

NSHEDB: Noise-Sensitive Homomorphic Encrypted Database Query Engine

BORAM JUNG, UC Riverside, United States

HUNG-WEI TSENG, UC Riverside, United States

YULIANG LI, Meta, United States

Homomorphic encryption (HE) enables computations directly on encrypted data, offering strong cryptographic guarantees for secure and privacy-preserving data storage and query execution. However, despite its theoretical power, practical adoption of HE in database systems remains limited due to extreme ciphertext expansion, memory overhead, and the computational cost of bootstrapping, which resets noise levels for correctness.

This paper presents NSHEDB¹, a secure query processing engine designed to address these challenges at the system architecture level. NSHEDB uses word-level leveled HE (LHE) based on the BFV scheme to minimize ciphertext expansion and avoid costly bootstrapping. It introduces novel techniques for executing equality, range, and aggregation operations using purely homomorphic computation, without transciphering between different HE schemes (e.g., CKKS/BFV $\leftrightarrow$ TFHE) or relying on trusted hardware. Additionally, it incorporates a noise-aware query planner to extend computation depth while preserving security guarantees.

We implement and evaluate NSHEDB on real-world database workloads (TPC-H) and show that it achieves 20 $\times$ –1370 $\times$ speedup and a 73 $\times$ storage reduction compared to state-of-the-art HE-based systems, while upholding 128-bit security in a semi-honest model with no key release or trusted components.

CCS Concepts: • **Security and privacy** $\rightarrow$ **Database and storage security**; • **Information systems** $\rightarrow$ **Database query processing**.

Additional Key Words and Phrases: Homomorphic Encryption, Database, Query Optimization

ACM Reference Format:

Boram Jung, Hung-Wei Tseng, and Yuliang Li. 2026. NSHEDB: Noise-Sensitive Homomorphic Encrypted Database Query Engine. *Proc. ACM Manag. Data* 4, 1 (SIGMOD), Article 57 (February 2026), 26 pages. <https://doi.org/10.1145/3786671>

1 Introduction

Homomorphic encryption (HE) offers a compelling foundation for building privacy-preserving database systems, enabling secure computation over encrypted data without revealing sensitive contents. It provides end-to-end data confidentiality in untrusted environments, eliminating the need for trusted execution environments (TEEs) or key escrow. However, despite this strong cryptographic guarantee, practical deployment of HE-based database systems remains elusive due to severe performance, storage, and usability constraints.

- **Performance:** Modern HE operations depend on a time-consuming bootstrapping process to reset the noise level in computed data, ensuring the correctness of subsequent computations. Existing HE-based database systems also rely on costly transciphering between various HE

¹Implementation available at <https://github.com/escalab/NSHEDB.git>

Authors' Contact Information: Boram Jung, UC Riverside, United States, bjung022@ucr.edu; Hung-Wei Tseng, UC Riverside, United States, htseng@ucr.edu; Yuliang Li, Meta, United States, yuliangli@meta.com.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

ACM 2836-6573/2026/2-ART57

<https://doi.org/10.1145/3786671>

schemes to support a wider range of database operations. These systems take at least several minutes to perform queries [55, 63, 76].

- **Memory and Storage:** HE schemes result in significant expansion in ciphertext volume as the schemes need to encode difficult-to-decrypt noises while preserving the budget for correct numerical computation. State-of-the-art HE libraries like OpenFHE [8] have to expand data volume by $10,000\times$ to support bit-level operations and meet the standard 128 security level. Storing HE-encrypted data on storage devices or main memory during runtime is economically unfeasible.
- **Functionality:** Although system designers can optionally reduce bootstrapping operations or coarsen the granularity of encoding to decrease the computation and space overhead, each compromise would sacrifice the functionality of the resulting system.

This paper introduces NSHEDB, Noise-Sensitive Homomorphically Encrypted Database, a secure query processing engine designed to make HE-based encrypted databases viable in untrusted execution settings. Rather than proposing new cryptographic primitives, NSHEDB focuses on addressing the system-level bottlenecks that have hindered the practicality of HE-backed query execution. It achieves this through a combination of optimized data encoding, HE-friendly query transformations, and noise-aware execution planning—all while preserving strict security guarantees under a well-defined semi-honest adversary model. NSHEDB is not simply a performance optimization of previous HE DBMS designs—it is a secure system architecture that makes principled design decisions to minimize leakage, avoid scheme-switching transciphering, and operate within the constraints of leveled HE.

The very first design decision of NSHEDB is to utilize the Leveled HE (LHE) [14] mechanism, based on the *Brakerski/Fan-Vercauteren (BFV) scheme* [15], and to encrypt data at the *word level*. The use of LHE at the word level offers several advantages to NSHEDB compared to traditional HE-based database systems. First, this design choice allows NSHEDB to significantly reduce the data expansion rate to only $28\times$ relative to raw data, making storage usage more economical and preserving valuable interconnect bandwidth. Second, by encrypting data at the word level, NSHEDB leverages batch encoding, where multiple data points are encrypted into a single ciphertext, unlocking data-level parallelism to further enhance performance. Finally, since LHE can execute a limited number of operations without bootstrapping, NSHEDB effectively circumvents the computational overhead in time-consuming bootstrapping operations. This makes query execution feasible within the computational budget of leveled HE, avoiding the need for expensive bootstrapping or compromise on security. While approximate encryption methods like CKKS [21] also offer large plaintext space and batch encoding, they are unsuitable for databases requiring exact query selection, as they rely on approximate arithmetic. Since exact comparisons are essential for database queries, BFV is a more suitable choice for NSHEDB.

However, simply implementing word-level LHE does not enable NSHEDB to perform database operations due to the lack of support for bitwise logical operations and arithmetic operations such as division. To address this, NSHEDB proposes a set of algorithms designed to efficiently execute critical database functions, including equality checks, range queries, averages, etc., without the need to re-encrypt data into bit-level HE on the fly. Another constraint of LHE is the limited number of operations it can perform. This paper has identified that the sequence in which operations are performed significantly impacts the growth of the noise level. Therefore, NSHEDB meticulously evaluates each query to minimize the need to reset the noise level of computation results. We have implemented and evaluated NSHEDB, comparing its performance with existing HE-based DB prototypes. NSHEDB is $20\times$ to $1370\times$ faster in a set of TPC-H queries. Furthermore, NSHEDB

only requires 1.4% of the space compared to its counterpart, or 28× relative to the raw data storage without encryption.

In summary, NSHEDB advances encrypted database systems through the following contributions:

- (1) **A novel query processing architecture for FHE constraints:** We identify how the constraints of word-level homomorphic encryption (HE) impose a new architectural objective that renders traditional index structures ineffective. Our novel scan-first database architecture embraces this constraint, pivoting the properties of HE into a strength by maximizing data-level parallelism
- (2) **Implementations of all major relational and arithmetic operators under FHE:** We provide the first comprehensive set of database-aware arithmetic algorithms implementing all major SQL operations (joins, aggregations, filters) entirely within a single homomorphic encryption scheme, eliminating costly encryption scheme conversions (e.g., BFV $\leftrightarrow$ TFHE or CKKS $\leftrightarrow$ TFHE).
- (3) **The first noise-aware query optimizer for encrypted databases:** Unlike traditional optimizers that minimize I/O or CPU, we introduce logical planning that treats multiplicative depth as the primary cost. Our novel optimizations—the *Predicate Pull-Up* and *Mask-Injection Tuning* rewrites—are fundamentally database transformations that restructure query plans to minimize noise and eliminate time-consuming bootstrapping
- (4) **A real-world implementation and benchmark evaluation demonstrating feasibility:** NSHEDB is rigorously evaluated using the *TPC-H* benchmark, demonstrating that secure query processing with leveled HE is feasible under practical workloads and a well-defined threat model.

2 Challenges of using HE in Database Systems

2.1 Homomorphic Encryption

Homomorphic Encryption (HE) is a state-of-the-art encryption technique that allows for the computation of encrypted data without decryption. HE ensures the confidentiality of sensitive information, as all inputs, including user inputs and data storage, remain encrypted throughout the computation process.

Example 2.1. The following example illustrates the basic concept behind HE using one-bit data. Assume the encryption function in an HE system is given by:

$$E(m) = m + 2r + as \quad (1)$$

where m is a one-bit raw data, s represents the secret key, $r \ll |s|$ is a small random noise, and a is a larger, random noise. If we have two data values m_1 and m_2 , and we encrypt m_1 and m_2 with different noise but the same key, then we can derive $E(m_1) = m_1 + 2r_1 + a_1s$ and $E(m_2) = m_2 + 2r_2 + a_2s$. If we directly perform $x = E(m_1) + E(m_2)$, then we can get another encrypted value

$$\begin{aligned} x &= E(m_1) + E(m_2) = (m_1 + m_2) + 2(r_1 + r_2) + (a_1 + a_2)s \\ &= (m_1 + m_2) + 2r' + a's. \end{aligned}$$

If someone has the secret key s , they can decrypt x by modulus s and 2 as

$$\begin{aligned} D(x) &= [(m_1 + m_2) + 2r' + a's] \mod s \mod 2 \\ &= [(m_1 + m_2) + 2r'] \mod 2 = m_1 + m_2 \end{aligned}$$

The same theory works for multiplications. By computing $y = E(m_1) \times E(m_2)$, we can get an encrypted value

$$\begin{aligned} y = E(m_1) \times E(m_2) &= m_1 m_2 + (a_1 E(m_2) + a_2 E(m_1))s \\ &\quad + 2(m_1 r_2 + r_1 m_2 + 2r_1 r_2) = m_1 m_2 + a''s + 2r'' \end{aligned}$$

With s , we can decrypt y and get $m_1 m_2$ through the same decryption function.

Through the above example, we can make several observations. First, computing x or y does not require s , the secret key. Second, anyone received $E(m_1)$, $E(m_2)$, and x cannot easily derive m_1 , m_2 and $D(x)$ without s . Third, even with s , the system cannot decrypt the encrypted data if the a' and r' or a'' and r'' are too large. Finally, when we compute x or y , the noise level in x and y grows as the computation aggregates or multiplies the noise from $E(m_1)$ and $E(m_2)$.

While the first two observations highlight the capability to compute on encrypted data without decryption, thereby ensuring security, the latter two present a key challenge for HE systems: managing noise growth. Modern HE approaches, such as Fully Homomorphic Encryption (FHE) and Leveled Homomorphic Encryption (LHE), offer different strategies to address this issue.

