem—Private Information Retrieval (PIR)—arises throughout computing:

- **Search engines:** Query logs reveal interests, health concerns, political views
- **Medical databases:** Searching for a condition reveals potential diagnosis
- **Financial systems:** Querying stock prices reveals investment strategy
- **Location services:** Map queries reveal physical movements

The naive solution—download the entire database—provides perfect privacy but is impractical for large databases. Can we do better?

1.2 Fundamental Lower Bounds

A celebrated result in theoretical computer science says no, at least for perfect information-theoretic security:

Theorem 1.1 (PIR Lower Bound [?]). *Any single-server information-theoretic PIR protocol requires $\Omega(n)$ communication, where n is the database size.*

This lower bound seems to doom practical PIR. To retrieve one item privately, we must essentially transfer the entire database. However, three approaches circumvent this bound:

1. **Multiple servers:** With k non-colluding servers, communication drops to $O(n^{1/k})$
2. **Computational assumptions:** With cryptographic hardness, single-server PIR achieves sublinear communication
3. **Approximate retrieval:** Accept controlled errors via Bernoulli types, enabling efficient protocols

This paper focuses on the third approach, showing how the Bernoulli framework from our companion papers [? ?] enables practical PIR systems.

1.3 Connection to Bernoulli and Oblivious Types

This paper builds on two companion works:

- **Bernoulli Types** [?]: Established latent/observed duality and error propagation for approximate computation
- **Oblivious Computing** [?]: Showed that uniformity provides privacy, and developed inverse-frequency encoding

PIR combines both perspectives:

- Queries are *hidden* values; server responses are *observable* traces
- Retrieval may be *approximate* (Bernoulli), accepting false positives for efficiency
- Query patterns should be *uniform* to hide access patterns

The key insight is that tuple encoding—representing conjunctive queries as single hash probes—achieves uniformity while hiding query structure.

1.4 Contributions

This paper makes the following contributions:

1. **PIR Foundations** (§2): We review PIR theory, lower bounds, and the three approaches to circumventing them.
2. **Boolean Query PIR** (§3): We develop error propagation for complex queries, showing how Bernoulli types enable approximate Boolean retrieval.

3. **Tuple Encoding** (§4): We present the key technical innovation: encoding query pairs to hide correlations.
4. **Multi-Level Obfuscation** (§5): We develop comprehensive query privacy through layered techniques.
5. **Oblivious Data Structures** (§6): We survey building blocks for PIR systems.
6. **Systems** (§7): We analyze deployed systems demonstrating practical feasibility.

2 PIR Foundations

We begin with the theoretical foundations of Private Information Retrieval.

2.1 Problem Definition

Definition 2.1 (Private Information Retrieval). A *Private Information Retrieval* scheme for database $\mathcal{D} = (x_1, \dots, x_n)$ consists of:

1. $\text{Query}(i) \rightarrow q$: Client generates query for index i
2. $\text{Answer}(\mathcal{D}, q) \rightarrow a$: Server computes answer from database and query
3. $\text{Decode}(a, i) \rightarrow x_i$: Client recovers desired item

with the *privacy requirement*: Server learns nothing about i from q .

The privacy requirement can be formalized at three levels (following Paper 2):

- **Perfect**: $\forall i, j : \mathbb{P}[q \mid i] = \mathbb{P}[q \mid j]$ (information-theoretic)
- **Statistical**: Distributions are negligibly close
- **Computational**: No efficient adversary can distinguish

2.2 The Information-Theoretic Lower Bound

Theorem 2.2 (Single-Server IT-PIR Lower Bound). *For any single-server information-theoretic PIR protocol, the communication complexity satisfies $\text{Comm} \geq n$ bits.*

Proof Sketch. Consider the server's view. For IT-security, the query distribution must be identical for all indices. But the server must respond correctly for any index, which requires enough information in the response to reconstruct any item. This forces $\text{Comm} \geq n$. ■

2.3 Multi-Server PIR

With multiple non-colluding servers, we can achieve sublinear communication:

Theorem 2.3 (Multi-Server PIR [?]). *With k non-colluding servers, PIR achieves $O(n^{1/k})$ communication.*

Construction 2.4 (Two-Server XOR-PIR). For database $\mathcal{D} \in \{0, 1\}^n$ and desired index i :

1. Client generates random subset $S \subseteq [n]$

2. To Server 1: Send S
3. To Server 2: Send $S \oplus \{i\}$ (symmetric difference)
4. Server j returns $\bigoplus_{k \in S_j} x_k$
5. Client XORs responses: $(\bigoplus_{k \in S} x_k) \oplus (\bigoplus_{k \in S \oplus \{i\}} x_k) = x_i$

Each server sees a random subset, learning nothing about i .

