*electronics*

MDPI

*Article*

# False Data Injection Screening for Temperature Monitoring of Distribution Transformers Based on Lightweight Blockchain and Winding Structural Fingerprinting

**Ying Guan [1,*], Sheng Ma [1], Junqi Yang [1], Weitao Meng [1], Bo Zhang [2] and Xingcheng Yu [2]**

[1] School of Computer Science and Technology, Shenyang Institute of Engineering, Shenyang 110136, China; masheng@sie.edu.cn (S.M.); 2024523202@sie.edu.cn (J.Y.); 2024523219@stu.sie.edu.cn (W.M.)
[2] School of Electrical Engineering, Shenyang University of Technology, Shenyang 110870, China; solomon@sut.edu.cn (B.Z.); yuxingcheng@smail.sut.edu.cn (X.Y.)
* Correspondence: guanying@sie.edu.cn

**Abstract**

Temperature monitoring underpins the protection and condition assessment of distribution transformers, whose dominant failure mode is overheating-induced insulation aging; this makes temperature data the prime target of false data injection (FDI) attacks. Because FDI acts on data rather than on the physical asset, this paper exploits the winding structure, which data injection cannot alter, as a physical trust anchor: a structural fingerprint is extracted from port measurements through frequency-dependent harmonic resistance; the systematic temperature-response differences between wire-wound and foil-wound windings build an independent physical cross-check channel; and a lightweight PBFT consortium blockchain notarizes the fingerprint baseline so that a compromised local database cannot rewrite it. Structural, physical, and command consistency indicators classify input-, output-, and model-parameter-tampering attacks. Finite-element simulations and 100-round randomized attack tests show a 9.6% structural temperature deviation that provides a clear physical screening margin: the three attack types are detected and classified within 4 s at a 100% detection rate (95% CI [96.4, 100]%) and a 0.3% false-alarm rate, while a stealthy ramp injection is caught after about 105 min and a sub-threshold constant offset is caught by trend monitoring in about 4 h. A seven-node chain implemented with real ECDSA-secp256k1 signatures achieves 0.42–0.90 s finality per notarization, 8.9–19.0 notarizations/s, and a 0.47 s view change, two orders of magnitude above the field notarization demand, meeting lightweight deployment requirements.

**Keywords:** distribution transformer; temperature monitoring; winding structural fingerprint; lightweight blockchain; false data injection; physical-layer verification

Check for updates

Academic Editor: José Matas

Received: 24 September 2026
Revised: 16 October 2026
Accepted: 21 October 2026
Published: 23 October 2026

**Copyright:** © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

## 1. Introduction

### 1.1. Background

Distribution transformers are critical assets in power distribution networks, and their operating-state awareness directly supports grid dispatching and protection [1,2]. Unlike bulk power transformers, whose typical failures are insulation breakdown and short-circuit impact, distribution transformers are small in capacity, subject to large load fluctuations, and operate in harsh environments; overheating and the resulting thermal aging of insulation dominate their failure statistics. Temperature monitoring data therefore constitute the core input for state awareness, protection decisions, overload management, and lifetime assessment. However, the

same networked upgrade that enables real-time monitoring also exposes temperature data to cyber threats [3]. False data injection (FDI) attacks manipulate temperature sensor readings, alarm thresholds, or thermal-model parameters through carefully crafted attack vectors, thereby bypassing conventional bad-data detection [4] and driving monitoring and protection systems to wrong decisions: forged high temperatures can induce false tripping and interrupt the power supply, whereas forged low temperatures can mask real overheating, accelerate insulation aging, and eventually lead to burnout; both consequences have been demonstrated in attack-and-defense studies [5,6].

A point that must be clarified is that FDI attacks target monitoring data rather than the physical asset. An attacker can alter any data reported over the network, but cannot change the real geometry and material properties of the winding through data injection, even though the winding structure alone determines the physical law of the temperature distribution. The physical structure of the equipment therefore forms an information channel that is independent of the network layer: the true structure cannot be shaken by false data, and the temperature behavior expected from the true structure is uniquely determined. Existing FDI detectors, including statistical anomaly detection, machine-learning classifiers, and state-estimation residual analysis [5,6], all operate in the cyber layer and essentially verify data with data. When a temperature attack is designed to be consistent with historical distributions or physical constraints, these methods can hardly distinguish a cyber attack from a genuine equipment fault [6]; more critically, the detection model and the data under inspection reside in the same trust domain, so once the thermal-model parameters or the verification baseline themselves are tampered with, the defense fails entirely [3,4]. What the monitoring system lacks is not a more sophisticated data algorithm, but a verification mechanism that is physically anchored, immutably anchored on-chain, and independently cross-checking.

The winding structure happens to provide the physical basis for such a mechanism. Distribution-transformer windings are mainly configured either as wire-wound (round or rectangular conductors) or as foil-wound (copper or aluminum sheets), and the two configurations differ significantly in electromagnetic loss and temperature response under harmonic excitation [7]. In foil windings, the large cross-sectional dimensions make the skin and proximity effects concentrate the current density towards the edges, producing local hot spots that average-temperature estimation cannot capture [8]. The temperature behavior expected from the true structure of a given transformer is uniquely determined, whereas a calculation based on a wrong structural assumption can deviate from the truth by nearly 10%. This fact has a twofold meaning: for temperature estimation it is an error source, but for security verification it is a valuable asset, because temperature data reported by SCADA that is inconsistent with the physical law dictated by the true structure can be identified as injected. Structural identification thus becomes the reconciliation benchmark for screening temperature data. Meanwhile, modern smart distribution-terminal procurement and maintenance specifications emphasize plug-and-play operation, and practices that rely on manual nameplate entry or offline database matching can hardly support large-scale deployment; structural knowledge must, therefore, be obtainable automatically from port electrical measurements alone and maintained immutably with the help of blockchain.

### *1.2. Related Work and Open Gaps*

Since Liu et al. [4] formalized false data injection attacks against power-system state estimation, a large body of detection research has emerged. Conventional residual-based bad-data detection can be bypassed by attack vectors that keep the measurements consistent [9]; learning-based methods, including deep reinforcement learning and graph neural networks, depend on historical training data and may fail when attack patterns

deviate from the training distribution [10]. Strategic protection approaches invalidate attacker knowledge by actively perturbing system parameters or protecting critical measurements against data-injection attacks [10], and distributed observer networks localize attacks through spatial measurement redundancy, but both require additional hardware or specific attack-model assumptions. On the application of blockchain, existing studies have applied it to audit trails, data-integrity protection, and lightweight security architectures in smart grids and the industrial Internet of Things [11,12]; however, such schemes notarize data records themselves, so if the on-chain baseline is polluted before notarization, immutability merely fossilizes the error. Most frameworks also fail to answer the fundamental question of where the trust anchor notarized on-chain comes from. The common limitation of existing work is that the verification of temperature monitoring data stays in the data domain: it is neither anchored to a physical attribute that attackers cannot forge, nor protected against tampering of the verification baseline itself.

For winding-structure identification, frequency response analysis (FRA) detects mechanical deformation by comparing frequency-response features [13], but requires offline excitation and specialized equipment. Harmonic-based methods distinguish winding types through the frequency-dependent resistance caused by eddy-current effects [7], but they target loss calculation and fault diagnosis rather than security; the fact that the true structure is a physical fact that temperature data injection cannot alter, and can therefore serve as a physical trust anchor notarized on-chain, has not been recognized.

For hot-spot temperature modeling, the IEC 60076-7 loading guide provides a standardized thermal model [14], and online estimation methods based on the thermal equivalent circuit can derive winding and hot-spot temperatures from electrical quantities [15]; however, both implicitly assume a uniform winding structure and thus cannot perform structure-sensitive reconciliation. Numerical thermal models and coupled-field simulations can predict the winding temperature distribution and local hot spots accurately [16,17], but require detailed geometry and material parameters that are unavailable in online scenarios; dynamic thermal models dedicated to distribution transformers [18] still ignore the non-uniform current distribution of foil windings. In summary, three gaps remain: first, there is no plug-and-play means of obtaining the true structure of a transformer from port electrical measurements alone; second, there is no mechanism to notarize structural knowledge on-chain as a tamper-proof verification baseline, because reconciliation is meaningless if the baseline itself can be altered; third, there is no mechanism that uses the physical law that structure determines temperature behavior to cross-check temperature monitoring data independently of the cyber layer.

### *1.3. Contributions*

To close these gaps, this paper fuses lightweight blockchain notarization with winding structural fingerprinting into an FDI screening framework for distribution-transformer temperature monitoring data. The building blocks—harmonic-resistance analysis, transformer thermal modeling, PBFT consensus, and blockchain integrity protection—are mature individually; the contribution is their specific combination, together with its validation:

(1) the winding fingerprint used as a physical trust reference, unforgeable without physical access to the equipment, requiring no internal geometric parameters;
(2) a structure-sensitive independent temperature cross-check channel whose correction model compresses the reconciliation-baseline error from 9.6% to 0.26%;
(3) blockchain notarization protecting the verification baseline rather than the data, anchored across mutually distrustful administrative domains and validated on a real seven-node test network with measured performance;

(4) a three-indicator attack classification—structural, physical, and command consistency—with randomized multi-round statistical evaluation including stealthy attack variants.

