

Referee Report

CasimirRFM: A Mathematica package for Riemann-flat compactifications with Casimir energies

arXiv:2607.26140

Synopsis

The manuscript presents **CasimirRFM**, a Wolfram Language package for compactifications of higher-dimensional field theories on compact flat manifolds represented as freely acting quotients T^k/Γ . Its public functions cover four linked tasks: construction and manipulation of affine finite groups; determination of invariant metrics, real cohomology, and spin structures; evaluation of one-loop Casimir potentials and local energy densities through Ewald summation; and computation of little-group characters for gravitons, forms, spinors, and Rarita–Schwinger fields. Section 2 develops the mathematical conventions and documents the functions. Section 3 applies them to Type IIB supergravity on a freely acting $T^6/\mathbb{Z}_8$ quotient, computing invariant moduli and spin structures, representation traces, a one-dimensional slice of the potential, and local “Casimir-brane” profiles.

The subject is worthwhile. A reliable package joining flat-manifold data, spin structures, representation characters, and rapidly convergent Casimir sums would be useful to the compactification community. The manuscript is also generally readable, and the worked Type IIB notebook is a good choice for illustrating the intended workflow. The main formal ingredients are not new: the Casimir reduction and much of the compactification framework are taken from Ref. [1], while standard results on flat manifolds, spin structures, finite-order integral matrices, characters, and Ewald summation are drawn from Refs. [2–22]. The principal new contribution is therefore the implementation and its integration into a usable package. For such a paper, correctness, versioned reproducibility, and broad numerical validation are the central criteria.

Overall assessment and recommendation

In its present form, the paper does not yet establish that the package produces the claimed physical quantities. I find several substantive inconsistencies between the equations and the released implementation, together with problems in the spin-lift group law, the overall Casimir sign, the handling of general twists, the localization of Casimir branes, and the claimed truncation accuracy. Some of these issues can change the signs and magnitudes of the Section 3 results, not merely their presentation.

Recommendation: major revision; the manuscript is not suitable for publication in JHEP or PRD in its current form. I would be willing to reconsider a substantially revised version after the authors correct the group-theoretic and numerical issues, recompute the example, and provide a versioned test suite and benchmark evidence. Because the software is the paper’s main original result, these checks are prerequisites rather than optional additions.

Major comments

M1. The spin lift is not specified as a representation of the space group.

Equation (2.4) stores an affine isometry as $\{\mathbf{D}_\gamma, \mathbf{b}_\gamma, \sigma_\gamma\}$, with $\sigma_\gamma \in \{0, \frac{1}{2}\}$ described as the sign of a lift. A choice of sign for each rotation is not by itself a lift of the group: after choosing a section $s : SO(k) \rightarrow Spin(k)$, multiplication obeys

$$s(D_1)s(D_2) = \epsilon(D_1, D_2)s(D_1D_2), \quad \epsilon(D_1, D_2) \in \{\pm 1\}.$$

The cocycle ϵ cannot in general be discarded. In particular, the p -th power of a reference lift of an order- p rotation need not be $+1$. A half-turn in one two-plane has a lift whose square is -1 , even though the vector rotation has order two. This is precisely the elementary case relevant to orientable screw manifolds.

Equation (2.20) instead assumes $(\pm \mathbf{D}_\gamma)^p = e^{2\pi i p \sigma_\gamma} \mathbf{1}$. It is missing the intrinsic central sign of $\mathbf{D}_\gamma^p$. Likewise, simply adding the σ ’s modulo one when composing elements does not implement the spin double cover. This affects both **GetGroup** and all fermion traces. With the principal-angle representatives used in Section 2.5, powers of the $\mathbb{Z}_8$ generator in Section 3 cross spin-cover branches; the displayed fermion trace 32 for the squared generator therefore requires rederivation from an explicit, consistent lift.

The authors should represent actual $Spin(k)$ elements, or implement and document the corresponding $\mathbb{Z}_2$

cocycle. They should verify every defining relation of the Bieberbach group in the spin representation and add tests for a half-turn space, a quarter-turn space, and the $\mathbb{Z}_8$ example. The corrected lift must then be used consistently in `GetGroup`, `SpinStructures`, `TraceSpinor`, and `TraceRS`.

M2. The stated and implemented spin-structure conditions do not agree.

Even within the simplified convention of Equation (2.20), the stated condition is

$$p(\sigma_\gamma - \mathbf{h} \cdot \mathbf{b}_\gamma) \in \mathbb{Z}.$$

The public implementation tests whether this quantity is an *even* integer. These are different conditions. For example, a pure half-period translation with $p = 2$, periodic boundary conditions, and $\sigma_\gamma = \frac{1}{2}$ satisfies the printed congruence but fails the even-integer test. More fundamentally, once the missing central sign in Comment M1 is restored, Equation (2.20) itself must be replaced by the full lifted space-group relation.

The revised paper should derive the condition from the action on a spinor, including the torus translation character and the chosen lift of the rotational part. The package should reject inconsistent generator data and return diagnostics identifying the failed group relation. The four spin structures reported in Section 3 should then be recomputed rather than assumed unchanged.

