QAE-BAC: Achieving **Q**uantifiable **A**nonymity and **E**fficiency in **B**lockchain-Based **A**ccess **C**ontrol with Attribute

Jie Zhang, *Student Member, IEEE*, Xiaohong Li, *Member, IEEE*, Mengke Zhang, Ruitao Feng, Shanshan Xu, Zhe Hou, and Guangdong Bai, *Member, IEEE*

***Abstract*—Blockchain-based Attribute-Based Access Control (BC-ABAC) offers a decentralized paradigm for secure data governance but faces two inherent challenges: the transparency of blockchain ledgers threatens user privacy by enabling re-identification attacks through attribute analysis, while the computational complexity of policy matching clashes with blockchain's performance constraints. Existing solutions, such as those employing Zero-Knowledge Proofs (ZKPs), often incur high overhead and lack measurable anonymity guarantees, while efficiency optimizations frequently ignore privacy implications. To address these dual challenges, this paper proposes QAE-BAC (Quantifiable Anonymity and Efficiency in Blockchain-Based Access Control with Attribute). QAE-BAC introduces a formal (r, t)-anonymity model to dynamically quantify the re-identification risk of users based on their access attributes and history. Furthermore, it features an Entropy-Weighted Path Tree (EWPT) that optimizes policy structure based on real-time anonymity metrics, drastically reducing policy matching complexity. Implemented and evaluated on Hyperledger Fabric, a superior balance between privacy and performance is demonstrated by QAE-BAC. Experimental results demonstrate effective mitigation of re-identification risks and outperforms state-of-the-art baselines, achieving up to an 11x improvement in throughput and an 87% reduction in latency, proving its practicality for privacy-sensitive decentralized applications.**

***Index Terms*—Anonymity Quantification, Attribute-Based Access Control, Blockchain, Entropy-Weighted Path Tree, Privacy Preservation.**

I. Introduction

The exponential growth of data generation and exchange in modern digital ecosystems has made robust access control a critical cybersecurity cornerstone. Attribute-Based Access Control (ABAC) [1] has emerged as the preeminent model for dynamic and distributed environments, offering superior flexibility and fine-grained control over traditional role-based models. However, conventional ABAC deployments often rely on centralized authorities, introducing vulnerabilities like single points of failure and ambiguous data sovereignty [2], [3].

Blockchain technology, with its core tenets of decentralization, immutability, and transparency, presents a compelling solution to this mismatch [4], [5], [6]. By executing ABAC policies through smart contracts, researchers have built decentralized access control systems that reduce reliance on a single trusted third party by shifting trust to a decentralized protocol and its immutable code [7], [8], [9], [10], [11]. This fusion, however, intensifies two fundamental and deeply intertwined challenges that threaten the viability of Blockchain-based Attribute-Based Access Control (BC-ABAC) in practice.

First, the **privacy-transparency paradox** becomes severe. Although the public records of blockchain provide transparency for auditability, this transparency also exposes sensitive attributes required for ABAC policy evaluation [12], [13]. The very attributes required for fine-grained policy evaluation (e.g., 'role', 'clearance', 'affiliation') are often sensitive. When recorded on an immutable ledger, they form a rich, permanent dataset for adversaries. Through sophisticated linkage attacks and frequency analysis, malicious actors can de-anonymize users, trace their behavior across transactions, and infer sensitive information [14], [15]. This risk is particularly acute in systems with many fine-grained attributes, where certain combinations can act as quasi-identifiers, uniquely pinpointing individuals within a small user pool [16].

Second, the **performance-complexity gap** is widened. ABAC inherently suffers from the "attribute explosion" problem [17], where the growing number and complexity of attributes and policies make the request-to-policy matching process computationally expensive, leading to increased authorization latency [18], [19]. Blockchain platforms, often characterized by lower transaction throughput and higher consensus latency compared to centralized systems, act as a performance bottleneck, dramatically amplifying this inherent complexity [20], [21]. Consequently, the combined system struggles to meet the low-latency, high-throughput demands of large-scale, real-world applications [22].

Limitations of Existing Work & Our Motivation: Existing research has made considerable strides but often addresses these challenges in isolation, leading to a fragmented landscape. On one hand, **privacy-focused approaches** frequently employ advanced cryptographic techniques like Zero-Knowledge Proofs (ZKPs) [11], [23], [24] or anonymization methods [16], [25], [26]. While these approaches can hide attribute values or identities, they often introduce substantial computational overhead, lack a mechanism for quantitatively measuring the achieved level of anonymity and may not ade-

This work is supported in part by the National Key Research and Development Program of China under Grant 2023YFB3107103, in part by the National Natural Science Foundation of China under Grant 62262073, 62332005.

Jie Zhang, Xiaohong Li and Mengke Zhang are with the College of Intelligence and Computing, Tianjin University, Tianjin, China. (email: {jackzhang, xiaohongli, mengkezhangcs}@tju.edu.cn).

Mengke Zhang is also with Shanghai ZTE Software Co., Ltd., China.

Ruitao Feng is with the Faculty of Science and Engineering, Southern Cross University, Australia (e-mail: ruitao.feng@scu.edu.au).

Shanshan Xu is with the School of Geographic Sciences, East China Normal University, Shanghai, China. (email: s.xu.ecnu@gmail.com).

Zhe Hou is with the School of Information and Communication Technology, Griffith University, Nathan, Australia. (email: z.hou@griffith.edu.au).

Guangdong Bai is with the Department of Computer Science, City University of Hong Kong, Hong Kong, China. (e-mail: baiguangdong@gmail.com).

Jie Zhang and Mengke Zhang are contributed equally to this work.

Ruitao Feng and Guangdong Bai are the corresponding authors.

quately protect against privacy leaks from dynamic access patterns [15]. On the other hand, **efficiency-focused approaches** optimize policy retrieval and matching [17], [27], [28], [29] but are fundamentally *privacy-agnostic*; their design does not consider whether optimizing for speed might inadvertently simplify an attacker's task of re-identifying users, potentially exacerbating the privacy risks they ignore. This clear dichotomy highlights a critical research gap: **The absence of a *holistic, co-designed framework* where *continuous, quantifiable anonymity assessment actively guides and informs performance optimization*. Without this synergy, systems are forced to choose between privacy and performance, or suffer the penalties of both.** Bridging this gap is the primary motivation for this work.

Our Approach and Novelty: This paper proposes QAE-BAC (**Q**uantifiable **A**nonymity and **E**fficiency in **B**lockchain-Based **A**ccess **C**ontrol with Attribute), a novel framework that breaks the prevailing privacy-efficiency trade-off through deep integration. The core novelty of QAE-BAC is established in its **closed-loop feedback system**, which actively uses real-time privacy metrics to govern performance optimization. This tight coupling ensures that the system is not just fast, but **responsibly fast**; it is not just private, but **efficiently private**. The proposed framework provides a foundational shift towards building scalable, efficient, and truly privacy-preserving decentralized data governance systems. A comprehensive version of this work, including detailed algorithms and extended analysis, is available in arXiv [30].

Contributions: The main contributions of this work are four-fold:

- The novel QAE-BAC framework is proposed to cohesively integrate continuous anonymity quantification with privacy-aware policy optimization for BC-ABAC. This deep integration establishes a feedback loop where live anonymity scores directly guide the optimization process, effectively breaking the prevailing privacy-efficiency trade-off.
- The privacy threat in BC-ABAC is formalized by defining a dynamic "credential subject space" and introducing an (r, t)-anonymity model. This model provides a quantitative, real-time metric for assessing re-identification risk [14], addressing a critical gap in existing privacy solutions that offer protection but no measure of its strength.
- The Entropy-Weighted Path Tree (EWPT) structure and a corresponding fast authorization algorithm are designed. The innovation of EWPT is that its weights and structure are derived from real-time anonymity metrics and access patterns, achieving a fundamental reduction in policy matching time complexity ($O(\log n)$) while ensuring that optimization does not create new privacy vulnerabilities.
- A prototype of QAE-BAC is implemented on Hyperledger Fabric and extensive experiments are conducted using real-world attribute datasets. Results demonstrate that QAE-BAC effectively maintains high subject anonymity under various conditions and outperforms state-of-the-art baselines, achieving up to an 11x improvement in throughput and an 87% reduction in latency.

II. Related Work

This section reviews related work in blockchain-based access control, privacy enhancements, and efficient policy management, highlighting both advancements and limitations that motivate the approach.

A. Blockchain-Based Access Control

Zhang et al. [7] and Liu et al. [8] pioneered blockchain-based ABAC implementations, demonstrating decentralized authorization but revealing on-chain storage and privacy challenges. Subsequent works proposed architectural optimizations: Xia et al. [31] and Yang et al. [32] developed hybrid storage models; Tong et al. [21] and Zhang et al. [33] explored sharding and multi-chain architectures for cross-domain scalability. Liu et al. [34] and Xing et al. [2] further advanced sharding-based approaches for cross-domain data sharing and decentralized access control, respectively. However, these approaches primarily addressed architectural concerns while overlooking privacy implications of attribute exposure. The persistent "attribute explosion" problem [35], [17] creates performance bottlenecks that existing solutions fail to adequately resolve.

