Below is a **formal peer review** suitable for a mathematical / computational topology periodical. I have treated the manuscript as a serious submission and evaluated it on originality, correctness, methodological soundness, and publication readiness.

Peer Review Report

Manuscript: *The Universal Coefficient Theorem as a Computational Primitive: An Empirical Study of Amortised Cohomology with Multiple Coefficient Systems*
Author: Richard Murdoch Montgomery
Venue: *Proceedings · Mathematical Sciences Series* (Scottish Science Society)
Montgomery_UCT_revised

---

1. Summary and Contribution

The manuscript advances a clear and focused thesis:

> The Universal Coefficient Theorem (UCT) should be understood not merely as a structural result in algebraic topology, but as an **algorithmic amortisation primitive** enabling efficient computation of cohomology across multiple coefficient systems.

The paper:

- Formalises two computational strategies:
  - **Direct route**: recomputation per coefficient system
  - **UCT route**: single integral homology + Hom/Ext post-processing
- Introduces a **cost lemma** quantifying asymptotic and practical trade-offs
  Montgomery_UCT_revised
- Provides **empirical benchmarks** demonstrating:
  - Near-constant cost for the UCT route in number of coefficients
  - Linear scaling for the direct route
  - Crossover at ~2–4 coefficient systems
- Positions the result within **persistent cohomology and multi-characteristic analysis**

Core contribution:
A conceptually simple but *useful reframing* of a classical theorem into an explicitly computational paradigm, backed by empirical validation.

---

2. Originality and Significance

Strengths

- The **amortisation perspective** is well articulated and pedagogically effective.
- While implicit in prior computational topology practice, it is:
  - Rarely formalised
  - Even more rarely **quantified empirically**
- The **cost lemma** provides a clean abstraction linking algebraic structure to computational complexity.

Limitations

- The conceptual novelty is **moderate rather than deep**:
  - The UCT's reuse property is well known in computational topology circles.
  - The contribution lies more in **synthesis + quantification** than in theoretical innovation.
- No fundamentally new algorithm is introduced; rather, an **interpretative and practical clarification.**

Verdict on significance

- **Appropriate for a specialised computational topology or applied algebra journal**
- Likely to be **useful and cited**, especially in:
  - Persistent homology
  - Multi-field computations
  - Arithmetic topology experiments

---

3. Mathematical Correctness

General assessment

The manuscript is **mathematically sound.**

- Definitions of chain/cochain complexes and UCT are standard and correct.
- The decomposition

$$H^n(X;G) \cong \mathrm{Hom}(H_n(X;\mathbb{Z}),G) \oplus \mathrm{Ext}^1(H_{n-1}(X;\mathbb{Z}),G)$$

  is correctly stated (non-natural splitting acknowledged) Montgomery_UCT_revised

Cost Lemma (Section 3)

- Formulation is correct at the level of **algorithmic modelling**
- Proof is informal but adequate for the intended audience
- The asymptotic statement:

$$T_{\mathrm{UCT}}(k) = T_{\mathrm{SNF}}(N) + o(k)$$

  is reasonable given bounded invariant factors

Minor issues

- The notation occasionally conflates:
  - k (panel size)
  - k_X (number of invariant factors)
- The omission of bit-complexity considerations in SNF discussion is acceptable but should be clarified.

Conclusion

No substantive mathematical errors detected.

---

4. Methodology and Experimental Design

Strengths

- Clear experimental objective: validate scaling behaviour
- Sensible comparison:
  - Fixed panel vs varying complex size
  - Fixed complex vs varying coefficient count
- Proper verification against known spaces:
  - $S^1, S^2, T^2, \mathbb{RP}^2, L(3,1)$ Montgomery_UCT_revised

Critical Weaknesses

(1) Implementation limitations

- Python + SymPy SNF is **not representative of state-of-the-art**
- No use of:
  - Sparse matrix optimisations
  - Advanced SNF libraries (LinBox, FLINT)
  - Persistent homology toolchains

This weakens external validity.

---

(2) Absence of modern baselines
The direct method does **not include**:

- Clearing optimisation (Chen–Kerber)
- Ripser-like optimisations
- Sparse reduction heuristics

This is explicitly acknowledged, but still a **major limitation.**

---

(3) Scale of experiments

- Maximum ~600 simplices is **very small**
- Modern TDA routinely handles:
  - 10^5–10^7 simplices

Thus:

- Results are **qualitative**, not industrially convincing

---

(4) Lack of statistical robustness

- No variance analysis
- No repeated trials
- No confidence intervals

Conclusion on methodology

The experiments are **directionally correct but insufficiently rigorous** for a high-impact venue.

---

5. Clarity and Exposition

Strengths

- Exceptionally well written
- Elegant and controlled prose
- Strong pedagogical narrative
- Good balance between:
  - Algebraic topology
  - Computational insight

The worked example correcting a previous error is commendably transparent
Montgomery_UCT_revised

Weaknesses

- Occasional rhetorical overstatement:
  - "order of magnitude faster" may overgeneralise beyond tested regime
- Some stylistic density may limit accessibility to non-specialists

6. Relation to Existing Literature

Adequate but incomplete

- Classical references (Hatcher, Spanier, etc.) are appropriately acknowledged
- Computational topology literature is cited conceptually

Missing / underdeveloped

- Explicit comparison with:
  - Persistent cohomology implementations
  - Multi-field homology studies
- No benchmarking against:
  - GUDHI
  - Ripser
  - Dionysus

---

7. Recommendations for Improvement

Essential revisions

1. **Benchmark against optimised implementations**
   - Include at least one modern library baseline
2. **Scale up experiments**
   - Larger complexes ($\geq 10^4$ simplices)
3. **Clarify complexity model**
   - Distinguish:
     - Arithmetic complexity
     - Bit complexity (for SNF)
4. **Refine cost lemma discussion**
   - Explicit dependence on:
     - torsion structure
     - sparsity

Desirable enhancements

- Include **runtime tables** (not only plots)
- Provide **code repository**
- Extend to **persistent setting (even preliminary)**

---

8. Final Verdict

Recommendation: Minor Revision (borderline accept)

Rationale:

- Conceptually sound and clearly written
- Empirically suggestive and useful
- However:
  - Experimental depth is limited
  - Comparison is not fully fair or modern

---

9. Editorial Summary

This manuscript does not revolutionise algebraic topology, but it does something arguably more valuable:

> It converts a classical theorem into a **practically actionable computational principle**, with clear implications for modern topological data analysis.

With modest strengthening of the experimental section, it should be a **worthy and citable contribution** to computational topology.