ASX Encryption Algorithm

Architecture, Diffusion Characteristics, Statistical Behavior and Experimental Evaluation

White Paper v2_0_04

TargetDeLab Group

September 2026

Abstract

This white paper presents the ASX encryption/decryption architecture and its experimentally motivated evolution into Bilateral Diffusion ASX. The original ASX combines APPEND, SWAP and XOR inside a segmented, repeated-pass chain and can generate high byte entropy, near-balanced bit populations and low short-lag correlation. Earlier experiments nevertheless identified a specific intrinsic-diffusion weakness: for incompressible random input, where compression does not first broaden a one-bit plaintext disturbance, the original architecture transports the disturbance over long distances but changes too few ciphertext positions to create a source-independent global avalanche. A first explicit diffusion enhancement greatly increased differential multiplication for favorable source positions, but a five-position experiment exposed strong positional dependence. The final Bilateral Diffusion ASX introduces a proprietary reversible whole-state diffusion stage. Its internal transformation sequence and inverse construction are intentionally not disclosed in this public white paper. Native MATLAB experiments show that this change fundamentally alters the measured random-source behavior. In a five-position bilateral diagnostic the mean ciphertext avalanche was 49.808% with 0.437% standard deviation and 100% propagation span. A final eight-position, five-pass experiment measured 47.272% mean avalanche, 99.997% mean normalized propagation span, 98.878% mean changed-byte density within that span, and all eight ciphertext regions affected in every trial. The paper directly compares the original and bilateral architectures and concludes that, for the experimentally identified diffusion objective, Bilateral Diffusion ASX is decisively superior: it supplies an explicit, reversible, source-independent propagation mechanism that removes the severe directional collapse and converts sparse random-source perturbations into dense global ciphertext differences. These results characterize diffusion and statistical behavior; they do not by themselves establish cryptographic security or resistance to attack.

Contents

1 Purpose and Scope **4**
- 1.1 What the principal measurements mean 4
- 1.2 From statistical quality to cryptographic security 4

2 Operational Data Path **5**

3 The Three Principal ASX Transformations **6**
- 3.1 APPEND: controlled state expansion 6
- 3.2 SWAP: data-dependent positional rearrangement 7
- 3.3 XOR: key-dependent value transformation 7
- 3.4 Architectural simplicity as a practical advantage 8
- 3.5 Why combine the three operations? 8

4 Experimental Metrics **8**
- 4.1 Differential avalanche 8
- 4.2 Entropy and bit balance 9
- 4.3 Correlation and spatial propagation 9

5 Experimental Methodology **9**

6 Segmentation and Differential Behavior **10**

7 Statistical Appearance of the Ciphertext **10**

8 Source-Class Dependence **11**

9 Intrinsic-Diffusion Stress Test **11**

10 Component-Level Interpretation **12**

11 Pass Accumulation and Performance Cost **13**

12 Where Does a One-Bit Difference Travel? **14**

13 Controlled APPEND-Size Sensitivity **15**

14 Evolution from Original ASX to Proprietary Bilateral Diffusion **17**
- 14.1 Why an enhancement was required 17
- 14.2 Intermediate enhancement and the positional limitation 17

15 Bilateral Diffusion ASX Architecture **18**
- 15.1 Public architectural description 18

15.2 Relationship to the original ASX architecture 19
15.3 Disclosure boundary 19

16 Bilateral ASX Experimental Evidence **19**
16.1 Five-position directional diagnostic 19
16.2 Pass-dependent buildup on incompressible random input 20
16.3 Final eight-position propagation map 21

17 Direct Architectural Comparison **23**
17.1 Like-for-like random-source evidence 23
17.2 Strict comparison by demonstrated property 23
17.3 Why Bilateral ASX is better for the measured diffusion objective 25

18 Engineering Trade-offs of the Bilateral Layer **26**

19 Validity Limits and Remaining Verification **27**

20 Original ASX Quality Assessment **27**
20.1 Demonstrated strengths 27
20.2 Demonstrated limitations 27
20.3 What the experiments do not prove 28

21 Final Conclusions **28**

22 References **30**

1 Purpose and Scope

Encryption algorithms are often introduced through implementation details, which makes it difficult for a new reader to understand what the transformation is trying to accomplish. This document takes the opposite approach. It first explains the mathematical role of the three principal ASX processes, then defines the measurements used to evaluate them, and finally presents the experimental evidence.

The objective is not to claim that statistical tests prove cryptographic security. The U.S. National Institute of Standards and Technology (NIST), in Special Publication 800-22 Rev. 1a, explicitly cautions that statistical testing cannot substitute for cryptanalysis [4]. Entropy, frequency balance, correlation and avalanche measurements illuminate different properties and must be interpreted separately. The central question is therefore: *how effectively does the tested ASX processing chain disperse information and differential changes, and under what input conditions?*

1.1 What the principal measurements mean

Before using these measurements, it is useful to state their meaning in plain terms and the theoretical targets normally sought from random-looking ciphertext. **Entropy** measures uncertainty in the observed symbols. For a byte-oriented stream there are 256 possible values, so the maximum Shannon entropy is 8 bits per byte; an ideally uniform byte stream approaches this value. **Frequency balance** asks whether symbols occur without systematic preference. At bit level the ideal is an equal population of zeros and ones, $P(0) = P(1) = 0.5$; at byte level an ideally uniform distribution gives each byte value probability $1/256$. **Correlation** measures simple linear dependence between values at different positions. For a random-looking ciphertext the desired nonzero-lag correlation is approximately zero. Zero linear correlation, however, does not prove statistical independence. **Avalanche** measures how strongly a controlled small input change affects the output. In a one-bit perturbation experiment, the ideal changed-bit fraction is approximately 50%; 100% is not a better avalanche result because it would indicate a deterministic complement relation rather than unbiased diffusion.

Finite samples normally fluctuate around these ideals; exact equality is neither expected nor required. Approaching these targets is desirable evidence of ciphertext quality, but none of them individually—or collectively—constitutes proof of cryptographic security.

1.2 From statistical quality to cryptographic security

A serious evaluation of an encryption/decryption algorithm proceeds through several distinct levels of evidence. First comes **functional correctness**: encryption and decryption must be mutually consistent and recover the original message exactly. Second comes **statistical**

Table 1: Plain-language interpretation and theoretical targets of the principal measurements.

Property	Plain meaning	Ideal theoretical target
Byte entropy	Uncertainty/uniformity of 256 byte values	8 bits/byte
Bit balance	Relative population of zeros and ones	50% / 50%
Byte frequency	Relative occurrence of each byte value	1/256 each
Nonzero-lag correlation	Simple linear dependence between separated values	0
One-bit avalanche	Output bits changed by one input-bit perturbation	50%

and diffusion characterization: entropy, frequency balance, correlations, avalanche and propagation measurements can reveal obvious structure or weak diffusion. Third comes **cryptanalytic evaluation**: the construction should be examined against attack models and techniques applicable to its design, including exhaustive key search, known-plaintext and chosen-plaintext analysis, chosen-ciphertext analysis where relevant, differential and linear techniques, related-key or structural attacks where applicable, and reduced-round analysis to estimate security margin. Fourth comes **implementation security**: timing, power, electromagnetic leakage, fault behavior and other side channels can compromise an otherwise sound algorithm. Finally, confidence in a new cipher normally requires sustained independent scrutiny, reproducible analysis and, where appropriate, formal standardization or validation processes. The public evaluation and standardization of AES illustrate this broader model rather than reliance on a single statistical test [5].

Cryptographic *security strength* is commonly expressed in bits as an estimate of the computational work needed for the best relevant attack. A nominal k-bit key does not automatically establish k bits of effective security: structural cryptanalysis may reduce the required work below exhaustive search. Consequently, this paper does **not** assign ASX or Bilateral Diffusion ASX a numerical security strength. The present evidence establishes measured statistical and differential properties and a comparative diffusion improvement. A general claim of cryptographic security would require a substantially broader body of independent attack analysis.

2 Operational Data Path

The evaluated complete machine is represented as

$$P \xrightarrow{\text{compression}} Z \xrightarrow{\text{ASX}} C, \tag{1}$$

where P is plaintext, Z is the compressed byte stream and C is ciphertext. For differential experiments a single source bit is changed,

$$P' = P \oplus e_i, \tag{2}$$

and both paths are processed independently:

$$P \to Z \to C, \qquad P' \to Z' \to C'. \tag{3}$$

This separation is important because compression can itself spread a one-bit source change over many compressed bits. The experiments therefore report the differential at the compression boundary and again at the final ciphertext.

