Acta Math., 206 (2011), 93–125
DOI: 10.1007/s11511-011-0060-4

Constructing integrable systems of semitoric type

by

Álvaro Pelayo

Institute for Advanced Study
Princeton, NJ, U.S.A.

Washington University in St. Louis
St. Louis, MO, U.S.A.

San Vũ Ngọc

Université de Rennes 1
Rennes, France

1. Introduction

The present paper is motivated by some remarkable results proven in the 1980s by Atiyah, Guillemin–Sternberg and Delzant, in the context of Hamiltonian torus actions. Indeed, Atiyah [1, Theorem 1] and Guillemin–Sternberg [13] proved that if an n-dimensional torus acts on a compact, connected symplectic manifold (M, ω) in a Hamiltonian fashion, then the image $\mu(M)$ under the momentum map $\mu:=(\mu_1, ..., \mu_n): M \to \mathbb{R}^n$ is a convex polytope. Delzant [5] showed that if the dimension n of the torus is half the dimension of M, this polytope, which in this case is called a *Delzant polytope* (i.e. a convex polytope with the property that at each vertex of it there are precisely n codimension-1 faces with normals which form a $\mathbb{Z}$-basis of the integral lattice $\mathbb{Z}^n$), determines the isomorphism type of M, and moreover, M is a toric variety. He also showed that starting from any Delzant polytope one can construct a symplectic manifold with a Hamiltonian torus action for which its associated polytope is the one we started with.

From the viewpoint of symplectic geometry, the situation described by the momentum polytope is, nevertheless, very rigid. It is natural to wonder whether any of these striking results persist in the case where the torus is replaced by a non-compact group acting Hamiltonianly. The seemingly simplest case happens when the group is $\mathbb{R}^n$, and the study of these $\mathbb{R}^n$-actions is precisely the goal of the theory of integrable systems. Building on previous work of the authors, and of many other authors, we shall present a "Delzant-type" classification for integrable systems, for which one component of the system is generated by a Hamiltonian circle action; these systems are called *semitoric*.

The first author was partially supported by an NSF post-doctoral fellowship. This work was done while the first author was at the Massachusetts Institute of Technology (2007–2008) and at the University of California, Berkeley (2008–2010).

Let (M,ω) be a connected, symplectic 4-dimensional manifold, where we do not assume that M is compact. Any smooth function f on M induces a unique vector field $\mathcal{X}_f$ on M which satisfies $\omega(\mathcal{X}_f,\cdot)=-df$. It is called the *Hamiltonian vector field induced by* f. An *integrable system* on M is a pair of real-valued smooth functions J and H on M, for which the Poisson bracket $\{J,H\}:=\omega(\mathcal{X}_J,\mathcal{X}_H)$ identically vanishes on M, and the differentials dJ and dH are almost-everywhere linearly independent. Of course, here $(J,H)\colon M\to\mathbb{R}^2$ is the analogue of the momentum map in the case of a torus action. In some local symplectic coordinates of M, (x,y,ξ,η), the symplectic form ω is given by $d\xi\wedge dx+d\eta\wedge dy$, and the vanishing of the Poisson brackets $\{J,H\}$ amounts to the partial differential equation

$$\frac{\partial J}{\partial \xi}\frac{\partial H}{\partial x}-\frac{\partial J}{\partial x}\frac{\partial H}{\partial \xi}+\frac{\partial J}{\partial \eta}\frac{\partial H}{\partial y}-\frac{\partial J}{\partial y}\frac{\partial H}{\partial \eta}=0.$$

This condition is equivalent to J being constant along the integral curves of $\mathcal{X}_H$ (or H being constant along the integral curves of $\mathcal{X}_J$).

A *semitoric integrable system on* M is an integrable system for which the component J is a proper momentum map for a Hamiltonian circle action on M, and the associated map $F:=(J,H):M\to\mathbb{R}^2$ has only non-degenerate singularities in the sense of Williamson, without real-hyperbolic blocks. We also use the term 4-*dimensional semitoric integrable system* to refer to the triple $(M,\omega,(J,H))$. Recall that the properness of J means that the preimage by J of a compact set is compact in M (which is immediate if M is compact), and the non-degeneracy hypothesis for F means that, if p is a critical point of F, then there exists a 2×2 matrix B such that, if we set $\widetilde{F}=B\circ(F-F(p))$, then one of the following situations holds in some local symplectic coordinates (x,y,ξ,η) centered at p (meaning that $x=y=\xi=\eta=0$ at p):

(1) $\widetilde{F}(x,y,\xi,\eta)=\big(\eta+\mathcal{O}(\eta^2),\,\frac{1}{2}(x^2+\xi^2)+\mathcal{O}((x,\xi)^3)\big)$;
(2) $\widetilde{F}(x,y,\xi,\eta)=\frac{1}{2}(x^2+\xi^2,y^2+\eta^2)+\mathcal{O}((x,\xi,y,\eta)^3)$;
(3) $\widetilde{F}(x,y,\xi,\eta)=(x\xi+y\eta,x\eta-y\xi)+\mathcal{O}((x,\xi,y,\eta)^3)$.

The first case is called a *transversally* (or *codimension*-1) *elliptic singularity*; the second case is an *elliptic-elliptic singularity*; and the last case is a *focus-focus singularity.* In [16, Theorem 6.2] the authors constructed, starting from a given semitoric integrable system on a 4-manifold, a collection of five symplectic invariants associated with it and proved that these completely determine the integrable system up to isomorphisms. The goal of the present is to complement that work, by providing a general method to construct *any* 4-dimensional semitoric integrable system starting from an abstract collection of ingredients. Both throughout [16] and the present paper we make a generic assumption on our semitoric systems; this is explained in §2.1.

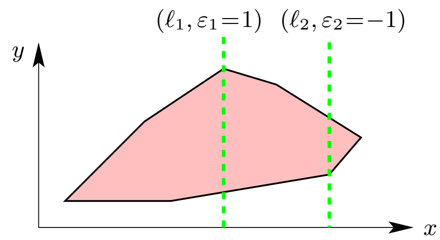

Figure 1.1. The weighted polygon $(\Delta, (\ell_1, \ell_2), (1, -1))$.

The symplectic invariants constructed in [16], for a given 4-dimensional semitoric integrable system, are the following: (i) *the number of singularities invariant*: an integer m_f counting the number of isolated singularities; (ii) *the singularity type invariant*: a collection of m_f infinite Taylor series in two variables which locally classifies the type of singularity; (iii) *the polygon invariant*: the equivalence class of a weighted rational convex([1]) polygon

$$(\Delta, \{\ell_j\}_{j=1}^{m_f}, \{\varepsilon_j\}_{j=1}^{m_f}).$$

Here Δ is a convex polygon domain in $\mathbb{R}^2$, the ℓ_j are vertical lines intersecting Δ and the ε_j are ± 1 signs giving each line ℓ_j an orientation; (iv) *the volume invariant*: m_f numbers measuring volumes of certain submanifolds at the singularities; (v) *the twisting index invariant*: m_f integers measuring how twisted the system is around singularities. This is a subtle invariant, which depends on the representative chosen in (iii). Here, we write m_f to emphasize that the singularities that m_f counts are focus-focus singularities. We then proved in [16] that two semitoric systems $(M, \omega_1, (J_1, H_1))$ and $(M, \omega_2, (J_2, H_2))$ are isomorphic if and only if they have the same invariants (i)–(v), where an isomorphism is a symplectomorphism $\varphi: M_1 \to M_2$ such that $\varphi^*(J_2, H_2) = (J_1, f(J_1, H_1))$ for some smooth function f such that $\partial f / \partial H_1$ vanishes nowhere.

We have found that some restrictions on these symplectic invariants must be imposed. Indeed, we call the following collection of items (i)–(v) the *semitoric list of ingredients*: (i) any integer number $0 \leqslant m_f < \infty$; (ii) an m_f-tuple of real formal power series in two variables, with vanishing constant term and first terms $\sigma_1 X + \sigma_2 Y$ with $\sigma_2 \in [0, 2\pi)$; (iii) a Delzant weighted polygon $(\Delta, \{\ell_j\}_{j=1}^{m_f}, \{\varepsilon_j\}_{j=1}^{m_f})$, of complexity m_f, where Δ is a polygon, the ℓ_j are again vertical lines intersecting Δ and the ε_j are ± 1 signs giving each line ℓ_j an orientation; here the Delzant property for Δ is not the standard one for polygons, but rather a more delicate one for weighted polygons which takes into account the presence of the lines ℓ_j; (iv) an m_f-tuple of positive real numbers $\{h_i\}_{i=1}^{m_f}$ such that $0 < h_i < \text{length}(\Delta \cap \ell_i)$ for each $i \in \{1, ..., m_f\}$. (v) an arbitrary collection of m_f integers

([1]) Generalizing the Delzant polygon and which may be viewed as a bifurcation diagram.

$\{k_i\}_{i=1}^{m_f}$. Our main theorem (Theorem 4.6) says that, starting from a semitoric list of ingredients one can construct a 4-dimensional semitoric integrable system $(M, \omega, (J, H))$ such that the list of its invariants is equal to this semitoric list. Moreover, M is compact if and only the polygon in item (iii) is compact.

With this in mind we may formulate the uniqueness theorem in [16] as: two systems constructed in this fashion are isomorphic if and only if the ingredients (i), (ii) and (iv) are identical for both systems and the ingredients (iii) and (v) are related by some simple transformation. This is why, when we formulate the existence theorem, the ingredients (iii) and (v) are given by orbits of weighted polygons and pondered weighted polygons, respectively, under the action of certain groups. Together with [16, Theorem 6.2], this gives the aforementioned classification (Theorem 4.7) .

While the construction of semitoric systems in the present paper is relatively self-contained, we are indebted to the articles of Delzant [5], Atiyah [1] and Guillemin–Sternberg [13], in the context of Hamiltonian torus actions, which served as an inspiration to study the more general situation of integrable systems with circular Hamiltonian symmetry. Furthermore, many papers have played an important role in our investigation of 4-dimensional semitoric systems, by serving as stepping stones to construct the symplectic invariants in [19] associated with semitoric systems; notably we used work of Dufour–Molino [6], Duistermaat [7], Eliasson [8], Miranda–Zung [15] and Vũ Ngọc [18], [19].

In this work, we are in a situation where the moment map (J, H) is a "torus fibration" with singularities, and its base space becomes endowed with a singular integral affine structure. These structures have been studied in the context of integrable systems (in particular by Zung [23]), but also became a central concept in the works by Symington [17], Symington–Leung [14], in the context of symplectic geometry and topology, and by Gross–Siebert [9], [10], [11] and [12], among others, in the context of mirror symmetry and algebraic geometry. In fact, our ingredients (i), (iii) and (iv) could have been expressed in terms of this affine structure. However, the ingredients (ii) and (v) do not appear in the affine structure. Nevertheless it is expected that these ingredients play an important role in the quantum theory of integrable systems. We hope to be able to explore these ideas in the future.

The paper is structured as follows: in §2 we recall how to construct a collection of symplectic invariants for a semitoric system, and state more precisely that two semitoric systems are isomorphic precisely when they have the same invariants; this was done in [16], and we need to review it here in order to state the existence theorem for semitoric systems. In §3 we explain the symplectic glueing construction (i.e. how to glue symplectic manifolds equipped with momentum maps). The last two sections of the paper are respectively devoted to state the main theorem and to prove it. One might argue that the proof is more informative than the statement, as it gives an *explicit* construction of all semitoric integrable systems in dimension 4.

Acknowledgments. We are very grateful to an anonymous referee for several interesting comments and remarks which have led to improvements. We thank Denis Auroux for offering several comments, and for pointing out the papers by Gross and Siebert.

2. Review of the uniqueness theorem for semitoric systems

We recall the definition of the invariants that we assigned to a semitoric integrable system in our previous paper [16], to which we refer for further details. Then we state the uniqueness theorem proved therein.

2.1. The Taylor series invariant

It was proven in [19] that a semitoric system $(M, \omega, F := (J, H))$ has finitely many focus-focus critical values $c_1, ..., c_{m_f}$, that if we write $B := F(M)$ then the set of regular values of F is $\mathrm{Int}(B) \setminus \{c_1, ..., c_{m_f}\}$, that the boundary of B consists of all images of elliptic singularities, and that the fibers of F are connected. The integer m_f was the first invariant that we associated with such a system. Let i be an integer, with $1 \leqslant i \leqslant m_f$.

We assume that the critical fiber $\mathcal{F}_{m_i} := F^{-1}(c_i)$ contains exactly one critical point m_i, which according to Zung [23] is a generic condition, and let $\mathcal{F}$ denote the associated singular foliation. Moreover, we make for simplicity an even stronger generic assumption: if m is a focus-focus critical point for F then m is the unique critical point of the level set $J^{-1}(J(m))$. A semitoric system is *simple* if this genericity assumption is satisfied. These conditions imply that the values $J(m_1), ..., J(m_{m_f})$ are pairwise distinct. We assume throughout the article that the critical values c_i are *ordered* by their J-values:

$$J(m_1) < J(m_2) < ... < J(m_{m_f}).$$

Vũ Ngọc proved in [18] that a semiglobal neighborhood of the ith focus-focus point is classified up to symplectic equivalence by a Taylor series $(S_i)^\infty \in \mathbb{R}[[X, Y]]$ in two variables with vanishing constant term. The series $(S_i)^\infty$ is called the *Taylor series invariant* of $(M, \omega, (J, H))$ at the focus-focus point c_i. For the purpose of this paper we do not need to know how to obtain this Taylor series from the integrable system; it suffices to know that it completely classifies symplectically a neighborhood of the singular fiber above the focus-focus point. A summary of the construction in [18] appeared in [16, §3.2].