B. Privacy Enhancements

Privacy research follows two main strands. Cryptographic approaches, particularly ZKP-based solutions by Wu et al. [23], Hu et al. [24], and Lin et al. [36], provide strong attribute hiding but incur substantial overhead and lack quantifiable anonymity metrics. Anonymization techniques, including Anonymous Access Control (AAC) schemes by Yuen et al. [25], Lanus et al. [16], and Fang et al. [26], decouple identity from access rights but often adopt static protection models. Recent work by Zhang et al. [15] considers access history impacts but does not fully address attribute distribution. Panda et al. [37] explored contextual anonymity, while Santis et al. [38] analyzed information leakage from policy structures. A common limitation is the absence of dynamic, measurable anonymity assessment.

C. Efficient Policy Management

Policy efficiency research addresses "attribute explosion" through various optimizations. Karimi et al. [27] investigated automated policy learning, Geng et al. [28] focused on conflict resolution, and Bai et al. [29] developed optimized data structures for policy retrieval. In blockchain contexts, Duan et al. [9] explored trusted execution, while Zhang et al. [11] proposed cross-domain optimizations. Hao et al. [4] surveyed blockchain access control architectures. These approaches improve efficiency but are predominantly privacy-agnostic, optimizing for performance without considering re-identification risks introduced by policy organization.

D. Comparison with State-of-the-Art

Tables I and II summarize comparisons between QAE-BAC and related works. QAE-BAC uniquely integrates quantifiable

TABLE I: Comparison of QAE-BAC with related works.

Work	Decentralized	Fine-Grained	Privacy Protection	Quantifiable Anonymity	Efficient Authorization
[7]	✓	✓	✗	✗	✗
[8]	✓	✓	✗	✗	✗
[21]	✓	✗	✗	✗	✓
[33]	✓	✓	✗	✗	✓
[23]	✓	✓	✓ (ZKP)	✗	✗
[24]	✓	✓	✓ (ZKP)	✗	✗
[26]	✓	✓	✓ (AAC)	✗	✗
[16]	✗	✓	✓ (AAC)	-	✗
[15]	✗	✓	✓(AAC)	✓	✗
[28]	✗	✓	✗	✗	✓
QAE-BAC	✔	✔	✔(AAC)	✔	✔

✓: supported, ✗: not supported, -: unknown, AAC: Anonymous Access Control

TABLE II: Qualitative and Quantitative Comparison of Privacy-Efficiency Trade-off

Approach	Privacy	Overhead	Complexity	Limitation
BC-ABAC [8]	Low	Low	$O(n)$	High re-identification risk, attributes exposed.
ZKP-based [24]	Very High (Value Hiding)	Very High	$O(n)$ + (ZKP cost)	High latency, no quantifiable measure.
AAC [15]	High (Issuance)	Medium	$O(n)$	Anonymity degrades with history.
Policy Mana. [28] Optimization	Low	Low	$\sim O(\log n)$	Privacy-agnostic, may reduce anonymity.
QAE-BAC	**High (Quant.)**	**Medium**	$O(\log n)$	**Balances privacy and performance.**

anonymity assessment with privacy-aware policy optimization, addressing limitations of existing approaches that either prioritize privacy at efficiency's expense or optimize performance while ignoring privacy implications.

III. Preliminaries

This section formalizes the core concepts underpinning the QAE-BAC framework, establishing the foundational terminology for anonymity quantification and access control mechanisms.

A. Formal Model of Attribute-Based Access Control

This subsection introduces the basic elements of attribute-based access control model [1].

Definition 1 (Attribute Space)**.** *The attribute space is partitioned into disjoint sets: subject attributes* $AS = \{a_1^s, a_2^s, \ldots, a_m^s\}$*, object attributes* $AO = \{a_1^o, a_2^o, \ldots, a_n^o\}$*, environment attributes* AE*, and operation attributes* OP*. The attribute space* $A = AS \cup AO \cup AE \cup OP$*. Each attribute* $a \in A$ *is defined as* $a = (t, w, V)$*, where* t *denotes type,* w *represents weight, and* V *is the value domain.*

Definition 2 (Subject and Object)**.** *A subject* s *is an entity initiating access requests, represented by a tuple of subject attribute values:* $s = (v_1^s, v_2^s, \ldots, v_m^s), \quad v_i^s \in V(a_i^s),\ a_i^s \in AS$*. An object* o *is a target resource, represented by a tuple of object attribute values:* $o = (v_1^o, v_2^o, \ldots, v_n^o), \quad v_j^o \in V(a_j^o),\ a_j^o \in AO$*.*

Definition 3 (Attribute Credential and Access Request)**.** *An attribute credential* $c \subseteq s$ *is a minimal subset of a subject's attributes presented for authentication. An access request* $req = (c, o, op, env)$ *is composed of the credential, target object, requested operation, and environmental context.*

Definition 4 (Access Policy)**.** *An access policy* $P = \{r_1, r_2, \cdots, r_n\}$ *governs resource permissions, where each rule* $r_i = \{(a_1 : v_1), (a_2 : v_2), \cdots, (a_m : v_m)\}$ *specifies attribute constraints, where each* $a_i \in AS \cup AO \cup AE \cup OP$ *is an attribute, and* v_i *is either a specific value in the domain* $V(a_i)$ *or a constraint (e.g., range, set) over that domain. A request is granted if it satisfies at least one rule.*

B. Anonymity Metrics and Quantification

This subsection formalizes identifiability [15] and anonymity quantification concepts [16].

Definition 5 (Subject Identifier)**.** *A subject identifier uniquely identifies a subject within* S*.* ***Explicit identifiers*** *(e.g., unique IDs) directly bind to identity, while* ***implicit identifiers*** *are attribute combinations that uniquely identify a subject despite individual attributes not being exclusive. A subject* s_i *can have* k *implicit identifiers, where* $k \in [0, 2^{|s_i|} - 1]$*, posing significant re-identification risks when present in access requests [15].*

Definition 6 (Credential Subject Space)**.** *Given credential* c*, its subject space* $\mathcal{SS}_c = \mathcal{S}_1^c \cup \mathcal{S}_2^c$ *contains subjects potentially linked to requests carrying* c*, where* $\mathcal{S}_1^c = \{s \mid c \subseteq s,\ s \in S\}$ *(generators) and* $\mathcal{S}_2^c = \{s \mid \exists req \in R, c \in req \wedge c \subseteq s\}$ *(historical users). Anonymity quantification depends on* $\mathcal{SS}_c$*'s distribution.*

Definition 7 (Request Probability Entropy)**.** *For request* req_c *with signed credential* $\sigma_s(c)$*, the request probability entropy* $\mathcal{E}_{req}(req_c) = -\sum_{s \in \mathcal{SS}_c} P(X = s) \log_2 P(X = s)$ *measures adversary uncertainty, where* X *represents the adversary's guess of subject origin and* $P(X = s)$ *is the probability estimation. Higher values indicate stronger anonymity protection.*

Definition 8 ((r,t)-Anonymity)**.** *Given subject attribute matrix* M*, the* (r,t)*-anonymity model [16] guarantees that for subjects* $a_t = \{s_i \mid t \leq |s_i|, s_i \in M\}$ *with at least* t *attributes, the minimum credential subject space size satisfies* $r = \min\{|\mathcal{S}_1^{s_i}| \mid s_i \in a_t\}$*, providing worst-case re-identification protection. A higher* r *indicates stronger anonymity.*

C. Cryptographic Foundation

The QAE-BAC framework is built upon a rigorous cryptographic foundation to ensure tamper-resistant and verifiable security guarantees. During registration, each subject s is assigned a public-private key pair (pk_s, sk_s). The system's security rests on the following standard computational hardness assumption:

Assumption 1 (Discrete Logarithm (DL) Assumption)**.** *Let* $\mathbb{G}$ *be a cyclic group of prime order* p *with generator* g*. For a uniformly random element* $h = g^a \in \mathbb{G}$ *where* $a \xleftarrow{\$} \mathbb{Z}_p$*, no Probabilistic Polynomial-Time (PPT) algorithm* $\mathcal{A}$ *can recover the exponent* a *with probability greater than* $\mathsf{negl}(\lambda)$*, where* λ *is the security parameter.*

A digital signature scheme that is Existentially Unforgeable under Chosen Message Attacks (EUF-CMA) [39] is employed, whose security is reducible to the DL assumption. This scheme enables subjects to generate unforgeable attestations for their attribute credentials.

IV. System Overview

This section delineates the overarching architecture and operational workflow of the proposed QAE-BAC framework, designed to mitigate identity re-identification risks and authorization inefficiencies in decentralized systems.

A. Threat Model and Design Goals

To rigorously evaluate the security of the QAE-BAC framework, a comprehensive threat model is first established, core design goals are defined, and formal security definitions are provided. These elements are discussed and analyzed in Section V and VI.