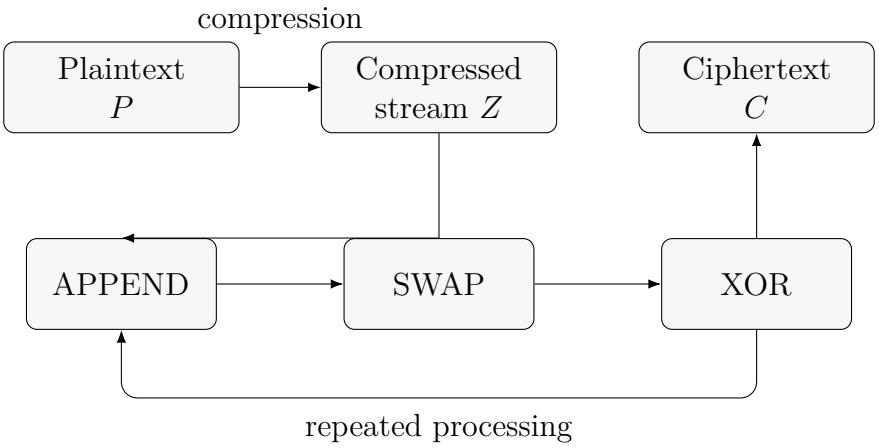


Figure 1: Conceptual ASX data path. Labels are separated from process boxes and connectors for readability. The figure describes algorithmic transformations rather than software control details.

3 The Three Principal ASX Transformations

3.1 APPEND: controlled state expansion

APPEND enlarges the working byte sequence by introducing additional bytes. If an input segment is

$$X = (x_1, x_2, \ldots, x_n), \tag{4}$$

APPEND constructs an expanded state

$$A(X) = (x_1, x_2, \ldots, x_n, r_1, r_2, \ldots, r_m), \tag{5}$$

where r_j denotes appended data generated during processing. When segmented operation is used, this expansion is applied repeatedly to the segment stream rather than merely once at the end of the complete message.

APPEND has two important consequences. First, it increases the state on which subsequent transformations operate. Second, repeated application can make the working representation substantially larger than the compressed source. The experiments later in this paper show that this process is strongly associated with the transition from a distributed compression-stage perturbation to a near-50% final differential for structured inputs. It is therefore more than padding in the ordinary storage sense: in the tested ASX chain it changes the geometry and size of the transformation state. The measurements establish the importance of APPEND as an active state-expansion mechanism. A subsequent controlled sensitivity experiment also varied the nominal APPEND amount independently while holding the remaining tested configuration fixed. As shown later, increasing the nominal APPEND amount from 16 to 256 bytes per segment increased ciphertext size but did not increase global avalanche on incompressible input. Thus APPEND size and APPEND participation must be distinguished: APPEND is important to the composite architecture, but simply appending more data does not by itself create strong differential multiplication.

3.2 SWAP: data-dependent positional rearrangement

SWAP operates from opposite ends of a sequence. Consider two symmetrically selected values x_i and x_j. A data-dependent parity decision is formed from their values. In simplified mathematical form,

$$s_{ij} = (x_i + x_j) \bmod 2. \tag{6}$$

When $s_{ij} = 1$, the two positions are exchanged; otherwise they remain in their current order. The pointers then move inward and the operation continues. Thus SWAP is a conditional permutation rather than a fixed reversal.

Its purpose is positional scrambling. It does not necessarily change the values of the bytes; instead it changes where selected values occur. Because the decision depends on the data being processed, a modified input can alter later positional relationships. In segmented operation the permutation is performed within the processing regions, making segment length an important architectural parameter.

3.3 XOR: key-dependent value transformation

The XOR process combines each processed data byte with a byte drawn from a protection-key sequence. For byte x_i and key byte $k_{q(i)}$,

$$y_i = x_i \oplus k_{q(i)}, \tag{7}$$

where $q(i)$ maps the data position into the key sequence. XOR is reversible:

$$y_i \oplus k_{q(i)} = x_i, \tag{8}$$

which makes it naturally suitable for a symmetric encryption/decryption chain. Its principal role is value masking: unlike SWAP, which primarily changes position, XOR directly changes byte values according to key material.

The ASX architecture also evolves the working key across repeated processing. Conceptually, this means that later transformations need not apply an identical mask to the state seen in earlier passes. The implementation details of those internal controls are intentionally outside the scope of this public architectural description.

3.4 Architectural simplicity as a practical advantage

A notable engineering advantage of ASX is the simplicity of the elementary actions required by its principal transformations. APPEND extends a sequence, SWAP exchanges selected positions, and XOR performs a bitwise reversible combination with key material. These actions are straightforward to understand and map naturally to software or digital hardware. The architectural interest lies not in the complexity of any single primitive, but in the repeated composition of simple state expansion, positional rearrangement and value transformation. This simplicity can ease implementation, verification and hardware realization; it must not, however, be interpreted as evidence of cryptographic strength by itself.

3.5 Why combine the three operations?

The processes address different dimensions of scrambling:

$$\underbrace{\text{APPEND}}_{\text{state expansion}} \quad + \quad \underbrace{\text{SWAP}}_{\text{position}} \quad + \quad \underbrace{\text{XOR}}_{\text{value/key interaction}} \quad . \tag{9}$$

Repeated processing allows these effects to accumulate. The experimental question is whether accumulation merely produces random-looking output or also causes a small input difference to multiply into a global output difference.

4 Experimental Metrics

4.1 Differential avalanche

For equal-length byte streams U and V, let $HD(U, V)$ denote bitwise Hamming distance. The changed-bit percentage is

$$A(U, V) = 100 \frac{HD(U, V)}{8|U|}. \tag{10}$$

We report

$$A_{ZIP} = A(Z, Z'), \qquad A_C = A(C, C'). \tag{11}$$

A value near 50% is used as a descriptive benchmark for broad differential mixing. This is not, by itself, a proof of the Strict Avalanche Criterion or of cryptographic security; avalanche-style criteria are conventionally treated as separate design properties rather than statistical certification [3].

4.2 Entropy and bit balance

Following the standard information-theoretic definition of entropy [1], byte entropy is estimated as

$$H(C) = -\sum_{b=0}^{255} p_b \log_2 p_b, \tag{12}$$

whose maximum is 8 bits per byte for a perfectly uniform byte distribution. Bit-plane balance measures the fraction of ones in each bit position; values near 0.5 are desirable for an unbiased-looking stream.

4.3 Correlation and spatial propagation

Short-lag Pearson correlations test whether nearby ciphertext bytes exhibit simple linear dependence. Differential propagation is additionally described by the first and last changed output byte, the span between them, the number of changed bytes, and the fraction of coarse output regions touched by the perturbation.

5 Experimental Methodology

The authoritative results in this paper were produced by the native MATLAB implementation. The original encryption/decryption implementation was held fixed; measurement code was external to the device under test. Paired differential experiments used the same operating configuration and controlled random conditions for the baseline and one-bit-perturbed input. For segmented comparisons, the segment size was derived from the baseline compressed length and held fixed within each pair so that a change in compressed length could not silently change the ASX configuration.

The campaign included text-like data, highly structured synthetic patterns, sparse data, incrementing bytes and incompressible random data. This variety is essential: an encryption chain that appears excellent on a compressible source may behave very differently when compression contributes almost no differential spreading.

6 Segmentation and Differential Behavior

A 64 KiB text-like source compressed to 401 bytes in the representative campaign. Thirty-two independent one-bit perturbations were evaluated at four segment regimes. Table 2 shows a highly non-monotonic dependence on segment size.

Table 2: Differential behavior versus segment size.