2.2. The semitoric polygon invariant

The plane $\mathbb{R}^2$ is equipped with its standard affine structure with origin at $(0,0)$, and orientation. Let $\mathrm{Aff}(2,\mathbb{R}){:=}\mathrm{SL}(2,\mathbb{R})\ltimes\mathbb{R}^2$ be the group of affine transformations of $\mathbb{R}^2$. Let $\mathrm{Aff}(2,\mathbb{Z}){:=}\mathrm{SL}(2,\mathbb{Z})\ltimes\mathbb{R}^2$ be the subgroup of *integral-affine* transformations. Let $\mathfrak{I}$ be the subgroup of $\mathrm{Aff}(2,\mathbb{Z})$ of those transformations which leave a vertical line invariant, or equivalently, an element of $\mathfrak{I}$ is a vertical translation composed with a matrix T^k, where $k\in\mathbb{Z}$ and

$$T^k := \begin{pmatrix} 1 & 0 \\ k & 1 \end{pmatrix} \in \mathrm{SL}(2,\mathbb{Z}). \tag{2.1}$$

Let $\ell\subset\mathbb{R}^2$ be a vertical line in the plane, not necessarily through the origin, which splits it into two half-spaces, and let $n\in\mathbb{Z}$. Fix an origin in ℓ. Let $t_\ell^n\colon\mathbb{R}^2\to\mathbb{R}^2$ be the identity on the left half-space, and T^n on the right half-space. By definition, t_ℓ^n is piecewise affine. Let ℓ_i be a vertical line through the focus-focus value $c_i=(x_i,y_i)$, where $1\leqslant i\leqslant m_f$, and for any tuple $\vec{n}{:=}(n_1,...,n_{m_f})\in\mathbb{Z}^{m_f}$ we set

$$t_{\vec{n}} := t_{\ell_1}^{n_1}\circ...\circ t_{\ell_{m_f}}^{n_{m_f}}.$$

The map $t_{\vec{n}}$ is piecewise affine.

Definition 2.1. A *convex polygonal set* Δ is the intersection in $\mathbb{R}^2$ of (finitely or infinitely many) closed half-planes such that on each compact subset of the intersection there is at most a finite number of corner points. We say that Δ is *rational* if each edge is directed along a vector with rational coefficients. For brevity, in this paper we usually write *"polygon"* instead of *"convex polygonal set"*, see Remark 2.2.

Remark 2.2. The word "polygon" is commonly used to refer to the convex hull of a finite set of points in $\mathbb{R}^2$ which is a compact set (this is not necessarily the case in algebraic geometry, cf. e.g. Newton polygons). Notice (see Ziegler [22, Theorem 1.2]) that a convex polygonal set Δ which is not a half-space has exactly two edges of infinite Euclidean length if and only if it is non-compact, and Δ has all of its edges of finite length if and only if it is compact.

Some authors call "2-dimensional convex polyhedron" what we call "convex polygonal domain"; we prefer not to use this terminology as historically the word "polyhedron" has referred to 3-dimensional objects.

Let $B_r{:=}\mathrm{Int}(B)\setminus\{c_1,...,c_{m_f}\}$, which is precisely the set of regular values of F. Given a sign $\varepsilon_i\in\{-1,1\}$, let $\ell_i^{\varepsilon_i}\subset\ell_i$ be the vertical half-line starting at c_i and extending in the direction of ε_i: upwards if $\varepsilon_i=1$, downwards if $\varepsilon_i=-1$. Let

$$\ell^{\vec{\varepsilon}} := \bigcup_{i=1}^{m_f} \ell_i^{\varepsilon_i}.$$

In [19, Theorem 3.8] it was shown that for $\vec{\varepsilon}\in\{-1,1\}^{m_f}$ there exists a homeomorphism onto its image $f=f_{\varepsilon}: B\to f(B)\subset\mathbb{R}^2$, modulo a left composition by a transformation in $\mathfrak{I}$, such that $f|_{B\setminus\ell^{\vec{\varepsilon}}}$ is a diffeomorphism into its image $\Delta:=f(B)$, which is a *rational convex polygon* (or more precisely, a rational convex polygonal domain), $f|_{B_r\setminus\ell^{\vec{\varepsilon}}}$ is affine (it sends the integral affine structure of B_r to the standard structure of $\mathbb{R}^2$) and f preserves J, i.e. $f(x,y)=(x,f^{(2)}(x,y))$.

Let us recall what we mean by the affine structure of B_r [16]. An *integral affine manifold* is a smooth manifold that admits an atlas where transition functions are in $\mathrm{Aff}(n,\mathbb{Z}):=\mathrm{SL}(n,\mathbb{Z})\ltimes\mathbb{R}^n$. On the tangent plane at any point m of an integral affine manifold M, there is a well-defined lattice, namely the preimage of $\mathbb{Z}^n$ by the differential $d_m\varphi$ of a chart $\varphi: U\subset M\to\mathbb{R}^n$. Of course $\mathbb{R}^n$ itself is canonically an integral affine manifold, endowed with the tangent lattice $\mathbb{Z}^n$.

It follows from the action-angle theorem [7] that any proper Lagrangian fibration $F: M\to B$ naturally defines an integral-affine structure on the base B. This affine structure can be characterized by the following fact: a local diffeomorphism $g:(B,b)\to(\mathbb{R}^n,0)$ is an integral-affine chart if and only if the Hamiltonian flows of the n coordinate functions of $g\circ F$ are periodic of primitive period equal to 2π. Therefore an integrable system with proper momentum map $F=(J,H)$ defines an integral-affine structure on the set B_r of regular values of F. In our case, this structure extends to the boundary of B_r in a natural way.

Although B_r is a subset of $\mathbb{R}^2$, the integral-affine structure of B_r is in general different from the induced canonical integral-affine structure of $\mathbb{R}^2$.

In order to arrive at Δ one cuts $(J,H)(M)\subset\mathbb{R}^2$ along each of the vertical half-lines $\ell_i^{\varepsilon_i}$. Then the resulting image becomes simply connected and thus there exists a global 2-torus action on the preimage of this set. The polygon Δ is just the closure of the image of a toric momentum map corresponding to this torus action. Of course, if the system is toric (i.e. $m_f=0$), then Δ is the usual Delzant polygon. In the general semitoric case, there are important differences to be noted. One of them is that the H-coordinates of the vertices are not all critical values of H (while, as in the toric case, the J-coordinates of the vertices are exactly the critical values of J). Another important fact is that this polygon is not unique. The choice of the "cut direction" is encoded in the signs ε_j, and there remains some freedom for choosing the toric momentum map. Precisely, the choices and the corresponding homeomorphisms f are the following:

(a) *an initial set of action variables f_0 of the form (J,K)* near a regular Liouville torus in [19, proof of Theorem 3.8, step 2]. If we choose f_1 instead of f_0, we get a polygon Δ' obtained by left composition with an element of $\mathfrak{I}$. Similarly, if we choose f_1 instead of f_0, we obtain f composed on the left with an element of $\mathfrak{I}$;

(b) *a tuple* $\vec{\varepsilon}$ *of* 1 *and* -1. If we choose $\vec{\varepsilon}\,'$ instead of $\vec{\varepsilon}$ we obtain $\Delta'=t_{\vec{u}}(\Delta)$ with $u_i=\frac{1}{2}(\varepsilon_i-\varepsilon_i')$, by [19, Proposition 4.1, formula (11)]. Similarly, instead of f we get $f'=t_{\vec{u}}\circ f$.

LEMMA 2.3. *Once* f_0 *and* $\vec{\varepsilon}$ *have been fixed as in* (a) *and* (b), *respectively, then there exists a unique toric momentum map* μ *on* $M_r:=F^{-1}\big(\operatorname{Int} B\setminus\big(\bigcup_{j=1}^{m_f}\ell_j^{\varepsilon_j}\big)\big)$ *which preserves the foliation* $\mathcal{F}$, *and coincides with* $f_0\circ F$ *where they are both defined. Then, necessarily, the first component of* μ *is* J, *and we have*

$$\overline{\mu(M_r)}=\Delta. \tag{2.2}$$

Proof. The uniqueness follows from the fact that $\operatorname{Int} B\setminus\big(\bigcup_{j=1}^{m_f}\ell_j^{\varepsilon_j}\big)$ is simply connected, and (2.2) follows directly from the construction of Δ in [19], since $\mu=f\circ F$. □

We sometimes call μ the (*generalized*) *momentum map* associated with the polytope Δ. Let $\operatorname{Polyg}(\mathbb{R}^2)$ be the space of rational convex polygons in $\mathbb{R}^2$. Let $\operatorname{Vert}(\mathbb{R}^2)$ be the set of vertical lines in $\mathbb{R}^2$. A *weighted polygon* of *complexity* s is a triple of the form

$$\Delta_{\mathrm{w}}=(\Delta,\{\ell_{\lambda_j}\}_{j=1}^s,\{\varepsilon_j\}_{j=1}^s),$$

where s is a non-negative integer, $\Delta\in\operatorname{Polyg}(\mathbb{R}^2)$, $\ell_{\lambda_j}\in\operatorname{Vert}(\mathbb{R}^2)$ and $\varepsilon_j\in\{-1,1\}$ for every $j\in\{1,...,s\}$,

$$\min_{s\in\Delta}\pi_1(s)<\lambda_1<...<\lambda_s<\max_{s\in\Delta}\pi_1(s),$$

where $\pi_1:\mathbb{R}^2\to\mathbb{R}$ is the canonical projection $\pi_1(x,y)=x$ and $\pi_1(\ell_{\lambda_j})=\lambda_j$. For any $s\in\mathbb{N}$, let $G_s:=\{-1,1\}^s$. Obviously, an element of the group $\mathfrak{I}$ sends a rational convex polygon to a rational convex polygon. It corresponds to the transformation described in (a). On the other hand, the transformation described in (b) can be encoded by the group G_s acting on the triple Δ_{w} by the formula

$$\{\varepsilon_j'\}_{j=1}^s\cdot(\Delta,\{\ell_{\lambda_j}\}_{j=1}^s,\{\varepsilon_j\}_{j=1}^s)=(t_{\vec{u}}(\Delta),\{\ell_{\lambda_j}\}_{j=1}^s,\{\varepsilon_j'\varepsilon_j\}_{j=1}^s), \tag{2.3}$$

where $\vec{u}=\big\{\frac{1}{2}(\varepsilon_i-\varepsilon_i')\big\}_{i=1}^s$. This, however, does not always preserve the convexity of Δ, as is easily seen when Δ is the unit square centered at the origin and $\lambda_1=0$. However, when Δ comes from the construction described above for a semitoric system (J,H), the convexity is preserved. Thus, we give the following definition.

Definition 2.4. A weighted polygon is *admissible* when the G_s-action preserves convexity. We denote by $\mathcal{W}\operatorname{Polyg}_s(\mathbb{R}^2)$ the space of all admissible weighted polygons of complexity s.

The set $G_s \times \mathfrak{I}$ is an abelian group, with the natural product action. The action of $G_s \times \mathfrak{I}$ on $\mathcal{W}\,\mathrm{Polyg}_s(\mathbb{R}^2)$ is

$$(\{\varepsilon'_j\}_{j=1}^s, T^k+\alpha)\cdot(\Delta, \{\ell_{\lambda_j}\}_{j=1}^s, \{\varepsilon_j\}_{j=1}^s) = (t_{\vec{u}}((T^k+\alpha)(\Delta)), \{\ell_{\lambda_j}\}_{j=1}^s, \{\varepsilon'_j\varepsilon_j\}_{j=1}^s),$$

where $\vec{u}=\{\frac{1}{2}(\varepsilon_i-\varepsilon'_i)\}_{i=1}^s$ and α denotes the vertical translation $(x,y)\mapsto(x,y+\alpha)$.

Definition 2.5. A *semitoric polygon* is the equivalence class of an admissible weighted polygon under the $(G_{m_f}\times\mathfrak{I})$-action.

Let Δ be a rational convex polygon obtained from the momentum image $(J,H)(M)$ according to the above construction of cutting along the half-lines $\ell_1^{\varepsilon_1}, ..., \ell_{m_f}^{\varepsilon_{m_f}}$.

Definition 2.6. The *semitoric polygon invariant of* $(M,\omega,(J,H))$ is the semitoric polygon equal to the $(G_{m_f}\times\mathfrak{I})$-orbit

$$(G_{m_f}\times\mathfrak{I})\cdot(\Delta, \{\ell_j\}_{j=1}^{m_f}, \{\varepsilon_j\}_{j=1}^{m_f}) \in \mathcal{W}\,\mathrm{Polyg}_{m_f}(\mathbb{R}^2)/(G_{m_f}\times\mathfrak{I}). \tag{2.4}$$

2.3. The volume invariant

Consider a focus-focus critical point m_i whose image by (J,H) is c_i, and let Δ be a rational convex polygon corresponding to the system $(M,\omega,(J,H))$. If μ is a toric momentum map for the system $(M,\omega,(J,H))$ corresponding to Δ, then the image $\mu(m_i)$ is a point in the interior of Δ, along the line ℓ_i. We proved in [16] that the vertical distance

$$h_i := \mu(m_i) - \min_{s\in\ell_i\cap\Delta} \pi_2(s) > 0 \tag{2.5}$$

is independent of the choice of momentum map μ. Here $\pi_2\colon\mathbb{R}^2\to\mathbb{R}$ is $\pi_2(x,y)=y$. See [16, §5.1] for an explanation of the usage of the word "volume".