2.1.1 Fully Homomorphic Encryption and Leveled Homomorphic Encryption.

Fully Homomorphic Encryption. Homomorphic operations introduce noise into ciphertexts, and if this noise exceeds a threshold, it can corrupt the decrypted results, limiting the number of computations that can be performed. Fully Homomorphic Encryption (FHE) resolves this by using bootstrapping, a process that resets the noise level, allowing for unlimited computations. This capability makes FHE suitable for evaluating arbitrary circuits and supporting complex applications [31].

Leveled Homomorphic Encryption. Leveled Homomorphic Encryption (LHE) manages noise growth by limiting the number of allowed multiplications, known as computation levels (or multiplicative depth). Each multiplication contributes to noise accumulation, and if the total noise exceeds a threshold, decryption errors may occur. Multiplicative depth measures the longest path of sequential multiplications from input to output.

Unlike Fully Homomorphic Encryption (FHE), which resets noise via bootstrapping to enable unlimited computation, LHE operates efficiently within a fixed multiplicative depth, reducing bootstrapping overhead. However, when computations exceed the predefined depth, LHE can selectively apply bootstrapping to enable deeper computations. Schemes like BFV [15], BGV [14], and CKKS [21] define the maximum allowable depth based on encryption parameters. While LHE aims to minimize bootstrapping, complex queries may still leverage it selectively when necessary.

2.1.2 Multiplicative Depth. In homomorphic encryption, ciphertexts accumulate *noise* with each operation, and deeper computations risk exceeding the noise budget—the maximum noise a ciphertext can tolerate before decryption fails. Multiplicative depth captures the longest sequence of multiplications from input to output. Formally, for two ciphertexts A and B with depths $\text{depth}(A)$ and $\text{depth}(B)$, their product $C = A \times B$ satisfies:

$$\text{depth}(C) = \max(\text{depth}(A), \text{depth}(B)) + 1.$$

As multiplications add new noise and amplify existing noise, computations with high multiplicative depth increase the risk of surpassing the noise budget and bootstrapping. However, careful parameter selection in LHE can ensure computations remain within a manageable noise budget, reducing or eliminating the need for bootstrapping.

In addition to the multiplicative depth that provides a straightforward measure of noise growth, the order of multiplications can potentially change the noise aggregation. Operations like *relinearization* or *rescaling* occur after each multiplication and can introduce additional noise, depending on the intermediate ciphertext states. Thus, even if two multiplication strategies ultimately implement the same expression (e.g., x^4) — one as $(x \times x) \times (x \times x)$ and another as $x \times (x \times (x \times x))$ — they can exhibit different intermediate noise behavior, potentially causing one strategy to reach the noise budget sooner than the other.

2.1.3 Bootstrapping. During homomorphic computations, noise accumulates in ciphertexts, and once it exceeds a threshold, it can corrupt decryption results. Bootstrapping resets the noise level, enabling continued computations without loss of accuracy. This technique applies the decryption algorithm homomorphically, reducing noise while keeping data encrypted. Bootstrapping is crucial in FHE for supporting unlimited computations. Although specifics may vary across HE schemes, the primary goal remains the same: to reset noise and ensure computational integrity.

2.1.4 Data Encryption Granularity of HE. The granularity of data encryption significantly affects the efficiency and capabilities of HE schemes. Two primary types of encryption schemes that support exact computations are: (1) bit-level encryption (e.g., TFHE [23], FHEW [29]), and (2) word-level encryption (e.g., BGV [14], BFV [15]).

Bit-Level Encryption. Bit-level encryption, based on Boolean circuits, operates on individual bits and supports logical operations like AND, OR, and XOR. It is useful for non-linear algorithms and comparison operations, but its efficiency is limited due to significant data expansion, as each ciphertext handles only a few bits of data, leading to high storage and computational costs.

Word-Level Encryption. In contrast, word-level encryption operates on integers using arithmetic circuits. It offers better efficiency for linear algorithms and supports batch encoding, where multiple integers are packed into a single ciphertext. This enables data-level parallelism, making word-level encryption more space- and computationally-efficient. However, it is less effective for operations requiring fine-grained bitwise logic.

2.1.5 Data Type Encoding in BFV. Since BFV operates on integers, non-integer data types require encoding. Floating-point numbers are encoded as fixed-point representations to maintain precision, while Boolean values map to integers (0 for False, 1 for True). For string data, we employ dictionary encoding, assigning sequential integer IDs to unique string values.

2.1.6 BFV Operations. The BFV scheme provides native support for arithmetic operations on packed ciphertexts. Addition and Subtraction operate element-wise between ciphertexts, as does Multiplication. Negation flips the sign of all packed values, while Rotation cyclically shifts elements by a specified offset.

Building on these primitives, HE systems implement two derived operations critical for database queries. Extract isolates a single element by multiplying with a basis vector (all zeros except a single 1 at the target position). Broadcast replicates a value across all slots via extraction followed by rotations and additions. These derived operations enable the fine-grained data manipulation required for query processing.

2.1.7 Basic logical comparisons.

Equality. HE systems commonly evaluate equality using Fermat's Little Theorem [18, 33, 37, 38, 72]. While one could theoretically achieve bitwise comparisons in BFV by setting the plaintext modulus to 2, this fundamentally changes the data representation. Integers must be bit-decomposed across multiple slots (e.g., 16 slots for a 16-bit integer), reducing capacity by 16 times. Moreover,

arithmetic operations would require complex bit-level circuits with rotations, eliminating BFV's efficient SIMD arithmetic advantages. Queries mixing bitwise and arithmetic operations would require maintaining dual encryption schemes with different moduli.

Therefore, arithmetic approaches using standard plaintext moduli are strongly preferred. These approaches leverage Fermat's Little Theorem to perform comparisons arithmetically. For prime p and non-zero a , Fermat's Little Theorem states:

$$a^{(p-1)} \equiv 1 \pmod{p} \quad (2)$$

This property enables equality testing: the expression $(x-y)^{(p-1)}$ equals 1 for non-zero differences and is undefined for zero. The system computes equality between x and y as:

$$EQ(x, y) = 1 - (x - y)^{(p-1)} \quad (3)$$

This computation yields encrypted 1 when values are equal, 0 otherwise.

Range Comparisons. Range comparisons extend the Fermat's Little Theorem approach by checking whether the difference $x - y$ falls within the negative range $[-\frac{p-1}{2}, -1]$. Under modular arithmetic, negative values map to the upper half of the field. The principle follows directly: $x < y$ holds if and only if $x - y$ yields a negative result. For instance, the comparison $3 < 5$ produces $3 - 5 = -2$, which falls in the negative range and thus evaluates to true, while $5 < 3$ produces $5 - 3 = 2$, which falls in the positive range and evaluates to false.

The system computes the less-than comparison by checking whether $x - y$ equals any value in the negative range:

$$LT(x, y) = \sum_{a=-\frac{p-1}{2}}^{-1} \left(1 - (x - y - a)^{(p-1)} \right) \quad (4)$$

This computation yields encrypted 1 when $x < y$, 0 otherwise. Systems can compute both equality and range comparisons simultaneously using the optimization from [37].

2.2 Challenges

Despite the properties of preserving the form of encrypted data during the whole computing process to ensure the privacy and security of user data, using HE in database systems is very challenging in the following aspects.

Performance. FHE performance characteristics create an entirely different optimization landscape for database systems. Bootstrapping—required to reset noise levels—dominates execution time [24, 28, 39, 57, 58, 62]. A breakdown of execution time for Boolean gate evaluations on 1-bit data using OpenFHE [8] on our testbed shows that bootstrapping contributes to 99.99% of the total execution time. Prior work [3] shows that CKKS takes 44 seconds to bootstrap 13-bit 32,768 elements (0.0013 seconds per bit), BGV takes 163 seconds to bootstrap 16-bit 1,024 elements (0.0099 seconds per bit). Recent OpenFHE implementations achieve TFHE bootstrapping at approximately 5ms per 1-bit gate on modern CPUs [8], with GPU and specialized hardware achieving even lower latencies. Recent studies on HE-based databases [55, 63, 76] further highlight that bootstrapping dominates query processing time, accounting for up to 90% of total execution time. This fundamentally changes what database systems must optimize:

- Traditional cost: I/O and CPU cycles for data movement and computation
- FHE cost: Multiplicative depth that accumulates noise toward bootstrapping thresholds
- Traditional heuristics become counterproductive: predicate pushdown can trigger bootstrapping, turning a cheap filter into an expensive operation

- Operation ordering dramatically affects total cost: the same query can vary from zero to multiple bootstraps depending on execution order

Database optimizers must be completely redesigned to minimize multiplicative depth rather than I/O patterns.

Space and Memory Efficiency. Encrypting data significantly expands the size of the resulting ciphertexts. When using a state-of-the-art HE library that supports bit-level HE [8] and meets the standard 128-bit security level, data expansion reaches 10,000×. Similarly, HE-based databases [55, 63, 76] that operate at the bit level experience data expansion of approximately 8,000×, leading to substantial storage and computational overhead. Storing data in bit-level FHE is economically infeasible. Even when assuming secure storage and converting data into HE form before offloading it to a third party, the ciphertexts impose significant pressure on network bandwidth, system interconnects, and memory capacity due to the enormous data volume [28, 32, 35, 54, 64].

Functionality. An HE system can reduce computation, memory, and storage overhead by encrypting data at the word level or limiting the depth of operations [47, 64, 76]. However, both approaches come at the cost of functionality. Word-level encryption lacks native bitwise operations, such as AND/OR, making it infeasible for fine-grained comparisons for database tasks. Without bit-level operations, database systems cannot directly compare ciphertexts, posing challenges for queries involving equality or range checks. Additionally, limiting the depth of operations means that complex computations will inevitably require bootstrapping to reset noise levels, ensuring the correctness of results. Without sufficient bootstrapping, deeper computations may become unreliable, restricting the types of queries the system can support.

2.3 Alternatives

Prior work has investigated other alternatives for making database systems secure or privacy-preserving. However, none achieve the same level as NSHEDB.