1) Threat Model: A powerful adversary $\mathcal{A}$ capable of both passive and active attacks in blockchain environments is considered. $\mathcal{A}$ can eavesdrop on network traffic, perform data analysis and linkage attacks, compromise limited subjects, and actively probe the system. However, $\mathcal{A}$ is computationally bounded and cannot break cryptographic primitives or control consensus. The primary security objective is protecting identity privacy by preventing re-identification and request linking. *Detailed threat analysis is available in the full version [30].*

2) Design Goals: The QAE-BAC framework is designed to achieve the following goals under the aforementioned threat model:

G1 Subject Anonymity: The system should ensure that an adversary cannot determine the real-world identity of a subject from its attribute credential c or a series of credentials used in access requests. This requires that for any credential c, the credential subject space $\mathcal{SS}_c$ is sufficiently large and non-unique.

G2 Unlinkability: Given two or more access requests, an adversary should not be able to determine with confidence whether they originated from the same subject, unless this is explicitly revealed by the policy logic itself.

G3 Fine-Grained & Efficient Access Control:The system must enforce fine-grained access control policies without compromising performance. Authorization decisions should be both accurate and efficient, even as the number of subjects, attributes, and policies scales.

G4 Resilience to Attribute Correlation Attacks: The system should be resilient against attacks that leverage the correlation between different attributes or between requests and background knowledge to reduce anonymity.

3) Security Definitions: Based on the threat model and design goals, the key security properties of QAE-BAC are formalized. The complete formal definitions and security games are provided in the full version [30].

Definition 9 (Request Anonymity)**.** *Let* Π *be the QAE-BAC framework and* $\mathcal{A}$ *a PPT adversary. In experiment* $\mathsf{Exp}^{\mathsf{Req\text{-}Anon}}_{\mathcal{A},\Pi}(1^\lambda)$*,* $\mathcal{A}$ *chooses subjects* s_0, s_1*, receives request* req_b *from random* s_b*, and outputs guess* b'*.* Π *provides Request Anonymity if* $\mathsf{Adv}^{\mathsf{Req\text{-}Anon}}_{\mathcal{A},\Pi} = |\Pr[b' = b] - \frac{1}{2}|$ *is negligible.*

Definition 10 (Request Unlinkability)**.** *Let* Π *be the QAE-BAC framework and* $\mathcal{A}$ *a PPT adversary. In experiment* $\mathsf{Exp}^{\mathsf{Unlink}}_{\mathcal{A},\Pi}(1^\lambda)$*,* $\mathcal{A}$ *observes request sequence from* s*, then distinguishes between new requests from* s *and random* s'*.* Π *provides Request Unlinkability if* $\mathcal{A}$*'s advantage is negligible.*

B. Core Modules

The architecture of QAE-BAC, illustrated in Fig. 1, consists of three core modules deployed on the blockchain. A key enhancement across these modules is the integration of cryptographically signed attribute credentials.

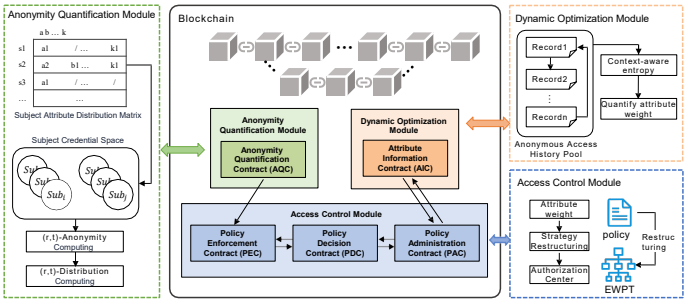

Fig. 1: Architecture of the QAE-BAC framework.

1) Anonymity Quantification Module: Implemented in the Anonymity Quantification Contract (AQC), this module evaluates request and subject anonymity. It first verifies signed credential $\sigma_s(c)$ against subject public key pk_s, then analyzes the subject attribute matrix, constructs credential subject spaces $\mathcal{SS}_c$, and applies (r,t)-anonymity to quantify anonymity. Requests meeting thresholds proceed to policy execution.

2) Dynamic Optimization Module: Operating via the Attribute Information Contract (AIC), this module introduces dynamic attribute weighting by maintaining an Anonymous Access History (AAH) Pool and computing context-aware entropy from AAH data and AQC anonymity scores. The output is an optimized attribute weight list prioritized by discriminatory power.

3) Access Control Module: This authorization engine refines traditional Policy Enforcement Contract (PEC), Policy Decision Contract (PDC), Policy Administration Contract (PDC) through three smart contracts. The PEC verifies signature $\sigma_s(c)$ before parsing requests, then dynamically reconstructs the EWPT via PDC for fast authorization through path existence checks via PDC.

C. Operational Workflow

The end-to-end workflow of QAE-BAC is composed of four sequential phases that incorporate cryptographic verification.

Phase 1: Anonymity Assessment. Executed by AQC: configures attribute space A; registers subjects/objects; for requests with $\sigma_s(c)$, verifies signature and computes request probability entropy and (r,t)-anonymity against thresholds.

Phase 2: Request Authentication. Handled by PEC: performs cryptographic validation of $\sigma_s(c)$ and confirms Phase 1 passage; rejects failures immediately.

Phase 3: Dynamic Authorization. For authenticated requests, PDC checks EWPT for path matching attribute sequence; grants authorization if valid path exists.

Phase 4: Weight Update. Operated periodically by AIC: updates attribute weights using context-aware entropy from AAH Pool and AQC scores; pushes updates to PDC for EWPT reorganization.

V. Detailed Design of QAE-BAC

This section presents the comprehensive algorithmic foundation of the QAE-BAC framework.The eight core algorithms that implement the three functional modules are detailed, integrating the cryptographic foundation from Section III-C with the information-theoretic anonymity metrics that form the core of the privacy preservation approach. The algorithms are designed to work in concert to achieve the four primary goals (G1-G4) in Section IV-A2.

A. Anonymity Quantification Module

The Anonymity Quantification Module provides formal, measurable guarantees against re-identification attacks. Its design is grounded in information theory, using entropy to quantify the uncertainty an adversary faces when attempting to identify subjects from their attribute credentials. The module comprises four algorithms that collectively ensure the achievement of **G1** and **G2** by rigorously measuring and enforcing anonymity levels. Algorithm 1 serves as the foundational step for all subsequent anonymity calculations. It first verifies the cryptographic signature on the credential to ensure its authenticity and integrity, preventing forgery and spoofing attacks. This step is critical for maintaining the trustworthiness of the system. It then constructs the *credential subject space* $\mathcal{SS}_c$ by combining two sets: subjects who can generate the credential ($\mathcal{S}_1^c$) and subjects who have used it historically ($\mathcal{S}_2^c$). The size of $\mathcal{SS}_c$ directly determines the theoretical upper bound of anonymity for the request. This algorithm directly contributes to **G1** and **G2** by ensuring that only valid credentials are processed and by defining the population of possible subjects for a given credential, which is essential for quantifying anonymity.

Algorithm 2 computes the request probability entropy $\mathcal{E}_{req}(req_c)$ based on the credential subject space $\mathcal{SS}_c$ generated by Algorithm 1. The entropy is calculated using the Shannon entropy formula, which measures the adversary's uncertainty about the subject's identity. A higher entropy value indicates greater anonymity. This algorithm is the core metric for evaluating **G1** at the request level. It also supports **G2** by ensuring that multiple requests from the same subject, when using different credentials, yield high entropy values, making linking difficult. The algorithm returns zero if $\mathcal{SS}_c$ has only one subject, indicating a complete loss of anonymity.

The (r, t)-anonymity model evaluates systemic anonymity by computing the minimum credential subject space size r across all subjects with at least t attributes. This provides a global anonymity guarantee, ensuring protection even for subjects with extensive attribute sets. Additionally, subject-level anonymity scores $\mathcal{A}_{sub}(s)$ are computed by aggregating request anonymity values across all possible credentials, weighted by their frequency. These computations ensure baseline anonymity protection (G1) and resilience to correlation attacks (G4). The detailed algorithmic implementations are provided in [30].

