STRINGQFT

String-Theoretic Construction
of Quantum Field Theory

VOLUME VII

Toric Singularities and Brane Tilings

Xingyang Yu

StringQFT v1.0

Drafted with AI assistance under the supervision of the author.
Manuscript free to read. All rights reserved.

Contents

VII.1 D3-branes at the conifold: the Klebanov–Witten theory 1

VII.1.1 Leaving the orbifold world . . . 1
VII.1.2 The conifold as a hypersurface . . . 3
VII.1.3 The quotient, the link, and the toric square . . . 4
VII.1.4 The D3 probe and the two-node quiver . . . 6
VII.1.5 The quartic superpotential and its F-terms . . . 9
VII.1.6 The one-brane mesonic moduli space . . . 10
VII.1.7 Marginality, R-charges, and symmetries . . . 12
VII.1.8 The square dimer . . . 15
VII.1.9 The holographic dual . . . 17
VII.1.10 What has been computed, and what is deferred . . . 18
Status and sources . . . 20
References . . . 22

VII.2 General toric Calabi–Yau threefold probes 23

VII.2.1 From one example to a general rule . . . 23
VII.2.2 The toric diagram and the (p,q)-web . . . 24
VII.2.3 The cone as a quotient: the GLSM . . . 27
VII.2.4 The D3 probe and the toric quiver . . . 29
VII.2.5 The mesonic moduli space recovers the cone . . . 31
VII.2.6 The inverse problem and toric duality . . . 33
VII.2.7 The bridge examples . . . 34
VII.2.8 What has been computed, and what is deferred . . . 37
Status and sources . . . 38
References . . . 40

VII.3 Brane tilings and dimer models 42

VII.3.1 One graph for the quiver and the superpotential . . . 42
VII.3.2 The dimer dictionary . . . 43
VII.3.3 The periodic quiver . . . 45
VII.3.4 The simplest fixture: the C^3 honeycomb . . . 47
VII.3.5 The Kasteleyn matrix and the toric diagram . . . 49
VII.3.6 Perfect matchings and the lattice points . . . 50
VII.3.7 Worked fixtures: the conifold square and dP_0 . . . 52
VII.3.8 Zig-zag paths and consistency . . . 55

VII.3.9 What has been computed, and what is deferred 58
Status and sources 59
References 62

VII.4 The NS5–D5 picture: brane tilings from T-duality 63
VII.4.1 Why the dimer is not just a drawing 63
VII.4.2 The setup: a D3 probe and the toric two-torus 65
VII.4.3 The two T-dualities and the brane dictionary 66
VII.4.4 The NS5 holomorphic curve and its Newton polygon 68
VII.4.5 The tiling as the skeleton of the projected curve 70
VII.4.6 Zig-zags, NS5 legs, and the (p, q)-web 72
VII.4.7 Worked anchors: the conifold and $\mathbb{C}^3$ 74
VII.4.8 The mirror route, named 77
VII.4.9 What has been constructed, and what is deferred 78
Status and sources 79
References 80

VII.5 Moduli spaces from dimers: perfect matchings and the Hilbert series 82
VII.5.1 From a tiling to a geometry 82
VII.5.2 Perfect matchings and the P-matrix 83
VII.5.3 The charge matrices and the toric diagram 85
VII.5.4 The conifold, worked through 87
VII.5.5 Perfect matchings carry multiplicities 88
VII.5.6 The master space and the mesonic–baryonic split 90
VII.5.7 The Hilbert series and the plethystics 93
VII.5.8 The worked bridge: $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ end-to-end 95
VII.5.9 What has been computed, and what is deferred 99
Status and sources 100
References 102

VII.6 Volume minimization and a-maximization 104
VII.6.1 Two extremization problems and one number 104
VII.6.2 Trial R-charges from the dimer data 105
VII.6.3 The principle: a-maximization 107
VII.6.4 The toric parametrization 109
VII.6.5 Volume minimization 110
VII.6.6 The match: a-maximization is volume minimization 112
VII.6.7 The conifold: completing the Chapter 1 preview 114
VII.6.8 $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$: the orbifold inheritance 115
VII.6.9 $Y^{2,1}$: the irrational case 116
VII.6.10 What has been computed, and what is deferred 119
Status and sources 120
References 122

VII.7 Toric duality as Seiberg duality 124

VII.7.1 One cone, several gauge theories . . . 124
VII.7.2 The rank-preserving toric window . . . 125
VII.7.3 The Seiberg-dual pair on the quiver . . . 127
VII.7.4 Seiberg duality as the dimer square move . . . 128
VII.7.5 The worked example: $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$ . . . 130
VII.7.6 The matched invariants . . . 134
VII.7.7 The del Pezzo phase counts and the duality tree . . . 136
VII.7.8 What has been computed, and what is deferred . . . 138
Status and sources . . . 139
References . . . 140

VII.8 Duality cascades and conifold deformation 142
VII.8.1 From the static brane to its dynamics . . . 142
VII.8.2 The fractional brane as a rank deformation . . . 143
VII.8.3 The three types of fractional brane . . . 146
VII.8.4 The duality cascade . . . 151
VII.8.5 The deformed-conifold infrared . . . 154
VII.8.6 The cascade for a general toric singularity . . . 157
VII.8.7 What has been computed, and what is deferred . . . 159
Status and sources . . . 160
References . . . 162

VII.9 Orientifolds of brane tilings 164
VII.9.1 The orientifold made combinatorial . . . 164
VII.9.2 The orientifold as an involution of the tiling . . . 165
VII.9.3 Fixed loci as O-planes: SO versus USp . . . 168
VII.9.4 Two-index matter and the crosscap . . . 170
VII.9.5 The anomaly-closing rank relation . . . 173
VII.9.6 Worked examples . . . 177
VII.9.7 What is and is not given . . . 180
Status and sources . . . 182
References . . . 183

VII.10 D1-branes at toric Calabi–Yau fourfolds: brane brick models and 2d $(0,2)$ theories 184
VII.10.1 From D3-branes to D1-branes . . . 184
VII.10.2 The D1-brane at a toric fourfold . . . 185
VII.10.3 The $(0,2)$ quiver: chiral and Fermi multiplets . . . 188
VII.10.4 The J-terms, the E-terms, and a Jacobi identity . . . 190
VII.10.5 $(0,2)$ triality as a cube move . . . 193
VII.10.6 The fixtures: $\mathbb{C}^4$ and the $\mathbb{C}^4/\mathbb{Z}_4$ orbifold . . . 195
VII.10.7 The sequence, and the outlook . . . 199
VII.10.8 What has been computed, and what is routed . . . 200
Status and sources . . . 202
References . . . 204

VII.11 Atlas: del Pezzo, $Y^{p,q}$, $L^{a,b,c}$, and brane-brick catalog 206
VII.11.1 What this atlas is, and how to read it 206
VII.11.2 The recipe, walked: the conifold, dP_0, and F_0 208
VII.11.3 Axis 1: the toric del Pezzo core 211
VII.11.4 Axis 2: the Sasaki–Einstein families 212
VII.11.5 Axis 3: the brane-brick CY4 sample 213
VII.11.6 Orientifold fixtures 215
VII.11.7 Cross-catalog patterns 216
VII.11.8 How to extend the atlas, and what is not in it 216
Status and sources 217

Chapter VII.1

D3-branes at the conifold: the Klebanov–Witten theory

VII.1.1 Leaving the orbifold world

The previous volume told its whole story through finite groups. A stack of D3-branes was placed at the tip of an orbifold $\mathbb{C}^3/\Gamma$, the gauge theory was read off the representation theory of the finite group Γ, the quiver was its McKay quiver, and the superpotential was inherited by projection from the $\mathcal{N}{=}4$ parent. Nodes were irreducible representations, arrows were counted by tensoring with the defining representation, and ranks were the dimensions of those representations. Everything followed from Γ.

Calabi–Yau singularities extend beyond finite quotients of flat space. The conifold is a concrete example: it has no finite group acting on $\mathbb{C}^3$ whose quotient gives the singularity, and therefore no McKay construction to apply directly. The task of this volume is to build the gauge theory on a D3-brane probe of a *general* toric Calabi–Yau threefold, and to find the combinatorial object that plays for the toric case the role the McKay quiver played for the orbifold case. That object is the brane tiling, or dimer, and the chapters that follow construct it. This chapter is the hinge: it works the simplest non-orbifold example in full, the conifold, and uses it to set up everything the tiling machinery will have to reproduce.

The conifold is the right first example for two reasons. It is iconic: its gauge theory, due to Klebanov and Witten, is the most studied $\mathcal{N}{=}1$ quiver after the orbifolds, and it reappears throughout the volume as the seed of the duality cascade (Chapter 8) and as the unit-cell example of every piece of tiling technology. And it is the cleanest possible illustration of the central contrast, because its defining equation is simple enough to study by hand yet its isolated singularity is not an orbifold. There is no finite $\Gamma \subset SU(3)$ with $\mathbb{C}^3/\Gamma$ equal to the conifold, and so there is no McKay quiver and no projection. Where, then, does the gauge theory come from?

Take-home message of this chapter

For an orbifold $\mathbb{C}^3/\Gamma$ the quiver and superpotential are dictated by the representation theory of Γ. The conifold has *no* finite Γ, hence no McKay correspondence to inherit from. The Klebanov–Witten construction supplies a candidate gauge theory. Its *geometry and symmetry* then give stringent checks: the flavour action, the one-brane mesonic quotient, and marginality at the interacting fixed point must fit together. Within the chosen two-node, four-bifundamental, single-trace ansatz the invariant quartic is fixed up to its coefficient. These checks do not establish uniqueness among all gauge theories. The same data is what the *dimer* of Chapter 3 will repackage as a graph on a torus. *Do not look for a finite group in this chapter; the geometry is the input.*

This chapter will not merely name the Klebanov–Witten theory. It will build it twice. First as a gauge theory: a two-node quiver whose superpotential, F-terms, anomalies, and one-brane moduli space are all computed by hand, with the moduli space turning out to be the conifold. Then as a brane tiling: the square dimer, whose faces, edges, and nodes encode the same quiver and superpotential, and whose toric output the later chapters learn to compute in general. The two constructions are the two halves of the volume, met here on their simplest example.

A word on what the chapter assumes and what it merely previews, so that a reader knows what to hold in mind and what may be deferred. The chapter assumes the probe construction and the basic four-dimensional $\mathcal{N}=1$ quiver vocabulary of Volume VI: gauge nodes, bifundamental arrows, the superpotential with its F-terms, the D-terms, and the gauge-invariant mesons. Toric diagrams, dimers, Kasteleyn matrices, and a-maximization appear here only as names and previews; each is built in the chapters that follow, and nothing in this chapter requires computing with them.

The answer, worked out below, is the Klebanov–Witten theory with two gauge nodes and four bifundamentals. We take its D3-brane identification from the string construction and show how its symmetries, quartic interaction, and moduli space fit the conifold. The symmetry argument constrains the interaction after the field content has been chosen; it does not derive an open-string spectrum from an arbitrary singularity. The familiar Seiberg move returns the conifold quiver in this toric phase, while other singularities have distinct toric phases (Chapter 7). The plan of the chapter is: the conifold as a hypersurface and a rank-one matrix (§VII.1.2); its presentation as a symplectic quotient, the link $T^{1,1}$, and the toric square (§VII.1.3); the D3 probe and the two-node quiver (§VII.1.4); the quartic superpotential and its F-terms (§VII.1.5); the one-brane mesonic moduli space (§VII.1.6); marginality, beta functions, R-charges, and the global symmetries (§VII.1.7); the square dimer fixture (§VII.1.8); the holographic dual (§VII.1.9); and a closing map of what was computed and what is deferred (§VII.1.10). Throughout, the contrast with Volume VI is the organizing thread: the same probe construction, but with the geometry, not a finite group, as the input.

VII.1.2 The conifold as a hypersurface and a rank-one matrix

The conifold is the simplest singular Calabi–Yau threefold beyond the orbifolds. It is the quadric hypersurface

$$\mathcal{C} = \{\, xy - zw = 0 \,\} \subset \mathbb{C}^4. \tag{VII.1.1}$$

To see that it is singular, and where, examine the defining function $f = xy - zw$ and its gradient,

$$\nabla f = (\partial_x f, \partial_y f, \partial_z f, \partial_w f) = (y,\, x,\, -w,\, -z). \tag{VII.1.2}$$

A point of $\mathcal{C}$ is a smooth point of the hypersurface where $\nabla f \neq 0$, and a singular point where $f = 0$ and $\nabla f = 0$ simultaneously. The gradient vanishes only at $x = y = z = w = 0$, which lies on $\mathcal{C}$, so the conifold has a single isolated singular point at the origin and is smooth everywhere else. It is a cone: if (x, y, z, w) solves (VII.1.1) then so does $(\lambda^2 x, \lambda^2 y, \lambda^2 z, \lambda^2 w)$ for any λ, so $\mathcal{C}$ is invariant under complex rescaling and is the affine cone over a projective base. With its Ricci-flat Calabi–Yau metric and the regular D3 projection, this background preserves four Poincare supercharges. The Calabi–Yau condition is a trivial canonical class, equivalently a nowhere-vanishing holomorphic three-form on the smooth locus, which here is the Poincaré residue

$$\Omega = \operatorname*{Res}_{xy-zw=0} \frac{\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z \wedge \mathrm{d}w}{xy - zw} = \frac{\mathrm{d}x \wedge \mathrm{d}z \wedge \mathrm{d}w}{x} \quad \text{(on the patch } x \neq 0\text{)}, \tag{VII.1.3}$$

the residue dividing by $\partial_y(xy - zw) = x$; it is nowhere vanishing on the smooth locus, and the patches glue across $\{x = 0\}$. We use only its existence below and return to the geometry in Part I.

The structure of the singularity is clearest if the four coordinates are packed into a 2×2 matrix,

$$\mathcal{Q} = \begin{pmatrix} x & z \\ w & y \end{pmatrix}, \qquad \det \mathcal{Q} = xy - zw, \tag{VII.1.4}$$

so that the defining equation is the single condition

$$\det \mathcal{Q} = 0, \qquad \text{equivalently} \qquad \operatorname{rank} \mathcal{Q} \leq 1. \tag{VII.1.5}$$

The conifold is the variety of rank-one (and rank-zero) 2×2 matrices. This rank-one form exposes the symmetry. A pair of $SU(2)$ matrices acting on the rows and columns, $\mathcal{Q} \mapsto L\,\mathcal{Q}\,R^\dagger$ with $L, R \in SU(2)$, preserves $\det \mathcal{Q}$, while an overall phase $\mathcal{Q} \mapsto e^{i\alpha}\mathcal{Q}$ preserves its zero locus, so $\mathcal{C}$ is invariant under

$$SU(2)_A \times SU(2)_B \times U(1)_R, \tag{VII.1.6}$$

of real dimension $3 + 3 + 1 = 7$. The two $SU(2)$ factors rotate the rows and the columns of $\mathcal{Q}$; the $U(1)$ rescales it by a phase, and it is the geometric realization of the superconformal R-symmetry once the gauge theory is switched on. This continuous symmetry constrains the interactions of the proposed probe theory. The choice of its gauge factors and charged fields also uses the D3-brane construction.

The action is concrete. Writing $\mathcal{Q}_{i\alpha}$ with row index i and column index α, the rule $\mathcal{Q} \mapsto L\,\mathcal{Q}\,R^\dagger$ has the left factor $L \in SU(2)_A$ act on the row index and the right factor $R \in SU(2)_B$ on the column index. Each column, $(x, w)^\mathsf{T}$ and $(z, y)^\mathsf{T}$, is therefore a doublet of $SU(2)_A$, and each row, (x, z) and (w, y), a doublet of $SU(2)_B$; the $U(1)_R$ multiplies all four coordinates by a common phase. In the outer-product parametrization of the next section, $\mathcal{Q} = v\,u^\mathsf{T}$, this is the statement that $v = (A_1, A_2)$ is an $SU(2)_A$ doublet and $u = (B_1, B_2)$ an $SU(2)_B$ doublet, so the mesons $M_{ij} = A_i B_j$ transform as $(\mathbf{2}, \mathbf{2})$, exactly how the bifundamentals of the gauge theory will. The geometry already contains the flavour structure of the quiver.

The global form depends on which objects the symmetry acts on. For the geometric action $\mathcal{Q} \mapsto e^{i\alpha} L\mathcal{Q}R^\dagger$, with α of period 2π, trivial action on all matrices requires $L = s_A\mathbb{1}$, $R = s_B\mathbb{1}$ and $e^{i\alpha} = s_A s_B$, where $s_A, s_B = \pm 1$. The rank-one matrices span the matrix space, so the same kernel is obtained by restricting to the conifold. Thus the faithful connected geometric action is

$$G_{\text{geom}} = \frac{SU(2)_A \times SU(2)_B \times U(1)_{\text{phase}}}{K}, \qquad K = \{(s_A\mathbb{1}, s_B\mathbb{1}, s_A s_B)\} \simeq \mathbb{Z}_2 \times \mathbb{Z}_2. \tag{VII.1.7}$$

In particular, the flavour subgroup alone has image $(SU(2)_A \times SU(2)_B)/\mathbb{Z}_{2,\text{diag}} = SO(4)$ on the mesons. The geometric phase is the R-rotation on mesonic scalar coordinates, which have charge one. This does not determine the faithful global group on all quantum operators. Elementary fields have R-charge $1/2$, baryons depend on N, and fermions and supercharges also transform. Their global identifications require the gauge-centre and fermion-parity information. Below we use the continuous flavour and R-symmetry algebra, without asserting a global quotient for the full interacting theory.

That this is not an orbifold is settled topologically in the next section: the conifold's link is the smooth $S^2 \times S^3$, never a quotient S^5/Γ.

VII.1.3 The symplectic quotient, the link, and the toric square

The rank-one description (VII.1.5) has a constructive counterpart that will return, almost word for word, as the gauge theory of the next sections. A rank-one 2×2 matrix is an outer product: $\mathcal{Q} = v\,u^\mathsf{T}$ for two-component vectors $v = (A_1, A_2)$ and $u = (B_1, B_2)$, so that

$$\mathcal{Q} = \begin{pmatrix} A_1B_1 & A_1B_2 \\ A_2B_1 & A_2B_2 \end{pmatrix}, \qquad x = A_1B_1,\ y = A_2B_2,\ z = A_1B_2,\ w = A_2B_1. \tag{VII.1.8}$$

Every rank-one $\mathcal{Q}$ is of this form, and the parametrization is redundant: rescaling $v \mapsto \lambda v$, $u \mapsto \lambda^{-1}u$ by $\lambda \in \mathbb{C}^*$ leaves $\mathcal{Q}$ unchanged. At zero stability parameter the conifold is the affine quotient of four variables A_1, A_2, B_1, B_2 modulo this $\mathbb{C}^*$,

$$(A_1, A_2;\, B_1, B_2) \;\sim\; (\lambda A_1, \lambda A_2;\, \lambda^{-1}B_1, \lambda^{-1}B_2), \qquad \lambda \in \mathbb{C}^*, \tag{VII.1.9}$$

where the double slash identifies closed orbit limits. In particular, the nonclosed orbits with $A = 0$ or $B = 0$ have the same image as the origin. Away from the

origin, quotienting a four-dimensional space by a one-dimensional group gives the three complex dimensions of $\mathcal{C}$. This is the conifold's presentation as a toric (symplectic, or GIT) quotient $\mathbb{C}^4//\mathbb{C}^*$, with the four coordinates carrying charges $(+1,+1,-1,-1)$. The invariants of (VII.1.9) are exactly the four products $M_{ij} = A_i B_j$ of (VII.1.8), and we will see in §VII.1.6 that the relation among them is precisely $\det \mathcal{Q} = 0$. The point to register now is that the conifold has a four-variable description with one $\mathbb{C}^*$ identification, and that this description is, field for field, the gauge theory of a single D3-brane probe.

The link. A cone is determined by its base, the link cut out at fixed radius. Pass to the compact (symplectic) version of the quotient by imposing the real moment-map constraint that goes with the $\mathbb{C}^*$ action and quotienting by its compact $U(1)$,

$$\mu \;=\; |A_1|^2 + |A_2|^2 - |B_1|^2 - |B_2|^2 \;=\; \zeta, \qquad (A,B) \sim (e^{i\theta}A, e^{-i\theta}B). \tag{VII.1.10}$$

At $\zeta = 0$, together with the radial sphere $|A_1|^2 + |A_2|^2 + |B_1|^2 + |B_2|^2 = r^2$, the two conditions force $|A_1|^2 + |A_2|^2 = |B_1|^2 + |B_2|^2 = r^2/2$, i.e. a three-sphere for the A's and a three-sphere for the B's, quotiented by the common circle $U(1)_H$:

$$T^{1,1} \;\simeq\; \big(S^3 \times S^3\big)/U(1)_H, \qquad U(1)_H :\; (A,B) \mapsto (e^{i\theta}A, e^{-i\theta}B). \tag{VII.1.11}$$

Its real dimension is $3+3-1=5$, as a link of a Calabi–Yau threefold must be. Writing each S^3 as $SU(2)$, this is the homogeneous space

$$T^{1,1} = \frac{SU(2)_A \times SU(2)_B}{U(1)_H}, \tag{VII.1.12}$$

whose topology is the smooth, simply connected $S^2 \times S^3$. (The S^2 is the base of the Hopf-like fibration of one $SU(2)$ by the quotient circle; the surviving S^3 is the other.) A point of notation that the geometry literature is strict about: the quotient subgroup $U(1)_H$ is the *stabilizer*, the relative-phase circle $(A,B) \mapsto (e^{i\theta}A, e^{-i\theta}B)$ embedded oppositely in the two $SU(2)$ Cartans, and it is the circle that becomes the gauged relative $U(1)$ in the field theory. It must not be confused with the Reeb circle $U(1)_R$ of (VII.1.6), the overall-phase isometry $\mathcal{Q} \mapsto e^{i\alpha}\mathcal{Q}$ that acts on the cone and becomes the superconformal R-symmetry; $U(1)_R$ is an isometry of $T^{1,1}$, not the subgroup divided out. The full isometry is the $SU(2)_A \times SU(2)_B \times U(1)_R$ of (VII.1.6), in the global form (VII.1.7). This is the sharpest form of the non-orbifold statement: an orbifold of flat space has link S^5/Γ, which for a free action is a spherical space form and otherwise an orbifold link with singular strata, and in neither case gives the conifold as a complex quotient singularity. For an isolated quotient singularity in complex dimension three without complex reflections, the action on the sphere is free. Its smooth link has finite covering sphere and hence zero rational second homology, whereas $H_2(S^2 \times S^3; \mathbb{Q}) = \mathbb{Q}$. A nonfree action would introduce a singular stratum away from the apex. Nontrivial complex reflections cannot lie in $SU(3)$, since their determinant would not be one.

The conifold is not an orbifold

There is no finite $\Gamma \subset SU(3)$ with $\mathbb{C}^3/\Gamma = \mathcal{C}$. The orbifold link is S^5/Γ; the conifold link is the smooth, simply connected $S^2 \times S^3$. They are not homeomorphic, so there is no McKay quiver and no orbifold projection. This is the structural break with all of Volume VI, and the reason a new construction is needed.

The constraint (VII.1.10) also names the conifold's resolution without constructing it. At $\zeta = 0$ the quotient is the singular cone. At $\zeta > 0$ the constraint forces $(A_1, A_2) \neq 0$, so the A's define a point of $\mathbb{P}^1$ and the singular point is replaced by that $\mathbb{P}^1$ (an S^2); at $\zeta < 0$ the roles of A and B swap and one finds the other $\mathbb{P}^1$. These are the two small resolutions, exchanged by the flop $\zeta \to -\zeta$. The alternative smoothing, the deformation $xy - zw = \epsilon$ that opens an S^3, is not visible in this real moment map; it is a complex-structure deformation. The small resolution is the Fayet–Iliopoulos branch and the home of the wrapped fractional branes of Chapter 8; the deformation is the geometric input to the Klebanov–Strassler infrared, also Chapter 8. The Sasaki–Einstein metric on $T^{1,1}$, and the resolution and deformation as geometry, belong to Part I; we use only the topology and the symmetry here.

The toric square. The same quotient is encoded combinatorially by the conifold's toric data, the lattice polygon dual to the cone. For the conifold this polygon is the **unit square**,

$$\Delta_{\mathcal{C}} = \mathrm{conv}\{(0,0),(1,0),(1,1),(0,1)\}, \tag{VII.1.13}$$

four boundary lattice points and no interior point. The general dictionary, lattice polygon to gauge theory, is the content of Chapter 2; here the unit square is quoted as the conifold's toric polygon, and only two integers read off it are used in this chapter,

$$2\,\mathrm{Area}(\Delta_{\mathcal{C}}) = \#\{\text{gauge nodes}\} = 2, \qquad d = \#\{(p,q)\text{ legs}\} = 4. \tag{VII.1.14}$$

The first, twice the normalized area, is the number of gauge nodes of the quiver; we quote this from the toric/dimer dictionary and defer its derivation to the tiling chapters. The second, the number of boundary lattice points, equals the number of external legs of the dual (p,q)-web and controls the baryonic symmetries of §VII.1.7. The (p,q)-web is the graph dual to the polygon: each edge of $\Delta_{\mathcal{C}}$ becomes a leg perpendicular to it, and the legs of an external (boundary) edge run off to infinity. The unit square has four boundary edges, so the conifold web has four external legs, two pairs of anti-parallel $(\pm 1, 0)$ and $(0, \pm 1)$ rays meeting at a vertex; this fourfold external structure is the geometric face of the single baryonic $U(1)_B$, $d - 3 = 1$. (The web is also the skeleton of the deformed and resolved conifold: splitting the vertex into two trivalent junctions, the two ways of splitting being the two small resolutions of §VII.1.3.) Figure VII.1.1 draws the polygon and its web.

VII.1.4 The D3 probe and the two-node quiver

Now place a stack of N D3-branes at the tip of the cone and ask for the four-dimensional theory on their worldvolume. The construction outputs the same kind of data as

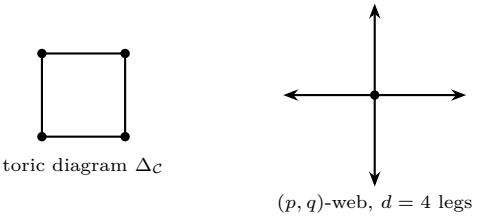


Figure VII.1.1: The conifold's toric data. Read the left panel as the toric input (the unit square, four boundary lattice points, no interior point) and the right panel as its perpendicular (p,q)-web dual (four external legs). Only two integers are used in this chapter: $2\,\mathrm{Area}=2$, quoted as the number of gauge nodes, and $d=4$, the number of boundary points that sets the count of baryonic symmetries.

in Volume VI, a quiver. Here we use the Klebanov–Witten proposal and test its compatibility with the geometry.

Begin with the colour structure. Each gauge node carries a colour space, and for the conifold there are two of them, as quoted in (VII.1.14),

$$V_1 \simeq \mathbb{C}^N, \qquad V_2 \simeq \mathbb{C}^N, \tag{VII.1.15}$$

with gauge group acting on each. The matter fields are linear maps between the colour spaces. The conifold isometry (VII.1.6) must be realized as a global symmetry, and the minimal field content carrying the two flavour doublets is two maps in each direction,

$$A_i:\ V_2 \to V_1, \qquad B_j:\ V_1 \to V_2, \qquad i,j=1,2, \tag{VII.1.16}$$

with the A_i a doublet of $SU(2)_A$ and the B_j a doublet of $SU(2)_B$. This is the field content of the Klebanov–Witten construction. Together with regular-brane ranks, it defines the toric phase used here. Its self-duality is a dynamical statement discussed in Chapter 7, not a consequence of the dimension of the mesonic quotient. In the language of representations, with the arrow convention that a field on the arrow $i \to j$ is a fundamental of node i and an antifundamental of node j,

$$A_i:\ 1 \to 2, \quad A_i \in (\mathbf{N}_1, \overline{\mathbf{N}}_2); \qquad B_j:\ 2 \to 1, \quad B_j \in (\overline{\mathbf{N}}_1, \mathbf{N}_2). \tag{VII.1.17}$$

The brane quiver retains the unitary Chan–Paton factors of the cited construction,

$$\prod_i U(N_i) = U(N)_1 \times U(N)_2, \qquad N_1 = N_2 = N, \tag{VII.1.18}$$

the equal, regular-brane ranks: N mobile D3-branes carry rank N at each node. Table VII.1.1 collects the quantum numbers and Figure VII.1.2 is the quiver.

The abelian factors. The unitary gauge group contains two abelian factors that it is essential to separate, because they have opposite fates. Decompose

$$U(N)_1 \times U(N)_2 \ \simeq\ \big[SU(N)_1 \times SU(N)_2 \times U(1)_+ \times U(1)_-\big]/(\text{finite centre}), \tag{VII.1.19}$$

field	gauge representation	$SU(2)_A$	$SU(2)_B$	$U(1)_B$
A_i $(i=1,2)$	$(\mathbf{N}_1, \overline{\mathbf{N}}_2)$	doublet	singlet	$+1$
B_j $(j=1,2)$	$(\overline{\mathbf{N}}_1, \mathbf{N}_2)$	singlet	doublet	-1

Table VII.1.1: Matter content of the Klebanov–Witten quiver and its quantum numbers under the flavour symmetry $SU(2)_A \times SU(2)_B$ and the baryonic symmetry $U(1)_B$ of §VII.1.7. The superconformal R-charge $R(A) = R(B) = \frac{1}{2}$ is derived in §VII.1.7.

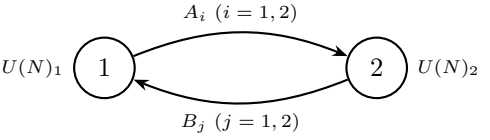


Figure VII.1.2: The Klebanov–Witten quiver. Two gauge nodes $U(N)_1 \times U(N)_2$; two bifundamentals A_i from node 1 to node 2 (an $SU(2)_A$ doublet) and two bifundamentals B_j from node 2 to node 1 (an $SU(2)_B$ doublet). The closed products A_iB_j are mesonic matrices; their traces are gauge invariant, and for one brane the products themselves are invariant coordinates.

where $U(1)_\pm$ are the sum and difference of the two central $U(1)$'s. The diagonal $U(1)_+$, under which A and B are both neutral, is the decoupled centre-of-mass gauge multiplet of the stack; it does not couple to the interacting fields. (Unlike the $\mathcal{N}$=4 case, the centre-of-mass *position* is not a scalar in this $U(1)_+$ vector multiplet; it is carried by the mesonic chiral coordinates, equivalently the trace part of the meson matrix, as §VII.1.6 makes explicit.) The relative $U(1)_-$ acts as

$$A_i \mapsto e^{+i\alpha} A_i, \qquad B_j \mapsto e^{-i\alpha} B_j, \tag{VII.1.20}$$

exactly the $\mathbb{C}^*$ of the conifold quotient (VII.1.9) at the abelian level. The relative gauge coupling runs to zero in the infrared description of this construction. When we discuss the interacting sector alone, we omit the abelian vector factors and use

$$SU(N) \times SU(N) \quad \text{with the global symmetry } U(1)_B. \tag{VII.1.21}$$

The baryonic transformation has the same charges as (VII.1.20), but is now a global symmetry of the SU theory. A gauge redundancy has not become a physical global symmetry inside the same quotient: the unitary brane moduli problem and the interacting SU theory are different descriptions. We keep the relative $U(1)$ and its D-term when constructing the one-brane geometric quotient below, but the construction's output is (VII.1.18), and we never write the brane gauge group as a bare SU.

Anomaly cancellation. The non-abelian gauge anomaly cancels at the equal ranks, and it is worth seeing that the cancellation is between conjugate representations, not a coincidence of counting. The cubic $SU(N)_1^3$ anomaly is the sum of the cubic indices of the chiral fermions charged under node 1, weighted by their multiplicity from the other

node:

$$\mathcal{A}\big(SU(N)_1\big) \;\propto\; \underbrace{2}_{A_1,A_2}\, N\,A(\mathbf{N}) \;+\; \underbrace{2}_{B_1,B_2}\, N\,A(\overline{\mathbf{N}}) \;=\; 2N\big[A(\mathbf{N})+A(\overline{\mathbf{N}})\big] = 0, \tag{VII.1.22}$$

because the cubic anomaly coefficient of a conjugate representation is the opposite, $A(\overline{\mathbf{N}}) = -A(\mathbf{N})$. Counted as fundamentals minus antifundamentals this is simply $\#\text{fund} - \#\text{antifund} = 2 - 2 = 0$, and the same holds at node 2 by the symmetry of the quiver. For $N \geq 3$ this is the perturbative cubic gauge-anomaly test. For $N = 2$ the local cubic tensor vanishes and each node has an even number $4N$ of doublets, so its mod-two gauge anomaly also vanishes. For $N = 1$ there is no nonabelian factor. The mixed $U(1)_B\,SU(N)^2$ anomaly cancels by the same opposite-charge count, since A and B carry baryonic charges $+1$ and -1; we return to $U(1)_B$ in §VII.1.7.

VII.1.5 The quartic superpotential and its F-terms

The matter content of (VII.1.17) is not yet a theory: an $\mathcal{N}$=1 quiver is defined by its quiver *and* its superpotential, a gauge-invariant holomorphic function of the chiral fields. Single-trace mesonic chiral operators are traces of closed oriented loops in the quiver, so the candidate superpotential terms are enumerated by loop length.

- *Length one.* A single bifundamental is an open arrow, not gauge invariant. No term.
- *Length two.* $\mathrm{Tr}(A_iB_j)$ closes the loop $1 \to 2 \to 1$ and is gauge invariant, but it transforms as $(\mathbf{2},\mathbf{2})$ of $SU(2)_A \times SU(2)_B$, not as a flavour singlet, so it is not allowed by the symmetry. It is also a mass term: including it would lift the bifundamentals and destroy the conifold fixed point. The conifold theory has no mass term.
- *Length three.* A loop of three arrows on a two-node quiver cannot close: each arrow switches nodes, and three switches end on the opposite node. No cubic term. (This is the sharp contrast with the orbifold quivers of Volume VI, where three-node loops gave a cubic superpotential.)
- *Length four.* $\mathrm{Tr}(A_iB_kA_jB_l)$ closes the loop $1 \to 2 \to 1 \to 2 \to 1$ and is gauge invariant. It is the shortest single trace that can be made a flavour singlet, and it is marginal at the interacting fixed point.

For two A doublets and two B doublets, the only $SU(2)_A \times SU(2)_B$ invariant tensors are the antisymmetric ε^{ij} on the A flavours and ε^{kl} on the B flavours. Within this single-trace quartic ansatz the superpotential is unique up to its overall coefficient,

$$W = \lambda\,\varepsilon^{ij}\varepsilon^{kl}\;\mathrm{Tr}(A_i\,B_k\,A_j\,B_l)\,. \tag{VII.1.23}$$

Each epsilon is nonzero only for the two orderings $(1,2)$ and $(2,1)$, so the double sum over the A flavours (i,j) and the B flavours (k,l) has $2\times 2 = 4$ terms. With $\varepsilon^{12} = -\varepsilon^{21} = 1$,

the sign of a term is the product of the two epsilons: $(i,j;k,l) = (1,2;1,2)$ and $(2,1;2,1)$ carry $+$, while $(1,2;2,1)$ and $(2,1;1,2)$ carry $-$,

$$W = \lambda\,\mathrm{Tr}\big(A_1B_1A_2B_2 - A_1B_2A_2B_1 - A_2B_1A_1B_2 + A_2B_2A_1B_1\big), \tag{VII.1.24}$$

and cyclicity of the trace identifies the first with the fourth and the second with the third, so

$$W = 2\lambda\,\mathrm{Tr}\big(A_1B_1A_2B_2 - A_1B_2A_2B_1\big). \tag{VII.1.25}$$

This is the two-term form that Chapter 3 will read directly off the dimer. It is also a determinant in disguise: collecting the mesonic operators $M_{ik} = A_iB_k$, each an adjoint-valued operator at node 1, the two terms are $M_{11}M_{22}$ and $M_{12}M_{21}$, so

$$W = 2\lambda\,\mathrm{Tr}\big(M_{11}M_{22} - M_{12}M_{21}\big), \tag{VII.1.26}$$

quadratic in the mesons, not cubic. For one brane, where the fields commute, this is literally the trace of the 2×2 determinant of the meson matrix; for nonabelian fields the entries do not commute, and it is better read as the antisymmetrized trace expression that reduces to that determinant on the abelian branch.

The F-terms. The supersymmetric vacua extremize W. Writing $h = 2\lambda$ for the two-term normalization (VII.1.25), the F-term of each field is the cyclic derivative of the trace, obtained by deleting that field from each word and reading the remainder cyclically:

$$\frac{\partial W}{\partial A_1} = h\big(B_1A_2B_2 - B_2A_2B_1\big), \qquad \frac{\partial W}{\partial A_2} = h\big(B_2A_1B_1 - B_1A_1B_2\big), \tag{VII.1.27}$$

$$\frac{\partial W}{\partial B_1} = h\big(A_2B_2A_1 - A_1B_2A_2\big), \qquad \frac{\partial W}{\partial B_2} = h\big(A_1B_1A_2 - A_2B_1A_1\big). \tag{VII.1.28}$$

Two features of these equations carry the physics of the moduli space. For a single brane, $N = 1$, the fields are commuting numbers and every right-hand side vanishes identically: for example $B_1A_2B_2 - B_2A_2B_1 = A_2B_1B_2 - A_2B_1B_2 = 0$. The one-brane F-terms impose no condition, and the moduli space is governed entirely by the gauge quotient, as the next section shows. For $N > 1$ the fields are matrices, the words do not commute, and the F-terms (VII.1.27)–(VII.1.28) are the non-trivial relations that, together with the gauge quotient, make the moduli space the symmetric product of conifold points on the generic commuting branch. It is because of this $N > 1$ content that the determinantal form (VII.1.26) was worth recording: the interaction constrains noncommuting matrices even though it vanishes for a single brane.

VII.1.6 The one-brane mesonic moduli space

The probe construction promises that the brane reconstructs its transverse geometry as a branch of its own moduli space of vacua. For Volume VI this recovered $\mathbb{C}^3/\Gamma$; here it must recover the conifold, and for one regular brane it does so as a clean ring computation that reproduces the quotient presentation of §VII.1.3.

Take $N = 1$, so the gauge group is $U(1)_1 \times U(1)_2$ and the fields A_i, B_j are complex numbers. As established, the F-terms are trivial, so the moduli space is the space of the four fields modulo the gauge group. The diagonal $U(1)_+$ acts trivially on the bifundamentals (the centre of mass), and only the relative $U(1)_-$ acts, with charges

	A_1	A_2	B_1	B_2
$U(1)_-$	$+1$	$+1$	-1	-1

(VII.1.29)

Combining this $U(1)$ gauge rotation with its D-term constraint complexifies it to the $\mathbb{C}^*$ of (VII.1.9), and, at $\zeta = 0$, the moduli space is the affine GIT quotient $\mathbb{C}^4/\!/\mathbb{C}^*$, i.e. the ring of $\mathbb{C}^*$-invariants.

The invariant ring. An invariant monomial in A_1, A_2, B_1, B_2 must have $U(1)_-$ charge zero, hence equal numbers of A's and B's. Any such monomial can be paired off into factors A_iB_j, so the invariant ring is generated by the four bilinears

$$M_{ij} = A_iB_j, \qquad i, j = 1, 2. \tag{VII.1.30}$$

These four generators are not independent. Compute the two products that use all four fields:

$$M_{11}M_{22} = (A_1B_1)(A_2B_2) = A_1A_2B_1B_2 = (A_1B_2)(A_2B_1) = M_{12}M_{21}, \tag{VII.1.31}$$

so $M_{11}M_{22} - M_{12}M_{21} = 0$. To show that there are no further relations, let $S = \mathbb{C}[m_{11}, m_{12}, m_{21}, m_{22}]$ map to the invariant ring by $m_{ij} \mapsto A_iB_j$. We proved surjectivity by pairing every charge-zero monomial. The image has transcendence degree three: the exponent vectors of A_1B_1, A_1B_2, A_2B_1 are independent, while all four vectors lie in the rank-three charge-zero lattice. The kernel is therefore a height-one prime in the polynomial UFD S. The polynomial $m_{11}m_{22} - m_{12}m_{21}$ is irreducible: it is primitive and linear in m_{22} over $\mathbb{C}[m_{11}, m_{12}, m_{21}]$, since m_{11} and $m_{12}m_{21}$ are coprime. The kernel contains this polynomial and hence is precisely its principal ideal. This proves the ring presentation over $\mathbb{C}$, rather than only checking a determinant at selected points. The invariant ring is therefore

$$\mathbb{C}[A_1, A_2, B_1, B_2]^{\mathbb{C}^*} = \mathbb{C}[M_{11}, M_{12}, M_{21}, M_{22}]/(M_{11}M_{22} - M_{12}M_{21})\,. \tag{VII.1.32}$$

Identifying the four mesons with the four conifold coordinates,

$$(x, y, z, w) = (M_{11}, M_{22}, M_{12}, M_{21}), \tag{VII.1.33}$$

the single relation is exactly $xy - zw = 0$, and the moduli space is

$$\mathcal{M}^{\text{mes}}_{N=1} = \text{Spec}\ \frac{\mathbb{C}[x, y, z, w]}{(xy - zw)} = \mathcal{C}. \tag{VII.1.34}$$

The one-brane mesonic moduli space is the conifold, recovered not as a dimension count but as an isomorphism of coordinate rings: it is the quotient presentation of §VII.1.3, now derived from the gauge theory rather than assumed. The dimension check is the

corollary: four generators with one relation, or four fields with one $\mathbb{C}^*$ removed, give the three complex dimensions of $\mathcal{C}$.

It is worth seeing that the gauge-theory derivation is, term by term, the geometry of §VII.1.3. The complexified $\mathbb{C}^*$ of (VII.1.9) is the relative $U(1)_-$ combined with its real D-term, and that D-term is precisely the moment map (VII.1.10),

$$D_- = |A_1|^2 + |A_2|^2 - |B_1|^2 - |B_2|^2 - \zeta = 0, \tag{VII.1.35}$$

with ζ the Fayet–Iliopoulos parameter of the relative $U(1)$. At $\zeta = 0$ the quotient is the singular conifold. This abelian quotient describes the probe geometry; it is not the interacting $SU(N)^2$ SCFT at $N = 1$; turning on $\zeta \neq 0$ resolves the tip exactly as in (VII.1.10), the small resolution being the gauge theory's Fayet–Iliopoulos (Kähler) deformation. At nonzero ζ, the quotient also includes stability: $A \neq 0$ for $\zeta > 0$, and $B \neq 0$ for $\zeta < 0$. The affine meson ring still describes the singular conifold, because the exceptional $\mathbb{P}^1$ is contracted by all mesons. Thus the resolved space and the affine spectrum of that ring are not the same variety. The relative FI parameter belongs to this unitary probe quotient, not to an FI deformation of a simple $SU(N)$ factor.

For N regular branes the F-terms (VII.1.27)–(VII.1.28) are no longer trivial. On the regular-brane mesonic branch at $\zeta = 0$, away from coincidences and the apex, the four mesonic matrices $M_{ij} = A_i B_j$ commute and are simultaneously diagonalizable, and the F-terms reduce, eigenvalue by eigenvalue, to the one-brane relation $\det M = 0$: each of the N eigenvalues traces out a copy of the conifold. The N branes are indistinguishable, so the closure of this position branch is described by the symmetric product

$$\mathcal{M}_N^{\text{mes}} = \text{Sym}^N(\mathcal{C}), \tag{VII.1.36}$$

the N unordered points being the positions of the D3-branes on the conifold, encoded in the eigenvalues of the meson matrix; for $N = 2$, for instance, the two D3-branes sit at two arbitrary points of $\mathcal{C}$. The diagonal $U(1)_+$ gauge multiplet decouples from this interacting mesonic dynamics. As in Volume VI, this is the classical geometric moduli space; reading it as the literal transverse space carries the α' and g_s probe caveats of Volume V. The baryonic branch and the full master space, the moduli space before the baryonic quotient, are the business of Chapter 5; here only the mesonic branch is needed, and it is the conifold.

VII.1.7 Marginality, beta functions, and R-charges

For $N \geq 2$, we use the interacting Klebanov–Witten fixed point and its unbroken exchange symmetry as physical premises. Its superconformal R-symmetry assigns

$$R(A_i) = R(B_j) = \tfrac{1}{2}. \tag{VII.1.37}$$

This section shows where that value comes from. Symmetry alone does not fix it; what fixes it here is a beta-function consistency computation at the assumed interacting fixed point. The general selection principle, a-maximization, is the deliverable of Chapter 6; for the conifold the symmetric answer is visible without it, and seeing it makes the word superconformal earned.

Why a quartic can be marginal. In a four-dimensional $\mathcal{N}$=1 superconformal theory a chiral primary of R-charge R has scaling dimension

$$\Delta = \tfrac{3}{2}R, \tag{VII.1.38}$$

and a superpotential term, integrated over $\mathrm{d}^2\theta$, must have R-charge 2 and hence dimension 3 to be marginal. A quartic of fields of dimension $\frac{3}{4}$ has dimension 3: this is possible only because the fields acquire anomalous dimensions at the interacting fixed point, and it is why a quartic superpotential, naively irrelevant by free-field counting, is marginal here.

The beta-function check. The anomalous dimension is fixed by demanding a fixed point. For the gauge factor $SU(N)_1$ with $N_f = 2N$ flavours (the two A doublets, each N from node 2, and likewise the B's), the NSVZ numerator is the one-loop coefficient corrected by the anomalous dimensions γ_A, γ_B of the matter,

$$\beta\big(SU(N)_1\big) \;\propto\; 3N - N(1-\gamma_A) - N(1-\gamma_B). \tag{VII.1.39}$$

By the $A \leftrightarrow B$ symmetry of the conifold point $\gamma_A = \gamma_B \equiv \gamma$, and the numerator is $N(1+2\gamma)$. A conformal fixed point requires it to vanish,

$$N(1+2\gamma) = 0 \quad\Longrightarrow\quad \gamma = -\tfrac{1}{2}. \tag{VII.1.40}$$

Here γ is normalized by $\Delta = 1 + \gamma/2$; conventions using $\Delta = 1 + \gamma$ give an anomalous dimension half as large. For the charged elementary fields, Δ denotes the fixed-point scaling assignment entering the beta function, not the dimension of a gauge-invariant local primary. The mesons are chiral primaries with $\Delta(M) = 3/2$, compatible with the unitarity bound. The scalar scaling assignment is then $\Delta = 1 + \frac{1}{2}\gamma = \frac{3}{4}$, and the superconformal relation (VII.1.38) gives

$$\Delta(A) = \Delta(B) = \tfrac{3}{4} \quad\Longrightarrow\quad R(A) = R(B) = \tfrac{1}{2}, \tag{VII.1.41}$$

which is (VII.1.37). As a consistency check, the quartic then has $\Delta(W) = 4 \cdot \frac{3}{4} = 3$ and $R(W) = 4 \cdot \frac{1}{2} = 2$, marginal, as required for the superpotential to be present at the fixed point.

The central charge: the field-theory side of a preview. With the R-charge in hand the conformal central charge a can be read off the R-symmetry 't Hooft anomalies. For an $\mathcal{N}$=1 superconformal theory

$$a = \frac{3}{32}\big(3\ \mathrm{Tr}\, R^3 - \mathrm{Tr}\, R\big), \tag{VII.1.42}$$

the traces running over the Weyl fermions: the two gauginos with $R = 1$ (each in the adjoint, of dimension N^2-1) and the four bifundamental fermions with $R = R(A) - 1 = -\frac{1}{2}$ (each of dimension N^2). In particular,

$$\mathrm{Tr}\, R = -2, \qquad \mathrm{Tr}\, R^3 = \tfrac{3}{2}N^2 - 2. \tag{VII.1.43}$$

Substituting,

$$a_{\mathrm{KW}} = \frac{3}{32}\Big(\tfrac{9}{2}N^2 - 4\Big) \;\xrightarrow{N\to\infty}\; \frac{27}{64}\,N^2, \tag{VII.1.44}$$

for the interacting $SU(N)^2$ theory with $N \geq 2$ and no accidental R-symmetry. The same anomaly formula gives

$$c = \frac{1}{32}(9\,\mathrm{Tr}\,R^3 - 5\,\mathrm{Tr}\,R) = \frac{27}{64}N^2 - \frac{1}{4}, \qquad a - c = -\frac{1}{8}. \tag{VII.1.45}$$

These are finite-N anomaly expressions for this sector. Adding free abelian vectors would change them, and setting $N = 1$ does not describe an interacting KW fixed point. The leading term in a is the number the holographic computation of §VII.1.9 reproduces from the volume of $T^{1,1}$. This evaluation uses the R-charge already fixed; the principle that *selects* that R-charge, by maximizing the trial a over the mixing below, is a-maximization.

This is a consistency computation using the symmetric assignment, not the general derivation. Marginality of W alone fixes only the sum,

$$R(W) = 2R(A) + 2R(B) = 2 \quad \Longrightarrow \quad R(A) + R(B) = 1, \tag{VII.1.46}$$

and leaves a one-parameter family that mixes the R-symmetry with the baryonic symmetry,

$$R(A) = \tfrac{1}{2} + s, \qquad R(B) = \tfrac{1}{2} - s, \tag{VII.1.47}$$

satisfying (VII.1.46) for every s. Here the $A \leftrightarrow B$ node-exchange symmetry forces $s = 0$, and the beta function confirms it. For a general toric singularity there is no such exchange and no beta function pins the mixing; the correct R-symmetry is then selected by a-maximization, the subject of Chapter 6, where the trial a-function is maximized over exactly this kind of mixing parameter.

The two global symmetries. Two global symmetries appear and must not be confused. One is the continuous flavour and R-symmetry algebra,

$$\mathfrak{su}(2)_A \oplus \mathfrak{su}(2)_B \oplus \mathfrak{u}(1)_R. \tag{VII.1.48}$$

Its action on mesonic scalars is the geometric action (VII.1.7); the global group on all quantum operators requires additional identifications, as explained there. The other is the non-R baryonic symmetry $U(1)_B$ of the interacting $SU(N)^2$ theory, under which

$$A_i \text{ has charge } +1, \qquad B_j \text{ has charge } -1. \tag{VII.1.49}$$

Table VII.1.2 keeps them apart.

symmetry	acts as	later use
$\mathfrak{su}(2)_A \oplus \mathfrak{su}(2)_B \oplus \mathfrak{u}(1)_R$	continuous symmetries of the cone; $U(1)_R$ superconformal	a-max and volume minimization (Ch. 6)
$U(1)_B$	$A: +1$, $B: -1$; non-R baryonic	baryonic Hilbert series, master space (Ch. 5); cascade (Ch. 8)

Table VII.1.2: The two global symmetries of the Klebanov–Witten theory. The first row lists the continuous symmetry algebra; its action on mesons is geometric. The second row gives the global baryonic symmetry after restricting to the interacting SU sector.

What $U(1)_B$ counts. The baryonic symmetry has charged operators, the baryons, gauge invariant in the $SU(N)^2$ theory. Antisymmetrizing N copies of a single bifundamental over the colour indices of both nodes gives, for each flavour,

$$\mathcal{B}_A^{(i)} = \varepsilon_{\alpha_1\cdots\alpha_N}\varepsilon^{\beta_1\cdots\beta_N}(A_i)^{\alpha_1}{}_{\beta_1}\cdots(A_i)^{\alpha_N}{}_{\beta_N}, \tag{VII.1.50}$$

and similarly $\mathcal{B}_B^{(j)}$ from the B_j. These are $SU(N)^2$ gauge invariant (the two epsilons saturate the colour indices of nodes 1 and 2) and carry baryonic charge

$$Q_B(\mathcal{B}_A) = +N, \qquad Q_B(\mathcal{B}_B) = -N. \tag{VII.1.51}$$

written for a single flavour; more generally, antisymmetrizing over colour while symmetrizing the N flavour indices, the A-baryons fill the symmetric spin-$N/2$ representation of $SU(2)_A$ and the B-baryons the spin-$N/2$ of $SU(2)_B$. They are not mesons, and they are precisely what the baryonic branch of the moduli space and the baryonic Hilbert series of Chapter 5 organize; the cascade of Chapter 8 changes their ranks. The number of continuous non-anomalous baryonic $U(1)$ symmetries in this isolated toric D3 construction is read from the external toric points: with d such points the standard dictionary gives $d-3$. We use it for the conifold and the isolated dP_0 example; nonisolated singularities require additional care with boundary points and extra sectors. For the conifold $d = 4$, so there is exactly one, the $U(1)_B$ above. This count is $d-3$ and not, say, one less than the number of gauge nodes: the conifold has two nodes and one baryonic symmetry, so the two coincide here, but already for dP_0 (a triangular diagram with $d = 3$) the count is $d - 3 = 0$, no continuous non-anomalous baryonic $U(1)$, while the node count would wrongly give two. The boundary-point count is the correct invariant.

A last word on the fixed point. We have worked at the symmetric Klebanov–Witten point. More generally the theory has exactly marginal gauge and superpotential couplings, and the fixed point sits on a conformal manifold; the present chapter does not analyze that manifold, needing only the symmetric point and the R-charge that makes the quartic marginal. The general technology for exactly marginal deformations travels with the superconformal-dynamics material of Part III, while Chapter 6 uses the toric data to perform the a-maximization that selects the R-symmetry.

VII.1.8 The square dimer as the first brane tiling

The cited D3-brane construction has now passed the geometric and field-theory consistency checks above. The combinatorial repackaging that the rest of the volume is built on is the brane tiling, or dimer: a bipartite graph on a two-torus that encodes the quiver and the superpotential at once. This section names the conifold's dimer, reads the gauge theory off it, and does the conifold's Kasteleyn computation once, as a preview of the machinery built in Chapters 3–5.

The conifold's brane tiling is the **square lattice**. In one fundamental domain of the torus it has one white node, one black node, four edges, and two faces,

$$(V, E, F) = (2, 4, 2), \qquad V - E + F = 0, \tag{VII.1.52}$$

the vanishing Euler characteristic of a cellular tiling of T^2. The dictionary, used for the whole volume, assigns a piece of gauge-theory data to each cell of the tiling,

$$\text{faces} \leftrightarrow \text{gauge nodes } U(N)_a, \qquad \text{edges} \leftrightarrow \text{bifundamentals}, \qquad \text{nodes} \leftrightarrow W \text{ terms}, \tag{VII.1.53}$$

with white nodes contributing to W with one sign and black nodes with the other. So the two faces are the gauge nodes $U(N)_1, U(N)_2$; the four edges are the four bifundamentals A_1, A_2, B_1, B_2; and the two graph nodes are the two superpotential terms of opposite sign, which by the stated sign rule reproduce the two-term quartic (VII.1.25). Each edge bounds exactly two faces and meets exactly two nodes, so each bifundamental appears in exactly two superpotential terms with opposite sign, the two-term structure characteristic of toric quivers. Figure VII.1.3 draws the tiling.

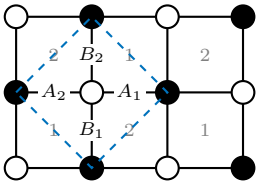


Figure VII.1.3: The conifold dimer, drawn as a patch of the square lattice on the universal cover of T^2 so the periodicity is manifest; the dashed parallelogram is one fundamental domain, with period vectors $(1,1)$ and $(1,-1)$. The vertices are black and white (the two superpotential nodes), checkerboard-coloured; the faces are the gauge nodes $1, 2$; the four edges of a cell are the bifundamentals A_1, A_2, B_1, B_2. One fundamental domain has one white and one black vertex, four edges, and two faces, so the Euler count is $V - E + F = 2 - 4 + 2 = 0$. Reading the two nodes as the two opposite-sign superpotential terms recovers the quartic (VII.1.25).

How to read the figure, since this is the first dimer in the book. First identify the faces: these are the two gauge factors. Then identify the edges: these are the four bifundamentals. Finally identify the black and white nodes: these are the two superpotential monomials, with opposite signs. The quiver and the superpotential are therefore not two separate pieces of input; both are encoded by the one graph on T^2, and this packaging is the whole point of the dimer.

It is worth extracting the superpotential by hand once. Each node of the tiling is a superpotential term, the product of the edges meeting it taken in cyclic order around the node: clockwise at a white node, counterclockwise at a black node, with a relative sign between the two colours. The white node has the four edges A_1, B_1, A_2, B_2 around it, giving the term $+\operatorname{Tr}(A_1 B_1 A_2 B_2)$; the black node has the same four edges, but the opposite cyclic reading A_1, B_2, A_2, B_1 gives $-\operatorname{Tr}(A_1 B_2 A_2 B_1)$. Their sum,

$$W = \operatorname{Tr}(A_1 B_1 A_2 B_2) - \operatorname{Tr}(A_1 B_2 A_2 B_1), \tag{VII.1.54}$$

is exactly the two-term quartic (VII.1.25) (up to the overall coupling), and the rule that each edge meets one white and one black node is why each field sits in two terms with opposite sign. The dimer has reproduced the superpotential that symmetry fixed in §VII.1.5, now as a property of a graph.

The Kasteleyn fixture. The toric output of a dimer is computed from its Kasteleyn matrix, a weighted, signed adjacency matrix between the white and black nodes in which an edge crossing the fundamental domain carries a power of the torus coordinates z, w recording its winding. For the conifold there is one white and one black node, so the Kasteleyn matrix is 1×1, a single entry summing the four edges with their windings $(0,0), (1,0), (0,1), (1,1)$,

$$K(z,w) = p_1 + p_2\, z + p_3\, w + p_4\, zw, \tag{VII.1.55}$$

with p_a independent nonzero signed edge weights, including any signs chosen in the adjacency convention. Only their nonvanishing and windings enter the present support calculation; no signed matching-count formula is inferred here. Its determinant is the Newton polynomial $P(z,w) = \det K(z,w) = K(z,w)$, whose support, the set of exponent vectors of its monomials, is

$$\{(0,0), (1,0), (0,1), (1,1)\} = \Delta_{\mathcal{C}}, \tag{VII.1.56}$$

the unit-square toric diagram (VII.1.13). The dimer thus reproduces the conifold's toric data. The perfect matchings, the sets of edges covering every node exactly once, are here just the individual edges (one white and one black node are covered by a single edge), so there are

$$\#\{\text{perfect matchings}\} = 4 = \#\{\text{lattice points of } \Delta_{\mathcal{C}}\}, \tag{VII.1.57}$$

one per corner of the square, each of multiplicity one. The equality of these two counts is special to this fixture: in a general dimer the Kasteleyn matrix is larger, the perfect matchings are more numerous, and the map from perfect matchings to lattice points is many-to-one, with multiplicities. Chapters 3 and 5 build that machinery; the conifold is the case where it is a single line. The dimer is thus a second, combinatorial presentation of the same theory: §VII.1.4–§VII.1.6 test the cited quiver and quartic against the geometry, and the square lattice encodes all of it at once, generalizing to singularities with no geometric shortcut.

VII.1.9 The holographic dual

For a Ricci-flat Calabi–Yau cone with Sasaki–Einstein link normalized by $\mathrm{Ric}(Y_5) = 4g$, the classical D3 supergravity solution has a near-horizon limit that is the product of five-dimensional anti-de Sitter space with the link,

$$N \text{ D3 at } C(Y_5) \quad \longrightarrow \quad AdS_5 \times Y_5. \tag{VII.1.58}$$

For the conifold $Y_5 = T^{1,1}$, so the dual is

$$AdS_5 \times T^{1,1}, \tag{VII.1.59}$$

the conifold analogue of the $AdS_5 \times S^5$ dual of $\mathcal{N}$=4 super Yang–Mills, and the first holographic dual with a non-spherical, non-orbifold internal space. The Sasaki–Einstein

metric on $T^{1,1}$ is the homogeneous metric on (VII.1.12),

$$\mathrm{d}s^2_{T^{1,1}} = \tfrac{1}{6}\sum_{i=1}^{2}\big(\mathrm{d}\theta_i^2 + \sin^2\theta_i\,\mathrm{d}\phi_i^2\big) + \tfrac{1}{9}\big(\mathrm{d}\psi + \textstyle\sum_{i=1}^{2}\cos\theta_i\,\mathrm{d}\phi_i\big)^2, \tag{VII.1.60}$$

quoted here with $0 \le \theta_i \le \pi$, $0 \le \phi_i < 2\pi$, and $0 \le \psi < 4\pi$. The volume form is $\sin\theta_1 \sin\theta_2\,\mathrm{d}\theta_1\mathrm{d}\phi_1\mathrm{d}\theta_2\mathrm{d}\phi_2\mathrm{d}\psi/108$. Integrating gives

$$\mathrm{Vol}(T^{1,1}) = \frac{16\pi^3}{27}, \tag{VII.1.61}$$

to be compared with $\mathrm{Vol}(S^5) = \pi^3$.

Central charge: a preview. The single number this geometry computes that the chapter can preview is the central charge a of the dual gauge theory, which at leading order for large N and in the classical supergravity regime $g_sN \gg 1$, $g_s \ll 1$, is inversely proportional to the volume of the internal space,

$$a(Y_5) = \frac{\pi^3}{4\,\mathrm{Vol}(Y_5)}\,N^2. \tag{VII.1.62}$$

For the round sphere this returns $a = \frac{1}{4}N^2$, the large-N central charge of $\mathcal{N}$=4; for the conifold it gives

$$a(T^{1,1}) = \frac{27}{64}\,N^2, \qquad \frac{a(T^{1,1})}{a(\mathcal{N}{=}4)} = \frac{27}{16}. \tag{VII.1.63}$$

This matches the field-theory anomaly $a_{\mathrm{KW}} \to \frac{27}{64}N^2$ computed in §VII.1.7, the two sides of the central charge that Chapter 6 will derive and match in general. This is a preview, not a derivation: the field-theory computation that selects the R-symmetry by maximizing the trial a-function over the mixing of (VII.1.47), and its matching to the geometric volume minimization, are the content of Chapter 6. One caveat guards a common confusion: the celebrated number 27/32 associated with the conifold is *not* the ratio (VII.1.63) to $\mathcal{N}$=4; it is the leading-large-N ratio of the conifold central charge to that of the $\mathbb{C}^2/\mathbb{Z}_2$ orbifold theory along the Klebanov–Witten renormalization-group flow, a decrease consistent with the a-theorem. The Kaluza–Klein spectrum, the subleading-in-N corrections, and the full a-maximization all belong to Chapter 6 and Part III.

VII.1.10 What has been computed, and what is deferred

Starting from the cited Klebanov–Witten D3 construction, the chapter has checked its geometric and field-theory consequences. The conifold $\{xy = zw\} = \{\det\mathcal{Q} = 0\}$ is the variety of rank-one matrices, the cone over the smooth Sasaki–Einstein space $T^{1,1} = (S^3 \times S^3)/U(1) = S^2 \times S^3$ with isometry $SU(2)_A \times SU(2)_B \times U(1)_R$, a toric Calabi–Yau threefold with the unit-square toric diagram, and it is not an orbifold. A

stack of N D3-branes at the tip carries the Klebanov–Witten theory: gauge group $U(N)_1 \times U(N)_2$ (interacting sector $SU(N) \times SU(N)$ with a baryonic $U(1)_B$), four bifundamentals in $SU(2)_A \times SU(2)_B$ doublets, the flavour-invariant single-trace quartic superpotential (VII.1.23), and its F-terms. The one-brane mesonic moduli space at zero FI is the conifold as a ring statement (VII.1.34); the beta-function condition at the assumed symmetric $N \geq 2$ fixed point gives $R = \frac{1}{2}$ (VII.1.41); the baryons carry $U(1)_B$ charge $\pm N$; the square dimer reproduces the quiver and quartic and returns the unit-square toric diagram through its Kasteleyn determinant; and the holographic dual is $AdS_5 \times T^{1,1}$ with the central charge previewed.

A reader leaving the chapter should be able to do the following without looking ahead: write the conifold as $\det \mathcal{Q} = 0$ and identify its isometry; present it as the $\mathbb{C}^*$ quotient (VII.1.9) and its link $(S^3 \times S^3)/U(1)$; draw the two-node $U(N)_1 \times U(N)_2$ quiver with its four bifundamentals and check that it is anomaly-free; enumerate the single-trace quiver loops and write the invariant quartic; compute its four F-terms and see them vanish for one brane; prove the one-brane mesonic moduli space is the conifold by the invariant ring; run the NSVZ check to $\gamma = -\frac{1}{2}$ and $R = \frac{1}{2}$; separate $U(1)_R$ from $U(1)_B$ and write the baryons; and read the square dimer, recovering both superpotential terms and the unit-square Newton polygon.

What has been deferred, and to where, is the map of the rest of the volume.

- The general selection of R-charges by a-maximization, and the matching to the Sasaki–Einstein volume minimization that the central-charge preview (VII.1.63) anticipates, are Chapter 6. The conifold's $R = \frac{1}{2}$ was fixed here by symmetry and the beta function only because the $A \leftrightarrow B$ exchange removed the mixing parameter.

- The dimer was named and its conifold Kasteleyn fixture computed once. The general Kasteleyn matrix, the enumeration of perfect matchings with multiplicities, and the forward algorithm from a tiling to a toric diagram are Chapters 3–5; the physical origin of the tiling as an NS5–D5 brane configuration obtained by two T-dualities is Chapter 4.

- Only the mesonic branch was computed. The baryonic branch and the master space, organized by $U(1)_B$ and its baryons, are Chapter 5.

- The holographic dual was named, its metric and volume quoted, and its central charge previewed; every gravity-side computation beyond (VII.1.63), and the precision spectrum, are Chapter 6 and Part III.

- The conifold's small resolution and deformation were named as its two smooth neighbours; their geometry and the Sasaki–Einstein metric are Part I. Fractional branes wrapping the resolution $\mathbb{P}^1$, the duality cascade, and the deformed-conifold infrared (confinement and chiral-symmetry breaking) are Chapter 8. This chapter carries only the regular brane, with equal ranks $N_1 = N_2 = N$, at the interacting conformal point.

The single lesson to carry forward is the one in the keybox at the start: with no finite group, the cited gauge theory was tested against the geometry and its symmetry, and

then encoded combinatorially by the dimer. The next chapter makes the first half of that statement general (geometry to gauge theory for an arbitrary toric diagram), and the chapters after it make the second half general (the dimer as the computational method).

Status and sources

The evidence-code tags ([WS] physical-worldsheet derivation in the untwisted string, [MATH] mathematical result with stated hypotheses and proof basis), the convention pointers, and the source anchors are collected here, keyed by section. The chapter develops the conifold as the first non-orbifold toric probe, after the gauge theory of Klebanov and Witten (1998), the conifold geometry of Candelas and Ossa (1990), the toric moduli-space derivation of Morrison and Plesser (1999), and the volume and metric of $T^{1,1}$ of Gubser (1999); it forward-references the brane-tiling corpus of Hanany and Kennaway (2005) and Franco et al. (2006) for the dimer, and Klebanov and Strassler (2000) for the cascade. It reuses the toric-diagram conventions (toric lattice orientation, the conifold unit-square fixture), the dimer/brane-tiling conventions (bipartite tiling conventions, the conifold square-lattice fixture and the superpotential sign rule), the superpotential store (superpotential normalization, the quartic), the gauge/rank data (quiver gauge groups and ranks, $\prod_i U(N_i) \to SU(N)^2 + U(1)_B$), the arrow convention, and the supercharge count. No new convention is introduced. The finite tiling calculations concern the conifold example.

§VII.1.1 Leaving the orbifold world. Framing: finite quotients as a class distinct from the conifold, the conifold as the simplest non-orbifold toric CY_3, the build-it-twice promise (gauge theory and dimer), and the reader contract. Files (the conifold is not an orbifold; the Klebanov–Witten quiver is not a McKay quiver; refines the stated distinction).

§VII.1.2 The conifold as a hypersurface. The quadric (VII.1.1), the singular locus from $\nabla f = 0$ (VII.1.2), the rank-one matrix form (VII.1.4)–(VII.1.5), the isometry $SU(2)_A \times SU(2)_B \times U(1)_R$ (VII.1.6), and the faithful geometric action (VII.1.7) with its four-element kernel (derived here, restricted to mesonic coordinates). (split mathematical and cited physical claims) (toric lattice orientation; Candelas and Ossa (1990), Klebanov and Witten (1998)). the toric diagram is the unit square with $2\,\mathrm{Area} = 2$ and $d = 4$, and is not the $\mathbb{C}^3$ triangle; the $\mathbb{Z}_2$ centre acts as $+1$ on the coordinates $M_{ij} = A_i B_j$ and as -1 on the elementary doublets.

§VII.1.3 The symplectic quotient, the link, and the toric square. The rank-one outer product (VII.1.8), the $\mathbb{C}^*$ quotient (VII.1.9), the real moment map (VII.1.10) naming the small resolution, the link $T^{1,1} = (S^3 \times S^3)/U(1) = (SU(2) \times SU(2))/U(1)$ (VII.1.11)–(VII.1.12) with base $S^2 \times S^3$, and the unit-square toric diagram (VII.1.13)–(VII.1.14). (split mathematical and cited physical claims) (toric lattice orientation; Candelas and Ossa (1990)). The conifold geometry, the Sasaki–Einstein metric, and the two smoothings are conifold geometry (home Vol. IV; the SE side recurs in Chapter 6), cited not constructed. the unit-square toric diagram.

§VII.1.4 The D3 probe and the two-node quiver. The colour spaces (VII.1.15), the maps $A_i : V_2 \to V_1$, $B_j : V_1 \to V_2$ (VII.1.16), the representations with bars (VII.1.17) (Table VII.1.1), the gauge group (VII.1.18), the abelian split $U(1)_+ \times U(1)_-$ with the centre of mass decoupling and the distinction between the relative gauge quotient and global $U(1)_B$ of the SU sector (VII.1.19)–(VII.1.21), and the conjugate-representation anomaly cancellation (VII.1.22). (physical classification unresolved; quiver gauge groups and ranks, fundamental-at-tail arrows, spinor and supersymmetry counting; Klebanov and Witten (1998), Morrison and Plesser (1999)). Files (the gauge group is not a bare SU; the baryonic $U(1)_B$ exists, with non-anomalous count $d - 3$). $G = 2$ gauge nodes, four-valent dimer, each node anomaly-free with $\#\mathrm{fund} - \#\mathrm{antifund} = 0$.

§VII.1.5 The quartic superpotential and its F-terms. The loop enumeration ruling out mass and cubic terms, the single-trace invariant quartic (VII.1.23), its expansion to the two-term form (VII.1.24)–(VII.1.25), the determinantal form (VII.1.26), and the four explicit F-terms (VII.1.27)–(VII.1.28) with their $N = 1$ vanishing. (physical classification unresolved; superpotential normalization; Klebanov and

Witten (1998)). W collapses to the meson determinant and is degree two not cubic; the four F-terms equal the cyclic derivatives entrywise; the $N = 1$ F-terms vanish for commuting fields.

§VII.1.6 The one-brane mesonic moduli space. The relative $U(1)_-$ charges (VII.1.29), the invariant ring generated by the four mesons (VII.1.30), the single relation (VII.1.31), the ring (VII.1.32), the identification with $xy - zw = 0$ and Spec (VII.1.33)–(VII.1.34), and Sym^N for N branes (VII.1.36). (physical classification unresolved; Klebanov and Witten (1998), Morrison and Plesser (1999)). Carries (the probe variety carries the α'/g_s caveats of Volume VI). The baryonic branch and master space are routed to Chapter 5. $\det M = 0$ identically for the rank-one $N = 1$ meson matrix, and $xy - zw = 0$; the forward algorithm gives $\dim\ker Q_t = 3$.

§VII.1.7 Marginality, beta functions, R-charges, and global symmetries. The SCFT relation $\Delta = \frac{3}{2}R$ (VII.1.38); the NSVZ condition at the assumed symmetric fixed point giving $\gamma = -\frac{1}{2}$ (VII.1.39)–(VII.1.40) and hence $R = \frac{1}{2}$ (VII.1.41); marginality $R(W) = 2$ and the trial mixing $R = \frac{1}{2} \pm s$ fixed by node exchange (a-maximization deferred to Chapter 6) (VII.1.46)–(VII.1.47); the geometric isometry (VII.1.48) versus the baryonic $U(1)_B$ (VII.1.49) (Table VII.1.2); the baryons (VII.1.50) with charge $\pm N$ (VII.1.51); and the continuous non-anomalous baryonic count $d - 3$ for the isolated examples. (physical classification unresolved; spinor and supersymmetry counting; Klebanov and Witten (1998)). The a-maximization derivation is Chapter 6. the NSVZ numerator $N(1 + 2\gamma)$ giving $\gamma = -\frac{1}{2}$, then $\Delta = \frac{3}{4}$, $R = \frac{1}{2}$, $\Delta(W) = 3$, $R(W) = 2$; the baryon charges $\pm N$; $R(W) = 2$ for any s; and the baryonic count $d - 3 = 1$, equal to $G - 1$ for the conifold but not in general (dP_0: $d - 3 = 0 \neq G - 1 = 2$).

§VII.1.8 The square dimer fixture. The square-lattice tiling: $(V, E, F) = (2, 4, 2)$, $V - E + F = 0$ (VII.1.52); the faces/edges/nodes dictionary (VII.1.53) reproducing the two-term quartic; the figure-reading bridge; the 1×1 Kasteleyn matrix (VII.1.55), its Newton polygon = the unit square (VII.1.56), and the four perfect matchings = the four corners (VII.1.57), with the general many-to-one multiplicity map guarded to Chapter 5. (physical classification unresolved; bipartite tiling conventions, the conifold square-lattice fixture; Klebanov and Witten (1998), Hanany and Kennaway (2005), Franco et al. (2006)). The Kasteleyn matrix, perfect-matching enumeration, and forward algorithm machinery are Chapters 3–5. $V - E + F = 0$ on T^2, both nodes four-valent, $\det K$ Newton polygon = the conifold toric diagram, four perfect matchings each of multiplicity one.

§VII.1.9 The holographic dual. The near-horizon $AdS_5 \times Y_5$ (VII.1.58), specialized to $AdS_5 \times T^{1,1}$ (VII.1.59); the quoted Sasaki–Einstein metric (VII.1.60) and volume $\mathrm{Vol}(T^{1,1}) = 16\pi^3/27$ (VII.1.61); and the central-charge *preview* $a(Y_5) = \pi^3 N^2/(4\mathrm{Vol})$ giving $a(T^{1,1}) = \frac{27}{64}N^2$ and the ratio 27/16 to $\mathcal{N}$=4 (VII.1.62)–(VII.1.63), with the leading-large-N 27/32 flow-ratio caveat. (split mathematical and cited physical claims) (Gubser (1999), Klebanov and Witten (1998)). The a-maximization derivation of the central charge, the Kaluza–Klein spectrum, and the subleading-N corrections are deferred to Chapter 6 and Part III; VII.1 previews the central charge from the volume and defers the derivation. evaluated(sanity, as a preview): $a(T^{1,1})/a(\mathcal{N}{=}4) = \mathrm{Vol}(S^5)/\mathrm{Vol}(T^{1,1}) = 27/16$ and $a(T^{1,1}) = \frac{27}{64}N^2$.

§VII.1.10 What is and is not given. Scope and limitations. the chapter builds the conifold geometry, the Klebanov–Witten quiver and quartic with F-terms, the mesonic moduli space as a ring, the conditional fixed-point R-charges, the baryons, the square dimer with its Kasteleyn fixture, and the named dual with a central-charge preview; it routes the a-maximization derivation to Chapter 6, the dimer machinery to Chapters 3–5, the baryonic branch to Chapter 5, the cascade and deformed-conifold infrared to Chapter 8, the conifold geometry and metric to Part I (conifold geometry), and the precision holographic spectrum to Part III.

Evidence and assumptions. WS denotes an actual physical-worldsheet derivation in the untwisted string; none is supplied by the algebraic checks in this chapter. The D3 identification and interacting fixed-point existence remain cited physical premises with null primary classification in the finite calculation; the calculations supply only the stated WS certificates. The invariant-ring result is MATH over $\mathbb{C}$ with the generation, height-one prime, and irreducibility argument in §VII.1.6. The cyclic-word identities are exact algebra once the quartic is specified. The anomaly coefficients (VII.1.43), and a, c obtained through the supersymmetric anomaly relations, carry ANOM with two coefficients $\mathrm{Tr}\, R$ and $\mathrm{Tr}\, R^3$ evaluated; this is not two independently matched physical constructions. Classical geometry and the leading gravitational comparison use SUGRA premises, not a worldsheet derivation. The companion grades three bounded submitted-answer families; it does not certify the physical premises.

Original scripts remain separate legacy computations, including the centre-action checks that do not establish the former combined global quotient.

Precise source locations. In Klebanov and Witten (1998), equations (10)–(16) describe the conifold metric, quotient and link, (17) gives the quartic, (18)–(20) the fixed-point conditions, and (34) the F-terms. Its coefficient $\lambda_{\mathrm{KW}}/2$ in front of the epsilon sum corresponds to $\lambda_{\mathrm{KW}} = 2\lambda$ in (VII.1.23); hence the present $h = 2\lambda$ is unchanged. In that source, the paragraphs "Global Structure Of The Symmetry Group" on page 10 concern the flavour $SO(4)$ image. The discussion on pages 7–9 supplies the physical fixed-point premise, not a worldsheet calculation. Its anomalous-dimension convention differs by a factor of two from (VII.1.39). The volume comparison and leading anomaly ratio are in Gubser (1999), section 2 and section 2.1, equations (15)–(19). The scalar mesonic quotient is also discussed in Morrison and Plesser (1999), section 4.2, Table 4 and Figure 3; this chapter's ring proof is given explicitly above. For the dimer support and matching construction see Hanany and Kennaway (2005), section 3, and Franco et al. (2006), section 2. The general finite-N a, c anomaly relations used here are stated in Intriligator and Wecht (2003), equation (1.2).

Related constructions. D3-branes at the conifold give the Klebanov–Witten quiver $U(N)_1 \times U(N)_2$ with its quartic superpotential. Its mesonic equations give the conifold $\det M = 0 = xy - zw$. The conformal R-charge is $R = \frac{1}{2}$, and the proposed gravity description is $AdS_5 \times T^{1,1}$. Chapter 3 gives the conifold square dimer and the Kasteleyn and perfect-matching construction. The separate orbifold comparison is developed in Chapters 2 and 5.

References

Candelas, P. and X. C. de la Ossa (1990). "Comments on Conifolds". In: *Nucl. Phys. B* 342, pp. 246–268. DOI: 10.1016/0550-3213(90)90577-Z.

Franco, S. et al. (2006). "Brane dimers and quiver gauge theories". In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Gubser, S. S. (1999). "Einstein manifolds and conformal field theories". In: *Phys. Rev. D* 59, p. 025006. DOI: 10.1103/PhysRevD.59.025006. arXiv: hep-th/9807164.

Hanany, A. and K. D. Kennaway (2005). "Dimer models and toric diagrams". In: arXiv: hep-th/0503149.

Intriligator, K. A. and B. Wecht (2003). "The Exact superconformal R symmetry maximizes a". In: *Nucl. Phys. B* 667, pp. 183–200. DOI: 10.1016/S0550-3213(03)00459-0. arXiv: hep-th/0304128.

Klebanov, I. R. and M. J. Strassler (2000). "Supergravity and a confining gauge theory: Duality cascades and chi SB resolution of naked singularities". In: *JHEP* 08, p. 052. DOI: 10.1088/1126-6708/2000/08/052. arXiv: hep-th/0007191.

Klebanov, I. R. and E. Witten (1998). "Superconformal field theory on three-branes at a Calabi-Yau singularity". In: *Nucl. Phys. B* 536, pp. 199–218. DOI: 10.1016/S0550-3213(98)00654-3. arXiv: hep-th/9807080.

Morrison, D. R. and M. R. Plesser (1999). "Nonspherical horizons. 1." In: *Adv. Theor. Math. Phys.* 3, pp. 1–81. DOI: 10.4310/ATMP.1999.v3.n1.a1. arXiv: hep-th/9810201.

Chapter VII.2

General toric Calabi–Yau threefold probes

VII.2.1 From one example to a general rule

The previous chapter built the Klebanov–Witten theory of a D3-brane probe at the conifold without ever invoking a finite group. The quiver, the superpotential, the moduli space, and the R-charges all followed from the geometry and the symmetry of a single non-orbifold cone. That worked because the conifold is exceptionally symmetric: a continuous $SU(2) \times SU(2) \times U(1)$ isometry did the work that the representation theory of Γ did in Volume V. The generic toric Calabi–Yau threefold has no such large symmetry. It has, at most, the abelian torus that every toric variety carries, and that alone does not pin down a quiver. A new organizing object is needed, one that is combinatorial rather than group-theoretic, and it is the subject of this chapter: the toric diagram, a convex lattice polygon, together with its planar dual, the (p, q)-web.

This chapter is the statement layer of the volume. It says *what* is to be constructed for an arbitrary toric cone and *why* the answer is not unique, leaving the machine that does the construction to the chapters that follow. The conifold of Chapter 1 is now read as the simplest non-trivial instance of the general rule rather than as a special miracle, and a second worked instance, the orbifold $\mathbb{C}^3/\mathbb{Z}_3$ presented for the first time as a toric polygon, ties the toric language back to the McKay quivers of Volume V. The two examples are the bridge: the conifold is the non-orbifold anchor, $\mathbb{C}^3/\mathbb{Z}_3$ is the orbifold one, and the same toric-diagram language organizes both (their diagrams differ, a square and a triangle; the *type* of geometric description is shared).

Take-home message of this chapter

A toric Calabi–Yau threefold is presented by a single combinatorial object, a convex lattice polygon Δ at height one (the toric diagram). For a consistent brane tiling describing regular D3-branes at that cone, the number of gauge nodes is $2\,\mathrm{Area}(\Delta)$; arrows are bifundamentals (including adjoints), and each field appears in one positive and one negative superpotential term. On the geometric mesonic branch, the moduli space of one regular brane reproduces the affine cone. The new phenomenon, not visible in the elementary McKay examples used here, is that this geometry-to-gauge map is *not one-to-one*: one polygon admits several inequivalent quivers, related by Seiberg duality. *The polygon fixes the geometry and the toric data; it does not fix a unique quiver.*

The plan is to make each clause of that box precise, then defer its proof to the right chapter. A word first on what is assumed and what is merely named. The chapter assumes the four-dimensional $\mathcal{N}{=}1$ quiver vocabulary of Volume VI and Chapter 1: gauge nodes, bifundamental arrows, the superpotential and its F-terms, gauge-invariant mesons, and the probe prescription that reads the transverse geometry from a brane's geometric moduli component. From toric geometry it uses the diagram-and-web presentation and a chamber-sensitive Cox/GIT quotient; the fan, polytope, and quotient foundations themselves are cited to Part I, not proved here. The dimer, the Kasteleyn matrix, perfect matchings, the Hilbert series, a-maximization, and the explicit inverse algorithm all appear in this chapter as *names and motivations*; each is built in the chapters that follow, and nothing here requires computing with them.

The chapter proceeds as follows. The toric diagram and its (p, q)-web are defined and their invariants extracted (§VII.2.2); the same data is recast as the gauged linear sigma model, the cone as a quotient of $\mathbb{C}^k$ (§VII.2.3); the D3 probe is placed and its quiver, gauge group, and toric superpotential are stated (§VII.2.4); the one-brane mesonic moduli space is shown to recover the cone (§VII.2.5); the inverse problem is posed and its non-uniqueness named as toric duality (§VII.2.6); the two bridge examples, $\mathbb{C}^3/\mathbb{Z}_3 = dP_0$ as a polygon and the conifold recalled, are worked (§VII.2.7); and a closing section maps what was computed and what is deferred (§VII.2.8). The organizing contrast with Volume VI runs throughout: there the input was a finite group and its character table, here it is a lattice polygon and its area.

VII.2.2 The toric diagram and the (p, q)-web

A toric Calabi–Yau threefold is an affine cone whose geometry is encoded by a fan in the three-dimensional lattice $\mathbb{Z}^3$. The cone is generated by a set of rays, and the single fact that makes the cone Calabi–Yau is a statement about where those rays point.

The height-one (Calabi–Yau) condition. Write the primitive generators of the extreme rays as $r_a \in \mathbb{Z}^3$. For a strongly convex, full-dimensional normal affine toric cone, its canonical divisor is Cartier and principal precisely when there exists an *integral* covector $n \in (\mathbb{Z}^3)^*$ with

$$n \cdot r_a = 1 \qquad \text{for every ray } a. \tag{VII.2.1}$$

Such a covector is automatically primitive: if $n = mn_0$ with $|m| > 1$, its pairing with an integral ray cannot be one. A unimodular basis change therefore takes n to $(0, 0, 1)$ and

r_a to $(v_a, 1)$. The associated character trivializes the canonical sheaf on the normal affine toric variety; on its smooth locus this is the nowhere-vanishing holomorphic three-form. The singular cone itself need not be smooth, and the holonomy/supercharge statement belongs to the D3 probe dictionary rather than to this lattice proof. Testing one chosen covector is only a necessary check for that covector: $r = (1, 1, 2)$ fails the *chosen* $(0, 0, 1)$ test, but one must solve for every integral n before concluding that height one is impossible. For example, $(0, 0, 1), (1, 0, 1), (3, 3, 2), (0, 1, 1)$ are four primitive extreme rays: their intersections with $z = 1$ form a strictly convex quadrilateral. The first, second and fourth rays force $n = (0, 0, 1)$, which pairs to two with the third. Thus no height-one covector exists. By contrast, adding $(1, 1, 2) = (1, 0, 1) + (0, 1, 1)$ to the first, second and fourth rays does not change their smooth cone; a redundant generator cannot obstruct its canonical class. Here (VII.2.1) is the working height-one condition; its canonical class proof is referred to Part I.

