

FLAT SURFACES AND STABILITY STRUCTURES

by F. HAIDEN, L. KATZARKOV, and M. KONTSEVICH

CONTENTS

| | |
|--|-----|
| 1. Introduction | 247 |
| 1.1. Spaces of stability structures and flat surfaces | 247 |
| 1.2. Fukaya categories of surfaces | 248 |
| 1.3. Finite vs. infinite area, finite length property | 249 |
| 1.4. Meromorphic connections on T^*C and Nevanlinna's <i>Kernpolygon</i> | 249 |
| 1.5. Comparison with the work of Bridgeland–Smith | 250 |
| 2. Moduli spaces of flat surfaces | 251 |
| 2.1. Grading of surfaces and curves | 251 |
| 2.2. Flat surfaces with conical singularities | 253 |
| 2.3. Saddle connections | 256 |
| 2.4. Horizontal strip decompositions | 259 |
| 2.5. Complex-analytic point of view | 260 |
| 2.6. Voronoi and Delaunay partitions | 267 |
| 3. The topological Fukaya category of a surface | 268 |
| 3.1. Preliminaries on A_∞ -categories | 268 |
| 3.2. Surfaces with marked boundary | 271 |
| 3.3. Minimal A_∞ -category of an arc system | 272 |
| 3.4. Formal generators | 276 |
| 3.5. Localization | 278 |
| 3.6. Cosheaf of categories | 280 |
| 4. Tameness and geometricity | 283 |
| 4.1. Twisted complexes from curves | 283 |
| 4.2. Representations of nets | 286 |
| 4.3. Minimal twisted complexes | 293 |
| 4.4. Classification of objects | 294 |
| 5. Comparison of moduli spaces | 297 |
| 5.1. Stability structures on categories | 297 |
| 5.2. Charge lattice | 298 |
| 5.3. Statement of main theorem | 301 |
| 5.4. Support property | 302 |
| 5.5. Harder–Narasimhan filtrations | 303 |
| 5.6. The map $\mathcal{M}(X) \rightarrow \text{Stab}(\mathcal{F}(X))$ | 306 |
| 6. Cluster-like structures | 308 |
| 6.1. S-graphs | 308 |
| 6.2. Example: Dynkin type A_n | 311 |
| 6.3. Example: affine type \tilde{A}_n | 313 |
| 7. Open problems and further directions | 315 |
| Acknowledgements | 315 |
| References | 316 |

1. Introduction

1.1. Spaces of stability structures and flat surfaces

The space of stability structures, $\text{Stab}(\mathcal{C})$, of a triangulated category \mathcal{C} was introduced in seminal work of Bridgeland [11] based on a proposal of M.R. Douglas on D-branes in super-conformal field theories. The data of a stability structure, σ , is a distinguished class of objects in \mathcal{C} , the *semistable* ones, each with a *phase* $\phi \in \mathbf{R}$, and an additive

map $Z : K_0(\mathcal{C}) \rightarrow \mathbf{C}$, the *central charge*. These have to satisfy a number of axioms, see Section 5.1. Although the definition seems somewhat strange at first from a mathematical point of view, it leads to the remarkable fact, proven by Bridgeland, that $\text{Stab}(\mathcal{C})$ naturally has the structure of a complex manifold, possibly infinite-dimensional. Furthermore, $\text{Stab}(\mathcal{C})$ comes with an action of $\text{Aut}(\mathcal{C})$ and the universal cover of $\text{GL}^+(2, \mathbf{R})$. Thanks to the efforts of a number of people, the structure of $\text{Stab}(\mathcal{C})$ is understood in many particular cases, see e.g. [7, 12, 13, 15, 29, 34].

Moduli spaces of very similar nature appear in the theory of complex curves with a quadratic differential, also known as *half-translation surfaces* or simply *flat surfaces*, as they carry a flat Riemannian metric with conical singularities. Each of these moduli spaces, $\mathcal{M}(S)$, like $\text{Stab}(\mathcal{C})$, comes with a wall-and-chamber structure and an action of $\text{GL}^+(2, \mathbf{R})$. An explanation for this similarity was proposed by the third named author and Soibelman in [31]: The spaces $\mathcal{M}(S)$ should be realized as moduli spaces of stability structures for a suitably defined Fukaya category of the surface S . The main achievement of the present paper is to make precise and verify this claim, leading to the following result (Theorem 5.3 in the main text).

Theorem. — *Let S be a marked surface of finite type, $\mathcal{M}(S)$ the space of marked flat structures on S , and $\mathcal{F}(S)$ the Fukaya category of S . Then there is a natural map*

$$(1.1) \quad \mathcal{M}(S) \rightarrow \text{Stab}(\mathcal{F}(S))$$

which is bianalytic onto its image, which is a union of connected components.

For now it remains an open question whether the image is in fact all of $\text{Stab}(\mathcal{F}(S))$. The method of proof is new and relies on a detailed understanding of the category $\mathcal{F}(S)$.

1.2. Fukaya categories of surfaces

In Section 3 we give a self-contained elementary definition of the Fukaya A_∞ -category $\mathcal{F}(S)$ of a surface, and provide arguments to show that it agrees with other definitions found in the literature. Broadly speaking, these are either based on Floer intersection theory (see [22], [44], and many others), sheaves on ribbon graphs [21, 36, 45], or more direct approaches [1]. Our own approach to define $\mathcal{F}(S)$ falls into the latter category and is based on explicit A_∞ -structures and contractibility of the classifying space of arc systems on a marked surface.

The most significant result about $\mathcal{F}(S)$ which we prove is a complete classification of objects (Theorem 4.3 in the main text).

Theorem. — *Isomorphism classes of objects in $\mathcal{F}(S)$ are classified by isotopy classes of certain immersed curves on S together with local system of vector spaces on them. In particular, these categories are of tame representation type.*

Another way to say this is that all objects are “geometric”. This is not just a direct consequence of the definition, since $\mathcal{F}(S)$ is required to be closed under taking cones. We reduce the problem to the classification of certain pairs of filtrations on vector spaces with subquotients identified.

An important tool, e.g. for proving the equivalence with other definitions and computing K-theory, are formal generators of $\mathcal{F}(S)$ with endomorphism algebras given by graded quivers with quadratic monomial relations. These exist only for some S , but the other cases are obtained by localization of $\mathcal{F}(S)$. They are graded versions of the *gentle algebras* introduced by Assem–Skowroński [3].

1.3. *Finite vs. infinite area, finite length property*

The flat surfaces of interest in ergodic theory are the ones with finite area, corresponding to quadratic differentials with arbitrary order zeros and simple poles. For a generic flat surface with infinite area, on the other hand, the structure of the horizontal foliation is very explicitly described by certain generalized ribbon graphs (see Section 6), and the horizontal measure is not ergodic. In fact, the wall-and-chamber structure on the space of flat surfaces of infinite area is described by mutations similar to those in the theory of cluster algebras.

Another central tool in the study of flat surfaces of infinite area, apart from the horizontal foliation, is the *core*, $\text{Core}(S)$, which is defined to be the convex hull of the conical points. It is a compact subset of S with geodesic boundary, and in fact S is completely determined by its core. Changes to the structure of the core are related to wall-crossings of the first kind, (central charges becoming colinear over \mathbf{R}) while mutations are associated with wall-crossings of the second kind (central charges becoming real).

In terms of the associated stability structures, this dichotomy manifests itself in the finite length property of the t-structure, i.e. whether every descending chain of objects is eventually constant. Any stability structure on a category \mathcal{C} determines a family of t-structures. If the surface has infinite area, the generic t-structure in this family is finite length with a finite number of isomorphism classes of simple objects. If the surface has finite area on the other hand, the t-structures are not finite length. Closely related is the question of the distribution of phases of stable objects, in particular density, which was investigated by Dimitrov and the authors in [18].

1.4. *Meromorphic connections on T^*C and Nevanlinna’s Kernpolygon*

For a meromorphic quadratic differential, φ , on a compact Riemann surface, C , there is a corresponding connection D on T^*C so that the scalar multiples of $\pm\sqrt{\varphi}$ are its flat sections. This connection has a simple pole at the points where φ has a zero or pole of any order. Note also that the monodromy of D is necessarily contained in $\{\pm 1\}$.

Conversely, given a meromorphic connection D on T^*C with monodromy $\{\pm 1\}$ and simple poles, we get a quadratic differential, up to scalar multiple.

If the connection D has higher order poles then the associated quadratic differentials are not meromorphic but locally of the form

$$(1.2) \quad \exp(f(z))g(z)dz^2$$

where f and g are meromorphic and f has a pole of order one greater than D . We call this an *exponential singularity* of a quadratic differential. In terms of flat geometry, this singularity gives rise to n infinite-angle singularities of the metric, where n is the order of the pole of f . Such a singularity is modeled on the metric completion of the universal cover of the punctured Euclidean plane, and can be regarded as a limit of the conical singularity when the cone angle goes to infinity. We show the following:

Theorem. — *Let S be a half-translation surface of finite type with conical singularities, possibly infinite-angle ones, then S comes from a compact Riemann surface with quadratic differential with exponential singularities.*

See Theorem 2.3 in the main text. Thus allowing exponential singularities is quite natural from a geometric point of view, and turns out to be necessary to get stability structures on *partially* wrapped Fukaya categories of surfaces.

The previous theorem is similar in nature and has some overlap with an old result of Nevanlinna [38] on the class of functions which have Schwarzian derivative a polynomial. Indeed, ideas of his are found in our proof. The analysis of the topological structure of the branched covering defined by these functions leads him to a certain combinatorial structure, the *Kernpolygon*, which translates into the t-structure of the stability structure for a quadratic differential $\exp(P(z))dz^2$.

1.5. Comparison with the work of Bridgeland–Smith

A connection between moduli spaces of quadratic differentials and stability structures is also made in the work of Bridgeland–Smith [15]. They do not use the Fukaya category of the surface directly, but a certain CY3 category which is defined in terms of quivers with potential attached to ideal triangulations. These categories are connected to cluster theory via their quivers, and were shown to be subcategories of Fukaya categories of certain non-compact 6-folds by Smith [46]. Being CY3, they have generalized Donaldson–Thomas invariants in the sense of [31]. The DT invariants and their wall-crossing had previously appeared in work of Gaiotto–Moore–Neitzke [23], who also introduce the “WKB triangulation”, a central tool in [15]. Indeed, one of the motivations of Bridgeland–Smith is to develop a mathematical theory for [23].

The methods of [15] require one to restrict to meromorphic quadratic differentials without higher order zeros, and, at a generic point in the moduli space, at least one higher

order pole. Some quadratic differentials with simple poles do appear in the limit. It is expected that their results extend in some form to holomorphic quadratic differentials with simple zeros, but the stability structures are not described by quivers with potential and would require different methods to construct. In this paper we do indeed construct stability structures from these quadratic differentials, though not on CY3 categories.

It is somewhat puzzling that the same geometry should give rise to stability structures on, at least superficially, very different categories. The relation between the two categories (the wrapped Fukaya category of the surface and the CY3 category) is still to be clarified. It seems to be simpler in the case when all poles have order ≥ 3 . One can show that there is a non-commutative divisor in the sense of Seidel [42] with total space the wrapped category and central fiber the CY3 category. We suspect that in the case of double poles an additional deformation is necessary.

More recently, an extension of the results of Bridgeland–Smith to certain polynomial quadratic differentials on \mathbf{C} was carried out in [14, 27]. There, $\text{CY}n$ categories appear for $n \geq 2$.

2. Moduli spaces of flat surfaces

In this section we define the moduli spaces $\mathcal{M}(\mathbf{X})$ of marked flat structures. Section 2.1 discusses some topological preliminaries which are mostly relevant later when we define Fukaya categories of surfaces, but also play a small role in the definition of $\mathcal{M}(\mathbf{X})$. Next, in Section 2.2 we define flat surfaces with conical singularities and their moduli spaces and state a theorem about their local structure. This theorem is proven in the case of finite area in Section 2.3, using triangulations by saddle connections. Section 2.4 deals with the case of infinite area, where the main tool is the horizontal strip decomposition. In Section 2.5 we show that our class of flat surfaces may be described in terms of quadratic differentials on compact Riemann surfaces. Finally, in Section 2.6 we review the dual Voronoi and Delaunay partitions.

2.1. Grading of surfaces and curves

The notion of grading discussed in this subsection is essentially a special case of the one for symplectic manifolds with almost complex structure and their Lagrangian submanifolds, due to the third named author and Seidel [30, 43]. In the case of surfaces, no symplectic geometry is involved.

We use the notation $\Pi_1(\mathbf{X}, x, y)$ for the set of homotopy classes of paths from x to y in a space \mathbf{X} . Let \mathbf{X} be a smooth oriented surface. A *grading*, η , on \mathbf{X} is a foliation, i.e. section of the projectivized tangent bundle, $\mathbf{P}(\mathbf{TX})$. The pair (\mathbf{X}, η) is a *graded surface*. A morphism of graded surfaces $(\mathbf{X}, \eta) \rightarrow (\mathbf{Y}, \theta)$ is a pair (f, \tilde{f}) where $f : \mathbf{X} \rightarrow \mathbf{Y}$ is an orientation preserving local diffeomorphism, and \tilde{f} is a homotopy from $f^*\theta$ to η , up to

homotopy, i.e. $\tilde{f} \in \Pi_1(\Gamma(\mathbf{X}, \mathbf{P}(\mathrm{TX})), f^*\theta, \eta)$. Composition is given by

$$(2.1) \quad (f, \tilde{f}) \circ (g, \tilde{g}) = (f \circ g, (g^*\tilde{f}) \cdot \tilde{g})$$

where $\alpha \cdot \beta$ denotes concatenation of paths. Every graded surface \mathbf{X} has a *shift* automorphism given by the pair $(\mathrm{id}_{\mathbf{X}}, \sigma)$ where σ restricts, for every $x \in \mathbf{X}$, to the generator of $\pi_1(\mathbf{P}(\mathrm{T}_x\mathbf{X}))$ given by the orientation of \mathbf{X} , i.e. the path which rotates a line counterclockwise by an angle of π . Anticipating the connection with triangulated categories, we write this automorphism as $[1]$ and its integer powers as $[n]$.

For an oriented surface \mathbf{X} observe that the set

$$(2.2) \quad \mathcal{G} := \pi_0(\Gamma(\mathbf{X}, \mathbf{P}(\mathrm{TX})))$$

is a torsor over $\mathbf{H}^1(\mathbf{X}, \mathbf{Z}) = [\mathbf{X}, \mathbf{S}^1]$. The set \mathcal{G} classifies gradings on \mathbf{X} up to graded diffeomorphism with underlying map on \mathbf{X} the identity. The set of gradings of \mathbf{X} up to isomorphism is thus a quotient of \mathcal{G} .

Example. — Consider the cylinder $\mathbf{X} = \mathbf{R} \times \mathbf{S}^1$. Then

$$(2.3) \quad \pi_0(\Gamma(\mathbf{X}, \mathbf{P}(\mathrm{TX}))) \cong \pi_0(\mathrm{Maps}(\mathbf{S}^1, \mathbf{S}^1)) \cong \mathbf{Z}.$$

The above isomorphism simply takes a foliation η and counts how many times it rotates along the loop $\{0\} \times \mathbf{S}^1 \subset \mathbf{X}$. Call this number $\mathrm{deg}(\eta)$, then the map $(x, \theta) \rightarrow (-x, \theta)$ from the cylinder to itself acts on the foliation by switching the sign of $\mathrm{deg}(\eta)$. From this one can see that $|\mathrm{deg}(\eta)| \in \mathbf{Z}_{\geq 0}$ is a complete invariant of the graded surface.

Next, we consider immersed curves in graded surfaces. A *graded curve* in a graded surface (\mathbf{X}, η) is a triple $(\mathbf{I}, c, \tilde{c})$ where \mathbf{I} is a 1-manifold, possibly disconnected and with boundary, $c : \mathbf{I} \rightarrow \mathbf{X}$ is an immersion, and $\tilde{c} \in \Pi_1(\Gamma(\mathbf{I}, c^*\mathbf{P}(\mathrm{TX})), c^*\eta, \dot{c})$. To spell this out, \tilde{c} is given by a homotopy class of paths in $\mathbf{P}(\mathrm{T}_{c(t)}\mathbf{X})$ from the subspace given by the grading to the tangent space of the curve, varying continuously with $t \in \mathbf{I}$. The pushforward of a graded curve as above by a graded morphism (f, \tilde{f}) is given by $(\mathbf{I}, f \circ c, (c^*\tilde{f}) \cdot \tilde{c})$. Note that graded curves are not naturally oriented.

A point of transverse intersection of a pair $(\mathbf{I}_1, c_1, \tilde{c}_1), (\mathbf{I}_2, c_2, \tilde{c}_2)$ of graded curves determines an integer as follows. Suppose $t_i \in \mathbf{I}_i$ with

$$(2.4) \quad c_1(t_1) = c_2(t_2) = p \in \mathbf{X}, \quad \dot{c}_1(t_1) \neq \dot{c}_2(t_2) \in \mathbf{P}(\mathrm{T}_p\mathbf{X}).$$

We have the following homotopy classes of paths in $\mathbf{P}(\mathrm{T}_p\mathbf{X})$: (1) $\tilde{c}_1(t_1)$ from $\eta(p)$ to $\dot{c}_1(t_1)$, (2) $\tilde{c}_2(t_2)$ from $\eta(p)$ to $\dot{c}_2(t_2)$, (3) κ from $\dot{c}_1(t_1)$ to $\dot{c}_2(t_2)$ given by counterclockwise rotation in $\mathrm{T}_p\mathbf{X}$ by an angle $< \pi$. Define the *intersection index* of c_1, c_2 at p

$$(2.5) \quad i_p(c_1, c_2) = \tilde{c}_1(t_1) \cdot \kappa \cdot \tilde{c}_2(t_2)^{-1} \in \pi_1(\mathbf{P}(\mathrm{T}_p\mathbf{X})) \cong \mathbf{Z}.$$




FIG. 1. — Smoothing an intersection with $i_p(c_1, c_2) = 1$

More precisely, this notation is correct only if c_1, c_2 pass through p exactly once, in particular if c_i are in general position, otherwise the index may depend on the t_i as well. We note that

$$(2.6) \quad i_p(c_1, c_2) + i_p(c_2, c_1) = 1$$

and

$$(2.7) \quad i_p(c_1[m], c_2[n]) = i_p(c_1, c_2) + m - n.$$

Another observation which will be used repeatedly, is that if graded curves c_1 and c_2 intersect in p with $i_p(c_1, c_2) = 1$, then we may perform a kind of smoothing near p which produces again a graded curve (see Figure 1). (If the intersection index is not 1 then the grading does not extend continuously along the modified curves.)

A graded surface (X, η) has a canonical double cover τ of X with fiber over $p \in X$ the set of orientations of $\eta(p)$. We will consider τ as a locally constant sheaf and write

$$(2.8) \quad \mathbf{Z}_\tau := \mathbf{Z} \otimes_{\mathbf{Z}/2} \tau$$

and similarly for \mathbf{R}_τ and \mathbf{C}_τ . Singular (co)homology with coefficients in any of these local systems is defined, as is deRham cohomology with coefficients in \mathbf{R}_τ or \mathbf{C}_τ . Integration provides a pairing

$$(2.9) \quad H_1(X, \partial X; \mathbf{Z}_\tau) \otimes_{\mathbf{Z}} H_{\text{dR}}^1(X, \partial X; \mathbf{C}_\tau) \rightarrow \mathbf{C}$$

where $H_{\text{dR}}^*(M, \partial M; E)$ is the cohomology of forms with coefficients in E which vanish when pulled back to ∂M .

2.2. Flat surfaces with conical singularities

The following types of structures on an orientable surface X are in one-to-one correspondence:

1. Flat Riemannian metrics together with a covariantly constant foliation η and orientation.

2. Complex structures together with a non-vanishing holomorphic quadratic differential (section of the square of the holomorphic cotangent bundle).
3. Pairs of pointwise linearly independent commuting vector fields, defined up to simultaneous change of sign, i.e. a section of the quotient of the frame bundle of TX by $\mathbf{Z}/2$.

On \mathbf{R}^2 the standard versions of these structures are: (1) the standard Euclidean metric with the foliation by horizontal lines, (2) the standard complex structure $\mathbf{R}^2 = \mathbf{C}$ and the quadratic differential dz^2 , (3) the pair of vector fields (∂_x, ∂_y) . A general surface with one of the structures above is locally isomorphic to standard \mathbf{R}^2 by charts obtained from: (1) the exponential map of the metric, (2) the complex path integral of a square-root of the quadratic differential, (3) the flows of the vector fields. The transition functions for such charts are of the form

$$(2.10) \quad v \mapsto \pm v + c.$$

Indeed, these are the symmetries of \mathbf{R}^2 with any of the standard structures above. A *smooth flat surface* is a surface with any of the above structures. The covariantly constant foliation makes it a graded surface and is called the *horizontal foliation*. Another more precise but longer term used for this kind of structure is *half-translation surface*, emphasizing the fact that the holonomy of the metric must be a subgroup of $\{\pm 1\}$.

The metric completion of a smooth flat surface may have fairly complicated singularities in general [10]. Here, we will only consider singularities which are conical with cone angle a positive integer multiple of π or infinite. They may be described as follows. For $n \in \mathbf{Z}$ let C_n be the space

$$(2.11) \quad C_n := (\mathbf{Z}/n) \times \mathbf{R} \times \mathbf{R}_{\geq 0} / \sim, \quad (k, x, 0) \sim (k+1, -x, 0) \quad \text{for } x \leq 0$$

and

$$(2.12) \quad C_\infty := \mathbf{Z} \times \mathbf{R} \times \mathbf{R}_{\geq 0} / \sim, \quad (k, x, 0) \sim (k+1, -x, 0) \quad \text{for } x \leq 0.$$

Let C be one of the above spaces and $0 \in C$ the distinguished point $(0, 0, 0) \sim (k, 0, 0)$. The standard flat structure on \mathbf{R}^2 gives $C \setminus 0$ the structure of a smooth flat surface and C is its metric completion. The metric is singular unless $C = C_2$. The singularity is called a *conical singularity* with cone angle $n\pi$ if $C = C_n$ and an *infinite angle singularity* if $C = C_\infty$. See Figure 2 for a picture of C_3 .

It is technically convenient to work not with the above spaces directly, but instead their real blow-up at the origin, denoted \widehat{C} . This construction produces a surface with boundary in which the cone point $0 \in C$ is replaced by the set of geodesics starting at that point, which is topologically S^1 in the case of C_n and \mathbf{R} in the case of C_∞ . The complex structure on $C \setminus 0$ does not extend to \widehat{C} , but the horizontal foliation does. Also, locally, any square-root of the quadratic differential extends as a closed complex-valued 1-form to \widehat{C} which vanishes on vectors tangent to the boundary. Indeed, if we identify \widehat{C}_∞ with $\mathbf{R}_{\geq 0} \times \mathbf{R}$ the 1-form is given by $e^{i\psi}(dx + ixdy)$.




FIG. 2. — Conical singularity with cone angle 3π (“monkey saddle”)

Definition. — Let Y be a smooth surface with boundary. The structure of a real blow-up of a flat surface on Y is given by the structure of a smooth flat surface on $Y \setminus \partial Y$ such that each component of ∂Y has a neighborhood diffeomorphic to a neighborhood of the boundary of some \widehat{C}_n or \widehat{C}_∞ by a diffeomorphism preserving the flat structure on the interior, and satisfying the following completeness condition: The space X obtained by identifying points lying in the same component of ∂Y is the metric completion of $Y \setminus \partial Y$.

Definition. — A flat surface is a complete metric space together with a partition $X = X_{\text{sm}} \cup X_{\text{sg}}$ into smooth and singular points and the structure of a smooth flat surface on X_{sm} (i.e. direction of horizontal foliation) such that near each $x \in X_{\text{sg}}$ the space is locally isometric to one of C_n or C_∞ .

Note that in the above definition $x \in X_{\text{sg}}$ is not necessarily a singular point of the metric if near x the model is C_2 . One can consider it as an arbitrary marked point in the metrically smooth part. As the terminology suggests, the real blow-up of a flat surface has the structure of a *real blow-up of a flat surface*. Conversely, one obtains a flat surface from the real blow-up by taking the metric completion of the interior. (Topologically, this contracts each boundary component to a point.)

Definition. — Fixing a graded surface S we consider the moduli space of marked flat surfaces $\mathcal{M}(S)$. An element of $\mathcal{M}(S)$ is represented by a real blow-up of a flat surface Y with underlying oriented surface S together with a morphism of graded surfaces $f : S \rightarrow Y$ with underlying map the identity on S . Two such pairs (Y, f) , (Z, g) are equivalent if there is a diffeomorphism $h : Y \rightarrow Z$ preserving the flat structure so that $h \circ f$ and g are isotopic as maps of graded surfaces. We give flat structures (before identification) the topology of pointwise convergence on the real blow-up, and $\mathcal{M}(S)$ the quotient topology.

Note that the grading on S fixes the number and types of singular points and marked points of any flat surfaces in $\mathcal{M}(S)$. Explicit descriptions of $\mathcal{M}(S)$ when S is a disk or annulus are found in Section 6.

If Y is the real blow-up of a flat surface then the square-roots of the quadratic differential on $Y \setminus \partial Y$ may be viewed as a closed 1-form α with values in \mathbf{C}_τ , and as such it extends to the boundary. In particular we get a class $[\alpha] \in H_{\text{dR}}^1(Y, \partial Y; \mathbf{C}_\tau)$ so using the pairing (2.9) we get a homomorphism $H_1(Y, \partial Y; \mathbf{Z}_\tau) \rightarrow \mathbf{C}$ which represents the *periods*

of the flat surface. We say that a graded surface Y is of *finite type* if Y and ∂Y have a finite number of connected components and $H_1(Y, \partial Y; \mathbf{Z}_\tau)$ is finitely generated.

Theorem 2.1. — *Let S be a graded surface of finite type. The map*

$$(2.13) \quad \Pi : \mathcal{M}(S) \rightarrow \text{Hom}(H_1(S, \partial S; \mathbf{Z}_\tau), \mathbf{C}),$$

assigning to each flat surface its periods, is a local homeomorphism.