2.4. The twisting-index invariant

The twisting-index at a focus-focus point c_i is a dynamical invariant which expresses the fact that in a neighborhood of c_i there is a privileged toric momentum map ν; this momentum map encodes the existence of a unique hyperbolic radial vector field in a neighborhood of the focus-focus fiber. We refer to [16, §5, step 4] for a detailed construction of the privileged momentum map; what is most relevant in the proof of the existence result in this paper (§5, second stage) is the fact that we have the *privileged* choice of momentum map ν. Since any semitoric polygon defines a (generalized) toric momentum map μ, we define the twisting-index at c_i as the integer $k_i\in\mathbb{Z}$ such that

$d\mu = T^{k_i}\, d\nu$. The twisting index invariant is essentially the m_f-tuple of values k_i, but one has to be more careful [16, §5, step 5], because as was the case with the semitoric polygon invariant, we need to factor out the actions of $\mathfrak{I}$ and G_s in order to get the actual well-defined invariant. For any integer s, consider the action of the product $G_s \times \mathfrak{I}$ on the space $\mathcal{W}\,\mathrm{Polyg}_s(\mathbb{R}^2)\times\mathbb{Z}^s$:

$$(\{\varepsilon'_j\}_{j=1}^s, T^k+\alpha)\star(\Delta, \{\ell_{\lambda_j}\}_{j=1}^s, \{\varepsilon_j\}_{j=1}^s, \{k_j\}_{j=1}^s) \\ = (t_{\vec{u}}((T^k+\alpha)(\Delta)), \{\ell_{\lambda_j}\}_{j=1}^s, \{\varepsilon'_j\varepsilon_j\}_{j=1}^s, \{k_j+k\}_{j=1}^s),$$

where $\vec{u} = \left\{\frac{1}{2}(\varepsilon_j - \varepsilon'_j)\right\}_{j=1}^s$, for all integers $j \in \{1, ..., s\}$.

Definition 2.7. The *twisting-index invariant of* $(M, \omega, (J, H))$ is the $(G_{m_f}\times\mathfrak{I})$-orbit of a weighted polygon pondered by twisting indices at the focus-focus singularities of the system given by

$$(G_{m_f}\times\mathfrak{I})\star(\Delta, \{\ell_j\}_{j=1}^{m_f}, \{\varepsilon_j\}_{j=1}^{m_f}, \{k_j\}_{j=1}^{m_f}) \in (\mathcal{W}\,\mathrm{Polyg}_{m_f}(\mathbb{R}^2)\times\mathbb{Z}^{m_f})/(G_{m_f}\times\mathfrak{I}).$$

2.5. The uniqueness theorem

To a semitoric system we assign the above list of invariants and state the main theorem in [16].

Definition 2.8. Let $(M, \omega, (J, H))$ be a 4-dimensional simple semitoric integrable system. The *list of invariants of* $(M, \omega, (J, H))$ consists of the following items.

(i) The integer number $0 \leqslant m_f < \infty$ of focus-focus singular points.

(ii) The m_f-tuple $\{(S_i)^\infty\}_{i=1}^{m_f}$, where $(S_i)^\infty$ is the Taylor series of the ith focus-focus point.

(iii) The semitoric polygon invariant, cf. Definition 2.6.

(iv) The volume invariant, i.e. the m_f-tuple $\{h_i\}_{i=1}^{m_f}$, where h_i is the height of the ith focus-focus point.

(v) The twisting-index invariant, cf. Definition 2.7.

THEOREM 2.9. ([16, Theorem 6.2]) *The two* 4*-dimensional simple semitoric integrable systems* $(M_1, \omega_1, (J_1, H_1))$ *and* $(M_2, \omega_2, (J_2, H_2))$ *are isomorphic if and only if their lists of invariants* (i)–(v), *as in Definition* 2.8, *coincide.*

3. The symplectic glueing construction

In this section we explain how to symplectically glue an arbitrary collection of symplectic manifolds $\{M_\alpha\}_{\alpha\in A}$ equipped with continuous, proper maps $F_\alpha\colon M_\alpha \to \mathbb{R}$ to form a new symplectic manifold M equipped with a continuous, proper map which restricted to M_α is equal to F_α, cf. Theorem 3.11. The results of this section, while perhaps well-known among experts, we could not find in the literature.

3.1. Glueing maps, glueing groupoids

Let A be an arbitrary set of indices, and let $\{M_\alpha\}_{\alpha\in A}$ be a family of sets. Recall that the *disjoint union of the sets* M_α, $\alpha\in A$, is the subset of $\left(\bigcup_{\alpha\in A} M_\alpha\right)\times A$ defined by

$$\bigsqcup_{\alpha\in A} M_\alpha := \{(x,\alpha) : x\in M_\alpha\}.$$

We denote by j_α, $\alpha\in A$, the natural inclusions:

$$\begin{aligned} j_\alpha\colon M_\alpha &\hookrightarrow \bigsqcup_{\alpha\in A} M_\alpha, \\ x &\longmapsto (x,\alpha). \end{aligned}$$

Notice that if $B\subset A$ then $\bigsqcup_{\alpha\in B} M_\alpha\subset\bigsqcup_{\alpha\in A} M_\alpha$. Of course, if all M_α's are pairwise disjoint, as sets, then there is a natural bijection bewteen $\bigsqcup_{\alpha\in A} M_\alpha$ and the usual union $\bigcup_{\alpha\in A} M_\alpha$.

If the M_α's are topological spaces, the disjoint union $\bigsqcup_{\alpha\in A} M_\alpha$ is endowed with the final topology: the finest topology that makes the inclusions j_α continuous. In particular $j_\alpha(M_\alpha)$ is an open set in $\bigsqcup_{\alpha\in A} M_\alpha$.

Definition 3.1. A *glueing map for the family* $\{M_\alpha\}_{\alpha\in A}$ is a homeomorphism

$$\varphi\colon U_\alpha \longrightarrow U_\beta,$$

where $(\alpha,\beta)\in A^2$ and $U_\alpha\subset M_\alpha$ and $U_\beta\subset M_\beta$ are open sets.

In this text we use the standard set-theoretical convention that the notation φ includes the source and target sets U_α and U_β; in particular the notation $\varphi(x)$ implies that $x\in U_\alpha$. When required, we use the notation U^s_φ and U^t_φ for the source and target sets of φ (assuming $U^t_\varphi=\varphi(U^s_\varphi)$).

Definition 3.2. Let $\mathcal{G}$ be a collection of glueing maps for $\{M_\alpha\}_{\alpha\in A}$. The associated *glueing groupoid* G is the groupoid generated by the set of all restrictions of all glueing maps $\varphi\in\mathcal{G}$ to open subsets of the source sets, with the natural groupoid law: $\varphi_2\circ\varphi_1$ exists whenever the image of the source set of φ_1 is included in the source set of φ_2.

Definition 3.3. We say that $\mathcal{G}$ is *free* when there is no non-trivial $\varphi\in G$ with both source and target in the same set M_α.

3.2. Topological glueing

We now define the general patching construction. Throughout this section, and unless otherwise stated, we do not require topological spaces to be paracompact or Hausdorff.

Definition 3.4. Let $\{M_\alpha\}_{\alpha\in A}$ be a collection of pairwise disjoint topological spaces, and G an associated glueing groupoid. From this we define the set M, called the *glueing of* $\{M_\alpha\}_{\alpha\in A}$ *along* G, as $M:=\bigsqcup_{\alpha\in A} M_\alpha/{\sim}$, where $\sim$ is the equivalence relation on $\bigsqcup_{\alpha\in A} M_\alpha$ defined by

$$(x,\alpha)\sim(x',\beta) \quad\Longleftrightarrow\quad x=x' \text{ or there exists } \varphi\in G \text{ with } x'=\varphi(x).$$

Let us check that $\sim$ is indeed an equivalence relation. The reflexivity is obvious. If $(x,\alpha)\sim(x',\beta)$ and $(x,\alpha)\neq(x',\beta)$ then $\varphi(x)=x'$ for some $\varphi\in G$. But G is a groupoid, so $\varphi^{-1}\in G$ and of course $x=\varphi^{-1}(x')$, and thus $(x',\beta)\sim(x,\alpha)$, which proves the symmetry property. Finally, if $(x,\alpha)\sim(x',\beta)$ and $(x',\beta)\sim(x'',\gamma)$ then there exist φ and φ' in G such that $\varphi(x)=x'$ and $\varphi'(x')=x''$. Therefore $\varphi'\circ\varphi$ is well defined on an open neighborhood of x, so $\varphi'\circ\varphi\in G$, and $(x,\alpha)\sim(x'',\gamma)$, and hence we have shown the transitivity property.

Here again we could have dropped the assumption that the M_α's are pairwise disjoint, or we could have used a standard union instead of a disjoint union.

The following lemma follows from the definition of the equivalence relation.

LEMMA 3.5. *Let* $\pi:\bigsqcup_{\alpha\in A} M_\alpha\to M$ *be the quotient map. For any subset* $K\subset M_\alpha$, *one has*

$$\pi^{-1}(y_\alpha(K))=j_\alpha(K)\cup\Bigg(\bigcup_{\varphi\in G} j_{\alpha(\varphi)}(\varphi(K\cap U_\varphi^s))\Bigg),$$

where it is assumed that the union is over all φ *whose source set* U_φ^s *intersects* K, *and* $\alpha(\varphi)$ *is the element in* A *such that* $U_\varphi^t\subset M_{\alpha(\varphi)}$.

LEMMA 3.6. *For the natural quotient topology on* M, *the maps* $y_\alpha=\pi\circ j_\alpha: M_\alpha\to M$, $\alpha\in A$, *are open and continuous. They are injective if and only if* $\mathcal{G}$ *is free.*

Proof. By definition of quotient topology, the map π is continuous. Hence $y_\alpha=\pi\circ j_\alpha$ is continuous. Finally, if $U\subset M_\alpha$ is open, then it follows from Lemma 3.5 that $\pi^{-1}(y_\alpha(U))$ is open in $\bigsqcup_{\alpha\in A} M_\alpha$. This means that $y_\alpha(U)$ is open in M.

Fix $\alpha\in A$. Let x and x' be elements of M_α. If $y_\alpha(x)=y_\alpha(x')$ then either $x=x'$ or $\varphi(x)=x'$ for some $\varphi\in G$. The latter is ruled out by the assumption that there is no non-trivial $\varphi\in G$ with both source and target in M_α. Thus in this case y_α is injective. If the condition is violated then there exist $x\neq x'$ in M_α with $j_\alpha(x)\sim j_\alpha(x')$ so y_α cannot be injective. □

3.3. Smooth glueing

LEMMA 3.7. *If all the M_α's are smooth manifolds, all $\varphi\in\mathcal{G}$ are diffeomorphisms and $\mathcal{G}$ is free, then there exists a unique smooth structure on M for which the maps y_α, $\alpha\in A$, are embeddings.*

Proof. Let $U\subset M_\alpha$ be open and let $g: U\to\mathbb{R}^n$ be a homeomorphism. By Lemma 3.6, y_α is a homeomorphism onto its image. Let $\widetilde{U}=y_\alpha(U)$ and $\tilde{g}=g\circ(y_\alpha|_U)^{-1}$. Then $\widetilde{U}$ is an open subset of M and $\tilde{g}:\widetilde{U}\to\mathbb{R}^n$ is a homeomorphism. This shows that any chart of M_α descends onto a chart of M. Obviously the union of a family of open covers of M_α for all $\alpha\in A$ descends to an open cover of M. In order to get an atlas on M, it remains to check the compatibility condition when an open set $\widetilde{V}_\alpha$ coming from an atlas of M_α intersects an open set $\widetilde{V}_\beta$ coming from an atlas of M_β. Thus, let (V_α, g_α), $V_\alpha\subset M_\alpha$, and (V_β, g_β), $V_\beta\subset M_\beta$, be local charts such that $y_\alpha(V_\alpha)=y_\beta(V_\beta)$ and $\alpha\neq\beta$. Now consider the following formula, given by Lemma 3.5,

$$j_\alpha(V_\alpha)\cup\biggl(\bigcup_{\varphi\in G} j_{\alpha(\varphi)}(\varphi(V_\alpha\cap U_\varphi^s))\biggr)=j_\beta(V_\beta)\cup\biggl(\bigcup_{\varphi\in G} j_{\alpha(\varphi)}(\varphi(V_\beta\cap U_\varphi^s))\biggr).$$

Because $\mathcal{G}$ is free, any φ whose source set intersects V_α *and* such that $\alpha(\varphi)=\alpha$ must be the identity. Hence, in the left-hand side one can omit all the φ's such that $\alpha(\varphi)=\alpha$. For the same reason, one may assume that all the $\alpha(\varphi)$'s are pairwise different. Of course the analogue observation holds for the right-hand side. Hence we can equate terms in the unions (up to permutation). In particular there must exist some φ with $\alpha(\varphi)=\beta$ and $j_\beta(\varphi(V_\alpha\cap U_\varphi^s))=j_\beta(V_\beta)$. Since j_β is injective, $\varphi(V_\alpha\cap U_\varphi^s)=V_\beta$. Let $x\in V_\beta$ and $x'=\varphi^{-1}(x)\in V_\alpha$. Then $y_\alpha(x')=y_\beta(x)$, i.e. $x'=y_\alpha^{-1}\circ y_\beta(x)$. Thus $(y_\alpha|_{V_\alpha})^{-1}\circ(y_\beta|_{V_\beta})=\varphi^{-1}|_{V_\beta}$. Hence the transition map for the charts $\tilde{g}_u:=g_u\circ(y_u|_{V_u})^{-1}$, $u=\alpha,\beta$, is equal to

$$\tilde{g}_\alpha\circ\tilde{g}_\beta^{-1}=g_\alpha\circ((y_\alpha|_{V_\alpha})^{-1}\circ(y_\beta|_{V_\beta}))\circ g_\beta^{-1}=g_\alpha\circ\varphi^{-1}\circ g_\beta^{-1}, \tag{3.1}$$

which is indeed a composition of local diffeomorphisms. Thus M has a natural smooth structure.

Consider now $y_\alpha: M_\alpha\hookrightarrow M$. Read in a chart $(\widetilde{V}_\alpha, \tilde{g}_\alpha)$ of M, with $\tilde{g}_\alpha:=g_\alpha\circ(y_\alpha|_{V_\alpha})^{-1}$, for some chart (V_α, g_α) on M_α, it becomes $\tilde{g}_\alpha\circ y_\alpha=g_\alpha|_{V_\alpha}$, which is a local diffeomorphism. Since we already know that y_α is a homeomorphism onto its image, it is an embedding.

