The Alphabet of Orbifold Boundary Conditions

A boundary condition is a representation;
the gauge algebra chooses the type, the global form the admissible lifts

Part IX — Parts I to VIII each begin by fixing a boundary condition and work downstream of it; this one asks what the choice is, on the orbifold of Parts VI to VIII

Carles Marín*

September 2, 2026

Abstract

An orbifold boundary condition in gauge theory is an action on the gauge fields, so what it defines is a homomorphism into the automorphisms of the gauge algebra; in the inner, zero-flux sector considered here it admits a linear lift and becomes a representation of the space group $\Gamma = \Lambda \rtimes \mathbb{Z}_m$ of the orbifold. After the lift, a boundary condition is a representation — and that single identification fixes what the classification is made of. The labels are the irreducible representations of Γ induced from characters of the translation lattice Λ whose stabiliser fixes only finitely many characters — the rest being the Wilson lines; the weight of a label is its dimension, which is the index of that stabiliser; and how many labels of each weight there are is given in closed form by a Möbius inversion over $|\det(I - A^k)|$, the number of fixed points of the k-th power of the rotation. That formula holds in any rank, not only the two the programme uses; and under the hypothesis that makes the fixed sets finite — $\det(I - A^k) \neq 0$ for $1 \leq k < m$, which is exactly the condition that every eigenvalue of A have order m — those determinants have themselves a closed arithmetic form, so the alphabet depends on the order of the rotation and the rank of the lattice and on nothing else. Latimer–MacDuffee then counts how many rotations that confuses — exactly $h(\mathbb{Q}(\zeta_m))$ of them — so the alphabet determines the rotation in every rank below 22 and not in rank 22, where three inequivalent order-23 rotations, on two orbifolds, carry one alphabet. The gauge algebra enters by choosing the type of the defining space: over $SO(N)$ a label weighs its real dimension and over $Sp(N)$ its quaternionic one, decided by the Frobenius–Schur type of the corresponding complex irreducible; the global form of the group enters only through which lifts are admissible. A boundary condition of rank N is then a multiset of labels of total weight N, and its class is the tuple of local data at the fixed points — the unbroken gauge group at each brane together with its embedding. Under the hypothesis that this local datum is a complete invariant — which we prove for G compact and simply connected on a sphere with cone points, and check a second way on each $T^2/\mathbb{Z}_m$ by counting the components of the published equivalence relations against the fibres of the marginal map — one argument reproduces the twelve closed-form counts published for this programme, gives the class count of $S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$, the orbifold of published models, and supplies the third type: the $Sp(N)$ alphabets, counts and generating functions on $T^2/\mathbb{Z}_3$, $T^2/\mathbb{Z}_4$ and $T^2/\mathbb{Z}_6$, with all twenty-one labels built explicitly as unitary symplectic matrices. The classification is finite for a reason that is also its scope: characters of trivial stabiliser give continuous families of irreducibles, and those are the Wilson lines, so what is classified is the sector of vanishing Wilson line and vanishing 't Hooft flux. The product orbifold whose count opens the paper is the one on which Parts VI to VIII of this series work, and the question this part asks is the one they all take as given.

*Independent researcher, karlesmarin@gmail.com. The computations were carried out and cross-checked with Claude (Anthropic) as an AI research assistant against a common machine-verifiable ground truth; every number in this paper regenerates from the ancillary scripts, each of which writes its own receipt.

The paper on one card

The identification. An orbifold boundary condition is an action on the gauge fields; in the inner, zero-flux sector of this paper it lifts, and *after the lift* it is a representation of the space group $\Gamma = \Lambda \rtimes \mathbb{Z}_m$. Its labels are the irreducibles of Γ induced from characters of Λ whose stabiliser fixes only finitely many characters — the rest are the Wilson lines; the weight of a label is its dimension, and how many labels of each weight there are is a Möbius count over the fixed-point numbers $|\det(I - A^k)|$. The gauge algebra enters by choosing whether that space is complex, real or quaternionic; its global form enters only through which lifts are admissible.

What follows from it.

- the twelve closed-form counts published for this programme, from one argument, with no free parameter (§4);
- the degree of each of them, read off the cone signature by (7), with no series expanded;
- the class count of $S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$ — the orbifold of Parts VI to VIII — which is 2809 at rank 4, with the gauge transformation written down (§5);
- the third type, $Sp(N)$: alphabets, counts, closed forms and all twenty-one labels built as matrices (§7);
- and a preflight that returns all of it from the rotation matrix and the gauge group alone (§8);
- the alphabet of $\mathbb{Z}^r \rtimes_A \mathbb{Z}_m$ in any rank (Theorem 1), and exactly when it determines the rotation up to conjugacy (§9).

What is not claimed. The completeness of the local datum is proved here in the generality the counts need, and *no priority is claimed for it* (Appendix A); the trace conservation laws, the local gauge groups, the classification of $\mathbb{Z}_N$ embeddings and the identification of the surviving class with the 't Hooft flux are all inherited; and so is *everything in §9 except one line* — the count of rotations an alphabet confuses is Latimer–MacDuffee and Charlap, not ours. Table 6 says which is which, line by line.

Scope. What is classified is the sector of vanishing Wilson line and vanishing 't Hooft flux. Both exclusions are consequences of the construction, not conventions imposed on it.

1 A question the literature does not answer

Gauge-Higgs models are built on a fixed orbifold, and the first thing one wants to know about the orbifold is how many inequivalent boundary conditions it carries at rank N: that number is the size of the model space one is choosing from. For $S^1/\mathbb{Z}_2$ the answer has been $(N+1)^2$ for twenty years [1], and for each $T^2/\mathbb{Z}_m$ there is now a closed form [2, 5, 6, 7]. For $S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$, which is the orbifold of the $SU(7)$ grand gauge-Higgs unification of Komori and Maru [9] among others, this paper obtains it from the same argument as all the rest (§5).

That orbifold is not chosen for the example. It is the one Parts VI, VII and VIII of this series work on: the proton-decay obstruction of [11], the certified ceiling on the compactification scale of [12], and the certified gap of [13] are all statements about the $SU(7)$ model of [9] on $S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$. Each of them — and Parts I to V before them — begins by fixing a boundary condition, taken from a published list, and works downstream of it: the bulk content, the potential, the spectrum, the scale. What none of them derives is how many inequivalent boundary conditions that orbifold has, because each takes its own as given; and Part V [10], whose subject is what the Higgs potential cannot distinguish, is a statement about the continuous directions that this classification deliberately excludes. This part turns round and asks about the input rather than the output.

That is a small question with a large one inside it. Every count above was obtained the same way: write down the boundary-condition matrices the orbifold admits, find the moves that relate them, and count what is left. The matrices are different in each case — signs on $S^1/\mathbb{Z}_2$, blocks of size up to three on $T^2/\mathbb{Z}_6$ — and so are the moves, so each orbifold has been a separate piece

of work, and a new one is a new piece of work. Before asking how many boundary conditions a product orbifold has, one has to ask what a boundary condition on it is *made of*.

This paper answers the second question, and the first then costs nothing. The pieces are not a list to be found per orbifold: they are irreducible representations of the space group — the discrete ones, in the sense §2 makes precise — their weights are their dimensions, and how many there are of each weight is a count of fixed points. From that one identification the twelve published counts follow by a single argument (§4), the opening question is answered — the number is 2809 at rank 4 (§5) — and the case the programme had not reached, because it is a third type of defining space rather than a third orbifold, the symplectic one, follows in a line (§7).

Established here: the question, and why it is a question about an alphabet. *Left open:* everything below.

2 Where the labels come from

What a boundary condition assigns to the space group $\Gamma = \Lambda \rtimes \mathbb{Z}_m$ — Λ the translation lattice, $\mathbb{Z}_m$ generated by the rotation ρ acting on Λ by an integer matrix A of order m — is an action on the gauge fields, so what it defines is a homomorphism into the automorphisms of the gauge algebra, and not into G itself. In the sector considered here, where the Wilson line and the 't Hooft flux both vanish (§10, Appendix A), that action is inner and admits a linear lift to the defining space: the twists are unitary matrices, determined up to a central phase, and a boundary condition becomes a finite-dimensional unitary representation of Γ once such a lift is fixed.

Taking the lift is not a formality, and the smallest case shows why: on $S^1/\mathbb{Z}_2$ the two twists are sign matrices, and a sign matrix of determinant -1 is not in $SU(N)$. Requiring the lift to land in $SU(N)$ would force a parity condition on the multiplicities and would not return the published count $(N+1)^2$; taking it in $U(N)$, with the central phase irrelevant because it acts trivially on the fields, does. The same distinction returns for the orthogonal groups in §3, where it is decided by measurement.

With the lift fixed, decomposing the defining representation into irreducibles of Γ writes any boundary condition as a multiset of irreducibles whose dimensions sum to N; those are the labels, and the weight of a label is its dimension. The alphabet is therefore a set of irreducible representations of Γ — not all of them, as §10 makes precise: the irreducibles induced from characters of trivial stabiliser come in continuous families and are the Wilson lines. What is used here is the finite set

$$\mathcal{A}(\Gamma) \;=\; \left\{ \mathrm{Ind}_{\Lambda\rtimes H}^{\Gamma}(\chi\otimes\sigma) \;:\; \chi\in\widehat{\Lambda},\; H=\mathrm{Stab}(\chi),\; \#\mathrm{Fix}(H)<\infty,\; \sigma\in\mathrm{Irr}(H) \right\},$$

the *discrete*, zero-Wilson-line alphabet, and the question is what it is. For the rotation groups of this section the condition reads simply "non-trivial stabiliser": under the hypothesis of (2) the fixed set of every ρ^k with $1 \le k < m$ is finite. It is stated in the stronger form because the point group of §5 contains reflections, whose fixed characters form a continuum; a character fixed by a reflection and by nothing else moves in a family, and its induced irreducibles are Wilson lines like the others.

Clifford theory answers it in the form we need. Let χ be a character of Λ and let $\mathbb{Z}_d \le \mathbb{Z}_m$ be its stabiliser under the action of the rotation on characters. The extensions of χ to $\Lambda \rtimes \mathbb{Z}_d$ exist and are d in number, since $\mathbb{Z}_d$ is cyclic, and inducing each of them to Γ gives an irreducible of dimension m/d. Every irreducible arises this way, and two arise from the same ρ-orbit of characters. Hence

$$\text{a } \rho\text{-orbit of characters with stabiliser } \mathbb{Z}_d \;\longmapsto\; d \text{ labels of weight } m/d, \tag{1}$$

and the weight of a label is the index $[\mathbb{Z}_m : \mathbb{Z}_d]$ of the stabiliser of the character it comes from.

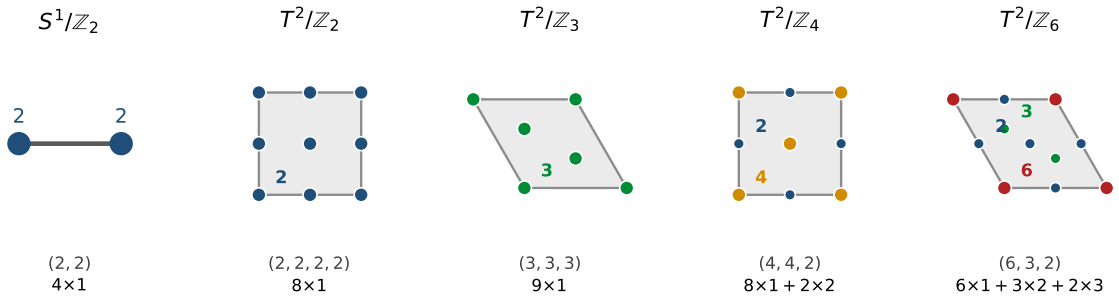


Figure 1: Where the alphabet comes from, drawn. Each panel is a fundamental domain of the translation lattice with every point of non-trivial stabiliser marked, coloured and labelled by the order of that stabiliser; points on an edge or corner are repeated at their lattice translates. Counting them by order is (2), and the cone points are their ρ-orbits, which is the signature in each title. The alphabet in the title is what (1) then returns. The figure is generated by the same code that derives the alphabet, and the number of points of each order is checked against $|\det(I - A^k)|$ before it is drawn.