Algorithm 1 Credential Verification and Subject Space Construction

Input: Signed credential $\sigma_s(c)$, subject space S, request history R, public key pk_s
Output: Credential subject space $\mathcal{SS}_c$ or ERROR

Verify Signature: result $\leftarrow$ Verify$(pk_s, c, \sigma_s(c))$
if result $=$ False **then**
  return ERROR ▷ Reject forged or tampered credential
end if
$\mathcal{S}_1^c \leftarrow \{s \mid c \subseteq s, s \in S\}$ ▷ Construct capability set
$\mathcal{S}_2^c \leftarrow \{s \mid \exists req \in R$ where $c \in req \wedge s$ initiated $req\}$ ▷ Construct usage set
$\mathcal{SS}_c \leftarrow \mathcal{S}_1^c \cup \mathcal{S}_2^c$
if $|\mathcal{SS}_c| = 0$ **then**
  return ERROR ▷ Invalid credential configuration
end if
return $\mathcal{SS}_c$

Algorithm 2 Request Anonymity Quantification

Input: Credential subject space $\mathcal{SS}_c$, request req_c
Output: Request anonymity metric $\mathcal{E}_{req}(req_c)$

$\mathcal{E}_{req} \leftarrow 0.0$
if $|\mathcal{SS}_c| = 1$ **then**
  return 0.0 ▷ Zero anonymity - implicit identifier
end if
for all subject $s_i \in \mathcal{SS}_c$ **do**
  $p_i \leftarrow \frac{\text{frequency of } s_i \text{ in } \mathcal{SS}_c}{\sum_{s_j \in \mathcal{SS}_c} \text{frequency of } s_j}$ ▷ Probability estimation
  $\mathcal{E}_{req} \leftarrow \mathcal{E}_{req} - p_i \cdot \log_2(p_i)$ ▷ Accumulate entropy
end for
return $\mathcal{E}_{req}$

B. Dynamic Optimization Module

The Dynamic Optimization Module adapts attribute weights based on both authorization patterns and anonymity considerations, creating a feedback loop that continuously improves system performance. It plays a key role in achieving **G3** and **G4** by optimizing policy evaluation efficiency while maintaining privacy.

1) Theoretical Foundation: Information gain from authorization decisions is combined with anonymity metrics to compute dynamic attribute weights. The information gain $I(D, A)$ for an attribute A measures how much it reduces uncertainty about authorization outcomes:

$$I(D, A) = H(D) - H(D|A) \tag{1}$$

Algorithm 3 Attribute Weight Optimization

Input: Attribute space A, AAH Pool H
Output: Sorted attribute weight list W

$W \leftarrow \emptyset$
for all attribute $a_i \in A$ **do**
  $I(D, a_i) \leftarrow \text{CalculateInformationGain}(H, a_i)$ ▷ Computing information gain from AAH Pool
  $(a_1, r) \leftarrow \text{CalculateRTAnonymity}(M, 1)$ ▷ $t = 1$
  $\mathcal{A}_{attr}(a_i) \leftarrow r$ ▷ Use min subject space size as anonymity measure
  $w_i \leftarrow I(D, a_i) + \mathcal{A}_{attr}(a_i)$ ▷ Combined metric
  $W \leftarrow W \cup \{(a_i, w_i)\}$
end for
$\text{sort}(W)$ by w_i descending
return W

where $H(D)$ is the entropy of authorization decisions:

$$H(D) = -\sum_{d \in \{\text{grant,deny}\}} P(d) \log_2 P(d) \tag{2}$$

and $H(D|A)$ is the conditional entropy:

$$H(D|A) = -\sum_{v \in V_A} P(A = v) \times \sum_{d \in \{\text{grant,deny}\}} P(d|A = v) \log_2 P(d|A = v) \tag{3}$$

The final weight combines information gain with the attribute's **individual anonymity contribution**, which is derived from the (r, t)-anonymity assessment with $t = 1$:

$$A_w = I(D, A) + \mathcal{A}_{attr}(a_i) \tag{4}$$

where $\mathcal{A}_{attr}(a_i)$ represents the anonymity contribution of attribute a_i when considered individually, computed as the minimum credential subject space size r from the (r, t)-anonymity model. The detailed implementation of this calculation is provided in [30].

2) Algorithm Implementation: The information gain $I(D, a)$ for each attribute a is computed based on the AAH Pool, measuring the reduction in uncertainty about authorization decisions when the attribute value is known. This entropy-based calculation identifies attributes most predictive of access outcomes, supporting efficient policy structuring (G3) and reducing correlation attack risks by prioritizing discriminative attributes (G4). The specific entropy calculation procedure is detailed in [30].

Algorithm 3 generates a sorted list of attribute weights by combining the information gain computed from the AAH Pool with the anonymity scores from the AQC. The combined weight w_i reflects both the attribute's decision-making power and its privacy impact. This algorithm is central to **G3**, as the weight list directly guides the construction of the EWPT, ensuring that the most discriminative and privacy-preserving attributes are checked first. It also enhances **G4** by dynamically adjusting weights to mitigate correlation risks based on current access patterns. This continuous, data-driven update mechanism enables QAE-BAC to naturally adapt to evolving

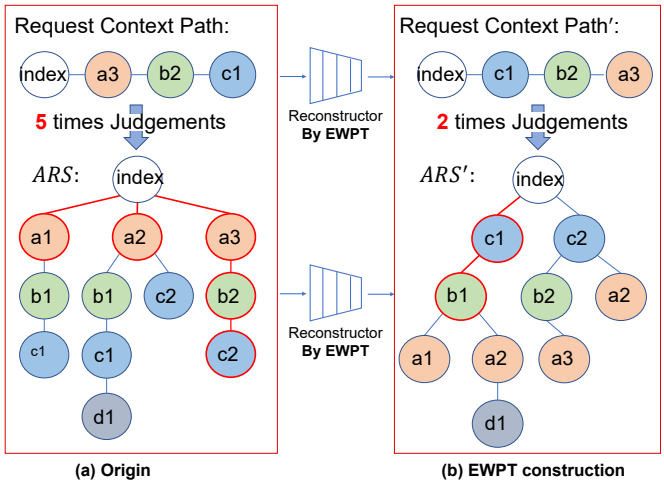

Fig. 2: Entropy-Weighted Path Tree structure: (a) initial policy organization; (b) optimized structure after dynamic weight adjustment. The tree structure enables $O(m)$ authorization time complexity where m is the number of attributes.

policies. When policies are added, removed, or modified, the resulting changes in access patterns are captured in the AAH Pool. Subsequent weight recalculations and EWPT reconstructions automatically incorporate these changes, ensuring that the authorization structure remains optimized for the current policy environment without manual intervention. The reconstruction is governed by a hybrid trigger mechanism combining periodic and event-driven conditions; in our experiments, a periodic update is performed every 100 blocks, while an event-driven update is triggered when the AAH Pool accumulates 500 new records.

C. Access Control Module

The Access Control Module implements efficient policy evaluation through the EWPT structure, incorporating cryptographic verification for security and path-based matching for efficiency. It is responsible for achieving **G3** by enabling fast authorization decisions and upholding **G1** and **G2** through integrated anonymity checks.

Algorithm 4 builds the EWPT using the policy rule set and the attribute weight list from Algorithm 3. The tree is constructed by sorting attributes in each rule by their weight descending, ensuring that high-weight attributes (those with high information gain and anonymity) form the shared prefixes of paths. This structure dramatically reduces the average depth of policy matching, enabling $O(m)$ time complexity where m is the number of attributes in a request. To facilitate comprehension, Fig.2a shows a simplified example of P_{W1} with the initial attribute weight $W_1 = [a, b, c, d]$, containing four access policy rules: $P = \{(a1, b1, c1), (a2, b1, c1, d1), (a2, c2), (a3, b2, c2)\}$. When the attribute weights are updated to $W_2 = [c, b, a, d]$, the policy P is reconstructed into P_{W2}, as shown in Fig.2b. For an anonymous request $R = (a3, b2, c1)$, the tree-based matching of origin P_{W1} requires 5 attempts, while the EWPT P_{W2} further optimizes the process, needing only 2 comparisons. This

Algorithm 4 EWPT Construction

Input: Policy rule set P, attribute weight list W
Output: Entropy-Weighted Path Tree T

$T \leftarrow$ CreateRootNode()
for all rule $r \in P$ **do**
  sorted_attrs $\leftarrow$ sort(r.attributes) by weight in W descending
  current $\leftarrow T$.root
  for all attribute $a_j \in$ sorted_attrs **do**
    if $a_j \notin$ current.children **then**
      current.children$[a_j] \leftarrow$ CreateNode(a_j)
    end if
    current $\leftarrow$ current.children$[a_j]$
  end for
  MarkAsLeaf(current) ▷ Complete policy path
end for
return T

Algorithm 5 Authorization Decision with Cryptographic Verification

Input: Request $req = (\sigma_s(c), o, op, env)$, EWPT T, weight list W, public key pk_s
Output: Authorization decision: GRANT or DENY

Step 1: Cryptographic Verification
valid $\leftarrow$ Verify($pk_s, c, \sigma_s(c)$)
if not valid **then**
  return DENY ▷ Reject unverifiable request
end if
Step 2: Anonymity Validation
$\mathcal{E}_{req} \leftarrow$ CalculateRequestAnonymity(req) ▷ Using Alg. 2
if $\mathcal{E}_{req} <$ threshold **then**
  return DENY ▷ Insufficient anonymity
end if
Step 3: Policy Path Evaluation
attr_sequence $\leftarrow$ ExtractAndSortAttributes(req, W)
current $\leftarrow T$.root
for all value $v_i \in$ attr_sequence **do**
  if $v_i \notin$ current.children **then**
    return DENY ▷ No matching path
  end if
  current $\leftarrow$ current.children$[v_i]$
end for
if current.isLeaf **then**
  return GRANT ▷ Valid path exists
else
  return DENY ▷ Incomplete path
end if

algorithm is pivotal for **G3**, as it optimizes the policy structure for efficient evaluation. The tree's design also supports **G4** by promoting the use of attributes that are less susceptible to correlation.