The toric diagram. For the affine cone, its extreme rays are recovered from the two-dimensional picture of the points $v_a \in \mathbb{Z}^2$. The convex hull of those points is the **toric diagram**

$$\Delta = \mathrm{conv}\{v_a\} \subset \mathbb{R}^2, \tag{VII.2.2}$$

a convex lattice polygon. Its lattice points come in two kinds, and both carry data: the boundary points (on the edges of Δ) and the interior points. The polygon is read up to the lattice automorphisms $GL(2, \mathbb{Z})$ and translations, which act as unimodular shears and shifts of $\mathbb{Z}^2$; two diagrams related by such a transformation describe the same cone. The quantities that survive this freedom are the genuine invariants, and three integers read off Δ control the entire chapter:

$$G = 2\,\mathrm{Area}(\Delta), \qquad d = \#\{\text{boundary points}\}, \qquad I = \#\{\text{interior points}\}. \tag{VII.2.3}$$

Here Area is the ordinary Euclidean area measured in lattice units, so the unit square has area 1 and the unit triangle area $\frac{1}{2}$ ($G = 1$); the integer invariant the gauge theory reads is the *normalized area* $2\,\mathrm{Area}(\Delta)$, twice that. The first integer becomes the number of faces, hence gauge nodes, in a consistent brane tiling of the toric phase. The second counts external web legs *with lattice multiplicity*; for the standard nondegenerate regular-brane theory it gives $d - 3$ non-anomalous baryonic currents. The third counts compact toric divisors in a crepant resolution that uses all interior lattice points as rays. These physical readings require the stated tiling and resolution premises; Pick alone establishes only the polygon identity. That these three are $GL(2, \mathbb{Z})$ invariants is what makes them legitimate gauge-theory data: G is a unimodular invariant because a shear of determinant one preserves area, and the boundary and interior point counts are preserved because a lattice automorphism maps lattice points to lattice points and the boundary to the boundary. The internal consistency of the two point counts with the area is Pick's theorem,

$$\mathrm{Area}(\Delta) = I + \tfrac{d}{2} - 1, \qquad \text{equivalently} \qquad G = 2I + d - 2, \tag{VII.2.4}$$

which every toric diagram satisfies and which ties the node count G to the point counts. Table VII.2.1 runs Pick's theorem on the three fixtures, so that the area, the two point counts, and the node count are seen to be consistent before any gauge theory is built:

in each row the independently counted $2I + d - 2$ equals the independently measured $G = 2\,\mathrm{Area}$.

polygon	Area	d	I	$2I + d - 2$	$G = 2\,\mathrm{Area}$
$\mathbb{C}^3$ unit triangle	$\frac{1}{2}$	3	0	1	1
conifold unit square	1	4	0	2	2
dP_0 local-$\mathbb{P}^2$ triangle	$\frac{3}{2}$	3	1	3	3

Table VII.2.1: Pick's theorem (VII.2.4) checked on the three example cones. The boundary count d and the interior count I are read off the polygon, the area is computed directly, and the two sides of $G = 2\,\mathrm{Area} = 2I + d - 2$ agree in every row. The unit triangle, unit square, and local-$\mathbb{P}^2$ triangle give $G = 1, 2, 3$.

The (p, q)-web. The polygon has a planar dual, the (p, q)-web, which is the geometric face of the gauge theory's external structure and the object that the brane picture of Chapter 4 will realize physically. The web is the *perpendicular* dual: to each *primitive segment* of the boundary of Δ (the unit pieces between consecutive boundary lattice points) it assigns an external leg running off to infinity along the *outward primitive normal*, so that the number of legs equals the number d of boundary lattice points, and to each triangle of a chosen triangulation of Δ a trivalent junction. The leg vector is the (p, q) charge of a five-brane in the dual configuration. One consistency condition is the closure of the web,

$$\sum_{\text{external legs}} (p_\ell, q_\ell) = (0, 0), \tag{VII.2.5}$$

the statement that the outward normals of a closed convex polygon, weighted by edge length, sum to zero, so that the total five-brane charge of the web cancels. This is genuinely a perpendicular, not a parallel, dual: each leg satisfies $(p_\ell, q_\ell) \cdot e_\ell = 0$ against its edge vector e_ℓ, where the parallel choice would give $e_\ell \cdot e_\ell \neq 0$. Per-edge perpendicularity rejects the parallel choice; primitive charges with the correct multiplicity and outward orientation select the external web. Closure alone does not: the oriented edges of a closed polygon also sum to zero, as do the inward normals. Equation (VII.2.5) checks net charge and detects an omitted leg.

How to read Figure VII.2.1, since the diagram-and-web pair recurs in every later chapter. The full lattice polygon is the geometric input; area and point counts are useful $GL(2, \mathbb{Z})$ invariants, but do not classify it. For example, the triangle with vertices $(0, 0), (4, 0), (0, 2)$ and the square with vertices $(0, 0), (2, 0), (2, 2), (0, 2)$ both have $(2\,\mathrm{Area}, d, I) = (8, 8, 1)$. Their different numbers of extreme vertices rule out an affine unimodular equivalence. The right web is the polygon's perpendicular dual, with one leg per primitive boundary segment; its legs close, and their directions are the five-brane charges that Chapter 4 will build out of NS5- and D5-branes. The polygon and the web are two presentations of the toric geometry. A consistent tiling supplies the additional data needed to read a particular gauge theory from it.

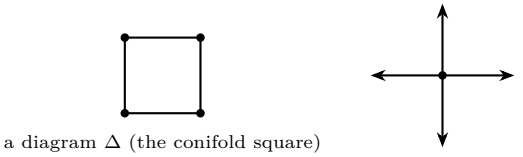


Figure VII.2.1: The toric diagram and its (p, q)-web, drawn for the conifold square as the running illustration. Read the left panel as the toric input, a convex lattice polygon; the boundary and interior points and the (normalized) area are its invariants. Read the right panel as its perpendicular dual: one external leg along the outward normal of each primitive boundary segment. The legs sum to zero (closure), which is the no-net-charge condition. The unresolved square web is drawn here as one four-valent junction with four external legs; choosing a triangulation, equivalently one of the two small resolutions, splits that four-valent vertex into two trivalent junctions joined by an internal segment. This chapter uses only the external legs; the resolved internal segment is deferred.

VII.2.3 The cone as a quotient: the gauged linear sigma model

The toric diagram has a constructive counterpart that returns directly as part of the gauge theory, exactly as the conifold's $\mathbb{C}^*$ quotient did in Chapter 1. An affine toric variety with k rays is the quotient of $\mathbb{C}^k$, one coordinate per ray, by an abelian group acting through the charges that the rays fail to be independent by. This is the gauged linear sigma model presentation, and it is the bridge from the polygon to the field theory.

Collect the ray generators as the columns of an integer matrix,

$$G_\Sigma = \begin{pmatrix} v_1 & v_2 & \cdots & v_k \\ 1 & 1 & \cdots & 1 \end{pmatrix} \in \mathbb{Z}^{3\times k}, \tag{VII.2.6}$$

the bottom row of ones being the height-one condition (VII.2.1). When the chosen rays span rank three, there are $k-3$ independent integral relations. Choose a *saturated integer basis* of this kernel as the rows of the GLSM charge matrix,

$$\mathrm{row}_{\mathbb{Z}}(Q) = \ker_{\mathbb{Z}} G_\Sigma, \qquad Q \in \mathbb{Z}^{(k-3)\times k}, \quad Q\,G_\Sigma^{\mathsf{T}} = 0. \tag{VII.2.7}$$

For a chosen fan or subdivision and FI parameter ζ in a stability chamber, the connected GLSM quotient is written

$$X_{\Sigma,\zeta} = \mathbb{C}^k /\!/_{Q,\zeta}\, U(1)^{k-3}, \tag{VII.2.8}$$

whose expected complex dimension is $k-(k-3)=3$. This formula suppresses a finite quotient if the selected rays generate a proper sublattice of the ambient lattice. At zero FI parameter the affine quotient gives the singular cone for the examples below; a nonzero chamber can give a crepant resolution. The quotient and its irrelevant set must therefore be specified before identifying a smooth space. The Calabi–Yau condition

reappears at the level of the charges as the statement that each $U(1)$ acts with charges summing to zero,

$$\sum_{a=1}^{k} Q_{Ia} = 0 \qquad \text{for each } I = 1, \dots, k-3, \tag{VII.2.9}$$

which is the bottom row of (VII.2.7): Q annihilates the all-ones vector because that vector is the height-one row of G_Σ. The conifold of Chapter 1 is the smallest example, $k = 4$ rays and one charge $Q = (+1, +1, -1, -1)$, summing to zero, and the quotient $\mathbb{C}^4 /\!/ \mathbb{C}^*$ recovered there is exactly (VII.2.8) with a single $U(1)$.

Solving two kernels by hand. The kernel is small enough to solve directly, and doing so once fixes the method. For the conifold, with rays $(0,0,1), (1,0,1), (1,1,1), (0,1,1)$ as the columns of G_Σ, a charge $q = (q_1, q_2, q_3, q_4)$ in the kernel must annihilate each row,

$$q_2 + q_3 = 0, \qquad q_3 + q_4 = 0, \qquad q_1 + q_2 + q_3 + q_4 = 0, \tag{VII.2.10}$$

the first two from the planar coordinates and the last from the height-one row. Solving gives $q_2 = -q_3$, $q_4 = -q_3$, and then $q_1 = q_3$, so one integer basis is

$$q = (1, -1, 1, -1), \tag{VII.2.11}$$

which is $(+1, +1, -1, -1)$ after reordering the corners, the conifold charge of Chapter 1. For dP_0, with rays $(1,0,1), (0,1,1), (-1,-1,1), (0,0,1)$, the same conditions read

$$q_1 - q_3 = 0, \qquad q_2 - q_3 = 0, \qquad q_1 + q_2 + q_3 + q_4 = 0, \tag{VII.2.12}$$

so $q_1 = q_2 = q_3$ and $q_4 = -3q_3$, giving

$$q = (1, 1, 1, -3). \tag{VII.2.13}$$

The charge defines a $\mathbb{C}^*$ action on the four coordinates (u_1, u_2, u_3, p), one per ray,

$$(u_1, u_2, u_3, p) \;\sim\; (\lambda u_1, \lambda u_2, \lambda u_3, \lambda^{-3} p), \qquad \lambda \in \mathbb{C}^*, \tag{VII.2.14}$$

and the resolved geometry is read off it: the locus $p = 0$ with $(u_1, u_2, u_3) \neq (0,0,0)$, quotiented by λ, is a $\mathbb{P}^2$, the compact divisor of the crepant resolution, and in the resolved chamber the cone is the total space of $\mathcal{O}_{\mathbb{P}^2}(-3)$ with that $\mathbb{P}^2$ as zero section. At the origin of the Fayet–Iliopoulos space the $\mathbb{P}^2$ collapses to the tip of the singular cone. The -3 of the interior ray is exactly the degree of the normal bundle of this divisor.

Table VII.2.2 runs the construction on the three fixtures, so the reader can build a ray matrix and read its kernel by hand. The one subtlety it makes explicit is dP_0: its toric diagram has three *external* vertices but *four* rays, because the interior lattice point $(0,0)$ is also a ray. That interior ray carries charge -3 in the resolved GLSM ($k = 4$, one connected $U(1)$). The singular cone already exists with just its three extreme rays: those rays span an index-three sublattice, so the corresponding affine Cox quotient has a residual $\mathbb{Z}_3$ action. The interior ray makes the crepant resolution visible; it is not a fourth extreme ray of the affine cone. Geometrically the interior point becomes an exceptional ray of the crepant resolution; the singular cone is recovered at the origin of the Fayet–Iliopoulos space, where that $\mathbb{P}^2$ divisor collapses.

cone	height-one rays (columns of G_Σ)	charge matrix $Q = \ker_{\mathbb{Z}} G_\Sigma$	check
$\mathbb{C}^3$	$(0,0,1)$, $(1,0,1)$, $(0,1,1)$	none ($k=3$)	no quotient; $X=\mathbb{C}^3$
conifold	$(0,0,1)$, $(1,0,1)$, $(1,1,1)$, $(0,1,1)$	$(+1,-1,+1,-1)$, i.e. $(+1,+1,-1,-1)$ up to corner order	$\sum_a Q_a = 0$; one $U(1)$
dP_0	$(1,0,1)$, $(0,1,1)$, $(-1,-1,1)$, $(0,0,1)$	$(+1,+1,+1,-3)$	$\sum_a Q_a = 0$; interior ray charge -3

Table VII.2.2: The gauged-linear-sigma-model data of the three fixtures, each computed as the integer kernel of the height-one ray matrix (VII.2.6). $\mathbb{C}^3$ has three independent rays and no quotient; the conifold has one charge $(+1,+1,-1,-1)$ (up to the ordering of the four corners); dP_0 has four rays, the three corners plus the interior point $(0,0)$, and the single charge $(+1,+1,+1,-3)$. Every charge row sums to zero, the Calabi–Yau condition (VII.2.9).

The point to register is structural. The polygon, ray matrix, and charges present related toric data, with the chosen subdivision and FI chamber deciding which variety the GLSM describes. Its homogeneous coordinates are auxiliary GLSM variables; they are not generally the D3 quiver arrows, and the D3 theory is not obtained by simply promoting these $U(1)$ factors to $U(N_i)$. A tiling or another physical construction is needed for that step. The toric-geometry theorems behind (VII.2.8), that this quotient is the affine toric variety of the fan and that its singularities are read off the cones of the fan, are foundational and are cited to Part I; here the quotient is used as the geometric target the brane must reproduce.

VII.2.4 The D3 probe: quiver, gauge group, and the toric superpotential

Place a stack of N D3-branes at the tip of the toric cone X_Δ and read off the four-dimensional $\mathcal{N}{=}1$ theory on their worldvolume. As in Chapter 1 the construction outputs a quiver, but now the polygon constrains a family of toric phases rather than selecting a unique quiver. A consistent brane tiling selects the arrows and superpotential of one phase.

The gauge group. For a consistent periodic bipartite tiling of a toric phase, the number of faces is the first invariant of (VII.2.3),

$$\#\{\text{gauge nodes}\} = G = 2\,\mathrm{Area}(\Delta), \tag{VII.2.15}$$

and each node carries a colour space $\mathbb{C}^{N_i}$. The gauge group is the product of unitary factors, the honest output of the Chan–Paton data,

$$\prod_{i=1}^{G} U(N_i), \qquad N_i = N \text{ at the regular brane}, \tag{VII.2.16}$$

the regular-brane ranks being equal, $N_i = N$, because N mobile D3-branes carry rank N at each node. Unequal ranks are fractional branes, rank-vector deformations

associated with the collapsed cycles; the geometry is resolved by the Fayet–Iliopoulos / GIT (Kähler) data, not by the rank vector itself, and the fractional-brane dynamics and cascades are Chapter 8. The overall diagonal $U(1)$ is the decoupled centre of mass, and the surviving relative $U(1)$'s split into the baryonic symmetries counted below and the combinations that become massive; the interacting sector is the product of the $SU(N_i)$'s, but the construction's output is the product of unitary groups (VII.2.16), and we never write the brane gauge group as a bare SU.

For the standard regular-brane toric theory with the nondegenerate external-leg lattice, the number of independent non-anomalous baryonic $U(1)$ currents is read off the boundary-point count, not off the number of nodes:

$$\#\{\text{baryonic } U(1)\} = d - 3. \tag{VII.2.17}$$

This is the count that the conifold ($d = 4$, one baryonic $U(1)_B$) of Chapter 1 illustrated, and it is worth keeping separate from $G - 1$. The two agree for the conifold, where $d - 3 = 1$ and $G - 1 = 1$, but they diverge in general: for the orbifold $\mathbb{C}^3/\mathbb{Z}_3$ of §VII.2.7, with a triangular diagram $d = 3$, the baryonic count is $d - 3 = 0$, no continuous non-anomalous baryonic $U(1)$ symmetry, while $G - 1 = 2$ would be wrong. Relative gauge $U(1)$ factors, anomalous combinations and possible discrete symmetries are different questions. Its physical derivation is in Chapter 5.

The matter and the toric superpotential. The matter fields are bifundamental chiral multiplets, the arrows of a directed quiver: a field X_a on the arrow $i \to j$ transforms as a fundamental of $U(N_i)$ and an antifundamental of $U(N_j)$,

$$X_a : \ i \to j, \qquad X_a \in (\mathbf{N}_i, \overline{\mathbf{N}}_j). \tag{VII.2.18}$$

An adjoint is the special case in which the two faces coincide, an arrow $i \to i$ in $(\mathbf{N}_i, \overline{\mathbf{N}}_i)$, as for the three fields X, Y, Z of the single-node $\mathbb{C}^3$ honeycomb; "bifundamental" covers this case, and the apparent tension with the one-node fixture is only that. For the brane-tiling toric phases used here, the defining combinatorial property of the superpotential is the following form. It is a sum of signed traces over closed oriented loops of the quiver,

$$W = \sum_{\alpha} \pm \operatorname{Tr}\big(X_{a_1} X_{a_2} \cdots X_{a_{n_\alpha}}\big), \tag{VII.2.19}$$

subject to the **toric condition**: every chiral field appears in *exactly two* terms of W, once with a plus sign and once with a minus sign,

$$\text{each } X_a \text{ appears in } \begin{cases} \text{one } (+) \text{ term} \\ \text{one } (-) \text{ term} \end{cases} \tag{VII.2.20}$$

This two-terms-per-field structure is necessary for the bipartite tiling description, but by itself does not prove that an arbitrary quiver probes the intended cone. It also explains why the bipartite graph of Chapter 3 can encode the superpotential combinatorially: a field that sits in exactly one positive and one negative term is an edge that bounds exactly one white node and one black node of the dimer. The conifold quartic of Chapter 1, $W = \operatorname{Tr}(A_1 B_1 A_2 B_2) - \operatorname{Tr}(A_1 B_2 A_2 B_1)$, is the smallest non-trivial instance:

each of A_1, A_2, B_1, B_2 appears once with each sign. The toric condition is a property one can read off any candidate superpotential, and a field appearing in only one term, or in two terms of the same sign, signals a non-toric theory.

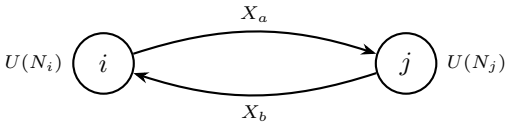


Figure VII.2.2: A piece of a toric quiver. Nodes are gauge factors (the interior label is the node index, the adjacent label its gauge group $U(N_i)$); arrows are bifundamental chiral fields. The full quiver of a consistent tiling has $G = 2\,\mathrm{Area}(\Delta)$ nodes. A chosen consistent tiling fixes the arrows and superpotential of a phase, with every field appearing in exactly two terms of opposite sign; the polygon alone does not choose that phase.

The node count is fixed under the consistent-tiling premise, while a choice of tiling supplies the directed matter and toric W. Distinct choices can give different phases of the same polygon; Chapter 3 packages one choice in a graph. What this section asserts, and what the next section verifies for one brane, is that this data is correct: that the gauge theory so constructed has the toric cone as its moduli space.

VII.2.5 The mesonic moduli space recovers the cone

The probe prescription requires that the brane reconstruct its transverse geometry as a branch of its own moduli space of supersymmetric vacua. For Volume VI this branch recovered $\mathbb{C}^3/\Gamma$; for the conifold it recovered the quadric $xy = zw$; and for a general toric cone it must recover X_Δ. This section states that identification, dimension-checks it, and shows it on the conifold as a ring computation, deferring the systematic perfect-matching derivation to Chapter 5.

The statement and its dimension. Take one regular brane, $N = 1$, so that each $U(N_i)$ is a $U(1)$ and the bifundamentals are complex numbers. The gauge-invariant chiral operators are the closed oriented loops of the quiver, the **mesons**, and the mesonic moduli space is the space of these operators subject to the F-term relations $\partial W = 0$ and the D-term (gauge) quotient. On the geometric mesonic component of a consistent regular-brane phase, at zero FI parameter, the probe identification is

$$\mathcal{M}^{\mathrm{mes}}_{N=1} = X_\Delta, \tag{VII.2.21}$$

a complex three-fold, so that one D3-brane is free to sit anywhere on the transverse Calabi–Yau. The dimension is the immediate check: a single D3-brane transverse to a Calabi–Yau threefold has three transverse complex directions, and the gauge theory must reproduce exactly that,

$$\dim_{\mathbb{C}} \mathcal{M}^{\mathrm{mes}}_{N=1} = 3. \tag{VII.2.22}$$

For a full-rank GLSM presentation, $k-(k-3)=3$ is a dimension consistency check on the geometric quotient. It does not establish that a proposed quiver has that component, nor exclude additional branches. The perfect-matching construction in Chapter 5 tests the quiver-to-cone identification; the two rings below provide direct examples.

The conifold as a ring, recalled. The cleanest illustration is the conifold, recalled here as the worked non-orbifold cone. With the four bifundamentals A_1, A_2, B_1, B_2 of Chapter 1 commuting at $N=1$, the gauge-invariant mesons are the four bilinears $M_{ij}=A_iB_j$, which assemble into a rank-one matrix and therefore satisfy

$$\det M = M_{11}M_{22}-M_{12}M_{21}=(A_1B_1)(A_2B_2)-(A_1B_2)(A_2B_1)=0, \qquad \text{(VII.2.23)}$$

identically. Identifying $(x,y,z,w)=(M_{11},M_{22},M_{12},M_{21})$, the single relation is $xy-zw=0$, and the four mesons modulo it cut out a three-fold, the conifold,

$$\mathcal{M}^{\text{mes}}_{N=1}=\operatorname{Spec}\frac{\mathbb{C}[x,y,z,w]}{(xy-zw)}, \qquad \dim_{\mathbb{C}}=4-1=3. \qquad \text{(VII.2.24)}$$

This is the cone recovered not by a dimension count but as an isomorphism of coordinate rings, four generators with one relation giving the three complex dimensions, exactly as in Chapter 1. For a consistent toric phase, the geometric mesonic coordinate ring is the semigroup algebra of lattice points in the dual cone. This is a physical identification for the selected quiver, not a consequence of the dimension count alone. Making this general statement into a systematic computation, by enumerating the perfect matchings of the dimer and reading the lattice points (with their multiplicities) off them, is the content of Chapter 5; here the conifold is the worked case and the dimension check (VII.2.22) stands for the general identification.

dP_0 as a second ring. The orbifold $\mathbb{C}^3/\mathbb{Z}_3$ supplies a second mesonic ring without running the perfect-matching machine, directly from invariant theory. At one regular brane the gauge-invariant mesons are the $\mathbb{Z}_3$-invariant monomials in z_1, z_2, z_3, those of total degree divisible by three,

$$\mathcal{M}^{\text{mes}}_{N=1}=\operatorname{Spec}\mathbb{C}[z_1,z_2,z_3]^{\mathbb{Z}_3}=\operatorname{Spec}\bigoplus_{n\geq 0}\operatorname{Sym}^{3n}\mathbb{C}^3, \qquad \text{(VII.2.25)}$$

the cone over the degree-three Veronese embedding of $\mathbb{P}^2$. The generators are the ten degree-three monomials $z_1^iz_2^jz_3^k$ with $i+j+k=3$, and the ring they generate, modulo the Veronese ideal, has Krull dimension three, matching (VII.2.22): three transverse directions for the one brane. The full Hilbert series, with the count of generators and relations, is Chapter 5's; here it suffices that the invariant ring is generated in degree three and cuts out the three-fold local-$\mathbb{P}^2$ cone of (VII.2.29). The conifold and dP_0 thus give the two ends of the bridge as explicit rings: a non-orbifold quadric and an orbifold Veronese cone, both three complex dimensions.

Two honesty caveats. Two qualifications travel with (VII.2.21). First, reading the classical mesonic branch as the literal transverse geometry carries the α' and g_s probe caveats inherited from Volume VI: it is the leading geometric identification, with stringy and quantum corrections set aside. Second, the effect of turning on FI parameters

depends on the chosen GIT chamber and triangulation; a smooth crepant resolution is obtained only when that chamber gives a unimodular subdivision. The chambers are named here and the explicit resolution is routed to Chapter 5 and Part I. The moduli space of one regular brane at the origin of Fayet–Iliopoulos space is the singular cone X_Δ, the superconformal point of the gauge theory, and that is what (VII.2.21) asserts.

VII.2.6 The inverse problem and toric duality

The geometry-to-probe statement is settled in principle: a toric cone gives a quiver gauge theory whose mesonic moduli space is the cone. The chapter's distinctive content is the inverse direction and its surprise.

The inverse problem. The **inverse problem** is to start from the geometry, a toric diagram together with its (p, q)-web, and reconstruct a quiver and toric superpotential whose mesonic moduli space is the given cone. A quiver-plus-superpotential pair that solves it is a **toric phase** of the cone. There is an algorithm that produces such a phase, the inverse algorithm of the brane-tiling literature, which routes the diagram through a brane tiling and reads the quiver off the tiling; this chapter names it and defers its construction to Chapters 3 through 5. The terminology is worth fixing now, because the two directions are easily confused: geometry to gauge theory is the *inverse* problem, and the *forward* algorithm is the other direction, dimer or quiver to toric data, which is Chapter 5's. This chapter uses inverse for geometry-to-gauge throughout.

The non-uniqueness. The new phenomenon beyond the elementary orbifold examples of Volume VI is that the inverse problem has *several* solutions. One toric cone admits several inequivalent toric phases: different quivers, with different numbers of arrows and different superpotentials, all of which have the same mesonic moduli space. The map from geometry to gauge theory is one-to-many,

$$\Delta \;\longmapsto\; \{\,\text{toric phase I},\ \text{toric phase II},\ \ldots\,\}, \tag{VII.2.26}$$

and choosing a phase is a choice of duality frame, not a property of the geometry. This is **toric duality**.

What the geometry does fix, and what all phases of one cone therefore share, is the toric data and its invariants. Two consistent toric phases of the same cone agree on the diagram (up to $GL(2,\mathbb{Z})$ and translation), on the node count, and on the external-leg count, and at the level of the protected spectrum on the mesonic Hilbert series,

$$\Delta \text{ and } G = 2\,\mathrm{Area}(\Delta) \text{ and } d \text{ and the mesonic Hilbert series are phase-invariant.} \tag{VII.2.27}$$

What they need not agree on is the quiver: the number of bifundamental fields can differ between phases. Not every cone has more than one phase, and the smallest do not. A local square move on an equal-rank tiling phase requires a suitable node with $N_f = 2N_c$; this rank condition alone does not ensure a *distinct* phase, as the self-dual conifold shows. Other constructions of toric phases are not classified by this local criterion. Flat $\mathbb{C}^3$ (one node, no dual), the conifold (each node $N_f = 2N_c$, self-dual back to the same quiver), and the orbifold $\mathbb{C}^3/\mathbb{Z}_3 = dP_0$ of the next section (each node

$N_f = 3N_c$, outside the window) are each *single-phase*. The simplest cone with two genuinely distinct toric phases is $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$, whose two four-node quivers share $G = 2\,\mathrm{Area} = 4$ and the same cone but differ in field content; the higher del Pezzo cones dP_2 and dP_3 have two and four phases. These multi-phase examples are worked in Chapter 7. The geometry pins the cone; it does not pin the field content.

A preview of the first genuinely two-phase cone fixes the idea before Chapter 7 performs the move. The diagram of $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$ is the square with vertices $(\pm 1, 0), (0, \pm 1)$ and one interior point $(0, 0)$, so

$$G = 2\,\mathrm{Area}(\Delta_{F_0}) = 4, \qquad d = 4, \qquad I = 1, \qquad d - 3 = 1, \tag{VII.2.28}$$

and these invariants, together with the mesonic Hilbert series, are shared by both phases. What differs is the quiver: F_0 has two inequivalent four-node toric phases with different numbers and arrangements of bifundamentals and different superpotentials, related by the toric-duality move on a node in the $N_f = 2N_c$ window. Both reproduce the same cone over F_0, the point already visible here: the polygon (VII.2.28) does not select a unique quiver. The explicit two quivers and the matching of their invariants are Chapter 7.

Three things the polygon does and does not fix

The toric diagram Δ *fixes* the geometry (the cone X_Δ), the node count $G = 2\,\mathrm{Area}(\Delta)$, the baryonic count $d - 3$, and the mesonic Hilbert series. It *does not fix* a unique quiver: one cone has several toric phases (toric duality), differing in their field content and superpotential. *Geometry to gauge theory is one-to-many; do not read the diagram as selecting a single quiver.*

Why duality. Calling the non-uniqueness a duality is a physical claim, not just a name. For consistent regular-brane phases related by Seiberg moves and admitting the stated superconformal limit, the duality identifies their infrared fixed point and protected spectrum. Equality of the mesonic moduli spaces alone does not establish this physical equivalence. On the combinatorial side the change of phase is a local move on the brane tiling, the urban-renewal move, and the matching of the invariants across the move, the toric data, the Hilbert series, and the central charge, is the content of Chapter 7. This chapter establishes that the non-uniqueness exists and names its meaning; the move, the Seiberg-duality proof, and the invariant-matching are Chapter 7, and the field-theory Seiberg-duality theorem in full generality is Part III. The lesson to carry forward is the one in the box: the inverse map is one-to-many, and a toric phase is a choice of duality frame for a fixed geometry.

VII.2.7 The bridge examples: $\mathbb{C}^3/\mathbb{Z}_3$ as a polygon, and the conifold recalled

The general statements above are made concrete by two examples that also tie the toric language to the rest of the book. The first is the orbifold $\mathbb{C}^3/\mathbb{Z}_3$, the dP_0 cone, now read as a toric polygon, which is the bridge from the McKay quivers of Volume VI

to the dimers of this volume. The second is the conifold of Chapter 1, recalled as the non-orbifold anchor.

$\mathbb{C}^3/\mathbb{Z}_3$ as a toric diagram. The orbifold $\mathbb{C}^3/\mathbb{Z}_3$ with the action $(z_1, z_2, z_3) \mapsto \omega(z_1, z_2, z_3)$, $\omega = e^{2\pi i/3}$, the weights $(1, 1, 1)$ of the directed-McKay quiver of Volume VI, is a toric cone because the orbifold group is abelian, and its toric diagram is the **local-$\mathbb{P}^2$ polygon**,

$$\Delta_{dP_0} = \mathrm{conv}\{(1,0), (0,1), (-1,-1)\}, \qquad \text{interior point } (0,0), \tag{VII.2.29}$$

the triangle with those three vertices and the single interior lattice point $(0, 0)$: three boundary points and one interior point, up to $GL(2, \mathbb{Z})$ and translation. This is the content of the orbifold action read in toric language. The invariant ring $\mathbb{C}[z_1, z_2, z_3]^{\mathbb{Z}_3}$ is generated by the $\mathbb{Z}_3$-invariant monomials, those of total degree divisible by three; the smallest invariant degree is three, and the ten degree-three monomials $z_1^i z_2^j z_3^k$ with $i + j + k = 3$ are the sections of $\mathcal{O}_{\mathbb{P}^2}(3)$, the anticanonical bundle of $\mathbb{P}^2$. These ten are not new elementary fields of the quiver; they are mesonic invariants of the one-brane theory, the coordinates on the quotient cone. The cone over $\mathbb{P}^2$ in this embedding is exactly the local-$\mathbb{P}^2$ geometry, and its toric diagram is (VII.2.29). It is emphatically not the side-length-three standard-lattice triangle, $\mathrm{conv}\{(0,0), (3,0), (0,3)\}$, which would carry spurious boundary points; the correct diagram has exactly three boundary and one interior lattice point, and its invariants are

$$G = 2\,\mathrm{Area}(\Delta_{dP_0}) = 3, \qquad d = 3, \qquad I = 1, \qquad d - 3 = 0, \tag{VII.2.30}$$

three gauge nodes, three external legs, one compact divisor, and no continuous non-anomalous baryonic $U(1)$ symmetry. The (p, q)-web legs are the outward normals of the three boundary edges, computed by hand from the vertices:

$$\begin{aligned} (1,0) \to (0,1) &: \text{edge } (-1,1) \;\Rightarrow\; \text{leg } (1,1), \\ (0,1) \to (-1,-1) &: \text{edge } (-1,-2) \;\Rightarrow\; \text{leg } (-2,1), \\ (-1,-1) \to (1,0) &: \text{edge } (2,1) \;\Rightarrow\; \text{leg } (1,-2), \end{aligned} \qquad \sum_{\ell} (p_\ell, q_\ell) = (0,0), \tag{VII.2.31}$$

each leg the inward edge vector rotated to the outward perpendicular, and the three close as the condition (VII.2.5) requires. This is the (p, q)-web of Figure VII.2.3 read as a reproducible operation rather than a picture.

The bridge is this. The same cone $\mathbb{C}^3/\mathbb{Z}_3$ was constructed in Volume VI from the representation theory of $\mathbb{Z}_3$: the directed-McKay quiver with three nodes, adjacency $a_{ij} = 3\,\delta_{j,i+1}$, nine bifundamentals, and a cubic superpotential. That quiver is now recognized as *one toric phase* of the local-$\mathbb{P}^2$ cone, the phase whose node count $G = 3$ is the area of (VII.2.29) and whose three nodes are the three irreducible representations of $\mathbb{Z}_3$. The orbifold construction of Volume VI and the toric construction of this volume are two routes to the same target, the McKay route through the character table and the toric route through the polygon, and the equivalence of their gauge theories is the crosslink recorded between the two volumes. In the standard equal-rank toric-phase catalog used here, the McKay quiver is the single dP_0 phase. Each node has $N_f = 3N_c$, outside the local $N_f = 2N_c$ square-move window. Dualizing a node is still a legitimate

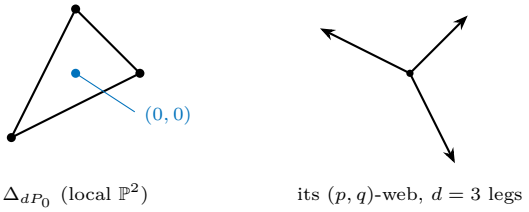


Figure VII.2.3: The orbifold $\mathbb{C}^3/\mathbb{Z}_3 = dP_0$ as a toric diagram and its (p, q)-web. The left polygon is the local-$\mathbb{P}^2$ triangle: three boundary points, the vertices $(1, 0), (0, 1), (-1, -1)$, and one interior point $(0, 0)$. Its area is $\frac{3}{2}$, so $G = 2\,\text{Area} = 3$, matching the three nodes of the Volume VI directed-McKay quiver of $\mathbb{C}^3/\mathbb{Z}_3$. The web has three legs $(-2, 1), (1, -2), (1, 1)$, the outward normals of the three edges, and they sum to zero. There are $d - 3 = 0$ continuous non-anomalous baryonic $U(1)$ symmetries.

field-theory operation, but it does not return an equal-rank regular-brane quiver, so no distinct equal-rank toric phase of dP_0 appears. Toric duality, with several phases for one cone, sets in only for larger cones such as $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$ and the higher del Pezzos, treated in Chapter 7.

The conifold recalled. The conifold of Chapter 1 is the non-orbifold anchor, and it is instructive to set the two side by side. Its toric diagram is the unit square $\text{conv}\{(0, 0), (1, 0), (1, 1), (0, 1)\}$, with four boundary points and no interior point, so

$$G = 2\,\text{Area} = 2, \qquad d = 4, \qquad I = 0, \qquad d - 3 = 1, \tag{VII.2.32}$$

two gauge nodes, four external legs, one baryonic $U(1)_B$, no compact divisor. The unit square is $GL(2, \mathbb{Z})$-inequivalent to both the unit triangle of flat $\mathbb{C}^3$ ($G = 1$) and the local-$\mathbb{P}^2$ triangle of $\mathbb{C}^3/\mathbb{Z}_3$ ($G = 3$), which is the toric face of the fact that the conifold is not an orbifold: there is no finite group $\Gamma \subset SU(3)$ with $\mathbb{C}^3/\Gamma$ the conifold, and most toric cones, the conifold being the simplest, are not orbifolds at all. For an affine height-one simplicial toric cone, a triangular diagram gives a finite abelian quotient of $\mathbb{C}^3$ by the finite lattice cokernel; conversely the diagonal abelian orbifolds considered in Volume VI give such triangles. The lattice *and* ray embedding fix the group, so triangle shape alone in ordinary Euclidean space is insufficient; such a triangle can carry many boundary and interior lattice points (a larger Γ gives a larger triangle with more points, as $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ already shows with its single interior point), but it has exactly three vertices. A polygon with more than three vertices, such as the conifold's unit square, is not an abelian orbifold, and the McKay reading does not extend to it. Table VII.2.3 collects the two bridge examples with flat $\mathbb{C}^3$ as the reference.

How to read Table VII.2.3. Each row is a cone, given by its toric diagram, and the columns are the $GL(2, \mathbb{Z})$ invariants that the gauge theory reads off the diagram. The node count G grows from one for flat $\mathbb{C}^3$ to three for $\mathbb{C}^3/\mathbb{Z}_3$; the baryonic count $d - 3$ is zero except for the conifold; and the orbifold column marks which cones are quotients, the two triangles being orbifold (one of them trivially) and the square not. The table is the chapter in miniature: the polygon is the input, its area and point counts are the output, and the McKay reading is the special case, not the rule.

cone	$G = 2\,\mathrm{Area}$	d	I	$d-3$	nodes	orbifold?
$\mathbb{C}^3$	1	3	0	0	1	yes (trivial)
conifold ($xy = zw$)	2	4	0	1	2	no
$\mathbb{C}^3/\mathbb{Z}_3 = dP_0$	3	3	1	0	3	yes ($\mathbb{Z}_3$)

Table VII.2.3: The toric invariants of the three example cones. The node count is the normalized area, twice the Euclidean area; d is the boundary-point count (external legs); I the interior-point count (compact divisors); $d-3$ the number of non-anomalous baryonic $U(1)$'s. The conifold is the non-orbifold anchor (a unit square, $GL(2,\mathbb{Z})$-inequivalent to the two triangles); $\mathbb{C}^3/\mathbb{Z}_3$ is the orbifold bridge to the Volume VI directed-McKay quiver. Note that $d-3$ is not $G-1$: dP_0 has $d-3=0$ but $G-1=2$.

VII.2.8 What has been computed, and what is deferred

The chapter has set up the toric-probe correspondence, computed its lattice data and two mesonic examples, and located the physical premises needed for the general construction. A toric Calabi–Yau threefold is presented by a convex lattice polygon Δ at height one, the Calabi–Yau condition being the height-one hyperplane of the rays; its (p,q)-web is the perpendicular dual, with legs that close. The same data is the gauged linear sigma model, a chamber-dependent connected quotient of $\mathbb{C}^k$ with charges a saturated basis of the ray-matrix kernel, with a finite quotient retained when the rays span a proper sublattice. A stack of N D3-branes carries an $\mathcal{N}{=}1$ quiver with $G = 2\,\mathrm{Area}(\Delta)$ gauge nodes $\prod_i U(N_i)$, bifundamental matter, and a toric superpotential in which every field appears in exactly two terms of opposite sign; the geometric mesonic component of a consistent one-brane phase recovers the affine cone at zero FI, shown as ring statements for the conifold and dP_0. The inverse problem, geometry to a toric phase, is named and shown to be one-to-many: one cone, several Seiberg-dual phases, which is toric duality, with the diagram, the node count, the baryonic count, and the mesonic Hilbert series the phase-invariants. The orbifold $\mathbb{C}^3/\mathbb{Z}_3$ is presented as the local-$\mathbb{P}^2$ polygon and its Volume VI directed-McKay quiver recognized as one toric phase, the conifold recalled as the non-orbifold anchor.

A reader leaving the chapter should be able to do the following without looking ahead: write a toric cone's height-one ray data and read its diagram Δ; extract the invariants $G = 2\,\mathrm{Area}$, d, and I, and check them against Pick's theorem; draw the perpendicular (p,q)-web and verify that its legs close; pass to the gauged linear sigma model by taking the integer kernel of the ray matrix and confirm that the charges sum to zero; write down the gauge group $\prod_i U(N_i)$ and the baryonic count $d-3$, keeping it distinct from $G-1$; state the toric condition on the superpotential and test it on a candidate W; dimension-check the mesonic moduli space at three and recover the conifold as $xy = zw$; explain why the geometry-to-gauge map is one-to-many and what toric duality means; and present $\mathbb{C}^3/\mathbb{Z}_3$ as the local-$\mathbb{P}^2$ triangle, identifying its three nodes with the Volume VI McKay quiver and its $d-3=0$ baryonic count.

What has been deferred, and to where, is the rest of the machinery the volume builds.

- The dimer, the bipartite graph on the two-torus that encodes the quiver and the toric superpotential in one object, is named here as the right organizing object (the two-terms-per-field structure is exactly what it packages) and built in Chapter 3; its physical origin as a Type IIB NS5–D5 brane configuration obtained by two T-dualities is Chapter 4, where the zig-zag paths of the tiling are matched to the external (p, q)-web legs of this chapter.

- The mesonic moduli space was stated to be the cone and dimension-checked; the systematic derivation, by enumerating the dimer's perfect matchings and reading the Newton-polygon lattice points off them with their multiplicities, together with the P, Q_F, Q_D, Q_t charge matrices, the master space, and the mesonic and baryonic Hilbert series, is the forward algorithm of Chapter 5, which also works the explicit $\mathbb{C}^3/\mathbb{Z}_3$ matching.

- Toric duality was named and its phase-invariants identified; the urban-renewal dimer move that relates the phases, the Seiberg-duality proof that they are physically dual, and the explicit matching of the toric data, Hilbert series, and central charge across phases are Chapter 7, while the field-theory Seiberg-duality theorem in full generality is Part III.

- The base of the cone is a Sasaki–Einstein five-manifold, and the superconformal R-charges and the central charge are computable from the toric data; the selection of the R-symmetry by a-maximization and its match to the Sasaki–Einstein volume minimization is Chapter 6.

- The toric-geometry foundations, the fans and polytopes, the quotient construction of (VII.2.8), and the height-one Gorenstein condition as a theorem of toric geometry, are cited to Part I (Volume IV), not proved here; the resolution of the cone by the Fayet–Iliopoulos chambers is Chapter 5 and Part I.

The single lesson to carry forward is the keybox at the start: the polygon fixes the toric geometry and, under the consistent-tiling premise, the node count, but not a unique quiver, and the chapters that follow build the dimer that constructs a phase and the algorithm that reads the geometry back out of it.

Status and sources

The evidence-code tags ([MATH] cited mathematical theorem; physical dictionary claims retain an explicit unclassified primary until a suitable derivation and evidence route are located), the convention pointers, and the source anchors are collected here, keyed by section. The chapter states the general toric-probe correspondence, after the toric D-brane gauge theories of Morrison and Plesser (1999) and Beasley et al. (2000), the toric construction and the discovery of toric duality of Feng et al. (2001), and the diagram-and-web presentation of Hanany and Iqbal (2002); it forward-references the inverse algorithm of Hanany and Vegh (2007) and the brane-tiling corpus of Hanany and Kennaway (2005) and Franco et al. (2006), recalls the conifold gauge theory of Klebanov and Witten (1998), and uses the abelian-orbifold toric Higgs branch of Douglas et al. (1997) for the $\mathbb{C}^3/\mathbb{Z}_3$ bridge. It reuses the toric-diagram conventions (toric lattice orientation, the height-one polygon, the perpendicular (p, q)-web, the GLSM charge matrix as the kernel of the ray matrix; VII.2 is a named first load-bearing use), the gauge/rank data (quiver gauge groups and ranks, $\prod_i U(N_i)$, the interacting sector and the

baryonic count), the toric superpotential store (superpotential normalization, the two-terms-per-field condition), the dimer/brane-tiling conventions (bipartite tiling conventions, named here, the machinery deferred to Chapters 3–5), the resolution/Fayet–Iliopoulos data (resolution and discrete data, named for the GIT chambers), the arrow convention (fundamental-at-tail arrows), and the supercharge count (spinor and supersymmetry counting, $\mathcal{N}$=1 from the toric $SU(3)$ holonomy). [WS] requires a controlled untwisted-string worldsheet derivation; neither a GLSM kernel nor a dimer identity supplies that. The physical probe, node-count, baryonic-count, and toric-duality claims here therefore retain source-backed but unclassified primary status pending claim-specific derivations and classification. Geometric comparisons and duality chains have their own physical premises. In particular, the Seiberg-duality premise needs its physical duality source, not a lattice identity. Mathematical identities carry [MATH] only with their stated lattice or ring hypotheses. The finite calculations use the displayed lattice and tiling data.

§VII.2.1 From one example to a general rule. Framing: the conifold as the special symmetric case, the toric diagram as the general combinatorial polygon, the build-promise (the polygon to the quiver, the one-to-many surprise), and the reader contract. In force: (the quiver is fixed by the structure, now toric-combinatorial rather than representation-theoretic) and (most toric CY_3 are not orbifolds; the McKay reading is the special case).

§VII.2.2 The toric diagram and the (p, q)-web. The height-one Calabi–Yau condition (VII.2.1), the diagram Δ (VII.2.2), the three $GL(2,\mathbb{Z})$ invariants $G = 2\,\mathrm{Area}$, d, I (VII.2.3) and Pick's theorem (VII.2.4), and the perpendicular (p, q)-web with its closure (VII.2.5). The lattice identities carry [MATH] under the hypotheses stated above; identifying the normal vectors as physical five-brane charges retains an unclassified physical primary (toric lattice orientation; Morrison and Plesser (1999), Hanany and Iqbal (2002)). The toric-geometry foundations (fans, polytopes, the Gorenstein condition) are toric geometry (home Vol. IV ch.4), cited not constructed. the three fixtures are height-one CY polygons with closing webs ($\sum$ legs $= 0$); Pick's theorem holds on each; the point counts are $GL(2,\mathbb{Z})$ invariants; the web closure detects an omitted leg. The separate dot-product check distinguishes a parallel edge from its normal. In the submitted polygon exercise, the outward primitive leg multiset and the net-charge sum are separate obligations: a closed parallel list passes the latter and fails the former.

§VII.2.3 The cone as a quotient: the GLSM. The ray matrix (VII.2.6), the GLSM charge matrix as its integer kernel (VII.2.7), the quotient $\mathbb{C}^k /\!/ U(1)^{k-3}$ (VII.2.8), and the Calabi–Yau charge-sum condition (VII.2.9). The integer-kernel and charge-sum identities carry [MATH]; the Cox/GIT quotient theorem is source-only with an exact proof locator still owed (toric lattice orientation; Morrison and Plesser (1999)). The quotient-construction theorems are toric geometry, cited not constructed. evaluated(as the kernel/height structure underlying the dimension count): the forward data gives $\dim\ker Q_t = 3$ and $Q_t \cdot \mathbf{1} = 0$ (the CY height) on every fixture; the fixed-covector pairings are exact. Its height-two example rejects only the chosen covector, not every integral covector or the underlying cone.

§VII.2.4 The D3 probe: quiver, gauge group, toric superpotential. The node count $G = 2\,\mathrm{Area}$ (VII.2.15), the gauge group $\prod_i U(N_i)$ (VII.2.16), the baryonic count $d - 3$ (VII.2.17), the bifundamental arrow convention (VII.2.18), and the toric superpotential with its two-terms-per-field condition (VII.2.19)–(VII.2.20). : geometric probe dictionary, physical primary classification pending (quiver gauge groups and ranks, superpotential normalization, fundamental-at-tail arrows, spinor and supersymmetry counting, Feng et al. (2001), Beasley et al. (2000)). In force: (the gauge group is not a bare SU; the baryonic count is $d - 3$, not $G - 1$). $G = 1/2/3$ and $d = 3/4/3$ for $\mathbb{C}^3$/conifold/dP_0; the toric superpotential is two-terms-per-field on every fixture dimer (each chiral on exactly one white and one black node); the baryonic count is $0/1/0$ (requiring two appearances for each field); $G = 2\,\mathrm{Area}$ by the shoelace determinant, $GL(2,\mathbb{Z})$-invariant.

§VII.2.5 The mesonic moduli space recovers the cone. The identification $\mathcal{M}^{\mathrm{mes}}_{N=1} = X_\Delta$ (VII.2.21), the dimension check at three (VII.2.22), the conifold as a ring $xy = zw$ (VII.2.23)–(VII.2.24), and the two honesty caveats. : geometric probe dictionary, physical primary classification pending (Feng et al. (2001), Beasley et al. (2000)). Carries (the classical mesonic branch carries the α'/g_s probe caveats of Volume VI). The systematic perfect-matching derivation and the master space are routed to Chapter 5; the resolution chambers (resolution and discrete data) to Chapter 5 and Part I. the mesonic moduli complex dimension is three on every fixture; the conifold meson matrix is rank one with $\det M = 0$

and the cone relation $xy - zw = 0$, with the expected dimension obtained as $4 - 1$. These operations check the determinant relation and dimension arithmetic; they do not compute an elimination ideal or prove it principal.

§VII.2.6 The inverse problem and toric duality. The inverse problem and the toric phase, the one-to-many map (VII.2.26), the phase-invariants (VII.2.27), the keybox, and the Seiberg-duality reading. : cited toric-phase and Seiberg-duality premises, physical primary classification pending (Feng et al. (2001), Hanany and Vegh (2007)). Files (the toric diagram does not determine a unique quiver; the inverse map is one-to-many, toric duality). The urban-renewal move, the Seiberg proof, and the matched invariants are Chapter 7; the field-theory theorem is Part III. The phase counts are imported catalog data. The comparison uses the catalogued counts for $\mathbb{C}^3$, the conifold, dP_0 and F_0; the finite check3b checks area and boundary counts under a shear of one F_0 polygon. Neither operation constructs both phases or proves their common infrared theory. The dP_2 and dP_3 phase counts in the body are sourced previews of Chapter 7, with exact catalog locators still owed here. The physical Seiberg-duality assertion is likewise source-only and unclassified.

§VII.2.7 The bridge examples. The $\mathbb{C}^3/\mathbb{Z}_3 = dP_0$ local-$\mathbb{P}^2$ polygon (VII.2.29) with its invariants (VII.2.30), the degree-three invariant-ring reading, the identification of the Volume VI directed-McKay quiver as one toric phase, and the conifold recalled with its invariants (VII.2.32) (Table VII.2.3). The polygon and invariant-ring calculations carry [MATH]; the McKay-to-toric physical identification retains an unclassified primary (toric lattice orientation; Douglas et al. (1997), Feng et al. (2001)). The directed-McKay and toric constructions give the same target in this example. The full dimer perfect-matching computation of dP_0 (Newton-polygon recovery) is Chapter 5. the chosen dP_0 triangle has vertices $(1,0), (0,1), (-1,-1)$, interior $(0,0)$, three boundary and one interior point, $G = 3$. the finite check4b compares that triangle to its chosen shear; it does not derive the orbifold lattice from the weights $(1,1,1)$. The orbifold identification uses the invariant-ring argument and the cited geometric premise; the conifold is the unit square, $G = 2$, $d = 4$, $GL(2,\mathbb{Z})$-inequivalent to both triangles ; the ten degree-three $\mathbb{Z}_3$-invariants are the sections of $\mathcal{O}_{\mathbb{P}^2}(3)$, no lower-degree invariant.

§VII.2.8 What is and is not given. Scope and limitations. the chapter builds the toric diagram and web, the GLSM quotient, the quiver and gauge group with the toric superpotential, the mesonic moduli space as a dimension-three cone (the conifold as a ring), the inverse problem and toric duality named, and the two bridge examples; it routes the dimer machinery to Chapters 3–4, the forward algorithm and master space to Chapter 5, the toric-duality move and invariant matching to Chapter 7, a-maximization and the Sasaki–Einstein volume to Chapter 6, the toric-geometry foundations to Part I (toric geometry), and the field-theory Seiberg-duality theorem and precision dynamics to Part III. The Sasaki–Einstein and R-charge claims are sourced previews of Chapter 6, outside the submitted computations here.

Related constructions. A D3 probe of a toric Calabi–Yau threefold gives an $\mathcal{N} = 1$ quiver with a toric superpotential, and the mesonic moduli space reconstructs the cone. The height-one diagram and its (p, q) web give the toric data. The inverse problem is nonunique. Chapter 7 discusses the Seiberg-duality interpretation. For $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$, the local-$\mathbb{P}^2$ polygon recovers the same target as the directed-McKay quiver in Volume VI, Chapters 2 and 6.

References

Beasley, C. E. et al. (2000). "D3-branes on partial resolutions of Abelian quotient singularities of Calabi-Yau threefolds". In: *Nucl. Phys. B* 566, pp. 599–640. DOI: 10.1016/S0550-3213(99)00766-0. arXiv: hep-th/9907186.

Douglas, M. R., B. R. Greene, and D. R. Morrison (1997). "Orbifold resolution by D-branes". In: *Nucl. Phys. B* 506, pp. 84–106. DOI: 10.1016/S0550-3213(97)00517-8. arXiv: hep-th/9704151.

Feng, B., A. Hanany, and Y.-H. He (2001). "D-brane gauge theories from toric singularities and toric duality". In: *Nucl. Phys. B* 595, pp. 165–200. DOI: 10.1016/S0550-3213(00)00429-7. arXiv: hep-th/0003085.

Franco, S. et al. (2006). "Brane dimers and quiver gauge theories". In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Hanany, A. and A. Iqbal (2002). "Quiver theories from D-branes". In: *JHEP* 04, p. 009. DOI: 10.1088/1126-6708/2002/04/009. arXiv: hep-th/0108137.

Hanany, A. and K. D. Kennaway (2005). "Dimer models and toric diagrams". In: arXiv: hep-th/0503149.

Hanany, A. and D. Vegh (2007). "Quivers, tilings, branes and rhombi". In: *JHEP* 10, p. 029. DOI: 10.1088/1126-6708/2007/10/029. arXiv: hep-th/0511063.

Klebanov, I. R. and E. Witten (1998). "Superconformal field theory on three-branes at a Calabi-Yau singularity". In: *Nucl. Phys. B* 536, pp. 199–218. DOI: 10.1016/S0550-3213(98)00654-3. arXiv: hep-th/9807080.

Morrison, D. R. and M. R. Plesser (1999). "Nonspherical horizons. 1." In: *Adv. Theor. Math. Phys.* 3, pp. 1–81. DOI: 10.4310/ATMP.1999.v3.n1.a1. arXiv: hep-th/9810201.

Chapter VII.3

Brane tilings and dimer models

VII.3.1 One graph for the quiver and the superpotential

The last two chapters built toric quiver gauge theories one at a time. For the conifold the quiver and its quartic superpotential were fixed by the geometry and its symmetry; for a general toric Calabi–Yau threefold the previous chapter organized the gauge-theory data using the toric diagram and its (p, q)-web, together with a choice of toric phase. In both cases the superpotential was a separate piece of work, written down by hand and then justified by marginality and the moduli-space check. The quiver said which fields existed; nothing in the quiver said how they coupled. For the conifold that gap was small, since the symmetry left only one quartic. For a del Pezzo cone or a $Y^{p,q}$ it is not small at all: there are many candidate gauge-invariant words, and getting the superpotential wrong gives the wrong moduli space and the wrong fixed point.

This chapter introduces the single combinatorial object that closes the gap. A brane tiling, or dimer, is a bipartite graph drawn on a two-torus. Its faces are the gauge nodes, its edges are the bifundamentals, and its two-coloured graph nodes are the superpotential terms, with a sign rule fixed by the colouring. One planar object therefore carries the quiver *and* the superpotential words and their relative signs at once, with the toric normalization of the couplings. This is the chapter's payoff, and it is exactly the role the McKay quiver played in Volume VI, transposed from representation theory to combinatorial geometry. There the organizing object was the McKay quiver of a finite group Γ, and the cubic superpotential was a separate inheritance from the $\mathcal{N}{=}4$ parent. Here the organizing object is a graph on T^2, and the superpotential is read off the same graph.

Take-home message of this chapter

A brane tiling is a bipartite graph on T^2. *Faces* are the gauge nodes $\prod_i U(N_i)$, *edges* are the bifundamentals, and the black and white *nodes* are the superpotential terms, combined by the sign rule $W = \sum_{\text{white}} \text{Tr}(\text{cw cycle}) - \sum_{\text{black}} \text{Tr}(\text{ccw cycle})$. The one graph fixes the quiver and the superpotential together. For a consistent toric phase the Newton polygon of its Kasteleyn determinant recovers the toric diagram. Zig-zag paths test the embedding; their winding balance alone does not establish an interacting superconformal theory. *The dimer is to this volume what the McKay quiver was to the last.*

The dimer is not only a bookkeeping device. Three computations come out of the same graph, and this chapter does each on worked examples. The dictionary itself, faces and edges and coloured nodes to gauge theory, is the first. The second is the route from the tiling back to the toric diagram: a determinant built from the graph, the Kasteleyn determinant, has a Newton polygon that, for a consistent toric phase, is the toric diagram of the singularity. The third is a consistency filter, because not every bipartite graph on T^2 is a physical gauge theory; the diagnostic is a family of closed paths on the graph, the zig-zag paths, whose winding classes must match the external legs of the (p, q)-web. Failure rules out the proposed identification with those web legs; passing the winding test still leaves the intersection conditions and the physical interpretation to check.

A word on what the chapter assumes and what it merely names, so that a reader knows what to hold in mind and what may be deferred. The chapter assumes the four-dimensional $\mathcal{N}$=1 toric-quiver vocabulary of the previous chapters: gauge nodes $\prod_i U(N_i)$, bifundamental arrows, the superpotential with its ± 1 toric couplings, and the toric diagram with its (p, q)-web of external legs. It treats the dimer as a physical brane configuration and names its Type IIB realization, but the derivation of that realization is the next chapter's, and the physical interpretation here uses it as an imported premise. The moduli-space construction, the forward construction that turns a tiling into the full toric data and the Hilbert series, is named and its inputs are prepared, but it is run in Chapter 5. The dimer moves that relate different tilings of one singularity are named and deferred to Chapter 7.

The plan of the chapter is: the dimer dictionary and its sign rule (§VII.3.2); the periodic quiver as the planar dual (§VII.3.3); the C^3 honeycomb worked in full as the simplest fixture (§VII.3.4); the Kasteleyn matrix and the route to the toric diagram (§VII.3.5); perfect matchings and the lattice points of the Newton polygon (§VII.3.6); the conifold square recalled and a third fixture, the dP_0 tiling, computed (§VII.3.7); zig-zag paths and the consistency that tests the geometric consistency of the tiling (§VII.3.8); and a closing map of what was computed and what is deferred (§VII.3.9).

VII.3.2 The dimer dictionary

A **brane tiling**, or **dimer model**, is a bipartite graph drawn on the two-torus T^2, with a fixed orientation of the surface and disk faces. The disk-face assumption makes this a cellular embedding; a list of edge windings alone does not specify it. Cyclic orders at the vertices are also required. Bipartite means the graph nodes come in two colours,

black and white, and every edge joins a black node to a white node. Drawn on T^2 the graph cuts the torus into faces, and the three kinds of cell, faces, edges, and nodes, are exactly the three kinds of data of an $\mathcal{N}$=1 quiver gauge theory.

$$\underbrace{\text{faces}}_{\text{gauge nodes } U(N_i)} \qquad \underbrace{\text{edges}}_{\text{bifundamentals } X_{ij}} \qquad \underbrace{\text{black/white nodes}}_{\text{superpotential terms}}. \tag{VII.3.1}$$

Each piece of the dictionary is read as follows.

Faces are gauge nodes. A fundamental domain of the tiling has some number G of faces, and each face is a gauge factor $U(N_i)$, so the gauge group is the product over faces,

$$G_{\text{gauge}} = \prod_{i=1}^{G} U(N_i). \tag{VII.3.2}$$

For a single stack of regular D3-branes the ranks are equal, $N_i = N$, and the construction's output is the unitary product; the interacting non-abelian sector is $\prod_i SU(N_i)$, and of the $G-1$ relative central $U(1)$'s the non-anomalous ones (the $d-3$ counted in Chapter 5) survive as baryonic global symmetries. We never write the brane gauge group as a bare SU.

Edges are bifundamentals. An edge of the tiling lies between two faces i and j, and carries a chiral bifundamental X_{ij} in the representation $(\mathbf{N}_i, \overline{\mathbf{N}}_j)$, a fundamental of the node on one side and an antifundamental of the node on the other. The orientation of the arrow $i \to j$ is fixed by the colouring through the rule below, so that the quiver, an oriented graph of nodes and arrows, is determined by the tiling.

Nodes are superpotential terms. Each graph node, black or white, is a single-trace superpotential term. The edges meeting a node, read in cyclic order around it, spell a closed loop of bifundamentals, hence a gauge-invariant word $\mathrm{Tr}(X_{i_1 i_2} X_{i_2 i_3} \cdots)$. The colour of the node fixes the orientation of the reading and the sign of the term. This is the convention of the whole volume, used throughout:

$$W = \sum_{\text{white nodes}} \mathrm{Tr}\big(\text{edges clockwise}\big) - \sum_{\text{black nodes}} \mathrm{Tr}\big(\text{edges counterclockwise}\big), \tag{VII.3.3}$$

with the displayed couplings normalized to ± 1 in the toric presentation. The colour is the surface normal by the right-hand rule: a white node is read clockwise and enters W with a plus sign, a black node is read counterclockwise and enters with a minus sign.

The sign rule is not free decoration. It is fixed so that the F-term equations of W reproduce the relations of the toric quiver, and the simplest tilings show it doing exactly that (§VII.3.4). Its single structural consequence is worth stating on its own, because the consistency conditions of §VII.3.8 lean on it. In a bipartite graph every edge joins one black node to one white node. Reading the dictionary backwards, every bifundamental therefore sits in exactly two superpotential terms, one from each colour, and because the two colours enter W with opposite signs,

$$\text{each field appears in exactly two } W\text{-terms, with opposite signs.} \tag{VII.3.4}$$

This is the dimer form of the toric-quiver fact that the McKay quivers of Volume VI also obeyed: the abelian orbifold superpotentials had each field in two terms of opposite sign. The dimer makes it automatic, a property of the graph.

F-terms are path equalities. The two-terms structure is what makes the sign rule more than bookkeeping: it turns each F-term into a geometric statement. Take any edge X. It bounds exactly one white node and one black node, so it sits in exactly two terms of W, the white one contributing $+\operatorname{Tr}(XP)$ with P the product of the other edges read around that node, the black one $-\operatorname{Tr}(XQ)$ with Q the product around it. Cyclic differentiation gives the F-term

$$\frac{\partial W}{\partial X} = P - Q = 0, \qquad \text{equivalently} \qquad P = Q, \tag{VII.3.5}$$

the statement that the two quiver paths P and Q, which run between the head and tail of X, are equal in the chiral ring. For this presentation every F-term is such an equality of two paths, one of each colour. These binomial relations are the starting point for the mesonic moduli construction, including its gauge quotient and choice of component. The two-occurrence condition alone is not a theorem that every bipartite graph has a Calabi–Yau threefold moduli space. The honeycomb of §VII.3.4 is the smallest instance: there $P = YZ$ and $Q = ZY$, so $\partial_X W = [Y, Z]$.

The arrow-crossing rule. The orientation of an arrow is fixed by demanding that (VII.3.3) come out consistently when the same edge is read from its two nodes. The rule, stated for use, is that the quiver arrows circulate *clockwise* around a white node and *counterclockwise* around a black node, equivalently that an arrow keeps the white node on its right. With this rule every edge inherits a single consistent orientation, and the cyclic word at each node is a genuine closed oriented loop, so each term of (VII.3.3) is gauge invariant. The arrow direction and the term sign are then two faces of the one colouring.

The Euler constraint. A cellular graph on T^2 obeys the surface Euler relation, so the cell counts are not independent. Writing V for the number of graph nodes, E for the edges, and F for the faces,

$$V - E + F = \chi(T^2) = 0. \tag{VII.3.6}$$

In gauge-theory terms V is the number of superpotential terms, E the number of fields, and F the number of gauge nodes, so this cellular torus presentation has $\#\{W \text{ terms}\} + \#\{\text{gauge nodes}\} = \#\{\text{fields}\}$. This is a cheap but real check on any proposed tiling, and it is the first thing the worked examples below satisfy.

VII.3.3 The periodic quiver

The dimer encodes the same data as the quiver-with-superpotential, but in a dual layout, and it is worth seeing the two side by side once. The bridge is the **periodic quiver**: the ordinary quiver, with its superpotential, drawn on the same torus T^2 (or on its universal cover, the plane $\mathbb{R}^2$, where the periodicity is manifest).

The periodic quiver is the planar dual of the tiling, with one geometric exchange:

$$\begin{aligned} \text{tiling face} &\longleftrightarrow \text{quiver node}, \\ \text{tiling edge} &\longleftrightarrow \text{quiver arrow}, \\ \text{tiling node} &\longleftrightarrow \text{quiver face (plaquette)}. \end{aligned} \tag{VII.3.7}$$

A node of the periodic quiver sits inside a face of the tiling; an arrow of the periodic quiver crosses a tiling edge, picking up its orientation from the arrow-crossing rule; and a minimal closed loop, a plaquette, of the periodic quiver surrounds a node of the tiling. Reading the plaquette as the product of the arrows around it gives a closed gauge-invariant word, which is precisely the superpotential term at the enclosed tiling node. The plaquettes of the periodic quiver are therefore the terms of W, with the plaquettes around white-encircling nodes entering with one sign and those around black with the other. Figure VII.3.1 shows the duality on the conifold square.

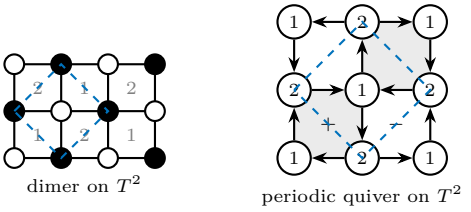


Figure VII.3.1: The tiling and its periodic quiver, on the conifold square, each drawn as a patch on the universal cover of T^2 with one fundamental domain dashed. The diamond has period vectors $(1,1)$ and $(1,-1)$; opposite sides are identified, preserving the checkerboard colors, gauge labels, and arrow pattern. Read the left panel as the dimer (faces $1,2$ are gauge nodes; the edges are the bifundamentals; the white and black vertices are the two superpotential terms); after boundary identification the dashed cell holds one white and one black vertex, four edges, and two faces. Read the right panel as its planar dual: the checkerboard nodes $1,2$ are the gauge groups (one per tiling face), the arrows are the bifundamentals (one per tiling edge, with the arrow-crossing orientation), and each square *plaquette* is a closed four-cycle of arrows, a quartic W-term, surrounding a tiling node. Following the stated rule of §VII.3.2: the plaquette surrounding a white tiling node is read clockwise and enters with the plus sign (the shaded plaquettes), the one around a black node counterclockwise with the minus sign (unshaded); each edge borders one of each, so every field sits in two W-terms of opposite sign. Both panels carry the quiver *and* the superpotential, the whole theory.

Reading one plaquette makes the duality concrete. Pick a dimer edge between faces 1 and 2; the quiver arrow crossing it is oriented with the white node on its right, by the arrow-crossing rule of §VII.3.2. Follow the four arrows around the plaquette surrounding the adjacent white tiling node, read clockwise: their product is the plus quartic word $\mathrm{Tr}(A_1B_1A_2B_2)$. Follow the four arrows around the adjacent black-node plaquette, counterclockwise: their product is the minus word $\mathrm{Tr}(A_1B_2A_2B_1)$. The same edge borders both plaquettes, so each field enters two W-terms of opposite sign, the structural fact (VII.3.4) read straight off the periodic quiver.

Two remarks set up the rest of the volume. First, the universal cover $\mathbb{R}^2$, not the single fundamental domain, is the natural drawing surface, both for the periodic quiver and for the zig-zag paths of §VII.3.8, because a closed loop on T^2 can wind around the torus and is seen as an open, periodically repeating path upstairs. Second, the periodic quiver is the object on which the dimer *moves* act: a local rearrangement of the tiling that preserves the toric singularity but changes the quiver is the combinatorial face of Seiberg duality, and it is performed in Chapter 7. This chapter takes a tiling as given and reads its data; it does not move it.

VII.3.4 The simplest fixture: the $\mathbb{C}^3$ honeycomb

The smallest brane tiling is the hexagonal one, and it is worth doing entirely by hand before any machinery, because every later computation specializes to it. It is the dimer of flat space $\mathbb{C}^3$, equivalently of $\mathcal{N}$=4 super Yang–Mills viewed as the trivial "orbifold" by $\Gamma = 1$, the parent theory of all of Volume V.

The honeycomb tiling of T^2 has, in one fundamental domain,

$$(V, E, F) = (2, 3, 1), \qquad V - E + F = 2 - 3 + 1 = 0, \tag{VII.3.8}$$

one white node, one black node, three edges, and one face, satisfying the Euler constraint (VII.3.6). The dictionary reads off the gauge theory immediately. The single face is one gauge node,

$$G_{\text{gauge}} = U(N), \tag{VII.3.9}$$

and since both ends of every edge land on the same face, the three edges are three *adjoint* chiral fields, which we call X, Y, Z. They are the three complex coordinates of $\mathbb{C}^3$, transverse to the D3-branes. There is one white and one black node, each meeting all three edges, so each is trivalent and each superpotential term is a cubic word.

Apply the sign rule (VII.3.3). The white node, read clockwise, gives the cyclic word $\mathrm{Tr}(XYZ)$ with a plus sign; the black node, read counterclockwise, gives the opposite cyclic order $\mathrm{Tr}(XZY)$ with a minus sign, so

$$W = \mathrm{Tr}(XYZ) - \mathrm{Tr}(XZY) = \mathrm{Tr}\ X[Y, Z]. \tag{VII.3.10}$$

This is the $\mathcal{N}$=4 cubic superpotential, the corresponding parent of the whole volume, recovered from a graph with two nodes and three edges. Each of X, Y, Z sits in both terms, with opposite signs, the structural fact (VII.3.4): the commutator $[Y, Z]$ is exactly the two terms of opposite sign sharing the edges Y and Z. The honeycomb is the cleanest illustration that the colouring *is* the sign rule: flipping the two colours would send $W \to -W$, and choosing the same colour for both nodes would give $\mathrm{Tr}(XYZ) + \mathrm{Tr}(XZY)$, not a commutator, and the wrong moduli space.

That the moduli space is right is a one-line check. The F-terms are the cyclic derivatives of (VII.3.10); taking the X-derivative term by term,

$$\frac{\partial}{\partial X}\,\mathrm{Tr}(XYZ) = YZ, \qquad \frac{\partial}{\partial X}\,\mathrm{Tr}(XZY) = ZY, \tag{VII.3.11}$$

so $\partial_X W = YZ - ZY = [Y, Z]$, and the same for the other two,

$$\partial_X W = [Y, Z], \qquad \partial_Y W = [Z, X], \qquad \partial_Z W = [X, Y], \tag{VII.3.12}$$

so the supersymmetric vacua have X, Y, Z commuting. For one brane ($N = 1$) the fields are numbers, the commutators vanish identically, and the three unconstrained scalars are the position of the brane in $\mathbb{C}^3$: the moduli space is $\mathbb{C}^3$, as it must be. For N branes the F-flat commuting triples must also satisfy the D-term equation $[X, X^\dagger] + [Y, Y^\dagger] + [Z, Z^\dagger] = 0$ and be divided by $U(N)$. Commuting complex matrices alone need not be diagonalizable: a nonzero nilpotent Jordan block and two zero matrices give a counterexample. To see what the D-term adds, simultaneously triangularize a commuting triple by a unitary change of basis. The first diagonal entry of the D-term is a sum of squared magnitudes of the off-diagonal entries in the first row. It vanishes only when all those entries vanish. Repeating on the remaining block makes the triple diagonal. The remaining permutations identify the eigenvalue triples, giving $\mathrm{Sym}^N(\mathbb{C}^3)$ as the classical vacuum space at zero FI parameter, the branch of N unordered D3-brane positions. The honeycomb is thus a fully worked dimer, not merely a named one.

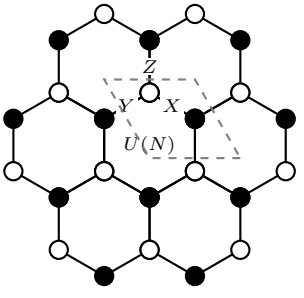


Figure VII.3.2: The C^3 honeycomb tiling, drawn as a patch of the hexagonal lattice on the universal cover of T^2 so the periodicity is manifest; the dashed rhombus is one fundamental domain. How to read it: each hexagonal face is the single gauge node $U(N)$; the three edge directions X, Y, Z are the three adjoint chiral fields; the white and black vertices are the two cubic superpotential terms of opposite sign, which combine into $\mathrm{Tr}\, X[Y, Z]$. One fundamental domain contains one white and one black vertex, the three edges, and one face; the windings $(0,0), (1,0), (0,1)$ of X, Y, Z are the displacements between the cell's vertices and their copies in the neighbouring cells, measured in the reflected ordered period basis $a'_1 = (-1.143, 0)$, $a'_2 = (-0.5715, 0.99)$ (rounded drawing coordinates). This basis has opposite handedness to the page axes; the stated clockwise reading is clockwise on the page.

Figure VII.3.2 draws the tiling. Two of its three edges wrap around the torus, and that winding is precisely the data the next section needs. The honeycomb is the case in which every count is one: one gauge node, one white and one black node, and, as §VII.3.6 shows, one perfect matching per lattice point. Everything that grows in a general tiling is one here.

VII.3.5 The Kasteleyn matrix and the toric diagram

The dictionary turns a tiling into a gauge theory. The reverse direction, from the tiling to the geometry it probes, is computed by a single determinant. This is the route that makes the dimer a way to compute the geometry.

The Kasteleyn matrix. Choose a fundamental domain and label the two cycles of T^2 by formal weights z and w. The **Kasteleyn matrix** $K(z,w)$ of a tiling is the weighted, signed adjacency matrix between the white and the black nodes:

$$K_{ab}(z,w) = \sum_{\text{edges } e:\, w_a \to b_b} \sigma_e \, z^{p_e} \, w^{q_e}, \tag{VII.3.13}$$

the rows a indexed by the white nodes, the columns b by the black nodes, the sum running over the edges joining them. Each edge contributes its Kasteleyn sign $\sigma_e = \pm 1$ and a monomial $z^{p_e} w^{q_e}$ recording how many times that edge crosses the two boundaries of the fundamental domain, that is, its translation $(p_e, q_e) \in \mathbb{Z}^2$ between chosen vertex lifts. An individual edge is open and has no intrinsic homology class. Closed paths and differences of perfect matchings have homology classes; changing vertex lifts changes the edge translations but only shifts the matching exponents together. An edge wholly inside the domain has winding $(0,0)$ and contributes a constant. For a balanced tiling the numbers of white and black nodes are equal, so K is square and its determinant is defined.

The signs σ_e obey a face rule. For each disk face with $2m$ boundary edges, counted with repetitions,

$$\prod_{e \in \partial f} \sigma_e = (-1)^{m+1}. \tag{VII.3.14}$$

Thus a square needs an odd number of negative signs and a hexagon an even number. For a finite balanced bipartite cellular torus graph this signing exists. With positive edge weights, matchings in each fixed height sector contribute with one sign, although different sectors can have different signs. A single torus determinant is therefore a signed generating function, not the unsigned partition function. The existence and same-sector sign statements are mathematical results used here.

The characteristic polynomial and its Newton polygon. The output of the construction is the determinant, a Laurent polynomial in z and w,

$$P(z,w) = \det K(z,w) = \sum_{(p,q)} c_{pq} \, z^p w^q, \tag{VII.3.15}$$

called the characteristic polynomial of the tiling (in the mathematics literature, the equation $P = 0$ cuts out the spectral curve). The geometry is read off its exponents. The **Newton polygon** of P is the convex hull of the set of exponent vectors (p,q) of its monomials,

$$\mathrm{Newt}(P) = \mathrm{conv}\{(p,q) : c_{pq} \neq 0\} \subset \mathbb{R}^2, \tag{VII.3.16}$$

For a consistent brane tiling representing the mesonic branch, the physical identification is that this polygon is the toric diagram of the Calabi–Yau threefold the brane probes,

$$\mathrm{Newt}(\det K) \;=\; \Delta_{\mathrm{toric}}, \tag{VII.3.17}$$

up to the unimodular changes of basis $GL(2,\mathbb{Z})$ and the translations that the toric diagram is defined modulo. The tiling, through one determinant, computes the singularity.

The honeycomb, worked. The honeycomb makes the route concrete with no determinant to expand, because with one white and one black node the matrix K is 1×1 and $\det K = K$ itself. Its single entry sums the three edges X, Y, Z with the windings of Figure VII.3.2, $(0,0)$, $(1,0)$, and $(0,1)$. A Kasteleyn orientation here is just a choice of three signs on the three edges; taking $(+,-,-)$ on (X,Y,Z) gives

$$K_{\mathbb{C}^3}(z,w) = 1 - z - w, \qquad P_{\mathbb{C}^3}(z,w) = 1 - z - w. \tag{VII.3.18}$$

Table VII.3.1 assembles this entry edge by edge, so the determinant is reproducible rather than quoted. The sign choice is not fitted to the answer: a different consistent choice, or multiplying K by an overall monomial $z^a w^b$ or by -1, only translates the Newton polygon or flips the polynomial's sign, neither of which changes the toric diagram (which is the polygon up to $GL(2,\mathbb{Z})$ and translation).

edge	winding (p,q)	Kasteleyn sign	monomial $\sigma\, z^p w^q$
X	$(0,0)$	$+$	$+1$
Y	$(1,0)$	$-$	$-z$
Z	$(0,1)$	$-$	$-w$

Table VII.3.1: The 1×1 Kasteleyn entry of the $\mathbb{C}^3$ honeycomb, edge by edge: each of the three edges contributes one signed, winding-weighted monomial, and their sum is $K_{\mathbb{C}^3} = 1 - z - w$ (VII.3.18). The signs are the Kasteleyn orientation; the windings are read off the fundamental domain of Figure VII.3.2.

The exponents of its three monomials are

$$\{(0,0),(1,0),(0,1)\}, \tag{VII.3.19}$$

whose convex hull is the unit right triangle. This is the toric diagram of $\mathbb{C}^3$, the unit triangle with no interior point, recovered from the tiling. The honeycomb has done in one line what the toric construction of the earlier chapters did by hand: the brane reads its transverse $\mathbb{C}^3$ off the graph.

Twice the area is the number of gauge nodes. One integer read off the Newton polygon links it back to the gauge theory and is a useful cross-check. For the consistent toric phases under discussion, the number of gauge nodes equals twice the Euclidean area (in lattice units) of the toric diagram,

$$G = \#\{\text{gauge nodes}\} = 2\ \mathrm{Area}\big(\mathrm{Newt}(\det K)\big). \tag{VII.3.20}$$

The honeycomb triangle has area $\frac{1}{2}$, so $G = 1$, the single $U(N)$; the conifold square, recalled below, has area 1, so $G = 2$. The same count came from the faces of the tiling directly, so (VII.3.20) ties the two ends of the dictionary together.

VII.3.6 Perfect matchings and the lattice points

The Newton polygon (VII.3.16) records only the convex hull of the exponents. The finer object, the one the moduli-space construction of Chapter 5 runs on, is the full set

of *lattice points* of the toric diagram together with how many ways the tiling realizes each. That counting is done by the perfect matchings of the graph.

A **perfect matching** of a dimer is a subset of its edges that covers every node, black and white, exactly once. Equivalently it is a way of pairing each white node with an adjacent black node, using each node once.

Why the determinant counts matchings. The link to the Kasteleyn matrix is the permutation expansion of the determinant, and it is worth running once. With the white nodes indexing the rows and the black nodes the columns, an $n \times n$ K has

$$\det K(z,w) = \sum_{\sigma \in S_n} \operatorname{sgn}(\sigma) \prod_{a=1}^{n} K_{a,\sigma(a)}, \tag{VII.3.21}$$

a sum over the $n!$ permutations. A term is nonzero only when every factor $K_{a,\sigma(a)}$ is a genuine edge, that is, when σ pairs each white node a with a black node $\sigma(a)$ it is actually joined to. When each entry $K_{a,\sigma(a)}$ is a single edge (as for dP_0 here) such a σ *is* a perfect matching; in general an entry can be a sum over several parallel edges, and one then expands each entry and chooses one edge per row, each choice a perfect matching weighted by $\operatorname{sgn}(\sigma)$ times the product of its edge Kasteleyn signs. Either way the product of the chosen edges carries $\operatorname{sgn}(\sigma) \prod_{e \in M} \sigma_e \, z^p w^q$ with (p,q) the matching's net winding, so the determinant is exactly the signed generating function of the perfect matchings,

$$\det K(z,w) = \sum_{\text{perfect matchings } M} \operatorname{sign}(M)\, z^{p(M)} w^{q(M)}, \tag{VII.3.22}$$

which is (VII.3.21) read combinatorially. The dP_0 tiling makes this concrete: its Kasteleyn matrix is 3×3, so (VII.3.21) runs over the $|S_3| = 6$ permutations, and all six give genuine edges of $K_{3,3}$, the six perfect matchings of §VII.3.7. Three of them (the three transpositions) land on the interior point with sign $-$, giving its coefficient -3; the other three (identity and the two 3-cycles) land on the three corners with sign $+$. A perfect matching therefore carries a net winding $(p,q) \in \mathbb{Z}^2$, its position in the plane of exponents, and the determinant collects matchings by position. The multiplicity of a lattice point is *defined* as the number of perfect matchings landing on it; the absolute coefficient $|c_{pq}|$ of (VII.3.15) equals that count whenever the matchings over (p,q) do not cancel in sign, which is the case for the corrected Kasteleyn signings used here. The unsigned count is instead obtained by the permanent of the adjacency matrix with positive edge weights. For example, the all-positive dP_0 adjacency matrix below has permanent $z^{-1}+w^{-1}+3z^{-1}w^{-1}+z^{-2}w^{-2}$ and determinant with interior coefficient -3. Both encode six matchings after signs are removed sector by sector. Merely matching the Newton support cannot test either the signs or the multiplicities. The map

$$\{\text{perfect matchings}\} \;\longrightarrow\; \{\text{lattice points of } \Delta_{\text{toric}}\} \tag{VII.3.23}$$

is the object of the moduli-space construction, and the single fact this chapter needs about it is that *it is not a bijection.* Several perfect matchings can land on the same lattice point, so a lattice point carries a *multiplicity*, the number of matchings over it.

This is the toric face of the statement that the perfect matchings are the homogeneous coordinates of a gauged linear sigma model whose vacuum moduli space is the singularity; assembling those coordinates and computing the moduli space and its Hilbert series is the forward construction of Chapter 5. Here the matchings are enumerated and tied to the lattice points, but the moduli space is not yet built.

The honeycomb and conifold are the bijective cases. For the honeycomb, with one white and one black node, a perfect matching is a single edge: there are three matchings, $\{X\}$, $\{Y\}$, $\{Z\}$, landing at the three corners $(0,0)$, $(1,0)$, $(0,1)$ of the triangle, one each, all of multiplicity one. The map (VII.3.23) is here a bijection. The same is true for the conifold square, with four matchings at the four corners of the unit square. Table VII.3.2 lists these matchings with the exponents they land on, so the bijection is read off directly. These are the special cases in which there are no interior lattice points and no multiplicities, the ones for which the dimer is simplest; the next section exhibits the first case where the map is genuinely many-to-one.

fixture	perfect matching	exponent (p,q)	multiplicity
$\mathbb{C}^3$	$\{X\}$	$(0,0)$	1
$\mathbb{C}^3$	$\{Y\}$	$(1,0)$	1
$\mathbb{C}^3$	$\{Z\}$	$(0,1)$	1
conifold	$\{A_1\}$	$(0,0)$	1
conifold	$\{B_1\}$	$(1,0)$	1
conifold	$\{A_2\}$	$(1,1)$	1
conifold	$\{B_2\}$	$(0,1)$	1

Table VII.3.2: The perfect matchings of the two bijective fixtures. With one white and one black node a matching is a single edge; the three honeycomb edges land at the three triangle corners and the four conifold edges at the four square corners, each lattice point hit exactly once. The map (VII.3.23) is a bijection here; dP_0 breaks it, with the interior point hit three times (§VII.3.7).

VII.3.7 Worked fixtures: the conifold square and dP_0

Two more tilings put the machinery through its paces. The first, the conifold square, was already met in Chapter 1, where its gauge theory was derived from the geometry; here it is recalled only as a dimer, to confirm that the dictionary reproduces it. The second, the dimer of $\mathbb{C}^3/\mathbb{Z}_3 = dP_0$, is the first example with a genuine multiplicity, and it is the bridge back to the orbifold construction of Volume V.

The conifold square (recalled). The conifold tiling is the square lattice: one white node, one black node, four edges, two faces,

$$(V,E,F) = (2,4,2), \qquad V - E + F = 0. \tag{VII.3.24}$$

The two faces are the gauge nodes $U(N)_1 \times U(N)_2$; the four edges are the bifundamentals A_1, A_2 from node 1 to node 2 and B_1, B_2 back. Both nodes are four-valent, so each superpotential term is a quartic word. The sign rule (VII.3.3) reads the white node

clockwise as $\mathrm{Tr}(A_1B_1A_2B_2)$ and the black node counterclockwise as $\mathrm{Tr}(A_1B_2A_2B_1)$, giving

$$W = \mathrm{Tr}(A_1B_1A_2B_2) - \mathrm{Tr}(A_1B_2A_2B_1) = \tfrac{1}{2}\varepsilon^{ij}\varepsilon^{kl}\,\mathrm{Tr}(A_iB_kA_jB_l), \qquad \text{(VII.3.25)}$$

the Klebanov–Witten quartic, with $\varepsilon^{12} = 1$. The factor $1/2$ compensates for the two cyclic copies of each trace in the double epsilon sum. Each of the four fields again in two terms of opposite sign. The Kasteleyn matrix is 1×1, summing the four edges with windings $(0,0), (1,0), (1,1), (0,1)$,

$$P_{\text{conifold}}(z,w) = 1 - z - w - zw, \qquad \text{(VII.3.26)}$$

assembled edge by edge in Table VII.3.3, whose four exponents are the corners of the unit square, the conifold toric diagram.

edge	winding (p,q)	Kasteleyn sign	monomial $\sigma\, z^p w^q$
1	$(0,0)$	$+$	$+1$
2	$(1,0)$	$-$	$-z$
3	$(1,1)$	$-$	$-zw$
4	$(0,1)$	$-$	$-w$

Table VII.3.3: The 1×1 Kasteleyn entry of the conifold square, edge by edge: the four edges (the bifundamentals A_1, A_2, B_1, B_2) carry the windings $(0,0), (1,0), (1,1), (0,1)$ and signs $(+,-,-,-)$, summing to $P_{\text{conifold}} = 1 - z - w - zw$ (VII.3.26). The four exponents are the corners of the unit square. The sign product is -1, as required on each square face. The polynomial $1 - z - w + zw$ has the same support but fails that face rule.

The four perfect matchings are the four single edges, one per corner, all of multiplicity one. The dimer reproduces the gauge group, matter, quartic words, and toric diagram discussed in Chapter 1. This is the point of the construction. For the conifold the geometry was a shortcut; the dimer needs no shortcut.

The dP_0 tiling and its first multiplicity. The orbifold $\mathbb{C}^3/\mathbb{Z}_3$, the singular cone associated with $dP_0 = \mathbb{P}^2$, whose crepant resolution is the total space of $\mathcal{O}_{\mathbb{P}^2}(-3)$, was built in Volume VI from the McKay quiver of $\mathbb{Z}_3$: three gauge nodes, nine bifundamentals Φ_i^a $(a = 1,2,3)$ running $i \to i+1$ cyclically, and a cubic superpotential. The same theory is a brane tiling, and seeing it as one is the cleanest statement of the bridge "the abelian orbifold is the abelian-toric special case of the dimer." Its tiling is the hexagonal lattice with an enlarged fundamental domain holding three of the honeycomb's cells; as a bipartite graph it is the complete bipartite graph $K_{3,3}$, with

$$(V,E,F) = (6,9,3), \qquad V - E + F = 6 - 9 + 3 = 0, \qquad \text{(VII.3.27)}$$

three white and three black nodes (the six cubic W-terms), nine edges (the nine bifundamentals), and three faces (the three gauge nodes $U(N)^3$). Each node is trivalent, so every W-term is cubic, and the sign rule assembles the six trivalent nodes into the cubic superpotential of the $\mathbb{Z}_3$ daughter, the six S_3 orderings of one field from each arrow class with alternating signs.