Counting them is counting fixed points. A character χ is fixed by ρ^k exactly when $(A^k - I)^{\mathsf{T}}\chi$ is integral, so the number of characters fixed by ρ^k is

$$\boxed{\#\{\chi : \rho^k\chi = \chi\} \;=\; \big|\det(I - A^k)\big| \;=\; \#\{\text{fixed points of } \rho^k\} \qquad (I - A^k \text{ nonsingular}),} \tag{2}$$

the last equality because the fixed points of ρ^k on the torus are the solutions of $(A^k - I)x \in \Lambda$, counted the same way. The hypothesis is not decorative: if A^k has eigenvalue 1 the fixed set is not finite and the determinant vanishes, which is not a count of anything. For the rotations of the crystallographic restriction and $1 \le k < m$ it holds.

Möbius inversion then turns "fixed by $\rho^{m/d}$" into "stabiliser exactly $\mathbb{Z}_d$", and carrying it out once gives the alphabet in closed form. Nothing in the argument is two-dimensional — it uses (2), (1) and orbit-stabiliser, none of which knows the rank of Λ — so we state it for $\Lambda = \mathbb{Z}^r$.

Theorem 1. *Let $A \in GL(r, \mathbb{Z})$ have order m and satisfy $\det(I - A^k) \neq 0$ for $1 \le k < m$, and let $\Gamma = \mathbb{Z}^r \rtimes_A \mathbb{Z}_m$. Then the discrete zero-Wilson-line alphabet $\mathcal{A}(\Gamma)$ consists of $M_{m/d}$ labels of weight m/d for each divisor $d \ge 2$ of m, where*

$$\boxed{M_{m/d} \;=\; \frac{d^2}{m} \sum_{d \,|\, e \,|\, m} \mu\!\left(\frac{e}{d}\right) \big|\det\left(I - A^{m/e}\right)\big|, \qquad d \mid m,\; d \ge 2.} \tag{3}$$

Moreover the hypothesis holds if and only if every eigenvalue of A has order exactly m, that is, if and only if the characteristic polynomial of A is Φ_m^s with $r = s\,\varphi(m)$; and then

$$\boxed{\big|\det(I - A^k)\big| \;=\; \Phi_{m/d}(1)^{\,s\,\varphi(m)/\varphi(m/d)}, \qquad d = \gcd(k, m),} \tag{4}$$

with $\Phi_n(1) = p$ when n is a power of a prime p and $\Phi_n(1) = 1$ when n has two or more prime factors. In particular $\mathcal{A}(\Gamma)$ depends only on m and r.

Proof. Write F_d for the number of characters whose stabiliser *contains* $\mathbb{Z}_d$, so $F_d = |\det(I - A^{m/d})|$ by (2), and E_d for those whose stabiliser is exactly $\mathbb{Z}_d$; then $F_d = \sum_{d|e|m} E_e$ and Möbius inversion gives E_d. An orbit with stabiliser $\mathbb{Z}_d$ contains m/d characters and yields d labels of weight m/d by (1), which is (3).

For the second half, A has finite order so its eigenvalues are roots of unity, and A^k has eigenvalue 1 exactly when some eigenvalue has order dividing k; the hypothesis therefore says that no eigenvalue has order $< m$, and since $A^m = I$ every order divides m, so every order is m. The eigenvalues are then the primitive m-th roots of unity, each with the same multiplicity s because the characteristic polynomial has integer coefficients, whence $r = s\varphi(m)$. Raising to the k-th power sends the primitive m-th roots onto the primitive (m/d)-th roots with $d = \gcd(k, m)$, each hit $\varphi(m)/\varphi(m/d)$ times, so $\det(I - A^k) = \big[\prod_\zeta (1-\zeta)\big]^{s\varphi(m)/\varphi(m/d)}$ over the primitive (m/d)-th roots, and that product is $\Phi_{m/d}(1)$. Since (4) names no matrix, neither does (3) through it. □

It returns the five alphabets of Table 1 exactly, and dropping the Möbius factor breaks it on $T^2/\mathbb{Z}_4$ and $T^2/\mathbb{Z}_6$ — and on those two only, since for m prime the correction is invisible. Outside dimension two it is checked against a direct count of character orbits: in rank 3 the involution gives 16×1, in rank 4 the order-3 rotation gives 27×1 and the order-8 one gives $16\times1+4\times2+6\times4$ — a weight-4 label, which no two-dimensional case produces.

What (3) is an instance of, and where it goes. The numbers F_d are the *marks* of the cyclic group on the character lattice, and inverting F to E is the inversion of its table of marks. Written that way the same scheme is formally available for a point group P that is not cyclic — we write P rather than F here, the letter F already being the marks — provided the fixed-character stratum it needs is finite:

$$\text{marks } F_H = \#\{\chi : H\chi = \chi\} \xrightarrow{\text{inverse table of marks}} E_H \xrightarrow{\text{Clifford}} \text{labels of weight } [P:H]. \quad (5)$$

The scheme is not yet a theorem for arbitrary P, and the reason is where a table of marks lives. Marks are defined for a finite P-set, and $\widehat{\Lambda}$ is a torus: the mark of a subgroup is $|\det(I - M)|$ only when that determinant does not vanish for every M in the subgroup, and a reflection has eigenvalue 1, so its fixed characters are a continuum and its mark is not a number. Where the fixed loci that carry the discrete sector are finite, they form a finite P-set and (5) is the ordinary inversion on it; separating that stratum from the positive-dimensional ones is what a general statement would have to do first, and this paper does not do it.

What does *not* obstruct (5) is worth saying, because in the general Clifford setting it would. The extensions of an H-invariant character to $\Lambda \rtimes H$ are counted by the irreducibles of a twisted group algebra, with a class in $H^2(H, U(1))$ deciding the twist. Here the extension $\Lambda \to \Gamma \to F$ *splits* — the space groups of this paper are symmorphic — and for a split extension the twist is trivial: if $\chi(h\lambda) = \chi(\lambda)$ for all $h \in H$ then $\rho(\lambda, h) = \chi(\lambda)\sigma(h)$ is a homomorphism for *every* $\sigma \in \mathrm{Irr}(H)$, so the extensions number $|\mathrm{Irr}(H)|$ on the nose, cyclic or not. A multiplier can appear for a non-symmorphic space group, where the extension does not split; that is a different generalisation from this one.

The orbifold of §5 is the first example, and it is an example rather than a proof: its point group $\mathbb{Z}_2\times\mathbb{Z}_2$ is generated by reflections and is not cyclic, and the finite stratum is visible because the geometry of this particular action is known. The two reflections have singular $I - M$ and contribute no discrete labels. The rotation $-I$ has mark 4; all four of those characters are fixed by the whole of P, so their stabiliser is P and its fixed set is those same four points, which is finite; and each contributes $|\mathrm{Irr}(\mathbb{Z}_2 \times \mathbb{Z}_2)| = 4$ labels of weight $[P : P] = 1$ — sixteen labels of weight one, which is the alphabet §5 uses, obtained there by a different route. What that shows is that the scheme reaches at least one point group that is not a rotation group; it does not show that it reaches all of them.

Two identities hold across the table and are checked rather than assumed. The first is $\sum_i (\dim V_i)^2 = m \cdot \#\{\chi \text{ with non-trivial stabiliser}\}$, which is the orthogonality relation restricted to the labels; on $T^2/\mathbb{Z}_6$ it reads $6 + 12 + 18 = 36 = 6 \cdot 6$. The second is that the derivation is sensitive to its own steps: dropping the Möbius subtraction returns no labels at all on $T^2/\mathbb{Z}_6$ instead of $6 \times 1 + 3 \times 2 + 2 \times 3$, and putting a rotation of order 4 in the $m = 6$ slot returns a

orbifold	$\lvert\det(I - A^k)\rvert$, $k = 1, \ldots, m-1$	alphabet, derived	alphabet, as published
$S^1/\mathbb{Z}_2$	2	4×1	4×1
$T^2/\mathbb{Z}_2$	4	8×1	8×1
$T^2/\mathbb{Z}_3$	3, 3	9×1	9×1
$T^2/\mathbb{Z}_4$	2, 4, 2	$8 \times 1 + 2 \times 2$	$8 \times 1 + 2 \times 2$
$T^2/\mathbb{Z}_6$	1, 3, 4, 3, 1	$6 \times 1 + 3 \times 2 + 2 \times 3$	$6 \times 1 + 3 \times 2 + 2 \times 3$

Table 1: The alphabet, derived from the rotation alone, against the alphabet the literature lists [1, 2, 5, 6]. Written as (multiplicity)×(weight). The row for $S^1/\mathbb{Z}_2$ has Λ of rank one and two independent reflections rather than a rotation, and the same count gives its four labels. The published column is a list of matrices; the derived column is a consequence of A, and no matrix from any paper enters it.

different alphabet. A derivation that gave the right answer with a step removed would not be deriving anything.

Why the alphabet is finite, and what that costs. A character with *trivial* stabiliser also induces an irreducible, of dimension m; but such characters come in continuous families, since the condition defining them is empty, and so do the induced representations. Those families are the Wilson lines. The alphabet is finite exactly because it consists of the characters with non-trivial stabiliser, and therefore what is classified here is the sector of vanishing Wilson line — a scope statement that is usually made by fiat and is here a property of the construction. The complementary statement, the potential on those moduli, is a different problem.

Prior art, and what is not claimed. The local half of this construction — an embedding of the orbifold's rotation into the gauge group, at one fixed point — was classified and tabulated long before the programme this paper reproduces, in another literature. Katsuki, Kawamura, Kobayashi, Ohtsubo and Tanioka [15] classify completely the unbroken gauge groups of $\mathbb{Z}_N$ orbifold models of the $E_8 \times E_8$ heterotic string, for $N = 3, 4, 6, 7, 8, 12$; the companion paper [14], with Ono, classifies the independent embeddings themselves, and Nilles, Ramos-Sánchez and Vaudrevange [16] extend the classification to $SO(32)$. Three things about the relation to this paper are worth stating precisely.

First, their gauge group is fixed and ours is not: what is enumerated there is which subgroups of $E_8 \times E_8$ survive, whereas the alphabet above is a function of the rotation and the choice of G enters only through §3. Second, they impose modular invariance, which is a string constraint with no counterpart here. Third, and most importantly, the object classified there is the embedding at a point; the classification this paper is about is of the *tuple* of local data over all the fixed points, which is §4. What is claimed here is the identification of the alphabet with the induced part of $\mathrm{Irr}(\Lambda \rtimes \mathbb{Z}_m)$, of the weight with the dimension, and of the multiplicities with (2) — not the classification of embeddings, which is theirs.

