



Spectral Watermarking Techniques for Closed-Loop Resilience in Electrified Multi-Echelon Logistics

Padma B¹

¹Research Scholar

How to Cite this Article:

B, P. (2026). Spectral Watermarking Techniques for Closed-Loop Resilience in Electrified Multi-Echelon Logistics. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(9), 1-9.
<https://doi.org/10.55041/ijcope.v2i9.170>

License:

This article is published under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

© The Author(s). Published by International Journal of Creative and Open Research in Engineering and Management.



<https://doi.org/10.55041/ijcope.v2i9.170>

Abstract—

This paper develops spectral watermarking as the integrity observer of a closed-loop architecture for multi-echelon resilience under electrified logistics. The plant remains stock coupled to pack energy. The state observer remains a digital twin plus a shift-invariant FKF channel. The new object is a mark written into reserved FKF coefficients of lead-time, order, and current series so that a later detector can ask whether the continuous trajectory still matches the trajectory that was committed. A ledger still answers who wrote the event. The watermark answers a different question. Together they close a class of spoofs that either channel misses alone.

Embedding uses exponential modulation and structured displacement in the (ω, τ) plane already used to watch bullwhip and electrical resonance. Detection is a windowed agreement statistic with an explicit sequential test for window length. The mark is semi-fragile by design: lawful delay appears as a translation; spliced dwell and replayed charger traces do not. The same scalar health index that prices spectral ringing in the green stage cost also reports watermark integrity into the horizon objective. Urban corridors, solar-assisted packs, discrete packing, and human override stay inside the loop. No new transform is claimed. The contribution is to write watermarking as a first-class observer output rather than as a media afterthought.

Keywords: spectral watermark; FKF modulation; digital twin; ledger integrity; electrified logistics; closed-loop resilience

I. INTRODUCTION

The operator who sits inside an Industry 5.0 loop already sees exceptions as first-class work. Spectral watermarking gives that operator a second light on the same console: not only whether a recommended slot is labor-legal, but whether the series that produced the recommendation still carries the mark that was written when the series was committed. That is a smaller claim than a new security product and a larger claim than a transform identity.

Complexity language still helps. Internal complexity grows with the dimension of the estimated state and with the number of roles allowed to override. External complexity grows with the measure of the disruption ensemble. A watermark that cannot be computed on the same clock as the twin increases internal complexity without lowering recovery time. The design constraint is therefore computational as well as spectral: Λ must be cheap beside (8).

Resilience writing describes absorption, adaptation, and recovery under internal and external complexity [26], [27]. Digital-transformation surveys add sensors, platforms, and visibility programs [13]. What those reviews rarely write down is a plant, an observer, and an actuator, and they almost never write down a watermark. Without those objects, integrity remains a checkbox on a dashboard.

The plant of this paper is a multi-echelon inventory-and-fleet system whose vehicles obey energy inequalities from plug-in hybrid sizing and later battery reviews [5], [10]. The observer is a digital twin [1] augmented by an FKF spectral channel so that periodic bullwhip and electrical resonance are visible [16], [17]. The actuators are a green AI policy [2], a quantum or hybrid discrete solver for unit-load and disruption assignment [8], [11], and a human override required by Industry 5.0 [24]. Measurements are attested on a ledger when identity matters [3].



Smart mobility work [9] and the IoT–AI–quantum stack [23] define the urban and sensing envelope. Power-line carrier telemetry [4] is kept as a reminder that the measurement channel need not be cellular. The swipe controller [6] is kept as a reminder that actuation includes a finger on a screen. Opportunity framing for quantum logistics [7] locates the discrete solver inside a live planning loop.

The gap is specific. A ledger can certify that a dwell value was committed by a named party. It cannot certify that the continuous lead-time series still has the same spectral shape as the series that was committed. Exponential modulation of FKF coefficients was already proposed as a watermark for digital supply-chain networks [22]. This paper takes that property out of a transform note and writes it as an observer block: host selection, embedding rule, detection statistic, threat model, and coupling to the closed-loop objective.

Host series L , i, q	FKF / FKL plane	Reserved bands
Exponential mark [22]	Twin + ledger [1], [3]	Detector Λ

Figure 1. Influence path from host series through the spectral mark to the detector.

II. Related Techniques and Why a Logistics Host Is Different

Capacity, invisibility, and robustness still trade off. On a lead-time host the capacity need is small: a few bits per window to bind a key and a commitment hash. Invisibility is process invisibility, not perceptual invisibility. Robustness is robustness to lawful delay and short outages, not to MP3. Those three substitutions are the entire methodological shift from media watermarking to logistics watermarking.