The strategy of the proof will depend on whether the area of the flat surface is finite or infinite and will be completed in the following two subsections.

Grading of curves on flat surfaces

Let X be a surface with flat structure defined in terms of a complex structure and non-vanishing holomorphic quadratic differential $\varphi \in \Gamma(X, (T^*X)^{\otimes 2})$. In this context, grading of curves and intersection index can be defined somewhat more concretely. Recall that the holomorphic and real tangent bundles may be identified, and the horizontal foliation is the unique section of $\mathbf{P}(TX)$ on which φ is real and positive. Moreover, we have an identification of $\mathbf{P}(TX)$ with the bundle with constant fiber $\mathbf{C}^*/\mathbf{R}_{>0} \cong U(1)$ under which this section corresponds to $1 \in U(1)$.

Given an immersed curve $\alpha : I \rightarrow X$ a grading is defined by a function $\phi : I \rightarrow \mathbf{R}$ such that

$$(2.14) \quad \varphi(\dot{\alpha}(t), \dot{\alpha}(t)) \in \mathbf{R}_{>0} e^{2\pi i \phi(t)}.$$

The corresponding path $\alpha^* \Omega \rightsquigarrow \dot{\alpha}$ is given by

$$(2.15) \quad s \mapsto e^{2\pi i \phi(t)s}, \quad s \in [0, 1]$$

under the identification as above. Note that α is geodesic if and only if ϕ is locally constant. Suppose we have curves α_1, α_2 with grading given by real-valued functions ϕ_1, ϕ_2 intersecting transversely in $p \in S$. Then the intersection index is

$$(2.16) \quad i_p(\alpha_1, \alpha_2) = \lceil \phi_1(p) - \phi_2(p) \rceil.$$

2.3. *Saddle connections*

Geodesics, in particular those of finite length, are fundamental in the study of flat surfaces. The smooth part, X_{sm} , of a flat surface X is a Riemannian surface, and by our completeness requirement any geodesic on X_{sm} can either be extended indefinitely or converges towards a point in the complement X_{sg} of X_{sm} . A *saddle connection* is a geodesic converging towards points in X_{sg} (“saddles”) in both directions. The endpoints are not required to be distinct.

A *geodesic arc system* on a flat surface is a collection of saddle connections which intersect at most in the endpoints.

Proposition 2.1. — *Let X be a flat surface of finite type, then the size of any geodesic arc system is bounded above by a constant depending only on the topology of the real blow-up $Y = \widehat{X}$.*

Proof. — Let \mathcal{A} be a finite geodesic arc system on X and $E = |\mathcal{A}|$ its cardinality. Each saddle connection in \mathcal{A} together with an arbitrary choice of orientation gives a non-zero class in $H_1(X, X_{\text{sg}}, \mathbf{Z})$. A complete set of relations between these classes is given by the polygons cut out by \mathcal{A} . Let F be the number of these polygons. Any bi-gon cut out by \mathcal{A} must contain a conical point with cone angle π in its interior, so their number, D , is bounded by the number of compact boundary components of Y . Assigning boundary edges to polygons we find that

$$(2.17) \quad F - D \leq \frac{2}{3}E.$$

By assumption, $B = \text{rk}H_1(X, X_{\text{sg}}, \mathbf{Z})$ is finite, and since $E - F \leq B$ we get

$$(2.18) \quad E \leq 3(D + B)$$

which completes the proof. □

Let \mathcal{A} be a maximal geodesic arc system on X . By the previous proposition it is finite. Let us study its local geometry. If $p \in X_{\text{sg}}$ and D is the corresponding boundary component of \widehat{X} , then the set of maximal geodesics starting at p is identified with D . A finite subset $E \subset D$ corresponds to (ends of) saddle connections in \mathcal{A} which converge to p . For components of $D \setminus E$ there are two possibilities: Either the maximal geodesics in these directions intersect a saddle connection in \mathcal{A} , or not. They cannot converge to a singularity by maximality of \mathcal{A} . In the first case there is a triangle cut out by \mathcal{A} with p as one of its vertices. In the second case there is an conical sector with cone angle $\geq \pi$, possibly infinite, again by maximality.

Proposition 2.2. — *Suppose \mathcal{A} is a maximal geodesic arc system on a flat surfaces X . Let K be the union of X_{sg} , the saddle connections in \mathcal{A} , and the triangles cut out by \mathcal{A} . Then K is the convex hull of X_{sg} , thus independent of \mathcal{A} .*

Proof. — From the local geometry above we see that any geodesic which leaves K stays in $X \setminus K$ for all subsequent times, so K is convex. It contains X_{sg} by definition. If K' is convex and contains X_{sg} , then it must also contain all saddle connections, in particular those in \mathcal{A} , and thus all triangles cut out by \mathcal{A} , so $K \subseteq K'$. □

We will refer to the convex hull of X_{sg} as the *core* of the flat surface X , denoted $\text{Core}(X)$. The previous proposition shows that, for X of finite type, it is covered by a finite number of singular points, saddle connections, and geodesic triangles, thus compact.

Proposition 2.3. — *Let X be a flat surface with the property that each component contains at least one singular point, and let K be the core of X . Then K is a deformation retract of X .*

Proof. — The assumption ensures that each $p \in X$ has finite distance to $\text{Core}(X)$. A deformation retraction can be defined by moving $p \in X$ along a geodesic path to the unique point in $\text{Core}(X)$ which is closest to p . \square

Let X be a connected flat surface with finite area and assume that X is not smooth, which excludes only the case of the torus. A maximal geodesic arc system on X gives a triangulation with some triangles possibly degenerate. By definition, a triangle is degenerate if two of its edges coincide, so that they meet in a conical singularity with cone angle π . Using triangulations by saddle connections, we can now give a proof of Theorem 2.1 in the case when the flat surface has finite area. This result is classical, see e.g. Veech [50], but we include a proof here for completeness. The proof in the case of infinite area (but still finite type) will be given in the next section.

Proof of Theorem 2.1, finite area case. — Fix a graded surface X of finite type admitting flat structures with finite total area. This means that X is compact as a surface with boundary. We also assume that there is at least one singularity, excluding the case of the torus which is easily handled directly. Let $\Gamma = H_1(X, \partial X; \mathbf{Z}_\tau)$ and $\Pi : \mathcal{M}(X) \rightarrow \text{Hom}(\Gamma, \mathbf{C})$ be the period map. We want to show that Π is a local homeomorphism.

Let $a_0 \in \mathcal{M}(X)$ and equip X with this flat structure. Pick a maximal collection \mathcal{A} of geodesic arcs on X . These arcs give a triangulation of X , since the flat metric has finite area by assumption. Fix a grading on each arc in \mathcal{A} , then each $\alpha \in \mathcal{A}$ gives a class $[\alpha] \in \Gamma$. A complete set of relations between these classes comes from the list of triangles. The number $\Pi(a_0)([\alpha]) \in \mathbf{C}$ records the length and slope of an arc $\alpha \in \mathcal{A}$.

Since the number of triangles is finite, there is a contractible neighborhood U of $\Pi(a_0) \in \text{Hom}(\Gamma, \mathbf{C})$ with the following property: For any $b \in U$ and any three classes $\alpha_1, \alpha_2, \alpha_3$ in Γ corresponding to the edges of a triangle of \mathcal{A} , the vectors $\Pi(b)([\alpha_i])$ form the edges of a non-degenerate triangle in \mathbf{C} . This ensures that for each $b \in U$ we get an actual flat structure on X essentially by direct construction. Namely, we cut X into its geodesic triangles, deform each of the triangles according to the values $\Pi(b)([\alpha])$, $\alpha \in \mathcal{A}$, then glue the triangles back together in the same combinatorial manner. In this way, we have defined a continuous section Ξ of Π over U .

To complete the proof, it remains to be shown that $\Xi(U)$ is an open neighborhood of a_0 . Note that there is a neighborhood W of $a_0 \in \mathcal{M}(X)$ such that all saddle connections in \mathcal{A} persist throughout W , i.e. do not break along a singularity. Making W sufficiently small, we can assume that $W \subset \Pi^{-1}(U)$. Then the same collection of arcs \mathcal{A} , up to homotopy, again form a maximal collection of geodesic arcs for any flat surface in W . This shows that $W \subset \Xi(U)$, so $\Xi(U)$ is a neighborhood of a_0 , in fact open by the same arguments. \square

2.4. *Horizontal strip decompositions*

Given a complex number $z \in \mathbf{C}$ with $\text{Im}(z) > 0$ the corresponding *horizontal strip of finite height* is the flat surface with boundary

$$(2.19) \quad \{w \in \mathbf{C} \mid 0 \leq \text{Im}(w) \leq \text{Im}(z)\}$$

with marked points $\{0, z\}$ on the boundary. The *horizontal strip of infinite height* is just the closure of the upper half-plane with the origin as marked point. It turns out that a generic surface of finite type with infinite area is obtained, in a canonical way, by gluing these types of pieces along their boundaries.

Proposition 2.4. — *Let X be flat surface of finite type with infinite area. Assume that X is connected and has at least one conical point. Then, after possibly rotating the horizontal direction, the leaves converging towards a conical point cut the surface into horizontal strips as above.*

Proof. — Assume that no horizontal leaf is a saddle connection. This can always be achieved by rotating the horizontal direction, as there are at most countably many slopes of saddle connections. In fact, the set of slopes of saddle connections is closed under the present assumptions, but this is not needed. Note that no leaf can be a closed geodesic either, as a cylinder foliated by closed geodesics is bounded by one or more saddle connections. Arguments of Strebel [47] show that there are then only two types of leaves. *Generic leaves* which are closed and intersect the core of X in a compact interval, and *Critical leaves* which converge towards a conical point in one direction and eventually leave the core in the other. Indeed, the boundary of the closure of a leaf remaining entirely in the core would be a union of leaves which are saddle connections, which is impossible. In particular, the closure of a critical leaf adds only a single conical point. The components of the union of all generic leaves are open parts of horizontal strips. The exponential maps at the various conical points give identifications of the closures of these open components with the standard horizontal strips of finite or infinite height. \square

Let us say a bit more about the horizontal strip decompositions appearing in the previous proposition. A conical point with finite cone angle lies on finitely many horizontal strips, while a conical point with infinite cone angle lies on infinitely many, all but finitely many of which will have infinite height. Each horizontal strip of finite height contains a unique saddle connection: the straight line segment connecting 0 and z . These are exactly the saddle connections $\alpha_1, \dots, \alpha_m$, which do not intersect any of the critical leaves. We will refer to them as *simple saddle connections*. Following the leaves and collapsing the horizontal strips of infinite height we construct a deformation retraction of the surface to the union of the α_i and the conical points. This shows that the α_i give a basis of $H_1(X, \partial X; \mathbf{Z}_\tau)$, after choosing gradings.

The graph formed by the α_i is identified with part of the Hausdorff version of the leaf space, L , of the horizontal foliation. We may realize L inside X by adding the positive

imaginary axis in each horizontal strip of infinite height. Combinatorially, \mathbf{L} is a bit of a generalization of a ribbon graph with some vertices connected to infinitely many half edges (those corresponding to infinite-angle singularities). More precisely, the set of half edges meeting a given vertex is a torsor over some cyclic group. This will be discussed in more detail in Section 6.1.

Proof of Theorem 2.1, infinite area case. — Let $a_0 \in \mathcal{M}(\mathbf{X})$ be a flat structure and assume first that there are no leaves which are saddle connections. By the considerations above, the critical leaves cut \mathbf{X} into horizontal strips. Let $\alpha_1, \dots, \alpha_m$ be the saddle connections corresponding to the horizontal strips of finite height. Also choose the grading on each α_i so that $\text{Im } \Pi(a_0)(\alpha_i) > 0$. This is possible since α_i is not horizontal and shifting the grading flips the sign of $\Pi(a_0)(\alpha_i)$. We get a basis $[\alpha_1], \dots, [\alpha_m]$ of $\Gamma = H_1(\mathbf{X}, \partial\mathbf{X}; \mathbf{Z}_\tau)$. Consider the open subset $U \subset \text{Hom}(\Gamma, \mathbf{C})$ of maps sending each $[\alpha_i]$ to the upper half plane. For each $b \in U$ we construct a flat surface $\Xi(b)$ with the same combinatorial type of horizontal strip decomposition as a_0 , but different slopes and lengths of the α_i given by b . The map Ξ defines a section of $\mathcal{M}(\mathbf{X}) \rightarrow \text{Hom}(\Gamma, \mathbf{C})$ over U . To see that $\Xi(U)$ is open, note that any flat surface sufficiently close to one in $\Xi(U)$ has the same horizontal strip decomposition.

If a_0 has horizontal saddle connections, we argue as follows. By assumption $\Gamma \neq 0$, so $\mathbf{R} \subset \widetilde{\text{GL}^+(2, \mathbf{R})}$ acts on $\mathcal{M}(\mathbf{X})$ by rotation of the horizontal direction and each orbit contains a flat structure with only critical and generic leaves considered before. As Π is equivariant with respect to the action it must be a local homeomorphism near a_0 as well. \square

2.5. Complex-analytic point of view

Meromorphic quadratic differentials

Let \mathbf{C} be a compact Riemann surface and φ a non-zero meromorphic quadratic differential on \mathbf{C} , i.e. a meromorphic section of the square of the canonical bundle. If $\mathbf{D} \subset \mathbf{C}$ is the set of zeros and poles of φ , then $\mathbf{X}_{\text{sm}} = \mathbf{C} \setminus \mathbf{D}$ has the structure of a smooth flat surface, possibly incomplete. The flat geometry near a zero or pole is analyzed by Strebel [47] by finding a local holomorphic coordinate in which the quadratic differential has a particularly simple form. Let us recall the results.

Zeros and simple poles. There is a local coordinate z such that $\varphi = z^n dz^2$, $n = -1, 1, 2, 3, \dots$. The metric has a conical singularity at $z = 0$ with cone angle $(n + 2)\pi$. In particular, if φ has no higher order poles, then the corresponding flat surface has finite total area.

Poles of order $n \geq 2$. If n is odd then there is a local coordinate z with $\varphi = z^{-n} dz^2$. However, if n is even then φ admits a square root near $z = 0$, and the residue $a \in \mathbf{C}$ of $\sqrt{\varphi}$ is an invariant up to sign. It turns out to be the only local invariant, so that there is a

coordinate z with

$$(2.20) \quad \sqrt{\varphi} = \left(z^{-n/2} + \frac{a}{z} \right) dz$$

when $n \geq 4$ and $\sqrt{\varphi} = \frac{a}{z} dz$ when $n = 2$. In any case, we note that the flat metric near a higher order pole has infinite area and is complete, i.e. higher order poles do not lead to any additional singularities of the metric.

Theorem 2.2. — *Let C be a compact Riemann surface with non-zero meromorphic quadratic differential φ with set of zeros and poles D . Then $|\varphi|$ gives $C \setminus D$ the structure of the smooth part of a flat surface of finite type with conical points which have cone angle an integer multiple of π . Conversely, any flat surface of finite type without infinite-angle conical points is obtained in such a way.*

Remark. — Let $\mathcal{T}_{g,n}$ be the Teichmüller space of compact Riemann surfaces of genus g with n marked points. The cotangent space to $\mathcal{T}_{g,n}$ at C is naturally identified with the space of quadratic differentials on C with at most simple poles at the marked points. Thus every point in the cotangent bundle of $\mathcal{T}_{g,n}$ corresponds to some flat surface of genus g with at most n conical points with cone angle π (and other conical points with bigger cone angle). The moduli spaces $\mathcal{M}(S)$ defined above correspond to strata in the cotangent bundle. The ribbon graph attached to a Strebel differential is a special case of the core of a flat surface introduced here.

Proof. — It remains to show the second part, so let X be a flat surface of finite type without infinite-angle conical points. First of all, we have a complex structure and non-vanishing holomorphic quadratic differential on the smooth part X_{sm} . We know that near a conical point the flat structure comes from a zero or first order pole of a quadratic differential, so the complex structure extends to all of X and φ extends meromorphically. If X has finite area, then it is compact and so $C = X$, $D = X_{sg}$, and we are done.

For the case of infinite area we need to find a compact $K \subset X$ such that the components of the complement $X \setminus K$ are isometric to some punctured neighborhoods of higher order poles. As a first step, take K to be the core of X . The boundary of K in X is a sequence of saddle connections meeting in exterior angles $\phi_p \geq \pi$, by convexity of K . The condition on the global monodromy of the metric (existence of horizontal foliation) implies that

$$(2.21) \quad \sum_p (\phi_p - \pi) = n\pi$$

where p runs over the corners of a fixed component of the boundary of K , and $n \in \mathbf{Z}_{\geq 0}$ depends on that component.

Points in the complement of K are parametrized by pairs (l, d) where l is a straight line starting at some point in the boundary of K , meeting K only in the one endpoint,

and meeting ∂K in angles $\geq \pi$, together with a positive distance $d \in \mathbf{R}_{>0}$ from K . In particular, each component of $X \setminus K$ is topologically a cylinder.

We enlarge K so that its boundary is a sequence of closed geodesics and straight edges (finite length geodesics) meeting in exterior angles $3\pi/2$, not necessarily at conical points. Consider some component B of the boundary of K . If B is a closed geodesic, then the corresponding component of $X \setminus K$ is identified with a flat cylinder which comes from a second order pole of a meromorphic quadratic differential. Otherwise, B is a union of an even number of edges with lengths a_1, \dots, a_{2n} in counterclockwise order. If we parallel transport some edge outwards from K by a distance $d \geq 0$ then the lengths of the adjacent edges increase by d . A simple proof by induction on n shows that when n is odd we can parallel transport the edges outwards so that they all have the same length l , which can be an arbitrary real number bigger than some constant depending on the original K . Hence, in this case the local geometry is completely determined by n , and so must coincide with the one coming from a pole of order $n + 2$.

When n is even we can arrange the sequence of lengths to be of the form $l_1, l_2, l_3, l_3, \dots, l_3$ after parallel transport of the edges, and add any positive real number to all elements at once. Without loss of generality, the first edge (of length a) is horizontal, then the integral of the square root of the quadratic differential around the cylinder is just $a = \pm(l_1 - l_2 + i(l_1 - l_3))$, so the geometry must be locally isometric to the one coming from the square of the abelian differential $(z^{-n-2} + a/z)dz$. \square

Exponential-type singularities

An *exponential-type singularity* is a transcendental singularity of a quadratic differential of the form

$$(2.22) \quad e^{f(z)}g(z)dz^2, \quad f, g \text{ meromorphic}$$

in some local coordinate z , where f has a pole of order n at $z = 0$. By change of coordinates we may assume that $f(z) = z^{-n}$. The following proposition describes the local flat geometry near such a singularity.

Proposition 2.5. — *Let*

$$(2.23) \quad \varphi = \exp(z^{-n})z^m g(z)dz^2$$

with $n \in \mathbf{Z}_{>0}$, $m \in \mathbf{Z}$, g holomorphic, non-vanishing on a neighborhood U of $0 \in \mathbf{C}$. Set

$$(2.24) \quad D = \{0 < |z| \leq r\} \subset U$$

then the completion of D with respect to $|\varphi|$ has n additional points, all of which are ∞ -angle singularities.

Lemma 2.1. — Let $\rho \geq 1$, $\pi/2 < \varepsilon < \pi$, $A = \{re^{i\phi} \mid r \geq \rho, |\phi| \leq \varepsilon\} \subset \mathbf{C}$, $h(z)$ a smooth function on A such that for some $C_1, C_2 > 0$, $\lambda \in \mathbf{R}$

$$(2.25) \quad C_1|z|^\lambda \leq h(z) \leq C_2|z|^\lambda$$

Then the completion of A with respect to the metric

$$(2.26) \quad (e^{-\operatorname{Re}(z)}h(z))^2|dz|^2$$

has a single additional point.

Proof. — 1. Define $\eta(z) = (\max\{1, |\operatorname{Im}(z)|\})^\lambda e^{-\operatorname{Re}(z)}$, then for sufficiently small $\varepsilon > 0$, the boundary of $\{\eta \geq \varepsilon\}$ in A is given by the curve

$$(2.27) \quad \operatorname{Re}(z) = \lambda \log(\max\{1, |\operatorname{Im}(z)|\}) - \log(\varepsilon).$$

Moreover, for distinct values of ε , these curves have positive distance with respect to the Euclidean metric.

2. We claim that $e^{-\operatorname{Re}(z)}h(z)$ is bounded below on $\{\eta \geq \varepsilon\}$ for any $\varepsilon > 0$. By assumption,

$$(2.28) \quad e^{-\operatorname{Re}(z)}h(z) \geq C_1|z|^\lambda e^{-\operatorname{Re}(z)}.$$

Case $\lambda \geq 0$: For $|\operatorname{Im}(z)| \geq 1$ we have

$$(2.29) \quad C_1|z|^\lambda e^{-\operatorname{Re}(z)} \geq C_1|\operatorname{Im}(z)|^\lambda e^{-\operatorname{Re}(z)}$$

$$(2.30) \quad = C_1\eta(z)$$

$$(2.31) \quad \geq C_1\varepsilon$$

Case $\lambda < 0$:

$$(2.32) \quad C_1|z|^\lambda e^{-\operatorname{Re}(z)} \geq C_1(|\operatorname{Re}(z)|e^{-\operatorname{Re}(z)/\lambda} + |\operatorname{Im}(z)|e^{-\operatorname{Re}(z)/\lambda})^\lambda$$

The first term is bounded above, as $\operatorname{Re}(z)$ is bounded above by $-\log(\varepsilon)$, the second term is bounded above by $\varepsilon^{1/\lambda}$.

3. We claim that for every $\varepsilon > 0$, the set $\{\eta \geq \varepsilon\}$ is complete with respect to the metric $g = (e^{-\operatorname{Re}(z)}h(z))^2|dz|^2$, thus any Cauchy sequence $z_j \in A$ without limit must satisfy $\eta(z_j) \rightarrow 0$. Namely, we have $g \geq Cg_{\text{eucl}}$ on $\{\eta \geq \varepsilon/2\}$ by the previous step, and the boundary curves $\{\eta = \varepsilon\}$, $\{\eta = \varepsilon/2\}$ have some positive distance δ , so

$$(2.33) \quad d(z_1, z_2) \geq C \min\{|z_1 - z_2|, 2\delta\}$$

for $z_1, z_2 \in \{\eta \geq \varepsilon\}$. Hence, any sequence which is Cauchy with respect to g is Cauchy for the standard metric.

4. We claim that $\text{diam}\{\eta \leq \varepsilon\} \rightarrow 0$ as $\varepsilon \rightarrow 0$. Assume $1 > \varepsilon > 0$ is sufficiently small so that $\{\eta \leq \varepsilon\}$ contains no z with $\text{Re}(z) \leq 0$, $|\text{Im}(z)| \leq \rho$, then for $z \in \{\eta \leq \varepsilon\}$ the path $\alpha(t) = z + t$, $t \in [0, \infty)$ is contained in $\{\eta \leq \varepsilon\}$. We compute the length of α ,

$$(2.34) \quad l(\alpha) = \int_0^\infty e^{-\text{Re}(z)-t} h(z+t) dt$$

$$(2.35) \quad \leq C_2 \int_0^\infty e^{-\text{Re}(z)-t} |z+t|^\lambda dt =: C_2 I$$

Case $\lambda > 0$:

$$(2.36) \quad I \leq \int_0^\infty (|\text{Re}(z) + t| e^{-(\text{Re}(z)+t)/\lambda} + |\text{Im}(z)| e^{-(\text{Re}(z)+t)/\lambda})^\lambda dt$$

As $\eta(z) \leq \varepsilon$, $\text{Re}(z) \geq -\log(\varepsilon) > 0$, the first summand is bounded above by $(-\log(\varepsilon) + t) \varepsilon^{1/\lambda} e^{-t/\lambda}$, and the second by $\varepsilon^{1/\lambda} e^{-t/\lambda}$, so

$$(2.37) \quad I \leq \varepsilon \int_0^\infty (1 - \log(\varepsilon) + t)^\lambda e^{-t} dt$$

$$(2.38) \quad \leq C \varepsilon e^{(1-\log(\varepsilon))/2}$$

$$(2.39) \quad = C \sqrt{e\varepsilon}$$

where $C = C(\lambda)$ is a constant.

Case $\lambda \leq 0$:

$$(2.40) \quad I \leq \int_0^\infty e^{-\text{Re}(z)-t} (\max\{1, |\text{Im}(z)|\})^\lambda dt$$

$$(2.41) \quad = \eta(z) \leq \varepsilon$$

To prove the claim, it remains to show that for different values of z_1, z_2 , the corresponding horizontal curves starting at z_1, z_2 have vanishing distance. Let $\beta(t) = x + it$, $t \in [y_1, y_2]$, $x > \rho$, then

$$(2.42) \quad l(\beta) = \int_{y_1}^{y_2} e^{-x} h(x+it) dt$$

$$(2.43) \quad \leq C_2 e^{-x} \int_{y_1}^{y_2} |x+it|^\lambda dt \xrightarrow{x \rightarrow \infty} 0$$

Which completes the proof of the lemma. \square

Proof of Proposition 2.5. — After a change of coordinates $w = e^{\pi i/n}/z$, φ is of the form

$$(2.44) \quad \varphi = e^{-w^n} h(w) dw^2$$

with h meromorphic in a neighborhood of ∞ . Let $B = D^{-1}$ and fix $\varepsilon \in (\frac{1}{2}, 1)$. Define sectors

$$(2.45) \quad V_n = \left\{ w \in B \mid \left| \arg(w) - \frac{2\pi k}{n} \right| < \frac{\pi \varepsilon}{n} \right\}$$

and their complement

$$(2.46) \quad A = B \setminus \bigcup_{k=0}^{n-1} V_n.$$

Then for $z = re^{i\phi} \in A$ we find $\cos(n\phi) \leq \cos(\pi \varepsilon) < 0$, hence

$$(2.47) \quad |\varphi| = e^{-r^n \cos(n\phi)} |hdz^2| \geq C|dz|^2$$

for some $C > 0$. Since the sets V_k have positive distance with respect to the standard metric, this shows that every Cauchy sequence in B without limit is eventually contained in one of the V_k .