One point of contact is worth recording because it is not a coincidence of vocabulary. Their classification is stated for *vanishing Wilson lines*, and so is this one, for the reason given above: the characters with trivial stabiliser are the Wilson lines, and dropping them is what makes the alphabet finite. The same restriction was imposed on the same objects, in a different literature, thirty-five years earlier. Their two ways of embedding, by an automorphism of the root lattice and by a shift, are also not the distinction between the labels of weight 1 and those of higher weight in Table 1: they state that a breaking through an automorphism can be equivalently realised through a shift, whereas a label of weight greater than 1 is an irreducible of dimension greater than 1 and no change of basis reduces it.

Established here: the alphabet of every case in this programme, from the fixed-point counts alone. *Left open:* nothing at this level; what remains an input is the orbifold's own rotation, which is its definition.

3 The gauge algebra chooses the type

Changing the gauge group changes which representations of Γ are available, because it changes the field over which the defining space is written. For the unitary groups that space is complex and the alphabet is the one of §2; for $SO(N)$ it is real and for $Sp(N)$ quaternionic. The three compact classical families are not three real forms of one complex group — they are three families, and $Sp(n)$ is a real form of $Sp(2n,\mathbb{C})$ rather than of $SL(n,\mathbb{C})$. What is common to them is the *type* of the defining representation, complex, real or quaternionic, and the passage between the three alphabets is Frobenius–Schur.

For each complex irreducible V consider the space of invariant bilinear forms, $\{B : \rho^{\mathsf{T}} B \rho = B \text{ for both generators}\}$; by Schur its dimension is at most one, and it decides the type: empty means V is not self-conjugate (*complex* type), one symmetric B means *real* type, one antisymmetric B means *quaternionic* type. Writing d for the complex dimension, the irreducible objects over the other two forms and the weight each carries are

type of V	over $SO(N)$ (weight = real dim.)	over $Sp(N)$ (weight = quaternionic dim.)
real	a real form of V, weight d	$V \otimes \mathbb{H}$, weight d
complex	$V \oplus \bar{V}$, weight $2d$	$V \oplus \bar{V}$, weight d
quaternionic	weight $2d$	V itself, weight $d/2$

Nothing else changes: the local datum is still the tuple of eigenvalue multiplicities at the fixed points, and a class is still a fibre of that map.

$O(N)$ or $SO(N)$, decided by the counts. A real irreducible preserves a symmetric form, which places it in $O(N)$; being in $SO(N)$ is the further condition $\det = +1$. That condition is not a count but a character: $\det\rho$ is itself a homomorphism $\Gamma \to \{\pm 1\}$, so a boundary condition lies in $SO(N)$ exactly when the product of the determinant characters of its constituents is trivial — one condition per generator of Γ, and in general several independent ones. How many is the rank of the subgroup of $\mathrm{Hom}(\Gamma,\{\pm1\})$ that the labels generate, and the three orbifolds show all three possibilities. On $T^2/\mathbb{Z}_3$ that subgroup is trivial: every label has determinant $+1$ and the two conventions agree. On $T^2/\mathbb{Z}_6$ the three labels of non-trivial determinant carry the *same* character, the subgroup is $\mathbb{Z}_2$, and there — and only there — the condition really is a single parity. On $T^2/\mathbb{Z}_4$ five labels have non-trivial determinant but they carry three distinct characters, generating $\mathbb{Z}_2 \times \mathbb{Z}_2$, so the condition is two independent ones.

That difference between the two orbifolds is not bookkeeping: the subgroup of $\mathrm{Hom}(\Gamma,\{\pm1\})$ generated by the determinant characters is $\mathbb{Z}_2 \times \mathbb{Z}_2$ on $T^2/\mathbb{Z}_4$ and $\mathbb{Z}_2$ on $T^2/\mathbb{Z}_6$, and the class count is what reads it. Which global form the twists are taken in leaves a finite arithmetic trace on a sequence of integers.

Imposing it changes the counts:

$$\begin{aligned} T^2/\mathbb{Z}_4: &\quad 1,4,14,36,83,168,316 &&\longrightarrow\ 1,1,6,9,28,42,96,\\ T^2/\mathbb{Z}_6: &\quad 1,2,6,12,25,44,77 &&\longrightarrow\ 1,1,4,6,15,22,44, \end{aligned}$$

and the published tables are the left-hand column. So the twists of the orthogonal classification are taken in $O(N)$, which is the same statement as the one made for the unitary case in §2: what is classified is the boundary condition as an automorphism of the algebra, and the determinant of a lift is not part of the datum. We write $SO(N)$ below, as the literature does, with that reading.

So one complex alphabet gives three:

Two consequences are worth isolating. First, the middle column is not an independent classification: applied to it, the argument of §4 returns the published orthogonal counts $1,1,5,5,15,15,34$ on $T^2/\mathbb{Z}_3$, $1,4,14,36,83,168,316$ on $T^2/\mathbb{Z}_4$ and $1,2,6,12,25,44,77$ on $T^2/\mathbb{Z}_6$. Second, the two weight rules are not interchangeable — assigning quaternionic weights where the real ones belong changes all three of those sequences already at $N=1$ — so the type is doing work rather than relabelling.

orbifold	$SU(N)$ (complex)	$SO(N)$ (real)	$Sp(N)$ (quaternionic)
$T^2/\mathbb{Z}_3$	9×1	$1\times 1+4\times 2$	5×1
$T^2/\mathbb{Z}_4$	$8\times 1+2\times 2$	$4\times 1+4\times 2$	$6\times 1+2\times 2$
$T^2/\mathbb{Z}_6$	$6\times 1+3\times 2+2\times 3$	$2\times 1+3\times 2+2\times 3+1\times 4$	$4\times 1+2\times 2+2\times 3$

Table 2: One alphabet, three types, written as (multiplicity)×(weight). The middle column is the alphabet of the published orthogonal classification [7], obtained here from the left one; the right column is §7. The three cases shown are those for which all three columns have been computed.

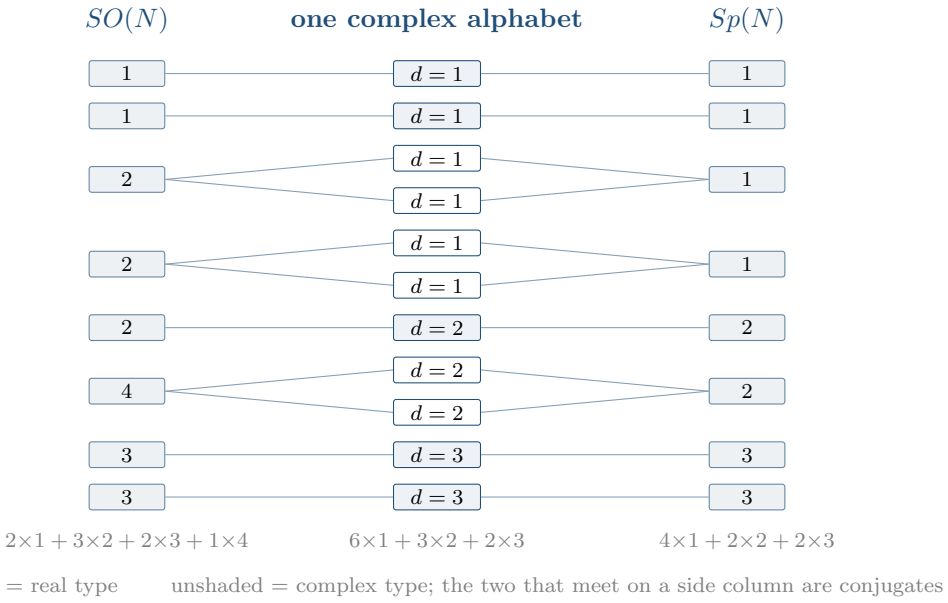


Figure 2: The three types on $T^2/\mathbb{Z}_6$, drawn. The middle column is the complex alphabet: eleven irreducibles, of which five are of real type and six form three conjugate pairs; no label is of quaternionic type, which is computed rather than assumed. Each side column carries the weight of the resulting label. A conjugate pair merges either way, and the two sides differ in what that costs: over $SO(N)$ the merged label weighs twice its complex dimension, over $Sp(N)$ it weighs the same, while a real label doubles into $V\otimes\mathbb{H}$ without changing its quaternionic weight. The three columns are the three rows of the table above, applied to one alphabet.

Where the redundancy comes from. Over the real and quaternionic types the spectrum is closed under complex conjugation, so $\mathrm{mult}(\zeta)=\mathrm{mult}(\bar\zeta)$ for every eigenvalue. A primitive root and its conjugate are then not two independent multiplicities but one, and the space in which the local datum lives shrinks accordingly: at a cone point of order m it carries $m-1$ free multiplicities over $SU(N)$ and $\lfloor m/2\rfloor$ over $SO(N)$. That is a property of the matrices, upstream of any equivalence relation, and it is why $T^2/\mathbb{Z}_6$ has a count of degree 8 over $SU(N)$ and of degree 5 over $SO(N)$ with the same orbifold.

The orthogonal and symplectic alphabets then share the degree but not the periodicity. On the three cases of Table 2 the order of the pole at $x=1$ is the same over $SO(N)$ and over $Sp(N)$ — 4, 6 and 6, in each case $1+\sum_i\lfloor m_i/2\rfloor$ over the cone orders of the orbifold — so the counts have the same degree. But over $Sp(N)$ all three are polynomials in N, while over $SO(N)$ they are quasi-polynomials of period 2, 2 and 4. Which of the two types is taken is invisible in the degree and visible in the residue of N.

Established here: one complex alphabet and three types, with the orthogonal alphabet recovered rather than assumed, and the group in which its twists live settled by the counts. The two $\mathbb{Z}_2$ quotients are not in the table above because for them the three columns coincide: every label of $S^1/\mathbb{Z}_2$ and $T^2/\mathbb{Z}_2$ is a character of order 2 and therefore of real type, so neither realification moves a weight and all three types carry the same alphabet and the same count — which is the

content of the two entries for them in Table 3, and which §8 reports again from the rotation matrix alone. *Left open:* the exceptional groups, where the same argument runs unchanged but the alphabet is not computed here.

4 The rule, and the counts it reproduces

With the alphabet fixed, a classification is a statement about multisets. A boundary condition of rank N is a multiset of labels whose weights sum to N; and the claim of this section is that its class is read off at the fixed points and nowhere else:

> **The rule**
>
> The class of a boundary condition is **the tuple of its local data**. At a cone point of order m the local datum is the multiset of eigenvalues of the local rotation over the N gauge indices — physically, the unbroken gauge group at that brane together with its embedding — and two boundary conditions are equivalent exactly when all their local data agree. Nothing about how the branes are joined survives.

