Teichmüller harmonic map flow

Harmonic maps, minimal immersions and their gradient flows

Peter Topping

** OUTLINE NOTES, NEITHER FINISHED NOR PROOF-READ **

** PLEASE DO NOT POST AND DO NOT DISTRIBUTE **

** PLEASE ALERT ME TO TYPOS **

January 12, 2016

http://www.warwick.ac.uk/~maseq

Contents

Overview

I hope in these lectures to give an introduction to some core topics in geometric analysis and differential geometry, and to combine them to give a tour of the so-called Teichmüller harmonic map flow, which is a geometric flow that tries to find branched minimal immersions. Topics that we will cover include:

1. The harmonic map energy; its properties for maps from surfaces, and links with minimal immersions.
2. The harmonic map flow of Eells-Sampson, and its theory in two dimensions. Bubbling.
3. Uniformisation/geometrisation of surfaces; preliminaries on the space of constant curvature metrics and its tangent space.
4. Definition of the Teichmüller harmonic map flow.
5. The structure of hyperbolic metrics and their degeneration.
6. A tour of the Teichmüller harmonic map flow - basic theory, asymptotics, singularities.

Much of the lectures consist of material that is central to a large amount of mathematics. In the final part, the Teichmüller harmonic map flow provides a topic where all these ideas are used in one theory. This latter theory was developed jointly with Melanie Rupflin, with contributions from Miaomiao Zhu and Tobias Huxol.

Prerequisites include basic Riemannian geometry and PDE theory. (No prior knowledge of Teichmüller theory is required.)

Chapter 1

The harmonic map energy

1.1 Definition of energy and area

Given a sufficiently smooth map $u : M \to N$ between smooth closed Riemannian manifolds (M, g) and $N = (N, G)$ of any dimension, we can define the harmonic map energy

$$E(u) = E(u, g) := \frac{1}{2}\int_M |du|_g^2 d\mu_g.$$

Smooth critical points are known as harmonic maps. By this, we mean that u is a harmonic map if it is smooth, and for any smooth one-parameter family of maps $u_\sigma : M \to N$, for σ in some neighbourhood of $0 \in \mathbb{R}$, with $u_0 = u$, we have

$$\left.\frac{d}{d\sigma}\right|_{\sigma=0} E(u_\sigma) = 0.$$

(In particular, we do not bother viewing E as a function on some infinite dimensional manifold of maps.)

Examples: if $M = S^1$, then nonconstant harmonic maps are closed geodesics in (N, G). Constant maps are always harmonic. In the special case that $(M, g) = (N, G)$, the identity map is always harmonic, as we will see later.

Later we will consider also complete, noncompact domains; we could also allow the energy to be infinite. You can still define harmonic maps by considering compactly supported variations of u, and asking that the energy restricted to the support of the variation does not change to first order. We could also consider complete noncompact targets; e.g. we could set $N = \mathbb{R}$, and Dirichlet's principle then tells us that harmonic maps are precisely the harmonic functions.

For most of the rest of these lectures we will only consider the case that the domain M is an oriented surface, and normally closed, say of genus γ.

In this case, near each point in the domain we can find *isothermal* coordinates x and y, i.e. so that the domain metric g can be written $\rho^2(dx^2 + dy^2)$ for some positive function $\rho : M \to (0, \infty)$. Then we can write $\frac{1}{2}|du|_g^2 = \frac{1}{2\rho^2}(|u_x|^2 + |u_y|^2)$, where $u_x := u_*(\frac{\partial}{\partial x})$ etc., and $d\mu_g = \rho^2 dxdy$, and in particular the energy integrand $|du|_g^2 d\mu_g = \frac{1}{2}(|u_x|^2 + |u_y|^2)dxdy$ does not depend on the conformal factor ρ, only the conformal structure.

A particularly significant point about two-dimensional domains is that the energy is closely related to the *area* functional $A(u)$, defined, for smooth enough u by

$$A(u) := \int_M du \wedge du,$$

where the ambiguous quantity $du \wedge du$ could be defined in our local isothermal coordinates to be

$$du \wedge du := |u_x \wedge u_y| dxdy$$

using the normalisation that if e_1 and e_2 are two orthogonal unit vectors in some tangent space in the target, then $|e_1 \wedge e_2| = 1$. When u is an immersion, we could write this as

$$A(u) = \int_M d\mu_{u^*G},$$

giving the area of the image of u, counted with multiplicity.

Smooth immersions u that are critical points of the area functional $A(u)$ are known as minimal immersions. You can consider them the two-dimensional counterpart of (one-dimensional) geodesics.

