

Validating the revised account of Exact Rational Computation Exchange

What the fixed-suite evidence supports, what the original audit required, and what v2 corrects

Corporate author: Haivian Technology

Version: HAIVIAN-DAP-VALIDATION-V2-20261016

Research manuscript examined: *Exact Rational Computation Exchange*, Richard Janutka, Qonduit Limited NZ; DAP-AI-PAPER-V2-20261016

Date: 16 October 2026

Publication status: Published in the Haivian Validation Journal; public release authorized 16 October 2026

Interest and independence disclosure: Haivian Technology and Qonduit Limited NZ share common ownership. The research author is Richard Janutka. Haivian's validation is therefore not external institutional review. The scientific work used procedurally separated Astra/high AI contexts within a shared project/provider environment, including independent initial assessments, direct challenge and a distinct final reviewer. Separate AI contexts are not external human peer review, an independent institution or a guarantee against correlated errors. The corporate author remains responsible for this account and its disclosures.

Canonical article: <https://haivian.com/validation-journal/validating-exact-rational-computation-exchange/>

Companion research article: <https://haivian.com/research-journal/exact-rational-computation-exchange/>

Rights: Copyright © 2026 Haivian Technology. Public article publication authorized; underlying code, data and private evidence are not released by this notice.

PUBLISHED ARTICLE

Canonical: <https://haivian.com/validation-journal/validating-exact-rational-computation-exchange/>

Companion research article: <https://haivian.com/research-journal/exact-rational-computation-exchange/>

Abstract

This validation manuscript distinguishes an audit of the original v1 paper from checking its revised v2 account. The fresh v1 audit supported the recorded byte arithmetic and independently matched all 348 archived result digests, while requiring material corrections to endpoint classification, verification provenance, oracle independence, domain specification and workload interpretation. The unchanged suite contains 24 scenarios in eight families, two arms and three deterministic repetitions. Candidate recorded frames total 86,490 bytes versus 100,800 for the baseline (-14.1964%); repeated context instead costs 7.6579% more. The arithmetic audit's 348 matching digests comprise 336 output successes and 12 correct first-false replies across 58 distinct task keys. Historical wire frames are absent, fault timing is incomplete, and the scored coefficient structure is unusually favorable to shared-denominator encoding. The v2 research manuscript explicitly incorporates these limits and corrects factual references. The separate correction-check report identifies the exact revised source bytes and determines closure of the listed corrections. Such closure validates the revised reporting at its stated scope; it does not create new experimental evidence or transfer blanket scientific credit from v1 to v2.

1. Versions, evidence and review scope

The submitted v1 paper, dated 12 October 2026, is preserved in exact-rational-computation-exchange-v1. The fresh audit completed on 16 October 2026 is preserved separately in audits/fresh-paper-review-20261016. Its conclusion was that the bounded numerical result survived scrutiny but the manuscript required material revision. The audit did not approve the original prose without qualification.

This article concerns the new research source **DAP-AI-PAPER-V2-20261016** and this distinct validation source **HAIVIAN-DAP-VALIDATION-V2-20261016**, both in papers/exact-rational-computation-exchange-v2. The

accompanying `VERSION_MANIFEST.json` supplies SHA-256 and byte lengths for their exact final files. `CORRECTION_CHECK_REPORT.md` separately identifies the sources actually checked. A version label alone is not evidence that later bytes were reviewed.

The original experiment remains `MXL-L1-SCORE-001`. Its source, raw observations, original reports and the completed fresh audit are unchanged. The revision adds scientific clarification and separately reviewed correction closure, not another scored invocation, a repaired experimental outcome or a new protocol implementation. The PDF/DOCX typesetting in this publication package has its own rendering inspection; source review alone does not certify the rendering.

2. Findings established by the fresh audit of v1

The audit combined complementary mathematics, study/evidence and primary-reference reviews. The first two reviewers assessed their questions independently before directly challenging decisive points. A distinct final reviewer inspected the synthesis, reran the fresh checkers and required a narrower historical-network claim. This division of work is recorded; it does not mean every reviewer repeated every task.