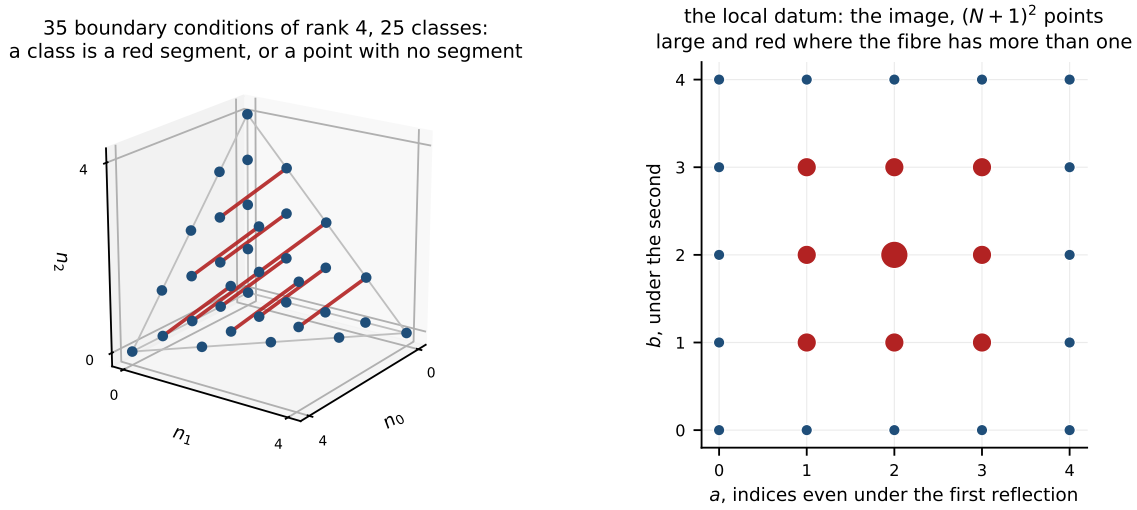


Figure 3: The rule, drawn whole, on the one case small enough to see. On $S^1/\mathbb{Z}_2$ a boundary condition of rank N is a multiset of the four labels, so it is a lattice point of the simplex $n_0 + n_1 + n_2 + n_3 = N$ — and three coordinates suffice, since the fourth is determined. Left: the 35 boundary conditions of rank 4. The local datum is the pair $(a, b) = (n_0 + n_1,\ n_0 + n_2)$, the number of gauge indices even under each of the two reflections; its fibres are the red segments, all of them in the single direction $(+1, -1, -1, +1)$, which is this orbifold's only relation. Right: the image, the $(N+1)^2 = 25$ classes, with the marker growing where the fibre has more than one point. The whole of §4 is this picture in higher dimension: the classes are the fibres, and the published relations are the directions the fibres run in.

The physical reading in the box is a translation and not an analogy: the eigenvalue multiplicities at a fixed point are the block sizes of the subgroup of G left unbroken by the boundary condition there, and the labelling of those blocks by eigenvalue is the embedding of that subgroup [8].

The rule has two halves, proved by different means and stated here once. That equivalent boundary conditions have equal local data is the necessary half: the conserved traces at the fixed points [6] are symmetric functions of exactly these multiplicities, so no equivalence can move them. That boundary conditions with equal local data are equivalent is the sufficient half; it is

Appendix A on a sphere with cone points, and §6 checks it a second and independent way by counting components of the published relations. What follows uses both.

The counts. Applying the rule to the alphabets of §§2–3 is a finite computation at each N, with no free parameter anywhere in it: the orbifold's cone signature fixes the alphabet, the alphabet fixes the local-datum map, and the class count is the number of distinct values that map takes on multisets of weight N. Carried out, it returns every closed-form count published for this programme.

orbifold	G	class count, from the rule	published as
$S^1/\mathbb{Z}_2$	$SU(N)$	$(N+1)^2$	[1]
$T^2/\mathbb{Z}_2$	$SU(N)$	$1, 8, 33, 96, 225, 456, 833, 1408$	[6], eq. (5.3)
$T^2/\mathbb{Z}_3$	$SU(N)$	$1, 9, 45, 163, 477, 1197, 2674$	[6], eq. (5.5)
$T^2/\mathbb{Z}_4$	$SU(N)$	$1, 8, 36, 120, 329, 784, 1680$ *(diagonal labels only)*	[6], eq. (5.7)
$T^2/\mathbb{Z}_4$	$SU(N)$	$1, 8, 38, 136, 403, 1040, 2412, 5136$	[6], §5
$T^2/\mathbb{Z}_6$	$SU(N)$	$\binom{N+5}{5}$ *(diagonal labels only)*	[6], §5
$T^2/\mathbb{Z}_6$	$SU(N)$	$1, 6, 24, 76, 207, 504, 1125$	[6], §5
$S^1/\mathbb{Z}_2$	$SO(N)$	$(N+1)^2$, as over $SU(N)$	[7], note 9
$T^2/\mathbb{Z}_2$	$SO(N)$	$1, 8, 33, 96, 225, 456, 833, 1408$, as over $SU(N)$	[7], §4.1
$T^2/\mathbb{Z}_3$	$SO(N)$	$1, 1, 5, 5, 15, 15, 34, 34, 65, 65, 111$	[7], Table 4
$T^2/\mathbb{Z}_4$	$SO(N)$	$1, 4, 14, 36, 83, 168, 316, 552, 917, 1452$	[7], Table 5
$T^2/\mathbb{Z}_6$	$SO(N)$	$1, 2, 6, 12, 25, 44, 77, 124, 196, 294, 433$	[7], Table 6

Table 3: Twelve counts, one rule. Each row is computed from the alphabet alone and agrees with the published count over the range in which that count is given; the two rows marked *diagonal labels only* are the same computation restricted to the labels of weight 1, which are the diagonal representatives of [3]. The last entry of the bottom row is the $SO(10)$ value.

A count that could have been insensitive. Twelve agreements are worth what the computation risked, so the same computation is run with one cone point deleted at a time. On $T^2/\mathbb{Z}_2$, $T^2/\mathbb{Z}_3$ and $T^2/\mathbb{Z}_4$ deleting any point changes the count — $1, 8, 33, 96, 225$ becomes $1, 8, 27, 64, 125$ on $T^2/\mathbb{Z}_2$, and so on — so the agreements are not an artefact of a map that would have matched whatever it was given. On $T^2/\mathbb{Z}_6$ the answer is different and is a statement about that orbifold: deleting the order-3 or the order-2 point changes nothing, and only the order-6 point matters. Its three cone points are not independent readings; the data at the two lesser ones are determined by the data at the first, which is the same fact as the alphabet of $T^2/\mathbb{Z}_6$ having six labels of weight 1 where $T^2/\mathbb{Z}_2$ has eight.

The degree is read off the signature. A closed form has a degree, and it need not be computed from the series. What fixes it is the rank of the lattice the labels span,

$$\deg \;=\; \operatorname{rank}_{\mathbb{Z}}\langle a_i \rangle - 1, \tag{6}$$

and that rank is bounded by the dimension of the space the local data live in, with equality exactly when the alphabet generates it. That ambient dimension is a count of coordinates — one per eigenvalue class at each cone point, one constraint per point, one for the total rank — so it is

$$\boxed{e_1 \;=\; 1 + \sum_i c(m_i), \qquad c(m) = m - 1 \;\text{ over } SU(N), \qquad c(m) = \lfloor m/2 \rfloor \;\text{ over } SO(N),} \tag{7}$$

over the cone orders m_i of the orbifold. The two rules for c differ for one reason, given in §3: over the real and quaternionic types a primitive root and its conjugate always carry the same multiplicity, so they are one coordinate and not two. Where the alphabet generates, $\operatorname{rank}\langle a_i \rangle = e_1$

and the count is a quasi-polynomial of degree $e_1 - 1$; that is the case for all ten counts of Table 3 whose alphabet is the full one, with $e_1 = 3, 5, 7, 8, 9$ over $SU(N)$ and $e_1 = 3, 5, 4, 6, 6$ over $SO(N)$ — degrees $2, 4, 6, 7, 8$ and $2, 4, 3, 5, 5$ — and no series expanded.

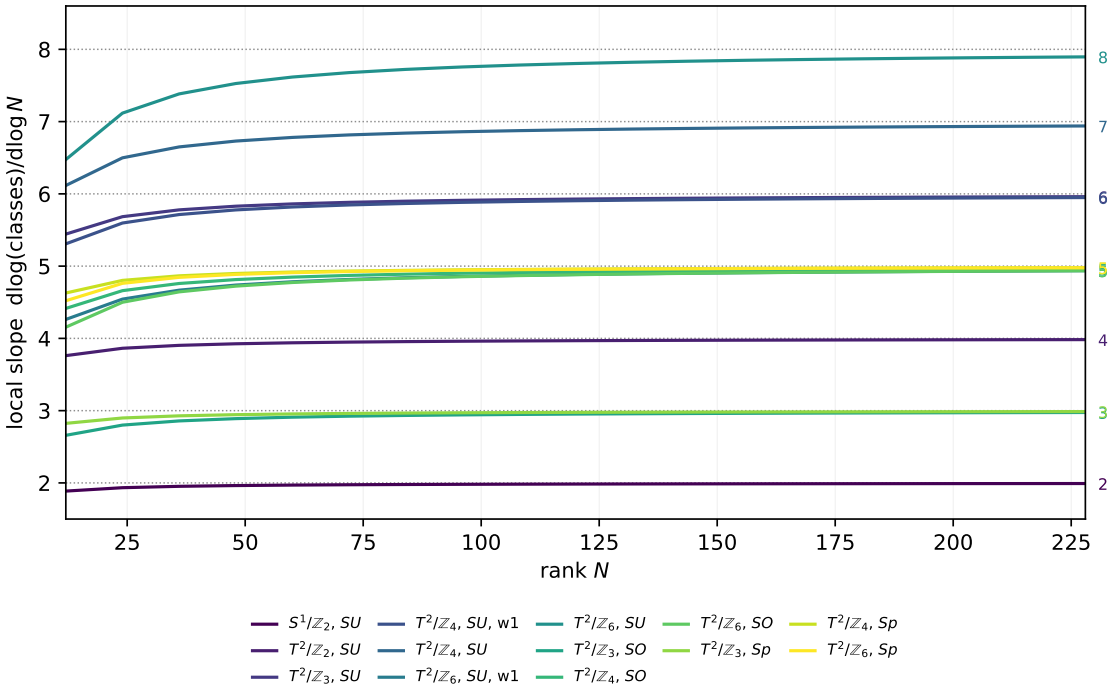


Figure 4: The degrees, measured. For each count of Table 3 — the two orthogonal rows that coincide with their unitary counterparts being the same curve, drawn once — and of Table 5, the local slope $\mathrm{d}\log(\text{classes})/\mathrm{d}\log N$, which tends to the degree of the quasi-polynomial. The dotted lines are the integers $2, 3, 4, 5, 6, 7, 8$ and the number at the right of each curve is the degree (7) predicts from the cone signature alone. Slopes are taken between multiples of 12, which is a multiple of every quasi-period occurring here, so the estimator cannot oscillate with the residue of N; the curves therefore approach their integers from below rather than around them.

The two remaining rows are what makes the hypothesis visible, and they are the reason (6) is stated first: the count sees the ambient dimension only if the alphabet *generates* it. The diagonal labels do not. On $T^2/\mathbb{Z}_4$ they span a 7-dimensional lattice inside the 8-dimensional one, and on $T^2/\mathbb{Z}_6$ a 6-dimensional lattice inside a 9-dimensional one, which is measured on the label vectors and is why those two rows have lower degree than (7) would give. It is also, read forwards, why labels of weight greater than one have to exist at all: without them the alphabet does not reach the local data the fixed points can carry.

Established here: one argument for the whole programme — five orbifolds, two gauge groups, twelve published counts, no free parameter — together with the degree of each, read from the cone signature. *Left open:* the completeness hypothesis itself, which is Appendix A; and the ranges here are ranges of N, not statements for all N, except where a closed form is exhibited.

5 The product orbifold

$S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$ is not of the form $T^2/\mathbb{Z}_m$, and the difference is not cosmetic. It is a product of quotients, its orbit space is a rectangle with mirror edges rather than a sphere with cone points, and the argument of Appendix A therefore runs in its disc version rather than its sphere one. The alphabet is unaffected — the space group is again of the form $\Lambda \rtimes F$ and §2 applies — and gives 16 labels of weight 1, one for each assignment of a parity under each of the four reflections,

two per circle. The four fixed points are the four ways of choosing one reflection from each circle, and each reads the corresponding pair of parities.