As for being harmonic, being minimal is a local property – we can talk about an immersion from a portion of a surface being minimal if locally supported smooth variations do not change the local area to first order. This property can be characterised by the image having zero mean curvature (if you know what that means – if not, don't worry).

Note that these quantities E and A depend on the metric G, but we will always consider that fixed. On the other hand, the energy E but not the area A will depend also on g.

1.2 Relationship between energy and area and their critical points

By estimating pointwise

$$|u_x \wedge u_y| \leq |u_x|.|u_y| \leq \frac{1}{2}(|u_x|^2 + |u_y|^2),$$

we see that

$$A(u) \leq E(u, g). \tag{1.2.1}$$

An immersion u will have equality in (1.2.1) if and only if it is *conformal* with respect to g, i.e iff there exists a positive function $f : M \to (0, \infty)$ such that $u^*G = fg$, i.e. iff orthogonal vectors are pushed forward to orthogonal vectors by u.

It is then easy to see (draw graph of E and A) that:

Lemma 1.2.1. *For two-dimensional domains, conformal harmonic maps are necessarily minimal immersions, and minimal immersions are (once the domain is equipped with the metric u^*G) conformal harmonic maps.*

For more general (smooth enough) maps u – not necessarily immersions – we have equality in (1.2.1) if and only if u is *weakly conformal* with respect to g, i.e iff there exists a nonnegative function $f : M \to [0,\infty)$ such that $u^*G = fg$

Definition 1.2.2. The map u is said to be a *branched minimal immersion* if it is a nonconstant weakly conformal harmonic map.

As for Lemma 1.2.1, a branched minimal immersion will be a critical point for A, but the the terminology is mainly justified by the following result of Gulliver-Osserman-Royden [11]:

Theorem 1.2.3. *A branched minimal immersion u in the sense above is a genuine minimal immersion except at* isolated points *where $du = 0$. For each of these isolated points $p \in M$, we can take isothermal coordinates x, y centred at p, and normal coordinates centred at $u(p)$, in which u takes the form*

$$u^1(z) = Re(z^k) + o(|z|^k)$$
$$u^2(z) = Im(z^k) + o(|z|^k)$$
$$u^\alpha(z) = o(|z|^k) \qquad \text{for } \alpha \geq 3,$$

for some $k \in \{2, 3, \ldots\}$, where $z = x + iy$.

1.3 Complex viewpoint

Summary: It will be convenient to 'complexify' the tangent bundle of our domain - i.e. consider elements $v_1 + iv_2$ where v_1 and v_2 are real tangent vectors, and also do the same for the cotangent bundle. Take local isothermal coordinates x and y, and define the corresponding local complex coordinate $z = x+iy$. There is a natural invariantly defined decomposition of the complex cotangent bundle into subbundles spanned by the complex covectors $dz := dx+idy$ and $d\bar{z} := dx-idy$. These are the so-called $(1,0)$ and $(0,1)$ parts of the complexified cotangent bundle. The dual basis to $\{dz, d\bar{z}\}$ is given by $\{\frac{\partial}{\partial z}, \frac{\partial}{\partial \bar{z}}\}$ where (beware signs!) $\frac{\partial}{\partial z} := \frac{1}{2}\left(\frac{\partial}{\partial x} - i\frac{\partial}{\partial y}\right)$ and $\frac{\partial}{\partial \bar{z}} := \frac{1}{2}\left(\frac{\partial}{\partial x} + i\frac{\partial}{\partial y}\right)$. We also abbreviate $f_z := \frac{\partial f}{\partial z}$ for a function $f : M \to \mathbb{C}$, and $u_z := u_*(\frac{\partial}{\partial z}) = u_x - iu_y$ for a map $u : M \to N$ etc.

Details: (..that we'll need also in other settings..)

It will pay to make a small investment in algebra. We can view a complex vector space of complex dimension n equivalently as a $2n$-dimensional real vector space V equipped with an endomorphism $J : V \to V$ such that $J^2 = -\text{Identity}$, i.e. a *complex structure*, corresponding to multiplication by i. Having eradicated the i, we can then complexify – i.e. take $V^c := V \otimes_{\mathbb{R}} \mathbb{C}$ – which amounts to considering elements of V with complex coefficients. The endomorphism J extends by complex linearity to V^c. The action of J naturally decomposes V^c into the eigenspaces $V^{1,0}$ and $V^{0,1}$ corresponding to the eigenvalues i and $-i$ respectively.

Note that $V^{1,0}$ is a complex vector space with complex structure given by multiplication by i, but by definition, equivalently by action of J.

We can write

$$V^{1,0} = \{v - iJv \ : \ v \in V\} \qquad \text{and} \qquad V^{0,1} = \{v + iJv \ : \ v \in V\}. \tag{1.3.1}$$

Complex conjugation is a real linear isomorphism from one to the other.