2.3.1 Trusted Hardware-Based Database. Conventional secure database systems rely on encrypted data storage, secure key exchange mechanisms, and trusted execution environments (TEEs) [5, 10, 45, 53, 56, 64, 67, 69, 71]. These designs assume the TEE can temporarily store the decryption key and perform the necessary decryption on the encrypted data to compute the plaintext data. If any component in the TEE is compromised, the whole system becomes insecure. In addition, the attacker can learn the data access patterns and gain access to the data due to the limited space within the TEEs [30]. Oblivious RAM (ORAM)-based database [26, 30] can alleviate the data access pattern leakage problem but still relies on the trustworthiness of the underlying hardware. In contrast, HE does not need to trust any third-party hardware, and the entire system remains intact even though we offload computation to compromised hardware.

2.3.2 Hybrid Cryptosystem. Hybrid cryptosystems integrate multiple cryptographic techniques, such as order-preserving encryption (OPE), deterministic encryption (DET), and partial homomorphic encryption (PHE) [52, 61]. While hybrid systems offer flexibility for different operations, they introduce vulnerabilities like information leakage and side-channel attacks, especially in DET and OPE schemes [27, 49, 50, 75]. These vulnerabilities compromise the system's security in ways that HE avoids.

2.3.3 Multi-Party Computation (MPC). MPC distributes a secret among parties, ensuring none possess the entire secret. Each party conducts computations on partial secrets, combined by the source for the desired outcome. Prior studies demonstrate MPC's feasibility in databases [11, 17, 36, 41, 43, 60, 68, 70].

MPC and HE serve different use cases and are complementary rather than competing approaches:

Different Trust Models: MPC requires trust assumptions about non-collusion between parties—if sufficient parties collude, security fails. HE relies solely on cryptographic hardness assumptions to ensure data confidentiality, typically assuming a semi-honest computing party.

Different Operational Models: MPC requires active participation and coordination among multiple parties throughout computation, with associated communication overhead. HE enables fully asynchronous computation by a single untrusted party, with no interaction after initial data encryption.

Different Regulatory and Architectural Fit: Single-provider custody requirements, common in healthcare and finance, may make distributed MPC architectures legally or practically infeasible. Conversely, scenarios requiring joint computation over distributed private inputs naturally fit MPC.

Performance Trade-offs: MPC achieves dramatically better computational performance—Secrecy [43] reports 6 seconds for TPC-H Q6 on 8 million rows, while HE-based systems such as HE³DB require 12,100 seconds for just 32K rows. This gap reflects fundamental architectural differences: MPC performs computations on secret shares with modest overhead, while FHE requires complex polynomial arithmetic on large ciphertexts.

We consider MPC-based databases as orthogonal to our work. Organizations choose between MPC and HE based on their specific trust model, regulatory constraints, and architectural requirements. When HE is necessary due to these constraints, our work demonstrates that careful system design can make HE-based databases practical, achieving orders-of-magnitude improvements over prior HE systems.

2.3.4 Existing HE-enabled DBMS. Before NSHEDB, researchers have explored the potential of using HE in database workloads. However, none is ideal in addressing the performance, space/memory, and functionality challenges.

Early proposals use somewhat Homomorphic Encryption (SHE) in retrieving encrypted records through conjunction queries [1, 12, 40, 48, 74]. However, as SHE only supports addition and multiplication, SHE-based systems can only achieve private information retrieval (PIR) but not general queries.

Recent projects [55, 63, 76] strategically combine various HE schemes to enhance the versatility of HE-based systems. These works store data using bit-level encryption (TFHE) and apply logical operations for comparisons in WHERE clauses. For arithmetic operations, all of them dynamically transcipher encrypted data from TFHE to BFV for HEDA [55] and to CKKS for HE³DB [63] and ArcEDB [76]. However, the transciphering process is as expensive as bootstrapping, as it involves re-encrypting data under a new HE scheme using the secret key of the new scheme [9, 20, 46]. HEDA further reports that 90% of the query evaluation time is spent on the Parameter-Lifting Bootstrapping procedure. Prior work has also explored encoding data into a vector of field elements (VFE) to leverage data-level parallelism, demonstrating the potential to manage comparison latency at the millisecond level [65]. However, these previous approaches do not encompass the broader domain of query evaluations addressed in this paper.

3 NSHEDB Design Overview

Figure 1 provides an overview of NSHEDB. It assumes the user application runs in a trusted zone (e.g., a private device or virtual machine) where only trusted software components can access the secret key. The application does not trust the system's interconnect, storage, or other resources. When NSHEDB receives a query to load or store data, it encrypts or decrypts batch ciphertext using the BFV scheme and the secret key. The query engine then translates operators into arithmetic operations supported by NSHEDB's encoding. NSHEDB evaluates the noise budget based on the computational




Fig. 1. The system overview of NSHEDB

chain and seeks opportunities to minimize multiplicative depth, effectively shortening the chain. By reducing depth where possible, NSHEDB postpones the need for bootstrapping, improving overall execution efficiency. Throughout this process, NSHEDB does not require any key exchanges with any resource outside the trusted zone and maintains data in HE formats.

At a high level, the NSHEDB design adheres to three key principles to address the challenges of modern HE schemes in database systems:

A Constraint-Driven Scan-First Architecture: Defining a New Optimization Objective. NSHEDB stores data homomorphically encrypted using word-level HE batch encoding, placing multiple data points into a single ciphertext to avoid $10,000\times$ data size expansion in bit-level. For instance, with an encryption parameter setting of a polynomial degree of 32,768, NSHEDB stores 32,768 data points in a single ciphertext instead of using 32,768 separate ciphertexts while enabling massive parallelism. However, simply applying word-level encoding creates a fundamental challenge: **no direct slot access**. The HE scheme operates in a SIMD-style, applying operations to all slots simultaneously. Isolating slot i requires a multi-step process involving creating a mask vector, multiplying the entire ciphertext by this mask, and using rotations to move the result.

This constraint fundamentally breaks traditional database assumptions and imposes a **new architectural objective**: Positional access through masking costs the same as full scans or leaks access patterns, demonstrating a shift in complexity from logarithmic to linear.

- **Index structures become ineffective**: While access to position i is technically possible through masking, it requires processing the values in the ciphertext—incurring the same cost as scanning all elements. B-trees and hash tables thus lose their efficiency advantage when $\text{array}[i]$ costs $O(n)$ instead of $O(1)$.
- **Data manipulation degrades to full rewrites**: Operations, such as swapping two elements or moving records between buckets, cannot be done in-place. Each "point update" requires creating masks, processing entire ciphertexts, and reconstructing the result—essentially rewriting the entire packed vector.
- **Security amplifies inefficiency**: Even if masking were cheap, using it to access specific positions would leak access patterns. To maintain security, we must process all positions anyway and mask the results.

Rather than attempting to preserve traditional access methods under this constraint, NSHEDB takes a different approach. Recognizing that conventional indexing provides no performance benefit when all operations require full ciphertext processing, we design a **scan-first architecture** that exploits the inherent parallelism of word-level encoding—processing 32,768 values per operation. This represents a specialized query processing architecture designed specifically for the unique characteristics of HE.

The implications of word-level HE constraints extend beyond simple data access to affect all traditional database optimizations. Materialized views, while technically feasible, suffer from the batching-versus-incremental tension. Consider a view with predicates (e.g., $\text{SUM}(\text{price}) \text{ WHERE price} > 1000$). Incremental maintenance requires evaluating predicates on single rows, paying

the full comparison cost without data-level parallelism from batching—the same computational cost for a single entry as for 32,768 entries. Materialized join results face prohibitive memory costs: a 1K×1K join produces 7.4GB of ciphertexts. Materialized views only benefit narrow cases: (1) predicate-free aggregations requiring single ciphertext additions, (2) complex views where full recomputation cost exceeds incremental update overhead by orders of magnitude, or (3) storing aggregated join results rather than full joins.

Table 1 quantifies how packing transforms database operations:

Operation	Plaintext	FHE with Packing
Point lookup	$O(1)$ or $O(\log n)$	$O(S)$ (process all slots)
Range scan	$O(\log n + k)$	$O(n)$ (scan + mask)
Hash join	$O(n + m)$	$O(n \cdot \lceil m/S \rceil)$
Optimization target	I/O, CPU	Multiplicative depth

Table 1. Packing forces scan-based behavior in all FHE operations

These fundamental shifts in complexity—from logarithmic to linear, from I/O-bound to depth-bound—demonstrate that FHE requires not adaptation but complete rearchitecting of database query processing.

Translating logical operations into arithmetic operations to avoid transciphering overhead. NSHEDB translates unsupported logical operations into arithmetic equivalents to avoid the need for transciphering between different HE schemes. Traditional HE-based systems often rely on transciphering between word-level and bit-level encryption to support logical operations such as AND, OR, and XOR. However, this process incurs significant latency and efficiency losses.

Specifically, transciphering introduces three major issues:

- (1) **Transciphering Latency:** Switching between word-level (e.g., BFV/CKKS) and bit-level (e.g., TFHE) encryption is costly, often involving re-encryption and bootstrapping—one of the most expensive HE operations. For instance, transciphering from LWE (TFHE format) to RLWE (BFV/CKKS format) is up to 100,000× more expensive than a single homomorphic arithmetic operation [9, 46], significantly slowing down query execution.
- (2) **Loss of Batch Encoding:** Word-level encryption (e.g., BFV/CKKS) supports batching (e.g., processing 32,768 elements in one ciphertext), enabling parallelism. Transciphering to bit-level encryption removes this capability, leading to significantly slower performance.
- (3) **Frequent Bootstrapping:** Bit-level encryption schemes lack the "levels" present in word-level schemes, necessitating frequent bootstrapping to control noise accumulation. This adds additional overhead and further increases the latency of bit-level operations.

To overcome these drawbacks, NSHEDB implements logical operations directly as arithmetic operations, eliminating the need for transciphering. While this increases the number of arithmetic operations, it retains the advantages of batch encoding, making the system far more efficient. Section 4.2 provides a detailed explanation of how NSHEDB translates essential logical operations in database workloads into arithmetic operations.

Managing Noise to Reduce Bootstrapping. Rather than relying on bootstrapping to control noise, NSHEDB uses LHE to manage the noise budget during query execution. Since bootstrapping can be 5,000× to 140,000× slower than standard homomorphic operations [3, 8, 22, 66], NSHEDB fundamentally reimagines query optimization for the FHE domain. Under leveled HE with packing, all operations are inherently scan-based—the optimizer cannot choose between access methods but

must instead minimize multiplicative depth while scanning. This leads to counterintuitive design decisions: NSHEDB inverts traditional heuristics through Predicate Pull-Up and Mask-Injection Tuning, deliberately delaying filter applications when it reduces depth. By decomposing complex predicates into independent subgraphs and carefully choosing where to inject masks in join trees, NSHEDB optimizes query plans to reduce multiplicative depth, thereby minimizing the frequency of bootstrapping. Section 4.3 provides further details on these noise-sensitive query optimizations.