Algorithm 5 is the culmination of the QAE-BAC framework, performing the final authorization decision. It integrates cryptographic verification, anonymity validation, and policy path checking in a three-step process. First, it verifies the digital signature on the credential to ensure authenticity (using the same method as Algorithm 1). Second, it checks the request's anonymity score against a threshold to ensure sufficient privacy (using Algorithm 2). Finally, it traverses the EWPT to check for a matching path. This algorithm directly enforces **G1** and **G2** by rejecting requests that fail anonymity or verification checks. It achieves **G3** through efficient path matching and contributes to **G4** by ensuring that only requests with safe attribute combinations are granted.

The integrated design of these eight algorithms creates a comprehensive framework that provides both strong privacy guarantees through information-theoretic anonymity metrics and efficient authorization through optimized policy structures, all while maintaining cryptographic security through digital signature verification, forming a framework that achieves the goals of subject anonymity, unlinkability, efficient access control, and resilience to attribute correlation attacks.

VI. Security Analysis

This section provides a formal reduction-based security analysis of the QAE-BAC framework. It is demonstrated that breaking the anonymity of QAE-BAC is computationally equivalent to solving well-established hard problems under the defined threat model (Section IV-A). The complete formal security proofs and probability calculations are provided in the extended version [30]

A. Assumptions and Security Analysis

The security reduction relies on the standard DL assumption (Assumption 1) and the EUF-CMA security of the digital signature scheme. The main security theorem is formally stated.

Theorem 11 (Security of QAE-BAC)**.** *Let λ be the security parameter. Let $\mathcal{A}$ be any PPT adversary against the Request Anonymity (Def. 9) or Request Unlinkability (Def. 10) of the QAE-BAC framwork Π, with advantage $\mathsf{Adv}^{\mathsf{Anon}}_{\mathcal{A},\Pi}(\lambda)$. If the digital signature scheme Σ is EUF-CMA secure and the DL assumption holds in group $\mathbb{G}$, then $\mathsf{Adv}^{\mathsf{Anon}}_{\mathcal{A},\Pi}(\lambda)$ is negligible. Formally, there exists PPT algorithms (simulators) $\mathcal{S}_1$ and $\mathcal{S}_2$ such that:*

$$\mathsf{Adv}^{\mathsf{Anon}}_{\mathcal{A},\Pi}(\lambda) \leq \mathsf{Adv}^{\mathsf{EUF-CMA}}_{\Sigma}(\mathcal{S}_1(\mathcal{A})) + \mathsf{Adv}^{\mathsf{DL}}_{\mathbb{G}}(\mathcal{S}_2(\mathcal{A})) + \mathsf{negl}(\lambda) \tag{5}$$

The proof of Theorem 11 is structured as a sequence of games. The core argument shows that any successful anonymity adversary $\mathcal{A}$ can be used to break the EUF-CMA security or compute discrete logarithms. It is first argued that the signature scheme ensures the authenticity and integrity of the attribute credential c, implying that $\mathcal{SS}_c$ is constructed from a valid credential. The adversary's advantage must stem from linking the signature to the private key. Two cases are distinguished:

1) **Case 1: Forgery.** If the adversary forges a signature, it contradicts EUF-CMA security.

2) **Case 2: Extraction.** If the adversary wins without forgery, it leverages signature information, allowing a simulator to extract discrete logarithms.

The complete proof is referred to [30] for intuition and provides probability calculations.

B. Discussion on the Security of Signature Schemes

The reduction in Case 2 is straightforward for deterministic signature schemes (e.g., EdDSA) or under the Random Oracle Model (ROM). Deterministic signatures prevent leakage through randomness, and security holds for probabilistic schemes under the ROM.

VII. Performance Evaluation

This section presents a comprehensive empirical analysis of the proposed QAE-BAC framework. To rigorously evaluate its effectiveness, the following three Research Questionss (RQss) are addressed:

RQ1: How do different system parameters (e.g., number of subjects, attributes) affect the anonymity guarantees provided by QAE-BAC's quantification module?

RQ2: Does QAE-BAC achieve significant performance improvements over state-of-the-art baselines? If so, what is the contribution of each key innovation (EWPT structure vs. dynamic optimization)?

RQ3: How does QAE-BAC perform under increasing system scale and complexity? Is it resilient to the "attribute explosion" problem?

The experimental setup is detailed, the design of the test cases is described, and a rigorous evaluation is presented to answer these RQs.

A. Experimental Setup

To thoroughly assess the performance of QAE-BAC, a simulation testbed was established on an Apple M1 Pro platform (16 GB RAM) running macOS Monterey v12.3. The environment was built on Hyperledger Fabric v2.2.1, with containerization managed by Docker v20.10.13 and Docker-compose v1.29.2. The smart contracts (chaincode) for the AQC, AIC, and PDC/PEC/PDC modules were developed in Golang v1.18 and deployed onto the Fabric network. The EdDSA signature scheme was employed using the Ed25519 curve, providing a security level of $\lambda = 128$ bits with SHA-256 as the hash function. To ensure the reliability and stability of the results, each experiment was repeated 10 times, and the average values are reported for analysis.

B. Test Case Design

To address **RQ1** and **RQ3**, QAE-BAC was evaluated using comprehensive configurations based on a real-world Internet of Things (IoT) dataset [40] collected via Zigbee Zolertia Z1 nodes. This dataset provides authentic physical-layer feature-based authentication scenarios, addressing the scarcity of realistic attribute complexity datasets for blockchain access control systems.

Following the One-Variable-At-A-Time (OVAT) principle [41], 15 test cases were designed in Table III covering seven key factors—entity counts, action volumes, control complexity, and attribute dimensionality. The selected scales (5K-15K subjects, 500K-1.5M requests) reflect realistic smart manufacturing and healthcare deployments, with synthetic data preserving real IoT statistical characteristics. **Note** that all synthetic data, including subjects, objects, requests, and policies, were generated to preserve the statistical characteristics observed in real IoT environments.

TABLE III: Parameters for the Designed Test Cases

Test Case	# Subjects	# Objects	# Requests	# Policies	Attr. Range	# Sub. Attrs.	# Obj. Attrs.
C1	**5K**	10K	1000K	100	4	4	2
C2	10K	10K	1000K	100	4	4	2
C3	**15K**	10K	1000K	100	4	4	2
C4	10K	**5K**	1000K	100	4	4	2
C5	10K	**15K**	1000K	100	4	4	2
C6	10K	10K	**500K**	100	4	4	2
C7	10K	10K	**1500K**	100	4	4	2
C8	10K	10K	1000K	**50**	4	4	2
C9	10K	10K	1000K	**150**	4	4	2
C10	10K	10K	1000K	100	**2**	4	2
C11	10K	10K	1000K	100	**6**	4	2
C12	**15K**	10K	1000K	100	4	**5**	2
C13	**15K**	10K	1000K	100	4	**3**	2
C14	10K	10K	1000K	100	**2**	4	**4**
C15	10K	10K	1000K	100	**2**	4	**3**

The test cases are grouped to analyze the impact of each factor, which is crucial for answering RQ1 and RQ3:

- **Subject Quantity (C1, C2, C3):** Tests system scalability w.r.t. user base size.
- **Object Quantity (C2, C4, C5):** Tests scalability w.r.t. resource base size.
- **Request Quantity (C2, C6, C7):** Tests performance under different load intensities.
- **Policy Quantity (C2, C8, C9):** Tests resilience to growing policy complexity.
- **Attribute Value Range (C2, C10, C11):** Tests impact of attribute granularity.
- **Subject Attribute Quantity (C3, C12, C13):** Tests resilience to subject attribute explosion.
- **Object Attribute Quantity (C10, C14, C15):** Tests resilience to object attribute explosion.

This structured design ensures a controlled and comparable basis for evaluating the impact of each variable. **Note** that the core innovation of QAE-BAC breaks the traditional privacy-efficiency trade-off. As demonstrated in Section VII-C and Section VII-D, QAE-BAC maintains high anonymity under diverse conditions while achieving significant performance gains. This proves that QAE-BAC not merely balances these dual objectives, but simultaneously delivers substantial improvements along both dimensions.

C. Anonymity Analysis

This subsection addresses **RQ1** by evaluating QAE-BAC's capability to preserve subject unlinkability under fine-grained policies. The consolidated results in Figures 3, 4, and 5 provide comprehensive insights into anonymity characteristics across all influencing factors.