Segment (B)	L_S/L_Z	Mean A_{ZIP} (%)	Mean A_C (%)	Mean $H(C)$
401	1.000	36.229	47.751	7.7161
200	0.499	36.229	49.500	7.9336
100	0.249	36.229	**49.922**	7.9907
64	0.160	36.229	**0.0976**	7.9939

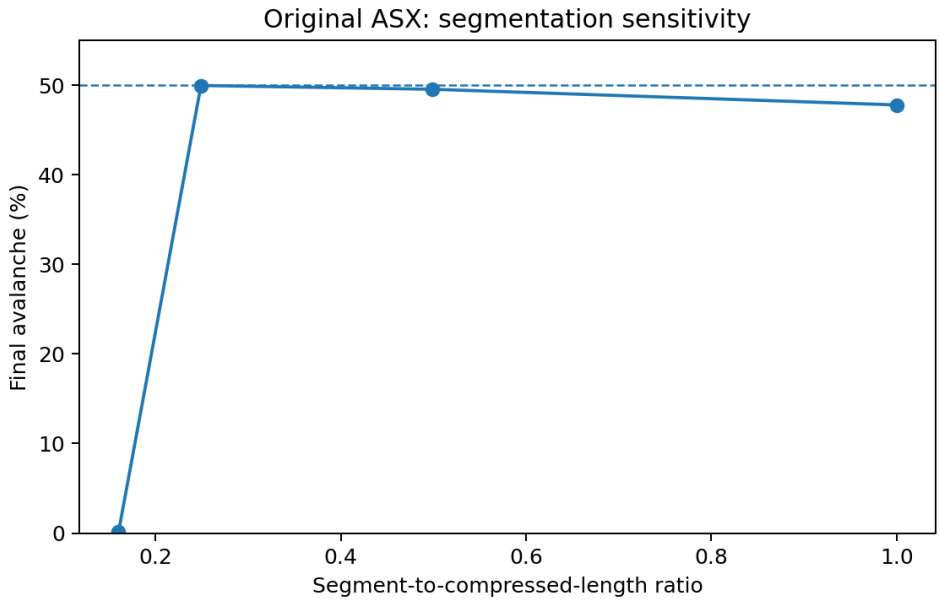


Figure 2: Final avalanche is strongly dependent on segmentation. High entropy at the smallest segment setting does not prevent differential collapse.

The 100-byte operating point produced $49.922\% \pm 0.066\%$ final avalanche, while the 64-byte case collapsed to about 0.098% despite slightly *higher* entropy. This is the first major warning against using entropy as a proxy for diffusion.

7 Statistical Appearance of the Ciphertext

Frequency measurements reinforce the same distinction. As segmentation became finer, the overall one fraction approached 0.5 and the maximum bit-plane bias decreased from about 4.13% to about 0.31%. At the 64-byte regime the mean entropy was approximately 7.99393 bits/byte and all 256 byte values were occupied in the tested outputs. Nevertheless, the differential result at that operating point was extremely weak.

Correlation measurements also improved with finer segmentation. Representative worst absolute short-lag correlation decreased from approximately 0.184 at the full-length segment to 0.00383 at the 64-byte regime. Thus the stream can simultaneously exhibit high entropy, balanced bits and very low short-range correlation while failing a one-bit global-diffusion test.

> **Key interpretation:** statistical randomness-like appearance and differential avalanche are different properties. Neither should be used as a substitute for the other.

8 Source-Class Dependence

The complete compression-plus-ASX chain was then exercised on several 65,536-byte source classes. For each source, Table 3 reports the strongest observed final differential among the tested segmentation regimes; the purpose is to show source dependence, not to assert a universal optimum.

Table 3: Representative best measured differential by source class.

Source	Compressed bytes	A_{ZIP} (%)	Final A_C (%)	Entropy
All-zero	214	17.285	43.654*	7.9851
All-one	214	17.583	49.771	7.9847
Alternating AA/55	217	17.127	49.788	7.9840
Incrementing bytes	715	39.000	49.750	7.9599
Repeated text	382	33.540	49.918	7.9938
Sparse impulse	340	38.256	49.903	7.9937
Random/incompressible	65,692	0.00671	**0.01684**	7.9968

*The all-zero class was unstable at its strongest mean operating point (large run-to-run spread), so the mean should not be interpreted as a stable near-50% result.

This experiment is central to understanding the complete machine. For structured data, a one-bit source change can alter a substantial fraction of the compressed representation. ASX then frequently drives the final output close to the 50% benchmark. For incompressible random data, the compressed stream is nearly the same length as the source and the one-bit change remains extremely sparse. Under that condition, the tested ASX settings do not independently create global avalanche.

9 Intrinsic-Diffusion Stress Test

To test whether ordinary parameter tuning could overcome the incompressible-input weakness, the random-source experiment varied pass count, key length, segmentation and a separate adaptive-segmentation branch. The compression-stage differential remained

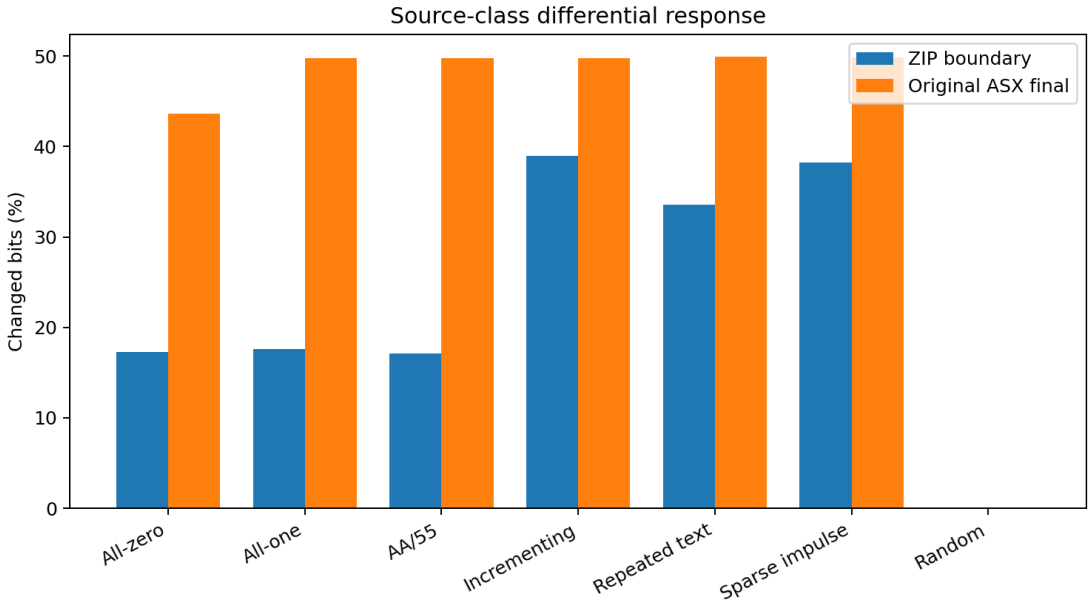


Figure 3: Compression-stage and final differential by source class. Structured/compressible sources can enter ASX with a broad differential; incompressible random input does not.

essentially constant at

$$A_{ZIP} = 0.00637\% \pm 0.00103\%. \tag{13}$$

Increasing passes from 1 to 20 changed the final mean avalanche only from approximately 0.0085% to 0.0139%. Key-length variation remained in the same order of magnitude. Finer fixed segmentation improved the result only modestly, with the strongest tested mean around 0.015%. Thus parameter tuning did not create a transition toward global avalanche.

This is a useful negative result: it indicates an architectural behavior rather than a poorly selected single parameter value.

10 Component-Level Interpretation

Ablation experiments removed major transformations one at a time while preserving the remainder of the processing chain. The strongest architectural observation concerned APPEND.

For incompressible random input, the full chain produced about 0.0112% final avalanche from a 0.00577% compression-stage differential. Without APPEND, the final differential fell to the same 0.00577% value. For repeated text, the full chain transformed approximately 34.3% compression-stage difference into approximately 49.8% final difference, whereas removal of APPEND left the final result near the incoming differential. At the same time, repeated-text ciphertext entropy fell from about 7.99 to about 4.34 bits/byte when state expansion was removed.

Removing SWAP or XOR individually did not produce the same collapse in the repeated-

text experiment. This does not mean those operations are cryptographically unnecessary; an ablation experiment measures behavior of the tested composite system and interactions can be nonlinear. It does show that APPEND is strongly associated with the observed state expansion and with the transition toward the high-diffusion statistical regime for structured input.

11 Pass Accumulation and Performance Cost

Repeated-pass measurements clarify how the structured-input differential develops. Table 4 compares the random and repeated-text sources with the full processing chain.

Table 4: Pass accumulation with APPEND active.