The cubic words can be recovered without guessing their order. Name the edge joining white vertex i to black vertex j by P_{ij}, with $i, j = 0, 1, 2$. Clockwise orders at the white vertices are (P_{i0}, P_{i1}, P_{i2}); at black vertex j use (P_{0j}, P_{2j}, P_{1j}). These orders give three hexagonal faces. With face labels chosen so that P_{00} goes from 0 to 1, the three rows of arrow tails are $(0, 1, 2)$, $(1, 2, 0)$ and $(2, 0, 1)$, with each head one step forward modulo three. Reading the black cycles counterclockwise therefore gives

$$W = \sum_{i=0}^{2} \mathrm{Tr}(P_{i0}P_{i1}P_{i2}) - \sum_{j=0}^{2} \mathrm{Tr}(P_{0j}P_{1j}P_{2j}). \tag{VII.3.28}$$

Each trace closes on the quiver, each field occurs once with either sign, and cyclic differentiation equates the two complementary length-two paths. The rotations specify which hexagonal embedding of $K_{3,3}$ is meant; the abstract graph alone would not determine these faces.

The new feature is in the toric output. The Kasteleyn matrix is now 3×3, and its determinant has a Newton polygon that is the toric diagram of dP_0, the local-$\mathbb{P}^2$ triangle (three times the area of the $\mathbb{C}^3$ one) with a single *interior* lattice point,

$$\Delta_{dP_0} \text{ has 3 boundary points and 1 interior point,} \tag{VII.3.29}$$

the interior point being the compact divisor $\mathbb{P}^2$ at the bottom of the resolution. Counting twice the area gives $2 \cdot \frac{3}{2} = 3$ gauge nodes, matching the three faces. The perfect-matching map (VII.3.23) is here genuinely many-to-one. Enumerating the perfect matchings of $K_{3,3}$ and reading their windings, the three boundary points each receive one matching, but the interior point receives

$$\text{multiplicity of the interior point} = 3, \tag{VII.3.30}$$

three distinct perfect matchings landing on the same lattice point, equal to the absolute coefficient of the interior monomial in $\det K$ (the exponent of that monomial depends on the torus-coordinate choice, the value of the coefficient does not). It is worth expanding the determinant once rather than quoting it. With the three white nodes indexing the rows and the three black nodes the columns, and the windings recorded in the torus weights z, w, the Kasteleyn matrix of the dP_0 tiling is

$$K_{dP_0}(z, w) = \begin{pmatrix} 1 & 1 & 1 \\ 1 & \dfrac{1}{wz} & \dfrac{1}{z} \\ 1 & \dfrac{1}{w} & \dfrac{1}{wz} \end{pmatrix}, \tag{VII.3.31}$$

a 3×3 matrix because there are three nodes of each colour. Expanding along the first row,

$$\det K_{dP_0} = \left(\tfrac{1}{w^2z^2} - \tfrac{1}{wz}\right) - \left(\tfrac{1}{wz} - \tfrac{1}{z}\right) + \left(\tfrac{1}{w} - \tfrac{1}{wz}\right) = z^{-1} + w^{-1} - 3\, z^{-1}w^{-1} + z^{-2}w^{-2}, \tag{VII.3.32}$$

whose four monomials are the four lattice points and whose interior coefficient has magnitude 3, the multiplicity, now read off an explicit determinant rather than quoted. This is the first time the convex hull undercounts: the Newton polygon has four lattice points, but the dimer has more than four perfect matchings, because the interior point is hit three times. The multiplicity is exactly the data the moduli machine of Chapter 5 needs and the convex hull alone does not see, and it is why perfect matchings, not just the toric diagram, are the right variables for the moduli space.

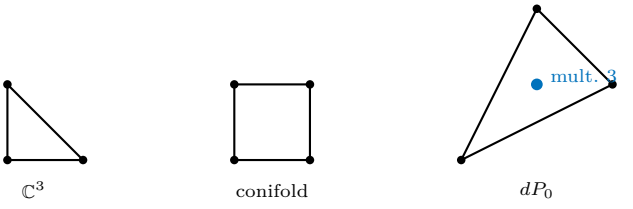


Figure VII.3.3: The three fixture toric diagrams, each computed as the Newton polygon of $\det K$. Read each as the lattice points (p, q) of the monomials of the Kasteleyn determinant. $\mathbb{C}^3$ is the unit triangle (three boundary points, no interior); the conifold is the unit square (four boundary points); dP_0 is the local-$\mathbb{P}^2$ triangle whose single interior lattice point (blue) is hit by three perfect matchings, the first genuine multiplicity. Diagrams are compared up to $GL(2, \mathbb{Z})$ and translation, so the absolute placement is not significant.

The three diagrams are collected in Figure VII.3.3. The dP_0 tiling is the honest first test of the full machinery: a non-trivial Kasteleyn determinant, a Newton polygon with an interior point, and a perfect-matching count that the convex hull cannot reproduce. That its three independent computations agree, the determinant's Newton polygon, the perfect-matching positions, and (next section) the zig-zag classes, checks several parts of the same embedded example. Homology data alone would not check its universal-cover intersections or establish its D3-brane interpretation.

VII.3.8 Zig-zag paths and consistency

An embedded bipartite graph supplies a quiver with a superpotential. To identify it with a toric D3-brane theory requires more: the embedding must satisfy appropriate consistency conditions, and the gauge theory must have the intended mesonic branch and physical limit. Zig-zags expose information that the determinant forgets, because they use the cyclic orders as well as the edge translations.

Winding balance is a necessary test

For a consistent toric phase, zig-zag homology matches the external web legs. The matching does not by itself test intersections in the universal cover or prove the existence of a superconformal fixed point. A graph can define a quiver and fail the proposed toric interpretation.

Zig-zag paths. A **zig-zag path** is a closed path on the tiling that turns as sharply as it can, and alternately. Arriving at a node along an edge, leave by the next edge in the cyclic order around that node, turning maximally left at every white node and maximally right at every black node (or the mirror convention, applied consistently). Because the turn alternates between the colours, the path zig-zags, and because the graph is finite on T^2 the path closes up. Every directed edge lies on exactly one zig-zag path (here a *directed* edge is one oriented traversal of an edge; each unoriented dimer edge is traversed by two zig-zags, one in each direction).

A closed path on T^2 has a homology class, its net winding around the two cycles of the torus,

$$[\,\eta\,] = (p, q) \in H_1(T^2) = \mathbb{Z}^2, \tag{VII.3.33}$$

and the slope q/p is the direction the zig-zag travels on the torus. The zig-zag classes are computed from the graph alone, by tracing the alternating turns and adding up the edge windings, with no reference to the geometry. This gives a separate combinatorial comparison once the cyclic orders and translations are fixed. If those data were chosen to fit the expected slopes, agreement is a consistency check of that choice, not an independent derivation of the physical model.

Zig-zags are the (p, q)-web legs. For a geometrically consistent toric phase the multiset of zig-zag homology classes agrees with the primitive outward normals to the boundary edges of the toric diagram, repeated according to each edge's lattice length,

$$\{\text{zig-zag classes}\} \;=\; \{\text{outward normals of } \partial\Delta_{\text{toric}}\} \;=\; \{(p, q)\text{-web external legs}\}, \tag{VII.3.34}$$

up to $GL(2, \mathbb{Z})$. The (p, q)-web is the graph dual to the toric diagram, with external legs perpendicular to the boundary edges, repeated with those lattice multiplicities; the zig-zag paths of the tiling are the combinatorial origin of those legs. Since the boundary edges of a polygon close up, their length-weighted primitive normals sum to zero. Zig-zag windings also sum to zero directly: the path permutation uses each directed edge once, so each translation occurs with both signs. Thus every such complete zig-zag decomposition, even an inconsistent one, obeys

$$\sum_{\text{zig-zags}} [\,\eta\,] = (0, 0). \tag{VII.3.35}$$

This correspondence is the precise sense in which the tiling and the geometry agree on the external structure of the singularity, the same external legs that count the baryonic symmetries ($d - 3$ of them for a diagram with d boundary points) in Chapter 1.

The three fixtures. On the honeycomb the three edges trace three zig-zag paths with classes

$$\{(0, -1), (1, 0), (-1, 1)\}, \qquad \text{sum} = (0, 0), \tag{VII.3.36}$$

the three outward normals of the $\mathbb{C}^3$ triangle (up to $GL(2, \mathbb{Z})$), summing to zero: three legs, the trivalent $\mathbb{C}^3$ web. These classes are not quoted but traced: each honeycomb zig-zag alternates two of the three edges, traversing one forward and the next backward,

so its homology class is the *difference* of the two edge windings of Table VII.3.1. Running the three alternations,

$$\begin{aligned} Y \text{ then } X &: (1,0) - (0,0) = (1,0), \\ Z \text{ then } Y &: (0,1) - (1,0) = (-1,1), \\ X \text{ then } Z &: (0,0) - (0,1) = (0,-1), \end{aligned} \tag{VII.3.37}$$

reproduces (VII.3.36) edge by edge, and the three differences sum to zero because the three windings cancel in the cyclic sum. The conifold has four zig-zags,

$$\{(0,-1),(1,0),(0,1),(-1,0)\}, \qquad \text{sum} = (0,0), \tag{VII.3.38}$$

the four normals of the unit square, the four-legged conifold web. The dP_0 tiling has three zig-zags whose primitive classes match the three outward normals of its local-$\mathbb{P}^2$ triangle, again summing to zero. In the translation basis of (VII.3.31), they are $(1,2)$, $(-2,-1)$ and $(1,-1)$. For example, one orbit reads $P_{00}^{-1}P_{01}P_{11}^{-1}P_{12}P_{22}^{-1}P_{20}$, where the exponent -1 denotes traversal from black to white, not an inverse field. Adding its signed translations gives $(1,2)$. The other two orbits use the remaining directed edges and give the other two vectors. In every case the classes are traced from the graph's turning data and the edge windings alone, then matched to the toric normals computed independently from the diagram; the agreement is the third leg of the self-consistency, alongside the Kasteleyn Newton polygon and the perfect-matching positions.

What consistency demands. The signed algebraic intersection of classes $u = (p,q)$ and $v = (r,s)$ is $ps - qr$. Its absolute value is a lower bound on the number of transverse crossings, not the signed intersection itself. A strong geometric condition, appropriate to the three regular lattice examples here, asks that lifted zig-zags be simple and that any two distinct lifted paths meet at most once. In particular, primitive parallel paths then have no crossings, and torus crossings realize the minimum $|ps - qr|$. This strong condition must be distinguished from weaker dimer consistency conditions formulated as the absence of repeated intersections with the same orientation. We do not identify these criteria for arbitrary embeddings.

A homology calculation cannot settle these questions. Two curves can acquire a pair of oppositely oriented crossings without changing either winding or the signed intersection. One must follow the actual lifted curves, including the translations at their shared edges. For the regular honeycomb and square drawings the paths can be drawn as non-self-intersecting train tracks with the stated minimal intersections; the enlarged honeycomb for dP_0 has the same local geometry. This geometric reading uses those regular embeddings, not just a list of winding vectors.

There are useful failure tests that require less information. Two edges between one white and one black vertex with identical translations give zero-winding zig-zags; they cannot supply the legs of a two-dimensional toric fan. Doubling the translations of the honeycomb without enlarging the cell gives nonprimitive windings. It fails the primitive-leg test for that proposed cell; it does not disprove the existence of an orbifold tiling with a correctly enlarged fundamental domain.

A sequence of consistency tests. First check that the graph is a cellular torus embedding and that the edge translations close around every face. The dictionary then

gives the signed cyclic words and the quiver. At equal regular ranks, count incoming and outgoing bifundamentals at each node to test the non-abelian gauge anomaly. An adjoint contributes no cubic gauge anomaly. The two-occurrence property of W follows separately from bipartiteness. Next trace the zig-zag permutation, retain zero windings if they occur, and compare raw winding vectors with the proposed primitive web legs. Finally examine intersections of lifted paths; this is an additional geometric test, not an output of winding balance.

An isoradial realization gives angles from which candidate charges $0 < R_e < 1$ can be assigned, with $\sum_{e \ni v} R_e = 2$ at each superpotential vertex and $\sum_{e \in \partial f}(1 - R_e) = 2$ at each gauge face. These are the marginality and vanishing gauge-beta-function constraints in this description. Establishing the physical superconformal theory and selecting its exact charges requires the additional dynamics discussed in Chapter 6. Neither positive trial charges nor the finite homology calculation is that dynamical argument.

VII.3.9 What has been computed, and what is deferred

The chapter has built the dimer as the organizing object of the volume and read three kinds of data off it, on three worked tilings, including the orbifold example dP_0. For a cellular embedding the dictionary supplies a quiver with superpotential; for the consistent D3-brane phases studied here it is the four-dimensional $\mathcal{N}$=1 toric quiver: faces are the gauge nodes $\prod_i U(N_i)$, edges are the bifundamentals, and the black and white nodes are the superpotential terms, combined by the surface-normal sign rule, so that one graph fixes the quiver and the superpotential at once and each field sits in exactly two terms of opposite sign. The periodic quiver is the planar dual of the tiling, with the W-terms as its plaquettes. For a consistent toric phase the Kasteleyn determinant $\det K(z, w)$ has a Newton polygon that is the toric diagram, so the tiling computes its own geometry; the perfect matchings are the lattice points of that diagram with multiplicities, the variables the moduli construction will run on. The zig-zag paths carry the homology classes that must equal the (p, q)-web legs, a necessary comparison whose success leaves the lifted intersection tests and the physical identification as separate obligations. The three fixtures, the C^3 honeycomb, the conifold square, and the dP_0 tiling, exhibit the whole route, the last with the first genuine perfect-matching multiplicity.

A reader leaving the chapter should be able to do the following without looking ahead: read the gauge group, matter, and superpotential off a given brane tiling, including the $\pm$ signs from the colouring; check the Euler relation $V - E + F = 0$ and the structural fact that each field is in two opposite-sign terms; write the Kasteleyn matrix of a tiling, take its determinant, and read the toric diagram off the Newton polygon; recover $\mathrm{Tr}\, X[Y, Z]$ from the honeycomb and the Klebanov–Witten quartic from the conifold square; enumerate the perfect matchings of a small dimer and see the dP_0 interior point hit three times; and trace the zig-zag paths, compute their homology classes, and test them against the length-weighted primitive toric normals, while distinguishing this

necessary comparison from the lifted intersection and dynamical questions.

What has been deferred, and to where, is the rest of cluster A and the volume.

- Why the dimer is a *physical* object and not a convenient drawing is asserted here and derived in Chapter 4: the tiling is the Type IIB NS5–D5 brane configuration obtained by two T-dualities of the D3-branes at the toric singularity, with the D5-branes filling the faces and a single NS5-brane wrapping a holomorphic curve. The bipartite graph is a limiting combinatorial description associated with the fivebrane system; a generic angular projection of the whole curve is not literally a one-dimensional graph. The zig-zag-to-web-leg correspondence is a brane statement there. The belief that the dimer is "just a drawing, and Type IIA" is corrected in Chapter 4.

- The moduli-space construction is named here and run in Chapter 5: the forward construction from a tiling to its gauged linear sigma model, the perfect matchings as homogeneous coordinates, the charge matrices and the master space, and the mesonic and baryonic Hilbert series. This chapter ties perfect matchings to Newton-polygon points with multiplicities, but it does not count the moduli space; the belief that perfect matchings are in bijection with lattice points is corrected there, the dP_0 interior point being the counterexample met above.

- The dimer *moves* that relate different tilings of one singularity, the combinatorial form of Seiberg duality, are named here and performed in Chapter 7. One toric Calabi–Yau threefold can have several consistent tilings; this chapter notes the non-uniqueness and defers the move.

- The selection of R-charges by a-maximization, of which the isoradiality and positive trial-charge discussion supplies geometric constraints, is Chapter 6, with the matching to the Sasaki–Einstein volume minimization.

- The Kasteleyn theorem itself, that the signed determinant weights perfect matchings sector by sector, and the dimer-model and spectral-curve mathematics of Kenyon and collaborators, are self-contained mathematics, cited and used but not proved.

- The non-orbifold targets beyond the conifold and dP_0, the $Y^{p,q}$, the $L^{a,b,c}$, and the higher del Pezzo cones, are named as the singularities the machinery serves; they are worked in the atlas of Chapter 11.

The single lesson to carry forward is the one in the first keybox: the quiver and the superpotential of a toric gauge theory are one graph on a torus, and that graph also computes the geometry and exposes geometric consistency questions. The next chapter explains why the graph is physical; the chapter after that turns it into the moduli-space construction.

Status and sources

The evidence-code tags ([WS] physical-worldsheet derivation in the untwisted string, [MATH] cited mathematical theorem), the convention pointers, and the source anchors are collected here, keyed by

section. The chapter introduces the brane tiling / dimer as the combinatorial object that encodes a toric $\mathcal{N}$=1 quiver and its superpotential at once, after Hanany and Kennaway (2005) and Franco, Hanany, et al. (2006) for the dictionary and the Kasteleyn route, Hanany and Vegh (2007) for the zig-zag paths and consistency, and Franco and Vegh (2006) for the moduli-space direction (run in Chapter 5); the dimer-model and Kasteleyn-theorem mathematics is Kenyon (2003) and Kenyon et al. (2003). The conifold gauge theory recalled in §VII.3.7 is Klebanov and Witten (1998) (owned by Chapter 1); the honeycomb / $\mathbb{C}^3$ orbifold bridge is Douglas et al. (1997); the closest textbook treatment is Johnson (2003). It reuses the dimer/brane-tiling conventions (bipartite tiling conventions, the surface-normal sign rule, the Kasteleyn matrix, and the two worked examples), the toric-diagram conventions (toric lattice orientation, the Newton polygon of $\det K$ as the toric diagram), the superpotential store (superpotential normalization, the ± 1 cyclic words), the gauge/rank data (quiver gauge groups and ranks, $\prod_i U(N_i)$), the arrow convention (fundamental-at-tail arrows, the arrow-crossing rule), and the supercharge count. No new convention is introduced. The finite calculations use the displayed tiling data and surface-normal signs.

Evidence and assumptions. No physical claim in this chapter is assigned WS merely because a quiver calculation or computer algebra agrees. WS means a controlled physical-worldsheet derivation in the untwisted string. Such a derivation is not supplied here; the brane interpretation is source-backed and its primary classification remains unassigned pending the required payload. MATH applies only to the stated mathematical domains: cyclic words modulo trace rotation, finite Laurent determinant expansions, and the cited Kasteleyn theorem with a cellular embedding and its face signing.

The companion contains three families of submitted exercises: the quiver and cyclic superpotential of fixed embedded graphs, exact signed determinants and unsigned matching counts, and zig-zag permutations with raw windings. It checks face closure, finite edge-sharing counts and primitivity as specified in its instructions. It does not prove universal-cover consistency, SCFT existence, or the physical brane dictionary.

The conifold polynomial here is $1-z-w-zw$. The polynomial $1-z-w+zw=(z-1)(w-1)$ has the same Newton support but violates the square-face signing rule. The two equal epsilon contractions give the factor $1/2$ in (VII.3.25). Neither correction changes the quiver or Newton polygon. For the face rule and the sign in a fixed height sector, see Kenyon et al. (2003), §§3.1.1–3.1.3 and Proposition 3.1. The dictionary and determinant construction in Franco, Hanany, et al. (2006) are §§2–3 (arXiv:hep-th/0504110, not hep-th/0505211). The geometric reading of zig-zags uses Hanany and Vegh (2007), §§3–4. Its §4.2 presents the web-leg correspondence as a conjecture; citing it is not a proof for arbitrary bipartite graphs. The legacy homology matcher does not implement those intersection conditions.

§VII.3.1 One graph for the quiver and the superpotential. Framing: the quiver fixes matter; the dimer also fixes cyclic words and relative signs, the volume's analog of the McKay quiver; the three computations (dictionary, Kasteleyn, zig-zag) and the reader contract. The opening repeats the conditional dictionary, geometry and consistency statements.

§VII.3.2 The dimer dictionary. The faces/edges/nodes dictionary (VII.3.1), the gauge group (VII.3.2), the surface-normal sign rule $W=\sum_{\text{white}}\operatorname{Tr}(\text{cw})-\sum_{\text{black}}\operatorname{Tr}(\text{ccw})$ (VII.3.3), the structural fact that each field sits in two opposite-sign terms (VII.3.4), the arrow-crossing rule, and the Euler constraint $V-E+F=0$ on T^2 (VII.3.6). (physical primary classification unassigned) (bipartite tiling conventions, quiver gauge groups and ranks, superpotential normalization, fundamental-at-tail arrows, spinor and supersymmetry counting; Hanany and Kennaway (2005), Franco, Hanany, et al. (2006)). Carries (the quiver is read off the dimer, the combinatorial organizer, not the shape of the singular space directly; refined toward the tiling). $V-E+F=0$ on T^2 for all three fixtures, and each edge meets exactly one white and one black node so each field is in two W-terms; the honeycomb and conifold W-words equal the canonical entry-6 store with each edge in two opposite-sign terms.

§VII.3.3 The periodic quiver. The planar duality tiling↔periodic quiver (VII.3.7) (faces→nodes, edges→arrows, nodes→plaquettes=W-terms), drawn on T^2 / its universal cover $\mathbb{R}^2$ (Figure VII.3.1), and the note that the dimer moves act on this object. (physical primary classification unassigned) (bipartite tiling conventions, fundamental-at-tail arrows; Hanany and Kennaway (2005), Franco, Hanany, et al. (2006)). The dimer moves / toric duality are routed to Chapter 7.

§VII.3.4 The C^3 honeycomb. The honeycomb cell counts $(V,E,F)=(2,3,1)$ (VII.3.8), the single $U(N)$ (VII.3.9), the three adjoints X,Y,Z, and the sign rule yielding $W=\operatorname{Tr}(XYZ)-\operatorname{Tr}(XZY)=\operatorname{Tr}\,X[Y,Z]$

(VII.3.10), the parent cubic (Figure VII.3.2). (physical identification imported; algebra computed) (bipartite tiling conventions, the honeycomb fixture; superpotential normalization, the parent cubic; Douglas et al. (1997), Hanany and Kennaway (2005)). $\mathrm{Tr}(XYZ) - \mathrm{Tr}(XZY) = \mathrm{Tr}\ X[Y,Z]$ exactly, the trivalent node valencies, and the canonical W-words.

§VII.3.5 The Kasteleyn matrix and the toric diagram. The Kasteleyn matrix $K(z,w)$ with winding fugacities (VII.3.13), the characteristic polynomial $P = \det K$ (VII.3.15), its Newton polygon (VII.3.16), identified in the consistent mesonic phase with the toric diagram up to $GL(2,\mathbb{Z})$ + translation (VII.3.17), the honeycomb $K_{\mathbb{C}^3} = 1 - z - w$ with support the unit triangle (VII.3.18)–(VII.3.19), and the area count $G = 2$ Area (VII.3.20). (mathematical and physical implications separated) (bipartite tiling conventions, toric lattice orientation; Hanany and Kennaway (2005), Franco and Vegh (2006); Kenyon (2003)). The Kasteleyn-signing existence theorem is dimer and Kasteleyn mathematics (self-contained mathematics; cited, not proved). $\det K$ for the honeycomb $= 1-z-w$ with support $\{(0,0),(1,0),(0,1)\}$, the legacy conifold $= 1 - z - w + zw$ with the unit-square support (the corrected signing is tested separately), both Newton polygons equal to the entry-7 toric diagram up to $GL(2,\mathbb{Z})$, and $G = 2$ Area on all three fixtures.

§VII.3.6 Perfect matchings and the lattice points. The perfect matching as an edge subset covering each node once, the determinant expansion (VII.3.22), the map perfect matchings→lattice points (VII.3.23) and the load-bearing fact that it is not a bijection (multiplicities), with the honeycomb and conifold as the bijective cases. (mathematical and physical implications separated) (bipartite tiling conventions; Franco and Vegh (2006), Kenyon (2003)). The moduli machine (GLSM coordinates, charge matrices, master space, Hilbert series) and the multiplicity *count* are routed to Chapter 5; the belief that perfect matchings biject with lattice points is corrected there (the stated distinction, owned by Chapter 5). the honeycomb / conifold are multiplicity-free, and the perfect-matching enumeration agrees with the $|\det K|$ coefficients on all three fixtures.

§VII.3.7 The conifold square and dP_0. The conifold square recalled: $(V,E,F) = (2,4,2)$ (VII.3.24), gauge group $U(N)_1 \times U(N)_2$, the Klebanov–Witten quartic from the sign rule (VII.3.25), and $\det K = 1 - z - w - zw$ with the unit-square diagram (VII.3.26); the $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ tiling as $K_{3,3}$ with $(V,E,F) = (6,9,3)$ (VII.3.27), the cubic $\mathbb{Z}_3$-daughter superpotential, the local-$\mathbb{P}^2$ Newton polygon with one interior point (VII.3.29), and the interior-point perfect-matching multiplicity 3 (VII.3.30) (Figure VII.3.3). (physical identification imported; algebra computed) (bipartite tiling conventions, the two worked examples; superpotential normalization, toric lattice orientation; Klebanov and Witten (1998), Hanany and Kennaway (2005), Franco, Hanany, et al. (2006); the $\mathbb{C}^3/\mathbb{Z}_3$ bridge Douglas et al. (1997)). The conifold gauge theory is owned by Chapter 1, recalled here as the fixture; the V→VI bridge $\mathbb{C}^3/\mathbb{Z}_3 = dP_0$ (McKay vs. dimer) is developed in Chapters 2 and 5. the conifold quartic from the sign rule, the legacy determinants and their Newton polygons (the candidate conifold signing is checked separately), the four-valent conifold versus the trivalent dP_0, and the dP_0 interior multiplicity 3 matching the $|\det K|$ coefficient and the perfect-matching enumeration.

§VII.3.8 Zig-zag paths and consistency. The zig-zag path (alternating maximal turns), its homology class $(p,q) \in H_1(T^2)$ (VII.3.33), the central correspondence zig-zag classes=(p,q)-web external legs= primitive toric outward normals with lattice-length multiplicities in the geometrically consistent class (VII.3.34), with their vanishing sum (VII.3.35), the three fixtures' classes (VII.3.36)–(VII.3.38), the consistency conditions (the strong geometric criterion for the regular embeddings, distinguished from weaker same-orientation criteria), and the sequence of necessary combinatorial tests and additional geometric and dynamical questions. (homology calculation only; physical classification unassigned) (bipartite tiling conventions, toric lattice orientation; Hanany and Vegh (2007)). Files (not every bipartite graph on T^2 represents a physical toric phase; the geometric and dynamical requirements are separate; refines the stated distinction). The R-charge selection by a-maximization is routed to Chapter 6. the zig-zag homology classes of all three fixtures, traced from the rotation system and windings alone, $GL(2,\mathbb{Z})$-match the toric outward normals and sum to zero, while an inconsistent (degenerate digon) tiling and a non-primitive (scaled-winding) tiling both fail the legacy homology test.

§VII.3.9 What is and is not given. Scope and limitations. the chapter builds the dimer dictionary, the periodic quiver, the Kasteleyn route to the toric diagram, the perfect-matching/lattice-point map with multiplicities, and the zig-zag consistency filter, worked on the $\mathbb{C}^3$ honeycomb, the conifold square, and dP_0; it routes the physical NS5–D5 origin of the tiling to Chapter 4 (the stated distinction, owned

there), the moduli-space machine and the multiplicity count to Chapter 5 (the stated distinction, owned there), the dimer moves / toric duality to Chapter 7, the R-charge selection by a-maximization to Chapter 6, the Kasteleyn theorem as self-contained mathematics to dimer and Kasteleyn mathematics, and the non-orbifold atlas targets to Chapter 11.

Related constructions. The bipartite graph on T^2 gives gauge faces, bifundamental edges and signed superpotential terms. The Newton polygon of $\det K$ gives the toric diagram for the consistent mesonic phase. Its comparison with zig-zag homology leaves the lifted-intersection condition separate. The $\mathbb{C}^3$ honeycomb gives $W = \operatorname{Tr} X[Y, Z]$ and the unit triangle. The conifold square gives the Klebanov–Witten quartic and the unit square. The dP_0 tiling illustrates multiplicity and gives the same target as the directed-McKay construction in Volume VI.

References

Douglas, M. R., B. R. Greene, and D. R. Morrison (1997). "Orbifold resolution by D-branes". In: *Nucl. Phys. B* 506, pp. 84–106. DOI: 10.1016/S0550-3213(97)00517-8. arXiv: hep-th/9704151.

Franco, S., A. Hanany, et al. (2006). "Brane dimers and quiver gauge theories". In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Franco, S. and D. Vegh (2006). "Moduli spaces of gauge theories from dimer models: Proof of the correspondence". In: arXiv: hep-th/0601063.

Hanany, A. and K. D. Kennaway (2005). "Dimer models and toric diagrams". In: arXiv: hep-th/0503149.

Hanany, A. and D. Vegh (2007). "Quivers, tilings, branes and rhombi". In: *JHEP* 10, p. 029. DOI: 10.1088/1126-6708/2007/10/029. arXiv: hep-th/0511063.

Johnson, C. V. (2003). *D-Branes*. Cambridge Monographs on Mathematical Physics. Cambridge University Press. DOI: 10.1017/CBO9780511606540.

Kenyon, R. (2003). "An introduction to the dimer model". In: arXiv: math/0310326 [math.CO].

Kenyon, R., A. Okounkov, and S. Sheffield (2003). "Dimers and amoebae". In: arXiv: math-ph/0311005.

Klebanov, I. R. and E. Witten (1998). "Superconformal field theory on three-branes at a Calabi-Yau singularity". In: *Nucl. Phys. B* 536, pp. 199–218. DOI: 10.1016/S0550-3213(98)00654-3. arXiv: hep-th/9807080.

Chapter VII.4

The NS5–D5 picture: brane tilings from T-duality

VII.4.1 Why the dimer is not just a drawing

The last chapter introduced the brane tiling as a bipartite graph on a two-torus and showed it encodes a quiver and a superpotential at once: faces are gauge nodes, edges are bifundamentals, the two colours of node are the two signs of the superpotential terms. It was a remarkably efficient piece of calculation. But efficiency is not the same as necessity, and a reasonable reader leaves Chapter 3 with a question the combinatorics cannot answer on its own. Why this graph? A bipartite graph on T^2 is a drawing; one could draw a different one. What forces the dimer of a given toric singularity to be the dimer it is, rather than a convenient mnemonic that happens to reproduce the right quiver?

This chapter answers that question through the Type IIB D5–NS5 realization of a chosen toric phase. Two transverse T-dualities take the D3 stack to D5-branes, while the degeneration of the toric fibre is represented by an NS5-brane on a holomorphic curve. The branes explain why faces carry gauge fields and why intersections carry charged matter. To recover a bipartite graph one also chooses compatible brane contours and a projection that preserves their intersections. The curve's full angular image is a region, not itself a graph. Thus the physical construction explains the tiling without asserting that an arbitrary polynomial with the right Newton polygon uniquely determines it.

Take-home message of this chapter

Take a D3-brane probe of a toric Calabi–Yau threefold. The two-circle duality relates it to a Type IIB D5–NS5 system, with the NS5 wrapping $\Sigma = \{P(z, w) = 0\}$. In a compatible brane-tiling phase, D5 regions carry gauge factors, open strings at adjacent regions carry chiral matter, and oriented junctions encode signed superpotential monomials. The Newton polygon fixes the asymptotic winding directions. The coefficients, the brane contours and their faithful projection are further data.

A correctness warning belongs at the very start, because it is the single most

common error about this construction. The configuration is *Type IIB*. The D5-brane is a Type IIB object, and two T-dualities of a Type IIB setup return Type IIB, because each T-duality exchanges IIA and IIB and an even number of them is the identity on the string theory. A reader who has met NS5–brane constructions in the Type IIA Hanany–Witten language must resist importing it here: there are no Type IIA D6–branes wrapping a curve, no Type IIA NS5 spanning a fixed set of directions; this is a Type IIB D5–NS5 system. Stating this once, sharply, guards the rest of the chapter.

Type IIB, not Type IIA (do not confuse)

The brane tiling lives in *Type IIB*. The chain of reasoning: the D3-brane probe is a Type IIB object; each T-duality flips the type IIA↔IIB; two T-dualities flip it twice, back to IIB; and the resulting D5-brane is a Type IIB brane, as it must be. The frequent claim that the dimer is "a Type IIA picture" is wrong, and so is the looser claim that it is "just a drawing." Its graph describes a compatible Type IIB D5–NS5 phase. The required contour and projection data matter as well as the curve.

The chapter develops one main line and names one alternative. The main line is the T-duality route: follow the D-brane dimensions through two T-dualities, identify the end configuration, and read the tiling off the projected NS5 curve. The alternative, named but not re-derived, is the mirror-symmetry route, in which suitable contours on the mirror curve recover the tiling; it recovers the chosen tiling with suitable brane contours and is cited as an alternative construction.

One boundary is worth fixing before starting. T-duality is used here as a *map between brane configurations*: the rule that a Dp transverse to a circle becomes a D$(p+1)$ and that the toric degeneration becomes an NS5. The worldsheet derivation of that rule, the Buscher transformation of the background fields, belongs to the foundations and is cited, not reproved. The amoeba and coamoeba of the curve Σ, and the holomorphic-curve technology that makes the projection precise, are named as the relevant mathematics and routed likewise. What this chapter computes is the dimension and type sequence and the curve geometry in explicit examples. The physical duality and the faithful projection of the chosen brane contours remain imported premises.

The plan is: the setup inherited from Chapters 1–2 and the toric T^2 one T-dualizes along (§VII.4.2); the two T-dualities and the brane dictionary, with the dimension and type calculation (§VII.4.3); the NS5 holomorphic curve, its Newton polygon, and the careful relation to the Kasteleyn determinant (§VII.4.4); the tiling as the skeleton of the projected curve, with faces, edges, and nodes read off the coamoeba (§VII.4.5); the three-way identification of zig-zags, NS5 legs, and (p,q)-web legs (§VII.4.6); the conifold and $\mathbb{C}^3$ worked anchors, with dP_0 named (§VII.4.7); the mirror route as a one-paragraph cross-reference (§VII.4.8); and the closing map of what was constructed and what is deferred (§VII.4.9).

VII.4.2 The setup: a D3 probe and the toric two-torus

The starting point is the configuration of Chapters 1 and 2, taken as given. A toric Calabi–Yau threefold cone X is encoded by its toric diagram, a convex lattice polygon $\Delta \subset \mathbb{Z}^2$, and a stack of N D3-branes sits at the tip of the cone, filling the four flat directions $\mathbb{R}^{1,3}$ and pointlike in the six transverse directions of X. The worldvolume theory is the $\mathcal{N}$=1 quiver gauge theory built in those chapters: gauge group the product of unitary factors over the faces,

$$\prod_a U(N_a), \qquad N_a = N \text{ for the regular brane,} \tag{VII.4.1}$$

bifundamentals on the arrows, and a toric superpotential whose terms are signed closed loops. The construction outputs the unitary factors; the interacting sector and the baryonic symmetries are read off as in Chapter 1, and we never write the gauge group as a bare SU.

What this chapter adds is a change of duality frame, and for that the relevant feature of X is its torus fibration. A toric threefold cone carries a $T^2 \times \mathbb{R}$ fibration over a three-dimensional base, the T^2 being the orbits of the two toric $U(1)$ isometries that act on the fibres. Away from the degeneration locus the fibre is a smooth two-torus; over special loci one or the other cycle of the T^2 pinches, and the way the cycles pinch over the boundary of the base is exactly the combinatorial data the toric diagram records. It is along the two independent cycles of this T^2 that we will T-dualize. Schematically,

$$X \;\simeq\; \big(T^2 \times \mathbb{R}\big) \text{ fibred over a 3d base,} \qquad T^2 = \{\text{orbits of the two toric } U(1)\text{'s}\}, \tag{VII.4.2}$$

with the fibre degenerating over the locus dual to the polygon. For a generic Calabi–Yau threefold probe the common supersymmetry is four real supercharges, or $\mathcal{N}$=1 in four dimensions. A Calabi–Yau condition does not exclude enhancement: the flat $\mathbb{C}^3$ D3 theory has sixteen supercharges and $\mathcal{N}$=4, written here in $\mathcal{N}$=1 multiplets. The two-circle construction retains the common four-supercharge description. Exact string T-duality preserves supersymmetry, but its realization as local Killing spinors in a dual background requires invariance along the dualized isometries. Neither the brane array nor the dimension count computes this spinor condition.

For flat $\mathbb{C}^3$ the familiar moment-map base is the positive octant. That is the base of a T^3 fibration: a coordinate angle shrinks on each facet. It should not be identified with the base of the $T^2 \times \mathbb{R}$ description in (VII.4.2). The latter uses the two-dimensional subtorus that preserves the holomorphic three-form. For example, on coordinates (x, y, t) of $\mathbb{C}^3$ it acts by $(x, y, t) \mapsto (e^{i\alpha}x, e^{i\beta}y, e^{-i(\alpha+\beta)}t)$. The product xyt is invariant. This distinguishes the torus used in the brane construction from the entire toric T^3, while its degeneration data still lead to the triangle and its three external directions.

Why two circles and not three. A toric Calabi–Yau threefold carries three angular $U(1)$ isometries, a T^3 moment-map fibration, and a reader meeting the construction may expect to T-dualize all three. The two-circle construction uses the kernel of the torus character by which the holomorphic three-form transforms. It therefore preserves

that form and the common $\mathcal{N}$=1 supercharges. The remaining angular action includes the R symmetry; on an irregular cone the Reeb flow need not even close to a circle. There is no universal splitting into this T^2 and a compact Reeb circle. Dualizing all three directions instead belongs to the mirror description, which has different branes. The two independent dual angles in the present description enter $(z, w) \in (\mathbb{C}^*)^2$.

What is reused and what is new. The toric diagram, the (p, q)-web dual to it, the quiver, and the superpotential are all inputs here, built in Chapters 1, 2, and 3. The geometry of the fibration (VII.4.2), in the precise sense of the metric and the special-Lagrangian structure, is part of the toric geometry routed to Part I; this chapter uses only that the T^2 exists with two cycles and degenerates over the boundary of the base. Everything new in this chapter is the consequence of T-dualizing along those two cycles, to which we now turn.

VII.4.3 The two T-dualities and the brane dictionary

T-duality along a circle acts on a D-brane by a simple rule of dimension. If the brane wraps the circle, T-duality removes that direction from the worldvolume and the brane loses a dimension; if the brane is transverse to the circle, sitting at a point on it, T-duality turns the dual winding into a worldvolume direction and the brane gains one. In symbols, for a Dp-brane and a circle it does not wrap,

$$\mathrm{D}p \xrightarrow{T \text{ on a transverse circle}} \mathrm{D}(p+1), \tag{VII.4.3}$$

and the opposite, $\mathrm{D}p \to \mathrm{D}(p-1)$, if the brane wraps the circle. Alongside the dimension, T-duality flips the type of the string theory: a single T-duality maps Type IIA to Type IIB and back. We keep both pieces of calculation, because both carry physics, and the type calculation is what guards the Type IIB statement of the introduction.

The D3 becomes a D5. The N D3-branes sit at a point of the toric T^2, transverse to both of its cycles. T-dualize along the first cycle: the branes are transverse to it, so by (VII.4.3) each D3 grows to a D4, now wrapping the first dual cycle. T-dualize along the second cycle: the D4 is still transverse to it, so it grows again, to a D5 wrapping both dual cycles, that is, filling the dual T^2. The chain is

$$\mathrm{D3} \xrightarrow{T_1} \mathrm{D4} \xrightarrow{T_2} \mathrm{D5}, \qquad 3 \;\to\; 3+1 \;\to\; 3+2=5, \tag{VII.4.4}$$

each step adding one to the dimension because the brane is transverse to the circle being dualized. At the end the N branes fill $\mathbb{R}^{1,3}$ and the two cycles of the dual T^2: they are N D5-branes wrapping $\mathbb{R}^{1,3} \times T^2$, spread out over the torus rather than localized at a point.

The type stays IIB. Each of the two T-dualities flips the type, and two flips return the starting type. Writing the type as a parity that is unchanged by an even number of flips,

$$\mathrm{IIB} \xrightarrow{T_1} \mathrm{IIA} \xrightarrow{T_2} \mathrm{IIB}, \qquad (\#\, T\text{-dualities}) \bmod 2 = 2 \bmod 2 = 0, \tag{VII.4.5}$$

so the end configuration is Type IIB, consistent with the D5-brane being a Type IIB object. This is the precise content of the introduction's warning: a Type IIA labeling would require an odd number of T-dualities, or a D5-brane in the wrong theory, and either is the dimension and type error under discussion. An odd chain, D3 $\to$ D4 alone, would land in Type IIA with a D4; a three-step chain would give a D6 and again Type IIA. The toric T^2 has exactly two cycles, so exactly two T-dualities are performed, and the type is preserved.

The degeneration becomes an NS5. The D-branes are not the whole story. The toric fibre degenerates over the boundary of the base, and that geometry T-dualizes into a brane as well. A locus where a cycle of the T^2 pinches is, in the T-dual frame, not a piece of geometry but an NS5-brane, the magnetic dual of the fundamental string and the natural partner of the D5 in Type IIB. A single NS5-brane appears. It wraps $\mathbb{R}^{1,3}$ and a two-real-dimensional surface Σ inside the $(\mathbb{C}^*)^2$ of the dual torus angles and their radial partners, and that surface is a holomorphic curve. This $(\mathbb{C}^*)^2$ is not wholly transverse to the D5's: the D5's fill the angular T^2, while Σ also runs in the radial directions, and the universally shared flat factor is $\mathbb{R}^{1,3}$; additional internal boundary intersections depend on the embedding. The next section makes Σ precise; here we only record that the end configuration of the two T-dualities is

$$\boxed{N \text{ D5 on } \mathbb{R}^{1,3} \times T^2 \quad + \quad \text{one NS5 on } \mathbb{R}^{1,3} \times \Sigma} \qquad \text{(Type IIB)}, \tag{VII.4.6}$$

with the D5 regions and NS5 embedding understood in the chosen brane-tiling limit. A smooth irreducible curve gives one connected NS5 component; reducible limits require a component analysis. Both branes share the four flat directions $\mathbb{R}^{1,3}$ where the gauge theory lives, and the imported duality identifies the low-energy theory with the same quiver, written in $\mathcal{N}$=1 language, (VII.4.1) as before, now realized on a D5–NS5 system. T-duality is a duality: it has changed the description, not the physics.

The brane array. Since the whole construction is a statement about dimensions, it pays to tabulate which brane occupies which directions. Split the ten directions into the gauge-theory spacetime $\mathbb{R}^{1,3}$, the two angles θ_1, θ_2 of the toric T^2 that are dualized, the two logarithmic moduli r_1, r_2 that complete those angles to the $(\mathbb{C}^*)^2$ of the next section, and a remaining transverse $\mathbb{R}^2$. The two five-branes are then six-dimensional in the two ways a six-dimensional worldvolume can sit here,

$$\dim(\text{D5}) = \underbrace{4}_{\mathbb{R}^{1,3}} + \underbrace{2}_{T^2} = 6, \qquad \dim(\text{NS5}) = \underbrace{4}_{\mathbb{R}^{1,3}} + \underbrace{2}_{\Sigma} = 6, \tag{VII.4.7}$$

The D5 fills $\mathbb{R}^{1,3}$ and the whole T^2. The NS5 fills $\mathbb{R}^{1,3}$ and the two-real-dimensional curve Σ inside the $(\mathbb{C}^*)^2$. Their common flat factor is $\mathbb{R}^{1,3}$; the internal boundaries of D5 regions can lie on Σ. Table VII.4.1 records the full occupation.

A note on what the dictionary is and is not. The rule (VII.4.3), the type flip, and the statement that a pinching cycle T-dualizes to an NS5 are quoted from the foundations; they follow from the Buscher transformation of the background and the action of T-duality on Ramond–Ramond and Neveu–Schwarz charges, which are not reproved here. What this chapter does with the dictionary is exactly the dimension and

	$\mathbb{R}^{1,3}$	θ_1	θ_2	r_1	r_2	transverse $\mathbb{R}^2$
N D3 (before)	=	·	·	·	·	·
N D5 (after)	=	=	=	·	·	·
one NS5 (after)	=	wraps $\Sigma \subset (\mathbb{C}^*)^2$				·

Table VII.4.1: The brane array before and after the two T-dualities. A filled direction is marked = and a point is marked ·. The N D3-branes (a point on the toric T^2) become N D5-branes filling $\mathbb{R}^{1,3}$ and the dual $T^2 = (\theta_1, \theta_2)$; the toric degeneration becomes one NS5 wrapping $\mathbb{R}^{1,3}$ and the holomorphic curve Σ inside the $(\mathbb{C}^*)^2 = (\theta_1, \theta_2, r_1, r_2)$. Both five-branes are six-dimensional (VII.4.7) and share $\mathbb{R}^{1,3}$, where the $\mathcal{N}$=1 quiver lives.

type calculation of (VII.4.4)–(VII.4.6): count the dimensions, track the type, identify the branes. The finite dimension and parity calculations test consequences of the dictionary; they do not establish the physical duality.

Areas and gauge couplings. The graph determines adjacency, not numerical gauge couplings. In a weakly curved D5 region D_a, with constant dilaton $e^{\phi} = g_s$, vanishing background two-form and flux, the quadratic D5 action gives

$$\frac{1}{g_{4,a}^2} = \frac{\mathrm{Area}(D_a)}{g_6^2}, \qquad g_6^2 = (2\pi)^3 g_s \alpha'. \tag{VII.4.8}$$

Indeed, the D5 tension without its dilaton factor is $T_5 = (2\pi)^{-5}(\alpha')^{-3}$. The F^2 term of the DBI expansion has coefficient $T_5(2\pi\alpha')^2/(4g_s)$; integrating over D_a gives (VII.4.8). Here the action convention is $-\frac{1}{4g_6^2}\int \mathrm{Tr}\, F^2$. The area and α' both have dimensions of length squared, so the resulting four-dimensional coupling is dimensionless. With nonconstant fields the same expansion involves a metric and dilaton weighted integral, with further changes from flux and boundary effects. A schematic torus drawing supplies none of these continuous data, and (VII.4.8) is not an evaluation of a coupling in the singular NS5 core.

VII.4.4 The NS5 holomorphic curve and its Newton polygon

The single NS5-brane wraps a curve Σ in the $(\mathbb{C}^*)^2$ formed by the two dual torus directions. This $(\mathbb{C}^*)^2$ is not transverse to the D5's: it contains the torus angles the D5's fill, and their common flat factor is $\mathbb{R}^{1,3}$. After the two T-dualities, the natural coordinates on it are a pair of $\mathbb{C}^*$ variables, the complexified versions of the two dual torus circles,

$$(z, w) \in (\mathbb{C}^*)^2, \qquad z = e^{r_1 + i\theta_1},\ w = e^{r_2 + i\theta_2}, \tag{VII.4.9}$$

the dimensionless angles θ_1, θ_2 having period 2π and $r_1, r_2 \in \mathbb{R}$ being dimensionless logarithmic moduli, not nonnegative physical radii. Physical lengths require an additional scale and metric. The curve the NS5 wraps is the zero locus of a single Laurent

polynomial in these variables,

$$\Sigma = \{\, P(z,w) = 0 \,\} \subset (\mathbb{C}^*)^2, \qquad P(z,w) = \sum_{(m,n)\in\Delta\cap\mathbb{Z}^2} c_{mn}\, z^m w^n, \tag{VII.4.10}$$

a sum of monomials whose exponent vectors (m,n) are lattice points of the toric diagram Δ. The polygon fixes the allowed exponents. Individual coefficients can vanish; we retain nonzero coefficients at every vertex so that the Newton polygon remains Δ.

The Newton polygon is the toric diagram. The set of exponent vectors of a Laurent polynomial is its support, and the convex hull of the support is its Newton polygon. For the NS5-curve polynomial (VII.4.10), nonzero vertex coefficients ensure

$$\mathrm{Newton}(P) = \mathrm{conv}\big(\mathrm{supp}\, P\big) = \Delta, \tag{VII.4.11}$$

the toric diagram itself. This is the central geometric fact of the section: the polynomial whose zero locus the NS5 wraps has the toric diagram as its Newton polygon. The asymptotic behaviour of Σ at large $|z|, |w|$ is governed by the boundary edges of this polygon, and that is the origin of the curve's external legs, taken up in §VII.4.6.

How an edge gives a leg. The legs are not an extra postulate; one balance computation forces them from the boundary edges of Δ. Write the curve in logarithmic coordinates $u = (\log|z|, \log|w|)$. A monomial $c_{mn}z^m w^n$, with exponent vector $\mathbf{m} = (m,n)$, has log-modulus $\mathbf{m}\cdot u + \log|c_{mn}|$. In a generic region of the u-plane one monomial dominates and $P \neq 0$; the curve can run to infinity only along a locus where two monomials $\mathbf{m}, \mathbf{m}'$ balance in modulus,

$$\mathbf{m}\cdot u + \log|c_{\mathbf{m}}| \;\simeq\; \mathbf{m}'\cdot u + \log|c_{\mathbf{m}'}| \quad\Longrightarrow\quad (\mathbf{m}-\mathbf{m}')\cdot u = \text{const}, \tag{VII.4.12}$$

an affine line normal to the difference vector $\mathbf{m}-\mathbf{m}'$. For a primitive edge, its two endpoint monomials control this limit. For an edge of lattice length $L > 1$, all its edge monomials can remain of the same order. After removing a monomial factor, the edge restriction is a polynomial of degree L in a primitive tangential monomial. If its L roots are nonzero and distinct, it gives L punctures with the same primitive outward normal. This derives the multiplicity without assuming that adjacent monomials dominate in pairs. Coincident roots change the boundary geometry and require a separate limit.

What the polygon counts. Assume that P is irreducible and nondegenerate with respect to Δ: for every face τ, including Δ itself, the equations $P_\tau = z\partial_z P_\tau = w\partial_w P_\tau = 0$ have no common solution in $(\mathbb{C}^*)^2$. Then its smooth toric compactification has the genus below, and its transverse boundary intersections give the punctures. Under these hypotheses two integers are read directly off Δ. The number of asymptotic legs is the number of primitive segments along the boundary $\partial\Delta$, and the genus of the smooth curve in its class is the number of interior lattice points,

$$\begin{aligned} \#\{\text{legs of } \Sigma\} &= \#\{\text{primitive boundary segments of } \Delta\},\\ g(\overline{\Sigma}) &= \#\{\text{interior lattice points of } \Delta\}, \end{aligned} \tag{VII.4.13}$$

the second a standard fact about a curve in the toric surface defined by Δ. The conifold square and the $\mathbb{C}^3$ triangle have no interior point, so their NS5 curves are rational, of

genus zero; the triangle of dP_0 has a single interior point, so its NS5 curve has genus one, the first sign of the compact divisor that distinguishes dP_0 from the conifold. Both counts return in the worked examples of §VII.4.7.

The careful relation to the Kasteleyn determinant. Chapter 3 produced, from the dimer, a characteristic polynomial $\det K(z, w)$ whose Newton polygon is also the toric diagram. It is tempting to declare $P = \det K$ outright, but that is more than the construction supports at this stage, and the precise statement is worth making. For a consistent tiling with suitable nonvanishing weights, the characteristic polynomial has the prescribed Newton polygon. Its coefficients depend on edge weights and their possible cancellations; the graph alone fixes neither the physical coefficients nor every occupied lattice point. The relation needed here is

$$\text{Newton}(P) \cong \text{Newton}(\det K) \cong \Delta \quad (\text{up to } GL(2,\mathbb{Z}) \text{ and translation}). \tag{VII.4.14}$$

Equality of convex hulls is weaker than equality of supports. For example, $z + w + (zw)^{-1} + c$ retains the local-$\mathbb{P}^2$ polygon when c vanishes, although its support loses the interior lattice point. Its smooth members still have genus one. A Kasteleyn polynomial is a weighted representative, not a unique canonical physical curve. In particular, neither the Newton polygon nor a choice of unweighted determinant proves that its coamoeba recovers a chosen tiling.

VII.4.5 The tiling as the skeleton of the projected curve

The curve Σ lives in $(\mathbb{C}^*)^2$, which is a T^2 (the angles θ_1, θ_2) times an $\mathbb{R}^2$ (the logarithmic moduli r_1, r_2). There are two natural projections, one onto each factor, and they carry the two pieces of toric data the volume has been using.

The logarithmic projection is the *amoeba*; the angular projection is the *coamoeba*. These are images of a real surface and generally have interiors. The amoeba's unbounded directions are fixed by the polygon. Its spine or an appropriate tropical degeneration gives a planar graph with these directions; finite edge lengths and subdivisions also depend on coefficients. A coamoeba need not retract onto the dimer of an arbitrarily chosen phase.

$$\Sigma \subset T^2 \times \mathbb{R}^2 : \qquad \begin{cases} \text{Log}(z, w) = (\log|z|, \log|w|), & \text{amoeba}, \\ \text{Arg}(z, w) = (\arg z, \arg w), & \text{coamoeba}. \end{cases} \tag{VII.4.15}$$

To obtain the tiling, use compatible contours on Σ whose intersections realize the brane configuration and whose angular projection preserves that intersection graph. Deforming these contours into the bipartite presentation then gives the chosen dimer. This additional condition is substantive: projected crossings alone do not imply intersections upstairs. The drawings in Figure VII.4.1 display this compatible limit, not a universal retraction theorem.

Faces, edges, nodes from the projection. In the compatible brane-tiling limit, the projected contours delimit regions, and the D5-branes filling the torus are divided

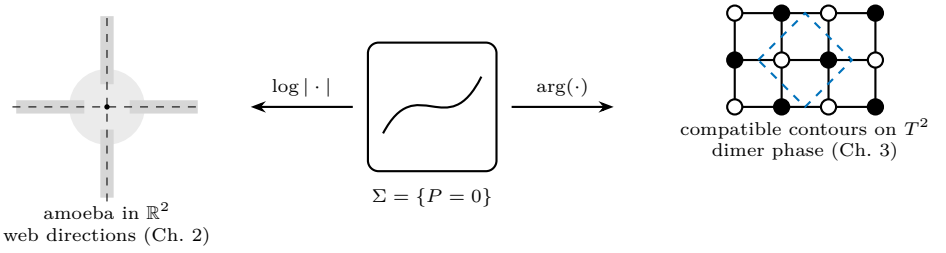


Figure VII.4.1: The logarithmic and angular projections of Σ, drawn schematically in a compatible brane-tiling phase. The left panel records the external web directions, not a computed amoeba shape or metric. The right panel is the bipartite graph obtained from suitable projected brane contours; it is not the full coamoeba. The dashed cell has square-grid periods $(1, 1)$ and $(1, -1)$. Coefficients and contour choices are needed in addition to the Newton polygon.

by the curve into the same regions. This reconciles the two ways of speaking: the D5 worldvolume covers the whole dual T^2, but the NS5 curve partitions it into face regions, and a regular stack puts N D5 sheets on each region, so each face carries one $U(N)$ factor (not a single D5 producing many nodes). Each region is a piece of D5 worldvolume, hence a gauge factor; the dictionary of Chapter 3 reads

$$\text{faces} \;\leftrightarrow\; \text{D5 regions} \;\leftrightarrow\; \text{gauge nodes } U(N_a), \tag{VII.4.16}$$

exactly the faces of the bipartite graph. At an interface between adjacent D5 regions, the brane construction supplies open-string chiral matter. Its two Chan–Paton endpoints determine the representation. The thin-tube picture represents this interface in the tiling limit, so

$$\text{edges} \;\leftrightarrow\; \text{curve pinches} \;\leftrightarrow\; \text{bifundamentals}, \tag{VII.4.17}$$

the edges of the graph, oriented by the arrow convention of the quiver. Concretely, a string across the pinch between D5 region a and D5 region b has one endpoint on each stack, so it transforms as a fundamental of $U(N_a)$ at one end and an antifundamental of $U(N_b)$ at the other, a bifundamental $(\mathbf{N}_a, \overline{\mathbf{N}}_b)$; the orientation of the string, fixed by the handedness of the junction it runs from, selects this representation over its conjugate $(\overline{\mathbf{N}}_a, \mathbf{N}_b)$, and that orientation is the arrow $a \to b$ of the quiver. The bipartite graph has vertices where these interfaces meet; a smooth holomorphic curve itself has no such singular vertices. In the graph, these come in two orientations, according to the handedness with which the curve sheets join, and the two orientations are precisely the black and white nodes of the bipartite graph,

$$\text{nodes} \;\leftrightarrow\; \text{NS5 junction orientations} \;\leftrightarrow\; W \text{ terms}, \quad \text{white } = +, \text{ black } = -, \tag{VII.4.18}$$

the two signs of the superpotential, matching the sign rule frozen in Chapter 3. In the projected bipartite skeleton the valency of a node, the number of edges meeting it, equals the length of the superpotential monomial it carries: a trivalent node is a cubic term, a four-valent node a quartic term, and so on. This is a property of the skeleton,

not a claim that the unresolved five-brane junction is locally trivalent (the conifold's nodes are genuinely four-valent in the dimer).

How to read the projection

In a compatible phase, read D5 regions as gauge nodes, interfaces as chiral matter, and the two orientations of graph vertices as the two signs of the superpotential. Node valency counts fields in the corresponding monomial. The graph describes selected brane contours in the angular projection. The entire coamoeba is a region, and an arbitrary projection need not preserve the brane intersections.

The graph is forced by a chosen configuration. This is the payoff promised in the introduction, stated with the right quantifier. The toric singularity fixes the Newton polygon of Σ (VII.4.11), and hence the asymptotic legs. But it does *not* fix a unique curve or a unique tiling. The coefficients of P are complex-structure moduli, and a singularity with several toric phases (Chapter 7) is realized by different compatible brane contour configurations with the same polygon and legs. A change of curve coefficients alone is not a derivation of Seiberg duality. The graph describes a *particular* D5–NS5 configuration and a faithful choice of projected contours. Once compatible contours and their projection are chosen, their D5 face data determine the graph. The number of faces, the pattern of pinches, and the orientations of the junctions are read off it, not chosen separately. A useful invariant that any such projection must respect is the Euler relation on the torus: for a cellular graph on T^2 with disk faces, V nodes, E edges, and F faces,

$$V - E + F = 0, \tag{VII.4.19}$$

the vanishing Euler characteristic of T^2. Any candidate projection that violates (VII.4.19), or that gives the wrong face count, or a junction orientation inconsistent with the superpotential sign rule, is not the projection of the actual curve and is rejected. These necessary conditions are not sufficient to select a physical tiling. The compatible brane contours supply the remaining construction data.

VII.4.6 Zig-zags, NS5 legs, and the (p, q)-web

The chapter's structural bridge ties three objects from three chapters into one. Chapter 2 gave the (p, q)-web, the external legs dual to the boundary edges of the toric diagram. Chapter 3 gave the zig-zag paths, the closed paths on the bipartite graph that turn maximally and alternately, each carrying a homology class in $H_1(T^2; \mathbb{Z}) = \mathbb{Z}^2$. This chapter gives the NS5 curve Σ, whose tentacles run off to infinity in definite asymptotic directions. The claim is that these three are the same data, read on three sides of the construction.

The NS5 legs are the (p, q)-web legs. The asymptotics of the curve $\Sigma = \{P = 0\}$ are governed by the boundary edges of its Newton polygon Δ. A boundary edge of length L with primitive outward normal (p, q) gives L punctures under the nondegeneracy assumptions above. Each asymptotic cylinder winds the angular torus in that direction. Here (p, q) labels a geometric winding, not the local RR/NS charge of a segment of the

same NS5. In the separate five-brane web frame one uses $D5 = (1, 0)$ and $NS5 = (0, 1)$; identifying its charge lattice with this winding lattice is part of the duality dictionary. Thus

$$\begin{aligned}\{\text{NS5-curve asymptotic legs}\} &= \{\text{external } (p,q)\text{-web legs}\} \\ &= \{\text{primitive outward normals of } \partial\Delta\}.\end{aligned} \tag{VII.4.20}$$

The web of Chapter 2 and the curve have the same asymptotic integer directions. A consistency requirement comes for free: a closed polygon has its directed boundary edges summing to zero, and for a counterclockwise edge $e_i = (dx_i, dy_i)$ of lattice length $L_i = \gcd(|dx_i|, |dy_i|)$, its primitive normal is $n_i = (dy_i, -dx_i)/L_i$. With L_i copies of each normal, their sum is the right-angle rotation of $\sum_i e_i = 0$, so

$$\sum_{\text{legs}} (p, q) = (0, 0), \tag{VII.4.21}$$

a lattice closure identity. In a five-brane web frame it expresses external charge balance; it is not by itself a compact-space RR tadpole cancellation calculation for the D5–NS5 tiling.

The zig-zags are the legs. The third side is the dimer. A zig-zag path on the bipartite graph, traced by turning to the next edge in the cyclic order at each node and alternating the sense of the turn at black and white nodes, is a closed curve on T^2 with a homology class $[a, b] \in \mathbb{Z}^2$. The identification this chapter makes is that the zig-zag homology classes are the asymptotic directions of the NS5 curve, hence the (p, q)-web legs,

$$\boxed{[a,b]_{\text{zig-zag (Ch.3)}} \;=\; (p,q)_{\text{NS5 leg (Ch.4)}} \;=\; (p,q)_{(p,q)\text{-web leg (Ch.2)}}} \tag{VII.4.22}$$

as multisets of primitive vectors, compared up to the $GL(2, \mathbb{Z})$ equivalence under which toric diagrams are quoted. This is the three-way identification: the combinatorial object of Chapter 3 (a path on a graph), the geometric object of Chapter 2 (a leg dual to a polygon edge), and the brane object of Chapter 4 (a tentacle of the wrapped NS5) are one and the same direction in $\mathbb{Z}^2$. Near a primitive puncture, the leading edge polynomial reduces to a binomial $z^{-q}w^p = C$. Its angular equation is $-q\theta_1 + p\theta_2 = \arg C$ modulo 2π, so its tangent cycle has winding (p, q). In a compatible contour construction these puncture cycles become the zig-zags. The winding calculation alone does not prove the required global intersections.

The counting that follows. Counting both sides gives a clean equality of integers,

$$\#\{\text{zig-zags}\} = \#\{\text{external legs}\} = \#\{\text{primitive boundary segments of } \Delta\}, \tag{VII.4.23}$$

the number of zig-zag paths on the dimer equal to the number of legs of the web equal to the number of primitive segments along the polygon boundary. This count is one of the sharpest checks of a candidate tiling, because it is computed two genuinely independent ways: trace the zig-zags on the graph and count them, or read the boundary of the polygon and count its primitive segments. When the two disagree, the tiling is not the projection of the right curve. Table VII.4.2 collects the three-way identification on the fixtures.

fixture	toric polygon	NS5 legs	zig-zags	junction valency
$\mathbb{C}^3$	unit triangle	3	3	trivalent (cubic W)
conifold	unit square	4	4	four-valent (quartic W)
dP_0	local-$\mathbb{P}^2$ triangle, one interior point	3	3	trivalent (cubic W)

Table VII.4.2: The three-way identification (VII.4.22) on the fixtures: the number of NS5 asymptotic legs, the number of zig-zag paths of the dimer, and the number of primitive boundary segments of the toric polygon agree, and the junction valency equals the superpotential monomial length. dP_0 has the same leg/zig-zag count as $\mathbb{C}^3$ (three) but a compact interior point, the first sign of its richer dimer.

VII.4.7 Worked anchors: the conifold and $\mathbb{C}^3$

The two worked examples of the volume, the conifold square and the $\mathbb{C}^3$ honeycomb, make the three-way identification concrete. Both are re-read here through the T-dual lens: the curve solves below reuse the Newton-polygon polynomials fixed in Chapters 1–3, now read as the asymptotics of the wrapped NS5.

The conifold square. The conifold's toric diagram is the unit square with vertices $(0,0),(1,0),(1,1),(0,1)$, and its NS5-curve polynomial has these as its support,

$$P_{\mathcal{C}}(z,w) = c_{00} + c_{10}\,z + c_{01}\,w + c_{11}\,zw, \qquad \text{Newton}(P_{\mathcal{C}}) = \text{unit square}, \quad \text{(VII.4.24)}$$

the four monomials of the 1×1 Kasteleyn determinant computed in Chapter 3, with generic coefficients. Assume all four coefficients are nonzero and $D = c_{00}c_{11} - c_{10}c_{01} \neq 0$. The compactification is a sphere; the curve in $(\mathbb{C}^*)^2$ is that sphere with four punctures. To see this, solving $P_{\mathcal{C}} = 0$ for z,

$$z = -\frac{c_{00} + c_{01}\,w}{c_{10} + c_{11}\,w}, \qquad \text{(VII.4.25)}$$

exhibits the curve as a rational (genus-zero) curve, with one puncture at each of $w = 0$, $w = \infty$, and the two points $c_{00} + c_{01}w = 0$ and $c_{10} + c_{11}w = 0$ where z vanishes or blows up. They are distinct precisely under the stated conditions. These four punctures are its four tentacles, one per edge of the square, in the four primitive outward-normal directions

$$\{(1,0),(0,1),(-1,0),(0,-1)\}, \qquad \text{(VII.4.26)}$$

which sum to zero as (VII.4.21) requires. These are the four legs of the conifold (p,q)-web of Chapter 1 (two pairs of antiparallel rays), and they are the four zig-zag classes of the square tiling: the conifold dimer has four zig-zags, in these four directions, matching the four legs leg for leg. For the compatible square phase, the projected brane contours delimit two faces of the tiling (the two gauge nodes $U(N)_1 \times U(N)_2$), with four pinches (the four bifundamentals A_1, A_2, B_1, B_2) and two four-valent junctions (the two quartic superpotential terms, of opposite sign). The four-valence of the junctions is the brane-side reason the conifold superpotential is quartic: a four-valent NS5 junction carries a length-four monomial. The face, edge, and node counts of the square tiling

close on the torus Euler relation (VII.4.19),

$$(V, E, F)_{\mathcal{C}} = (2, 4, 2), \qquad V - E + F = 2 - 4 + 2 = 0, \tag{VII.4.27}$$

two nodes (one white, one black), four edges, two faces. Figure VII.4.3 collects the conifold's four legs as web, zig-zags, and NS5 tentacles.

The excluded case is informative. For $c_{00} = c_{10} = c_{01} = c_{11} = 1$,

$$P = (1 + z)(1 + w), \qquad \partial_z P = 1 + w, \qquad \partial_w P = 1 + z. \tag{VII.4.28}$$

The curve is the union of $z = -1$ and $w = -1$, crossing at $(-1, -1)$. The rational expression (VII.4.25) cancels to $z = -1$ and misses the other component. Its Newton polygon is still a square: polygon agreement alone cannot establish smoothness or a connected NS5 curve. As $D \to 0$, two punctures in the smooth parametrization collide. This is a degeneration, not another generic four-punctured sphere.

The $\mathbb{C}^3$ honeycomb. $\mathbb{C}^3$ is the simplest check. Its toric diagram is the unit triangle $(0, 0), (1, 0), (0, 1)$, and its NS5-curve polynomial is the three-monomial

$$P_{\mathbb{C}^3}(z, w) = c_{00} + c_{10}\, z + c_{01}\, w, \qquad \text{Newton}(P_{\mathbb{C}^3}) = \text{unit triangle}, \tag{VII.4.29}$$

a genus-zero curve with three tentacles when $c_{00}c_{10}c_{01} \neq 0$. Solving $P_{\mathbb{C}^3} = 0$ for z is immediate,

$$z = -\frac{c_{00} + c_{01}\, w}{c_{10}}, \tag{VII.4.30}$$

a rational (genus-zero) curve with punctures at $w = 0$, $w = \infty$, and $c_{00} + c_{01}w = 0$ (where $z = 0$), its three tentacles, one per edge of the triangle, the same solve as the conifold's (VII.4.25) but with one monomial fewer. Taking the chosen representative of the triangle's three primitive outward normals, the legs are

$$\{(0, -1), (-1, 0), (1, 1)\}, \tag{VII.4.31}$$

again summing to zero, and these are the three legs of the $\mathbb{C}^3$ junction of Chapter 2 and the three zig-zags of the honeycomb tiling. For the compatible honeycomb phase, the projected brane contours delimit its single face (the one gauge node $U(N)$ of $\mathbb{C}^3$, whose adjoint matter X, Y, Z are the three edges) and two trivalent junctions (the two cubic terms of the superpotential $\operatorname{Tr} X[Y, Z]$, of opposite sign). The trivalence of the honeycomb junctions is the brane-side reason the $\mathbb{C}^3$ superpotential is cubic: a trivalent junction carries a length-three monomial. Its tiling closes on the same torus Euler relation,

$$(V, E, F)_{\mathbb{C}^3} = (2, 3, 1), \qquad V - E + F = 2 - 3 + 1 = 0, \tag{VII.4.32}$$

two nodes, three edges, one face. The contrast with the conifold is exactly the contrast between a triangle and a square: three boundary segments and trivalent junctions versus four boundary segments and four-valent junctions. Figure VII.4.2 draws the honeycomb and its three legs.

dP_0, named. The orbifold $\mathbb{C}^3/\mathbb{Z}_3$, equal to dP_0, is the bridge back to Volume VI, where it was built as a McKay quiver. Its toric diagram is the triangle of local $\mathbb{P}^2$,

with three boundary segments and one interior point, so a nondegenerate NS5 curve has genus one and three asymptotic legs, while its chosen dimer has three zig-zags, matching the three (p, q)-web legs of the dP_0 web. The interior point, absent for $\mathbb{C}^3$ and the conifold, signals a compact divisor and a more elaborate dimer (a nine-edge graph with three faces); the polygon and web are introduced in Chapter 2, the full tiling and its Kasteleyn determinant in Chapter 3, and the perfect-matching and moduli-space computation in Chapter 5. Here dP_0 is named only for its leg and zig-zag count, three each, to record that the same orbifold has a Volume VI description as a McKay quiver and a Volume VII description as a brane tiling, the two descriptions of the same probe.

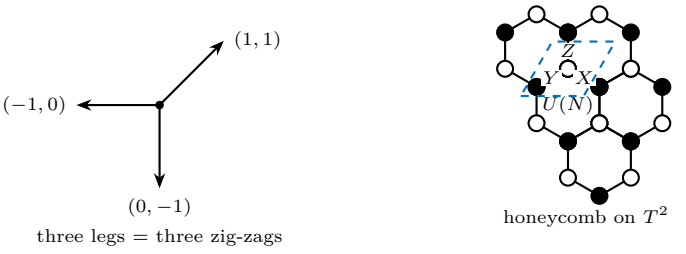
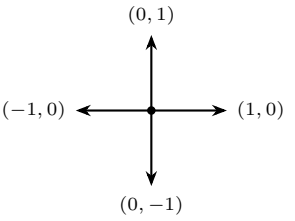
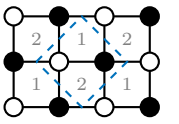


Figure VII.4.2: The $\mathbb{C}^3$ honeycomb, read three ways. Left: the three external legs $(0, -1), (-1, 0), (1, 1)$, equally the three NS5 tentacles and the three zig-zags, summing to zero. Right: a patch of the honeycomb tiling on the universal cover of T^2, one fundamental domain dashed. Each hexagonal face is the single gauge node $U(N)$; the three edge directions X, Y, Z are the three adjoint fields; the white and black vertices are the two cubic superpotential terms. One fundamental domain holds one white and one black vertex, the three edges, and one face, giving the two trivalent junctions of the cubic $\mathrm{Tr}(XYZ) - \mathrm{Tr}(XZY) = \mathrm{Tr}\ X[Y, Z]$.



four legs = four zig-zags



square tiling on T^2

Figure VII.4.3: The conifold's four legs, read three ways. Left: the four external (p, q)-web legs $(1, 0), (0, 1), (-1, 0), (0, -1)$, equally the four asymptotic tentacles of the NS5 curve $\Sigma_{\mathcal{C}}$, summing to zero. Right: a patch of the square brane tiling on the universal cover of T^2, one fundamental domain dashed (faces $1, 2$ are the gauge nodes; the black and white vertices are the two quartic superpotential terms; the four edges of a cell are the bifundamentals). The dashed cell has square-grid periods $(1, 1)$ and $(1, -1)$. Its four zig-zag paths carry exactly the four leg directions on the left. Read the left panel as the asymptotics of the projected curve and the right panel as the chosen compatible brane-tiling phase, not a literal plot of the full coamoeba.

VII.4.8 The mirror route, named

There is a second derivation of the same tiling, equivalent to the T-duality route and worth naming so the reader meets it under the right heading later. Mirror symmetry maps the toric Calabi–Yau threefold X to a mirror geometry built from the same Newton-polygon polynomial. In the Hori–Vafa presentation the mirror is the hypersurface

$$\check{X} = \{\, uv = P(z, w) \,\} \subset \mathbb{C}^2 \times (\mathbb{C}^*)^2, \qquad P(z, w) = \sum_{(m,n)\in\Delta\cap\mathbb{Z}^2} c_{mn}\, z^m w^n, \quad \text{(VII.4.33)}$$

a conic bundle $uv = P$ over the (z, w)-cylinder. Over a generic point of the cylinder $P(z, w) \neq 0$, so the fibre $uv = P$ is a smooth conic, a copy of $\mathbb{C}^*$. Over the curve $P(z, w) = 0$ the right-hand side vanishes, and the conic degenerates to the pair of lines $u = 0$ and $v = 0$ meeting at $u = v = 0$. The locus of degeneration is therefore exactly $\{P(z, w) = 0\}$, the same polynomial whose zero locus the T-dual NS5 wraps. The two descriptions therefore share this algebraic curve,

$$\begin{aligned} \text{T-duality:} \quad & \text{D3 on } X_\Delta \;\longrightarrow\; \text{D5/NS5}, \quad \Sigma = \{P(z, w) = 0\}, \\ \text{mirror:} \quad & \check{X} = \{uv = P(z, w)\} \;\longrightarrow\; \text{degeneration curve } P = 0, \\ \text{common output:} \quad & \text{chosen contours} \rightsquigarrow \text{dimer}, \qquad \text{amoeba}(P{=}0) \rightsquigarrow (p, q)\text{-web}, \end{aligned} \quad \text{(VII.4.34)}$$

so the conic degeneration and the NS5 locus are represented by the same equation. The fibre statement follows directly from $uv = P$: for $P \neq 0$ one can use $u \in \mathbb{C}^*$ and set $v = P/u$, whereas $P = 0$ gives two intersecting coordinate lines. Its interpretation as mirror symmetry is additional physics. Likewise, obtaining the tiling requires compatible brane cycles and a faithful angular projection; the two algebraic equations alone do

not establish that condition. The mirror construction supplies a useful independent description of the brane data, but neither route here computes the worldsheet amplitudes or the background metric. The detailed mirror map and contour construction belong to the foundations.

VII.4.9 What has been constructed, and what is deferred

The chapter relates the D3 probe to a Type IIB D5–NS5 description and computes the resulting dimension sequence and finite curve geometry. The physical T-duality dictionary is an input. For a chosen compatible phase, D5 regions, open-string intersections and oriented graph vertices reproduce gauge factors, chiral matter and superpotential monomials. The common four-supercharge presentation allows enhanced examples such as $\mathbb{C}^3$.

The Newton polygon fixes asymptotic directions with lattice multiplicities. Irreducibility and toric nondegeneracy give the smooth genus and puncture counts. Equality of Newton polygons does not fix coefficients, occupied interior points, a unique toric phase, or the full coamoeba. The conifold computation makes the distinction concrete: its generic curve is a four-punctured sphere, while $(1 + z)(1 + w) = 0$ is reducible with the same polygon. The triangle gives the three-punctured $\mathbb{C}^3$ curve; the local-$\mathbb{P}^2$ triangle has one interior point and gives genus one in the nondegenerate case.

A reader leaving the chapter should be able to do the following without looking ahead: run the dimension calculation D3 $\to$ D4 $\to$ D5 through two transverse T-dualities and check that the type stays Type IIB; state the end configuration as N D5 on $\mathbb{R}^{1,3} \times T^2$ plus one NS5 on $\mathbb{R}^{1,3} \times \Sigma$; write the NS5-curve polynomial with the prescribed Newton polygon, and distinguish that polygon from the support and coefficients; read faces, edges, and nodes in a compatible projected brane configuration and match them to the gauge nodes, bifundamentals, and superpotential terms; identify the zig-zag classes with the (p, q)-web legs and the NS5 tentacles, and check $\sum(p, q) = (0, 0)$; and count the conifold's four legs and $\mathbb{C}^3$'s three, with the matching junction valencies.

What has been deferred, and to where, is the rest of the volume and the foundations.

- The worldsheet derivation of T-duality on D-branes, the Buscher transformation and the Dp $\leftrightarrow$ D$(p \pm 1)$ dictionary, and the statement that a pinching cycle T-dualizes to an NS5, are foundations: they are quoted here and routed to the T-duality material of Volume III (T-duality and NS5 mirror geometry). This chapter checks the dimension and type calculation; it does not reprove the dictionary.

- The NS5-brane worldvolume theory and the holomorphic-curve, amoeba, and coamoeba technology that makes the projection precise are used with compatible contours and a faithful angular projection; the worldvolume analysis and the precise curve geometry are routed to the NS5 and brane-curve material (T-duality and NS5 mirror geometry, Volume II and Volume III) and to Part I.

- The mirror-symmetry derivation of the same tiling, the dimer from compatible brane contours on the mirror curve, is named as the equivalent alternative route

and routed to the mirror-symmetry volume; this chapter gives the T-duality route only.

- The dimer itself, its Kasteleyn matrix, perfect matchings, and zig-zag combinatorics, is Chapter 3, used here, not rebuilt; the forward algorithm from a tiling to a toric diagram, the GLSM charge matrices, the master space, and the Hilbert series are Chapter 5. The (p, q)-web and the toric diagram are Chapter 2.

- The identification "dimer = projected NS5 curve" is the leading brane-engineering picture and is topological and combinatorial: it fixes the leg directions and, with compatible brane contours, a graph, but not the metric or gauge couplings. Like the probe construction of Volume VI it carries the α' and g_s caveats; the precise complex structure of Σ, the isoradial and consistency refinements, and the metric data are geometry-dependent and are routed (the consistency filters themselves are Chapter 3).

- The dynamics that the construction supports, a-maximization (Chapter 6), toric duality as Seiberg duality realized as a move on the tiling (Chapter 7), and fractional branes and cascades (Chapter 8), all build on the physical tiling established here; the holographic spectra and the field-theory machinery beyond the construction are Part III. Orientifolds of the tiling are Chapter 9, and the higher-dimensional brane-brick analogue is Chapters 10–11.

The physical lesson is that the tiling describes a brane configuration with specified contour and phase data. The curve geometry gives stringent checks of that description, while gauge couplings and a physical-worldsheet derivation require information beyond the finite graph. The next chapter computes the moduli space from the quiver and superpotential on this stated footing.

Status and sources

[WS] means a controlled physical-worldsheet derivation in the untwisted string, not arbitrary worldvolume algebra. No such derivation is supplied here. The physical premises are supplied by the cited sources. [MATH] applies only to the stated mathematical domain and proof basis; a sourced premise is distinct from a calculation using it.

Setup and duality, §VII.4.2–VII.4.3. The D3 probe, Chan–Paton interpretation and T-duality rules are imported from Klebanov and Witten 1998; Polchinski 1998; Johnson 2003; Uranga 1999; Aganagic et al. 2000. The physical tiling construction is described in Hanany and Kennaway 2005; Franco et al. 2006; the NS5–D5 route through the mirror circle is discussed explicitly in Kennaway 2007, §4.6.2. Dimension arithmetic in (VII.4.4) and parity in (VII.4.5) are conditional on that dictionary. They do not prove supersymmetry preservation, localization at a degenerate fibre, or decoupling. The flat-$\mathbb{C}^3$ enhancement to sixteen supercharges is retained. Equation (VII.4.8) is the quadratic D5 action in the displayed constant-background approximation, using the D-brane action conventions in Polchinski 1998; Johnson 2003. Metric, dilaton and core corrections are not computed.

Curve and projections, §VII.4.4–VII.4.5. Newton polygon equality is weaker than support equality. The genus formula (VII.4.13) uses irreducibility and toric nondegeneracy over $\mathbb{C}$; the cited mathematical basis is Castryck and Voight 2009, §1–2, especially Proposition 1.7 and the adjunction discussion. The puncture count also requires transverse boundary intersections. The curve/dimer contour construction is in Feng et al. 2008, §5.3, §6.1 and §7.1–7.2. In particular, §7.2 assumes a

faithful projection of the intersection graph. A universal coamoeba retraction statement is not used: Forsgård 2019, Theorem 1.2 and §5, gives a criterion involving maximal complement components for its specified shell construction and examples obstructing a general assertion. That criterion is not tested here. Euler's relation assumes a cellular torus embedding. The three preserved drawings are schematic phase descriptions; they are not computed amoeba/coamoeba images.

Legs and examples, §VII.4.6–VII.4.7. The weighted normal sum is the rotated polygon boundary sum. Asymptotic cylinder winding is derived in Feng et al. 2008, equations (7.26)–(7.31). The geometric winding lattice is distinguished from the RR/NS charge lattice in a web duality frame. This is not a tadpole calculation. The conifold parametrization, its nonzero determinant condition, the factorization (VII.4.28), and the triangle parametrization are elementary algebra in this chapter. The smooth compactification and puncture statement has the explicitly stated coefficient domain. The honeycomb, square tiling and dP_0 counts are taken from the examples of Chapter 3; their physical identification is not established by comparing stored polygons.

Mirror and scope, §VII.4.8–VII.4.9. The mirror presentation and physical brane cycles follow Feng et al. 2008, §2 and §5. The fibre $uv = P$ is checked algebraically in the text; the mirror equivalence is cited, not proved. The foundation debt is T-duality and NS5 mirror geometry (T-duality, NS5 and mirror technology), with toric geometry for toric geometry. The dynamics destinations remain Chapters 6–8 and Part III, orientifolds Chapter 9, and brane bricks Chapters 10–11.

Conventions and calculation scope. We retain toric lattice orientation (web orientation and its charge convention), bipartite tiling conventions (dimer and signs), quiver gauge groups and ranks ($\prod U(N_a)$), superpotential normalization (cyclic words), fundamental-at-tail arrows (arrows), and spinor and supersymmetry counting (supercharges). The derivations give dimensions, parity, polygon normals and zig-zag homology for the stated examples. A coamoeba, a Killing spinor and a gauge coupling require separate calculations.

Exercise scope. The three exercise families cover brane orientation under fixed circle dualities; polygon area, boundary lengths and weighted normals; and rational parametrization with a reducible limit. See `benchmark/tasks.json` and `benchmark/PROTOCOL.md` for the public structured-answer format. The exercises recompute finite instances of (VII.4.4), (VII.4.5), (VII.4.21), (VII.4.25), (VII.4.30) and (VII.4.28). The exercises take the physical premises as inputs.

References

Aganagic, M. et al. (2000). "Mirror symmetries for brane configurations and branes at singularities". In: *Nucl. Phys. B* 569, pp. 277–302. DOI: 10.1016/S0550-3213(99)00674-3. arXiv: hep-th/9903093.

Castryck, W. and J. Voight (2009). "On nondegeneracy of curves". In: *Algebra & Number Theory* 3.3, pp. 255–281. arXiv: 0802.0420 [math.AG].

Feng, B. et al. (2008). "Dimer models from mirror symmetry and quivering amoebae". In: *Adv. Theor. Math. Phys.* 12, pp. 489–545. DOI: 10.4310/ATMP.2008.v12.n3.a2. arXiv: hep-th/0511287.

Forsgård, J. (2019). "On dimer models and coamoebas". In: *Annales de l'Institut Henri Poincaré D* 6, pp. 199–219. DOI: 10.4171/AIHPD/69. arXiv: 1602.01826 [math.AG].

Franco, S. et al. (2006). "Brane dimers and quiver gauge theories". In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Hanany, A. and K. D. Kennaway (2005). "Dimer models and toric diagrams". In: arXiv: hep-th/0503149.

Johnson, C. V. (2003). *D-Branes*. Cambridge Monographs on Mathematical Physics. Cambridge University Press. DOI: 10.1017/CBO9780511606540.

Kennaway, K. D. (2007). "Brane Tilings". In: arXiv: 0706.1660 [hep-th].
Klebanov, I. R. and E. Witten (1998). "Superconformal field theory on three-branes at a Calabi-Yau singularity". In: *Nucl. Phys. B* 536, pp. 199–218. DOI: 10.1016/S0550-3213(98)00654-3. arXiv: hep-th/9807080.
Polchinski, J. (1998). *String Theory. Vol. 2: Superstring Theory and Beyond*. Cambridge Monographs on Mathematical Physics. Cambridge University Press. DOI: 10.1017/CBO9780511618123.
Uranga, A. M. (1999). "Brane configurations for branes at conifolds". In: *JHEP* 01, p. 022. DOI: 10.1088/1126-6708/1999/01/022. arXiv: hep-th/9811004.

Chapter VII.5

Moduli spaces from dimers: perfect matchings and the Hilbert series

VII.5.1 From a tiling to a geometry

The four chapters before this one have, between them, assembled a complete object. The conifold chapter showed that a D3-brane at a toric singularity carries a quiver gauge theory with a superpotential, and that the same data can be drawn as a bipartite graph on a torus, a brane tiling or dimer. The general toric chapter ran that correspondence backwards, reading a quiver off a lattice polygon. The tiling chapter built the dimer itself, its Kasteleyn matrix and its consistency conditions; and the brane-configuration chapter explained where the tiling comes from physically, as the projection of an NS5–D5 system obtained by two T-dualities. We now hold a dimer, and with it a quiver and a superpotential, and we have so far taken it on faith that this data reconstructs the transverse geometry. We compute the classical abelian quotient from those data. Its interpretation as the transverse D3-brane geometry still uses the physical correspondence established in the preceding chapters.

The question is sharp. Given a dimer, with E edges (the bifundamental fields), G faces (the gauge nodes), and a superpotential read off its black and white vertices, what is the classical moduli space of vacua of the gauge theory, and how does one extract from it the toric Calabi–Yau threefold that the brane is probing? The answer is an algorithm, the *forward algorithm*: from the dimer one enumerates its perfect matchings, packages them into a small number of integer charge matrices, and reads the toric diagram off the kernel of those matrices. It runs in the reverse direction to the general toric chapter, which went from the polygon to the quiver; here the gauge theory is the input and the geometry is the output.