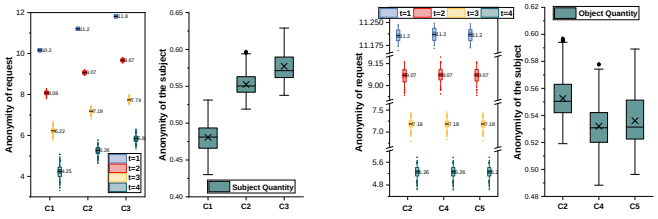

(a) **Subject Quantity** (b) **Object Quantity**

Fig. 3: The Impact of **Subject(Fig.3a)/ Object(Fig.3b) Quantity** on Anonymity Distribution.

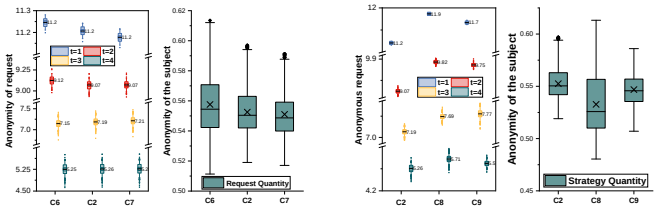

(a) **Request Quantity** (b) **Strategy Quantity**

Fig. 4: The Impact of **Request (Fig.4a) / Strategy(Fig.4b) Quantity** on Anonymity Distribution.

1) Graphical Interpretation: The results for each influencing factor group are presented in paired figures (e.g., Fig. 3a). For each group:

- The **left subfigure** in Fig. 3a shows the request anonymity distribution across different (r, t)-pairs. The x-axis represents the test cases, the y-axis represents the request anonymity ($\mathcal{E}_{req}$), and different-colored legends represent different t values. Comparisons between bars reveal the factor's impact, while comparisons within legends show the effect of varying t.
- The **right subfigure** in Fig. 3a displays the corresponding subject anonymity distribution ($\mathcal{A}_{sub}$) for the same test case group, presented using box plots to illustrate the central tendency and dispersion of anonymity scores across the subject population.

2) Key Findings (Answering RQ1): Analysis of the integrated anonymity distributions reveals four fundamental patterns:

(i) **Inverse Correlation with Attribute Count:** A strong negative correlation exists between attribute count t and request anonymity across all test scenarios (Figs. 3–5). As t decreases, the credential subject space $\mathcal{SS}_c$ expands significantly, increasing adversary uncertainty. This validates the (r, t)-anonymity model's effectiveness in capturing the core privacy-information trade-off.

(ii) **Positive Scaling with Subject Population:** Both request and subject anonymity scale positively with subject population size (Fig. 3). Larger subject spaces enable greater attribute combination reuse, expanding $\mathcal{SS}_c$ and enhancing anonymity guarantees for large-scale deployments.

(iii) **Indirect Influence of System Parameters:** Variations in object count, request volume, and policy complexity (Figs. 3, 4) indirectly affect anonymity by altering historical access pattern distributions. These factors influence $\mathcal{S}_2^c$ composition within $\mathcal{SS}_c$, demonstrating the system's adaptability to operational dynamics.

(iv) **Combinatorial Challenges with Attribute Complexity:** Attribute-related factors—value range, subject/object attribute counts (Fig. 5)—show clear negative correlation due to combinatorial explosion effects. This underscores the "attribute explosion" challenge and highlights the importance of QAE-BAC's dynamic optimization in maintaining practical anonymity levels.

Summary for RQ1: QAE-BAC's anonymity quantification provides consistent, measurable privacy guarantees primarily governed by subject population and credential attributes. The information-theoretic approach is validated by the framework's consistent performance in balancing privacy requirements against information disclosure under diverse operational conditions.

D. Efficiency Analysis

This subsection addresses **RQ2** and **RQ3** by evaluating the performance and scalability of QAE-BAC.

1) Baseline Selection: To thoroughly evaluate the efficiency and practicality of QAE-BAC, two baseline schemes were selected for comparison:

- **Fabric-IoT** [8]: A representative ABAC model implemented on Hyperledger Fabric. It serves as a benchmark for traditional, non-optimized ABAC performance within the same blockchain environment.
- **QAE-BAC-Static:** An **ablated variant** of QAE-BAC where the **dynamic optimization module is removed**. The EWPT structure is built using a fixed, initial attribute weight ordering. Comparing with QAE-BAC-Static helps isolate and quantify the performance contribution of the dynamic weight update mechanism itself, answering RQ2.

2) Performance Comparison (Answering RQ2 and RQ3): This paper evaluated the three schemes—Fabric-IoT, QAE-BAC-Static, and QAE-BAC—across all 15 test cases, measuring system throughput (Transactions Per Second, TPS) and average authorization latency. The results for the seven factor groups are summarized in Fig. 6.

The analysis of results provides clear answers to the RQs:

(i) Answer to RQ2 (Performance Advantage & Source): QAE-BAC achieves substantial performance improvements.

- **EWPT Contribution:** QAE-BAC-Static consistently and significantly outperforms Fabric-IoT (e.g., $\sim$10$\times$ higher throughput, $\sim$90% lower latency). This confirms that the EWPT structure itself, by reducing authorization to an $O(m)$ path traversal, is the primary source of performance gain.
- **Dynamic Optimization Contribution:** QAE-BAC further enhances performance beyond QAE-BAC-Static (e.g., $\sim$11$\times$ throughput gain over Fabric-IoT). This demonstrates the synergistic effect of the dynamic optimization module, which continuously adapts the EWPT to current access patterns, ensuring optimal performance.

(ii) Answer to RQ3 (Scalability and Resilience): QAE-BAC demonstrates excellent scalability and resilience.

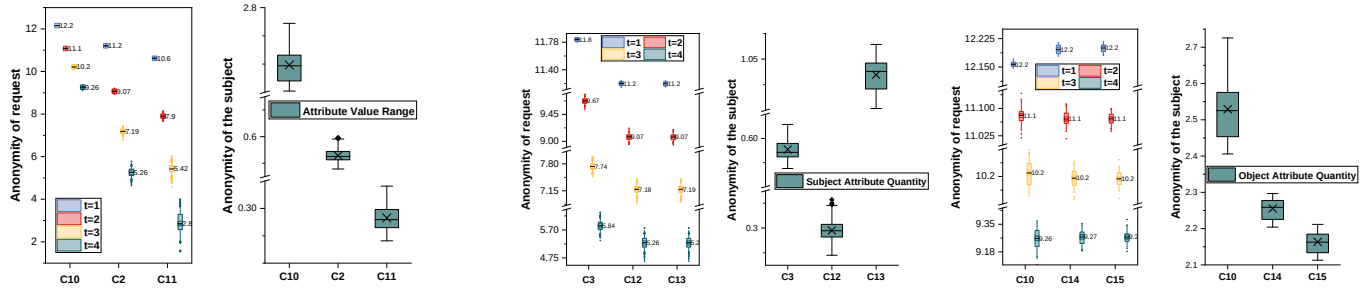

(a) **Attribute Value Range** (b) **Subject Attribute Quantity** (c) **Object Attribute Quantity**

Fig. 5: The Impact of **Attribute Value Range(Fig.5a)/ Subject Attribute Quantity(Fig.5b)/ Object Attribute Quantity(Fig.5c)** on Anonymity Distribution.

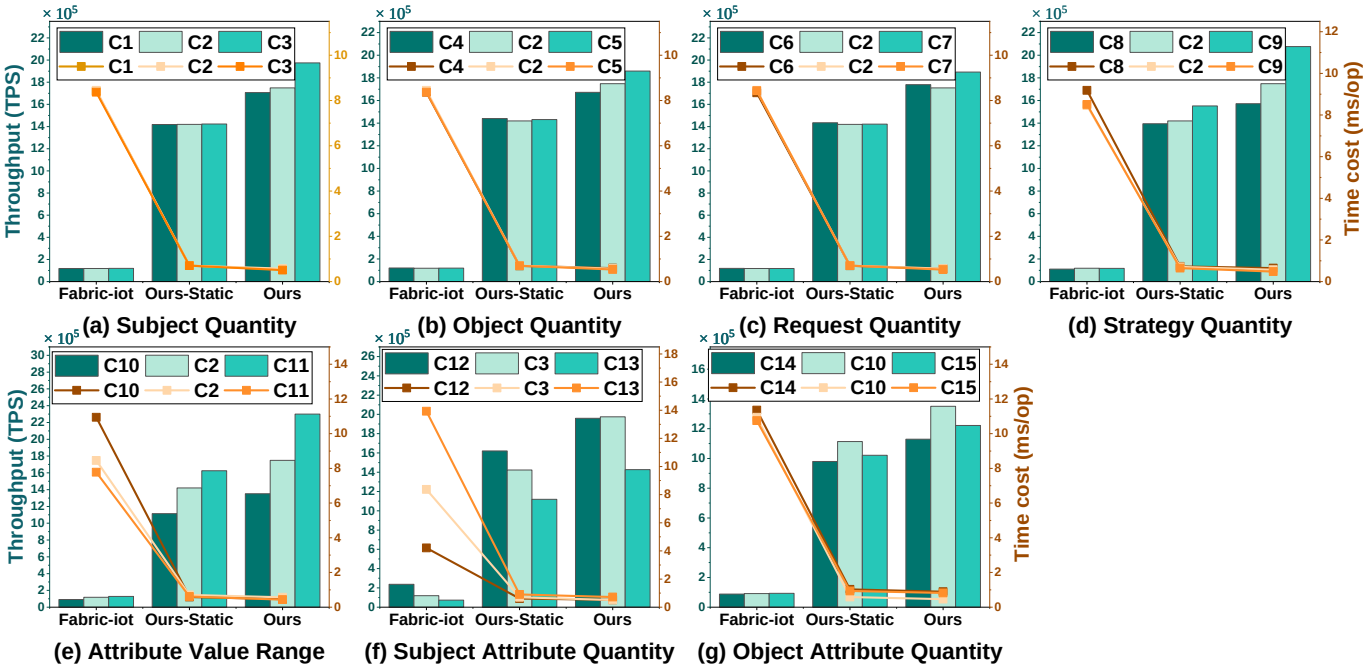

Fig. 6: Performance comparison across the seven influencing factor groups. All subfigures (a)-(g) share a common coordinate system: the left y-axis (green) represents **Throughput (TPS)** and the right y-axis (orange) represents **Latency (ms)**. The results demonstrate the performance gain from the EWPT structure (QAE-BAC-Static vs. Fabric-IoT) and the additional gain from dynamic optimization (QAE-BAC vs. QAE-BAC-Static).