Passes	Random A_C (%)	Repeated-text A_C (%)	Repeated-text bytes	Entropy
1	0.0108	31.121	638	6.1568
2	0.0155	40.902	1,086	7.1273
5	0.0223	48.581	5,246	7.9059
10	0.0200	49.870	67,966	7.9913
20	0.0201	49.979	11,242,174	7.9943

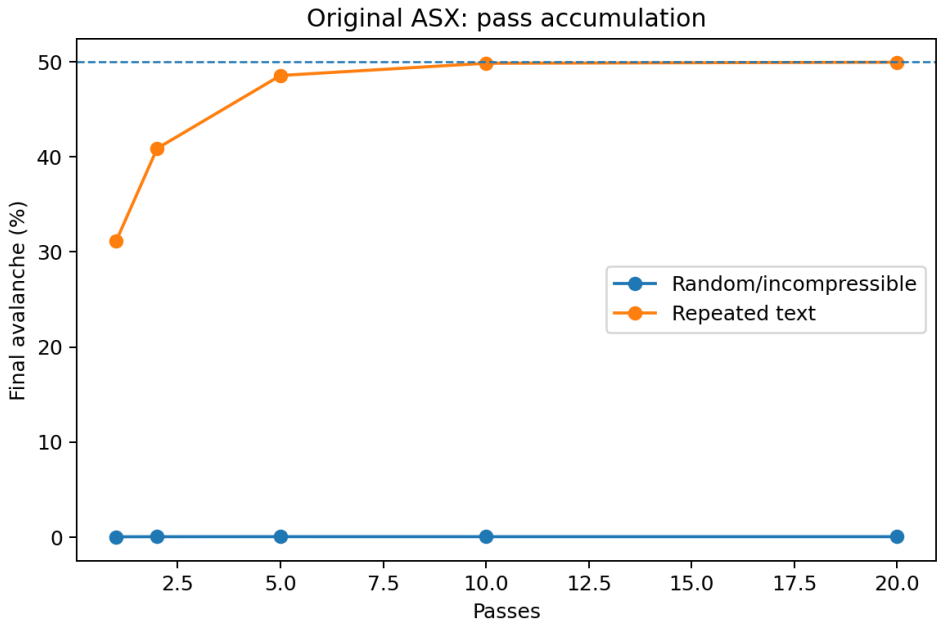


Figure 4: Pass accumulation strongly improves the structured-input differential but does not create global avalanche for incompressible random input.

The structured input approaches 50% as passes accumulate, but the cost can become extreme because state expansion compounds. The 382-byte compressed repeated-text representation grew to roughly 11.24 MB at 20 passes. The long-pass diagnostic campaign

was intentionally terminated after sufficient evidence had been collected because the computational and storage growth no longer justified the marginal information gained.

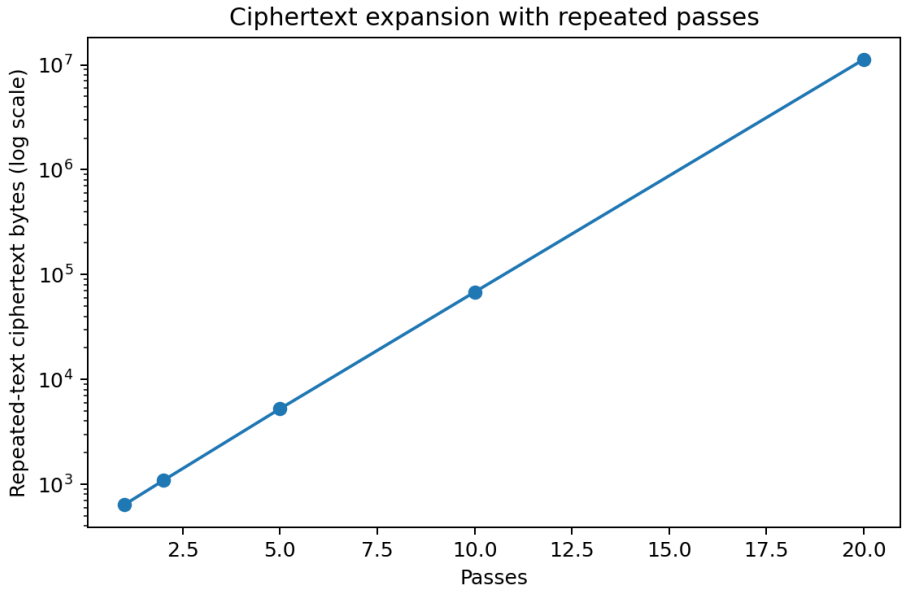


Figure 5: Measured ciphertext expansion for the repeated-text source. The vertical axis is logarithmic.

This is an important practical limitation. A high-pass configuration may improve statistical mixing for a compressible source while simultaneously imposing a severe expansion and execution-time penalty. Performance therefore cannot be separated from diffusion when selecting an operating region.

12 Where Does a One-Bit Difference Travel?

A final spatial experiment used only incompressible random input, five passes and eight one-bit perturbations distributed from the beginning to the end of the 65,536-byte source. The mean final avalanche remained only

$$A_C = 0.01580\% \pm 0.00595\%, \tag{14}$$

corresponding to an average of only 22.1 changed ciphertext bytes. Yet the mean span between the first and last changed output byte was approximately 33,800 bytes, or 50.3% of the ciphertext length.

The first source-bit perturbation changed only 28 output bytes, but those changes extended from ciphertext byte 47 to byte 66,849, spanning 99.37% of the output. Conversely, the perturbation at the final source byte changed 18 output bytes confined to the final 789-byte span (1.17% of the output). This positional asymmetry is consistent with predominantly forward propagation through the tested processing chain.

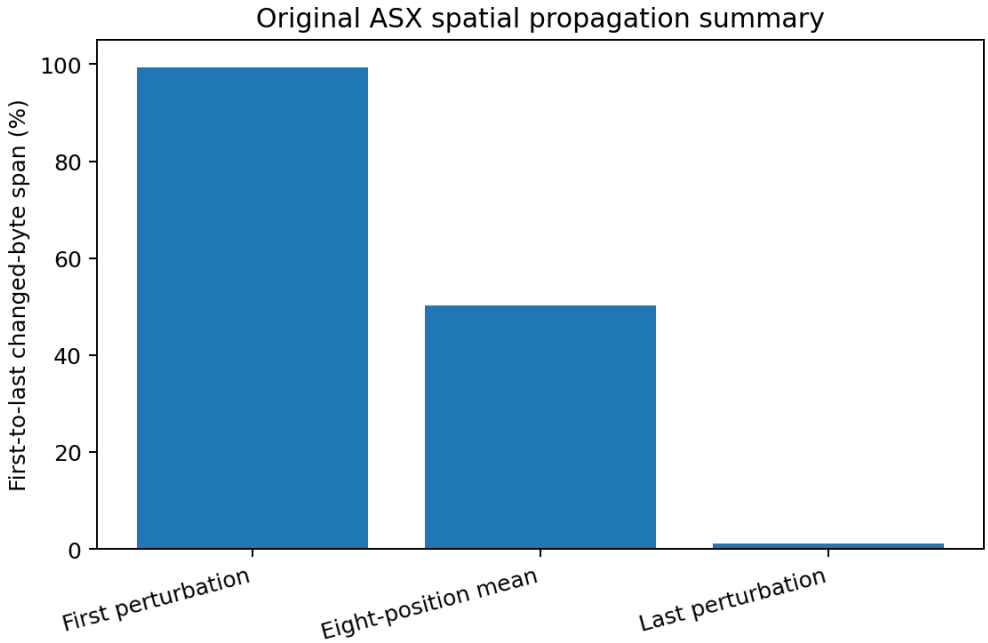


Figure 6: Spatial differential propagation summary. The first perturbation spans nearly the full ciphertext, the eight-position mean span is about one half of the ciphertext, and the final perturbation remains localized near the end.

The important distinction is therefore

$$\boxed{\text{long-range differential transport} \neq \text{strong differential multiplication.}} \tag{15}$$

ASX can carry a disturbance over a long distance, but for incompressible input it changes too few output bytes to form a global avalanche.

13 Controlled APPEND-Size Sensitivity

A dedicated sensitivity experiment tested whether the weak differential response of incompressible input could be improved simply by increasing the amount introduced by APPEND. Unlike the pass-count experiment, only the nominal APPEND amount was varied. The incompressible source, key, five-pass processing depth, segmentation and remaining transformation configuration were held fixed. Three paired one-bit perturbation trials were executed at each setting.