2.4 Computational PIR

With computational assumptions, single-server sublinear PIR is possible:

Theorem 2.5 (Computational PIR). *Under the Quadratic Residuosity, Decisional Diffie-Hellman, or Learning With Errors assumptions, single-server PIR achieves $O(n^\epsilon)$ communication for any $\epsilon > 0$.*

The practical overhead remains significant. LWE-based constructions achieve $O(\sqrt{n})$ or better but with substantial constants.

2.5 Approximate PIR via Bernoulli Types

Our approach uses Bernoulli types to enable efficient approximate retrieval:

Definition 2.6 (Approximate PIR). An (α, β) -approximate PIR allows:

- False positives: Return x_j for $j \neq i$ with probability at most α
- False negatives: Fail to return x_i with probability at most β

This relaxation enables dramatic efficiency improvements. A Bloom filter-based PIR achieves $O(k \log n)$ communication at the cost of $\alpha = (1 - e^{-kn/m})^k$ false positive rate.

3 Boolean Query PIR

Real queries are rarely single keywords. Users search for “privacy AND encryption,” “cancer OR tumor,” and “machine learning NOT neural networks.” We extend PIR to Boolean queries using Bernoulli type composition.

3.1 Query Types

Definition 3.1 (Boolean Query). A *Boolean query* over keyword set \mathcal{K} is built from:

- Atomic queries: $k \in \mathcal{K}$
- Conjunction: $q_1 \wedge q_2$
- Disjunction: $q_1 \vee q_2$
- Negation: $\neg q$

3.2 Error Propagation for Boolean Operations

When queries are approximate (Bernoulli), errors propagate through Boolean operations:

Theorem 3.2 (Conjunction Error). *For Bernoulli queries \tilde{q}_1, \tilde{q}_2 with false positive rates α_1, α_2 :*

$$\alpha(\tilde{q}_1 \wedge \tilde{q}_2) = \alpha_1 \cdot \alpha_2 \quad (1)$$

Conjunction reduces *false positives*—both must err independently.

Proof. A false positive for $q_1 \wedge q_2$ requires both \tilde{q}_1 and \tilde{q}_2 to false positive. For independent observations: $\mathbb{P}[\text{FP}] = \alpha_1 \cdot \alpha_2$. \blacksquare

Theorem 3.3 (Disjunction Error). *For Bernoulli queries \tilde{q}_1, \tilde{q}_2 :*

$$\alpha(\tilde{q}_1 \vee \tilde{q}_2) = \alpha_1 + \alpha_2 - \alpha_1 \cdot \alpha_2 = 1 - (1 - \alpha_1)(1 - \alpha_2) \quad (2)$$

Disjunction increases *false positives*—either error suffices.

Theorem 3.4 (Negation Error). *Negation swaps error types:*

$$\alpha(\neg \tilde{q}) = \beta(\tilde{q}), \quad \beta(\neg \tilde{q}) = \alpha(\tilde{q}) \quad (3)$$

3.3 Error Rate Bounds

These propagation rules let us compute error bounds for arbitrary Boolean queries:

Corollary 3.5 (Query Error Bounds). *For a Boolean query with c conjunctions and d disjunctions over base queries with FPR α :*

- *Pure conjunction ($d = 0$): $\alpha \leq \alpha^c$ (exponentially decreasing)*
- *Pure disjunction ($c = 0$): $\alpha \leq 1 - (1 - \alpha)^d$ (approaching 1)*
- *Mixed: Depends on structure; conjunctions help, disjunctions hurt*

3.4 XOR versus OR Constructions

Two Boolean operations appear similar but have different privacy properties:

Proposition 3.6 (XOR Privacy). *XOR-based constructions (e.g., XOR-PIR) provide perfect privacy but require multiple servers or computational assumptions.*

Proposition 3.7 (OR Privacy). *OR-based constructions (e.g., Bloom filters) enable single-server approximate PIR but leak information through false positive patterns.*

The choice depends on the security model. XOR is algebraically cleaner (forms a group), while OR aligns with set semantics and Bernoulli types.