The remainder of this paper is organized as follows. Section 2 presents the materials and methods. Section 3 reports the finite-element simulations and attack-scenario tests. Section 4 discusses the findings, and Section 5 concludes the paper.

## 2. Materials and Methods

### *2.1. Framework Overview and Threat Model*

#### 2.1.1. Overall Framework

The proposed framework is organized around the four-layer logic of "physical trust anchor, independent cross-check channel, on-chain tamper-proof baseline, and three-indicator verification", as shown in Figure 1. The physical layer consists of smart distribution terminals equipped with harmonic-analysis modules, which extract the structural fingerprint from port measurements and independently compute the expected temperature behavior from the true structure. The edge layer deploys lightweight blockchain consortium nodes, which notarize the fingerprint baseline on-chain and perform consistency verification. The application layer serves the SCADA system and the cybersecurity monitoring center with physically cross-validated temperature data and alarms. In this framework, the data sources of the structural fingerprint and the physical temperature channel are completely independent of the SCADA temperature monitoring link, forming a second information channel that is parallel to, and mutually distrustful of, the network layer; the baseline used for verification is maintained by blockchain consensus and is independent of terminal-local storage. To forge temperature data while deceiving both channels and rewriting the on-chain baseline, an attacker would need network intrusion, local penetration, and the ability to break the consortium consensus simultaneously, which raises the attack cost dramatically.

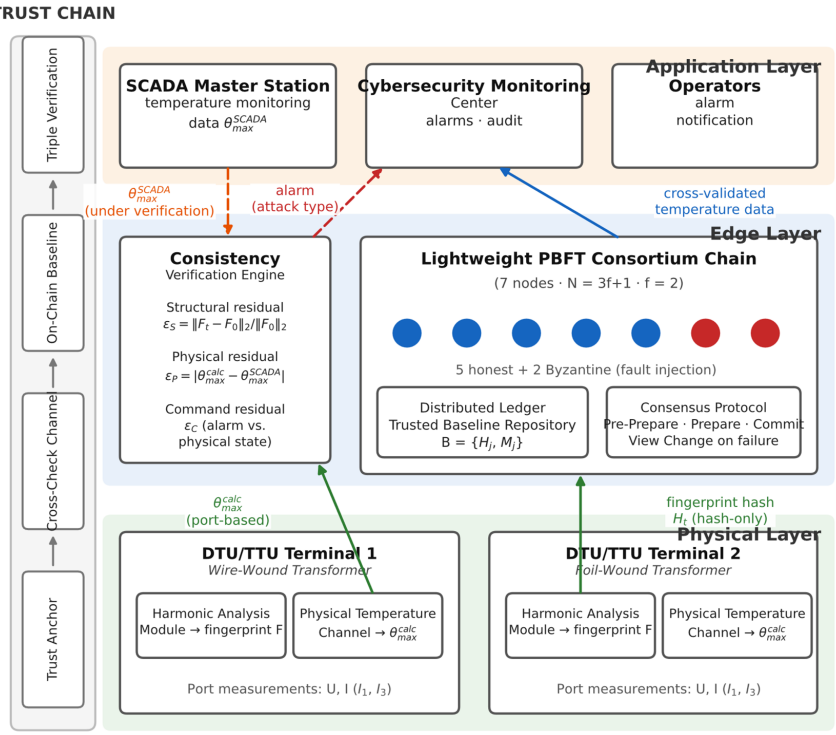


**Figure 1.** Overall architecture of the proposed framework: physical trust anchor, independent cross-check channel, on-chain tamper-proof baseline, and three-indicator consistency verification.

2.1.2. Threat Model and Trust Assumptions

The framework assumes a network attacker able to: (i) passively eavesdrop and actively modify, inject, and replay messages on the SCADA temperature monitoring link; (ii) write to the terminal-local database; and (iii) submit forged or replayed transactions to the consortium chain. The trust assumptions bound the attacker: (a) the winding geometry and material cannot be altered without a physical maintenance event; (b) private keys reside in the terminal secure element and cannot be extracted remotely; (c) no more than f = 2 of the 7 consortium nodes are controlled, so no single administrative domain can rewrite on-chain history; and (d) the measurement front-end (VT/CT inputs), the harmonic-analysis module, and the terminal firmware behave as specified—terminal compromise is covered by secure boot and remote attestation, recommended as deployment hardening and listed as future work. Under these assumptions, the fingerprint is unforgeable without physical access: data-only manipulation changes either the reported values (caught by physical consistency) or the baseline copy (caught by the on-chain integrity check), but cannot make the port-measured fingerprint match a forged structure. Throughout the paper, "unforgeable" is to be read as "unforgeable under the stated threat model".

### *2.2. Structural Fingerprinting Based on Harmonic Resistance*

The effectiveness of the structural fingerprint rests on the skin effect, a physical law determined by conductor geometry and material properties and independent of network state; the fingerprint therefore cannot be forged by data means [7,19]. When an alternating current flows through a transformer winding, the time-varying magnetic field induces eddy currents in the conductor and redistributes the current density towards the surface. For foil windings with a large aspect ratio, the current-density distribution $J(y)$ across the width direction follows the hyperbolic-cosine solution of the one-dimensional eddy-current equation, symmetric about the conductor center and approximately exponential towards each edge:

$$J(y) = J_0 \frac{\cosh(y/\delta)}{\cosh(h/(2\delta))} \tag{1}$$

where $y$ is the coordinate from the conductor center across the width, $y \in [-h/2,\ h/2]$, $J_0$ is the value at the center, $h$ is the conductor width, and $\delta$ is the skin depth at the operating frequency:

$$\delta = \sqrt{\frac{2}{\omega \mu \sigma}} \tag{2}$$

where $\omega$ is the angular frequency, $\mu$ is the permeability, and $\sigma$ is the conductivity. The non-uniform current distribution makes the effective AC resistance higher than the DC resistance, and the enhancement grows with frequency and conductor size.

The structural fingerprint is defined as a feature vector derived from the harmonic resistance characteristics measured at the transformer ports; its extraction, shown in Figure 2, covers four stages: data acquisition, harmonic decomposition, equivalent-resistance calculation, and fingerprint construction. First, under normal load conditions, the built-in harmonic-analysis module of the smart terminal acquires the port voltage and current waveforms. Fast Fourier transform or wavelet decomposition is then applied to extract the fundamental $(I_1,\ U_1)$ and third-harmonic $(I_3,\ U_3)$ components. The equivalent winding resistance at each frequency is subsequently computed by the instantaneous vector method:

$$R_{eq}(f) = \mathrm{Re}\left\{ \frac{U_1 - U_2/k}{I_1} \right\} \tag{3}$$

where $U_1$ and $I_1$ are the primary-side voltage and current phasors at frequency $f$, $U_2$ is the secondary-side voltage phasor, and $k$ is the turns ratio. The structural fingerprint vector $F$ is finally constructed as:

$$k_{RF} = \frac{R_{eq}(f_3)}{R_{eq}(f_1)} \tag{4}$$

$$F = [k_{RF}, R_{eq}(f_1), R_{eq}(f_3), THD_I, THD_U]^T \tag{5}$$

where $k_{RF} = \frac{R_{eq}(f_3)}{R_{eq}(f_1)}$ is the harmonic resistance coefficient, and THD denotes the total harmonic distortion. The harmonic-resistance ratios ($k_{\mathrm{RF}}$, $k_{\mathrm{RF5}}$) are ratios of quantities measured at the same instant and are therefore insensitive to slow load changes; they serve as the primary classification features. $THD_{\mathrm{I}}$ and $THD_{\mathrm{U}}$ are retained only as auxiliary context features with a wide tolerance inside $\tau_{\mathrm{S}}$. Since the fingerprint originates from the real electromagnetic response at the equipment ports, an attacker cannot forge a valid fingerprint without physically replacing the winding, and such replacement is a maintenance event that the verification mechanism itself captures (see the security analysis in Section 2.7).

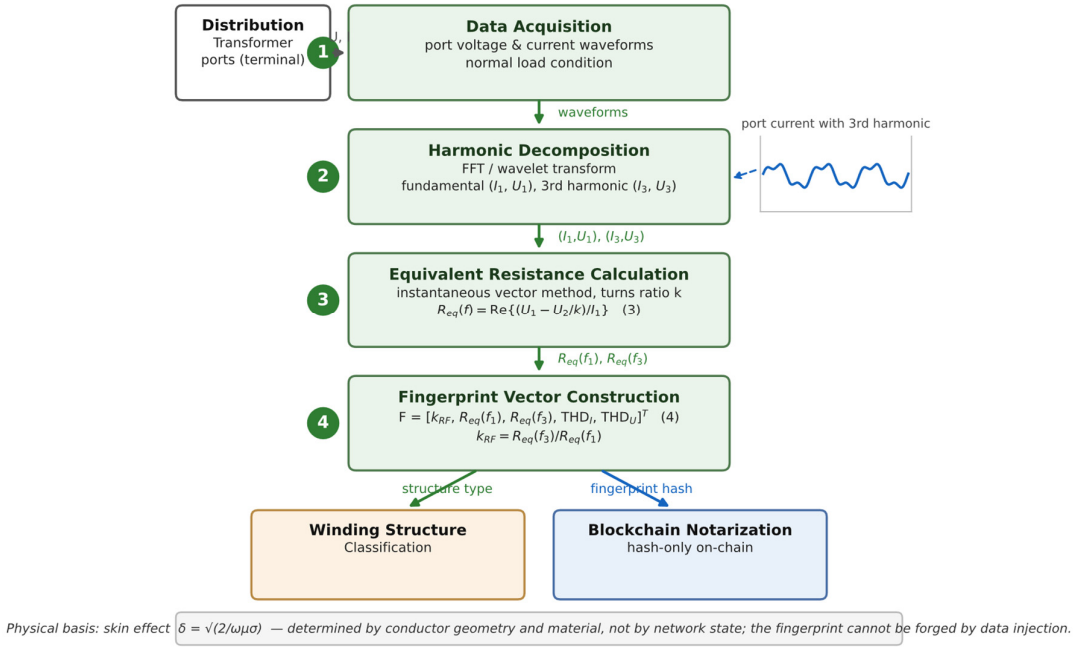


**Figure 2.** Structural fingerprint extraction procedure: data acquisition, harmonic decomposition, equivalent-resistance calculation, and fingerprint vector construction. Arrows indicate the data flow: the classified winding structure feeds the temperature cross-check channel, and the fingerprint hash is notarized on the consortium chain.