Across all five settings the mean compression-stage differential was 0.013891%, while ciphertext entropy remained essentially unchanged at approximately 7.997 bits/byte. Increasing nominal APPEND size by a factor of sixteen therefore increased the processed output size, but did not produce a corresponding increase in final avalanche. The mean final avalanche remained of order 10^{-2}%, and the changed-byte count stayed near only 46–53 bytes. With three trials per point, the small downward drift in the mean should

Table 5: Controlled APPEND-size sensitivity for incompressible input.

APPEND (bytes)	Mean A_C (%)	Std. A_C (%)	Changed bytes	Span (%)	Cipher size (bytes)
16	0.03802	0.02466	52.7	68.45	66,076
32	0.03480	0.02769	52.7	72.54	66,460
64	0.03254	0.02803	49.0	61.40	67,228
128	0.03339	0.02541	50.3	62.04	68,764
256	0.02952	0.02070	46.3	62.60	71,836

not be interpreted as a statistically established monotonic trend. The robust conclusion is narrower: **within the tested range, increasing APPEND size alone did not cure the weak intrinsic avalanche of incompressible input.**

Although very few ciphertext bytes changed, the mean first-to-last changed-byte span remained approximately 61–73% for most APPEND settings. This independently reinforces the propagation-map result: ASX can transport a sparse differential over long distance without multiplying it densely. The controlled sweep therefore supports the interpretation of APPEND primarily as a state-expansion and transport-enabling mechanism in the tested architecture, rather than as a sufficient differential-multiplication mechanism by itself.

14 Evolution from Original ASX to Proprietary Bilateral Diffusion

14.1 Why an enhancement was required

The original experiments did not motivate a redesign because of entropy, bit balance or short-lag correlation; those statistics were already strong in several operating regions. The redesign was motivated by a narrower differential observation. For incompressible random input the compression boundary carried only a sparse disturbance, and the original ASX architecture did not multiply that disturbance into a dense ciphertext difference. At five passes the earlier random-source campaign measured only about 0.0223% final avalanche. Spatially, an early perturbation could travel across nearly the entire ciphertext while only a small number of output bytes changed. Thus the engineering requirement was not merely to move a disturbance farther, but to introduce a reversible mechanism capable of multiplying a sparse disturbance throughout the processed state without relying on compression to perform that task first.

The development sequence was therefore empirical: Original ASX was measured, an intermediate diffusion enhancement was evaluated, a strong position-dependent limitation was observed, and a final Bilateral Diffusion ASX architecture was developed in response. The internal construction of the enhanced diffusion mechanism is intentionally proprietary and is not disclosed in this public white paper. The paper documents its externally measurable behavior, integration role and experimental consequences without publishing the transformation sequence, analytical recurrence, inverse relation, primitive operations, or implementation recipe.

14.2 Intermediate enhancement and the positional limitation

The first enhanced implementation demonstrated that an explicit diffusion stage could convert a sparse incoming random-source differential into a much larger ciphertext differential. A five-position experiment, however, exposed severe positional asymmetry. The measured ciphertext avalanche decreased from 48.108% for a perturbation at the beginning of the source to 3.572% for a perturbation at the end. Intermediate positions produced 8.319%, 6.474% and 5.294%, giving a mean of 14.353%.

The experiment therefore established a clear design requirement: the final enhancement had to provide strong differential response across the complete source-position range rather than only in favorable locations. This requirement was subsequently verified experimentally, as shown in the following sections.

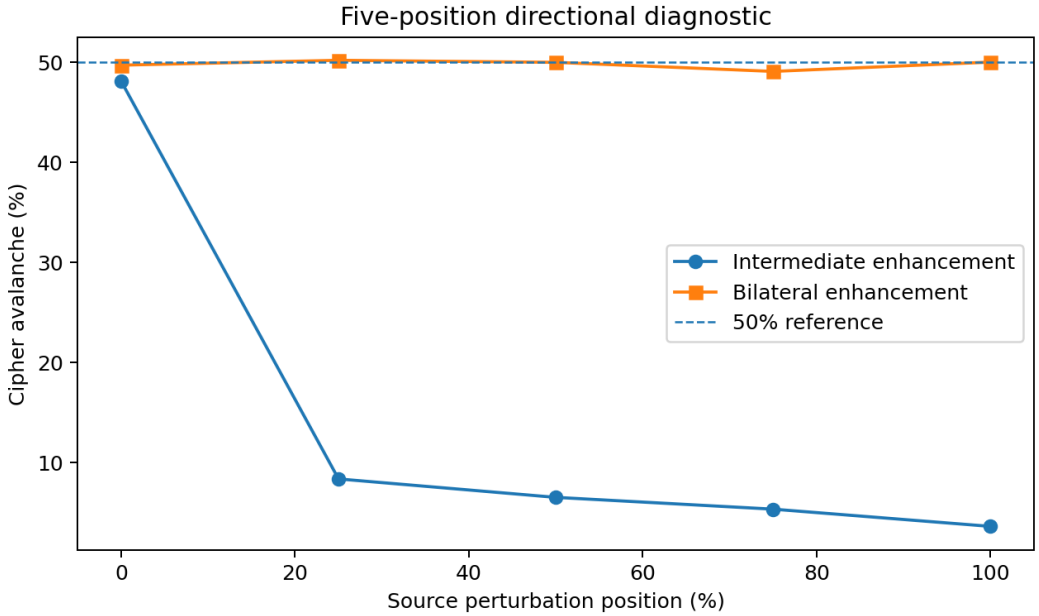


Figure 7: Five-position diagnostic. The intermediate enhancement exhibits strong position dependence, whereas the final Bilateral Diffusion ASX removes the observed positional collapse. Internal transformation details are intentionally omitted.

15 Bilateral Diffusion ASX Architecture

15.1 Public architectural description

Bilateral Diffusion ASX introduces a **proprietary, reversible, whole-state diffusion stage** into the repeated-pass ASX architecture. Its purpose is to provide deterministic differential propagation before the established ASX transformations continue processing the state. The stage is global with respect to the current pass and is designed so that a local source perturbation can influence the processed state independently of its original source position.

For intellectual-property protection, this document treats the bilateral stage as a black-box transformation. No internal equations, transformation ordering, primitive-operation sequence, inversion procedure, indexing rule, or implementation pseudocode is published. The experimentally relevant public interface is simply

$$\text{current pass state} \longrightarrow \boxed{\text{Proprietary Bilateral Diffusion}} \longrightarrow \text{continued ASX processing.} \tag{16}$$

The corresponding decryption implementation contains the private inverse required for exact recovery, but its construction is likewise outside the scope of this public document.



Internal transformation sequence intentionally undisclosed;
experimental input/output behavior is reported instead.

Figure 8: Public black-box representation of the Bilateral Diffusion ASX enhancement.

15.2 Relationship to the original ASX architecture

The enhancement does not replace the established ASX processing chain. APPEND continues to provide state expansion, SWAP provides positional rearrangement, and the original key-dependent value transformation remains part of the original ASX logic. The new bilateral stage has a distinct role: it supplies the explicit whole-state differential-propagation capability whose absence was exposed by the incompressible-input experiments.

This separation is important experimentally. The paired avalanche tests deliberately use the same initial pseudorandom seed for baseline and perturbed encryptions so that independently generated APPEND randomness cannot manufacture an artificial avalanche result. The improvement attributed to Bilateral Diffusion ASX is therefore evaluated through controlled differential experiments rather than inferred from random output appearance.

15.3 Disclosure boundary

The term *bilateral* in this white paper describes the experimentally observed design objective and behavior: strong propagation is obtained for perturbations distributed across the source-position range, without the severe positional collapse of the intermediate enhancement. It must not be interpreted as disclosure of the internal mechanism used to obtain that behavior. Reproduction of the private bilateral transformation requires implementation information that is intentionally absent from this document.

16 Bilateral ASX Experimental Evidence

16.1 Five-position directional diagnostic

The decisive first bilateral experiment repeated the positional stress test on incompressible random data. Across five source positions the compression-stage differential remained extremely small, with mean $A_{ZIP} = 0.006432\%$. The bilateral ciphertext results were 49.719%, 50.206%, 50.000%, 49.092% and 50.024%. Thus

$$\boxed{\overline{A_C} = 49.808\%, \qquad s = 0.437\%.} \tag{17}$$

Every trial exhibited 100% first-to-last changed-byte span, and the mean changed-byte density was 99.445%. The last-position perturbation, which produced only 3.572% under the unilateral enhancement, produced 50.024% under the bilateral construction.