Meanwhile the dual space V^* inherits a complex structure from V, defined by $(JL)(v) := L(Jv)$ for all $L \in V^*, v \in V$. Again, we can complexify V^* to give V^{*c}, and we see that $(V^*)^c = (V^c)^*$. And again, we can decompose V^{*c} into its two eigenspaces $V_{1,0}$ and $V_{0,1}$ corresponding to eigenvalues i and $-i$ respectively of the extended J.

Each element in $V_{1,0}$ has all of $V^{0,1}$ in its kernel, and similarly elements in $V_{0,1}$ kill all of $V^{1,0}$.

We can then decompose arbitrary tensor products of V and V^*, and their complexification, into tensor products of these subspaces. Of particular interest for us in a moment will be to consider the part of the decomposition of elements of $V^* \otimes V^*$ and its complexification that lie in $V^{1,0} \otimes V^{1,0}$ – this would be the so-called $(2,0)$ part of the element.

Any inner product g on V can be extended by complex linearity to a bilinear form on V^c, and we typically still call this g, or write $(\cdot, \cdot)$.

This extended g also induces a Hermitian inner product on V^c defined by $\langle v, w \rangle = g(v, \overline{w})$.

The decomposition above of complex vectors and covectors will automatically decompose tensor fields on manifolds equipped with an *almost complex structure* (i.e. a smoothly varying complex structure on each fibre) by making the decomposition above on each fibre.

Our first (but not only) application of the above abstraction is very simple. Given an oriented Riemannian surface, e.g. our domain (M, g), there is a natural complex structure J on each tangent space, defined intuitively by anticlockwise rotation by $90°$, or more precisely so that $\{v, Jv\}$ is a positively oriented orthonormal basis for each unit tangent vector v. By taking local isothermal coordinates x and y so $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ are positively oriented, we can define local complex coordinates $z = x + iy$ and turn the Riemannian surface into a Riemann surface, at which point we have

$$J\left(\frac{\partial}{\partial x}\right) = \frac{\partial}{\partial y} \qquad \text{and} \qquad J\left(\frac{\partial}{\partial y}\right) = -\frac{\partial}{\partial x},$$

and (beware signs!)

$$Jdx = -dy \qquad \text{and} \qquad Jdy = dx,$$

on each tangent space. (Only the Riemann surface structure, not the full Riemannian metric, will be relevant in the discussion of this section. That is, we use the complex structure, but not the conformal factor ρ.)

The $(1,0)$ and $(0,1)$ parts of each (complexified) *cotangent* space – often called the holomorphic cotangent space and the antiholomorphic cotangent space respectively – are spanned by $dz = dx + idy$ and $d\bar{z} = dx - idy$ respectively, cf. (1.3.1). These make up a basis for the (complexified) cotangent space, and so we can take its dual basis, which will be $\{\frac{\partial}{\partial z}, \frac{\partial}{\partial \bar{z}}\}$ where (beware signs!) $\frac{\partial}{\partial z} := \frac{1}{2}\left(\frac{\partial}{\partial x} - i\frac{\partial}{\partial y}\right)$ and $\frac{\partial}{\partial \bar{z}} := \frac{1}{2}\left(\frac{\partial}{\partial x} + i\frac{\partial}{\partial y}\right)$. That is, $dz\left(\frac{\partial}{\partial z}\right) = 1$ and $dz\left(\frac{\partial}{\partial \bar{z}}\right) = 0$ etc. We see that $\frac{\partial}{\partial z}$ and $\frac{\partial}{\partial \bar{z}}$ span the $(1,0)$ and $(0,1)$ parts of each (complexified) *tangent* space – also known as the holomorphic tangent space and the antiholomorphic tangent space respectively.

Note that these elements are 'isotropic', e.g. $(\frac{\partial}{\partial z}, \frac{\partial}{\partial z}) = 0$. Also, $\mathrm{tr}[dz \otimes dz] = 0$.

On the other hand, we have $(\frac{\partial}{\partial z}, \frac{\partial}{\partial \bar{z}}) = \frac{\rho^2}{2}$, and $\mathrm{tr}[d\bar{z} \otimes dz] = \frac{2}{\rho^2}$, etc.

We also abbreviate $f_z := \frac{\partial f}{\partial z}$ for a function $f : M \to \mathbb{C}$, and $u_z := u_*(\frac{\partial}{\partial z}) = u_x - iu_y$ for a map $u : M \to N$ etc.