### *2.3. Winding-Structure Classification Criteria*

Based on a finite-element parametric sweep over conductor cross-section dimensions, the following classification criteria are established: wire-wound windings correspond to $k_{\mathrm{RF}} \in [1.00, 1.05]$; foil-wound windings correspond to $k_{\mathrm{RF}} > 1.10$; when $k_{\mathrm{RF}}$ falls in the uncertain region $(1.05, 1.10]$, higher-order harmonic analysis (5th and 7th) is required for further discrimination. The thresholds already account for manufacturing tolerances, temperature effects, and measurement uncertainty [7]. For transformers in the uncertain region, the fifth-harmonic resistance ratio $k_{RF5} = \frac{R_{eq}(f_5)}{R_{eq}(f_1)}$ provides additional discrimination, because foil windings typically exhibit $k_{\mathrm{RF5}} > 1.15$ whereas wire windings remain below 1.08. The criteria are summarized in Table 1. Under insufficient harmonic content ($THD_{\mathrm{I}} < 2\%$, e.g., at very light load), classification is deferred: the terminal holds the last valid fingerprint and raises a low-harmonic-content flag instead of producing an unreliable classification.

**Table 1.** Winding-structure classification criteria.

| $k_{RF}$ Range | Classification | Subsequent Processing |
|---|---|---|
| [1.00, 1.05] | Wire-wound | Enter the temperature cross-check channel directly |
| (1.05, 1.10] | Uncertain region | Discriminate with 5th/7th-harmonic resistance ratios |
| >1.10 | Foil-wound | Apply the structural correction factor before cross-checking |

### *2.4. Structure-Aware Independent Temperature Cross-Check Channel*

The discrimination capability of the physical cross-check channel stems from the strong dependence of temperature response on structure: the same set of injected false temperature data produces residuals of very different magnitudes relative to the physically computed values under different true structures. To ensure that the residual $\varepsilon_P$ reflects the authenticity of the temperature data rather than model error, the structure-related correction must be introduced into the physical computation. For transformers confirmed as foil-wound by fingerprinting, the hot-spot temperature rise is corrected from the average winding temperature rise through the harmonic resistance coefficient:

$$\Delta T_{max} = k_{RF} \alpha \Delta T_{avg} \tag{6}$$

where $\Delta T_{avg}$ is the average temperature rise computed from resistance variation, $\alpha$ is the heat-dissipation coefficient determined by the cooling mode (typically $\alpha = 1.0$ for natural oil flow), and $k_{RF}$ serves as the structure-related correction factor that compensates for the edge current concentration. The corrected hot-spot temperature is:

$$\theta_{max} = \theta_{amb} + \Delta T_{oil} + H \Delta T_{max} \tag{7}$$

where $\theta_{amb}$ is the ambient temperature, $\Delta T_{oil}$ is the top-oil temperature rise, and $H$ is the hot-spot factor (typically 1.1 for distribution transformers); for five-limb core transformers, an additional core-structure correction factor of 0.9 accounts for the reduced distance from the winding to the top oil. It must be emphasized that the purpose of the correction model is not general temperature-measurement accuracy, but compressing the inherent error of the reconciliation baseline to a level far below the injected temperature deviation (the simulations show a reduction from 9.6% to 0.26%), so that $\varepsilon_P$ becomes a sensitive indicator of data authenticity.

Physical justification of Equation (6). Under steady-state thermal conditions, the hot-spot temperature rise is approximately proportional to the local winding loss. In foil windings, the skin and proximity effects concentrate the current towards the conductor edges, scaling the local loss by the same factor that raises the AC resistance; the harmonic-resistance ratio $k_{RF}$ measures precisely this AC-resistance enhancement at harmonic frequencies. Hence the ratio of edge hot-spot rise to average rise is approximated by $k_{RF}$—an engineering first-order relation whose accuracy is verified by the finite-element results reported below (baseline error reduced from 9.6% to 0.26%) and by the monotonically increasing edge-to-average heat-flux ratio observed in the conductor parametric sweep.

Coefficient determination. $\alpha$ is set by the cooling mode ($\alpha$ = 1.0 for natural oil flow, fitted by cooling configuration in deployment); $H$ is the conventional hot-spot factor for distribution transformers ($H$ = 1.1) following IEC 60076-7; the 0.9 factor for five-limb cores follows from the finite-element comparison of the two core structures: at equal top-oil temperature, the three-limb model gives $\theta_{hs} = \theta_{oil} + 1.1\Delta T = 83.5$ °C, while the five-limb model, whose winding ends sit closer to the top oil, gives 81 °C by the same expression versus a simulated 75 °C, i.e., $75/83.5 \approx 0.9$.

Data provenance. Every quantity in Equations (6) and (7) is obtained locally: $\theta_{amb}$ from the terminal-mounted ambient sensor; $\Delta T_{oil}$ estimated on-site from the port load current through the standard top-oil thermal model of IEC 60076-7 (not from the SCADA register); $\Delta T_{avg}$ from the winding resistance variation. Table 2 summarizes the provenance and network exposure of each variable.

**Table 2.** Data provenance of the independent physical channel.

| Variable | Source | Exposed to Network? |
|---|---|---|
| θ_amb | Terminal-mounted ambient sensor | No |
| ΔT_oil | On-site estimate from port current (IEC 60076-7 top-oil model) | No |
| ΔT_avg | Winding resistance variation (port measurement) | No |
| θ_maxˆSCADA | SCADA link | Yes (verification target) |

### *2.5. Lightweight Blockchain Notarization Mechanism*

The credibility of the reconciliation mechanism depends on the immutability of the verification baseline itself: if an attacker can modify the fingerprint or the thermal-model correction coefficients in the local database of the terminal, and then carefully craft false temperature data consistent with the corrupted model, any reconciliation becomes useless. A lightweight blockchain notarization is therefore introduced, whose protection target is the baseline rather than the temperature data. Considering the limited computation, storage, and bandwidth of distribution IoT terminals, a practical Byzantine fault tolerance (PBFT) variant optimized for IoT scenarios is adopted to build a permissioned consortium chain [20,21], whose computational overhead is far lower than proof-of-work public chains and which provides deterministic transaction finality.

The network and consensus structure are as follows. Consortium nodes are composed of edge servers operated by utility operators together with security-screened distribution terminals; the deployment size is $N = 3f + 1$ (tolerating $f$ Byzantine nodes; the experiments in this paper use seven nodes, i.e., $f = 2$). Consensus follows the PBFT three-phase flow (pre-prepare, prepare, commit): a transaction is proposed by the primary node and reaches finality once at least $2f + 1$ node signatures are collected, with no fork or rollback; when the primary times out or misbehaves, a view change elects the next primary. Inter-node communication uses TLS-encrypted channels, transactions are signed with ECDSA (secp256k1) key pairs, and private keys are kept in the secure element of the terminal.

The transaction and block structure is designed as a hash-only notarization mode. Each transaction contains the structural-fingerprint hash digest, the winding-type classification, the correction coefficients, the timestamp, and the terminal signature; it contains no raw electrical parameters or geometric information. A block is generated every 15 min or immediately upon a fingerprint-update event (re-registration after maintenance or replacement). On the storage side, each notarization contributes only a 32-byte hash digest (about 180 bytes per full transaction including metadata and signature); at a 15-min period, one terminal generates about 6.3 MB per year, and a hundred-terminal deployment stays below 0.7 GB, far below the capacity constraints of edge servers.

The notarization flow, shown in Figure 3, comprises three stages: registration, on-chain storage, and baseline-repository construction. When a new transformer is commissioned, the terminal extracts the initial structural fingerprint $F_0$ and computes the SHA-256 digest $H_0 = SHA - 256(F_0 \| T_0 \| W_0)$, where $T_0$ is the winding type and $W_0$ is the set of correction coefficients. The digest $H_0$, together with metadata (transformer ID, commissioning date, terminal ID), is packaged into a transaction and submitted to the blockchain network; after PBFT consensus, it is immutably recorded on the distributed ledger. The accumulated

on-chain records form the trusted baseline repository $B = \{H_i, M_i\}_{i=1}^{N}$, where $M_i$ denotes the metadata of transformer $i$. The hash-only on-chain design keeps sensitive parameters (precise resistance values and geometry) locally at the terminal, addressing data-privacy concerns by design. The construction of the blockchain test network and the performance and security experiments are detailed in Section 3.5.