On V_k we perform a change of coordinates $u = w^n$, so

$$(2.48) \quad \varphi = e^{-2u} f(u) du^2$$

and there are $C_1, C_2 > 0$ with

$$(2.49) \quad C_1 |u|^\lambda \leq \sqrt{|f(u)|} \leq C_2 |u|^\lambda, \quad \lambda = \frac{m-n+1}{2n}$$

where $m = \text{ord}_\infty h$. Applying the previous lemma, the first part of the theorem follows.

By the first part of the proof,

$$(2.50) \quad \bar{B} =: B \cup \{a_1, \dots, a_n\}$$

is a complete metric space. Let $a = a_k$, then we must show that a is an ∞ -angle singularity of \bar{B} .

As V_k is simply connected, we may choose a holomorphic $f : V_k \rightarrow \mathbf{C}$ with $(df)^2 = \varphi$. Choose $r > 0$ such that $D_{2r}^*(0) \subset V_k$. Then for any path α in D_{2r}^* with endpoint $z_1, z_2 \in D_r^*(0)$ we compute

$$(2.51) \quad |f(z_2) - f(z_1)| = \left| \int_{z_1}^{z_2} \sqrt{\varphi} \right| \leq \int_\alpha \sqrt{|\varphi|}$$

hence

$$(2.52) \quad |f(z_2) - f(z_1)| \leq d(z_1, z_2)$$

so that f extends to a continuous map on $V_k \cup \{a\}$.

Without loss of generality, $f(a) = 0$. We claim that f restricts to a covering $D_r^*(a) \rightarrow D_r^*(0)$. By (2.52), the image of $D_r^*(a)$ under f is contained in $D_r(0)$. The radius of injectivity at $z \in D_r^*(a)$ is $r_1 = d(z, a)$. It cannot be larger, as the singular point a would have to be contained in any disk of radius $> r_1$ centered at z , and it cannot be smaller by completeness of \bar{B} , and since any geodesic may be extended until it hits a or leaves $D_{2r}(a)$. We conclude that $D_{r_1}(z)$ is mapped isometrically to $D_{r_1}(f(z))$ by f , and thus $r_1 = |f(z)|$. This shows that f restricts to a covering $D_r^*(a) \rightarrow D_r^*(0)$.

From the proof of the lemma, it follows that a admits a fundamental system of neighborhoods W with the property that $W \setminus \{a\}$ is simply-connected. Thus $D_r^*(a)$ must be simply connected, and $f : D_r^*(a) \rightarrow D_r^*(0)$ is a universal covering. \square

Next we prove a kind of converse to the previous proposition, which is the main result of this subsection.

Theorem 2.3. — *Let X be a flat surface of finite type. Then there is a compact Riemann surface C with quadratic differential φ of exponential type giving $C \setminus D(\varphi)$ the structure of the smooth part of a flat surface isomorphic to X_{sm} .*

Proof. — Let $K := \text{Core}(X)$, which we may assume to contain more than one point. Define a *boundary walk* to be a piecewise geodesic path which follows the boundary of K so that $S \setminus K$ lies to the right (and possibly also to the left), takes the rightmost possible direction at every singularity, and is maximal with these properties. By rightmost direction we mean the following. From each singularity there is a finite set E of directions which lie in ∂K , and they are either cyclically or totally ordered. So the direction from which we approach the singularity has some successor, by definition the rightmost direction, or is a maximal element, and the boundary walk ends. Hence a boundary walk is either closed or starts and ends at infinite-angle singularities, and there is a unique boundary walk starting/ending at each infinite-angle singularity. If the boundary walk starts and ends at an infinite-angle singularity then t cuts X into two pieces, one of which is topologically a disk lying completely outside the core K . This follows from the description of X in terms of the core above.

Let $s_i \in X_{\text{sg}}$, $i \in \mathbf{Z}/k$ be a cyclic sequence of infinite-angle singularities so that there is a boundary path from s_i to s_{i+1} .

Choose a closed embedded curve α on X_{sm} which is a smoothing of the cyclic sequence of boundary paths and cuts X into two pieces X' , X'' so that X' contains the s_i and no other singularities. The smooth part of X' is topologically an annulus. It consists, up to modifications coming from the smoothing, of several pieces of the complement of K which are topological disks. The smoothing cases these disks to be glued together at infinite-angle singularities to form X' . To complete the proof, it suffices to show that X' corresponds to some punctured neighborhood of an exponential singularity.

Fix some direction r_i from the singularity s_i . For each integer $n > 0$ consider the pair of geodesics L_i, R_i which start at s_i in the directions $r_i + n\pi/2$ and $r_i - n\pi/2$ respec-

tively. For n sufficiently large, all geodesics L_i, R_i can be extended indefinitely in X' . We construct a flat surface X_n as follows. Cut X' along L_i, R_i and remove the (contractible) pieces left of each L_i and right of each R_i . Then glue each L_i to R_i to obtain X_n . Note that X_n has k conical singularities, all with cone angle $n\pi$. Let $U_n \subset X_n$ be the complement of the geodesic rays along which X_n was glued, including the singularities. If we also denote the corresponding subset of X' by U_n , then

$$(2.53) \quad U_n \subset U_{n+1}, \quad \bigcup_n U_n = X'_{\text{sm}}.$$

By the previous results about flat surfaces without infinite-angle singularities, there is a biholomorphic $f_n : X_n \rightarrow \mathbf{D}^*$ and meromorphic quadratic differential φ_n on \mathbf{D} with a pole at 0 and k zeros, such that f_n becomes an isometry of flat surfaces. Since $\mathbf{D} \subset \mathbf{C}$ is bounded, the f_n form a normal family and, after passing to a subsequence, converge to a holomorphic $f : X'_{\text{sm}} \rightarrow \mathbf{C}$. For degree reasons, f cannot be constant and consequently has image contained in \mathbf{D}^* . By a standard application of Hurwitz's theorem, f is injective as the limit of injective functions. We wish to show that the image of f is all of \mathbf{D}^* , and that the φ_n converge to a quadratic differential with exponential singularity at the origin.

Let $g_n(z)dz^2 = \varphi_n(z)$ and consider $h_n = (\partial g_n / \partial z) / g_n$, then h_n is meromorphic on the unit disk with $k + 1$ simple poles, one of which is at the origin. Convergence of f_n implies convergence of h_n , possibly after passing to a subsequence, to a meromorphic function h with $k + 1$ poles, counted with multiplicity. As X' has k infinite-angle singularities we must have a single pole of order $k + 1$ at the origin, so that φ_n converge to φ of the form $e^{p(z)}q(z)dz^2$ with p, q meromorphic and p having a pole of order k at the origin. Completeness of X requires the image of f to be \mathbf{D}^* . \square

2.6. Voronoi and Delaunay partitions

In this subsection we review the Voronoi and Delaunay partitions of a flat surface, cf. Masur–Smillie [35]. They will be used later to construct convergent sequences of flat surfaces.

Let X be a flat surface, not necessarily of finite type, but so that every component contains a conical point or marked point (i.e. a point in X_{sg}). For $p \in X$ let $\rho_p = d(p, X_{\text{sg}}) \geq 0$ be the distance to the closest conical point. Further, define $\mu_p \geq 1$ to be the number of directions from p in which the geodesic of length ρ_p starting at p ends at a conical point if $\rho_p > 0$, and $\mu_p = 1$ if $\rho_p = 0$. In other words, ρ_p is the radius of the largest isometrically immersed disk centered at p , and μ_p is the number of points on its boundary mapped to X_{sg} . The Voronoi cells are the connected components of the level sets of μ_p . Voronoi 2-cells, 1-cells, and 0-cells are components with $\mu_p = 1, \mu_p = 2,$ and $\mu_p \geq 3$ respectively. The 2-cells are contractible and in bijection with X_{sg} . The 1-cells are geodesically embedded open intervals, possibly of infinite length. The union of the 0-cells is a discrete subset of X .

The Delaunay polygonal subdivision is in a sense dual to the Voronoi partition. Given $p \in X$ let D be the largest isometrically immersed disk centered at p , and let $C(p) \subset X$ be the image of the convex hull (in D) of the points on the boundary of D which map to X_{sg} . Note that when p is contained in a Voronoi 2-cell, 1-cell, or 0-cell, then $C(p)$ is a point, edge, or polygon respectively, and $C(p)$ only depends on the component of the Voronoi partition. The sets $C(p)$ provide a polygonal subdivision of $\text{Core}(X)$. For a generic flat surface all polygons are triangles. Although the Voronoi and Delaunay cell decompositions are combinatorially dual, a Delaunay edge need not always intersect its Voronoi dual.

For a flat surface X of finite type consider the quantity

$$(2.54) \quad \rho(X) := \sup\{\rho(p) \mid p \in \text{Core}(X)\} = d_H(X_{\text{sg}}, \text{Core}(X))$$

where d_H denotes the Hausdorff distance of closed subsets. We use $\rho(X)$ to get an upper bound on the lengths of Delaunay edges.

Lemma 2.2. — *Let X be a flat surface of finite type, and $C \geq 0$ the maximal length of any saddle connection on the boundary of $\text{Core}(X)$. Then the length of any Delaunay edge is bounded above by $\max(C, 2\rho(X))$.*

Proof. — Let e be a Delaunay edge and $p \in X$ with $C(p) = e$. The geodesic e is a chord of the largest immersed disk, D , centered at p . If $p \in \text{Core}(X)$, then $\text{length}(e) \leq 2\rho_p \leq 2\rho(X)$. On the other hand, if $p \notin \text{Core}(X)$, then the preimage of the boundary of $\text{Core}(X)$ is either one or two chords in D , one of which has length greater than or equal to $\text{length}(e)$, and less than or equal to C . \square

3. The topological Fukaya category of a surface

In this section we give an elementary definition of the (topological) Fukaya categories of surfaces, and prove several results about them. First, in Section 3.1, we collect some definitions concerning A_∞ -categories. In the following subsection we introduce a class of marked surfaces S which is convenient for the definition of $\mathcal{F}(S)$, which is given in Section 3.3 in terms of explicit A_∞ -structures. In Section 3.4 we describe $\mathcal{F}(S)$ in terms of graded quivers with quadratic monomial relations, a central tool in what follows. This is possible only for some S , but the remaining categories are obtained by localization, as is discussed in Section 3.5. In the final subsection, we show that $\mathcal{F}(S)$ is the category of global sections of a certain cosheaf of categories on a ribbon graph in S .

3.1. Preliminaries on A_∞ -categories

We will need some basic definitions of the A_∞ language, for more details see e.g. [32, 44].

An A_∞ -category, \mathcal{A} , consists of a set $\text{Ob}(\mathcal{A})$ of objects, for each pair $X, Y \in \text{Ob}(\mathcal{A})$ a \mathbf{Z} -graded vector space $\text{Hom}(X, Y)$, and structure maps

$$\mu^n : \text{Hom}(X_{n-1}, X_n) \otimes \cdots \otimes \text{Hom}(X_0, X_1) \rightarrow \text{Hom}(X_0, X_n)[2 - n]$$

for each $n \geq 1$, satisfying the A_∞ -relations

$$(3.1) \quad \sum_{i+j+k=n} (-1)^{\|a_k\|+\cdots+\|a_1\|} \mu^{i+1+k}(a_n, \dots, a_{n-i+1}, \mu^j(a_{n-i}, \dots, a_{k+1}), a_k, \dots, a_1) = 0$$

where $\|a\| = |a| - 1$ is the reduced degree.

An A_∞ -category \mathcal{A} is *strictly unital* if for every object X there is a $1_X \in \text{Hom}^0(X, X)$ such that

$$(3.2) \quad \mu^1(1_X) = 0$$

$$(3.3) \quad \mu^2(a, 1_X) = (-1)^{|a|} \mu^2(1_Y, a) = a, \quad a \in \text{Hom}(X, Y)$$

$$(3.4) \quad \mu^k(\dots, 1_X, \dots) = 0 \quad \text{for } k \geq 3$$

Strictly unital A_∞ -categories with $\mu^k = 0$ for $k \geq 3$ correspond to small dg-categories with

$$(3.5) \quad da = (-1)^{|a|} \mu^1(a), \quad ab = (-1)^{|b|} \mu^2(a, b).$$

Indeed, in this case the first three A_∞ -relations correspond to $d^2 = 0$, the Leibniz rule, and associativity of the product.

An A_∞ -functor $F : \mathcal{A} \rightarrow \mathcal{B}$ is given by a map $\text{Ob}\mathcal{A} \rightarrow \text{Ob}\mathcal{B}$ and multilinear maps

$$(3.6) \quad F^d : \text{Hom}_{\mathcal{A}}(X_{d-1}, X_d) \otimes \cdots \otimes \text{Hom}_{\mathcal{A}}(X_0, X_1) \rightarrow \text{Hom}_{\mathcal{B}}(FX_0, FX_d)[1 - d]$$

for every $d \geq 1$, $X_0, \dots, X_d \in \text{Ob}\mathcal{A}$, satisfying the relations

$$(3.7) \quad \sum_{r \geq 1} \sum_{s_1 + \cdots + s_r = d} \mu_{\mathcal{B}}^r(F^{s_r}(a_d, \dots, a_{d-s_r+1}), \dots, F^{s_1}(a_{s_1}, \dots, a_1))$$

$$(3.8) \quad = \sum_{m, n} (-1)^{\|a_n\|+\cdots+\|a_1\|} F^{d-m+1}(a_d, \dots, a_{n+m+1}, \mu_{\mathcal{A}}^m(a_{n+m}, \dots, a_{n+1}), a_n, \dots, a_1).$$

We will mostly consider strict A_∞ -functors, i.e. those with $F^d = 0$ for $d > 1$.

Twisted complexes

We briefly recall the construction of the category of twisted complexes, $\text{Tw}\mathcal{A}$, over an A_∞ -category \mathcal{A} . Twisted complexes can be thought of as formal vector bundles with flat connection.

The first step is to form $\text{add}\mathbf{Z}\mathcal{A}$ whose objects are formal sums of tensor products of the form

$$(3.9) \quad V = \bigoplus_{X \in \text{Ob}\mathcal{A}} V_X \otimes X$$

with V_X finite-dimensional graded vector spaces, zero for all but finitely many X . Morphism spaces are defined by

$$(3.10) \quad \text{Hom}(V \otimes X, W \otimes Y) = \text{Hom}(V, W) \otimes \text{Hom}(X, Y)$$

and additivity. When extending the μ^k a Koszul sign gets introduced:

$$(3.11) \quad \mu^k(\phi_k \otimes a_k, \dots, \phi_1 \otimes a_1) = (-1)^{\sum_{i < j} |\phi_i| |a_j|} \phi_k \cdots \phi_1 \otimes \mu^k(a_k, \dots, a_1)$$

Note that $\text{add}\mathbf{Z}\mathcal{A}$ has a natural shift functor and formal finite direct sums. Strictly speaking, $\text{Ob}(\text{add}\mathbf{Z}\mathcal{A})$ is not a set, but since the category of finite-dimensional vector spaces is essentially small, this is a non-issue.

An object in $\text{Tw}\mathcal{A}$ is given by a pair (V, δ) with $V \in \text{add}\mathbf{Z}\mathcal{A}$ and $\delta \in \text{Hom}^1(V, V)$. The first condition is that there is a direct sum decomposition of V (i.e. of $\bigoplus V_X$ as an $\text{Ob}\mathcal{A} \times \mathbf{Z}$ -graded vector space) so that δ is strictly upper triangular. This ensures that $\mu^k(\delta, \dots, \delta) = 0$ for k big, so that the second condition, which is

$$(3.12) \quad \sum_{k \geq 1} \mu^k(\delta, \dots, \delta) = 0$$

makes sense. Morphism spaces are just

$$(3.13) \quad \text{Hom}((V, \delta), (W, \epsilon)) = \text{Hom}(V, W)$$

and the structure maps are

$$(3.14) \quad \mu^k(a_k, \dots, a_1) = \sum_{n_0, \dots, n_k \geq 0} \mu^{k+n_0+\dots+n_k}(\underbrace{\delta_k, \dots, \delta_k}_{n_k \text{ times}}, a_k, \dots, a_1, \underbrace{\delta_0, \dots, \delta_0}_{n_0 \text{ times}})$$

where $a_i \in \text{Hom}((V_{i-1}, \delta_{i-1}), (V_i, \delta_i))$. Again, this sum is actually finite by our requirement on the δ_i .

The main fact we need about $\text{Tw}\mathcal{A}$ is that its homotopy category, $H^0(\text{Tw}\mathcal{A})$, is triangulated. When we write $K_0(\text{Tw}\mathcal{A})$, we always mean the K_0 -group of this triangulated category.

3.2. Surfaces with marked boundary

In this subsection we define a class of surfaces which will be convenient for our definition of the Fukaya category, as well as cut and paste constructions. Topologically, they are related to the real blow-ups of flat surfaces by gluing several infinite ends.

We will work with surfaces with corners, S , and denote their 0- and 1-dimensional strata by $\partial_0 S$ and $\partial_1 S$ respectively. The boundary is $\partial S = \partial_0 S \sqcup \partial_1 S$ set-theoretically. A *marked surface* is a surface S with corners together with a subset $M \subset \partial S$ which is the closure of a collection of components of $\partial_1 S$ such that if B is a connected component of ∂S then the following holds: If B is smooth, then it belongs entirely to M , and if B contains corners, then every other component of $\partial_1 S \cap B$ belongs to M , i.e. $\partial_0 S$ is exactly the boundary of M in ∂S . In particular, if S is compact then each component of ∂S is either a circle which belongs entirely to M or a sequence of intervals which alternately belong to M and its complement.

An *arc* in a marked surface S is an embedded closed interval intersecting M transversely in the endpoints, and not isotopic to an embedded interval in M (the isotopy keeping endpoints in M). We generally identify isotopic arcs. A *boundary arc* is an arc which is isotopic to the closure of a component of $\partial S \setminus M$. Arcs which are not boundary arcs are referred to as *internal arcs*. An *arc system* in (S, M) is a collection of pairwise disjoint and non-isotopic arcs. The arc system is *full* if it includes all the boundary arcs and cuts the surface into contractible, relatively compact components (polygons).

A map of marked surfaces $f : (S_1, M_1) \rightarrow (S_2, M_2)$ is an orientation preserving immersion with $f(M_1) \subset M_2$ mapping boundary arcs of S_1 to disjoint non-isotopic arcs in S_2 . The condition insures that a full arc system in S_2 which includes all the images of boundary arcs under f can be lifted to a full arc system on S_1 . It is however not closed under composition in general.

Combinatorial description: graded ribbon graphs

Let (S, M, η) be a graded marked surface and A a full system of graded arcs on it. Such a collection of arcs is dual to a graph Γ on S having an edge for each element of A and an n -valent vertex for each $2n$ -gon cut out by A (see Figure 3). By definition, a *half-edge* is an edge of Γ with orientation pointing towards a vertex of the graph (but not a boundary arc). Since the graph is embedded in an oriented surface, there is a clockwise cyclic order on the set of half-edges pointing towards a given vertex (i.e. Γ has the structure of a *ribbon graph*). Moreover, for any half-edge h the grading determines an integer $d(h)$ as follows. Let $\sigma(h)$ be the half edge following h in the cyclic order. Then $h, \sigma(h)$ correspond to a pair of arcs α, β in A . Let $c : [0, 1] \rightarrow M$ be the embedded curve which starts on α , ends on β , follows the boundary with S to the right, and bounds the polygon corresponding to the vertex that $h, \sigma(h)$ point to. Also choose an arbitrary grading on c . Then

$$(3.15) \quad d(h) = i_{c(0)}(\alpha, c) - i_{c(1)}(\beta, c)$$




FIG. 3. — Full system of arcs in a marked surface and dual ribbon graph (*dashed*)

is independent of the choice of grading on c . The numbers $d(h)$ have the property that if v is a vertex of Γ , then

$$(3.16) \quad \sum_h d(h) = \text{val}(v) - 2$$

where the sum is over all half-edges pointing to v and $\text{val}(v)$ is the valency of v .

Conversely, given a ribbon graph together with integers $d(h)$ for every half-edge h satisfying (3.16) one constructs a graded marked surface with arc system by gluing polygons with suitable grading foliation which is tangent to the arcs. This is a convenient way to specify any compact graded marked surface by a finite amount of data.

3.3. Minimal A_∞ -category of an arc system

Fix a field of scalars \mathbf{K} throughout. Let $S = (S, M, \eta)$ is a graded marked surface with system of graded arcs A . We define a strictly unital A_∞ -category $\mathcal{F}_A(S)$ with $\mu^1 = 0$.

- *Objects*: The set of arcs in A .
- *Morphisms*: A *boundary path* is a non-constant path in M which follows the reverse orientation of the boundary (i.e. the surface lies to the right). Given arcs X and Y , a basis of morphisms from X to Y is given by boundary paths, up to reparameterization, starting at an endpoint of X and ending at an endpoint of Y , as well as the identity morphism if $X = Y$. The degree of a boundary path a from p to q joining arcs X and Y is by definition

$$(3.17) \quad |a| = i_p(X, a) - i_q(Y, a)$$

for arbitrary grading of a .

- *Composition*: Let a, b be boundary paths defining morphisms from X to Y and Y to Z respectively. If a and b are composable then $(-1)^{|a|} \mu^2(b, a) = a \cdot b$, otherwise $\mu^2(b, a) = 0$.

- *Higher operations:* Consider a marked surface (S', M') which is topologically a closed disk with M' having $n \geq 3$ components. Let a_1, \dots, a_n be the distinct boundary paths (components of M') ending in boundary arcs in clockwise order. Given a map $f : (S', M') \rightarrow (S, M)$ sending all boundary arcs of (S', M') to arcs in A we get a sequence $f \circ a_1, \dots, f \circ a_n$ of boundary paths in (S, M) , and call any such sequence a *disk sequence*.

We define higher A_∞ -operations so that if a_1, \dots, a_n is a disk sequence then

$$(3.18) \quad \mu^n(a_n, \dots, a_1 b) = (-1)^{|b|} b$$

for basis morphisms b with $a_1 b \neq 0$, and

$$(3.19) \quad \mu^n(b a_n, \dots, a_1) = b$$

for paths b with $b a_n \neq 0$, and μ^n vanishes on all sequences of paths not of the above forms. The lemma below ensures that this is really well-defined.

Lemma 3.1. — *For a sequence of composable basis morphisms a_n, \dots, a_1 there is at most one factorization $a_1 = a'_1 b$ with a'_1, a_2, \dots, a_n a disk sequence. If such a factorization exists with b not an identity, then there is no factorization $a_n = c a'_n$ with $a_1, \dots, a_{n-1}, a'_n$ a disk sequence and c not an identity. The dual statement also holds.*

Proof. — To see the first statement note that if we have such a factorization and α_i is the arc on which a_i ends, then the concatenation $a'_1 \cdot \alpha_1 \cdots a_n \cdot \alpha_n$ is a null-homotopic loop. Hence, b must be homotopic (relative endpoints) to the concatenation $a_1 \cdot \alpha_1 \cdots a_n \cdot \alpha_n$ and is thus uniquely determined as a morphism.

For the second statement, assume such a factorization exists. Then $c = b$ by the same argument as before, but b and a_n end at different endpoints of α_n , contradicting $a_n = c a'_n$. \square

Proposition 3.1. — *With the structure defined above, $\mathcal{F}_A(S)$ is a strictly unital A_∞ -category.*

Proof. — Let a_n, \dots, a_1 be a composable sequence of morphisms corresponding to boundary paths. The claim is that the A_∞ -equation (3.1) holds. For $n = 1, 2, 3$ there is nothing to prove, so we may assume $n \geq 4$. The following types of non-zero terms cancel pairwise, ignoring signs for the moment:

$$\begin{aligned} & \mu^2(\mu^i(a_n, \dots, a_1 b), c) \\ & \mu^i(a_n, \dots, a_2, \mu^2(a_1 b, c)) \\ & \mu^2(a, \mu^i(b c_n, \dots, c_1)) \\ & \mu^i(\mu^2(a, b c_n), c_{n-1}, \dots, c_1) \\ & \mu^i(a_i, \dots, a_{j+1}, \mu^k(a_j b_k, \dots, b_1), a_{j-1}, \dots, a_1 c) \end{aligned}$$

$$\begin{aligned} & \mu^{i+k-2}(a_i, \dots, a_j b_k, \dots, \mu^2(b_1, a_{j-1}), a_{j-2}, \dots, a_1 c) \\ & \mu^i(a_i, \dots, a_{j+1}, \mu^k(b_k, \dots, b_1 a_j), a_{j-1}, \dots, a_1 c) \\ & \mu^{i+k-2}(a_i, \dots, a_{j+2}, \mu^2(a_{j+1}, b_k), \dots, b_1 a_j, a_{j-1}, \dots, a_1 c) \end{aligned}$$

as well as variants of the previous two with c on the other side, and

$$\begin{aligned} & \mu^i(a_i, \dots, a_2, \mu^j(a_1 b c_j, c_{j-1}, \dots, c_1)) \\ & \mu^j(\mu^i(a_i, \dots, a_2, a_1 b c_j), c_{j-1}, \dots, c_1) \end{aligned}$$

All non-zero terms which can appear must indeed belong to one of the above pairs. To check that the signs are in fact opposite, one uses that

$$(3.20) \quad \sum \|a_i\| = -2$$

for any disk sequence a_1, \dots, a_n . □

There are two special cases for which the above description needs to be somewhat amended, as disks also contribute to μ^2 . The first is that of a square, in which the two boundary arcs are isotopic. They represent the same object X , which has $\text{End}(X) = \mathbf{K}$. The second case is when S is a cylinder and $M = \partial S$. Up to isotopy there is again only one arc, and the corresponding object X has $\text{End}(X) = \mathbf{K}[z^{\pm 1}]$, with the degree of z depending on Ω .

Remark. — The A_∞ -category $\mathcal{F}_A(S)$ was considered by Bocklandt [9], at least in the case when S has only marked boundary circles. He shows that it computes the partially wrapped Fukaya category of S as defined by Abouzaid–Seidel [2]. We expect the proof to extend to the case of partially wrapped Fukaya categories of Auroux [5].

Certain maps between graded marked surfaces induce functors between the respective Fukaya categories.