General tensors on surfaces can be decomposed in the bases above. Of particular interest will be when we take a symmetric tensor in the complexification of $T^*M \otimes T^*M$, in which case it can be written in terms of dz^2, $d\bar{z}^2$ and $|dz|^2$, where dz^2 is shorthand for $dz \otimes dz$, $d\bar{z}^2 := d\bar{z} \otimes d\bar{z}$ and $|dz|^2 := \frac{1}{2}(dz \otimes d\bar{z} + d\bar{z} \otimes dz) = dx^2 + dy^2$.

Such tensors that have only $(2, 0)$ part – i.e. that can be written locally as ωdz^2 for some locally defined complex valued function ω, are known as *quadratic differentials*. We will shortly define what it means for a tensor to be holomorphic, but in this case it will amount to the local function ω being holomorphic. Beware that some sources use the term quadratic differential to mean holomorphic quadratic differential.

The Levi-Civita connection of (M, g) can also be extended complex linearly to the complexification of the tangent and cotangent bundles of the surface and their tensor products and then decomposed as

$$\nabla = \bar{\partial} + \partial,$$

where ∂ consists of applying the connection ∇ and then taking only the $(1, 0)$ part of the new T^*M component of the tensor, and $\bar{\partial}$ is the same for the $(0, 1)$ part.

For example, applied to functions $f : M \to \mathbb{C}$, we have $\partial f = f_z dz$ and $\bar{\partial} f = f_{\bar{z}} d\bar{z}$ (so $d = \nabla = \partial + \bar{\partial}$). Or we could write

$$\partial f := (df)^{(1,0)} = \frac{1}{2}(df - iJdf),$$

where $(df)^{(1,0)}$ refers to the projection of df in the cotangent bundle (within the complexified cotangent bundle) onto the $(1, 0)$ part of that bundle.

Also, a short calculation shows that $\partial dz = dz \otimes \nabla_{\frac{\partial}{\partial z}} dz = -\frac{2\rho_z}{\rho} dz^2$ and $\bar{\partial} dz = 0$.

We say a tensor T is holomorphic if $\bar{\partial} T = 0$. Applying this to quadratic differentials ωdz^2, and using the fact that $\bar{\partial}(\omega dz^2) = \omega_{\bar{z}} dz^2$ because $\bar{\partial} dz = 0$, we see (as claimed earlier) that a holomorphic quadratic differential is one for which the locally-defined function ω is always holomorphic.

Consider now the divergence operator $\delta := -\text{tr}\nabla$ acting on sections of tensor products of T^*M, and hence also on sections of the *complexification* of tensor products of T^*M. That is, we write $\delta T_{i_1 \ldots i_k} = -\nabla_{i_1} T_{i_1 \ldots i_k}$ (see e.g. [33, §2.1]). But we saw above that $\nabla dz = \frac{2\rho_z}{\rho} dz^2$, and $\text{tr} dz^2 = 0$, and hence

$$\delta dz = 0, \tag{1.3.2}$$

and moreover

$$\delta dz^2 = (\delta dz) \otimes dz - \text{tr}[dz \otimes \nabla dz] = 0, \tag{1.3.3}$$

and thus

$$\delta(\phi dz^2) = -\text{tr}(\phi_z dz^3 + \phi_{\bar{z}} d\bar{z} \otimes dz^2) = -\frac{2\phi_{\bar{z}}}{\rho^2} dz. \tag{1.3.4}$$

1.4 The Hopf differential

Given a *real* symmetric tensor $h \in \Gamma(\mathrm{Sym}^2 T^*M)$ we can write it

$$h = \frac{1}{4}\phi dz^2 + a|dz|^2 + \frac{1}{4}\bar{\phi} d\bar{z}^2, \tag{1.4.1}$$

where the constants have been chosen to match future notation. We see that $a = 2h(\frac{\partial}{\partial z}, \frac{\partial}{\partial \bar{z}})$ and $\phi = 4h(\frac{\partial}{\partial z}, \frac{\partial}{\partial z})$.

Let's consider the special case that $h = u^*G$. Then, (extending G by complex linearity to the complexified tangent bundle as usual), we have

$$a = 2u^*G\left(\frac{\partial}{\partial z}, \frac{\partial}{\partial \bar{z}}\right) = \frac{1}{2}|du|^2$$

and

$$\phi = 4u^*G\left(\frac{\partial}{\partial z}, \frac{\partial}{\partial z}\right) = \left(|u_x|^2 - |u_y|^2 - 2i\langle u_x, u_y\rangle\right).$$

Thus, the coefficient of the $|dz|^2$ component is the energy density of u with respect to the local coordinates. The dz^2 component is of particular consequence for us. Indeed, ϕdz^2 is known as the Hopf differential, and we often write it $\Phi(u, g)$ when suppressing coordinates.