Conversely, if y_α, $\alpha\in A$, are embeddings for some smooth structure on M, then any local chart on M_α is sent by y_α to a local chart on M. Thus, necessarily, we obtain the same charts on M as the ones we have just constructed. □

Remark 3.8. The smooth manifold M given in Lemma 3.7 *is not necessarily a Hausdorff space*. The definition of manifold in Bourbaki [2] does not require M to be a Hausdorff topological space, nor a paracompact space. These are, however, conditions most frequently required. It follows from Bourbaki [2] that M is Hausdorff if and only if for any two smooth charts $\varphi: U \subset M \to \mathbb{R}^n$ and $\psi: V \subset M \to \mathbb{R}^n$ constructed as in the proof of Lemma 3.7, we have that the graph of $\psi \circ \varphi^{-1}: \varphi(U \cap V) \to \psi(U \cap V)$ is closed in $\varphi(U) \times \psi(V) \subset \mathbb{R}^n \times \mathbb{R}^n$.

3.4. Symplectic glueing

Unlike in the previous two sections, we shall be assuming that the M_α, $\alpha \in A$, are Haudorff, paracompact smooth manifolds. Moreover, we will be assuming that there exist continuous, proper maps $F_\alpha: M \to \mathbb{R}^n$ which can be glued together to give rise to a proper map $F: M \to \mathbb{R}$. With the aid of F we will show that the Hausdorff and paracompactness properties of the M_α are inherited by M.

LEMMA 3.9. *If, for each $\alpha \in A$, M_α is symplectic with symplectic form ω_α, and if all $\varphi \in \mathcal{G}$ are symplectomorphisms* (*and $\mathcal{G}$ is free*), *then there exists a unique symplectic structure ω on M such that $y_\alpha^* \omega = \omega_\alpha$, $\alpha \in A$.*

Proof. The following facts hold:

(1) all the y_α's are embeddings;

(2) $\bigcup_{\alpha \in A} y_\alpha(M_\alpha) = M$;

(3) when $y_\alpha(M_\alpha)$ intersects $y_\beta(M_\beta)$, $\alpha \neq \beta$, then $y_\beta^{-1} \circ (y_\alpha) = \varphi$ for some $\varphi \in G$ with $\varphi^* \omega_\beta = \omega_\alpha$.

Then, the formula $y_\alpha^* \omega = \omega_\alpha$ defines a unique symplectic form ω on M. ☐

We can finally apply this technique in our case.

PROPOSITION 3.10. *Let $\{M_\alpha\}_{\alpha \in A}$ be a collection of symplectic manifolds, each equipped with a map $F_\alpha: M_\alpha \to \mathbb{R}^n$. For any $\alpha, \beta \in A$, let $D_{\alpha\beta} := F_\alpha(M_\alpha) \cap F_\beta(M_\beta)$ and assume that*

(1) *$U_\alpha := F_\alpha^{-1}(D_{\alpha\beta})$ and $U_\beta := F_\beta^{-1}(D_{\alpha\beta})$ are open;*

(2) *$\varphi_{\alpha\beta}: U_\alpha \to U_\beta$ is a symplectomorphism such that $\varphi_{\alpha\beta}^* F_\beta = F_\alpha$;*

(3) *if $D_{\alpha\beta\gamma} := F_\alpha(M_\alpha) \cap F_\beta(M_\beta) \cap F_\gamma(M_\gamma) \neq \varnothing$, we have $\varphi_{\beta\gamma} \circ \varphi_{\alpha\beta} = \varphi_{\alpha\gamma}$* (*restricted to $F_\alpha^{-1}(D_{\alpha\beta\gamma})$*).

Then the smooth manifold M obtained by glueing the collection $\{M_\alpha\}_{\alpha \in A}$ along the set $\{\varphi_{\alpha\beta}\}_{\alpha, \beta \in A}$ is symplectic, and there exists a unique map $F: M \to \mathbb{R}^n$ satisfying $F_\alpha = F \circ y_\alpha$, where $y_\alpha: M_\alpha \hookrightarrow M$, $\alpha \in A$, are the natural symplectic embeddings.

Proof. The third assumption (the cocycle condition) implies that the corresponding glueing groupoid is free. □

THEOREM 3.11. (Symplectic glueing) *Let* $\{M_\alpha\}_{\alpha\in A}$ *be a collection of symplectic manifolds, each equipped with a continuous, proper map* $F_\alpha: M_\alpha\to V_\alpha\subset\mathbb{R}^n$, *where* V_α *is open. For any* $\alpha,\beta\in A$, *let* $D_{\alpha\beta}:=V_\alpha\cap V_\beta$ *and assume that*

(1) $\varphi_{\alpha\beta}: F_\alpha^{-1}(D_{\alpha\beta})\to F_\beta^{-1}(D_{\alpha\beta})$ *is a symplectomorphism such that* $\varphi_{\alpha\beta}^*F_\beta=F_\alpha$;

(2) *when* $V_\alpha\cap V_\beta\cap V_\gamma\neq\varnothing$, *we have* $\varphi_{\beta\gamma}\circ\varphi_{\alpha\beta}=\varphi_{\alpha\gamma}$.

Then the smooth manifold M *obtained by glueing the collection* $\{M_\alpha\}_{\alpha\in A}$ *along the set* $\{\varphi_{\alpha\beta}\}_{\alpha,\beta\in A}$ *is Hausdorff, paracompact* (*in other words, a smooth manifold in the usual sense*) *and symplectic, and there exists a unique continuous, proper map*

$$F: M\longrightarrow \bigcup_{\alpha\in A} V_\alpha\subset\mathbb{R}^n$$

satisfying $F_\alpha=F\circ y_\alpha$, *where* $y_\alpha: M_\alpha\hookrightarrow M$, $\alpha\in A$, *are the natural symplectic embeddings.*

Proof. The main statement is a corollary of Proposition 3.10 since

$$F^{-1}(V_\alpha\cap V_\beta)=F^{-1}(F(M_\alpha)\cap F(M_\beta))$$

and thus the right-hand side is automatically open.

Next we show that M is Hausdorff. Let $\bar z,\bar w\in M$, where $z,w\in\bigsqcup_{\alpha\in A}M_\alpha$. There are two possibilities, that $F(\bar z)=F(\bar w)$ or that $F(\bar z)\neq F(\bar w)$. If $F(\bar z)=F(\bar w)$, then by the definition of F (i.e. $F_\alpha=F\circ y_\alpha$), there exists $\alpha\in A$ such that $z\in M_\alpha$ and $w\in M_\alpha$. Here we are viewing M_α as a subset of $\bigsqcup_{\alpha\in A}M_\alpha$, under the canonical identification y_α. Because M_α is Hausdorff, there exist open sets $U_z\subset M_\alpha$ and $U_w\subset M_\alpha$, with $z\in U_z$, $w\in U_w$ and $U_z\cap U_w=\varnothing$. Since M_α is open in $\bigsqcup_{\alpha\in A}M_\alpha$, by Lemma 3.6 we have that $\pi(U_z)$ and $\pi(U_w)$ are open subsets of M. By construction, $\bar z\in\pi(U_z)$ and $\bar w\in\pi(U_w)$. It follows from the definition of π as the quotient map $\bigsqcup_{\alpha\in A}M_\alpha\to M=\bigsqcup_{\alpha\in A}M_\alpha/\sim$ that

$$\pi(U_z)\cap\pi(U_w)=\pi(U_z\cap U_w)=\pi(\varnothing)=\varnothing.$$

On the other hand, suppose that $F(\bar z)\neq F(\bar w)$. Since $F(\bar z)\in\mathbb{R}^n$ and $F(\bar w)\in\mathbb{R}^n$, and since $\mathbb{R}^n$ is Hausdorff, there exist open sets W_z and W_w in $\mathbb{R}^n$ such that $F(\bar z)\in W_z$, $F(\bar w)\in W_w$ and $W_z\cap W_w=\varnothing$. As F is continuous, $F^{-1}(W_z)$ and $F^{-1}(W_w)$ are open. Also, by construction, $\bar z\in F^{-1}(W_z)$ and $\bar w\in F^{-1}(W_w)$. Of course,

$$F^{-1}(W_z)\cap F^{-1}(W_w)=F^{-1}(W_z\cap W_w)=\varnothing.$$

Let us show that F is proper. Let $V:=\bigcup_{\alpha\in A}V_\alpha$. Let $K\subset V$ be compact in V. Since K is compact, there exists a finite number of open balls B_i of radius $\varepsilon>0$ that cover K

and such that any $\overline{B}_i$ is included in some $V_{\alpha(i)}$, $\alpha(i)\in A$. Let $\{O_\beta\}_{\beta\in B}$ be an open cover of $F^{-1}(K)$. For any i, the set $\overline{B}_i$ is compact in $V_{\alpha(i)}$, and hence $F_\alpha^{-1}(\overline{B}_i)$ is compact in M_α. Thus $y_\alpha(F_\alpha^{-1}(\overline{B}_i))$ is compact in M, and hence there exists a finite subset $B_i\subset B$ such that $\bigcup_{\beta\in B_i} O_\beta \supset y_\alpha(F_\alpha^{-1}(\overline{B}_i))$. We can conclude, using the fact that

$$y_\alpha(F_\alpha^{-1}(U)) = F^{-1}(U) \quad \text{for all } U\subset V_\alpha, \tag{3.2}$$

that $F^{-1}(K)\subset\bigcup_i\bigcup_{\beta\in B_i} O_\beta$, which shows that $F^{-1}(K)$ is indeed compact.

To complete the properness proof we must show that equality (3.2) holds. Indeed, the inclusion $y_\alpha(F_\alpha^{-1}(U))\subset F^{-1}(U)$ follows directly from the equality $F\circ y_\alpha = F_\alpha$. For the converse, we come back to the definition of M. If $\bar z\in F^{-1}(U)$ there must exist some $z_\beta\in M_\beta$ such that $\pi(z_\beta)=\bar z$ (π is the quotient map of Lemma 3.5). Thus $F_\beta(z_\beta)=F(\bar z)$. This means that $V_\alpha\cap V_\beta$ is not empty, and there is a symplectomorphism $\varphi_{\beta\alpha}$ such that $z_\alpha:=\varphi_{\beta\alpha}(z_\beta)\in M_\alpha$. This implies that $\pi(z_\alpha)=\pi(z_\beta)=\bar z$. Thus $F(\bar z)=F_\alpha(z_\alpha)$, which proves that $F^{-1}(U)\subset y_\alpha(F_\alpha^{-1}(U))$.

We are left to show that M is a paracompact space. We have already shown that $F: M\to V$ is a proper map, so in particular, the fibers of F are compact. On the other hand, for each $\alpha\in A$, M_α is a manifold in the usual sense, and hence it is locally compact, which then implies that $\bigsqcup_{\alpha\in A} M_\alpha$ is locally compact. We claim that M is locally compact. Indeed, let $\bar z\in M$, where $z\in M_\alpha$ for some α. Because M_α is locally compact, there is a compact neighborhood K_z of z in M_α containing an open set U_z, with $z\in U_z$. Since π is continuous, $\pi(K_z)$ is compact. As π is open, $\pi(U_z)$ is open, and hence $\pi(K_z)$ is a compact neighborhood of $\bar z$, and we have shown that M is locally compact.

On the other hand, a continuous, proper map between locally compact Hausdorff spaces is closed([2]); see [4, Proposition 3, p. 16]. We have already shown that M is Hausdorff and locally compact. Hence, since $F: M\to V$ is a proper map, it is a also a closed map.

Next we deduce the paracompactness of M from the following result [20, §20, p. 153], [3, Theorem 1]: if $f: X\to Y$ is a continuous, closed surjective mapping between topological spaces with compact fibers, and Y is paracompact, then X is paracompact as well. We may apply this result with X equal to M, Y equal to $F(M)\subset\mathbb{R}^n$ and f equal to $F: M\to F(M)$. The map $F: M\to F(M)$ is continuous, closed, and it has compact fibers, and $F(M)$, as a subset of $\mathbb{R}^n$, is paracompact. Hence M is paracompact. This concludes the proof of the proposition. □

([2]) Let $f: X\to Y$ be such a map. Let A be closed and let $y\in\overline{F(A)}$. Since Y is Hausdorff, $\{y\}$ is the intersection of closed neighborhoods of y. Since Y is locally compact, we may assume that one of these neighborhoods is compact. As f is continuous and proper, $A\cap f^{-1}(y)$ is a decreasing intersection of non-empty closed subsets of a compact set, and hence is non-empty. Hence $y\in f(A)$ and $f(A)$ is closed.

4. Main theorem: statement

Again we equip $\mathbb{R}^2$ with its standard affine structure with origin at $(0,0)$, and orientation.

4.1. Delzant semitoric polygons

Let $\Delta \in \mathrm{Polyg}(\mathbb{R}^2)$ be a rational convex polygon in $\mathbb{R}^2$, as in Definition 2.1. Recall that in our terminology, Δ is not necessarily compact. We call a point in the boundary $\partial\Delta$ where the meeting edges are not colinear a *vertex of* Δ. We shall make the following assumption:

(a1) The intersection of Δ with a vertical line is compact (it may be empty).

Consider such a vertical line intersecting the polytope. If the intersection is not just a point, then it is a vertical segment. The top end of this segment is said to belong to the *top-boundary* of Δ.

With each vertex z of Δ we associate a couple $\mathcal{B}_z$ of primitive integral vectors starting at z and extending along the direction of the edges meeting at z, in the order that makes them oriented. Then $\mathcal{B}_z$ defines a $\mathbb{Z}$-basis of $\mathbb{Z}^2 \subset \mathbb{R}^2$ when, viewed as a 2×2 matrix, its determinant is equal to 1.

Let $s \in \mathbb{N}^*$ and let $(\lambda_1, ..., \lambda_s) \in \mathbb{R}^s$ with $\lambda_1 < ... < \lambda_s$. As before, ℓ_{λ_j} is the vertical line $\{(x,y) : x = \lambda_j\}$. We are interested only in the following case:

(a2) The vertical lines ℓ_{λ_j}, $j=1, ..., s$, intersect the top-boundary of Δ.