4 Tuple Encoding

The key technical contribution of this paper is *tuple encoding*: a technique for hiding correlations in Boolean queries.

4.1 The Correlation Problem

Consider a conjunctive query “privacy AND encryption.” A naive approach sends two separate hash probes:

$$h(\text{“privacy”}), \quad h(\text{“encryption”}) \quad (4)$$

An adversary observing both probes learns:

- The query involves two terms (structure)
- These specific terms are correlated in the user’s interest
- Repeated queries for the same conjunction are identifiable

This leakage violates the uniformity principle from Paper 2.

4.2 The Tuple Encoding Solution

Definition 4.1 (Tuple Encoding). For query terms (t_1, t_2) , the *tuple encoding* is:

$$\widehat{(t_1, t_2)} = h((t_1, t_2)) \quad (5)$$

where the pair is hashed as a single unit.

Theorem 4.2 (Correlation Hiding). *Tuple encoding hides the relationship between t_1 and t_2 :*

1. *An adversary cannot determine whether a probe represents a single term or a tuple*
2. *Repeated probes for (t_1, t_2) and (t_1, t_3) appear uncorrelated*
3. *The AND relationship between terms is hidden*

Proof. The hash function treats (t_1, t_2) as an atomic string. Without knowing the hash function’s internal state, the adversary cannot distinguish $h((t_1, t_2))$ from $h(t_3)$ for any single term t_3 . The correlation is hidden inside the hash input. ■

4.3 Frequency-Proportional Sizing

For uniformity (Paper 2), we size encoding sets inversely to frequency:

Definition 4.3 (Uniform Tuple Encoding). For tuple (t_1, t_2) with frequency $\text{Freq}(t_1, t_2)$:

$$|\text{Valid}((t_1, t_2))| = \frac{c}{\text{Freq}(t_1, t_2)} \quad (6)$$

where c is a normalizing constant.

Theorem 4.4 (Tuple Uniformity). *With frequency-proportional sizing, tuple probes are uniformly distributed regardless of query distribution.*

Proof. Following the uniform encoding construction from Paper 2:

$$\mathbb{P}[\text{Probe} = p] = \sum_{(t_1, t_2)} \text{Freq}(t_1, t_2) \cdot \frac{|\{p \in \text{Valid}((t_1, t_2))\}|}{|\text{HashRange}|} \quad (7)$$

With $|\text{Valid}((t_1, t_2))| \propto 1/\text{Freq}(t_1, t_2)$, each term contributes equally, yielding uniform output. ■

4.4 N-gram Indexing

Tuple encoding extends to phrases via n-gram indexing:

Definition 4.5 (N-gram Index). An *n-gram index* pre-computes tuple encodings for frequent term sequences:

- Bigrams: (t_i, t_{i+1}) for adjacent terms
- Trigrams: (t_i, t_{i+1}, t_{i+2}) for three-term sequences
- General n-grams: Sequences of n terms

Input: Query terms (t_1, \dots, t_k) , frequency threshold τ

Output: Encoded query

$\text{encoded} \leftarrow \emptyset;$

for $i \leftarrow 1$ **to** $k - 1$ **do**

```

if  $\text{Freq}(t_i, t_{i+1}) > \tau$  then
   $\text{encoded} \leftarrow \text{encoded} \cup \{\widehat{(t_i, t_{i+1})}\};$ 
   $i \leftarrow i + 1;$                                 // Skip next term (already encoded)
end
else
   $\text{encoded} \leftarrow \text{encoded} \cup \{h(t_i)\};$ 
end

```

end

return $\text{encoded};$

Algorithm 1: Adaptive Tuple Encoding

The frequency threshold τ balances privacy (encode more tuples) against index size (fewer pre-computed tuples).