Security Analysis. NSHEDB assumes a semi-honest adversary who correctly executes protocols but may attempt to infer information from encrypted data—the standard model for encrypted databases [6, 52, 63, 76]. NSHEDB’s security relies on BFV’s IND-CPA (Indistinguishability under Chosen Plaintext Attack) guarantee, ensuring ciphertexts reveal nothing about plaintexts, and RLWE (Ring Learning With Errors), which makes decryption computationally infeasible without the secret key.

NSHEDB exposes the following leakage, denoted as L :

- **Schema information:** Table structure, column names, foreign key constraints
- **Static metadata:** Base table sizes, dictionary sizes for encoded columns
- **Query structure:** Operator DAG shape (but not predicate values or selectivities)
- **Output structure:** Fixed n -slot ciphertext vectors for all operators
Operator-specific leakage includes:
 - **SELECT:** Masks non-matches with encrypted zeros; ciphertext structure remains unchanged, concealing selectivity. Zeros are dropped on the client-side.
 - **GROUP BY:** Groups only on foreign key columns and always produces one aggregate per parent key ($|P|$ ciphertexts, where $|P|$ is already leaked via schema).
 - **JOIN:** Behaves identically to GROUP BY—outputs a fixed number of ciphertexts based on key cardinality. Equality masks are computed in bulk, revealing only output count (already leaked via FK constraints), not which specific key pairs matched.
 - **ORDER BY:** Scans the entire dictionary with fixed iterations, concealing value-frequency patterns.
- **Aggregations:** Return fixed-size results.

All operations utilize IND-CPA secure BFV ciphertexts, ensuring no information leaks beyond L . This leakage profile aligns with prior encrypted databases [51, 52] while uniquely avoiding the need for trusted hardware and the overhead of scheme switching.

4 NSHEDB

NSHEDB enhances the efficiency of HE for database workloads by leveraging the LHE BFV scheme, which offers benefits such as space efficiency, avoidance of bootstrapping, and batch encoding. However, simply applying the LHE BFV scheme is insufficient for DBMSs, as the LHE BFV scheme lacks the critical support of operators necessary for general-purpose database queries. This paper makes two significant contributions to the realization of LHE database systems. First, NSHEDB presents an innovative approach that efficiently supports the missing comparison operations, which are critical for database queries but are not natively supported by the BFV scheme. Second, NSHEDB introduces a novel noise-sensitive query optimization technique that manages noise accumulation during query execution while minimizing the invocation of bootstrapping. This section will describe both of these innovations in detail.

Table 2. NSHEDB Operator Support and Complexity

Operator	Support	Complexity
Selection (WHERE)	✓	$O(n/S)$ ops
Projection	✓	$O(1)$
COUNT/SUM	✓	$\log S$ depth
AVG	Partial [†]	returns (SUM, COUNT)
GROUP BY	✓	$O(G \cdot n/S)$ ops
JOIN	✓	$O(N \cdot \lceil M/S \rceil)$ ops
ORDER BY	✓	$O(D \cdot n/S)$ ops
Boolean AND/OR/NOT	✓	via arithmetic [‡]
B-tree/Hash index	×	degrades to scan
Native division	×	use workarounds

[†] Client-side division. [‡] AND= ab , OR= $a+b-ab$, NOT= $1-a$.

4.1 Data Storage and Encoding

NSHEDB encrypts data using the BFV scheme with batch encoding, packing 32,768 values into a single ciphertext. This significantly improves storage efficiency compared to encrypting values individually—storing one ciphertext per batch instead of 32,768 separate ciphertexts. For example, fully packing and encrypting 32,768 entries results in 27 times the raw data size—0.27 MB of raw data expands to a 7.4 MB ciphertext.

4.2 Supporting Database Queries

Unlike existing alternatives [55, 63, 76] that rely on transciphering between word-level and bit-level encryption schemes, NSHEDB builds database operators entirely within BFV’s arithmetic framework. By leveraging the comparison operations described in Section 2.1.7, NSHEDB avoids the costly scheme conversions that dominate execution time in prior systems. This section describes how NSHEDB combines these arithmetic building blocks to construct efficient database query operators.

4.2.1 Operator Support Overview. Table 2 summarizes NSHEDB’s operator support. All operations work within BFV’s word-level encryption, avoiding transciphering overhead. These complexities reflect packing constraints: **JOIN** reduces $N \times M$ comparisons to $N \cdot \lceil M/S \rceil$ ciphertext operations (e.g., $1K \times 1K$ join = 1000 ops vs 10^6 scalar). **GROUP BY** requires $|G|$ equality checks processing S values in parallel—expensive for high-cardinality columns. **ORDER BY** enumerates all $|D|$ domain values—practical only for small categorical domains. Each equality needs ≈ 16 multiplications (Section 2.1.7), high but below bootstrapping cost.

4.2.2 Supporting Database Query Operators. In this section, we describe how NSHEDB supports major SQL / relational algebra operators using HE.

SELECT operator. NSHEDB supports the SELECT operator (σ) by leveraging its unique comparison algorithms across HE-encrypted columns of interests. Since the third-party computing resource cannot decrypt the numerical results of HE computations to evaluate equality or inequality conditions, NSHEDB multiplies the encrypted columns with the results of equality or inequality evaluations under the SELECT condition.

For example, consider the following query:

```
1 SELECT A.price FROM A WHERE A.ID == A.key;
```




Fig. 2. The process of JOIN in NSHEDB

NSHEDB performs the following HE computation on the elements corresponding to the same record in the HE-encrypted columns:

$$A.price \times EQ(A.ID, A.key) \quad (5)$$

where EQ denotes the equality function from Section 2.1.7. For records where $A.ID == A.key$, the comparison function $EQ(A.ID, A.key)$ returns a ciphertext with encrypted "1"s in the corresponding positions. For records that do not match the condition, $EQ(A.ID, A.key)$ returns encrypted "0"s, effectively zeroing out those elements in the ciphertext. By multiplying the resulting ciphertext from the equality evaluation with the HE-encrypted $A.price$, only the exact values of $A.price$ that meet the condition are retained in the final ciphertext for the SELECT query.

Aggregation: COUNT, SUM, and AVG. NSHEDB also supports aggregations over the results of a SELECT operator in the following ways.

COUNT. Supporting COUNT requires overcoming the inability to directly access individual elements within batch-encoded ciphertexts. NSHEDB uses repeated HE rotation operations with a doubling pattern: rotate by 1 and add, then by 2 and add, then 4, 8, and so on.

For example, with 8 slots $[v_0, \dots, v_7]$:

- Step 1: Rotate by 4, add $\Rightarrow [v_0+v_4, v_1+v_5, v_2+v_6, v_3+v_7, \dots]$
- Step 2: Rotate by 2, add $\Rightarrow [v_0+v_4+v_2+v_6, v_1+v_5+v_3+v_7, \dots]$
- Step 3: Rotate by 1, add $\Rightarrow [\sum_i v_i, \sum_i v_i, \dots, \sum_i v_i]$

This rotate-sum reduction requires only $\log n$ steps because each step doubles the values accumulated. For COUNT, this applies to comparison result ciphertexts containing encrypted 1s and 0s. The final ciphertext contains the exact count in every slot.

SUM. NSHEDB applies a similar approach for SUM as it does for COUNT. However, instead of accumulating comparison results, NSHEDB multiplies the encrypted column of interest with the comparison result ciphertext and accumulates the actual values of the records using the same doubling rotation pattern.

AVG. Since BFV/BGV schemes lack division operations, NSHEDB employs three workarounds for AVG queries: (1) *Query Rewriting*: Transform predicates to avoid division—e.g., $\text{salary} > \text{AVG}(\text{salary})$ becomes $\text{salary} \times \text{COUNT}(\text{salary}) > \text{SUM}(\text{salary})$. (2) *Propagating Aggregates*: Pass (SUM, COUNT) pairs through the query plan instead of computing averages in intermediate results. (3) *Final Computation*: Perform actual division client-side after decryption using the SUM and COUNT components.

Range and set queries: IN and BETWEEN.

IN (set membership). NSHEDB evaluates the IN operation by applying equality functions on the given set of numbers and adds all equality results. If $\mathbb{S}$ represents the given set of numbers, NSHEDB implements the following mathematical function to fulfill the evaluation of whether x is in $\mathbb{S}$.

$$IN(x, \mathbb{S}) = \sum_{y \in \mathbb{S}} EQ(x, y) \quad (6)$$

Despite the execution time being linear to the number of elements in $\mathbb{S}$, the latency is still significantly lower than that of mechanisms relying on transciphering.

BETWEEN. NSHEDB computes the BETWEEN operation as the product of two comparison results, deriving $\leq$ and $\geq$ outcomes by combining the comparison functions from Section 2.1.7.

JOIN and GROUP BY. NSHEDB implements the JOIN (i.e., equi-joins $\bowtie$) and GROUP BY operators using equality comparisons across the entries in the targeted attribute(s).

GROUP BY. NSHEDB implements the GROUP BY operation by performing equality checks on each value in the group. First, NSHEDB retrieves metadata to determine the number of distinct values, n , in the grouping column. Since all values are already encoded into a contiguous integer range from 1 to n , NSHEDB directly operates on these unique IDs, ensuring consistent processing.

NSHEDB then performs equality checks for each distinct value in the target column, resulting in n ciphertexts, each containing encrypted “1”s for matching elements and encrypted “0”s for non-matching elements. For multi-column GROUP BY, NSHEDB generates permutations of the grouped values. For example, with columns $|G1| = 2$ and $|G2| = 3$, sorting by GROUP BY $G1, G2$ yields permutations like (1,1), (1,2), (1,3), (2,1), (2,2), (2,3). Reversing the order to GROUP BY $G2, G1$ results in different permutations.

NSHEDB then performs equality checks on the columns for each permutation and multiplies the results to identify elements that meet the specified conditions.

JOIN. NSHEDB supports (i) two-way equi-joins, (ii) k -way *star* joins ($k \geq 2$), and (iii) semi-/nested joins expressed with EXISTS/IN. Every variant is built from three HE steps: (1) Extract+Broadcast the reference-table key; (2) homomorphic equality to mask the fact table; and (3) Multiply the mask with projected attributes followed by Add across rows.