Let T be the linear transformation acting as the matrix

$$T := T^1 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}.$$

Definition 4.1. Let z be a vertex of the polygon Δ and $(u,v) = \mathcal{B}_z$. Then z is called

- a *Delzant corner* when there is no vertical line ℓ_{λ_j} through it and $\det(u,v)=1$;
- a *hidden Delzant corner* when there is a vertical line ℓ_{λ_j} through it, it belongs to the top-boundary and $\det(u,Tv)=1$;
- a *fake corner* when there is a vertical line ℓ_{λ_j} through it, it belongs to the top-boundary and $\det(u,Tv)=0$.

For the following lemma recall the definition of admissible weighted polygon, see Definition 2.4.

LEMMA 4.2. *Let Δ be a convex rational polygon equipped with a set of vertical lines $(\ell_{\lambda_1}, ..., \ell_{\lambda_s})$, such that the assumptions* (a1) *and* (a2) *are satisfied. Suppose also that*

- *any point in the top-boundary that belongs to some vertical line ℓ_{λ_j} is either a hidden Delzant corner or a fake corner*;
- *any other vertex of Δ is a Delzant corner.*

Then the triple $(\Delta, \{\ell_{\lambda_j}\}_{j=1}^s, (1, ..., 1))$ *is an admissible weighted polygon.*

Proof. We need to show that the convexity is preserved under the G_s-action. This amounts to show that, for any $j=1, ..., s$, the polygon $t_{\vec{e}_j}(\Delta)$ is convex, where $(\vec{e}_1, ..., \vec{e}_s)$ is the canonical basis of $\mathbb{Z}^s$. Since $t_{\vec{e}_j}$ is affine on both half-spaces delimited by the vertical line ℓ_{λ_j}, it suffices to show that $t_{\vec{e}_j}(\Delta)$ is locally convex near the points where ℓ_{λ_j} meets the boundary $\partial\Delta$.

We let $\{a, z\}=\ell_{\lambda_j}\cap\partial\Delta$ and assume that z lies on the top boundary. By assumption, z is either a hidden Delzant corner or a fake corner. Let us consider the vectors $(u, v)=\mathcal{B}_z$. Because z belongs to the top-boundary, the vector u must be directed to the left-hand side of z and v to the right-hand side. Since the transformation $t_{\vec{e}_j}$ acts only on the right half-space (and there it acts as T), the transformed edges of $t_{\vec{e}_j}(\Delta)$ at z are directed along (u, Tv). By assumption, $\det(u, Tv)$ is either 0 or 1, which implies the local convexity at z.

Now consider the "bottom boundary" at the point a. By assumption, the polygon is already locally convex at a (which means that $\det(u, v)\geqslant 0$), and a quick calculation shows that the action of $t_{\vec{e}_j}$ may only make it even "more" convex. □

It is easy to see that the properties of the lemma are preserved by the $\mathfrak{I}$-action. Thus we can state the following definition.

Definition 4.3. Let $[\Delta_{\mathrm{w}}]$ be a semitoric polygon as in Definition 2.5, and suppose that Δ_{w} is a representative of the form $(\Delta, \{\ell_{\lambda_j}\}_{j=1}^s, \{\varepsilon_j\}_{j=1}^s)$ with all the ε_j equal to 1. Then we say that $[\Delta_{\mathrm{w}}]$ is a *Delzant semitoric polygon* (of *complexity* s) if the polygon Δ equipped with the vertical lines ℓ_{λ_j} satisfies the hypothesis of Lemma 4.2.

We denote by $\mathcal{D}\mathrm{Polyg}_s(\mathbb{R}^2)\subset\mathcal{W}\mathrm{Polyg}_s(\mathbb{R}^2)/G_s\times\mathfrak{I}$ the space of Delzant semitoric polygons of complexity s, where $s<\infty$.

The following observation is a consequence of the construction of the homeomorphism f in §2.2.

Lemma 4.4. *The semitoric polygon in Definition* 2.8 (iii) *is a Delzant semitoric polygon.*

In addition, note also that for any representative Δ of the semitoric polygon $[\Delta_{\mathrm{w}}]$ in Definition 2.8, and for each $i\in\{1, ..., m_f\}$ as in Definition 2.8 (iv), the height h_i satisfies the inequality

$$0<h_i<\mathrm{length}(\Delta\cap\ell_i). \tag{4.1}$$

This is because, by (2.5), we have $h_i:=\mu(m_i)-\min_{s\in\ell_i\cap\Delta}\pi_2(s)$, where μ is a toric momentum map for the system $(M, \omega, (J, H))$ corresponding to Δ. Now, since $\mu(m_i)$ is a point in the interior of Δ, along the line ℓ_i, the expression (4.1) follows.

4.2. The main theorem

The following definition describes a collection of abstract ingredients. As we will see in the theorem following the definition, each such list of elements determines one and only one integrable system on a symplectic 4-manifold (which is not necessarily a compact manifold, but we can characterize precisely when it is, in terms of one of the ingredients of the list). Moreover, this integrable system is of semitoric type.

In the definition, the term $\mathbb{R}[[X,Y]]$ refers to the algebra of real formal power series in two variables, and $\mathbb{R}[[X,Y]]_0$ is the subspace of such series with vanishing constant term, and first term $\sigma_1 X + \sigma_2 Y$ with $\sigma_2 \in [0, 2\pi)$.

Definition 4.5. A *semitoric list of ingredients* consists of the following items:
(i) An integer number $0 \leqslant m_f < \infty$.
(ii) An m_f-tuple of Taylor series $\{(S_i)^\infty\}_{i=1}^{m_f} \in (\mathbb{R}[[X,Y]]_0)^{m_f}$.
(iii) A Delzant semitoric polygon $[\Delta_{\mathrm{w}}]$ of complexity m_f, as in Definition 4.3.
We denote the representative Δ_{w} of $[\Delta_{\mathrm{w}}]$ by $(\Delta, \{\ell_{\lambda_j}\}_{j=1}^{m_f}, \{\varepsilon_j\}_{j=1}^{m_f})$.
(iv) An m_f-tuple of numbers $\{h_j\}_{j=1}^{m_f}$ with $0 < h_j < \mathrm{length}(\Delta \cap \ell_i)$ for $j \in \{1, ..., m_f\}$.
(v) A $(G_{m_f} \times \mathfrak{I})$-orbit of $(\Delta_{\mathrm{w}}, \{k_j\}_{j=1}^{m_f})$, where $\{k_j\}_{j=1}^{m_f}$ is a collection of integers.

Now we are ready to state the main theorem, the proof of which is contructive and, in view of §2 and Lemma 4.4, gives a recipe to construct all semitoric integrable systems up to isomorphisms.

THEOREM 4.6. *For each semitoric list of ingredients, as in Definition* 4.5, *there exists a* 4*-dimensional simple semitoric integrable system* $(M, \omega, (J,H))$, *such that the list of invariants* (i)–(v) *of* $(M, \omega, (J,H))$ *as in Definition* 2.8 *is equal to this list of ingredients. Moreover*, M *is compact if and only if the polygon in* (iii) *is compact.*

4.3. Classification of 4-dimensional semitoric systems

Consequently, putting Theorem 4.6 together with Theorem 2.9 proved in [16], we obtain the classification of integrable systems in symplectic 4-manifolds.

THEOREM 4.7. (Classification of 4-dimensional semitoric integrable systems) *For each semitoric list of ingredients, as in Definition* 4.5, *there exists a* 4*-dimensional simple semitoric integrable system with list of invariants equal to this list of ingredients* (*cf. Definition* 2.8). *Moreover, two* 4*-dimensional simple semitoric integrable systems are isomorphic if and only if they are constructed from the same list of ingredients.*

5. Proof of the main theorem

Let $(\Delta, \{\ell_{\lambda_j}\}_{j=1}^s, \{\varepsilon_j\}_{j=1}^s)$ be a representative of $[\Delta_{\rm w}]$ with all the ε_j's equal to 1. The strategy is to use the glueing procedure of §3 in order to obtain a semitoric system by constructing a suitable singular torus fibration above $\Delta \subset \mathbb{R}^2$.

For $j=1, ..., m_f$, let $c_j \in \mathbb{R}^2$ be the point with coordinates

$$c_j = (\lambda_j, h_j + \min \pi_2(\Delta \cap \ell_{\lambda_j})). \tag{5.1}$$

Because of the assumption on h_j, all points c_j lie in the interior of the polygon Δ. We call these points *nodes*. We denote by ℓ_j^+ the vertical half-line through c_j pointing upwards. We call these half-lines *cuts*.

We have divided the proof of the theorem in a preliminary step, three intermediate steps and a conclusive step. In the preliminary step we construct a convenient covering of the polygon Δ.

Then we proceed as follows. First we construct a "semitoric system" over the part of the polygon away from the sets in the covering that contain the cuts ℓ_j^+; then we attach to this "semitoric system" the focus-focus fibrations, i.e. the models for the systems in a small neighborhood of the nodes. Third, we continue to glue the local models in a small neighborhood of the cuts. The "semitoric system" is given by a proper toric map only in the preimage of the polygon away from the cuts. We use the results of §3 as a stepping stone throughout.

Finally we recover the smoothness of the system and observe that the invariants of the system are precisely the ingredients we started with.

Preliminary stage: a convenient covering.

We construct an open cover of the polygon. Because of the discreteness of the set of vertices of the polygon, and the local compactness of $\mathbb{R}^2$, we can find an open cover $\{\Omega_\alpha\}_{\alpha \in A}$ of Δ such that the following three properties hold: there exists $\varrho > 0$ such that all the Ω_α's are integral-affine images of the open cube $C := I^2$ with $I =: (-\varrho, \varrho)$, i.e. for every $\alpha \in A$ there exists $R_\alpha \in \mathrm{Aff}(2, \mathbb{Z})$ such that $\Omega_\alpha = R_\alpha(C)$; each vertex of the polygon, and each node, is contained in only one open set Ω_α; two open sets containing a vertex or a node never intersect with each other. In fact, if

$$C_{\rm e} := \{(x,y) \in C : y \geqslant 0\} \quad \text{and} \quad C_{\rm ee} := \{(x,y) \in C : x \geqslant 0 \text{ and } y \geqslant 0\},$$

one may assume that, for any $\alpha \in A$, (1) if Ω_α intersects $\partial\Delta$ but does not contain any vertex then

$$\Omega_\alpha \cap \Delta = R_\alpha(C_{\rm e}),$$

and (2) if Ω_α contains a Delzant corner, then

$$\Omega_\alpha \cap \Delta = R_\alpha(C_{\mathrm{ee}}).$$

The first case holds since along any edge one can find a primitive vector, and complete it to a $\mathbb{Z}$-basis of $\mathbb{Z}^2$. It remains to compose by a suitable translation to position the image of C_{e} at the right place. The second case is similar, since at a Delzant corner the primitive vectors of the meeting edges form a $\mathbb{Z}$-basis of $\mathbb{Z}^2$; cf. Definition 4.1.

First stage: away from the cuts.

Let $A' \subset A$ be the subset obtained by removing all indices intersecting the cuts. We construct a semitoric system above $\bigcup_{\alpha\in A'} \Omega_\alpha$, by glueing the following local models. Let $\mathbb{D}$ be the open disk in $T^*\mathbb{R}=\mathbb{R}^2$ of radius $\sqrt{2\varrho}$, centered at the origin. Consider the following models: the *regular model*

$$M_{\mathrm{r}} := \mathbb{T}^2 \times C \subset T^*\mathbb{T}^2,$$

with momentum map

$$F_{\mathrm{r}}(x_1, x_2, \xi_1, \xi_2) := (\xi_1, \xi_2);$$

the *tranversally elliptic model* $M_{\mathrm{e}} := (\mathbb{T}^1 \times I) \times \mathbb{D} \subset T^*\mathbb{T}^1 \times T^*\mathbb{R}$, with momentum map

$$F_{\mathrm{e}}(x_1, \xi_1, x_2, \xi_2) := \big(\xi_1, \tfrac{1}{2}(x_2^2+\xi_2^2)\big);$$

and the *elliptic-elliptic model* $M_{\mathrm{ee}} := \mathbb{D}\times\mathbb{D} \subset T^*\mathbb{R}\times T^*\mathbb{R}$, with momentum map

$$F_{\mathrm{ee}}(x_1, \xi_1, x_2, \xi_2) := \big(\tfrac{1}{2}(x_1^2+\xi_1^2), \tfrac{1}{2}(x_2^2+\xi_2^2)\big).$$

Observe that $F_{\mathrm{r}}(M_{\mathrm{r}})=C$, $F_{\mathrm{e}}(M_{\mathrm{e}})=C_{\mathrm{e}}$ and $F_{\mathrm{ee}}(M_{\mathrm{ee}})=C_{\mathrm{ee}}$. Notice also that these models are all toric, in the sense that the momentum maps generate an effective Hamiltonian $\mathbb{T}^2$-action. What is more, these momentum maps are proper for the topology induced on their images.

Given any Ω_α, $\alpha\in A'$, we obtain a (singular) proper Lagrangian momentum map over Ω_α, whose image is precisely $\Omega_\alpha\cap\Delta$, and which defines an effective Hamiltonian $\mathbb{T}^2$-action, by the following simple rule:

(a) *If Ω_α contains no boundary point of Δ and no node*, then we choose $M_\alpha := M_{\mathrm{r}}$, with momentum map $F_\alpha := R_\alpha \circ F_{\mathrm{r}}$. Then F_α is a regular Lagrangian torus fibration.

(b) *If Ω_α intersects the boundary $\partial\Delta$ but does not contain any vertex of Δ*, we choose $M_\alpha := M_{\mathrm{e}}$, with momentum map $F_\alpha := R_\alpha \circ F_{\mathrm{e}}$. The set of singular values of F_α is the bounded open line segment $R_\alpha(\{(x,y)\in C_{\mathrm{e}}:y=0\})$. The two components J_α and H_α of the momentum map F_α have colinear differentials on the preimage

$$U:=F_{\mathrm{e}}^{-1}(\{(x,y)\in C_{\mathrm{e}}:y=0\})$$

under F_α of the set of singular values: if $(a,b)\in\mathbb{R}^2$ is a vector colinear to this segment, then we have $a\,dH_\alpha=b\,dJ_\alpha$ on U.