4.5 Security Analysis

Theorem 4.6 (Tuple Encoding Security). *Under the random oracle model, tuple encoding achieves:*

1. **Term privacy:** Individual terms are hidden within tuples
2. **Structure privacy:** Query structure (conjunctions) is hidden
3. **Frequency privacy:** With uniform sizing, query frequency is hidden

Proof. In the random oracle model, $h((t_1, t_2))$ is indistinguishable from $h(t_3)$ for any t_3 of the same length. The adversary's view is a sequence of random-looking probes with no correlation structure. ■

5 Multi-Level Query Obfuscation

Tuple encoding hides query structure but not query frequency. A comprehensive approach layers multiple obfuscation techniques.

5.1 The Three-Level Strategy

Definition 5.1 (Multi-Level Obfuscation). A *multi-level obfuscation* strategy combines:

1. **Level 1 (Encoding):** Tuple encoding hides query structure
2. **Level 2 (Noise):** Fake queries provide plausible deniability
3. **Level 3 (Distribution):** Shifting queries mask overall patterns

5.2 Noise Injection

Definition 5.2 (Noise Query). A *noise query* is a fake query indistinguishable from real queries but whose result is discarded.

Proposition 5.3 (Noise Privacy). *With n real queries and m noise queries, an adversary's confidence in identifying a real query is at most $n/(n+m)$.*

```

Input: Real query  $q$ , noise budget  $B$ 
Output: Obfuscated query stream
 $Q_{\text{real}} \leftarrow \text{TupleEncode}(q);$ 
 $Q_{\text{noise}} \leftarrow \emptyset;$ 
for  $i \leftarrow 1$  to  $B$  do
     $q_{\text{fake}} \leftarrow \text{SamplePlausibleQuery}();$ 
     $Q_{\text{noise}} \leftarrow Q_{\text{noise}} \cup \{\text{TupleEncode}(q_{\text{fake}})\};$ 
end
return  $\text{Shuffle}(Q_{\text{real}} \cup Q_{\text{noise}});$ 

```

Algorithm 2: Query Stream with Noise

5.3 Distribution Shifting

Even with noise, query frequency distribution may leak information. Distribution-shifting queries flatten the observed distribution:

Definition 5.4 (Distribution Shifting). *Distribution-shifting queries* are selected to make the observed query distribution approach uniform:

$$\mathbb{P}[\text{Observed query} = q] \approx \frac{1}{|Q|} \quad (8)$$

Input: Query history H , target distribution U (uniform)

Output: Shifting queries to issue

```

 $P_{\text{observed}} \leftarrow \text{EstimateDistribution}(H);$ 
 $\text{deficit} \leftarrow U - P_{\text{observed}};$ 
 $Q_{\text{shift}} \leftarrow \emptyset;$ 
foreach  $q$  with  $\text{deficit}[q] > 0$  do
     $n_q \leftarrow \lceil \text{deficit}[q] \cdot |H| \rceil;$ 
    Add  $n_q$  copies of  $q$  to  $Q_{\text{shift}};$ 
end
return  $Q_{\text{shift}};$ 

```

Algorithm 3: Distribution-Aware Obfuscation

5.4 Complete Obfuscation Protocol

Input: Real query q , noise budget B , history H

Output: Fully obfuscated query stream

```
// Level 1: Structure hiding
 $Q_1 \leftarrow \text{TupleEncode}(q);$ 
// Level 2: Plausible deniability
 $Q_2 \leftarrow \text{GenerateNoiseQueries}(B);$ 
// Level 3: Distribution flattening
 $Q_3 \leftarrow \text{GenerateShiftingQueries}(H);$ 
// Combine and shuffle
 $Q_{\text{all}} \leftarrow Q_1 \cup Q_2 \cup Q_3;$ 
return  $\text{Shuffle}(Q_{\text{all}});$ 
```

Algorithm 4: Comprehensive Query Obfuscation

5.5 Privacy Analysis

Theorem 5.5 (Multi-Level Privacy). *Comprehensive obfuscation achieves:*

1. *Structure privacy from tuple encoding*
2. *Query privacy from noise injection (factor $n/(n+m)$ reduction)*
3. *Distribution privacy from shifting (approaches uniform)*

These protections compose: total leakage is bounded by the product of individual leakage factors.