Hybrid media pipelines that cascade a wavelet with a cosine transform are instructive only as a reminder that two bases can share a host. The pair in this paper is FKF for analysis and FKL for inversion, already present in the cited transform series. No third basis is introduced.

A. Classical spectral marks

Image and audio watermarking already treat the host in a transform. Discrete cosine marks ride mid-frequency AC coefficients because codecs keep them. Discrete Fourier marks split magnitude and phase. Wavelet marks add locality. Spread-spectrum embedding adds a scaled payload to selected coefficients. Quantization-index modulation snaps coefficients onto a payload-dependent lattice. Patchwork forces two coefficient

subsets apart. These families are the media baseline this paper does not reinvent.

Time–frequency marks embed only in selected patches of a joint plane. Fractional and polar-complex transforms add an order parameter. Zero-watermarking registers a feature instead of writing into the host. Fragile marks are meant to die under edit; robust marks are meant to live under compression. Semi-fragile marks survive an allowed class of processing and break on a hostile class. The logistics mark of this paper is semi-fragile on purpose.

B. What a supply-chain host changes

A lead-time series is not a song. Lawful operations include delay, short outages, solar-driven changes in pack current [12], and corridor restrictions from intelligent-transport policy [9]. Unlawful operations include splicing a forged dwell interval, replaying a clean charger trace over a ringing dock, and rebuilding a twin on a stolen Hamiltonian [8]. The allowed class is exactly the class the FKF time-shift property already names [21]. The hostile class is everything that creates a new resonant mode or a new residue while the ledger still looks consistent [16], [19].

Battery chemistry and BMS limits change feasible traction power [10], [28]. They must not be allowed to look like watermark attacks, and watermark embedding must not be allowed to look like a thermal-margin violation. That coupling is absent from media watermarking and present here because the plant is electrified.

III. Host, Transform, and Reserved Bands

Three hosts are not three papers. Inventory residuals, lead times, and pack currents share a clock and a twin. A mark on only one host leaves the other two as unmarked side channels. The architecture therefore treats embedding as a vector operation across the triple (L_k, i_v, q_k) with a shared key and distinct reserved sets Ω_m . Cross-host disagreement is itself a detector.

Forbidden bands must be maintained as the network changes. A new weekly cycle at echelon three can migrate energy toward Ω_m . The twin is responsible for announcing that migration and for rotating the reserved set, just as it rebuilds a Hamiltonian after a major event. A static mask is a stale observer.

A. Material and energy hosts

Let $q_{k,t}$ be on-hand stock at echelon k and $o_{k,t}$ the order placed upstream. A standard balance is

$$q_{k,t+1} = q_{k,t} + o_{k,t-L} - d_{k,t}(I)$$

Lead time L_k is not a constant. Shift-invariant FKF analysis treats $L_k(t)$ as the signal of interest [17].



Demand may be physical consumption or a shipment onto a vehicle whose usable pack energy evolves as

$$E_{v,t+1} = E_{v,t} + \eta_c P_{\text{chg},t} \Delta t + \eta_s P_{\text{sol},t} \Delta t - \frac{P_{\text{trac},t} \Delta t}{\eta_d} \quad (2)$$

Sizing papers give the integral form of the same balance [5]. Solar-assisted architectures add P_{sol} as a designed source [12], [25]. Materials and BMS reviews constrain E and P by temperature and state of health [10], [28]. A shipment in (1) is infeasible if every vehicle that could carry it violates the floor implied by (2). Current traces $i_v(t)$ drawn from (2) are the second watermark host. Inventory residuals from (1) are the third.

Energy floors remain constraints; embedding must not violate them.

$$E_v(t) \geq E_{\text{floor}}(P_{\text{trac}}, \text{SOH}, T) \quad (19)$$

B. Working integral

The FKF transform of a host $x(t) \in \{L_k, i_v, q_k\}$ is the working integral of the spectral framework for chain dynamics [18]:

$$X(\omega, \tau) = \int_w x(t) K(\omega, \tau; t) dt \quad (3)$$

Distributional extension admits stockout impulses [14]. Sequential convergence in the testing-function space tells the implementer when the window is long enough to trust a mark [15]. FKL properties give an alternative inversion route [20]. Differential identities describe how residues move when a buffer changes [19]. Residue extraction is shared with hybrid-vehicle electrical traces [16].