M3. The overall sign of the Casimir energy is inconsistent between Equations (2.21)–(2.22) and the package.

Equations (2.21) and (2.22) attach an overall minus sign to the bosonic lattice sum. The released `CasimirPotential` and `CasimirEnergyDensity` implement the same positive prefactor but do not implement this overall minus sign; fermions are subtracted separately. For the simplest check—a periodic massless scalar on a circle—this reverses the standard Casimir energy. Consequently, the positive numerical values and the claimed maximum at $c_2 = 0$ in Section 3 may instead have the opposite physical sign and interpretation.

The authors should write Equations (2.21)–(2.22) explicitly as a supertrace, including $(-1)^F$, and state the effective-action convention that fixes the remaining overall sign. They should add analytic unit tests for a scalar and a fermion on S^1 , a rectangular torus, and a supersymmetric spectrum with periodic boundary conditions. The implementation, documentation, notebooks, and every Section 3 plot and numerical value must use one convention.

M4. The advertised support for general twisted boundary conditions is not currently correct.

For general real $\mathbf{h}$ and shift $\mathbf{c}$, the lattice sum in Equation (2.24) is complex. The implementation accumulates real and imaginary parts but then adds them as ordinary real numbers rather than returning $\text{Re } S + i \text{Im } S$. This happens separately in the direct and reciprocal sums. The result is harmless only in special cases, such as the $\mathbb{Z}_2$ phases for which the imaginary parts vanish by symmetry. It is incompatible with the repeated claim that `Ewald` accepts general phase vectors.

The complex arithmetic should be corrected, and the authors should test non- $\mathbb{Z}_2$ twists against direct high-precision sums in one and two dimensions. If only real $\mathbb{Z}_2$ spin structures are intended, the API and all claims of general twisted boundary conditions should instead be restricted explicitly.

M5. The cutoff analysis does not justify the stated accuracy parameter.

There are both algebraic and conceptual problems in the discussion following Equation (2.24).

- In Equation (2.32), the incomplete gamma function inside the integral must depend on the integration variable, $\Gamma(k/2 - s, \pi^2 \rho / \alpha)$, rather than on the fixed lower limit. After integration, the first numerator term is proportional to α^s , not α^2 . The implementation uses α^s , so the published derivation and the code disagree.
- The inequalities preceding Equations (2.29) and (2.31) are written for complex phased sums without absolute values. An ordering relation “ $< \varepsilon/2$ ” is not defined for a complex number.
- Replacing a discrete lattice tail by a continuum volume estimate is an approximation, not an upper bound. It can be poor for anisotropic metrics, shifted lattices, short dual vectors, and modest cutoff

radii. A fixed fallback cutoff when root finding fails is also unrelated to the requested ε .

Thus $\varepsilon = 10^{-p}$ has not been shown to target p reliable digits. The authors should either derive rigorous majorants for both tails or rename the option as a heuristic tolerance and quantify its behavior. A validation table should cover dimensions 1 through 8, several values of s , shifted and twisted sums, highly anisotropic positive-definite metrics, and comparison with direct arbitrary-precision summation. Results should be shown as ε , α , and working precision are varied.

M6. CasimirBranes omits the affine shift that determines the location.

Section 2.4 defines a Casimir brane by $\mathbf{D}_\gamma \mathbf{z} = \mathbf{z}$, and the corresponding function uses only $\mathbf{I} - \mathbf{D}_\gamma$. The local density in Equation (2.21), however, depends on

$$\mathbf{c}_\gamma(\mathbf{z}) = (\mathbf{I} - \mathbf{D}_\gamma)\mathbf{z} - \mathbf{b}_\gamma \pmod{\mathbb{Z}^k}.$$

The transverse localization loci therefore depend on the transverse projection of $\mathbf{b}_\gamma$. A freely acting element has no full fixed locus precisely because its invariant component of $\mathbf{b}_\gamma$ is nonzero; the effective wrapped locus should be defined after separating invariant and transverse components, not by discarding $\mathbf{b}_\gamma$. The Section 3 example hides the problem because its shift lies entirely along an invariant direction.

The authors should give a coordinate-invariant definition of these effective loci, solve the affine congruence with the correct right-hand side (for example, using Smith normal form), and test covariance under a change of origin. The name “brane” should also be qualified: this is a contribution to a one-loop image sum, not an independently introduced localized source.

M7. The reduction to the invariant lattice needs a derivation and integer-lattice tests.

Equation (2.22) depends on the projected lattice Ξ_γ , its covolume, the induced metric, the projected shift, and the dual phase. The manuscript states the result but does not explain how the implementation constructs a primitive basis of $\ker(\mathbf{I} - \mathbf{D}_\gamma) \cap \mathbb{Z}^k$. A rational nullspace basis is not in general a primitive integer-lattice basis, and using unnormalized left and right nullspaces as if they were mutually dual can change the measure, shift, and metric in the reduced sum.