The class count follows:

$$\boxed{G(x) \;=\; \frac{1+7x+20x^2+28x^3+7x^4+x^5}{(1-x)^9}, \qquad 1,\ 16,\ 128,\ 688,\ \mathbf{2809},\ 9376,\ \ldots} \tag{8}$$

the denominator recording that the local datum spans a nine-dimensional lattice: 16 labels carrying $16-9=7$ independent relations among their images.

Both halves, on this orbifold. The lower bound is the necessary half again: each class lies inside one fibre of the local-datum map, so there are at least as many classes as fibres. The upper bound has to be constructed, and on a disc the construction is explicit. A gauge transformation is a smooth $\Omega : \mathbb{R}^2 \to SU(N)$ satisfying the four orbifold relations; evaluating each relation on its own fixed line collapses all four onto the four *sides* of the fundamental rectangle, leaving the interior unconstrained, and conversely any map on the rectangle meeting those four conditions extends to $\mathbb{R}^2$ by the reflection rules. The problem is then finite in shape:

- each *corner* carries two conditions at once, so the two local parities there must be simultaneously conjugated — which for sign matrices is exactly the conserved traces;
- each *side* asks for a path inside a centraliser $S(U(a)\times U(b))$, which is connected;
- the *interior* is a disc, so filling it is a null-homotopy of the boundary loop and the obstruction lies in $\pi_1(SU(N)) = 0$.

Taking the gauge transformation to be $SU(N)$-valued is the vanishing-flux assumption of §2 in the place where it does work, and not a second convention: the group that carries the 't Hooft flux of an $\mathfrak{su}(N)$ theory is the adjoint one, with $\pi_1(PSU(N)) = \mathbb{Z}_N$ (Appendix A), and a non-trivial class there is exactly what would obstruct the filling below.

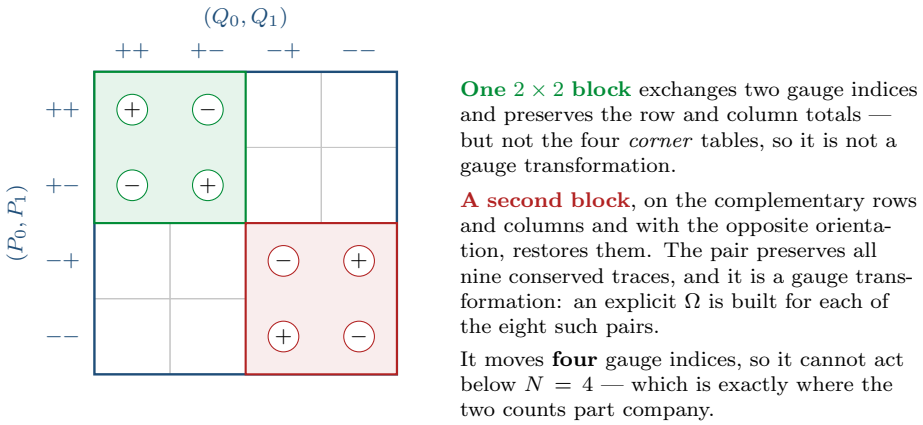
Proposition 1. *Let two boundary conditions of $SU(N)$ on $S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$ have the same local datum at each of the four fixed points. Then they are gauge equivalent.*

Proof. Equal local data at a corner means the two twists there are conjugate, so choose a conjugator C_j at each of the four corners. Along a side one parity P is fixed, and the conjugators taking it to its partner form a *coset* of the centraliser: if C_0 and C_1 are the choices at the two ends of that side then $C_0^{-1}C_1 \in C_G(P) = S(U(a)\times U(b))$, which is connected, so there is a path $c(t)$ in it from the identity to $C_0^{-1}C_1$ and $C_0 c(t)$ is a path of conjugators along the side. Four such paths concatenate to a loop $\gamma : S^1 \to SU(N)$ on the boundary of the rectangle. Since $\pi_1(SU(N)) = 0$, γ extends over the disc, which is the interior; and the extension propagates to $\mathbb{R}^2$ by the reflection rules, because the four relations were evaluated on their own fixed lines. The result is the required Ω. □

The construction is carried out explicitly for each of the eight moves that need it — four corner matrices, four side paths of the form $\exp(tX)$ with X in the centraliser the side requires, and the closed loop, to a worst deviation of 1.3×10^{-15} — which is a check on the construction and not a substitute for it. Both bounds are then the number of fibres, and (8) is the class count.

The degree, and what it means physically. By §6 the fibres of this orbifold connect at degree 4 and not below: eight independent moves relabel four gauge indices each, and no composition of moves touching two indices reaches them. That degree has a reading. A move is a half-period shift of the Wilson line along one circle, applied to a subset S of the gauge indices; the shift is a legal gauge transformation exactly when the fixed-point data survive it, which is three linear conditions per flipped parity. Counting the legal shifts by the size of S gives 8 on two

indices and 80 on four, and none at all on an odd number, since one index alone cannot balance a trace. *The degree at which the fibres connect is the number of gauge indices a legal half-period shift has to touch*, which is why nothing happens below $N = 4$: there are not four indices to move.



One 2×2 block exchanges two gauge indices and preserves the row and column totals — but not the four *corner* tables, so it is not a gauge transformation.

A second block, on the complementary rows and columns and with the opposite orientation, restores them. The pair preserves all nine conserved traces, and it is a gauge transformation: an explicit Ω is built for each of the eight such pairs.

It moves **four** gauge indices, so it cannot act below $N = 4$ — which is exactly where the two counts part company.

Figure 5: A degree-4 move on $S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$, drawn on the alphabet. The sixteen labels are the cells of a 4×4 grid: a label is a parity under each of the four reflections, two per circle, so a row fixes (P_0, P_1) and a column fixes (Q_0, Q_1). A boundary condition of rank N is N counters distributed over the cells, and the local datum at a fixed point is one of the four 2×2 tables obtained by summing the counters over a pair of parities, one from each circle. Neither shaded block is a legal move by itself; together they are, and by §6 no move of lower degree does the same job.

Higher products. For $(S^1/\mathbb{Z}_2)^k$ the same construction of the alphabet gives 4^k labels of weight 1, and the local-datum map has rank 3^k, both because label, fixed point and readout are products over the k circles and rank is multiplicative under the Kronecker product. So the count is a quasi-polynomial of degree $3^k - 1$ carried by 4^k generators. We do not give a closed form for general k here; what the pair $(4^k, 3^k)$ implies about the shape of that generating function is the subject of the companion note.

Established here: the class count of the product orbifold, with the gauge transformation constructed, and the degree at which its fibres connect. *Left open:* the closed form for $k \geq 3$, and whether the degree-4 behaviour is shared by every product of orbifolds.

6 The relations generate the fibres

The rule of §4 has two halves and they are proved by different means. That two boundary conditions in the same class carry the same local datum is the necessary half, and it is the trace conservation law [6]: the class cannot move what the traces at the fixed points fix. That two boundary conditions with the same local datum *are* in the same class is the sufficient half, and on the sphere it is Appendix A. Here we check it a second way, by asking whether the equivalence relations already in print are enough to walk from any boundary condition to any other with the same local datum.

The question is a finite one at each N. A boundary condition of weight N is a multiset of labels, an equivalence relation is a move on those multisets, and one may simply count the connected components of the move graph and compare them with the fibres of the local-datum map.

How much is needed. Independently of what is in print, one can enumerate *all* moves of degree at most d that preserve every conserved trace, and ask for the smallest d that connects

orbifold	relations	degree	components = fibres = published count
$T^2/\mathbb{Z}_2$	(3.13), three	2	$1, 8, 33, 96, 225, 456, 833, 1408, \ldots$ $(N \le 10)$
$T^2/\mathbb{Z}_3$	(3.23), two	3	$1, 9, 45, 163, 477, 1197, 2674, 5454, \ldots$ $(N \le 9)$
$T^2/\mathbb{Z}_4$	(3.37), one	4	$1, 8, 36, 120, 329, 784, 1680, 3312, \ldots$ $(N \le 10)$
$T^2/\mathbb{Z}_6$	none	—	$\binom{N+5}{5}$ $(N \le 10)$

Table 4: The sufficient half, counted. Equation numbers are those of [6]; the degree of a move is the number of gauge indices it relabels. In each row the three numbers coincide at every N in the range. On $T^2/\mathbb{Z}_6$ the diagonal sets are not connected by any relation and the local-datum map is injective, so the row is an identity rather than a check.

the fibres. The answer is 2, 3 and 4 for $T^2/\mathbb{Z}_2$, $T^2/\mathbb{Z}_3$ and $T^2/\mathbb{Z}_4$ — in each case the degree of the relations in Table 4. Nothing of lower degree suffices and nothing of higher degree is needed.

A scan of that kind is only worth its answer if it can return a different one, so it is calibrated on the orbifold of §5, where the two counts part company. There the moves of degree at most 2 give $1, 16, 128, 688, 2817$ while the fibres are $1, 16, 128, 688, 2809$; degree 3 changes nothing; degree 4 closes the gap. The scan therefore distinguishes "enough" from "not enough" on a case where the distinction is real, and it reports "enough" on the four above.

What the count rests on. Four properties are verified on the move vectors rather than taken from the prose that introduces them: every move annihilates every fixed-point marginal, so no move can join two fibres; the number of components is therefore never below the number of fibres, and is not; at $N = 1$ no move applies, so both equal the number of labels; and removing a single move from the three of $T^2/\mathbb{Z}_2$ raises the component count above the fibre count already at $N = 2$, which is what makes the agreement in Table 4 informative rather than automatic.

Established here: the sufficient half a second time, by counting, and the degree at which each set of relations operates. *Left open:* these are ranges of N, not proofs for all N; the statement that holds for all N is Appendix A.

7 The missing type: $Sp(N)$

Of the three compact classical families – three types of defining space, not three real forms of one group – two appear in this programme: the unitary case, and the orthogonal case [7]. We have not found the symplectic case treated, and by §3 it costs one line, because the type enters only through the weights.

Take the complex alphabet of §2 and read the Frobenius–Schur type of each irreducible. A complex irreducible V of complex dimension d that is of real type has $V \otimes \mathbb{H}$ as its symplectic realisation, of quaternionic dimension d; one of complex type pairs with its conjugate, and $V \oplus \bar{V}$ is again of quaternionic dimension d; one of quaternionic type is already symplectic, of quaternionic dimension $d/2$. Since none of the eleven irreducibles of $T^2/\mathbb{Z}_6$, and none of those of $T^2/\mathbb{Z}_3$ or $T^2/\mathbb{Z}_4$, is of quaternionic type — which is computed, not assumed — only the first two rules are used.

All three numerators factor as $\prod_i (1 - x^{d_i})$, with degrees $[3]$, $[4, 4]$ and $[6, 6]$ respectively; what that factorisation means about the underlying semigroup is the subject of the companion note and is not needed here. One consequence is worth recording, because it is visible already at this level: $T^2/\mathbb{Z}_3$ carries a single degree-3 factor over $Sp(N)$ where its orthogonal counterpart carries one of degree 6. The degree is a property of the alphabet's weights, not of the orbifold alone.