C. Quiet bands and forbidden bands

A scalar observer already used as a constraint is

$$\rho(X) = \frac{\int w(\omega, \tau) |X(\omega, \tau)|^2 d\omega d\tau}{\int |X(\omega, \tau)|^2 d\omega d\tau + \varepsilon} \quad (4)$$

Weight w down-weights known bullwhip and electrical-resonance bands [16], [17]. Those bands are forbidden for embedding. The watermark lives in a reserved quiet set Ω_m adjacent to, but not inside, the bands that ρ already prices. Embedding must not move ρ enough to look like a new mode. That is the logistics analogue of a perceptual mask.

The mark is admissible only if it does not create a new mode in ρ .

$$\Delta \rho_m = \rho(X^m) - \rho(X), \quad |\Delta \rho_m| \leq \varepsilon_\rho \quad (15)$$

IV. Embedding by Exponential Modulation

Gain α is chosen so that the marked series does not move fill-rate or thermal-margin reports outside their ordinary noise. That is an operational invisibility test, and it should be published beside PSNR when a toy twin is released. If α must be raised until Λ is stable, the reserved set is too small or the window is too short.

The energy written into reserved bands is kept small so that process metrics do not move.

$$E_m = \int_{\Omega_m} |\alpha M^2(\omega, \tau) d\omega d\tau| \quad (14)$$

The key κ is a twin secret. It is not a ledger field. Publishing κ would let an adversary write a consistent mark on a spliced host. Rotating κ after a confirmed compromise is a governance act in the sense of digital-transformation resilience programs, not a signal-processing act.

Exponential modulation of FKF coefficients supports both growth analysis and integrity watermarks [22]. Write the marked transform as

$$X^m(\omega, \tau) = X(\omega, \tau) + \alpha M(\omega, \tau) 1_{\Omega_m}(\omega, \tau) \quad (5)$$

or, in the multiplicative form that modulation language suggests,

$$X^m(\omega, \tau) = X(\omega, \tau) \left(1 + \alpha \exp(j \theta(\omega, \tau; \kappa)) \right) \quad \text{on } \Omega_m \quad (6)$$

Here α is a small gain, κ is a key known to the twin, and $\theta(\cdot; \kappa)$ is a structured phase. The time-shift property [21] says a lawful delay $t \mapsto t - t_0$ moves X by a structured displacement in (ω, τ) . The mark is written so that the same displacement carries M with X . A splice that is not a delay does not.

A lawful delay is a structured displacement of the same plane that carries the mark.

$$T(x(t - t_0)) = X^{\text{shift}}(\omega, \tau - t_0) \quad (16)$$

A QIM variant is available when capacity matters more than smoothness:

$$X^m = Q_b(\kappa)(X) \quad \text{on } \Omega_m \quad (7)$$

with Q_b a payload-dependent quantizer. For lead-time hosts the additive or multiplicative rule (5)–(6) is preferred because QIM on a sparse impulse-ready kernel [14] can look like a stockout. For current hosts either rule is admissible.

Inversion uses the FKF path or the FKL path [20]. The marked time series $x^m = T^{\{-1\}}(X^m)$ is what the twin stores beside the ledger commitment. Embedding is done after the twin estimator

$$\hat{x}_{t+1} = f(\hat{x}_t, u_t) + K(y_t - h(\hat{x}_t)) \quad (8)$$

so that the mark rides the estimated host, not a raw spoofable y [1]. Blockchain-attested fields in y still reduce the incentive to spoof the innovation [3]. Power-line carrier remains an admissible measurement path [4].

V. Detection, Windows, and Dual Integrity

False-alarm and miss rates should be stated on disruption ensembles drawn from the twin, not on a single week of clean data. A miss on a spliced current while residues ring is the costly error. A false alarm on



a lawful solar transient is the operational nuisance. Threshold λ is therefore a joint choice with the energy model, not a universal constant.

Miss and false-alarm rates are defined on twin ensembles, not on a single clean week.

$$P_{\text{miss}} = \Pr(\Lambda \geq \lambda \mid \text{splice}), \quad P_{\text{fa}} = \Pr(\Lambda < \lambda \mid \text{lawful shift})(17)$$

When radios fail and only a power-line carrier delivers a low-rate state, the watermark window lengthens. Sequential test (10) is what makes that lengthening explicit instead of implicit. A dock that can carry state but cannot carry a short window is still a valid sensor if n^* is allowed to grow.