- **Scalability:** Under increasing system scale (Fig. 6a,b,c), QAE-BAC maintains stable or even improving performance, while the baselines show degradation. This advantage stems from its ability to dynamically utilize historical access information to continuously optimize the authorization path structure within the EWPT, thereby enhancing decision-making efficiency. The system's feedback loop becomes more effective at identifying optimal paths as more data becomes available.
- **Resilience to Policy Growth:** With increasing policy complexity (Fig. 6d), QAE-BAC shows superior robustness (only $\sim$10% drop vs. $\sim$40% for Fabric-IoT).
- **Resilience to Attribute Explosion:** Against increasing attribute dimensions (Fig. 6e,f,g), QAE-BAC exhibits the smallest performance decline ($<$13% throughput drop). The dynamic prioritization in the EWPT effectively mitigates the performance impact of attribute explosion.

While the current experimental scale comprehensively covers typical IoT deployment scenarios [40], the scalability of QAE-BAC under even higher loads warrants discussion. The EWPT structure fundamentally reduces the policy matching complexity to $O(m)$, where m is the number of attributes in a request, making it independent of the total number of policies. This theoretical guarantee, combined with the observed trends in Fig. 6a,b,c where performance remains stable or improves with increasing numbers of subjects, objects, and requests, indicates a strong inherent scalability, suggesting that QAE-BAC is well-positioned to handle larger-scale deployments.

Summary for RQ2 & RQ3: QAE-BAC achieves significant performance gains through EWPT and dynamic optimization, demonstrating excellent scalability and resilience. The framework's synergy between anonymity and efficiency—where highly anonymous attributes naturally form efficient matching paths—creates a virtuous cycle rather than trade-off,

ensuring continuous optimization in dynamic environments. The approach to breaking the privacy-efficiency trade-off is quantitatively validated by these results.

VIII. Discussion

This section discusses key limitations of the anonymity metric and threat model, and presents future directions including enhanced metrics, trusted execution, cross-domain adaptation, and accountable anonymity via puncturable ring signatures.

A. Limitations

1) Limitations of the Anonymity Metric: The $\mathcal{E}_{req}$ metric's foundational assumption of uniform adversarial priors proves vulnerable to real-world attackers with auxiliary knowledge, including non-uniform prior distributions, exploitable attribute correlations within credentials, and behavioral/temporal side-channels in request patterns. Although increasing the anonymity threshold can alleviate this risk, this approach imposes significant throughput costs, highlighting a fundamental privacy-performance tradeoff.

2) Threat Model Implications: Relaxing blockchain trust assumptions reveals critical vulnerabilities: compromised nodes accessing sensitive state (AAH Pool, attribute matrix M), malicious administrators bypassing anonymity safeguards despite signature protections, and timing-based deanonymization through transaction metadata. These threats define precise security boundaries for QAE-BAC: while the framework provides robust protection against external passive adversaries, addressing advanced internal threats requires orthogonal mitigation strategies with substantial performance tradeoffs beyond its current scope.

B. Practical Deployment Considerations

This section further discusses the practical usability and integration of QAE-BAC in real-world deployments. The modular chaincodes AQC/AIC/PDC in QAE-BAC are developed in Go, successfully deployed on the widely adopted Hyperledger Fabric platform, demonstrating excellent compatibility. The deployment overhead primarily consists of initial attribute space configuration and access control policy downloads during operation, without incurring additional costs. In the integration of QAE-BAC with existing ABAC systems on Fabric (e.g., [8]), backward compatibility is achieved by replacing the existing authorization logic with the proposed contracts. Additionally, the automated execution feature of chaincodes ensures that no supplementary overhead is incurred. Therefore, the proposed system is endowed with genuine deployability.

C. Future Work

In **Enhanced Anonymity Metrics**, future research will develop Bayesian anonymity metrics for privacy quantification under adversarial knowledge, alongside integrating differential privacy with calibrated noise injection into AAH pools and $\mathcal{SS}_c$ spaces to strengthen guarantees against auxiliary information. In **Trusted System Implementation**, trusted execution environments will be explored to secure critical computations and ZKPs will be investigated for enabling credential-based authorization without identity exposure, thereby eliminating the need for on-chain anonymity calculations. In **Cross-Domain Adaptation**, research directions include designing federated anonymity estimation for multi-domain scenarios without data sharing, and implementing online learning mechanisms to dynamically adapt to concept drift and evolving adversarial tactics. In **Accountable Anonymity**, future research will explore puncturable ring signatures to reconcile strong unlinkability with conditional traceability, enabling designated authorities to trace malicious actors under specific conditions without compromising day-to-day anonymity guarantees.

IX. Conclusion

This paper has presented QAE-BAC, a novel framework that effectively tackles the dual challenges of privacy preservation and authorization efficiency in blockchain-based attribute-based access control. By introducing a quantifiable (r, t)-anonymity model for continuous privacy assessment and an EWPT for privacy-aware policy optimization, QAE-BAC achieves a breakthrough balance between these traditionally conflicting goals. Extensive experimental results demonstrate that this framework not only maintains strong anonymity guarantees but also significantly enhances performance, yielding up to 11× higher throughput and 87% lower latency compared to state-of-the-art alternatives, thereby enabling practical and secure fine-grained access control for decentralized applications.

References

[1] V. C. Hu, D. R. Kuhn, and D. F. Ferraiolo, "Attribute-based access control," *Computer*, vol. 48, no. 2, pp. 85–88, 2015.

[2] X. Xing, Y. Liu, Q. Wu *et al.*, "Multi-committee ABE based decentralized access control with sharding blockchain for web 3.0," *IEEE Trans. Dependable Secur. Comput.*, vol. 22, no. 3, pp. 2533–2549, 2025.

[3] J. Zhang, X. Li, M. Zheng *et al.*, "Verifuzzy: A dynamic verifiable fuzzy search service framework for encrypted cloud data," *IEEE Trans. Serv. Comput.*, vol. 19, no. 1, pp. 780–793, 2026.

[4] X. Hao, W. Ren, Y. Fei *et al.*, "A blockchain-based cross-domain and autonomous access control scheme for internet of things," *IEEE Trans. Serv. Comput.*, vol. 16, no. 2, pp. 773–786, 2023.

[5] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.

[6] Y. Lv, X. Li, K. Chen *et al.*, "Medexchain: Enabling secure and efficient PHR sharing across heterogeneous blockchains," *IEEE Internet Things J.*, vol. 12, no. 16, pp. 32 932–32 950, 2025.

[7] Y. Zhang, M. Yutaka, M. Sasabe *et al.*, "Attribute-based access control for smart cities: A smart-contract-driven framework," *IEEE Internet Things J.*, vol. 8, no. 8, pp. 6372–6384, 2021.

[8] H. Liu, D. Han, and D. Li, "Fabric-iot: A blockchain-based access control system in iot," *IEEE Access*, vol. 8, pp. 18 207–18 218, 2020.

[9] J. Duan, L. Wang, W. Wang *et al.*, "TRCT: A traceable anonymous transaction protocol for blockchain," *IEEE Trans. Inf. Forensics Secur.*, vol. 18, pp. 4391–4405, 2023.

[10] H. Mao, J. Zhang, Y. Zhang *et al.*, "A cross-domain data sharing scheme based on federated blockchain," in *19th TASE, Limassol, Cyprus, July 14-16, 2025*. Springer, 2025, pp. 285–302.

[11] M. Zhang, X. Li, J. Zhang *et al.*, "Swiftguard: Enhanced privacy and efficiency in blockchain-based fine-grained access control for cross-domain healthcare collaboration," in *28th CSCWD, Compiegne, France, May 5-7, 2025*. IEEE, 2025, pp. 1863–1868.

