



The Noumerica Encryption Paradigm:
Salting via Nonlocal Operations

Patent application:

A. Sowa, F. Bui, J. Norton, G. Harris McEwen, E. Hassani Sangani (co-inventors)
Noumerica Data Security Inc. (applicant)

Title: *Dirichlet Convolution-based Cryptography Methods and Systems*

Filed under the Patent Cooperation Treaty (PCT)

International Application: **PCT/CA2024/050722**

Filing date: 30 May 2024

Priority date: 5 September 2023

Status: Entered national phase (US, Canada), March 2026

Collaboration: Interdisciplinary team spanning mathematics, computer engineering, and tele-medicine

Research status: The mathematical foundations and information-theoretic properties of the DCM construction are developed in a manuscript prepared for submission.

Traditional cryptography relies either on computational hardness (in the public-key setting) or on carefully engineered local nonlinearity combined with iteration (in the symmetric setting, such as Advanced Encryption Standard). Noumerica-style cryptography (symmetric) instead relies on nonlocal structure and salting, so that no part of the data carries meaning on its own.

The One-Time Pad: Ideal Security, Structural Limits

Encryption as a group operation: $c = x + k$ in a finite abelian group G .

It becomes information-theoretically secure when:

- The key k is selected uniformly at random from G .
- Every element of G is a valid plaintext (no structural bias).
- The ciphertext reveals no statistical information: $I(x, c) = 0$.

Important nuance: In the classical OTP, $G = \bigoplus \mathbb{Z}_2$. Security is perfect — but only if the key is truly random and never reused.

Key takeaways: Utilizing a group operation for encryption is very attractive. However:

Perfect secrecy via OTP requires a key as long as the message.

Ciphertext length reveals message length.

Key reuse collapses security.

OTP is mathematically ideal, but its large-scale use is constrained by key generation, distribution, storage, and strict non-reuse.

From Local OTP to Nonlocal Structural Encryption

Noumerica retains group operation not unlike of OTP — but changes the structure.

Key structural differences:

- The plaintext and key are arrays over a finite field (characteristic 2 or higher).
- Encryption uses an abelian group operation $\zeta = \kappa * \epsilon$ (Dirichlet convolution).
- The operation is *nonlocal*: each ciphertext component depends on multiple components of the plaintext and the key, according to the multiplicative/divisor geometry of the index set.

Important nuance: Higher characteristics (e.g. $p = 257$) are natural for structured data such as images, allowing direct encryption of pixel values without binary fragmentation.

Key takeaway: Encryption becomes structured, nonlocal mixing rather than bitwise-local masking. The dependency pattern is determined by multiplicative structure, not by geometric proximity.

Dirichlet Convolution as Structured Mixing

Key sequence: $\kappa = [\kappa_1, \dots, \kappa_N]$.

Dirichlet convolution corresponds to multiplication by a structured matrix:

$$\kappa \xleftrightarrow{\quad} K = \begin{pmatrix} \kappa_1 & \cdot & \cdot & \cdot & \cdot & \cdots \\ \kappa_2 & \kappa_1 & \cdot & \cdot & \cdot & \cdots \\ \kappa_3 & \cdot & \kappa_1 & \cdot & \cdot & \cdots \\ \kappa_4 & \kappa_2 & \cdot & \kappa_1 & \cdot & \cdots \\ \kappa_5 & \cdot & \cdot & \cdot & \kappa_1 & \cdots \\ \kappa_6 & \kappa_3 & \kappa_2 & \cdot & \cdot & \kappa_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad \epsilon = \begin{pmatrix} \epsilon_1 \\ \epsilon_2 \\ \epsilon_3 \\ \epsilon_4 \\ \epsilon_5 \\ \epsilon_6 \\ \vdots \\ \vdots \end{pmatrix}$$

Encryption:

$$\zeta = \kappa * \epsilon = K\epsilon$$

Important observation: Each ciphertext component aggregates multiple plaintext components.

Key property: If $\kappa_1 \neq 0$, the matrix is invertible.

Such operators admit a quantum interpretation: J. Fransson, B. Sanders, A. Sowa, *Macroscopically distinguishable superposition in infinitely many degrees of freedom*, International Journal of Theoretical Physics (Springer), 65:16 (2026)

Under the Hood: Structural Encryption and Salting

Key selection and encryption.

A secret key $\kappa \in \mathcal{L}_N^*$ is chosen uniformly at random (binary or more generally over a finite field $p > 2$).

Encryption:

$$\zeta = \kappa * \epsilon$$

Decryption:

$$\epsilon = \kappa^{*-1} * \zeta$$

Important structural difference from OTP: The operation is *nonlocal*. Each output component depends on multiple input components.

Salting mechanism: Before encryption, the plaintext is extended:

$$[\mu_1, \dots, \mu_k, \epsilon_1, \dots, \epsilon_{N-k}],$$

where μ_j are uniformly random.

After decryption, the random prefix is discarded.

Security implications: Salting randomizes the effective input before the nonlocal transformation. The injected randomness propagates through the multiplicative structure of Dirichlet convolution, producing fresh ciphertexts under repeated use of the same key. The key length need not match plaintext length (e.g. shape the plaintext as a square, act on columns and rows).

Why Nonlocality Changes the Attack Geometry

Without salting:

If an attacker knows ϵ and ζ , then $\kappa = \epsilon^{*-1} * \zeta$.