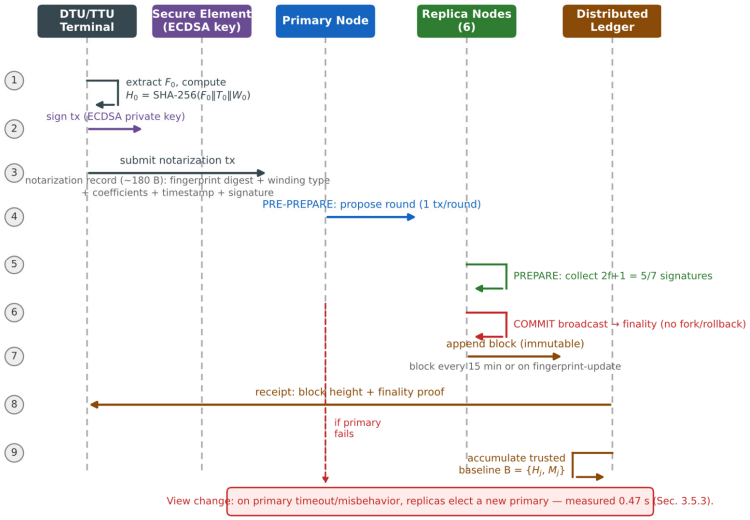


**Figure 3.** Sequence of the structural fingerprint notarization: registration, three-phase PBFT consensus, and baseline-repository construction, with view change on primary failure. The circled numbers indicate the protocol steps in temporal order.

Why a consortium chain is necessary. If the baseline were kept in an ordinary database or a locally signed audit log, the notarization authority and the monitored system would belong to the same trust domain: an attacker compromising that domain could re-notarize a corrupted baseline—the "garbage-in, garbage-out" problem. The baseline must therefore be anchored across mutually distrustful administrative domains (multiple utility operators, vendors, and auditors); PBFT consensus makes rewriting the baseline require collusion of more than f independent parties, which is exactly the property that defeats baseline-tampering attacks. A secure-element-signed log, by contrast, protects records against retroactive alteration but remains within one trust domain and cannot prevent a wrong baseline from being notarized in the first place.

Authorized re-registration. A legitimate re-registration (after maintenance or winding replacement) follows a two-factor authorization: an operator ticket triggers the procedure, and the new fingerprint is notarized only when signed by both the utility operator key and the on-site terminal key. The new record references the previous baseline on-chain, preserving full history; the verification engine switches to the new baseline only after both signatures validate and the notarization transaction is final; otherwise, the previous baseline remains in force.

### *2.6. Three-Indicator Consistency Verification and FDI Attack Classification*

During normal operation, the terminal periodically extracts the real-time structural fingerprint $F_t$ and first verifies the integrity of the locally stored baseline: the SHA-256 digest recomputed over the local baseline vector $F_0$ must equal the on-chain digest. This check is binary and guarantees only that $F_0$ has not been altered; the chain stores no raw parameters, and the raw baseline vector never leaves the terminal. The structural consistency residual is then defined on fingerprint vectors as

$$\varepsilon_s = \|F_t - F_0\|_2 / \|F_0\|_2 \quad (8)$$

If $\varepsilon_S > \tau_S$, a structural inconsistency alarm is raised, indicating that the reconciliation baseline may have been tampered with, or that an unauthorized winding replacement has occurred; even if the local baseline copy is altered, the verdict is unaffected, because its integrity is anchored to the on-chain digest, which the attacker cannot rewrite without controlling more than 2f + 1 consortium nodes. The hot-spot temperature computed independently at the port is cross-checked against the temperature reported by SCADA, and the physical consistency residual is:

$$\varepsilon_P = \left| \theta_{\max}^{calc} - \theta_{\max}^{SCADA} \right| \quad (9)$$

Because the two residuals alone identify the attacked asset but not the attack surface, a third indicator, the command-consistency flag $\varepsilon_C$, is introduced: $\varepsilon_C = 1$ if a control or alarm action contradicts the physical state—an alarm or trip asserted while the physically computed temperature does not support it ($\theta_{max}{}^{calc}$ below the legitimate threshold), or a configuration write (e.g., alarm threshold) without an authorized maintenance session—and $\varepsilon_C = 0$ otherwise. Threshold selection follows $\tau = \max(\mu_n + 5\sigma_n, \varepsilon_{model+margin})$, where $\mu_n$ and $\sigma_n$ are the mean and standard deviation of the residual under normal operation with randomized load and noise (for $\varepsilon_P$, $\mu_n = 0.2$ °C and $\sigma_n \approx 0.05$ °C), and $\varepsilon_{model} = 0.26\%$ is the corrected baseline model error; this gives $\tau_P = 3$ °C and $\tau_S = 0.05$, keeping the normal-operation false-alarm rate at 0.3% with a minimum detectable injection of about 4 °C. Table 3 summarizes the classification, and Figure 4 shows the decision logic.

**Table 3.** Classification of temperature FDI attacks.

| Attack Type | $\varepsilon_S$ | $\varepsilon_P$ | $\varepsilon_C$ | Interpretation |
|---|---|---|---|---|
| Input-type FDI | Normal | Anomalous | 0 | Temperature data manipulated; physical channel and on-chain baseline intact |
| Output-type FDI | Normal | Normal | 1 | Alarm threshold or other control commands spoofed; measurements valid |
| Model-parameter tampering | Anomalous | (Anomalous) | (0/1) | Reconciliation baseline attacked; on-chain integrity check fails |

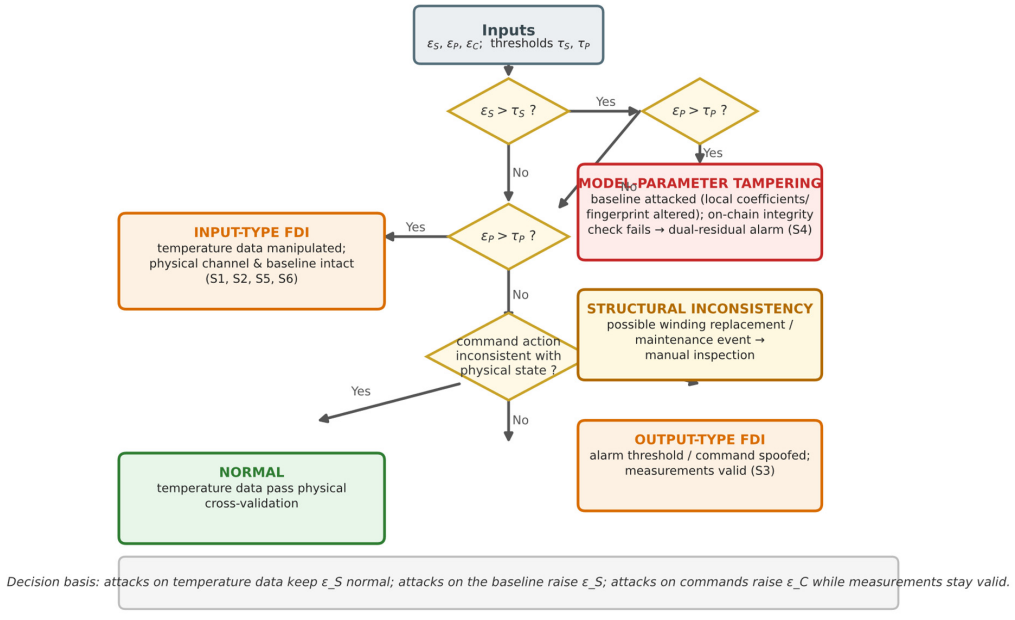


**Figure 4.** Decision logic for classifying temperature FDI attacks based on the structural, physical, and command consistency indicators.

### *2.7. Security Analysis*

The security of the framework rests on four core properties. First, the on-chain immutability of the baseline: once the structural baseline is notarized, it cannot be retroactively modified; even if an attacker completely compromises the local database of the terminal, the on-chain baseline remains intact, and model-parameter tampering targeting the baseline is immediately captured by the structural consistency check. Second, the physical unforgeability of the fingerprint under the threat model of Section 2.1.2: the harmonic resistance characteristic is determined by the real geometry and material of the winding, so no data-only means can forge a valid fingerprint; the only way is to physically replace the winding, which belongs to physical intrusion beyond the boundary of network-based FDI attacks, and the replacement event itself changes the fingerprint and triggers an alarm. Third, the independence of the verification channels: the physical temperature computation relies only on port electrical quantities, and its data source, computation model, and on-chain baseline are all independent of the SCADA temperature link; compromising a single point is insufficient to deceive both residual checks simultaneously. Fourth, lightweight feasibility: with a seven-node deployment, PBFT consensus achieves finality within seconds, and the storage overhead of hash-only notarization is 32 bytes per transformer per notarization, meeting the resource constraints of distribution terminals (measured performance and fault-injection tests are reported in Section 3.5).