As an illustrative example, consider the simplified SQL query demonstrating a basic JOIN scenario:

```
1 SELECT size, brand FROM A, B WHERE A.key = B.key;
```

Figure 2 illustrates the detailed process: ❶ NSHEDB applies the Extract and Broadcast operations on each element from $B.key$, and ❷ performs equality checks between $A.key$ and each ciphertext created by Broadcast from $B.key$. ❸ This creates a set of HE ciphertexts indicating matching records in table A. Next, ❹ NSHEDB applies the Extract and Broadcast operations to the size and brand columns in table B, and ❺ multiplies these extracted attributes with the equality check results from step ❸. This process produces ciphertexts containing the matched entries for each attribute. Finally, ❻ NSHEDB aggregates these ciphertexts to yield the final JOIN output. Recognizing memory constraints, NSHEDB optimizes JOIN operations by merging intermediate JOIN and aggregation steps when possible, reducing intermediate storage requirements and improving computational efficiency. This optimization is particularly beneficial when handling more complex JOIN patterns such as multi-way equi-joins or nested JOIN structures.

Recognizing memory challenges in joining attributes from different tables, NSHEDB fuses the join and aggregation processes to minimize memory consumption wherever possible. This approach reduces data duplication and intermediate storage, streamlining operations and enhancing efficiency, particularly for aggregation-focused joins.

Operation	Multiplicative Depth
Equality Check	$\lceil \log(p - 1) \rceil$
Between/In	$\lceil \log(p - 1) \rceil + \lceil \frac{\log(k)}{p} \rceil$
Aggregation	$\frac{\log(n)}{p}$
Join	$\lceil \log(p - 1) \rceil + 1$
Group by/Order by	$\lceil \log(p - 1) \rceil$

Table 3. Multiplicative Depth Analysis of Operations in NSHEDB Implementation.

4.2.3 ORDER BY. NSHEDB sorts encrypted data homomorphically by using equality checks on all possible values in the encrypted domain. Since direct comparisons are not feasible in homomorphic encryption, NSHEDB generates equality masks for each distinct value, marking matching records with encrypted "1"s and others with encrypted "0"s. For each value, NSHEDB applies this mask to extract and aggregate matching records sequentially, constructing the final sorted sequence. This process is repeated in increasing order for ascending sorts and in reverse order for descending sorts. For multi-column sorting, NSHEDB first sorts by the highest-priority column, then applies the same process to resolve ties based on the next column.

4.3 Noise-Sensitive Query Optimization

NSHEDB optimizes query execution by minimizing noise accumulation and reducing the need for costly bootstrapping. Since multiplications contribute to noise growth the most, while additions increase noise more gradually—by a factor of $\frac{1}{p}$, where p represents the number of distinct values the system can encode—NSHEDB reduces multiplication depth by reordering operations and breaking down multiplication chains to limit noise growth at both the intra- and inter-operator levels.

4.3.1 The Noise Escalation of Operators and Optimizations. NSHEDB optimizes each query operator by evaluating multiplications independently, avoiding chained dependencies on previous results. Table 3 summarizes the noise levels for each operator. Below, we detail the optimizations for several critical operators.

Equality checks. NSHEDB translates equality checks into arithmetic operations using Equation 3. To minimize noise, NSHEDB reduces the number of required multiplications to $\lceil \log(p - 1) \rceil$ by applying the exponentiation by squaring method [34].

BETWEEN/IN. NSHEDB implements BETWEEN/IN by evaluating multiple equality checks independently, which limits noise escalation to $\lceil \log(p - 1) \rceil$ multiplications. During aggregation, NSHEDB employs a divide-and-conquer algorithm to perform additions independently, avoiding the reliance on sequential results like in a traditional chain of operations. further reducing noise to $\lceil \frac{\log(k)}{p} \rceil$, where k represents the size of the set for IN or the range for BETWEEN. This optimization limits the worst-case multiplication depth to $\lceil \log(p - 1) \rceil + \lceil \frac{\log(k)}{p} \rceil$.

AGGREGATION Operations. NSHEDB implements aggregation operations, including SUM and COUNT, through rotations and additions. In our HE system, the noise level introduced by rotations is negligible [59]. NSHEDB applies a divide-and-conquer algorithm to efficiently summarize vectors after different shift amounts. The noise level increases by $\frac{\log(n)}{p}$ as rotations cycle through the padding elements across all records within the column, where n represents the maximum number of elements stored in a ciphertext.




Fig. 3. Noise optimization example. Plan (b) executes the SELECT and JOIN operators independently and then intersects their outputs, saving roughly one equality check in multiplicative depth compared with the sequential execution in plan (a).

JOIN. NSHEDB implements the JOIN operation in two steps. First, it identifies matching entries, producing a ciphertext of matching indexes. Then, it multiplies the joining column with this resulting ciphertext. The first stage uses equality checks, which escalate the noise level by $\lceil \log(p-1) \rceil$ multiplications. The second stage performs one HE multiplication. Therefore, in the worst case, the noise level escalation in NSHEDB’s JOIN operation equals $\lceil \log(p-1) \rceil + 1$ multiplications.

GROUP BY/ORDER BY. NSHEDB implements the GROUP BY/ORDER BY operation by performing equality checks on permutations of possible combinations and counting the matching records. NSHEDB performs equality checks, using $\lceil \log(p-1) \rceil$ multiplications to generate the resulting ciphertexts.

4.3.2 Inter-Operator Optimizations. NSHEDB’s goal is to cap noise *growth* by restructuring the query plan. Instead of executing operators in a deep, sequential pipeline, NSHEDB splits them into independent paths, thereby reducing multiplicative depth and avoiding unnecessary bootstraps.

Predicate Pull-Up & Mask-Injection Tuning. Classic predicate push-down shrinks intermediate data early. Under homomorphic encryption, however, filtering leaves ciphertext sizes unchanged; irrelevant tuples are merely multiplied by *predicate masks* and turned into encrypted zeros. The dominant cost therefore shifts from I/O to the *multiplicative depth* that provokes noise growth and bootstrapping. Given a mask, such as a date filter ($\sigma_{1994 \leq O.date < 1995}$) or a join equality ($n_regionkey = r_regionkey$), the planner must decide where to inject it: *early*, pruning later work but lengthening the deepest multiplication chain, or *late*, duplicating multiplications while shortening that chain.

To balance this trade-off, NSHEDB proposes two optimizations, *Predicate Pull-Up* and *Mask-Injection Tuning*, where predicate pull-up postpones mask multiplications and mask-injection tuning injects mask multiplications only at the latest depth that still satisfies the noise budget. Every filter or comparison runs in its own sub-graph to yield a mask ciphertext, and operators such as GROUP BY execute directly on the original columns without prior filtering. In an FK-PK *chain*, the same mask can be injected at *any* nesting depth—from the outermost loop down to the innermost loop. Injecting lower shortens the longest path (consuming less noise budget) but forces the mask to be extracted and broadcast at each additional level. Inside a single join tree, the planner faces the same depth-versus-work trade-off that predicate pull-up solves across subgraphs.

Cost-and-Decision Model. Consider a plan fragment whose longest path consists of m homogeneous stages (filters, joins, or a mix). Let d_s denote the multiplicative depth *per stage* ($\approx$

$\lceil \log(p - 1) \rceil + 1$ for an equi-join, 1 for a single multiplication, etc.). If the mask is injected after stage i ($0 \leq i \leq m$),² then

$$D_i = (m - i)d_s, \quad \Delta MUL_i = i C_{mul}.$$

The incremental cost is

$$\text{Cost}(i) = (m - i)C_{mul} + \Delta MUL_i + \left[D_i > B_{\text{noise}} \right] C_{\text{boot}},$$

with $C_{\text{boot}} \gg C_{mul}$. The planner chooses

$$i^* = \begin{cases} \max\{i \mid D_i \leq B_{\text{noise}}\}, & \text{if feasible;} \\ m, & \text{otherwise (pay one bootstrap).} \end{cases}$$

A single knob thus governs both *predicate pull-up* ($m = 1$) and *mask-injection tuning* (larger m).

Noise-Optimized Plan Construction. A single depth-first traversal of the algebraic DAG applies three rewrite rules:

- (1) **R1 Mask isolation** — rewrite every $\sigma_p(R)$ as $\langle R, M_p \rangle$ to compute the predicate once and store its mask.
- (2) **R2 Independent evaluation** — when two masks touch disjoint attributes, split $\sigma_{p \wedge q}(R)$ into $\sigma_p(R) \cap \sigma_q(R)$, halving the longest chain.
- (3) **R3 Late injection** — push each mask past joins until the deepest path length is $\leq B_{\text{noise}}$ (the cost model yields i^*); inject once, then intersect or aggregate without further multiplications. Repeated R1 + R2 realize *Predicate Pull-Up*; R3 at depth i^* realizes *Mask-Injection Tuning*. The whole pass runs in $O(|V| + |E|)$ time and short-circuits the longest multiplication chains while keeping extra multiplications cheap.

Example (TPC-H Q4). Figure 3 visualizes the impact of our noise optimization on TPC-H Q4. In the baseline plan (Fig. 3 (a)), a sequential JOIN-WHERE pipeline incurs

$$D_{\text{orig}} = 3\lceil \log(p - 1) \rceil + \frac{\log n}{p} + 2.$$

After applying *Predicate Pull-Up* and *Mask-Injection Tuning* with $i^* = 1$ (Fig. 3 (b)), the date filter is injected only *after* the join, reducing the depth to

$$D_{\text{opt}} = \lceil \log(p - 1) \rceil + \frac{\log n}{p} + 2,$$

thereby eliminating two bootstraps whenever $D_{\text{opt}} \leq B_{\text{noise}}$.

Take-away. *Predicate pull-up* splits a plan into independent sub-graphs, while *mask-injection tuning* reorders multiplications inside a join path. Together, they respect the noise budget across hardware and parameter sets, choosing extra multiplications whenever they are cheaper than bootstrapping.

5 Result

We implemented NSHEDB and evaluated it on a real system using TPC-H queries. The prototype delivers up to a **659× speed-up** over the strongest prior HE database and consumes only **1.3 %** of its storage space. This section describes the evaluation platform, workload, and results.

² $i = 0$: inject immediately; $i = m$: inject at the end.