Proposition 3.2. — *Let $f : S_1 \rightarrow S_2$ be a map of graded marked surfaces and A_i and arc system on S_i such that f maps graded arcs in A_1 to graded arcs in A_2 . Then f induces a strict A_∞ -functor f_* from $\mathcal{F}_{A_1}(S_1)$ to $\mathcal{F}_{A_2}(S_2)$.*

Proof. — Since f sends arcs in A_1 to arcs in A_2 and sends boundary arcs in S_1 to boundary arcs in S_2 , it is clear how to define f_* on objects and morphisms (and f_* is faithful, but not full in general). To see that f_* is a strict A_∞ -functor note that if D is an immersed disk in S_2 whose boundary consists of arcs and boundary curves, then if ∂D lifts to S_1 , so does D . □

The simplest example

Consider the case when $S = (S, M, \eta)$ is a graded marked surface which is topologically a disk and M has $n \geq 3$ components, and A is the arc system containing exactly

the boundary arcs, with some arbitrary grading. The A_∞ -category $\mathcal{F}_A(S)$ has objects E_k , $k \in \mathbf{Z}/n$, and morphisms $a_k : E_k \rightarrow E_{k+1}$ of some degrees $|a_k| \in \mathbf{Z}$ with $\sum |a_k| = n - 2$ which, together with the identity morphisms, form a basis of all morphisms. The only non-zero A_∞ -terms come from strict unitality and

$$(3.21) \quad \mu^n(a_{k+n-1}, \dots, a_k) = 1_{E_k}.$$

See also [36] for a discussion of these categories.

We make a simple observation which will be essential in what follows. Namely, that the twisted complex

$$E_1 \xrightarrow{a_1} E_2[\|a_1\|] \longrightarrow \dots \xrightarrow{a_{n-2}} E_{n-1}[\|a_1\| + \dots + \|a_{n-2}\|]$$

is isomorphic to $E_n[-|a_n|]$, the inverse isomorphisms being given by a_{n-1} and a_n . In other words, $n - 1$ of the boundary arcs already generate all of $\text{Tw}(\mathcal{F}_A(S))$. This also shows that $\text{Tw}(\mathcal{F}_A(S))$ is the bounded derived category of an A_{n-1} quiver with all arrows oriented in the same direction. To set up this equivalence send the objects E_1, \dots, E_{n-1} to the simple objects corresponding to vertices of the quiver, up to some shift.

Morita invariance

Fix a graded marked surface S . If $A \subset B$ are full arc systems, i.e. all arcs in A also belong to B , then it is evident from the definition that $\mathcal{F}_A(S)$ is a full subcategory of $\mathcal{F}_B(S)$.

Lemma 3.2. — *The inclusion functor $\mathcal{F}_A(S) \rightarrow \mathcal{F}_B(S)$ for full arc systems $A \subset B$ is a Morita equivalence.*

Proof. — We can assume that B is obtained from A by adding a single graded arc Y . Let D be one of the disks cut out by B and with Y on its boundary. Considering D as a graded marked surface with minimal full arc system, we have a morphism $f : D \rightarrow S$ inducing a functor of A_∞ -categories. From the discussion above, we know that in $\mathcal{F}(D)$ the object corresponding to Y is isomorphic to a twisted complex over the other boundary arcs, Y_1, \dots, Y_n , of D . As A is a full arc system, there is exactly one boundary of D which gets mapped to Y , hence Y_1, \dots, Y_n get mapped to arcs in A . Applying the functor f_* we thus find that Y is isomorphic to a twisted complex over objects in $\mathcal{F}_A(S)$. In particular, the inclusion functor $\mathcal{F}_A(S) \rightarrow \mathcal{F}_B(S)$ is a Morita equivalence. \square

Proposition 3.3. — *The Morita equivalence class of $\mathcal{F}_A(S)$ is independent of A . The equivalences are canonical in the higher categorical sense (explained in the proof).*

Proof. — Suppose A, B are full systems of graded arc. As established in the previous lemma we get a canonical Morita equivalence when $A \subset B$. The same is also clearly true when A and B just differ by grading shift.

Consider the set \mathcal{A} of full arc systems up to isotopy on (S, M) , partially ordered by inclusion. If we view \mathcal{A} as a category in the usual way, then the arguments above show that we get a functor $A \mapsto \mathcal{F}_A(S)$ from \mathcal{A} to the category of strictly unital A_∞ categories and Morita equivalences. From contractability of the classifying space $|\mathcal{A}|$ of full arc systems or ribbon graphs (see [25, 26]), it now follows that the various categories $\mathcal{F}_A(S)$ are canonically Morita equivalent. \square

Thus, $\text{Tw}\mathcal{F}_A(S)$ is essentially independent of A and we may drop it from the notation and just write $\mathcal{F}(S)$, the (*topological*) *Fukaya category* of S .

3.4. Formal generators

Recall that an A_∞ -algebra A is *formal* if it is quasi-isomorphic to its cohomology H^*A . A generator G of an A_∞ -category is *formal* if its endomorphism algebra $\text{End}(G)$ is. We will see in this subsection that $\mathcal{F}(S)$ has a formal generator whenever S has at least one boundary arc (unmarked part of ∂S) in each connected component. The generator is just a direct sum of arcs, and its endomorphism algebra is the path algebra of a quiver with quadratic monomial relations.

To set up some notation, a *graded quiver* Q with *quadratic monomial relations* is given by

- Q_0 —set of vertices
- Q_1 —set of arrows
- $\partial_0, \partial_1 : Q_1 \rightarrow Q_0$ —source and target maps
- $|\cdot| : Q_1 \rightarrow \mathbf{Z}$ —grading of arrows
- $R \subset Q_1 \times_{Q_0} Q_1$ —quadratic monomial relations, i.e. a collection of pairs of composable arrows

Given a ground field \mathbf{K} we may form the path category $\mathbf{K}Q = \mathbf{K}Q/R$, which is a graded linear category with set of objects Q_0 and basis of morphisms given by paths of arrows not containing any of the relations.

Let S be a graded marked surface, compact for simplicity. Note first that if A is any arc system on S , then the graded category obtained from $\mathcal{F}_A(S)$ by forgetting the higher μ^k is of the form $\mathbf{K}Q$ for some graded quiver with quadratic monomial relations. Namely, take as arrows the boundary paths which start and end at arcs of A but do not cross any other endpoints, and quadratic relations coming from composable arrows which do not correspond to composable paths. Furthermore, the higher A_∞ operations of $\mathcal{F}_A(S)$ vanish under the following condition on A : Any disk cut out by A is bounded by a boundary arc of S which does not belong to A . We call such systems of arcs *formal*. Thus, under this condition, the category $\mathcal{F}_A(S)$ is of the form $\mathbf{K}Q$ for some graded quiver with quadratic monomial relations Q . Call a formal system of arcs *full* if it cuts S into disks each of which has exactly a single boundary arc (of itself) not belonging to A . For a full formal system of arcs A the category $\text{Tw}(\mathcal{F}_A(S))$ is quasi-equivalent to $\mathcal{F}(S)$.

Lemma 3.3. — *Let S be a compact and connected graded marked surface. If S has at least one boundary arc, then it has a full formal system of arcs.*

Proof. — This follows from the fact that we can find a system of arcs on S which cuts it into a single disk, the boundary of which must include all boundary arcs of S . This in turn is seen from the dual ribbon graph. If it has more than one vertex we can contract some edges to reduce the number of vertices to one eventually. \square

We say a graded quiver with relations is of *type F1* if it arises as above from a full formal system of arcs. It is not difficult to see that a graded quiver with quadratic monomial relations is of type F1 if and only if

1. There are no cycles $a_1, \dots, a_n, n \geq 1$, with $a_i a_{i+1} = 0$ for all $i \in \mathbf{Z}/n$.
2. Each vertex has at most two incoming and outgoing arrows.
3. Let $a, b \neq c$ be arrows. If ab, ac are defined, then $ab = 0$ or $ac = 0$, but not both. If ba, ca are defined, then $ba = 0$ or $ca = 0$, but not both.

Resolution of the diagonal

Let Q be a graded quiver with quadratic monomial relations and $A = \mathbf{K}Q/R$ be its path algebra. We want to investigate regularity properties of A . Denote the constant path at the vertex i by e_i . There is a splitting of A

$$(3.22) \quad A = \bigoplus_{i \in Q_0} Ae_i, \quad A = \bigoplus_{i \in Q_0} e_i A$$

as a left (resp. right) module over A .

Define $A^{op} \otimes A$ -modules

$$(3.23) \quad M_n = \bigoplus_{\substack{\alpha_1, \dots, \alpha_n \in Q_1 \\ (\alpha_i, \alpha_{i+1}) \in R}} Ae_{\partial_0 \alpha_1} \otimes e_{\partial_1 \alpha_n} A, \quad n \geq 1$$

$$(3.24) \quad M_0 = \bigoplus_{i \in Q_0} Ae_i \otimes e_i A$$

connected by maps $f_n : M_n \rightarrow M_{n-1}$ such that for $a \otimes b$ in the $(\alpha_1, \dots, \alpha_n)$ -component of M_n

$$(3.25) \quad f_n(a \otimes b) = a\alpha_1 \otimes b + (-1)^n a \otimes \alpha_n b$$

and $f_0 : M_0 \rightarrow A, f_0(a \otimes b) = ab$.

Proposition 3.4. — *The sequence of bimodules*

$$(3.26) \quad \dots \longrightarrow M_2 \xrightarrow{f_2} M_1 \xrightarrow{f_1} M_0 \xrightarrow{f_0} A \longrightarrow 0$$

is exact.

A proof can be found in [6].

Proposition 3.5. — *Let Q be a graded quiver with quadratic monomial relations, $A = \mathbf{K}Q/R$ its path algebra.*

1. *Suppose there are no cyclic paths $\alpha_1 \cdots \alpha_n$ in Q with $(\alpha_i, \alpha_{i+1}) \in R$, $i \in \mathbf{Z}/n$, then A is homologically smooth.*
2. *Suppose there are no cyclic paths $\alpha_1 \cdots \alpha_n$ in Q with $(\alpha_i, \alpha_{i+1}) \notin R$, $i \in \mathbf{Z}/n$, then A is proper.*

Proof. — 1. Under the stated condition on cyclic paths, we see that $M_n = 0$ for $n \gg 0$, hence A is perfect as an $A^{\text{op}} \otimes A$ -module.

2. The condition implies that only finitely many paths are non-zero in A , so A has finite rank over \mathbf{K} . \square

We apply these results to graded linear models of $\mathcal{F}(S)$. In this case, the first condition of Proposition 3.5 is always satisfied, while the second is satisfied if S has no boundary components without corners.

Corollary 3.1. — *Let S be a compact graded marked surface without boundary components diffeomorphic to S^1 , then $\mathcal{F}(S)$ is homologically smooth and proper.*

Question. — If S is a compact graded marked surface which has a full arc system that gives an acyclic quiver, then the category $\mathcal{F}(S)$ has a full strong exceptional collection formed by these arcs. Are there simple conditions on S which allows one to find such an arc system?

3.5. Localization

We begin by recalling a version of Drinfeld's construction for strictly unital A_∞ -categories, studied in detail in [33]. Let \mathcal{A} be a strictly unital A_∞ -category over a field \mathbf{K} and $E \in \text{Ob}(\mathcal{A})$. It will be convenient to use the notation $\mathcal{A}(X, Y)$ for $\text{Hom}(X, Y)$ in this subsection. We will define the quotient category $\mathcal{A}/E = \mathcal{B}$ of \mathcal{A} by E , which will again be a strictly unital A_∞ -category. Informally, \mathcal{B} is obtained by freely adjoining a morphism $\varepsilon \in \mathcal{B}^{-1}(E, E)$ with $\mu^1(\varepsilon) = 1_E$. Set

$$(3.27) \quad \mathcal{B}(X, Y) = \mathcal{A}(X, Y) \oplus \left(\mathcal{A}(E, Y) \otimes \mathbf{K}[1] \otimes \left(\bigoplus_{n \geq 0} (\mathcal{A}(E, E) \otimes \mathbf{K}[1])^{\otimes n} \right) \otimes \mathcal{A}(X, E) \right)$$

as a graded vector space. Write generators of the summand

$$\mathcal{A}(E, Y) \otimes \mathbf{K}[1] \otimes (\mathcal{A}(E, E) \otimes \mathbf{K}[1])^{\otimes n-2} \otimes \mathcal{A}(X, E)$$

as

$$a_n \varepsilon a_{n-1} \varepsilon \cdots \varepsilon a_1$$

with $a_n \in \mathcal{A}(E, Y)$, $a_1 \in \mathcal{A}(X, E)$, and $a_i \in \mathcal{A}(E, E)$ for $2 \leq i \leq n - 1$, then

$$(3.28) \quad |a_n \varepsilon \cdots \varepsilon a_1| = |a_n| + \cdots + |a_1| - n + 1$$

$$(3.29) \quad \|a_n \varepsilon \cdots \varepsilon a_1\| = \|a_n\| + \cdots + \|a_1\|.$$

Structure maps are given by

$$(3.30) \quad \begin{aligned} \mu^r(a_{n_r} \varepsilon \cdots \varepsilon a_{n_r-1+1}, \dots, a_{n_1} \varepsilon \cdots \varepsilon a_1) \\ := \sum_{\substack{i+j+k=n_r \\ j+k>n_r-1 \\ k<n_1}} (-1)^{\|a_k\|+\cdots+\|a_1\|} a_{n_r} \varepsilon \cdots \varepsilon \mu^j(a_{j+k}, \dots, a_{k+1}) \varepsilon a_k \varepsilon \cdots \varepsilon a_1 \end{aligned}$$

where $0 = n_0 < n_1 < \cdots < n_r$, $r \geq 1$. All this generalizes in a straightforward manner to quotients by full subcategories $\mathcal{E} \subset \mathcal{A}$.

We return to the setting of surfaces with marked boundary where we consider the following modification. If S is a marked surface and E a boundary arc, we get a new marked surface S' by adding E to the marked part of the boundary and smoothing the two corners on which E ends. We do not quite get a morphism $S \rightarrow S'$, since E would need to map to an arc which is isotopic to a path in M' . For this reason we allow such arcs, which we call *null*, in this subsection. The definition of the A_∞ -category of a system of arcs works as before with the following small change: The category is no longer minimal, and the definition of μ^1 is modeled on the one for the higher μ^i . This means that the boundary path which is isotopic to a null arc E is a disk-sequence of length 1 and has differential the identity morphisms of that arc.

Proposition 3.6. — *Let A be an arc system on a graded marked surface S and $E \in A$ a boundary arc. By the modification as above we get S' with arc system A' containing a null arc. Then there is a natural equivalence of A_∞ -categories*

$$(3.31) \quad \mathcal{F}_A(S)/E \cong \mathcal{F}_{A'}(S')$$

under $\mathcal{F}_A(S)$.

Proof. — Comparing definitions, one notices an evident strict A_∞ -functor

$$(3.32) \quad G : \mathcal{F}_{A'}(S') \rightarrow \mathcal{F}_A(S)/E$$

compatible with the functors from $\mathcal{F}_A(S)$. It is obtained by factoring a boundary path in M' into its pieces which are alternately contained in M and E . The functor G is not

an isomorphism of A_∞ -categories, but becomes one once we drop E from both the target and the source category. By general properties, $E \in H^*(\mathcal{F}_A(S)/E)$ is a zero-object, and it is easily seen that the same is true in $H^*(\mathcal{F}_{A'}(S'))$, implying that G is a quasi-equivalence. \square

Remark. — Localization of topological Fukaya categories of surfaces is also studied in the recent works by Dyckerhoff [20] and Pascaleff–Sibilla [40]. The localizations they consider depend more generally on an open subgraph of a ribbon graph on S .

3.6. Cosheaf of categories

Let S be a graded marked surface with full system of graded arcs A . Dual to A is a ribbon graph G on S which is a deformation retract of the pair $(S, \partial S \setminus M)$. A constructible cosheaf \mathcal{E} of A_∞ -categories on G is given, concretely, by assigning an A_∞ -category \mathcal{C}_v (resp. \mathcal{C}_e) to each internal vertex v (resp. edge e) of G , together with functors $\mathcal{C}_e \rightarrow \mathcal{C}_v$ for every half-edge. We understand this as a cosheaf with values in the higher category of A_∞ -categories with equivalences the Morita-equivalences, so when taking global sections one needs to compute a homotopy colimit.

There is a canonical constructible cosheaf on G , defined as follows. For an edge e , \mathcal{C}_e is the category with a single object corresponding to the arc in A dual to e and endomorphism algebra \mathbf{K} . For a vertex v , let D be the dual polygon cut out by A , then \mathcal{C}_v is the category $\mathcal{F}(D)$ where we take the arc system consisting just of boundary arcs. Thus, \mathcal{C}_v is Morita-equivalent to the path-algebra of an A_n quiver, where $n+1$ is the valency of v . The functors $\mathcal{C}_e \rightarrow \mathcal{C}_v$ are fully-faithful inclusions of boundary arcs of D . By construction, we also have functors $\mathcal{C}_e \rightarrow \mathcal{F}_A(S)$ and $\mathcal{C}_v \rightarrow \mathcal{F}_A(S)$, and all these are compatible. Hence, by universality there is a functor from global sections of \mathcal{E} , $\Gamma(G, \mathcal{E}) \rightarrow \mathcal{F}_A(S)$. The main result of this subsection is that this is a Morita-equivalence.

Theorem 3.1. — *Let S be a graded marked surface with ribbon graph G dual to a full graded system of arcs. Then $\mathcal{F}(S)$ represents global sections of the constructible cosheaf \mathcal{E} on G defined above.*

The category of global sections $\Gamma(G, \mathcal{E})$ is by definition the homotopy colimit of the diagram formed by the $\mathcal{C}_e, \mathcal{C}_v$, and the functors between them. Thus we need to show that $\mathcal{F}(S)$ represents this homotopy colimit. The strategy of the proof will be to first treat the case with boundary arcs where we have a formal generator, then deduce the general case by localization.

Let Q_n be the quiver with $n \geq 1$ vertices, arrows $\alpha_{i,i+1}$ from the i -th to the $(i+1)$ -st vertex of degrees $|\alpha_{i,i+1}| \in \mathbf{Z}$, and quadratic relations $\alpha_{i,i+1}\alpha_{i-1,i} = 0$ for $1 < i < n$. Suppose that $Q = Q_{n_1} \sqcup \cdots \sqcup Q_{n_k}$ is a finite disjoint union quivers of this type, and that P is a finite set with maps $f_1, f_2 : P \rightarrow Q_0$ such that $f_1 \sqcup f_2 : P \sqcup P \rightarrow Q$ is injective. Construct a new quiver with relations R from Q by identifying the vertices $f_1(p)$ and $f_2(p)$ for

each $p \in \mathbf{P}$, and keeping the same arrows and relations. Passing to the associated path-categories $\mathcal{P}, \mathcal{Q}, \mathcal{R}$, where \mathbf{P} is thought of as a quiver without arrows, we get a coequalizer diagram

$$(3.33) \quad \mathcal{P} \begin{array}{c} \xrightarrow{f_1} \\ \xrightarrow{f_2} \end{array} \mathcal{Q} \longrightarrow \mathcal{R}$$

in the naive (1-categorical) sense.

Proposition 3.7. — *The diagram (3.33) is a homotopy coequalizer in the category of dg-categories.*

Proof. — To show this, we will make a cofibrant replacement $\tilde{\mathcal{Q}}$ of \mathcal{Q} making the diagram

$$(3.34) \quad \mathcal{P} \begin{array}{c} \xrightarrow{f_1} \\ \xrightarrow{f_2} \end{array} \tilde{\mathcal{Q}}$$

cofibrant by the condition on $f_1 \sqcup f_2$, and verify the coequalizer of (3.34) is quasi-isomorphic to \mathcal{R} .

The graded linear path-category \mathcal{A} of \mathcal{Q}_n has a cofibrant replacement given by the bar-cobar resolution $\mathcal{B} = \Omega\mathcal{B}\mathcal{A}$. The dg-category \mathcal{B} is described explicitly by a dg-quiver with n vertices, arrows β_{ij} from the i -th to the j -th vertex, $1 \leq i < j \leq n$, with degrees

$$(3.35) \quad |\beta_{ij}| = 1 + \sum_{i \leq k < j} (|\alpha_{k,k+1}| - 1)$$

and differential

$$(3.36) \quad d\beta_{ij} = \sum_{i < k < j} (-1)^{|\beta_{ij}|} \beta_{kj} \beta_{ik}$$

extended to paths by the graded Leibniz rule. The functor $\mathbf{P} : \mathcal{B} \rightarrow \mathcal{A}$ sends $\beta_{i,i+1}$ to $\alpha_{i,i+1}$ and all other arrows to zero. The right inverse \mathbf{I} sending $\alpha_{i,i+1}$ to $\beta_{i,i+1}$ is not multiplicative, but an inverse of \mathbf{P} on the level of cohomology, as the homotopy \mathbf{H} defined by

$$(3.37) \quad \mathbf{H}(\beta_{i_{k-1},i_k} \cdots \beta_{i_1,i_2}) = \begin{cases} (-1)^{|\beta_{i_{k-1},i_k}|} \beta_{i_{k-2},i_k} \cdots \beta_{i_1,i_2} & \text{if } i_k = i_{k-1} + 1, k \geq 3 \\ 0 & \text{else} \end{cases}$$

satisfying $d\mathbf{H} + \mathbf{H}d = 1 - \mathbf{IP}$, shows.

We have a cofibrant replacement $\tilde{\mathcal{Q}} = \Omega\mathcal{B}\mathcal{Q}$ which is described by a disjoint union of the dg-quivers above. The coequalizer, $\tilde{\mathcal{R}}$ of the diagram (3.34) is associated with the dg-quiver we get by identifying pairs of vertices. There is a dg-functor $\tilde{\mathcal{R}} \rightarrow \mathcal{R}$, which we claim is a quasi-isomorphism. To define the homotopy, let β_1, \dots, β_n be non-identity paths in the quiver of $\tilde{\mathcal{Q}}$ which are not composable, but become so in $\tilde{\mathcal{R}}$, so that $b := \beta_n \cdots \beta_1$ is a morphism in $\tilde{\mathcal{R}}$. Note that any non-identity path in the quiver of $\tilde{\mathcal{R}}$ canonically breaks up in such a way. Also define the length of a path, $l(\beta)$, so that $l(\beta_j) = j - i$ and $l(\beta_2\beta_1) = l(\beta_2) + l(\beta_1)$. We define H extending (3.37) in the following way. If $l(\beta_k) = 1$ for $1 \leq k \leq n$, then $H(\beta_n \cdots \beta_1) = 0$. Otherwise, there exists a largest k with $l(\beta_k) > 1$ and

$$(3.38) \quad H(\beta_n \cdots \beta_1) := (-1)^{|\beta_n| + \cdots + |\beta_{k+1}|} \beta_n \cdots H(\beta_k) \cdots \beta_1.$$

Let us check the identity $dH + Hd = 1 - \text{IP}$ in $\tilde{\mathcal{R}}$. If $l(\beta_k) = 1$ for all k , then evidently both sides vanish. Otherwise, we have a largest k with $l(\beta_k) > 1$ and the right-hand side vanishes. On the other hand,

$$(3.39) \quad (dH)(b) = (dH)(\beta_k) + \sum_{j < k} (-1)^{|\beta_k| + \cdots + |\beta_{j-1}|} \beta_n \cdots H(\beta_k) \cdots d\beta_j \cdots \beta_1$$

and

$$(3.40) \quad (Hd)(b) = (Hd)(\beta_k) + \sum_{j < k} (-1)^{|\beta_k| + \cdots + |\beta_{j-1}|} \beta_n \cdots H(\beta_k) \cdots d\beta_j \cdots \beta_1$$

so that the left-hand side also vanishes.

We conclude that the original coequalizer diagram (3.33) is quasi-isomorphic to a cofibrant one, and hence also a homotopy coequalizer. \square

Proof of Theorem 3.1. — Let S be a graded marked surface, which we may as well assume to be connected. Suppose first that S has at least one boundary arc. Then S has a full formal system of arcs A with dual ribbon graph G . Recall that $\Gamma(G, \mathcal{E})$ is the homotopy colimit of a diagram with the various categories $\mathcal{C}_e, \mathcal{C}_v$ and functors between them. Removing the categories \mathcal{C}_e for e an edge ending in a boundary arc from the diagram does not change the colimit, since such \mathcal{C}_e are being included as subcategories into some \mathcal{C}_v only once. In the reduced diagram there is now, by formality of the arc system, at least one object in each \mathcal{C}_v which is not in the image of any functor $\mathcal{C}_e \rightarrow \mathcal{C}_v$. We remove such an object from each \mathcal{C}_v , producing a Morita-equivalent diagram. Note that the resulting diagram is now a diagram of graded linear categories. It is essentially a diagram of the form (3.33) where \mathcal{P} is the coproduct of the \mathcal{C}_e (with \mathcal{P} identified with the set of arcs in A) and \mathcal{Q} is the coproduct of the \mathcal{C}_v . The coequalizer \mathcal{R} is just the graded linear model for $\mathcal{F}(S)$, and we know it is a homotopy coequalizer by the previous proposition.

The previous arguments also show that $\Gamma(G, \mathcal{E})$ is independent of G , and hence that the theorem holds in fact for *any* full system of arcs on a surface with boundary arc. Namely, it suffices to check that global sections do not change when an edge of G is contracted. This can be checked in a neighborhood of that edge, so we can assume that G has only one internal edge and is dual to a full arcs system A on the disk D with single internal arc. Let G' be the graph obtained by contracting the internal edge in G , and \mathcal{E}' its cosheaf of categories. As it is already established that both $\Gamma(G, \mathcal{E})$ and $\Gamma(G', \mathcal{E})$ are represented by $\mathcal{F}(D)$, they must be equivalent.

What remains is the case when S has no boundary arcs. As discussed in Section 3.5, the category $\mathcal{F}(S)$ is a localization of a category $\mathcal{F}(S')$ where S' is obtained from S by inserting a boundary arc on some boundary component. So to complete the proof of the theorem it suffices to check that the corresponding categories of global sections have the same relation.