A. Agreement statistic

Given a later series z and a claimed window W , form $Z = T(z|_W)$ and the detector

$$\Lambda(Z, X^m, \kappa) = \frac{\text{Re} \langle Z - X, M(\kappa) \rangle_{\Omega_m}}{\|M(\kappa)\|_{\Omega_m} \|Z\|_{\Omega_m} + \varepsilon}(9)$$

Accept the watermark if $\Lambda \geq \lambda$ and the ledger identity for the same window is valid [3]. Reject if either test fails. An adversary who forks neither the ledger nor the kernel still fails if the two disagree. That is the sentence this paper exists to make operational.

B. Sequential window test

Sequential convergence in the testing-function space [15] is the rule that decides when W is long enough. Write

$$n^* = \min(n : |\Lambda_n - \Lambda_{n-1}| \leq \delta \text{ and } n \geq n_{\min})(10)$$

A quiet window that is too short must not explode either ρ or Λ . The same ε that stabilizes (4) stabilizes (9). Cross-listed sequential statements [15] are treated as one claim, not two replications.

C. What the two channels answer

The ledger answers who committed this event. The watermark answers whether this time series still looks like the series that was committed [22]. Physical-layer telemetry can still carry a low-rate state on existing conductors [4]. Charging cabinets and powered docks in electrified fleets recreate that geometry [9], [12]. The architecture does not assume ubiquitous 5G at every supplier.

Event y	Twin host $\hat{x}$ [1]
Ledger commit [3]	Mark (5)–(6) [22]
Identity check	Λ and n^* (9)–(10)

Figure 2. Dual integrity: ledger identity and spectral agreement.

VI. Threat Model

Out of scope, and named so that they stay out: cryptanalysis of the ledger, physical theft of a pack, and type-approval fraud. Those are real risks. They are not watermark risks. In scope is the adversary who already has write access to a published series and wants the twin to treat that series as the series that was committed.

A delay-mask adversary is the most interesting case. The time-shift property says a pure delay is structured. A delay plus a new mode is not. The detector should therefore be run on the residual after a shift estimate, not on the raw window. That residual test is how watermarking reuses, rather than duplicates, shift-invariant lead-time analysis.

A lawful delay is a structured displacement of the same plane that carries the mark.

$$T(x(t - t_0)) = X^{\text{shift}}(\omega, \tau - t_0)(16)$$

Four hostile acts are in scope. Splice: replace an interval of L_k by a forged dwell while keeping ledger events. Replay: publish a clean historical i_v over a ringing charger. Delay-mask: hide a new resonant mode inside a true delay. Twin-theft: rebuild (8) and the Hamiltonian of the discrete actuator on stale or stolen data [8], [11].

Lawful acts are out of scope as attacks. A berth loss that translates L_k is a time-shift [21]. A solar transient that changes P_{sol} is a designed source [12], [25]. A planner override that rejects two slots is a human actuator [24], not a spoof. A corridor ban at 08:00 is a missing actuator [9], not a late shipment.

The detector is allowed to be informed. The twin stores X^m and κ . Blind detection is a later engineering option, not a claim of this paper. Near-term discrete hardware limits [7], [11] bound how often the Hamiltonian can be rebuilt; they do not bound how often Λ can be computed.

VII. Coupling to Actuators and to the Horizon Objective

Pricing a failed watermark inside the stage cost is what keeps integrity from becoming a parallel bureaucracy. If Λ is only a dashboard lamp, dispatchers will ignore it under time pressure. If Λ enters ℓ_t , the green policy and the discrete solver see it as they see carbon and energy. That is the difference between an architecture and a poster.

The integrity bit is the product of ledger identity and spectral agreement.

$$I_t = 1_{\text{ledger}} \cdot 1_{\{\Lambda \geq \lambda\}}(20)$$

Human override remains legitimate when the mark is healthy. Industry 5.0 does not require the operator to



rubber-stamp a high- Λ recommendation. It requires that the override be logged. A logged override on a healthy mark is labor policy. A logged override on a failed mark is incident response.

A. Green stage cost

Bibliometric work on AI for sustainable logistics organizes methods around routing, demand, and emissions [2]. A representative stage cost is

$$\ell_t = c_{\text{inv}} + c_{\text{en}} + c_{\text{CO}_2} + \mu \rho_t + \nu 1_{\{\Lambda_t < \lambda_j\}}(11)$$

The ρ term prices ringing. The new ν term prices a failed watermark. Without both, a dispatcher can look green on paper while ringing the lead-time spectrum or while running on a spliced host [18], [2].

B. Discrete and human actuators

When the decision is packing or disruption re-slotting, quantum annealing encodes a Hamiltonian as in ULD configuration studies [8]. Opportunity and survey papers locate that actuator inside a hybrid loop [7], [11]. The twin rebuilds the Hamiltonian after each major event [1]. A failed Λ is a reason not to rebuild from the suspect window.