5.1 Experimental Methodology

Hardware and cryptographic parameters. All experiments run on a 16-core Intel Raptor-Lake CPU (5.1 GHz boost) with 128 GB RAM. Homomorphic operators use Microsoft SEAL [59]. Parameters follow the HE Standard [4]: polynomial degree $n=32768$, coefficient modulus $\log Q=881$, and plaintext modulus $p=65537$, yielding 128-bit security.

Baselines. We compare NSHEDB with the hybrid-HE engines HE³DB [63] and ArcEDB [76].

Dataset and query set. All eight TPC-H tables at scale factor 1 are loaded; the largest, LINEITEM, is sampled to 32 K rows and related tables are scaled proportionally. Note that this limited table size is primarily due to the space requirements of HE³DB and ArcEDB, whereas NSHEDB can support larger workloads. We benchmark nine widely studied queries—Q1, 4, 5³, 6, 8, 12, 14, 17, 19—that together span the key access patterns in TPC-H:

- **Scan + aggregation** (Q1, Q6, Q14) – pipelines with SUM, AVG, COUNT, GROUP BY, and ORDER BY.
- **Inner and multi-way equi-joins** (Q5, Q8, Q12, Q14, Q19) – two-table stars up to a six-table join, all with predicate push-down.
- **Semi- / nested joins** (Q4, Q17) – patterns expressed with EXISTS and correlated subqueries.
- **Complex predicates** (Q19) – disjunctions, BETWEEN, and IN tests across joined tables.

The workload therefore exercises scans, filters, arithmetic aggregates, semi-joins, three- to six-way joins, correlated subqueries, and multi-predicate filters—providing a balanced view of NSHEDB's performance without redundant cases.

5.2 Single-Threaded Performance and Noise-Sensitive Query Optimization

5.2.1 Primitive Operations Performance. Table 4 presents the absolute execution times for primitive database operations, demonstrating the efficiency of NSHEDB at the operator level. NSHEDB achieves sub-millisecond performance for COUNT, SUM, equality, and IN operations, while comparison operations require only 3.66ms. These improvements stem from avoiding bit-level transciphering and leveraging batch encoding.

Operation	HE ³ DB (ms)	ArcEDB (ms)	NSHEDB (ms)
COUNT	1.27	1.27	0.04
SUM	1.27	1.27	0.04
Equality (=)	283.33	16.00	0.09
Comparison (<, >)	150.83	16.00	3.66
Comparison ($\leq$, $\geq$)	143.07	16.00	3.66
BETWEEN	287.35	33.69	7.32
IN	283.33	16.00	0.09
GROUP BY	283.33	16.00	0.09

Table 4. Execution time of primitive database operations

³We do not show Q5 in the later result figures because the theoretical runtimes of the two state-of-the-art baselines cannot finish this query within reasonable timeframe (more than a year) on our testbed. NSHEDB achieves 13,000 $\times$ speedup overall the projected runtimes of HE³DB and ArcEDB. However, since the runtimes of the baselines are projected, we opt out to list Q5 in result figures for fairness.




Fig. 4. Speedup for basic database operations on NSHEDB compared to baselines. NSHEDB achieves speedups up to 3,148× over HE³DB.




Fig. 5. NSHEDB's Multi-threaded Performance. It achieves a speedup of 1.74×, 2.73×, and 4.11× with 2, 4, and 8 threads.




Fig. 6. Execution time (seconds) for TPC-H queries on 32K records comparing HE³DB, ArcEDB, and NSHEDB with and without noise optimization.

Figure 4 visualizes these improvements as speedup ratios, showing NSHEDB consistently outperforms both baselines across all operations. The dramatic speedups for equality-based operations (3,148× over HE³DB and 41× over ArcEDB) highlight the efficiency of our arithmetic-only approach.

5.2.2 TPC-H Query Performance. Figure 6 presents a single-threaded performance comparison of NSHEDB against HE³DB and ArcEDB across various TPC-H queries. Unlike HE³DB and ArcEDB, which rely on transcribing between bit-level and arithmetic HE operations, NSHEDB exclusively employs arithmetic HE, eliminating the overhead of costly HE scheme conversions.

The figure shows two configurations of NSHEDB: without and with noise-sensitive optimizations. Even without optimization, NSHEDB significantly outperforms both baselines. For example, Q1 completes in 477 seconds compared to 14,454 seconds for HE³DB (30× speedup) and 4,748 seconds

for ArcEDB (10 $\times$ speedup). Across all queries, NSHEDB without optimization achieves speedups between 17 $times$ and 96 $\times$ over HE³DB (average 41 $\times$).

With noise-sensitive optimizations enabled, performance improves dramatically. These optimizations restructure query plans to minimize multiplicative depth, often avoiding bootstrapping entirely. Q8 execution time drops from 8,423 seconds (HE³DB) and 3,351 seconds (ArcEDB) to just 178 seconds with NSHEDB—speedups of 47 $\times$ and 19 $\times$ respectively. Overall, NSHEDB with optimization achieves speedups ranging from 20 $times$ to 1,370 $times$ over HE³DB (average 659 $times$) and from 5 $times$ to 78 $times$ over ArcEDB (average 34 $times$).

Table 5 breaks down Q6 execution time to reveal the source of these improvements. Bootstrapping dominates in HE-based systems, consuming 97.5% of execution time in HE³DB (11,509s out of 11,802s) and 84.5% in ArcEDB (2,753s out of 3,257s). In stark contrast, NSHEDB eliminates bootstrapping entirely through noise-aware planning, completing the same query in just 590 seconds with only filtering and minimal aggregation overhead.

System	Boot.	Filter	Conv.	Agg.	Total
HE ³ DB	11,509 (97.5%)	251 (2.1%)	42 (0.4%)	0.01 (<0.01%)	11,802
ArcEDB	2,753 (84.5%)	430 (13.2%)	74 (2.3%)	0.21 (<0.01%)	3,257
NSHEDB	0 (0%)	589 (99.8%)	0 (0%)	1.41 (0.2%)	590

Table 5. Execution time breakdown (seconds) for TPC-H Q6 on 32K records.

5.2.3 Scalability Within Packing Limits. Table 6 demonstrates how packing affects scaling for Q6 with varying row counts up to the packing limit ($S = 32,768$):

Table 6. Q6 execution time scaling within packing limits

Rows	HE ³ DB (s)	NSHEDB (s)	Speedup
4K	1,541	590	2.6 $\times$
8K	3,044	590	5.2 $\times$
16K	6,514	590	11.0 $\times$
32K	11,802	590	20.0 $\times$

NSHEDB’s runtime remains constant at 590 seconds for all sizes $\leq S$ (fitting within one ciphertext), while HE³DB scales linearly with data size. This demonstrates packing’s efficiency: all operations within a single ciphertext have the same cost regardless of how many slots are filled. Beyond S rows, NSHEDB would exhibit stepwise growth—requiring 2 ciphertexts for $S + 1$ to $2S$ rows, 3 for $2S + 1$ to $3S$ rows, and so on.

5.3 Multi-Threaded Performance

NSHEDB exploits parallelism at multiple levels. At the data level, batch encoding enables simultaneous operations on multiple data points per ciphertext. At the query execution level, we applied OpenMP to assign independent predicates to threads, demonstrating NSHEDB’s parallel potential despite being orthogonal to our main contributions.

Figure 5 illustrates performance scaling with increasing thread counts, achieving speedups of 1.74 $\times$, 2.73 $\times$, and 4.11 $\times$ with 2, 4, and 8 threads, respectively. Scaling is not linear due to three factors:

Limited parallelism: Queries with few predicates (e.g., Q1, Q6) leave threads idle beyond 4 cores.

Bandwidth saturation: Inequality checks generate intermediate ciphertexts and require lookup table accesses (BSGS optimization), creating memory pressure. Equality tests use simple squaring with minimal overhead.

Operation dependence: Equality-heavy queries (Q4, Q12) scale better than inequality-heavy ones (Q8, Q19) due to lower memory pressure. Joins outscale filters despite similar complexity, using more equality checks.

Future optimizations could explore memory-efficient HE schemes and predicate-aware scheduling to improve scalability beyond eight threads.

5.4 Memory and Storage Efficiency




Fig. 7. Storage and memory efficiency of NSHEDB. (a) Storage overhead across different record counts. NSHEDB reduces storage overhead by 73 $\times$, requiring only 1.4% of the space used by HE³DB and ArcEDB. (b) Breakdown of in-memory object types. NSHEDB reduces memory overhead by up to 1.6 $\times$ compared to the other systems.

NSHEDB optimizes memory and storage efficiency through improved data encoding schemes and reduced runtime memory usage. Figure 7(a) compares the storage footprint of NSHEDB with HE³DB and ArcEDB. To store 32K rows (16-bit integers), NSHEDB requires only 7.4 MB, whereas HE³DB and ArcEDB require 537 MB—73 $\times$ more than NSHEDB.

Figure 7(b) presents a breakdown of memory usage during Q6 execution. NSHEDB reduces overall memory overhead by 1.6 $\times$ despite slightly higher intermediate memory usage. This reduction stems from two key factors: 1) Bit-level storage incurs significantly higher overhead, expanding the total memory footprint. 2) HE³DB and ArcEDB require secret keys for both TFHE and CKKS, as well as CKKS scheme integration with TFHE, along with additional context data for both HE schemes. By

eliminating transciphering and maintaining a more efficient memory footprint, NSHEDB enables more scalable and practical homomorphic query execution.

6 Other Related Work

6.1 Parallel Query Optimizations

NSHEDB's query optimization shares similarities with existing parallel optimization techniques but differs by focusing on reducing noise rather than raw performance.

Existing works [2] and [73] introduce bushy-tree-based optimizations for JOIN operations, aiming to reduce the number of MapReduce jobs and reorganize the sequence of operations. Similarly, prior research [16] proposes a query chopping method for multi-threaded environments, where tasks are distributed across threads for concurrent processing.

However, these approaches primarily target raw data, while NSHEDB must contend with encrypted data and noise accumulation. As a result, NSHEDB introduces additional computational steps to minimize noise, a trade-off that slightly increases processing time but avoids the costly bootstrapping process, ultimately leading to better overall efficiency in a homomorphic encryption context.

6.2 Multiplicative Depth Optimization

Several works focus on minimizing multiplicative depth in homomorphic encryption, which is critical for optimizing performance. For boolean circuits, Aubry et al. [7] propose methods to minimize depth by rewriting sets of operations like XOR and AND, but these optimizations are specific to bit-level HE and may not translate well to word-level encryption schemes like BFV used in NSHEDB.