This is the cleanest evidence that the bilateral modification solved the directional limitation identified by the unilateral experiment. The conclusion is deliberately specific: the tested bilateral construction removed the measured position-dependent collapse. It is not a claim of the strict avalanche criterion over every input bit and every key.

16.2 Pass-dependent buildup on incompressible random input

A second campaign examined how the bilateral response accumulates with repeated passes. Table 6 gives the mean of two paired random-source trials at each pass count, both with normal APPEND and with APPEND disabled.

Table 6: Bilateral random-source avalanche versus pass count.

Passes	FULL A_C (%)	NO APPEND A_C (%)	FULL bytes	FULL entropy
1	4.940	4.960	65,948	7.99755
2	9.640	12.855	66,268	7.99706
5	46.090	49.023	67,228	7.99715
10	50.650	50.433	68,828	7.99705
20	49.920	49.847	72,028	7.99742

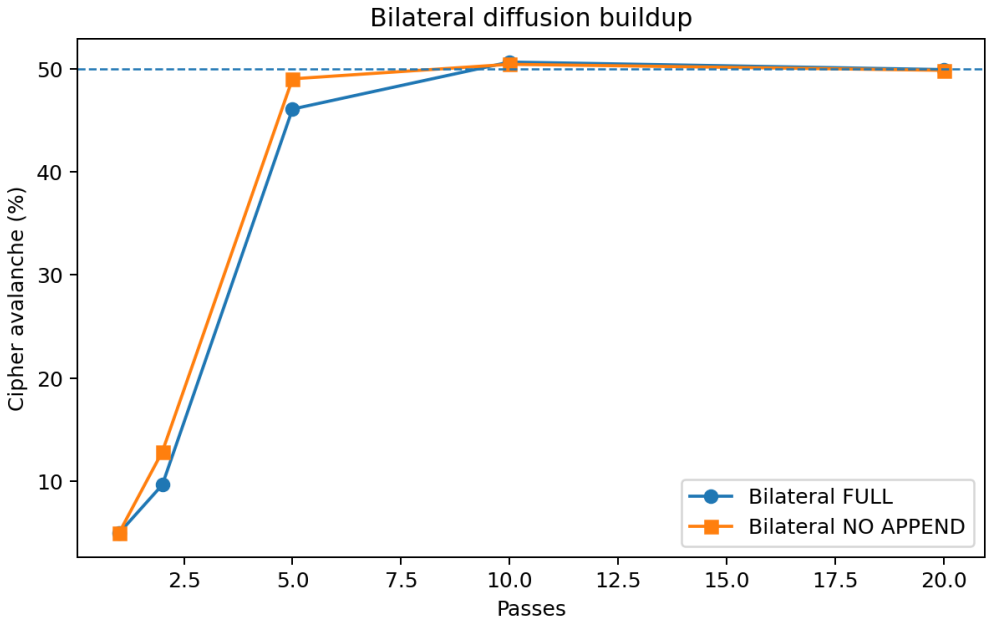


Figure 9: Bilateral diffusion builds progressively with repeated passes on incompressible random input. In this small sample, ten passes reached the near-50% regime and twenty passes produced no evident avalanche benefit.

Two conclusions follow. First, the bilateral mechanism changes the intrinsic random-source behavior by orders of magnitude compared with original ASX. Second, APPEND is not required for the bilateral design to reach the high-avalanche regime after sufficient passes: at ten passes the FULL and NO APPEND means were 50.650% and 50.433%, respectively. This does not make APPEND useless; it shows that APPEND and bilateral diffusion serve different functions.

The same campaign also demonstrates why pass count cannot be selected from avalanche alone. For repeated text, normal APPEND expanded the 382-byte compressed representation to 638 bytes at one pass, 5,246 bytes at five passes, 67,966 bytes at ten passes and 11,242,174 bytes at twenty passes. APPEND-on-key at ten passes produced 31,272,382 bytes in the tested configuration. Those repeated-text differential values are not used as clean intrinsic-avalanche evidence because the baseline and perturbed ciphertext lengths differed in all paired repeated-text cases, making position-aligned Hamming percentages sensitive to shifts as well as to diffusion.

16.3 Final eight-position propagation map

The final campaign was intentionally restricted to incompressible random input so that compression could not dominate the result. Eight one-bit perturbations were distributed uniformly from the first to the last source bit. Five bilateral passes, a 200-byte key, fixed segmentation at one quarter of the compressed length and normal APPEND were used. Baseline and perturbed compressed lengths were identical in all eight trials (65,692 bytes), and baseline and perturbed ciphertext lengths were also identical in every trial (67,228 bytes). The final Hamming measurements are therefore free of the unequal-length alignment problem that affected the repeated-text campaign.

Table 7: Final eight-position bilateral propagation experiment.

Position (%)	A_{ZIP} (%)	A_C (%)	Changed-byte density (%)	Regions
0.0	0.005899	51.567	99.558	8/8
14.3	0.005518	49.665	99.484	8/8
28.6	0.007421	35.144	94.655	8/8
42.9	0.006660	49.172	99.442	8/8
57.1	0.005138	49.424	99.552	8/8
71.4	0.006279	49.958	99.454	8/8
85.7	0.007040	42.551	99.406	8/8
100.0	0.005518	50.692	99.469	8/8

The aggregate result is

$$\overline{A_{ZIP}} = 0.006184\%, \tag{18}$$

$$\overline{A_C} = 47.272\%, \qquad s_{A_C} = 5.615\%, \tag{19}$$

$$\text{mean normalized span} = 99.9974\%, \tag{20}$$

$$\text{mean changed-byte density in span} = 98.878\%, \tag{21}$$

$$\text{mean affected regions} = 8/8. \tag{22}$$

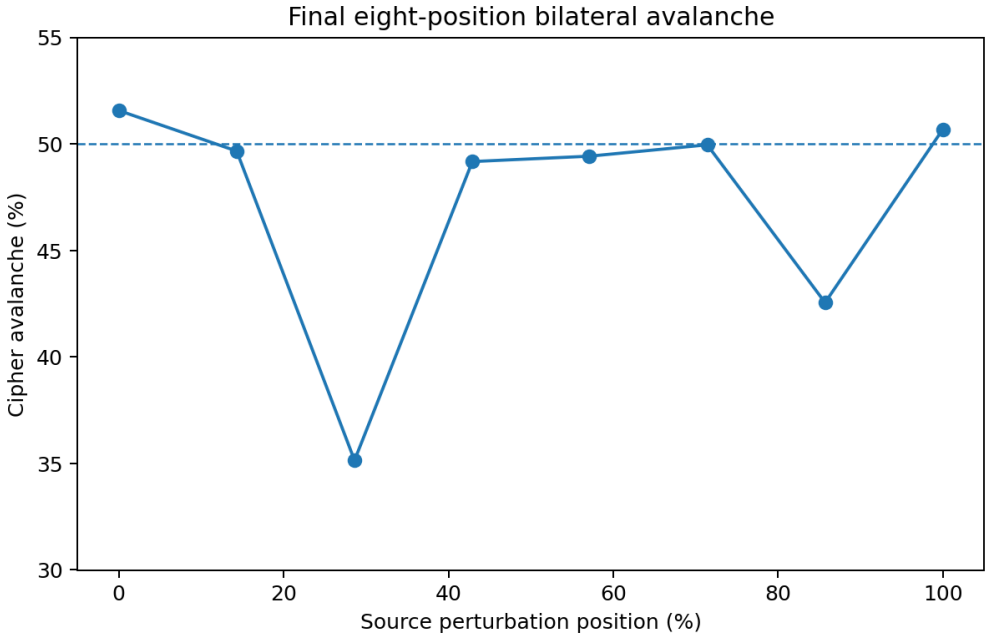


Figure 10: Final eight-position bilateral avalanche. The response is globally strong but not uniformly 50%; two positions produced 35.144% and 42.551%.