The record-only checker independently reproduced 24 scenarios, eight families, 144 scenario-arm-repetition cells, 48 scenario-arm cells and 378 raw rows. It found 348 recorded harness-PASS rows, 30 recorded typed detections and 18 retained permitted clean recoveries. All 126 deterministic keys had matching recorded request/response/result identities over three repetitions. These repetitions and subexchanges are dependent, not independent samples from a task population.

Recorded simulated frame bytes	Candidate	Baseline	Candidate difference
Complete fixed suite	86,490	100,800	-14.1964%
Repeated-context family	37,578	34,905	+7.6579%
Largest repeated-context case, R03	27,651	21,270	+30%

The byte figures include dirty attempts and declared recoveries. Candidate harness-PASS components sum to 81,159 bytes; dirty rows contribute 5,331 more. Adapter overhead is counted in the totals, although dirty records lack its component split. Inspected source counts the declared request/response frames. A valid intermediate response used to construct its corrupted replacement is internal preparation under the declared fault model, not a second modeled send. The inspected measurement path generates local byte strings and contains no network transport; these figures are not measurements of network traffic.

A separately authored mathematical checker transcribed the frozen input-generation specification and derived expected answers with explicit integer-pair arithmetic. It did not import the original Fraction-based oracle or scored harness. All 348 expected-result digests matched the log: 336 output responses and 12 correct first-false responses, corresponding to 58 scenario/revision task keys across two arms and three repeats. S03 first fails at index 3 and O03 at index 12. O03's `false_last` label does not establish that only its last obligation, index 15, fails.

This result adds independent arithmetic evidence beyond the historical review's flag counting and hash comparisons. It remains conditional on the specified input generation and faithful recording of result digests. Historical inputs and wire payloads were not independently captured by the fresh checker, and absent frames cannot be re-decoded.

Fresh post hoc controls examined 134 small contracts, 1,027 signed-integer roundtrips and ten malformed/canonicity cases without unexpected mismatch. They are limited new correctness checks, not additional scored repetitions. Two intentionally wrong but well-formed responses passed the native structural decoder and failed independent arithmetic: a wrong value and a later false index. The distinction is decisive. Canonicity and request binding do not prove computational truth; the separate arithmetic comparison supplies that check in the measured harness.

3. Why v1 required revision

The numerical agreement did not justify a blanket PASS for the paper. The original text misclassified secondary byte measurements as primary endpoints and implied that the historical read-only review had independently rederived semantics. Its “independent oracle” wording did not explain the shared Fraction library, affine specification and validation dependencies. The mathematical description omitted material admission bounds.

The paper also needed to disclose that every scored non-chain row already shares one coefficient denominator, while chain rows use integers. This is particularly relevant to a shared-denominator representation. Fresh mixed-denominator controls support limited correctness, not a measured mixed-denominator size advantage. Adverse G03 (+5.1254%) and C03 (+10.4866%) results further show why family and aggregate sums do not establish a universal ranking.

Candidate and baseline faults target corresponding failure functions using different malformed objects. For example, truncated native responses and float-placeholder pairs are not equal corruptions or equal preparatory computation. All 30 dirty-attempt rows lack timing despite the planned complete-cost endpoint. Quantized CPU and stage-local allocation measurements cannot establish comparative process performance, and missing child telemetry remains unknown.

Finally, the bibliography misidentified Chiesa's M.Eng. thesis as doctoral and linked a 2008 X.690 record while naming the 2021 edition. The prior-art text conflated logical proof-carrying-code safety evidence with cryptographic verification. Manifest matching was also too easily read as stronger authentication, and an untested release checklist as demonstrated reproducibility sufficiency. These were reporting and specification defects even though the recorded totals were correct.

4. V2 corrections and the separate closure check

The revised research source changes the account in the following concrete ways; the separate reviewer checks those changes rather than inheriting v1 acceptance.