The discrete actuator remains an Ising-family Hamiltonian rebuilt only after a passing integrity bit.

$$H(s) = \sum_i^N h_i s_i + \sum_{i < j} J_{ij} s_i s_j(18)$$

Industry 5.0 treats the planner as a legitimate actuator [24]. Formally

$$u_t = a_t u^{\text{rec}} + (1 - a_t) u^{\text{hum}}(12)$$

where $a_t = 1$ means the recommendation is accepted. Interface ancestry includes swipe-style controls [6]. Override rate remains a first-class log. A planner who rejects a slot because Λ failed is doing integrity work, not noise.

C. Horizon problem

Collecting the blocks, the planning problem on a horizon H is

$$\min \sum_{t=1}^H \ell_t + \psi T_{\text{rec}}(13)$$

subject to (1)–(3), energy floors from (2), spectral floors on (4), assignment feasibility, the accept/override rule (12), and $\Lambda_t \geq \lambda$ on committed windows. Recovery time T_{rec} is estimated on disruption ensembles drawn from the twin, which is how resilience reviews asked measurement to be done [13], [26]. Internal complexity is the dimension of $\hat{x}$ plus the number of roles that may set a_t [27]. Adding sensors without adding (4) and (9) often raises internal complexity faster than it lowers T_{rec} [13].

VIII. Mobility, Energy, and Electrical Residues

Weather-dependent solar power changes the right-hand side of the energy balance on a clock that delivery windows already use. Aligning irradiance with those windows is therefore both a routing duty and a

watermark duty. A mark written at noon on a dark dock will not look like a mark written at noon on a bright dock if the reserved set was chosen on current amplitude rather than on current shape. Shape-first embedding is the safer rule.

Classroom sizing still matters. Students who can size a pack for a known track can be shown why chance constraints replace that track, and why a watermark on current must survive the same chance constraints. The algebra of (2) does not change when the measure on P becomes stochastic. The detector must be written so that it does not either.

Smart mobility surveys list connected corridors, charging, and policy as first-class constraints on urban freight [9]. A van that cannot enter a zone at 08:00 is a missing actuator. Solar-assisted design frameworks [12] and contest architectures [25] change the right-hand side of (2) in weather-dependent ways that routing AI [2] must see. Battery reviews change feasible traction power [10], [28].

Residue methods watch whether the electrical plant is ringing while the material plant looks calm [16]. The watermark on i_v is the integrity twin of that residue test. If residues say the charger is ringing and Λ says the published current is clean, the two observers disagree and the window is hostile. Formula-style sizing remains the cleanest classroom reduction of (2) [5]. Logistics replaces the track with a stochastic corridor [9] and certainty with chance constraints. The algebra does not change; the measure on P does.

The same Hamiltonian family used for ULD packing [8] can serve charger and gate assignment if the twin exports a common interface [7], [11]. A failed watermark on the dwell series is a reason not to trust that interface for the current window.

IX. Two Worked Narratives

End-to-end simulation of either narrative is future work. Every object in both narratives is already cited. The missing object is a shared clock file: y, L_k, i_v , the ledger bit, Λ, ρ , override, thermal margin, irradiance. Until that file exists, claims about commutation remain architectural.

A third miniature helps teaching. Take a clean weekly cycle, embed (6), apply a pure delay, and show Λ holding. Then splice four hours of forged dwell and show Λ falling while the ledger still lists a named party. That pair of figures is the smallest honest demonstration of dual integrity.

A coastal depot loses a berth. The twin [1] raises dwell in L_k . FKF shift analysis [17], [21] shows the delay is a translation, not a new resonant mode, so ρ dips only



mildly and Λ stays above λ if the mark rode the shift. Green AI [2] reroutes inland. Several ULDs must be rebuilt; the annealer returns a configuration [8]. A planner rejects two slots that would break a labor rule [24]. Meanwhile a solar-hybrid tractor [12] reports a current harmonic; residues [16] tag the depot charger, not the pack chemistry [10]. The quality event is committed [3] and watermark-checked [22]. Recovery time estimated on the twin enters (13).

Now change one fact. The published current is clean and Λ is high, but residues on the dock are not. Dual integrity fails even though the ledger is consistent. That is the spoof class watermarking exists to catch.