## 3. Results

### *3.1. Experimental Platform and Finite-Element Setup*

The experimental platform consists of three parts:

(1) an ANSYS Maxwell/Fluent (Ansys 2021 R2; Ansys, Canonsburg, PA, USA) electromagnetic-thermal coupling simulation environment that generates the physical ground truth of the temperature distribution for different winding structures;
(2) an emulated SCADA temperature monitoring link in which the terminal and the master station exchange temperature and alarm data over Modbus TCP, with the attack injection point located at a man-in-the-middle position on this link (temperature data are chosen as the attack object because overheating and the resulting insulation aging are the dominant failure modes of distribution transformers, making temperature data the attacker's preferred target);
(3) blockchain test network of seven node processes on an isolated test host (corresponding to empty $N = 3f + 1$, $f = 2$ in Section 2.5) running the lightweight PBFT consensus implementation, detailed in Section 3.5.

Two transformer models are built: Model A is a wire-wound winding with round copper conductors of 2 mm diameter and 1000 turns; Model B is a foil-wound winding with copper foil of 1.5 mm thickness, 1400 mm height, and 10 turns. The two models share the same core structure (three-phase three-limb), rated capacity (1000 kVA), and cooling conditions (natural oil flow, ambient temperature 22 °C). The electromagnetic-thermal coupling procedure is as follows: a 2D eddy-field analysis in Maxwell computes the frequency-dependent winding losses; the loss distribution is exported to Fluent for thermal-fluid coupling analysis; and the procedure iterates until convergence of the temperature-dependent resistivity, as shown in Figure 5. The two models represent two typical categories in the structural fingerprint space and are used to examine the distinguishability of the physical anchor, the sensitivity of the cross-check channel, and the effectiveness of attack classification. Model equivalence and verification. The two models share the core, rated capacity (1000 kVA, type S-1000/10; load loss 11.6 kW, no-load loss 1.8 kW), voltage class, and cooling condition; the turn-count difference (1000 vs. 10) is inherent to the two winding

technologies, and equivalence is established at the level of rated quantities, core flux density, and loss. Material properties follow the model database (copper $\sigma = 5.8 \times 10^7$ S/m; CRGO core; transformer oil per the database). Mesh independence was verified with three mesh densities (hot-spot variation below 1%), and the coupling loop iterates until the temperature-dependent resistivity update is below 0.1%. The hot-spot temperatures reported below as "FEM reference" values are finite-element simulation results.

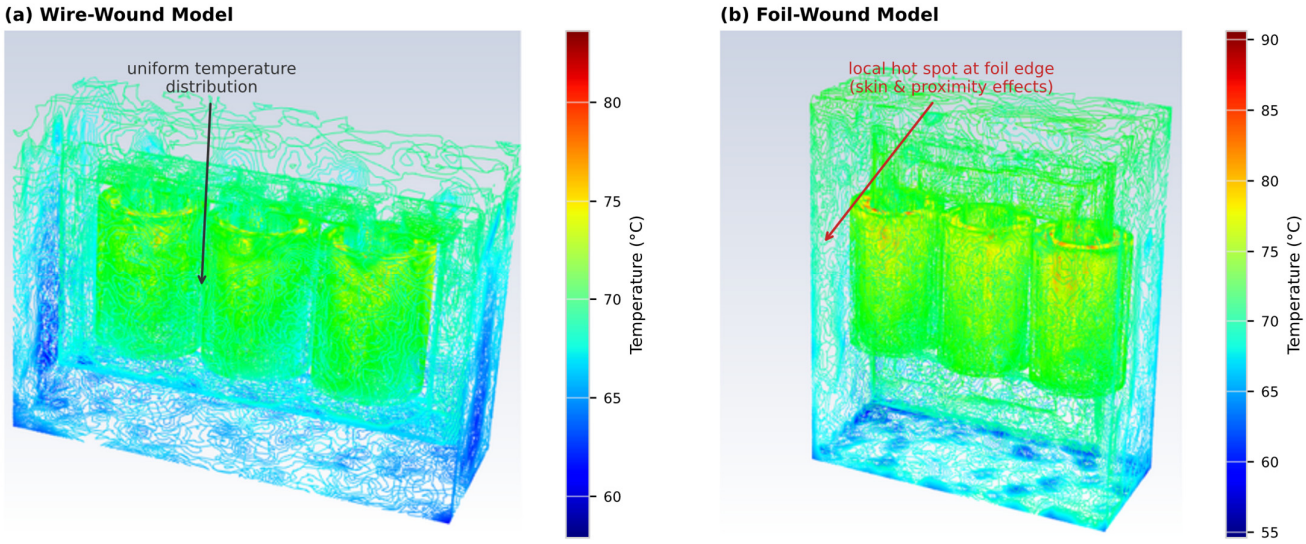


**Figure 5.** Finite-element thermal simulation of the two winding structures at rated load: (**a**) wire-wound model with a uniform temperature distribution; (**b**) foil-wound model with local hot spots at foil edges caused by skin and proximity effects.

### *3.2. Distinguishability of the Structural Fingerprint*

The physical trust anchor works only if the fingerprints of different winding structures are clearly separable. Table 4 lists the simulated equivalent resistances of the two winding types at 50 Hz and 150 Hz. The resistance of the wire-wound model increases by only 5% at the third harmonic, consistent with the weak skin effect of thin conductors, whereas the foil-wound model shows an increase of 25.9%, clearly separating the two structures. The classification threshold $k_{RF} = 1.10$ separates the two classes with ample margin, and the conductor-dimension sweep of the harmonic resistance ratio is shown in Figure 6. The results show that the fingerprint is mathematically separable and, more importantly, that its physical origin guarantees the feature cannot be forged by network attacks: tampering with any temperature monitoring message does not change the $k_{RF}$ value measured at the ports.

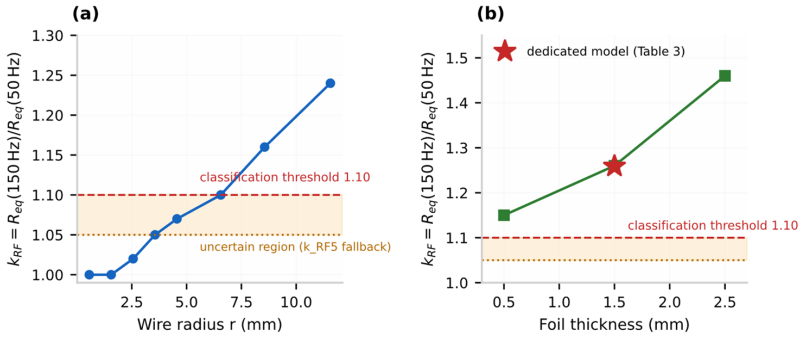


**Figure 6.** Harmonic resistance ratio $k_{RF} = R_{eq}(150\ Hz)/R_{eq}(50\ Hz)$ versus conductor dimensions: (**a**) wire radius; (**b**) foil thickness. The dashed line marks the classification threshold $k_{RF} = 1.10$.

**Table 4.** Equivalent resistances at different frequencies.

| Winding Type | $R_{eq}$(50 Hz) [mΩ] | $R_{eq}$(150 Hz) [mΩ] | $k_{RF}$ |
|---|---|---|---|
| Wire-wound | 7.787 | 8.177 | 1.050 |
| Foil-wound | 7.787 | 9.807 | 1.259 |

To examine the generalizability of the $k_{RF}$ thresholds beyond the two models of Table 4, a parametric sweep over conductor cross-section dimensions is performed in the two-dimensional eddy-field model. For foil windings, the third-harmonic resistance ratio is computed at thicknesses of 0.5, 1.5, and 2.5 mm (height 1400 mm, 10 turns, loaded state); for wire windings, at radii from 0.56 to 11.56 mm. Table 5 summarizes the results. The 1.5 mm foil point ($k_{RF}$ = 1.26) is consistent with the dedicated model of Table 4 (1.259 at full precision). Within the conductor range typical of distribution transformers (foil thickness 0.5–2.5 mm; round wires up to 3.75 mm in diameter), the two families remain separated by a clear margin around the adopted thresholds, and the edge-to-average heat-flux ratio of the sweep grows monotonically with $k_{RF}$. For oversized round conductors (radius above roughly 3 mm), the skin effect strengthens, and the wire family approaches the uncertain region, where the higher-harmonic ratio $k_{RF5}$ is used as the fallback. Classification is performed under normal load; on-site validation across a larger transformer population remains future work.

**Table 5.** $k_{RF}$ of the conductor parametric sweep (loaded state; foil: thickness at 1400 mm height, 10 turns; wire: conductor radius).

| Foil Thickness [mm] | $k_{RF}$ | Wire Radius [mm] | $k_{RF}$ |
|---|---|---|---|
| 0.5 | 1.15 | 0.56 | 1.00 |
| 1.5 | 1.26 | 1.56 | 1.00 |
| 2.5 | 1.46 | 2.56 | 1.02 |
| | | 3.56 | 1.05 |

### *3.3. Structure-Temperature Coupling Strength and Baseline Calibration*

The physical basis for screening false temperature data is the strength of the structure-temperature coupling: if the temperature-response difference caused by the true structure were comparable to the injected deviation, reconciliation could not distinguish the two. Table 6 quantifies this coupling strength and calibrates the reconciliation baseline. For the foil-wound transformer, the temperature computed under the wrong structural assumption (uniform current distribution) deviates from the true hot spot by 9.6%; this magnitude constitutes the physical margin within which injected temperature data "betray" themselves. For the wire-wound transformer, the current distribution is inherently uniform, so a structural misjudgment introduces a deviation of only about 0.4%, showing that the discrimination capability of the cross-check mechanism adapts to the structural difference. On this basis, the structural correction model compresses the baseline error of the foil-wound case to 0.26%, reducing the model-error share of the residual to a few percent, so that an exceeded residual indicates injected temperature data rather than model inaccuracy. For wire-wound transformers, the correction factor $k_{RF}$ is close to one, and the model automatically degenerates to the conventional estimation, so framework compatibility is unaffected.