Audit requirement	Revised scientific account
Restore endpoint hierarchy	§4 identifies exactness/response/detection/false acceptance as primary, bytes/recovery/costs as secondary and replay as a control.
Separate evidence layers	§6 distinguishes original instrumented observations, historical record review, fresh integer-pair verification and v2 correction checking.
Explain oracle dependence	§3.3 states shared Fraction/specification and baseline validator/canonicalization dependencies; structural response validation is separate from arithmetic checking.
State admitted semantics	§2 gives dimensions, numeric and framing bounds, normalization, primitive obligations, equality sign, sparse uniqueness, order and duplicate-position rules.
Preserve actual outcomes and coverage	§§4-5 state homogeneous denominators, limited generated inputs/obligations, 336/12 outcome split and O03 index12 rather than a sole final-obligation test.
Clarify bytes, comparator and adverse cases	§§3 and 5 describe the concrete baseline, format-specific faults and simulated-traffic accounting; retain family/scenario reversals and avoid causal-ablation claims.
Report measurement deviations	§5.4 identifies missing fault timing as deviation; retains CPU/allocation limits and unknown child telemetry.
Bound provenance and reproduction	§§4,6,8 attribute freeze/history to internal records, limit hashes to integrity, separate reproduction targets and call the release list a proposed minimum.
Correct references and disclose interests	§3.4 and references correct proof traditions, degree/edition details and precise prior-art roles; §8 names the human author and common ownership, procedural AI separation and pending release rights.

CORRECTION_MAP.md maps all 27 fresh-audit claim dispositions to the revised sections.

CORRECTION_CHECK_REPORT.md is the independent check of the actual revised scientific wording and relevant calculations, with any correction rounds and residual limits retained. The research author is Richard Janutka, Qonduit Limited NZ; Haivian Technology is the corporate author of this validation manuscript. These identities do not turn a common-owner process into external review.

The v2 correction check adds one precise arithmetic clarification to the earlier audit: for the specified O03 input, obligations 12-14 are false and the final obligation, index 15, holds. It confirms the logged first-false index 12

while showing why the `false_last` generator label is misleading. This is a separately identified revision-check derivation, not a new scored observation or a result retroactively attributed to the v1 audit.

The exact reviewed versions must be used together with the report and final manifest. Any later substantive change, including altered captions or numbers during typesetting, needs a corresponding check; clean rendering of v1 gives no automatic rendering or scientific approval to v2.

5. What remains unestablished after correction closure

Reporting corrections cannot recover the historical frame payloads, prove an exhaustive original execution history or remove the fixed workload's structural biases. No unseen-fault rate, general correctness proof, optimal-comparator result, cryptographic soundness, sender authentication, AI-communication advantage, token/cost saving or production reliability is established. The repeated-context reversal remains part of the result.

The proposed public bundle has not been independently exercised. Recalculating archived totals differs from regenerating encodings and behavior. Public release would require suitable schemas/code/fixtures, versions and evidence, together with contributor authority and explicit licenses. An originality assertion is not independently completed IP clearance. This article does not certify ownership or patent freedom to operate.

6. Conclusion

The fresh audit supports the bounded recorded-result arithmetic and size comparison while requiring material correction of v1's account. V2 incorporates those corrections and is subject to its own explicit source-level closure check. The defensible validation is that the checked revised account accurately states supported observations, adverse results and unresolved evidence limits. It is not a new scored result, an external institutional endorsement or blanket scientific certification. Publication of this validation article is authorized; any public evidence release remains separately controlled.

Evidence references

1. Janutka R. *Exact Rational Computation Exchange*. Qonduit Limited NZ. DAP-AI-PAPER-V2-20261016.
<https://haivian.com/research-journal/exact-rational-computation-exchange/>
2. *Fresh substantive audit: Exact Rational Computation Exchange*. 16 October 2026. Internal review record retained by the publisher; not included in this public package.
3. Fresh executable audit evidence and separate rerun receipts. Internal records concern the audit of v1 and unchanged archived data; they are not a v2 experiment and are not included in this public package.
4. Separate v2 correction map, correction-check report, preservation records and version manifest. Internal source-control records retained by the publisher.

Primary literature is cited in the research manuscript. The public package preserves exact final article identities without releasing private audit evidence.