The lower 35.144% trial is scientifically useful rather than inconvenient. Its disturbance still occupied 99.993% of the first-to-last span, changed bytes with 94.655% density inside that span and reached all eight ciphertext regions. Therefore the experiment establishes an important distinction:

$$\boxed{\text{global spatial propagation} \neq \text{exact } 50\% \text{ bit avalanche.}} \tag{23}$$

Bilateral ASX can propagate a disturbance globally even when the detailed bit-change population for a particular perturbation falls below one half.

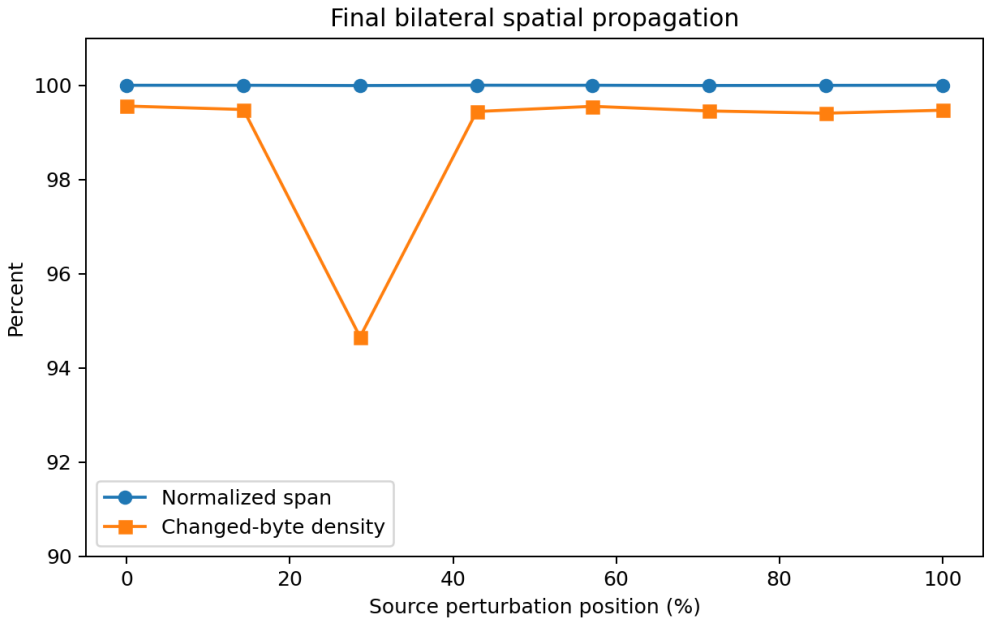


Figure 11: Spatial measures in the final bilateral experiment. Despite variation in bit avalanche, the disturbance reached essentially the full ciphertext and all eight regions in every trial.

17 Direct Architectural Comparison

17.1 Like-for-like random-source evidence

The strongest comparison is the incompressible random-source case because it minimizes help from compression. In the original architecture, five passes produced approximately 0.0223% final avalanche in the pass-accumulation experiment. Under Bilateral ASX, the corresponding five-pass random-source campaign produced 46.090% mean avalanche with normal APPEND; the final eight-position campaign independently produced a 47.272% mean across the source with equal ciphertext lengths in every pair.

Using the representative five-pass means, the change from 0.0223% to 46.090% is more than three orders of magnitude in measured changed-bit fraction. The exact ratio should not be treated as a universal constant because the campaigns use finite trial sets and different experimental instances; the robust conclusion is the regime change from a sparse differential of order $10^{-2}\%$ to a dense differential of order $10^{1}\%$.

17.2 Strict comparison by demonstrated property

Table 8 summarizes only properties supported by the experiments or by the architecture itself.

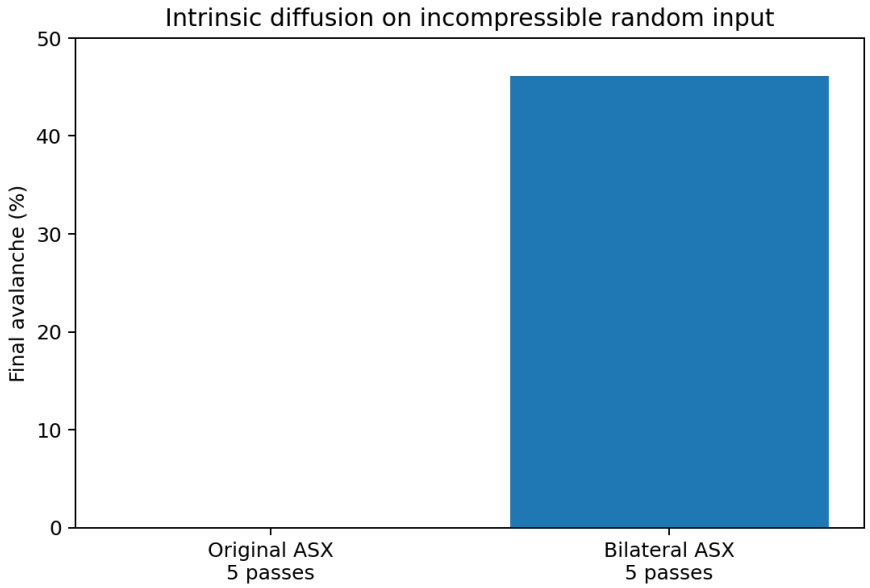


Figure 12: Representative five-pass random-source comparison. The bilateral architecture changes the intrinsic-diffusion regime rather than merely improving an already strong statistic.

Table 8: Original ASX versus Bilateral Diffusion ASX.

Property	Original ASX	Bilateral Diffusion ASX
Explicit global diffusion layer	None. Diffusion emerges from interaction of compression, APPEND, SWAP, XOR, segmentation and passes.	Yes. A proprietary reversible whole-state diffusion stage is integrated into repeated processing; internal construction is undisclosed.
Random-source intrinsic avalanche	Weak in tested architecture; five-pass result about 0.0223%.	Strong in tested architecture; five-pass mean 46.090% in pass sweep and 47.272% across eight final positions.
Directional dependence	Original spatial test showed early changes could travel far while late changes had little span.	Five-position bilateral diagnostic produced 49.092–50.206% across the entire source and 100% span in every trial.
Sparse differential multiplication	Weak: long-range transport occurred with only tens of changed bytes.	Dense: final eight-position test averaged 66,472 changed bytes out of 67,228 and 98.878% density inside the propagation span.

Property	Original ASX	Bilateral Diffusion ASX
Dependence on compression for strong avalanche	Strong complete-chain results occurred mainly when compression first broadened the differential.	Random/incompressible source itself reaches dense global differences; compression broadening is no longer required for the demonstrated diffusion effect.
Role of APPEND	Important for state expansion and for several original structured-source mixing results. Increasing APPEND size alone did not cure random-source weakness.	Still provides random redundant state, but strong random-source diffusion persists when APPEND is disabled after sufficient passes.
Entropy/statistical appearance	Often near 8 bits/byte and low short-lag correlation, even when avalanche is weak.	Remains near 8 bits/byte in tested random-source campaigns while differential behavior is also strong.
Implementation cost	Simpler pass path; cost dominated by original transformations and state expansion.	Adds a proprietary diffusion stage and therefore additional processing cost; exact internal operation count and implementation structure are intentionally undisclosed.
Security conclusion	Statistical evidence does not prove cryptographic security.	Same limitation: improved diffusion does not by itself prove cryptographic security.

17.3 Why Bilateral ASX is better for the measured diffusion objective

For the specific engineering objective established by the original experiments, the bilateral architecture is superior for four strict reasons.

First, it fixes the identified failure mode rather than optimizing a secondary statistic. Original ASX already had high entropy in the random-source case, yet its avalanche remained extremely weak. Bilateral ASX directly targets propagation and changes that measured behavior from sparse to dense.

Second, it removes the experimentally observed positional collapse. The intermediate design showed that the first enhancement remained strongly position dependent; the final

bilateral architecture produced near-50% response at both the beginning and the end of the source in the five-position diagnostic. The mechanism producing this behavior is intentionally undisclosed.

Third, it separates diffusion from random expansion. The NO APPEND bilateral measurements show that, after sufficient passes, strong random-source avalanche can persist without APPEND. APPEND can therefore retain its original state-extension role instead of being implicitly relied upon to manufacture diffusion.

Fourth, the final eight-position test demonstrates spatial globality directly. Every perturbation reached all eight output regions and essentially the complete ciphertext span. Even the two lower bit-avalanche cases retained dense spatial coverage. The improvement is therefore not an artifact of one favorable source position.