6 Oblivious Data Structures

PIR systems require data structures that hide access patterns. We survey oblivious variants of fundamental structures.

6.1 The Access Pattern Problem

Definition 6.1 (Access Pattern). The *access pattern* of a computation is the sequence of memory locations accessed during execution.

Even with encrypted data, access patterns leak information:

- Which records are accessed together
- Frequency of access to each location
- Temporal patterns (when accesses occur)

6.2 Oblivious RAM (ORAM)

ORAM is the foundational tool for hiding access patterns:

Definition 6.2 (Oblivious RAM). An *Oblivious RAM* scheme transforms memory accesses so that for any two access sequences of the same length, the resulting physical access patterns are computationally indistinguishable.

Theorem 6.3 (ORAM Lower Bound [?]). *Any ORAM scheme requires $\Omega(\log N)$ bandwidth overhead per access, where N is memory size.*

Construction 6.4 (Path ORAM [?]). Path ORAM achieves $O(\log N)$ overhead using a binary tree structure:

1. Data blocks stored in tree nodes
2. Each block has a random “path” assignment
3. Access reads entire path, rewrites with fresh randomness
4. Stash handles overflow during tree traversal

6.3 Oblivious Arrays

The simplest oblivious data structure:

Proposition 6.5 (Oblivious Array Access). *For an array of N elements:*

- *Linear scan: $O(N)$ time, perfect obliviousness*
- *ORAM-backed: $O(\log N)$ time, computational obliviousness*

Linear scan is preferred for small arrays; ORAM for large ones.

6.4 Oblivious Trees

Binary search trees require care—the access path reveals the search key:

Construction 6.6 (Oblivious Binary Search). To search obliviously:

1. Access all $O(\log N)$ levels
2. At each level, read both children
3. Use oblivious selection to choose the correct path
4. Return result without revealing which path was taken

Overhead: $O(\log N)$ nodes accessed (vs. $O(\log N)$ for normal BST, but all accessed regardless of key).

6.5 Oblivious Hash Tables

Hash tables present unique challenges—the hash function reveals bucket selection:

Construction 6.7 (Oblivious Cuckoo Hashing). Cuckoo hashing with ORAM-backed buckets:

1. Two hash functions h_1, h_2 determine possible locations
2. Access both locations via ORAM
3. Obliviously select the correct item

Overhead: $O(\log N)$ from ORAM, constant number of hash evaluations.

6.6 Composition of Oblivious Structures

Theorem 6.8 (Oblivious Composition). *For oblivious data structures A and B :*

1. *Sequential composition: If A then B , overhead is $O(T_A + T_B)$*
2. *Nested composition: A containing B , overhead is $O(T_A \cdot T_B)$*
3. *Parallel composition: A and B independently, overhead is $O(\max(T_A, T_B))$*

This theorem enables building complex oblivious systems from simple primitives.

7 Systems Case Studies

Theory becomes convincing when deployed at scale. We examine systems serving millions to billions of users.

7.1 Signal: Private Contact Discovery

Signal's contact discovery allows users to find which of their contacts use Signal without revealing their entire contact list to the server.

Construction 7.1 (Signal Contact Discovery).

1. Client hashes phone numbers to identifiers
2. Server maintains set of registered user identifiers
3. Protocol uses Private Set Intersection (PSI) variant
4. Intel SGX provides hardware-level isolation

Scale: Billions of potential phone numbers, hundreds of millions of users.

Technique: Combines PIR with secure enclaves for practical efficiency.

7.2 CryptDB: Encrypted SQL

CryptDB enables SQL queries over encrypted databases using layered “onion” encryption:

Construction 7.2 (CryptDB Onion Encryption). Each data item is encrypted with multiple layers:

1. Outermost: Random encryption (maximum security)
2. Middle: Deterministic encryption (enables equality checks)
3. Inner: Order-preserving encryption (enables range queries)

Layers are peeled as needed for query types.