[12] K. Qin and D. Gu, "To share or hide: Confidential model compilation as a service with privacy-preserving transparency," in *43rd SRDS, Charlotte, NC, USA, 2024*. IEEE, 2024, pp. 126–138.

[13] J. Koo, G. Kang, and Y. Kim, "Access control framework for cross-platform interoperability in the industrial internet of things," *IEEE Trans. Ind. Informatics*, vol. 21, no. 1, pp. 801–810, 2025.

[14] W. Xia, Y. Liu, Z. Wan *et al.*, "Enabling realistic health data re-identification risk assessment through adversarial modeling," *J. American Medical Informatics Association*, pp. 744–752, 2021.
[15] R. Zhang, G. Liu, H. Kang *et al.*, "Anonymity in attribute-based access control: Framework and metric," *IEEE Trans. Dependable Secur. Comput.*, vol. 21, no. 1, pp. 463–475, 2024.
[16] E. Lanus, C. J. Colbourn, and G. Ahn, "Guaranteeing anonymity in attribute-based authorization," *J. Inf. Secur. Appl.*, p. 103895, 2024.
[17] G. Fedrecheski, L. C. C. D. Biase, P. C. Calcina-Ccori *et al.*, "Smartabac: Enabling constrained iot devices to make complex policy-based access control decisions," *IEEE Internet Things J.*, pp. 5040–5050, 2022.
[18] W. Zhang, X. Huo, and Z. Bao, "A secure and efficient multi-domain data sharing model on consortium chain," *J. Supercomput.*, vol. 79, no. 8, pp. 8538–8582, 2023.
[19] Y. Lv, R. Feng, M. Ma *et al.*, "Reinventing multi-user authentication security from cross-chain perspective," *IEEE Trans. Inf. Forensics Secur.*, vol. 19, pp. 8908–8923, 2024.
[20] S. S. Ullah, V. A. Oleshchuk, and H. S. G. Pussewalage, "A survey on blockchain envisioned attribute based access control for internet of things: Overview, comparative analysis, and open research challenges," *Comput. Networks*, vol. 235, p. 109994, 2023.
[21] W. Tong, X. Dong, Y. Shen *et al.*, "A blockchain-driven data exchange model in multi-domain iot with controllability and parallelity," *Future Gener. Comput. Syst.*, vol. 135, pp. 85–94, 2022.
[22] L. Liu, J. Li, J. Lv *et al.*, "Privacy-preserving and secure industrial big data analytics: A survey and the research framework," *IEEE Internet Things J.*, pp. 18 976–18 999, 2024.
[23] N. Wu, L. Xu, and L. Zhu, "A blockchain based access control scheme with hidden policy and attribute," *Future Gener. Comput. Syst.*, vol. 141, pp. 186–196, 2023.
[24] Q. Hu, C. Huang, G. Zhang *et al.*, "Towards accountable and privacy-preserving blockchain-based access control for data sharing," *J. Inf. Secur. Appl.*, vol. 85, p. 103866, 2024.
[25] T. H. Yuen, J. K. Liu, M. H. Au *et al.*, "k-times attribute-based anonymous access control for cloud computing," *IEEE Trans. Computers*, vol. 64, no. 9, pp. 2595–2608, 2015.
[26] H. Fang, L. Xu, G. Nan *et al.*, "Accountable distributed access control with privacy preservation for blockchain-enabled internet of things systems: A zero-trust security scheme," *IEEE Internet Things J.*, vol. 12, no. 11, pp. 17 936–17 947, 2025.
[27] L. Karimi, M. Aldairi, J. Joshi *et al.*, "An automatic attribute-based access control policy extraction from access logs," *IEEE Trans. Dependable Secur. Comput.*, vol. 19, no. 4, pp. 2304–2317, 2022.
[28] X. Geng, Y. Wen, Z. Mo *et al.*, "An access control framework for multilayer rail transit systems based on trust and sensitivity attributes," *Applied Sciences*, vol. 13, no. 23, p. 12904, 2023.
[29] L. Bai, K. Fan, Y. Bai *et al.*, "Cross-domain access control based on trusted third-party and attribute mapping center," *J. Syst. Archit.*, vol. 116, p. 101957, 2021.
[30] J. Zhang, X. Li, M. Zhang *et al.*, "QAE-BAC: **Q**uantifiable **A**nonymity and **E**fficiency in **B**lockchain-based **A**ccess **C**ontrol with attribute," *CoRR*, vol. abs/2507.10927, 2025.
[31] Q. Xia, E. B. Sifah, A. Smahi *et al.*, "BBDS: blockchain-based data sharing for electronic medical records in cloud environments," *Inf.*, vol. 8, no. 2, p. 44, 2017.
[32] D. Yang and W. Tsai, "An optimized encryption storage scheme for blockchain data based on cold and hot blocks and threshold secret sharing," *Entropy*, vol. 26, no. 8, p. 690, 2024.
[33] J. Zhang, X. Li, R. Feng *et al.*, "From isolation to integration: A reputation-backed auditable model for cohort data sharing," *IEEE Trans. Dependable Secur. Comput.*, vol. 23, no. 1, pp. 637–654, 2026.
[34] Y. Liu, X. Xing, Z. Tong *et al.*, "Secure and scalable cross-domain data sharing in zero-trust cloud-edge-end environment based on sharding blockchain," *IEEE Trans. Dependable Secur. Comput.*, vol. 21, no. 4, pp. 2603–2618, 2024.
[35] P. Biswas, R. S. Sandhu, and R. Krishnan, "Attribute transformation for attribute-based access control," in *CODASPY, Scottsdale, Arizona, USA, March 24, 2017*. ACM, 2017, pp. 1–8.
[36] X. Lin, Y. Zhang, C. Huang *et al.*, "An access control system based on blockchain with zero-knowledge rollups in high-traffic iot environments," *Sensors*, vol. 23, no. 7, p. 3443, 2023.
[37] S. Panda, S. Sahoo, R. Halder *et al.*, "Contextual attribute-based access control scheme for cloud storage using blockchain technology," *Softw. Pract. Exp.*, vol. 54, no. 10, pp. 2042–2062, 2024.
[38] A. D. Santis, A. L. Ferrara, B. Masucci *et al.*, "An information-theoretic approach to anonymous access control," in *ISIT, Athens, Greece, July 7-12, 2024*. IEEE, 2024, pp. 3326–3331.
[39] M. Bellare, C. Namprempre, and G. Neven, "Security proofs for identity-based identification and signature schemes," *J. Cryptol.*
[40] K. I. Ahmed, M. Tahir, S. L. Lau *et al.*, "Dataset for authentication and authorization using physical layer properties in indoor environment," *Data in Brief*, p. 110589, 2024.
[41] A. S. Al-doori, A. H. Aboud, and Z. Sedrah, "Optimization and characterization of dextransucrase production by local leuconostoc mesenteroides," *International Journal of Scientific Engineering and Applied Science*, vol. 1, no. 6, pp. 476–483, 2015.

BIOGRAPHY SECTION

Jie Zhang (Student Member, IEEE) received the M.S. degree in computer application technology from the Yunnan Normal University in 2022. He is currently working toward the Ph.D. degree in computer science and technology with the Tianjin University. His research interests include efficient and secure data sharing within blockchain systems.

Xiaohong Li (Member, IEEE) received the Ph.D. degree in computer application technology from Tianjin University in 2005. She is currently a Full Tenured Professor with the Department of Cyber Security, College of Intelligence and Computing, Tianjin University. Her research interests include knowledge engineering, trusted computing, and security software engineering.

Mengke Zhang received a bachelor's degree from Chongqing University in 2022 and M.S. degree in computer science and technology from Tianjin University in 2025. She is currently affiliated with the Baseband and Computing Department, Shanghai ZTE Software Co., Ltd. Her research interests include efficient and secure access control and data sharing in the blockchain system.

Ruitao Feng is a Lecturer at Southern Cross University, Australia. He received the Ph.D. degree from the Nanyang Technological University. His research centers on security and quality assurance in software-enabled systems, particularly AI4Sec & SE. This encompasses learning-based intrusion/anomaly detection, malicious behavior recognition for malware, and code vulnerability detection.

Shanshan Xu received the M.S. degree in computer application technology from the Yunnan Normal University in 2022. She is currently working toward the Ph.D. degree in physical geography with the School of Geographic Sciences, East China Normal University. Her research interests include Atmospheric vapor, machine learning and model building.

Zhé Hóu is a Senior Lecturer at Griffith University, Australia. He obtained his Ph.D. degree from the Australian National University in 2015. His research mainly focuses on automated reasoning, formal methods, AI, quantum computing and blockchain.

Guangdong Bai (Member, IEEE) received the B.S. and M.S. degrees in computing science from Peking University in 2008 and 2011, respectively, and the Ph.D. degree in computing science from the National University of Singapore in 2015. He is currently an Associate Professor with the Department of Computer Science, City University of Hong Kong, China. His research interests include cyber security, software engineering, and machine learning.