These points justify a strict architectural statement: **within the tested configurations and with respect to intrinsic differential diffusion, Bilateral Diffusion ASX is experimentally superior to Original ASX.** This statement is narrower than "more secure." Security superiority would require attack-oriented cryptanalysis beyond the present evidence.

18 Engineering Trade-offs of the Bilateral Layer

The enhancement is not free. The proprietary bilateral stage adds processing and memory activity to the original ASX path. The five-position bilateral diagnostic required roughly 205 seconds for a perturbed encryption in the tested MATLAB R2016a environment, approximately twice the earlier intermediate-enhancement diagnostic order of magnitude. This timing is implementation specific and is reported only as reference-software evidence; the internal operation count and transformation sequence are intentionally not disclosed.

A hardware realization of the proprietary bilateral stage would introduce additional datapath, storage/control and bandwidth considerations beyond Original ASX. Detailed mapping strategies are intentionally omitted because they would reveal implementation characteristics of the private transformation. No final RTL implementation of the bilateral layer was benchmarked for this white paper.

Pass count also matters. The random-source sweep indicates that five passes already enter the high-diffusion regime, ten passes reached approximately 50%, and twenty passes produced no evident additional avalanche benefit in the small sample. Meanwhile APPEND-driven state growth can become extreme on highly compressible data. A practical implementation should therefore select pass count from a joint diffusion/expansion/throughput requirement rather than maximizing passes by default.

19 Validity Limits and Remaining Verification

The evidence base is intentionally transparent. Several bilateral campaigns use small trial counts; the final eight-position map samples eight perturbations rather than all possible input bits. Near-50% average avalanche is not a proof of the strict avalanche criterion. Some earlier structured-source bilateral experiments produced unequal ciphertext lengths between baseline and perturbed paths, so their position-aligned Hamming percentages are descriptive but not clean intrinsic-diffusion measurements. The final random-source eight-position experiment avoids that issue because all paired lengths are equal.

The bilateral layer is designed to be exactly reversible, while its inverse construction remains proprietary. Before a production release or security claim, an explicit automated byte-for-byte round-trip regression over representative configurations should be retained as a release criterion:

$$\boxed{D_K(E_K(P)) = P.} \tag{24}$$

This requirement is separate from avalanche testing: diffusion can be excellent while an implementation contains a decryption defect, and exact reversibility can hold while diffusion is weak.

20 Original ASX Quality Assessment

20.1 Demonstrated strengths

The native measurements support several positive conclusions. ASX also has an implementation-oriented advantage: its principal transformations rely on simple elementary operations, which makes the conceptual data path comparatively easy to understand and potentially economical to realize in software or digital logic. In favorable segmentation regimes ASX produces high byte entropy, near-balanced bit planes and low short-lag correlations. For multiple structured/compressible source classes, the complete compression-plus-ASX chain transforms a substantial compression-stage differential into a final changed-bit fraction close to 50%. Repeated passes provide a clear progression toward that regime. The architecture also demonstrates long-range movement of differential effects, and its reversible XOR/value transformation combined with positional rearrangement and state expansion provides a richer transformation than a single masking operation alone.

20.2 Demonstrated limitations

The most important limitation is equally clear. Strong complete-chain avalanche is source dependent. When an incompressible source enters ASX with an extremely sparse compression-stage differential, the tested architecture does not multiply that disturbance into a global

avalanche. Increasing passes, changing key length and changing segmentation did not remove this limitation. High entropy and low correlation persist even in cases with extremely weak differential propagation, so those statistics cannot be interpreted as evidence of strong avalanche.

A second limitation is performance. Repeated state expansion can cause very large ciphertext growth for highly compressible inputs at deep pass counts. The controlled APPEND-size sweep strengthens this concern: increasing the nominal APPEND amount from 16 to 256 bytes increased ciphertext size from 66,076 to 71,836 bytes, yet the measured mean avalanche did not improve. Larger APPEND values therefore imposed additional state growth without curing the weak incompressible-input differential in this experiment. At deep pass counts the expansion can become much more severe and can dominate execution time and storage. Practical configurations therefore require a compromise between mixing behavior and expansion cost.

20.3 What the experiments do not prove

These experiments are a characterization of measured diffusion and statistical behavior, not a proof of resistance to cryptanalysis. They do not establish semantic security, indistinguishability under chosen-plaintext attack, resistance to differential or linear cryptanalysis, key-search complexity beyond nominal key-space considerations, side-channel resistance, or authenticated-encryption properties. A production security claim would require formal design analysis and independent cryptanalytic review in addition to the measurements presented here.

21 Final Conclusions

The experimental program documents a complete engineering progression from observation to architectural correction. Original ASX demonstrated attractive statistical behavior and strong complete-chain diffusion for several structured/compressible inputs, but incompressible random data exposed a fundamental limitation in intrinsic differential multiplication. A one-bit disturbance could travel far without changing enough ciphertext bytes. Increasing APPEND size, changing segmentation, changing key length and increasing pass count did not remove that limitation.

The first enhancement demonstrated that an explicit diffusion stage could multiply a sparse differential, but its measured response remained severely position dependent. That negative result directly motivated the final Bilateral Diffusion ASX architecture. The final design introduces a proprietary reversible whole-state propagation mechanism while preserving the established ASX processing framework. Its internal equations, operation sequence and inverse construction are intentionally not disclosed.

The measured improvement is substantial. In the five-position bilateral diagnostic the mean random-source avalanche was 49.808% with only 0.437% standard deviation, 100% propagation span and 99.445% changed-byte density. In the final eight-position five-pass experiment, where all baseline/perturbed lengths were equal, the mean avalanche was 47.272%, the mean normalized span was 99.9974%, the mean changed-byte density was 98.878%, and every perturbation affected all eight ciphertext regions. Individual avalanche values ranged from 35.144% to 51.567%, which prevents an unsupported claim of uniform 50% avalanche but does not alter the observation of global spatial propagation.

The direct comparison therefore supports a clear engineering conclusion: **Bilateral Diffusion ASX is superior to Original ASX for the experimentally defined objective of source-independent intrinsic differential diffusion.** The enhancement removes the original random-source sparse-avalanche regime, removes the intermediate design's positional collapse, and makes strong diffusion less dependent on compression and APPEND. It does so at the cost of additional processing and runtime; internal implementation costs are intentionally not decomposed in this public document.

No stronger security conclusion is drawn. Entropy, avalanche, bit balance, correlation and propagation measurements are necessary characterization tools but do not establish semantic security, chosen-plaintext resistance, differential/linear cryptanalytic resistance, side-channel resistance or authenticated-encryption properties. The present result should therefore be read as a rigorous experimental diffusion comparison and an architectural improvement, not as a completed cryptographic certification.

Reader Summary

Original ASX uses simple reversible mechanisms—APPEND, SWAP and XOR—inside a segmented repeated-pass architecture. It can create highly scrambled output and performs strongly when preprocessing already supplies a broad differential, but its original form does not independently multiply a sparse random-source perturbation into a global avalanche. The bilateral enhancement adds a proprietary reversible whole-state diffusion stage whose internal construction is intentionally undisclosed. Native measurements show that this stage changes the random-source diffusion regime from sparse to dense and removes the severe positional dependence observed in the intermediate design. For diffusion, the bilateral architecture is therefore the experimentally stronger design. Cryptographic security remains a separate question requiring independent attack analysis.

22 References

References

[1] C. E. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal*, vol. 27, pp. 379-423 and 623-656, 1948.

[2] C. E. Shannon, “Communication Theory of Secrecy Systems,” *Bell System Technical Journal*, vol. 28, no. 4, pp. 656-715, 1949.

[3] A. F. Webster and S. E. Tavares, “On the Design of S-Boxes,” in *Advances in Cryptology - CRYPTO ’85*, Lecture Notes in Computer Science, vol. 218, Springer, 1986, pp. 523-534.

[4] L. E. Bassham III et al., *A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications*, NIST Special Publication 800-22 Rev. 1a, National Institute of Standards and Technology, 2010. doi:10.6028/NIST.SP.800-22r1a.

[5] National Institute of Standards and Technology, *Advanced Encryption Standard (AES)*, FIPS PUB 197, updated 2023. doi:10.6028/NIST.FIPS.197.