**Table 6.** Hot-spot temperature estimation accuracy and reconciliation-baseline calibration.

| Method | Winding Type | Computed $T_{max}$ [°C] | FEM Reference $T_{max}$ [°C] | Error [%] |
|---|---|---|---|---|
| Conventional (uniform assumption) | Foil-wound | 59.61 | 65.96 | 9.6 |
| Proposed (corrected) | Foil-wound | 65.79 | 65.96 | 0.26 |
| Conventional (uniform assumption) | Wire-wound | 68.2 | 68.5 | 0.4 |
| Proposed (corrected) | Wire-wound | 68.2 | 68.5 | 0.4 |

### *3.4. Attack-Scenario Implementation and Detection Performance*

#### 3.4.1. Attack Scenarios and Injection Implementation

To systematically evaluate the detection capability, seven attack scenarios are implemented on the emulated SCADA temperature monitoring link and on the terminal-local database, each repeated for 100 batches. The attack vectors are constructed as follows:

(1) Input-type temperature FDI (scenario S1): the attacker intercepts, at a man-in-the-middle position on the Modbus TCP link, the top-oil-temperature message (register address 0x000C) sent by the terminal to the master station, adds 10 °C to the register value, and forwards it; the hot-spot temperature computed by the master station is raised accordingly. Variants S1a/S1b set the injection amplitude to +5 °C and +15 °C, respectively, for the amplitude sweep.

(2) Stealthy ramp injection (scenario S2): a latent attack that evades threshold detection is emulated by raising the oil-temperature reading by 0.5 °C every 15 min (one block period), keeping each single-period injection below empty $\tau_P$ so that the deviation accumulates slowly under the cover of noise; the test records the number of periods required to cross $\tau_P$ and the total detection time.

(3) Output-type FDI (scenario S3): the attacker impersonates the master station and issues a remote write command (function code 0x10) to the terminal, lowering the hot-spot alarm threshold from 85 °C to 70 °C so as to induce false tripping. The attack is detected by the command-consistency indicator $\varepsilon_C$: the unauthorized configuration write is rejected at receipt, and whenever the lowered threshold issues an alarm or trip action while the physically computed temperature (66–68.5 °C at rated load, Table 6) is below the legitimate threshold, $\varepsilon_C = 1$ is asserted. At rated load, no trip occurs (66 < 70); the false-tripping risk materializes at peak load, when the true hot spot approaches 70–85 °C, and the forged threshold converts a normal state into a trip—which is precisely the condition $\varepsilon_C$ detects.

(4) Model-parameter tampering (scenario S4): the attacker penetrates the terminal-local database, modifies the foil-winding correction coefficient $k_{RF}$ from 1.259 to 1.050, and simultaneously forges temperature data consistent with the corrupted model, attempting to make the false data and the damaged model corroborate each other; this scenario directly tests whether the on-chain baseline can resist "model-data collusion" attacks.

(5) Controlled experiment (scenario S5): the injected temperature data are made exactly equal to the values computed under the wrong structural assumption (wire-wound model), so that the false data remain fully self-consistent under statistical residual tests; the scenario verifies that conventional cyber-layer bad-data detection fails while the proposed physical channel still detects the injection.

(6) Negative injection (scenario S6): the attacker lowers the reported top-oil temperature by 10 °C to mask real overheating—the physically symmetric counterpart of S1; the test verifies that the screening margin works in the safety-critical direction.

(7) Sub-threshold constant offset (scenario S7): a constant +2.5 °C offset, below $\tau_P$, is injected continuously to remain invisible to the single-shot residual check; detection relies on the periodic re-estimation of the baseline residual statistics, in which the persistent offset shifts the mean of $\varepsilon_P$ and is flagged by the trend monitor, quantifying the stealth–latency trade-off.

#### 3.4.2. Detection and Classification Results

Statistical protocol. Each attack scenario is repeated for 100 batches, randomizing (i) the measurement noise seed, (ii) the injection start time (±15 min), and (iii) the injection amplitude (±10% around the nominal value); normal operation is tested over 1000 additional batches with the load level randomly varied within the rated range. Table 7 reports the mean response time with the 95% Clopper–Pearson confidence interval of the detection rate; Table 8 gives the classification confusion matrix over the four operating conditions. The conventional baseline used for comparison in S5 is a $\chi^2$ residual bad-data test with the threshold set for a 1% false-alarm probability.

**Table 7.** Detection results of temperature FDI attacks (100 randomized batches per scenario; normal operation: 1000 batches).

| Scenario | Indicators ($\varepsilon_S/\varepsilon_P/\varepsilon_C$) | Detection Rate [95% CI] [%] | Mean Response |
|---|---|---|---|
| Normal | 0.01/0.2 °C/0 | — (3 false alarms, 0.3%) * | — |
| S1 input +10 °C | 0.01/10.3 °C/0 | 100 [96.4, 100] | 2.7 s |
| S1a +5 °C | 0.01/5.1 °C/0 | 100 [96.4, 100] | 2.7 s |
| S1b +15 °C | 0.01/14.8 °C/0 | 100 [96.4, 100] | 2.8 s |
| S2 ramp 0.5 °C/period | 0.01/3.5 °C (7th)/0 | 100 [96.4, 100] | ≈105 min |
| S3 output (85 → 70 °C) | 0.01/0.3 °C/1 | 100 [96.4, 100] | 1.8 s |
| S4 parameter tampering | 0.17/6.5 °C/0 | 100 [96.4, 100] | 3.5 s |
| S5 structure-consistent | 0.01/6.2 °C/0 | 100 [96.4, 100] | 2.6 s |
| S6 negative −10 °C | 0.01/10.2 °C/0 | 100 [96.4, 100] | 2.2 s |
| S7 constant +2.5 °C | 0.01/2.5 °C/0 | 100 [96.4, 100] | 4.1 ± 0.6 h (trend) |

* The normal-operation row reports false alarms rather than detections: 3 false alarms in 1000 normal batches (0.3%).

**Table 8.** Classification confusion matrix over the four operating conditions.

| Actual\\Verdict | Normal | Input | Output | Tampering |
|---|---|---|---|---|
| Normal | 997 | 3 | 0 | 0 |
| Input-type | 0 | 100 | 0 | 0 |
| Output-type | 0 | 0 | 100 | 0 |
| Parameter tampering | 0 | 0 | 0 | 100 |

S4 activates both the structural alarm and the on-chain integrity-check failure; the recomputed fingerprint distance reaches $\varepsilon_S = 0.17$ ($= |1.050 - 1.259| / 1.259$), far above $\tau_S = 0.05$, and the locally tampered coefficient cannot influence the verdict. The S5 control experiment confirms that when the injected temperature data are self-consistent with the wire-wound assumption, the $\chi^2$ detector misses all 100 batches, whereas the proposed method—anchored to the true foil-wound structure—exhibits the 6.2 °C gap caused by the 9.6% structural error and detects all 100 batches, demonstrating that the authenticity of temperature monitoring data must be arbitrated by physical law rather than by data self-consistency.

### *3.5. Blockchain Test-Network Performance and Security Experiments*

#### 3.5.1. Implementation and Setup

The test network is implemented in Python 3.12 (about 480 lines) following the PBFT protocol of Castro and Liskov [20]: three-phase voting, view change, and vote retransmission are retained; the smart-contract layer is removed and signature verification is batched, and the permissioned node architecture follows enterprise-grade consortium-chain implementations [22]. One notarization record (~180 bytes) is submitted per consensus round, and the PRE-PREPARE message carries the record itself; all messages are signed with real ECDSA over secp256k1, whose verification (0.57–1.55 ms per operation on the test host) is the dominant per-message cost. Seven node processes run on an isolated test host (Intel Xeon Platinum, 2 cores available to the container, loopback networking); two nodes exhibit Byzantine behavior (randomly delayed votes, 20% conflicting prepare messages). The client keeps up to eight rounds in flight. The code will be released in a public repository upon acceptance.

#### 3.5.2. Performance

The measured performance is listed in Table 9. With seven nodes (including two Byzantine), the mean finality latency is 899.1 ms (P95 = 967.4 ms) per notarization round, and the sustained notarization rate is 8.9 rounds/s; after the two Byzantine nodes are removed (phase B: 500 rounds), the latency drops to 419.0 ms and the rate rises to 19.0 rounds/s—two orders of magnitude above the aggregate field demand of about 0.1 notarizations/s (a hundred terminals at a 15-min period). Node resource consumption during phase A averages 21.0% CPU (39.4% peak) with a 44.0 MB memory peak and 306 KB/s loopback traffic. Figure 7 summarizes the per-phase latency and notarization rate. A resource-limited terminal runs a verify-only light client: at the field rate, even a 10× slower ARM terminal spends a negligible fraction of one core on verification.