Choose some ribbon graph G' on S' dual to a full system of graded arcs. To get a ribbon graph G for S we just need to remove the edge ending on the unique boundary arc. This means that in the diagram computing global sections we change some \mathcal{C}_v from type A_n to type A_{n-1} , which is indeed just localization by the boundary arc. Finally, taking the quotient by some object is a special case of a homotopy push out, thus commutes with colimits, so that $\Gamma(G, \mathcal{E})$ is the quotient of $\Gamma(G', \mathcal{E}')$ by the boundary arc. \square

4. Tameness and geometricity

In this section we deal with the problem of classifying objects in $\mathcal{F}(S)$. First, in Section 4.1, we assign objects in $\mathcal{F}(S)$ to certain immersed curves in S with a local system of vector spaces. The purpose of the rest of the section is to prove that we get all objects in this way. In Section 4.2 we introduce *nets* and study their representations. Section 4.3 discusses a minimality condition on twisted complexes which ensures uniqueness up to isomorphism. The proof of the classification is completed in Section 4.4.

4.1. Twisted complexes from curves

Fix a graded marked surface S and a ground field \mathbf{K} . An immersed curve c in S is *unobstructed* if it does not bound an immersed teardrop, which is a map from the closed disk D to the surface which takes ∂D to c and which is a smooth immersion at every point except one point of ∂D (see Figure 4). An *admissible curve* is an unobstructed graded curve c such that one of the following holds:

1. The domain of c is S^1 , the image of c is disjoint from ∂S , and c represents a primitive class in $\pi_1(S)$.
2. The domain of c is $[0, 1]$, c intersects ∂S transversely in M and only in the endpoints, and c is not homotopic relative endpoints to a path in M .




FIG. 4. — Curve bounding a teardrop

Suppose that S is compact so that we have a category $\mathcal{F}(S)$, well defined up to canonical equivalence. The purpose of this subsection is to show that an admissible curve together with a local system of finite dimensional \mathbf{K} -vector spaces (on its domain) gives an equivalence class of objects in $\mathcal{F}(S)$.

First, in the case when S is topologically a disk any admissible curve c is a graded arc. Thus we can find a full arc system A which includes c , so that c is an object of $\mathcal{F}_A(S)$ by definition. The isomorphism class of that object is clearly well-defined in $\mathcal{F}(S)$ independently of the arc system.

Returning to the case of general S , we will first deal with admissible curves c which have domain $[0, 1]$. For any such c we can find a graded marked surface S' which is of disk-type and with a map f to S so that c is the image of an admissible curve \tilde{c} under f . To see this, consider the universal cover \tilde{S} of S and lift a full arc system A on S to \tilde{A} on \tilde{S} . Lift c to \tilde{c} on \tilde{S} and take as S' a closed disk which is cut out by arcs in \tilde{A} and which contains \tilde{c} . Now, \tilde{c} gives an equivalence class of objects in $\mathcal{F}(S')$, and the image under the functor $\mathcal{F}(S') \rightarrow \mathcal{F}(S)$ is independent of the choice of (S') . This follows from the fact that if we have S', S'' as above, then the maps $S' \rightarrow S, S'' \rightarrow S$ both factor through a third $S''' \rightarrow S$ as can be seen by looking at the universal cover again.

Suppose now instead that c is an admissible curve with domain S^1 and local system V of finite-dimensional vector spaces on it. We will follow the same strategy as before and assume first that (S, M) is of annular type, i.e. topologically a compact annulus with corners on each boundary component. Choose a cyclic sequence of disjoint non-isotopic arcs $X_i, i \in \mathbf{Z}/n$ so that at least one connects the two components of ∂S and such that every component of M contains either exactly two endpoints of the arcs, belonging X_i, X_{i+1} for some i , or none of the endpoints (see Figure 5). Thus we get a sequence $a_i, i \in \mathbf{Z}/n$ of distinct boundary paths so that a_i connects endpoints of X_i, X_{i+1} . If we follow $X_0, a_0, X_1, a_1, \dots$ we get a path which after suitable smoothing near the intersection points becomes a simple closed loop isotopic to c . It is possible to choose grading and a local systems on X_i, a_i , so that the smoothed path is isotopic to c as a graded curve with local system. As a result, each a_i will be morphism of degree 1 either from X_i to X_{i+1} or in the other direction. Further we get a vector space V_i of sections of over X_i and parallel transport $T_i : V_i \rightarrow V_{i+1}$ along a_i . Consider the twisted complex

$$(4.1) \quad \left(\bigoplus V_i \otimes X_i, \sum T_i^{\pm 1} \otimes a_i \right)$$




FIG. 5. — Example of allowed sequence of arcs (*dashed*)

where the signs are determined by the direction of the morphisms a_i . This is the object of $\mathcal{F}(S, M, \Omega)$ we assign to c .

Lemma 4.1. — *The equivalence class of the twisted complex constructed depends only on the isotopy class of the graded curve c with local system.*

Proof. — We claim first that we can replace X_1, \dots, X_{n-1} by a single arc Y with the pair X_0, Y giving an isomorphic twisted complex. Indeed, cutting S along X_0 we get a surface of disk type S' in which the sequence of arcs X_1, \dots, X_{n-1} concatenates to a single graded arc Y . The corresponding isomorphism of twisted complexes formed from X_1, \dots, X_{n-1} and Y respectively was established in Section 3.3, and the claimed isomorphism of twisted complexes formed from X_0, \dots, X_{n-1} and X_0, Y follows.

We have reduced the problem to the case of two arcs. Any two pairs of arcs satisfying our requirements are related by Dehn twists (automorphisms) of S . It is clear that the isotopy class of c is invariant under Dehn twists. To finish the proof we need to show that the twisted complex associated with c is invariant under the induced autoequivalence, up to isomorphism. This can be checked by direct computation, or by using the equivalence

$$(4.2) \quad H^0(\text{Tw}(\mathcal{F}_{X_0, Y}(S))) = D^b(\mathbf{P}^1)$$

under which X_0, Y correspond to $\mathcal{O}, \mathcal{O}(1)[-1]$, the Dehn-twist to $\otimes \mathcal{O}(1)$, and the twisted complex assigned to c to a torsion sheaf. \square

We have excluded above the case when one or both boundary components are entirely contained in M . These cases can be handled fairly easily directly, or alternatively by the localization construction of the previous section. The case of general S is handled by finding maps from surfaces of annular type. Here the argument uses the annular covering associated with c instead of the universal covering.

4.2. Representations of nets

In this subsection we classify representations of certain combinatorial structures, called *nets*, based on ideas in the work of Nazarova–Roiter [37] on a class of tame linear algebra problems. We will later use this result to derive the classification of objects in Fukaya categories of surfaces with boundary. The authors are reluctant to introduce yet another combinatorial structure, but it gives the necessary flexibility for the proof of the classification.

A *net* is a quadruple $\mathcal{X} = (A, \alpha, B, \beta)$ where

- A is a finite set,
- α is a fixed-point-free involution on A ,
- B is a finite set with partition $B = \bigsqcup_{i \in \Lambda} B_i$ into totally ordered sets,
- β is a fixed-point-free involution on $\text{Dom}(\beta) \subset B$.

A morphism of nets $f : (A, \alpha, B, \beta) \rightarrow (A', \alpha', B', \beta')$ is given by maps $f_1 : A \rightarrow A', f_2 : B \rightarrow B'$ such that

(4.3)
$$f_1 \circ \alpha = \alpha' \circ f_1$$

(4.4)
$$f_2(\text{Dom}\beta) \subset \text{Dom}\beta'$$

(4.5)
$$f_2(B \setminus \text{Dom}\beta) \subset B' \setminus \text{Dom}\beta'$$

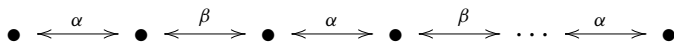
(4.6)
$$f_2|_{\text{Dom}\beta} \circ \beta = \beta' \circ f_2|_{\text{Dom}\beta}$$

(4.7)
$$f_2(B_i) \subset B'_{f_1(i)} \quad \text{and} \quad f_2|_{B_i} \text{ is increasing}$$

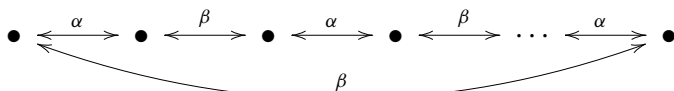
The *height* of a net $\mathcal{X} = (A, \alpha, B, \beta)$ is defined as

(4.8)
$$h(\mathcal{X}) = \max_{i \in \Lambda} |B_i|.$$

Nets of height 1 are disjoint unions of the following types of nets:

(4.9) 

and

(4.10) 

where we have indicated elements of $A = B$ by dots and the action of α, β by arrows. In both cases $|A| \geq 2$.

A *representation of a net* (A, α, B, β) is given by

- a finite-dimensional vector space V with a direct sum decomposition

$$V = \bigoplus_{i \in A/\alpha} V_i,$$

- for each $i \in A$ an increasing exhaustive filtration $\{F_j\}_{j \in B_i}$ on V_i (i.e. $j \leq k \implies F_j \subset F_k$, $F_{\max B_i} = V_i$),
- isomorphisms

$$\phi_i : \text{gr}_i V \rightarrow \text{gr}_{\beta(i)} V, \quad i \in \text{Dom}(\beta)$$

with $\phi_{\beta(i)} = \phi_i^{-1}$. Here $\text{gr}_i V = F_i/F_{<i}$ is the associated graded.

In short, a representation of a net is a collection of vector spaces, each with two filtrations, and with certain subquotients identified. Representations of a net form a linear category with morphisms $(V, F_i, \phi_i) \rightarrow (W, G_i, \psi_i)$ being linear maps $f : V \rightarrow W$ such that $f(F_i) \subset G_i$ and

$$\begin{array}{ccc} \text{gr}_i V & \xrightarrow{\text{gr}_i f} & \text{gr}_i W \\ \downarrow \phi_i & & \downarrow \psi_i \\ \text{gr}_{\beta(i)} V & \xrightarrow{\text{gr}_{\beta(i)} f} & \text{gr}_{\beta(i)} W \end{array}$$

commutes for all $i \in \text{Dom}(\beta)$. We note that the category of representations of a net as in (4.9) is the category of finite dimensional vector spaces, while for (4.10) we get the category of finite dimensional representations of \mathbf{Z} (up to equivalence).

Given a morphism of nets $f = (f_1, f_2) : \mathcal{X} = (A, \alpha, B, \beta) \rightarrow \mathcal{X}' = (A', \alpha', B', \beta')$ with f_2 injective on each B_i , there is an induced *pushforward* functor f_* on the corresponding categories of representations. If $V = (V, F_i, \phi_i)$ is a representation of \mathcal{X} , then we produce a representation $f_* V = W = (W, G_i, \psi_i)$ of \mathcal{X}' with $W = V$ as vector spaces and

$$(4.11) \quad W_j = \bigoplus_{i \in A/\alpha, f_1(i)=j} V_i, \quad j \in A'/\alpha'$$

$$(4.12) \quad G_l = \bigoplus_{i \in A, l \in B'_{f_1(i)}} F_k, \quad \text{where } k = \max\{r \in B_i \mid f_2(r) \leq l\}$$

for which we have

$$(4.13) \quad \text{gr}_j W = \bigoplus_{i \in B, f_2(i)=j} \text{gr}_i V$$

so that we can define

$$(4.14) \quad \psi_j = \bigoplus_{i \in B, f_2(i)=j} \phi_i.$$

If f_2 fails to be injective on some B_i 's, then one can still define f_*V given additional choices on V (i.e. in a non-functorial way). Namely, for each $i \in A$ and $j \in B'_{f_1(i)} \cap \text{Dom}(\beta')$ with

$$(4.15) \quad (f_2|_{B_i})^{-1}(j) \cong \{1, \dots, n\}, \quad n \geq 2$$

choose a splitting of the inclusions $F_1V \subset \dots \subset F_nV$. This gives us an isomorphism

$$(4.16) \quad \text{gr}_1V \oplus \dots \oplus \text{gr}_nV \cong F_nV/F_{<1}V = \text{gr}_jW$$

allowing us to define ψ_j as direct sums of $\psi_1 \oplus \dots \oplus \psi_n$ for various i . Different choices of splittings give isomorphic f_*V .

Theorem 4.1. — *Let \mathcal{X} be a net, V a representation of \mathcal{X} . Then there is a net \mathcal{X}' of height 1, a morphism $f : \mathcal{X}' \rightarrow \mathcal{X}$, and representation V' of \mathcal{X}' such that $f_*V' = V$.*

As a preliminary to the proof of the theorem, we discuss diagrams

$$\begin{array}{ccc} & A & \\ & \curvearrowright & \\ X & \xrightarrow{\quad} & Y \\ & \curvearrowleft & \\ & B & \end{array}$$

of parallel surjective linear maps between finite dimensional vector spaces. First, we have two increasing filtrations on X :

$$(4.17) \quad A^{-1}0 \subset A^{-1}BA^{-1}0 \subset A^{-1}BA^{-1}BA^{-1}0 \subset \dots \subset \underbrace{A^{-1}(BA^{-1})^{k-1}0}_{=:F_k} \subset \dots$$

$$(4.18) \quad B^{-1}0 \subset B^{-1}AB^{-1}0 \subset B^{-1}AB^{-1}AB^{-1}0 \subset \dots \subset \underbrace{B^{-1}(AB^{-1})^{k-1}0}_{=:G_k} \subset \dots$$

They necessarily stabilize with $F_\infty = F_N, G_\infty = G_N$ for N big. In general, these filtrations are not exhaustive, but

$$(4.19) \quad F_\infty = G_\infty.$$

Furthermore, there are inverse isomorphisms

$$\begin{array}{ccc} & B^{-1}A & \\ & \xrightarrow{\quad} & \\ M_{i,j} & \xleftrightarrow{\quad} & M_{i-1,j+1} \\ & \xleftarrow{\quad} & \\ & A^{-1}B & \end{array}$$

for $i \geq 0, j \geq 1$, where

$$(4.20) \quad M_{i,j} = \frac{F_i \cap G_j}{(F_{i-1} \cap G_j) + (F_i \cap G_{j-1})}$$

for $i, j \geq 1$. In particular,

- $\text{Ker}(\mathbf{B})$ has an exhaustive filtration $F \cap \text{Ker}(\mathbf{B})$ with associated graded pieces $M_{i,1}$,
- $\text{Ker}(\mathbf{A})$ has an exhaustive filtration $G \cap \text{Ker}(\mathbf{A})$ with associated graded pieces $M_{1,j}$,
- there are isomorphisms $M_{i,1} \longrightarrow M_{1,i}$.

We proceed with the proof of the theorem.

Proof. — Let $(V, \{F_i\}, \phi_i)$ be a representation of a net $\mathcal{X} = (A, \alpha, B, \beta)$. The proof is by induction over

$$(4.21) \quad \sum_{|B_i| \geq 2} \dim V_i$$

for all nets simultaneously. If $h(\mathcal{X}) = 1$ we are done, so let us assume that $h(\mathcal{X}) \geq 2$. We can also assume, by passing to subsets of A and the B_i , that all associated graded $\text{gr}_i V$ are non-zero and all B_i are non-empty.

Let $r \in A$ with $|B_r| = h(\mathcal{X})$. The goal is to separate a “strand” of the representation going through r , thus decreasing the quantity (4.21). Let $n = \max B_r$ and find the unique $k \in B_{\alpha(r)}$ such that

$$(4.22) \quad F_{<k} + F_{<n} \subsetneq V_r, \quad F_k + F_{<n} = V_r.$$

Set $X_1 = F_{<k} + F_{<n}$, $X_2 = F_k \cap (F_{<k} + F_{<n})$ and consider the refinements of the two filtrations of V_r :

$$(4.23) \quad \cdots \subset F_{<n} \subset X_1 \subset F_n = V_r, \quad \cdots \subset F_{<k} \subset X_2 \subset F_k \subset \cdots$$

Note that the identity map induces an isomorphism of associated graded

$$(4.24) \quad F_n/X_1 \longrightarrow F_k/X_2$$

which follows from the general isomorphism theorem $S/(S \cap T) \cong (S + T)/T$.

Case $\beta(k) \neq n$. The first step is to modify the net and its representation to decrease (4.21). The second step is to show that the pushforward of the modified representation gives the original one.

- If $n \in \text{Dom}(\beta)$, $\beta(n) \in \mathbf{B}_p$, can refine the filtration on V_p to

$$\cdots \subset F_{<\beta(n)} \subset X_3 \subset F_{\beta(n)} \subset \cdots$$

where X_3 is the preimage of $\phi_n(X_1/F_{<n})$ under the quotient map $F_{\beta(n)} \twoheadrightarrow \text{gr}_{\beta(n)}X$.

- Similarly, if $k \in \text{Dom}(\beta)$, $\beta(k) \in \mathbf{B}_p$, can refine the filtration on V_p to

$$\cdots \subset F_{<\beta(k)} \subset X_4 \subset F_{\beta(k)} \subset \cdots$$

where X_4 is the preimage of $\phi_k(X_2/F_{<k})$ under the quotient map $F_{\beta(k)} \twoheadrightarrow \text{gr}_{\beta(k)}X$.

We construct a modified net $\mathcal{X}' = (A', \alpha', B', \beta')$ with

$$(4.25) \quad A' = A \sqcup \{1, 2\}, \quad \alpha'(1) = 2, \quad \alpha'|_A = \alpha$$

and $B'_1 = \{1\}$, $B'_2 = \{2\}$. Moreover, if $n \in \text{Dom}\beta$ we insert an element “3” into B just before $\beta(n)$ with $\beta'(1) = 3$, and similarly if $k \in \text{Dom}\beta$ we insert an element “4” before $\beta(k)$ with $\beta'(2) = 4$. Note that the new elements of B correspond to the additional pieces of the filtrations described above. There is a morphism of nets $f : \mathcal{X}' \rightarrow \mathcal{X}$ with $1 \mapsto n$, $2 \mapsto k$, $3 \mapsto \beta(n)$, $4 \mapsto \beta(k)$.

V determines a representation $V' = (V', F_i, \psi_i)$ of \mathcal{X}' with

$$(4.26) \quad V'_1 = V'_2 = F_n/X_1 \cong F_k/X_2, \quad V'_r = V'_{\alpha(r)} = X_1$$

where we suppress the canonical isomorphism $F_n/X_1 \cong F_k/X_2$ from now on. The filtrations on V'_r are obtained from those on V_r by restriction, and the filtrations containing F_3, F_4 by refinement. The isomorphisms $\psi_n, \psi_k, \psi_1, \psi_2$ and their inverses are induced by ϕ_n, ϕ_k and their inverses respectively. We claim that f_*V' is isomorphic to V . Choose complements

$$(4.27) \quad X_1/F_{<n} \oplus Y_1 = \text{gr}_n V$$

$$(4.28) \quad X_2/F_{<k} \oplus Y_2 = \text{gr}_k V.$$

If $n \in \text{Dom}\beta$, then the choice of Y_1 is equivalent to a choice of complement of $X_3/F_{<\beta(n)} \subset \text{gr}_{\beta(n)}V$, and similarly for Y_2 . We use these complements in the definition of f_*V' .

We claim that we can find a $U \subset F_k$ such that

$$(4.29) \quad F_n \twoheadrightarrow \text{gr}_n V \quad \text{induces an isomorphism } U \rightarrow Y_1$$

$$(4.30) \quad F_k \twoheadrightarrow \text{gr}_k V \quad \text{induces an isomorphism } U \rightarrow Y_2$$

Namely, choose a $U \subset F_k$ with

$$(4.31) \quad X_2 \oplus U = F_k$$

such that (4.30) holds. We also have a direct sum $X_1 \oplus U = V_r$, but (4.29) may fail. However, we can ensure (4.29) after shearing U in the directions of $F_{<k}$, which does not affect (4.30).

The choice of U gives an identification

$$(4.32) \quad V_r = X_1 \oplus U = X_1 \oplus (F_n/X_1) = V'_r \oplus V'_1$$

which we use to define the isomorphism $V \rightarrow f_*V'$. Compatibility with filtrations follows from $U \subset F_k$, and compatibility with ϕ_n, ϕ_k since they are block-diagonal with respect to the chosen splittings of associated graded.

Case $\beta(k) = n$. The overall strategy is similar to the previous case. The main difference is that both F_n/X_1 and F_k/X_2 are now subquotients of the same vector space and hence can have some overlap. As a consequence, we cannot just excise the subquotient $F_n/X_1 \cong F_k/X_2$ as before, but must subdivide further.

We have a diagram of parallel surjections:

$$\begin{array}{ccc} \mathrm{gr}_n V & \xrightarrow{p_1} & F_n/X_1 \\ \cong \downarrow \phi_n & & \downarrow \cong \\ \mathrm{gr}_k V & \xrightarrow{p_2} & F_k/X_2 \end{array}$$

By the discussion preceding the proof we get first of all a pair of non-exhaustive filtrations

$$(4.33) \quad 0 = G_0 \subset G_1 \subset \cdots \subset G_m, \quad 0 = H_0 \subset H_1 \subset \cdots \subset H_m = G_m$$

of $\mathrm{gr}_n V$, where $G_1 = \mathrm{Ker}(p_1)$ and $H_1 = \phi_n^{-1}(\mathrm{Ker}(p_2))$. For the associated double graded M_{ij} , $1 \leq i, j, i+j \leq m+1$, we have isomorphisms $M_{ij} \rightarrow M_{i+1, j-1}$ induced by ϕ_n . Further, we have restricted filtrations on $\mathrm{Ker}(p_1), \mathrm{Ker}(p_2)$ lifting to refinements of the filtrations on V_r :

$$(4.34) \quad \cdots \subset F_{<n} = X_{1,0} \subset X_{1,1} \subset \cdots \subset X_{1,m} = X_1 \subset F_n$$

$$(4.35) \quad \cdots \subset F_{<k} = X_{2,0} \subset X_{2,1} \subset \cdots \subset X_{2,m} = X_2 \subset F_k \subset \cdots$$

with

$$(4.36) \quad X_{1,i}/X_{1,i-1} \cong M_{1,i}, \quad X_{2,i}/X_{2,i-1} \cong M_{i,1}$$

We construct a modified net $\mathcal{X}' = (A', \alpha', B', \beta')$ with

$$(4.37) \quad A' = A \sqcup \{1, 2\} \sqcup \{(l, i, j) \mid l \in \{1, 2\}, i \geq 2, j \geq 2, i+j \leq m+1\}$$

$$(4.38) \quad \alpha'(1) = 2, \quad \alpha'(1, i, j) = (2, i, j), \quad \alpha'|_A = \alpha$$

$$(4.39) \quad B'_1 = \{1\}, \quad B'_2 = \{2\}, \quad B'_{l,i,j} = \{(l, i, j)\}$$

$$(4.40) \quad B'_r = (B_r \setminus \{n\}) \sqcup \{(1, 1, 1) < (1, 1, 2) < \cdots < (1, 1, m)\}$$

$$(4.41) \quad B'_{\alpha(r)} = (B_{\alpha(r)} \setminus \{k\}) \sqcup \{(2, 1, 1) < (2, 2, 1) < \cdots < (2, m, 1)\}$$

$$(4.42) \quad \beta'(1) = 2, \quad \beta'(1, i, j+1) = (2, i+1, j), \quad i \geq 1, j \geq 1, i+j \leq m$$

where in B'_r (resp. $B'_{\alpha(r)}$) the new elements replace n (resp. k) in the total order. There is a morphism of nets $f: \mathcal{X}' \rightarrow \mathcal{X}$ with $1, (1, i, j) \mapsto n$, and $2, (2, i, j) \mapsto k$.

Let G be the preimage of $G_m = H_m$ under the projection $F_n \rightarrow \text{gr}_n V$. V determines a representation $V' = (V', F_i, \psi_i)$ of \mathcal{X}' with

$$(4.43) \quad V'_r = X_1, \quad V'_1 = F_n/G, \quad V'_{1,i,j} = M_{i,j},$$

the restrictions of the filtrations (4.34) to X_1 , and $\psi_1, \psi_{1,i,j}$ induced by ϕ_n . We claim that $f_* V'$ is isomorphic to V . Choose a complement Y to $G \subset F_n$ with $Y \subset F_k$. Further, let $Y_{1,1}, \dots, Y_{1,m} \subset \text{gr}_n V$ with

$$(4.44) \quad X_{1,j}/F_{<n} = Y_{1,1} \oplus \cdots \oplus Y_{1,j}$$

and define $Y_{i+1,j} = \phi_n(Y_{i,j+1})$. We get $Y_{ij} \cong M_{ij}$ and

$$(4.45) \quad G/X_1 = \bigoplus_{\substack{ij \geq 2 \\ i+j < m+1}} Y_{ij}.$$

Also we can use the $X_{1,j}$ and $X_{i,1}$ in the definition of $f_* V'$. Choose a complement Z to $X_1 \subset G$ with $Z \subset F_k$, allowing us to lift Y_{ij} to $Z_{ij} \subset F_k$. Combining the various splittings we obtain an isomorphism

$$(4.46) \quad X_1 \oplus (F_n/G) \oplus \bigoplus_{\substack{ij \geq 2 \\ i+j < m+1}} M_{ij} \rightarrow V_r$$

which we use to identify $f_* V'$ with V . □

Call a net as in (4.10) a *cycle*. Note that for any cycle and any $n \geq 2$ there is a morphism from another cycle which is $n:1$, i.e. an n -fold “covering”. These, and isomorphisms, are the only morphisms between connected nets of height 1. We use this to formulate a strengthening of the previous theorem.

Theorem 4.2. — *Let \mathcal{X} be a net, V an indecomposable representation of \mathcal{X} . Then there exists a connected net \mathcal{X}' of height 1, a morphism $f: \mathcal{X}' \rightarrow \mathcal{X}$ which does not factor through a covering, and an indecomposable representation V' of \mathcal{X}' such that $f_* V' = V$. Moreover, for any triple \mathcal{X}'', f'', V'' with these properties there is an isomorphism of nets $g: \mathcal{X}' \rightarrow \mathcal{X}''$ over \mathcal{X} such that $g_* V' \cong V''$.*

Proof. — Existence follows from Theorem 4.1, which gives us \mathcal{X}', f, V' , with $f_*V' = V$. Since V is indecomposable, V' must be as well, and can only be supported on a single component of \mathcal{X}' so we can take \mathcal{X}' to be connected. If f factors through some covering, we just push V' forward along it.