The labels, built. A table of weights is an inference until somebody writes the matrices, so we write them. Let V be a complex irreducible with generators ρ_0, ρ_1 and invariant bilinear form B, and let $\varepsilon = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$.

orbifold	$Sp(N)$ alphabet	$G(x)$	first terms
$T^2/\mathbb{Z}_3$	5×1	$\dfrac{1-x^3}{(1-x)^5}$	1, 5, 15, 34, 65, 111, 175, 260
$T^2/\mathbb{Z}_4$	$6\times 1+2\times 2$	$\dfrac{(1-x^4)^2}{(1-x)^6(1-x^2)^2}$	1, 6, 23, 68, 169, 370, 735, 1352
$T^2/\mathbb{Z}_6$	$4\times 1+2\times 2+2\times 3$	$\dfrac{(1-x^6)^2}{(1-x)^4(1-x^2)^2(1-x^3)^2}$	1, 4, 12, 30, 66, 132, 245, 428

Table 5: The symplectic case. The alphabet is written as (multiplicity)×(weight), the weight being the quaternionic dimension of the label. Each count is the number of distinct local data of total weight N, which is the number of equivalence classes under the hypothesis of Appendix A; that hypothesis applies here verbatim, since $Sp(n)$ is simply connected.

Proposition 2. *Each label of Table 5 is realised by the pair (R_0, R_1) and the form J below, which is unitary, satisfies the orbifold's consistency conditions, and satisfies $R^{\mathsf{T}}JR = J$ with J nondegenerate and antisymmetric:*

$$\begin{array}{lll} \textit{real type} & R_i = \rho_i \otimes I_2, & J = B \otimes \varepsilon, \\ \textit{complex type} & R_i = \rho_i \oplus \bar{\rho}_i, & J = \left(\begin{smallmatrix} 0 & I \\ -I & 0 \end{smallmatrix}\right), \\ \textit{quaternionic type} & R_i = \rho_i, & J = B. \end{array}$$

Proof. Unitarity and the consistency conditions are inherited from ρ in each case. For the first, $(\rho\otimes I)^{\mathsf{T}}(B\otimes\varepsilon)(\rho\otimes I) = (\rho^{\mathsf{T}}B\rho)\otimes\varepsilon = B\otimes\varepsilon$, and B symmetric against ε antisymmetric is antisymmetric. For the second, ρ unitary gives $\rho^{\mathsf{T}}\bar{\rho} = \rho^{\mathsf{T}}(\rho^{-1})^{\mathsf{T}} = (\rho^{-1}\rho)^{\mathsf{T}} = I$, which is exactly the off-diagonal block of $R^{\mathsf{T}}JR = J$. For the third, B is antisymmetric by definition of the type. □

A unitary matrix preserving a nondegenerate antisymmetric form lies in a conjugate of the compact symplectic group, of quaternionic dimension equal to half the complex one; so Proposition 2 assigns to each label exactly the weight Table 5 gives it. The twenty-one labels — five, eight and eight — were constructed and the four conditions verified on the matrices, to a worst deviation of 3.7×10^{-15}; their local data agree with the data the alphabet assigns them, and recounting from the constructed matrices returns the three sequences of Table 5.

Established here: the third type — alphabet, counts, closed forms, and every label written down. *Left open:* the change of basis to the standard symplectic form, which would replace "a conjugate of $Sp(n)$" by $Sp(n)$ and is bookkeeping; and the exceptional groups, where the argument of §3 runs unchanged but the alphabet is not computed here.

8 A preflight, and who should run it

If §2 and §4 are right, then classifying the boundary conditions of a new orbifold is not a project but a computation, and its only inputs are the rotation and the gauge group. That is a falsifiable claim about the shape of the work, so we make it operative: the ancillary script `bc_preflight.py` takes an integer matrix A of finite order acting on Λ and one of $SU(N)$, $SO(N)$, $Sp(N)$, and returns

- the cone signature — the orbits of the rotation on its own fixed points, with the order of each — which is derived and not supplied;
- the alphabet: how many labels of each weight, with each label built as explicit matrices from the induction of §2;
- the local datum of every label at every cone point, and hence the class count at each rank;

- the degree of that count, predicted from the signature alone by (7).

Run on the four two-dimensional rotations, it returns the signatures $(2,2,2,2)$, $(3,3,3)$, $(4,4,2)$ and $(6,3,2)$; the alphabets of Table 2; every count of Table 3 over $SU(N)$ and $SO(N)$; the three symplectic counts of Table 5, which reach them by a route independent of §7, since here they come from the induced matrices rather than from the Frobenius–Schur bookkeeping; and the degrees. Nothing in that list is entered into the script: the signature, the alphabet, the weights and the data all come out of A. A matrix of infinite order is refused rather than classified, which is the crystallographic restriction appearing as a property of the input; in rank r the orders that can occur at all are those with $\Phi(m) \le r$ for the additive totient Φ [28], and the ones this paper's hypothesis admits are the smaller set with $\varphi(m) \mid r$. For a rotation of prime order the possible inputs are themselves classified [25, Thm. 74.3],[26], so in that case the preflight can be run over all of them in a given rank rather than over one orbifold at a time.

One case is worth naming because the run makes it visible. On $T^2/\mathbb{Z}_2$ every label is a character of order 2, hence of real type, so *all three* types carry the same alphabet and the same count — the published statement that the orthogonal and unitary classifications agree there [7] is the $SO(N)$ half of a three-way coincidence, and it fails as soon as the orbifold has a cone point of order greater than 2.

Who should run it. Anyone about to classify boundary conditions on an orbifold or for a gauge group that has not been done. Before a single twist matrix is written by hand, the preflight gives the number of labels and their weights, which is the size of the problem; the degree of the answer, which is what a closed form must have; and the local data, which is what any equivalence relation must preserve. A proposed relation that moves a local datum is wrong, and a proposed count of the wrong degree is missing labels — both checks cost seconds and neither requires the classification to be finished.

Established here: the paper's operative consequence as a runnable check. *Left open:* orbifolds with mirror boundary. There the point group contains reflections, so the fixed-character strata are not all finite and the count is not $|\det(I - A^k)|$; separating the discrete stratum is exactly what §2 says this paper does not do.

9 How much of the rotation does the alphabet remember?

Theorem 1 says the alphabet is a function of (m, r) alone. Read backwards that is a question: two orbifolds with the same (m, r) have the same alphabet, so how much of the rotation survives in it?

The answer is exact, it is arithmetic rather than geometric, and none of it is ours. It is worth setting out in full, because between Theorem 1 and the literature the question is already closed.

Under the hypothesis the characteristic polynomial is Φ_m^s, so Λ is a torsion-free module of rank s over $\mathbb{Z}[\zeta_m]$, and two rotations are conjugate in $GL(r,\mathbb{Z})$ exactly when those modules are isomorphic. Four classical results then say everything.

- For $s = 1$ the characteristic polynomial is irreducible and the conjugacy classes of integer matrices with a given monic irreducible characteristic polynomial f are in bijection with the ideal classes of $\mathbb{Z}[\alpha]$, $f(\alpha) = 0$: **Latimer and MacDuffee** [24], 1933. Since $\mathbb{Z}[\zeta_m]$ is the maximal order, those are the ideal classes of $\mathbb{Q}(\zeta_m)$. For $s \ge 2$ the same follows once one knows that integral similarity carries no invariant beyond the module: the hypothesis makes Λ a module over $\mathbb{Z}[\zeta_m]$, which is the maximal order and hence a Dedekind domain, so Λ is torsion-free of rank s over it, two rotations are conjugate exactly when the modules are isomorphic, and Steinitz classifies those as $\mathbb{Z}[\zeta_m]^{s-1} \oplus I$ by the class of I alone.

- The $\mathbb{Z}G$-lattices of a cyclic group of prime order are classified — there are exactly $2h_p + 1$ indecomposable ones, and the h_p in the middle are the ideal classes: **Diederichsen** and **Reiner**, in the form of [25, Thm. 74.3].
- Two of them give the *same* group $\Lambda \rtimes \mathbb{Z}_m$ exactly when their ideal classes are **Galois conjugate**, since replacing the generator ρ by ρ^k twists the class by σ_k. That is **Charlap**'s classification [26, Ch. IV, Thm. 6.2], 1965, through his notion of semi-linear isomorphism.
- And for prime p, $h_p = 1$ if and only if $p \leq 19$: **Montgomery** and **Uchida** [27], 1971. Turning that into a statement about the *dimension* is the higher-dimensional crystallographic restriction, $\Phi(m) \leq n$ with Φ the additive totient: **Hiller** [28], 1985.

So the count of rotations is $h(\mathbb{Q}(\zeta_m))$ and the count of *space groups* is the number of Galois orbits, and the two differ. The equivalence that separates them is the one Charlap uses: integral conjugacy compares matrices, while *semilinear* equivalence — an isomorphism of lattices intertwining the two actions up to an automorphism of $\mathbb{Z}_m$ — compares the symmorphic space groups $\Lambda \rtimes \mathbb{Z}_m$ themselves, and allows ρ to be replaced by ρ^k. At $m = 23$, $r = 22$ — the first rank where anything happens — $h = 3$ and $h^+(\mathbb{Q}(\zeta_{23})) = 1$, so the class group is its own minus part and complex conjugation acts by inversion, leaving the orbits $\{1\}$ and $\{[\mathfrak{a}], [\mathfrak{a}]^{-1}\}$: *three rotations, two orbifolds.* The threshold itself — nothing below rank 22 — is the same one Nery reaches for a different invariant [29, Cor. 1.3], by the same two steps.

What this paper adds to that is one line, and it is a statement about our own construction rather than about lattices:

> **What the alphabet forgets**
>
> By Theorem 1 the alphabet is a function of (m, r). It is therefore constant on each of the families the classical results above enumerate, and
>
> *the alphabet remembers the rotation exactly as far as $h(\mathbb{Q}(\zeta_m)) = 1$ and no further.*
>
> Below rank 22 that is everything; at rank 22 it is three rotations, on two orbifolds, carrying one alphabet. The classification of the rotations is not ours, the class numbers are not ours, and the threshold is not ours; what is ours is which of our objects survives the identification, and the answer is: all of the alphabet, none of the arithmetic.

Everything above is inside the class of Theorem 1 — characteristic polynomial Φ_m^s — and says nothing about a rotation with several cyclotomic blocks, for which Λ is not a module over a single $\mathbb{Z}[\zeta_m]$ and the argument does not run. Within that class two consequences are worth separating. The negative one is that the alphabet is not a complete invariant: in rank 22 there are three inequivalent order-23 rotations, on two orbifolds, whose boundary conditions are indexed by the same letters with the same weights, and any classification built from the alphabet alone cannot tell them apart. The positive one is that below that rank it *is* complete, which includes every case in this paper and every case the programme has treated: for $m \in \{2, 3, 4, 6\}$ the alphabet recovers the rotation.

Established here: the alphabet as a function of (m, r), and hence that it does not separate what the classical classification separates. *Inherited:* everything else in this section. *Left open:* whether the *classification* — alphabet together with local data and the fibres of §4 — separates the two orbifolds of rank 22 that the alphabet alone does not; and the same question for a rotation whose characteristic polynomial has more than one cyclotomic factor, where even the alphabet is not covered by Theorem 1.

10 Reflections

What a weight is. A weighted alphabet looks like a counting convenience — some letters cost more than others — and it is not. The weight of a label is the dimension of an irreducible

representation of Γ, and by Clifford theory that dimension is the index of the stabiliser of a character of Λ. So the statement "this orbifold has labels of weight 3" is the statement "this orbifold has a rotation power whose fixed-point set is small", and the weights are a reading of the singular geometry rather than an accounting device. The extreme case is $T^2/\mathbb{Z}_6$, where the order-6 point has a single fixed point and carries labels of weight up to 3, while the two lesser cone points carry no independent information at all.