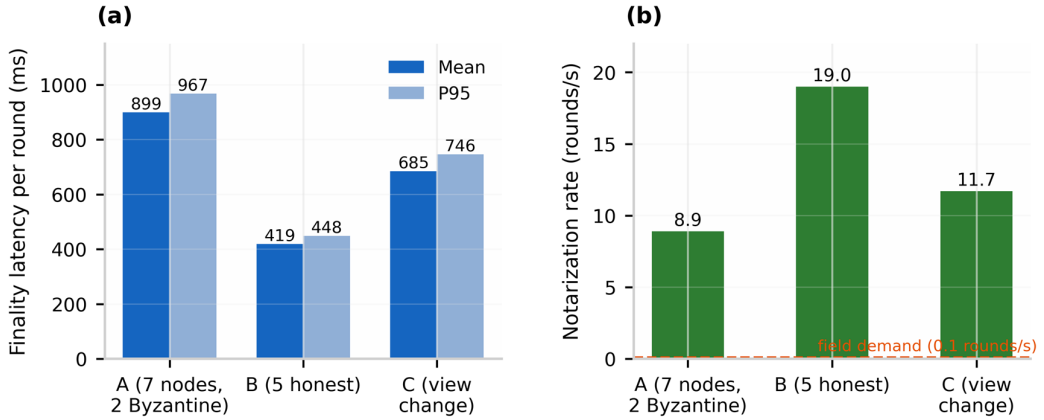


**Figure 7.** Measured performance per phase: (**a**) finality latency per notarization round (mean and P95); (**b**) notarization rate. Phase A: seven nodes with two Byzantine nodes; phase B: after the two Byzantine nodes are killed; phase C: after the primary is killed (view change).

**Table 9.** Measured performance of the consortium chain.

| Metric | Measured Value | Remark |
|---|---|---|
| Finality latency, phase A (7 nodes, 2 Byzantine) | 899.1 ms mean/967.4 ms P95 | per notarization round |
| Finality latency, phase B (5 honest) | 419.0 ms mean/448.5 ms P95 | per notarization round |
| Notarization rate, phase A/B | 8.9/19.0 rounds/s | vs. field demand 0.1 rounds/s |

**Table 9.** *Cont.*

| Metric | Measured Value | Remark |
|---|---|---|
| View-change duration (primary killed) | 0.47 s | kill to NEW-VIEW |
| Failover recovery (kill to first confirmation) | 0.54 s | phase C |
| Signature verification cost | 0.573–1.545 ms/op | secp256k1 |
| Node CPU/memory peak (phase A) | 21.0% avg/44.0 MB | per node process |
| Loopback traffic (phase A) | 306 KB/s | aggregate |
| Byzantine tolerance | f = 2 | 7-node configuration |
| Ledger growth per terminal | ≈6.3 MB/year | 15-min period, ≈180 B per record |

#### 3.5.3. Fault Injection and Security Tests

Closing the two Byzantine nodes (phase B) leaves five honest nodes, which continue to confirm every round with no loss. In phase C, the Byzantine nodes are restarted, view change is enabled, and the primary is killed mid-phase: the view change completes in 0.47 s and the first confirmation after the kill arrives 0.54 s later, after which all remaining rounds confirm normally. Security tests target three attack paths: (a) an attacker rewrites the local baseline copy and correction coefficients—the integrity check against the on-chain digest fails immediately (S4, response 3.5 s); (b) a confirmed notarization transaction is resent verbatim and is rejected because of the expired sequence number; (c) altered content carrying the original signature fails ECDSA verification and is dropped at the pre-prepare stage. All three paths are blocked.

### *3.6. Noise Robustness and Threshold Sensitivity*

The noise robustness of the physical channel determines its field applicability. Gaussian noise is superimposed on the port voltage and current measurements, with the noise level defined precisely as the ratio of the added noise standard deviation to the RMS amplitude of the fundamental component. As shown in Table 10, the structural classification remains correct up to the 20% level. The physical residual grows with the noise level: it stays within 2.1 °C at 5% (below $\tau_P$ = 3 °C), reaches 3.8 °C at 10%—marginally above the threshold—and 5.7 °C at 20%. The framework therefore requires instrument transformers of class 0.5 or better (ratio error $\leq$ 0.5%), whose effective noise is far below the 5% level; under this operating condition, a joint sensitivity analysis of $\tau_S$ (0.02–0.08) and $\tau_P$ (2–5 °C) shows that all attack scenarios are correctly classified with a false-alarm rate no higher than 0.5%, indicating that the detection performance is insensitive to the threshold selection.

**Table 10.** Measurement noise robustness.

| Noise Level [%] | $k_{RF}$ (Foil-Wound) | Classification | $\varepsilon_P$ [°C] |
|---|---|---|---|
| ∞ (ideal) | 1.259 | Foil-wound | 0.3 |
| 5 | 1.247 | Foil-wound | 2.1 |
| 10 | 1.232 | Foil-wound | 3.8 |
| 20 | 1.198 | Foil-wound | 5.7 |

### *3.7. Ablation and Baseline Comparison*

Table 11 reports an ablation over three structural variants of the framework on the same attack batches, together with two baselines: a conventional $\chi^2$ residual bad-data detector (threshold set for 1% false-alarm probability) and a data-driven detector (a one-dimensional convolutional autoencoder trained on normal-operation temperature sequences, decision by reconstruction error). The autoencoder represents the learning-based detection direction, including hybrid CNN–ViT architectures recently proposed for energy-site anomaly detection [23].

**Table 11.** Ablation and baseline comparison (detection rate, %).

| Method | S1 (+10 °C) | S4 | S5 | S6 (−10 °C) |
|---|---|---|---|---|
| Full method | 100 | 100 | 100 | 100 |
| Without blockchain (local baseline) | 100 | 0 | 100 | 100 |
| Without structural correction | 100 | 100 | 100 | 100 |
| —same, at +4/+5/+6 °C * | 40/75/92 | — | — | — |
| Temperature residual only (no fingerprint) | 100 | 0 | 0 | 100 |
| $\chi^2$ residual bad-data detection | 100 | 0 | 0 | 100 |
| Autoencoder (data-driven) | 100 | 0 | 0 | 100 |

* The normal-operation row reports false alarms rather than detections: 3 false alarms in 1000 normal batches (0.3%).

Two conclusions follow. First, each component carries its weight: removing the on-chain baseline defeats S4; removing the structural correction masks small injections (the 9.6% structural error, ≈6.3 °C on this model, swallows a +4 °C injection); removing the fingerprint defeats S5. Second, the physical anchor is irreplaceable by data-driven detection: the autoencoder—like the $\chi^2$ test—accepts structure-consistent false data (S5) and output-type attacks that leave the measurements unchanged, because both are consistent with the training distribution.

## 4. Discussion

### *4.1. Comparison with Existing Methods*

The fundamental difference between the proposed method and existing FDI detection paradigms lies in the trust root. Pure machine-learning detectors [5,6] root trust in the statistical distribution of historical data, which the attacker can probe and mimic; the proposed method roots trust in the physical structure of the equipment, which temperature data injection cannot alter, so there is no adversarial weakness of "failure once the training distribution is figured out", and no historical training data are required, making the method applicable to newly commissioned transformers. Compared with conventional blockchain-IoT audit schemes [11,12], the protection target differs: conventional schemes use the blockchain to fossilize data records, risking a "garbage-in, garbage-out" baseline pollution, whereas this work notarizes only the physical-fingerprint digest on-chain while the temperature data themselves are still reconciled physically off-chain; the role of the blockchain shifts from "data safe" to "notary of the physical anchor". The measured 8.9–19.0 notarizations/s—two orders of magnitude above the aggregate field demand—indicates that the role change from "data safe" to "notary of the physical anchor" incurs no substantial deployment cost.

The structure-related temperature correction model serves security rather than general thermometry: its purpose is to compress the reconciliation-baseline error to one-tenth of the injected temperature deviation, making the physical residual a sensitive indicator of data authenticity. This "accuracy in the service of security" positioning also explains why the framework remains effective for wire-wound transformers: when the true structure is wire-wound, the correction factor is approximately one, and the physical channel still provides a reconciliation baseline independent of SCADA.

### *4.2. Limitations*

Several limitations deserve attention. First, the physical channel depends on the quality of port harmonic measurements: under light-load or outage conditions, the harmonic content is insufficient, the SNR drops, and the reconciliation sensitivity degrades accordingly; adaptive thresholds coordinated with load conditions are needed, which is the inherent cost of physical verification relative to pure data methods. Second, the detection time of a stealthy ramp attack (S2) depends on the attacker's cumulative injection rate; when the single-period injection is extremely small, and the duration spans the boundary of normal aging and seasonal temperature variation, trend analysis of the temperature rise must be combined to shorten the detection time. Third, the current fingerprint space covers only the winding structure; if an attacker combines physical winding replacement with temperature data injection, the composite attack may exceed the defense boundary of FDI; future work will explore extending the fingerprint to the core structure through flux-transfer characteristics, increasing the physical forgery cost. Fourth, PBFT consensus relies on the honesty of consortium nodes, and consortium governance and key management require careful mechanisms to guard against collusion of more than f nodes. Fifth, experimental validation on a physical transformer or a hardware-in-the-loop platform is required before field claims; this is the first item of future work. Sixth, the stealth–latency trade-off is fundamental: a sub-threshold constant offset (S7) is detected only by trend monitoring (about 4 h), and the 105-min figure of S2 suits monitoring-data validation but not fast protection; the framework complements—rather than replaces—primary protection.