Please derive the exact map from the original integral in Equation (2.21) to Equation (2.22), including all determinant factors. The implementation should use a primitive integer basis obtained by Smith or Hermite normal form and its proper dual. For several non-diagonal finite-order matrices, compare `ReducedLatticeSum` with a direct numerical integration and unreduced image sum. This is especially important because the claimed efficiency comes from this reduction.

M8. Several auxiliary functions require correctness fixes and regression tests.

Two examples affect physical outputs. First, extracting the moduli-space metric from Equation (2.15) by taking coefficients of $\delta\phi_i\delta\phi_j$ overcounts off-diagonal entries by a factor of two: the coefficient of a mixed monomial is already $2g_{ij}$. The robust construction is one half of the Hessian with respect to the formal differentials. Second, the vector-step branch of `RFMHessian` does not recognize a numerical vector with the advertised interface. These should be covered by tests that reconstruct the quadratic form from the returned metric and compare finite differences with analytic derivatives.

More broadly, the package paper needs a versioned, automated test suite. At a minimum it should test affine group closure and freeness, all lifted group relations, invariant metrics, Betti numbers for standard flat manifolds, known representation dimensions and characters, analytic lattice sums, boson–fermion cancellations, basis-change covariance, and the complete Section 3 notebook. The exact package version should be archived with a DOI or immutable commit identifier and included with the paper. A moving repository branch is not sufficient for reproducible journal results.

M9. The scope and novelty of the finite-order classification should be stated more carefully.

Section 2.2 correctly notes that a rational conjugacy class can split into multiple integral conjugacy classes, but then says that the block-diagonal cyclotomic representatives describe the distinct lattice actions. `FiniteOrderMatrices` actually returns representatives only up to $GL(k, \mathbb{Q})$ conjugacy. It therefore does not classify all $GL(k, \mathbb{Z})$ actions and cannot by itself enumerate distinct cyclic flat manifolds. The claim

should be narrowed, or the missing integral classes should be implemented.

Since most formal results are inherited from Ref. [1] and standard literature, the revised manuscript should also compare the package’s functionality with GAP, CARAT, and existing Ewald/lattice-sum tools, explaining precisely which steps are new, which are automated combinations of known algorithms, and where the package offers a demonstrated performance or usability advantage.

Minor comments

- m1.** In the prose before Equation (2.6), the equivalence transformation is missing a factor of the shift vector. It should read $\mathbf{b}_\gamma \mapsto \mathbf{D}_\gamma^\dagger \mathbf{b}_\gamma + (\mathbf{I} - \mathbf{D}_\gamma)\mathbf{u}$, subject to the precise equivalence relation intended.
- m2.** The bottom-right entry of the companion matrix in Equation (2.13) should be $-a_{d-1}$, not $+a_{d-1}$. The package definition appears to use the standard minus sign.
- m3.** Distinguish consistently between the finite affine group acting on the torus and its rotational image (the point or holonomy group). Calling both objects Γ obscures possible pure translations and is particularly unhelpful in the spin-lift discussion.
- m4.** In Section 2.5, a generic element should be written as conjugate to $\exp(i \sum_j \theta_j H_j)$. The expression involving a single root α and E_α^3 is not a general Cartan parametrization, and the angles are not literally the eigenvalues of the vector matrix.
- m5.** Equation (2.33) is a subgroup/branching statement and would be clearer as $SO(D-2) \supset SO(k) \times SO(d-2)$, together with the restriction $k \leq D-2$. State separately what the package does when this condition fails.
- m6.** Define G whenever it denotes $\det \mathbf{G}$, especially in Equations (2.21), (2.24), and (2.29)–(2.32). At present the same symbol is used for a matrix and its determinant.
- m7.** State the mathematical domain supported by `Ewald`, including the conditions on s , treatment of the pole at $s = k/2$, the excluded zero-distance term, and whether analytic continuation is implemented.
- m8.** The table description of `RFMGrad` suggests a fixed vector step $10^{-4}\mathbf{1}$, whereas the implementation uses a relative step proportional to a nonzero coordinate and an absolute step only at zero. The documentation and function signature should agree. Both gradient and Hessian routines should report a convergence estimate rather than a single unqualified finite-difference value.
- m9.** In the Type IIB discussion, label the two gravitini and two dilatini consistently, state the Majorana–Weyl counting convention, and explain why discarding the chirality label is safe for every element used. The argument following the spectrum definition should be tied explicitly to the existence of an invariant vector for each freely acting holonomy element.
- m10.** Add a compact performance table. The paper repeatedly emphasizes efficiency and parallelization but gives no timings, memory use, scaling with dimension or tolerance, or comparison with a direct box sum. Such data are important for judging the practical significance of the package.
- m11.** Correct small presentation issues, including “supplementary” in Section 3, inconsistent punctuation in function tables, and several overlong function-description sentences. A final copy edit would make the documentation substantially easier to consult.

Conclusion

The package addresses a useful problem and could become a valuable research tool. At present, however, the spin-lift construction and several numerical and normalization details are not reliable enough to support the physical claims or the Type IIB example. A corrected implementation, an immutable release, and a substantial analytic and numerical validation suite would materially strengthen both the paper’s correctness and its significance. I therefore recommend major revision and a new round of review.