Why the classification is finite, and what that excludes. Only the characters of Λ that are isolated in the fixed set of their own stabiliser give rise to the labels counted here — for a rotation group, exactly those of non-trivial stabiliser. A character whose stabiliser has an infinite fixed set induces irreducibles as well, but they come in a continuous family, and those families are the Wilson lines. The classification is finite precisely because it drops them, and by Appendix A it drops the 't Hooft flux for the same kind of reason. So what all these counts count is the boundary conditions at vanishing Wilson line and vanishing flux; the continuous directions are not absent from the physics but are the other half of it, the half on which the Hosotani potential is a function. A classification and a potential are the discrete and continuous parts of one problem, and this paper is entirely about the first.

Where the rule stops. On the five simple quotients the local datum is a complete invariant and the counting is closed. On the product orbifold it is still complete — the disc argument applies — but the published relations no longer generate the fibres, and the moves that close them have degree 4 (§6). The place where the rule needs more is therefore the place where a list of relations obtained by hand falls short, and both are the same feature of the orbifold: its orbit space is a rectangle with four corners rather than a sphere with cone points, and a legal half-period shift must move four gauge indices at once.

11 Conclusion

After the lift of §2, an orbifold boundary condition is a representation of the space group. Everything in this paper is that sentence, applied: the labels are the irreducibles of $\Lambda \rtimes \mathbb{Z}_m$ induced from characters of non-trivial stabiliser, the weight of a label is its dimension, how many there are of each weight is a Möbius count over $|\det(I - A^k)|$, and the class of a boundary condition is the tuple of local data at the fixed points — the unbroken gauge group at each brane with its embedding. The gauge algebra enters that description only by choosing the type of the defining space — which does two things, not one: it fixes the real or quaternionic alphabet, and it fixes how much redundancy the local data carry, since over the real and quaternionic types a primitive root and its conjugate always share a multiplicity and a cone point of order m contributes $\lfloor m/2 \rfloor$ free coordinates instead of $m - 1$. The global form of the group enters where the lifts do, which is §5 and Appendix A.

What follows from it is a programme's worth of arithmetic obtained in one place: the twelve published counts, over five orbifolds and two gauge groups, with the degree of each read off the cone signature; the class count of $S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$, with the gauge transformation that closes it written down; the degree at which the fibres of each orbifold connect; and the symplectic column, which the programme had not reached because it is not another orbifold but another type of defining space. The identification also says where its own limits are: what is classified is the sector of vanishing Wilson line and vanishing 't Hooft flux, and the completeness it rests on is proved here for a compact simply connected group on a sphere with cone points, and over $SU(N)$ on a disc with mirror edges.

The consequence we would most like to be used is the smallest one. If a boundary condition is a representation, then the alphabet, the weights, the local data, the count and its degree are computable from the rotation matrix and the gauge group before any classification is attempted,

and §8 is that computation. An orbifold nobody has classified is not a research programme; it is an input.

The companion note [23] takes the same alphabet and asks the next question about it. The division is clean enough to be worth stating: what is settled here, by Möbius inversion and Clifford theory, is which generators exist; what is settled there, by toric and Markov-basis arguments, is which relations hold among them.

That is also what this part returns to the ones before it. Parts I to VIII each fix a boundary condition and ask what follows — which fermion content is admissible, what the potential selects, where the compactification scale can sit. The choice was always made from a list, and a list does not say how long it is. On $S^1/\mathbb{Z}_2 \times S^1/\mathbb{Z}_2$, the orbifold of Parts VI to VIII [11, 12, 13], it is 2809 at rank 4; and on any orbifold and gauge group those parts might move to next, the same number is now a computation rather than a search.

12 Ledger

Table 6: What each result owes.

Item	Inherited (cite)	Ours
classification of orbifold bundles by isotropy data	Furuta–Steer for $U(1)$, $SU(2)$, $U(2)$ [17, 18]; the equivariant-bundle line for abelian structure group [21] and for a graph base [19]; and, with the same two-invariant structure, Yalçinkaya [20] over S^2	nothing — the case we need is Appendix A, and we claim no priority
trace conservation on the boundary-condition matrices, and the equivalence relations counted in Table 4	Takeuchi–Inagaki [5, 6, 7]	nothing
local gauge groups at the fixed points	Asaka–Buchmüller–Covi [8]	nothing
conjugacy classes of $\mathbb{Z}_N$ embeddings, and their tables	Katsuki *et al.* [15, 14]; Nilles–Ramos-Sánchez–Vaudrevange [16]	nothing
't Hooft flux as π_1 of the adjoint	't Hooft [31]; von Gersdorff [32]; Steinberg [22]	nothing
Clifford–Mackey induction	standard	nothing
the alphabet as $\mathcal{A}(\Lambda \rtimes \mathbb{Z}_m)$, the induced part of $\mathrm{Irr}(\Lambda \rtimes \mathbb{Z}_m)$; weight = dimension, counted by fixed points	—	**ours**
that the gauge algebra enters through the type of the defining space — alphabet and local-datum redundancy both — and that on $T^2/\mathbb{Z}_2$ the three types coincide	—	**ours**

Item	Inherited (cite)	Ours
Theorem 1: the alphabet of $\mathbb{Z}^r \rtimes_A \mathbb{Z}_m$ in any rank, the hypothesis as a condition on eigenvalue orders, and (4)	**the whole of §9 except one line is inherited**: Latimer–MacDuffee [24] for the bijection with ideal classes, Steinitz for $s \geq 2$, Diederichsen–Reiner and [25, Thm. 74.3] for the lattices, Charlap [26] for when two of them give the same group, Montgomery–Uchida [27] and Masley–Montgomery [30] for the class numbers, Hiller [28] for the crystallographic restriction that turns them into a statement about the rank	**ours**: the alphabet as a function of (m, r), hence that it is constant on those classes — a statement about our construction and not about lattices
the class count of the product orbifold, with the gauge transformation constructed	—	**ours**
that $(S^1/\mathbb{Z}_2)^k$ has 4^k labels and a local-datum map of rank 3^k, hence a count of degree $3^k - 1$	—	**ours** (no closed form for general k is given here)
the degree at which the fibres connect, in each case	—	**ours**
the $Sp(N)$ alphabets, counts and constructed labels	—	**ours**

A Completeness of the local datum

Nothing in this appendix is claimed as new; the closing paragraph says where it stands. It is included because the rule of §4 needs its sufficient half in the generality in which the paper uses it, and because the proof locates exactly which hypothesis each case consumes.

The setting. A boundary condition is a Γ-equivariant structure on the trivial bundle $\mathbb{R}^2 \times G$, with $\Gamma = \Lambda \rtimes \mathbb{Z}_m$; a gauge transformation is an isomorphism of the underlying bundles intertwining two such structures. Isomorphism classes are therefore isomorphism classes of G-bundles over the orbit space, in the orbifold sense, and the local datum at a singular point is the isotropy representation there — for a cone point of order m, a homomorphism $\mathbb{Z}_m \to G$ up to conjugacy, which is the multiset of eigenvalue multiplicities used throughout this paper.

Theorem 2. *Let the orbit space be a sphere with k cone points and let G be compact, connected and simply connected. Then two boundary conditions are gauge equivalent if and only if their isotropy representations agree at every cone point.*

Proof. Necessity is the conservation of the fixed-point data. For sufficiency, remove an open disc D_i around each cone point and write P for the complement.

Each D_i retracts equivariantly onto its centre, so the structure over it is determined by the isotropy representation ρ_i. The complement P is homotopy equivalent to a wedge of $k-1$ circles, and a G-bundle over a graph is trivial for G connected, so the bundle underlying the structure over P is trivial and all the remaining information sits in the boundary gluing maps modulo

changes of trivialisation. Once trivialisations are chosen on P and on each D_i, the class is the tuple $(\rho_i; [g_1], \dots, [g_k])$ with $[g_i] \in \pi_1(G)$ the homotopy class of the transition map on ∂D_i. There are exactly two kinds of change of trivialisation, and they behave differently.

Over P, a map $h : P \to G$ sends $g_i \mapsto g_i \cdot (h|_{\partial D_i})^{-1}$. Since $\pi_1(G)$ is abelian, h acts only through $h_* : H_1(P) \to \pi_1(G)$, and $H_1(P) = \mathbb{Z}^k/(\sum = 0)$ is free of rank $k-1$; so the achievable shifts are exactly the tuples $(a_1, \dots, a_k)$ with $\sum_i a_i = 0$, and all of them are achievable.

Lemma 1. *Fix the isotropy ρ_i of finite order at a cone point. The equivariant transition functions on ∂D_i, modulo those that extend over D_i, form a torsor under $\pi_1(G)$.*

Proof. An equivariant $g : \partial D_i \to G$ satisfies $g(\omega z) = \rho_i\, g(z)\, \rho_i^{-1}$, so it is determined by its restriction to one fundamental arc of the $\mathbb{Z}_m$-action on the circle, which is a path in G from $x = g(z_0)$ to $\rho_i x \rho_i^{-1}$; conversely any such path defines an equivariant g. Two equivariant maps with the same endpoint x therefore differ by a loop at x, and their difference class in $\pi_1(G)$ is well defined because G is connected and π_1 is abelian, so the base point does not matter.

The endpoint itself can be normalised, and this is a step and not a remark. Let β be a path in G from x to e, which exists because G is connected, and set

$$\gamma_t \;=\; \big(\beta|_{[0,t]}\big)^{-1} \cdot \gamma \cdot \big(\rho_i\, \beta|_{[0,t]}\, \rho_i^{-1}\big).$$

Each γ_t runs from $\beta(t)$ to $\rho_i \beta(t) \rho_i^{-1}$, so each defines an equivariant map on ∂D_i; the family is continuous in t, $\gamma_0 = \gamma$ and γ_1 has endpoint e. A path of equivariant maps is an equivariant homotopy of transition functions, and by the standard clutching-homotopy argument homotopic equivariant transition functions define equivariantly isomorphic gluings — equivalently, the homotopy is realised on a collar — so γ and γ_1 represent the same class. We may therefore take $x = e$, and then γ is a loop at e and the class in question is $[\gamma]$.

It remains to see that g extends equivariantly over D_i exactly when $[\gamma] = 0$, and an ordinary null-homotopy is not by itself an equivariant one. If $[\gamma] = 0$ then γ is homotopic to the constant rel endpoints; transporting that homotopy by the $\mathbb{Z}_m$-action gives an equivariant homotopy from g to the constant map e on ∂D_i, which the homotopy extension property realises in a collar, the constant map filling the disc inside it. Conversely let Φ extend g equivariantly, and write $c = \Phi(0)$, which lies in $C_G(\rho_i)$ because the centre is fixed by the action. The value at the centre can be normalised *without touching the boundary*, and the two normalisations must not be made by the same move: multiplying Φ on the right by the constant c^{-1} would send g to gc^{-1} and undo the endpoint normalisation just made. Instead use a *radial* one. Since G is compact simply connected and ρ_i has finite order, $C_G(\rho_i)$ is connected [22]; choose a path $u : [0,1] \to C_G(\rho_i)$ with $u(0) = c^{-1}$ and $u(1) = e$, and set $\Phi'(z) = \Phi(z)\, u(|z|)$. This is still equivariant, since u takes values in $C_G(\rho_i)$ and therefore commutes past the conjugation; it has $\Phi'(0) = e$; and on ∂D_i it is $\Phi' = g\, u(1) = g$, unchanged.