With salting:

The true input has the form $[\mu_1, \dots, \mu_k, \epsilon_1, \dots, \epsilon_{N-k}]$, where μ_j are unknown random components.

Important consequence:

The attacker does not know the effective input, so direct group inversion is no longer available. More importantly, the random prefix is propagated nonlocally through the ciphertext.

Key insight:

Security is not merely a larger search problem. Fresh randomness is transported through the multiplicative dependency structure, disrupting deterministic alignment between plaintext and ciphertext coordinates.



Random data masks the structure of the input.

From Local to Global Masking

Setting: Encryption is performed over $\mathbb{F}_p$, $p > 2$, with fresh uniformly random salt and nonzero key coefficients under the hypotheses of the masking theorem.

Local masking: A protected ciphertext coordinate ζ_n is uniformly distributed in $\mathbb{F}_p$ and reveals no information about the corresponding plaintext symbol:

$$I(\zeta_n; \epsilon_n) = 0.$$

Stronger result: For coordinates reached by the propagated salt, the same ciphertext coordinate is independent of the entire information-bearing plaintext vector ϵ :

$$I(\zeta_n; \epsilon) = 0.$$

Interpretation: The masking is not merely coordinate-to-coordinate. Nonlocal Dirichlet mixing propagates fresh randomness so that an individual protected ciphertext coordinate carries no statistical information about the plaintext as a whole.

Precise hypotheses and coordinate ranges are given in the accompanying DCM manuscript.

Key Reuse with Fresh Salts

Setup:

Multiple encryptions use the same secret key, but each encryption receives a fresh, independent random salt.

What the masking result gives:


For every encryption separately, protected ciphertext coordinates retain the information-theoretic masking guaranteed by the theorem.

Why fresh salt matters:

Repeated use does not reproduce the same effective input. Each encryption introduces new randomness before the nonlocal Dirichlet transformation.

Key takeaway:


The construction is designed to support key reuse with fresh salting. Claims about repeated observations should be understood under the precise hypotheses and scope of the DCM security results, rather than inferred from marginal uniformity alone.



A fresh random frame changes the effective input at every encryption.

Flat Histogram: Empirical Illustration

Encryption over $\mathbb{F}_p$ produces uniform symbol distribution.



Observation: The encrypted image appears noise-like, and its histogram is statistically flat.

Interpretation: The histogram is consistent with uniform marginal symbol frequencies in $\mathbb{F}_p$. A flat histogram alone does not establish statistical independence between ciphertext coordinates.

Security Beyond Brute-Force Key Size

Large key spaces are available.

For example, a length-512 key with 256 admissible values per coordinate already gives a nominal space of

$$256^{512} = 2^{4096}.$$

But key-space size is not the central security argument.

The distinctive mechanism is structural:

- fresh randomness is inserted into the effective plaintext,
- Dirichlet convolution propagates that randomness nonlocally,
- masking results quantify what selected ciphertext coordinates reveal about the plaintext.

Key takeaway: Security analysis should focus on information leakage and structural attacks, not merely on the cost of exhaustive key search.

Efficient Implementation via Dirichlet Structure

Goal: Compute $\zeta = \kappa * \epsilon$ efficiently.

Key idea: Exploit divisor structure to achieve *quasi-linear complexity* $O(N \log N)$.

Lifting scheme (Dirichlet convolution)

Initialize $\zeta = [0, \dots, 0]$.

```
for  $k = 1 : N$   
   $\zeta(k : k : N) += \kappa(k) \cdot \epsilon(1 : \lfloor N/k \rfloor)$   
end
```

Interpretation: Each coefficient $\kappa(k)$ propagates along its multiples — a structured, nonlocal update pattern.

Origin: This structure arises from number-theoretic matrix factorization.

A. Sowa, *Factorizing matrices by Dirichlet multiplication*, Linear Algebra Appl. 438 (2013), 2385–2393

Cross-Platform Performance

512×512 Grayscale Image

- Python encode: 80.03 ms
- Rust encode: 23.00 ms
- CUDA encode: 56.90 ms

1024×1024 Grayscale Image

- Python encode: 153.44 ms
- Rust encode: 107.00 ms
- CUDA encode: 140.25 ms

Key observation: Rust implementation consistently achieves the best performance.

Medical Imaging Use Case (DICOM)

673×1102, 28-frame, 3-channel DICOM file

- Python encryption: 22.73 ms
- Rust (single-thread): 4.42 ms
- Rust (12-core): 3.45 ms
- CUDA: 4.23 ms

Key takeaway: Multithreaded Rust achieves sub-4 ms encryption for complex medical data.

The Noumerica Paradigm: From Primitive to Architecture

Structured nonlocal cryptography

Dirichlet convolution over a finite field provides invertible mixing governed by multiplicative/divisor geometry rather than geometric locality.

Information-theoretic masking

For $p > 2$, fresh salting can produce protected ciphertext coordinates that are uniform and, under the theorem hypotheses, independent of the information-bearing plaintext.

Efficient implementation

The divisor structure supports quasi-linear algorithms, with systems-level implementations already demonstrated on image and medical-data workloads.

Composable keys

Keys combine multiplicatively under Dirichlet convolution, enabling composite cryptographic transformations without requiring a retained master key.

Toward Security through Structural Silos (S^3)

These properties motivate a broader security architecture in which information and authority can be distributed across independently protected domains, limiting the consequences of compromise or privilege escalation.

Direction: DCM provides the cryptographic foundation; S^3 explores the system architecture built upon it.