A second narrative is quieter. No berth is lost. Weekly order cycles amplify at echelon three. Shift-invariant spectra [17] catch the amplification before fill rate moves. Differential FKF properties [19] predict that a modest buffer at echelon two will detune the mode. The twin tests that prediction [1]. AI does not need to discover the buffer by reinforcement if the spectrum already named the pole [18]. Sequential window tests [15] decide when the quiet narrative has enough samples. FKL inversion [20] is available if the FKF window is awkward. The watermark on L_k should survive the buffer change because the change is a lawful residue move, not a splice.

Disruption + lawful shift	Quiet mode	weekly
Λ holds, ρ mild	Buffer	detunes pole [19]
Spliced current vs residue	Dual check fails	

Figure 3. Lawful narratives keep Λ ; spliced current against residues does not.

X. What This Watermark Refuses

These refusals are scoped to the cited program. Other transforms, other QPUs, and other chemistries remain legitimate. The claim is only that a loop which omits a continuous-series check is an open loop, regardless of how many dashboards it drives.

It refuses a twin without an observer channel that can see oscillations [1], [18]. It refuses green AI that cannot see energy floors [2], [5]. It refuses quantum calls on stale Hamiltonians [8], [11]. It refuses ledgers that never look at the continuous signal [3], [22]. It refuses automation that hides override [24]. It refuses mobility programs that treat the pack as an infinite bus [10], [12]. It refuses resilience dashboards that never appear in a constraint [13], [26].

It also refuses a watermark written into the same bands that ρ must watch, a mark that dies under a lawful time-shift [21], and a mark that ignores sequential window length [15]. FKL remains an admissible inversion [20]. New batteries change numbers in (2), not the block diagram [28]. New ITS policies change the feasible set of (12), not the existence of a human actuator [9], [6]. Contest solar work remains an existence proof that the right-hand side of (2) is designable [25]. Weaving-resilience language still names internal and external complexity [27].

XI. Research Agenda and Limitations

Humanitarian and perishable settings will need different weights in (4) and different labor rules in (12). They should not need a different philosophy of loop closure. Time-shift and differential identities are the first instruments those settings should reuse. Quantum opportunity notes and transportation surveys remain the discrete-actuator reading list. Sequential convergence and FKL inversion remain the observer reading list.

Commensurability remains a risk. A testing-function space and a shop-floor override live in different categories. Green-AI and twin reviews remain bibliometric and survey-grade evidence. Contest architecture and classroom sizing are existence proofs for the energy balance, not fleet statistics.

Publish a toy twin with exportable y and L_k so that FKF code [14], [18] and annealer embeddings [8] can be compared on the same clock as Λ . Report ρ next to fill rate. Report override rate next to automation rate [24]. Report pack thermal margin next to on-time percentage [10]. Put solar irradiance on the same clock as delivery windows [12], [9]. Report Λ next to the ledger accept/reject bit.

Prove, or fail to prove, that pricing ρ and Λ in (11) reduces bullwhip amplification relative to carbon-only cost [2], [18]. Prove, or fail to prove, that watermark-plus-ledger catches a class of spoofs that either control alone misses [3], [22]. Measure whether power-line side channels [4] still matter on electrified docks. Keep teaching sizing [5] and simple interfaces [6] so that students who will meet (3) and the discrete Hamiltonian still know a watt and a button.

No new field experiment is reported here. Equations (1)–(13) reorganize published constructs. Quantum sections inherit near-term device limits [7], [11]. Spectral sections inherit the assumptions of the FKF/FKL series [18], [20]. Mobility sections inherit battery and solar coverage [10], [12]. Early hardware citations [4], [6] are conceptual, not a 2026 bill of



materials. A testing-function space [15] and a shop-floor override [24] live in different categories. The architecture claims only that both appear in the same loop. Single-author corpus bias is inherited from the reference list. Distributional companions [14] should not be counted as independent replications of later numberings of the same kernel.

A one-hour classroom reduction uses only (2), (3), (6), and (9). Size a pack for a known track. Take a synthetic lead-time sinusoid. Embed a known phase. Delay it. Compute Λ . Splice it. Compute Λ again. The four equations are enough to show why a ledger name does not replace a continuous-series check, and why a lawful delay is not an attack. Swipe-style interfaces can display Λ as a single lamp. Students who will later meet a Hamiltonian still leave the hour knowing a watt, a button, and a mark.

XI-A. Teaching reduction

The same release should state α , the rule that built Ω_m , $n_{\min}$ and δ in (10), and λ . Without those four numbers the detector cannot be reproduced. With them, independent FKF code and independent annealer embeddings can be compared without arguing about narratives.