Take-home message of this chapter

For a connected, nondegenerate, consistent toric tiling and one regular brane, the dimer encodes a quiver and a superpotential. Its *perfect matchings* (the ways to cover every vertex of the bipartite graph by disjoint edges) are the natural homogeneous coordinates of the gauge linear sigma model. Expressing each field as a monomial in the perfect matchings collapses the F-terms and D-terms into integer charge matrices P, Q_F, Q_D, Q_t, and the kernel of Q_t is the toric diagram. *The geometry is computed from the gauge theory.* The map from perfect matchings to lattice points is many-to-one, so points carry multiplicities. The Q_F quotient describes the coherent component of the F-flat master space, of dimension $G+2$; the additional Q_D quotient gives the three-dimensional mesonic cone at zero Fayet–Iliopoulos parameters.

Two corrections to the naive picture must be flagged at the outset, because both are easy to get wrong and both are computed in full below. First: perfect matchings are *not* in one-to-one correspondence with the lattice points of the toric diagram. Each matching maps to a lattice point, but several matchings can map to the same point, which then carries a *multiplicity*. Equality of the two counts is the exception, not the rule, realized only in the simplest cases. Second: the two quotients describe different spaces. The Q_F quotient retains the coherent F-flat component, including relative abelian gauge orbits; the full Q_t quotient removes these orbits to give the mesonic threefold. The full F-flat variety can have further components. The count of non-anomalous baryonic global symmetries in the interacting nonabelian theory is a separate question. Both points have a precise statement and a worked example here.

A word on scope, so the reader knows what to hold and what is deferred. This chapter consumes the dimer as built in the tiling chapter; it does not re-derive the bipartite graph, the Kasteleyn matrix, the sign rule, or the zig-zag consistency filters. It computes the *classical* moduli space and its counting data, for one regular brane in the toric phase; it does not do the moduli-space dynamics. In particular the Hilbert series here is graded by a counting degree, not by the superconformal R-charge: the principle that fixes the physical R-charge, a-maximization, is the next chapter, and is only named here. The plan is: the perfect matchings and the P-matrix (§VII.5.2); the charge matrices Q_F, Q_D, Q_t and the toric diagram as a kernel (§VII.5.3); the conifold worked through (§VII.5.4); the perfect-matching multiplicities (§VII.5.5); the master space and its mesonic/baryonic anatomy (§VII.5.6); the Hilbert series and the plethystics (§VII.5.7); the worked Volume VI bridge $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ end-to-end (§VII.5.8); and a closing map of what was computed and what is deferred (§VII.5.9). The organizing thread is the one in the keybox: the dimer is the McKay quiver's successor, and this is the chapter that turns it into geometry.

VII.5.2 Perfect matchings and the P-matrix

Fix a dimer: a bipartite graph on the two-torus, with white vertices w, black vertices b, and edges between them, each edge an oriented bifundamental field of the quiver. A **perfect matching** p_α is a subset of the edges such that every vertex, white and black, is the endpoint of exactly one chosen edge. Combinatorially the perfect matchings

index the expanded terms of the permanent of the Kasteleyn matrix, equivalently the expanded terms of $\det K$ (before coefficients are collected) built in the tiling chapter; here we take them as given and count them. Write

$$E = \#\{\text{edges}\} = \#\{\text{fields}\}, \qquad G = \#\{\text{faces}\} = \#\{\text{gauge nodes}\},$$
$$c = \#\{\text{perfect matchings}\}. \tag{VII.5.1}$$

For the conifold, the four edges are themselves the four matchings ($c = 4$); for the honeycomb of $\mathbb{C}^3$, $c = 3$; for the dP_0 tiling computed in §VII.5.8, $c = 6$.

The reason perfect matchings are the right variables is that they are the homogeneous coordinates of the gauge linear sigma model whose vacuum moduli space is the brane's geometry. The dictionary between the original fields and the matchings is one matrix.

Principle VII.5.1 (The P-matrix)**.** For a dimer with E edges and c perfect matchings, the *perfect-matching matrix* P is the $E \times c$ matrix of zeros and ones,

$$P_{a\alpha} = \begin{cases} 1, & \text{edge } a \text{ belongs to matching } p_\alpha, \\ 0, & \text{otherwise,} \end{cases} \qquad X_a = \prod_{\alpha=1}^{c} p_\alpha^{P_{a\alpha}}. \tag{VII.5.2}$$

The rows are indexed by the quiver fields X_a, the columns by the perfect matchings p_α, and the second equation expresses each field as the monomial in the matchings that contain its edge. The p_α are the gauge-linear-sigma-model fields; P is the change of variables from the quiver fields to them.

We assume a connected, nondegenerate, consistent toric brane tiling describing one regular brane, with the standard signed superpotential and its coherent toric component. These hypotheses exclude arbitrary bipartite graphs. Under them the matrix has rank

$$\operatorname{rank} P = G + 2, \tag{VII.5.3}$$

a fact verified on every fixture below and central to the dimension counts: it says that, of the c perfect matchings, only $G + 2$ combinations are independent in expressing the original fields, and the remaining $c - G - 2$ kernel directions are redundancies of these homogeneous coordinates. They are not a count of polynomial generators of the F-term ideal. And the map (VII.5.2) is multiplicative: a product of fields is the product of the corresponding monomials, so products of fields become monomials in the matchings. For an edge a, the two terms in $\partial W/\partial X_a$ come from its two endpoints. Each matching covers each endpoint once. If it contains a, its exponent is zero on both differentiated terms; if it does not, its exponent is one on both. Thus the two monomials agree after substitution. This argument uses commuting fields for one regular brane; it is not a matrix F-term proof.

For the two simplest fixtures the P-matrix is already explicit and trivial. Each matching is a single edge, so P is an identity matrix,

$$P_{\mathbb{C}^3} = \mathbb{1}_3 \quad (E = c = 3), \qquad P_{\text{conifold}} = \mathbb{1}_4 \quad (E = c = 4), \tag{VII.5.4}$$

each field equal to its own matching, $X = p_X, Y = p_Y, Z = p_Z$ for $\mathbb{C}^3$ and $A_1 = p_1, B_1 = p_2, A_2 = p_3, B_2 = p_4$ for the conifold. Here $\operatorname{rank} P = c = G + 2$ leaves no

F-term relations, the degenerate case worked in §VII.5.4. The first nontrivial P-matrix, 9×6 of rank 5 with a one-dimensional kernel, is dP_0's in §VII.5.8; it is where the change of variables does real work.

We can see directly why (VII.5.2) is a sensible change of variables and not a formal trick. The original fields X_a are subject to the F-term relations $\partial W = 0$, which are not independent and which a direct attack on the moduli space must impose by hand. The perfect matchings, by contrast, are unconstrained coordinates: every assignment of them gives fields satisfying the abelian F-terms by the endpoint argument above. The closure of this monomial image is the coherent component under the stated tiling hypotheses; additional components of the full F-flat variety need separate treatment. Passing to the p_α therefore trades a constrained problem in the X_a for a free polynomial ring in the p_α, quotiented by the linear charges of the next section. The work of solving $\partial W = 0$ has been done once and for all by the combinatorics of the matchings.

VII.5.3 The charge matrices and the toric diagram

With the perfect matchings as coordinates, the moduli space is the space of the p_α modulo two sets of integer charges: the F-term charges, which record the relations among the matchings, and the D-term charges, which record the gauge action. We build them in turn, then stack them.

The F-term charges Q_F. The original fields are not freely related to the matchings: the map (VII.5.2) can be non-injective on the level of charges, and the kernel of P records the torus action that leaves all original fields unchanged. An integer vector $q \in \mathbb{Z}^c$ assigns a charge q_α to each matching p_α; under it a field X_a acquires charge $\sum_\alpha P_{a\alpha} q_\alpha = (Pq)_a$. Demanding that every field be neutral, so that the change of variables is consistent, is the condition $Pq = 0$. The F-term charge matrix is therefore a basis of the integer kernel of P itself,

$$Q_F = \ker_{\mathbb{Z}}(P) \subset \mathbb{Z}^c, \qquad \text{shape } (c - \operatorname{rank} P) \times c = (c - G - 2) \times c. \tag{VII.5.5}$$

This is the kernel of P, not of its transpose: with P taken $E \times c$ as in (VII.5.2), the F-flatness relations are the integer combinations of *columns* (matchings) that the fields cannot see, and those are the vectors q with $Pq = 0$. The number of such rows, $c - G - 2$ by (VII.5.3), is the rank of the redundant torus action on the matchings, rather than the number of defining equations in the original fields. For $\mathbb{C}^3$ and the conifold, $c = G + 2$ exactly and Q_F is empty; the first nonempty Q_F appears at dP_0, with a single row.

The D-term charges Q_D. The gauge group is $\prod_i U(N_i)$, one unitary factor per node, and the matchings carry charges under it. The overall diagonal $U(1)$, under which everything is neutral and which is the decoupled centre of mass, is removed, leaving $G - 1$ independent abelian charges. The matching charges Q_D, of shape $(G-1) \times c$, are fixed by requiring them to reproduce, through the P-matrix, the known gauge charges of the original fields. Let δ be the reduced quiver incidence matrix, the $(G-1) \times E$ matrix whose entry is the charge of field X_a under the i-th node $U(1)$, with the diagonal $U(1)$ dropped (each field has charge $+1$ at its tail node and -1 at its head node, and

one of the G node-rows is redundant). Then Q_D is an integral lift satisfying

$$\delta = Q_D\, P^{\mathsf{T}}, \qquad \text{shape of } Q_D = (G-1)\times c, \tag{VII.5.6}$$

which says: the charge of the field X_a (a column of δ, equivalently a row of δ^{T}) must equal the charge $\sum_\alpha Q_{D,i\alpha}P_{a\alpha}$ that the matchings induce on it through P. This is a reproducible prescription, not a guess: given the dimer's incidence and its P-matrix, (VII.5.6) is a linear system for Q_D, solved over the integers on the toric fixtures. For $\mathbb{C}^3$ there is a single node, $G=1$, and Q_D is empty; for the conifold, $G=2$, Q_D is one row, the baryonic charge ($A:+1$, $B:-1$) of the conifold chapter.

An integral basis matters. Clearing denominators in a rational nullspace gives integer vectors, but it need not give the whole integer kernel. For example, the primitive vectors $(-1,2,0)$ and $(-1,0,2)$ annihilate $(2,1,1)$, yet miss $(-1,1,1)$. For a full-row-rank integer matrix B whose rows span the desired rational kernel, the gcd of its maximal minors is the index in its saturation. Requiring that gcd to be one certifies the full lattice: in Smith normal form the nonzero invariant factors multiply to this gcd, and an index-one sublattice equals its saturation. Integer row changes must be unimodular. For Q_D, adding an integer combination of Q_F changes no field charge; changing the gauge-row basis by $U \in GL(G-1,\mathbb{Z})$ requires changing δ to $U\delta$ as well. The total charges present an affine torus quotient at zero FI parameters; a nonzero FI chamber requires stability data in addition to the kernel polygon.

The total charge matrix and the kernel. Stacking the two charge matrices gives the total gauge-linear-sigma-model charge matrix

$$Q_t = \begin{pmatrix} Q_F \\ Q_D \end{pmatrix}, \qquad \text{shape } (c-3)\times c, \tag{VII.5.7}$$

of $(c-G-2)+(G-1)=c-3$ rows. The toric diagram is the integer kernel of Q_t,

$$G_t = \ker_{\mathbb{Z}}(Q_t), \qquad \text{shape } 3\times c, \tag{VII.5.8}$$

a $3\times c$ matrix whose c columns are the lattice points $G_\alpha \in \mathbb{Z}^3$ at which the matchings sit. That the kernel is three-dimensional, for any toric Calabi–Yau threefold, is the statement that Q_t has rank $c-3$; the three rows of G_t are read as one height row (constant, the Calabi–Yau condition) and a two-dimensional toric polygon, exactly the orientation convention of the toric chapters. The Newton polygon of $\det K$, computed independently from the Kasteleyn matrix, is the convex hull of these same points, and the two agree on every fixture: the forward algorithm and the dimer's characteristic polynomial are two routes to one polygon.

The forward algorithm, in six lines

(1) Enumerate the perfect matchings $p_1,\dots,p_c$ of the dimer. **(2)** Build P, the $E\times c$ matrix with $P_{a\alpha}=1$ iff edge $a\in p_\alpha$, so $X_a=\prod_\alpha p_\alpha^{P_{a\alpha}}$. **(3)** $Q_F=\ker_{\mathbb{Z}}(P)$, shape $(c-G-2)\times c$ (redundant homogeneous-coordinate charges). **(4)** Q_D from $\delta = Q_D P^{\mathsf{T}}$, shape $(G-1)\times c$ (the gauge charges, the diagonal $U(1)$ removed). **(5)** $Q_t=(Q_F;\,Q_D)$, shape $(c-3)\times c$. **(6)** The toric diagram is $\ker_{\mathbb{Z}}(Q_t)$, a $3\times c$ matrix of lattice points; its convex hull is the polygon, equal to the Newton polygon of $\det K$.

The shapes are worth holding in mind as a running consistency check, since they are fixed entirely by G and c: P is $E \times c$ of rank $G+2$; Q_F has $c-G-2$ rows, Q_D has $G-1$ rows, Q_t has $c-3$ rows, and the kernel of Q_t has the three dimensions of the Calabi–Yau. We turn to the smallest nontrivial instance.

VII.5.4 The conifold, worked through

The trivial control example: $\mathbb{C}^3$. Before the conifold, flat $\mathbb{C}^3$ is the zero-charge control example on which every matrix is empty. Its honeycomb has $E=3$ edges, $G=1$ face, and $c=3$ perfect matchings, one per edge, so the P-matrix is the 3×3 identity $P=\mathbb{1}_3$. Then $\operatorname{rank} P = 3 = G+2$, the F-term matrix $Q_F = \ker_{\mathbb{Z}}(P)$ has $c-G-2=0$ rows, the D-term matrix Q_D has $G-1=0$ rows, and the total Q_t is empty: with no charges to impose, the three perfect matchings are themselves the three lattice points, the unit triangle, and the moduli space is $\mathbb{C}^3$ with no quotient. This is the control example against which the conifold's single D-term charge is the first nontrivial step.

The conifold. The conifold is the cleanest first run of the algorithm with a genuine quotient: its dimer is the square lattice, with one white and one black vertex, four edges, and two faces, so $E=4$, $G=2$. Its four perfect matchings are the four edges themselves, each a single edge covering the one white and the one black vertex,

$$p_1=\{A_1\},\quad p_2=\{B_1\},\quad p_3=\{A_2\},\quad p_4=\{B_2\},\qquad c=4. \tag{VII.5.9}$$

Because each matching is a single edge, the P-matrix is the 4×4 identity,

$$P=\begin{pmatrix}1&0&0&0\\0&1&0&0\\0&0&1&0\\0&0&0&1\end{pmatrix},\qquad \operatorname{rank} P = 4 = G+2, \tag{VII.5.10}$$

so each field equals its own matching, $A_1=p_1$, $B_1=p_2$, $A_2=p_3$, $B_2=p_4$. The rank is $G+2=4$, equal to c, so there are no F-term relations and Q_F is empty; this is the statement that the conifold's F-terms, which vanished identically for one brane in the conifold chapter, impose nothing on the matchings either.

The single D-term charge is the baryonic charge. The reduced incidence has one row, the node-2 charge of each field with the diagonal removed, and solving $\delta = Q_D P^{\mathsf{T}}$ with P the identity returns δ itself,

$$Q_D=\begin{pmatrix}-1 & 1 & -1 & 1\end{pmatrix}, \tag{VII.5.11}$$

the charges $(A_1,B_1,A_2,B_2)=(-1,+1,-1,+1)$, that is $A:-1$, $B:+1$, the relative $U(1)$ of the conifold quotient up to an overall sign. With Q_F empty, the total charge matrix is this single row,

$$Q_t = Q_D = \begin{pmatrix}-1 & 1 & -1 & 1\end{pmatrix},\qquad \text{shape } 1\times 4 = (c-3)\times c. \tag{VII.5.12}$$

Its integer kernel is three-dimensional, with a row basis

$$G_t = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix}, \tag{VII.5.13}$$

read as the height-one row on top and the two-dimensional polygon below: the four lattice points $(0,0),(1,0),(1,1),(0,1)$, the unit square, the conifold's toric diagram. In the signing of Chapter 3, the same square is the support of $\det K = 1 - z - w - zw$ of the conifold's 1×1 Kasteleyn matrix. The four edge signs have product -1, as required by the square-face rule; only the Newton support is compared here. The algorithm has reproduced the geometry: the four matchings are the four corners of the unit square, each of multiplicity one, and the invariant ring of the baryonic charge action is the conifold coordinate ring.

The algorithm even reconstructs the coordinate *ring*, not just the diagram. The mesonic moduli space is the GIT quotient of the four matchings by the single charge,

$$\mathcal{M}^{\text{mes}} = \operatorname{Spec}\Big(\mathbb{C}[p_1,p_2,p_3,p_4]^{\mathbb{C}^*_{(-1,1,-1,1)}}\Big), \tag{VII.5.14}$$

the ring of charge-zero monomials. With p_1, p_3 of charge -1 and p_2, p_4 of charge $+1$, the neutral monomials are generated by the four products

$$x = p_1 p_2, \quad y = p_3 p_4, \quad z = p_1 p_4, \quad w = p_3 p_2, \tag{VII.5.15}$$

Indeed, charge neutrality equates the number of factors drawn from p_1, p_3 with the number drawn from p_2, p_4. Pair the factors across these two sets in any order. Every pair is one of x, y, z, w, so every neutral monomial is a product of the displayed four. These generators satisfy the relation

$$xy = (p_1 p_2)(p_3 p_4) = p_1 p_2 p_3 p_4 = (p_1 p_4)(p_3 p_2) = zw. \tag{VII.5.16}$$

This is $xy = zw$: the forward algorithm has reproduced not only the conifold's toric diagram but its coordinate ring, now from the GLSM charge Q_t rather than from the geometry of Chapter 1.

VII.5.5 Perfect matchings carry multiplicities

The conifold and $\mathbb{C}^3$ are deceptively clean: their matchings are in one-to-one correspondence with the lattice points of their toric diagrams, three corners for $\mathbb{C}^3$ and four for the conifold, each point hit once. This is the exception. In general the map sending a matching p_α to its lattice point G_α is *many-to-one*: several matchings can land on the same point, which then carries a multiplicity equal to the number of matchings projecting to it.

Do not confuse: matchings versus lattice points

The perfect matchings of a dimer are *not* in bijection with the lattice points of the toric diagram. The map (matching $\mapsto$ lattice point) is many-to-one, so

$$\#\{\text{perfect matchings}\} \;\geq\; \#\{\text{lattice points}\}, \tag{VII.5.17}$$

with equality only in simple multiplicity-free fixtures such as $\mathbb{C}^3$ and the conifold. The number of matchings landing on a point is that point's *multiplicity*. For the consistent (minimal) tilings used here, corner (vertex) points always have multiplicity one, while internal and some boundary points generically have multiplicity greater than one. The mesonic Hilbert series of the next section *collapses* the multiplicities, seeing only the toric cone; a count retaining the homogeneous matching variables distinguishes their labels. A master-space Hilbert series can retain additional charge information, but need not reconstruct a unique perfect-matching presentation. Even flavour refinement of the mesonic ring does not recover arbitrary GLSM multiplicities.

The multiplicities are not chosen, they are computed, and computed two independent ways that must agree. The first is by enumeration: list the matchings, map each to its lattice point via the net winding it accumulates around the two cycles of the torus, and count how many land on each point. The multiplicity is *defined* by that enumeration count. The second route is the Kasteleyn determinant: the support of $\det K$ is the toric diagram, and the *integer coefficient* at each lattice point equals, up to sign, the multiplicity, provided the matchings over the point do not cancel in sign (as on every fixture here). A point of multiplicity m then appears in $\det K$ with coefficient $\pm m$. The agreement of these two counts, the combinatorial enumeration and the polynomial coefficient, is a nontrivial consistency check on the dimer, and it holds on every fixture.

The first example with a genuine multiplicity is dP_0, worked in full in §VII.5.8: it has six perfect matchings but only four lattice points, three corners of multiplicity one and a single interior point of multiplicity three. Its Kasteleyn determinant, anticipating the computation,

$$\det K = z^{-1} + w^{-1} - 3\, z^{-1}w^{-1} + z^{-2}w^{-2}, \tag{VII.5.18}$$

shows the interior coefficient -3 at once: three matchings sit at the interior point, and the coefficient of that monomial counts them. The three corners carry coefficient $+1$, multiplicity one each, and the total $3 \times 1 + 1 \times 3 = 6$ matches the matching count. This is the cleanest illustration of the multiplicity, and the reason a bijection is the wrong picture. Table VII.5.1 accumulates the six matchings onto the four lattice points, with the $\det K$ coefficient alongside the enumeration count, and Figure VII.5.1 draws the same many-to-one map.

The general combinatorics of which points carry which multiplicities, for an arbitrary toric threefold, is catalogued for the families in the atlas chapter; here the principle and one worked instance suffice.

lattice point	matchings landing there	multiplicity	$\det K$ coefficient
$(0,0)$ corner	p_1	1	$+1$
$(1,2)$ corner	p_4	1	$+1$
$(2,1)$ corner	p_5	1	$+1$
$(1,1)$ interior	p_2, p_3, p_6	3	-3

Table VII.5.1: The dP_0 multiplicities, by enumeration and by determinant. The six perfect matchings (VII.5.35) accumulate onto four lattice points: the three corners once each, the interior point $(1,1)$ three times. The multiplicity (the enumeration count) equals the absolute value of the coefficient of that monomial in $\det K$ (VII.5.18), the interior -3 counting the three matchings p_2, p_3, p_6. The two routes agree, $3\times1+1\times3=6$.

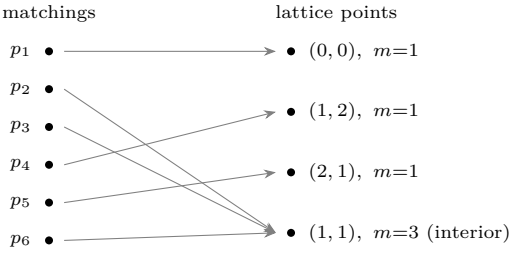


Figure VII.5.1: The many-to-one matching-to-point map of dP_0, the source of the multiplicities. The six perfect matchings $p_1, \dots, p_6$ (left) map to the four lattice points of the toric diagram (right): p_1, p_4, p_5 land on the three corners, one each (multiplicity $m = 1$), while p_2, p_3, p_6 all land on the single interior point $(1,1)$, which therefore has multiplicity $m = 3$, so $3 \times 1 + 1 \times 3 = 6$. The interior coefficient -3 of $\det K$ (VII.5.18) counts those three. The map is many-to-one; $\mathbb{C}^3$ and the conifold, with no interior point, are the bijective exceptions. (The lattice points are one $GL(2,\mathbb{Z})$ representative of the local-$\mathbb{P}^2$ diagram; the matchings $p_1, \dots, p_6$ and their points are worked in full in §VII.5.8.)

VII.5.6 The master space and the mesonic–baryonic split

The forward algorithm contains two successive quotients. Their difference is already visible for the conifold: four unconstrained abelian fields form $\mathbb{C}^4$, while their gauge-invariant products describe a three-dimensional quadric.

Principle VII.5.2 (The master space and its coherent component)**.** Write $I_F = \langle \partial W/\partial X_a \rangle \subset \mathbb{C}[X_1, \dots, X_E]$. The full master space $\mathcal{F}^\flat = V(I_F)$ is the F-flat variety before the gauge quotient. Its coherent component $\mathcal{F}^\flat_{\mathrm{coh}}$ is the closure of the locus where all fields are nonzero. For the consistent toric phases considered here, perfect matchings parametrize this component, and

$$\dim_{\mathbb{C}} \mathcal{F}^\flat_{\mathrm{coh}} = c - \operatorname{rank} Q_F = G + 2. \tag{VII.5.19}$$

The full F-flat variety can be reducible. The dimension formula and matching quotient

refer to the coherent component, not to each component of $V(I_F)$.

The distinction has a direct algebraic test. Set $s = X_1 \cdots X_E$. The closure of the F-flat locus in the algebraic torus is defined by

$$I_{\rm coh} = I_F : s^\infty = \big(\langle I_F, 1 - us\rangle \cap \mathbb{C}[X_1, \ldots, X_E]\big). \tag{VII.5.20}$$

Here u is an auxiliary variable eliminated after imposing $us = 1$. The equality follows by localizing the ring at s and contracting back. Calling this ideal the coherent component assumes the torus locus is irreducible and reduced, as in our toric phases. The alternative elimination of $X_a - \prod_\alpha p_\alpha^{P_{a\alpha}}$ computes the closure of the matching image. Agreement with the saturated F-term ideal is the precise comparison; a dimension count alone would not establish it.

For one regular brane, the gauge group is $U(1)^G$. Its diagonal acts trivially. At zero FI parameters the affine mesonic quotient is

$$\mathcal{M}^{\rm mes} = \operatorname{Spec} \mathbb{C}[\mathcal{F}^\flat_{\rm coh}]^{(\mathbb{C}^*)^{G-1}}, \tag{VII.5.21}$$

and, on the dense locus where the effective torus acts with generic orbit dimension $G - 1$,

$$\dim_{\mathbb{C}} \mathcal{M}^{\rm mes} = c - \operatorname{rank} Q_t = 3. \tag{VII.5.22}$$

Thus $G + 2 = 3 + (G - 1)$ is a dimension difference under a quotient, not a canonical product or an embedding of the quotient as a branch of the master space. Nonzero FI parameters select GIT chambers and may resolve the cone. Higher-rank $\prod_i U(N_i)$ vacua require matrix F-terms and are not computed by this abelian construction.

The two fixtures with a quotient make the counts concrete,

$$\begin{aligned} \text{conifold } (G = 2):&\ \dim \mathcal{F}^\flat = 4,\ \dim \mathcal{M}^{\rm mes} = 3,\ G - 1 = 1,\ d - 3 = 1, \\ dP_0\ (G = 3):&\ \dim \mathcal{F}^\flat = 5,\ \dim \mathcal{M}^{\rm mes} = 3,\ G - 1 = 2,\ d - 3 = 0, \end{aligned} \tag{VII.5.23}$$

the conifold having one baryonic direction that is also a non-anomalous symmetry, dP_0 having two baryonic master-space directions but no non-anomalous baryonic symmetry, the distinction made sharp in the next keybox.

F-flat space and gauge quotient

The master space is an F-flat parent space. Its coherent component has dimension $G + 2$; the effective abelian gauge quotient has dimension three. The $G - 1$ relative gauge orbits are often called baryonic directions in the master-space description. Determinants of bifundamentals at $N > 1$ are invariant under the special-unitary factors, but carry charges under relative $U(1)$ factors. They are not gauge-invariant baryons of the full $\prod_i U(N_i)$ theory without further qualifications. Conserved baryonic global symmetries refer to the interacting nonabelian theory after the abelian factors are treated separately.

So far so symmetric: $G - 1$ baryonic directions in the geometry, $G - 1$ D-term charge rows in Q_D. But the count of physical, conserved baryonic symmetries of the interacting fixed point is a *different* number, and this is the sharpest distinction of the chapter.

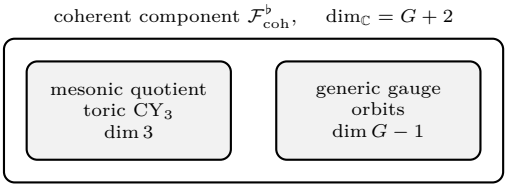


Figure VII.5.2: Dimensions in the abelian quotient. The boxes record the dimension of the mesonic quotient and the generic relative gauge orbit inside the coherent F-flat component; they do not depict a product decomposition or a mesonic subvariety. At zero FI parameters the quotient is the toric cone. The separate $d-3$ count concerns non-anomalous baryonic global symmetries of the interacting nonabelian theory.

Do not confuse: $G-1$ versus $d-3$

The forward algorithm uses $G-1$ reduced D-term (baryonic) charge rows, where G is the number of gauge nodes. These are the gauge-linear-sigma-model charges, and they enter Q_D and the dimension count. But the number of physical, *non-anomalous* baryonic global $U(1)$ symmetries of the superconformal theory is

$$\#\{\text{non-anomalous } U(1)_B\} = d - 3, \tag{VII.5.24}$$

For the isolated toric cones used here, d counts the boundary lattice points, which are the primitive corners; equivalently the number of external legs of the (p,q)-web. Within this class $d-3 \le G-1$, and the anomalous baryonic charges refine the gauge-linear-sigma-model count but are not conserved symmetries of the interacting theory. Conifold: $d=4$, so $d-3=1$, and $G-1=1$ too, they coincide. $\mathbb{C}^3$: $G=1$, no quotient and no baryonic charge. dP_0: $d=3$, so $d-3=0$, no continuous non-anomalous baryonic symmetry, even though $G-1=2$. Never read the baryonic count off the node number.

The conifold is the case where the two counts coincide, which is exactly why it can mislead. With $G=2$ nodes and $d=4$ boundary points, $G-1=d-3=1$, and the relative charge assignment becomes the baryonic global symmetry when the abelian gauge factors are separated from the interacting special-unitary theory. The first example separating the counts is dP_0: its triangular toric diagram has $d=3$ boundary points and no interior boundary structure, so $d-3=0$ and there is no continuous non-anomalous baryonic symmetry, while the node count would wrongly suggest $G-1=2$. The boundary-point count $d-3$ is the correct invariant, read off the geometry, and the $G-1$ D-term rows of the gauge-linear-sigma-model are an overcount of the physical symmetries whenever $d-3<G-1$. The baryonic-branch dynamics, the role of the baryon operators in confinement and the duality cascade, are the business of later chapters and of Part III; here the directions are identified and counted, and the two counts kept apart.

VII.5.7 The Hilbert series and the plethystics

The toric diagram fixes the moduli space; the **Hilbert series** counts the operators on it. It is the generating function for the dimensions of the graded pieces of the coordinate ring,

$$H(t) = \sum_{n \geq 0} \dim(R_n)\, t^n, \tag{VII.5.25}$$

where the grading is positive, $R_0 = \mathbb{C}$, and each R_n is finite-dimensional. For the mesonic ring, R_n counts classical gauge-invariant chiral operators of degree n. For a finitely generated positively graded coordinate ring of a toric Calabi–Yau threefold H is a rational function whose pole at $t = 1$ has order equal to the complex dimension, three. The series can be mesonic, counting only the toric-cone operators, or refined by baryonic and flavour fugacities; the mesonic series is the main object here.

Declare the grading. A Hilbert series is meaningless until one says what degree the variable t counts, because the same ring admits more than one grading. The conifold coordinate ring $\mathbb{C}[x, y, z, w]/(xy - zw)$, with four generators and one relation at twice the generator degree, has the series

$$H_{\mathcal{C}}(t) = \frac{1 - t^2}{(1 - t)^4} = \frac{1 + t}{(1 - t)^3} \qquad \text{(generators of degree 1)}, \tag{VII.5.26}$$

the four generators contributing $(1-t)^{-4}$ and the single quadratic relation the numerator $1 - t^2$. If instead the four generators are taken to be the mesons $M_{ij} = A_i B_j$, which are degree two in the elementary A, B fields, one sets $t \to t^2$ and obtains

$$H_{\mathcal{C}}(t) = \frac{1 - t^4}{(1 - t^2)^4} \qquad \text{(mesons of degree 2)}. \tag{VII.5.27}$$

These are the *same ring*, the relation $xy = zw$ at twice the generator degree in both, and the two gradings differ only by $t \mapsto t^2$; the check passes or fails on which grading is declared, so the grading is stated, not assumed. (The physical grading, by the superconformal R-charge, is the one a-maximization fixes in the next chapter; here the degree is the elementary counting grade.) Expanding (VII.5.27) in even powers,

$$H_{\mathcal{C}}(t) = \sum_{n \geq 0} (n + 1)^2\, t^{2n} = 1 + 4t^2 + 9t^4 + 16t^6 + \cdots, \tag{VII.5.28}$$

and the coefficient $(n + 1)^2$ is the dimension of the spin-$(n/2, n/2)$ representation of $SU(2) \times SU(2)$: the mesonic operators of the conifold organize into the tensors symmetric separately in the n A-indices and the n B-indices, exactly the geometry's isometry seen on the operator spectrum.

The same series from the quotient. The coefficients $(n+1)^2$ are not only recognized; they are computed directly from the gauge-linear-sigma-model quotient, as a Molien integral over the baryonic $U(1)$. The four matchings p_1, p_3 carry charge -1 and p_2, p_4 charge $+1$ under $Q_t = (-1, 1, -1, 1)$ of (VII.5.12). Weighting each matching by t and

its charge by a fugacity u, and projecting onto the neutral (gauge-invariant) sector by integrating u around the unit circle,

$$H_{\mathcal{C}}(t) = \oint_{|u|=1} \frac{du}{2\pi i\, u} \frac{1}{(1-t\,u)^2\,(1-t\,u^{-1})^2} = \sum_{n\geq 0} (n+1)^2\, t^{2n}, \tag{VII.5.29}$$

the two charge-$(+1)$ matchings contributing $(1-tu)^{-2}$, the two charge-(-1) matchings $(1-tu^{-1})^{-2}$, and the contour extracting the u^0 term. That term is the diagonal of the two squared geometric series, $\sum_{a,b}(a{+}1)(b{+}1)\,t^{a+b}\,u^{a-b}\big|_{a=b} = \sum_a (a{+}1)^2 t^{2a}$, reproducing (VII.5.28) and hence (VII.5.27), now derived from the charge Q_t rather than from counting generators. The Molien integral is the GLSM-side route to the same Hilbert series the ring-theoretic count gave.

The plethystic exponential and logarithm. The Hilbert series of a freely generated ring is built from its generators by the **plethystic exponential**, and the generators and relations of a non-free ring are read back out by its inverse, the **plethystic logarithm**. For a formal series $f(t)$ with $f(0) = 0$, including a generator series,

$$\mathrm{PE}[f](t) = \exp\left(\sum_{k\geq 1} \frac{f(t^k)}{k}\right) \tag{VII.5.30}$$

defines a formal series. When f counts actual independent generators with nonnegative integer multiplicities, this is the Hilbert series of their symmetric algebra. Signed f also occur, for example to subtract a regular sequence of relations. Its inverse on series with $H(0) = 1$ is the plethystic logarithm,

$$\mathrm{PL}[H](t) = \sum_{k\geq 1} \frac{\mu(k)}{k}\, \log H(t^k), \tag{VII.5.31}$$

with μ the Möbius function. For a complete intersection with a homogeneous regular sequence, it records generator degrees minus relation degrees. More general rings require care: higher coefficients combine relations among relations and cannot be read term by term as minimal generators or syzygies. For the conifold, in the degree-one grading,

$$\mathrm{PL}\big[H_{\mathcal{C}}\big](t) = 4t - t^2, \tag{VII.5.32}$$

the four positive units the four generators x, y, z, w of degree one, and the single negative term $-t^2$ the one quadratic relation $xy - zw$ at degree two. Applying the exponential to this finite expression reproduces the full series, $\mathrm{PE}[4t - t^2] = H_{\mathcal{C}}$, by the formal identity $\mathrm{PE}[at^d] = (1-t^d)^{-a}$. This establishes the full rational identity; agreement of a finite Taylor expansion alone would not do so. The contrast with the free ring is instructive: $\mathbb{C}^3$ has the freely generated mesonic ring,

$$H_{\mathbb{C}^3}(t) = \frac{1}{(1-t)^3}, \qquad \mathrm{PL}\big[H_{\mathbb{C}^3}\big](t) = 3t, \tag{VII.5.33}$$

three generators and *no* negative term, the signature of a complete intersection with no relation. Here the polynomial-ring structure is already known from the quotient; it is

not inferred from finitely many vanishing coefficients. For a complete intersection the plethystic logarithm terminates and is read literally as generators minus relations, as for $\mathbb{C}^3$ and the conifold here; for a non-complete-intersection ring (such as dP_0 below) later coefficients combine higher relation data and are not individually graded Betti numbers. A Hilbert series alone is not a presentation of the ring.

Three uses of the plethystic exponential, kept apart. The plethystic exponential is used in three distinct ways in this subject, and conflating them is a common source of error. First, $H = \mathrm{PE}[\mathrm{PL}[H]]$ reconstructs the single-regular-brane coordinate-ring series. Second, exponentiating a positive generator series builds its symmetric algebra; relations require additional information. These are formal operations on graded rings. Third, it appears in multi-brane and multi-trace counting, where the single-trace operators are exponentiated to the full spectrum at large N; that use carries the symmetric-product and large-N qualifications and is only named here, not developed. The first two are what this chapter uses; the third belongs with the dynamics.

VII.5.8 The worked bridge: $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ end-to-end

The chapters of Volume VI built the quiver of $\mathbb{C}^3/\mathbb{Z}_3$ from the representation theory of the finite group $\mathbb{Z}_3$, by the McKay correspondence. The same singularity, the cone over $\mathbb{P}^2$ written as dP_0 (the zeroth del Pezzo, the local $\mathbb{P}^2$), at its affine contraction is a toric Calabi–Yau threefold, and so its quiver and geometry can equally be computed from its dimer by the forward algorithm. Running both and matching them is the bridge between the two volumes: the McKay quiver and the dimer-computed quiver are the same quiver, and the toric diagram the dimer returns is the local-$\mathbb{P}^2$ polygon. This section does the dimer computation end-to-end and checks it against the corresponding Volume VI data.

The dimer and the quiver. The chosen embedded dP_0 tiling has three white and three black vertices, every white joined to every black, so nine edges. Its imported face-incidence data specify three gauge nodes and identify the nine edges as bifundamentals. With those data, the quiver arrows run cyclically, three from each node to the next,

$$a_{ij} = 3\,\delta_{j,\,i+1} \pmod 3, \qquad \begin{pmatrix} 0 & 3 & 0 \\ 0 & 0 & 3 \\ 3 & 0 & 0 \end{pmatrix}, \tag{VII.5.34}$$

the adjacency matrix with three arrows $0 \to 1$, three $1 \to 2$, three $2 \to 0$. This is exactly the $\mathbb{C}^3/\mathbb{Z}_3$ McKay quiver of Volume VI, three nodes of equal rank with $3\,\delta$ chiral arrows, recovered here from the chosen embedded tiling and its face incidence. The three black and three white trivalent vertices give six signed cubic monomials. With the imported black/white superpotential dictionary, these become the six trace terms of the determinantal cubic W inherited from the $\mathcal{N}{=}4$ parent.

The perfect matchings. The complete bipartite graph on $3+3$ vertices has six perfect matchings, one for each way of pairing the three white vertices bijectively with the

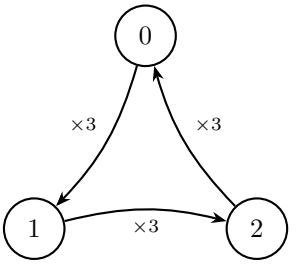


Figure VII.5.3: The $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ quiver, read from the chosen embedded tiling and its imported face-incidence data. Three gauge nodes $\prod_i U(N_i)$ of equal rank; three chiral bifundamentals $0 \to 1$, three $1 \to 2$, three $2 \to 0$, the adjacency $a_{ij} = 3\,\delta_{j,i+1}$ of (VII.5.34) (nodes indexed $0, 1, 2$ as in Volume VI). Read it as the node interiors carrying the gauge factors and each labelled arrow as a triple of fields. This is the Volume VI McKay quiver for this embedding.

three black, that is the six permutations of the symmetric group S_3. Each is a triple of edges, one per white vertex, covering all three black vertices,

$$\begin{aligned} &p_1 = \{e_{00}, e_{11}, e_{22}\}, \quad p_2 = \{e_{00}, e_{12}, e_{21}\}, \quad p_3 = \{e_{01}, e_{10}, e_{22}\}, \\ &p_4 = \{e_{01}, e_{12}, e_{20}\}, \quad p_5 = \{e_{02}, e_{10}, e_{21}\}, \quad p_6 = \{e_{02}, e_{11}, e_{20}\}, \end{aligned} \qquad c = 6, \quad \text{(VII.5.35)}$$

where e_{ij} is the edge from white vertex i to black vertex j. With $E = 9$ fields and $c = 6$ matchings, the P-matrix is 9×6,

$$P = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \end{pmatrix}, \qquad \operatorname{rank} P = 5 = G + 2, \qquad \text{(VII.5.36)}$$

its rows the nine fields $e_{00}, e_{01}, e_{02}, e_{10}, \ldots$ and its columns the six matchings; its rank is $G + 2 = 5$, as the algorithm requires.

The charge matrices. The kernel of P is one-dimensional, $c - G - 2 = 6 - 5 = 1$, so Q_F is a single row,

$$Q_F = \begin{pmatrix} -1 & 1 & 1 & -1 & -1 & 1 \end{pmatrix}, \qquad \text{(VII.5.37)}$$

a primitive generator of the redundant charge lattice; one verifies $P\,Q_F^{\mathsf{T}} = 0$ and that its entries have gcd one. This checks the whole rank-one kernel, not a polynomial equation imposed on the six matchings. The reduced incidence has $G - 1 = 2$ rows, and solving $\delta = Q_D P^{\mathsf{T}}$ returns the two D-term charge rows

$$Q_D = \begin{pmatrix} 0 & -1 & 1 & 0 & 0 & 0 \\ 1 & -1 & -2 & 1 & 1 & 0 \end{pmatrix}. \qquad \text{(VII.5.38)}$$

This is one convenient integral representative: unimodular gauge-row operations (with the same change to δ), sign flips, and integer shifts by the Q_F row give equivalent Q_D's presenting the same quotient, so a later computation that picks a different basis is not in conflict. Stacking gives the total charge matrix, of $c-3=3$ rows,

$$Q_t = \begin{pmatrix} -1 & 1 & 1 & -1 & -1 & 1 \\ 0 & -1 & 1 & 0 & 0 & 0 \\ 1 & -1 & -2 & 1 & 1 & 0 \end{pmatrix}, \tag{VII.5.39}$$

and its integer kernel is three-dimensional, the toric diagram. Reading off the kernel, the six matchings sit at the lattice points

$$(0,0),\ (1,1),\ (1,1),\ (1,2),\ (2,1),\ (1,1), \tag{VII.5.40}$$

so that three of the matchings, p_2, p_3, p_6, land on the single interior point, which therefore carries multiplicity three, while the three corner points $(0,0),(1,2),(2,1)$ are each hit once. The distinct lattice points are

$$\big\{(0,0),\ (1,2),\ (2,1)\big\} \text{ corners}, \qquad (1,1) \text{ interior, multiplicity } 3, \tag{VII.5.41}$$

four points in all, three corners and one interior, the local-$\mathbb{P}^2$ toric diagram. Assembling these six columns with a height-one row on top into the 3×6 matrix G_t, one verifies the kernel relation $Q_t\, G_t^{\mathsf{T}} = 0$ directly: every charge row annihilates the lattice points. That one line is the payoff of the forward algorithm, the toric geometry read out of the gauge theory.

The toric diagram, and what it is not. Up to a $GL(2,\mathbb{Z})$ transformation and a translation, (VII.5.41) is the standard local-$\mathbb{P}^2$ polygon with vertices $(1,0),(0,1),(-1,-1)$ and the single interior point $(0,0)$,

$$\Delta_{dP_0} = \mathrm{conv}\{(1,0),(0,1),(-1,-1)\}, \qquad \text{interior point } (0,0), \tag{VII.5.42}$$

twice its area equal to three (so $2\,\mathrm{Area} = 3 = G$, the three nodes) and exactly four lattice points. It is emphatically *not* the side-length-three standard triangle $\mathrm{conv}\{(0,0),(3,0),(0,3)\}$: that triangle has area $9/2$, hence $2\,\mathrm{Area} = 9$, and ten lattice points (six extra boundary points), so it is not $GL(2,\mathbb{Z})$-equivalent to the local-$\mathbb{P}^2$ diagram, whose area is $3/2$ and whose $2\,\mathrm{Area} = 3$ gives the three gauge nodes. The correct dP_0 diagram is the minimal one with three vertices and one interior point. The same polygon is the support of the Kasteleyn determinant (VII.5.18), whose interior coefficient -3 records the multiplicity three of the interior point, the two independent computations agreeing.

The baryonic count, and a warning. The dP_0 run is also the cleanest illustration of the $d-3$ versus $G-1$ distinction. The toric diagram has $d=3$ external (boundary) lattice points, the three corners, with no boundary point between them, so

$$\#\{\text{non-anomalous } U(1)_B\} = d-3 = 0, \qquad \text{whereas} \qquad G-1 = 2. \tag{VII.5.43}$$

dP_0 has *no continuous non-anomalous* baryonic global symmetry, despite having three gauge nodes and two D-term charge rows in Q_D. This is the difference from the conifold,

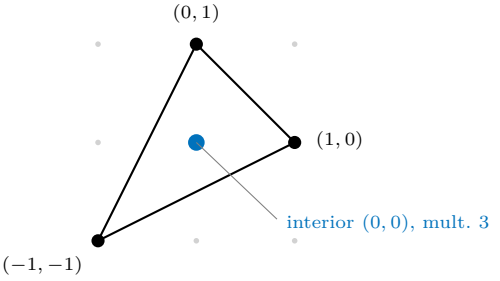


Figure VII.5.4: The $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ toric diagram, the local-$\mathbb{P}^2$ polygon. Read the three filled vertices as the corner lattice points (each the image of one perfect matching) and the central point as the interior point, here marked $\times 3$ because three of the six perfect matchings project to it (multiplicity three, equal to the absolute value of the interior coefficient -3 of $\det K$). Four lattice points, twice-area three. It is not the side-length-three triangle, which would carry six spurious boundary points.

where the two counts coincided at one; the bridge must not claim a baryonic $U(1)$ for dP_0. The master space is still five-dimensional, $G + 2 = 5$, three mesonic plus two baryonic directions in the gauge-linear-sigma model, but the two baryonic directions are not conserved global symmetries of the fixed point.

The dP_0 Hilbert series. The run is completed by counting the mesonic operators. Grading by the anticanonical degree of local $\mathbb{P}^2$ (the degree of $\mathcal{O}_{\mathbb{P}^2}(3)$, equivalently the $\mathbb{Z}_3$-invariant $\mathbb{C}^3$ monomials of degree $3n$), the mesonic Hilbert series is the rational function

$$H_{dP_0}(t) = \sum_{n\geq 0} h^0\big(\mathbb{P}^2, \mathcal{O}(3n)\big)\, t^n = \sum_{n\geq 0} \binom{3n+2}{2} t^n = \frac{1 + 7t + t^2}{(1-t)^3}, \tag{VII.5.44}$$

whose plethystic logarithm

$$\mathrm{PL}\big[H_{dP_0}\big] = 10\,t - 27\,t^2 + 105\,t^3 - \cdots \tag{VII.5.45}$$

reads off ten degree-one generators, the ten sections of $\mathcal{O}(3)$ (the $\mathbb{Z}_3$-invariant cubics $z_1^i z_2^j z_3^k$, $i + j + k = 3$), and, from the $-27\,t^2$, twenty-seven independent quadratic relations among them. Indeed, there are $\binom{11}{2} = 55$ quadratic products of the ten generators, whereas the degree-two piece has $\binom{8}{2} = 28$ monomials. Their difference is 27. Since there are no linear relations, these quadrics are minimal. The ring has codimension $10 - 3 = 7$, so it cannot be a complete intersection. This dimension argument identifies the first relations without treating the later PL coefficients as a minimal resolution. As for the conifold, the grading must be declared (here the anticanonical degree); grading instead by the elementary $\mathbb{C}^3$ coordinates keeps only the degrees divisible by three, the same ring counted with a rescaled fugacity.

What the bridge establishes. The dimer of dP_0, run through the forward algorithm with no group theory, has reproduced the Volume VI McKay quiver (VII.5.34) and

the $\mathbb{C}^3/\mathbb{Z}_3$ toric diagram (VII.5.42), the same singularity reached in Volume VI by the representation theory of $\mathbb{Z}_3$. For a finite diagonal abelian subgroup $\Gamma \subset SU(3)$, the quotient $\mathbb{C}^3/\Gamma$ is a toric Calabi–Yau threefold. In the present example, the McKay construction and the dimer construction are two computations of one geometry, and they agree. The matching is the bridge from the orbifold volume to the tiling volume, providing a common geometric description of the two constructions.

VII.5.9 What has been computed, and what is deferred

For the connected, nondegenerate, consistent toric phases and one regular brane considered here, the forward algorithm has produced the zero-FI abelian geometry. From a brane tiling with E edges, G faces, and c perfect matchings, the perfect matchings were enumerated as the homogeneous coordinates of the gauge linear sigma model; the P-matrix (VII.5.2) expressed the fields in them with rank $G+2$; the F-term charge matrix $Q_F = \ker_{\mathbb{Z}}(P)$ and the D-term charge matrix Q_D fixed by $\delta = Q_D P^{\mathsf{T}}$ were stacked into Q_t, and the toric diagram came out as the three-dimensional kernel $\ker_{\mathbb{Z}}(Q_t)$, agreeing with the Newton polygon of $\det K$. The map from matchings to lattice points was shown to be many-to-one, with multiplicities read both by enumeration and from the Kasteleyn coefficients. The Q_F quotient described the coherent F-flat component of dimension $G+2$; the additional abelian quotient left three mesonic dimensions by removing generic gauge orbits of dimension $G-1$, with the physical non-anomalous baryonic symmetries counted by $d-3$ and not $G-1$. The Hilbert series was defined with its grading declared, the conifold and $\mathbb{C}^3$ series computed, and the plethystic exponential and logarithm used to read off generators and relations. Finally $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ was run end-to-end and matched to the Volume VI McKay quiver and the local-$\mathbb{P}^2$ toric diagram.

A reader leaving the chapter should be able to do the following without looking ahead: enumerate the perfect matchings of a dimer and write its P-matrix; compute Q_F as the integer kernel of P, Q_D from the reduced incidence by $\delta = Q_D P^{\mathsf{T}}$, and Q_t by stacking; read the toric diagram off $\ker_{\mathbb{Z}}(Q_t)$ and check it against the Newton polygon of $\det K$; explain why the matching-to-point map is many-to-one and read a point's multiplicity from the Kasteleyn coefficient; distinguish the full F-flat master space from its coherent component of dimension $G+2$, and form the latter's abelian mesonic quotient; count the non-anomalous baryonic symmetries as $d-3$ and not $G-1$; write the conifold Hilbert series in a declared grading and read its generators and relation from the plethystic logarithm; and run dP_0 from its dimer to its toric diagram, recovering the Volume VI quiver.

What has been deferred, and to where, is the map of the rest of the volume and of Part I.

- The Hilbert series here is graded by an elementary counting degree, not by the physical superconformal R-charge. The principle that selects the R-charge, a-maximization, and with it the true R-graded Hilbert series and the central charge, is the next chapter; this chapter supplies the perfect-matching and global-symmetry data that the trial R-charges of that computation are parametrized by,

and names the link.

- The inverse direction, from a toric diagram to a quiver, and the non-uniqueness of that inverse (toric duality), are the general toric chapter and the toric-duality chapter; this chapter establishes the invariants, the toric diagram and the mesonic Hilbert series, that the duality chapter matches across Seiberg-dual phases.
- Only the classical moduli space was computed. The baryonic-branch dynamics, the role of the baryon operators in confinement and in the duality cascade, route to the cascade chapter and to Part III.
- The dimer itself, its Kasteleyn matrix, its sign rule, and its zig-zag consistency filters were consumed, not re-derived, from the tiling chapter; the physical NS5–D5 origin of the tiling is the brane-configuration chapter.
- The identification of the integral kernel data of Q_t with a presentation of the toric variety, the gauge-linear-sigma-model as a symplectic or geometric-invariant-theory quotient equal to the toric variety from its fan, is cited to Part I (toric geometry), not proved; this chapter constructs the diagram and dimension-checks the mesonic moduli space. The Hilbert-series and plethystic commutative algebra, the graded-ring theory, the Molien–Weyl integral, and the syzygy structure the plethystic logarithm reads off, are likewise used as tools and cited to Part I (Hilbert-series methods).
- The leading identification of the mesonic moduli space with the toric threefold carries the α' and g_s probe caveats of Volume VI; the chapter dimension-checks the mesonic branch, it does not assert the metric. The forward algorithm here is for one regular brane in the toric phase; higher-rank and non-toric cases are named and routed.

The single lesson to carry forward is the one in the opening keybox: the dimer organizes the computation, and the forward algorithm turns it into geometry, the count of perfect matchings into a toric diagram with multiplicities, and the gauge theory's vacua into a master space whose coherent component has the threefold as its abelian mesonic quotient. The next chapter grades that data by the physical R-charge and extracts the central charge.

Status and sources

The evidence-code tags ([WS] controlled physical-worldsheet derivation in the untwisted string, [MATH] mathematical statement with explicit hypotheses and proof basis), the convention pointers, and the source anchors are collected here, keyed by section. The chapter develops the forward algorithm and the moduli-space machinery of the brane tiling, after the dimer corpus of Hanany and Kennaway (2005) and Franco et al. (2006), the perfect-matching and amoeba treatment of Feng et al. (2008), the Hilbert-series and plethystic counting of Benvenuti et al. (2007) and Butti et al. (2007), and the master space of Forcella et al. (2008); it reuses the orbifold quiver of Hanany and He (1999) at the dP_0 bridge, the conifold gauge theory of Klebanov and Witten (1998), and the textbook and review treatments of Johnson (2003) and Yamazaki (2008). It reuses the dimer/brane-tiling conventions (bipartite tiling conventions, the perfect-matching → lattice-point-with-multiplicity rule and the height orientation),

the toric-data conventions (toric lattice orientation, the toric diagram as the kernel of the charge matrix), the gauge/rank data (quiver gauge groups and ranks, $\prod_i U(N_i)$, the diagonal $U(1)$ removed), the Fayet–Iliopoulos/resolution chambers, the superpotential store (superpotential normalization, the F-terms Q_F encodes), the McKay multiplicity (McKay multiplicities, only at the dP_0 bridge), the arrow convention, and the supercharge count. The forward-algorithm matrix names (P, Q_F, Q_D, Q_t), the Hilbert-series grading, and the plethystic notation are local conventions; they are used here as the literature does. The finite calculations use the stated toric and dimer data.

§VII.5.1 From a tiling to a geometry. Framing: the dimer as a quiver plus superpotential, the forward algorithm as the reverse of the general toric chapter, the two corrections (multiplicities, master space), and the reader contract.

§VII.5.2 Perfect matchings and the P-matrix. The perfect matchings as gauge-linear-sigma-model coordinates (VII.5.1), the P-matrix $X_a = \prod_\alpha p_\alpha^{P_{a\alpha}}$ (VII.5.2), and the rank $\operatorname{rank} P = G + 2$ (VII.5.3). physical classification unassigned (bipartite tiling conventions, superpotential normalization; Franco et al. (2006), Hanany and Kennaway (2005)). P is $E \times c$ with $\operatorname{rank} P = G + 2$ on all three fixtures.

§VII.5.3 The charge matrices and the toric diagram. The F-term charges $Q_F = \ker_{\mathbb{Z}}(P)$ of shape $(c - G - 2) \times c$ (VII.5.5), the D-term charges Q_D of shape $(G - 1) \times c$ fixed by $\delta = Q_D P^{\mathsf{T}}$ (VII.5.6), the stacked Q_t of shape $(c - 3) \times c$ (VII.5.7), and the toric diagram as the three-dimensional kernel $G_t = \ker_{\mathbb{Z}}(Q_t)$ (VII.5.8), with the six-line recipe. physical classification unassigned; bounded algebra and cited toric premise (toric lattice orientation, quiver gauge groups and ranks, resolution and discrete data, fundamental-at-tail arrows; Franco et al. (2006), Feng et al. (2008)). The sourced toric interpretation of the integral kernel data is toric geometry (home Part I), cited. The displayed calculations give kernel inclusion $PQ_F^{\mathsf{T}} = 0$ and rational rank, not general integral saturation; $\delta = Q_D P^{\mathsf{T}}$; the shapes $(c - G - 2)$, $(G - 1)$, $(c - 3)$; Q_t annihilates the all-ones height row and the perfect-matching positions; $\dim \ker Q_t = 3$.

§VII.5.4 The conifold, worked through. The four single-edge matchings (VII.5.9), the identity P-matrix (VII.5.10), the empty Q_F, the single baryonic $Q_D = Q_t = (-1, 1, -1, 1)$ (VII.5.11)–(VII.5.12), and the kernel giving the unit-square toric diagram (VII.5.13). physical classification unassigned (toric lattice orientation, bipartite tiling conventions; Klebanov and Witten (1998), Franco et al. (2006)). the conifold forward run validates with $\operatorname{rank} P = 4 = G + 2$, Q_F empty, Q_D one row, $\dim \ker Q_t = 3$, and the kernel matches the entry-7 unit square up to $GL(2, \mathbb{Z})$.

§VII.5.5 Perfect matchings carry multiplicities. The many-to-one map and the inequality #pm $\geq$ #points (VII.5.17), the two independent computations of the multiplicity (enumeration and the Kasteleyn coefficient), and the dP_0 example with the interior coefficient -3 (VII.5.18). physical classification unassigned; finite combinatorics (bipartite tiling conventions; Hanany and Kennaway (2005), Feng et al. (2008)). Files (perfect matchings are not in bijection with lattice points; they map with multiplicities). $\mathbb{C}^3$ and the conifold are the clean case (#pm = #points, all multiplicity one); dP_0 has six matchings, four points, the interior point multiplicity three; the perfect-matching multiplicity multiset equals the Kasteleyn $|\det K|$ coefficient multiset on all three fixtures.

§VII.5.6 The master space and the mesonic–baryonic split. The full master space is F-flat before the gauge quotient; its coherent component has dimension $G + 2 = 3 + (G - 1)$ in (VII.5.19); the abelian mesonic moduli space as the effective $(\mathbb{C}^*)^{G-1}$ quotient of dimension three (VII.5.22); and the $d - 3$ versus $G - 1$ distinction (VII.5.24) (conifold $d - 3 = G - 1 = 1$; dP_0 $d - 3 = 0 \neq G - 1 = 2$; $\mathbb{C}^3$ none). physical classification unassigned (quiver gauge groups and ranks, resolution and discrete data; Forcella et al. (2008), Butti et al. (2007)). Files (the gauge group is $\prod_i U(N_i)$ not bare SU, and the non-anomalous baryonic count is $d - 3$, shared with the conifold chapter) and (the F-flat parent and its mesonic quotient must be distinguished); carries (the probe caveats). The baryonic-branch dynamics route to the cascade chapter and Part III. master dimension $c - \operatorname{rank} Q_F = G + 2$ ($\mathbb{C}^3$ 3, conifold 4, dP_0 5); mesonic $c - \operatorname{rank} Q_t = 3$, strictly smaller for $G > 1$; the baryonic count $d - 3$ (conifold $1 = G - 1$, dP_0 $0 \neq G - 1 = 2$, $\mathbb{C}^3$ none).

§VII.5.7 The Hilbert series and the plethystics. The Hilbert series (VII.5.25); the declared conifold grading, $(1 - t^2)/(1 - t)^4$ generators degree one and $(1 - t^4)/(1 - t^2)^4$ mesons degree two (VII.5.26)–(VII.5.27), the same ring with coefficients $(n + 1)^2$ (VII.5.28); the plethystic exponential and logarithm (VII.5.30)–(VII.5.31), the conifold PL $= 4t - t^2$ (VII.5.32) and the free $\mathbb{C}^3$ PL $= 3t$ (VII.5.33); and

the three uses of the plethystic exponential kept apart. exact graded-ring identities under the stated grading (superpotential normalization; Benvenuti et al. (2007), Butti et al. (2007)). The Hilbert-series/plethystic commutative algebra is Hilbert-series methods (home Part I), cited; the grading and plethystic notation are local conventions. The a-maximization that fixes the physical R-grading is the next chapter. the two gradings agree under $t \to t^2$; the pole at $t = 1$ has order three (Krull dimension three); the coefficients are $(n+1)^2$; $\mathrm{PL}[H_{\mathcal{C}}] = 4t - t^2$ and $\mathrm{PE}[4t - t^2]$ round-trips; $\mathbb{C}^3$ gives $\mathrm{PL} = 3t$ with no relation.

§VII.5.8 The worked bridge: $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$. The complete bipartite dimer and the recovered McKay quiver $a_{ij} = 3\,\delta_{j,i+1}$ (VII.5.34); the six S_3 perfect matchings (VII.5.35), the 9×6 P-matrix (VII.5.36), the single-row Q_F (VII.5.37), the two-row Q_D (VII.5.38), the stacked Q_t (VII.5.39); the kernel points (VII.5.40) with the interior multiplicity three, the local-$\mathbb{P}^2$ diagram (VII.5.42) (not the side-three triangle), and the $d - 3 = 0$ baryonic count (VII.5.43). physical classification unassigned (McKay multiplicities, bipartite tiling conventions; Hanany and He (1999), Hanany and Kennaway (2005)). The target agrees with the Volume VI McKay projection. the dP_0 dimer's quiver adjacency equals the frozen Volume VI McKay quiver $a_{ij} = 3\,\delta_{j,i+1}$ ($G = 3$, nine fields); the forward algorithm reproduces the local-$\mathbb{P}^2$ toric data (four lattice points, twice-area three) up to $GL(2,\mathbb{Z})$ and translation, and not the side-three triangle. The dimer's perfect matchings and windings were selected jointly so that $\det K$ is the local $\mathbb{P}^2$ and Q_t annihilates the matching positions (an imported algebraically consistent fixture; the search itself is not an independent physical derivation).

§VII.5.9 What is and is not given. Scope and limitations. the chapter builds the forward algorithm, the charge matrices and the toric diagram as a kernel, the multiplicities, the master space and its mesonic/baryonic split with the $d-3$ count, the Hilbert series with declared grading and the plethystics, and the dP_0 bridge; it routes a-maximization and the physical R-grading to the next chapter, the inverse algorithm to the general toric chapter, toric duality to the toric-duality chapter, the baryonic-branch dynamics and the cascade to the cascade chapter and Part III, the dimer machinery to the tiling and brane-configuration chapters, and the toric/Hilbert commutative algebra to Part I (toric geometry, Hilbert-series methods).

Calculation scope and exercises. The physical primary classifications remain unassigned where the physical derivation is absent. Worldvolume and polynomial algebra leave open [WS]. The full master space, its coherent component and the zero-FI abelian quotient are distinguished throughout this candidate; the component and matching descriptions are sourced to Forcella et al. (2008), Sections 2.1.3, 2.1.6 and 2.1.7, especially equations (2.27)–(2.28) and (2.34). Saturation in (VII.5.20) follows by localization and contraction; the component identification still uses the stated toric hypotheses. The companion has three exercise families: enumerate the fixed dP_0 graph's matchings and multiplicities; submit its P, Q_F, Q_D, Q_t, G_t with integral basis choices; and compute the conifold quotient ideal and the exact conifold and dP_0 Hilbert series in named gradings. The conifold relation is recomputed by elimination, and rational functions are compared exactly. Finite expansions through orders 8 or 10 leave all-degree termination open. The general saturation calculation, full primary decomposition, nonabelian moduli and physical string derivation remain outside these exercises.

Related constructions. The forward algorithm takes dimer, quiver and superpotential data to the linear sigma model matrices P, Q_F, Q_D and Q_t, then obtains the toric diagram from perfect matchings. The master space, mesonic cone and $d-3$ nonanomalous baryonic directions have distinct roles in Hilbert-series counting. The worked $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ example gives the same target as the McKay-projection construction in Volume VI.

References

Benvenuti, S. et al. (2007). "Counting BPS operators in gauge theories: Quivers, syzygies and plethystics". In: *JHEP* 11, p. 050. DOI: 10.1088/1126-6708/2007/11/050. arXiv: hep-th/0608050.

Butti, A. et al. (2007). “Counting chiral operators in quiver gauge theories”. In: *JHEP* 11, p. 092. DOI: 10.1088/1126-6708/2007/11/092. arXiv: 0705.2771 [hep-th].

Feng, B. et al. (2008). “Dimer models from mirror symmetry and quivering amoebae”. In: *Adv. Theor. Math. Phys.* 12, pp. 489–545. DOI: 10.4310/ATMP.2008.v12.n3.a2. arXiv: hep-th/0511287.

Forcella, D. et al. (2008). “The Master Space of N=1 Gauge Theories”. In: *JHEP* 08, p. 012. DOI: 10.1088/1126-6708/2008/08/012. arXiv: 0801.1585 [hep-th].

Franco, S. et al. (2006). “Brane dimers and quiver gauge theories”. In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Hanany, A. and Y.-H. He (1999). “NonAbelian finite gauge theories”. In: *JHEP* 02, p. 013. DOI: 10.1088/1126-6708/1999/02/013. arXiv: hep-th/9811183.

Hanany, A. and K. D. Kennaway (2005). “Dimer models and toric diagrams”. In: arXiv: hep-th/0503149.

Johnson, C. V. (2003). *D-Branes*. Cambridge Monographs on Mathematical Physics. Cambridge University Press. DOI: 10.1017/CBO9780511606540.

Klebanov, I. R. and E. Witten (1998). “Superconformal field theory on three-branes at a Calabi-Yau singularity”. In: *Nucl. Phys. B* 536, pp. 199–218. DOI: 10.1016/S0550-3213(98)00654-3. arXiv: hep-th/9807080.

Yamazaki, M. (2008). “Brane Tilings and Their Applications”. In: *Fortsch. Phys.* 56, pp. 555–686. DOI: 10.1002/prop.200810536. arXiv: 0803.4474 [hep-th].

Chapter VII.6

Volume minimization and a-maximization

VII.6.1 Two extremization problems and one number

Chapters 1–5 built the construction. A stack of D3-branes was placed at the tip of a toric Calabi–Yau threefold, the quiver and its superpotential were read from the dimer (Chapters 2 and 3), the brane tiling was given its physical origin as a configuration of NS5 and D5 branes (Chapter 4), and the perfect-matching, moduli-space, and Hilbert-series machinery was assembled (Chapter 5). At the end of all that the theory has a name, a quiver, a superpotential, a moduli space, and a global symmetry, but it does not yet have a single number that a physicist would call dynamical. This chapter computes the first one: the conformal central charge a of the infrared superconformal field theory, the quantity that counts its degrees of freedom and that decreases along unitary four-dimensional renormalization-group flows between conformal fixed points.

The central charge is worth this chapter for two reasons. It is the cleanest invariant the construction can compute on its own, without leaving the world of toric data and 't Hooft anomalies, and it is the place where the gauge theory and the dual geometry meet most sharply. On the gauge side, a is fixed by a field-theory extremization principle: a-maximization, due to Intriligator and Wecht. Among all the candidate R-symmetries consistent with the anomalies, the true superconformal one *locally maximizes* a certain cubic in the R-charges. On the geometry side, the same number is fixed by a dual extremization: volume minimization, due to Martelli, Sparks, and Yau. The Sasaki–Einstein metric on the link is the one whose Reeb vector *minimizes* a volume functional, and a is inversely proportional to that minimal volume. Butti and Zaffaroni related the two problems after extremizing the baryonic mixing directions, for the toric charge and multiplicity prescription used here. The conclusion is the slogan of the volume: *the geometry computes the central charge.*

Take-home message of this chapter

For regular-brane toric quivers with an interacting fixed point and no omitted accidental symmetries, the leading large-N central charge is compared on the gauge and geometry sides. Gauge side: maximize the trial a-function locally, $a(R) = \frac{3}{32}(3\,\mathrm{Tr}\,R^3 - \mathrm{Tr}\,R)$ over the R-symmetries the anomalies allow (a-maximization). Geometry side: minimize the Martelli–Sparks–Yau volume function $Z(b)$ over the Reeb vector b, and read $a = \pi^3 N^2/(4\,\mathrm{Vol}(Y))$. The toric dictionary relates the two extrema after baryonic extremization. We compute both sides for the conifold and dP_0; for $Y^{2,1}$ we compute the geometry and import the gauge-side result.

A word on what this chapter does and does not establish, because the boundary is sharp and is part of the design of the volume. It *states* the Intriligator–Wecht principle (that the superconformal R is the one maximizing a) and the Martelli–Sparks–Yau principle (that the Sasaki–Einstein Reeb vector minimizes the volume), and it *works* the resulting finite-dimensional extremization on toric examples. It does *not* prove either principle. The field-theory derivation that a is the maximized quantity, the a-theorem, and the unitarity-bound subtleties belong to the superconformal-dynamics material of Part III. The existence of the Sasaki–Einstein metric realizing the minimizing Reeb vector belongs to the geometry of Part I. What is load-bearing here is the calculation, not a proof of the theorems, and the worked toric matches are the evidence. Anything beyond the leading central charge is Part III: the Kaluza–Klein spectrum, the subleading corrections in N, and the operator-by-operator holographic dictionary.

The plan is to build each side as a self-contained calculation and then collide them. First we assign trial R-charges from the dimer data and impose the two linear constraints they must satisfy (§VII.6.2). We state the a-maximization principle and write the cubic that is extremized (§VII.6.3). We give the Butti–Zaffaroni toric parametrization, which turns the constrained maximization into a clean problem on the corners of the toric diagram (§VII.6.4). We write the Martelli–Sparks–Yau volume function from the toric legs and state volume minimization (§VII.6.5). We identify the two extremizations as one and record $a \propto 1/\,\mathrm{Vol}(Y)$ (§VII.6.6). Then we run all of it on three examples: the conifold (§VII.6.7), $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ (§VII.6.8), and the irrational $Y^{2,1}$ (§VII.6.9); the chapter closes with what was and was not done (§VII.6.10). The conifold completes the preview left open in Chapter 1. dP_0 makes the Volume VI orbifold-inheritance relation geometric. And $Y^{2,1}$ is the historic test: the central charge is an irrational number, and the geometry reproduces it exactly.

VII.6.2 Trial R-charges from the dimer data

Everything begins with the superconformal R-symmetry. In a four-dimensional $\mathcal{N}$=1 superconformal theory the R-symmetry sits inside the superconformal algebra, and it fixes the dimensions of all chiral primaries through $\Delta = \frac{3}{2}R$ and the central charges through its 't Hooft anomalies. The construction does not hand us the superconformal R directly; it hands us a whole family of *candidate* R-symmetries, and the work of the chapter is to select the right one. This section sets up the family.

A trial R-symmetry assigns an R-charge R_a to each chiral field X_a of the quiver

(the gauginos always carry $R = 1$, fixed by supersymmetry). Not every assignment is admissible. Two linear conditions cut the candidates down to the physically allowed ones.

The superpotential constraint. The superpotential is an integral over half of superspace, $\int \mathrm{d}^2\theta\, W$, and $\mathrm{d}^2\theta$ carries R-charge -2, so W must carry R-charge $+2$ for the term to be R-invariant. Since W is a sum of gauge-invariant traces, this is a condition on each term: if a monomial of W is the product of fields $X_{a_1} \cdots X_{a_k}$, then

$$R(W\text{-term}) = \sum_{a \in \text{term}} R_a = 2. \tag{VII.6.1}$$

For the conifold quartic $\mathrm{Tr}(A_i B_k A_j B_l)$ this reads $2R(A) + 2R(B) = 2$; for the dP_0 cubic $\mathrm{Tr}(XYZ)$ it reads $R(X) + R(Y) + R(Z) = 2$. The constraint is the same linear equation that marginality of W imposed in Chapter 1.

The anomaly-free node constraint. A trial R-symmetry must also be free of the mixed R–gauge anomaly at every gauge node, or it cannot be a symmetry of the quantum theory. The mixed $U(1)_R\, SU(N_i)^2$ anomaly at node i is the sum over the fermions charged under that node of their R-charge times their Dynkin index. The gaugino is in the adjoint with R-charge 1 and index N_i; a bifundamental fermion is in the fundamental with index $\frac{1}{2}$, R-charge $R_a - 1$, and the partner from the far node supplies the multiplicity. For N regular branes, with all interacting nonabelian ranks equal to N, this is proportional to the Novikov–Shifman–Vainshtein–Zakharov (NSVZ) numerator at a conformal point, using $\gamma_a = 3R_a - 2$, and demanding it vanish gives the clean condition

$$2 + \sum_{a \text{ bifund. at node } i} \big(R_a - 1\big) = 0, \tag{VII.6.2}$$

the 2 being twice the gaugino contribution and the sum running over the bifundamentals incident to node i, each once at the endpoint it touches. The weights are index-weighted. A bifundamental enters with weight one. An adjoint chiral would enter with weight *two*, because its Dynkin index $T(\text{adj}) = N$ is twice the $T(\text{fund})\, N = \frac{1}{2}N$ that a bifundamental supplies (the fundamental index $\frac{1}{2}$ times the far-node multiplicity N). So the general node condition reads

$$2 + \sum_{a \text{ bifund.}} \big(R_a - 1\big) + 2 \sum_{\phi \text{ adj.}} \big(R_\phi - 1\big) = 0. \tag{VII.6.3}$$

For the adjoint-only $\mathbb{C}^3$ (the $\mathcal{N}$=4 node, three adjoints at $R = \frac{2}{3}$) this reads $2 + 2 \cdot 3\big(\frac{2}{3} - 1\big) = 0$, the adjoint weight 2 load-bearing. The toric quivers worked below carry only bifundamentals, so (VII.6.2) is the operative form: for a node with n_i incident bifundamentals all of the same R-charge R it collapses to the single linear equation $2 + n_i(R - 1) = 0$, i.e. $R = 1 - 2/n_i$.

Work the two fixtures. The conifold node carries four incident bifundamentals (A_1, A_2, B_1, B_2, two arrows in and two out), so $n_i = 4$ and the node condition is

$$2 + 4(R - 1) = 0 \implies R = \tfrac{1}{2}, \tag{VII.6.4}$$

while the superpotential is the quartic $\mathrm{Tr}(A_iB_kA_jB_l)$, four fields per term, so (VII.6.1) reads $4R = 2$ and gives the same $R = \frac{1}{2}$. At a dP_0 node there are six incident arrows in total, three incoming (from one neighbour) and three outgoing (to the other), so $n_i = 6$ ($N_f = 3N$) and

$$2 + 6(R - 1) = 0 \implies R = \tfrac{2}{3}, \tag{VII.6.5}$$

while the superpotential is the cubic $\mathrm{Tr}(XYZ)$, three fields per term, so (VII.6.1) reads $3R = 2$ and again gives $R = \frac{2}{3}$. In both cases the two linear conditions agree. This verifies the mixed R–gauge constraints on these symmetric assignments; cancellation of a pure gauge anomaly is a separate condition on the quiver representations. A general toric quiver need not admit one common R-charge solving both sets of equations.

What the constraints leave. The two conditions (VII.6.1) and (VII.6.2) are linear, so they leave an affine space of trial R-charges. The dimension of that space is the number of non-anomalous abelian flavor symmetries the R-charge can mix with. For a toric quiver this count is read off the toric diagram. For the isolated toric cones considered here, the maximal abelian subgroup of the continuous non-R global symmetry of the interacting theory has rank $d - 1$ and is

$$U(1)^{d-1} \quad \big(d = \#\{\text{external legs / corner perfect matchings of the toric diagram}\}\big), \tag{VII.6.6}$$

The count is $(d - 3) + 2$: $d - 3$ baryonic directions and two mesonic flavor Cartan directions. Enhanced nonabelian flavor groups can be present. This count is not obtained by removing a diagonal gauge $U(1)$ from d gauge nodes. The brane construction supplies $\prod_i U(N)$; its decoupled or massive abelian gauge factors are distinguished here from global baryonic symmetries of the interacting sector. A trial R-symmetry is the superconformal R plus an arbitrary mixture with these $U(1)^{d-1}$, so after imposing $R(W) = 2$ the trial space is $(d - 1)$-dimensional. For the conifold $d = 4$, leaving three parameters; for dP_0 a triangle has $d = 3$, leaving two. The superconformal R is one point in this space, and the next two sections say which point.

This is the seam with Chapter 5. The perfect matchings, the external legs / corners, and the non-anomalous global symmetry data are exactly what the moduli-space machine of Chapter 5 produced. There the symmetries graded the Hilbert series as a *counting* device; here the same data parametrizes the trial R, and a-maximization fixes the one R-charge that is the *superconformal* grading. The chapter does not recompute the matchings or the Hilbert series; it uses their global-symmetry output.

VII.6.3 The principle: a-maximization

Among the $(d - 1)$-parameter family of admissible trial R-symmetries, which one is the superconformal R of the infrared fixed point? The answer is the Intriligator–Wecht principle.

The trial central charge. To each trial R-symmetry assign the cubic 't Hooft-anomaly combination

$$a(R) = \frac{3}{32}\big(3\,\mathrm{Tr}\,R^3 - \mathrm{Tr}\,R\big), \tag{VII.6.7}$$

where the traces run over the Weyl fermions of the theory: each gaugino contributes with $R = 1$, and the fermion in a chiral field of scalar R-charge R_a contributes with $R = R_a - 1$. At the superconformal point $a(R)$ is the genuine conformal central charge, one of the two Weyl-anomaly coefficients of the stress tensor (the combination $3\,\mathrm{Tr}\,R^3 - \mathrm{Tr}\,R$ is the one that computes a; a different combination computes the second coefficient c). Off the superconformal point, evaluated on a wrong trial R, $a(R)$ is just a number built from anomalies, and that is what makes it useful as an objective function.

The principle. Assuming the full set of infrared abelian flavor currents is included, the exact superconformal R-symmetry is a local maximum of $a(R)$ on the non-anomalous trial space:

$$R_{\text{superconformal}} = \underset{R \text{ non-anomalous}}{\arg\,\text{local max}}\; a(R), \qquad \frac{\partial a}{\partial R}\Big|_{R_*} = 0, \quad \frac{\partial^2 a}{\partial R^2}\Big|_{R_*} < 0. \tag{VII.6.8}$$

The extremum is a local maximum: the second variation is negative definite on the constraint surface, which both selects the correct critical point when there are several and is the consequence of positivity of flavor-current two-point functions in a unitary SCFT. This is the precise content the chapter *states*. The proof that the maximizing combination is the superconformal R, and the separate a-theorem (that a decreases between fixed points along unitary flows), are field-theory results derived in Part III; here (VII.6.8) is the input, justified there and tested by the matches below.

Why a maximum, concretely. To see that $a(R)$ has the right shape, write it as a function of the trial R-charges. Using $R = R_a - 1$ for the chiral fermions, the chiral contribution to a is $\frac{3}{32}\sum_a \big[3(R_a-1)^3 - (R_a-1)\big]$, a sum of cubics, and the gauginos add a constant $\frac{3}{32}\cdot 2G$ for G gauge nodes (each adjoint fermion at $R = 1$ contributing $3\cdot 1 - 1 = 2$). Thus, at leading order in N and in units of N^2,

$$\frac{a(R)}{N^2} = \frac{3}{32}\Big[\,2G + \sum_a \big(3(R_a-1)^3 - (R_a-1)\big)\Big]. \tag{VII.6.9}$$

Restricted to the constraint surface it is a polynomial of degree at most three. Nondegeneracy must be checked at the proposed extremum; a generic cubic is not globally bounded above. To see the shape directly, take a theory with two field types of trial charges R_1, R_2 tied by one superpotential relation, say $R_1 + R_2 = 1$ so $R_2 = 1 - R_1$, and multiplicities m_1, m_2. Then a/N^2 is a cubic in the single free charge R_1, and its derivative,

$$\frac{\partial}{\partial R_1}\frac{a}{N^2} = \frac{3}{32}\Big[\,9\,m_1(R_1-1)^2 - m_1 - 9\,m_2 R_1^2 + m_2\Big]\Big|_{R_2 = 1-R_1} = 0, \tag{VII.6.10}$$

has degree at most two. Its roots and the sign of its second derivative depend on m_1, m_2; the cubic sign alone does not establish a maximum. In the conifold, $R(A_i) = r$, $R(B_j) = 1 - r$ preserves $SU(2)_A \times SU(2)_B$ and obeys every superpotential and mixed node constraint. This leaves baryonic mixing unfixed. Substitution with $m_1 = m_2 = 2$ gives

$$\frac{a(r)}{N^2} = \frac{27}{16}r(1-r), \qquad r_* = \frac{1}{2}, \qquad \frac{a(r_*)}{N^2} = \frac{27}{64}, \qquad \frac{\mathrm{d}^2(a/N^2)}{\mathrm{d}r^2} = -\frac{27}{8}. \tag{VII.6.11}$$

This is an actual constrained maximum. If exchange symmetry is imposed first, $r = 1/2$ is already fixed and its location cannot test the objective function. The value and curvature still discriminate the anomaly combination. Allowing all four individual charges with $R_4 = 2 - R_1 - R_2 - R_3$ gives the Hessian $-\frac{27}{32}(I_3 + \mathbf{1}\mathbf{1}^T)$ at $R_a = 1/2$, with eigenvalues $-27/8, -27/32, -27/32$. Thus the remaining flavor Cartan directions are also locally maximal. These algebraic checks use the specified fermion content; the identification with the superconformal current uses the Intriligator–Wecht result.

VII.6.4 The toric parametrization

The constrained maximization of §VII.6.3 can be run brute-force, solving (VII.6.1)–(VII.6.2) for the free parameters and extremizing. For toric quivers there is a much cleaner parametrization, due to Butti and Zaffaroni, that organizes the trial R-charges by the corners of the toric diagram and makes the match to the geometry transparent. It is the gauge-side half of the dictionary that the next two sections complete.

One charge per external leg. Assign a single trial parameter a_i to each external leg / boundary perfect matching of the toric diagram, $i = 1, \ldots, d$ (in the primitive examples of this chapter, where every boundary lattice point is a polygon corner, these are the corners; non-corner boundary points, the $\mathbb{C}^2/\mathbb{Z}_n$ lines of Chapter 8, are absent here), subject to

$$a_i \geq 0, \qquad \sum_{i=1}^{d} a_i = 2. \tag{VII.6.12}$$

The R-charge of a chiral field is then the sum of the a_i over the legs the field "sweeps". Concretely, each bifundamental of a toric quiver corresponds to a chain of consecutive external legs (the legs between the two zig-zag paths bounding the field), and

$$R_a = \sum_{i \,\in\, \text{legs swept by } a} a_i, \tag{VII.6.13}$$

a sum of *consecutive* external-leg charges. The assignment is geometric: the same zig-zag / (p, q)-leg data that Chapter 4 attached to the fields now distributes the corner charges among them.

Vertex and face constraints. A boundary perfect matching covers every dimer vertex once. Thus the edges in each superpotential term receive every corner charge once, and $\sum_i a_i = 2$ gives the vertex condition $\sum_{e \in \partial v} R_e = 2$. A gauge group corresponds instead to a *face*, whose mixed R anomaly requires $\sum_{e \in \partial f} R_e = n_f - 2$, where n_f counts its incident edges with multiplicity. The face edges do not generally partition the corners. For the conifold, four incident fields of charge $1/2$ sum to $2 = 4 - 2$; for dP_0, six fields of charge $2/3$ sum to $4 = 6 - 2$, with each of its three corner charges occurring twice. The Butti–Zaffaroni charge prescription is used here for a consistent physical toric phase whose face conditions hold; vertex coverage alone does not establish those conditions for an arbitrary tiling. Their Section 5.4 checks the face condition in examples and leaves its general proof open. Within this specified assignment the normalization $\sum_i a_i = 2$

leaves d nonnegative a_i, hence

$$\#\{\text{free trial parameters}\} = d - 1, \tag{VII.6.14}$$

the same count (VII.6.6) arrived at from the global symmetry. The a-maximization of §VII.6.3 becomes a maximization of a, now a function of the a_i, over the $(d-1)$-simplex (VII.6.12): a clean, finite-dimensional problem posed entirely on the toric diagram.

The corner charges of the two fixtures. Make the parametrization concrete on the symmetric examples, where the count and the constraint fix the corners outright. The conifold square has $d = 4$ corners, so $\sum_{i=1}^{4} a_i = 2$ leaves $d - 1 = 3$ free parameters; the $SU(2)_A \times SU(2)_B$ symmetry together with $A \leftrightarrow B$ exchange equates all four, and the normalization gives

$$a_1 = a_2 = a_3 = a_4 = \tfrac{1}{2}, \qquad \sum_{i=1}^{4} a_i = 4 \cdot \tfrac{1}{2} = 2, \tag{VII.6.15}$$

with every bifundamental sweeping a single corner so that $R_a = a_i = \frac{1}{2}$, recovering (VII.6.4). The dP_0 triangle has $d = 3$ corners, so $\sum_{i=1}^{3} a_i = 2$ leaves $d - 1 = 2$ free parameters; the $SU(3)$ flavor symmetry equates the three corner charges, and

$$a_1 = a_2 = a_3 = \tfrac{2}{3}, \qquad \sum_{i=1}^{3} a_i = 3 \cdot \tfrac{2}{3} = 2, \tag{VII.6.16}$$

each of X, Y, Z sweeping a single corner so that $R_a = \frac{2}{3}$, recovering (VII.6.5). In both cases the symmetry collapses the $(d-1)$-simplex to its center and no maximization is needed to locate the symmetric point. Equation (VII.6.11) retains a mixing direction to test the objective; the later $Y^{2,1}$ calculation keeps the two Reeb coordinates and gives an irrational extremum.

The shape of the answer. Because the a_i live on the corners of the polygon and the volume function of the next section also sums over those corners, the toric charge prescription allows the baryonic directions to be eliminated. The resulting two-variable function equals the inverse volume with its stated normalization. This identity is an additional result, not a consequence of both expressions being sums over corners. The next section builds the geometry side in the same language.

VII.6.5 Volume minimization

Cross to the dual geometry. A stack of N D3-branes at the tip of a toric Calabi–Yau cone $C(Y)$, when the link admits the required Sasaki–Einstein metric, has a near-horizon $AdS_5 \times Y$ with Y a five-dimensional Sasaki–Einstein space (Chapter 1 named this for the conifold, $Y = T^{1,1}$). The central charge of the dual theory is holographic and inversely proportional to the volume of Y,

$$a = \frac{\pi^3 N^2}{4\,\mathrm{Vol}(Y)}, \tag{VII.6.17}$$

at leading order in large N, with $\mathrm{Ric}(Y) = 4g_Y$ and $\mathrm{Vol}(S^5) = \pi^3$, in the regime admitting the classical gravity description. This is the formula previewed in Chapter 1. To use it we need $\mathrm{Vol}(Y)$, and the remarkable fact, due to Martelli, Sparks, and Yau, is that the volume can be computed from the toric data alone, by a second extremization that needs no metric.