For uniqueness, consider the inductive procedure in the proof of Theorem 4.1. After choosing a total order on A , the output \mathcal{X}', f, V' depends, up to isomorphism, only on the isomorphism class of V . Moreover, if V is of the form f_*V' for some connected net of height 1, $f : \mathcal{X}' \rightarrow X$ not factoring through a covering, and V' indecomposable, then the output of the procedure applied to V is isomorphic to the same (\mathcal{X}', f, V') . \square

In applications of this theorem below it will be natural to consider nets which are not finite (only individual B_i are). Still, all our results extend to this case, as finite dimensional representations are supported on a finite subnet.

4.3. Minimal twisted complexes

Let \mathcal{A} be a graded linear category which is *augmented* in the sense that there are splittings

$$(4.47) \quad \mathcal{A}(X, Y) = \mathcal{A}_e(X, Y) \oplus \mathcal{A}_r(X, Y)$$

such that $\mathcal{A}_e(X, Y)$ is $\mathbf{K}1_X$ for $X = Y$ and zero otherwise, and $\mathcal{A}_r(X, Y)$ are closed under composition. We view \mathcal{A}_r as a non-unital category with the same objects as \mathcal{A} . Assume for the rest of the subsection that \mathcal{A}_r is nilpotent in the sense that there exists an integer $N > 0$ such that any composition of N morphisms vanishes. Recall that a functor F is said to *reflect isomorphisms* if f is invertible whenever $F(f)$ is.

Lemma 4.2. — *The functor $T : \text{add}\mathbf{Z}(\mathcal{A}) \rightarrow \text{add}\mathbf{Z}(\mathcal{A}_e)$ induced by the augmentation reflects isomorphisms.*

Proof. — Let $M, N \in \text{add}\mathbf{Z}(\mathcal{A})$ and $\phi : M \rightarrow N$ such that $T(\phi) = \bar{\phi}$ is an isomorphism. Thus we have graded vector spaces $M(X), N(X)$ for each $X \in \text{Ob}(\mathcal{A})$, and components of ϕ

$$\begin{aligned} \phi_X &\in \text{Hom}(M(X), N(X)) \oplus (\text{Hom}(M(X), N(X)) \otimes \mathcal{A}_r(X, X)) \\ \phi_{X,Y} &\in \text{Hom}(M(X), N(Y)) \otimes \mathcal{A}_r(X, Y), \quad X \neq Y \end{aligned}$$

Let $\bar{\phi}_X$ be the component of ϕ_X in $\text{Hom}(M(X), N(X))$. By assumption, the $\bar{\phi}_X$ are isomorphisms. Composing ϕ with the morphism $\bar{\phi}^{-1}$ with components $\bar{\phi}_X^{-1}$, we may assume that $M(X) = N(X)$ and all $\bar{\phi}_X$ are identity morphisms. Thus,

$$(4.48) \quad \phi = 1 - \epsilon \in \text{End}^0(M), \quad \epsilon^N = 0$$

which clearly has an inverse. \square

A two-sided twisted complex over \mathcal{A} is given by a pair (M, δ) with $M \in \text{Ob}(\text{add}\mathbf{Z}\mathcal{A})$ and $\delta \in \text{End}^1(M)$ with $\delta^2 = 0$. We say that (M, δ) is *minimal* if the image $\bar{\delta}$ of δ under the functor $\text{add}\mathbf{Z}\mathcal{A} \rightarrow \text{add}\mathbf{Z}\mathcal{A}_e$ vanishes, i.e. δ has components in (the degree 1 part of)

$$(4.49) \quad \text{Hom}(M(X), M(Y)) \otimes \mathcal{A}_r(X, Y).$$

Proposition 4.1.

1. Every twisted complex $A \in \text{Tw}\mathcal{A}$ is isomorphic to a direct sum $A = A_m \oplus A_c$ with A_m minimal and A_c contractible.
2. Any homotopy equivalence between minimal twisted complexes is an isomorphism.

Proof. — 1. Let (M, δ) be a twisted complex over \mathcal{A} , $(M, \bar{\delta})$ its image in $\text{Tw}(\mathcal{A}_e)$. By semisimplicity of \mathcal{A}_e , $(M, \bar{\delta})$ is isomorphic to a direct sum $B_c \oplus B_m$ where B_c is contractible and B_m has trivial differential. Therefore, (M, δ) is isomorphic to a twisted complex of the form

$$\left(\mathbf{K} \oplus \mathbf{K}[-1] \oplus \mathbf{L}, \begin{pmatrix} \delta_{11} & \delta_{12} & \delta_{13} \\ 1_{\mathbf{K}} + \delta_{21} & \delta_{22} & \delta_{23} \\ \delta_{31} & \delta_{32} & \delta_{33} \end{pmatrix} \right)$$

with $\bar{\delta}_{ij} = 0$. Using elementary row and column operations (automorphisms of $\mathbf{K} \oplus \mathbf{K}[-1] \oplus \mathbf{L}$) one reduces the matrix to the form

$$\begin{pmatrix} 0 & 0 & 0 \\ 1_{\mathbf{K}} & 0 & 0 \\ 0 & 0 & \delta_{\mathbf{L}} \end{pmatrix}$$

with $\bar{\delta}_{\mathbf{L}} = 0$, thus providing the direct sum decomposition.

2. Suppose (M, δ^M) and (N, δ^N) are minimal twisted complexes, $f : M \rightarrow N$ and $g : N \rightarrow M$ closed morphisms of degree 0, $G : M \rightarrow M$ and $H : N \rightarrow N$ morphisms of degree -1 such that

$$(4.50) \quad 1_M - gf = \delta^M G + G \delta^M, \quad 1_N - fg = \delta^N H + H \delta^N.$$

Using minimality, $\bar{g}f = 1_M$ and $\bar{f}g = 1_N$, i.e. \bar{f}, \bar{g} are inverse isomorphisms. Hence, the claim follows from the previous lemma. \square

4.4. Classification of objects

We are now ready to state and prove the main result of this section.

Theorem 4.3. — Let S be a compact graded marked surface, then the construction of Section 4.1 sets up a bijection between isomorphism classes of indecomposable objects in $H^0(\mathcal{F}(S))$ and isotopy classes of admissible curves with indecomposable local system.

Remark. — In the special case when $H^0(\mathcal{F}(S))$ is the bounded derived category of a gentle algebra the classification of indecomposable objects was found by Drozd–Burban [16] and Bekkert–Merklen [8]. To classify objects in $\mathcal{F}(S)$ in general we need to consider \mathbf{Z} -graded algebras and the language of twisted complexes instead of projective resolutions. Consider for example the Dynkin quiver with two parallel arrows which do not have the same degree.

Proof. — Let Q be a graded quiver with quadratic monomial relations of type F1, i.e. associated with some systems of arcs A on a compact graded marked surface (S, M, Ω) . The graded linear category $\mathbf{K}Q$ has a natural augmentation with $(\mathbf{K}Q)_r$ generated by non-identity paths. Assume that sufficiently long paths are zero in $\mathbf{K}Q$, equivalently that no components of M are diffeomorphic to S^1 , so that $(\mathbf{K}Q)_r$ is nilpotent and we can apply the results of the previous subsection.

Consider first the case when Q is of the form

$$(4.51) \quad \bullet \longleftarrow \bullet \longleftarrow \dots \longleftarrow \bullet$$

without relations and arbitrary grading. Let us number the vertices from left to right by $\{1, \dots, n\}$ and let α_i denote the arrow from $i + 1$ to i . A minimal twisted complex over $\mathbf{K}Q$ is then given by finite-dimensional graded vector spaces V_1, \dots, V_n , and for $1 \leq i < j \leq n$ a linear map

$$(4.52) \quad \delta_{ij} : V_j \rightarrow V_i, \quad \deg(\delta_{ij}) = \deg(\alpha_i) + \dots + \deg(\alpha_{j-1}) - 1$$

such that the matrix $\delta = (\delta_{ij})$ has square zero.

Consider the direct sum

$$(4.53) \quad V = \bigoplus_{i=1}^n V_i,$$

which has the filtration

$$(4.54) \quad V_1 \subset V_1 \oplus V_2 \subset \dots \subset V_1 \oplus \dots \oplus V_n.$$

The endomorphism δ gives a three-step filtration

$$(4.55) \quad \text{Im}\delta \subset \text{Ker}\delta \subset V$$

and an isomorphism of associated graded $V/\text{Ker}\delta \cong \text{Im}\delta$. We find that we have a representation of the net

$$(4.56) \quad (\{1, 2\} \times \mathbf{Z}, (12) \times \text{id}, (\{1, \dots, n\} \sqcup \{1, 2, 3\}) \times \mathbf{Z}, (13) \times \text{id})$$

with total order so that $(i, k + d(\alpha_i) - 1) < (i + 1, k)$ on $\{1, \dots, n\} \times \mathbf{Z}$ and $(1, n) < (2, n) < (3, n)$ on $\{1, 2, 3\} \times \mathbf{Z}$. In fact, we get an embedding of the category of minimal twisted complexes into the category of representations of the above net as a full subcategory. We do not get an equivalence of categories, since we only consider δ which decrease the filtration (4.54).

Turn now to the case of general \mathbf{Q} . Let \mathbf{D} be the set of maximal non-zero paths in \mathbf{KQ} . Then for each element of \mathbf{D} there is a sub-quiver of \mathbf{Q} which is of the simple form above, and \mathbf{Q} is obtained from their disjoint union by identifying some pairs of vertices. The corresponding net is now given by

$$(4.57) \quad \mathbf{A} = \mathbf{D} \times \{1, 2\} \times \mathbf{Z}$$

$$(4.58) \quad \alpha = \text{id} \times (12) \times \text{id}$$

$$(4.59) \quad \mathbf{B} = (\{(v, d) \in \mathbf{Q}_0 \times \mathbf{D} \mid v \text{ on } d\} \sqcup (\mathbf{D} \times \{1, 2, 3\})) \times \mathbf{Z} =: \mathbf{B}_g \sqcup \mathbf{B}_a$$

with partial order defined so that for an arrow α in $d \in \mathbf{D}$ we have $(\partial_1(\alpha), n + d(\alpha) - 1) < (\partial_0(\alpha), n)$ and $(d, 1, n) < (d, 2, n) < (d, 3, n)$. Further,

$$(4.60) \quad \beta = (\tau \sqcup (\text{id} \times (13))) \times \text{id}$$

where τ sends a pair $(v, d_1) \in \mathbf{Q}_0 \times \mathbf{D}$ with v on $d_1, d_2 \in \mathbf{D}$, $d_1 \neq d_2$ to (v, d_2) . The definitions are set up so that the following proposition holds. The proof is just chasing through the definitions.

Proposition 4.2. — *There is a one-to-one correspondence between isotopy classes of admissible curves on \mathbf{S} with local system on the one hand, and morphisms from nets of height 1 with indecomposable representation to the net $(\mathbf{A}, \alpha, \mathbf{B}, \beta)$ constructed above on the other. Under this correspondence, nets of the type (4.9) (resp. (4.10)) correspond to admissible curves with domain $[0, 1]$ (resp. \mathbf{S}^1).*

There is a fully-faithful functor from the category of minimal twisted complexes over \mathbf{Q} to the category of representations of $(\mathbf{A}, \alpha, \mathbf{B}, \beta)$. Theorem 4.2 gives us a classification of indecomposable minimal twisted complexes over \mathbf{KQ} , equivalently isomorphism classes in $\mathbf{H}^0(\mathcal{F}(\mathbf{S}))$ by Proposition 4.1.

It remains to deal with the case when \mathbf{M} has components diffeomorphic to \mathbf{S}^1 . Let \mathbf{S}' be the surface obtained by adding two corners and a boundary arc on each such component of \mathbf{M} . Then $\mathbf{H}^0(\mathcal{F}(\mathbf{S}))$ is a localization of the triangulated category $\mathbf{H}^0(\mathcal{F}(\mathbf{S}'))$, and the localization functor is essentially surjective. Moreover, if the images of $\mathbf{E}, \mathbf{F} \in \mathbf{H}^0(\mathcal{F}(\mathbf{S}))$ corresponding to admissible curves become isomorphic under localization, then \mathbf{E} and \mathbf{F} differ by extensions by boundary arcs, so map to isotopic curves in \mathbf{S} . \square

5. Comparison of moduli spaces

This section contains our main result, identifying $\mathcal{M}(S)$ with an open and closed subset of $\text{Stab}(\mathcal{F}(S))$. We begin by reviewing Bridgeland’s axioms in Section 5.1. In Section 5.2 we compute $K_0(\mathcal{F}(S))$ in terms of singular homology. The main theorems are stated in Section 5.3 and proven in the remaining subsections. Sections 5.4 and 5.5 deal with the support property and Harder–Narasimhan filtrations, respectively. In the final Section 5.6 we prove that the map on moduli spaces is complex bianalytic onto an open and closed subset.

5.1. Stability structures on categories

Fix a triangulated category \mathcal{C} and a homomorphism $\text{cl} : K_0(\mathcal{C}) \rightarrow \Gamma$ to a finitely generated abelian group Γ . A *stability structure* (cf. [11], [31], also called a *stability condition*) on \mathcal{C} is given by

- for each $\phi \in \mathbf{R}$ a full additive subcategory $\mathcal{C}^\phi \subset \mathcal{C}$ of *semistable objects of phase ϕ* , and
- an additive map $Z : \Gamma \rightarrow \mathbf{C}$, the *central charge*.

This data has to satisfy the following axioms.

1. $\mathcal{C}^\phi[1] = \mathcal{C}^{\phi+1}$
2. If $E \in \mathcal{C}^{\phi_1}$, $F \in \mathcal{C}^{\phi_2}$, $\phi_1 > \phi_2$, then $\text{Hom}(E, F) = 0$.
3. Every $E \in \mathcal{C}$ has a *Harder–Narasimhan filtration*: A tower of triangles

$$\begin{array}{ccccccc}
 0 = E_0 & \longrightarrow & E_1 & & \cdots & & E_{n-1} \longrightarrow E_n = E \\
 & & \downarrow & & & & \downarrow \\
 & & A_1 & & & & A_n \\
 & \swarrow \text{dotted} & & & & & \swarrow \text{dotted}
 \end{array}$$

with $0 \neq A_i \in \mathcal{C}^{\phi_i}$ and $\phi_1 > \phi_2 > \cdots > \phi_n$. The A_i are called the *semistable components* of E .

4. If $0 \neq E \in \mathcal{C}^\phi$ then $Z(E) := Z(\text{cl}([E])) \in \mathbf{R}_{>0} e^{\pi i \phi}$.
5. The *support property*: For some norm $\|\cdot\|$ on $\Gamma \otimes \mathbf{R}$ and $C > 0$ we have an estimate

(5.1) $\quad \|\text{cl}(E)\| \leq C|Z(E)|$

for $E \in \mathcal{C}^\phi$.

The set $\text{Stab}(\mathcal{C}, \Gamma)$ of all stability structures has a natural topology which is induced by the metric

(5.2)
$$d(\sigma_1, \sigma_2) = \sup_{0 \neq E \in \mathcal{C}} \left\{ |\phi_{\sigma_2}^-(E) - \phi_{\sigma_1}^-(E)|, |\phi_{\sigma_2}^+(E) - \phi_{\sigma_1}^+(E)|, \left| \log \frac{m_{\sigma_2}(E)}{m_{\sigma_1}(E)} \right| \right\}$$

cf. [11]. Here, for an object $E \in \mathcal{C}$ with semistable components A_1, \dots, A_n with phases $\phi_1 > \dots > \phi_n$ one defines $m(E) = \sum |Z(A_i)|$, $\phi^-(E) = \phi_n$, and $\phi^+(E) = \phi_1$. The main result of [11] (see also [31]) is that the projection

$$\text{Stab}(\mathcal{C}, \Gamma) \rightarrow \text{Hom}(\Gamma, \mathbf{C})$$

is a local homeomorphism, hence $\text{Stab}(\mathcal{C}, \Gamma)$ has the structure of a complex manifold of dimension $\text{rk}(\Gamma)$, if it is non-empty.

5.2. Charge lattice

Recall that a graded surface has a double cover τ given by the orientations of the foliation lines, and $\mathbf{Z}_\tau = \mathbf{Z} \otimes_{\mathbf{Z}/2} \tau$ is the associated local system of abelian groups. The following theorem describes \mathbf{K}_0 of the Fukaya category in terms of homology with coefficients in \mathbf{Z}_τ .

Theorem 5.1. — *Let $S = (S, M, \eta)$ be a compact graded marked surface, then there is a natural isomorphism of abelian groups*

$$\mathbf{K}_0(\mathcal{F}(S)) \cong H_1(S, M; \mathbf{Z}_\tau).$$

Proof. — Choose a full system A of graded arcs on S . Each arc defines a class in $H = H_1(S, M; \mathbf{Z}\sqrt{\Omega})$. From cellular homology it follows that H is the group generated by these arcs and with relations of the form $\pm X_1 \pm \dots \pm X_n = 0$ where X_1, \dots, X_n are the arcs bounding a disk cut out by A . Correspondingly, it is also clear that the arcs in A generate

$$(5.3) \quad \mathbf{K} := \mathbf{K}_0(\mathcal{F}(S)) = \mathbf{K}_0(\text{Tw}(\mathcal{F}_A(S)))$$

and satisfy the same relations, which follows from the observation in Section 3.3. By construction we get a surjective homomorphism $H \rightarrow \mathbf{K}$ which is independent of the choice of arcs and natural with respect to maps of graded marked surfaces. It remains to show that no additional relations are needed to present \mathbf{K} .

Suppose first that M has no components diffeomorphic to S^1 . Choosing a formal collection A of arcs, we can present $\mathcal{F}(S)$ as the category $\text{Tw}(\mathbf{KQ})$ of twisted complexes over the path category of a graded quiver Q with quadratic monomial relations. Note that H is freely generated by A , so we need to verify that \mathbf{K} is freely generated by the vertices Q_0 of Q . One way to see this is by using the explicit resolution of the diagonal of \mathbf{KQ} , showing that, as a bimodule, \mathbf{KQ} is a repeated cone of Yoneda bimodules. As a consequence, $\text{Tw}(\mathbf{KQ}) \cong \text{Perf}(\mathbf{KQ})$, and an inverse of the map

$$(5.4) \quad \mathbf{Z}^{Q_0} \rightarrow \mathbf{K}_0(\text{Perf}(\mathbf{KQ}))$$

sending a vertex to the corresponding simple module, is induced by the dimension vector of a module.

For a general graded marked surface S there is a S' obtained by adding two corners to each component of ∂S diffeomorphic to S^1 . Let N denote set of boundary arcs of S' created in this process, and \mathcal{N} the corresponding full triangulated subcategory of $H^0(\mathcal{F}(S'))$ generated by N . Each arc in N has endomorphism algebra of the form $\mathbf{K}[x]/x^2$ with $|x| \in \mathbf{Z}$, so \mathcal{N} is a product of categories of the form $H^0(\text{Tw}(\mathbf{K}[x]/x^2))$. We get a commutative diagram

$$(5.5) \quad \begin{array}{ccccccc} \mathbf{Z}^N & \longrightarrow & H_1(S', M', \mathbf{Z}\sqrt{\Omega'}) & \longrightarrow & H_1(S, M, \mathbf{Z}\sqrt{\Omega}) & \longrightarrow & 0 \\ \downarrow & & \downarrow & & \downarrow & & \\ K_0(\mathcal{N}) & \longrightarrow & K_0(\mathcal{F}(S')) & \longrightarrow & K_0(\mathcal{F}(S)) & \longrightarrow & 0 \end{array}$$

where we use Proposition 3.6 to get the bottom row. We claim that the rows are exact. For the top row this is just the exact sequence of a triple. For the bottom row this uses the fact that \mathcal{N} is idempotent complete, thus a thick subcategory so that Proposition 3.1 in [24] can be applied. Since the left and middle horizontal maps are isomorphisms, so is the right one. \square

When a flat surface S comes from a quadratic differential with higher order poles, the category $\mathcal{F}(S)$ includes some objects which have infinite length since

$$(5.6) \quad \int_0^1 x^{-n/2} dx = \infty$$

for $n \geq 2$. We need to pass to the subcategory of $\mathcal{F}(S)$ which does not include these objects. Formally, we want to allow $S = (S, M)$ to have boundary circles which are unmarked, i.e. not belonging to M . Thus, define $\mathcal{F}(S, M)$ as the full subcategory of $\mathcal{F}(S, M')$ of objects corresponding to immersed curves avoiding $M' \setminus M$, where M' is the union of M and all boundary circles.

Example. — Consider the case when S is an annulus. By homological mirror symmetry, $\mathcal{F}(S) \cong D^b\text{Coh}(\mathbf{G}_m)$. Objects in $\text{Coh}(\mathbf{G}_m)$ are direct sums of line bundles and torsion sheaves. The relevant subcategory is just the one of torsion sheaves. Coincidentally, “finite-length” in the sense of abelian categories corresponds in this case to finite length with respect to the flat metric.

Lemma 5.1. — *The full subcategory $\mathcal{F}(S, M) \subset \mathcal{F}(S, M')$ is triangulated, i.e. closed under shifts and cones.*

Proof. — For each component B of the boundary we will introduce a functor I_B from $\mathcal{F}(S, M')$ to the category of complexes of vector spaces (possibly unbounded), such that $E \in \mathcal{F}(S, M)$ if and only if $I_B(E) = 0$ for all components B of M' which are not in M , which implies the claim. Roughly speaking, I_B is $\text{Hom}(B, _)$, but in general the curve B does not have a grading and so is not an object of the Fukaya category as defined here. As a workaround we define this functor explicitly by hand.

Fix a boundary component $B \subset \partial S$. The grading and vectors tangent to the boundary define sections s_1, s_2 of $\mathbf{P}(\text{TS}|_B)$. If $B \cong S^1$, then since S is oriented, B and $\mathbf{P}(\text{TS}|_B)$ carry natural orientations, and so the intersection number, m , between s_1 and s_2 is well defined. If $B \cong \mathbf{R}$ we take $m = 0$. Let $N = |m|$. If $N = 0$, then B has a grading in the sense of Section 2.1, and if $N > 0$ then B still has a \mathbf{Z}/N -grading: A lift of s_2 to the fiberwise N -fold covering of $\mathbf{P}(\text{TS}|_B)$. In either case, choose an appropriate grading on B . Then for every graded curve A ending on B transversely, we get an intersection number $i(B, A) \in \mathbf{Z}/N$.

The cohomology of $I_B(E)$ will be bounded if $N = 0$ and N -periodic if $N > 0$. It suffices to define I_B on generators of $\mathcal{F}(S, M')$. Let A be an arc in S . If both endpoints of A do not lie on B , then $I_B(A) = 0$. If one endpoint of A lies on B , then

$$(5.7) \quad I_B(A) = \bigoplus_{\substack{j \in \mathbf{Z} \\ j \equiv i(B, A)}} \mathbf{K}[-j].$$

Finally, if both endpoints of A lie on B , then $I_B(A)$ is a sum of two graded vector spaces as above. The map on morphisms is defined in the obvious way so that $I_B(a) = 0$ if a is a boundary path which is not in B . One checks that this gives a well defined A_∞ -functor. \square

Proposition 5.1. — *Let S be a compact graded marked surface which is allowed to have unmarked boundary circles. Then there is a natural map*

$$(5.8) \quad K_0(\mathcal{F}(S)) \longrightarrow H_1(S, M; \mathbf{Z}_\tau).$$

The group $K_0(\mathcal{F}(S))$ is often infinitely generated, and so the above map will fail to be injective. It may also fail to be surjective, e.g. for the cylinder with non-standard grading, where $\mathcal{F}(S)$ is trivial if both boundary circles are unmarked.

Proof. — Let S' be the graded marked surface obtained from S by replacing each unmarked boundary circle with one marked and one unmarked boundary arc. Then $\mathcal{F}(S)$ is also a subcategory of $\mathcal{F}(S')$, as follows from localization. Note that

$$(5.9) \quad \iota : H_1(S, M; \mathbf{Z}_\tau) \longrightarrow H_1(S', M'; \mathbf{Z}_\tau) \cong K_0(\mathcal{F}(S'))$$

is an inclusion and the image of ι contains the image of $K_0(\mathcal{F}(S)) \longrightarrow K_0(\mathcal{F}(S'))$. Hence we have the desired natural map. \square

5.3. *Statement of main theorem*

Let \widehat{X} be the real blow-up of a flat surface X of finite type. It contains a compact surface with marked boundary $S \subset \widehat{X}$, canonical up to isotopy, as a deformation retract. The complement $\widehat{X} \setminus S$ is a disjoint union of punctured open disks, one for each higher order pole of the quadratic differential, and rectangles $\mathbf{R}_{>0} \times [0, 1]$, one for each infinite-angle singularity (more precisely, non-cyclic boundary walk). The marked boundary of S is the part of ∂S in $\partial \widehat{X}$. Each removed punctured open disk is bounded by an unmarked boundary circle, while each removed rectangle is bounded on one end by an unmarked part of ∂S . Thus, the Fukaya category $\mathcal{F}(X) = \mathcal{F}(S)$ is defined. We also note that the inclusion of pairs $(S, M) \subset (\widehat{X}, \partial \widehat{X})$ is a homotopy equivalence, so both give isomorphic relative homologies.

Theorem 5.2. — *Let X be a flat surface of finite type. Define subcategories $\mathcal{C}^\phi \subset H^0(\mathcal{F}(X))$ so that isomorphism classes of indecomposable objects correspond to unbroken graded geodesics of phase ϕ with indecomposable local system, and let*

$$(5.10) \quad Z : K_0(\mathcal{F}(S)) \longrightarrow H_1(X, \partial X; \mathbf{Z}_\tau) =: \Gamma \rightarrow \mathbf{C}$$

be the period map. Then this data satisfies the axioms of a stability structure. In particular, for a graded surface of finite type X there is a continuous map

$$(5.11) \quad \mathcal{M}(X) \longrightarrow \text{Stab}(\mathcal{F}(X))$$

over $\text{Hom}(\Gamma, \mathbf{C})$.

Remark. — Let us emphasize that we put no restrictions on the ground field \mathbf{K} of $\mathcal{F}(S)$. In fact this category is defined over any ring, but already $\text{Perf}(\mathbf{Z})$ does not admit stability structures.