Trade-off: Functionality requires weaker encryption. Deterministic encryption leaks equality; OPE leaks order.

Lesson: Bernoulli types could provide an alternative—accept approximate results for stronger privacy.

7.3 Tor: Anonymous Routing

Tor provides anonymous internet access through onion routing:

Construction 7.3 (Tor Circuit). 1. Client selects three relays (guard, middle, exit)

2. Establishes encrypted channel to each
3. Traffic wrapped in three encryption layers
4. Each relay removes one layer, forwards to next

Scale: Millions of daily users, thousands of relays.

Connection to PIR: Directory servers use PIR-like techniques for relay selection.

7.4 Oblivious DNS (ODoH)

DNS queries reveal browsing patterns. Oblivious DNS hides them:

Construction 7.4 (Oblivious DoH). 1. Client encrypts DNS query to resolver's public key

2. Sends via proxy (which sees client IP but not query)
3. Resolver sees query but not client IP
4. Neither party has full information

Deployment: Supported by major DNS providers (Cloudflare, Apple).

7.5 Lessons from Deployed Systems

1. **Hybrid approaches work:** Combine cryptography, hardware, and approximation
2. **Trade-offs are unavoidable:** Perfect privacy is too expensive; practical systems accept leakage
3. **Bernoulli types help:** Approximate computation makes privacy practical
4. **Scale is achievable:** Billions of users with reasonable overhead

8 Implementation and Optimization

We conclude with practical guidance for building PIR systems.

8.1 Hardware Acceleration

- **Intel SGX:** Secure enclaves for isolated computation (used by Signal)
- **GPUs:** Parallel hash computation for Bloom filter operations
- **FPGAs:** Custom circuits for cryptographic operations

Hardware can provide orders of magnitude improvement over software-only solutions.

8.2 Batching and Caching

- **Query batching:** Amortize ORAM overhead across multiple queries
- **Result caching:** Cache obliviously (fixed-size cache, random eviction)
- **Preprocessing:** Offline computation to reduce online latency

8.3 Performance Guidelines

Based on deployed systems:

- **Latency target:** 100-500ms for interactive queries
- **Throughput:** 1000+ queries per second per server
- **Space overhead:** 10-100 \times for full obliviousness; 2-5 \times for Bernoulli-based
- **Accuracy:** 99%+ for practical acceptability

8.4 Open Challenges

1. **Tighter bounds:** Close the gap between upper and lower bounds
2. **Adaptive security:** Handle adversaries that adapt based on observations
3. **Composition:** Better understand leakage in complex pipelines
4. **Usability:** Make privacy-preserving systems easy to deploy

9 Related Work

Private Information Retrieval. PIR was introduced by Chor et al. [?], with the single-server lower bound. Computational PIR based on various assumptions followed [?]. Recent work focuses on practical efficiency through preprocessing and batching.

Oblivious RAM. ORAM was introduced by Goldreich and Ostrovsky [?]. Path ORAM [?] achieved practical constants. Recent work explores differential obliviousness for relaxed security.

Searchable Encryption. Song et al. [?] introduced practical encrypted search. Extensive work on leakage analysis [?] shows that even “secure” schemes leak significant information through access patterns.

Bernoulli Types. Our companion Paper 1 develops the theoretical foundations. Paper 2 connects approximation to privacy through uniform distributions.

10 Conclusion

We have bridged PIR theory and practice through the lens of Bernoulli types. Key contributions:

- **Tuple encoding:** Hide Boolean AND correlations via $h((x, y))$
- **Frequency-proportional sizing:** Achieve uniform output distribution
- **Multi-level obfuscation:** Layer encoding, noise, and distribution shifting

- **Oblivious data structures:** Building blocks for PIR systems
- **Real systems:** Evidence that these techniques work at scale

The message is optimistic: practical PIR is achievable. The Bernoulli framework provides the theoretical foundation; tuple encoding provides the technical innovation; deployed systems provide the evidence. Privacy-preserving computation is not just a theoretical possibility but a practical reality.

References