### *4.3. Practical Deployment Considerations*

For deployment in existing distribution networks, the proposed framework can be integrated into smart terminals through firmware updates, reusing their existing harmonic-analysis capability without additional sensors; the blockchain layer can be implemented as a private consortium chain among utility operators, and according to the measured storage and computation overheads, ordinary edge servers are sufficient. The hash-only on-chain strategy keeps sensitive transformer parameters locally at the terminal, addressing data-privacy concerns. The framework is compatible with both winding types, so deployment does not require distinguishing the structural types of the existing transformer population, meeting plug-and-play requirements at the terminal-integration level (no internal geometric parameters and no additional sensors beyond port measurements and a terminal-mounted ambient sensor).

## 5. Conclusions

Aiming at false data injection in distribution-transformer temperature monitoring, this paper proposed a verification framework combining winding structural fingerprinting with lightweight blockchain notarization, under an explicit threat model. The winding fingerprint serves as a physical trust reference that data injection cannot forge without physical access; the structure-determined temperature behavior builds a physical reconciliation channel fed entirely by local quantities; and a PBFT consortium chain anchors the verification baseline across mutually distrustful domains. Three consistency indicators—structural, physical, and command—place input-, output-, and model-parameter-tampering attacks in distinguishable regions.

Finite-element simulations and randomized 100-round attack tests show a 9.6% structural temperature deviation providing the physical screening margin, a corrected baseline error of 0.26%, 100% detection (95% CI [96.4, 100]) of the seven attack variants with a 0.3% false-alarm rate, and response times within 4 s for instantaneous attacks. Ablation confirms that each component—on-chain baseline, structural correction, fingerprint—is necessary,

and that data-driven detectors miss the structure-consistent and output-type attacks by construction. On a real seven-node implementation with ECDSA-secp256k1 signatures, finality per notarization is 0.42–0.90 s, the notarization rate is 8.9–19.0 rounds/s, a view change completes in 0.47 s, and failover recovery takes 0.54 s—two orders of magnitude above the field demand. All findings are stated within the scope of the tested scenarios; hardware-in-the-loop validation, core-geometry fingerprint extension, and federated cross-terminal detection are the next steps.

**Author Contributions:** Conceptualization, Y.G. and S.M.; methodology, Y.G. and X.Y.; software, J.Y. and B.Z.; validation, Y.G., S.M., W.M. and B.Z.; formal analysis, Y.G.; investigation, J.Y. and B.Z.; resources, S.M. and X.Y.; data curation, W.M.; writing—original draft preparation, Y.G.; writing—review and editing, Y.G. and S.M.; visualization, J.Y.; supervision, Y.G. and X.Y.; project administration, S.M. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by the Liaoning Provincial Department of Education through the 2025 Basic Scientific Research Project of Higher Education Institutions (grant number LJ242511632022).

**Data Availability Statement:** The blockchain test code and the experiment configuration will be released in a public repository upon acceptance; the finite-element models and the attack-test data are available from the corresponding author on request.

**Acknowledgments:** The authors thank the research team of Shenyang Institute of Engineering for their support.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## Abbreviations

| Abbreviation | Full Form |
|---|---|
| FDI | False Data Injection |
| SCADA | Supervisory Control and Data Acquisition |
| PBFT | Practical Byzantine Fault Tolerance |
| ECDSA | Elliptic Curve Digital Signature Algorithm |
| THD | Total Harmonic Distortion |
| FFT | Fast Fourier Transform |
| SNR | Signal-to-Noise Ratio |
| DTU | Distribution Terminal Unit |
| TTU | Transformer Terminal Unit |
| TPS | Transactions Per Second |
| IEC | International Electrotechnical Commission |

## References

1. Pawar, R.R.; Deosarkar, S.B. Health condition monitoring system for a distribution transformer using Internet of Things (IoT). In *Proceedings of the 2017 International Conference on Computing Methodologies and Communication (ICCMC), Erode, India, 18–19 July 2017*; IEEE: Piscataway, NJ, USA, 2017; pp. 117–122.
2. Gezegin, C.; Ozgonenel, O.; Dirik, H. A monitoring method for average winding and hot-spot temperatures of single-phase, oil-immersed transformers. *IEEE Trans. Power Deliv.* **2021**, *36*, 3196–3203. [CrossRef]
3. Mo, Y.; Kim, T.H.-J.; Brancik, K.; Dickinson, D.; Lee, H.; Perrig, A.; Sinopoli, B. Cyber–physical security of a smart grid infrastructure. *Proc. IEEE* **2012**, *100*, 195–209. [CrossRef]
4. Liu, Y.; Ning, P.; Reiter, M.K. False data injection attacks against state estimation in electric power grids. *ACM Trans. Inf. Syst. Secur.* **2011**, *14*, 1–33. [CrossRef]
5. He, Y.; Mendis, G.J.; Wei, J. Real-time detection of false data injection attacks in smart grid: A deep learning-based intelligent mechanism. *IEEE Trans. Smart Grid* **2017**, *8*, 2505–2516. [CrossRef]
6. Boyaci, O.; Narimani, M.R.; Davis, K.R.; Ismail, M.; Overbye, T.J.; Serpedin, E. Joint detection and localization of stealth false data injection attacks in smart grids using graph neural networks. *IEEE Trans. Smart Grid* **2022**, *13*, 807–819. [CrossRef]
7. Dowell, P.L. Effects of eddy currents in transformer windings. *Proc. Inst. Electr. Eng.* **1966**, *113*, 1387–1394. [CrossRef]

8. Skillen, A.; Revell, A.; Iacovides, H.; Wu, W. Numerical prediction of local hot-spot phenomena in transformer windings. *Appl. Therm. Eng.* **2012**, *36*, 96–105. [CrossRef]
9. Kim, J.; Tong, L.; Thomas, R.J. Subspace methods for data attack on state estimation: A data-driven approach. *IEEE Trans. Signal Process.* **2015**, *63*, 1102–1114. [CrossRef]
10. Kim, T.T.; Poor, H.V. Strategic protection against data injection attacks on power grids. *IEEE Trans. Smart Grid* **2011**, *2*, 326–333. [CrossRef]
11. Zhuang, P.; Zamir, T.; Liang, H. Blockchain for cybersecurity in smart grid: A comprehensive survey. *IEEE Trans. Ind. Inform.* **2021**, *17*, 3–19. [CrossRef]
12. Fernández-Caramés, T.M.; Fraga-Lamas, P. A review on the use of blockchain for the Internet of Things. *IEEE Access* **2018**, *6*, 32979–33001. [CrossRef]
13. *IEC 60076-18:2012*; Power Transformers—Part 18: Measurement of Frequency Response. IEC: Geneva, Switzerland, 2012.
14. *IEC 60076-7:2018*; Power Transformers—Part 7: Loading Guide for Mineral-Oil-Immersed Power Transformers. IEC: Geneva, Switzerland, 2018.
15. Swift, G.; Molinski, T.S.; Lehn, W. A fundamental approach to transformer thermal modeling. I. Theory and equivalent circuit. *IEEE Trans. Power Deliv.* **2001**, *16*, 171–175. [CrossRef]
16. Susa, D.; Lehtonen, M.; Nordman, H. Dynamic thermal modelling of power transformers. *IEEE Trans. Power Deliv.* **2005**, *20*, 197–204. [CrossRef]
17. Torriano, F.; Chaaban, M.; Picher, P. Numerical study of parameters affecting the temperature distribution in a disc-type transformer winding. *Appl. Therm. Eng.* **2010**, *30*, 2034–2044. [CrossRef]
18. Susa, D.; Lehtonen, M.; Nordman, H. Dynamic thermal modeling of distribution transformers. *IEEE Trans. Power Deliv.* **2005**, *20*, 1919–1929. [CrossRef]
19. Cheng, D.K. *Field and Wave Electromagnetics*, 2nd ed.; Addison-Wesley: Reading, MA, USA, 1989.
20. Castro, M.; Liskov, B. Practical Byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI), New Orleans, LA, USA, 22–24 February 1999*; USENIX Association: Berkeley, CA, USA, 1999; pp. 173–186.
21. Dorri, A.; Kanhere, S.S.; Jurdak, R. Towards an optimized blockchain for IoT. In *Proceedings of the 2017 IEEE/ACM Second International Conference on Internet-of-Things Design and Implementation (IoTDI), Pittsburgh, PA, USA, 18–21 April 2017*; ACM: New York, NY, USA, 2017; pp. 173–178.
22. Androulaki, E.; Barger, A.; Bortnikov, V.; Cachin, C.; Christidis, K.; De Caro, A.; Enyeart, D.; Ferris, C.; Laventman, G.; Manevich, Y.; et al. Hyperledger Fabric: A distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference (EuroSys), Porto, Portugal, 23–26 April 2018*; ACM: New York, NY, USA, 2018; pp. 1–15.
23. Darban, A.A. Anomaly detection in solar farms with a hybrid CNN-ViT model to enhance accuracy and efficiency. *Int. J. Ambient Energy* **2025**, *46*, 2541270. [CrossRef]


**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.