Theorem 5.3 (Main theorem). — *Let X be a graded surface of finite type, then the map $\mathcal{M}(X) \rightarrow \text{Stab}(\mathcal{F}(X))$ is injective and its image open and closed, thus a union of components of $\text{Stab}(\mathcal{F}(X))$.*

Remark. — When C is compact, then necessarily $g(C) = 1$. In this case one can use homological mirror symmetry for elliptic curves [41] and Atiyah’s classification [4] to derive a version of Theorem 5.2. The computation of $\text{Stab}(D^b(E))$ in [11] shows that $\mathcal{M}(X) = \text{Stab}(X)$.

In the case where C is a punctured torus, the mirror is a nodal elliptic curve E . The space of stability conditions on $D^b(E)$ was computed by Burban–Kreuzler [17]. Our theorem gives a different description of the same space.

The proofs of Theorem 5.2 and Theorem 5.3 will occupy the rest of this section.

5.4. *Support property*

Let X be a flat surface of finite type and $Y = \widehat{X}$ its real blow-up. We have as before

$$(5.12) \quad \Gamma = H_1(Y, \partial Y; \mathbf{Z}_\tau), \quad \Gamma_{\mathbf{R}} = \Gamma \otimes_{\mathbf{Z}} \mathbf{R} = H_1(Y, \partial Y; \mathbf{R}_\tau)$$

Proposition 5.2. — *There is a norm $\|\cdot\|$ on $\Gamma_{\mathbf{R}}$, $C > 0$ such that for $\gamma \in \Gamma$ the class of a finite geodesic on X there is an estimate*

$$(5.13) \quad \|\gamma\| \leq C|Z(\gamma)|$$

where $Z : \Gamma \rightarrow \mathbf{C}$ is the period map.

Proof. — The locally constant sheaf \mathbf{R}_τ is given by flat sections of a flat line bundle with metric for which we use the same notation. We have

$$(5.14) \quad \Gamma_{\mathbf{R}}^* = H_1(Y, \partial Y; \mathbf{R}_\tau)^*$$

$$(5.15) \quad = H^1(Y, \partial Y; \mathbf{R}_\tau)$$

$$(5.16) \quad = H_{dR}^1(Y, \partial Y; \mathbf{R}_\tau)$$

Choose forms $\omega_1, \dots, \omega_n$ representing a basis of $H_{dR}^1(Y, \partial Y; \mathbf{R}_\tau)$. Let \widehat{K} be the closure of the preimage of $K = \text{Core}(X)$ under the inclusion $Y \setminus \partial Y \rightarrow X$. Define a norm on $\Gamma_{\mathbf{R}}$ by

$$(5.17) \quad \|\gamma\| = \sum_i |[\omega_i](\gamma)|$$

and set

$$(5.18) \quad C = \sum_i \|\omega_i|_{\widehat{K}}\|_\infty$$

where the norm is taken with respect to the flat metric. Let α be a finite geodesic on X with class $\gamma \in \Gamma$, then

$$(5.19) \quad l(\alpha) = |Z(\gamma)|$$

where $l(\alpha)$ is the length. Thus

$$(5.20) \quad \|\gamma\| = \sum_i |[\omega_i](\gamma)|$$

$$(5.21) \quad = \sum_i \left| \int_\alpha \omega_i \right|$$

$$(5.22) \quad \leq \sum_i l(\alpha) \|\omega_i|_{\widehat{K}}\|_\infty$$

$$(5.23) \quad = C|Z(\gamma)|. \quad \square$$

5.5. Harder–Narasimhan filtrations

We continue with the proof of Theorem 5.2. The strategy is to first deal with the special case of a disk and annulus, then deduce the general case by passing to the universal and annular coverings. The idea for the proof of the Harder–Narasimhan property is to use the fact that arbitrary indecomposable objects are represented by immersed curves which we can deform to their geodesic representative: a closed loop or concatenation of saddle connections. Each semistable component of the HN filtration corresponds to several such finite length geodesics.

Let X be a flat surface with X_{sm} diffeomorphic to a disk and all conical singularities with infinite angle. The corresponding graded marked surface S is homeomorphic to a closed disk with $n + 1 = |X_{\text{sg}}|$ marked and unmarked boundary arcs. Consider the data of a stability structure on $\mathcal{F}(X)$ as in the statement of theorem. We first check that $\text{Hom}(A, B) = 0$ for semistable objects $A, B \neq 0$ with $\phi(A) > \phi(B)$. Note that, by the long exact sequence for Hom , it suffices to do this for A, B stable, so that they correspond to graded arcs. We consider the various cases.

1. If A and B differ only by a grading shift, the claim follows from $\text{Ext}^{<0}(A, A) = 0$.
2. If A and B are disjoint, then $\text{Ext}(A, B) = 0$.
3. If A and B intersect in a smooth point p of the surface, then $\text{Ext}(A, B)$ is concentrated in degree

$$(5.24) \quad i_p(A, B) = \lceil \phi(A) - \phi(B) \rceil$$

so $\phi(A) > \phi(B)$ implies $\text{Ext}^0(A, B) = 0$. This is clear from the alternative definition of the Fukaya category where morphisms are given by intersection points, and can be verified here by direct computation for the Fukaya category of a disk with four marked boundary components.

4. If A and B meet in an ∞ -angle singularity p in an angle $\phi > 0$, then

$$(5.25) \quad i_p(A, B) = \phi(A) - \phi(B) + \frac{\phi}{\pi}$$

and $\text{Ext}(A, B)$ is either concentrated in degree $i_p(A, B)$ or zero, depending on the order in which they meet p .

Next we check the HN-property. It suffices to do this for indecomposable objects E . By the classification, E corresponds to a graded arc α connecting two ∞ -angle singularities. A geodesic representative of α is a concatenation of graded saddle connections $\alpha_1, \dots, \alpha_k$ corresponding to objects A_1, \dots, A_k in $\mathcal{F}(X)$. Let $\phi_1 > \dots > \phi_l$ be the distinct phases of the α_i (generically, $l = k$). For $1 \leq i \leq l$ there is a semistable B_i given by a twisted complex

$$(5.26) \quad B_i = \left(\bigoplus_{\phi(A_j)=\phi_i} A_j, \delta_i \right)$$

where δ_i has non-zero coefficients corresponding to the singular points where the geodesic representative of α passes through without changing phase, which is equal to ϕ_i . Also,

$$(5.27) \quad E = \left(\bigoplus_j A_j, \delta \right)$$

where δ has non-zero coefficients corresponding to all the singular points the geodesic representative passes through (i.e. where α_i meets α_{i+1}). Note that this really gives morphisms of degree 1 since the concatenation of the saddle connections smooths to a graded curve. To show that the twisted complex representation of E in terms of the B_i is a HN-tower, we need to check that components of δ increase phase. For this we use that fact that the concatenation of the α_i is a geodesic, and so α_i, α_{i+1} necessarily meet at a singular point in an angle $\phi \geq \pi$, thus $\phi > \pi$ if the phase jumps. The corresponding morphism, without loss of generality from A_i to A_{i+1} , has degree 1, and so

$$(5.28) \quad 1 = \phi(A_i) - \phi(A_{i+1}) + \frac{\phi}{\pi}$$

which implies $\phi(A_i) < \phi(A_{i+1})$.

Suppose now that X is a general flat surface of finite type. The first claim is that $\text{Hom}(A, B) = 0$ for a pair of graded saddle connections A, B with $\phi(A) > \phi(B)$. To see this, let Y be the flat surface so that its real blow-up \widehat{Y} is the universal cover of \widehat{X} . Note that all singularities of Y have infinite cone angle. Lift A, B to graded curves $\widetilde{A}, \widetilde{B}$ on Y with the same phases. We have

$$(5.29) \quad \text{Hom}(A, B) = \bigoplus_{g \in \pi_1(\widehat{X})} \text{Hom}(\widetilde{A}, g\widetilde{B})$$

and hence it suffices to prove the claim for Y . But any two saddle connections in Y are objects in some subcategory of $\mathcal{F}(Y)$ which is the Fukaya category of some surface of disk-type as above, and so our previous arguments can be applied.

Second, we claim that the HN-property holds for admissible curves c with domain $[0, 1]$. To see this, lift c to a graded curve \tilde{c} in the universal cover Y . We work in the category generated by all the saddle connections in a geodesic representative of \tilde{c} . It corresponds to some surface of disk-type in the universal cover. We get a HN-tower for \tilde{c} in this category and push it forward to $\mathcal{F}(X)$ to get a tower for c in that category.

We turn to the case when X is of annular type, i.e. X_{sm} is diffeomorphic to an annulus with grading such that a simple closed loop around it is gradable, and all singularities have infinite cone angles. Let $S \subset \widehat{X}$ be the corresponding compact graded marked surface. The category $\mathcal{F}(X)$ is the bounded derived category of finite-dimensional representations of a quiver which is a cyclic chain of arrows, all in degree zero. The number of arrows oriented in one or the other way is equal to the number of marked boundary arcs

on the two components of ∂S , respectively. In arguments below we omit straightforward computations of $\text{Ext}^*(A, B)$ in $\mathcal{F}(X)$.

First we check that $\text{Hom}(A, B) = 0$ for stable A, B with $\phi(A) > \phi(B)$. It remains to consider the cases when either A or B , or both, are closed geodesics with simple local system.

1. If A and B both correspond to loops with simple local system, then they either differ by a shift or else are orthogonal, i.e. $\text{Ext}^*(A, B) = \text{Ext}^*(B, A) = 0$. In the first case it suffices to note that $\text{Ext}^{<0}(A, A) = 0$ since A lies in an abelian category.
2. If one of A or B is a geodesic loop and the other a saddle connection tending in both directions to the same component of ∂S , then A and B are orthogonal.
3. If one of A or B is a geodesic loop and the other a saddle connection from one component of ∂S to the other, then A and B intersect in a unique point p and $\text{Ext}^*(A, B)$ is concentrated in degree $i_p(A, B) = \lceil \phi(A) - \phi(B) \rceil > 0$.

Recall that an annular covering of a surface is a covering corresponding to a subgroup of π_1 generated by a simple closed loop (cf. [47]). By passing to the annular covering (of \widehat{X}) we see, as in the case of two saddle connections, that the statement is true for general X when at least one of A or B is a saddle connection.

For annular X there is, up to isotopy and grading, a unique admissible curve c in X with domain S^1 . We need to check that HN-filtrations exist for the corresponding indecomposable objects X . If X has a geodesic loop (necessarily isotopic to c), then c is semistable, so we are done. If not, then c has a geodesic representative which is a concatenation of saddle connections. The corresponding collection of arcs satisfies the properties we required when assigning a twisted complex to c in Section 3.1. Note that in this case X cannot contain a semi-infinite flat cylinder, so S has no unmarked boundary circles. The same argument as in the disk case shows that this twisted complex gives a HN-filtration for X . Again, passing to the annular covering, we find that HN-filtration exist for objects corresponding to admissible curves with local system and domain S^1 on general X . Since all isomorphism classes of indecomposable objects in $H^0(\mathcal{F}(X))$ come from admissible curves, this completes the proof of the HN-property.

Let X now be a general flat surface of finite type. It remains to show that for two objects $A, B \in \mathcal{F}(X)$ corresponding to closed geodesics α, β with local systems and $\phi(A) > \phi(B)$ we have $\text{Hom}(A, B) = 0$. We may also assume that α and β are non-isotopic as simple closed curves. Let Z denote the flat cylinder foliated by closed geodesics isotopic to α . If Z has infinite height then α and β cannot intersect at all, so assume Z has finite height with singular points on each boundary component, see Figure 6. The other loop β can cross Z any number of times, but does so always with the same slope. Then we can represent B as an extension of graded saddle connections with local systems S_i , which are contained in Z and connect the two components of its boundary, and arcs E_i which are disjoint from the interior of Z . Moreover, these saddle connections can be chosen so




FIG. 6. — The cylinder Z , cut along a *dotted line*, closed geodesics α, β , saddle connections S_i , arcs E_i

that $\phi(S_i) < \phi(B)$. Namely, there is an isotopy moving β to the concatenation of the S_i and E_i , and it should move any point of intersection with ∂Z so that Z lies to the right. From what we have shown in the annular case, it follows that $\text{Hom}(A, S_i) = 0$ for all i . Further, A is orthogonal to any E_i , so $\text{Hom}(A, B) = 0$ follows. This completes the proof of Theorem 5.2.

5.6. *The map $\mathcal{M}(X) \rightarrow \text{Stab}(\mathcal{F}(X))$*

We consider a construction for general categories with a stability structure. Let \mathcal{C} be a triangulated category with stability structure σ , and $F : \mathcal{C} \rightarrow \text{D}(\text{Mod}(\mathbf{K}))$ an exact functor to the derived category of chain complexes over a field \mathbf{K} . We combine the two to define

$$(5.30) \quad C_{F,\sigma} = \bigcup_{\substack{E \in \mathcal{C}^{\text{ss}}, k \in \mathbf{Z} \\ H^k(F(E)) \neq 0}} [\log |Z(E)|, +\infty) \times \{\pi(\phi(E) + k)\} \subset \mathbf{R}^2$$

which is a union of horizontal rays starting at points of the form $\log(Z(E))$, E semistable. The support property ensures that the central charges $Z(E) \in \mathbf{C}$, E semistable, form a discrete subset, hence $C_{F,\sigma} \subset \mathbf{R}^2$ is closed.

Lemma 5.2. — For $\sigma_1, \sigma_2 \in \text{Stab}(\mathcal{C})$, \mathcal{C} and F as above, we get

$$(5.31) \quad d_H(C_{F,\sigma_1}, C_{F,\sigma_2}) \leq d(\sigma_1, \sigma_2)$$

where d_H is the Hausdorff distance on closed subsets induced by the max-metric on \mathbf{R}^2 .

Proof. — Let $\epsilon = d(\sigma_1, \sigma_2)$, then by definition

$$(5.32) \quad |\phi_{\sigma_1}^-(E) - \phi_{\sigma_2}^-(E)| \leq \epsilon, \quad |\phi_{\sigma_1}^+(E) - \phi_{\sigma_2}^+(E)| \leq \epsilon$$

$$(5.33) \quad |\log(m_{\sigma_1}(E)) - \log(m_{\sigma_2}(E))| \leq \epsilon$$

for all $0 \neq E \in \mathcal{C}$. For a subset A of a metric space let $B_\epsilon(A)$ be the set of points with distance $\leq \epsilon$ to A . To show that $C_{F,\sigma_2} \subseteq B_\epsilon(C_{F,\sigma_1})$ it suffices to find, for every σ_2 -semistable object E with $H^k(F(E)) \neq 0$, a σ_1 -semistable object E' with $H^k(F(E')) \neq 0$ and

$$(5.34) \quad |\phi_{\sigma_2}(E) - \phi_{\sigma_1}(E')| \leq \epsilon, \quad \log |Z_{\sigma_1}(E')| \leq \log |Z_{\sigma_2}(E)| + \epsilon.$$

Let A_1, \dots, A_n be the σ_1 -semistable components of E , $n \geq 1$, then by exactness of F there is an m with $H^k(F(A_m)) \neq 0$, and we can take $E' = A_m$.

By symmetry we also have $C_{F,\sigma_1} \subseteq B_\epsilon(C_{F,\sigma_2})$, thus $d_H(C_{F,\sigma_1}, C_{F,\sigma_2}) \leq \epsilon$. \square

Proof of Theorem 5.3. — 1. Let $B \subset \partial X$ be a connected component of the boundary. As in the proof of Lemma 5.1 we choose a grading on B to get a functor $I_B : \mathcal{F}(X) \rightarrow D(\text{Mod}(\mathbf{K}))$. Recall that if B is gradable in the usual sense, then $I_B(E)$ is perfect for every $E \in \mathcal{F}(S)$. On the other hand, if B is only \mathbf{Z}/n -gradable, n maximal, then $I_B(E)$ is n -periodic for every $E \in \mathcal{F}(S)$.

Suppose X has a flat structure with corresponding stability structure $\sigma \in \text{Stab}(\mathcal{F}(X))$. If B is a boundary component corresponding to a conical point s , then geometrically $C_{I_B,\sigma}$ is just the complement of the domain of the exponential map at s , $\text{Dom}(\exp_s)$, in logarithmic coordinates. The Voronoi cell, V_s , around s is contained in $\text{Dom}(\exp_s)$, and in fact V_s is determined by $\text{Dom}(\exp_s)$ alone, since by definition $x \in V_s$ if and only if the disk centered at x and with s on the boundary is contained in the domain. This shows that the flat structure is recovered from σ , proving injectivity.

2. Let (F_i) be a sequence of flat structures on X so that the corresponding sequence of stability structure (σ_i) converges to $\sigma \in \text{Stab}(\mathcal{F}(X))$. We want to show that the sequence $(F_i)_i$ converges in $\mathcal{M}(X)$. The idea is to show that the corresponding sequence of Voronoi diagrams eventually stabilizes, so we can then just take the limits of the lengths of the edges.

Let $\beta(F_i)$ be the maximal length of a boundary edge of $\text{Core}(F_i)$. The first step is to show that $\beta(F_i) \leq \beta$ for some $\beta > 0$. Let $E \in \mathcal{F}(X)$ be the direct sum of objects corresponding to all the boundary edges of $\text{Core}(F_i)$ with arbitrary fixed grading. An edge should appear twice in the sum if it does not bound the interior of $\text{Core}(F_i)$. The mass $m_{\sigma_i}(E)$ is then just the circumference of $\text{Core}(F_i)$. Since $\sigma_i \rightarrow \sigma$, we have an upper bound on $m_{\sigma_i}(E) \geq \beta(F_i)$, and the claim follows.

The second step is to show that $\rho(F_i)$ is bounded. Fix a conical point s of S . We will show that $d(s, x)$ is bounded for $x \in V_{s,i} \cap \text{Core}(F_i)$, where $V_{s,i}$ is the Voronoi cell around s in F_i , and since

$$(5.35) \quad \rho(F_i) = \max_{s \in (F_i)_{\text{sg}}} d_H(\{s\}, V_{s,i} \cap \text{Core}(F_i))$$

this will imply the claim. Consider the functor $I : \mathcal{F}(X) \rightarrow D(\text{Mod}(\mathbf{K}))$ associated with s as before. We have closed subsets $C_i = C_{I,\sigma_i} \subset \mathbf{R}^2$ converging to $C = C_{I,\sigma}$ in the Hausdorff topology. Let \mathcal{D} denote the set of directions (of geodesics) from s . The choice of

grading of marked boundary as above gives an identification $\mathcal{D} \cong \mathbf{R}/n\pi$ when s has cone angle $n\pi$ and $\mathcal{D} \cong \mathbf{R}$ when s has infinite cone angle. By the *support* of C_i , $\text{Supp}(C_i)$, we mean the closure of its projection to \mathcal{D} . It follows from compactness of $\text{Core}(F_i)$ that each C_i is compactly supported. By convergence, there is compact $J \subset \mathcal{D}$ containing all $\text{Supp}(C_i)$ and $\text{Supp}(C)$. Let $r_i(\phi) \in \mathbf{R}_{\geq 0} \cup \{+\infty\}$ be the length of the maximal geodesic starting at s in direction $\phi \in \mathcal{D}$ and which is entirely contained in $V_{s,i}$. Note that the functions r_i are determined by C_i , thus σ_i alone, are continuous, and converge pointwise to $r: \mathcal{D} \rightarrow \mathbf{R}_{\geq 0} \cup \{+\infty\}$ determined by σ . Define $q_i(\phi) \in \mathbf{R}_{\geq 0}$ similarly as the length of the maximal geodesic starting at s in direction $\phi \in \mathcal{D}$ and which is entirely contained in $V_{s,i} \cap \text{Core}(F_i)$. We need to show that q_i are uniformly bounded. By definition, $q_i(\phi) \leq r_i(\phi)$, and

$$(5.36) \quad (\phi - \pi, \phi + \pi) \cap J = \emptyset \implies q_i(\phi) = 0.$$

Suppose, for contradiction, that there is a sequence $\phi_i \in \mathcal{D}$ with $q_i(\phi_i) \rightarrow \infty$. By (5.36) the ϕ_i are contained in a compact subset of \mathcal{D} and we may assume $\phi_i \rightarrow \phi$ after passing to a subsequence. Now since $q_i(\phi_i) \rightarrow \infty$ we must have $r(\phi) = \infty$, which means that the interval $(\phi - \pi, \phi + \pi)$ is disjoint from $\text{Supp}(C)$. Thus also

$$(5.37) \quad \left[\phi - \frac{\pi}{2}, \phi + \frac{\pi}{2} \right] \cap \text{Supp}(C_i) = \emptyset$$

for sufficiently large i . The sector bounded by directions $\phi - \pi/2$ and $\phi + \pi/2$ in $\text{Dom}(\exp_s)$ cannot map entirely to $\text{Core}(F_i)$, as this contradicts $\rho(F_i) < \infty$. Hence the part of the sector which maps to $\text{Core}(F_i)$ is a triangle Δ or just the point s . The length of the side of the triangle opposite the vertex s is bounded above by β . Thus also $q_i|_{(\phi - \pi/2, \phi + \pi/2)}$ has an upper bound in terms of β , which is independent of i , a contradiction.

We have established that $\rho(F_i)$ and $\beta(F_i)$ have upper bounds independent of i . By Lemma 2.2 the length the Delaunay edges on all the F_i is bounded above. Thus, there are only a finite number of polygonal subdivisions up to isotopy arising as Delaunay partitions in the sequence F_i . Passing to a subsequence, we may hence assume that the Delaunay subdivisions for the F_i are all isotopic, i.e. combinatorially equivalent. Since $\sigma_i \rightarrow \sigma$, lengths of edges converge to positive numbers as $i \rightarrow \infty$. It may still happen that some polygons become degenerate in the limit, i.e. their area goes to zero. If we glue back together the set of non-degenerate limiting polygons, we get a flat structure F which is the limit of the F_i and corresponds to σ . \square

6. Cluster-like structures

6.1. S -graphs

In this section we specialize to the case of flat surfaces with infinite area. It turns out that for generic choice of horizontal direction the heart of the t-structure of the cor-

responding stability structure is Artinian with a finite number of simple objects and described by a certain type of finite graph. Tilting of the t-structure leads to mutation of the graph.

Suppose X is a flat surface of finite type with infinite area. Recall from Section 2.4 that, for generic choice of horizontal direction, the leaves converging to the conical singularities cut the surface into horizontal strips of finite or infinite height. The combinatorial structure of this horizontal strip decomposition can be encoded in terms of a finite graph, G , with set of vertices X_{sg} and set of edges the horizontal strips of finite height. An edge is attached to the singularities which lie on the boundary of the corresponding horizontal strip. We also need to record two additional pieces of data:

1. For each $v \in X_{sg}$ with finite (resp. infinite) cone angle a cyclic (resp. total) order on the set H_v of half-edges meeting that vertex given by the counter-clockwise order of horizontal strips around v .
2. For each pair of successive half-edges $a < b \in H_v$ a positive integer, $d(a, b)$, so that $d(a, b) - 1$ is the number of horizontal strips of infinite height between a and b as we go around v .

We call a graph as above, i.e. a finite graph with orders on H_v and integers $d(a, b)$, an *S-graph*. In order to record the entire flat structure we attach a number $Z(e)$ in the upper half-plane, \mathbf{H} , to each edge e , given by the vector between the singularities on the boundary of the corresponding horizontal strip. We summarize the construction in the following proposition.

Proposition 6.1. — *There is a bijective correspondence between (isomorphism classes of)*

1. *Flat surfaces of finite type with infinite area with the property that there are no horizontal leaves of finite length, and*
2. *S-graphs with choice of number $Z(e) \in \mathbf{H}$ for each edge e .*

The data of an S-graph G may be encoded in terms of a graded quiver with relations Q . Vertices of Q are edges of G . Arrows of Q are given by pairs (h_1, h_2) of successive half edges $h_1 < h_2 \in H_v$ in degree $d(h_1, h_2)$. Relations are quadratic: $(h_1, h_2) \cdot (h_3, h_4) = 0$ when $h_2 \neq h_3$ belong to the same edge. Note that in the graded linear path category $\mathcal{A} = \mathbf{K}Q$ we have

$$(6.1) \quad \text{Hom}^{<0}(\mathbf{A}, \mathbf{B}) = 0, \quad \text{Hom}^0(\mathbf{A}, \mathbf{B}) = \begin{cases} \mathbf{K} & \text{if } \mathbf{A} = \mathbf{B} \\ 0 & \text{else} \end{cases}$$

These properties ensure that the extension closure of the objects of \mathcal{A} in $\text{Tw}\mathcal{A} = \text{Perf}\mathcal{A}$ is the heart of an Artinian bounded t-structure, and objects of \mathcal{A} are precisely the simple ones.

Proposition 6.2. — *Let X be a flat surface with infinite area such that no horizontal leaves have finite length (which is true for generic choice of horizontal direction), and \mathcal{A} be the heart of the t -structure for the stability structure on $\mathcal{F}(X)$ determined by the flat metric, then \mathcal{A} is Artinian with a finite number of simple objects. In terms of the S -graph, \mathcal{A} is described by the graded quiver with relations considered above.*

Proof. — Choose a family X_t , $t \in [0, 1]$ of flat structures so that $X_0 = X$, none of the X_t have finite horizontal leaves, and the simple saddle connections of X_1 are vertical. All the X_t have horizontal strip decompositions described by the same S -graph, only the parameters $Z(e)$ change, and are in $i\mathbf{R}_{>0}$ for X_1 . The corresponding stability structures σ_t all have the same heart of the t -structure, as none of them have semistable objects with real central charge (by the condition on the horizontal leaves), i.e. there is no wall-crossing of the second kind.

The flat surface X_1 has a purely one-dimensional core, which is just the union K of the simple saddle connections. This is clear, since K contains all singularities, and exterior angles are $\geq \pi$. We conclude that X_1 has no closed geodesics and all saddle connections are simple. Thus, stable objects of σ_1 in the heart of the t -structure are just the simple geodesics. Also, the central charge of σ_1 takes values in $i\mathbf{R}$, so simple objects in the heart of the t -structure are the same as stable ones.