Bonte et al. [13] present an optimized solution for homomorphic string search with constant multiplicative depth, focusing on equality checks. However, this approach is limited to specific query types and does not address more general database operations.

Other approaches, like the homomorphic sorting algorithm in [19], prioritize depth optimization for sorting but do not extend beyond the sorting operation to subsequent computations.

Some neural network frameworks in [25, 42, 44], introduce depth-optimized methods for specific machine learning operations like convolution layers. However, these techniques focus on neural networks rather than general-purpose database systems, making them challenging to adapt for NSHEDB's query processing needs.

7 Conclusion

Our work introduces NSHEDB, a secure query processing engine that rethinks how homomorphic encryption (HE) can be deployed in practical database systems. Rather than relying on costly transciphering or trusted hardware, NSHEDB uses word-level HE and introduces novel algorithms for executing common database operations entirely within the encrypted domain. It promotes a new approach to HE-based query execution by translating logical operations into efficient arithmetic equivalents, while prioritizing noise growth control over simply minimizing operation counts.

By integrating noise-aware query planning with carefully optimized encodings, NSHEDB avoids excessive bootstrapping and significantly improves computational efficiency. Our evaluation on TPC-H workloads shows that NSHEDB achieves up to 1370× speedup over state-of-the-art HE-based systems, with only 1.4% storage overhead, and no reliance on trusted execution environments, key release, or scheme-switching.

This work demonstrates that secure, privacy-preserving query processing with HE is not only theoretically sound but also practical and performant. NSHEDB sets a new direction for building encrypted database systems grounded in principled cryptographic design and system-level innovation.

Acknowledgments

The authors would like to thank the anonymous reviewers for their helpful comments. This work was sponsored by the National Science Foundation (NSF) award, CNS-2231877, and Intel.

References

- [1] Ishtiyaque Ahmad, Divyakant Agrawal, Amr El Abbadi, and Trinabh Gupta. Pantheon: Private retrieval from public key-value store. *Proceedings of the VLDB Endowment*, 16(4):643–656, 2022.
- [2] Rafi Ahmed, Rajkumar Sen, Meikel Poess, and Sunil Chakkappen. Of snowstorms and bushy trees. *Proceedings of the VLDB Endowment*, 7(13):1452–1461, 2014.
- [3] Ahmad Al Badawi and Yuriy Polyakov. Demystifying bootstrapping in fully homomorphic encryption. *Cryptology ePrint Archive*, 2023.
- [4] Martin Albrecht, Melissa Chase, Hao Chen, Jintai Ding, Shafi Goldwasser, Sergey Gorbunov, Shai Halevi, Jeffrey Hoffstein, Kim Laine, Kristin Lauter, Satya Lokam, Daniele Micciancio, Dustin Moody, Travis Morrison, Amit Sahai, and Vinod Vaikuntanathan. Homomorphic encryption standard. *Cryptology ePrint Archive*, Paper 2019/939, 2019. <https://eprint.iacr.org/2019/939>.
- [5] Panagiotis Antonopoulos, Arvind Arasu, Kunal D. Singh, Ken Eguro, Nitish Gupta, Rajat Jain, Raghav Kaushik, Hanuma Kodavalla, Donald Kossmann, Nikolas Ogg, Ravi Ramamurthy, Jakub Szymaszek, Jeffrey Trimmer, Kapil Vaswani, Ramarathnam Venkatesan, and Mike Zwillig. Azure sql database always encrypted. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’20, page 1511–1525, New York, NY, USA, 2020. Association for Computing Machinery.
- [6] Arvind Arasu, Ken Eguro, Manas Joglekar, Raghav Kaushik, Donald Kossmann, and Ravi Ramamurthy. Transaction processing on confidential data using cipherbase. In *2015 IEEE 31st International Conference on Data Engineering*, pages 435–446. IEEE, 2015.
- [7] Pascal Aubry, Sergiu Carpov, and Renaud Sirdey. Faster homomorphic encryption is not enough: Improved heuristic for multiplicative depth minimization of boolean circuits. In *Topics in Cryptology—CT-RSA 2020: The Cryptographers’ Track at the RSA Conference 2020, San Francisco, CA, USA, February 24–28, 2020, Proceedings*, pages 345–363. Springer, 2020.
- [8] Ahmad Al Badawi, Jack Bates, Flavio Bergamaschi, David Bruce Cousins, Saroja Erabelli, Nicholas Genise, Shai Halevi, Hamish Hunt, Andrey Kim, Yongwoo Lee, Zeyu Liu, Daniele Micciancio, Ian Quah, Yuriy Polyakov, Saraswathy R.V., Kurt Rohloff, Jonathan Saylor, Dmitriy Sponitsky, Matthew Triplett, Vinod Vaikuntanathan, and Vincent Zucca. Openfhe: Open-source fully homomorphic encryption library. *Cryptology ePrint Archive*, Paper 2022/915, 2022. <https://eprint.iacr.org/2022/915>.
- [9] Youngjin Bae, Jung Hee Cheon, Jaehyung Kim, Jai Hyun Park, and Damien Stehlé. Hermes: Efficient ring packing using mlwe ciphertexts and application to transciphering. In *Annual International Cryptology Conference*, pages 37–69. Springer, 2023.
- [10] Sumeet Bajaj and Radu Sion. Trustedd: a trusted hardware based database with privacy and data confidentiality. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, pages 205–216, 2011.
- [11] Johes Bater, Gregory Elliott, Craig Eggen, Satyender Goel, Abel Kho, and Jennie Rogers. Smcql: secure querying for federated databases. *arXiv preprint arXiv:1606.06808*, 2016.
- [12] Dan Boneh, Craig Gentry, Shai Halevi, Frank Wang, and David J Wu. Private database queries using somewhat homomorphic encryption. In *Applied Cryptography and Network Security: 11th International Conference, ACNS 2013, Banff, AB, Canada, June 25–28, 2013. Proceedings 11*, pages 102–118. Springer, 2013.
- [13] Charlotte Bonte and Ilia Iliashenko. Homomorphic string search with constant multiplicative depth. In *Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop*, pages 105–117, 2020.
- [14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)*, 6(3):1–36, 2014.
- [15] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. *SIAM Journal on computing*, 43(2):831–871, 2014.
- [16] Sebastian Breß, Henning Funke, and Jens Teubner. Robust query processing in co-processor-accelerated databases. In *Proceedings of the 2016 International Conference on Management of Data*, pages 1891–1906, 2016.

- [17] Martin Burkhart, Mario Strasser, Dilip Many, and Xenofontas Dimitropoulos. {SEPIA}:: {Privacy-Preserving} aggregation of {Multi-Domain} network events and statistics. In *19th USENIX Security Symposium (USENIX Security 10)*, 2010.
- [18] Gizem S Cetin, Yarkin Doroz, Berk Sunar, and William J Martin. Arithmetic using word-wise homomorphic encryption. *Cryptology ePrint Archive*, 2015.
- [19] Gizem S Çetin, Yarkin Doröz, Berk Sunar, and Erkey Savaş. Depth optimized efficient homomorphic sorting. In *Progress in Cryptology–LATINCRYPT 2015: 4th International Conference on Cryptology and Information Security in Latin America, Guadalajara, Mexico, August 23–26, 2015, Proceedings 4*, pages 61–80. Springer, 2015.
- [20] Hao Chen, Wei Dai, Miran Kim, and Yongsoo Song. Efficient homomorphic conversion between (ring) lwe ciphertexts. In *International Conference on Applied Cryptography and Network Security*, pages 460–479. Springer, 2021.
- [21] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. Homomorphic encryption for arithmetic of approximate numbers. In *Advances in Cryptology–ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3–7, 2017, Proceedings, Part I 23*, pages 409–437. Springer, 2017.
- [22] Seonyoung Cheon, Yongwoo Lee, Dongkwan Kim, Ju Min Lee, Sunchul Jung, Taekyung Kim, Dongyoon Lee, and Hanjun Kim. {DaCapo}: Automatic bootstrapping management for efficient fully homomorphic encryption. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 6993–7010, 2024.
- [23] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Tfhe: fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, 2020.
- [24] David Bruce Cousins, Kurt Rohloff, and Daniel Sumorok. Designing an fpga-accelerated homomorphic encryption co-processor. *IEEE Transactions on Emerging Topics in Computing*, 5(2):193–206, 2016.
- [25] Eric Crockett. A low-depth homomorphic circuit for logistic regression model training. *Cryptology ePrint Archive*, 2020.
- [26] Natacha Crooks, Matthew Burke, Ethan Cecchetti, Sitar Harel, Rachit Agarwal, and Lorenzo Alvisi. Obladi: Oblivious serializable transactions in the cloud. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 727–743, 2018.
- [27] Shujie Cui, Xiangfu Song, Muhammad Rizwan Asghar, Steven D Galbraith, and Giovanni Russello. Privacy-preserving dynamic symmetric searchable encryption with controllable leakage. *ACM Transactions on Privacy and Security (TOPS)*, 24(3):1–35, 2021.
- [28] Leo de Castro, Rashmi Agrawal, Rabia Yazicigil, Anantha Chandrakasan, Vinod Vaikuntanathan, Chiraag Juvekar, and Ajay Joshi. Does fully homomorphic encryption need compute acceleration? *arXiv preprint arXiv:2112.06396*, 2021.
- [29] Léo Ducas and Daniele Micciancio. Fhew: bootstrapping homomorphic encryption in less than a second. In *Advances in Cryptology–EUROCRYPT 2015: 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26–30, 2015, Proceedings, Part I 34*, pages 617–640. Springer, 2015.
- [30] Saba Eskandarian and Matei Zaharia. Oblidb: Oblivious query processing for secure databases. *arXiv preprint arXiv:1710.00458*, 2017.
- [31] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 169–178, 2009.
- [32] Alvin Oliver Glova, Itir Akgun, Shuangchen Li, Xing Hu, and Yuan Xie. Near-data acceleration of privacy-preserving biomarker search with 3d-stacked memory. In *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 800–805. IEEE, 2019.
- [33] Charles Gouert, Dimitris Mouris, and Nektarios Tsoutsos. Sok: New insights into fully homomorphic encryption libraries via standardized benchmarks. *Proceedings on privacy enhancing technologies*, 2023.
- [34] Shay Gueron. Efficient software implementations of modular exponentiation. *Journal of Cryptographic Engineering*, 2:31–43, 2012.
- [35] Saransh Gupta and Tajana Šimunić Rosing. Accelerating fully homomorphic encryption with processing in memory. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pages 1335–1338. IEEE, 2021.
- [36] Zhian He, Wai Kit Wong, Ben Kao, David Wai Lok Cheung, Rongbin Li, Siu Ming Yiu, and Eric Lo. Sdb: A secure query processing system with data interoperability. *Proceedings of the VLDB Endowment*, 8(12):1876–1879, 2015.
- [37] Ilia Iliashenko and Vincent Zucca. Faster homomorphic comparison operations for bgv and bfv. *Proceedings on Privacy Enhancing Technologies*, 2021(3):246–264, 2021.
- [38] Myungsun Kim, Hyung Tae Lee, San Ling, Shu Qin Ren, Benjamin Hong Meng Tan, and Huaxiong Wang. Better security for queries on encrypted databases. *Cryptology ePrint Archive*, 2016.
- [39] Sangpyo Kim, Jongmin Kim, Michael Jaemin Kim, Wonkyung Jung, John Kim, Minsoo Rhu, and Jung Ho Ahn. Bts: An accelerator for bootstrappable fully homomorphic encryption. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*, pages 711–725, 2022.