The Reeb vector and the volume function. A Sasaki–Einstein structure on Y determines a *Reeb vector* b, a vector in the interior of the toric cone that picks out the contact / R-symmetry direction. For a toric threefold the cone is described by its toric diagram, a height-one lattice polygon: the d corners are lifted to three-vectors

$$v_a = (1, x_a, y_a), \qquad a = 1, \ldots, d, \tag{VII.6.18}$$

listed counterclockwise, and the Reeb vector is $b = (b_1, b_2, b_3)$ with its first component fixed,

$$b_1 = 3, \tag{VII.6.19}$$

by normalizing the holomorphic three-form to have Reeb charge 3. The height-one convention makes this the first component. Height one alone does not fix the scale of an arbitrary Reeb vector. The Martelli–Sparks–Yau volume function is the rational sum over the corners

$$Z(b) = \sum_{a=1}^{d} \frac{\langle v_{a-1}, v_a, v_{a+1} \rangle}{\langle v_{a-1}, v_a, b \rangle \, \langle v_a, v_{a+1}, b \rangle}, \tag{VII.6.20}$$

where $\langle \cdot, \cdot, \cdot \rangle$ is the 3×3 determinant and the corner labels are cyclic ($v_0 = v_d$, $v_{d+1} = v_1$). The normalized volume of the trial Sasaki space is $\mathrm{Vol}(Y; b) = (\pi^3/3)\, Z(b)$. Each term is the contribution of one corner, built from the corner and its two neighbours; the denominator pairs the corner with the Reeb vector, which requires $\langle v_a, v_{a+1}, b \rangle > 0$ for every consecutive pair. Equivalently, $(b_2, b_3)/3$ lies strictly inside the polygon. The volume diverges at its boundary; an exterior point need not be a pole, but the rational continuation there is not a Sasaki volume. The relation $\mathrm{Vol} = (\pi^3/3)Z$ is used only on the slice $b_1 = 3$.

The principle. The volume-minimization principle is that the *actual* Sasaki–Einstein metric is the one whose Reeb vector minimizes Z:

$$b_* = \operatorname*{arg\,min}_{b \text{ interior},\ b_1 = 3} Z(b), \qquad \mathrm{Vol}(Y) = \frac{\pi^3}{3}\, Z(b_*), \qquad a = \frac{3}{4}\frac{N^2}{Z(b_*)}, \tag{VII.6.21}$$

the last equality combining (VII.6.17) with $\mathrm{Vol}(Y) = (\pi^3/3)Z(b_*)$. The minimization is over the two free components (b_2, b_3), b_1 being fixed; Z is a convex function on the normalized interior Reeb slice, with a unique minimum for the strictly convex good toric cones in question. A minimum of this function by itself does not prove existence of a Sasaki–Einstein metric. As with a-maximization, the chapter *states* this principle and *uses* it. The derivation of the volume functional from the Sasaki–Einstein metric, the localization formula behind (VII.6.20), and the existence of a metric realizing b_* are the geometry of Part I, routed there and cited, not constructed; the function (VII.6.20) needs only the leg data, so the chapter can minimize it without ever writing a metric.

The structural relation of the gauge side is already visible: both (VII.6.9) and (VII.6.20) are rational/polynomial functions summed over the corners of the same polygon, one maximized, one minimized. The next section states the relation precisely.

VII.6.6 The match: a-maximization is volume minimization

The two extremizations become equivalent after eliminating baryonic mixing. The Butti–Zaffaroni algebraic identity assumes the stated toric assignment of field charges and multiplicities. Its physical application also assumes an interacting SCFT with no omitted accidental currents. These are distinct premises.

The dictionary. Identify the gauge-side corner charges a_i with the geometry-side Reeb data. The clean statement is that the corner R-charge of the i-th external leg is, at the Reeb vector, that leg's share of the volume function. Write the i-th summand of the volume function (VII.6.20) as

$$\ell_i(b) = \frac{\langle v_{i-1}, v_i, v_{i+1}\rangle}{\langle v_{i-1}, v_i, b\rangle \langle v_i, v_{i+1}, b\rangle}, \qquad Z(b) = \sum_{i=1}^{d} \ell_i(b), \tag{VII.6.22}$$

so that Z is the sum of the corner contributions ℓ_i. The Butti–Zaffaroni dictionary then assigns the corner charge

$$a_i = \frac{2\,\ell_i(b)}{\sum_j \ell_j(b)} = \frac{2\,\ell_i(b)}{Z(b)}, \tag{VII.6.23}$$

a *ratio* of the volume-function summands, not a single determinant. The normalization $\sum_i a_i = 2$ is then automatic. The point is that the gauge-side corner charges are built from the *same* determinants $\langle v_{a-1}, v_a, b\rangle$ as the geometry-side volume function. This is the Butti–Zaffaroni identification. The $d-1$ trial charges have $d-3$ baryonic directions in addition to two mesonic directions. On $a_i = 2\ell_i/Z$ the baryonic derivatives vanish; after this elimination the trial coefficient obeys $a(a_i(b))/N^2 = 3/(4Z(b))$. Thus the two-dimensional Reeb domain parametrizes the baryonically extremized surface, not the full trial space. The identity is cited; it is not inferred from a term-by-term identification of two sums. On the symmetric fixtures (VII.6.22)–(VII.6.23) reproduce the corner charges directly: the conifold's four equal summands $\ell_i = \frac{4}{9}$ at $b_* = (3, \frac{3}{2}, \frac{3}{2})$ give $a_i = 2 \cdot \frac{4}{9}/\frac{16}{9} = \frac{1}{2}$, and dP_0's three equal $\ell_i = \frac{1}{3}$ at $Z = 1$ give $a_i = \frac{2}{3}$, matching (VII.6.15)–(VII.6.16).

The consequence. Locally maximizing a over the physical trial charges, including baryonic directions, therefore reduces to minimizing Z on its normalized domain. The maximizing corner charges correspond, under the dictionary, to the minimizing Reeb vector, and the central charges agree:

$$\max_{a_i}\, a(a_i) \iff \min_{b}\, Z(b), \qquad a = \frac{\pi^3 N^2}{4\,\mathrm{Vol}(Y)} \;\propto\; \frac{1}{\mathrm{Vol}(Y)}. \tag{VII.6.24}$$

This is the precise sense of "geometry computes the central charge": a purely field-theory quantity, fixed by the anomalies and a maximization, equals a purely geometric quantity, fixed by a volume and a minimization.

The match, computed both ways. The load-bearing content is that the two computations return the same number. Write the gauge side and the geometry side as the two functions whose extrema we compare,

$$\frac{a}{N^2}\Big|_{\text{gauge}} = \frac{3}{32}\Big[\,2G+\sum_a\big(3(R_a-1)^3-(R_a-1)\big)\Big], \qquad \frac{a}{N^2}\Big|_{\text{geom}} = \frac{3}{4\,Z(b_*)}, \tag{VII.6.25}$$

the second following from $a=\pi^3N^2/(4\,\mathrm{Vol})$ with $\mathrm{Vol}=(\pi^3/3)Z(b_*)$, so that $a/N^2=\pi^3/(4\,\mathrm{Vol})=3/(4Z)$. On the conifold ($G=2$, four fields at the maximizing $R=\frac{1}{2}$, so $R-1=-\frac{1}{2}$, and $Z(b_*)=16/9$) the two sides read

$$\frac{3}{32}\Big[\,2\cdot 2+4\big(3(-\tfrac{1}{2})^3-(-\tfrac{1}{2})\big)\Big] = \frac{3}{32}\cdot\frac{9}{2} = \frac{27}{64} \overset{!}{=} \frac{3}{4\cdot\frac{16}{9}} = \frac{27}{64}. \tag{VII.6.26}$$

On dP_0 ($G=3$, nine fields at $R=\frac{2}{3}$, so $R-1=-\frac{1}{3}$, and $Z(b_*)=1$) they read

$$\frac{3}{32}\Big[\,2\cdot 3+9\big(3(-\tfrac{1}{3})^3-(-\tfrac{1}{3})\big)\Big] = \frac{3}{32}\cdot 8 = \frac{3}{4} \overset{!}{=} \frac{3}{4\cdot 1} = \frac{3}{4}. \tag{VII.6.27}$$

The agreement is the test. A word on what is and is not claimed: the Butti–Zaffaroni dictionary is the *principle* that identifies the variational problems, and the chapter's checks verify that the two computed central charges *agree* (here 27/64 both ways, 3/4 both ways), not that a literal term-by-term map of summands was carried out. It is the equality of the answers, on rational and (below) irrational cones alike, that is the evidence.

a is an invariant of the cone. One corollary is worth stating, because the next chapter rests on it. The central charge a depends only on $\mathrm{Vol}(Y)$, a property of the geometry, so it is an invariant of the toric cone. Two different quivers presenting the same cone, different dimers with the same Newton polygon, have the same leading coefficient of a when they are consistent regular-brane phases with the same rank N and flow to the stipulated SCFT. This is what makes a phase-independent across Seiberg-dual toric phases, and Chapter 7 uses exactly that invariance as a check of toric duality. Here the point is registered, not exploited: volume minimization fixes a from the cone, so any field theory dual to that cone inherits the same number.

The conifold and dP_0 run both sides explicitly: rational and symmetric, the maximization collapses and the match is arithmetic. For $Y^{2,1}$ the geometry side (the Z-minimization) is computed here and the gauge-side equality is invoked through Butti–Zaffaroni; the match is the historic test, between an irrational metric volume and an irrational central charge.

Table VII.6.1 collects the dictionary: a-maximization and volume minimization are one variational problem written in two languages.

Gauge theory	Geometry
trial R-symmetry	trial Reeb vector b
constraints $R(W)=2$, node anomaly	normalization $b_1=3$, interior chamber
external-leg charges a_i	volume summands $\ell_i(b)$
maximize $a(R)$	*minimize* $Z(b)$
central charge from 't Hooft anomalies	central charge from $1/\operatorname{Vol}(Y)$

Table VII.6.1: The gauge-side / geometry-side dictionary (Butti–Zaffaroni): a-maximization and volume minimization have matching extrema at leading large N, after baryonic extremization using the toric charge and multiplicity prescription. The Reeb coordinates do not parametrize all $d-1$ trial directions.

VII.6.7 The conifold: completing the Chapter 1 preview

Chapter 1 fixed the conifold R-charge $R(A)=R(B)=\frac{1}{2}$ by the $A\leftrightarrow B$ exchange symmetry and a beta-function check, and previewed the central charge from the volume of $T^{1,1}$. This section recovers the same numbers from the two extremization principles, the rational, symmetry-fixed warm-up.

Gauge side. The conifold quiver has two nodes and four bifundamentals A_1, A_2, B_1, B_2, with the quartic superpotential $\operatorname{Tr}(A_iB_kA_jB_l)$. The toric diagram is the unit square, $d=4$, so the trial space is three-dimensional before symmetry. The $SU(2)_A\times SU(2)_B$ flavor symmetry equates the two A charges and the two B charges separately. Adding $A\leftrightarrow B$ exchange forces a common charge R, and the superpotential constraint (VII.6.1) is $4R=2$, so

$$R(A)=R(B)=\tfrac{1}{2}, \tag{VII.6.28}$$

with nothing left to maximize: the symmetry collapses the trial space to a point. (The node constraint (VII.6.2), $2+4(R-1)=0$, gives the same $R=\frac{1}{2}$, the consistency of §VII.6.2.) Evaluating the trial central charge (VII.6.9) at $G=2$ and four fermions of $R-1=-\frac{1}{2}$,

$$\frac{a}{N^2}=\frac{3}{32}\Big[\,2\cdot 2+4\big(3(-\tfrac{1}{2})^3-(-\tfrac{1}{2})\big)\Big]=\frac{3}{32}\Big[\,4+4\cdot\tfrac{1}{8}\Big]=\frac{27}{64}, \tag{VII.6.29}$$

the Klebanov–Witten central charge, matching the field-theory anomaly of Chapter 1. Although the symmetry pins $R=\frac{1}{2}$ for any objective function, the *value* 27/64 is sensitive to using the correct combination: the swapped function $3\operatorname{Tr}R-\operatorname{Tr}R^3$ at the same $R=\frac{1}{2}$ gives $-9/64$, not 27/64, so it is genuinely (VII.6.7) that the geometry will reproduce.

Geometry side. The conifold cone has the unit-square diagram, with corners (counterclockwise) $(0,0),(1,0),(1,1),(0,1)$ lifted to $v_a=(1,x_a,y_a)$. The Reeb vector is $b=(3,b_2,b_3)$, and inserting the four corners into the volume function (VII.6.20) gives the explicit rational function

$$Z(b_2,b_3)=\frac{9}{b_2\,b_3\,\big(b_2b_3-3b_2-3b_3+9\big)}, \tag{VII.6.30}$$

the simplest polygon worked all the way into Z. Setting $\partial_{b_2} Z = \partial_{b_3} Z = 0$ and using the $b_2 \leftrightarrow b_3$ symmetry of the square gives the unique interior critical point

$$b_* = \big(3, \tfrac{3}{2}, \tfrac{3}{2}\big), \qquad Z(b_*) = \frac{16}{9}, \tag{VII.6.31}$$

manifestly interior: the projected Reeb point $(b_2/b_1, b_3/b_1) = (\frac{1}{2}, \frac{1}{2})$ sits at the centre of the unit-square diagram. The minimal volume is then

$$\mathrm{Vol}(T^{1,1}) = \frac{\pi^3}{3}\, Z(b_*) = \frac{\pi^3}{3} \cdot \frac{16}{9} = \frac{16\pi^3}{27}, \tag{VII.6.32}$$

the value quoted in Chapter 1, to be compared with $\mathrm{Vol}(S^5) = \pi^3$. The central charge from the volume is

$$a = \frac{\pi^3 N^2}{4\,\mathrm{Vol}(T^{1,1})} = \frac{3}{4}\frac{N^2}{Z(b_*)} = \frac{3}{4} \cdot \frac{9}{16}\, N^2 = \frac{27}{64}\, N^2, \tag{VII.6.33}$$

exactly (VII.6.29). The two sides meet: the gauge-side a-maximization (here trivial, fixed by symmetry) and the geometry-side volume minimization give the same $27/64$, completing the preview Chapter 1 left open.

VII.6.8 $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$: the orbifold inheritance

The second rational example is dP_0, the cone over $\mathbb{P}^2$, equivalently the orbifold $\mathbb{C}^3/\mathbb{Z}_3$. It is the bridge to Volume VI: the central charge it computes is the geometric face of the orbifold-inheritance relation that Volume VI derived from the parent $\mathcal{N}{=}4$ theory.

Gauge side. The dP_0 quiver has three nodes and nine bifundamentals X, Y, Z (three of each), with the cubic superpotential $\mathrm{Tr}(XYZ)$ antisymmetrized over the flavor indices. The toric diagram is a triangle, $d = 3$. The $SU(3)$ flavor symmetry and cyclic node symmetry select the symmetric assignment, with common corner charge $a_i = \frac{2}{3}$, hence a common field R-charge, and the superpotential constraint (VII.6.1) reads $3R = 2$, so

$$R(X) = R(Y) = R(Z) = \tfrac{2}{3}, \tag{VII.6.34}$$

again with nothing left to maximize (the node constraint $2 + 6(R-1) = 0$, six fields at a node, gives the same $R = \frac{2}{3}$). The trial central charge (VII.6.9) at $G = 3$ and nine fermions of $R - 1 = -\frac{1}{3}$ is

$$\frac{a}{N^2} = \frac{3}{32}\Big[\,2\cdot 3 + 9\big(3(-\tfrac{1}{3})^3 - (-\tfrac{1}{3})\big)\Big] = \frac{3}{32}\Big[\,6 + 9\cdot\tfrac{2}{9}\Big] = \frac{3}{32}\cdot 8 = \frac{3}{4}. \tag{VII.6.35}$$

Geometry side. The dP_0 cone is the orbifold $\mathbb{C}^3/\mathbb{Z}_3$, whose link is the quotient $S^5/\mathbb{Z}_3$ of the round five-sphere by a free $\mathbb{Z}_3$ action. The volume of a free quotient is the volume divided by the order of the group,

$$\mathrm{Vol}(S^5/\mathbb{Z}_3) = \frac{\mathrm{Vol}(S^5)}{3} = \frac{\pi^3}{3}, \tag{VII.6.36}$$

which is also what volume minimization gives directly: inserting the three triangle corners $v_a = (1, x_a, y_a)$ into the volume function (VII.6.20) gives a $\mathbb{Z}_3$-symmetric $Z(b)$ whose interior minimum sits at $b_* = (3, 0, 0)$ with $Z(b_*) = 1$, so $\mathrm{Vol} = (\pi^3/3) \cdot 1 = \pi^3/3$. The central charge from the volume is

$$a = \frac{\pi^3 N^2}{4 \cdot (\pi^3/3)} = \frac{3}{4}\, N^2, \tag{VII.6.37}$$

matching (VII.6.35).

The inheritance. This number makes the Volume VI orbifold-daughter relation geometric. The parent $\mathcal{N}$=4 super Yang–Mills with gauge group $SU(M)$ has, at large M, central charge

$$a\big(\mathcal{N}{=}4,\ SU(M)\big) = \frac{\dim G}{4} = \frac{M^2 - 1}{4} \ \xrightarrow{M\to\infty}\ \frac{M^2}{4}, \tag{VII.6.38}$$

(a vector multiplet plus three adjoint chirals at $R = \frac{2}{3}$, evaluated through (VII.6.7)). The dP_0 theory is the $\mathbb{Z}_3$ orbifold daughter of the parent $SU(3N)$ theory, keeping the $1/|\Gamma| = \frac{1}{3}$ slice, so

$$a(dP_0)\big|_{N^2} = \frac{3}{4}\, N^2 = \frac{1}{3}\, a\big(\mathcal{N}{=}4,\ SU(3N)\big)\big|_{N^2} = 3\, a\big(\mathcal{N}{=}4,\ SU(N)\big)\big|_{N^2}, \tag{VII.6.39}$$

since $\frac{1}{3} \cdot \frac{(3N)^2}{4} = \frac{3}{4}N^2$ and $3 \cdot \frac{N^2}{4} = \frac{3}{4}N^2$. The rank must be held explicit. At fixed rank N, dividing the link volume by three multiplies the leading central charge by three. Relative to the parent rank $3N$, the daughter instead has one-third of the leading central charge. These statements are consistent because the numerator of the holographic formula is the square of the flux, hence of the rank.

For the interacting $SU(N)^3$ sector, retaining $\dim SU(N) = N^2 - 1$ in the gaugino contribution gives

$$a_{dP_0} = \frac{3}{4}N^2 - \frac{9}{16}, \qquad \frac{1}{3}a_{\mathcal{N}=4,SU(3N)} = \frac{3}{4}N^2 - \frac{1}{12}. \tag{VII.6.40}$$

The two differ by $-23/48$. Similarly the interacting conifold sector has $a = 27N^2/64 - 3/8$. Free abelian multiplets would add their own contributions; these formulas specify the nonabelian interacting sector, for ranks where that SCFT description applies. The classical volume comparison concerns the N^2 term, not these constant terms.

VII.6.9 $Y^{2,1}$: the irrational case

The conifold and dP_0 are rational because they are highly symmetric. The decisive test of the a-maximization / volume-minimization correspondence came from a family where the symmetry is too small to fix the R-charges by inspection, and the answer is irrational: the $Y^{p,q}$ Sasaki–Einstein manifolds. This section works the simplest nontrivial member, $Y^{2,1}$: the geometry-side volume minimization explicitly, with the Butti–Zaffaroni equivalence identifying the gauge-side a-maximization result.

The cone. $Y^{2,1}$ is a Sasaki–Einstein metric on $S^2 \times S^3$, a member of the $Y^{p,q}$ family (the conifold is the degenerate $Y^{1,0}$ limit). Its toric diagram is the quadrilateral

$$\Delta_{Y^{2,1}} = \big\{(0,0),(2,1),(1,2),(0,1)\big\}, \tag{VII.6.41}$$

with $d = 4$ corners, one interior lattice point, and twice-area $4 = 2p$ (so the quiver has $G = 2 \cdot \text{Area} = 4$ gauge nodes). Figure VII.6.1 draws the lattice polygon; the minimizing Reeb vector is computed below. The field theory is a four-node toric quiver (the Benvenuti–Franco–Hanany–Martelli–Sparks theory); its R-charges are exactly what the corner parametrization (VII.6.12)–(VII.6.13) produces once the maximization is done.

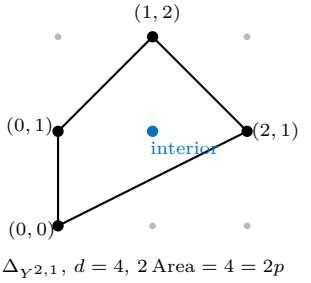


Figure VII.6.1: The $Y^{2,1}$ toric diagram, a lattice quadrilateral with four corners and one interior lattice point. Read it as the toric input to both extremizations: the four corners carry the trial R-charges a_i of the gauge-side parametrization and index the four facets of the geometry-side volume function. Twice the Euclidean lattice area, $4 = 2p$, is the number of gauge nodes ($G = 4$). The Reeb vector $b = (3, b_2, b_3)$ that minimizes the volume sits at an *irrational* interior point (schematic; the figure shows the lattice, not b, which is not a lattice vector).

The computation. Run volume minimization on (VII.6.41). Lift the four corners to $v_a = (1, x_a, y_a)$ and the Reeb vector to $b = (3, b_2, b_3)$; the volume function (VII.6.20) is the explicit rational function

$$Z(b_2, b_3) = \frac{6\,(2b_2 - b_3 + 9)}{b_2\,(b_2 - 2b_3)\,(b_2 - b_3 + 3)\,(b_2 + b_3 - 9)}, \tag{VII.6.42}$$

whose four denominator factors are the four corner pairings $\langle v_{a-1}, v_a, b\rangle$ with the Reeb vector. Its interior critical point solves the stationarity equations $\partial_{b_2} Z = \partial_{b_3} Z = 0$. Set $b_3 = 3$ as a candidate slice and require both gradient components to vanish. On this slice their numerators have the common factor $b_2^2 + 2b_2 - 12$. An unreduced polynomial condition can be written as

$$(b_2 - 6)\,(b_2^2 + 2b_2 - 12) = 0, \tag{VII.6.43}$$

whose factor $b_2 = 6$ is a boundary pole of Z, where a corner denominator $\langle v_{a-1}, v_a, b\rangle$ vanishes (a spurious root left by clearing denominators, not an interior critical point),

while the genuine interior root is $b_2 = \sqrt{13} - 1$. Unlike the symmetric fixtures the answer is *not* a lattice point but carries a surd,

$$b_* = \big(3,\ \sqrt{13} - 1,\ 3\big), \qquad Z(b_*) = \frac{23}{54} + \frac{13\sqrt{13}}{108} \approx 0.85992747, \tag{VII.6.44}$$

an algebraic irrational. Substitution in both derivatives verifies stationarity; the positive-definite Hessian and strict convexity on the normalized Reeb chamber select the minimum. The four chamber inequalities are $b_2 > 0$, $2b_3 - b_2 > 0$, $9 - b_2 - b_3 > 0$, and $b_2 - b_3 + 3 > 0$. Thus the projected Reeb point $(b_2, b_3)/3$ lies inside the quadrilateral, and the irrationality follows from the nonsquare discriminant 13. The minimal volume follows,

$$\mathrm{Vol}(Y^{2,1}) = \frac{\pi^3}{3} Z(b_*) \approx 0.28664249\,\pi^3. \tag{VII.6.45}$$

The number is not a numerical accident: it agrees to all computed digits with the Gauntlett–Martelli–Sparks–Waldram closed-form volume, which for general (p, q) is

$$\frac{\mathrm{Vol}(Y^{p,q})}{\pi^3} = \frac{q^2\big(2p + \sqrt{4p^2 - 3q^2}\big)}{3p^2\big(3q^2 - 2p^2 + p\sqrt{4p^2 - 3q^2}\big)}, \tag{VII.6.46}$$

for coprime integers $0 < q < p$, which give the smooth family used here. It is not irrational for every admissible pair; the discriminant can be a square. The displayed quotient is $0/0$ at $q = 0$, but rationalizing first gives $\mathrm{Vol}\,/\pi^3 = (2p + \sqrt{4p^2 - 3q^2})^2/[9p^2(p + \sqrt{4p^2 - 3q^2})]$, whose continuous $q \to 0$ limit is $16/(27p)$. For $p = 1$ it recovers $T^{1,1}$. At $(p, q) = (2, 1)$ the discriminant is $4p^2 - 3q^2 = 16 - 3 = 13$, so the closed form carries an explicit $\sqrt{13}$,

$$\frac{\mathrm{Vol}(Y^{2,1})}{\pi^3} = \frac{1 \cdot \big(4 + \sqrt{13}\big)}{3 \cdot 4 \cdot \big(3 - 8 + 2\sqrt{13}\big)} = \frac{4 + \sqrt{13}}{12\,\big(2\sqrt{13} - 5\big)} \approx 0.28664249, \tag{VII.6.47}$$

an algebraic irrational, matching the symbolic minimum (VII.6.45) to all computed digits. The central charge follows from (VII.6.17),

$$\frac{a(Y^{2,1})}{N^2} = \frac{\pi^3}{4\,\mathrm{Vol}(Y^{2,1})} = \frac{3}{4\,Z(b_*)} \approx 0.872167, \tag{VII.6.48}$$

itself irrational. The contrast with the first two examples is the whole point:

$$\frac{a}{N^2} = \frac{27}{64}\ (\text{conifold}), \quad \frac{a}{N^2} = \frac{3}{4}\ (dP_0), \quad \frac{a}{N^2} = \frac{3}{4Z(b_*)} \notin \mathbb{Q}\ (Y^{2,1}), \tag{VII.6.49}$$

the symmetric cones giving rational central charges and $Y^{2,1}$ giving an irrational one carrying the same $\sqrt{13}$ on the gauge side and the geometry side.

Why this was decisive. On the field-theory side the gauge a-maximization gives the same irrational a and the same irrational R-charges: the Butti–Zaffaroni equivalence ties it to the Z-minimization computed here, and the four-node $Y^{2,1}$ quiver maximization was carried out in the literature, with the $\sqrt{13}$ appearing as the root of the constrained

cubic. This is the strongest evidence the construction offers that the two sides describe the same physics. This field theory's central charge is an irrational number, with no reason to be guessable. Yet it is reproduced exactly by minimizing a volume function read off a lattice quadrilateral. The field-theory calculation is imported from the literature, where the full $Y^{p,q}$ quiver charges and multiplicities are specified. No independent $Y^{2,1}$ gauge maximization is carried out in this chapter. Two cautions belong with the result. The match is exact at leading order in N; the irrationality is a feature of the leading central charge, not a subleading correction. And $Y^{2,1}$ is chosen, like the conifold and dP_0, to have no accidental symmetry that would invalidate the bare toric a-maximization; the general treatment of accidental symmetries is named in the next section and routed to Part III.

VII.6.10 What has been computed, and what is deferred

The chapter has compared the leading central charge using explicit gauge and geometry calculations for the conifold and dP_0. For $Y^{2,1}$ the geometric extremum is calculated and the gauge-side equality is imported. On the gauge side, trial R-charges were assigned from the dimer data subject to two linear constraints (the superpotential carries R-charge 2, every node is anomaly-free), parametrized cleanly by one charge per corner of the toric diagram, and the superconformal R was selected by maximizing the cubic a-function $a(R) = \frac{3}{32}(3\,\mathrm{Tr}\,R^3 - \mathrm{Tr}\,R)$. On the geometry side, the same number was obtained by minimizing the Martelli–Sparks–Yau volume function $Z(b)$ over the Reeb vector and reading $a = \pi^3 N^2/(4\,\mathrm{Vol}(Y))$. The Butti–Zaffaroni dictionary identified the two extremizations as one. The conifold recovered $R = \frac{1}{2}$ and $\mathrm{Vol}(T^{1,1}) = 16\pi^3/27$, completing the Chapter 1 preview; dP_0 gave $R = \frac{2}{3}$ and $a = \frac{3}{4}N^2 = \frac{1}{3}\,a(\mathcal{N}{=}4, SU(3N))$, making the Volume VI orbifold inheritance geometric; and $Y^{2,1}$ produced an irrational central charge matching the closed-form Sasaki–Einstein volume, the historic test of the correspondence.

A reader leaving the chapter should be able to do the following without looking ahead: write the two linear constraints on a trial R-symmetry and count the $d-1$ free parameters from the toric diagram; assign the Butti–Zaffaroni corner charges and see $\sum_i a_i = 2$ make the superpotential constraint automatic; write the cubic a-function from the fermion content and state that the superconformal R maximizes it; write the Martelli–Sparks–Yau volume function from the corners and minimize it with $b_1 = 3$ fixed; recover $\mathrm{Vol}(T^{1,1})$ and $a = \frac{27}{64}N^2$ for the conifold and $\mathrm{Vol}(S^5/\mathbb{Z}_3)$ and $a = \frac{3}{4}N^2$ for dP_0; explain why $a(dP_0)$ is one-third of the $SU(3N)$ parent; and explain why $Y^{2,1}$ gives an irrational answer on both sides.

What has been deferred, and to where, is the boundary that the volume is built around.

- The *proof* that the superconformal R maximizes a, the a-theorem, and the treatment of accidental symmetries and the unitarity bound (when a gauge-invariant chiral primary reaches $\Delta = 1$, or a trial assignment would put it below that bound, requiring a free sector and an accidental current) are field-theory

results of Part III. The toric examples here were chosen so that no accidental symmetry intervenes and the bare toric a-maximization is exact; the chapter does not claim that bare toric a-maximization is always exact.

- The *existence* of the Sasaki–Einstein metric realizing the minimizing Reeb vector, and the explicit $Y^{p,q}$ and $L^{a,b,c}$ metrics, are the geometry of Part I, cited not constructed. The chapter used only the volume *functional* (VII.6.20), which needs the leg data and no metric.
- The match was exhibited at leading order in N, for the central charge $a = \pi^3 N^2/(4\,\mathrm{Vol})$. The subleading-in-$N$ corrections, the Kaluza–Klein and string spectrum, and the operator-by-operator holographic dictionary are Part III. The leading classical central charge carries the α'/g_s probe caveats inherited from Volume V.
- The central charge a was shown to depend only on $\mathrm{Vol}(Y)$, hence to be an invariant of the cone. Chapter 7 uses that invariance as a check across Seiberg-dual toric phases (the toric-duality move itself, urban renewal of the dimer, is Chapter 7's); this chapter names the invariance and does not run the duality.
- The central charge was fixed at the conformal fixed point. The renormalization-group flow, the duality cascade along which a decreases, and the deformed-conifold infrared are Chapter 8.

At leading large N, with the SCFT and toric-prescription hypotheses retained, the lesson to carry forward is the keybox at the start: the field theory's central charge, an anomaly computed by maximizing a cubic, equals the geometry's, a volume computed by minimizing a rational function, and for the toric quivers considered here the two extrema are related after baryonic extremization. The next chapter turns the resulting invariance of a into a test of duality.

Status and sources

The evidence-code tags ([WS] controlled physical-worldsheet derivation in the untwisted string, [MATH] mathematical result with stated hypotheses and proof basis), the convention pointers, and the source anchors are collected here, keyed by section. The chapter opens the construction-intrinsic dynamics chapters (VII.6–VII.8) by computing the infrared central charge of a toric-quiver fixed point two ways and matching them: gauge-side a-maximization after Intriligator and Wecht (2003), geometry-side volume / Z-minimization after Martelli et al. (2006) and Martelli et al. (2008), with the toric corner-charge parametrization and the equivalence of the two extremizations after Butti and Zaffaroni (2005). The worked examples are the conifold (recalled from Chapter 1, after Klebanov and Witten (1998) and Gubser (1999)), $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ (the Volume VI orbifold-inheritance bridge), and the irrational $Y^{2,1}$ (after Benvenuti et al. (2005) for the gauge theory and Gauntlett et al. (2004) for the Sasaki–Einstein metric). It introduces the new convention a-maximization and volume minimization (the a-maximization / volume-minimization normalization; this is its first load-bearing use) and reuses the toric-diagram conventions (toric lattice orientation, the height-one polygon and external legs), the dimer conventions (bipartite tiling conventions, the corner perfect matchings and external legs (distinct sets with the same count in these examples)), the superpotential store (superpotential normalization, the $\mathrm{R}(W) = 2$ constraint), the gauge/rank data (quiver gauge groups and ranks, $\prod_i U(N_i)$, the $d-3$ baryonic $U(1)$'s), the supercharge count (spinor and supersymmetry counting, $\mathcal{N}$=1), and the arrow convention. The trial-charge and volume comparisons use the stated finite toric data.

Evidence and assumptions. The legacy Q1–Q5 groupings below identify source topics, not a single evidence class. No physical-worldsheet derivation is supplied by the anomaly polynomial, a rational volume function or its algebraic extremum. These calculations leave physical primary classifications null where a the physical derivation is absent. Local mathematical results have the precise finite polynomial or Reeb chamber hypotheses stated in the text. The Intriligator–Wecht principle is imported from Intriligator and Wecht (2003), equations (1.5)–(1.7); the toric volume and normalized minimization from Martelli et al. (2006), section 3.3; and baryonic elimination and inverse-volume equality from Butti and Zaffaroni (2005), equations (4.11)–(4.16), appendices A.3–A.4. The $Y^{p,q}$ gauge-side answer is imported from Benvenuti et al. (2005), section 4.2, equations (4.10)–(4.13); it is not recomputed here. The legacy suites remain unchanged and test their encoded objects. Their passing results neither supply independent physical evidence nor establish full trial-space coverage. In particular a numerical noninteger is not a proof of irrationality.

Executable exercises. The chapter-local companion supplies three structured-answer families: the conifold constrained polynomial with value and Hessian; the exact conifold Reeb extremum with chamber and normalization; and the asymmetric $Y^{2,1}$ geometry in $\mathbb{Q}(\sqrt{13})$. Each family grades new submitted quantities. Current text and source identities are distinct from these mathematical checks. The companion does not parse arbitrary printed formulas, establish metric existence, or supply the missing gauge-side $Y^{2,1}$ derivation.

§VII.6.1 Two extremization problems and one number. Framing: a as the cleanest construction-intrinsic invariant; the two dual extremization principles (a-maximization on the gauge side, volume minimization on the geometry side) and the Butti–Zaffaroni equivalence; the reader contract and the dynamics boundary.

§VII.6.2 Trial R-charges from the dimer data. The superpotential constraint $\sum R_a = 2$ (VII.6.1), the anomaly-free node constraint $2 + \sum(R_a - 1) = 0$ (VII.6.2) and its $R = 1 - 2/n_i$ solution, the consistency of the two on the symmetric examples, and the maximal non-R abelian subgroup $U(1)^{d-1}$ leaving $(d-1)$ free trial parameters (VII.6.6); the seam with the Chapter 5 moduli/global-symmetry machine. physical classification pending (a-maximization and volume minimization, superpotential normalization, quiver gauge groups and ranks; Intriligator and Wecht (2003), Butti and Zaffaroni (2005)). Carries (the baryonic count is $d-3$, the gauge group $\prod_i U(N_i)$ not bare SU). on the conifold and dP_0 the W-constraint and the node-anomaly condition give the same symmetric R-charge $(\frac{1}{2}, \frac{2}{3})$, and the corner count gives $d - 1 = 3, 2$ free parameters.

§VII.6.3 The principle: a-maximization. The trial central charge $a(R) = \frac{3}{32}(3\,\mathrm{Tr}\,R^3 - \mathrm{Tr}\,R)$ over the fermions (VII.6.7), the Intriligator–Wecht maximization principle with the negative-definite Hessian (VII.6.8), and the explicit cubic (VII.6.9). Legacy Q2 (a-maximization and volume minimization; Intriligator and Wecht (2003)). The proof that the maximizing combination is the superconformal R, and the a-theorem, are the dynamics of a-maximization (Part III), stated and cited not proved. Maximizing $a(R)$ over the constrained trial R returns the known R-charges and central charges in the examples. The actual constrained Hessian is displayed in (VII.6.11), and the swapped combination $3\,\mathrm{Tr}\,R - \mathrm{Tr}\,R^3$ gives the wrong a.

§VII.6.4 The toric parametrization (Butti–Zaffaroni). One corner charge $a_i \geq 0$ per external leg with $\sum_i a_i = 2$ (VII.6.12); each field's R-charge a sum of consecutive a_i (VII.6.13); the dimer-vertex W-constraint follows from perfect-matching coverage, while each gauge-face anomaly condition must hold for the specified consistent charge assignment; the $(d-1)$ free parameters (VII.6.14). physical classification pending (a-maximization and volume minimization, toric lattice orientation, bipartite tiling conventions; Butti and Zaffaroni (2005)). $\sum_i a_i = 2$ and the $d-1$ count on the fixtures.

§VII.6.5 Volume minimization (Martelli–Sparks–Yau). The holographic $a = \pi^3 N^2/(4\,\mathrm{Vol}(Y))$ (VII.6.17); the Reeb vector with $b_1 = 3$ fixed by the holomorphic three-form normalization in the height-one basis (VII.6.19); the MSY corner-sum volume function $Z(b)$ (VII.6.20); the minimization principle and $\mathrm{Vol}(Y) = (\pi^3/3)Z(b_*)$ (VII.6.21). Legacy Q3 (a-maximization and volume minimization, toric lattice orientation; Martelli et al. (2006), Martelli et al. (2008)). The Sasaki–Einstein metric existence, the localization derivation of the volume functional, and the geometric side of the minimization theorem are Sasaki–Einstein metric existence (Vol III), cited not constructed. $Z(b)$ minimizes (with $b_1 = 3$ fixed and load-bearing, $Z(b_1{=}3) \neq Z(b_1{=}2)$) to $\mathrm{Vol}(S^5) = \pi^3$ ($\mathbb{C}^3$), $16\pi^3/27$ ($T^{1,1}$), $\pi^3/3$ ($S^5/\mathbb{Z}_3$), at

interior Reeb vectors ($b_* = (3, \frac{3}{2}, \frac{3}{2})$ for the conifold); F2 compares the non-minimizing but interior point $(b_2, b_3) = (1, 1)$, while the separate scale check changes b_1.

§VII.6.6 The match: a-maximization is volume minimization. The baryonically extremized corner-R ↔ Reeb dictionary $a_i = 2\ell_i(b)/Z(b)$ with ℓ_i the i-th volume-function summand (VII.6.22)–(VII.6.23) (the corner charge a ratio of the same determinants as Z, not a single determinant), the Butti–Zaffaroni identification of the two variational problems (the equality of the computed central charges being the test, beyond the dictionary), the equivalence (VII.6.24) and $a \propto 1/\operatorname{Vol}(Y)$, and the corollary that a is an invariant of the cone (phase-independent, Chapter 7). Legacy Q4 (a-maximization and volume minimization; Butti and Zaffaroni (2005)). the gauge-side a-max and the geometry-side $3/(4Z_{\min})$ agree on the conifold and dP_0, the BZ corner charges $a_i = 2\ell_i/Z$ reproduce $\frac{1}{2}$ (conifold) and $\frac{2}{3}$ (dP_0), and the constant $a\operatorname{Vol}/N^2 = \pi^3/4$ is the same across fixtures; the wrong volume normalization (π^3 instead of $\pi^3/3$) breaks the match.

§VII.6.7 The conifold. The symmetry-fixed $R = \frac{1}{2}$ (VII.6.28), the gauge-side $a = \frac{27}{64}N^2$ (VII.6.29), the geometry-side $b_* = (3, \frac{3}{2}, \frac{3}{2})$, $Z = 16/9$, and $\operatorname{Vol}(T^{1,1}) = 16\pi^3/27$ (VII.6.31)–(VII.6.32), and the match (VII.6.33), completing the Chapter 1 preview. physical classification pending (a-maximization and volume minimization; Klebanov and Witten (1998), Gubser (1999)). the conifold a-max returns $R = \frac{1}{2}$, $a/N^2 = \frac{27}{64}$, and the Z-min gives $b_* = (3, \frac{3}{2}, \frac{3}{2})$, $\operatorname{Vol} = 16\pi^3/27$.

§VII.6.8 $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$. The flavor-symmetric $R = \frac{2}{3}$ (VII.6.34), $a = \frac{3}{4}N^2$ (VII.6.35); $\operatorname{Vol}(S^5/\mathbb{Z}_3) = \operatorname{Vol}(S^5)/3 = \pi^3/3$ (VII.6.36)–(VII.6.37); the leading-N^2 orbifold-inheritance relation $a(dP_0) = \frac{1}{3}a(\mathcal{N}{=}4, SU(3N)) = 3a(\mathcal{N}{=}4, SU(N))$ (VII.6.38)–(VII.6.39), the Volume VI (VI.10) bridge made geometric. physical classification pending. Volume VI, Chapter 10 gives the same leading central charge through planar inheritance. the dP_0 a-max returns $R = \frac{2}{3}$, $a/N^2 = \frac{3}{4}$; the Z-min gives $\operatorname{Vol} = \pi^3/3$; the orbifold relation $a(dP_0) = 3a(\mathcal{N}{=}4, SU(N)) = \frac{1}{3}a(\mathcal{N}{=}4, SU(3N))$ holds at leading order, and the factor is $\frac{1}{3}$ not 1.

§VII.6.9 $Y^{2,1}$: the irrational case. The $Y^{2,1}$ toric quadrilateral with one interior point and $G = 4$ (VII.6.41) (Figure VII.6.1); the irrational volume minimum (VII.6.45), its agreement with the Gauntlett–Martelli–Sparks–Waldram closed form (a surd in $\sqrt{13}$) (VII.6.46), and the irrational central charge (VII.6.48); the historic point that the gauge-side and geometry-side surds agree. Legacy Q5 (a-maximization and volume minimization; Benvenuti et al. (2005), Gauntlett et al. (2004)). The SE metric existence is Sasaki–Einstein metric existence (Vol III). Files (toric central charges / R-charges are not always rational, and a-maximization needs no gravity dual; the match to $\operatorname{Vol}(Y)$ is a check, not an input). The finite toric calculation Z-min volume matches the GMSW closed form, and the central charge is irrational, not rational.

§VII.6.10 What is and is not given. Scope and limitations. the chapter computes both sides for two symmetric cones and computes the $Y^{2,1}$ geometry while importing its gauge-side match; it routes the a-theorem and the accidental-symmetry treatment to Part III (the dynamics of a-maximization), the Sasaki–Einstein metric existence to Vol III (Sasaki–Einstein metric existence), the subleading-N and Kaluza–Klein spectrum to Part III, the toric-duality use of the invariance to Chapter 7, and the renormalization-group flow / cascade to Chapter 8.

Related constructions. For toric quivers the Butti–Zaffaroni comparison relates the maximized gauge-theory a-function to the minimized Sasaki–Einstein volume, giving $a = \pi^3 N^2/(4\operatorname{Vol}(Y))$ at leading order. For dP_0 this yields $a = \frac{1}{3}a(\mathcal{N} = 4, SU(3N))$, also obtained by planar inheritance in Volume VI, Chapter 10. Part III discusses the holographic spectrum.

References

Benvenuti, S. et al. (2005). "An Infinite family of superconformal quiver gauge theories with Sasaki-Einstein duals". In: *JHEP* 06, p. 064. DOI: 10.1088/1126-6708/2005/06/064. arXiv: hep-th/0411264.

Butti, A. and A. Zaffaroni (2005). "R-charges from toric diagrams and the equivalence of a-maximization and Z-minimization". In: *JHEP* 11, p. 019. DOI: 10.1088/1126-6708/2005/11/019. arXiv: hep-th/0506232.

Gauntlett, J. P. et al. (2004). "Sasaki-Einstein metrics on S**2 x S**3". In: *Adv. Theor. Math. Phys.* 8, pp. 711–734. DOI: 10.4310/ATMP.2004.v8.n4.a3. arXiv: hep-th/0403002.

Gubser, S. S. (1999). "Einstein manifolds and conformal field theories". In: *Phys. Rev. D* 59, p. 025006. DOI: 10.1103/PhysRevD.59.025006. arXiv: hep-th/9807164.

Intriligator, K. A. and B. Wecht (2003). "The Exact superconformal R symmetry maximizes a". In: *Nucl. Phys. B* 667, pp. 183–200. DOI: 10.1016/S0550-3213(03)00459-0. arXiv: hep-th/0304128.

Klebanov, I. R. and E. Witten (1998). "Superconformal field theory on three-branes at a Calabi-Yau singularity". In: *Nucl. Phys. B* 536, pp. 199–218. DOI: 10.1016/S0550-3213(98)00654-3. arXiv: hep-th/9807080.

Martelli, D., J. Sparks, and S.-T. Yau (2006). "The Geometric dual of a-maximisation for Toric Sasaki-Einstein manifolds". In: *Commun. Math. Phys.* 268, pp. 39–65. DOI: 10.1007/s00220-006-0087-0. arXiv: hep-th/0503183.

Martelli, D., J. Sparks, and S.-T. Yau (2008). "Sasaki-Einstein manifolds and volume minimisation". In: *Commun. Math. Phys.* 280, pp. 611–673. DOI: 10.1007/s00220-008-0479-4. arXiv: hep-th/0603021.

Chapter VII.7

Toric duality as Seiberg duality

VII.7.1 One cone, several gauge theories

Chapter 2 built the dictionary from a toric Calabi–Yau threefold to its probe gauge theory and flagged, at once, that the dictionary is not a function. The forward map (a brane tiling, hence a quiver and a superpotential, gives a toric diagram) is well defined, but its inverse is not: a single toric diagram is the output of several inequivalent tilings. One cone, several gauge theories. Chapter 2 named this non-uniqueness *toric duality* and recorded that its simplest instance is $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$, whose centred-square diagram supports two genuinely distinct toric phases, while the conifold and dP_0 each support exactly one. It then deferred the question of what the non-uniqueness *means*, and the combinatorial move that relates the phases, to this chapter.

For the consistent regular-brane phases considered here, the physical identification is Seiberg duality. The different gauge theories are different ultraviolet descriptions of the same infrared physics. They are *Seiberg-dual*: distinct ultraviolet Lagrangians that flow to one and the same infrared superconformal field theory, the theory dual to $AdS_5 \times Y$ for the Sasaki–Einstein base Y of the cone. Toric duality is Seiberg duality, seen on the brane tiling. The combinatorial non-uniqueness of the inverse algorithm is exactly the infrared equivalence of Seiberg's electric and magnetic frames, and the field-theory operation that takes one phase to another, the gauging of a node's electric–magnetic duality, is realized on the tiling as a strictly local move on a single face. We call this move the *square move* (the dimer literature also calls it urban renewal or the spider move), and exhibiting it, with its invariants, is the work of the chapter.

This is worth pausing on, because it sharpens a contrast carried the whole way from Volume V. There, two different constructions could land on the same geometry: an orbifold read through the McKay correspondence and the same orbifold read through its toric dimer gave the same theory in two languages, a crosslink, not a new theory. Here the two phases are *genuinely different* gauge theories. In the F_0 example they have different numbers of fields, different quivers, and different superpotentials. And yet they flow to the same infrared fixed point and probe the same cone. The slogan "toric duality is Seiberg duality" is the precise statement that the inverse algorithm's

combinatorial ambiguity is the physical ambiguity of choosing an electric or a magnetic description of one infrared theory.

A word on what the chapter proves and what it cites, since the line matters. The chapter *states* the Seiberg-duality statement, the electric/magnetic pair for an $\mathcal{N}$=1 gauge factor, and uses it; the field-theory evidence (the matching of 't Hooft anomalies and chiral rings, the analysis of the conformal window) is the business of Part III, cited here, not derived. What the chapter *does*, the construction-intrinsic content that is its own, is twofold: exhibit the local square move on the tiling, use its sourced cone-preservation property, and compute necessary invariant comparisons, conditional on the physical duality. Matching a few invariants does not establish infrared equivalence. The single example worked in full is F_0; the next del Pezzo cases, dP_2 and dP_3, are reported as catalogued phase counts with their duality tree, the richest small example being dP_3.

The plan is: the rank-preserving toric window $N_f = 2N_c$ that makes a Seiberg move land back on a toric theory (§VII.7.2); the Seiberg-dual pair stated and read on the quiver (§VII.7.3); the square move as the dimer realization of the move (§VII.7.4); the worked example $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$, its two phases and the move between them (§VII.7.5); the invariants matched across phases, toric diagram, mesonic Hilbert series, and central charge (§VII.7.6); the del Pezzo phase counts and the duality tree (§VII.7.7); and the closing map of what was computed and what is deferred (§VII.7.8). Throughout, the organizing fact is the one Chapter 2 set up: the inverse map is one-to-many. The worked F_0 pair lies in one square-move orbit; for a general cone, reachable phases may form more than one component.

VII.7.2 The rank-preserving toric window

Before the move, the arithmetic that makes it possible. We restrict the dualized node to an SQCD-like factor with bifundamental quarks, no adjoint at that node, and balanced fundamental and antifundamental flavor counts. Neighboring groups are first treated as flavor symmetries; their gauging and the superpotential deformation are restored using the physical Seiberg-duality prescription. We take $N \geq 2$ for this nonabelian discussion; the single-brane mesonic ring below is a separate abelian construction. Seiberg duality acts on a single gauge node, and the reason a toric theory can be dual to another toric theory, rather than to something off the toric list entirely, is a constraint on which nodes can be dualized without changing the ranks. That constraint is the content of this section, and it is pure counting.

Recall the Seiberg-dual pair for an $\mathcal{N}$=1 gauge factor, stated here and discussed in Part III. An electric $SU(N_c)$ with N_f fundamental flavours is dual, in the infrared, to a *magnetic* theory with gauge group

$$SU(N_c) \text{ with } N_f \;\longrightarrow\; SU(N_f - N_c) \text{ with } N_f, \tag{VII.7.1}$$

the same number of flavours but a shifted rank, together with a gauge-singlet meson field M and a cubic coupling that we write down in §VII.7.3. The magnetic rank is $N_f - N_c$. For a generic node this is not equal to N_c, and the dual theory carries unequal ranks: it is no longer a regular-brane (equal-rank) configuration, hence in general no longer a toric phase of the same cone.

The toric phases of one cone are the equal-rank configurations, $N_i = N$ at every node, the regular-brane point. For the dual of a node to be again an equal-rank toric phase, the dualized node's rank must be preserved,

$$N_f - N_c \stackrel{!}{=} N_c \quad \Longleftrightarrow \quad N_f = 2N_c. \tag{VII.7.2}$$

At the equal-rank point $N_c = N$ this is the condition $N_f = 2N$, and the dual rank is

$$N_f - N_c = 2N - N = N = N_c, \tag{VII.7.3}$$

the rank is preserved. Remaining in a physical toric phase additionally requires the consistent square move and the transformed superpotential. This is the *toric (rank-preserving) window*: a node sits in the window when it carries exactly $N_f = 2N_c$ flavours, and only window nodes can be Seiberg-dualized within the toric class. It is the single fact behind the whole chapter, so it deserves to be stated as the reason the move exists at all: a square move can connect equal-rank phases because some nodes carry $N_f = 2N_c$ and can be dualized without leaving the equal-rank toric family.

Why a Seiberg move can stay toric

Here a toric phase is a consistent brane tiling with its superpotential at equal regular ranks, $N_i = N$. Equal ranks alone are insufficient. Seiberg duality sends a node $SU(N_c)$ with N_f flavours to $SU(N_f - N_c)$. The new rank equals the old one, $N_f - N_c = N_c$, if and only if $N_f = 2N_c$. Only such *window* nodes can be dualized inside the toric class; dualizing a node with $N_f \neq 2N_c$ changes the ranks and leaves the toric phase (it is the rank-shifting step of the cascade, Chapter 8). The window $N_f = 2N_c$ is the dimer's four-valent face.

The window has a combinatorial face on the tiling, and it is the link to the geometry that will run the rest of the chapter. The flavour count N_f at a node is set by how much matter the node sees, and in the dimer dictionary the matter charged under a face is carried by the edges bounding that face. A node with $N_f = 2N_c$ at equal rank is, on the tiling, a face bounded by exactly four edges: two incoming and two outgoing bifundamentals, each an $N \times N$ block, give $N_f = 2N$ flavours against $N_c = N$ colours. In general a face of valency $2k$ has k incoming and k outgoing bifundamentals, hence $N_f = kN$ flavours (one counts the flavour pairs, not every incident edge separately), so the window $N_f = 2N_c$ is exactly $k = 2$. The window node is the *four-valent* (square, quadrilateral) face. A face with more than four edges carries $N_f > 2N_c$ and is outside the window; a face with fewer carries $N_f < 2N_c$. So the question "which nodes can be Seiberg-dualized while staying toric" is, on the dimer, the question "which faces are quadrilaterals," and the move will act on a quadrilateral face.

The two boundary cases that Chapter 2 fixed are now explained, rather than merely recorded, by the window. The conifold's two nodes are each four-valent, $N_f = 2N_c$: they are in the window, and a Seiberg move on either is allowed, but, as we will see, it returns the same quiver, so the conifold is *self-dual* and has a single phase. The cone $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ has three nodes, each carrying $N_f = 3N_c$ (its adjacency is the directed triangle with three arrows per step, six edges per node), strictly outside the window: a Seiberg move on a dP_0 node sends $N_c = N$ to $N_f - N_c = 2N \neq N$, changes the ranks,

and produces no second equal-rank toric phase. dP_0 has no outgoing rank-preserving move from this phase. Its single-phase classification is a catalogued result, requiring more than that local observation. The local boundary arithmetic is

$$\text{conifold: } N_f = 2N_c \text{ (window, self-dual)}; \quad dP_0 : N_f = 3N_c \text{ (outside the window)}, \tag{VII.7.4}$$

and these two anchors will bracket the genuinely two-phase example, F_0, where a window node exists, the move is non-trivial, and the second phase is real.

VII.7.3 The Seiberg-dual pair on the quiver

With the window in hand, write the Seiberg move itself, first as field theory and then as an operation on the quiver, deferring the dimer realization to the next section. The electric/magnetic statement (VII.7.1) comes with matter, a superpotential coupling, and a prescription for the global charges; the toric setting specializes all three.

The magnetic theory of an electric $SU(N_c)$ with N_f flavours $q, \tilde{q}$ (the quarks and antiquarks) has, besides the dual quarks of the magnetic gauge group, a gauge-singlet *meson* field

$$M = \tilde{q}\, q, \tag{VII.7.5}$$

the electric mesonic bilinear promoted to an elementary field of the magnetic theory, with a holomorphic matching scale absorbed in the definition of M, and a cubic superpotential coupling it to the dual quarks $\hat{q}, \hat{\tilde{q}}$,

$$W_{\text{mag}} \;\supset\; M\, \hat{q}\, \hat{\tilde{q}}. \tag{VII.7.6}$$

At an interacting fixed point the superpotential has R-charge two, hence dimension three. The claim that the magnetic theory reproduces the electric 't Hooft anomalies and chiral ring belongs to Seiberg duality, stated and used here, supported by field-theory arguments discussed in Part III.

On a quiver node the quarks and antiquarks are the bifundamentals attached to the dualized node k. The duality acts on the interacting $SU(N_c)$ factor of the construction's $U(N_c)$ at that node; the decoupled diagonal $U(1)$ is a spectator, as throughout the volume. At a window node k the incoming arrows from neighbouring nodes are the electric quarks and the outgoing arrows the antiquarks (or the reverse; the move is symmetric in the two). At the regular ranks $N_i = N$, a four-valent window node carries two incoming and two outgoing $N \times N$ bifundamentals, so the flavour and colour counts and the magnetic rank are

$$N_c = N, \qquad N_f = 2N = 2N_c, \qquad N_f - N_c = 2N - N = N = N_c, \tag{VII.7.7}$$

the magnetic gauge group $SU(N_f - N_c) = SU(N)$ equal to the electric one: the node's rank survives the move, which is why the dual is again a regular-brane phase of the same cone. The Seiberg move then rewrites the quiver in three steps.

1. *Reverse the arrows at k.* Every arrow into k becomes an arrow out of k, and every arrow out of k an arrow in; the dual quarks of the magnetic node $SU(N_f - N_c)$ point the opposite way. On the adjacency matrix a_{ij} (the number of arrows $i \to j$) this exchanges the k-th row and the k-th column.

2. *Add the magnetic mesons.* For each incoming arrow $i \to k$ and outgoing arrow $k \to j$ of the electric quiver, the meson $M = \tilde{q}\, q$ is a composite that becomes an elementary field on a new arrow $i \to j$. The number of meson arrows added from i to j is the product $a_{ik}\, a_{kj}$ of the in- and out-multiplicities at k.

3. *Integrate out the massive fields.* The meson superpotential (VII.7.6) is cubic, and where a new meson arrow M on $i \to j$ runs antiparallel to an existing arrow X on $j \to i$, the two form a two-cycle. Such a pair can be massive only if the transformed superpotential contains a nondegenerate gauge-invariant quadratic

$$W \;\supset\; M\, X, \tag{VII.7.8}$$

a mass term for the pair (M, X). The F-terms of a massive pair set both fields to their solutions,

$$\frac{\partial W}{\partial M} = X + \cdots = 0, \qquad \frac{\partial W}{\partial X} = M + \cdots = 0, \tag{VII.7.9}$$

so M and X are eliminated algebraically (integrated out), removing one arrow in each direction and feeding the remaining quartic of the meson coupling back into the surviving fields. When the quadratic mass matrix has full rank, on the adjacency this cancels the matched arrows of an antiparallel pair, $a_{ij} \to a_{ij} - m$, $a_{ji} \to a_{ji} - m$ with $m = \min(a_{ij}, a_{ji})$.

The output is a new quiver, with a new field content and a new superpotential, for the same cone. The rules are entirely local to the dualized node and its neighbours, and they are involutive: a second move at the same node undoes the first, the magnetic of the magnetic being the electric, which the explicit example will confirm. Each surviving field still sits in exactly two superpotential terms, the toric two-terms-per-field condition, because the integrate-out step removes precisely the quadratic terms and the meson coupling supplies the cubic ones the dictionary requires.

The adjacency records possible mesons, but a two-cycle alone does not determine a mass. Both bifundamental and adjoint mesons require the actual superpotential. In F_0, the first move has no possible mass partner, and the second move has a full-rank quadratic pairing that we derive below. For the conifold the mesons are adjoints of the spectator node. The electric quartic becomes quadratic in these mesons; eliminating them recovers the conifold quartic after field redefinitions. Thus its self-duality is a sourced statement about W, not a consequence of the rank count or a catalog entry.

VII.7.4 Seiberg duality as the dimer square move

The quiver mutation of the previous section has a sharper and more local form on the brane tiling, where a gauge node is a *face* and Seiberg duality on a four-valent face is a rearrangement of the four edges around it. This is the local *square move* (the dimer literature also calls it urban renewal or the spider move), and it is the construction-intrinsic content the chapter owns: a combinatorial operation that takes

one consistent tiling of a cone to another consistent tiling of the same cone, with the field content changed and the toric diagram fixed.

Recall the dictionary, used throughout the volume: faces are gauge nodes, edges are bifundamentals, the black and white vertices are the two signs of superpotential terms, and the convex hull of the support of the Kasteleyn determinant is the toric diagram. A four-valent face is a quadrilateral bounded by four edges, the dimer face of a window node. The square move acts on such a face by replacing its four edges with a smaller square of dual-quark edges and four spokes carrying mesons. The new square has alternating black and white vertices, and each spoke joins opposite colors. The outer connections to the surrounding tiling remain attached. This creates the four magnetic cubic terms. If a superpotential vertex becomes bivalent, its quadratic term pairs two fields; solving their F-terms removes those fields and merges their other superpotential terms. The resulting bipartite graph must retain a consistent torus embedding.

Two properties make this a duality move rather than an arbitrary edit, with different evidential roles in the worked example: quiver involutivity is calculated below, whereas preservation of the toric polygon is imported from the square-move construction. First, the move is *involutive* at the face: repeating the move and eliminating the massive pairs undoes it, the dimer face of magnetic^2 = electric. Second, the move *preserves the Newton polygon* of the Kasteleyn determinant. The toric diagram is computed from the tiling as the support of $\det K(z,w)$, and the square move, being a local rearrangement that conserves the winding content of the edges around the face, leaves that convex hull invariant up to the $GL(2,\mathbb{Z})$ freedom of the toric lattice. The two tilings therefore have $GL(2,\mathbb{Z})$-equivalent toric diagrams: they probe the same cone. A move that changed the Newton polygon would not be a toric-duality move at all, since it would not preserve the geometry; the Newton-polygon invariance is a necessary geometric condition; physical admissibility also requires the superpotential and tiling consistency.

The square move

On a brane tiling, Seiberg duality at a four-valent (square) face is the local *square move*: replace the face by the dual square and four meson spokes, and eliminate any massive pair associated to a two-valent vertex. The move is involutive at the face and preserves the Newton polygon of $\det K$ (hence the cone). The dimer literature also calls it urban renewal or the spider move. This pins the *move* on the dimer dictionary; the static dictionary (faces, edges, vertices, Kasteleyn) is unchanged.

The calculation is exact on the adjacency, and it is worth doing the move once symbolically before the worked example renders it as a tiling, since the matrix step is what the finite matrix calculation checks. Take the square-necklace phase of F_0, the doubled four-cycle $0 \to 1 \to 2 \to 3 \to 0$, whose adjacency $a_{ij} = \#(i \to j)$ is the cyclic-shift-by-one matrix with entries 2,

$$a^{\mathrm{I}} = \begin{pmatrix} 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \\ 2 & 0 & 0 & 0 \end{pmatrix}, \qquad \#\text{fields} = \sum_{ij} a^{\mathrm{I}}_{ij} = 8. \tag{VII.7.10}$$

Dualizing the four-valent node $k = 1$ runs the three steps of §VII.7.3 on this matrix.

Step (i), *reverse the arrows at* $k=1$, exchanges the row and the column of node 1: the two $0\to1$ become two $1\to0$ ($a_{01}=2\mapsto a_{10}=2$) and the two $1\to2$ become two $2\to1$ ($a_{12}=2\mapsto a_{21}=2$). Step (ii), *add the magnetic mesons* $i\to1\to j$, deposits $a^{\mathrm{I}}_{01}\,a^{\mathrm{I}}_{12}=2\cdot2=4$ new arrows on $0\to2$, the only source–target pair flowing through node 1. Step (iii), *integrate out the two-cycles*, finds none, since $a^{\mathrm{I}}_{20}=0$ leaves the new $0\to2$ mesons no antiparallel partner to pair with. The output is the phase-II adjacency

$$a^{\mathrm{I}}\ \xrightarrow{\text{mutate } k=1}\ a^{\mathrm{II}}=\begin{pmatrix}0&0&4&0\\2&0&0&0\\0&2&0&2\\2&0&0&0\end{pmatrix},\qquad \#\text{fields}=\sum_{ij}a^{\mathrm{II}}_{ij}=12, \tag{VII.7.11}$$

the doubled four-cycle with node 1's arrows reversed and the quadruple meson arrow $0\to2$ added: eight fields become twelve. Reading the row and column sums of a^{II} gives the per-node (in,out) valencies

$$(\text{in},\text{out})=(4,4),\ (2,2),\ (4,4),\ (2,2)\ \text{at nodes } 0,1,2,3, \tag{VII.7.12}$$

so nodes $0,2$ are now eight-valent ($N_f=4N_c$, outside the window) and nodes $1,3$ stay four-valent; every node keeps incoming equal to outgoing, so the rank-weighted anomaly $(a^{\mathrm{II}}-a^{\mathrm{II}\mathsf{T}})\,\mathbb{1}$ vanishes and phase II is a regular-brane quiver, as a toric phase must be. Applying the same matrix move a second time at node 1 returns a^{I} exactly (the four $0\to2$ mesons of the first move pair with the four $2\to0$ mesons of the second and integrate out, $m=\min(4,4)=4$), so the move is involutive,

$$\mathrm{mutate}\big(\mathrm{mutate}(a^{\mathrm{I}},1),1\big)=a^{\mathrm{I}}, \tag{VII.7.13}$$

the matrix form of magnetic^2 = electric. This is the move the figure of §VII.7.5 draws on the tiling; the adjacency computation here is an exact finite calculation, and it is what the involutivity and field-count checks verify.

The post-move graph must still be a *consistent* tiling to be a physical phase, the dimer condition (each field in two terms of opposite sign, balanced zig-zag paths) that Chapter 3 imposed; a move whose output failed that filter would not produce a toric gauge theory, only a graph. In the worked example the inherited signed support polynomial has the F_0 Newton polygon, but is not certified as a physical Kasteleyn characteristic. The move is verified involutive at the quiver level; the second phase probes the same cone, so it shares that Newton polygon. The move is local, involutive, polygon-preserving, and consistency-respecting: the next section calculates the ranks, arrows, and superpotential for F_0, using the sourced tiling identification for the geometry.

VII.7.5 The worked example: $F_0=\mathbb{P}^1\times\mathbb{P}^1$

The simplest cone with two genuinely distinct toric phases is F_0, the zeroth Hirzebruch surface $\mathbb{P}^1\times\mathbb{P}^1$, whose complex cone is the Calabi–Yau singularity at hand. Its toric diagram is the centred square

$$\Delta_{F_0}=\mathrm{conv}\{(1,0),(0,1),(-1,0),(0,-1)\},\qquad \text{interior point } (0,0), \tag{VII.7.14}$$

four boundary lattice points and a single interior point. The two integers the chapter uses are read off it exactly as in Chapter 2,

$$G = 2\,\mathrm{Area}(\Delta_{F_0}) = 4 = \#\{\text{gauge nodes}\}, \qquad d = 4 = \#\{\text{boundary points}\}, \tag{VII.7.15}$$

so every phase of F_0 has four gauge nodes and four external (p,q)-web legs. The number $G = 4$ is a property of the diagram, shared by both phases; what differs between the phases is the field content, the number and arrangement of the bifundamentals on those four nodes.

Phase I: the doubled square necklace. The first phase is the four-node quiver whose arrows run once around a square with each arrow doubled,

$$a^{\mathrm{I}}_{01} = a^{\mathrm{I}}_{12} = a^{\mathrm{I}}_{23} = a^{\mathrm{I}}_{30} = 2, \tag{VII.7.16}$$

a closed necklace $0 \to 1 \to 2 \to 3 \to 0$ with two bifundamentals on each step, eight fields in all. Every node has two incoming and two outgoing arrows, hence is four-valent, $N_f = 2N_c$, and *every node of phase I sits in the window*. The gauge group is the construction's output $\prod_i U(N)_i = U(N)_0 \times U(N)_1 \times U(N)_2 \times U(N)_3$ at the regular-brane ranks, with interacting sector $\prod_i SU(N)_i$ and the decoupled diagonal $U(1)$; we never write it as a bare SU. The superpotential is the toric necklace word built from the doubled edges. Writing the two $i \to i{+}1$ fields as $X^{(1)}_{i,i+1}$ and $X^{(2)}_{i,i+1}$ (indices mod 4), it is the flavour-antisymmetrized sum of quartic loops around the square,

$$W_{\mathrm{I}} = \epsilon_{ac}\,\epsilon_{bd}\,\mathrm{Tr}\big(X^{(a)}_{01} X^{(b)}_{12} X^{(c)}_{23} X^{(d)}_{30}\big), \qquad a,b,c,d = 1,2, \tag{VII.7.17}$$

the two ϵ tensors pairing opposite edges, (01) with (23) and (12) with (30). Set $A_a = X^{(a)}_{01}$, $B_b = X^{(b)}_{12}$, $C_c = X^{(c)}_{23}$, $D_d = X^{(d)}_{30}$ and $\epsilon_{12} = 1$. Expanding the contraction gives four terms,

$$\begin{aligned} W_{\mathrm{I}} = {} & \mathrm{Tr}(A_1B_1C_2D_2) - \mathrm{Tr}(A_1B_2C_2D_1) \\ & - \mathrm{Tr}(A_2B_1C_1D_2) + \mathrm{Tr}(A_2B_2C_1D_1), \end{aligned} \tag{VII.7.18}$$

the doubled-square analogue of the conifold quartic (indeed F_0 is the $\mathbb{Z}_2$ quotient of the conifold, and W_{I} is the orbifold image of the Klebanov–Witten quartic). Each of the eight fields sits in exactly two of these terms with opposite sign, the toric two-terms-per-field condition that the brane tiling enforces. The standard square tiling identifies this superpotential with F_0. A separate inherited signed Laurent polynomial used for the support comparison is

$$P_{\mathrm{supp}}(z,w) = 2 + z + \tfrac{1}{z} - w + \tfrac{1}{w}, \qquad \mathrm{conv}(\mathrm{supp}\,P_{\mathrm{supp}}) = \Delta_{F_0}, \tag{VII.7.19}$$

the centred square (VII.7.14). The inherited eight-edge signed fixture has four central perfect matchings with contributions $+1, +1, -1, +1$, which cancel to the displayed coefficient 2. It therefore does not establish a physical Kasteleyn characteristic or the tiling's perfect-matching multiplicities. The physical tiling identification uses the sourced standard F_0 construction, separately from this support comparison.

The move: dualize node 1. Node 1 is in the window, so apply the Seiberg move of §VII.7.3 to it. This is exactly the adjacency mutation computed in §VII.7.4: reversing node 1's arrows ($0 \to 1 \Rightarrow 1 \to 0$, $1 \to 2 \Rightarrow 2 \to 1$) and adding the $a^{\mathrm{I}}_{01}\, a^{\mathrm{I}}_{12} = 4$ magnetic mesons $0 \to 2$ (no two-cycle to cancel) sends a^{I} of (VII.7.10) to the phase-II adjacency a^{II} of (VII.7.11), twelve fields, with the per-node valencies (VII.7.12) making nodes $0, 2$ eight-valent and leaving $1, 3$ in the window; phase II is again regular-brane anomaly-free. Figure VII.7.1 draws the two quivers; Figure VII.7.2 draws the square move on the square face.



Figure VII.7.1: The two toric phases of $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$, related by one Seiberg move on node 1. Phase I (left) is the doubled square necklace $0 \to 1 \to 2 \to 3 \to 0$, every node four-valent, eight fields; every node is in the window. Phase II (right) is obtained by dualizing node 1: its arrows reverse ($0 \to 1 \Rightarrow 1 \to 0$, $1 \to 2 \Rightarrow 2 \to 1$) and four magnetic mesons (blue, $2 \cdot 2$) appear on $0 \to 2$, making nodes $0, 2$ eight-valent and giving twelve fields. The node-label convention is the irrep/gauge index inside, the gauge factor $U(N)$ understood at each node. The schematic suppresses the multiplicities' precise routing.

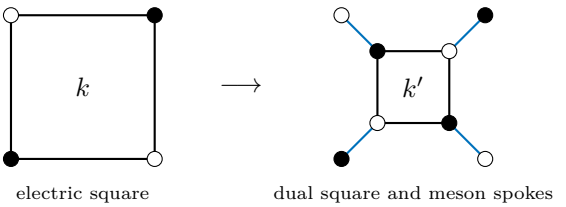


Figure VII.7.2: Local square move before eliminating any massive pair. Black and white dots are superpotential vertices, the inner square carries dual quarks, and the blue spokes carry mesons. The surrounding edges attached to the four outer vertices are suppressed. A global torus embedding and its winding data are required to determine a Kasteleyn matrix; this local drawing alone supplies neither.

Table VII.7.1 reads the move off field by field. Write $M_{ab} = A_a B_b : 0 \to 2$, $\alpha_a : 1 \to 0$, and $\beta_b : 2 \to 1$. A choice of dual-quark signs and meson normalization gives

the complete magnetic potential

$$W_{\rm II} = \sum_{a,b,c,d} \epsilon_{ac}\epsilon_{bd}\,\mathrm{Tr}(M_{ab}C_cD_d) + \sum_{a,b} \mathrm{Tr}(M_{ab}\beta_b\alpha_a). \tag{VII.7.20}$$

Every cubic is a closed path, $0 \to 2 \to 3 \to 0$ or $0 \to 2 \to 1 \to 0$. No quadratic term occurs on this first move. The displayed basis fixes the mesons literally as A_aB_b. To display opposite signs for each field, set $\alpha_a = s_a\widetilde{\alpha}_a$ and $\beta_b = -s_b\widetilde{\beta}_b$, with $s_1 = 1, s_2 = -1$. Since $L_{ab} = s_as_bC_{3-a}D_{3-b}$, the potential becomes $\sum_{ab} s_as_b\,\mathrm{Tr}[M_{ab}(C_{3-a}D_{3-b} - \widetilde{\beta}_b\widetilde{\alpha}_a)]$. Each field then occurs twice with opposite signs. We keep the original basis for elimination, since invertible sign changes do not change its F-term content.

arrow	phase I	phase II
$0 \to 1$ / $1 \to 0$	2 ($0 \to 1$)	2 ($1 \to 0$)
$1 \to 2$ / $2 \to 1$	2 ($1 \to 2$)	2 ($2 \to 1$)
$2 \to 3$	2	2
$3 \to 0$	2	2
magnetic mesons $0 \to 2$	0	4
total	8	12

Table VII.7.1: Field content across the node-1 Seiberg move on F_0. The arrows into and out of the dualized node 1 reverse, four magnetic mesons appear on $0 \to 2$, and no field is integrated out (no two-valent vertex), so eight fields become twelve. The superpotential gains the meson cubics; a second move on node 1 brings back phase I by the integrate-out (VII.7.21).

Involutivity, read on the superpotential. The adjacency computation (VII.7.13) already showed that a second move at node 1 returns phase I; it is worth seeing *why* the cancellation in that step is forced, since it is the two-cycle integrate-out and not a coincidence. The second move reverses node 1's arrows again (restoring $0 \to 1$, $1 \to 2$) and adds $2 \cdot 2 = 4$ mesons M' on $2 \to 0$, the source-to-target pair through node 1. Phase II already carries the four $0 \to 2$ fields M from the first move, so each new M' runs antiparallel to an M and the superpotential acquires the mass term

$$W \;\supset\; \sum_{r=1}^{4} \mathrm{Tr}\big(M_r\,M'_r\big), \qquad \frac{\partial W}{\partial M_r} = M'_r + \cdots = 0, \quad \frac{\partial W}{\partial M'_r} = M_r + \cdots = 0, \tag{VII.7.21}$$

four two-cycles with an invertible mass pairing. The full elimination keeps the terms hidden by the dots. Define $L_{ab} = \sum_{cd} \epsilon_{ac}\epsilon_{bd}C_cD_d$ and let $A'_a : 0 \to 1$, $B'_b : 1 \to 2$ be the twice-dual quarks. With $P_{ba} : 2 \to 0$ the new mesons, choose their cubic coupling with a minus sign,

$$W'' = \sum_{ab} \mathrm{Tr}(M_{ab}L_{ab} + M_{ab}P_{ba} - P_{ba}A'_aB'_b). \tag{VII.7.22}$$

The sign is absorbed in a dual-quark field redefinition. The cyclic derivatives give

$$P_{ba} = -L_{ab}, \qquad M_{ab} = A'_a B'_b, \qquad W''_{\text{eff}} = \sum_{abcd} \epsilon_{ac}\epsilon_{bd} \operatorname{Tr}(A'_a B'_b C_c D_d). \tag{VII.7.23}$$

Substitution cancels the ML and MP terms; the last term supplies the recovered quartic. This is why all four two-cycles may be removed in this fixture. It is a calculation in cyclic noncommuting words, so it applies to arbitrary matrix size with compatible quiver endpoints, not just commuting fields. Twelve fields return to eight,

$$\text{mutate}\big(\text{mutate}(\text{phase I}, 1), 1\big) = \text{phase I}, \tag{VII.7.24}$$

where equality includes the recovered superpotential up to field redefinitions. The algebra is conditional on the Seiberg-duality prescription; it does not derive that physical prescription.

VII.7.6 The matched invariants

The two descriptions have different matter counts, yet a dual pair must agree on infrared quantities. We compare their toric geometry, single-brane mesonic ring, and leading central charge. These are necessary consistency conditions. Their agreement alone is not a sufficient criterion for Seiberg duality.

The toric diagram. Both phases probe F_0, so both have the same toric diagram up to the $GL(2,\mathbb{Z})$ and translation freedom of the lattice. This is guaranteed by construction (the square move preserves the Newton polygon, §VII.7.4) and compared with the inherited signed support fixture. The polynomial of (VII.7.19) has

$$P_{\text{supp}}(z,w) = 2 + z + \tfrac{1}{z} - w + \tfrac{1}{w}, \qquad \text{conv}(\text{supp}\, P_{\text{supp}}) = \Delta_{F_0}, \tag{VII.7.25}$$

the centred square. As above, this signed polynomial is only a Newton-support comparison. The phase-II dimer is a different tiling, hence a different Kasteleyn characteristic, but the square move moved the winding content of the edges only locally, so the convex hull of the new support is the same polygon up to $GL(2,\mathbb{Z})$. We take this from the preservation duality statement (§VII.7.4) rather than print a second Kasteleyn determinant; no independent phase-II determinant is computed in the inherited checks. The shared diagram fixes the shared counts (VII.7.15), $G = 4$ nodes and $d = 4$ external legs, in both phases. The toric data is a phase-independent invariant, the most elementary one: it says the two theories live on the same geometry.

The mesonic Hilbert series. The graded ring of mesonic gauge-invariant operators, the coordinate ring of the cone, is a property of the geometry alone, not of the gauge-theory frame used to compute it. Its generating function, the mesonic Hilbert series, is therefore the same in every phase: the two phases have the isomorphic single-brane mesonic coordinate rings under the imported cone identification, so the same rational Hilbert series in a declared grading. It is computable from the shared toric diagram alone. The coordinate ring is the semigroup algebra of the *dual cone*. For rays $(1,0,1), (0,1,1), (-1,0,1), (0,-1,1)$, a dual lattice point (u,v,k) satisfies $k \pm u \geq 0$ and

$k \pm v \geq 0$. At height $k \geq 0$ it therefore lies in the square $[-k,k]^2$, giving $(2k+1)^2$ monomials. Consequently

$$H_{\text{mes}}(t) = \sum_{k\geq 0}(2k+1)^2 t^k = \frac{1+6t+t^2}{(1-t)^3} = 1+9t+25t^2+49t^3+\cdots . \qquad \text{(VII.7.26)}$$

Here degree one corresponds to an anticanonical section of $\mathbb{P}^1 \times \mathbb{P}^1$, of bidegree $(2,2)$; its nine sections independently recover the coefficient nine. Using $\sum kt^k = t/(1-t)^2$ and $\sum k^2 t^k = t(1+t)/(1-t)^3$ derives the full rational function. By contrast, the primal diamond has $2k^2+2k+1$ lattice points and Ehrhart series $(1+2t+t^2)/(1-t)^3$. That is a different count. The coordinate ring is determined by the cone identification, imported for each phase; we have not independently eliminated both phases' F-term ideals. Nor does equality of a Hilbert series by itself prove isomorphism of rings. At $k=0$ both counts give one constant, so this degenerate degree does not discriminate the two constructions; degree one does. The grading here is for one regular brane, with baryonic operators omitted.

The central charge. The conformal central charge a, equivalently the volume of the Sasaki–Einstein base Y, is fixed by the infrared fixed point, so it too is phase-independent. We can compute its value directly in the phase-I frame, which is all this chapter needs; the general a-maximization / volume-minimization *machinery* is Chapter 6. Phase I has $G=4$ gauge nodes and eight bifundamental chirals, and the $SU(2)^2$ flavor symmetry equates the members of each doublet. The symmetric conformal solution, including the discrete node symmetries and a-maximization, assigns all eight $R=\frac{1}{2}$, so at leading order in large N the trial central charge $a(R) = \frac{3}{32}(3\,\mathrm{Tr}\,R^3 - \mathrm{Tr}\,R)$ evaluates to

$$\frac{a(F_0)}{N^2} = \frac{3}{32}\Big[\,2\cdot 4 + 8\big(3(-\tfrac{1}{2})^3 - (-\tfrac{1}{2})\big)\Big] = \frac{27}{32}. \qquad \text{(VII.7.27)}$$

The geometry side returns the same number through the holographic relation of Chapter 6, which reads a off the base volume,

$$a = \frac{\pi^3 N^2}{4\,\mathrm{Vol}(Y)}, \qquad \mathrm{Vol}(Y_{F_0}) = \frac{8\pi^3}{27} \implies \frac{a(F_0)}{N^2} = \frac{\pi^3}{4}\cdot\frac{27}{8\pi^3} = \frac{27}{32}, \qquad \text{(VII.7.28)}$$

with $\mathrm{Vol}(Y_{F_0}) = 8\pi^3/27$ the Sasaki–Einstein volume of the F_0 base (the MSY volume function of Chapter 6 evaluated on the F_0 diagram, not a value Chapter 6 itself records). The conifold base $T^{1,1}$ has $\mathrm{Vol} = 16\pi^3/27$, twice the F_0 volume, so by the same relation

$$\frac{a(F_0)}{a(\text{conifold})} = \frac{\mathrm{Vol}(T^{1,1})}{\mathrm{Vol}(Y_{F_0})} = \frac{16\pi^3/27}{8\pi^3/27} = 2, \qquad a(F_0) = 2\cdot\tfrac{27}{64}N^2 = \tfrac{27}{32}N^2, \quad \text{(VII.7.29)}$$

a clean reflection of F_0 being the $\mathbb{Z}_2$ quotient of the conifold that doubles the relevant degrees of freedom. The point is that one number serves both phases: a is fixed by the shared cone, so phase I and phase II must return the same a under a-maximization, a necessary field-theory comparison for the stated duality. A separate extremization in phase II is not performed here.

What is shared and what differs across phases

Shared (the infrared SCFT / the cone): the toric diagram up to $GL(2,\mathbb{Z})$ (hence $G = 4$, $d = 4$), the mesonic Hilbert series, and the leading large-N central charge a / volume $\mathrm{Vol}(Y)$. *Differs* (the ultraviolet frame): the field content (for F_0, 8 vs 12 fields), the quiver, and the superpotential. Seiberg-dual phases must match the first list. This list alone is not sufficient to establish the duality.

The three comparisons share physical premises; they are three faces of the single statement that both ultraviolet theories flow to the dual of $AdS_5 \times Y$. The toric diagram says they live on the same cone, the Hilbert series compares the dimensions of the graded mesonic spaces, and the central charge says they have the same number of degrees of freedom at the fixed point. Field content is the one thing that may differ, and for F_0 it does, 8 versus 12, which is exactly why the two are distinct ultraviolet descriptions rather than a relabelling.

VII.7.7 The del Pezzo phase counts and the duality tree

F_0 is the first cone with more than one phase; the toric del Pezzo family continues the pattern, and the phases of a fixed cone organize into a graph. This section records the catalogued counts and records four phase nodes for the richest small example, dP_3, reporting rather than deriving every phase, since the full atlas of phase structure is the business of Chapter 11.

The phases of a cone are the nodes of its *duality tree*: a graph whose nodes are the toric phases and whose edges are square moves. The del Pezzo counts, recorded as literature data (Feng–Hanany–He–Uranga), are collected with their window status in Table VII.7.2. The two anchors have a single catalogued phase. The window arithmetic of §VII.7.2 distinguishes their local moves, while their phase counts remain sourced. The conifold sits in the window, $N_f = 2N_c$, so the rank is preserved, $N_f - N_c = N_c$, but the move returns the same quiver (the mesons are adjoints removed by the quartic W), so its tree is one self-dual node. The cone dP_0 has $N_f = 3N_c$ at every node, and a Seiberg move there shifts the rank off the equal-rank locus,

$$N_f - N_c = 3N_c - N_c = 2N_c \neq N_c, \tag{VII.7.30}$$

so this move produces no second equal-rank toric phase. The single-node catalog requires a separate classification input. For F_0, the first nontrivial tree has two nodes joined by the move of §VII.7.5 and its inverse.

The tree is built by a single rule. Its nodes are the toric phases; from any phase, dualizing each four-valent (window) face produces a neighbour, and an edge is drawn for each such move. A move that returns the same quiver (a self-dual node, as for the conifold) adds no new node, and a node outside the window ($N_f \neq 2N_c$) is not dualized at all, since that rank-shifting move leaves the equal-rank toric class (it is a possible rank-changing step; an RG cascade also requires running couplings, Chapter 8). Enumerating the phases is then a breadth-first search: start from one tiling, apply

cone	nodes G	#{phases}	window status
conifold	2	1	$N_f = 2N_c$ at every node, but self-returning
$dP_0 = \mathbb{C}^3/\mathbb{Z}_3$	3	1	$N_f = 3N_c$, no node in the window
$F_0 = \mathbb{P}^1 \times \mathbb{P}^1$	4	2	a window node, one non-trivial move
dP_1	4	1	also $G = 4$, but a single toric phase
dP_2	5	2	window nodes, two phases
dP_3	6	4	window nodes, four phases

Table VII.7.2: Toric phase counts for the small del Pezzo cones (catalogued literature data, Feng–Hanany–He–Uranga). No monotone rule relating phase count to the number of window nodes is implied; the conifold and dP_0 are the single-phase boundary anchors, for the two distinct reasons of (VII.7.4). That dP_1 shares $G = 4$ with F_0 yet has a single phase shows that the node number G alone does not fix the count. No exhaustiveness claim follows from the window count.

the square move at each window face, keep the inequivalent tilings with their cyclic superpotentials, and repeat until no new phase appears. The conifold and dP_0 terminate at once (no window move yields a new node); F_0 closes after the single move; dP_2 and dP_3 run a few rounds before closing on two and four phases.

> Duality tree (breadth-first). (1) Seed a queue with one consistent tiling. (2) Pop a phase; for each four-valent (window) face apply the square move to produce a neighbour. (3) Discard any neighbour equivalent as an embedded tiling with superpotential to a phase already seen, queue the rest, and draw an edge per move. Repeat until the queue empties; this enumerates the reachable component only. Polygon equivalence cannot distinguish phases of one cone, and exhaustiveness requires a separate classification argument.

dP_3, with four phases, is the richest example small enough to draw. Its six-node diagram has enough four-valent faces that dualizing different window nodes reaches four inequivalent toric phases, and the square moves connect them into a closed graph: every phase is reachable from any other by a sequence of moves, so the four phases form one orbit of toric duality. The detailed phase identifications (which four-valent node, dualized in which phase, lands on which other phase) are the business of the phase atlas, Chapter 11; here we record the count and the connectivity. Figure VII.7.3 draws the catalogued phase nodes, with only the F_0 edge exhibited here.

A caution closes the section, to keep the static tree of this chapter apart from the dynamics of the next. The duality tree here is a graph of *equal-rank* phases, connected by the rank-preserving window moves of §VII.7.2. Allowing the moves that change the ranks, dualizing a node outside the window, can, with appropriate running couplings and scales, produce a flow: a sequence of Seiberg dualities with shifting ranks, the duality cascade, which runs the theory down an energy scale rather than relabelling

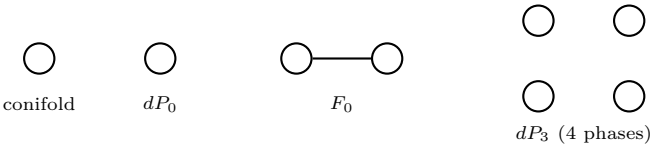


Figure VII.7.3: Duality trees: nodes are toric phases, edges are square (Seiberg) moves. The conifold and dP_0 are single-phase (one node, no edge). F_0 has two phases joined by one move. The four dP_3 dots record the catalog count; no edge pattern is asserted for them in this drawing. The counts (Table VII.7.2) are catalogued literature data; the static tree is shown here, the iterated *rank-shifting* version (the cascade) is Chapter 8, and the full phase atlas is Chapter 11.

it at a fixed one. The single equal-rank move and the static tree are this chapter's; the iterated, rank-shifting move and the renormalization-group flow it generates are Chapter 8.