Directly from the definition, we see that u is weakly conformal if and only if the Hopf differential vanishes everywhere.

In particular, a nonconstant harmonic map that has vanishing Hopf differential is a branched minimal immersion in the sense described earlier.

1.5 First variation of energy

Given a map u and metric g on the domain surface as above, we wish to know the rate of change of energy E as we move u and g in arbitrary directions that we denote $\frac{\partial u}{\partial t}$ and $\frac{\partial g}{\partial t}$.

Lemma 1.5.1.

$$dE_{(u,g)}\left(\frac{\partial u}{\partial t}, \frac{\partial g}{\partial t}\right) = -\int_M \left(\langle \tau_g(u), \frac{\partial u}{\partial t}\rangle + \frac{1}{4}\langle Re(\Phi(u,g)), \frac{\partial g}{\partial t}\rangle\right) d\mu_g.$$

Here $\Phi(u, g) = \phi dz^2 := 4(u^*G)^{(2,0)}$ is the Hopf differential of u as above, and $\tau_g(u) := \mathrm{tr}\nabla du \in \Gamma(u^*(TN))$ is the tension field [8]. We can write the tension field in local coordinates $\{y^\alpha\}$ on the target as

$$\tau_g(u) = \left(\Delta_{(M,g)} u^\alpha +^N \Gamma^\alpha_{\beta\gamma} \frac{\partial u^\beta}{\partial x^i}\frac{\partial u^\gamma}{\partial x^j} g^{ij}\right) \frac{\partial}{\partial y^\alpha}(u).$$

i.e.

$$\tau_g(u) = g^{ij}\left(\frac{\partial^2 u^\alpha}{\partial x^i \partial x^j} -^M \Gamma^k_{ij}\frac{\partial u^\alpha}{\partial x^k} +^N \Gamma^\alpha_{\beta\gamma}\frac{\partial u^\beta}{\partial x^i}\frac{\partial u^\gamma}{\partial x^j}\right)\frac{\partial}{\partial y^\alpha}(u).$$

Note that this last formulation of the tension shows immediately that the identity map is always harmonic, because we can compute in normal coordinates. (This was promised at the start of the course.)

Think of the tension as the right notion of Laplacian of u.

From this, we can see that if we fix g and consider critical points of E with respect to variations in u – i.e. consider harmonic maps – then the Euler-Lagrange equation is simply

$$\tau_g(u) = 0.$$

On the other hand to reduce E as quickly as possible by varying g (more precisely with respect to the L^2 inner product) we should move g in the direction of the real part of the Hopf differential.

Of course, the real part of a quadratic differential is well defined, and does not involve using coordinates. Ultimately a quadratic differential is just a real section of $T^*M \otimes T^*M$ plus i times another section of the same bundle, so the real part is just the first of these sections.

On the other hand, by the particular form of quadratic differentials, we see that they are zero if and only if their real part is zero, which may be counterintuitive if you have only ever thought about complex *functions*. For example, working with respect to a local complex coordinate z, we can write the differential as $(a + ib)dz^2$ for locally defined real functions a and b, and so

$$Re[(a + ib)dz^2] = Re[(a + ib)(dx^2 - dy^2 + 2idxdy)] = a(dx^2 - dy^2) - 2b\,dxdy.$$

This tensor can only be zero if both a and b are zero.

An important but immediate first consequence of the first variation formula, Lemma 1.5.1, is:

Lemma 1.5.2. *Critical points of E with respect to variations both of u and g are branched minimal immersions, when they are not constant maps.*

Much of these lectures are devoted to finding such objects.

Another useful consequence of this first variation formula:

Lemma 1.5.3. *A harmonic map from a surface has holomorphic Hopf differential ϕdz^2. More generally, given a C^2 (say) map u, and writing $g = \rho^2|dz|^2$, we have*

$$\phi_{\bar{z}} = 2\rho^2(\tau_g(u), u_z).$$

We've given the formula in coordinates because it's quickest to assimilate without ambiguity in conventions. In more geometric language it would be

$$\mathrm{tr}\bar{\partial}\Phi = 4(\tau_g(u), \partial u),$$

where ∂u, as you can guess, is obtained by taking the $(1, 0)$ part of the T^*M part of du.

Proof. Diffeomorphism invariance in Riemannian geometry leads to conservation laws (e.g. second Bianchi identity). That is what is going on here too.