The quiver determined by the S -graph is a special case of the graded quiver with relations describing $\text{End}(G)$ of a formal generator of $\mathcal{F}(X)$ (Section 3.4). \square

Mutation

We have assumed above that there are no horizontal saddle connections. Points in the moduli space where this does happen, i.e. some stable object has real central charge, are by definition on walls of the second kind. Let us consider a generic path $\gamma : (-\epsilon, \epsilon) \rightarrow \mathcal{M}(X)$ which crosses one of these walls at time $t = 0$. There are two possibilities, either the foliation of $\gamma(0)$ has a single horizontal saddle connection, e , or some of the leaves are closed loops foliating a cylinder. The latter is a more complicated kind of wall-crossing, and we will restrict attention to the former case.

Making ϵ smaller, we may assume that the horizontal foliations of $\gamma(t)$, $t \neq 0$, have no finite-length leaves. Let G_- , G_+ be the S -graphs for $\gamma(t)$ for $t < 0$, $t > 0$ respectively. The saddle connection e becomes simple for small non-zero values of t , and so can be identified with an edge of both G_- and G_+ . The graphs G_- , G_+ are related by mutation, which we want to describe explicitly. There are two cases: *Left mutation* if the phase of e is moving clockwise, *right mutation* if it is moving counterclockwise. Let us assume the former. Looking at the horizontal strip decompositions one sees that S_+ is obtained from S_- as follows. For a half-edge a denote by $S(a)$ and $S^{-1}(a)$ its successor and predecessor respectively.

1. Let a, b be the half-edges of e . Decrement $d(a, S(a))$ and $d(b, S(b))$ by one.
2. If $d(a, S(a)) = 0$, then slide the end of $S(a)$ along e so that it becomes the new predecessor, c , of b with $d(c, b) = 1$. Otherwise, increment $d(S^{-1}(b), b)$ by one.
3. Do the previous step with a and b switched.




FIG. 7. — Left mutation of the S-graph

The process is illustrated in Figure 7. Right mutation is then simply the inverse operation.

Remark. — In the special case when all zeros are simple we get an ideal triangulation (for generic horizontal direction) by picking a single horizontal leaf from the interior of each horizontal strip as in [23]. A mutation of the S-graph corresponds then to a flip of the triangulation. In general however the mutations described here are not mutations in the sense of cluster theory.

6.2. Example: Dynkin type A_n

Let S be a graded surface with interior diffeomorphic to a disk and with $n + 1 \geq 2$ marked boundary components. Such a graded surface is unique up to isomorphism. The Fukaya category $\mathcal{F}(S)$ is equivalent to the bounded derived category of any A_n -type quiver.

Theorem 6.1. — For S as above of A_n -type,

$$(6.2) \quad \mathcal{M}(S) = \text{Stab}(\mathcal{F}(S)) \cong \mathbf{C}^n.$$

The case $n = 2$ of this theorem is contained in [14].

Proof. — The proof consists of two parts, first to show that $\mathcal{M}(S) \cong \mathbf{C}^n$ as complex manifolds, and then to show that the map $\mathcal{M}(S) \rightarrow \text{Stab}(\mathcal{F}(S))$ is surjective.

1. We may assume that the interior of S is identified with \mathbf{C} , and the grading given by the standard horizontal foliation on \mathbf{C} . Define an analytic map $u : \mathbf{C}^n \rightarrow \mathcal{M}(S)$ by

$$(6.3) \quad (a_0, \dots, a_{n-1}) \mapsto (\exp(z^{n+1} + a_{n-1}z^{n-1} + \dots + a_0)dz^2, a_0).$$

This requires some explanation. By the results of Section 2.5, the quadratic differentials above define a flat structure on \mathbf{C} , the completion of which is, as a topological space, independent of the parameters a_i . We assume that S is the real blow-up of this completion, so that we get a flat structure on S . To determine a point in $\mathcal{M}(S)$ we also need a

homotopy class of paths from the horizontal foliation of the quadratic differential to the standard horizontal foliation on z , which is what a_0 is used for. Since S is contractible, such a homotopy class of paths is determined by its restriction to a single point, which we take to be the origin, where a_0 defines a path $\exp((1-t)a_0)dz^2$ in $(T_0^*\mathbf{C})^{\otimes 2}$, $t \in [0, 1]$.

An analytic automorphism of \mathbf{C} acting on the set of polynomials of the form $z^{n+1} + a_{n-1}z^{n-1} + \dots$ must be of the form $z \mapsto \zeta z$, $\zeta^{n+1} = 1$. All these extend to automorphisms of S , but only for $\zeta = 1$ is it isotopic to the identity. This implies injectivity of u . For surjectivity, let S' be a contractible flat surface with $n+1$ conical points, all with infinite cone angle. By the results of Section 2.5, S' comes from a quadratic differential φ with exponential singularities on a compact Riemann surface C . Since S' is contractible, C must be \mathbf{CP}^1 and φ must have a single exponential singularity. After applying a suitable Möbius transformation, the exponential singularity is at ∞ and φ is of the form $\exp(z^{n+1} + a_{n-1}z^{n-1} + \dots + a_0)dz^2$.

2. Let $\sigma \in \text{Stab}(\mathcal{F}(S))$ be a stability structure, \mathcal{A} the heart of its t-structure. As $\mathcal{F}(S)$ has only a finite number of indecomposable objects up to shift, \mathcal{A} must be Artinian with $n = \text{rk}K_0(\mathcal{F}(S))$ simple objects E_1, \dots, E_n , up to isomorphism. Each E_i corresponds to an isotopy class of graded arcs in S , which we denote by the same symbol. Since $\text{Ext}^{\leq 1}(E_i, E_j) = 0$ for $i \neq j$, we can choose representatives of the E_i which intersect only at the endpoints. Because the E_i give a basis of $H_1(S, \partial S; \mathbf{Z}_\tau)$, they necessarily form a full formal system of graded arcs. Its structure is given by an S-graph which is a tree with n edges. Rotating σ slightly by an angle ϵ , we may assume that no central charge lies in \mathbf{R} . Using Proposition 6.1 we get a flat surface X corresponding to σ , after rotating the horizontal direction back by $-\epsilon$. We already showed in the first part that X belongs to $\mathcal{M}(S)$. \square

Proposition 6.3. — *Let S be the graded marked surface of A_n -type, $n > 1$, then the canonical map $\text{MCG}(S) \rightarrow \text{Aut}(\mathcal{F}(S))$ from the mapping class group of graded automorphisms to the group of autoequivalences up to natural equivalence, is an isomorphism. The group is abelian, isomorphic to $\mathbf{Z}^2/\mathbf{Z}(2, n+1)$.*

In the case $n = 1$ this map is surjective with kernel $\mathbf{Z}/2$.

Proof. — The assumption $n > 1$ guarantees that distinct graded boundary arcs of S correspond to distinct isomorphism classes of objects in $\mathcal{F}(S)$. Since any element of $\text{MCG}(S)$, which are all represented by rotations, is determined by its action on graded boundary arcs, this implies injectivity.

Note that (isomorphism classes of) objects $E \in \mathcal{F}(S)$ corresponding to graded boundary arcs are intrinsically characterized, e.g. by being indecomposables with $\mathcal{F}(S)/E$ equivalent to the bounded derived category of an A_{n-1} quiver. Thus an automorphism Φ of $\mathcal{F}(S)$ induces a map on the set \mathcal{B} of graded boundary arcs. \mathcal{B} has a natural total




FIG. 8. — Possible topologies of \mathbf{K} when $n = 3$ for a generic stability structure. Note that the number of stable objects up to shift is 6, 5, 4, 3 respectively

order, isomorphic to the total order on \mathbf{Z} , so that $\dim \text{Ext}(E_i, E_j) = 1$ if $j = i + 1$, and 0 otherwise. This implies that the induced map on \mathcal{B} is a shift, hence comes from the action of an element of $\text{MCG}(\mathbf{S})$. Finally, any autoequivalence of $\mathcal{F}(\mathbf{S})$ which fixes all objects is seen to be naturally isomorphic to the identity functor. \square

Remark. — Fix a grading on each boundary arc of \mathbf{S} and let $\mathbf{B}_1, \dots, \mathbf{B}_{n+1}$ be the Yoneda-duals of the corresponding objects in $\mathcal{F}(\mathbf{S})$. Their sum is a functor $\mathbf{B} : \mathcal{F}(\mathbf{S}) \rightarrow \mathcal{C} \cong \text{Perf}(\mathbf{K}^{n+1})$. In order to make the statement of the above proposition work for $n = 1$ one needs to replace $\text{Aut}(\mathcal{F}(\mathbf{S}))$ by the group of pairs of autoequivalences, one of $\mathcal{F}(\mathbf{S})$, one of \mathcal{C} , which are compatible with \mathbf{B} , up to natural equivalence. Something like this also seems to be needed in order to generalize the proposition to other marked surfaces. In the context of Fukaya–Seidel categories, \mathbf{B} is the restriction to the distinguished fiber.

Corollary 6.1. — *Let \mathbf{S} be the graded marked surface of \mathbf{A}_n -type, $n > 1$, then*

$$(6.4) \quad \text{Stab}(\mathcal{F}(\mathbf{S})) / \text{Aut}(\mathcal{F}(\mathbf{S})) = \mathcal{M}(\mathbf{S}) / \text{MCG}(\mathbf{S})$$

$$(6.5) \quad = \{e^{\mathbf{P}(z)} dz^2 \mid \deg \mathbf{P} = n + 1\} / \text{Aut}(\mathbf{C})$$

$$(6.6) \quad = \{e^{z^{n+1} + a_{n-1}z^{n-1} + \dots + a_0} dz^2\} / \mathbf{Z} / (n + 1)\mathbf{Z}$$

as complex orbifolds.

In order to gain a better understanding of the wall and chamber structure on $\mathcal{M}(\mathbf{S}) = \text{Stab}(\mathcal{F}(\mathbf{S}))$ one needs to study possible types of cores and S-graphs, giving chambers between walls of the first and second kind respectively. We already know that S-graphs which occur for \mathbf{S} are those with underlying graph a tree with n edges and total order on each set of half-edges. Possible types of cores, even generic ones, are more difficult to enumerate except for small values of n . Figure 8 gives some examples for $n = 3$.

6.3. *Example: affine type $\tilde{\mathbf{A}}_n$*

Fix positive integers p, q and let $n + 1 = p + q$. Let \mathbf{S} be a graded surface with interior an annulus and p (resp. q) marked boundary components, each diffeomorphic

to an open interval, on the left (resp. right) end, and grading so that a simple closed loop around the annulus has vanishing Maslov index, i.e. is gradable. Such S is uniquely determined up to isomorphism. For concreteness we take S to be the real blow-up of the completion of \mathbf{C}^* with respect to the metric $|\varphi|$, where $\varphi = \exp(z^p + z^{-q})dz^2$, and grading given by the horizontal foliation of φ . The Fukaya category $\mathcal{F}(S)$ is equivalent to the bounded derived category of any quiver with underlying graph the extended Dynkin diagram \tilde{A}_n and p (resp. q) arrows directed counterclockwise (resp. clockwise).

Theorem 6.2. — For S as above of \tilde{A}_n -type,

$$(6.7) \quad \mathcal{M}(S) = \text{Stab}(\mathcal{F}(S)) \cong \mathbf{C}^{n+1}.$$

When $n = 1$ the category $\mathcal{F}(S)$ is equivalent to the bounded derived category of coherent sheaves on \mathbf{P}^1 and it was proven by Okada [39] that $\text{Stab}(\mathbf{D}^b(\mathbf{P}^1)) \cong \mathbf{C}^2$. For $n = 2$ it was shown by different methods in [19] that $\text{Stab}(\mathcal{F}(S))$ is contractible. By homological mirror symmetry $\mathcal{F}(S)$ is equivalent to the bounded derived category of the weighted projective line with weights (p, q) (see e.g. [45]) so our result gives a complete description of their spaces of stability conditions.

Proof. — As in the Dynkin case the proof consists of two parts, first to show that $\mathcal{M}(S) = \mathbf{C}^{n+1}$ as complex manifolds, and then to show that the map $\mathcal{M}(S) \rightarrow \text{Stab}(\mathcal{F}(S))$ is surjective.

1. To a point $(a_{p-1}, \dots, a_{-q}) \in \mathbf{C}^{n+1}$ we assign the flat surface given by the quadratic differential

$$(6.8) \quad \exp(z^p + a_{p-1}z^{p-1} + \dots + a_{-q+1}z^{-q+1} + \exp(a_{-q})z^{-q})dz^2$$

on \mathbf{C}^* . Since \mathbf{C}^{n+1} is contractible and the dependence on the parameters continuous, we can resolve the MCG(S)-ambiguity and get a map $u : \mathbf{C}^{n+1} \rightarrow \mathcal{M}(S)$. Arguments as in the proof of Theorem 6.1 show that u is an isomorphism.

2. Let $\sigma \in \text{Stab}(\mathcal{F}(S))$. Unlike the Dynkin case, the heart of the t-structure of σ may not be Artinian. We will show however that this happens non-generically and the t-structure becomes Artinian after possibly rotating σ slightly. Let $\Gamma = \mathbf{K}_0(\mathcal{F}(S)) \cong \mathbf{Z}^{n+1}$, and let $\mathbf{N} \subset \Gamma$ be the subgroup generated by a simple curve around the annulus. Equivalently, \mathbf{N} is the kernel of the Euler form on Γ . Let $\mathbf{E} \subset \Gamma$ be the set of classes of indecomposable objects. It follows from the classification of indecomposable objects that $\mathbf{E}/\mathbf{N} \subset \Gamma/\mathbf{N}$ is finite. Hence, the image of \mathbf{E} under $Z : \Gamma \rightarrow \mathbf{C}$ has finitely many orbits under translation by the vectors $Z(\mathbf{N})$, and thus phases of stable objects can only accumulate towards $\text{Arg}(Z(\mathbf{N}))$, or not at all if $Z(\mathbf{N}) = \{0\}$. In particular, after possibly rotating σ , 0 is not in the closure of the set of phases of stable objects, and thus the

heart of the t-structure is Artinian. From here we can proceed as in the proof of Theorem 6.1. \square

7. Open problems and further directions

Moduli of objects

For a partially wrapped Fukaya category \mathcal{F} of a surface S our classification describes $\text{Ob}(\mathcal{F})$ as a set in terms of curves on S . However, $\text{Ob}(\mathcal{F})$ has a much richer structure of a derived stack.

Problem. — Describe the geometry of $\text{Ob}(\mathcal{F})$.

Also, the case of Fukaya categories of closed surfaces remains. This would be an A -side generalization of Atiyah's classification for the elliptic curve. However, for $g > 1$ the category is only $\mathbf{Z}/2$ -graded and there are no t-structures or stability structures.

Problem. — Classify objects of $\mathcal{F}(S)$ when S is a closed surface with $g(S) > 1$, assuming this is a tame problem.

Higher dimensions

There is at present no complete conjectural picture extending the results of this paper to higher dimensional symplectic manifolds. Investigations into this direction were started by Thomas [48] and Thomas–Yau [49] with the aim of finding an algebraic notion of stability describing special Lagrangian submanifolds and their mean curvature flow. More recently, their proposal has been updated by Joyce [28], taking into account recent developments on Fukaya categories, mean curvature flow, and Bridgeland's axiomatics. The conjecture is that given a compact Calabi–Yau M one has a stability structure on its Fukaya category, $\mathcal{F}(M)$, such that singular special Lagrangians in M , with additional data to make them objects in $\mathcal{F}(M)$, are (all?) stable objects. The central charge is given by integration of the holomorphic volume form. Mean curvature flow applied to an object in $\mathcal{F}(S)$, at least with sufficiently small phase variation, should converge to the components of the Harder–Narasimhan filtration. It is however still an open problem to describe (a component of) $\text{Stab}(\mathcal{F}(M))$ in general, even conjecturally.

Problem. — Find a geometric description of $\text{Stab}(\mathcal{F}(M))$ for higher dimensional M .

Acknowledgements

The idea that quadratic differentials with exponential-type singularities can be identified with stability structures was suggested to the authors by Alexander Polishchuk.

We also thank Mohammed Abouzaid, Denis Auroux, Tom Bridgeland, Igor Burban, Tobias Dyckerhoff, Alexander Efimov, Alexander Goncharov, Dominic Joyce, Mikhail Kapranov, Gabriel Kerr, Andrew Neitzke, Tony Pantev, Yan Soibelman, Pawel Sosna, and Zachary Sylvan for discussions relating to this paper. Our thanks also go to the anonymous referees for carefully reading the manuscript and providing helpful suggestions. The authors were partially supported by the FRG grant DMS-0854977, the research grants DMS-0854977, DMS-0901330, DMS-1201321, DMS-1201475, DMS-1265230, OISE-1242272 PASI from the National Science Foundation, the FWF grant P24572-N25, by an ERC GEMIS grant, and by an RF Government grand, ag. No. 14.641.31.0001. We are grateful to the IHES for providing perfect working conditions. The authors were supported by two Simons Foundation grants.

REFERENCES

1. M. ABOUZAIID, On the Fukaya categories of higher genus surfaces, *Adv. Math.*, **217** (2008), 1192–1235.
2. M. ABOUZAIID and P. SEIDEL, An open string analogue of Viterbo functoriality, *Geom. Topol.*, **14** (2010), 627–718.
3. I. ASSEM and A. SKOWROŃSKI, Iterated tilted algebras of type \tilde{A}_n , *Math. Z.*, **195** (1987), 269–290.
4. M. F. ATIYAH, Vector bundles over an elliptic curve, *Proc. Lond. Math. Soc.*, **3** (1957), 414–452.
5. D. AUROUX, Fukaya categories and bordered Heegaard–Floer homology, in *Proceedings of the International Congress of Mathematicians, vol. II*, pp. 917–941, Hindustan Book Agency, Gurugram, 2010.
6. M. BARDZELL, The alternating syzygy behavior of monomial algebras, *J. Algebra*, **188** (1997), 69–89.
7. A. BAYER, E. MACRÌ and Y. TODA, Bridgeland stability conditions on threefolds I: Bogomolov-Gieseker type inequalities, *J. Algebraic Geom.*, **23** (2014), 117–163.
8. V. BEKKERT and H. A. MERKLEN, Indecomposables in derived categories of gentle algebras, *Algebr. Represent. Theory*, **6** (2003), 285–302.
9. R. BOCKLANDT, Noncommutative mirror symmetry for punctured surfaces, *Trans. Am. Math. Soc.*, **368** (2016), 429–469.
10. J. P. BOWMAN and F. VALDEZ, Wild singularities of flat surfaces, *Isr. J. Math.*, **197** (2013), 69–97.
11. T. BRIDGELAND, Stability conditions on triangulated categories, *Ann. Math.*, **166** (2007), 317–345.
12. T. BRIDGELAND, Stability conditions on K3 surfaces, *Duke Math. J.*, **141** (2008), 241–291.
13. T. BRIDGELAND, Stability conditions and Kleinian singularities, *Int. Math. Res. Not.*, **21** (2009), 4142–4157.
14. T. BRIDGELAND, Y. QIU and T. SUTHERLAND, Stability conditions on the A_2 quiver, [arXiv:1406.2566](https://arxiv.org/abs/1406.2566).
15. T. BRIDGELAND and I. SMITH, Quadratic differentials as stability conditions, *Publ. Math. IHÉS*, **121** (2015), 155–278.
16. I. BURBAN and Y. DROZD, On derived categories of certain associative algebras, in *Representations of Algebras and Related Topics*, Fields Inst. Commun., vol. 45, pp. 109–128, Am. Math. Soc., Providence, 2005.
17. I. BURBAN and B. KREUSSLER, Derived categories of irreducible projective curves of arithmetic genus one, *Compos. Math.*, **142** (2006), 1231–1262.
18. G. DIMITROV, F. HAIDEN, L. KATZARKOV and M. KONTSEVICH, Dynamical systems and categories, in *The Influence of Solomon Lefschetz in Geometry and Topology: 50 Years of Mathematics at CINVESTAV*, Con. Math., vol. 621, pp. 133–170, Am. Math. Soc., Providence, 2014.
19. G. DIMITROV and L. KATZARKOV, Stability conditions on the acyclic triangular quiver, [arXiv:1410.0904](https://arxiv.org/abs/1410.0904).
20. T. DYCKERHOFF, A1-homotopy invariants of topological Fukaya categories of surfaces, [arXiv:1505.06941](https://arxiv.org/abs/1505.06941).
21. T. DYCKERHOFF and M. KAPRANOV, Triangulated surfaces in triangulated categories, [arXiv:1306.2545](https://arxiv.org/abs/1306.2545).
22. K. FUKAYA, Y.-G. OH, H. OHTA and K. ONO, *Lagrangian Intersection Floer Theory, Anomaly and Obstruction, Parts I, II*, AMS/IP Studies in Adv. Math., vol. 46.1, 2009.
23. D. GAIOTTO, G. W. MOORE and A. NEITZKE, Wall-crossing, Hitchin systems, and the WKB approximation, *Adv. Math.*, **234** (2013), 239–403.
24. A. GROTHENDIECK, Groupes de classes des categories abeliennes et triangulees. Complexes parfaits (Redige par I. Bucur), in *Semin. Geom. Algebr. Bois-Marie 1965–1966, SGA 5*, Lect. Notes Math., vol. 589, Expose No. VIII, pp. 351–371, 1977.

25. J. L. HARER, Stability of the homology of the mapping class groups of orientable surfaces, *Ann. Math. (2)*, **121** (1985), 215–249.
26. J. L. HARER, The virtual cohomological dimension of the mapping class group of an orientable surface, *Invent. Math.*, **84** (1986), 157–176.
27. A. IKEDA, Stability conditions on CY_n categories associated to A_n -quivers and period maps, [arXiv:1405.5492](https://arxiv.org/abs/1405.5492).
28. D. JOYCE, Conjectures on Bridgeland stability for Fukaya categories of Calabi–Yau manifolds, special Lagrangians, and Lagrangian mean curvature flow, [arXiv:1401.4949](https://arxiv.org/abs/1401.4949).
29. H. KAJIURA, K. SAITO and A. TAKAHASHI, Matrix factorization and representations of quivers. II. Type ADE case, *Adv. Math.*, **211** (2007), 327–362.
30. M. KONTSEVICH, Homological algebra of mirror symmetry, in *Proceedings of the International Congress of Mathematicians, vols. 1, 2, Zürich, 1994*, pp. 120–139, Birkhäuser, Basel, 1995.
31. M. KONTSEVICH and Y. SOIBELMAN, Stability structures, motivic Donaldson-Thomas invariants and cluster transformations, [arXiv:0811.2435](https://arxiv.org/abs/0811.2435).
32. M. KONTSEVICH and Y. SOIBELMAN, Notes on A_∞ -algebras, A_∞ -categories and non-commutative geometry, in *Homological Mirror Symmetry*, Lecture Notes in Phys., vol. 757, pp. 153–219, Springer, Berlin, 2009.
33. V. LYUBASHENKO and S. OVSIENKO, A construction of quotient A_∞ -categories, *Homol. Homotopy Appl.*, **8** (2006), 157–203.
34. E. MACRÌ, Stability conditions on curves, *Math. Res. Lett.*, **14** (2007), 657–672.
35. H. MASUR and J. SMILLIE, Hausdorff dimension of sets of nonergodic measured foliations, *Ann. Math. (2)*, **134** (1991), 455–543.
36. D. NADLER, Cyclic symmetries of A_n -quiver representations, [arXiv:1306.0070](https://arxiv.org/abs/1306.0070).
37. L. A. NAZAROVA and A. V. ROÏTER, A certain problem of I.M. Gel’fand, *Funk. Anal. Prilozh.*, **7** (1973), 54–69.
38. R. NEVANLINNA, Über Riemannsche Flächen mit endlich vielen Windungspunkten, *Acta Math.*, **58** (1932), 295–373.
39. S. OKADA, Stability manifold of \mathbf{P}^1 , *J. Algebraic Geom.*, **15** (2006), 487–505.
40. J. PASCALEFF and N. SIBILLA, Topological fukaya category and mirror symmetry for punctured surfaces, [arXiv:1604.06448](https://arxiv.org/abs/1604.06448).
41. A. POLISHCHUK and E. ZASLOW, Categorical mirror symmetry: the elliptic curve, *Adv. Theor. Math. Phys.*, **2** (1998), 443–470.
42. P. SEIDEL, Fukaya A_∞ structures associated to Lefschetz fibrations. II, [arXiv:1404.1352](https://arxiv.org/abs/1404.1352).
43. P. SEIDEL, Graded Lagrangian submanifolds, *Bull. Soc. Math. Fr.*, **128** (2000), 103–149.
44. P. SEIDEL, *Fukaya Categories and Picard-Lefschetz Theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zürich, 2008.
45. N. SIBILLA, D. TREUMANN and E. ZASLOW, Ribbon graphs and mirror symmetry, *Sel. Math. New Ser.*, **20** (2014), 979–1002.
46. I. SMITH, Quiver algebras as Fukaya categories, [arXiv:1309.0452](https://arxiv.org/abs/1309.0452).
47. K. STREBEL, *Quadratic Differentials*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 5, Springer, Berlin, 1984.
48. R. P. THOMAS, Moment maps, monodromy and mirror manifolds, in *Symplectic Geometry and Mirror Symmetry, Seoul, 2000*, pp. 467–498, World Scientific, River Edge, 2001.
49. R. P. THOMAS and S.-T. YAU, Special Lagrangians, stable bundles and mean curvature flow, *Commun. Anal. Geom.*, **10** (2002), 1075–1113.
50. W. A. VEECH, Flat surfaces, *Am. J. Math.*, **115** (1993), 589–689.

F. H.

Department of Mathematics, Science Center,
 Harvard University,
 One Oxford Street,
 Cambridge, MA 02138, USA
haiden@math.harvard.edu

L. K.
Fakultät für Mathematik,
Universität Wien,
Oskar-Morgenstern-Platz 1,
1090 Wien, Austria
lkatzarkov@gmail.com
and
HSE Moscow,
Moscow, Russia

M. K.
Institut des Hautes Études Scientifiques,
35 route de Chartres,
91440 Bures-sur-Yvette, France
maxim@ihes.fr

Manuscrit reçu le 23 mai 2016
Manuscrit accepté le 14 novembre 2017
publié en ligne le 21 novembre 2017.