VII.7.8 What has been computed, and what is deferred

The chapter has given the non-uniqueness of Chapter 2 its physical content. A single toric Calabi–Yau threefold admits, in general, several inequivalent toric phases, and those phases are Seiberg-dual gauge theories, distinct ultraviolet descriptions of one infrared superconformal field theory. The rank-preserving toric window $N_f = 2N_c$, the four-valent dimer face, is what lets a Seiberg move stay inside the equal-rank toric family; the move itself is the square move, a local, involutive, Newton-polygon-preserving rearrangement of a square face; and the duality is supported by three necessary invariant comparisons, the toric diagram, the mesonic Hilbert series, and the central charge a / $\mathrm{Vol}(Y)$, while the field content differs. The ranks and superpotential were calculated explicitly on $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$, whose two phases (8 and 12 fields) are related by one move on node 1, with dP_2 and dP_3 reported as catalogued phase counts and its four catalogued phase nodes drawn without claiming an edge list.

A reader leaving the chapter should be able to do the following without looking ahead: state the Seiberg-dual pair $SU(N_c)$ with N_f flavours $\to SU(N_f - N_c)$ with a meson M and the coupling $M\hat{q}\hat{\tilde{q}}$, and identify the rank-preserving window $N_f = 2N_c$ as the four-valent dimer face; run the three-step quiver mutation (reverse the arrows, add the $a_{ik}a_{kj}$ magnetic mesons, eliminate pairs with nondegenerate quadratic terms) on a window node; perform the square move on a window face of a tiling and see that it is involutive and Newton-polygon-preserving; build F_0's phase I, dualize node 1 to phase II ($8 \to 12$ fields), and check that a second move returns phase I; separate the sourced conifold self-duality and catalogued dP_0 count from the window arithmetic; and list the three necessary invariant comparisons without treating them as a proof of infrared equivalence.

What has been deferred, and to where, is the rest of the duality story.

- The Seiberg-duality *duality statement* itself, the matching of 't Hooft anomalies and chiral rings and the conformal-window analysis, is stated and used here but discussed in Part III (the field-theory derivation of Seiberg duality). The chapter's claim is the construction-intrinsic content, the move and the invariants, not the field-theory duality statement.

- The central charge a / Vol(Y) used as the "same a" invariant is computed locally here (the phase-I evaluation at the symmetric conformal charges, $a(F_0)/N^2 = 27/32$); the general a-maximization / volume-minimization machinery is Chapter 6. The chapter asserts only its phase-independence, since both phases flow to the same fixed point.

- The *iterated* move with rank shifts, the duality cascade, the fractional branes, and the deformed-conifold infrared, are Chapter 8. This chapter gives only the single equal-rank move and the static duality tree; the cascade turns that tree into a renormalization-group flow.

- The square move is, mathematically, a quiver *mutation*, the cluster-algebra and derived-category (tilting of the exceptional collection) reading of the move; that homological machinery is named here and developed in Volumes III–IV (it extends the derived-McKay material, Volume VI's the threefold derived McKay correspondence).

- Which phase is the correct description in which region of Kähler/Fayet–Iliopoulos moduli space, the marginal-stability and wall-crossing story, is named and routed to Part III and Volume V; the chapter does not analyze the stability of the dual phases.

- The full atlas of phase structure across cones, beyond the del Pezzo counts recorded here, is Chapter 11; the conifold's adjoint-meson self-duality, imported from its superpotential dualization rather than inferred from the rank count, is the one case where the W-dependent integrate-out is not tracked by the adjacency alone.

The single lesson to carry forward is the one in the opening: the inverse map from a cone to its gauge theory is one-to-many. The worked F_0 phases form one orbit of square moves; the general phase set is a union of reachable components unless connectivity is established separately. The next chapter lets the move change the ranks and watches the resulting cascade run.

Status and sources

The chapter's physical duality and toric-geometry identifications are source-supported premises. The chapter derives the finite algebraic consequences of those premises. [WS] means a controlled physical-worldsheet derivation in the untwisted string; quiver arithmetic and worldvolume F-term algebra do not supply that derivation. [MATH] is reserved here for the explicitly stated finite algebra and lattice-count implications with their hypotheses, not the physical Seiberg duality.

Sections 7.1–7.3. The Seiberg prescription is used as physical input, following Seiberg 1995. The local application to a gauged quiver node, including deformation of W, is explained in Franco et al.

2006, section 6.1. Toric duality is developed in Feng, Hanany, and He 2001a; Beasley and Plesser 2001; Feng, Hanany, He, and Uranga 2001. Rank subtraction and flavor counting are conditional arithmetic consequences, not a proof of duality. Conventions 2, 5, 6, 17 and The square move uses the tiling and gauge conventions stated above.

Sections 7.4–7.5. The physical square move uses Franco et al. 2006, sections 5–6 and Figure 14. The corrected opposite-edge contraction in (VII.7.17) agrees with Feng, Hanany, He, and Uranga 2001, section 4.1, equation (4.1), after the relabeling $(X, Y, Z, W) = (A, B, C, D)$; that source calls the eight-field phase II, whereas this chapter calls it I. Equations (VII.7.20)–(VII.7.23) give a direct cyclic-word derivation in a declared sign convention. The inherited signed polynomial's Newton support is checked in the original calculations. Its four central matching contributions cancel to coefficient two; a physical Kasteleyn characteristic, perfect-matching multiplicities, torus embedding, and an independent phase-II determinant are not established by those checks. The physical cone identification is imported from the standard construction. Convention 7 sets the lattice and convention 18 the dimer dictionary.

Section 7.6. The dual-cone inequalities and summation in (VII.7.26) prove the Hilbert series of the stated semigroup algebra. Identifying this algebra with the physical single-brane mesonic ring is a separate premise. The separate perfect-matching count gives lattice points in the primal diamond, giving $(1 + 2t + t^2)/(1 - t)^3$; its successful arithmetic does not verify the corrected mesonic series. Both phase F-term eliminations remain unperformed. Equations (VII.7.27)–(VII.7.29) use the leading large-N anomaly formula and the classical holographic volume relation, with Chapter 6's convention 21 and Intriligator and Wecht 2003. For the interacting $SU(N)^4$ theory at the stated charges the finite-rank correction is $a = 27N^2/32 - 3/4$; the displayed holographic ratio is only the leading term. Neither a phase-II extremization nor equality of all physical observables is calculated here.

Sections 7.7–7.8. The phase counts in Table VII.7.2 are catalogued, not enumerated by this chapter. In particular Feng, Hanany, He, and Uranga 2001, sections 6.1–6.3 and Figure 12, supports the two F_0, two dP_2, and four dP_3 phases; the earlier Feng, Hanany, and He 2001b alone does not provide that complete dP_3 catalog. A complete phase graph, global exhaustiveness, stability chambers and the geometric interpretation of mutations are left to the phase atlas and later treatments, including Cachazo et al. 2001. the field-theory derivation of Seiberg duality routes the Seiberg-duality evidence to Part III; the threefold derived McKay correspondence routes the derived-category extension. The atlas is Chapter 11; rank-shifting cascades are Chapter 8 and require additional running-coupling dynamics.

Calculations and companion. The exercises cover three submitted-answer families: fixed-node rank and matter counts, cyclic magnetic superpotential elimination, and dual-cone lattice counts with a rational generating function. It binds the named equations and records all other formulas as text identities only. The exercise calculations leave the physical premises separate.

Related constructions. The local meson signs above follow from the chosen field definitions. Chapter 2 discusses why distinct toric phases can give the same geometry.

References

Beasley, C. E. and M. R. Plesser (2001). "Toric duality is Seiberg duality". In: arXiv: `hep-th/0109053`.

Cachazo, F. et al. (2001). "A Geometric unification of dualities". In: arXiv: `hep-th/0110028`.

Feng, B., A. Hanany, and Y.-H. He (2001a). "D-brane gauge theories from toric singularities and toric duality". In: *Nucl. Phys. B* 595, pp. 165–200. DOI: `10.1016/S0550-3213(00)00429-7`. arXiv: `hep-th/0003085`.

Feng, B., A. Hanany, and Y.-H. He (2001b). "Phase structure of D-brane gauge theories and toric duality". In: arXiv: `hep-th/0104259`.

Feng, B., A. Hanany, Y.-H. He, and A. M. Uranga (2001). "Toric duality as Seiberg duality and brane diamonds". In: arXiv: hep-th/0109063.

Franco, S. et al. (2006). "Brane dimers and quiver gauge theories". In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Intriligator, K. A. and B. Wecht (2003). "The Exact superconformal R symmetry maximizes a". In: *Nucl. Phys. B* 667, pp. 183–200. DOI: 10.1016/S0550-3213(03)00459-0. arXiv: hep-th/0304128.

Seiberg, N. (1995). "Electric - magnetic duality in supersymmetric nonAbelian gauge theories". In: *Nucl. Phys. B* 435, pp. 129–146. DOI: 10.1016/0550-3213(94)00023-8. arXiv: hep-th/9411149.

Chapter VII.8

Duality cascades and conifold deformation

VII.8.1 From the static brane to its dynamics

Every chapter of this volume so far has placed a *regular* brane. A stack of N D3-branes at a toric Calabi–Yau singularity carries the same rank at every gauge node. That rank is $N_i = N \dim R_i$ for the orbifolds of Volume VI, and $N_i = N$ at each face of a brane tiling here. The regular brane is the mobile one, the one whose mesonic moduli space reconstructs the transverse cone, and its equal-rank quiver satisfies the inherited nonabelian rank-balance equation for every N. The whole construction, geometry to quiver to dimer to moduli space, was built on it.

This chapter turns the static construction into dynamics by relaxing exactly one assumption: the ranks need not be equal. An allowed *unequal*-rank assignment of the same quiver records fractional-brane charge and can set the gauge theory in motion. The anomaly constraint and the physical brane dictionary determine which assignments can enter this description. The two gauge couplings of the conifold quiver, tied together at the symmetric Klebanov–Witten point of Chapter 1, come unstuck when the ranks differ; one node runs to strong coupling, and a useful description of the same infrared physics is obtained by trading it for its Seiberg dual. That is the square move of Chapter 7, the single combinatorial step that replaces one dimer of a toric singularity by another. Run once, it is a duality. Run again and again as the couplings flow, it is a *duality cascade*: a staircase of Seiberg dualities whose ranks step down by a fixed amount at each landing.

For the specified confining conifold branch with $N = kM$, $k \geq 1$ and $M \geq 2$, the cascade has a finite stopping stage. Its relation to the deformed geometry is the physical identification developed in this chapter. At the bottom the effective rank reaches the scale where the gauge theory confines, and the geometry responds: the conifold *deforms*, its singular tip replaced by a three-sphere of finite size. Confinement and chiral-symmetry breaking are named by reading that deformed geometry, and the deformation parameter is fixed in scale by the gaugino condensate. This is the

Klebanov–Strassler theory, and it is the destination of the forward references that the conifold chapter left open.

The calculation begins with the anomaly-free rank lattice and the flavor count at a strongly coupled node. These determine candidate ranks, but they do not determine a vacuum. For the conifold we use the Klebanov–Strassler physical description to connect the repeated rank step to confinement and a deformed geometry. The fractional-brane types and the logarithmically running supergravity charge are additional physical inputs. We will separate these inputs from the finite calculations that make their consequences explicit. The warped metric, flux solution, and glueball, baryon, and Kaluza–Klein spectra belong to Part III.

Take-home message of this chapter

A regular brane has equal node ranks $N_i = N$. The anomaly-free fractional-rank classes form $\ker_{\mathbb{Z}}(a - a^{\mathsf{T}})/\mathbb{Z}(1, \ldots, 1)$. For the conifold, one primitive class gives ranks $(N+M, N)$. When the Seiberg description applies, the flavor count $N_f = 2N$ yields $(N, M) \mapsto (N - M, M)$. For $N = kM$ with $k \geq 1$ and $M \geq 2$, this arithmetic stops at $(2M, M)$, where the physical quantum constraint and the choice of confining branch replace the ordinary rank step. The Klebanov–Strassler identification relates that branch to the deformed conifold and its finite S^3.

The plan: the fractional brane as an anomaly-free rank deformation, with the anomaly-free condition derived and counted (§VII.8.2); the three geometric types, deformation, $\mathcal{N}$=2, and dynamical-supersymmetry-breaking, classified by the wrapped cycle (§VII.8.3); the duality cascade as iterated Seiberg duality, with the conifold rank step computed and its leading-log running stated (§VII.8.4); the deformed-conifold infrared, confinement and chiral-symmetry breaking identified by the geometry, with the deformation set apart from the resolution (§VII.8.5); the cascade for a general toric singularity, named and given its dimer-move picture (§VII.8.6); and a closing map of what was computed and what is deferred (§VII.8.7).

VII.8.2 The fractional brane as an anomaly-free rank deformation

A quiver gauge theory with gauge group $\prod_i U(N_i)$ is specified, beyond its arrows and superpotential, by its rank vector $N = (N_1, \ldots, N_G)$. The construction outputs the regular brane, the equal ranks

$$N_i = N \qquad (\text{all } i), \tag{VII.8.1}$$

for a tiling with G faces. A *fractional brane* is a different choice of rank vector, an integer deviation from (VII.8.1). Not every deviation is allowed: an unequal assignment generically suffers a gauge anomaly, and the consistency condition that survives is the nonabelian condition used in Volume V.

The anomaly-free condition. Record the quiver by its adjacency matrix, a_{ij} the number of arrows from node i to node j. A chiral bifundamental on the arrow $i \to j$ is a

fundamental of $SU(N_i)$ and an antifundamental of $SU(N_j)$. The cubic gauge anomaly of the factor $SU(N_i)$ counts the chiral fermions in its fundamental against those in its antifundamental, each weighted by the rank of the node at the other end of the arrow,

$$\mathcal{A}\big(SU(N_i)\big) \;\propto\; \sum_j a_{ij}N_j \;-\; \sum_j a_{ji}N_j \;=\; \big[(a-a^{\mathsf{T}})\,N\big]_i, \tag{VII.8.2}$$

the first sum the outgoing arrows (fundamentals of node i, multiplicity N_j) and the second the incoming arrows (antifundamentals). For occupied $SU(N_i)$ nodes with $N_i \geq 3$, no fundamental flavor sectors, and only bifundamental or adjoint chiral matter, this perturbative cubic anomaly vanishes precisely when the rank vector lies in the kernel of the *antisymmetric* part of the adjacency,

$$(a-a^{\mathsf{T}})\,N = 0. \tag{VII.8.3}$$

The regular brane (VII.8.1) is the all-ones vector $N = (1,\dots,1)$ times N; it always solves (VII.8.3), because $\sum_j (a_{ij}-a_{ji})$ is the net number of arrows at node i, which vanishes for the balanced quivers that the dimer construction produces (each node has equal incoming and outgoing arrows, the toric condition). The anomaly-free *fractional* branes are the *other* integer solutions of (VII.8.3). Counting them is a kernel dimension:

$$\#\{\text{independent fractional-brane directions}\} = \dim\ker(a-a^{\mathsf{T}}) - 1, \tag{VII.8.4}$$

where the dimension is taken over $\mathbb{Q}$. More precisely, put $K = \ker_{\mathbb{Z}}(a-a^{\mathsf{T}})$ and $r = (1,\dots,1)$. The fractional-rank lattice is $K/\mathbb{Z}r$. The vector r is primitive in K: if $kv = nr$ for an integer vector v, then all components of v are the same integer. Thus the quotient has no torsion and its rank is the right-hand side of (VII.8.4). Physical ranks must also be nonnegative. This lattice condition alone does not enforce stability, the full string charge constraints, or the existence of a supersymmetric vacuum. For rank-one, rank-two, or absent nodes, we use it as the inherited rank constraint, rather than an equivalence with all gauge anomalies; abelian and global anomalies require separate analysis.

Throughout this chapter *fractional brane* means an anomaly-free rank deformation of the local four-dimensional quiver, a solution of (VII.8.3) other than the regular direction. This is narrower than the existence of a fractional-brane *charge* in the stringy charge lattice, the orbifold / irrep sense of Volume VI, in which $\mathbb{C}^3/\mathbb{Z}_3$ certainly carries fractional-brane charges. A rank vector with $(a-a^{\mathsf{T}})N \neq 0$ may still label a charge sector in a compact or flavour-completed setup; it is just not, by itself, a consistent isolated local gauge theory.

The kernel of an antisymmetric integer matrix is computed by exact integer linear algebra, and the fractional-brane direction records part of a fractional D-brane charge in the local charge lattice of the singularity. In the ALE examples of Volume VI that charge lattice was visible as $H_0 \oplus H_2$, the point charge plus the wrapped two-cycles. The regular brane was the marks-weighted combination, carrying *pure* H_0. For a toric Calabi–Yau threefold the charge lives more naturally in the local K-theory / even-cohomology data, and it may carry higher even-dimensional (H_4, i.e. D7) components. So it is *not* simply $H_0 \oplus H_2$ in general. The conifold is the clean case: its fractional

brane is a D5-brane wrapping the vanishing two-sphere of the small resolution, the familiar H_2 charge. What the quiver anomaly calculation (VII.8.3) sees is the part of that charge data the rank vector records. The rank-deviation and the wrapped charge are the same object seen from the gauge theory and the geometry. But for a general toric cone the full charge is richer than the $H_0 \oplus H_2$ of the ALE case.

The three fixtures. Work the condition on the volume's three running examples. Table VII.8.1 collects the adjacencies, their antisymmetric parts, and the resulting counts.

singularity	G	$a - a^{\mathsf{T}}$	fractional directions
$\mathbb{C}^3$	1	(0)	0
conifold	2	$\left(\begin{smallmatrix} 0 & 0 \\ 0 & 0 \end{smallmatrix}\right)$	1
$dP_0 = \mathbb{C}^3/\mathbb{Z}_3$	3	$\left(\begin{smallmatrix} 0 & 3 & -3 \\ -3 & 0 & 3 \\ 3 & -3 & 0 \end{smallmatrix}\right)$	0

Table VII.8.1: The anomaly-free fractional rank-deformation count for the three fixtures, from $\dim\ker(a - a^{\mathsf{T}}) - 1$. $\mathbb{C}^3$ has one node and no rank-deformation direction. The conifold has $a - a^{\mathsf{T}} = 0$, so its kernel is all of $\mathbb{Z}^2$ and there is exactly one fractional direction beyond the regular one. dP_0 has a rank-two antisymmetric adjacency, so its kernel is the regular direction alone: *no* anomaly-free rank-deformation direction.

$\mathbb{C}^3$. One gauge node, no arrows between distinct nodes, $a = a^{\mathsf{T}} = (0)$. The kernel is one-dimensional and is the regular direction; there is no fractional rank deformation, as there is no cycle to wrap.

The conifold. The Klebanov–Witten quiver has two nodes with $a_{12} = a_{21} = 2$ (the two A_i and the two B_j), so the adjacency is already symmetric and

$$a - a^{\mathsf{T}} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \qquad \ker(a - a^{\mathsf{T}}) = \mathbb{Z}^2, \tag{VII.8.5}$$

the full rank lattice. *Every* rank assignment (N_1, N_2) is anomaly-free: the conifold quiver satisfies the inherited nonabelian balance at any ranks. The regular direction is $(1, 1)$; one independent fractional direction remains, the relative shift

$$(1, 0) \pmod{\mathbb{Z}(1, 1)}, \qquad (N_1, N_2) = (N + M,\, N), \tag{VII.8.6}$$

M primitive fractional charges on top of N regular ones. Indeed, the homomorphism $(a, b) \mapsto a - b$ has kernel $\mathbb{Z}(1, 1)$ and maps $(1, 0)$ to one. The allowed vector $(1, -1)$ maps to two, so it cannot replace $(1, 0)$ as an integral quotient generator. This is the single fractional class of the conifold, and it is the brane that drives the Klebanov–Strassler cascade.

dP_0. The directed-McKay quiver of $\mathbb{C}^3/\mathbb{Z}_3$ has three nodes with $a_{ij} = 3\,\delta_{j,\,i+1}$ (three arrows from each node to the next, cyclically), the chiral, oriented quiver of Volume V.

Its antisymmetric adjacency is the circulant

$$a - a^{\mathsf{T}} = \begin{pmatrix} 0 & 3 & -3 \\ -3 & 0 & 3 \\ 3 & -3 & 0 \end{pmatrix}, \tag{VII.8.7}$$

whose rank is two, not three; the unique zero eigenvector is the all-ones direction $(1, 1, 1)$. The kernel is the regular brane *alone*, one-dimensional, so by (VII.8.4) the number of independent fractional-brane directions is $\dim \ker(a - a^{\mathsf{T}}) - 1 = 1 - 1 = 0$. *The dP_0 singularity admits no anomaly-free fractional-brane rank deformation of the isolated local four-dimensional quiver*, without adding extra sectors or tadpole cancellation. Any rank shift away from the regular vector violates the inherited balance equation; when all occupied nodes have rank at least three it turns on a perturbative cubic gauge anomaly.

This last result is worth pausing on, because it is the sharpest illustration of what a fractional brane is. One might guess that any unequal-rank assignment of a quiver is a fractional brane, that fractional ranks are freely adjustable knobs. They are not. A fractional brane is a *constrained* object, a vector in a geometry-determined kernel, and for dP_0 that kernel happens to contain only the regular brane. The reason is chirality: the strongly oriented dP_0 quiver has so much net arrow flow between nodes that no unequal assignment can balance the cubic anomaly. The conifold, with its symmetric two-way arrows, has no net flow, so it tolerates any ranks. dP_0, maximally chiral, tolerates none. The cascade worked examples of this chapter are therefore the conifold (one deformation-type fractional brane, the Klebanov–Strassler spine) and the other del Pezzo and Sasaki–Einstein cones of §VII.8.6, where a nonzero fractional direction does exist. dP_0 is the cautionary fixture, not a cascading one.

VII.8.3 The three types of fractional brane

A nonzero anomaly-free fractional direction exists; the question is what the brane wrapping it *does* in the infrared. The answer is not uniform. Following Franco–Hanany–Saad–Uranga, an anomaly-free fractional brane falls into one of three types, distinguished by the geometry of the cycle it wraps and by the dynamics that geometry triggers. The classification is the organizing statement of the chapter's middle, and the conifold realizes exactly one of the three; we work a concrete example of each, the suspended pinch point for the $\mathcal{N}$=2 type and the cone over dP_1 for the DSB type.

What there is to classify. Before naming the types, fix how many fractional directions there are to classify, because the count is itself the kernel computation of §VII.8.2, read for the two fixtures that matter here. The conifold's symmetric two-way adjacency $a_{12} = a_{21} = 2$ has

$$a - a^{\mathsf{T}} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \qquad \ker(a - a^{\mathsf{T}}) = \mathbb{Z}^2 = \langle (1,1),\, (1,0) \rangle_{\mathbb{Z}}, \tag{VII.8.8}$$

so beyond the regular $(1, 1)$ there is exactly one primitive fractional class, represented

by $(1,0)$: *one* brane to classify. The maximally chiral dP_0 has instead

$$a - a^{\mathsf{T}} = \begin{pmatrix} 0 & 3 & -3 \\ -3 & 0 & 3 \\ 3 & -3 & 0 \end{pmatrix}, \qquad \text{rank} = 2, \qquad \ker(a - a^{\mathsf{T}}) = \langle (1,1,1) \rangle, \quad \text{(VII.8.9)}$$

the regular direction alone: *zero* branes to classify (the cautionary fixture, no anomaly-free rank shift exists). Table VII.8.2 tabulates the count and, for the conifold's one direction, the type. The physical classification concerns specified brane configurations and their dynamics. It is not a linear labeling of a lattice basis: mixtures of supersymmetric fractional branes can have different dynamics. The conifold class supplies the deformation example.

singularity	$\dim \ker(a - a^{\mathsf{T}})$	fractional directions	type(s)
conifold	2	1 (the class $[(1,0)]$)	deformation
$dP_0 = \mathbb{C}^3/\mathbb{Z}_3$	1	0	none (no rank deformation)

Table VII.8.2: The anomaly-free fractional directions to classify, from $\dim \ker(a - a^{\mathsf{T}}) - 1$ (VII.8.8)–(VII.8.9), and their geometric type. The conifold has one direction, of *deformation* type (a smoothable S^3, confining); dP_0 has none, so there is nothing to classify. Other toric cones have further directions; their types require the geometry and superpotential.

The three types, read from the geometry of the cycle each anomaly-free direction wraps, are the following.

1. **Deformation type.** The collapsed cycle the brane wraps is one the geometry can smooth by a *complex-structure deformation*, opening a finite cycle of one real dimension higher (an S^3 for the conifold). The gauge theory *confines*, and the deformation is the geometric face of confinement. The conifold's fractional brane is the standard example: the deformed conifold and the Klebanov–Strassler infrared of §VII.8.5.

2. **$\mathcal{N}$=2 type.** The brane sits along a *non-isolated* locus of the singularity, a line of $\mathbb{C}^2/\mathbb{Z}_n$ singularities rather than an isolated point. Along such a line the local geometry has eight supercharges, not four, and the fractional brane probes a local $\mathcal{N}$=2 sub-sector. Its dynamics opens an $\mathcal{N}$=2 Coulomb branch, a family of supersymmetric vacua along which the brane can move. A noncompact flat direction is not by itself a potential-driven runaway. The contrast that names this type is the supercharge count: $\mathcal{N}$=1 generically, $\mathcal{N}$=2 along the non-isolated locus.

3. **Dynamical-supersymmetry-breaking (DSB) type.** In the examples in this class, the tree superpotential and strong gauge dynamics are incompatible with a supersymmetric vacuum at finite field values. An Affleck–Dine–Seiberg term can cause a runaway once the relevant baryonic or FI directions are allowed to vary.

Too few flavors alone does not prove this behavior; pure supersymmetric Yang–Mills is a counterexample to that inference. A runaway also does not establish a stable nonsupersymmetric ground state. This is the type Franco–Hanany–Saad–Uranga identified; the supersymmetry-breaking model-building that uses it is a pointer, out of this volume's scope.

A worked example of each type. Each type has a concrete realization. The quiver and geometry identify the candidate configuration; its dynamical label uses the physical analysis of that configuration. The deformation type is the conifold itself (§VII.8.4–§VII.8.5); here are the other two.

The $\mathcal{N}$=2 type: the suspended pinch point. A toric Calabi–Yau *two*-fold $\mathbb{C}^2/\mathbb{Z}_n = A_{n-1}$ has a *one*-dimensional toric diagram, a segment of lattice length n. Its fingerprint inside a threefold is therefore an *edge* of the toric polygon of lattice length greater than one: a boundary lattice point that is *not* a corner marks a transverse $\mathbb{C}^2/\mathbb{Z}_n$ line, an embedded A_{n-1} surface running along that edge. The conifold and dP_0 have none (every boundary point is a corner, the singularity isolated). When the *whole* diagram is one thickened segment, $\mathrm{conv}\{(0,0),(n,0),(0,1)\}$, the geometry is $\mathbb{C}^2/\mathbb{Z}_n \times \mathbb{C}$, which is globally an $\mathcal{N}$=2 theory; the representative case is instead an $\mathcal{N}$=1 host with *one* long edge. The standard example (Franco–Hanany–Saad–Uranga, Section 4) is the suspended pinch point $xy = zw^2$,

$$\text{SPP}: \quad \mathrm{conv}\{(0,0),(2,0),(1,1),(0,1)\}, \qquad 2\,\text{Area} = 3 \text{ nodes}, \tag{VII.8.10}$$

a genuine quadrilateral, not a thickened segment, hence globally $\mathcal{N}$=1; its non-corner boundary point $(1,0)$, interior to the length-two bottom edge, is one line of $\mathbb{C}^2/\mathbb{Z}_2 = A_1$ singularities. Its three-node quiver, with the superpotential of Franco–Hanany–Saad–Uranga, has the adjacency

$$a = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}, \qquad a - a^{\mathsf{T}} = 0, \tag{VII.8.11}$$

the off-diagonal entries the three bifundamental pairs and the *adjoint* $a_{11} = 1$ at node 1. The quiver is non-chiral, $a - a^{\mathsf{T}} = 0$, yet the theory is $\mathcal{N}$=1. The $\mathcal{N}$=1 versus $\mathcal{N}$=2 distinction is the toric one above, not chirality. Since $a - a^{\mathsf{T}} = 0$, the whole rank lattice is anomaly-free, so $n_{\text{frac}} = \dim\ker -1 = 3 - 1 = 2$. There are two fractional branes, $(1,0,0)$ and $(0,1,0)$. The $\mathcal{N}$=2 one is $(1,0,0)$, the node carrying the adjoint X_{11}; that adjoint is precisely the modulus that slides the brane along the curve of A_1 singularities, a supersymmetric $\mathcal{N}$=2 Coulomb branch. The other, $(0,1,0)$, is a deformation brane, so the SPP realizes two of the three types at once. What the $\mathcal{N}$=2 brane sees transversally is the eight-supercharge local model $\mathbb{C}^2/\mathbb{Z}_2 \times \mathbb{C} = \mathbb{C}^3/\mathbb{Z}_2$, the A_1 sub-sector.

The adjoint interpretation can be tested directly in the superpotential. In the arrow convention of (VII.8.11), take

$$\begin{aligned} W_{\text{SPP}} = {} & \mathrm{Tr}(X_{21}X_{12}X_{23}X_{32} - X_{32}X_{23}X_{31}X_{13}) \\ & + \mathrm{Tr}(X_{13}X_{31}X_{11} - X_{12}X_{21}X_{11}). \end{aligned} \tag{VII.8.12}$$

Every monomial is a closed path. Differentiation with respect to the adjoint gives

$$\partial_{X_{11}} W = X_{13}X_{31} - X_{12}X_{21} = 0. \tag{VII.8.13}$$

On the isolated rank sector $(M, 0, 0)$ all bifundamentals are absent and W vanishes. The remaining vector and adjoint have the field content of pure $\mathcal{N} = 2$ gauge theory. Their identification with the moving fractional brane also uses its kinetic couplings and the local transverse geometry. Indeed, for $f = xy - zw^2$, the equations $f = \partial f = 0$ give $x = y = w = 0$ with z free. Near $z \neq 0$ a local rescaling of w puts the transverse surface in the form $xy = w'^2$, the A_1 surface. This distinguishes SPP from an arbitrary symmetric adjacency with an adjoint added to it.

The DSB type: the cone over dP_1. The first del Pezzo cone is isolated like the conifold (its toric diagram $\mathrm{conv}\{(1,0),(0,1),(-1,-1),(0,-1)\}$ is all-corner, $2\,\mathrm{Area} = 4$ nodes, and $GL(2,\mathbb{Z})$-distinct from F_0), but its complex-structure deformation is *obstructed* (Franco–Hanany–Saad–Uranga): it is *not* of deformation type. Its quiver (Figure 2 of Franco–Hanany–Saad–Uranga) has the adjacency

$$a = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 2 & 1 & 0 & 0 \end{pmatrix}, \qquad \ker(a - a^{\mathsf{T}}) = \big\langle (1,1,1,1),\, (0,3,1,2) \big\rangle, \tag{VII.8.14}$$

so beyond the regular brane there is exactly one fractional direction (the $d-3 = 4-3 = 1$ count of §VII.8.2), the rank vector $(0, 3, 1, 2)$. The anomaly-free cascade ranks are

$$(N,\, N+3M,\, N+M,\, N+2M) = N\,(1,1,1,1) + M\,(0,3,1,2), \tag{VII.8.15}$$

and as the cascade lowers N the smallest node (rank N) reaches zero first. To read the fractional sector, identify matrix rows and columns $0, 1, 2, 3$ in (VII.8.14) with source nodes $1, 2, 3, 4$, respectively. At $N = 0$ the source's cyclic potential (3.1) restricts to its (3.2),

$$W_{dP_1}^{\mathrm{frac}} = \mathrm{Tr}\big[X_{42}\left(X_{23}^1 X_{34}^2 - X_{23}^2 X_{34}^1\right)\big]. \tag{VII.8.16}$$

Both terms follow the closed path $4 \to 2 \to 3 \to 4$. Node 2 has $N_c = 3M$ and $N_f = 2M$, counted either from the rank-$2M$ node 4 through X_{42} or from the two rank-M arrows $X_{23}^{1,2}$. Thus $N_f < N_c$. The flavor inequality permits an ADS term; it does not prove supersymmetry breaking. The sourced analysis combines this tree potential with the nonperturbative term and finds no supersymmetric vacuum at finite fields in the stated model. With dynamical FI or baryonic directions the result is a runaway, rather than a stable nonsupersymmetric minimum. The detailed infrared analysis belongs to Part III.

Table VII.8.3 collects the three types as a diagnostic: how each is read off the geometry and off the quiver / rank data, and what infrared it drives.

In these examples the deformation type gives the confining branch, the pure $\mathcal{N}$=2 configuration gives a supersymmetric Coulomb branch, and the DSB configuration can run away when the stated moduli are dynamical. The slogan "a fractional brane triggers confinement" is the conifold special case, not the general rule, and the type must be read from the cycle.

type	geometric signal	dimer / quiver signal	infrared	here
deformation	a complex-structure deformation (finite S^3)	rank vector compatible with confinement	confines (S^3)	conifold
$\mathcal{N}$=2	a non-isolated $\mathbb{C}^2/\mathbb{Z}_n$ line	the specified $\mathcal{N}$=2 sub-sector	Coulomb branch	SPP
DSB	obstructed smoothing in the stated model	incompatible tree and dynamical potentials	runaway with dynamical moduli	dP_1

Table VII.8.3: The three types of anomaly-free fractional brane, read off the geometry of the wrapped cycle and off the quiver / rank data, with the infrared behavior of each specified configuration. Each is realized by a worked example whose physical label is supplied by the dynamics: the conifold (deformation), the suspended pinch point ($\mathcal{N}$=2), and the cone over dP_1 (DSB). The displayed deformation example confines; the $\mathcal{N}$=2 and DSB infrared dynamics route to Part III (§VII.8.7).

The conifold has only the deformation type. The conifold's single fractional direction (VII.8.6) is the deformation type, and one can see why from the geometry already in hand. The conifold is an *isolated* singularity: its only singular point is the tip, and there is no line of $\mathbb{C}^2/\mathbb{Z}_n$ singularities for an $\mathcal{N}$=2 brane to live on. So the $\mathcal{N}$=2 type is absent. The collapsed cycle at the tip is the two-cycle of the small resolution, and the conifold admits a complex-structure deformation that replaces the tip by a three-sphere (the next section); the fractional brane wrapping the vanishing cycle is therefore of deformation type. And the resulting gauge theory has a supersymmetric, confining vacuum, so the DSB type is absent. The conifold realizes one type, the deformation type, and its infrared is confinement. Other toric cones can have several types of fractional brane, including mixtures whose dynamics differs from that of either constituent. The conifold is the simplest deformation example.

Three types, read from the cycle

Deformation: the wrapped cycle smooths by a complex-structure deformation (a finite S^3); the theory **confines**. The conifold's brane. $\mathcal{N}$=2: the brane lives on a non-isolated $\mathbb{C}^2/\mathbb{Z}_n$ line (eight supercharges locally); an $\mathcal{N}$=2 **Coulomb branch**. *DSB*: a configuration with no finite supersymmetric vacuum in the stated dynamics; an Affleck–Dine–Seiberg **runaway** when the relevant moduli vary. Only the deformation type confines; the type uses both the wrapped geometry and the dynamics of the specified configuration.

VII.8.4 The duality cascade

Now set the fractional brane in motion. Take the conifold with N regular and M fractional branes, the rank vector (VII.8.6), so the gauge group is

$$U(N+M)_1 \times U(N)_2, \tag{VII.8.17}$$

with the four bifundamentals $A_i : 1 \to 2$, $B_j : 2 \to 1$ and the conifold quartic

$$W = \lambda\, \varepsilon^{ij}\varepsilon^{kl}\, \mathrm{Tr}\big(A_i B_k A_j B_l\big) \tag{VII.8.18}$$

of Chapter 1. Its physical presence is an imported input. Given that quartic and $\varepsilon^{12} = 1$, expanding the epsilon contractions yields twice the difference of the two cyclic quartics. The two ranks differ, and that difference drives the dynamics.

Why the couplings run apart. The two gauge couplings are no longer tied. For an $\mathcal{N}$=1 gauge factor $SU(N_c)$ with N_f flavours of anomalous dimension γ, the NSVZ numerator is

$$\beta\big(SU(N_c)\big) \;\propto\; 3N_c - N_f(1-\gamma). \tag{VII.8.19}$$

At the symmetric conifold point $\gamma = -\frac{1}{2}$ and $N_f = 2N_c$, and the numerator vanishes: both couplings are marginal, the Klebanov–Witten fixed point. Shift the ranks to (VII.8.17). Node 1, now $SU(N+M)$, still sees $N_f = 2N$ flavours from node 2, while node 2, $SU(N)$, sees $N_f = 2(N+M)$ from node 1. With $\gamma = -\frac{1}{2}$ held to leading order for $M/N \ll 1$, the two numerators become

$$\beta_1 \;\propto\; 3(N+M) - \tfrac{3}{2}(2N) = 3M, \qquad \beta_2 \;\propto\; 3N - \tfrac{3}{2}\,2(N+M) = -3M, \tag{VII.8.20}$$

equal and opposite, of order M. The inverse squared couplings run linearly in $\ln\mu$ with opposite slopes: one node strengthens as the other weakens. Run the theory down in energy and the larger-rank node, $SU(N+M)$, becomes strongly coupled first for the usual cascade trajectory. The description in terms of (VII.8.17) breaks down there, and the theory must be continued.

The Seiberg step. The continuation is the Chapter 7 move: Seiberg-dualize the strongly-coupled node. An $SU(N_c)$ gauge theory with N_f fundamental flavours is dual to an $SU(N_f - N_c)$ theory with the same flavours plus mesons. Apply it to node 1, $SU(N+M)$ with $N_f = 2N$:

$$N_c = N+M, \quad N_f = 2N \;\;\Longrightarrow\;\; N_f - N_c = 2N - (N+M) = N - M. \tag{VII.8.21}$$

Node 1 becomes $SU(N-M)$, and the gauge group is now

$$U(N+M)_1 \times U(N)_2 \;\longrightarrow\; U(N)_2 \times U(N-M)_1, \tag{VII.8.22}$$

the same conifold quiver, with ranks $(N, N-M)$ in place of $(N+M, N)$: the higher rank has dropped from $N+M$ to N, and the lower from N to $N-M$. On the dimer this is a single square move, exactly the operation of Chapter 7; the electric quartic becomes a nondegenerate quadratic term for the four magnetic mesons, while the magnetic superpotential couples those mesons cubically to dual quarks. For nonzero quartic

coupling the mesons integrate out, returning the two-node conifold quiver with the same quartic. The net effect on the construction data is the rank step

$$(N+M,\,N)\;\longrightarrow\;(N,\,N-M), \tag{VII.8.23}$$

the lower rank decreased by M. This ordinary step requires $N > M$; the positive magnetic rank condition alone does not determine the coupling trajectory. We use interacting SU factors for the strong dynamics. The displayed U groups record the brane ranks; abelian factors do not obey this SQCD duality rule unchanged.

The cascade. After the step the roles swap: node 2 is now the larger and runs to strong coupling, and the same move applied to it drops the ranks again by M. Iterating, the ranks descend by M at each duality step,

$$(N+M,\,N)\;\to\;(N,\,N-M)\;\to\;(N-M,\,N-2M)\;\to\;\cdots, \tag{VII.8.24}$$

a *duality cascade*: the same conifold theory at every step, with the ranks marching down the staircase by M per landing. Figure VII.8.1 draws the descent. Equivalently, the *effective number of regular branes* decreases as the energy scale drops. The five-form flux that counts the regular D3-charge is not constant in the cascade background; it runs logarithmically, and the field-theory interpretation is the running of the effective rank,

$$N_{\rm eff}(\mu)\;\sim\;\frac{3}{2\pi}\,g_sM^2\,\ln\frac{\mu}{\Lambda}, \tag{VII.8.25}$$

with g_s the string coupling and an additive constant suppressed. This is the Klebanov–Tseytlin asymptotic Maxwell D3 charge, with radial scale identified with μ to leading logarithmic accuracy. It is not the integer Page charge, whose changes by M under large B_2 gauge transformations describe successive rank frames. The coefficient comes from the normalized flux solution, not from $N_f - N_c$. Its supergravity use requires small string coupling and small curvature (in particular $g_sM \gg 1$ for the smooth throat), and the asymptotic logarithm is not an infrared formula at the tip. In this regime, each factor of e in the scale changes the continuous charge by $\mathcal{O}(g_sM^2)$ in the effective rank, which is the cascade descent (VII.8.24) read continuously. The number of steps from the ultraviolet down to the confinement scale is of order N/M, and the cascade is therefore *finite*: it has a bottom, reached when the effective rank runs out.

The quantum bottom is explicit for $N = kM$ with $k \geq 1$, $M \geq 2$, on the specified confining branch. The descent (VII.8.24) then reaches the last nontrivial stage

$$SU(2M)\times SU(M), \tag{VII.8.26}$$

where the strongly-coupled $SU(2M)$ sees $N_f = 2M$ flavours from the $SU(M)$ node, exactly $N_f = N_c$. At $N_f = N_c$ the ordinary rank recurrence stops. The SQCD quantum constraint is $\det\mathcal{M} - B\widetilde{B} = \Lambda_{2M}^{4M}$, not a new rank-zero magnetic gauge factor. On the baryonic branch, $\mathcal{M} = 0$ and $B\widetilde{B} = -\Lambda_{2M}^{4M}$; the remaining $SU(M)$ sector confines. The complete branch includes baryonic moduli; the M condensate phases below count the confining sector, not all its continuous vacua. This quantum constraint and the identification with the deformed geometry are physical inputs from

the Klebanov–Strassler analysis. This is the gauge-theory origin of the finite S^3 of §VII.8.5. So "the rank runs out" means the regular D3-charge has cascaded away and the remaining fractional-brane dynamics confines. It does not mean that all gauge dynamics disappears. The bottom is the subject of the next section.

For $N = kM$ with $k \geq 1$ and $M \geq 2$, precisely $k - 1$ ordinary steps reach $(2M, M)$. For $N = kM + p$ with $0 < p < M$, the same subtraction instead reaches $(M + p, p)$; its larger node has $N_f = 2p < N_c = M + p$, so the preceding ordinary step cannot be repeated. Residual probe sectors and the nonperturbative infrared must then be specified. At $M = 0$ there is no rank descent. At $k = 0$ the starting $(M, 0)$ is a pure fractional sector, not a passage through $(2M, M)$; at $M = 1$ the residual $SU(1)$ is trivial. The ordinary subtraction applies to positive ranks where its flavor and magnetic-rank conditions hold, but this physical bottom requires the stated interacting factors. These stopping rules distinguish the exact integer recurrence from the large-rank approximation used to estimate the running couplings.

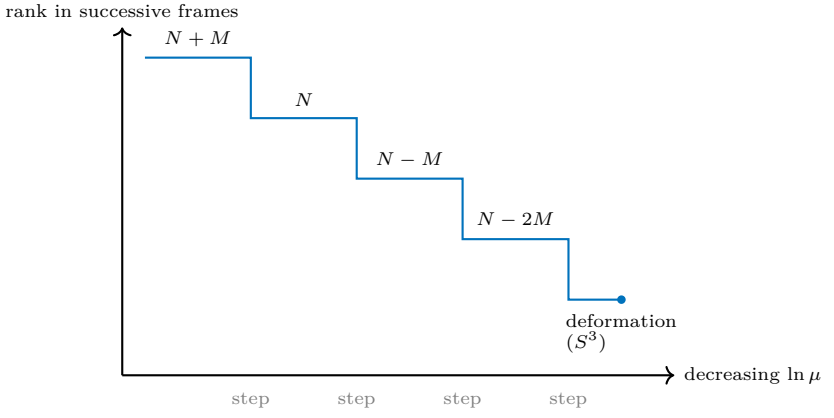


Figure VII.8.1: The duality cascade as a schematic staircase of integer rank frames. Reading right (down in energy), each landing represents one permitted Seiberg duality and the relabeled larger rank drops by M: $N + M \to N \to N - M \to \cdots$ (VII.8.24). Equal drawn widths do not specify the full RG intervals. The continuous logarithm (VII.8.25) is the asymptotic Maxwell charge, a distinct quantity from the Page charge associated with these rank frames. For $N = kM$, $k \geq 1$, $M \geq 2$, the specified confining branch ends at a finite S^3 (§VII.8.5).

The flavor count and the assumed Seiberg prescription give the rank step (VII.8.23). The leading-log charge (VII.8.25) additionally uses the supergravity flux normalization. These are separate calculations with different regimes. The precise matching of couplings across each Seiberg duality beyond the leading order, the scheme-dependence of the running, and the holographic c-function profile are the precision renormalization-group questions; this chapter carries the leading cascade and routes the precision matching to Part III.

VII.8.5 The deformed-conifold infrared

The cascade ends, and what it ends in is read off the geometry. As the effective rank runs to the bottom of (VII.8.24), the gauge theory reaches the scale where the strongly-coupled dynamics confines. The geometry's response is a *deformation* of the conifold, and this section identifies the confining infrared by that deformation, sets it apart from the resolution it must not be confused with, and fixes the deformation parameter in scale.

Two smoothings of one cone. The singular conifold $\mathcal{C} = \{xy - zw = 0\}$ has two distinct smooth neighbours, and the distinction is the heart of this section. Chapter 1 named them; here they are kept rigidly apart, because the cascade ends in one and not the other.

The *small resolution* blows the singular point up to a two-sphere. In the symplectic quotient of Chapter 1, turning on the Fayet–Iliopoulos parameter $\zeta > 0$ of the relative gauge $U(1)$ forces $(A_1, A_2) \neq 0$, so the A's parametrize a $\mathbb{P}^1$ and the tip is replaced by that $\mathbb{P}^1 = S^2$,

$$\text{small resolution:} \quad \text{tip} \;\to\; \mathbb{P}^1 = S^2 \text{ (real dimension 2)}, \qquad \mathrm{vol}(S^2) \sim \zeta, \quad \text{(VII.8.27)}$$

a *Kähler* deformation, controlled by the Fayet–Iliopoulos modulus ζ, and toric (it is a chamber of the toric data). This is the branch the wrapped fractional branes of §VII.8.2 live on.

The *deformation* is a different operation, on the complex structure rather than the Kähler class. Deform the defining equation,

$$xy - zw = \epsilon^2, \qquad \text{equivalently} \qquad \sum_{a=1}^{4} z_a^2 = \epsilon^2, \qquad \text{(VII.8.28)}$$

after the linear change of coordinates $x = z_1 + iz_2$, $y = z_1 - iz_2$, $z = z_3 + iz_4$, $w = -(z_3 - iz_4)$ that diagonalizes the quadric. Count the dimension first. The deformed variety is one polynomial relation in $\mathbb{C}^4$, so as a complex hypersurface it has

$$\dim_{\mathbb{C}}\{\textstyle\sum_a z_a^2 = \epsilon^2\} = (\text{ambient}) - (\text{relations}) = 4 - 1 = 3, \qquad \text{(VII.8.29)}$$

a Calabi–Yau threefold, the same complex dimension as the singular cone. It is now *smooth*: for $\epsilon \neq 0$ the gradient $\nabla(\sum z_a^2 - \epsilon^2) = 2(z_1, z_2, z_3, z_4)$ vanishes only at the origin, which is no longer on the variety ($\sum 0 = 0 \neq \epsilon^2$), so there is no singular point. The minimal compact cycle is now a *three-sphere*. Intersect the variety with the real slice $\mathrm{Im}\, z_a = 0$: the real parts satisfy

$$\sum_{a=1}^{4} (\mathrm{Re}\, z_a)^2 = \epsilon^2 \quad (\epsilon \text{ real}), \qquad \text{(VII.8.30)}$$

a round S^3 of coordinate radius $|\epsilon|$ in $\mathbb{R}^4$, the vanishing cycle that shrinks to the singular tip as $\epsilon \to 0$. So the deformation replaces the tip by a finite three-sphere,

$$\text{deformation:} \quad \text{tip} \;\to\; S^3 \text{ (real dimension 3)}, \qquad \text{coordinate radius} = |\epsilon|, \quad \text{(VII.8.31)}$$

This is a coordinate statement, not a proper radius in the Ricci-flat or warped metric. For a complex nonzero ϵ, rotate the coordinates by its phase before taking the real slice. The small resolution (VII.8.27) instead put an S^2 of one real dimension lower: the two smoothings differ in the very dimension of the cycle they open, 3 against 2. The deformation is a *complex-structure* deformation, controlled by the modulus ϵ, and *not* toric (a generic complex-structure deformation breaks the torus action that fixes the S^2 resolution chamber). Table VII.8.4 sets the two apart.

	small resolution	deformation
operation	Kähler (blow-up)	complex structure
local model	$xy = zw$, $\zeta > 0$	$xy - zw = \epsilon^2$
finite cycle	$\mathbb{P}^1 = S^2$	S^3
real dimension	2	3
modulus	Fayet–Iliopoulos ζ	complex structure ϵ
toric?	yes (a toric chamber)	no
the cascade IR?	no	**yes** (confinement)

Table VII.8.4: The two smoothings of the conifold are distinct geometries with distinct infrareds. The cascade ends in the *deformation* (the S^3, the confining / chiral-symmetry-breaking branch), not the resolution. The deformation parameter ϵ is a complex-structure modulus, not a Fayet–Iliopoulos parameter; the two have different real dimensions of the finite cycle (3 vs 2) and different toric status.

Figure VII.8.2 draws the contrast. The two smoothings are exchanged by a geometric transition: the cascade does not resolve the singularity, it deforms it, and the chapter is careful with the word. The resolution is the FI/Kähler branch of the probe quotient; the *deformation* is the confining infrared.

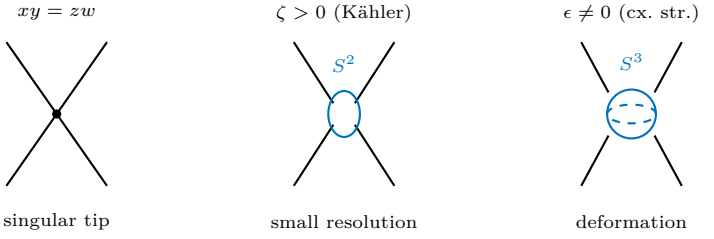


Figure VII.8.2: The singular conifold (left) and its two smoothings: the small resolution (middle), a Kähler blow-up replacing the tip with a two-sphere $S^2 = \mathbb{P}^1$ of area proportional to ζ; and the deformation (right), a complex-structure smoothing replacing the tip with a three-sphere S^3 of coordinate radius $|\epsilon|$. The spheres are drawn schematically (an S^3 is suggested by a sphere with an equatorial circle); only the dimensions and moduli are exact. The cascade ends in the deformation, the confining branch.

The infrared, identified. The deformation is the geometric input from which the chapter reads the infrared physics, and it reads three things. First, *confinement*: the cascade has run the effective rank down to the deformation scale, the gauge theory has no further Seiberg step to take, and the strongly-coupled dynamics confines. Second, *chiral-symmetry breaking*: the deformed conifold has a finite S^3 where the singular cone had none, and the deformation parameter ϵ is an order parameter; the geometric breaking of the would-be symmetry of the tip is the chiral-symmetry breaking of the confining gauge theory. Third, the *scale*: the deformation parameter is set by the strong-coupling scale through the gaugino condensate,

$$\epsilon^2 \;\propto\; \langle\lambda\lambda\rangle \;\sim\; \Lambda_{\rm IR}^3, \tag{VII.8.32}$$

the dimension-three condensate of the residual confining sector, with proportionality factors depending on coordinate and holomorphic scale conventions. At the conformal Klebanov–Witten point the mesons have scaling dimension 3/2, so the determinant has dimension three. Along the cascade that conformal assignment is not an exact running engineering dimension; it motivates the matching and does not derive it. The deformed conifold's S^3 is finite precisely because the gaugino has condensed; $\epsilon \to 0$ returns the singular quadric as a geometric limit. At fixed nonzero $\Lambda_{\rm IR}$ this is not a vacuum on the confining branch.

What chiral symmetry breaks can be named precisely. The classical conifold theory has a $U(1)_R$; with M fractional branes it is anomalous, and only a discrete $\mathbb{Z}_{2M}$ survives the anomaly. The gaugino condensate $\langle\lambda\lambda\rangle$, of R-charge two, breaks this $\mathbb{Z}_{2M}$ further to $\mathbb{Z}_2$, leaving

$$\mathbb{Z}_{2M} \xrightarrow{\langle\lambda\lambda\rangle\neq 0} \mathbb{Z}_2, \qquad \#\{\text{condensate phases}\} = M, \tag{VII.8.33}$$

M inequivalent condensate phases in the confining sector, which geometrically are the phases of the deformation parameter $\epsilon^2 \sim \Lambda^3$. This is the precise content of chiral-symmetry breaking in the deformed-conifold infrared: a discrete chiral $\mathbb{Z}_{2M}$ spontaneously broken to $\mathbb{Z}_2$ by the condensate. The discrete-anomaly analysis itself is the Klebanov–Strassler field theory, stated here and routed to Part III.

What the geometry says about the IR

On the confining branch described above, the cascade ends at the **deformation**, $xy - zw = \epsilon^2$, a finite S^3. Reading this geometry: the gauge theory **confines**; chiral symmetry is **broken**, with ϵ the order parameter; and the deformation scale is fixed by the gaugino condensate, $\epsilon^2 \sim \Lambda^3 \sim \langle\lambda\lambda\rangle$. The chapter *names* these from the deformed geometry. It does *not* construct the Klebanov–Strassler supergravity solution (the warped metric, the three- and five-form fluxes) or compute any glueball, baryon, or Kaluza–Klein spectrum; those are routed to Part III (§VII.8.7).

This is the dynamics boundary at its sharpest. The construction *identifies* the infrared: it names confinement and chiral-symmetry breaking by reading the deformed geometry, and it fixes the deformation parameter's scale by the condensate (VII.8.32). The Klebanov–Strassler *supergravity solution*, the warped deformed-conifold metric that realizes this infrared as a smooth gravity dual, its three-form and five-form fluxes, and

the dual glueball, baryon, and Kaluza–Klein spectrum, are not constructed here; they are the dual-gravity foundations, routed to Part III. The chapter names the infrared and the deformation that produces it, and stops at the boundary of the construction data.

A last point of language, to keep the gravity side honestly routed. In the singular cascade background, before the deformation, the running flux produces a naked singularity (the Klebanov–Tseytlin singularity). It is the deformation, the smoothing to the finite S^3, that *removes* that naked singularity; the chapter says "removed by the deformation" and not "resolved", reserving "resolution" for the Kähler S^2 branch that is a different geometry entirely. The supergravity statement that the warped deformed throat is smooth is the Part III deliverable; here the deformation is named as what cures the singularity, and the gravity-side construction is routed.

VII.8.6 The cascade for a general toric singularity

The cascade is not special to the conifold. Other toric quivers admit sequences of Seiberg descriptions with changing ranks. Existence of an anomaly-free fractional class does not by itself choose such a trajectory or its endpoint. Deformation-type branes can lead to confining endings; DSB configurations such as dP_1 can also have a cascade in the ultraviolet, with a different infrared. This section names the generalization, gives its dimer-move picture, and defers its richer dynamics.

The structural statement is that a cascade uses successive toric dual descriptions. A permitted square move connects neighboring tilings; a sequence of such moves connects phases within a demonstrated reachable component of the phase graph. The F_0 pair below is one actual move. With unequal ranks, a cascade additionally requires a valid Seiberg-duality frame at each chosen node and a running-coupling trajectory that reaches it. Only then is the cascade a *walk* on the phase graph, with a rank change set by the node's flavor count. The conifold is the simplest periodic example: its two-node quiver returns at each step while the ranks descend uniformly.

A non-conifold example makes the walk concrete, and it is a worked kernel computation, the same one of §VII.8.2 on a four-node quiver. Take $F_0 = \mathbb{P}^1 \times \mathbb{P}^1$ in its phase I dimer: four nodes on a directed necklace, two arrows between each consecutive pair, $a_{01} = a_{12} = a_{23} = a_{30} = 2$ (eight bifundamentals in all). Its adjacency and antisymmetric part are

$$a = \begin{pmatrix} 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \\ 2 & 0 & 0 & 0 \end{pmatrix}, \qquad a - a^{\mathsf{T}} = \begin{pmatrix} 0 & 2 & 0 & -2 \\ -2 & 0 & 2 & 0 \\ 0 & -2 & 0 & 2 \\ 2 & 0 & -2 & 0 \end{pmatrix}. \tag{VII.8.34}$$

Unlike the maximally chiral dP_0, this antisymmetric matrix has rank two, so its kernel is *two*-dimensional,

$$\ker(a - a^{\mathsf{T}}) = \big\langle\, (1,0,1,0),\ (0,1,0,1) \,\big\rangle, \qquad n_{\text{frac}} = \dim\ker -1 = 2 - 1 = 1, \tag{VII.8.35}$$

the regular direction being the sum $(1,1,1,1) = (1,0,1,0) + (0,1,0,1)$ and a primitive fractional class represented by $(1,0,1,0)$. The alternating vector $(1,-1,1,-1)$ represents

twice that class; together with the regular vector it spans an index-two sublattice. So F_0 *does* admit one anomaly-free fractional brane where dP_0 admitted none. Here F_0 is used only to exhibit the anomaly-free rank direction and the dimer-walk mechanism; its detailed infrared type and cascade endpoint are not analyzed in this chapter (the dynamics route to Part III).

For a symmetric parametrization of an even relative-rank difference, take: the rank vector is

$$N_{\text{rank}} = N\,(1,1,1,1) + M\,(1,-1,1,-1) = (N+M,\ N-M,\ N+M,\ N-M), \quad \text{(VII.8.36)}$$

anomaly-free by construction. One Seiberg step makes the descent concrete. Choose couplings for which a higher-rank node becomes strongly coupled. Dualizing node 0, which sees $N_f = 2(N-M)$ flavours (two outgoing arrows to node 1 and two incoming arrows from node 3, both at rank $N-M$) against $N_c = N+M$ colours, sends it to

$$N_f - N_c = 2(N-M) - (N+M) = N - 3M, \quad \text{(VII.8.37)}$$

for $N > 3M$, the formal rank step $SU(N+M) \to SU(N-3M)$ at that node, with the square move carrying the quiver to phase II when this duality frame applies. At $N = 3M+1$ the magnetic $SU(1)$ is trivial, so the subtraction is rank arithmetic rather than a two-factor interacting cascade. The chosen F_0 node loses $4M$ colors. For comparison, the dualized conifold node loses $2M$ colors, although the relabeled lower rank decreases by M. Thus one must specify both the rank coordinate and the chosen node when comparing steps. A single subtraction establishes neither periodicity nor its absence. F_0 also has two distinct toric phases (the two dimers of Chapter 7), and the square move on the strongly-coupled node carries the eight-field phase I to a twelve-field phase II,

$$\text{phase I: 8 bifundamentals} \xrightarrow{\text{dimer move}} \text{phase II: 12 bifundamentals}, \quad \text{(VII.8.38)}$$

both toric phases of the same F_0. This single transition does not establish an alternating two-phase cascade for arbitrary couplings and ranks. Further steps require the new flavor counts, superpotential, and RG trajectory. The eight-to-twelve field transition is imported from the regular-brane toric phase construction; the unequal-rank theory uses its quiver and superpotential, not the regular-brane conformal-window equality $N_f = 2N_c$. Here we compute only the kernel and one possible rank step. The detailed F_0 infrared remains outside the chapter.

The general cascade can be richer than the conifold's clean periodic descent, and here the chapter names two phenomena and defers their dynamics. The rank steps can *accumulate*: the energy intervals between successive Seiberg dualities can shrink geometrically, so that infinitely many steps pile up at a finite energy scale, a *duality wall*. The $Y^{p,q}$ Sasaki–Einstein cones supply further cascading examples. In particular, the $Y^{p,p-1}$ and $Y^{p,1}$ families have self-similar rank sequences after node permutations in the trajectories studied by Herzog–Ejaz–Klebanov. This does not establish a universal slowly drifting or quasi-periodic sequence for every $Y^{p,q}$. The detailed beta functions, choice of trajectory and possible wall accumulation require additional dynamics and are routed to Part III; the examples also appear in the Chapter 11 atlas. The point this

section establishes is structural and not dynamical: when a suitable sequence of Seiberg descriptions exists, its rank changes can be computed node by node. The conifold supplies the explicit periodic example; a general anomaly-free lattice does not guarantee one.

VII.8.7 What has been computed, and what is deferred

The chapter has turned the static toric construction into dynamics by relaxing the equal-rank assumption, and has read the endpoint of that dynamics off the geometry. The finite calculations and their physical inputs fit together as follows. A fractional brane is an anomaly-free deviation from the regular rank vector, an integer solution of $(a - a^{\mathsf{T}})N = 0$ (VII.8.3), and the number of independent fractional directions is $\dim\ker(a - a^{\mathsf{T}}) - 1$ (VII.8.4): one for the conifold, zero for $\mathbb{C}^3$, and, the honest result, zero for dP_0, whose maximal chirality leaves no room for an anomaly-free rank shift. The anomaly-free fractional branes fall into three geometric types, deformation, $\mathcal{N}$=2, and dynamical-supersymmetry-breaking, classified by the wrapped cycle, with only the deformation type confining and the conifold carrying that one type. The conifold cascade is the iterated Seiberg-duality move with $N_f = 2N$ at the dualized node, the rank step $(N+M, N) \to (N, N-M)$ (VII.8.23), the descent by M per landing (VII.8.24), and the sourced asymptotic Maxwell charge $N_{\text{eff}} \sim (3/2\pi) g_s M^2 \ln \mu$ (VII.8.25). The cascade ends in the deformation, $xy - zw = \epsilon^2$, a finite S^3 (VII.8.28)–(VII.8.31), distinct from the resolution's S^2, and from that deformed geometry the chapter identifies confinement and chiral-symmetry breaking and imports the holomorphic scale matching $\epsilon^2 \propto \langle\lambda\lambda\rangle \sim \Lambda_{\text{IR}}^3$ (VII.8.32).

A reader leaving the chapter should be able to do the following without looking ahead: write the cubic-anomaly condition $(a - a^{\mathsf{T}})N = 0$ and compute the fractional-brane count for a given quiver; see that the conifold admits one fractional direction and dP_0 none; name the three fractional-brane types and say which confines; set up the conifold cascade $U(N+M) \times U(N)$, run the Seiberg step to $U(N) \times U(N-M)$ via $N_f = 2N$, and iterate the descent; and distinguish the deformation (S^3, complex structure, ϵ, the confining infrared) from the resolution (S^2, Kähler, ζ).

What has been deferred, and to where, completes the volume's deferral map.

- The Klebanov–Strassler *supergravity solution*, the warped deformed-conifold metric, its three-form and five-form fluxes and the logarithmically running five-form flux, and the dual glueball, baryon, and Kaluza–Klein spectrum, are Part III (Klebanov–Strassler dynamics and gravity). The chapter names the infrared and the deformation; the dual gravity solution and its spectrum are foundations. This is the chapter's primary deferral, the dynamics boundary at its sharpest.

- The gaugino-condensation *dynamics* that fixes $\epsilon^2 \sim \Lambda^3$, the nonperturbative gauge dynamics behind the deformation parameter, is Part III (Klebanov–Strassler dynamics and gravity). The chapter states the scale and the order parameter; the dynamics that produces it is routed.

- The precise step-matching of the cascade beyond the leading log, the exact matching of couplings across each Seiberg duality, the scheme-dependence, and the holographic c-function profile, are Part III. The chapter carries the leading-order cascade.
- The Seiberg-duality *move* is Chapter 7's, used here and iterated; the field-theory proof of Seiberg duality is Chapter 7's deferral (Part III), cited.
- The $\mathcal{N}$=2 and DSB fractional-brane types are named and read off the geometry; their detailed dynamics, the $\mathcal{N}$=2 Coulomb branch and the DSB Affleck–Dine–Seiberg runaway with its supersymmetry-breaking model-building, are routed (Part III and a phenomenology pointer).
- The general toric cascade is named with its dimer-move picture; the duality-wall accumulation and the specific $Y^{p,q}$ trajectories are routed to Part III (and revisited in the Chapter 11 atlas).
- The deformed-conifold and Sasaki–Einstein *geometry* as differential and algebraic geometry, the deformation as a complex-structure smoothing and the metric on $T^{1,1}$, is conifold geometry (Part I), cited not constructed.
- The leading beta-function estimate requires $M \ll N$; the exact rank subtraction and its stopping rule do not require that approximation, and reading its moduli as literal transverse geometry carries the α' and g_s probe caveats inherited from Volume VI ; the finite-N and sub-leading questions are routed.

The lesson to carry out of the chapter is the one in the opening keybox: a sequence of permitted dimer moves can supply successive infrared descriptions when a valid duality frame and RG trajectory exist. On the specified conifold branch the deformation at the cascade bottom identifies confinement. The quiver computes allowed ranks and candidate steps. The physical duality and nonperturbative dynamics identify the trajectory and its confining branch. Their geometric description connects this construction to the gravity analysis in Part III.

Status and sources

The evidence-code tags ([WS] physical-worldsheet derivation in the untwisted string, [MATH] mathematical result with stated hypotheses and proof basis), the convention pointers, and the source anchors are collected here, keyed by section. The chapter closes Chapters 6–8 of Volume VII by turning the static toric construction into dynamics: the fractional brane as an anomaly-free rank deformation (Franco, Hanany, Saad, et al. (2006)), the duality cascade as iterated Seiberg duality (Klebanov and Strassler (2000), Klebanov and Nekrasov (2000), Klebanov and Tseytlin (2000)), and the deformed-conifold infrared (Klebanov and Strassler (2000), Candelas and Ossa (1990)). It reuses the gauge/rank data (quiver gauge groups and ranks, $\prod_i U(N_i)$, the fractional brane as a deviation from the regular rank vector), the resolution/deformation chambers (resolution and discrete data, the Fayet–Iliopoulos ζ versus the complex-structure ϵ), the adjacency and anomaly kernel (McKay multiplicities, fundamental-at-tail arrows, the Chapter 7 result $(a - a^{\mathsf{T}})N = 0$), the conifold quartic (superpotential normalization), the dimer move (bipartite tiling conventions, Chapter 7's square move), the toric diagram, and the supercharge count (spinor and supersymmetry counting, the $\mathcal{N}$=2 type). No new convention is introduced. The finite calculations use the stated quiver and cascade data.

§VII.8.1 From the static brane to its dynamics. Framing: the regular brane relaxed to unequal ranks, the cascade as the iterated Chapter 7 move, the deformed-conifold endpoint, and the dynamics boundary drawn at its sharpest. No new claim-atom.

§VII.8.2 The fractional brane as an anomaly-free rank deformation. The cubic gauge anomaly as $[(a-a^{\mathsf{T}})N]_i$ (VII.8.2), the anomaly-free condition $(a-a^{\mathsf{T}})N=0$ (VII.8.3), the fractional-direction count $\dim\ker(a-a^{\mathsf{T}})-1$ (VII.8.4), the wrapped-cycle reading, and the three fixtures (Table VII.8.1): $\mathbb{C}^3$ zero, the conifold one (the primitive class $[(1,0)]$, ranks $(N+M,N)$ (VII.8.6)), and dP_0 *zero* (the rank-two antisymmetric adjacency (VII.8.7); its maximal chirality leaves only the regular direction). (physical primary status unassigned here) (quiver gauge groups and ranks, McKay multiplicities, fundamental-at-tail arrows; Franco, Hanany, Saad, et al. (2006), the Volume VI anomaly result, the Volume VI charge lattice). Records (a fractional brane is a *constrained* anomaly-free rank deformation, not an arbitrary rank shift; the dP_0 zero-count sharpens it, the stated distinction). the fractional-direction counts $\mathbb{C}^3\to 0$, conifold $\to 1$, $dP_0\to 0$; the conifold admits $(1,-1)$ as a nonprimitive quotient class and the regular brane is anomaly-free; the dP_0 kernel is the regular $(1,1,1)$ alone.

§VII.8.3 The three types of fractional brane. The classification deformation / $\mathcal{N}$=2 / DSB by the wrapped cycle, with a worked example of each: the conifold (deformation, isolated and supersymmetric, the chapter spine), the suspended pinch point ($\mathcal{N}$=2, the Franco–Hanany–Saad–Uranga SPP quiver with the fractional brane $(1,0,0)$ on the A_1 line and its adjoint Coulomb modulus), and the cone over dP_1 (DSB, the Franco–Hanany–Saad–Uranga fractional brane $(0,3,1,2)$ and its runaway). Only the deformation type confines. (physical primary status unassigned here) (spinor and supersymmetry counting; Franco, Hanany, Saad, et al. (2006), Cachazo et al. (2001), Bertolini et al. (2001)). Records (not all fractional branes confine; the IR depends on the type). The $\mathcal{N}$=2 and DSB dynamics are routed (Part III / a phenomenology pointer). The quiver gives the multiplicities and ranks. The physical types require the stated dynamics.

§VII.8.4 The duality cascade. The cascade theory $U(N+M)\times U(N)$ with the conifold quartic (VII.8.17)–(VII.8.18); the NSVZ split $\beta_1\propto 3M$, $\beta_2\propto -3M$ (VII.8.19)–(VII.8.20); the Seiberg dual rank $N_f-N_c=2N-(N+M)=N-M$ (VII.8.21), the rank step $(N+M,N)\to(N,N-M)$ (VII.8.23), the descent by M (VII.8.24) (Figure VII.8.1), and the leading-log asymptotic Maxwell charge $N_{\text{eff}}\sim(3/2\pi)g_sM^2\ln\mu$ (VII.8.25). (physical primary status unassigned here) (quiver gauge groups and ranks, superpotential normalization, bipartite tiling conventions; Klebanov and Strassler (2000), Klebanov and Nekrasov (2000), Klebanov and Tseytlin (2000), the Chapter 7 move). Records (the specified $N=kM$, $k\geq 1$, $M\geq 2$ conifold branch has a finite cascade ending in confinement, not an infinite IR tower). The precise step-matching beyond the leading log is routed (Part III). $N_f=2N$ at the dualized node, the dual rank $N-M$, the descent $(N+M,N)\to(N,N-M)\to(N-M,N-2M)$, and the leading-log coefficient $(3/2\pi)g_sM^2$.

§VII.8.5 The deformed-conifold infrared. The two smoothings: the small resolution $\mathbb{P}^1=S^2$, Kähler, FI ζ, toric (VII.8.27); and the deformation $xy-zw=\epsilon^2$, S^3, complex structure, non-toric (VII.8.28)–(VII.8.31) (Table VII.8.4, Figure VII.8.2); the infrared identified, confinement and chiral-symmetry breaking named via the geometry, the scale $\epsilon^2\sim\Lambda^3\sim\langle\lambda\lambda\rangle$ (VII.8.32); the deformation (not "resolution") removing the Klebanov–Tseytlin naked singularity. (physical primary status unassigned here) (resolution and discrete data; Klebanov and Strassler (2000), Candelas and Ossa (1990)). Records (the deformed and resolved conifold are different IRs; ϵ is not an FI parameter; refines the stated distinction). The KS supergravity solution, its spectrum, and the gaugino-condensation derivation are routed to Part III (Klebanov–Strassler dynamics and gravity); the deformed-conifold geometry is conifold geometry (Part I). Legacy supplied geometry records compare the deformed conifold's S^3 (real dim 3) versus the resolution's S^2 (real dim 2), ϵ a complex-structure modulus not FI, the deformation non-toric, the deformed threefold complex dimension 3, $\epsilon^2\sim\Lambda^3$ a (mass)^3. The supergravity spectrum is *not* computed.

§VII.8.6 The cascade for a general toric singularity. The cascade as a walk through permitted neighboring toric phases within a demonstrated reachable component, conditional on duality frames and RG flow; the F_0 example (one fractional direction, the two phases with eight and twelve fields exchanged by the dimer move, the leading rank step); the duality wall and the self-similar $Y^{p,p-1}$ and $Y^{p,1}$ examples named, dynamics deferred. (physical primary status unassigned here) (bipartite tiling conventions; Franco, Hanany, He, et al. (2003), Herzog, Ejaz, et al. (2005), the Chapter 7 move). The duality-wall

accumulation and the detailed $Y^{p,q}$ trajectories are routed (Part III; the $Y^{p,q}$ atlas is Chapter 11). F_0 admits a fractional brane (one direction) and the local rank subtraction. The imported regular-brane phase construction has one move between the eight- and twelve-field phases. The check does not establish that an unequal-rank RG trajectory traverses that edge.

§VII.8.7 What is and is not given. Scope and limitations. the chapter builds the anomaly-free fractional-brane count, the three-type classification, the conifold cascade with its rank step and leading-log running, and the deformed-conifold infrared with confinement, chiral-symmetry breaking, and the deformation scale identified; it routes the Klebanov–Strassler supergravity solution and its spectrum to Part III (Klebanov–Strassler dynamics and gravity), the gaugino-condensation dynamics to Part III (Klebanov–Strassler dynamics and gravity), the precision step-matching to Part III, the Seiberg-duality proof to Chapter 7 / Part III, the $\mathcal{N}$=2 and DSB dynamics and the general-cascade duality wall to Part III, the deformed-conifold geometry to Part I (conifold geometry), and the finite-N / sub-leading questions to Part III.

Related constructions. Anomaly-free fractional ranks lie in $\ker(a - a^{\mathsf{T}})$. Their cycle interpretations distinguish deformation, $\mathcal{N} = 2$ and dynamical supersymmetry-breaking types. The repeated move of Chapter 7 reduces rank by M per cascade step, with logarithmically running effective rank. In the Klebanov–Strassler example the confining infrared geometry is $xy - zw = \epsilon^2$, with an S^3 and $\epsilon^2 \sim \Lambda^3$. Part III discusses the supergravity solution and its spectrum. This toric charge problem extends the wrapped-brane picture developed for $H_0 \oplus H_2$ in Volume VI. **Evidence and assumptions.** No

calculation in this chapter is a physical-worldsheet derivation in the untwisted string; worldvolume anomaly arithmetic does not supply the tag [WS]. The physical primary classifications remain source-based here. Their physical inputs are source-supported, while the integral quotient, rank recurrence and specified polynomial identities have separate mathematical statements and hypotheses. Franco–Hanany–Saad–Uranga Franco, Hanany, Saad, et al. (2006), Sections 2, 3.1–3.2.3 and 4, supply the type distinctions, the dP_1 potential (3.1)–(3.2) and DSB analysis, the dynamical-modulus qualification of its runaway, and the SPP potential (4.1). Klebanov–Strassler Klebanov and Strassler (2000), Section 3 and Appendix (114)–(116), supply the cascade and the quantum constraint and branch distinction. The coefficient of the asymptotic charge is imported from the normalized flux construction of Klebanov and Tseytlin (2000), Klebanov and Nekrasov (2000), and Herzog, Klebanov, et al. (2001); no rank subtraction derives it. The distinction between Maxwell and Page charge follows Marolf (2000), Sections 3–4; the normalized conifold flux and large B_2 transformations are discussed in Herzog, Klebanov, et al. (2001). The original calculation files remain unchanged. Their conifold $(1, -1)$ membership check does not test whether it generates the integral quotient; it does not. Their dimension and type comparisons also do not establish confinement, supersymmetry breaking, or a metric. The companion exercises test submitted integral bases, allowed finite cascade sequences, and the SPP potential, adjoint equation and singular locus. Remaining displayed formulas are tracked as text and retain the evidence and limitations stated here. The scale relation uses the confining sector and convention-dependent holomorphic matching; the M phases do not count the full continuous baryonic branch. The $\mathcal{N} = 2$ branch is a flat direction, not a compulsory runaway. All overview statements in this chapter use these same scopes.

References

Bertolini, M., P. Di Vecchia, and R. Marotta (2001). "N=2 four-dimensional gauge theories from fractional branes". In: arXiv: `hep-th/0112195`.

Cachazo, F. et al. (2001). "A Geometric unification of dualities". In: arXiv: `hep-th/0110028`.

Candelas, P. and X. C. de la Ossa (1990). "Comments on Conifolds". In: *Nucl. Phys. B* 342, pp. 246–268. DOI: `10.1016/0550-3213(90)90577-Z`.

Franco, S., A. Hanany, Y.-H. He, et al. (2003). "Duality walls, duality trees and fractional branes". In: arXiv: `hep-th/0306092`.

Franco, S., A. Hanany, F. Saad, et al. (2006). "Fractional branes and dynamical supersymmetry breaking". In: *JHEP* 01, p. 011. DOI: 10.1088/1126-6708/2006/01/011. arXiv: hep-th/0505040.

Herzog, C. P., Q. J. Ejaz, and I. R. Klebanov (2005). "Cascading RG flows from new Sasaki-Einstein manifolds". In: *JHEP* 02, p. 009. DOI: 10.1088/1126-6708/2005/02/009. arXiv: hep-th/0412193.

Herzog, C. P., I. R. Klebanov, and P. Ouyang (2001). "Remarks on the Warped Deformed Conifold". In: arXiv: hep-th/0108101.

Klebanov, I. R. and N. A. Nekrasov (2000). "Gravity duals of fractional branes and logarithmic RG flow". In: *Nucl. Phys. B* 574, pp. 263–274. DOI: 10.1016/S0550-3213(00)00006-8. arXiv: hep-th/9911096.

Klebanov, I. R. and M. J. Strassler (2000). "Supergravity and a confining gauge theory: Duality cascades and chi SB resolution of naked singularities". In: *JHEP* 08, p. 052. DOI: 10.1088/1126-6708/2000/08/052. arXiv: hep-th/0007191.

Klebanov, I. R. and A. A. Tseytlin (2000). "Gravity duals of supersymmetric $SU(N) \times SU(N+M)$ gauge theories". In: *Nucl. Phys. B* 578, pp. 123–138. DOI: 10.1016/S0550-3213(00)00206-7. arXiv: hep-th/0002159.

Marolf, D. (2000). "Chern-Simons terms and the Three Notions of Charge". In: arXiv: hep-th/0006117.

Chapter VII.9

Orientifolds of brane tilings

VII.9.1 The orientifold made combinatorial

Volume VI built the orientifold of an orbifold probe. A stack of D3-branes was placed at $\mathbb{C}^3/\Gamma$, the worldsheet-parity quotient $\Omega(-1)^{F_L}\sigma$ was gauged with a geometric involution σ normalizing the finite group Γ, and the Chan–Paton factors were projected by a matrix γ_Ω. The result folded the McKay quiver: fixed nodes became SO or USp factors, paired nodes were identified, and the chiral matter on the self-identified nodes turned into symmetric or antisymmetric tensors. The worked example was the $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ orientifold, whose cubic gauge anomaly closed on a rank relation between the SO rank and the surviving unitary rank.

That construction leaned, at every step, on the representation theory of Γ. The involution σ acted on the McKay irreducibles; the node folding was a permutation of those irreducibles; the projection γ_Ω lived in the Chan–Paton space the irreducibles organize. As in the rest of Volume VI, the finite group did the work.

This volume has spent eight chapters learning to do without Γ. The brane tiling, or dimer, replaced the McKay quiver as the object that holds the gauge theory of a general toric Calabi–Yau threefold. It is a bipartite graph on a two-torus: faces are gauge nodes, edges are bifundamental fields, and nodes are superpotential terms. It encodes the same quiver and superpotential that the orbifold projection produced in the special case, but it is defined for any toric singularity, orbifold or not. The natural question, and the content of this chapter, is what the orientifold becomes in this language. Where Volume VI's σ was a geometric involution acting through the representation theory of Γ, here it is an *involution of the tiling*, a symmetry of the bipartite graph on the torus. The face and edge identifications become combinatorial. Their physical interpretation still requires the Chan–Paton action, the field parities, and a consistent string lift; the map alone does not determine a normalized Ramond–Ramond charge.

Take-home message of this chapter

For an orbifold the orientifold acts through the representation theory of Γ. For a toric singularity described by a brane tiling, it acts on the faces, edges, and superpotential vertices of the dimer. Point reflections exchange the two vertex colours; line reflections preserve them. The fixed loci and projection signs determine which factors and fields can survive, once the physical orientifold dictionary is supplied. For the uniform dP_0 cubic the nonzero interaction correlates the choices: SO accompanies antisymmetric tensors and $N_0 = N_{12} - 4$, while USp accompanies symmetric tensors and $N_0 = N_{12} + 4$. These are local cubic-anomaly conditions for the stated unflavoured spectrum. They do not settle Ramond–Ramond tadpoles or global consistency.

The chapter uses the scoped projection machinery of Volume VI and makes its sign conditions explicit. Four results carry over from the orbifold case, recalled in a box and applied: the assignment of SO to a fixed node with one O-plane sign and USp with the other; the iff-chain that ties the sign to the matrix γ_Ω; the symmetric and antisymmetric tensor indices that feed the anomaly; and the requirement that the orientifold preserve the superpotential. What is new is how they read on the dimer. For dP_0 we import the embedded fixed-point tiling action of Franco et al., Section 3.5.1, Figure 7 (cited in Status and sources). We check its labelled face and edge fold and cubic interaction, without reconstructing the complete periodic vertex map. Fixed loci are enumerated as points or lines and read as O-planes carrying signs. Projection signs are attached to the fixed loci; a charge vector in a physical RR basis needs further input. And the anomaly-closing rank relation drops out of the folded spectrum of the supplied labelled involution, recovering the field-theory parent as a local check of the imported tiling action.

The plan is the five steps of the construction, each worked on the three parent dimers. First, the orientifold as a tiling involution, and the node, arrow, and superpotential folding it induces (§VII.9.2). Second, the fixed loci as O-planes and the sign that fixes SO versus USp (§VII.9.3). Third, the two-index matter on the self-identified nodes and the crosscap data (§VII.9.4). Fourth, the cubic gauge anomaly and the rank relation it closes on, with the dP_0 orientifold as the anchor (§VII.9.5). Fifth, the three worked examples: the $\mathbb{C}^3$ honeycomb, the conifold square, and the dP_0 tiling (§VII.9.6). The chapter closes with an honest map of what is constructed and what is deferred (§VII.9.7). The boundary of the chapter is the boundary of Volume VI. The local, perturbative projection is worked out for the stated imported tilings. The O-plane's backreacted geometry, its F-theory lift, and the cancellation of its charge over a compact Calabi–Yau are routed to the later volumes.

VII.9.2 The orientifold as an involution of the tiling

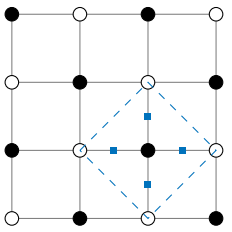
A brane tiling is a bipartite graph on a two-torus T^2. Its faces are the gauge nodes, its edges are the bifundamental chiral fields, and its vertices, two-coloured black and white, are the superpotential terms: a white vertex contributes the clockwise product of the fields on its incident edges with a plus sign, a black vertex the counter-clockwise product with a minus sign. The orientation of the torus and this two-colouring are the

data the superpotential reads.

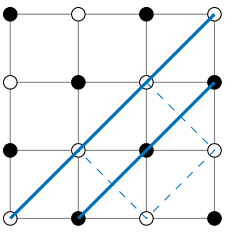
An orientifold is the worldsheet-parity quotient $\Omega(-1)^{F_L}\sigma$, with σ a geometric involution of the Calabi–Yau. On the dimer the combined operation is an *involution of the tiling*: an isometry σ of T^2 of order two,

$$\sigma^2 = \mathrm{id}, \tag{VII.9.1}$$

that is a symmetry of the bipartite graph. Being a symmetry means σ maps the graph to itself preserving incidence: faces go to faces, edges to edges, vertices to vertices, and an edge keeps its endpoints. Worldsheet parity Ω reverses the open-string orientation. In the tiling dictionary, this is compatible with two distinct geometric actions. A point reflection $x \mapsto -x$ preserves the orientation of T^2 and exchanges black and white vertices. A line reflection reverses the orientation of T^2 and preserves their colours. Both induce reversal of the quiver arrows. The fixed points or lines depict orientifold loci in the tiling frame; they are not a literal drawing of an O3- or O7-plane in the original Calabi–Yau. Figure VII.9.1 distinguishes these two operations.



point reflection about $(\frac{1}{2}, 0)$
black $\leftrightarrow$ white



line reflection $(x, y) \mapsto (y, x)$
colours preserved

Figure VII.9.1: Two involutions of the square tiling. The dashed cell has periods $(1, 1)$ and $(1, -1)$; black and white vertices alternate at integer lattice points. Left: rotation by π about $(\frac{1}{2}, 0)$ exchanges the colours. The four blue markers are inequivalent fixed points modulo these periods, not four copies of one corner. Right: reflection in $x = y$ preserves colours and fixes the two circles represented by $x - y = 0, 1$. These actions have different face and edge orbits; the fixed-line action supplies the $SO \times USp$ example below.

The content of (VII.9.1) and the graph-automorphism condition is that the orientifold folding is not a free choice but a constrained one. Three permutations are induced, and they must be mutually consistent:

$$\sigma_F : \{\text{faces}\} \to \{\text{faces}\}, \quad \sigma_E : \{\text{edges}\} \to \{\text{edges}\}, \quad \sigma_V : \{\text{vertices}\} \to \{\text{vertices}\}, \tag{VII.9.2}$$

each an involution, with σ_E compatible with σ_F (the image of an edge between faces i and j is an edge between $\sigma_F(i)$ and $\sigma_F(j)$) and σ_V exchanging colours for point reflections or preserving them for line reflections. The face permutation σ_F is the node folding of Volume VI; the edge permutation σ_E is the arrow folding; the vertex permutation σ_V is the superpotential folding. The single combinatorial map carries

what was, in the orbifold case, three separate statements about the quiver, the matter, and the superpotential.

The quiver-level test. A necessary test of the graph action is visible already in the quiver adjacency. Let a_{ij} be the number of arrows from node i to node j, read from the oriented edges of the tiling. The orientifold acts on an arrow by the face permutation together with the orientation reversal Ω induces, sending the arrow $i \to j$ to the arrow $\sigma_F(j) \to \sigma_F(i)$ (the reversal swaps head and tail). For the involution to be a symmetry of the dimer, the adjacency must be invariant under this action,

$$a_{ij} = a_{\sigma_F(j)\,\sigma_F(i)} \qquad \text{for all } i, j. \tag{VII.9.3}$$

A face permutation that fails (VII.9.3) maps an arrow to a non-arrow, or two arrows to a single one: it is not a symmetry of the tiling, hence not an orientifold. This is the sharp form of the warning that folding the graph in half is a cartoon. To see that it has teeth, consider a three-node quiver with $a_{01} = 1$, $a_{12} = 1$, $a_{20} = 2$ and the candidate involution swapping nodes 1 and 2 while fixing node 0. Equation (VII.9.3) then demands $a_{20} = a_{\sigma(0)\sigma(2)} = a_{01}$, that is $2 = 1$, which fails; the swap is rejected. The condition is not automatic, and it is checked on every fixture before any orientifold data is read.