Take an arbitrary vector field X on M, with compact support. We can use this to generate a family of diffeomorphisms ψ_s for $s \in (-\varepsilon, \varepsilon)$ say, with ψ_0 the identity. The point is that the energy will not change if we pull back u and g by a diffeomorphism, and in particular by ψ_s. Differentiating with respect to s at $s = 0$, this amounts to saying that

$$dE_{(u,g)}\,(du(X), \mathcal{L}_X g) = 0.$$

Applying the first variation formula, we find that

$$\begin{aligned} 0 &= \int_M \langle \tau_g(u), du(X)\rangle + \frac{1}{4}\langle Re(\Phi(u,g)), \mathcal{L}_X g\rangle \\ &= \int_M \langle \tau_g(u), du(X)\rangle + \frac{1}{2}\langle Re(\delta(\phi dz^2)), X\rangle \end{aligned} \tag{1.5.1}$$

since if we consider $\delta : \Gamma(\mathrm{Sym}^2 T^*M) \to \Gamma(T^*M)$, i.e. the divergence operator $\delta := -\mathrm{tr}\nabla$ restricted to *symmetric* $(0,2)$ tensors then its formal adjoint $\delta^* : \Gamma(T^*M) \to \Gamma(\mathrm{Sym}^2 T^*M)$ is given by $\omega \mapsto \frac{1}{2}\mathcal{L}_{\omega^\#} g$. But X was arbitrary, so recalling (1.3.4), we have

$$\begin{aligned} 0 &= \langle \tau_g(u), du\rangle + \frac{1}{2} Re(\delta(\phi dz^2)) \\ &= \langle \tau_g(u), du\rangle - \frac{1}{\rho^2} Re(\phi_{\bar z} dz) \\ &= (\tau_g(u), u_z dz + u_{\bar z} d\bar z) - \frac{1}{2\rho^2}(\phi_{\bar z} dz + \phi_z d\bar z). \end{aligned} \tag{1.5.2}$$

Taking the dz part completes the proof. □

We will spend quite some time later describing the space of holomorphic quadratic differentials. However, one simple fact to get going with is:

Lemma 1.5.4. *The only holomorphic quadratic differential on S^2 is identically zero.*

Proof. A low-tech proof is as follows. We can find a complex coordinate z on the whole sphere minus one point. Indeed, if we remove the south pole, stereographic projection will give us such a coordinate $z \in \mathbb{C}$. In these coordinates we write the holomorphic quadratic differential Ψ as ψdz^2. On the other hand, we can take the 'opposite' chart obtained by removing the north pole, to get another complex coordinate $w = 1/z$. This way, we have $\Psi = \psi \frac{dw^2}{w^4}$, and by considering what happens as $w \to 0$, we see that $\frac{\psi}{w^4}$ must remain bounded, and in particular, $\psi \to 0$ as $w \to 0$, i.e. as $z \to \infty$.

We deduce that ψ is a bounded holomorphic function on the plane – and is therefore constant by Liouville's theorem – and since it converges to zero, that constant must be zero. □

Combining Lemmas 1.5.3 and 1.5.4, we find:

Lemma 1.5.5. *Any harmonic map from S^2 to an arbitrary Riemannian manifold must be weakly conformal, and thus a branched minimal immersion.*

Thus to find branched minimal spheres, we only have to find harmonic spheres - i.e. harmonic maps from S^2. One way of finding harmonic maps is by performing a gradient flow on energy, and that is what we discuss in the next section.

Chapter 2

The harmonic map flow

2.1 Definition of the flow

The harmonic map flow was introduced by Eells-Sampson in 1964; their work could be considered the start of the field of geometric flows. The flow takes a map u between closed Riemannian manifolds (M,g) and (N,G) of any dimension, and deforms it to make its energy decrease as quickly as possible; more precisely it performs L^2 gradient flow. Although we are temporarily considering domains of arbitrary dimension, with fixed domain metric g, the first variation formula from Lemma 1.5.1 extends to this case, and we end up with the equation

$$\frac{\partial u}{\partial t} = \tau_g(u) \tag{2.1.1}$$

for the flow. For smooth solutions, the energy then decreases according to

$$\frac{dE}{dt} = -\int_M |\tau_g(u)|^2 d\mu_g.$$

2.2 Nonpositively curved targets: The theorem of Eells-Sampson

The original work of Eells-Sampson [8] proved that if we impose an extra condition on the target, then the flow behaves very well, and finds a harmonic map.

Theorem 2.2.1 ([8, 13]). *Suppose that (N,G) is a closed Riemannian manifold of nonpositive sectional curvature, and (M,g) is any closed Riemannian manifold, of any dimension. Then given a smooth map $u_0 : M \to N$, there exists a unique smooth solution $u : M \times [0,\infty) \to N$ to the harmonic map flow (2.1.1) with $u(0) = u_0$, and there exists a harmonic map $u_\infty : M \to N$ such that*

$$u(t) \to u_\infty \qquad \text{in } C^k \text{ for any } k \text{ as } t \to \infty.$$

Here $u(t)$ is shorthand for $u(\cdot,t)$.