A release that wants to be compared should publish, on one clock: the raw measurement y ; the estimated hosts L_k, i_v, q_k ; pack energy E_v and thermal margin; solar irradiance; corridor and zone indicators; p ; Λ ; ledger accept/reject; recommended u ; accepted a_t ; fill rate; on-time percentage; and estimated T_{rec} on a declared disruption ensemble. Missing any of those columns makes commutation unverifiable. Extra columns are welcome. Missing Λ is not.

XI-B. Reporting template for a toy twin

If the discrete actuator is called in the same window, the Hamiltonian is rebuilt only after step seven returns a pass. If the human rejects a slot, the rejection is logged with the same pair. If residues on a charger disagree with a high Λ on published current, the window is flagged as dual-integrity failure regardless of the ledger name. That seven-step list is not a protocol standard. It is the minimum script that makes the block diagram executable on a toy twin.

The implementer can treat the following as a single procedure on each committed window. First, the twin updates the estimated host by (8) from attested measurements. Second, the FKF integral (3) is computed on the current window. Third, forbidden bands are read from the current weight w in (4) and the reserved set Ω_m is the complement inside a designer mask. Fourth, the mark is written by (6) with gain α

and key κ . Fifth, the marked coefficients and a hash of $(\Omega_m, \alpha, \text{window id})$ are stored beside the ledger commitment. Sixth, at audit time the detector forms (9) and the sequential length test (10). Seventh, the pair (ledger bit, $\Lambda \geq \lambda$) is written into the stage cost (11) and, if required, shown to the human actuator (12).

XII. Conclusion

Multi-echelon resilience under electrified logistics can be written as a control loop whose integrity observer is a spectral watermark plus a ledger. The plant is stock plus pack energy. The state observer is a twin plus an FKF channel. The watermark is an exponential modulation of reserved quiet coefficients in the same plane that already watches bullwhip and electrical resonance. Detection is an agreement statistic with an explicit window test. Lawful delay is a translation. Spliced dwell and replayed current are not.

References

- [1] A. Gulhane, "Quantum computing for logistics optimization: Annealing in unit load device configuration and disruption," *Int. J. Res. Publ. Rev.*, vol. 7, no. 6, pp. 4643–4648, Jun. 2026.
- [2] A. Gulhane, "Smart mobility and intelligent transportation systems: Emerging trends, technologies, and challenges," *IJSREM*, vol. 9, no. 11, Nov. 2025.
- [3] A. Gulhane, "Recent advances in electric vehicle battery technologies: Materials, management systems, and sustainability perspectives," *IJSREM*, vol. 9, no. 7, Jul. 2025.
- [4] A. Gulhane, "Quantum computing in transportation and logistics: Current state, industrial applications, and future directions," *IJSREM*, vol. 9, no. 9, Sep. 2025.
- [5] A. Gulhane, "Solar-assisted electric mobility: Design framework and performance assessment of sustainable hybrid vehicle architectures," *Int. J. Eng. Res. Sci. Technol.*, vol. 21, no. 3, pp. 362–368, 2025.
- [6] A. Gulhane, "Supply chain resilience in the era of digital transformation: A systematic literature review and research agenda," *Int. J. Eng. Res. Sci. Technol.*, vol. 21, no. 4, pp. 1058–1068, 2025.
- [7] A. Gulhane, "Distributional extension of the FKF transform and its relevance to supply-chain systems," *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026.
- [8] A. Gulhane, "Sequential convergence in the testing function space of the two-sided FKF



transform,” *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026.

[9] A. Gulhane, “Spectral analysis for multi-scale supply chain dynamics and residue-based resonance extraction in hybrid-vehicle systems,” *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026. DOI: <https://doi.org/10.55041/ijcope.v2i8.141>

[10] A. Gulhane, “Shift-invariant spectral analysis of multi-echelon supply-chain lead times via the FKF transform,” *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026. DOI: <https://doi.org/10.55041/ijcope.v2i8.136>.

[11] A. Gulhane, “Spectral framework based on the FKF transform for supply chain dynamics and resilience,” *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026. DOI: <https://doi.org/10.55041/ijcope.v2i8.140>.

[12] A. Gulhane, “Spectral differential properties of the distributional FKF transform with applications to resilient supply-chain systems,” *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026. DOI: <https://doi.org/10.55041/ijcope.v2i8.139>.

[13] A. Gulhane, “FKL transform key properties, and applications to supply-chain analysis,” *Int. J. Eng. Sci. Adv. Technol.*, vol. 25, no. 12, pp. 596–605, 2025. DOI: <https://doi.org/10.64771/ijesat.2025.v25.i12.3816>.