The three fixtures of this volume each pass. For the $\mathbb{C}^3$ honeycomb the single face is necessarily fixed and (VII.9.3) is trivial; for the conifold square the involution that fixes both faces and reverses the four arrows is adjacency-invariant; and for the dP_0 tiling the involution

$$\sigma_F : \quad 0 \mapsto 0, \qquad 1 \leftrightarrow 2, \tag{VII.9.4}$$

which fixes node 0 and swaps nodes 1 and 2, is the one we use throughout. It is worth running (VII.9.3) once as an explicit matrix calculation rather than only asserting the test passes. The dP_0 quiver has three arrows on each of the directed edges $0 \to 1$, $1 \to 2$, $2 \to 0$, so its adjacency and the candidate involution are

$$A_{dP_0} = \begin{pmatrix} 0 & 3 & 0 \\ 0 & 0 & 3 \\ 3 & 0 & 0 \end{pmatrix}, \qquad \sigma = (0)(12), \tag{VII.9.5}$$

the rows and columns indexed by nodes $0, 1, 2$ and the permutation fixing 0 and swapping 1 and 2. The orientifold action $a_{ij} \mapsto a_{\sigma(j)\,\sigma(i)}$ reflects the matrix across its anti-diagonal in the swapped block and transposes head with tail; carrying it out entry by entry,

$$A^{\Omega\sigma}_{ij} = A_{\sigma(j)\,\sigma(i)} = \begin{pmatrix} 0 & 3 & 0 \\ 0 & 0 & 3 \\ 3 & 0 & 0 \end{pmatrix} = A_{dP_0}, \tag{VII.9.6}$$

the same matrix. For instance the arrow $1 \to 2$ ($a_{12} = 3$) maps to $\sigma(2) \to \sigma(1) = 1 \to 2$, returning $a_{12} = 3$; the arrow $0 \to 1$ ($a_{01} = 3$) maps to $\sigma(1) \to \sigma(0) = 2 \to 0$, returning $a_{20} = 3$; and the chain $0 \to 1 \to 2 \to 0$ goes to $0 \to 2 \to 1 \to 0$ traversed backwards, every multiplicity matching. The adjacency is invariant. This verifies the quiver-level condition, but a dimer symmetry also requires an edge map, vertex incidence, cyclic

order, and compatibility with W. For this example the complete embedded fixed-point dimer action is imported from Franco et al., Section 3.5.1, Figure 7. Below we check the labelled face and edge map and the flavour-preserving cubic transformation; the periodic vertex incidence and cyclic order are not independently reconstructed. (That it is the tiling realization of Volume VI's dP_0 orientifold, which fixed McKay node 0 and paired 1 with 2, is the bridge this chapter builds.)

Fixed and paired faces. The face permutation σ_F partitions the gauge nodes into two kinds. A face with $\sigma_F(i) = i$ is *fixed*: it is invariant as a Chan–Paton sector, and its gauge factor is projected using the O-plane sign of the fixed locus passing through or associated with that sector, so its unitary factor will project to an orthogonal or symplectic one. (The face is invariant as a set, which is weaker than every point of it lying on the fixed locus.) A face with $\sigma_F(i) = j \neq i$ is *paired* with its image, and the two unitary factors $U(N_i)$ and $U(N_j)$ are identified into a single surviving $U(N_i)$ with $N_i = N_j$. Every face is either fixed or in exactly one paired pair, so the face count balances,

$$\#\{\text{fixed faces}\} + 2\,\#\{\text{paired pairs}\} = G, \tag{VII.9.7}$$

with G the total number of gauge nodes. For the three fixtures this reads $1 = 1$ for $\mathbb{C}^3$ (one fixed face), $2 + 0 = 2$ for the conifold (both faces fixed), and $1 + 2 \cdot 1 = 3$ for dP_0 (node 0 fixed, the pair $\{1, 2\}$). The surviving gauge group has one factor per fixed face and one per paired pair; the involution does not create gauge factors, it removes them.

The superpotential tests information that the adjacency has forgotten. Edges with the same tail and head may belong to different cyclic words. One must give their individual images and transform those words, including field parities and transposition. A permutation that preserves arrow counts can still fail this test. We will perform it explicitly for the dP_0 cubic. The existence of the parent dimer and its physical interpretation are separate inputs to that calculation.

VII.9.3 Fixed loci as O-planes: SO versus USp

The fixed loci of the involution are the orientifold planes. A point-reflection involution leaves isolated points of T^2 fixed; a line-reflection involution leaves lines fixed. Each fixed locus carries an O-plane sign, the tiling counterpart of the orbifold-case γ_Ω data, and on a self-identified face that sign fixes the gauge group.

The fixed set of each involution class is a one-line topology computation on the torus $T^2 = \mathbb{R}^2/\mathbb{Z}^2$, worth doing once. For the point reflection $(x, y) \mapsto (-x, -y)$ a fixed point satisfies $(x, y) \equiv (-x, -y) \bmod \mathbb{Z}^2$, that is

$$(x, y) = (-x, -y) \bmod \mathbb{Z}^2 \implies 2x,\, 2y \in \mathbb{Z}, \tag{VII.9.8}$$

so x and y each lie in $\{0, \frac{1}{2}\}$ and the fixed set is the four *half-period* points $(0, 0)$, $(\frac{1}{2}, 0)$, $(0, \frac{1}{2})$, $(\frac{1}{2}, \frac{1}{2})$: four isolated O-plane points. For the line reflection $(x, y) \mapsto (-x, y)$ only the first coordinate is constrained,

$$(x, y) = (-x, y) \bmod \mathbb{Z}^2 \implies 2x \in \mathbb{Z} \;\; (x = 0 \text{ or } \tfrac{1}{2}), \tag{VII.9.9}$$

with y free, so for this rectangular torus the fixed set is the two circles $x = 0$ and $x = \frac{1}{2}$: two O-plane lines on the torus. The half-period points of (VII.9.8) and the two lines of (VII.9.9) give the same fixed-set counts illustrated in Figure VII.9.1. On a rhombic torus a reflection can instead have one fixed circle, so the two-circle answer uses the specified lattice.

The rule is the iff-chain of Volume VI, recalled here and applied, not re-ruled. The orientifold projects the gauge field on a fixed node by

$$A_\mu \;\longmapsto\; -\,\gamma_\Omega\, A_\mu^{\mathsf{T}}\, \gamma_\Omega^{-1}, \tag{VII.9.10}$$

and the projection keeps the subalgebra of generators that survive. Whether γ_Ω is symmetric or antisymmetric is fixed by the O-plane sign at the fixed locus, and the surviving algebra follows:

$$\begin{aligned} \gamma_\Omega^{\mathsf{T}}\gamma_\Omega^{-1} = +1 \;&\Longleftrightarrow\; \gamma_\Omega \sim \mathbb{1} \;\Longleftrightarrow\; SO(N_i), \\ \gamma_\Omega^{\mathsf{T}}\gamma_\Omega^{-1} = -1 \;&\Longleftrightarrow\; \gamma_\Omega \sim J \;\Longleftrightarrow\; USp(N_i), \end{aligned} \tag{VII.9.11}$$

where $\sim$ means a Chan–Paton basis choice and J is the standard nondegenerate antisymmetric form. The compact gauge groups require the usual Hermitian reality condition. Here N_i denotes defining-representation dimension, so $USp(N_i) = USp(2k)$ requires even N_i; its Lie rank is k. We label the two cases by the *Chan–Paton projection sign*

$$\gamma_\Omega^{\mathsf{T}}\gamma_\Omega^{-1} = s_{\text{proj}}, \qquad s_{\text{proj}} = +1 \;\; (SO), \quad s_{\text{proj}} = -1 \;\; (USp), \tag{VII.9.12}$$

the sign of the symmetry of γ_Ω that the iff-chain reads. This is *not* the same symbol as the Ramond–Ramond crosscap charge q_O that the O-plane carries in the covering-space tadpole below: s_{proj} is the projection sign that decides SO versus USp, while q_O is the physical O-plane charge, and the standard names O^- and O^+ refer to the latter (RR-charge) convention. Keeping the two apart is the same discipline, at the level of signs, that keeps the gauge anomaly and the crosscap tadpole apart later. The whole content of (VII.9.11) is taken from Volume VI; the only new statement here is that s_{proj} is a sign the dimer attaches to each fixed locus.

That the chain is genuine is a fact about generators, and it is worth seeing the count once. The projection (VII.9.10) keeps the matrices M with $M^{\mathsf{T}}\gamma_\Omega^{-1} + \gamma_\Omega^{-1}M = 0$. For $\gamma_\Omega = \mathbb{1}$ these are the antisymmetric matrices, the algebra $so(N)$ of dimension $N(N-1)/2$; for $\gamma_\Omega = J$ they are the symplectic matrices, the algebra $usp(N)$ of dimension $N(N+1)/2$. At $N = 4$ the symmetric choice keeps $\frac{4\cdot 3}{2} = 6$ generators ($so(4)$) and the antisymmetric choice $\frac{4\cdot 5}{2} = 10$ generators ($usp(4)$), the two dimensions an explicit generator count confirms. The dimensions differ, the symmetry of γ_Ω decides which, and the O-plane sign decides the symmetry of γ_Ω: the sign on the dimer is load-bearing.

Reading the gauge group off the involution. Putting (VII.9.7) together with (VII.9.11), the surviving gauge group is assembled from the face permutation and the per-fixed-locus signs:

$$G_{\text{orient}} = \Big[\prod_{i \text{ fixed}} \big(SO \text{ or } USp\big)(N_i)\Big] \times \Big[\prod_{\{i,j\} \text{ paired}} U(N_i)\Big], \tag{VII.9.13}$$

each fixed face contributing $SO(N_i)$ when its projection sign is $s_{\rm proj} = +1$ and $USp(N_i)$ when $s_{\rm proj} = -1$, and each paired pair contributing one diagonal $U(N_i)$. On the conifold square, the involution fixes both faces; assigning $s_{\rm proj} = +1$ to the first and $s_{\rm proj} = -1$ to the second gives

$$G_{\rm orient}^{\rm conifold} = SO(N_0) \times USp(N_1), \tag{VII.9.14}$$

a worked instance of (VII.9.13) with one fixed node of each projection sign and no surviving unitary factor. On the dP_0 tiling, node 0 is fixed and the pair $\{1, 2\}$ is identified; with $s_{\rm proj} = +1$ on node 0,

$$G_{\rm orient}^{dP_0} = SO(N_0) \times U(N_{12}), \tag{VII.9.15}$$

one orthogonal factor and one surviving unitary factor $U(N_{12})$ for the diagonal of the two paired nodes. These two readings, $SO \times USp$ for the conifold and $SO \times U$ for dP_0, are the structural difference that the next two sections turn into a difference in the matter and the anomaly: the conifold orientifold keeps no unitary node and so has no cubic anomaly to close, while the dP_0 orientifold keeps one and does.

The fixed-point signs obey a useful necessary consistency condition. Write $\epsilon_p = \pm 1$ for the four signs in the dimer convention, where a positive sign on a face gives SO and a positive sign on an invariant edge gives a symmetric tensor. For the fixed-point class the sign rule is

$$\prod_{p=1}^{4} \epsilon_p = (-1)^{N_W/2}, \tag{VII.9.16}$$

where N_W is the even number of superpotential vertices in a fundamental cell. The honeycomb has $N_W = 2$ and dP_0 has $N_W = 6$, so both require product -1. This rule compares a mesonic path transported around the torus with the signs from crossing fixed points and using F-term relations. It is a restriction on an already specified dimer action, not a construction of a full string quotient from four numbers. It is not the rule for fixed-line signs, and ϵ_p is not a normalized RR charge.

For the uniform dP_0 coupling considered here, let the fixed-node form have symmetry $s = \pm 1$ and each tensor have symmetry $t = \pm 1$. The nonzero cubic will require $st = -1$. In dimer language its four signs are (s, t, t, t), whose product $st^3 = -1$ also obeys (VII.9.16). More general flavour actions can give mixed tensor symmetries and different couplings; the uniform epsilon contraction below does not classify them. The atlas lists representatives rather than a complete catalogue.

A surviving $U(N_{12})$ includes an abelian factor, but this factor is not automatically a free centre-of-mass photon. The Q and T fields carry charges -1 and $+2$ in the convention below. Abelian anomalies, RR couplings, and possible Stueckelberg masses need their own analysis. We compute the $SU(N_{12})^3$ coefficient while retaining $U(N_{12})$ in the group notation; this does not discard its abelian consistency conditions.

VII.9.4 Two-index matter and the crosscap

The matter folds with the gauge group. An edge of the tiling is a bifundamental field, and the edge permutation σ_E sorts the edges into orbits under the orientifold $\mathbb{Z}_2$. Three kinds of orbit occur, and each gives a definite kind of surviving field.

Self-identified edges: two-index tensors. A self-image arrow from a face to its image has two indices of the surviving unitary factor and projects to a symmetric or antisymmetric tensor. An adjoint loop on a fixed face instead becomes a two-index representation of the orthogonal or symplectic factor. In either case the holomorphic projection has the form

$$X = \eta_X \gamma_\Omega X^{\mathsf{T}} \gamma_\Omega^{-1}, \qquad \eta_X = \pm 1. \tag{VII.9.17}$$

For a fixed face, lower the second index with γ_Ω and set $B = X\gamma_\Omega$. If $\gamma_\Omega^{\mathsf{T}} = s\gamma_\Omega$, then $B^{\mathsf{T}} = s\eta_X B$. Thus the tensor symmetry depends on both the CP form and the field parity. In particular, the adjoint projection has $\eta_X = -1$: it is antisymmetric for SO and symmetric after index lowering for USp. The $SU(N)$ indices $N \pm 4$ below apply to tensors on the surviving unitary node, not to arbitrary tensors at an SO or USp node. For the $\mathbb{C}^3$ honeycomb, whose single face carries three adjoint chirals X, Y, Z (the edges of the trivalent tiling, all with tail equal to head), the involution fixes the face and self-identifies each adjoint into a two-index tensor of $SO(N_0)$ or $USp(N_0)$:

$$\{X, Y, Z\} \;\longrightarrow\; 3 \text{ two-index tensors of } SO(N_0) \text{ (or } USp(N_0)), \tag{VII.9.18}$$

For the O3 choice preserving $\mathcal{N} = 4$, all three have $\eta_X = -1$ and are adjoints of the projected group. Other parity choices describe different orientifolds. The involuted honeycomb gives three tensors and no bifundamental, a simple possible instance: the parent theory has no chiral bifundamentals to fold, so the entire spectrum is two-index tensors.

Paired edges: folded bifundamentals. An edge and its distinct image under σ_E are not two independent fields; the orientifold identifies them into one surviving field. This is the lesson the orbifold case taught and the dimer makes visible: a paired orbit $\{(i \to j), (\sigma j \to \sigma i)\}$ of two arrows folds to a *single* folded bifundamental. Counting the two arrows as independent fields is the double-count the orientifold removes. An edge running between a fixed face and a paired face folds to a bifundamental of the SO or USp factor times the surviving U factor.

The dP_0 tiling exhibits both. Under the involution (VII.9.4) the nine arrows of the dP_0 quiver, three on each directed edge $0 \to 1$, $1 \to 2$, $2 \to 0$, sort into two kinds of orbit. The three arrows $1 \to 2$ run between the two paired faces; each is self-identified by the fold (the image of $1 \to 2$ is $\sigma(2) \to \sigma(1) = 1 \to 2$), so they project to three two-index tensors of the surviving $U(N_{12})$. The three arrows $0 \to 1$ pair with the three arrows $2 \to 0$ (the image of $0 \to 1$ is $\sigma(1) \to \sigma(0) = 2 \to 0$), folding to three bifundamentals running between the fixed node 0 and the surviving node $\{1, 2\}$. The folded spectrum is therefore

$$3 \text{ two-index tensors of } U(N_{12}) \;+\; 3 \text{ bifundamentals } SO(N_0)\text{–}U(N_{12}), \tag{VII.9.19}$$

six surviving fields out of the nine the parent quiver had: folding identifies fields, it does not create them. The involuted dP_0 dimer gives exactly this, three tensors and three folded bifundamentals. Table VII.9.1 records the edge-orbit map that produces them.

It is worth writing the folded superpotential as an explicit stepwise derivation rather than only naming it, in three steps. *Step one*, the parent. The dP_0 cubic is the

parent dP_0 edge orbit	orientifold image	surviving field
$1 \to 2$ arrows	self-identified	three two-index tensors T^a of $U(N_{12})$
$0 \to 1$ arrows	paired with $2 \to 0$ arrows	three bifundamentals Q^a between $SO(N_0)$ and $U(N_{12})$

Table VII.9.1: The dP_0 edge-orbit map under the involution (VII.9.4). The three self-identified $1 \to 2$ arrows (image $\sigma(2) \to \sigma(1) = 1 \to 2$) project to three two-index tensors T^a of the surviving $U(N_{12})$; the three $0 \to 1$ arrows pair with the three $2 \to 0$ arrows (image $\sigma(1) \to \sigma(0) = 2 \to 0$) and fold to three bifundamentals Q^a between $SO(N_0)$ and $U(N_{12})$, the flavour index $a = 1, 2, 3$ running over the three arrows on each edge. Nine parent arrows fold to six surviving fields: folding identifies, it does not create.

determinant-like contraction of the three flavoured bifundamentals,

$$W_{dP_0} = \epsilon_{abc}\, \mathrm{Tr}\big(X^a_{01}\, X^b_{12}\, X^c_{20}\big), \tag{VII.9.20}$$

with a, b, c the three flavours on each edge and ϵ_{abc} the $SU(3)$ invariant. *Step two*, the field identifications dictated by the edge-orbit map of Table VII.9.1: the self-identified $1 \to 2$ edges become the two-index tensors T^a on the surviving $U(N_{12})$, and the fixed-paired $\{0 \to 1,\ 2 \to 0\}$ orbits fold to the single bifundamental Q between $SO(N_0)$ and $U(N_{12})$, its two legs appearing as Q^a and its transpose,

$$X^a_{12} \;\longmapsto\; T^a, \qquad X^a_{01} \;\sim\; Q^a, \qquad X^a_{20} \;\sim\; s(Q^a)^{\mathsf{T}} G^{-1}, \quad G^{\mathsf{T}} = sG, \tag{VII.9.21}$$

the second leg of the paired orbit appearing as the transpose because the orientifold image of the $0 \to 1$ leg is the $2 \to 0$ leg. *Step three*, substitute (VII.9.21) into (VII.9.20); the projected superpotential is

$$W_{\mathrm{proj}} \;\sim\; \epsilon_{abc}\, Q^a_{Iu} (T^b)^{uv} Q^c_{Jv}\, H^{IJ}, \qquad H^{IJ} = s(G^{-1})^{JI}. \tag{VII.9.22}$$

Here I, J are fixed-node indices, u, v are unitary indices, Q is an antifundamental of $U(N_{12})$, and T has two fundamental indices. The inverse form contracts the fixed-node indices; it cannot be suppressed when comparing the orthogonal and symplectic cases. There are three independent Q fields, not three Q fields plus three new images.

To check the sign, set $H^{\mathsf{T}} = sH$ and $(T^b)^{\mathsf{T}} = tT^b$. Exchange the two Q slots. The flavour epsilon contributes -1, the fixed-node form contributes s, and the tensor contributes t. Since chiral superfield components commute, the polynomial equals $-st$ times itself. A nonzero coupling therefore requires

$$st = -1: \qquad SO \Rightarrow t = -1, \qquad USp \Rightarrow t = +1. \tag{VII.9.23}$$

For $st = +1$ the displayed cubic vanishes identically. The parent matrix action can also be checked before projection: for $A = X_{01}$, $B = X_{12}$, $C = X_{20}$ choose

$$A \mapsto GC^{\mathsf{T}}, \quad B \mapsto tB^{\mathsf{T}}, \quad C \mapsto sA^{\mathsf{T}} G^{-1}, \qquad W \mapsto -stW. \tag{VII.9.24}$$

This action squares to the identity because $G^{\mathsf{T}} = sG$. Transpose reverses the matrix order and the exchange of the first and third flavours reverses epsilon. Thus $st = -1$ preserves this parent polynomial. The geometric flavour action, superspace measure, and full string lift remain physical inputs; polynomial invariance alone does not establish them.

The crosscap data. A crosscap is an orientation-reversing identification on a world-sheet, not a boundary. O-planes couple to closed strings through crosscap states; open-string boundaries describe D-branes. On the dimer the crosscap data is the per-fixed-locus O-plane sign together with the covering-space Ramond–Ramond charge the O-plane carries. The amplitude that computes that charge, the Klein-bottle and Möbius-strip diagrams of the unoriented string, is foundations and is cited, not re-derived; what the dimer supplies is which fixed loci there are and what sign each carries. That charge enters a tadpole condition of its own, kept carefully separate from the gauge anomaly of the next section. The covering-space Ramond–Ramond tadpole is a charge-balance condition between the D-brane charge Q_D and the O-plane charge Q_O, which we denote schematically by

$$Q_D + Q_O = 0. \tag{VII.9.25}$$

The O-plane charge Q_O is in general nonzero. The fixed-locus projection signs do not alone determine its magnitude or its components in twisted RR sectors. The precise normalization of Q_D and Q_O on a common node/charge basis, and the compact completion that cancels the total Ramond–Ramond charge over a closed Calabi–Yau, are *not* derived here; they are foundations and global model-building, routed to the Status apparatus (orientifold backreaction and compact completion). What *is* computed explicitly in this chapter is the downstairs gauge-anomaly sum of the next section, and (VII.9.25) must not be conflated with it. The tadpole is an upstairs covering-space charge balance involving the O-plane source. The gauge anomaly is a downstairs field-theory condition on the folded matter. The two are read from different data on the dimer. Keeping them apart is the discipline of the next section.

VII.9.5 The anomaly-closing rank relation

The folded theory must be consistent, and the sharpest consistency condition is that its cubic non-abelian gauge anomalies vanish. This is where the chapter recovers its field-theory parent: the anomaly of a surviving unitary node, fed by the folded two-index and bifundamental matter, closes on a relation between ranks, and for dP_0 that relation is the one Volume VI derived from the McKay quiver.

The ingredients are the cubic anomaly indices of the surviving representations, all taken from Volume V. A fundamental contributes $A(\square) = +1$ and an antifundamental -1; the two-index tensors carry

$$A(\text{symmetric}) = N + 4, \qquad A(\text{antisymmetric}) = N - 4, \qquad A(\text{fundamental}) = 1, \tag{VII.9.26}$$

for $SU(N)$ with $N \geq 3$, normalized to the fundamental. Conjugate representations reverse the sign. The real or pseudoreal representations occurring at the fixed nodes

have zero local cubic trace. This representation statement avoids an exception to the blanket claim that every orthogonal algebra has no cubic invariant: $so(6) \simeq su(4)$ has complex spinor representations with a cubic anomaly, which are absent here.

The index formulas can be derived without assuming them. For a traceless diagonal $SU(N)$ generator with charges q_i, put $S_3 = \sum_i q_i^3$. The ordered tensor square has

$$\sum_{i,j} (q_i + q_j)^3 = 2N S_3, \qquad \sum_i (2q_i)^3 = 8 S_3. \tag{VII.9.27}$$

The first equality uses $\sum_i q_i = 0$. Half the sum and difference of these two expressions gives the symmetric and antisymmetric traces, $(N+4)S_3$ and $(N-4)S_3$. Their sum $2N$ is the cubic index of $\Box \otimes \Box$, not of the adjoint. The adjoint weights $q_i - q_j$ occur in opposite pairs, so

$$A(\text{adj}) = 0. \tag{VII.9.28}$$

For example, at $N = 3$ the antisymmetric is the antifundamental and its index is -1; at $N = 4$ it is real and its index is zero. At $N = 2$ there is no local $SU(2)^3$ invariant, so one must not infer a rank equation by extrapolating the normalization.

The dP_0 sum. Assemble the cubic $SU(N_{12})^3$ anomaly of the one surviving unitary node of the dP_0 orientifold from the folded matter (VII.9.19), reading each contribution off Table VII.9.2. The three two-index tensors on $U(N_{12})$ each contribute their index $N_{12} \pm 4$ (the upper sign for the symmetric projection, the lower for the antisymmetric), and the three folded bifundamentals to the fixed $SO/USp(N_0)$ node each contribute the rank N_0 of the other node weighted by the orientifold chirality $s_Q = \pm 1$ (the sign of whether the surviving leg is a fundamental or an antifundamental). The fixed-node representations have zero cubic trace of their own.

matter at $U(N_{12})$	multiplicity	cubic index contribution
two-index tensor T	3	$3\,(N_{12} \pm 4)$
bifundamental Q to $SO/USp(N_0)$	3	$3\,s_Q\,N_0$

Table VII.9.2: The local cubic $SU(N_{12})^3$ anomaly sum of the dP_0 orientifold. The three two-index tensors contribute their index $N_{12} \pm 4$ each (the $+4$ for the symmetric projection, the -4 for the antisymmetric); the three folded bifundamentals to the fixed $SO/USp(N_0)$ node contribute the other rank N_0 weighted by the chirality $s_Q = \pm 1$ each; the fixed-node representations have zero local cubic trace. The sum, set to zero, closes on the rank relation $N_0 = N_{12} \pm 4$. For the uniform nonzero cubic the minus branch is SO with antisymmetric tensors and the plus branch is USp with symmetric tensors.

Summing the two rows of Table VII.9.2, the anomaly is

$$\mathcal{A}\big(SU(N_{12})^3\big) = 3\,s_Q\,N_0 + 3\,(N_{12} \pm 4), \qquad s_Q = \pm 1, \tag{VII.9.29}$$

the factor of three counting the three tensors and the three bifundamentals. Setting it to zero is a single linear condition on the two ranks. Taking the standard chirality $s_Q = -1$,

$$\mathcal{A}\big(SU(N_{12})^3\big) = -3\,N_0 + 3\,(N_{12} \pm 4) = 0, \tag{VII.9.30}$$

which solves to the rank relation

$$N_0 = N_{12} \pm 4, \tag{VII.9.31}$$

the $+4$ for the symmetric tensor and the -4 for the antisymmetric, with no added flavour matter. For the nonzero uniform coupling, (VII.9.23) assigns the $+4$ branch to $USp(N_0)$ and the -4 branch to $SO(N_0)$. Here $N_{12} \geq 3$, $N_0 \geq 0$, and symplectic N_0 must be even. A zero fixed-node dimension is a degenerate limit, not an interacting fixed gauge factor. This is the conditional relation Volume VI derived for the dP_0 orientifold from the folded McKay quiver, recovered here from the labelled face and edge orbits of the imported point-reflection dimer. Assembling the sum (VII.9.29) from those field orbits and solving it gives (VII.9.31) for both signs, the symmetric branch returning $N_0 = N_{12} + 4$ from the anomaly $-3N_0 + 3N_{12} + 12$ and the antisymmetric branch $N_0 = N_{12} - 4$ from $-3N_0 + 3N_{12} - 12$. Figure VII.9.2 draws the whole fold, the parent dP_0 quiver on the left and the folded orientifold theory on the right.

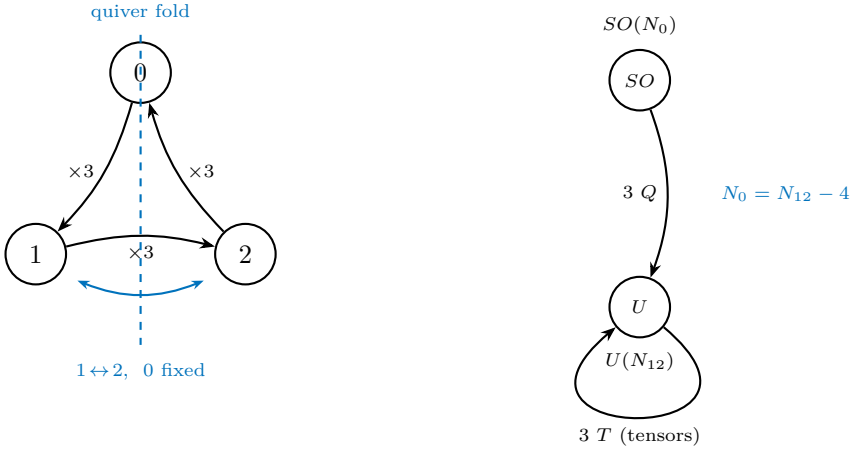


Figure VII.9.2: The dP_0 orientifold fold, the chapter's anchor. *Left*: the parent dP_0 quiver, three gauge nodes $0, 1, 2$ (the node interiors carry the gauge factors) with three chiral bifundamentals on each directed edge $0 \to 1$, $1 \to 2$, $2 \to 0$ (each labelled $\times 3$). The dashed line marks the quiver face map (VII.9.4): it fixes node 0 and swaps nodes 1 and 2; it is not an axis of the embedded dimer action. *Right*: the folded orientifold theory. The fixed node 0 projects to $SO(N_0)$; the paired nodes $1, 2$ merge into one surviving $U(N_{12})$; the three self-identified $1 \to 2$ arrows become three two-index tensors T of $U(N_{12})$ (the self-loop box), and the $\{0 \to 1,\ 2 \to 0\}$ arrow orbits fold into three independent bifundamentals Q between $SO(N_0)$ and $U(N_{12})$; the image leg is the same field's transposed representative, not another set of three. The cubic $SU(N_{12})^3$ anomaly closes on $N_0 = N_{12} - 4$ for the SO branch with three antisymmetric tensors. The alternative uniform branch is USp with symmetric tensors and $N_0 = N_{12} + 4$. This is a schematic quiver cartoon; the data is the surviving gauge group, the folded matter, and the closing relation.

The branch is fixed by the tensor symmetry. The sign in (VII.9.31) tracks the tensor symmetry and the chosen field parities. For our uniform cubic it is correlated with the fixed gauge group by (VII.9.23). A symmetric-tensor orientifold closes on $N_0 = N_{12} + 4$, an antisymmetric-tensor orientifold on $N_0 = N_{12} - 4$, and assigning the wrong branch to a given projection is an inconsistency the sum detects. The other consistency the sum enforces is that the anomaly (VII.9.29) is genuinely non-zero: the dP_0 orientifold is chiral, its cubic anomaly does not vanish identically, and the rank relation is forced. This is the property that distinguishes it from the next case.

The conifold case without a cubic rank constraint. Not every tiling orientifold forces a rank relation, and it is important to say so plainly. The conifold orientifold of (VII.9.14) fixes both faces, giving $SO(N_0) \times USp(N_1)$ with no surviving unitary node. There is no $SU(N)^3$ cubic anomaly to close, because its defining fixed-node representations have zero cubic trace. The folded matter is two bifundamentals running between the two fixed nodes, and the quiver is non-chiral, with antisymmetric adjacency part $a - a^{\mathsf{T}} = 0$. There is therefore *no* forced local anomaly rank relation analogous to the dP_0 relation (VII.9.31): there is no surviving unitary node and no rank relation from this cubic test for the conifold and the $\mathbb{C}^3$ orientifolds, the case without a cubic rank constraint. A unitary-node variant of a conifold-type orientifold with self-identified two-index matter must be examined case by case; it is not a universal dP_0-like rank shift, and the chapter does not assert one. The rank constraint follows when the specified matter gives a nonzero rank-dependent $SU(N)^3$ coefficient. A unitary node with tensors need not do so in every spectrum.

Two equations, not one. The rank relation (VII.9.31) and the crosscap tadpole (VII.9.25) answer different questions and must not be conflated, a point Volume VI was careful about and this chapter keeps. The rank relation is the field-theory gauge anomaly of the surviving unitary node, computed downstairs from the folded matter; the crosscap tadpole $Q_D + Q_O = 0$ is the covering-space Ramond–Ramond charge balance, a statement upstairs about the O-plane source. They involve different data (the cubic indices (VII.9.26) versus the D-brane and O-plane charges) and bind different quantities. For the dP_0 orientifold the anomaly closes on $N_0 = N_{12} \pm 4$, while the covering tadpole is the separate charge-balance statement (VII.9.25) with its own normalization and compact completion not derived here. Treating them as one equation is the standard error, and the dimer keeps them apart because they are read from different data on it.

For several surviving unitary nodes one writes a system of cubic anomaly equations. Some equations can vanish identically or be dependent; their rank is a result of the chosen spectrum. Additional flavours change the equations. The chapter works one representative involution per example and names the catalogue of inequivalent involutions, which belongs to the atlas chapter; the calculation here reproduces the anomaly-closing rank relation from the supplied labelled fold and its physical projection data.

The absence of a local cubic anomaly leaves further consistency conditions. At a USp factor, an odd number of fundamental Weyl fermions can produce a mod-two global anomaly. The conifold has $2N_0$ such fundamentals, which is even. The uniform dP_0 USp branch has $3N_{12}$; its even $N_0 = N_{12} + 4$ also makes N_{12} even. These elementary counts do not compute the full K-theory charge or a compact RR tadpole. For instance,

adding flavour branes can repair a local anomaly while changing both the superpotential and the global charge problem.

VII.9.6 Worked examples

The three parent dimer examples of this volume each carry an orientifold, and worked end to end they illustrate three outcomes of the construction: a pure two-index spectrum with no unitary node, an $SO \times USp$ pairing with no local cubic rank constraint, and the chiral dP_0 orientifold that recovers the field-theory parent. Table VII.9.3 collects them; the prose works each.

tiling	involution	gauge group	folded matter	rank relation
$\mathbb{C}^3$ (honeycomb)	node 0 fixed	$SO(N_0)$	3 tensors	none (no U node)
conifold (square)	both fixed	$SO(N_0) \times USp(N_1)$	2 bifund.	none (non-chiral)
dP_0 (triangle)	0 fix., $1 \leftrightarrow 2$	$SO(N_0) \times U(N_{12})$	3 anti.+3 bif.	$N_0 = N_{12} - 4$

Table VII.9.3: The three worked tiling orientifolds. The $\mathbb{C}^3$ honeycomb folds to a pure two-index spectrum on a single SO (or USp) node; the conifold square folds to a non-chiral $SO \times USp$ pair with no surviving unitary node and so no rank relation (the case without a cubic rank constraint); the dP_0 tiling folds to $SO \times U$ with chiral two-index matter, whose cubic anomaly closes on $N_0 = N_{12} - 4$ for this SO branch. The USp alternative has symmetric tensors and $N_0 = N_{12} + 4$, with even N_0.

The $\mathbb{C}^3$ honeycomb

The $\mathbb{C}^3$ tiling is the hexagonal honeycomb: a single gauge node, three adjoint chirals X, Y, Z on the three edges of the trivalent graph, and the cubic superpotential $W = \mathrm{Tr}\, X[Y, Z]$ read from the two vertices of the unit cell. The orientifold is a point reflection fixing the single face. With one face there is no pairing to do; the face is fixed, its $U(N_0)$ projects to $SO(N_0)$ for the $+1$ sign (or $USp(N_0)$ for -1), and the three adjoints self-identify into three adjoint two-index tensors of that factor, as in (VII.9.18). The folded spectrum is three tensors and no bifundamental; there is no surviving unitary node, so no cubic $SU(N)^3$ anomaly and no rank relation. The projected superpotential is the σ-invariant image of $\mathrm{Tr}\, X[Y, Z]$, the orientifold projection of the $\mathcal{N}$=4 cubic, written in the tensors. This is a simple tiling orientifold and the warm-up for the others: the orientifold of $\mathcal{N}$=4 super Yang–Mills on $\mathbb{C}^3$, seen on the dimer (Figure VII.9.3).

The conifold square

The conifold tiling is the square dimer: two gauge nodes, two arrows A_1, A_2 from node 0 to node 1 and two arrows B_1, B_2 from node 1 to node 0, and the quartic superpotential $W = \mathrm{Tr}(A_1 B_1 A_2 B_2 - A_1 B_2 A_2 B_1)$ read from the two vertices of the square cell. The orientifold is a colour-preserving line reflection fixing both faces as sets. Each fixed face



Figure VII.9.3: The $\mathbb{C}^3$ honeycomb orientifold, a simple tiling case. *Left*: the hexagonal $\mathbb{C}^3$ dimer under a point reflection (cross) about the centre of its single face; the face is fixed and projects to $SO(N_0)$ for the $+1$ sign (or $USp(N_0)$ for -1), the point reflection exchanging the white and black vertices, the two superpotential terms. *Right*: the folded theory. With one face there is no pairing to do: the three adjoint chirals X, Y, Z self-identify into three antisymmetric adjoint tensors T of $SO(N_0)$ (the self-loop box), and there is no surviving unitary node, hence no cubic $SU(N)^3$ anomaly and no rank relation. This is the orientifold of $\mathcal{N}$=4 super Yang–Mills, seen on the dimer. The figure is a schematic; the data is the single fixed face, its projection, and the three-tensor folded spectrum.

carries an O-plane sign; taking opposite signs gives

$$G_{\text{orient}}^{\text{conifold}} = SO(N_0) \times USp(N_1), \tag{VII.9.32}$$

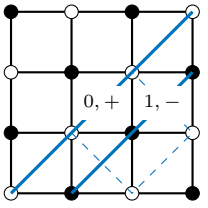
as in (VII.9.14). The four arrows fold under the involution: the two A's pair with the two B's (the image of $0 \to 1$ is $\sigma(1) \to \sigma(0) = 1 \to 0$, a B-type arrow), so the folded matter is two bifundamentals running between the $SO(N_0)$ and $USp(N_1)$ nodes. The quiver is non-chiral, $a - a^{\mathsf{T}} = 0$, the defining representations at both fixed factors have zero cubic trace, and there is no surviving unitary node: there is therefore no cubic anomaly and no forced rank relation. With a nondegenerate antisymmetric form J at the USp node, the image identification can be written $B_i = JA_i^{\mathsf{T}}$ up to the chosen field phases. Substitution gives

$$W_{\text{coni}} \sim \operatorname{Tr}\big(A_1 J A_1^{\mathsf{T}} A_2 J A_2^{\mathsf{T}} - A_1 J A_2^{\mathsf{T}} A_2 J A_1^{\mathsf{T}}\big)\,. \tag{VII.9.33}$$

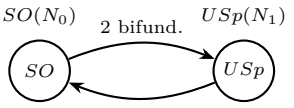
The form contracts the symplectic indices, and the orthogonal indices are traced. This displays the quartic of the two surviving matrices; its identification with the fixed-line string construction uses the orientifold dictionary. The conifold orientifold is the unconstrained-cubic case: a clean $SO \times USp$ probe theory, with no rank relation from this local cubic test. Global gauge anomalies and physical RR tadpoles still have to be checked independently (see Figure VII.9.4).

The dP_0 orientifold

The $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ tiling is the bipartite graph with three white and three black vertices, nine edges arranged so that the quiver has three arrows on each directed edge $0 \to 1$,



two fixed lines, opposite signs



non-chiral, $a - a^{\mathsf{T}} = 0$
no forced rank relation

Figure VII.9.4: The conifold orientifold, the unconstrained-cubic case. *Left*: the conifold square dimer under a line reflection; both faces are *fixed*, not swapped, and carry opposite O-plane signs (+ on one, − on the other). *Right*: the folded quiver. The two fixed faces project to $SO(N_0)$ (the + sign) and $USp(N_1)$ (the − sign), and the four arrows fold into two bifundamentals running between them. The quiver is non-chiral ($a - a^{\mathsf{T}} = 0$), the representations at both fixed factors have zero cubic trace, and there is no surviving unitary node, so this cubic test imposes no rank relation. It does not decide global anomalies or RR tadpoles. The figure is a schematic; the data is the both-fixed structure with opposite signs and the resulting $SO \times USp$ pairing.

$1 \to 2$, $2 \to 0$, and a cubic superpotential built from the six trivalent vertices. The orientifold is the involution (VII.9.4), fixing node 0 and swapping nodes 1 and 2. The compatible embedded point reflection is the published dP_0 dimer action of Franco et al., Section 3.5.1, Figure 7. Starting from that action, the labelled fold and local consistency calculation run as follows:

- The face permutation fixes node 0 and pairs $\{1, 2\}$, so the gauge group is $SO(N_0) \times U(N_{12})$ of (VII.9.15), one orthogonal factor on the fixed O-plane and one surviving unitary factor on the diagonal of the paired nodes.
- The edge permutation self-identifies the three $1 \to 2$ arrows into three two-index tensors of $U(N_{12})$ and pairs the three $0 \to 1$ arrows with the three $2 \to 0$ arrows into three folded $SO(N_0)$–$U(N_{12})$ bifundamentals, the folded spectrum (VII.9.19).
- For the supplied flavour-preserving field map, the six cubic words transform as (VII.9.24). The projected W is written in the surviving tensors and bifundamentals, and its nonzero uniform coupling requires $st = -1$.
- The cubic $SU(N_{12})^3$ anomaly (VII.9.29), fed by the three tensors of index $N_{12} - 4$ and three antifundamental bifundamentals, closes on $N_0 = N_{12} - 4$ for the displayed SO branch. The USp alternative requires symmetric tensors and gives $N_0 = N_{12} + 4$.

The resulting gauge group, matter, projected superpotential, and rank relation match the field-theory parent: $SO(N_0) \times U(N_{12})$, three antisymmetric tensors, three independent bifundamentals, and $N_0 = N_{12} - 4$. This is the bridge between Volume VI's McKay fold and the labelled fold of the imported point-reflection dimer. Figure VII.9.5 shows the local half turn and its face and colour action; the full torus embedding remains the cited input.

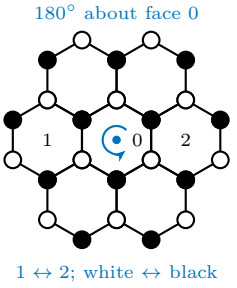


Figure VII.9.5: Local honeycomb patch of the $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ tiling. The blue dot is the isolated half-turn centre in face 0; rotation by 180° fixes that face and exchanges faces 1 and 2. It also sends the plotted white vertex at $(0, 0.66)$ to the black vertex at $(2.286, -0.66)$, with the other plotted vertices paired in the same way. Thus the displayed rotation has the correct colour action for the fixed-point class. The complete periodic torus embedding, edge and vertex action and remaining fixed points are imported from Franco et al., Section 3.5.1, Figure 7, not proved by this finite patch. Figure VII.9.2 gives the labelled quiver fold and the local SO rank relation.

VII.9.7 What is and is not given

This chapter uses the established tiling involutions to determine what their labelled folds and projection data compute. For dP_0, the complete embedded point-reflection action is imported from Franco et al., Section 3.5.1, Figure 7. The local picture checks its half-turn colour and face action; the supplied field map and cubic are checked separately. The chapter enumerates the fixed-locus classes and reads the SO versus USp assignment off their signs; it folds the matter into two-index tensors and bifundamentals on the self-identified and paired nodes; it distinguishes projection signs from normalized crosscap charges; and it assembles the cubic gauge anomaly and the conditional rank relation for dP_0. For the uniform unflavoured cubic the SO branch has antisymmetric tensors and $N_0 = N_{12} - 4$; the USp branch has symmetric tensors and $N_0 = N_{12} + 4$, with even N_0. The calculation requires $N_{12} \geq 3$ and nonnegative CP dimensions. The conifold and $\mathbb{C}^3$ orientifolds are worked as cases without a cubic rank constraint, with no rank relation from the local cubic test, and the construction is laid out for one representative involution per fixture.

What is deferred, and to where, is the boundary of Volume VI. The O-plane is represented here by a fixed locus and projection signs; its normalized covering-space charge is additional physical data; its *backreaction*, the tension and charge sources and the geometry they produce, is not solved. The orientifold is the local, perturbative quotient; its *F-theory lift*, the realization of the O7-plane as a degeneration of an elliptically fibered geometry in the Sen limit, is not constructed. And the crosscap charge is carried locally; the *compact tadpole cancellation*, the gluing of the local O-plane charge into a compact Calabi–Yau with total Ramond–Ramond neutrality and the brane and flux content that cancels it globally, is not done. These three, the backreaction, the F-theory lift, and the global completion, are the chapter's load-bearing honest gap,

routed to the F-theory and flux-compactification volumes. The amplitude analysis that determines crosscap couplings, using the Klein bottle together with the annulus and Möbius strip, is foundations of the unoriented string and is cited, not re-derived; so are the precise normalization of the covering-space tadpole $Q_D + Q_O = 0$ and its compact completion. The chapter names the O-plane charge and keeps it separate from the gauge anomaly, but does not derive the tadpole normalization or assume the charge vanishes.

The projection machinery itself, the SO versus USp assignment, the iff-chain tying it to γ_Ω, the two-index tensor indices, and the W-compatibility constraint, is taken from Volume VI and made combinatorial here. The anomaly-closing rank relation $N_0 = N_{12} \pm 4$ is recovered from the supplied labelled fold as a local consistency check: the imported dimer action and projection data specify the folded spectrum, then the cubic indices give the sum. This does not reconstruct the full periodic edge and vertex action or its string lift. The atlas collects representative involutions and rank relations; it is not a complete classification; this chapter works one representative per fixture. The supersymmetric Calabi–Yau examples are chosen to preserve four supercharges; the flat O3 example has the enhanced sixteen-supercharge theory; the wider O-plane-type catalogue and the two-dimensional Spin(7) orientifolds of brane brick models are named and routed, not carried. And, as throughout the volume, the leading classical projection statements carry the probe α' and g_s caveats of Volume VI: the orientifold is a leading-order, local, probe statement, with subleading string effects and compact questions routed onward. The representation indices and conditional rank shifts computed here are exact at the stated finite ranks; they are not a large-rank approximation.

Exercises: reconstruct the projection

1. Starting from the labelled dP_0 arrows $A_a : 0 \to 1$, $B_a : 1 \to 2$, $C_a : 2 \to 0$ and $\sigma_F = (0)(12)$, take $A_a \leftrightarrow C_a$ and $B_a \mapsto B_a$. List the face and individual edge orbits, recompute $A_{\sigma(j)\sigma(i)}$, and check every arrow endpoint. Repeat for the conifold with $A_i \leftrightarrow B_i$ and for the three fixed honeycomb loops. Solve the point- and line-reflection congruences on $\mathbb{R}^2/\mathbb{Z}^2$. Which extra information would be needed to turn these answers into an embedded dimer action?

2. Expand the six terms in $\epsilon_{abc}\,\mathrm{Tr}(A_a B_b C_c)$. Apply (VII.9.24), reversing each transposed word before using cyclicity. Compute its overall parity for three choices: $(s,t) = (+1,-1)$, $(-1,+1)$ and $(+1,+1)$. Independently expand the projected contraction with $H = I_2$ or the canonical antisymmetric 2×2 form. Determine when it is a nonzero polynomial. Evaluate the four-sign product for (s,t,t,t) and compare with $N_W = 6$. Explain why agreement of these sign tests does not supply RR tadpole cancellation.

3. Using $q = (1, \ldots, 1, 1 - N)$, compute the symmetric, antisymmetric and adjoint cubic indices by summing weights. For each tensor branch combine them with the edge-orbit multiplicities to find the affine anomaly as a function of N_0 and its zero. Test $N = 3, 4, 6, 8$ as applicable, enforcing nonnegative CP dimensions and even symplectic dimension. Identify the degenerate antisymmetric $N = 4, N_0 = 0$ case. What changes if one counts a field and its orientifold image separately?

See `benchmark/tasks.json` and `benchmark/PROTOCOL.md` for the public structured-answer format. Their calculations concern specified permutations, polynomials and representations; the source-supported identification with a string background remains an additional premise.

Status and sources

Evidence and conventions are collected here. [WS] means a controlled physical-worldsheet derivation in the untwisted string; worldvolume algebra and a passing symbolic calculation are not WS. [MATH] requires stated hypotheses and a proof basis. The present physical dimer and orientifold identifications are source-backed premises with unassigned primary classification here. No physical-worldsheet derivation is supplied here. The finite orbit calculations, matrix projection algebra, weight-sum identities and conditional anomaly equation are separated from those premises in the companion inventory. The source premises are stated separately from the finite calculations.

§VII.9.1 and §VII.9.2: the dimer action. The parent dictionary uses bipartite tiling conventions, McKay multiplicities, fundamental-at-tail arrows, superpotential normalization, as in Chapter 3. The construction sources are Hanany and Kennaway 2005; Franco, Hanany, Kennaway, et al. 2006. For the orientifold action see Franco, Hanany, Krefl, et al. 2007, §2.3, §3.1 footnote 5 and §4.1 footnote 10: point reflections exchange colours, line reflections preserve them. Adjacency invariance is necessary, not sufficient for a symmetry of the embedded dimer. The finite face and edge exercises check explicit maps; they do not reconstruct the embedding. Figure VII.9.1 uses the square lattice with periods $(1,1),(1,-1)$. Figure VII.9.5 directly shows a half turn of a finite honeycomb patch: its marked centre fixes face 0, pairs the plotted faces $1,2$ and exchanges vertex colours. The complete periodic fixed-point dimer action is imported from Franco, Hanany, Krefl, et al. 2007, §3.5.1, Figure 7. The local patch does not establish torus periods, all fixed points, or a complete edge/vertex permutation.

§VII.9.3: fixed sets and signs. The rectangular-torus fixed sets follow by solving the congruences in (VII.9.8)–(VII.9.9). The one-circle rhombic alternative is described in Franco, Hanany, Krefl, et al. 2007, §4.1, Figure 12. The CP algebra reuses orientifold projection and the corrected VI.8 candidate, `eq:V8-Aproj` and its explicit I, J calculations; physical projection sources include Gimon and Polchinski 1996; Witten 1998. The sign rule (VII.9.16) is imported from Franco, Hanany, Krefl, et al. 2007, §3.1.1 and §3.4, equation (3.12). Its parity evaluation is elementary arithmetic; the physical rule is a cited premise. Dimer signs, CP signs and RR charges are not interchangeable. The $U(1)$ is retained per quiver gauge groups and ranks; its dynamics and anomaly cancellation are not computed.

§VII.9.4: fields and interaction. The spectrum uses the supplied flavour-preserving image map and the local CP projection, orientifold projection. Three image pairs leave three independent Q fields. The uniform cubic uses superpotential normalization and the explicit form G; (VII.9.23) and (VII.9.24) are algebraic statements over characteristic zero with $G^{\mathsf{T}} = sG$, G invertible and $T^{\mathsf{T}} = tT$. They reproduce the corrected VI.8 candidate's `eq:V8-tensorcorrelation` and `eq:V8-dP0Winv`. The physical dP_0 examples appear in Franco, Hanany, Krefl, et al. 2007, §3.5.1, Table 5 and §5.3.1, equations (5.12)–(5.13). Other rows of Table 5 have mixed tensor symmetries; they are not ruled out by this chapter's restriction to a uniform epsilon coupling. Crosscap-state foundations are Gimon and Polchinski 1996; Johnson 2003; no normalized twisted RR vector is computed. The schematic (VII.9.25) needs that vector and a common charge basis.

§VII.9.5: cubic coefficients and ranks. The mathematical proof basis for (VII.9.26) is the traceless weight sum (VII.9.27), for $SU(N)$, $N \geq 3$, and characteristic zero. The adjoint has zero cubic index, while $A(\square \otimes \square) = 2N$. The physical spectrum is an imported premise; its local coefficient is one $SU(N_{12})^3$ coefficient, $-3N_0 + 3(N_{12} + 4t)$, with three Q and three T and no extra flavours. This is not an anomaly match between independently constructed backgrounds. The rank relation is conditional on this matter, the cubic domain and, for the uniform nonzero interaction, $st = -1$. Fixed-node statements concern the real or pseudoreal representations here, not all representations of every orthogonal algebra. Mod-two anomalies and RR/K-theory consistency are separate obligations. For sources of additional flavours and tadpoles see Franco, Hanany, Krefl, et al. 2007, §3.1.2, §3.5.1 equation (3.13), and §5.4; local singularity constructions are also discussed in Aldazabal et al. 2000.

§VII.9.6 and §VII.9.7: scope. The flat O3 example has three adjoint chirals and sixteen supercharges: Franco, Hanany, Krefl, et al. 2007, §3.1.1, cases 1–2, and Witten 1998. The $SO \times USp$ conifold is a fixed-line construction. See Franco, Hanany, Krefl, et al. 2007, §4.3.2, the conifold discussion around (4.13), and the direct two-face fold here. Equation (4.9) describes the wider orbifold family; (4.10) is its three-factor $N = 2$ example, not the two-face conifold. The point-reflection conifold in §3.2.2 has different matter. The SO dP_0 branch has antisymmetric tensors and $N_0 = N_{12} - 4$; the uniform symmetric branch instead has USp and $N_0 = N_{12} + 4$. Figures VII.9.2–VII.9.5 and Table VII.9.3 carry these same restrictions. Backreaction, F-theory lifts and compact completion remain orientifold backreaction and compact completion; crosscap amplitudes remain orientifold and crosscap amplitudes. The atlas is a list of representatives, not a full orientifold classification. The source-only physical premises include the string lift, preserved supersymmetry, abelian consistency, normalized RR charges and global/K-theory completion.

Projection and anomaly scope. The fold and anomaly calculation require the representation and superpotential projections together. The coefficient $2N$ concerns the displayed matter representations. The RR shift q_O requires sector data independent of the finite-rank anomaly calculation.

Related constructions. The orientifold projection and folded-quiver conventions are those of Volume VI, Chapter 8, applied here to a brane tiling. A graph fold requires additional projection data. The local probe uses the stated $\alpha\prime$ and g_s limits. The finite-rank cubic coefficients are exact in their stated representation domain.

References

Aldazabal, G. et al. (2000). "D-branes at singularities: A Bottom up approach to the string embedding of the standard model". In: *JHEP* 08, p. 002. DOI: 10.1088/1126-6708/2000/08/002. arXiv: hep-th/0005067.

Franco, S., A. Hanany, K. D. Kennaway, et al. (2006). "Brane dimers and quiver gauge theories". In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Franco, S., A. Hanany, D. Krefl, et al. (2007). "Dimers and Orientifolds". In: *JHEP* 09, p. 075. DOI: 10.1088/1126-6708/2007/09/075. arXiv: 0707.0298 [hep-th].

Gimon, E. G. and J. Polchinski (1996). "Consistency conditions for orientifolds and D-manifolds". In: *Phys. Rev. D* 54, pp. 1667–1676. DOI: 10.1103/PhysRevD.54.1667. arXiv: hep-th/9601038.

Hanany, A. and K. D. Kennaway (2005). "Dimer models and toric diagrams". In: arXiv: hep-th/0503149.

Johnson, C. V. (2003). *D-Branes*. Cambridge Monographs on Mathematical Physics. Cambridge University Press. DOI: 10.1017/CBO9780511606540.

Witten, E. (1998). "Baryons and branes in anti-de Sitter space". In: *JHEP* 07, p. 006. DOI: 10.1088/1126-6708/1998/07/006. arXiv: hep-th/9805112.

Chapter VII.10

D1-branes at toric Calabi–Yau fourfolds: brane brick models and 2d $(0,2)$ theories

VII.10.1 From D3-branes to D1-branes

The dimer has turned a D3-brane at a toric Calabi–Yau threefold into a 4d $\mathcal{N}$=1 quiver. Its faces give gauge groups, its edges give chiral fields, and its vertices give superpotential terms. What replaces this picture when the probe is a D1-brane and the transverse cone is a Calabi–Yau fourfold? For generic $SU(4)$ holonomy the brane now supports a 2d $(0,2)$ theory. Its matter includes both chiral and Fermi multiplets, so a graph that records only one kind of edge can no longer carry the full gauge-theory data.

The corresponding geometric object is a **brane brick model**, a tessellation of T^3. Bricks give gauge nodes, faces give the two matter species, and edges give their holomorphic interactions. Face orientation distinguishes a chiral from a Fermi multiplet. This extra choice is the essential change from the dimer on T^2. For each Fermi multiplet the interactions come as a pair (E_a, J_a), constrained by $\sum_a \mathrm{Tr}(E_a J_a) = 0$. We will see that the condition becomes the Jacobi identity for the flat fourfold $\mathbb{C}^4$, while an orbifold projection builds a four-node example on $\mathbb{C}^4/\mathbb{Z}_4$.

Take-home message of this chapter

A D3-brane at a toric threefold gives a 4d $\mathcal{N}$=1 theory encoded by a dimer on T^2. A D1-brane at a toric fourfold with generic $SU(4)$ holonomy gives a 2d $(0,2)$ theory encoded by a brane brick model on T^3. The oriented and unoriented faces carry chiral and Fermi multiplets. Their E- and J-terms obey $\sum_a \mathrm{Tr}(E_a J_a) = 0$; in the $\mathbb{C}^4$ model, the Jacobi identity gives this cancellation. A local order-three triality has a class and rank cycle constrained by the nonabelian gauge-anomaly condition.

We first count the preserved supercharges and read a quiver from the cells of the brick model. The chiral and Fermi fields then determine the gauge anomaly and the

(E, J) condition. The flat model makes both calculations explicit; its $\mathbb{Z}_4$ orbifold shows how the same data project to several nodes. We then examine the local rank and field-class cycle of triality and place the dimer and brick constructions in the graded-quiver sequence. The elliptic genus, c-extremization, and infrared dynamics require a separate treatment in Part III.

VII.10.2 The D1-brane at a toric fourfold and its brane brick model

Place a stack of N D1-branes at the tip of a toric Calabi–Yau fourfold cone and ask for the two-dimensional theory on their worldvolume. The logic is the probe logic of Volume VI and of this volume's Chapter 1: the brane's transverse space is the cone, the brane's gauge theory must reconstruct that cone as a branch of its moduli space, and the unbroken supersymmetry is fixed by the holonomy of the transverse geometry.

Count the supercharges first, because the count names the theory. A D1-brane in flat ten-dimensional type IIB preserves sixteen real supercharges on its two-dimensional worldvolume, the 2d $\mathcal{N}{=}(8,8)$ content of the flat-space D-string. Putting the brane transverse to a Calabi–Yau fourfold breaks the transverse rotation group to the holonomy $SU(4) \subset SO(8)$, and the assumed Ricci-flat Calabi–Yau metric has covariantly constant spinors associated with its $SU(4)$ holonomy. In this holonomy setting the trivial canonical class supplies the corresponding holomorphic volume form; by itself it does not fix the spinors of an arbitrary metric or at the singular tip. The holonomy-singlet spinors determine the surviving supercharges. A Calabi–Yau fourfold with generic $SU(4)$ holonomy keeps two of the sixteen, both real and of the same two-dimensional chirality, and the surviving algebra is the two-dimensional $(0,2)$ supersymmetry algebra:

$$\text{D1@CY}_4 \ (SU(4) \text{ holonomy}) \quad \longrightarrow \quad 2d \ (0,2) \quad (Q_+ = 2). \tag{VII.10.1}$$

The notation $(0,2)$ records the split by worldsheet chirality: zero left-moving and two right-moving supercharges. Special geometries can preserve more. The flat $\mathbb{C}^4$ example retains the D1-brane's full $(8,8)$ supersymmetry, although we write its fields in $(0,2)$ multiplets below. This is the fourfold analogue of the statement that a flat D3-brane starts with sixteen worldvolume supercharges, and an $SU(3)$-holonomy threefold leaves four, the 4d $\mathcal{N}{=}1$ algebra. One complex dimension up in the transverse geometry is two real worldvolume dimensions down in the brane, and half the supercharges that survive the holonomy.

Why exactly two survive is worth one sentence of representation theory, not a full spinor derivation. The holonomy $SU(4)$ acts on the eight transverse real directions, and the preserved supersymmetries are the spinors invariant under that $SU(4)$. A generic $SU(4)$ holonomy leaves one complex invariant spinor in the transverse spin representation; after the D1 projection this becomes two real worldsheet supercharges of one chirality, hence $(0,2)$. The same counting at $SU(3)$ holonomy on the six transverse directions of the threefold leaves the larger invariant subspace that gives the four supercharges of 4d $\mathcal{N}{=}1$; the fourfold's larger holonomy keeps fewer.

It is worth putting the two cases, and the next, on one line, because the pattern is the spine of the chapter's closing synthesis. Write m for the rung. A D$(5-2m)$-brane

probing a toric Calabi–Yau of complex dimension $n = m + 2$ has worldvolume real dimension $d = 6 - 2m$, and its theory is organized by a periodic tessellation of T^{m+1}:

m	Dp (p=5−2m)	CY$_n$ (n=m+2)	d=6−2m	tessellation
1	D3	CY$_3$	4	T^2 (dimer)
2	D1	CY$_4$	2	T^3 (brick)
3	D(−1)	CY$_5$	0	T^4

(VII.10.2)

The m=1 row is the D3-brane at a threefold, the 4d $\mathcal{N}$=1 dimer story of Chapters 1–9. The m=2 row is this chapter. The m=3 row, a D(−1)-instanton at a fivefold giving a zero-dimensional theory, is named here and constructed nowhere in this volume; it returns in the closing synthesis only as an outlook pointer. The sequence stops at m=3 because $d = 6 - 2m$ reaches zero there: a worldvolume of negative dimension is not a field theory. These integer relations, $p = 5 - 2m$, $d = 6 - 2m$, $n = m + 2$, are the dimension relations developed in the closing section; for now they place the D1-brane at the m=2 rung, with worldvolume dimension two and transverse complex dimension four.

The transverse fourfold and the surviving $(0, 2)$ supersymmetry are the input. The output is a two-dimensional $(0, 2)$ gauge theory, read off the combinatorial object we now introduce.

The brane brick model is the fourfold analogue of the dimer. A dimer is a bipartite graph drawn on T^2; a brane brick model is a tessellation of T^3, a cell decomposition of the three-torus into solid bricks meeting along faces, the faces meeting along edges. The dimensional structure of the dictionary is the whole content of the comparison, so state it as a table.

cell of the tessellation	dimer on T^2 (m=1)	brick model on T^3 (m=2)
top cell	face (2d) $\leftrightarrow$ gauge node	brick (3d) $\leftrightarrow$ gauge node
codim 1	edge (1d) $\leftrightarrow$ matter	face (2d) $\leftrightarrow$ matter
codim 2	node (0d) $\leftrightarrow$ W-term	edge (1d) $\leftrightarrow$ J/E-term

(VII.10.3)

Read the table by codimension and the parallel is mechanical. The top-dimensional cell is the gauge node: a face of the dimer, a brick of the brick model. The codimension-one cell is matter: a dimer edge separates two faces, so a bifundamental sits between two gauge nodes; a brick-model face separates two bricks, so again a bifundamental sits between two gauge nodes. The codimension-two cell is a piece of the superpotential: a dimer node is a W-term, a brick edge is a J- or E-term. So the gauge group is

$$\prod_i U(N_i), \qquad \text{one factor per brick,} \tag{VII.10.4}$$

the honest unitary output of the Chan–Paton data, never written as a bare SU, with the regular ranks $N_i = N$ for a stack of N D1-branes.

The chiral/Fermi split. Here the brick model says something the dimer cannot. A dimer edge carries exactly one species of matter, a chiral bifundamental. A brick-model

face carries one of *two* species, and which one is fixed by the *orientation* of the face. The two-dimensional degeneration rule is the standard brane-tiling orientation rule:

$$\text{oriented face} \;\leftrightarrow\; \text{chiral } X_{ij} \in (\mathbf{N}_i, \overline{\mathbf{N}}_j), \qquad \text{unoriented face} \;\leftrightarrow\; \text{Fermi } \Lambda_a. \tag{VII.10.5}$$

An oriented face carries an arrow from brick i to brick j and is a chiral multiplet X_{ij}, in the fundamental of $U(N_i)$ and the antifundamental of $U(N_j)$, exactly as the dimer's oriented edge convention assigns a bifundamental to a directed arrow. An unoriented face, by contrast, is a Fermi multiplet Λ_a; in the standard brick dictionary the chiral faces are oriented and the Fermi faces are the unoriented *square* (four-sided) faces, the shape distinguishing the two species in the actual tessellation. Following the standard brane-brick colour convention we draw chiral faces gray and Fermi faces red, so that colour reads off the species at a glance (Figure VII.10.1). The orientation of a two-dimensional face is a genuine new distinction that a one-dimensional dimer edge does not have, and it is precisely the room the brick model needs to encode two matter species at once. This is the central structural fact of the chapter: where the dimer had one kind of arrow, the brick model has two kinds of face, and the difference is the $(0, 2)$ chiral/Fermi distinction.

A brane brick model is therefore not a three-dimensional dimer or a graph drawn in three dimensions. Its top cells are solid three-dimensional bricks, not faces; its matter lives on two-dimensional faces, not one-dimensional edges; and the orientation of those faces carries information with no two-dimensional analogue. Figure VII.10.1 draws the single brick of the $\mathbb{C}^4$ model, a truncated octahedron, with its faces coloured by species.

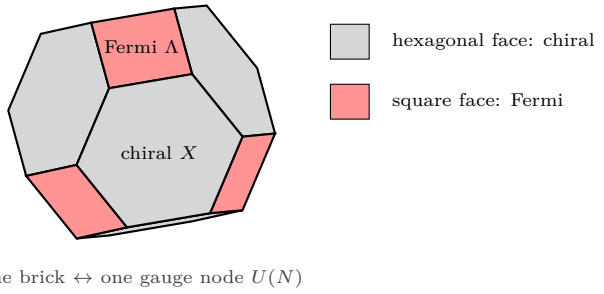


Figure VII.10.1: The single brick of the $\mathbb{C}^4$ brane brick model: a *truncated octahedron*, the Voronoi cell that tessellates T^3, following Figure 6 of (Franco, Lee, Seong, and Vafa 2016). Its faces carry the matter, coloured by species in the standard brane-brick convention (Franco, Lee, Seong, and Vafa 2016; Franco and Yu 2021): the eight *hexagonal* faces are gray *chiral* faces and the six *square* faces are red *Fermi* faces. Opposite faces are identified across the three-torus, so the eight hexagons give four chiral fields $X_1, \ldots, X_4$ and the six squares give three Fermi fields $\Lambda_1, \Lambda_2, \Lambda_3$ (Figure VII.10.2); the brick itself is the gauge node $U(N)$. The even-sided chiral faces and the four-sided Fermi faces are the two $(0, 2)$ matter species, the distinction a one-dimensional dimer edge cannot carry.

The toric data is read off the brick model as in the dimer case, lifted one dimension. The dimer's Kasteleyn determinant produced a Newton polygon, a lattice polygon in

two dimensions, the toric diagram of the threefold. The brick model produces, by the fourfold analogue of that route, a lattice *polytope* in three dimensions: a height-one lattice polytope whose three-dimensional shape is the toric diagram of the Calabi–Yau fourfold. For $\mathbb{C}^4$ this polytope is the unit tetrahedron,

$$\Delta_{\mathbb{C}^4} = \mathrm{conv}\{(0,0,0),(1,0,0),(0,1,0),(0,0,1)\}, \tag{VII.10.6}$$

the fourfold analogue of the $\mathbb{C}^3$ triangle, with four vertices and no interior point. The toric diagram of a fourfold is genuinely three-dimensional, as the diagram of a threefold was genuinely two-dimensional; this is the lift of the Newton-polygon route, and it is the construction-level invariant the duality of §VII.10.5 preserves.

The object that carries out this map is the fourfold analogue of a perfect matching, the **brick matching**. A perfect matching of a dimer was a subset of edges covering each node exactly once, and the perfect matchings mapped to the lattice points of the Newton polygon (with multiplicities) in Chapter 5. A brick matching is the same kind of object one rung up: it is the GLSM homogeneous coordinate of a brane brick model, a choice of chiral and Fermi data that covers the J- and E-term structure in the prescribed way. The brick matchings map to the lattice points of the three-dimensional toric CY_4 polytope, generally with multiplicities, exactly as the perfect matchings mapped to the points of the two-dimensional Newton polygon. They are the combinatorial bridge from the T^3 tessellation to the fourfold geometry, and the formal definition and the matching-to-polytope algorithm belong to the toric and GLSM foundations, cited and not built here.

VII.10.3 The $(0,2)$ quiver: chiral and Fermi multiplets

The brick dictionary returns a quiver with two arrow species, and it is worth fixing the two-dimensional $(0,2)$ multiplet vocabulary before drawing it, because it is the vocabulary the rest of the chapter computes in. A two-dimensional $(0,2)$ theory is built from three multiplets.

- The *chiral multiplet* X holds a complex scalar ϕ and a right-moving Weyl fermion ψ_+. It is the dimensional counterpart of the 4d $\mathcal{N}{=}1$ chiral multiplet: a propagating boson with its right-moving superpartner. A chiral field on the arrow $i \to j$ is a bifundamental, $X_{ij} \in (\mathbf{N}_i, \overline{\mathbf{N}}_j)$.

- The *Fermi multiplet* Λ holds a single left-moving Weyl fermion λ_- and *no* propagating scalar. It has no four-dimensional $\mathcal{N}{=}1$ analogue; it exists because two dimensions allow a multiplet whose only on-shell degree of freedom is a left-mover. Its superspace constraint contains the E-term below.

- The *vector multiplet*, the gauge sector, holds the gauge field and a *left*-moving adjoint gaugino. Its field-strength multiplet is itself a Fermi-type object, but it is the gauge degree of freedom, not matter, and it carries the kinetic term and the field strength. The left-moving chirality of the gaugino is load-bearing: it enters

the gauge anomaly with the same sign as the Fermi multiplets, and without it the chiral content alone would not be anomaly-free.

The two matter species are the point. Where the dimer's quiver had one kind of arrow, one chiral bifundamental per edge, the brick model's quiver has two. A chiral arrow $i \to j$ is an oriented face. A Fermi multiplet Λ_a is an unoriented face; conventionally it is drawn as an undirected segment between the two bricks it separates. Its gauge representation and the orientation of its E- and J-paths still matter. Figure VII.10.2 draws the single-node $\mathbb{C}^4$ quiver: one gauge node $U(N)$, four chiral self-loops (the four transverse complex directions of $\mathbb{C}^4$, all adjoints of the single node), and three Fermi self-loops, the dashed segments.

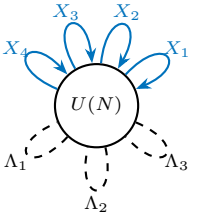


Figure VII.10.2: The $\mathbb{C}^4$ $(0,2)$ quiver. One gauge node $U(N)$ (a single brick); four chiral multiplets $X_1, \ldots, X_4$ (blue, directed self-loops, the four $\mathbb{C}^4$ directions, all adjoints of the single node); three Fermi multiplets $\Lambda_1, \Lambda_2, \Lambda_3$ (dashed, undirected self-loops, no propagating boson). The two arrow species, directed-chiral and undirected-Fermi, are the brick model's two face orientations. The number of chirals exceeds the number of Fermis here by exactly one. The E- and J-terms of the Fermis become commutator equations for the nonabelian matrix fields; the abelian $\mathbb{C}^4$ coordinates themselves have no algebraic relations.

The counts already say something. The $\mathbb{C}^4$ model has four chiral fields and three Fermi fields: four oriented faces and three unoriented faces. Four is the number of transverse complex directions of $\mathbb{C}^4$, the $(0,2)$ fields associated with its four coordinates. The three Fermi multiplets belong to the $(0,2)$ decomposition of the D1 worldvolume theory. Their E- and J-terms impose commutator equations on matrix-valued fields, as the next section makes explicit. The four coordinates of the abelian $\mathbb{C}^4$ geometry have no algebraic relation among themselves. For a stack of N D1-branes the four chirals and three Fermis are all adjoints of the single $U(N)$. This is the $(0,2)$ description of the maximally supersymmetric D1-brane theory. For $N > 1$ its matrix fields and gauge interactions make it an interacting nonabelian theory; "flat" refers to the transverse geometry.

The gauge anomaly of a two-dimensional $(0,2)$ theory is the two-dimensional analogue of the 4d cubic gauge anomaly, and the relevant count is right-movers minus left-movers. A chiral multiplet carries a right-moving fermion and enters with a plus sign; a Fermi multiplet carries a left-moving fermion and enters with a minus sign; and the vector multiplet carries the *left*-moving adjoint gaugino and so also enters with a minus sign.

The coefficient for the nonabelian factor at a gauge node is therefore

$$\sum_{\text{chiral } X} T(R_X) \;-\; \sum_{\text{Fermi } \Lambda} T(R_\Lambda) \;-\; T(\text{adj}) \;=\; 0, \tag{VII.10.7}$$

the chiral contribution against the Fermi and the adjoint-gaugino contributions, with $T(R)$ the Dynkin index of the representation. For the $\mathbb{C}^4$ node the four chirals, the three Fermis, and the vector are all adjoints of the single $U(N)$, so every index is $T(\text{adj})$ and the condition reads

$$4\,T(\text{adj}) \;-\; 3\,T(\text{adj}) \;-\; T(\text{adj}) \;=\; (4-3-1)\,T(\text{adj}) \;=\; 0. \tag{VII.10.8}$$

The vector contribution is what makes this work: the four chirals against the three Fermis alone would leave $4-3=1$ in units of $T(\text{adj})$; the left-moving gaugino closes the count. This is a nontrivial test of the $SU(N)$ factor for $N \geq 2$. At the single $U(1)$ node the adjoint fields are neutral, so the same numerical coefficient does not test an abelian anomaly. We check the nonabelian content balance in the examples and leave abelian anomaly cancellation and inflow to the general treatment.

One normalization point will reappear when §VII.10.5 writes the same anomaly in the flavour-counting form $N_1 + N_2 - N_3 - 2N_c = 0$. In the Gadde–Gukov–Putrov node convention the flavour numbers N_1, N_2, N_3 count fundamental flavour multiplicities with both orientations included, so the adjoint gaugino contributes $2N_c$ in the same normalization; this is the local-node version of the $-T(\text{adj})$ term in (VII.10.7). The vector enters as $T(\text{adj})$ when the matter is written by Dynkin index and as $2N_c$ when it is written by flavour count at a $U(N_c)$ node, the two forms of one accounting.

VII.10.4 The J-terms, the E-terms, and a Jacobi identity

A four-dimensional $\mathcal{N}$=1 quiver is defined by its arrows *and* its superpotential, one holomorphic gauge invariant. A two-dimensional $(0,2)$ quiver needs two holomorphic data, and they are attached to the Fermi multiplets, which is why the chiral/Fermi split of the previous section forces the structure of this one. Each Fermi multiplet Λ_a comes with a pair of holomorphic functions of the chiral fields:

$$\begin{aligned} &E_a(X) \text{ enters the Fermi constraint} && \bar{D}_+\Lambda_a = E_a, \\ &J_a(X) \text{ is the superpotential partner} && \mathcal{L} \supset \int d\theta^+\, \Lambda_a J_a. \end{aligned} \tag{VII.10.9}$$

The E-term is the right-hand side of the Fermi multiplet's defining constraint: its fermion's superpartner is set equal to a holomorphic function E_a of the chiral fields. The J-term is the genuine $(0,2)$ superpotential coupling, the holomorphic interaction $\int d\theta^+\, \Lambda_a J_a$. The two are exchanged under the multiplet conjugation $\Lambda_a \leftrightarrow \bar{\Lambda}_a$, so E and J are not independent kinds of object but the two faces of one structure, the $(0,2)$ generalization of the single 4d W. In the brick model each is read off the edges bounding

the Fermi face: J_a and E_a are each a difference of two oriented plaquettes, two terms each,

$$J_a = J_a^+ - J_a^-, \qquad E_a = E_a^+ - E_a^-, \tag{VII.10.10}$$

the two-dimensional analogue of the dimer fact that each edge sits in exactly two superpotential terms of opposite sign. This is the entry that says there are two holomorphic data, not one: the single four-dimensional superpotential is the special case where the Fermi multiplets and their E-terms are absent, the m=1 rung where matter is purely chiral. Figure VII.10.3 draws the local picture: a single square Fermi face whose four bounding edges are the chiral monomials, the two opposite edges of one pair combining to $J = J^+ - J^-$ and the two of the other pair to $E = E^+ - E^-$.

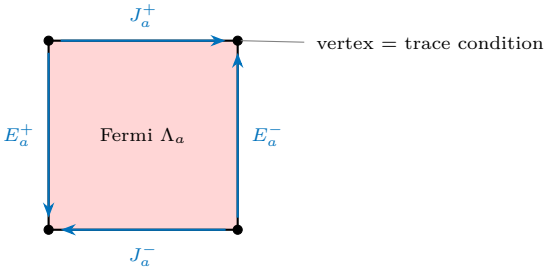


Figure VII.10.3: The Fermi face of a brane brick model and its E/J data. Every Fermi face is a *square* (the Fermi multiplet Λ_a, an unoriented red face). Its four *edges* are the codimension-two objects of the brick model, each associated with an oriented chiral monomial; opposite edges carry the same term with opposite signs, so the top and bottom edges are the J pair $J_a = J_a^+ - J_a^-$ and the left and right edges are the E pair $E_a = E_a^+ - E_a^-$, the two-term structure of (VII.10.10) (the $(0,2)$ analogue of the dimer fact that an edge sits in two superpotential terms of opposite sign). The four *vertices* (codimension three) are the trace-condition relations: this codimension dictionary, bricks $\to$ gauge groups, faces $\to$ matter, edges $\to$ J/E monomials, vertices $\to$ trace conditions, is that of the brane-brick and BFT_2 construction (Franco, Lee, Seong, and Vafa 2016; Franco and Yu 2021). The geometry of the tessellation, not a choice, supplies the structure.

The consistency condition. The J- and E-terms cannot be chosen without a compatibility condition. Supersymmetry of the $(0,2)$ action requires that, summed over all Fermi multiplets and traced over the gauge indices, the product of each Fermi's E- and J-terms vanish:

$$\boxed{\sum_a \mathrm{Tr}\big(E_a\, J_a\big) = 0} \tag{VII.10.11}$$

This is the central consistency equation of the brick model. In the brane-brick construction, the local edge and vertex structure supplies the cancellations needed for this trace condition. The explicit $\mathbb{C}^4$ calculation below shows one such cancellation. The corresponding statement for a general tessellation uses the geometric construction and is not derived from the single-brick example.

The $\mathbb{C}^4$ model: a Jacobi identity. For the flat $\mathbb{C}^4$ example the condition (VII.10.11) reduces to a Jacobi identity. The four chiral fields are the matrices X_1, X_2, X_3, X_4, all adjoints of the single $U(N)$, and the three Fermi multiplets carry the E- and J-terms

$$E_a = [X_4,\, X_a], \qquad J_a = [X_{a+1},\, X_{a+2}], \qquad a = 1,2,3 \ \ \text{(indices cyclic in } \{1,2,3\}), \tag{VII.10.12}$$

each a commutator of two of the four chiral matrices: the E-term pairs X_4 with X_a, and the J-term pairs the cyclically next two of X_1, X_2, X_3. Each is manifestly a two-term object as in (VII.10.10), $E_a = X_4X_a - X_aX_4$ and $J_a = X_{a+1}X_{a+2} - X_{a+2}X_{a+1}$. Form the sum in (VII.10.11). Using cyclicity of the trace to move X_4 to the front of each term,

$$\sum_{a=1}^{3} \mathrm{Tr}\big([X_4, X_a]\,[X_{a+1}, X_{a+2}]\big) = \mathrm{Tr}\Big(X_4 \sum_{a=1}^{3} \big[X_a,\, [X_{a+1}, X_{a+2}]\big]\Big), \tag{VII.10.13}$$

where the inner sum is the cyclic sum of nested commutators of X_1, X_2, X_3. That cyclic sum is the left-hand side of the Jacobi identity, which vanishes identically for any three matrices:

$$\big[X_1, [X_2, X_3]\big] + \big[X_2, [X_3, X_1]\big] + \big[X_3, [X_1, X_2]\big] = 0. \tag{VII.10.14}$$

Hence $\sum_a \mathrm{Tr}(E_aJ_a) = \mathrm{Tr}(X_4 \cdot 0) = 0$, and the $\mathbb{C}^4$ brick model satisfies its $(0,2)$ consistency condition by the Jacobi identity, for any gauge group $U(N)$ and any value of the matrices. The cancellation follows from cyclicity of the trace and the matrix Jacobi identity for any size. Another pairing need not satisfy it. Pairing J_a with $[X_a, X_{a+1}]$ instead of the cyclic complement, for instance, does not produce the Jacobi structure and the trace does not vanish. The condition is therefore not vacuous: it selects the correct E/J assignment from the many wrong ones.

Why two holomorphic data, not one

A 4d $\mathcal{N}$=1 theory has one superpotential W because it has one matter species, the chiral multiplet. A 2d $(0,2)$ theory has *two*, the J-terms and the E-terms, because the Fermi multiplet is constrained: its defining equation $\bar{D}_+\Lambda_a = E_a$ fixes the E-term, and the superpotential coupling $\int \Lambda_a J_a$ fixes the J-term. The two are tied by $\sum_a \mathrm{Tr}(E_aJ_a) = 0$, the $(0,2)$ trace condition. For $\mathbb{C}^4$ this condition is the Jacobi identity. *Do not look for a single superpotential in a $(0,2)$ theory; look for the (E_a, J_a) pairs and check their trace closes.*

The Jacobi identity is the marquee derivation of the chapter, and it is worth registering what it buys. It shows that the brick model's central consistency condition is not an external constraint imposed by hand but a property the construction enforces automatically, in the simplest case as the most basic identity of matrix algebra. A general brick model satisfies (VII.10.11) by the closure of its T^3 tessellation, the higher-dimensional version of this cancellation; the $\mathbb{C}^4$ case is the one where the closure reduces to a single named identity.

VII.10.5 $(0,2)$ triality as a cube move

One toric threefold can admit several brane tilings, the toric phases related by the urban-renewal dimer move of Chapter 7 (the **square move**, acting on a quadrilateral $N_f = 2N_c$ face), which is Seiberg duality acting on the gauge theory. One toric fourfold, likewise, can admit several brane brick models, and the move that relates them is the two-dimensional analogue of Seiberg duality: $(0,2)$ **triality**, due to Gadde, Gukov, and Putrov, realized on the brick model as a **cube move**. In the bipartite-field-theory description, the cube move acts locally on a suitable brick, in analogy with the dimer square move. The field-content cycle and its rank constraint can be written explicitly. Establishing the complete move on a given periodic brick model requires its mesons, E/J data, and toric diagram.

The decisive difference from the four-dimensional case is the *order* of the move. Seiberg duality is an involution: dualize a node twice and the theory returns. Triality is order three:

$$\mathcal{T}^3 = \mathrm{id} \quad (\text{up to relabeling}), \qquad \mathcal{T}^2 \neq \mathrm{id}, \tag{VII.10.15}$$

applying the triality move at a node three times returns the original theory, up to relabeling, and twice does not. Three brick models are cyclically related at each dualized node, where four dimensions had two. The move acts on the matter at the node by the Gadde–Gukov–Putrov transformation: the rank of the dualized node changes by the local flavour rule (the anomaly-balanced node rule, N_c' fixed by the flavours entering the node, not by a vague total count), and the three classes of field at the node, the incoming chirals, the outgoing chirals, and the Fermis, are cyclically reshuffled into one another. Schematically, with the node's field classes written as a triple,

$$\mathcal{T}: \ (\text{incoming chiral}) \ \to \ (\text{outgoing chiral}) \ \to \ (\text{Fermi}) \ \to \ (\text{incoming chiral}), \tag{VII.10.16}$$

a three-cycle on the three classes, which is exactly why the move has order three: the three-cycle returns to the identity after three applications and not before. Figure VII.10.4 draws the cycle, the three classes carried into one another by $\mathcal{T}$, making $\mathcal{T}^3 = \mathrm{id}$ visually plain. New composite "meson"-type fields are generated at the dualized node, and the (E, J) data is rewritten, with fields that become massive integrated out through their J/E equations, the $(0,2)$ analogue of the dimer move's integration of the two-valent fields it generates.

The local node and the rank rule. The rank rule is made precise on the Gadde–Gukov–Putrov local model, the $(0,2)$ analogue of the $N_f = 2N_c$ window that fixed the Seiberg dual in Chapter 7. Isolate a $U(N_c)$ node with N_1 fundamental chiral multiplets, N_2 antifundamental chiral multiplets, N_3 fundamental Fermi multiplets, and the adjoint vector. For $N_c \geq 2$, the nonabelian $SU(N_c)$ gauge-anomaly coefficient of (VII.10.7), chiral against Fermi against adjoint, vanishes when

$$N_1 + N_2 - N_3 - 2N_c = 0, \tag{VII.10.17}$$

the $(N_1 + N_2)$ chiral fundamentals against the N_3 Fermi fundamentals and the $2N_c$ from the adjoint gaugino, the same chiral-minus-Fermi-minus-adjoint count that gave $4 - 3 - 1 = 0$ for $\mathbb{C}^4$. This condition alone does not cancel the central $U(1)$ anomaly of a $U(N_c)$

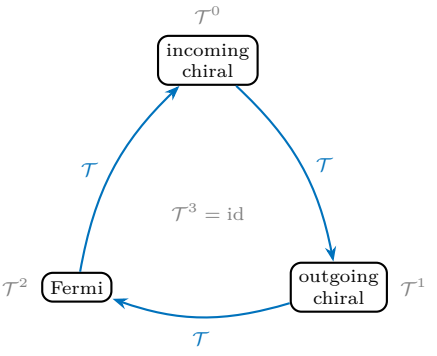


Figure VII.10.4: $(0,2)$ triality as an order-three cycle on the three field classes at a dualized node. The triality move $\mathcal{T}$ carries the incoming chirals to the outgoing chirals to the Fermis and back, a three-cycle. Reading one class around the triangle from $\mathcal{T}^0$ to $\mathcal{T}^1$ to $\mathcal{T}^2$ and back to $\mathcal{T}^0$ makes the order-three class cycle $\mathcal{T}^3 = \mathrm{id}$ visually plain: three schematic permutations return the original classes, where Seiberg duality at the dimer rung is an involution, $\mathcal{T}^2 = \mathrm{id}$. The rank changes by the rule (VII.10.18) at each step.

theory. The Gadde–Gukov–Putrov $U(N_c)$ construction includes two additional Fermi multiplets $\Omega_{1,2}$ in determinant representations for that abelian cancellation; neither their contribution nor a complete $U(N_c)$ anomaly is calculated here. Triality cyclically permutes the three flavour classes and changes the rank. In this node convention it sends the $U(N_c)$ node with flavours $(N_1, N_2, N_3) = (\text{fund chiral},\ \text{antifund chiral},\ \text{fund Fermi})$ to a $U(N_c')$ node with the cyclically permuted flavours (N_2, N_3, N_1) and the new rank

$$N_c' = N_2 - N_c. \tag{VII.10.18}$$

This value is not arbitrary; it is *fixed* by requiring the dual $SU(N_c')$ coefficient to vanish in the same convention when both ranks are at least two. The dual node, with flavours (N_2, N_3, N_1), has vanishing nonabelian coefficient when $N_2 + N_3 - N_1 - 2N_c' = 0$, and using the original node's anomaly condition $N_1 + N_2 - N_3 - 2N_c = 0$, that is $N_3 = N_1 + N_2 - 2N_c$, to eliminate N_3 gives

$$\begin{aligned} N_2 + N_3 - N_1 - 2N_c' &= N_2 + (N_1 + N_2 - 2N_c) - N_1 - 2N_c' \\ &= 2N_2 - 2N_c - 2N_c' = 0 \quad \Longrightarrow \quad N_c' = N_2 - N_c, \end{aligned} \tag{VII.10.19}$$

the rank rule (VII.10.18). The vanishing nonabelian coefficient is preserved by the cyclic reshuffle when its gauge factors are nontrivial. For the integer recurrence, retain the formal tuple $(N_1, N_2, N_3, N_c) = (3, 3, 4, 1)$, which satisfies $3+3-4-2 = 0$. Here $SU(1)$ is trivial: the polynomial identity at rank one is only formal. The six charged chirals and four charged Fermis in the displayed field list leave a $U(1)$ coefficient $6 - 4 = 2$ before the omitted Ω sector. One formal step permutes the flavours to $(N_2, N_3, N_1) = (3, 4, 3)$ and sets the rank to $N_2 - N_c = 3 - 1 = 2$:

$$(3,3,4,1) \xrightarrow{\mathcal{T}} (3,4,3,2) \xrightarrow{\mathcal{T}} (4,3,3,2) \xrightarrow{\mathcal{T}} (3,3,4,1), \tag{VII.10.20}$$

the polynomial vanishing at each state ($3{+}4{-}3{-}4=0$, $4{+}3{-}3{-}4=0$, $3{+}3{-}4{-}2=0$) and the third returning the original integers. This is an order-three recurrence, not a physical triality of an anomaly-free $U(1)$ theory. The unbalanced tuple $(2,2,4,1)$ does not close under this recurrence; the dual rank above is fixed by the stated nonabelian balance equations.