Remark 2.2.2. In fact, the theorem above incorporates improved asymptotics due to Hartman [13]. The original theorem proved convergence $u(t_i) \to u_\infty$ for some sequence of times $t_i \to \infty$.

This theorem has been highly influential. The general method gives a way of transforming a very general object – here an arbitrary map – into a very special one – here a harmonic map. Similar flows now exist for many other geometric objects. For example, instead of deforming maps, Ricci flow deforms Riemannian metrics and curve shortening flow deforms curves within some ambient Riemannian manifold. We might see both of these flows in these lectures.

A key idea in the proof of Theorem 2.2.1 is to consider the evolution of the energy density $e(u) := \frac{1}{2}|du|^2$. Once we have observed that the harmonic map flow is a *strictly parabolic* partial differential equation, and will admit solutions over some, possibly short, time interval $[0, \varepsilon)$, we mainly have to worry about the flow developing a singularity in finite time. The theory of parabolic regularity theory (e.g. using theory in [15]) will tell us that a singularity can only stop us if the energy density becomes unbounded. However, the energy density itself satisfies a nice parabolic PDE, according to the parabolic Bochner formula. When we inspect that PDE, we see that $(\frac{\partial}{\partial t} - \Delta)e(u)$ can be bounded linearly in terms of $e(u)$ plus a term that could be of order $e(u)^2$. Now that could lead to a blow-up of $e(u)$ as surely as the ODE $\frac{da}{dt} = a^2$ will blow up in general. On the other hand, in the case that the sectional curvature of the target is nonpositive, we find that this 'bad' term is negative, and will work to *prevent* energy blow-up rather than the other way round.

In the early days of the harmonic map flow, there was (anecdotally) a naive hope that the flow would *never* develop singularities, even if the target was allowed to break the hypothesis of nonpositive curvature. However, there are good geometric reasons why this could not be the case, as we will see.

2.3 Two dimensional domains: The theory of Struwe

After this brief interlude where we allowed the domain to have arbitrary dimension, we now return to the case that (M, g) is a closed surface. (For this part of the theory, it is not important that the surface is oriented.) The theory of the harmonic map flow is particularly elegant and well-developed in this case.

In order to be able to handle nonsmooth maps most easily, it will be convenient to embed the target manifold (N, g) isometrically in some Euclidean space $\mathbb{R}^n$, as we know we can by the Nash embedding theorem. This way, the energy could be written

$$E(u) := \frac{1}{2}\int_{\mathcal{M}} |\nabla u|^2 d\mu_g,$$

i.e. in terms of the sum of the Dirichlet energies of the n components u^i. The harmonic map flow can then be seen to be governed by the equation

$$\frac{\partial u}{\partial t} = (\Delta u)^T, \tag{2.3.1}$$

where the superscript T means the projection onto $T_{u(x)}N \subset \mathbb{R}^n$, and one can compute that this leads to the system of n PDE

$$\frac{\partial u}{\partial t} = \Delta u + g^{ij} A(u)\left(\frac{\partial u}{\partial x^i}, \frac{\partial u}{\partial x^j}\right) \tag{2.3.2}$$

where $A(u)$ is the second fundamental form of the embedding $N \hookrightarrow \mathbb{R}^n$. Taking this viewpoint makes it easier to consider *weak* solutions $u \in W^{1,2}(M \times [0,\infty), N)$ (that is, $u \in W^{1,2}(M \times [0,\infty), \mathbb{R}^n)$ with $u(x) \in N \hookrightarrow \mathbb{R}^n$ for almost all $x \in M$).

Note that $W^{1,2}$ represents the Sobolev space of weakly differentiable maps with just enough regularity – first derivatives in L^2 – to be able to make sense of the energy E. On the other hand, the topology of $W^{1,2}$ is just weak enough so $W^{1,2}$ convergence does not imply C^0 convergence.

The following theorem relies on the fact that the domain is two dimensional. In it, we denote by $u(t)$ the restriction $u(\cdot, t)$.