(c) *If* Ω_α *contains exactly one Delzant corner of* Δ, we choose $M_\alpha:=M_{\mathrm{ee}}$, with momentum map $F_\alpha:=R_\alpha\circ F_{\mathrm{ee}}$. Then the set of singular values of F_α is the "corner" $R_\alpha(\partial C_{\mathrm{ee}})$. If $\{(a,b),(c,d)\}$ is an oriented $\mathbb{Z}$-basis of $\mathbb{R}^2$, where (a,b) and (c,d) are directed along the respective directions of the edges of this corner, then the first component of F_α is

$$J_\alpha=\tfrac{1}{2}a(x_1^2+\xi_1^2)+\tfrac{1}{2}c(x_2^2+\xi_2^2).$$

In other words, the numbers a and c are the isotropy weights of the S^1-action induced by J_α. This is useful to notice, since this component J_α of F_α will not be modified, near the Delzant corner, by the rest of the construction process.

We describe now the transition functions: when $\Delta_{\alpha\beta}:=\Omega_\alpha\cap\Omega_\beta\neq\varnothing$, we want to define a symplectomorphism

$$\varphi_{\alpha\beta}\colon F_\alpha^{-1}(\Delta_{\alpha\beta})\longrightarrow F_\beta^{-1}(\Delta_{\alpha\beta})\quad\text{such that } \varphi_{\alpha\beta}^*F_\beta=F_\alpha. \tag{5.2}$$

For this we use the following notation: when $R\in\mathrm{Aff}(2,\mathbb{Z})$, we denote by $\widetilde{R}$ the symplectomorphism $\widetilde{R}\colon\mathbb{T}^2\times\mathbb{R}^2(=T^*\mathbb{T}^2)\to\mathbb{T}^2\times\mathbb{R}^2$ given by $(x,\xi)\mapsto(({}^t dR)^{-1}x,R\xi)$, where dR is the linear part of R. Remark that $\xi\circ\widetilde{R}=R\circ\xi$.

Case 1. If both F_α and F_β are regular models, we let

$$\varphi_{\alpha\beta}:=\widetilde{R}_\beta^{-1}\circ\widetilde{R}_\alpha. \tag{5.3}$$

Then $F_\beta\circ\varphi_{\alpha\beta}=R_\beta\circ F_{\mathrm{r}}\circ\varphi_{\alpha\beta}=F_{\mathrm{r}}\circ\widetilde{R}_\beta\circ\varphi_{\alpha\beta}=F_{\mathrm{r}}\circ\widetilde{R}_\alpha=F_\alpha$, i.e. (5.2) holds.

Case 2. If F_α is regular and F_β is transversally elliptic, we introduce the symplectomorphism (symplectic polar coordinates)

$$\begin{aligned}\varphi_{\mathrm{re}}:M_{\mathrm{r}}\cap(\mathbb{T}^1\times\mathbb{R})\times(\mathbb{T}^1\times\mathbb{R}_+^*)&\longrightarrow(\mathbb{T}^1\times\mathbb{R})\times(\mathbb{R}^2\setminus\{0\})\cap M_{\mathrm{e}}\\ (x_1,\xi_1,x_2,\xi_2)&\longmapsto(x_1,\xi_1,\sqrt{2\xi_2}\cos x_2,-\sqrt{2\xi_2}\sin x_2).\end{aligned}$$

Notice that $\varphi_{\mathrm{re}}^*F_{\mathrm{e}}=F_{\mathrm{r}}$. Thus we can define

$$\varphi_{\alpha\beta}:=\varphi_{\mathrm{re}}\circ\widetilde{R}_\beta^{-1}\circ\widetilde{R}_\alpha. \tag{5.4}$$

We have $F_\beta\circ\varphi_{\alpha\beta}=R_\beta\circ F_{\mathrm{e}}\circ\varphi_{\mathrm{re}}\circ\widetilde{R}_\beta^{-1}\widetilde{R}_\alpha=R_\beta\circ F_{\mathrm{r}}\circ\widetilde{R}_\beta^{-1}\widetilde{R}_\alpha=F_{\mathrm{r}}\circ\widetilde{R}_\alpha=F_\alpha$, i.e. (5.2) holds.

Case 3. Similarly, if F_α is regular and F_β is elliptic-elliptic, we introduce the symplectomorphism

$$\begin{aligned}\varphi_{\mathrm{ree}}: M_{\mathrm{r}}\cap(\mathbb{T}^1\times\mathbb{R}_+^*)\times(\mathbb{T}^1\times\mathbb{R}_+^*)&\longrightarrow(\mathbb{R}^2\setminus\{0\})\times(\mathbb{R}^2\setminus\{0\})\cap M_{\mathrm{ee}}\\ (x_1,\xi_1,x_2,\xi_2)&\longmapsto\begin{pmatrix}\sqrt{2\xi_1}\cos x_1,-\sqrt{2\xi_1}\sin x_1\\ \sqrt{2\xi_2}\cos x_2,-\sqrt{2\xi_2}\sin x_2\end{pmatrix}.\end{aligned}$$

Again $\varphi_{\mathrm{ree}}^*F_{\mathrm{ee}}=F_{\mathrm{r}}$, and if we define

$$\varphi_{\alpha\beta}:=\varphi_{\mathrm{ree}}\circ\widetilde{R}_\beta^{-1}\circ\widetilde{R}_\alpha, \tag{5.5}$$

then (5.2) holds.

Case 4. If both F_α and F_β are transversally elliptic models, then we have that the affine map $R_{\alpha\beta}:=R_\beta^{-1}R_\alpha$ is an oriented transformation that preserves the upper half-plane. Thus the horizontal axis is globally preserved, and the vector $e_1=(1,0)$ is an eigenvector of $dR_{\alpha\beta}$. Since $dR_{\alpha\beta}\in\mathrm{SL}(2,\mathbb{Z})$ it is of the form

$$T_k:=\begin{pmatrix}1&k\\0&1\end{pmatrix}$$

for some $k\in\mathbb{Z}$. Therefore, $R_{\alpha\beta}=\tau_u\circ T_k$, where τ_u is the translation by a horizontal vector $u=(u_1,0)$. Consider the symplectomorphism $\overline{R}_{\alpha\beta}(x_1,\xi_1,x_2,\xi_2):=(x_1',\xi_1',x_2',\xi_2')$ of $T^*\mathbb{T}^1\times T^*\mathbb{R}$ given by

$$\begin{cases}x_1'=x_1,\\ \xi_1'=\xi_1+\frac{1}{2}k(x_2^2+\xi_2^2)+u_1,\\ (x_2'+i\xi_2')=e^{ikx_1}(x_2+i\xi_2).\end{cases}$$

Observe that $F_{\mathrm{e}}\circ\overline{R}_{\alpha\beta}=R_{\alpha\beta}\circ F_{\mathrm{e}}$. Now we define

$$\varphi_{\alpha\beta}:=\overline{R}_{\alpha\beta}|_{F_\alpha^{-1}(\Delta_{\alpha\beta})}, \tag{5.6}$$

and we verify that $F_\beta\circ\overline{R}_{\alpha\beta}=R_\beta\circ F_{\mathrm{e}}\circ\overline{R}_{\alpha\beta}=R_\beta\circ R_{\alpha\beta}\circ F_{\mathrm{e}}=R_\alpha\circ F_{\mathrm{e}}=F_\alpha$, and hence (5.2) holds.

Case 5. If F_α is a transversally elliptic model, while F_β is elliptic-elliptic, then, as in the previous case, the intersection $\Delta_{\alpha\beta}$ contains a portion of an edge, but not the vertex itself. This edge is mapped by R_β from either the horizontal or the vertical positive axis. Suppose for simplicity that it is the horizontal axis. As before, the affine map $R_{\alpha\beta}$ defined in Case 4 is an oriented transformation that preserves the upper half-plane, and thus one can construct a symplectomorphism $\bar{R}_{\alpha\beta}$ of $T^*\mathbb{T}^1\times T^*\mathbb{R}$ such that $F_{\mathrm{e}}\circ\bar{R}_{\alpha\beta}=R_{\alpha\beta}\circ F_{\mathrm{e}}$. Introduce the symplectomorphism

$$\begin{aligned}\varphi_{\mathrm{eee}}\colon M_{\mathrm{e}}\cap(\mathbb{T}^1\times\mathbb{R}^*_+)\times\mathbb{R}^2 &\longrightarrow (\mathbb{R}^2\setminus\{0\})\times\mathbb{R}^2\cap M_{\mathrm{ee}},\\ (x_1,\xi_1,x_2,\xi_2) &\longmapsto (\sqrt{2\xi_1}\cos x_1, -\sqrt{2\xi_1}\sin x_1, x_2, \xi_2).\end{aligned}$$

Notice that $F_{\mathrm{ee}}\circ\varphi_{\mathrm{eee}}=F_{\mathrm{e}}$ and, whenever both are defined, $\varphi_{\mathrm{eee}}=\varphi_{\mathrm{ree}}\circ\varphi_{\mathrm{re}}^{-1}$. We define

$$\varphi_{\alpha\beta}:=\varphi_{\mathrm{eee}}\circ\bar{R}_{\alpha\beta}, \tag{5.7}$$

and verify now routinely that $F_\beta\circ\varphi_{\alpha\beta}=F_\alpha$, i.e. (5.2) also holds in this case.

We have defined the transition maps $\varphi_{\alpha\beta}$ in the five cases (5.3)–(5.7), and verified that equation (5.2) holds for each of them. In fact one should also mention that for the non-symmetric cases (5.4), (5.5) and (5.7), we let $\varphi_{\beta\alpha}:=\varphi_{\alpha\beta}^{-1}$ (this is automatic for the symmetric cases (5.3) and (5.6)). Then it is easy to verify that the cocycle condition is fulfilled. Namely, when the triple intersection $\Omega_{\alpha\beta}\cap\Omega_{\beta\gamma}\cap\Omega_{\gamma\alpha}$ is non-empty, then

$$\varphi_{\gamma\alpha}\circ\varphi_{\beta\gamma}\circ\varphi_{\alpha\beta}=\mathrm{Id}.$$

Thus we can apply the glueing construction, cf. Theorem 3.11, and obtain a symplectic manifold $M_{A'}$ with a surjective map

$$F_{A'}\colon M_{A'}\longrightarrow\bigcup_{\alpha\in A'}\Omega_\alpha\subset\mathbb{R}^2,$$

and, for each $\alpha\in A'\subset A$, there is a symplectic embedding $\iota_\alpha\colon M_\alpha\hookrightarrow M_{A'}$ such that $\iota_\alpha^*F_{A'}=F_\alpha$. Since all F_α are proper smooth toric momentum maps, so is $F_{A'}$.

Second stage: attaching focus-focus fibrations.

Fix an integer i, with $1\leqslant i\leqslant m_f$. Using the classification result of [18], one can construct a focus-focus model associated with an arbitrary Taylor series invariant. Precisely, for each node c_i, there exists a symplectic manifold M_i equipped with a smooth map $F_i\colon M_i\to C$ such that the symplectic invariant of the induced singular foliation is precisely the Taylor series S^∞. Using the result of [19], one can construct a continuous map $\mu_i\colon M_i\to D_i$, where $D_i\subset\mathbb{R}^2$ is some simply connected open set around the origin, that is a smooth proper toric momentum map outside $\mu_i^{-1}(\ell)$, where $\ell:=\{(0,y):y\geqslant 0\}$. In fact $\mu_i=g_i\circ F_i$, for some homeomorphism $g_i\colon C\to D_i$ which is smooth outside ℓ, and which preserves the first component: it is of the form

$$g_i(x,y)=(x,f_i(x,y)).$$

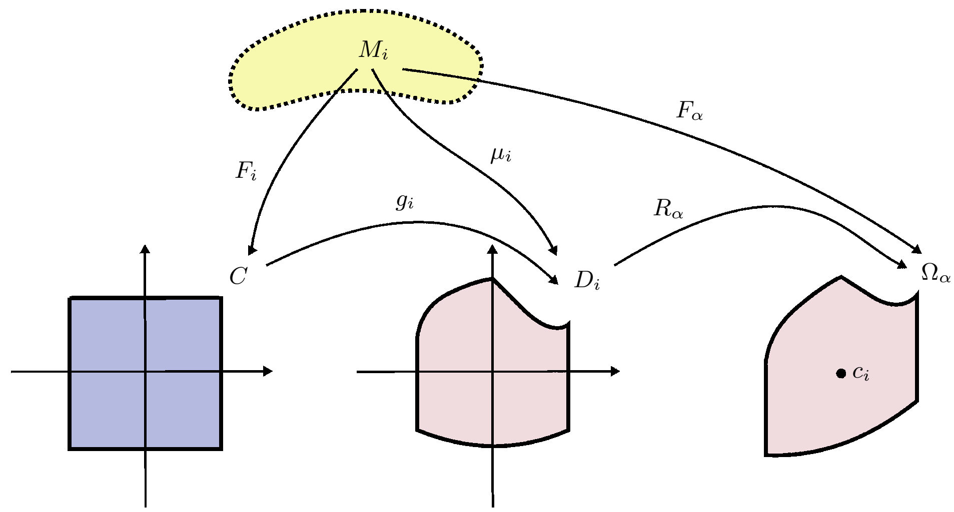

Figure 5.1. The pieces M_i and the chart diagrams for F_α, F_i, g_i and R_α.

This construction depends on the choice of a local toric momentum map for the fibration over $C\setminus\ell$. Here we choose the privileged momentum map as defined in §2.4. We are now in position to add to the index set A' all the indices $\alpha\in A$ corresponding to the nodes, and thus defining a new index set A''. If Ω_α contains the node c_i, we let R_α be the matrix T_{k_i} left-composed by the translation from the origin to the node c_i. Here k_j is the integer given as ingredient (v) in the list. We may assume that $\Omega_\alpha = R_\alpha(D_i)$. Then we choose $M_\alpha := M_i$ with momentum map $F_\alpha := R_\alpha \circ \mu_i$.