- [40] Nirajan Koirala, Jonathan Takeshita, Jeremy Stevens, and Taeho Jung. Summation-based private segmented membership test from threshold-fully homomorphic encryption. *Proceedings on Privacy Enhancing Technologies*, 2024.
- [41] Peeter Laud and Alisa Pankova. Privacy-preserving record linkage in large databases using secure multiparty computation. *BMC medical genomics*, 11:33–46, 2018.
- [42] Eunsang Lee, Joon-Woo Lee, Junghyun Lee, Young-Sik Kim, Yongjune Kim, Jong-Seon No, and Woosuk Choi. Low-complexity deep convolutional neural networks on fully homomorphic encryption using multiplexed parallel convolutions. In *International Conference on Machine Learning*, pages 12403–12422. PMLR, 2022.
- [43] John Liagouris, Vasiliki Kalavri, Muhammad Faisal, and Mayank Varia. Secrecy: Secure collaborative analytics on secret-shared data. *arXiv preprint arXiv:2102.01048*, 2021.
- [44] Qian Lou and Lei Jiang. Hemet: A homomorphic-encryption-friendly privacy-preserving mobile neural network architecture. In *International conference on machine learning*, pages 7102–7110. PMLR, 2021.
- [45] Di Lu, Minqiang Shi, Xindi Ma, Ximeng Liu, Rui Guo, Tianfang Zheng, Yulong Shen, Xuewen Dong, and Jianfeng Ma. Smaug: A tee-assisted secured sqlite for embedded systems. *IEEE Transactions on Dependable and Secure Computing*, 20(5):3617–3635, 2023.
- [46] Wen-jie Lu, Zhicong Huang, Cheng Hong, Yiping Ma, and Hunter Qu. Pegasus: bridging polynomial and non-polynomial evaluations in homomorphic encryption. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1057–1073. IEEE, 2021.
- [47] Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE symposium on security and privacy (SP)*, pages 19–38. IEEE, 2017.
- [48] Muhammad Haris Mughees, Hao Chen, and Ling Ren. Onionpir: Response efficient single-server pir. In *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security*, pages 2292–2306, 2021.
- [49] Muhammad Naveed, Seny Kamara, and Charles V Wright. Inference attacks on property-preserving encrypted databases. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 644–655, 2015.
- [50] Simon Oya and Florian Kerschbaum. Hiding the access pattern is not enough: Exploiting search pattern leakage in searchable encryption. In *30th USENIX security symposium (USENIX Security 21)*, pages 127–142, 2021.
- [51] Rishabh Poddar, Tobias Boelter, and Raluca Ada Popa. Arx: An encrypted database using semantically secure encryption. *Cryptology ePrint Archive*, 2016.
- [52] Raluca Ada Popa, Catherine MS Redfield, Nickolai Zeldovich, and Hari Balakrishnan. Cryptdb: Protecting confidentiality with encrypted query processing. In *Proceedings of the twenty-third ACM symposium on operating systems principles*, pages 85–100, 2011.
- [53] Christian Priebe, Kapil Vaswani, and Manuel Costa. Enclavedb: A secure database using sgx. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 264–278, 2018.
- [54] Dayane Reis, Jonathan Takeshita, Taeho Jung, Michael Niemier, and Xiaobo Sharon Hu. Computing-in-memory for performance and energy-efficient homomorphic encryption. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 28(11):2300–2313, 2020.
- [55] Xuanle Ren, Le Su, Zhen Gu, Sheng Wang, Feifei Li, Yuan Xie, Song Bian, Chao Li, and Fan Zhang. Heda: Multi-attribute unbounded aggregation over homomorphically encrypted database. *Proceedings of the VLDB Endowment*, 16(4):601–614, 2022.
- [56] Pedro S Ribeiro, Nuno Santos, and Nuno O Duarte. Dbstore: A trustzone-backed database management system for mobile applications. In *ICETE (2)*, pages 562–569, 2018.
- [57] Sujoy Sinha Roy, Furkan Turan, Kimmo Jarvinen, Frederik Vercauteren, and Ingrid Verbauwhede. Fpga-based high-performance parallel architecture for homomorphic computing on encrypted data. In *2019 IEEE International symposium on high performance computer architecture (HPCA)*, pages 387–398. IEEE, 2019.
- [58] Nikola Samardzic, Axel Feldmann, Aleksandar Krastev, Srinivas Devadas, Ronald Dreslinski, Christopher Peikert, and Daniel Sanchez. F1: A fast and programmable accelerator for fully homomorphic encryption. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 238–252, 2021.
- [59] Microsoft SEAL (release 4.1). <https://github.com/Microsoft/SEAL>, January 2023. Microsoft Research, Redmond, WA.
- [60] Maryam Sepehri, Stelvio Cimato, and Ernesto Damiani. Privacy-preserving query processing by multi-party computation. *The Computer Journal*, 58(10):2195–2212, 2015.
- [61] Hossein Shafagh, Anwar Hithnawi, Andreas Dröschner, Simon Duquennoy, and Wen Hu. Talos: Encrypted query processing for the internet of things. In *Proceedings of the 13th ACM conference on embedded networked sensor systems*, pages 197–210, 2015.
- [62] Reza Shokri and Vitaly Shmatikov. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pages 1310–1321, 2015.
- [63] Haowen Pan, Ran Mao, Zian Zhao, Yier Jin, Song Bian, Zhou Zhang, and Zhenyu Guan. He³db: An efficient and elastic encrypted database via arithmetic-and-logic fully homomorphic encryption. In *Proceedings of the 2023 ACM SIGSAC*

- Conference on Computer and Communications Security (CCS '23)*, November 26–30, 2023, Copenhagen, Denmark. ACM, 2023.
- [64] Yuanyuan Sun, Sheng Wang, Huorong Li, and Feifei Li. Building enclave-native storage engines for practical encrypted databases. *Proceedings of the VLDB Endowment*, 14(6):1019–1032, 2021.
 - [65] Benjamin Hong Meng Tan, Hyung Tae Lee, Huaxiong Wang, Shuqin Ren, and Khin Mi Mi Aung. Efficient private comparison queries over encrypted databases using fully homomorphic encryption with finite fields. *IEEE Transactions on Dependable and Secure Computing*, 18(6):2861–2874, 2020.
 - [66] Michiel Van Beirendonck, Jan-Pieter D’Anvers, Furkan Turan, and Ingrid Verbauwhede. Fpt: A fixed-point accelerator for torus fully homomorphic encryption. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 741–755, 2023.
 - [67] Dhinakaran Vinayagamurthy, Alexey Gribov, and Sergey Gorbunov. Stealthdb: a scalable encrypted database with full sql query support. *Proc. Priv. Enhancing Technol.*, 2019(3):370–388, 2019.
 - [68] Nikolaj Volgushev, Malte Schwarzkopf, Ben Getchell, Mayank Varia, Andrei Lapets, and Azer Bestavros. Conclave: secure multi-party computation on big data. In *Proceedings of the Fourteenth EuroSys Conference 2019*, pages 1–18, 2019.
 - [69] Sheng Wang, Yiran Li, Huorong Li, Feifei Li, Chengjin Tian, Le Su, Yanshan Zhang, Yubing Ma, Lie Yan, Yuanyuan Sun, et al. Operon: An encrypted database for ownership-preserving data management. *Proceedings of the VLDB Endowment*, 15(12):3332–3345, 2022.
 - [70] Yilei Wang and Ke Yi. Secure yannakakis: Join-aggregate queries over private data. In *Proceedings of the 2021 International Conference on Management of Data*, pages 1969–1981, 2021.
 - [71] Yongzhi Wang, Yulong Shen, Cuicui Su, Jiawen Ma, Lingtong Liu, and Xuewen Dong. Cryptsqlite: Sqlite with high data security. *IEEE Transactions on Computers*, 69(5):666–678, 2020.
 - [72] Andrew Wiles. Modular elliptic curves and fermat’s last theorem. *Annals of mathematics*, 141(3):443–551, 1995.
 - [73] Sai Wu, Feng Li, Sharad Mehrotra, and Beng Chin Ooi. Query optimization for massively parallel data processing. In *Proceedings of the 2nd ACM Symposium on Cloud Computing*, pages 1–13, 2011.
 - [74] Xun Yi, Mohammed Golam Kaosar, Russell Paulet, and Elisa Bertino. Single-database private information retrieval from fully homomorphic encryption. *IEEE Transactions on Knowledge and Data Engineering*, 25(5):1125–1134, 2012.
 - [75] Shufan Zhang, Xi He, Ashish Kundu, Sharad Mehrotra, and Shantanu Sharma. Secure normal form: Mediation among cross cryptographic leakages in encrypted databases. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pages 5560–5573. IEEE, 2024.
 - [76] Zhou Zhang, Song Bian, Zian Zhao, Ran Mao, Haoyi Zhou, Jiafeng Hua, Yier Jin, and Zhenyu Guan. Arcedb: An arbitrary-precision encrypted database via (amortized) modular homomorphic encryption. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 4613–4627, 2024.

Received July 2025; revised October 2025; accepted November 2025