[14] A. Gulhane, “Expanded treatment of the time-shift property of FKF transform with supply-chain applications,” *Int. J. Eng. Sci. Adv. Technol.*, vol. 25, no. 12, pp. 606–616, 2025. DOI: <https://doi.org/10.64771/ijesat.2025.v25.i12.3817>.

[15] A. Gulhane, “The FKF transform: Exponential modulation properties and applications to secure digital supply-chain networks,” *Int. J. Eng. Sci. Adv. Technol.*, vol. 25, no. 12, pp. 617–623, 2025. DOI: <https://doi.org/10.64771/ijesat.2025.v25.i12.3818>.

[16] A. Gulhane, “Spectral differential properties of the distributional FKF transform with applications to resilient supply-chain systems,” *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026. DOI: 10.55041/ijcope.v2i8.139.

[17] A. Gulhane, “Sequential convergence in the testing function space of the two-sided FKF

transform,” *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026. DOI: 10.55041/ijcope.v2i8.138.

[18] A. Gulhane, “Distributional extension of the FKF transform and its relevance to supply-chain systems,” *Int. J. Creative Open Res. Eng. Manage.*, vol. 2, no. 8, Aug. 2026. DOI: 10.55041/ijcope.v2i8.137.

[19] A. Gulhane, “Internet of Things, artificial intelligence, and quantum computing: A convergent framework for smart logistics ecosystems,” *IJESAT*, vol. 24, no. 12, pp. 297–317, 2024. DOI: 10.64771/ijesat.2024.v24.i12.3372.

[20] A. Gulhane, “Industry 5.0 and intelligent logistics: Transforming supply chain operations through human-centric automation,” *IJESAT*, vol. 24, no. 12, pp. 287–296, 2024. DOI: 10.64771/ijesat.2024.v24.i12.3371.

[21] A. Gulhane, “The future of vehicles: Solar-powered battery electric hybrid vehicle architecture,” *Tech Briefs Contest*, 2020.

[22] A. Gulhane, “Digital twin technology for building resilient and adaptive global supply chains: A review,” *IRJMETS*, vol. 6, no. 12, 2024. DOI: 10.56726/IRJMETS64992.

[23] A. Gulhane, “Artificial intelligence applications in sustainable logistics and green supply chain management: A bibliometric review,” *IRJMETS*, vol. 6, no. 12, 2024. DOI: 10.56726/IRJMETS65006.

[24] A. Gulhane, “Artificial intelligence and blockchain integration for enhancing supply chain transparency, traceability, and resilience,” *IRJMETS*, vol. 6, no. 12, 2024. DOI: 10.56726/IRJMETS65016.

[25] A. Gulhane, “Power line carrier communication based anti-theft system,” *IJRIM*, vol. 4, no. 12, pp. 1–11, 2014.

[26] A. Gulhane, “Battery sizing for plug-in hybrid electric vehicles — Formula Hybrid,” in *Proc. IEEE ICPCSI*, 2017, pp. 368–372. DOI: 10.1109/ICPCSI.2017.8392317.

[27] A. Gulhane, “Swipe controller,” *IJREAS*, vol. 4, no. 12, pp. 1–7, 2014.

[28] A. Gulhane, “Navigating the quantum revolution in logistics: Opportunities and practical applications in supply chain management,” *IJET*,



vol. 12, pp. 553–556, 2026. DOI:
10.29126/ijet.2026.v12.i3.553.

[29] S. M. Harle et al., “Digital twins and artificial intelligence for sustainable textile raw material processing,” IRJMETS, vol. 8, pp. 14–30, Jul. 2026.

[30] S. M. Harle, “Artificial intelligence-driven framework for village-level water conservation and groundwater sustainability,” IRJMETS, vol. 8, pp. 31–49, Jul. 2026.

[31] S. M. Harle et al., “Artificial intelligence approaches for strength prediction and optimization of composite concrete mixtures: A systematic review and future research framework,” IRJMETS, vol. 8, pp. 81–97, Jul. 2026.

[32] H. M. Ozaktas, Z. Zalevsky, and M. A. Kutay, The Fractional Fourier Transform. Wiley, 2001.

[33] S. C. Pei and M. H. Yeh, “The discrete fractional cosine and sine transforms,” IEEE Trans. Signal Process., vol. 49, pp. 1198–1207, 2001.

[34] A. H. Zemanian, Distribution Theory and Transform Analysis. New York, NY, USA: McGraw-Hill.

[35] A. H. Zemanian, Generalized Integral Transformations. New York, NY, USA: Interscience.