What triality preserves. A triality move on a full brane brick model is expected to preserve the toric fourfold diagram, the height-one lattice polytope of §VII.10.2, as in the cited construction. Comparing those polytopes requires both periodic brick models. The local flavour cycle and rank orbit above do not supply them. The geometric comparison parallels Chapter 7's statement that the two toric phases of F_0 share their Newton polygon: the field content changes, the toric moduli space does not. Table VII.10.1 places triality beside Seiberg duality in the graded-quiver sequence.

rung m	theory	n-ality move	order $m{+}1$	invariant preserved
1	4d $\mathcal{N}{=}1$	Seiberg / toric duality	2 (involution)	toric CY_3 diagram
2	2d $(0,2)$	$(0,2)$ triality	3	toric CY_4 diagram
3	0d	quadrality	4	toric CY_5 diagram

Table VII.10.1: The n-ality move of each rung of the sequence (VII.10.2), of order $m{+}1$. Seiberg duality (order two) at the dimer rung, $(0,2)$ triality (order three) at the brick rung, and quadrality (order four) at the matrix-model rung are one uniform family, not unrelated phenomena. The volume constructs the $m{=}1$ move (Seiberg / toric duality) in Chapter 7 and calculates the local $m{=}2$ class/rank cycle here; the full cube move and toric comparison are cited, while the $m{=}3$ row is an outlook pointer (§VII.10.7). Toric-diagram preservation for full moves is the cited construction-level result.

The local calculation establishes the formal rank orbit satisfying the stated polynomial and its order-three closure. The cube move on a complete periodic brick model and the equality of the associated toric diagrams are results of the cited construction. Infrared triality would require a further comparison, for example of elliptic genera. That quantum question belongs to Part III.

VII.10.6 The fixtures: $\mathbb{C}^4$ and the $\mathbb{C}^4/\mathbb{Z}_4$ orbifold

Two worked examples anchor the construction, the $(0,2)$ analogues of Chapter 3's $\mathbb{C}^3$ honeycomb and conifold square: the flat $\mathbb{C}^4$ model and the orbifold $\mathbb{C}^4/\mathbb{Z}_4$, the first multi-node fourfold example here.

$\mathbb{C}^4$, end to end. The flat fourfold has a single brick in a fundamental domain of the T^3 tessellation. Collecting the pieces of the previous sections: one gauge node $U(N)$ (one brick); four chiral multiplets $X_1,\dots,X_4$ (four oriented faces, the four $\mathbb{C}^4$ directions) and three Fermi multiplets $\Lambda_1,\Lambda_2,\Lambda_3$ (three unoriented faces); the (E_a,J_a) data of (VII.10.12), with $E_a=[X_4,X_a]$ and $J_a=[X_{a+1},X_{a+2}]$; the consistency condition $\sum_a \mathrm{Tr}(E_aJ_a)=0$ satisfied by the Jacobi identity (VII.10.14); and the toric fourfold diagram the unit tetrahedron (VII.10.6). This is the flat-space example, and every

piece of the brick dictionary is exercised on it: the bricks give the gauge group, the two face species give the two matter species, the edges give the (E, J) data, and the closure of the tessellation gives the consistency condition, here as the most elementary identity of matrix algebra. The $\mathbb{C}^4$ model plays for the fourfold story the role $\mathcal{N}{=}4$ super Yang–Mills and the $\mathbb{C}^3$ honeycomb played for the threefold: the parent from which the orbifold example is projected.

$\mathbb{C}^4/\mathbb{Z}_4$, worked. The first multi-node example here is the orbifold $\mathbb{C}^4/\mathbb{Z}_4$, an abelian toric fourfold, the brick-model analogue of $dP_0 = \mathbb{C}^3/\mathbb{Z}_3$ as the bridge from orbifold to general toric. Its brick model is built the same way the Douglas–Moore orbifold quiver of Volume VI was built from $\mathcal{N}{=}4$ super Yang–Mills: the brane brick model of an abelian orbifold *is* the orbifold projection of the parent. Take the $\mathbb{C}^4$ model and project by the $\mathbb{Z}_4$ action that assigns the weight 1 to each of the four chiral coordinates,

$$\mathbb{Z}_4 : \ (X_1, X_2, X_3, X_4) \ \mapsto \ (\zeta X_1, \zeta X_2, \zeta X_3, \zeta X_4), \qquad \zeta = e^{2\pi i/4}, \tag{VII.10.21}$$

each chiral coordinate carrying the weight one, with $1+1+1+1 = 4 \equiv 0 \pmod 4$, the condition $\sum_k w_k \equiv 0 \pmod 4$ being the Calabi–Yau requirement, exactly what keeps the action inside $SU(4)$ and makes the Fermi charges consistent below. The projection acts on the regular Chan–Paton representation of N D1-branes on the quotient. It begins with $4N$ covering-space Chan–Paton indices and retains a product over the four $\mathbb{Z}_4$ representations:

$$U(4N) \ \longrightarrow \ \prod_{j=0}^{3} U(N_j), \qquad N_j = N, \tag{VII.10.22}$$

four bricks, $j = 0, 1, 2, 3$, one per irreducible representation. A chiral coordinate of weight one maps the brick j to $j+1$, and a Fermi multiplet carries the charge of its E-term $E_a = [X_4, X_a]$, namely $w_4 + w_a = 1 + 1 = 2$, mapping j to $j+2$. The projected field content is read straight off these charges:

$$\begin{aligned} &\text{chirals:} \quad X_k : \ j \to j+1, \quad k = 1, \dots, 4, \ j \in \mathbb{Z}_4 \quad \Rightarrow \quad 4 \times 4 = 16 \text{ chiral fields;} \\ &\text{Fermis:} \quad \Lambda_a : \ j \to j+2, \quad a = 1, 2, 3, \ j \in \mathbb{Z}_4 \quad \Rightarrow \quad 3 \times 4 = 12 \text{ Fermi fields,} \end{aligned} \tag{VII.10.23}$$

four chirals per node going $j \to j+1$ and three Fermis per node going $j \to j+2$, a four-node $(0,2)$ quiver with both species. Table VII.10.2 collects the node-by-node data and one representative E- and J-term.

The orbifold quiver is drawn in Figure VII.10.5: the four bricks as a cyclic square, the chiral arrows stepping $j \to j+1$ and the Fermi arrows crossing $j \to j+2$.

The (E, J) data descends from the parent's. Writing $X_k^{(j)}$ for the orbifold block of X_k on the arrow $j \to j+1$, the parent assignment $E_a = [X_4, X_a]$, $J_a = [X_{a+1}, X_{a+2}]$ projects onto the $j \to j+2$ Fermi sector as products of two consecutive chiral blocks. For the node-0 Fermi Λ_1 the representative pair is

$$E_1^{(0)} = X_4^{(0)} X_1^{(1)} - X_1^{(0)} X_4^{(1)}, \qquad J_1^{(0)} = X_2^{(2)} X_3^{(3)} - X_3^{(2)} X_2^{(3)}, \tag{VII.10.24}$$

Here matrices transform as $X^{(j)} \mapsto g_j X^{(j)} g_{j+1}^{-1}$, so the products follow their paths from left to right. Each expression is a difference of two oriented chiral paths as in (VII.10.10).

node j	chiral arrows X_k ($j \to j{+}1$)	Fermi arrows Λ_a ($j \to j{+}2$)
0	$0 \to 1$ (X_1, X_2, X_3, X_4)	$0 \to 2$ $(\Lambda_1, \Lambda_2, \Lambda_3)$
1	$1 \to 2$ (X_1, X_2, X_3, X_4)	$1 \to 3$ $(\Lambda_1, \Lambda_2, \Lambda_3)$
2	$2 \to 3$ (X_1, X_2, X_3, X_4)	$2 \to 0$ $(\Lambda_1, \Lambda_2, \Lambda_3)$
3	$3 \to 0$ (X_1, X_2, X_3, X_4)	$3 \to 1$ $(\Lambda_1, \Lambda_2, \Lambda_3)$

Table VII.10.2: The $\mathbb{C}^4/\mathbb{Z}_4$ brane brick model, the $(0,2)$ orbifold projection of the $\mathbb{C}^4$ parent by the diagonal action $(1,1,1,1)$. Four bricks $j = 0,1,2,3$ ($\prod_j U(N_j)$), sixteen chiral fields (four per node, $j \to j{+}1$), and twelve Fermi fields (three per node, $j \to j{+}2$, the E-/J-charge $1{+}1 = 2$). The node-by-node arrow data; a representative (E, J) pair for the node-0 Fermi is displayed in the body, (VII.10.24).

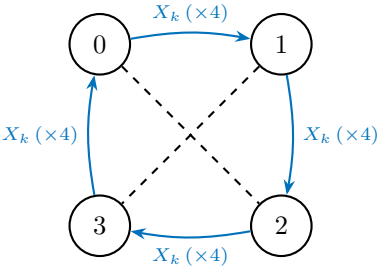


Figure VII.10.5: The $\mathbb{C}^4/\mathbb{Z}_4$ $(0,2)$ quiver, the orbifold projection of the single-node $\mathbb{C}^4$ quiver (Figure VII.10.2) by the diagonal action $(1,1,1,1)$. Four gauge nodes $0,1,2,3$ ($\prod_j U(N_j)$) on a cyclic square. On each edge a bundle of four directed chiral arrows $X_k : j \to j+1$ (blue, $4 \times 4 = 16$ chirals in all); each dashed diagonal carries two opposite triples of Fermi arrows $\Lambda_a : j \to j+2$ (one triple for each source node, hence $2 \times 6 = 12$ Fermis in all), the E-/J-charge $1+1 = 2$ stepping two nodes. In words: the orbifold projection turns the one node and seven self-loops of $\mathbb{C}^4$ into four nodes carrying sixteen chirals and twelve Fermis, exactly the counts of (VII.10.23).

The endpoints check the charges: $E_1^{(0)}$ is a difference of two $0 \to 1 \to 2$ paths (each chiral block steps one node, so the product steps two, $0 \to 2$, matching the Λ_1 charge), and $J_1^{(0)}$ is a difference of the complementary paths returning $2 \to 3 \to 0$, the $2 \to 0$ leg that closes the E-/J-pair into a gauge invariant. The consistency condition holds on these blocks,

$$\sum_a \mathrm{Tr}\big(E_a\, J_a\big) = 0 \qquad \text{on the } \mathbb{C}^4/\mathbb{Z}_4 \text{ blocks,} \tag{VII.10.25}$$

inherited directly from the $\mathbb{C}^4$ Jacobi identity (VII.10.14): the projected $X_k^{(j)}$ are specific block matrices, and the Jacobi cancellation, holding for any matrices, holds in particular for them. It is not vacuous, because the projected commutators are themselves non-vanishing on the orbifold blocks.

The per-node gauge anomaly. The orbifold projection preserves the $\mathbb{C}^4$ nonabelian

anomaly balance $4-3-1=0$ node by node at equal ranks $N \geq 2$, and the check is a clean one-line count. Take any node j with equal ranks N. The chiral matter charged at the node is four outgoing arrows ($j \to j+1$) and four incoming ($j-1 \to j$), so eight chiral legs; the Fermi matter is three outgoing ($j \to j+2$) and three incoming ($j-2 \to j$), so six Fermi legs; and the vector multiplet's adjoint gaugino contributes 2 in the same per-node convention, the $2N_c$ normalization of (VII.10.7). The anomaly coefficient at the node is therefore

$$8-6-2=0, \tag{VII.10.26}$$

chiral legs against Fermi legs against the adjoint vector, so each nonabelian $SU(N)$ factor has zero coefficient at equal ranks $N \geq 2$. The abelian factors require separate treatment. The eight-against-six-against-two of the orbifold is the four-against-three-against-one of the parent doubled by counting both arrow orientations at the node, the same $(4-3-1)$ balance carried through the projection to every $\mathbb{Z}_4$ image.

The toric diagram of $\mathbb{C}^4/\mathbb{Z}_4$ is the three-dimensional lattice polytope of the quotient geometry. The cited brane-brick construction also describes triality-related models. This chapter does not derive a dual periodic brick or compare its polytope. Figure VII.10.6 draws the two fourfold toric diagrams side by side: the flat $\mathbb{C}^4$ tetrahedron with no interior point, and the $\mathbb{C}^4/\mathbb{Z}_4$ tetrahedron with one interior lattice point, the fourfold analogue of the dP_0 triangle's single interior point.

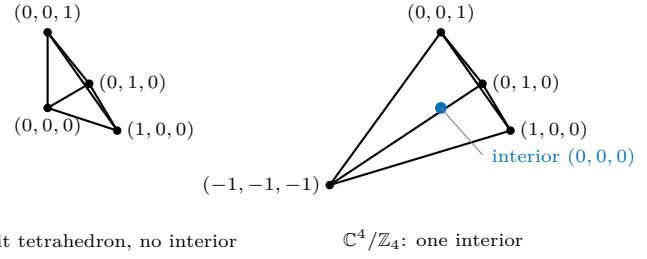


Figure VII.10.6: The toric diagrams of the two worked fourfolds, height-one lattice polytopes in three dimensions (drawn with the same oblique linear projection). Left: $\mathbb{C}^4$, the unit tetrahedron $\mathrm{conv}\{(0,0,0),(1,0,0),(0,1,0),(0,0,1)\}$ of (VII.10.6), four vertices and no interior point, the fourfold analogue of the $\mathbb{C}^3$ triangle. Right: $\mathbb{C}^4/\mathbb{Z}_4$, whose crepant resolution is local $\mathbb{P}^3$, has the tetrahedron $\mathrm{conv}\{(1,0,0),(0,1,0),(0,0,1),(-1,-1,-1)\}$ with the single interior lattice point $(0,0,0)$ (blue), the fourfold analogue of the dP_0 triangle's interior point. The normalized volume of the right tetrahedron is 4, matching the four gauge nodes of the orbifold projection (VII.10.22), just as dP_0's normalized area 3 matched its three nodes.

Together the two examples show the scope of the construction: the flat example shows the dictionary reproduces the known content and the trivial diagram, and the orbifold example shows it produces a multi-node $(0,2)$ quiver with a consistent (E,J) superpotential and a nontrivial toric diagram; the full brick data of the worked fourfolds and the explicit triality duals are in the cited brane-brick literature.

VII.10.7 The sequence, and the outlook

The dimer and brick models share a dimensional pattern. For $m = 1$ and $m = 2$, respectively, a D$(5 - 2m)$-brane probes a toric Calabi–Yau of complex dimension $n = m+2$, producing a theory in $d = 6-2m$ worldvolume dimensions organized on T^{m+1}. In the graded-quiver language, grade m assigns an arrow a degree $c \in \{0, 1, \ldots, m\}$. Conjugation pairs degrees by

$$\text{arrow degrees } \{0, 1, \ldots, m\}, \qquad \text{conjugation } \; c \;\longleftrightarrow\; m - c. \tag{VII.10.27}$$

At m=1 the degrees are $\{0, 1\}$ and the conjugation swaps them: the ordinary quiver has chiral arrows (degree 0) and their conjugates (degree 1), the 4d $\mathcal{N}$=1 dimer of Chapters 1–9. At m=2 the degrees are $\{0, 1, 2\}$ and the conjugation fixes the middle degree 1: the brick model has chiral arrows (degree 0, with conjugate degree 2) and Fermi arrows (degree 1, self-conjugate), the 2d $(0, 2)$ theory of this chapter. The self-conjugate middle degree at m=2 accounts for the unoriented Fermi face in the graded description. It does not remove the Fermi field's gauge representation or the oriented E/J paths. The conjugation $c \leftrightarrow m - c$ is an involution on the degree set for every m, with a self-conjugate middle precisely when m is even.

The associated dimensions are the relations of (VII.10.2):

$$p = 5 - 2m \;(3, 1, -1), \qquad d = 6 - 2m \;(4, 2, 0), \qquad n = m + 2 \;(3, 4, 5), \qquad m = 1, 2, 3, \tag{VII.10.28}$$

Here p is the spatial brane dimension, $d = p + 1$ its worldvolume dimension, and n the complex dimension of the transverse Calabi–Yau. Eliminating $m = n - 2$ gives $d = 10 - 2n$. The $m = 3$ entry reaches $d = 0$ and is an outlook to a matrix model. A formal $m = 4$ entry would give $d = -2$, outside this family of worldvolume theories.

One sequence, one family of dualities

The dimer and the brick model are the m=1 and m=2 rungs of a single sequence: a D$(5 - 2m)$-brane at a toric CY$_{m+2}$, a grade-m graded quiver, a tessellation of T^{m+1}, in $d = 6-2m$ dimensions. Each rung has its own n-ality move, of order m+1: Seiberg duality (order 2), $(0, 2)$ triality (order 3), quadrality (order 4). These are not three unrelated dualities but one uniform family, the order-$(m$+1) move of the grade-m graded quiver. *The sequence does not stop at two dimensions, and quadrality is the same kind of operation as Seiberg duality, one rung up.*

The graded-quiver framework groups Seiberg duality, $(0, 2)$ triality, and quadrality into an order-$(m + 1)$ family. Here we calculated the local triality rank cycle at $m = 2$; the quadrality transformation at $m = 3$ and the derived or A_∞ interpretation are supplied by the cited graded-quiver literature and developed in Volume V and Part III.

The outlook: pointers beyond the sequence. The synthesis closes by naming, as pointers with references and not as constructions, the topics one rung further that a quantum-field-theory construction volume does not build.

- The m=3 rung, the D(-1)-instanton at a toric Calabi–Yau fivefold, gives a zero-dimensional theory, a supersymmetric matrix model with the minimal zero-dimensional supersymmetry, and its n-ality move is **quadrality**, the order-four

duality of the grade-three graded quiver. A zero-dimensional matrix model is not a field theory, and this volume constructs neither it nor the quadrality action on its field content; the rung is named because it completes the order-$(m{+}1)$ pattern, and its construction and dynamics are routed.

- The brane brick models admit **orientifolds** with Spin(7) holonomy, which project the 2d $(0,2)$ theory to a 2d $(0,1)$ theory, the special-holonomy branch off the brick rung. This is an exotic corner for a construction volume, and the chapter names it without building the Spin(7) projection or the $(0,1)$ spectrum.

These pointers, and the adjacent topics of BPS counting, D0-brane quiver quantum mechanics, and crystal melting, have their homes in Volume V and Part III; the chapter names them and routes their dynamics. The pedagogical value of the sequence, the uniform n-ality family and the one-spine reading of the volume's whole machinery, is kept; the frontier constructions are pointed to, not built.

VII.10.8 What has been computed, and what is routed

The chapter has done the following. A D1-brane at a toric Calabi–Yau fourfold preserves 2d $(0,2)$ supersymmetry (VII.10.1) and is the $m{=}2$ rung of the sequence (VII.10.2). Its theory is read off a brane brick model, a tessellation of T^3 whose bricks are gauge nodes, whose faces are matter fields with the orientation fixing the species (oriented chiral, unoriented Fermi), and whose edges are the J- and E-terms (VII.10.3)–(VII.10.5). The matter has two species, chiral and Fermi, the single new structural fact; the superpotential is the pair (E_a, J_a) per Fermi multiplet, constrained by $\sum_a \mathrm{Tr}(E_a J_a) = 0$ (VII.10.11), which for $\mathbb{C}^4$ is the Jacobi identity (VII.10.14). The toric fourfold diagram is the three-dimensional lattice polytope read off the brick model, the unit tetrahedron for $\mathbb{C}^4$ (VII.10.6). For triality, the local order-three cycle is computed here, while equality of the toric diagrams of full phases follows the cited construction (VII.10.15). The flat $\mathbb{C}^4$ and the $\mathbb{C}^4/\mathbb{Z}_4$ orbifold are worked, and the whole construction is placed on the m-graded-quiver sequence, with the n-ality moves a uniform order-$(m{+}1)$ family and the rung beyond named as an outlook (VII.10.27)–(VII.10.28).

A reader leaving the chapter should be able to do the following without looking ahead: count the supercharges of a D1-brane at a fourfold and name the theory 2d $(0,2)$; read a brane brick model's dictionary, assigning bricks to gauge nodes, oriented and unoriented faces to chiral and Fermi multiplets, and edges to J/E-terms; check the 2d $(0,2)$ gauge anomaly $\sum_{\text{chiral}} T(R) - \sum_{\text{Fermi}} T(R) - T(\text{adj}) = 0$ at a node, with the vector's left-moving gaugino load-bearing ($4{-}3{-}1 = 0$ for $\mathbb{C}^4$); write the (E_a, J_a) pair of a Fermi multiplet and check the trace condition $\sum_a \mathrm{Tr}(E_a J_a) = 0$, proving it for $\mathbb{C}^4$ by the Jacobi identity; build the $\mathbb{C}^4/\mathbb{Z}_4$ orbifold brick model by the $(0,2)$ orbifold projection and read off its four nodes, sixteen chirals, and twelve Fermis; describe $(0,2)$ triality as an order-three cube move, write the GGP node anomaly $N_1 + N_2 - N_3 - 2N_c = 0$, and name the toric fourfold diagram as its invariant; and place the dimer and the brick model on the sequence, reading off the relations $p = 5 - 2m$, $d = 6 - 2m$, $n = m + 2$ and the n-ality order $m{+}1$.

What has been routed, and to where, is the chapter's honest gap, the boundary of the local calculation and the exact analogue of Chapter 8's deferral of the Klebanov–Strassler dynamics.

- The two-dimensional $(0,2)$ *quantum dynamics*, the elliptic genus (the flavoured $(0,2)$ Witten index), c-extremization (the determination of the infrared R-symmetry and central charge by extremizing a trial c-function, the two-dimensional analogue of a-maximization), and the $(0,2)$ renormalization-group flow to its infrared conformal field theory with its AdS_3 dual, are Part III, with the two-dimensional $(0,2)$ foundations in Volume 0. This is the chapter's primary deferral. Note the contrast with Chapter 6: there the toric a-maximization was a construction-intrinsic calculation the dimer performed; here c-extremization stays fully out, because it determines the infrared R-symmetry and is dynamics, not a brick-model read-off. The chapter names c-extremization as the two-dimensional analogue of a-maximization and computes nothing.

- The deeper geometry, the T^3 tessellation and periodic-quiver mathematics, the fourfold's toric and GLSM data, and the Sasaki–Einstein seven-geometry of the link, are Volume IV; the chapter uses the brick combinatorics and the toric polytope and cites these foundations.

- The local triality rank cycle is calculated. The full cube move and its toric-diagram invariant are taken from the cited construction, and triality is not verified at the level of the infrared conformal field theory (the elliptic-genus matching across triality phases), which is the routed dynamics.

- The full m-graded-quiver homological-algebra theory, the graded quiver as a derived or A_∞ structure and quadrality as a derived equivalence, is Volume V and Part III; the chapter states the sequence and uses it as language.

- The outlook constructions, the m=3 zero-dimensional matrix model with quadrality and the Spin(7) orientifolds giving 2d $(0,1)$ theories, are named as pointers with references and constructed nowhere here; their construction and dynamics are routed.

- The probe interpretation carries the α' and g_s regime conditions discussed in Volume V. We do not compute quantum or string corrections here. The matrix Jacobi identity, as an algebraic statement, holds for any N independently of that physical regime.

The single lesson to carry forward is the one in the closing keybox: the dimer and the brick model are two rungs of one sequence, the matter species and the dualities of each rung fixed by its grade, and the whole machinery of this volume, the periodic tessellation, the toric data, and the n-ality move, is read off that sequence. The next dimension up is named and pointed to; the construction this volume makes precise is the D-brane probe of a toric singularity, told here at its second rung.

Status and sources

The evidence-code tags ([WS] controlled perturbative derivation in the untwisted physical string worldsheet, [MATH] mathematical theorem with hypotheses and proof status, [ANOM] anomaly-coefficient evidence with stated power), the convention pointers, and the source anchors are collected here, keyed by section. Worldvolume and superspace algebra do not by themselves establish [WS]. The cited physical dictionaries below are premises; no local untwisted-worldsheet derivation is supplied. Their classifications remain unresolved, without implying that the physical statements are false. Mathematical classifications cover only the stated algebraic consequences, not their imported physical interpretations. The chapter develops the D1-brane at a toric Calabi–Yau fourfold and its brane brick model, after Franco, Lee, Seong, and Vafa (2016) (the founding brane-brick paper: the T^3 tessellation, bricks = gauge, faces = chiral/Fermi, edges = J/E, the toric fourfold diagram) and Franco, Ghim, et al. (2015) (the 2d $(0,2)$ quiver gauge theories from D-branes, the $\mathbb{C}^4/\mathbb{Z}_4$ worked model), Gadde et al. (2014) (the original $(0,2)$ triality and the J/E structure of 2d $(0,2)$ gauge theory), and Franco, Lee, and Seong (2016) (triality as a cube move); the graded-quiver sequence follows Franco and Hasan (2019) and Closset et al. (2019), and the outlook points to Franco, Lee, and Seong (2017) (quadrality) and Franco, Mininno, et al. (2022) (Spin(7) orientifolds). It recalls the dimer corpus of Hanany and Kennaway (2005) and Franco, Hanany, et al. (2006) as the m=1 rung. It reuses the brane-tiling orientation rule (bipartite tiling conventions, the BFT_2 chiral/Fermi face split lifted to T^3), the superpotential store extended to the $(0,2)$ (E,J) data (superpotential normalization, the J/E 2d extension and $\sum_a \mathrm{Tr}(E_a J_a) = 0$), the gauge/rank data (quiver gauge groups and ranks, $\prod_i U(N_i)$), the arrow convention, the supercharge count (spinor and supersymmetry counting, the 2d $(0,2)$ row), and the toric-diagram lift (toric lattice orientation, the height-one fourfold polytope). The new brane-brick / 2d $(0,2)$ machinery is reserved as an extension of bipartite tiling conventions and superpotential normalization for this brick construction. The finite calculations use the displayed brick models and graded-quiver data.

§VII.10.1 From D3-branes to D1-branes. Framing: the dimer-to-brick lift, the chiral/Fermi split as the single new structural fact, the build-the-theory-then-route-the-dynamics promise, and the reader contract. Files (a brane brick model is not a 3d dimer or a graph in 3d: it is a T^3 tessellation with two matter species, refines the stated distinction).

§VII.10.2 The D1-brane at a toric fourfold and its brane brick model. The supercharge count, the CY_4 holonomy $SU(4)$ keeping two of sixteen supercharges (with the one-sentence representation-theory bridge: a generic $SU(4)$ holonomy leaves one complex invariant spinor, two real $(0,2)$ supercharges after the D1 projection), the 2d $(0,2)$ theory (VII.10.1), and the sequence placement (VII.10.2) at the m=2 rung, $\mathrm{D}(5-4) = \mathrm{D1}$, $\mathrm{CY}_{2+2} = \mathrm{CY}_4$, $d = 6-4 = 2$. Then the T^3 tessellation, the by-codimension dictionary (VII.10.3) (bricks = gauge nodes, faces = matter, edges = J/E-terms), the gauge group $\prod_i U(N_i)$ (VII.10.4), the chiral/Fermi face-orientation split (VII.10.5) (the Fermi faces the unoriented square faces), the brick matching as the fourfold analogue of the perfect matching (the GLSM homogeneous coordinate mapping to the toric polytope points), and the toric fourfold diagram as a height-one 3d lattice polytope, the unit tetrahedron for $\mathbb{C}^4$ (VII.10.6). use cited physical premises; the toric-diagram identifications remain sourced mathematical assertions without a local proof or precise theorem anchor; their WS classification is unresolved (spinor and supersymmetry counting; bipartite tiling conventions, the BFT_2 orientation rule; quiver gauge groups and ranks; toric lattice orientation; Franco, Lee, Seong, and Vafa (2016), Franco and Hasan (2019)). Files the stated distinction. The 2d $(0,2)$ SUSY foundations and the dimensional-sequence supercharge table route to Volume I (the 2d row of the sequence); the T^3 tessellation and CY_4 toric/GLSM foundations are toric fourfold and brane brick geometry (home Vol. IV); both cited, not re-derived. The integer relations use the displayed data, and the supersymmetry statement uses the brane construction: the m=2 rung is D1 @ $\mathrm{CY}_4 \to$ 2d $(0,2)$ on T^3 with $p = 5-2m$, $d = 6-2m$, $n = m+2$; and chirals and Fermis are distinct species (every chiral an (i,j) face, every Fermi an (E,J) pair), four chiral and three Fermi on the $\mathbb{C}^4$ fixture, the toric diagram the unit tetrahedron with vertex $(0,0,0)$.

§VII.10.3 The $(0,2)$ quiver: chiral and Fermi multiplets. The 2d $(0,2)$ multiplet content (chiral = scalar + right-mover; Fermi = left-mover, no scalar; vector = gauge sector with a *left*-moving adjoint gaugino), the two arrow species (oriented chiral, undirected Fermi), the $\mathbb{C}^4$ quiver (one node, four chiral, three Fermi, all adjoint), and the 2d $(0,2)$ gauge-anomaly condition $\sum_{\text{chiral}} T(R) - \sum_{\text{Fermi}} T(R) - T(\text{adj}) = 0$

(VII.10.7) with the $\mathbb{C}^4$ check $4-3-1=0$ (VII.10.8), the vector contribution load-bearing, and the normalization bridge to the GGP flavour-count form (the adjoint gaugino as $2N_c$ when matter is counted by flavour multiplicity). uses sourced multiplet and chirality premises; WS is unresolved. The anomaly formula and its single nonabelian coefficient carry [ANOM] (fundamental-at-tail arrows, quiver gauge groups and ranks; Gadde et al. (2014), Franco, Lee, Seong, and Vafa (2016)). The 2d $(0,2)$ multiplet vocabulary is defined at first use and the systematic 2d $(0,2)$ SUSY theory routed to Volume I / two-dimensional supersymmetric dynamics; the detailed 2d anomaly inflow is foundations. the $\mathbb{C}^4$ model has both species, four chiral and three Fermi, and has vanishing nonabelian $SU(N\geq 2)$ gauge-anomaly coefficient by $4-3-1=0$ (the vector's left-moving gaugino load-bearing: $4-3\neq 0$ without it).

§VII.10.4 The J-terms, the E-terms, and a Jacobi identity. The (E_a, J_a) pair per Fermi multiplet (VII.10.9), the two-term structure (VII.10.10) with the local E-/J-patch (Figure VII.10.3, the Fermi face with its four chiral paths, $J=J^+-J^-$ and $E=E^+-E^-$), the consistency condition $\sum_a \mathrm{Tr}(E_a J_a)=0$ (VII.10.11), the $\mathbb{C}^4$ assignment $E_a=[X_4,X_a]$, $J_a=[X_{a+1},X_{a+2}]$ (VII.10.12), and its vanishing by the Jacobi identity (VII.10.13)–(VII.10.14). has unresolved WS status for the superspace and brick dictionary. The trace expansion and Jacobi identity carry [MATH] for associative square matrices with a cyclic trace and the displayed pairing; the chapter gives the algebraic proof (superpotential normalization, the J/E 2d extension; Franco, Lee, Seong, and Vafa (2016), Gadde et al. (2014)). Files (a $(0,2)$ theory has two holomorphic data, not one W; triality permutes the J/E data). The cyclic-word calculation gives the trace identity. The $\mathbb{C}^4$ sum $\sum_a \mathrm{Tr}([X_4,X_a][X_{a+1},X_{a+2}])=0$ identically on symbolic matrices (every entry a free symbol, the Jacobi identity), and it is anti-vacuous (the commutators are not identically zero, and the wrong, non-cyclic pairing $J_a=[X_a,X_{a+1}]$ does not vanish).

§VII.10.5 $(0,2)$ triality as a cube move. Triality as the cube move (VII.10.15), order three (vs Seiberg's order two), the Gadde–Gukov–Putrov field-content three-cycle (VII.10.16) (Figure VII.10.4), the GGP local model with the nonabelian $SU(N_c\geq 2)$ coefficient condition $N_1+N_2-N_3-2N_c=0$ (VII.10.17) (the chiral-minus-Fermi-minus-adjoint accounting, the same count as $\mathbb{C}^4$'s $4-3-1=0$; the GGP $U(N_c)$ theory also requires its cited determinant-Fermi abelian sector), the explicit rank rule $N_c'=N_2-N_c$ (VII.10.18) derived from the dual-node anomaly (VII.10.19) and the worked order-three orbit $(3,3,4,1)\to(3,4,3,2)\to(4,3,3,2)\to(3,3,4,1)$ (VII.10.20) as a formal integer recurrence, the polynomial relation preserved by the cyclic reshuffle, the source-supported rewriting of (E,J) and toric-diagram equality (Table VII.10.1, the n-ality sequence). imports the full physical triality move, whose WS classification is unresolved. Its rank elimination and finite integer orbit carry [MATH] conditional on the stated permutation and balance polynomial; the nonabelian coefficient carries [ANOM] with one independent coefficient, while the rank-one substitutions remain formal (Gadde et al. (2014), Franco, Lee, and Seong (2016); the m=1 analogue is Chapter 7). Files the stated distinction. The IR-CFT verification of triality (the elliptic-genus matching) is the 2d quantum dynamics, routed to Part III / two-dimensional supersymmetric dynamics. the encoded local field-class permutation has order three. The toric CY_4 diagram is named as a source-supported invariant. The GGP nonabelian condition $N_1+N_2-N_3-2N_c=0$ holds for balanced ranks and fails for unbalanced ones; and the rank rule $N_c'=N_2-N_c$ preserves the node polynomial and closes the formal orbit at order three for $(3,3,4,1)$, failing to close on an unbalanced tuple (the finite calculation; check, the move is not order two = Seiberg).

§VII.10.6 The fixtures: $\mathbb{C}^4$ and the $\mathbb{C}^4/\mathbb{Z}_4$ orbifold. The $\mathbb{C}^4$ model end to end (one node, four chiral and three Fermi, the (E,J) data with $\sum_a \mathrm{Tr}(E_aJ_a)=0$, the unit-tetrahedron toric diagram) and the worked $\mathbb{C}^4/\mathbb{Z}_4$ orbifold, built as the $(0,2)$ orbifold projection of the $\mathbb{C}^4$ parent by the diagonal action $(1,1,1,1)$ (VII.10.21): four bricks $\prod_j U(N_j)$ (VII.10.22), sixteen chirals ($j\to j+1$) and twelve Fermis ($j\to j+2$) (VII.10.23), Table VII.10.2 with the quiver Figure VII.10.5, the displayed representative $(E_1^{(0)},J_1^{(0)})$ (VII.10.24), the inherited $\sum_a \mathrm{Tr}(E_aJ_a)=0$ (VII.10.25), the per-node gauge anomaly $8-6-2=0$ (VII.10.26), the 3d toric polytope (Figure VII.10.6, local $\mathbb{P}^3$ with one interior point, normalized volume 4). A specific triality dual is not worked in this chapter. uses the cited orbifold dictionary as a premise with unresolved WS status. The determinant congruence, displayed paths, and inherited Jacobi consequence carry [MATH] under the stated block convention. The equal-rank node coefficient carries [ANOM]: four gauge-factor coefficients are evaluated, but cyclic symmetry reduces them to one independent coefficient expression. No abelian coefficient is matched (Franco, Lee, Seong,

and Vafa (2016), Franco, Ghim, et al. (2015), Franco, Lee, and Seong (2016)). Carries (the probe variety carries the α'/g_s caveats of Volume VI). The full brick data of the worked fourfolds and their explicit triality duals are in the founding literature, cited; the orbifold-projection construction is worked here. The matter lists and matrix identities give the $\mathbb{C}^4$ fixture end to end (one node, four chiral, three Fermi, $\sum_a \mathrm{Tr}(E_a J_a) = 0$, the toric diagram the unit tetrahedron) and the $\mathbb{C}^4/\mathbb{Z}_4$ projection (four nodes, sixteen chiral, twelve Fermi, the CY_4 weight condition $\sum w_k \equiv 0 \bmod 4$, $\sum_a \mathrm{Tr}(E_a J_a) = 0$ on the orbifold blocks, checked anti-vacuously, and the per-node anomaly $8 - 6 - 2 = 0$ at every node, the parent $4 - 3 - 1 = 0$ balance preserved through the projection) (the finite calculation; check, a non-CY_4 weight is rejected).

§VII.10.7 The sequence, and the outlook. The m-graded-quiver sequence (VII.10.2), the graded-quiver degree structure $\{0, \ldots, m\}$ with the conjugation $c \leftrightarrow m - c$ (VII.10.27) (the self-conjugate middle at m=2 = the Fermi self-conjugacy), the dimension relations (VII.10.28), the order-(m+1) n-ality family, and the outlook (the m=3 quadrality rung and the Spin(7) $\to$ 2d $(0,1)$ orientifolds as pointers). imports the physical sequence and outlook with unresolved WS status. The displayed affine integer relations and finite conjugation map carry [MATH] under the stated degree range; they do not prove the physical duality family (Franco and Hasan (2019), Closset et al. (2019); the outlook to Franco, Lee, and Seong (2017) and Franco, Mininno, et al. (2022)). Files (the sequence does not stop at 2d; the n-ality moves are one order-(m+1) family, not unrelated dualities). The full graded-quiver homological-algebra theory is graded-quiver homological algebra (Vol. V / Part III); the 0d matrix-model and Spin(7)/$(0,1)$ dynamics are quadrality and Spin(7) constructions (Part III). the relations $p = 5 - 2m$ (3,1,-1), $d = 6 - 2m$ (4,2,0), $n = m + 2$ (3,4,5) and the n-ality order m+1 (2,3,4) hold across $m = 1, 2, 3$; the conjugation $c \leftrightarrow m - c$ is an involution with a self-conjugate middle iff m even; the sequence stops at m=3 ($d = -2 < 0$ at m=4).

§VII.10.8 What is and is not given. This boundary is a scope declaration, not a WS derivation. Scope and limitations. the chapter builds the D1 @ CY_4 setup, the brane brick model on T^3, the 2d $(0,2)$ quiver with two matter species, the (E, J) superpotential with $\sum_a \mathrm{Tr}(E_a J_a) = 0$ (the $\mathbb{C}^4$ Jacobi identity), the local triality rank cycle, the two fixtures, and the sequence synthesis with its outlook. The full cube move and toric-diagram comparison are imported from the cited construction. The chapter routes the 2d $(0,2)$ quantum dynamics (the elliptic genus, c-extremization, the IR CFT) to Part III and Volume I (two-dimensional supersymmetric dynamics), the T^3/CY_4 deeper geometry to Vol. IV (toric fourfold and brane brick geometry), the graded-quiver homological algebra to Vol. V / Part III (graded-quiver homological algebra), and the 0d/quadrality and Spin(7)/$(0,1)$ constructions to the outlook pointers and Part III (quadrality and Spin(7) constructions). This is the boundary of the local calculation, the Chapter 8 analogue: VII.10 names the $(0,2)$ theory and its triality, and the 2d quantum dynamics is dynamics.

Related constructions. The brane brick model encodes the D1 probe of a toric Calabi–Yau fourfold as a T^3 tessellation. The two-dimensional $(0,2)$ quiver has chiral and Fermi multiplets and obeys $\sum_a \mathrm{Tr}(E_a J_a) = 0$. The Gadde–Gukov–Putrov triality is the order-three move at $m = 2$, with a matching toric diagram across triality phases. The D($5 - 2m$) probe sequence on CY_{m+2} connects the $m = 1$ dimer of Chapter 3 and its order-two Seiberg move in Chapter 7 to the $m = 2$ brick. The selected examples are $\mathbb{C}^4$ and $\mathbb{C}^4/\mathbb{Z}_4$. An explicit dual brick for the orbifold, zero-dimensional quadrality on Calabi–Yau fivefolds and the Spin(7) orientifold construction lie beyond these examples.

References

Closset, C. et al. (2019). "Graded quivers and B-branes at Calabi-Yau singularities". In: *JHEP* 03, p. 053. DOI: 10.1007/JHEP03(2019)053. arXiv: 1811.07016 [hep-th].

Franco, S., D. Ghim, et al. (2015). "2d (0,2) Quiver Gauge Theories and D-Branes". In: *JHEP* 09, p. 072. DOI: 10.1007/JHEP09(2015)072. arXiv: 1506.03818 [hep-th].

Franco, S., A. Hanany, et al. (2006). "Brane dimers and quiver gauge theories". In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Franco, S. and A. Hasan (2019). "Graded quivers, generalized dimer models and toric geometry". In: *JHEP* 11, p. 104. DOI: 10.1007/JHEP11(2019)104. arXiv: 1904.07954 [hep-th].

Franco, S., S. Lee, and R.-K. Seong (2016). "Brane brick models and 2d (0, 2) triality". In: *JHEP* 05, p. 020. DOI: 10.1007/JHEP05(2016)020. arXiv: 1602.01834 [hep-th].

Franco, S., S. Lee, and R.-K. Seong (2017). "Quadrality for Supersymmetric Matrix Models". In: *JHEP* 07, p. 053. DOI: 10.1007/JHEP07(2017)053. arXiv: 1612.06859 [hep-th].

Franco, S., S. Lee, R.-K. Seong, and C. Vafa (2016). "Brane Brick Models, Toric Calabi-Yau 4-Folds and 2d (0,2) Quivers". In: *JHEP* 02, p. 047. DOI: 10.1007/JHEP02(2016)047. arXiv: 1510.01744 [hep-th].

Franco, S., A. Mininno, et al. (2022). "2d $\mathcal{N} = (0,1)$ gauge theories and Spin(7) orientifolds". In: *JHEP* 03, p. 150. DOI: 10.1007/JHEP03(2022)150. arXiv: 2110.03696 [hep-th].

Franco, S. and X. Yu (2021). "BFT_2: a General Class of 2d $\mathcal{N} = (0,2)$ Theories, 3-Manifolds and Toric Geometry". In: arXiv: 2107.00667 [hep-th].

Gadde, A., S. Gukov, and P. Putrov (2014). "(0, 2) trialities". In: *JHEP* 03, p. 076. DOI: 10.1007/JHEP03(2014)076. arXiv: 1310.0818 [hep-th].

Hanany, A. and K. D. Kennaway (2005). "Dimer models and toric diagrams". In: arXiv: hep-th/0503149.

Chapter VII.11

Atlas: del Pezzo, $Y^{p,q}$, $L^{a,b,c}$, and brane-brick catalog

VII.11.1 What this atlas is, and how to read it

A toric cone can support several quiver descriptions of the same D3-brane probe theory. Comparing them requires more than the polygon: one must specify the brane construction, a consistent tiling, its ordered superpotential, and the regime in which the gauge theory flows to an interacting fixed point. Chapters 1–10 have worked through these ingredients. The outward normals give the (p, q)-web; the tilings describe matter and interactions; the mesonic quotient returns the cone. The central charge, fractional ranks, and duality moves then distinguish questions about a conformal theory from questions about its deformations.

This atlas brings those examples together. Its three axes compare the toric del Pezzo cones, the Sasaki–Einstein families, and the D1-brane theories on fourfolds. A short orientifold table adds the folded examples. The conifold, dP_0, and F_0 will carry the calculation from geometry to charges and ranks before the larger tables appear. Central-charge entries mean the coefficient of N^2 at large N; finite-N corrections will be displayed where they affect a comparison.

Different columns answer different questions. Areas, prescribed anomaly coefficients, and selected volumes can be calculated directly. Identifying a quiver as the theory of a physical brane requires the construction in the earlier chapter. Phase counts are literature catalog entries, even when a square move exhibits two of those phases. Collecting them in a table does not change their basis. In particular, a successful polynomial calculation does not establish a string worldsheet derivation or the existence of an infrared fixed point.

The three axes. The catalog is organized along the volume's natural divisions.

- **Axis 1**, the toric del Pezzo core (Table VII.11.2): the finite geometric list. The smooth toric del Pezzo surfaces, whose complex cones are the toric Calabi–Yau

threefold brane-probe theories, plus the conifold as the iconic non-del-Pezzo lead. The finite list is the standard classification of smooth toric del Pezzo surfaces; it is not a classification of their gauge phases.

- **Axis 2**, the Sasaki–Einstein families (Table VII.11.3): the infinite, sampled tail. The $Y^{p,q}$ and $L^{a,b,c}$ cones, where the volume can give rational or irrational central charges. We tabulate a representative sample with the general closed-form volume as the rule.

- **Axis 3**, the brane-brick CY4 sample (Table VII.11.4): the fourfold construction. A representative sample of two-dimensional $(0, 2)$ theories from D1-branes at toric fourfolds, with the chiral / Fermi quiver and the J/E pair, but no central charge (c-extremization is routed).

Axes 2 and 3 are sampled, not classified; axis 1 is the one complete list. To these the chapter adds a fourth, short subtable of **orientifold fixtures** (Table VII.11.5): the same toric cones folded by the dimer involution of Chapter 9, recording the surviving real gauge groups and rank relations. It is a small fixture set, the orientifold catalogue Chapter 9 routes here, not a classification of inequivalent involutions.

What "complete" means here. The catalog is complete relative to a stated rule, not in any absolute sense. On axis 1 the rule is the toric condition, and it gives the standard geometric list: the smooth toric del Pezzo surfaces are exactly $\{\mathbb{P}^2, F_0, dP_1, dP_2, dP_3\}$, because dP_n for $n \geq 4$ is non-toric (its automorphism group does not contain a two-torus), and these together with the conifold are the axis-1 entries. On axes 2 and 3 the families are infinite and the landscape is vast; we tabulate a sample and record the general formula, naming the full classification as out of scope. One reading is worth separating at the outset: a complete classification of toric Sasaki–Einstein manifolds, or of brane-brick models, is not what this atlas is, and the rows present are representatives, not a closure.

The recipe in one box. The following procedure organizes the toric-threefold rows. Its physical inputs are the brane and SCFT identifications of the preceding chapters.

From a toric diagram to its catalog row.

1. Draw the toric diagram (the lattice polygon Δ); read its outward normals, the legs of the (p, q)-web. The number of gauge nodes is $G = 2\,\mathrm{Area}(\Delta)$.
2. Choose a toric phase: one dimer (a bipartite tiling of T^2) with valid face signs and the required physical interpretation, whose Kasteleyn determinant has Δ as its Newton polygon. Matching support alone is insufficient. The inverse is not unique; the displayed choices represent phases in a demonstrated reachable component (Chapter 7).
3. Read off the periodic quiver: the gauge group $\prod_i U(N_i)$, the bifundamentals X_{ij}, the superpotential W (white vertices clockwise positive, black counterclockwise negative).
4. Compute the classical mesonic quotient. For a specified interacting SCFT, maximize over all admissible non-anomalous trial currents. The leading holographic comparison is $a = \pi^3 N^2/(4\,\mathrm{Vol}(Y)) + O(1)$, with the Sasaki–Einstein metric and holographic regime assumed.
5. Explore Seiberg-dual toric phases (a square move at an $N_f = 2N_c$ face); find the anomaly-free fractional ranks $\ker(a - a^{\mathsf{T}})$ minus the regular direction, and the cascade rank step where one runs.

The next section walks this procedure on the conifold, with dP_0 and F_0 as the second and third examples. The rational core volumes and rank kernels are recomputed from fixed polygons and adjacencies. The dP_2 central charge is checked by substituting into a sourced trial polynomial. The phase catalogs, the physical interpretation of the rank directions, and the general $L^{a,b,c}$ family retain their source dependence. This division lets one use the tables without mistaking a list of known examples for an exhaustive dynamical construction.

VII.11.2 The recipe, walked: the conifold, dP_0, and F_0

We run three examples through the recipe. The conifold is the iconic lead, the one that appears in the most chapters; dP_0 is the smallest chiral case; F_0 is the simplest worked cone with two distinct toric phases. Together they exhibit every column the catalog carries. Table VII.11.1 collects the three rows first, so the data shape is visible before the prose; the paragraphs below explain their different origins.

The conifold. The toric diagram is the unit square $\Delta = \{(0,0), (1,0), (1,1), (0,1)\}$, with $G = 2\,\mathrm{Area} = 2$ gauge nodes and four outward (p, q)-legs. The dimer is the square tiling of T^2; its quiver is the Klebanov–Witten theory, $U(N_1) \times U(N_2)$ with four bifundamentals A_1, A_2 (node $1 \to 2$) and B_1, B_2 (node $2 \to 1$) and the quartic

Table VII.11.1: The three worked comparisons. G counts gauge nodes; R is the symmetric scalar R-charge; a/N^2 denotes the leading coefficient. Phase counts are catalogued. The fractional count is the rank of the anomaly-free kernel modulo the regular direction, without a conclusion about its dynamics.

cone	Δ	G	W type	R	a/N^2	phases	frac.
conifold	unit square	2	quartic	1/2	27/64	1	1
dP_0	$\mathbb{P}^2$ triangle	3	cubic	2/3	3/4	1	0
F_0	square (+1 int.)	4	quartic	1/2	27/32	2	1

superpotential

$$W = \epsilon^{ij}\epsilon^{kl}\,\mathrm{Tr}(A_i B_k A_j B_l)\,, \tag{VII.11.1}$$

with the overall coupling suppressed. In its two cyclic terms each field appears once with positive and once with negative sign. This is the toric superpotential condition, not a uniqueness theorem for theories with two nodes. At equal ranks, the gauge-face condition is $\sum_{e\in f}(1-R_e)=2$, while each superpotential vertex obeys $\sum_{e\ni v}R_e=2$. For the conifold both give $R=1/2$ on the symmetric ansatz. For dP_0 the cubic vertices and six-edge faces instead give $3R=2$ and $6(1-R)=2$. Polygon corner charges used in volume minimization are yet another set of variables; their sum two should not be imposed on every gauge face.

The trial central charge $a(R)/N^2=\frac{3}{32}\big[2G+\sum_a(3(R_a-1)^3-(R_a-1))\big]$ evaluated at $R=\frac{1}{2}$ gives, equivalently by the Martelli–Sparks–Yau volume minimization over the unit square,

$$\frac{a_{\text{conifold}}}{N^2}=\frac{3}{4\,Z_{\min}}=\frac{27}{64},\qquad Z_{\min}=\frac{16}{9},\quad \mathrm{Vol}(T^{1,1})=\frac{\pi^3}{3}Z_{\min}=\frac{16\pi^3}{27}. \tag{VII.11.2}$$

The conifold is single-phase (one toric phase: it is self-dual under the square move). The antisymmetrized adjacency $a-a^{\mathsf{T}}$ of the two-node quiver has a two-dimensional kernel, so there is one anomaly-free fractional direction beyond the regular brane: the deformation-type fractional brane that drives the Klebanov–Strassler cascade. Written on the interacting (nonabelian) sector, the diagonal $U(1)$'s decoupled, the cascade is one Seiberg step

$$SU(N{+}M)\times SU(N)\;\longrightarrow\;SU(N)\times SU(N{-}M), \tag{VII.11.3}$$

The larger-rank coordinate decreases by M after relabelling; the dualized factor itself changes from $N+M$ to $N-M$. This is an integer rank relation under the Seiberg prescription, in a regime where the dual frame exists. It is neither the continuous asymptotic Maxwell charge nor a derivation of the quantized Page charge. At $N=M$ the confining $SU(2M)$ factor has $N_f=N_c$ and one uses its quantum constraint instead of interpreting a zero magnetic rank as another step. The usual confining endpoint requires $N=kM$, $k\geq 1$, $M\geq 2$, and a choice of branch (Chapter 8). These are the conifold's six catalog fields: toric diagram, dimer-quiver $+W$, $a=27/64$, one phase, one fractional direction, the cascade step.

$dP_0 = \mathbb{C}^3/\mathbb{Z}_3$. The toric diagram is the triangle $\{(1,0),(0,1),(-1,-1)\}$ (the local $\mathbb{P}^2$ diagram, one interior lattice point), with $G = 2\,\mathrm{Area} = 3$ nodes and three legs. The quiver is the chiral $U(N)^3$ theory with a directed triangle of three bifundamentals each and the cubic $W = \epsilon_{abc}\,\mathrm{Tr}(X^a_{01}X^b_{12}X^c_{20})$. Each cubic term carries one of each field type, so the symmetric R-charge is $R = \frac{2}{3}$, and

$$\frac{a_{dP_0}}{N^2} = \frac{3}{4\,Z_{\min}} = \frac{3}{4}, \qquad Z_{\min} = 1, \quad \mathrm{Vol}(S^5/\mathbb{Z}_3) = \frac{\pi^3}{3}. \tag{VII.11.4}$$

Unlike the conifold, dP_0 has $N_f = 3N_c$ at each node, no $N_f = 2N_c$ face, hence no rank-preserving square move: the displayed regular phase has no such outgoing move. Its catalog count is one. Its 3×3 antisymmetric adjacency has rank two, so $\ker(a - a^{\mathsf{T}})$ is one-dimensional, the all-ones regular direction only: dP_0 has *zero* anomaly-free fractional directions. This excludes unequal anomaly-free ranks in this unflavoured quiver. The conifold has deformation-type fractional branes; dP_1 supplies a different, DSB-type example whose possible runaway depends on dynamical moduli.

$F_0 = \mathbb{P}^1 \times \mathbb{P}^1$. The toric diagram is the square $\{(1,0),(0,1),(-1,0),(0,-1)\}$ (one interior point), with $G = 2\,\mathrm{Area} = 4$ nodes. F_0 is the worked cone with a non-self-equivalent move at an $N_f = 2N_c$ face: applying the dimer square move (urban renewal) at that face takes phase I (eight bifundamentals) to a distinct phase II (twelve fields), the two sharing the toric diagram, the Hilbert series, and the central charge but differing in field content. F_0 therefore has at least these two phases; its catalogued count is two. Its leading central coefficient is the doubled conifold value,

$$\frac{a_{F_0}}{N^2} = \frac{27}{32} = 2 \cdot \frac{27}{64} = 2\,\frac{a_{\mathrm{conifold}}}{N^2}, \tag{VII.11.5}$$

where both values use the leading large-N normalization. For the interacting $SU(N)^G$ theories in these symmetric examples, subtracting the missing abelian vector contributions gives

$$a_{\mathrm{conifold}} = \frac{27}{64}N^2 - \frac{3}{8}, \qquad a_{F_0} = \frac{27}{32}N^2 - \frac{3}{4}, \tag{VII.11.6}$$

so these particular SU polynomials also double. This statement uses the same prescribed matter and exact symmetric charges; it is not a generic finite-N holographic relation or a statement about arbitrary choices of abelian sectors. The F_0 doubled four-cycle has adjacency $a_{i,i+1} = 2$ with indices modulo four. The equation $(a - a^{\mathsf{T}})n = 0$ says $n_0 = n_2$, $n_1 = n_3$. Thus its kernel has rank two and the quotient by $(1,1,1,1)$ has one fractional direction.

Two further comparisons prevent geometrically similar expressions from being confused. In the fixed conifold signing of Chapters 3 and 5 the polynomial is $1-z-w-zw$. The all-positive expression has the same four lattice exponents, but $1 + z + w + zw = (1+z)(1+w)$ is reducible. For $P = a+bz+cw+dzw$ the four-punctured rational curve in $(\mathbb{C}^*)^2$ requires $abcd(ad - bc) \neq 0$; Newton support alone does not guarantee smoothness or a faithful physical coamoeba. The generic weighted expression in Chapter 1 only compares support. Similarly the inherited signed F_0 polynomial in Chapter 7 was only a support calculation, not a valid Kasteleyn signing.

The F_0 mesonic series uses the dual cone. At degree k its integral points obey $-k \le u, v \le k$, giving $(2k+1)^2$ monomials. Hence

$$H_{F_0}(t) = \sum_{k\ge 0} (2k+1)^2 t^k = \frac{1+6t+t^2}{(1-t)^3}. \tag{VII.11.7}$$

Counting lattice points in dilations of the primal diamond would instead give $(1+2t+t^2)/(1-t)^3$. That is a different counting problem. The physical identification of the dual semigroup with the mesonic ring remains the one imported from Chapters 5 and 7.

VII.11.3 Axis 1: the toric del Pezzo core

For a toric Calabi–Yau threefold cone, the number of gauge nodes is $G = 2\,\mathrm{Area}(\Delta)$ in lattice units, the central charge is the a-maximization value equal to the Martelli–Sparks–Yau volume minimization, $a/N^2 = 3/(4Z_{\min})$ with $\mathrm{Vol}(Y) = (\pi^3/3)Z_{\min}$, and the number of anomaly-free fractional directions is $\dim\ker(a - a^{\mathsf{T}}) - 1$. Not every entry is rational: the symmetric cones (conifold, dP_0, F_0, dP_3) carry rational leading central coefficients recomputed here, while the less symmetric dP_1 and dP_2 have irrational a, the a-maximization landing on a surd exactly as on axis 2. The leading doubling of F_0 and the conifold is one useful comparison, not a criterion for identifying a physical quiver.

Table VII.11.2: Axis 1: the smooth toric del Pezzo list, with the conifold added. a/N^2 denotes the leading coefficient. A_2 is the exact dP_2 value in (VII.11.10). Phase counts are sourced catalogs, not exhaustive results of the local move calculation. "Frac." counts the nonabelian anomaly kernel modulo the regular direction. "Cited" leaves the corresponding rank and dynamical data with the literature rather than a local calculation.

cone	G	a/N^2	phases	frac.	cascade	derived in
conifold ($T^{1,1}$)	2	$27/64$	1	1	yes	VII.1, VII.8
$dP_0 = \mathbb{C}^3/\mathbb{Z}_3$	3	$3/4$	1	0	no	VII.2, VII.5
dP_1	4	$13\sqrt{13} - 46$	1	1	DSB type	VII.8
dP_2	5	A_2	2	cited	cited	VII.7
dP_3	6	$9/8$	4	cited	cited	VII.7
$F_0 = \mathbb{P}^1 \times \mathbb{P}^1$	4	$27/32$	2	1	cited	VII.7

The phase column compares known presentations. A quadrilateral face with $N_f = 2N_c$ permits a rank-preserving square move, but the conifold shows that the result can be self-equivalent. The F_0 calculation exhibits different field counts, eight and twelve. The numbers two and four for dP_2 and dP_3 belong to the published phase catalogs. Local moves explore a reachable component; they do not prove that every possible presentation of the cone has been found.

The fractional column asks a different question. For an unflavoured chiral quiver, regular ranks must first lie in $K = \ker_{\mathbb{Z}}(a - a^{\mathsf{T}})$. Since the all-ones vector is primitive,

$K/\mathbb{Z}(1,\dots,1)$ is free of rank $\dim_{\mathbb{Q}}\ker(a-a^{\mathsf{T}})-1$. The conifold, dP_1, and F_0 each have one such direction, while dP_0 has none. A nonzero count is necessary for these rank deformations, not sufficient for a cascade, a stable brane, or a supersymmetric endpoint. In particular the dP_1 fractional brane is of DSB type, unlike the conifold deformation brane. Chapter 8 explains the additional superpotential and strong-coupling premises and when a dynamical FI modulus permits runaway.

The central-charge column supplies an exact contrast. The symmetric hexagon for dP_3 gives $a/N^2 = 9/8$; $dP_1 = Y^{2,1}$ gives $13\sqrt{13}-46$. For dP_2, the sourced four-current trial space gives a quadratic surd, not a cubic root. Write $A = a/N^2$ and take

$$\begin{aligned} A(x,y,z,w) = \frac{9}{32}\big[&5 + (x-1)^3 + (y-1)^3 + (z-1)^3 + (w-1)^3 \\ &+ (w-x-1)^3 + (x+z-1)^3 + 2(x+y-w-1)^3 \\ &+ (1-x-y)^3 + 2(1-x-y-z)^3\big]. \end{aligned} \tag{VII.11.8}$$

The multiplicities and the admissible trial-current space are physical inputs. Within that space the point

$$(x,y,z,w) = \left(\frac{\sqrt{33}-5}{2}, \frac{9-\sqrt{33}}{4}, \frac{3(19-3\sqrt{33})}{16}, \frac{17-\sqrt{33}}{16}\right) \tag{VII.11.9}$$

has zero gradient and negative-definite Hessian. Substitution yields

$$A_2 = \frac{243}{1024}(-59+11\sqrt{33}), \qquad \frac{\mathrm{Vol}(Y_{dP_2})}{\pi^3} = \frac{59+11\sqrt{33}}{486}. \tag{VII.11.10}$$

The second formula uses the leading holographic relation. The nonzero coefficient of $\sqrt{33}$ establishes irrationality exactly. A decimal extremum, by itself, would not establish this. Negative curvature in this trial space proves a strict local maximum, conditional on that space; identifying it as the exact physical R-current also assumes an interacting SCFT and no omitted accidental currents.

VII.11.4 Axis 2: the Sasaki–Einstein families

The $Y^{p,q}$ spaces, for coprime positive integers $p > q$, give an infinite family of Sasaki–Einstein five-manifolds with topology $S^2 \times S^3$. The three-integer family $L^{a,b,c}$ includes them as $Y^{p,q} = L^{p-q,p+q,p}$. Their leading central charges need not be rational, but they need not be irrational either. The distinction is controlled by the square root in the Gauntlett–Martelli–Sparks–Waldram volume,

$$\mathrm{Vol}(Y^{p,q}) = \frac{\pi^3\, q^2\left(2p+\sqrt{4p^2-3q^2}\right)}{3\,p^2\left(3q^2-2p^2+p\sqrt{4p^2-3q^2}\right)}, \tag{VII.11.11}$$

with the metric normalized as in Chapter 6. Rationalizing the expression gives

$$\frac{\mathrm{Vol}(Y^{p,q})}{\pi^3} = \frac{8p^3-9pq^2+(4p^2-3q^2)^{3/2}}{27p^2(p^2-q^2)}. \tag{VII.11.12}$$

For $0 < q < p$ the coefficient of the square root is nonzero, so the volume and its nonzero reciprocal are irrational precisely when $4p^2 - 3q^2$ is not an integer square. For example, $(p, q) = (7, 3)$ gives $169 = 13^2$ and

$$\frac{\mathrm{Vol}(Y^{7,3})}{\pi^3} = \frac{81}{980}, \qquad \frac{a_{Y^{7,3}}}{N^2} = \frac{245}{81}. \tag{VII.11.13}$$

This is a rational member with $q \neq 0$. The continuous formula has limit $16/(27p)$ as $q \to 0$; $p = 1$ recovers $T^{1,1}$. This limit of the formula is distinct from the positive coprime domain used to label the smooth family.

Table VII.11.3: Axis 2: selected Sasaki–Einstein volumes from (VII.11.11) and the leading relation $a/N^2 = \pi^3/(4\mathrm{Vol})$. The rational member $Y^{7,3}$ prevents the irrational examples from being read as a universal rule. $L^{a,b,c}$ is a family pointer; no general member is computed here.

member	$\mathrm{Vol}(Y)/\pi^3$	rational?	a/N^2
$T^{1,1} = Y^{1,0}$	$16/27$	yes	$27/64$
$Y^{2,1}$	$\frac{23}{162} + \frac{13\sqrt{13}}{324}$	no	$13\sqrt{13} - 46$
$Y^{3,2}$	$\frac{4}{45} + \frac{16\sqrt{6}}{405}$	no	$\frac{27\sqrt{6}}{4} - \frac{243}{16}$
$Y^{3,1}$	$\frac{7}{72} + \frac{11\sqrt{33}}{648}$	no	$\frac{297\sqrt{33}}{4} - \frac{1701}{4}$
$Y^{7,3}$	$81/980$	yes	$245/81$
$L^{a,b,c}$	general (a, b, c) family	varies	$Y^{p,q} = L^{p-q,p+q,p}$

The table is the historic AdS/CFT test made into a catalog row. The $Y^{p,q}$ quivers were constructed from the toric diagrams, their a-maximization performed in the gauge theory, and the result matched against the Sasaki–Einstein volume computed from the metric: agreement for the specified physical SCFT and metric. The atlas computes the geometric value and imports the gauge-side identification from Chapter 6, and notes the closed form (VII.11.11) as the rule that generates every member. The $L^{a,b,c}$ row carries the family and its $Y^{p,q}$ subfamily, the general closed-form volume being the entry the catalog points to rather than tabulating member by member. What the atlas does *not* carry is the full landscape: the classification of toric Sasaki–Einstein manifolds, and the existence theory behind Z-minimization, are named and routed, not catalogued.

VII.11.5 Axis 3: the brane-brick CY4 sample

For the fourfold construction, a D1-brane at a toric Calabi–Yau fourfold gives a two-dimensional $(0, 2)$ gauge theory, encoded by a brane-brick model: a tessellation of a three-torus whose bricks are gauge nodes, whose oriented faces are chiral multiplets, and whose unoriented faces are Fermi multiplets (Chapter 10). Each Fermi Λ_a carries a pair of holomorphic data, an E_a-term and a J_a-term, and the $(0, 2)$ consistency condition is

$$\sum_a \mathrm{Tr}(E_a J_a) = 0, \tag{VII.11.14}$$

the two-dimensional analog of the dimer's term-cancellation. The atlas tabulates a representative *sample*, not a catalog: a complete brane-brick atlas is out of scope, the toric fourfold landscape being vast.

Table VII.11.4: Axis 3: two sourced D1-brane constructions. The multiplet counts refer to $(0,2)$ notation; flat $\mathbb{C}^4$ has enhanced $(8,8)$ supersymmetry. The displayed E/J assignments satisfy a global cyclic trace identity for matrices of any size. This algebraic fact is distinct from the physical brane dictionary and the nonabelian gauge-anomaly condition. Neither row supplies an infrared central charge or a newly constructed triality-dual brick.

CY4	G	chirals	Fermis	$(0,2)$ data
$\mathbb{C}^4$	1	4	3	$E_a = [X_4, X_a]$, $J_a = [X_{a+1}, X_{a+2}]$
$\mathbb{C}^4/\mathbb{Z}_4$	4	16	12	E_a, J_a projected; $\sum_a \mathrm{Tr}(E_a J_a) = 0$

The $\mathbb{C}^4$ brick is the $(0,2)$ analog of the $\mathbb{C}^3$ honeycomb: a single gauge node, four chirals $X_1, \ldots, X_4$ for the four complex directions, and three Fermis $\Lambda_1, \Lambda_2, \Lambda_3$ with

$$E_a = [X_4, X_a], \qquad J_a = [X_{a+1}, X_{a+2}] \quad \text{(cyclic in } 1,2,3), \tag{VII.11.15}$$

for which $\sum_a \mathrm{Tr}(E_a J_a) = \mathrm{Tr}\big(X_4 \sum_{\text{cyclic}} [X_1, [X_2, X_3]]\big) = 0$ is the Jacobi identity. The proof uses cyclicity and the Jacobi identity in an associative algebra. Noncommuting matrix examples test the formulas but are not the proof. The toric diagram is the $\mathbb{C}^4$ tetrahedron (the unit three-simplex in the height-one slice). The row carries no central charge: in two dimensions the relevant extremization is c-extremization, and the elliptic genus and infrared data of these theories are the Chapter 10 deferral, routed to Part III.

The second row is the nontrivial fourfold Chapter 10 works, the orbifold $\mathbb{C}^4/\mathbb{Z}_4$ with the weights $(1,1,1,1)$. It is the $(0,2)$ analog of the Douglas–Moore orbifold quiver: the single $\mathbb{C}^4$ node unfolds into four nodes $\prod_{j=0}^{3} U(N_j)$, the four chirals $X_1, \ldots, X_4$ project to sixteen arrows ($X_k : j \to j{+}1$, one per node), and the three Fermis project to twelve, for the $(16, 12)$ content of the table. The $E_a = [X_4, X_a]$ and $J_a = [X_{a+1}, X_{a+2}]$ split across the nodes by their orbifold charges, but the cancellation is the same parent Jacobi identity carried through the projection, so $\sum_a \mathrm{Tr}(E_a J_a) = 0$ still holds after summing the trace over all nodes. It need not vanish at each node. To see the distinction, let each X_k be a 4×4 matrix whose only nonzero entries are $(X_k)_{j,j+1} = x_{k,j}$, with indices modulo four, and take

$$(x_{k,j}) = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 2 & 0 & 1 & 3 \\ 0 & 1 & 4 & 2 \\ 1 & 3 & 2 & 1 \end{pmatrix}. \tag{VII.11.16}$$

Direct multiplication with (VII.11.15) gives diagonal entries $(-45, 36, 5, 4)$ for $\sum_a E_a J_a$. Their sum is zero, although every entry is nonzero. Tensoring these arrows with the identity gives the same distinction at higher equal ranks. These off-shell matrices test an algebraic identity, not a vacuum or the existence of a physical gauge theory.

The nonabelian gauge anomaly is a separate sum over fermion representations. With $T(\square) = 1/2$ and equal $N \geq 2$, each orbifold node has eight incident chirals, six incident

Fermis, and one adjoint vector, giving

$$\mathcal{A}_j = \frac{N}{2}(8-6) - N = 0. \tag{VII.11.17}$$

The four equal-node evaluations test one independent coefficient expression. Chirality, representations, and the D1 projection are physical inputs to this calculation; abelian and global anomalies are separate. Chapter 10 discusses the general order-three triality prescription. No explicit triality-dual brick of this particular quotient is supplied here. The two rows compare the flat model with its nontrivial projection, with infrared dynamics left to Part III.

VII.11.6 Orientifold fixtures

Chapter 9 folds a brane tiling by an involution of the dimer, fixed loci carrying O-plane signs, and reads off the surviving real gauge group and the projected matter. That chapter routes its catalogue of inequivalent involutions to the atlas; we record it here as a short fixture set, three specified toric-cone projections, not a classification. Each row is the involution of Chapter 9 applied to the cone's dimer, its gauge group the fixed-face projection $U(N) \to SO(N)$ or $USp(N)$ on a self-mapped node and $U(N)$ on a paired orbit, its rank relation (when one is forced) the closure of the projected cubic anomaly.

Table VII.11.5: Scoped orientifold representatives from Chapter 9. The flat O3 example has three adjoint chirals; the conifold uses the colour-preserving fixed-line projection with two independent bifundamentals. The uniform unflavoured dP_0 cubic has the two branches shown. Ranks refer to fundamental dimensions, with USp dimension even. These local data do not assert RR tadpole or global consistency.

cone	gauge	matter	rank
$\mathbb{C}^3$	$SO(N_0)$	3 adjoints	n/a
conifold	$SO(N_0) \times USp(N_1)$	2 bifund.	none
dP_0	$SO(N_0) \times U(N)$	$3Q + 3\Lambda^2$	$N_0 = N - 4$
dP_0	$USp(N_0) \times U(N)$	$3Q + 3\mathrm{Sym}^2$	$N_0 = N + 4$

The flat $\mathbb{C}^3$ O3 example preserves the three adjoint chirals of $\mathcal{N} = 4$ Yang–Mills. An arbitrary choice of three tensors would not specify that theory. The conifold fixed-line projection leaves $SO \times USp$ and two independent bifundamentals. It has no surviving unitary node on which an SU cubic anomaly could impose a rank equation. This does not eliminate global or RR constraints on a complete string background.

The dP_0 example retains one unitary factor. Let its rank be $N \geq 3$, with three Q fields antifundamental under $U(N)$ and three tensors of symmetry $t = \pm 1$. In fundamental cubic-index units,

$$\mathcal{A}_{SU(N)^3} = -3N_0 + 3(N + 4t), \qquad N_0 = N + 4t. \tag{VII.11.18}$$

The uniform epsilon-contracted cubic requires $st = -1$, where $s = +1$ for the orthogonal fixed-node form and $s = -1$ for the symplectic form. Thus the orthogonal branch has

antisymmetric tensors and $N_0 = N - 4$; the symplectic branch has symmetric tensors and $N_0 = N + 4$. Require a positive fixed-node dimension, an even symplectic dimension, and no additional flavours. Low-rank degeneracies and abelian anomalies need their own treatment. The ± 4 follows from tensor cubic indices; it is not a measurement of O-plane RR charge. Chapter 9 and the Volume VI orientifold construction supply the physical projection.

VII.11.7 Cross-catalog patterns

The tables compare several structures without identifying them with one another.

- A suitable quadrilateral face permits a rank-preserving square move. The conifold move is self-equivalent, whereas the worked F_0 move changes eight fields to twelve. Published phase counts contain information beyond this local test.
- The leading central coefficients are rational for the symmetric core and irrational for dP_1 and dP_2. In the positive coprime $Y^{p,q}$ family the discriminant $4p^2 - 3q^2$ decides rationality: $Y^{7,3}$ is rational even though $q \neq 0$. No universal irrationality statement applies to $L^{a,b,c}$.
- The fractional-direction count is $\dim \ker(a - a^{\mathsf{T}}) - 1$ for the specified unflavoured quiver with a regular null vector. Conifold and F_0 give one, dP_0 and $\mathbb{C}^3$ zero. Dynamics requires more information than this count; the dP_1 DSB-type direction is not a conifold deformation brane.
- The leading equality $a(F_0)/N^2 = 2a(\text{conifold})/N^2$ follows from the displayed spectra and volumes. Finite-rank comparisons require a choice of gauge sector and its vector contribution, as in (VII.11.6).
- For a physically appropriate consistent tiling, external web legs are identified with zigzag windings. Finite winding closure alone neither proves the general correspondence nor constructs the physical coamoeba.

These comparisons explain why a row needs several columns. A polygon fixes an area and a dual cone, but it does not by itself specify the interaction, the full set of toric presentations, or the infrared trajectory. The brane construction supplies the physical dictionary; the finite calculations then test consequences of that dictionary.

VII.11.8 How to extend the atlas, and what is not in it

Each axis has a different boundary. Axis 1 is complete: the smooth toric del Pezzo surfaces are exactly $\{\mathbb{P}^2, F_0, dP_1, dP_2, dP_3\}$, and the conifold completes the worked comparisons. The first cone outside it, the non-toric dP_4, is named as the boundary: dP_n for $n \geq 4$ has too small an automorphism group to be toric, so the construction does not reach it as a toric probe, and it is a pointer only. Axis 2 is sampled: to add a member

one runs the same Z-minimization on its toric diagram, and the closed form (VII.11.11) (and its $L^{a,b,c}$ generalization) gives the volume in closed form; the full classification of toric Sasaki–Einstein manifolds, and the existence theory behind Z-minimization, are cited, not catalogued. Axis 3 is a sample: a complete brane-brick atlas of toric fourfolds is out of scope, and the sample is flat $\mathbb{C}^4$ and its first nontrivial projection $\mathbb{C}^4/\mathbb{Z}_4$, the fixtures the brane-brick chapter constructs. The orientifold fixtures are likewise a short fixture set, three folded cones, not a classification of inequivalent involutions; that classification is named and routed.

Three boundaries sit firmly outside this volume. The toric and dimer foundations (toric geometry, the gauged linear sigma model, the Kasteleyn combinatorics) are used here and developed in Part I. The classification theorems for toric Sasaki–Einstein manifolds, and the existence of the Z-minimizing Reeb vector, are cited to Volume III. And the strong-coupling dynamics beyond the construction-intrinsic, including the Klebanov–Strassler supergravity spectrum, the holographic c-function, the two-dimensional elliptic genus, the duality-wall accumulation, c-extremization, belong to Part III. This chapter gathers the specified toric catalog and brane-brick sample entries into a reference. Their available quiver, central-charge or volume, phase and fractional data are presented with calculated and source-dependent columns distinguished.

Status and sources

It indexes the earlier constructions without reclassifying them. [WS] means a physical-worldsheet derivation in the untwisted string, not worldvolume algebra. No new WS payload is supplied here. Physical primary classifications remain unassigned where the source does not supply a derivation. [MATH] below refers to specified exact algebra under explicit hypotheses; it is distinct from a cited physical premise. [ANOM] records a specified anomaly coefficient with its representation and chirality inputs, not full consistency. [DEC] means a stated decoupling-limit assumption.

The conventions used are spinor and supersymmetry counting, fundamental-at-tail arrows, superpotential normalization, toric lattice orientation, McKay multiplicities, quiver gauge groups and ranks, bipartite tiling conventions, resolution and discrete data, and a-maximization and volume minimization. The individual constructions are developed in Chapters 1–10. The exercises use finite central-charge, mesonic-ring, rank-kernel and local anomaly calculations under the same source premises.

§VII.11.1–VII.11.2. The brane, consistent-tiling, mesonic-ring and SCFT identifications are imported from current Chapters 1–7. The physical conifold source is Klebanov and Witten (1998); tiling constructions are Franco, Hanany, Kennaway, et al. (2006) and Hanany and Kennaway (2005). The corrected signed conifold support, the smooth-curve condition, and the distinction between dual-cone Hilbert counting and primal Ehrhart counting follow the current Chapters 3–5 and 7, with their stated limits. Conifold and dP_0 charges, the F_0 kernel and the leading coefficients are exact finite calculations. A consistent brane interpretation, complete trial-current space without accidental symmetries, and a normalized Sasaki–Einstein metric are separate premises for the physical conclusions. Volume minimization is Martelli et al. (2006) and Martelli et al. (2008); neither a local Hessian test nor a table establishes those physical premises.

§VII.11.3. The five smooth toric del Pezzo surfaces are the imported classification in Kasprzyk (2006), Section 5.1, Corollary 5.1.3, Proposition 5.1.4 and Figure 5-1 (printed pages 55–57). The proposition classifies terminal Fano polygons, the corollary identifies their smooth toric surfaces, and the figure names the five surfaces. The companion does not reproduce or check the classification proof. The phase catalogs are Feng, Hanany, and He (2001) and Feng, Hanany, He, and Uranga (2001), respectively hep-th/0109063 (Sections 4.1 and 6.1–6.3 for F_0 and dP_0–dP_3) and hep-th/0104259 (Section 5). These chapter-local keys identify works by eprint; their spellings differ from the Chapter 7 bibliography. Their catalog scope is inherited; the four dP_3 phases are presented as evidence, not an exhaustive classification

theorem. The dP_2 trial polynomial, stationary point and value are Bertolini et al. (2004), equations (3.26)–(3.29). Exact substitution and the Hessian establish a strict local maximum in the stated four-current trial space; the SCFT and current space are imported. Fractional-brane interpretation and the DSB/runaway qualifications follow Chapter 8 and Franco, Hanany, Saad, et al. (2006), Sections 2, 3.1–3.2.3 and 4. Non-toric dP_{4+} is outside the toric list. Geometry and GLSM foundations remain conifold geometry and toric geometry.

§VII.11.4. The volume and rational subclass are Gauntlett et al. (2004), equations (3.7)–(3.8) and the rational-roots discussion (PDF pages 11–12). Gauntlett et al. (2006) supplies the related metric construction. The gauge theories are Benvenuti et al. (2005) and Franco, Hanany, Martelli, et al. (2006), with the gauge-side SCFT identification imported as in Chapter 6. The $L^{a,b,c}$ family pointer is Cvetic et al. (2005); no general-$L^{a,b,c}$ volume calculation is supplied. The nonsquare criterion and $Y^{7,3}$ substitution are exact mathematics within the displayed domain. Metric existence and the holographic interpretation remain separate; Sasaki–Einstein metric existence retains the broader geometric questions.

§VII.11.5. The D1/brick and orbifold dictionary is imported from current Chapter 10 and Franco, Lee, et al. (2016) and Franco, Lee, et al. (2017); the first entry is arXiv:1510.01744. The physical primary status is unassigned. Chapter 10's current conditional [MATH] cyclic-trace results remain distinct from its [ANOM] coefficient results. For the quotient there is one independent nonabelian anomaly expression at equal ranks, not four independent tests. The universal trace proof is cyclicity plus Jacobi; (VII.11.16) disproves the stronger nodewise assertion. The atlas gives the fourfold projection drawing. The specific triality-dual construction remains open. General triality, abelian anomalies and infrared quantities retain the Chapter 10 source and deferral scope, two-dimensional supersymmetric dynamics and toric fourfold and brane brick geometry.

§VII.11.6. Current Chapter 9 supplies the scoped projection, uniform cubic and rank conditions. The primary orientifold source is Franco, Hanany, Krefl, et al. (2007), Section 3.5.1 for dP_0 and Sections 4.1, 4.3.2 near (4.13) for the fixed-line conifold rules. The latter identification combines general rules and the conifold discussion; that location does not print a standalone two-factor spectrum table. The tensor indices and $-3N_0 + 3(N + 4t)$ calculation are conditional algebra. No RR charge, K-theory or global completion follows. Volume VI, Chapter 8 develops the corresponding folded quiver.

§VII.11.7–VII.11.8. Repeated summaries use the same scopes: leading central coefficients, a nonzero fractional kernel without guaranteed dynamics, catalogued phases, and conditional physical dictionaries. Yamazaki (2008) is the survey reference, not a replacement proof for the individual constructions. The existing boundaries

References

Benvenuti, S. et al. (2005). "An Infinite family of superconformal quiver gauge theories with Sasaki-Einstein duals". In: *JHEP* 06, p. 064. DOI: 10.1088/1126-6708/2005/06/064. arXiv: hep-th/0411264.

Bertolini, M., F. Bigazzi, and A. L. Cotrone (2004). "New checks and subtleties for AdS/CFT and a-maximization". In: *JHEP* 12, p. 024. DOI: 10.1088/1126-6708/2004/12/024. arXiv: hep-th/0411249.

Cvetic, M. et al. (2005). "New Einstein-Sasaki spaces in five and higher dimensions". In: *Phys. Rev. Lett.* 95, p. 071101. DOI: 10.1103/PhysRevLett.95.071101. arXiv: hep-th/0504225.

Feng, B., A. Hanany, and Y.-H. He (2001). "Phase structure of D-brane gauge theories and toric duality". In: arXiv: hep-th/0104259.

Feng, B., A. Hanany, Y.-H. He, and A. M. Uranga (2001). "Toric duality as Seiberg duality and brane diamonds". In: arXiv: hep-th/0109063.

Franco, S., A. Hanany, K. D. Kennaway, et al. (2006). "Brane dimers and quiver gauge theories". In: *JHEP* 01, p. 096. DOI: 10.1088/1126-6708/2006/01/096. arXiv: hep-th/0504110.

Franco, S., A. Hanany, D. Krefl, et al. (2007). "Dimers and orientifolds". In: *JHEP* 09, p. 075. DOI: 10.1088/1126-6708/2007/09/075. arXiv: 0707.0298 [hep-th].

Franco, S., A. Hanany, D. Martelli, et al. (2006). "Gauge theories from toric geometry and brane tilings". In: *JHEP* 01, p. 128. DOI: 10.1088/1126-6708/2006/01/128. arXiv: hep-th/0505211.

Franco, S., A. Hanany, F. Saad, et al. (2006). "Fractional branes and dynamical supersymmetry breaking". In: *JHEP* 01, p. 011. DOI: 10.1088/1126-6708/2006/01/011. arXiv: hep-th/0505040.

Franco, S., S. Lee, et al. (2016). "Brane Brick Models, Toric Calabi-Yau 4-Folds and 2d (0,2) Quivers". In: *JHEP* 02, p. 047. DOI: 10.1007/JHEP02(2016)047. arXiv: 1510.01744 [hep-th].

Franco, S., S. Lee, et al. (2017). "Brane Brick Models in the Mirror". In: *JHEP* 02, p. 106. DOI: 10.1007/JHEP02(2017)106. arXiv: 1609.01723.

Gauntlett, J. P. et al. (2004). "Sasaki-Einstein metrics on S**2 x S**3". In: *Adv. Theor. Math. Phys.* 8, pp. 711–734. DOI: 10.4310/ATMP.2004.v8.n4.a3. arXiv: hep-th/0403002.

Gauntlett, J. P. et al. (2006). "A New infinite class of Sasaki-Einstein manifolds". In: *Adv. Theor. Math. Phys.* 8, pp. 987–1000. DOI: 10.4310/ATMP.2004.v8.n6.a3. arXiv: hep-th/0403038.

Hanany, A. and K. D. Kennaway (2005). "Dimer models and toric diagrams". In: arXiv: hep-th/0503149.

Kasprzyk, A. M. (Mar. 2006). "Toric Fano Varieties and Convex Polytopes". PhD thesis. University of Bath.

Klebanov, I. R. and E. Witten (1998). "Superconformal field theory on three-branes at a Calabi-Yau singularity". In: *Nucl. Phys. B* 536, pp. 199–218. DOI: 10.1016/S0550-3213(98)00654-3. arXiv: hep-th/9807080.

Martelli, D., J. Sparks, and S.-T. Yau (2006). "The Geometric dual of a-maximisation for Toric Sasaki-Einstein manifolds". In: *Commun. Math. Phys.* 268, pp. 39–65. DOI: 10.1007/s00220-006-0087-0. arXiv: hep-th/0503183.

Martelli, D., J. Sparks, and S.-T. Yau (2008). "Sasaki-Einstein manifolds and volume minimisation". In: *Commun. Math. Phys.* 280, pp. 611–673. DOI: 10.1007/s00220-008-0479-4. arXiv: hep-th/0603021.

Yamazaki, M. (2008). "Brane Tilings and Their Applications". In: *Fortsch. Phys.* 56, pp. 555–686. DOI: 10.1002/prop.200810536. arXiv: 0803.4474 [hep-th].