By making ϱ small enough, one may assume that all Ω_β, $\beta\in A'$, intersecting an open set Ω_α containing a node carry regular models. Thus we need to define transition functions between a regular model and a focus-focus model. On $\Delta_{\alpha\beta} := \Omega_\alpha \cap \Omega_\beta$, both momentum maps F_α and F_β are regular. Contrary to all the previous cases, the focus-focus model F_α is not explicit, and we cannot simply provide an elementary formula for the transition map $\varphi_{\alpha\beta}$. However, since $C\setminus\ell$ is simply connected and a set of regular values of F_i, we can invoke the Liouville–Mineur–Arnold action-angle theorem and assert that there exists a symplectomorphism

$$\varphi_i: F_i^{-1}(C\setminus\ell) \longrightarrow \mathbb{T}^2 \times C' \subset T^*\mathbb{T}^2 = \{(x,\xi) \in \mathbb{T}^2 \times \mathbb{R}^2\}$$

such that

$$F_i = \varphi_i^*(h_i(\xi)) \quad \text{for some diffeomorphism } h_i: C' \longrightarrow C\setminus\ell.$$

Then $\mu_i = \varphi_i^*(g_i \circ h_i(\xi))$. Since both μ_i and ξ are toric momentum maps for the same foliation, there exists a transformation $H_i \in \mathrm{Aff}(2,\mathbb{Z})$ such that $g_i \circ h_i = H_i$.

Thus, if F_α is focus-focus and F_β is regular, we introduce the symplectomorphism

$$\varphi_{\alpha\beta} := \widetilde{R}_\beta^{-1} \circ \widetilde{R}_\alpha \circ \widetilde{H}_i \circ \varphi_i: F_\alpha^{-1}(\Delta_{\alpha\beta}) \longrightarrow F_\beta^{-1}(\Delta_{\alpha\beta}). \tag{5.8}$$

We verify that

$$F_\beta \circ \varphi_{\alpha\beta} = F_r \circ \widetilde{R}_\beta \circ \varphi_{\alpha\beta} = R_\alpha \circ H_i \circ F_r \circ \varphi_i = R_\alpha \circ \mu_i = F_\alpha,$$

so we have shown (5.2).

We can now include these nodal pieces in the symplectic glueing construction using Theorem 3.11, which defines a symplectic manifold $M_{A''}$ and a proper map

$$F_{A''}: M_{A''} \longrightarrow \bigcup_{\alpha \in A''} \Omega_\alpha \subset \mathbb{R}^2.$$

However $F_{A''}$ is not smooth everywhere, but it is a smooth toric momentum map outside the preimages of the cuts ℓ_j^+, $j=1, ..., m_f$.

Third stage: filling in the gaps.

Here we add the open sets Ω_α that were covering the cuts ℓ_i by switching these lines on the other side. Let $t_i := t_{\ell_{\lambda_i}}$ as in §2.2. The cut ℓ_i^+ is invariant under t_i. The open sets $t_i(\Omega_\alpha)$, $\alpha \in A \setminus A''$, form a cover of $\ell_i \cap t_i(\Delta)$. Within the geometry of the new polygon $t_i(\Delta)$, each of these open sets can be associated with either a regular model, a transversally elliptic model, or an elliptic-elliptic model (indeed, under the transformation t_i, a fake corner disappears, and a hidden Delzant corner unhides itself).

Thus we can add these to our glueing data, which amounts to equip each such open set Ω_α with the model $(M_\alpha, t_i^{-1} \circ F_\alpha)$, where (M_α, F_α) is determined as before, but for the transformed polygon $t_i(\Delta)$.

The transition maps are defined with the same formulas as before, taking into account that the map R_α is now a piecewise affine transformation. The cocycle conditions remain valid as well.

Doing this for all indices i, because all the F_α's are continuous and proper, by Theorem 3.11, we obtain a smooth symplectic manifold $M = M_A$ equipped with a proper, continuous map $\mu = F_A$,

$$\mu: M \longrightarrow \bigcup_{\alpha \in A} \Omega_\alpha \subset \mathbb{R}^2, \tag{5.9}$$

whose image is precisely Δ.

However, the map μ is a proper toric momentum map only outside the cuts ℓ_i. In other words, μ fails to be smooth along the cuts ℓ_i. (Note that in the symplectic glueing construction, Theorem 3.11, we did not make any smoothness assumption on the F_α, nor made any conclusion on the smoothness of F.)

Fourth and final stage: recovering smoothness.

In this step we compose the final momentum map μ in (5.9) on the left by a suitable homeomorphism in order to make it smooth. Let Ω_α be the open set containing the node c_i. Let $h_i = g_i^{-1}: D_i \to C$. The map h_i is a bilipschitz homeomorphism fixing the origin and a smooth diffeomorphism outside the positive vertical axis. It is of the form

$$h_i(x,y) = (x, \eta_i(x,y)).$$

Since h_i is orientation preserving, $(\partial\eta_i/\partial y)(x,y) > 0$ for all $(x,y) \in D_i$. Let $\delta_i > 0$ be such that $[-2\delta_i, 2\delta_i]^2 \subset D_i$ and consider the vertical half-strip $\mathcal{S}_{\delta_i} := [-\delta_i, \delta_i] \times [-\delta_i, \infty)$.

Claim 5.1. *There exists a function* $\tilde{\eta}_i: D_i \to C$ *such that*

(1) $\tilde{\eta}_i(x,y) = \eta_i(x,y)$ *for all* $(x,y) \in D_i \cap \mathcal{S}_{\delta_i}$;

(2) $\tilde{\eta}_i(x,y) = y$ *for all* $(x,y) \in D_i \setminus \mathcal{S}_{2\delta_i}$;

(3) $(\partial\tilde{\eta}/\partial y)(x,y) > 0$ *for all* $(x,y) \in D_i$.

In order to show this recall that if $f: A \to \mathbb{R}$ is smooth and $A \subset U \subset \mathbb{R}^2$ is closed, then f has a smooth extension to $\tilde{f}: U \to \mathbb{R}$, where U is open, see for example [21, Lemma 5.58 and the remark below it]. Let us apply this fact in our situation. Let

$$A_{\delta_i} := (D_i \cap \mathcal{S}_{\delta_i}) \cup (D_i \setminus \mathrm{Int}(\mathcal{S}_{3\delta_i/2})),$$

which is a closed subset of $D_i \subset \mathbb{R}^2$, and let $\hat{\eta}_i: A_{\delta_i} \to \mathbb{R}$ be the smooth function given by

$$\hat{\eta}_i(x,y) = \begin{cases} \eta_i(x,y), & \text{if } (x,y) \in D_i \cap \mathcal{S}_{\delta_i}, \\ y, & \text{if } (x,y) \in D_i \setminus \mathrm{Int}(\mathcal{S}_{3\delta_i/2}). \end{cases} \tag{5.10}$$

Because $A_{\delta_i} \subset D_i$, and D_i is bounded, there exists a constant $0 < c_i < 1$ such that $\partial\eta_i/\partial y > c$ on A_{δ_i} and hence $\partial\hat{\eta}_i/\partial y > c_i$ on A_{δ_i}. Let $\zeta_i := \partial\hat{\eta}_i/\partial y - c_i: A_{\delta_i} \to \mathbb{R}$, which by assumption is strictly positive. By the above fact, ζ_i extends to a smooth function $G_i: D_i \to \mathbb{R}$. Because the proof of this fact preserves non-negativity, and $\zeta_i > 0$, we have that $G_i \geqslant 0$. By possibly shrinking the size of D_i, we may assume that D_i is a disk of radius $r_i > 0$ centered at the origin. Let

$$X_{\delta_i} := \left[-r_i, -\tfrac{3}{2}\delta_i\right] \cup \left[\tfrac{3}{2}\delta_i, r_i\right], \quad Y_{\delta_i} := [-\delta_i, \delta_i], \quad Z_{\delta_i} := \left[-\tfrac{3}{2}\delta_i, -\delta_i\right] \cup \left[\delta_i, \tfrac{3}{2}\delta_i\right],$$

and let $\nu_1^i: X_{\delta_i} \to \mathbb{R}$ and $\nu_2^i: Y_{\delta_i} \to \mathbb{R}$ be the functions given by $\nu_1^i(x) := -\hat{\eta}_i(x,0)$ and

$$\nu_2^i(x) := \hat{\eta}_i\left(x, -\frac{3\delta_i}{2}\right) - \int_0^{-3\delta_i/2} (G_i(x,t) + c_i)\, dt,$$

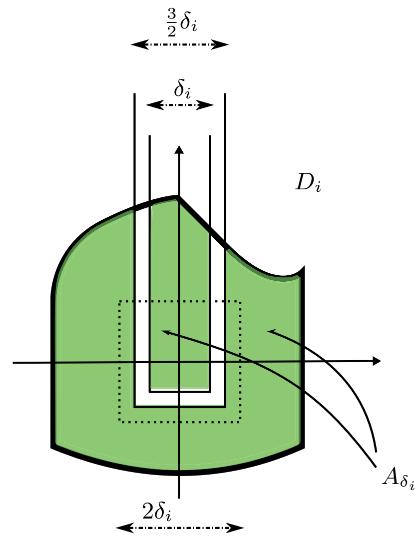

Figure 5.2. The set $A_{\delta_i} := (D_i \cap \mathcal{S}_{\delta_i}) \cup (D_i \setminus \mathrm{Int}(\mathcal{S}_{3\delta_i/2}))$, on which $\hat{\eta}_i$ is defined.

where we are using the convention $\int_a^b h = -\int_b^a h$ when $a > b$. Because $\hat{\eta}_i$ and G_i are smooth functions, ν_1^i and ν_2^i are also smooth. Let $\beta^i : [-r_i, r_i] \to \mathbb{R}$ be a smooth extension of the function $X_{\delta_i} \cup Y_{\delta_i} \to \mathbb{R}$ defined by ν_1^i on X_{δ_i} and by ν_2^i on Y_{δ_i}, which again exists by a partition of unity argument.

Consider the function $\tilde{\eta}_i : D_i \to \mathbb{R}$ given by

$$\tilde{\eta}_i(x, y) := \beta^i(x) + \int_0^y (G_i(x, t) + c_i)\, dt.$$

Because β is a smooth extension of ν_1^i and ν_2^i, and G is smooth, $\tilde{\eta}_i$ is smooth. We claim that $\tilde{\eta}_i|_{A_{\delta_i}}(x, y) = \hat{\eta}_i(x, y)$ if $(x, y) \in A_{\delta_i}$. First assume that $x \in Y_{\delta_i}$, and moreover that $-r_i \leqslant y \leqslant -\frac{3}{2}\delta_i$. Because G_i is an extension of g_i we have that

$$\tilde{\eta}_i|_{A_{\delta_i}}(x, y) = \nu_2^i(x) + \int_0^{-3\delta_i/2} (G_i(x, t) + c_i)\, dt + \int_{-3\delta_i/2}^y \frac{\partial \hat{\eta}_i}{\partial y}(x, t)\, dt,$$

and hence, by the fundamental theorem of calculus, and using the definition of ν_2^i, we obtain that

$$\tilde{\eta}_i|_{A_{\delta_i}}(x, y) = \nu_2^i(x) + \int_0^{-3\delta_i/2} (G_i(x, t) + c_i)\, dt + \hat{\eta}_i(x, y) - \hat{\eta}_i\left(x, -\frac{3\delta_i}{2}\right) = \hat{\eta}_i(x, y). \quad (5.11)$$

The remaining subcases within the case of $x \in Y_{\delta_i}$ are when $-\delta_i \leqslant y \leqslant 0$, which follows by the same reasoning as in (a) using the formula for ν_1^i instead of ν_2^i, the case where

$0\leqslant y\leqslant r_i$, which is trivial because the extension is defined by the original function therein, and the case where $-\frac{3}{2}\delta_i\leqslant y\leqslant -\delta_i$, in which $(x,y)\notin A_{\delta_i}$, so there is nothing to prove. The case where $x\in X_{\delta_i}$ follows by the same type of argument as the case of Y_{δ_i}. The case of $x\in Z_{\delta_i}$ is immediate because the extension is defined by the original function therein.

Applying again the fundamental theorem of calculus, because the functions ν_1^i, ν_2^i and β^i do not depend on y, we have that

$$\frac{\partial\tilde{\eta}_i}{\partial y}=G_i+c_i, \tag{5.12}$$

which is strictly positive, since $G_i\geqslant 0$ and $c_i>0$. Because (5.11) and (5.12) hold, we in turn have, in view of the definition (5.10) of $\hat{\eta}$, that properties (1), (2) and (3) are satisfied. This concludes the proof of Claim 5.1.

Let $\Omega_i:=D_i\cup\{(x,y):y<2\delta_i\}$. Because of properties (1), (2) and (3) of $\tilde{\eta}_i$, the map

$$\tilde{h}_i\colon (x,y)\longmapsto (x,\tilde{\eta}_i(x,y))$$

coincides with h_i in $\mathcal{S}_{\delta_i}$, while it is equal to the identity outside $\mathcal{S}_{2\delta_i}$. Thus we can extend it to Ω_i by letting it be the identity outside $D_i\cup\mathcal{S}_{2\delta_i}$. We call this extension $\tilde{h}_{\Omega_i}$. Consider the map

$$\check{h}_{\Omega_i}:=\tilde{h}_{\Omega_i}\circ t_0^{-1},$$

where t_0 is the piecewise affine map t_ℓ with ℓ being the positive vertical axis. In $t_0(\Omega\cap\mathcal{S}_{\delta_i})$, it is equal to $h_i\circ t_0^{-1}$, which is now smooth outside the negative vertical axis (this follows from [19, Theorem 3.8]; also from the fact that it is the homeomorphism that one obtains in the construction of the generalized momentum map $t_0\circ g_i\circ F_i=t_0\circ\mu_i$: this amounts to switching the cut downwards). Using the claim at the beginning of this step upside-down we can modify $\check{h}_{\Omega_i}$ in $\{(x,y)\in\Omega_i:y>\delta_i\}$ in such a way that we can then extend it to be smooth on $t_0(\{(x,y):y>\delta_i\})$. We obtain a homeomorphism of $\mathbb{R}^2$ that we denote $(\check{h}_{\mathbb{R}^2})_i$.