Now the fundamental sector $0 \le \theta \le 2\pi/m$ is a disc, and Φ' restricted to it exhibits its boundary loop $\alpha \cdot \gamma \cdot (\rho_i \alpha \rho_i^{-1})^{-1}$ as null-homotopic, where α is the value of Φ' along the radial edge — a genuine loop at e, since $\Phi'(0) = e$ and $\Phi'(z_0) = g(z_0) = e$. Conjugation by ρ_i is homotopic to the identity through maps fixing e, because G is connected, so $[\rho_i \alpha \rho_i^{-1}] = [\alpha]$ and the null-homotopy leaves $[\gamma] = 0$. □

The proof of Theorem 2 uses that torsor structure and nothing finer: what varies at a cone point is a class in $\pi_1(G)$, not a component of a twisted loop group.

Over D_i, an equivariant automorphism is a map $\varphi : D_i \to G$ with $\varphi(\omega z) = \rho_i\, \varphi(z)\, \rho_i^{-1}$, sending $g_i \mapsto (\varphi|_{\partial D_i}) \cdot g_i$. Its value at the centre lies in the centraliser $C_G(\rho_i)$, which for G compact simply connected and ρ_i of finite order is connected [22], so the choice of that value contributes nothing to the class and no $\pi_0(C_G(\rho_i))$ obstruction arises. And φ is defined on the whole disc, so $\varphi|_{\partial D_i}$ bounds in G and $[\varphi|_{\partial D_i}] = 0$. The cone contributes no freedom at all.

The single invariant surviving the gluing is therefore $\sum_i [g_i] \in \pi_1(G)$, which the shifts over P can redistribute but not change. For G simply connected it vanishes, and the isotropy representations are a complete invariant. □

The disc case, with mirrors. $S^1/\mathbb{Z}_2$ and the products $(S^1/\mathbb{Z}_2)^k$ have orbit spaces with boundary rather than cone points: the orbit space is contractible, there is no closed surface to carry a characteristic class, and the gluing data live on arcs. The statement proved here for them is Proposition 1 of §5, where the gauge transformation is written down rather than shown to exist — corner matrices, side paths inside the centralisers, and a null-homotopy of the boundary loop. Its two ingredients are that the centraliser of a boundary-condition matrix is connected, which for a simply connected group is Steinberg's theorem [22], and that $\pi_1(SU(N)) = 0$ lets the interior be filled. Both are hypotheses on the group, so what is proved is completeness for these orbifolds over $SU(N)$; the shape of the argument is what would be needed for a general compact connected G, and the two places it would have to be checked are exactly those two.

What the surviving class is. When π_1 of the group acting on the fields does not vanish the class $\sum_i [g_i]$ is not anonymous: it is the 't Hooft flux [31]. Its home is the fundamental group of the *adjoint* group, as von Gersdorff states [32], reached there through holonomy rather than through bundles — so for the algebra $\mathfrak{su}(N)$ the relevant group is $\pi_1(PSU(N)) = \mathbb{Z}_N$ and not $\pi_1(SU(N)) = 0$, which is the same distinction between the group and its automorphisms that §2 begins with. This fixes the scope of everything in this paper, including the counts of §4: the published classifications are those of the vanishing-flux sector, stated as such by their authors [4], and so are ours. Together with the exclusion of the continuous families of §2, which are the Wilson lines, the object being counted is precisely the zero-Wilson-line, zero-flux sector.

On priority

The classification of orbifold bundles by isotropy data is established for $U(1)$, $SU(2)$ and $U(2)$ by Furuta and Steer [17, 18], and the equivariant-bundle line gives it for abelian structure group [21] and for a base which is a graph [19]. Closest in shape is Yalçinkaya [20], who classifies Γ-equivariant principal G-bundles over S^2 for $\Gamma \leq SO(3)$ finite and G compact connected, by the isotropy representations at the singular set *together with* a Chern class — which is the two-invariant structure of Theorem 2, with $\sum_i [g_i] \in \pi_1(G)$ in the second place. His base is a spherical orbifold, so it is not one of ours: the orbit spaces of $T^2/\mathbb{Z}_m$ have signatures $(2,2,2,2)$, $(3,3,3)$, $(4,4,2)$ and $(6,3,2)$, which are Euclidean and are not quotients of S^2 by a finite subgroup of $SO(3)$. The case used here — G compact and simply connected over a sphere with cone points — we have not located in that form, and we do not claim it as new: Theorem 2 is included so that the paper's sufficient half is available at the generality the counts require, and its proof is written to make visible which hypothesis is doing the work in each geometry.

B Reproducibility

Every number in this paper is produced by one of the ancillary scripts, each of which writes its own receipt and states its own controls; a script that cannot fail a check is not listed as evidence for anything. The correspondence between claim and script:

script	what it establishes	§
`alphabet_from_group_gate.py`	the alphabet from the fixed-point counts; the three types	2, 3
`marks_gate.py`	the determinant characters, and (5) on a non-cyclic point group	2, 3

script	what it establishes	§
`dimension_gate.py`	Theorem 1 outside dimension two: the hypothesis as a condition on eigenvalue orders, (4), and the multiplicities against a direct orbit count in ranks 3, 4	2, 9
`alphabet_formula_gate.py`	the closed form (3), with its decoy; and which orthogonal group carries the twists	2, 3
`so_n_z2_gate.py`	the two $\mathbb{Z}_2$ quotients over $SO(N)$	3, 4
`cone_rule_vs_ti.py`	the rule against the diagonal counts; the sensitivity control	4
`offdiagonal_labels.py`, `offdiagonal_z6.py`	the counts with the weighted alphabet	4
`so_n_rule.py`, `so_n_z4_z6.py`, `so_n_z6.py`	the three orthogonal counts	4
`minimal_form.py`	the degree from the signature, and what cancels	4
`gauge_realization.py`	the eight moves realised by an explicit Ω	5
`wilson_shift.py`	the half-period shifts that preserve the fixed-point data, counted by the size of their support: 8 on two gauge indices, 80 on four, none on an odd number	5
`km_general_k_table.py`	the $(4^k, 3^k)$ structure of $(S^1/\mathbb{Z}_2)^k$	5
`sphere_moves_gate.py`	components against fibres, and the degree scan	6
`sp_n_prediction_gate.py`, `sp_n_construction_gate.py`	the symplectic alphabets, counts and matrices	7
`figure_table_gate.py`	that the figures and the tables of the two notes carry the same numbers, read out of the sources themselves	all
`bc_preflight.py`	all of the above from the rotation matrix alone	8

The last line is the one to run first: it recomputes the signatures, alphabets, local data, counts and degrees of every case in this paper from the rotation matrix and the gauge group, with no published value entered anywhere, and compares the result with what is in print.

References

[1] N. Haba, Y. Hosotani, Y. Kawamura, *Classification and dynamics of equivalence classes in $SU(N)$ gauge theory on the orbifold $S^1/\mathbb{Z}_2$*, Prog. Theor. Phys. **111** (2004) 265.

[2] Y. Kawamura, T. Miura, *Equivalence classes of boundary conditions in $SU(N)$ gauge theory on 2-dimensional orbifolds*, Prog. Theor. Phys. **122** (2009) 847, arXiv:0905.4123.

[3] Y. Kawamura *et al.*, *On diagonal representatives in boundary condition matrices on orbifolds*, arXiv:2009.10958.

[4] Y. Kawamura *et al.*, arXiv:2211.00877.

[5] K. Takeuchi, T. Inagaki, arXiv:2401.09809.

[6] K. Takeuchi, T. Inagaki, *Classification of $T^2/\mathbb{Z}_m$ orbifold boundary conditions*, arXiv:2404.19411.

[7] K. Takeuchi, T. Inagaki, *Classification of $T^2/\mathbb{Z}_m$ orbifold boundary conditions in $SO(N)$ gauge theory*, PTEP **2025** 043B03, arXiv:2501.05849.

[8] T. Asaka, W. Buchmüller, L. Covi, hep-ph/0108021.

[9] Y. Komori, N. Maru, *$SU(7)$ grand gauge-Higgs unification*, arXiv:2503.04090.

[10] C. Marín, *What the Higgs Potential Cannot See: Bulk Matter That Cannot Help Select a Boundary Condition*, Part V of this series, doi:10.5281/zenodo.21727094.

[11] C. Marín, *Proton Decay in $SU(7)$ Grand Gauge-Higgs Unification: An Obstruction, Its Minimal Repair*, Part VI of this series, doi:10.5281/zenodo.22033302.

[12] C. Marín, *An Upper Bound on the Compactification Scale of $SU(7)$ Grand Gauge-Higgs Unification*, Part VII of this series, doi:10.5281/zenodo.22087251.

[13] C. Marín, *A Certified 2.68 TeV Gap in the Closed-Form Map of the Compactification Scale in $SU(7)$ Grand Gauge-Higgs Unification*, Part VIII of this series, doi:10.5281/zenodo.22159036.

[14] Y. Katsuki, Y. Kawamura, T. Kobayashi, N. Ohtsubo, Y. Ono, K. Tanioka, *$\mathbb{Z}_N$ orbifold models*, Nucl. Phys. B **341** (1990) 611.

[15] Y. Katsuki, Y. Kawamura, T. Kobayashi, N. Ohtsubo, K. Tanioka, *Gauge groups of $\mathbb{Z}_N$ orbifold models*, Prog. Theor. Phys. **82** (1989) 171.

[16] H. P. Nilles, S. Ramos-Sánchez, P. K. S. Vaudrevange, *Exploring the $SO(32)$ heterotic string*, arXiv:hep-th/0603086.

[17] M. Furuta, B. Steer, *Seifert fibred homology 3-spheres and the Yang–Mills equations on Riemann surfaces with marked points*, Adv. Math. **96** (1992) 38.

[18] B. Nasatyr, B. Steer, *Orbifold Riemann surfaces and the Yang–Mills–Higgs equations*, alg-geom/9504015.

[19] I. Hambleton, J.-C. Hausmann, *Equivariant bundles and isotropy representations*, arXiv:0704.2763.

[20] E. Yalçinkaya, *Equivariant principal bundles over the 2-sphere*, arXiv:1902.06293 (v3, 2024).

[21] R. K. Lashof, J. P. May, G. B. Segal, *Equivariant bundles with abelian structural group*, Contemp. Math. **19** (1983) 167.

[22] R. Steinberg, *Endomorphisms of linear algebraic groups*, Mem. Amer. Math. Soc. **80** (1968), Theorems 8.1 and 9.1(a).

[23] C. Marín, *The Semigroup of Orbifold Boundary Conditions*, companion note to this part, doi:10.5281/zenodo.22254863.

[24] C. G. Latimer, C. C. MacDuffee, *A correspondence between classes of ideals and classes of matrices*, Ann. of Math. **34** (1933) 313.

[25] C. W. Curtis, I. Reiner, *Representation Theory of Finite Groups and Associative Algebras*, Wiley (1962), Theorem 74.3.

[26] L. S. Charlap, *Bieberbach Groups and Flat Manifolds*, Springer (1986), Chapter IV, Theorems 6.2 and 6.3; the classification is Charlap, Ann. of Math. **81** (1965) 15.

[27] K. Uchida, *Class numbers of imaginary abelian number fields III*, Tôhoku Math. J. **23** (1971) 573; H. L. Montgomery, independently.

[28] H. Hiller, *The crystallographic restriction in higher dimensions*, Acta Cryst. A **41** (1985) 541.

[29] G. J. Nery, *Profinite genus of fundamental groups of compact flat manifolds with holonomy group of prime order*, arXiv:2006.11810.

[30] L. C. Washington, *Introduction to Cyclotomic Fields*, 2nd ed., Graduate Texts in Mathematics 83, Springer (1997).

[31] G. 't Hooft, *A property of electric and magnetic flux in non-abelian gauge theories*, Nucl. Phys. B **153** (1979) 141.

[32] G. von Gersdorff, arXiv:0705.2410.