Define the map $\varphi_i\colon\mathbb{R}^2\to\mathbb{R}^2$ by

$$\varphi_i:=R_\alpha\circ(\check{h}_{\mathbb{R}^2})_i\circ t_0\circ R_\alpha^{-1}.$$

Because φ_i is a composition of homeomorphisms, it is a homeomorphism. Moreover, outside of $\mathcal{S}_{2\delta_i}$ we have that

$$\varphi_i=R_\alpha\circ(\check{h}_{\mathbb{R}^2})_i\circ t_0\circ R_\alpha^{-1}=R_\alpha\circ(\tilde{h}_{\Omega_i}\circ t_0^{-1})\circ t_0\circ R_\alpha^{-1},$$

and since $\tilde{h}_{\Omega_i}$ is the identity outside of $\mathcal{S}_{2\delta_i}$ we conclude that φ_i is the identity map outside $\mathcal{S}_{2\delta_i}$. Now let $\varphi\colon\mathbb{R}^2\to\mathbb{R}^2$ be the piecewise defined map

$$\varphi(x,y):=\begin{cases}\varphi_i(x,y), & \text{if } (x,y)\in\mathcal{S}_{2\delta_i},\\ (x,y), & \text{otherwise.}\end{cases} \tag{5.13}$$

Since each φ_i is a homeomorphism, and equal to the identity outside of $\mathcal{S}_{2\delta_i}$, the formula (5.13) defines a homeomorphism.

CLAIM 5.2. *The map* $\widetilde{F}: M \to \mathbb{R}^2$ *defined by* $\widetilde{F} := \varphi \circ \mu$ *is proper, and smooth everywhere.*

The properness claim is immediate, since φ is a homeomorphism and μ is proper.

In order to show that $\widetilde{F}$ is smooth, consider the map $\widetilde{F}_i: M \to \mathbb{R}^2$ defined as a composition $\widetilde{F}_i := \varphi_i \circ \mu$, where we recall that μ is the map (5.9). By the definition of φ, we have that $\widetilde{F}|_{\mathcal{S}_{\delta_i}} = \widetilde{F}_i$, and hence to prove the claim it suffices to show that each $\widetilde{F}_i$ is smooth. To prove this, we distinguish three cases.

Case 1. (In a neighborhood of c_i) In the neighborhood Ω_α of c_i sent by R_α^{-1} into $[-\delta_i, \delta_i]^2$, we have that

$$(\check{h}_{\mathbb{R}^2})_i \circ t_0 \circ R_\alpha^{-1} = \check{h}_{\Omega_i} \circ t_0 \circ R_\alpha^{-1} = \tilde{h}_{\Omega_i} \circ t_0^{-1} \circ t_0 \circ R_\alpha^{-1} = h_i \circ R_\alpha^{-1}.$$

Recall that $y_\alpha^* \mu = F_\alpha = R_\alpha \circ \mu_i$. Therefore one can write, in the preimage by μ of this neighborhood, $y_\alpha^*(\widetilde{F}_i) = y_\alpha^*(h_i \circ \mu_i) = F_i$. Since F_i is smooth, it follows that $\widetilde{F}_i$ is smooth in Ω_α.

Case 2. (Away from the cut ℓ_i) Let $\Lambda_i := \bigcup_{j \neq i} \mu^{-1}(\ell_j) \subset \mathbb{R}^2$. We have that

$$(\check{h}_{\mathbb{R}^2})_i \circ t_0 \circ R_\alpha^{-1} = \check{h}_\Omega \circ t_0 \circ R_\alpha^{-1} = \tilde{h}_\Omega \circ R_\alpha^{-1} \quad \text{on the set } (R_\alpha \circ t_0^{-1})\big(\big\{(x,y) : y < -\tfrac{1}{2}\delta_i\big\}\big),$$

which by construction is smooth on this set. Thus $\widetilde{F}_i$ has the same degree of smoothness as μ on the set

$$\mu^{-1}\big((R_\alpha \circ t_0^{-1})\big(\big\{(x,y) : y < -\tfrac{1}{2}\delta_i\big\}\big)\big).$$

Note that this set does not contain $\mu^{-1}(\ell_i)$. The same argument applies to the analogue subsets of M corresponding to the regions $\big\{(x,y) : x < -\frac{1}{2}\delta_i\big\}$ and $\big\{(x,y) : x > \frac{1}{2}\delta_i\big\}$. On the subset of M corresponding to the region $\big\{(x,y) : y > \frac{1}{2}\delta_i\big\}$, the map $(\check{h}_{\mathbb{R}^2})_i$ is smooth by construction. Hence the map $\widetilde{F}_i$ is smooth on $M \setminus \Lambda_i$.

Case 3. (Along the cut ℓ_i, away from c_i) Remark that $t_0 \circ R_\alpha^{-1} = R_\alpha^{-1} \circ t_i$. By the construction of μ above the open sets Ω_β covering the cut ℓ_i, we have that $y_\beta^* \mu = t_i^{-1} \circ F_\beta$. Hence

$$y_\beta^*((\check{h}_{\mathbb{R}^2})_i \circ t_0 \circ R_\alpha^{-1} \circ \mu) = y_\beta^*((\check{h}_{\mathbb{R}^2})_i \circ F_\beta) \quad \text{on the set } \mu^{-1}(\Omega_\beta),$$

and this expression defines a smooth map. Thus $\widetilde{F}_i$ is smooth.

Hence putting cases 1, 2 and 3 together we have shown that $\widetilde{F}_i$ is smooth on $\mu^{-1}(\Omega_\beta)$ for all Ω_β covering the cut ℓ_i, and elsewhere, $\widetilde{F}_i$ is as smooth as μ. This concludes the proof of Claim 5.2.

Write $\widetilde{F} := (J, H)$. We then have the following conclusive claim.

CLAIM 5.3. *The symplectic manifold* (M,ω), *equipped with* J *and* H, *is a semitoric integrable system. Furthermore, the list of invariants* (i)–(v) *of the semitoric integrable system* $(M,\omega,(J,H))$ *is equal to the list of ingredients* (i)–(v) *that we started with. Finally,* M *is a compact manifold if and only* Δ *is compact.*

Let us prove this claim. We know from Claim 5.2 that $\widetilde{F}$ is smooth. Since the first component J is obtained from glueing proper maps, it follows from Theorem 3.11 that J is proper. Moreover, the Hamiltonian flow of J is everywhere periodic of period 2π because this is true in any local piece M_α. Clearly $\{J,H\}=0$, since it is a local property. It is also easy to see that the only singularities of $\widetilde{F}$ come from the singularities of the models F_α, as the glueing procedure does not create any additional singularities. Now, near any elliptic critical value, the homeomorphism μ is a local diffeomorphism, so $\widetilde{F}$ has the same singularity type as the elliptic model F_α. Finally, near a node we have checked in the proof of Claim 5.2 that $\widetilde{F}$ is precisely equal to the model F_i, and hence possesses a focus-focus singularity. Thus, provided we show that M is connected, (J,H) is a semitoric system.

Let us now consider its invariants (the connectedness of M will follow).

(i) As we mentioned, the singularities of $\widetilde{F}$ are only elliptic, except for the nodes $c_1,...,c_{m_f}$ above each of which we have constructed a focus-focus singularity. Hence we have m_f focus-focus singularities.

(ii) Each focus-focus singularity was constructed by glueing a semilocal model with prescribed Taylor series invariant $(S_i)^\infty$. Since this Taylor series is precisely a semilocal symplectic invariant, it is unchanged in the glued system $(M,\widetilde{F})$.

(iii) Thus we have a completely integrable system on M that defines an integral affine structure (with boundary) on the image of $\widetilde{F}$, except at the nodes c_i. For any choice of vertical half cuts (ℓ_i,ε_i), the generalized momentum polygon is the image of the affine developing map. But the momentum map μ, outside the focus-focus fibers, is precisely such a developing map and its image, by the glueing procedure, is the polygon Δ. Hence the semitoric polygon invariant of $\widetilde{F}$ is the orbit of $\Delta_{\rm w}$. (See Lemma 2.3.)

Notice that this shows that the image of μ is connected, which implies that the total space M, obtained by glueing above the image of μ, is connected as well.

(iv) It follows directly from (iii) above and the definition of the nodes c_j in (5.1) that the volume invariant defined in (2.5) is equal to $(h_1,...,h_{m_f})$.

(v) We calculate the twisting indices of our semitoric system with respect to the fixed polygon Δ or, which amounts to the same, with respect to the toric momentum map μ. By definition, the jth twist is the integer $\tilde{k}_j$ such that

$$d\mu=T^{\tilde{k}_j}\circ d\mu_j,$$

where μ_j is the privileged momentum map of the focus-focus fibration above c_j. From the second stage of the construction, we know that

$$\mu = F_\alpha = R_\alpha \circ \mu_j = \tau \circ T^{k_j} \circ \mu_j,$$

where τ is some translation. Hence $d\mu = T^{k_j} \circ d\mu_j$, and thus $\tilde{k}_j = k_j$.

Thus we see that we could prove the second part of the claim because our construction is by symplectically glueing local pieces with the appropriate ingredients as in Definition 4.5. This is an advantage of constructing by glueing local pieces together, rather than, for example, a global reduction on a larger space.

This concludes the proof of Claim 5.3, and hence the proof of the theorem.

References

[1] Atiyah, M. F., Convexity and commuting Hamiltonians. *Bull. London Math. Soc.*, 14 (1982), 1–15.

[2] Bourbaki, N., *Éléments de mathématique.* Fasc. XXXIII. *Variétés différentielles et analytiques. Fascicule de résultats* (Paragraphes 1 à 7). Actualités Scientifiques et Industrielles, 1333. Hermann, Paris, 1967.

[3] Brandsma, H., Paracompactness, covers and perfect maps, in *Topology Explained.* Topology Atlas, York University, Toronto, ON, 2003. Available at `http://at.yorku.ca/p/a/c/a/00.htm`.

[4] Daverman, R. J., *Decompositions of Manifolds.* AMS Chelsea Publishing, Providence, RI, 2007.

[5] Delzant, T., Hamiltoniens périodiques et images convexes de l'application moment. *Bull. Soc. Math. France*, 116 (1988), 315–339.

[6] Dufour, J.-P. & Molino, P., Compactification d'actions de $\mathbf{R}^n$ et variables action-angle avec singularités, in *Symplectic Geometry, Groupoids, and Integrable Systems* (Berkeley, CA, 1989), Math. Sci. Res. Inst. Publ., 20, pp. 151–167. Springer, New York, 1991.

[7] Duistermaat, J. J., On global action-angle coordinates. *Comm. Pure Appl. Math.*, 33 (1980), 687–706.

[8] Eliasson, L. H., Normal forms for Hamiltonian systems with Poisson commuting integrals—elliptic case. *Comment. Math. Helv.*, 65 (1990), 4–35.

[9] Gross, M. & Siebert, B., Affine manifolds, log structures, and mirror symmetry. *Turkish J. Math.*, 27 (2003), 33–60.

[10] — Mirror symmetry via logarithmic degeneration data. I. *J. Differential Geom.*, 72 (2006), 169–338.

[11] — Mirror symmetry via logarithmic degeneration data, II. *J. Algebraic Geom.*, 19 (2010), 679–780.

[12] — From real affine geometry to complex geometry. Preprint, 2007. `arXiv:math/0703822 [math.AG]`.

[13] Guillemin, V. & Sternberg, S., Convexity properties of the moment mapping. *Invent. Math.*, 67 (1982), 491–513.

[14] Leung, N. C. & Symington, M., Almost toric symplectic four-manifolds. *J. Symplectic Geom.*, 8 (2010), 143–187.

[15] Miranda, E. & Zung, N. T., Equivariant normal form for nondegenerate singular orbits of integrable Hamiltonian systems. *Ann. Sci. École Norm. Sup.*, 37 (2004), 819–839.
[16] Pelayo, A. & Vũ Ngọc, S., Semitoric integrable systems on symplectic 4-manifolds. *Invent. Math.*, 177 (2009), 571–597.
[17] Symington, M., Four dimensions from two in symplectic topology, in *Topology and Geometry of Manifolds* (Athens, GA, 2001), Proc. Sympos. Pure Math., 71, pp. 153–208. Amer. Math. Soc., Providence, RI, 2003.
[18] Vũ Ngọc, S., On semi-global invariants for focus-focus singularities. *Topology*, 42 (2003), 365–380.
[19] — Moment polytopes for symplectic manifolds with monodromy. *Adv. Math.*, 208 (2007), 909–934.
[20] Willard, S., *General Topology*. Dover, Mineola, NY, 2004.
[21] Wloka, J. T., Rowley, B. & Lawruk, B., *Boundary Value Problems for Elliptic Systems*. Cambridge University Press, Cambridge, 1995.
[22] Ziegler, G. M., *Lectures on Polytopes*. Graduate Texts in Mathematics, 152. Springer, New York, 1995.
[23] Zung, N. T., Symplectic topology of integrable Hamiltonian systems. I. Arnold–Liouville with singularities. *Compositio Math.*, 101 (1996), 179–215.

Álvaro Pelayo
School of Mathematics
Institute for Advanced Study
Princeton, NJ 08540
U.S.A.
apelayo@math.ias.edu

Washington University in St. Louis
Department of Mathematics
Campus box 1146
One Brookings Drive
St. Louis, MO 63130
U.S.A.
apelayo@math.wustl.edu

San Vũ Ngọc
Institut de Recherches
Mathématiques de Rennes
Université de Rennes 1
Campus de Beaulieu
FR-35042 Rennes cedex
France
san.vu-ngoc@univ-rennes1.fr

Received April 3, 2009