Theorem 2.3.1. *(Struwe [29].) Given $u_0 \in W^{1,2}(M,N)$, for example a smooth map $u_0 : M \to N$, there exists a weak solution $u \in W^{1,2}_{loc}(M \times [0,\infty), N)$ to (2.3.2) which is smooth in $M \times [0,\infty)$ except possibly on a finite set of singular points $S_1 \subset M \times (0,\infty)$, and has the properties:*

1. *$u(0) = u_0$;*
2. *$E(u(t))$ is a (weakly) decreasing function of t on $[0,\infty)$;*
3. *If the flow is smooth for $t \in [0,T)$, then it is the unique smooth solution over this time interval with the given initial data;*
4. *There exist a sequence of times $t_i \to \infty$ and a smooth harmonic map $u_\infty : M \to N$ such that*
   - *(a) $u(t_i) \rightharpoonup u_\infty$ in $W^{1,2}(M,N)$, and*
   - *(b) $u(t_i) \to u_\infty$ in $W^{2,2}_{loc}(M \backslash S_2, N)$, and in particular in $C^0_{loc}(M \backslash S_2, N)$, where S_2 is a finite set of point in M.*

   as $i \to \infty$.

This theorem takes a general map u_0 and deforms it to a harmonic map u_∞, somewhat like in Theorem 2.2.1 of Eells-Sampson. Amongst the differences are that we allow the possibility of singularities on the finite sets S_1 and S_2. We cannot hope to do away with both of these sets because then the flow would homotop an arbitrary map to a harmonic map, and this is impossible in general:

Theorem 2.3.2 (Eells-Wood [9]). *There does not exist any harmonic map of degree 1 from a torus T^2 to a sphere S^2, whichever metrics we take on the domain and target.*

We can deduce from this that the flow organises itself in such a way as to jump homotopy class at one of the singular points in S_1 and/or S_2. We see how it goes about this in the next section.

2.4 Bubbling

At each of the finite-time singular points in S_1 and infinite time singular points in S_2 we get *bubbling*. (We can assume, without loss of generality, that any points in S_1 or S_2 that can be removed without breaking Theorem 2.3.1 have been removed. That is, all points in S_1 and S_2 are genuine singular points.)

The first thing to remark is that at each of these singular points we have energy concentrating in the following senses.

Lemma 2.4.1. *If $(x,T)\in S_1\subset M\times(0,\infty)$ is a singular point, then energy concentrates in the sense that*

$$\lim_{\nu\downarrow 0}\limsup_{t\uparrow T} E(u(t),B_\nu(x))\neq 0.$$

Similarly, if $x\in S_2\subset M$ is a singular point, then

$$\lim_{\nu\downarrow 0}\limsup_{t\to\infty} E(u(t),B_\nu(x))\neq 0.$$

To clarify, we are using here the notation $E(u,\Omega):=\frac{1}{2}\int_\Omega|\nabla u|^2$. Also, $B_\nu(x)$ is the geodesic ball centred at x, of radius ν.

To see what is going on at a singularity, where energy is concentrating, we have to *rescale*. When we do, we find a 'bubble'.

Theorem 2.4.2. *Suppose $u\in W^{1,2}_{loc}(M\times[0,\infty),N)$ is a flow from Theorem 2.3.1. Then without loss of generality, the sequence of times t_i can be chosen so that the following is also true. For each singular point $x_0\in S_2\subset M$, when we view u as a map in local isothermal coordinates on the domain with x_0 corresponding to the origin of the coordinates, there exist sequences $a_i\to 0\in\mathbb{R}^2$ and $\lambda_i\downarrow 0$ as $i\to\infty$, and a nonconstant harmonic map $\omega:S^2\to N$ (which we view as a map $\mathbb{R}^2\cup\{\infty\}\to N$) such that*

$$u(a_i+\lambda_i y,t_i)\to\omega(y),$$

as functions of y, in $W^{2,2}_{loc}(\mathbb{R}^2,N)$ as $i\to\infty$.

It is the map ω arising as a blow-up limit near the singularity that we call a bubble. We can also find these bubbles at finite time singularities:

Theorem 2.4.3. *(Slight variant of Struwe [29].) Suppose $u\in W^{1,2}_{loc}(M\times[0,\infty),N)$ is a flow from Theorem 2.3.1, and $T\in(0,\infty)$ is a singular time. Then there exist times $t_i\uparrow T$ with the property that*

$$\|\tau(u(t_i))\|^2_{L^2(M)}(T-t_i)\to 0 \tag{2.4.1}$$

as $i\to\infty$. Moreover, for every singular point $(x_0,T)\in S_1\subset M\times(0,\infty)$ at time T, when we view u as a map in local isothermal coordinates on the domain with x_0 corresponding to the origin, there exist sequences $a_i\to 0\in\mathbb{R}^2$ and $\lambda_i\downarrow 0$ with $\lambda_i(T-t_i)^{-\frac{1}{2}}\to 0$ as $i\to\infty$, and a nonconstant harmonic map $\omega:S^2\to N$ (which we view as a map $\mathbb{R}^2\cup\{\infty\}\to N$) such that

$$u(a_i+\lambda_i y,t_i)\to\omega(y),$$

as functions of y, in $W^{2,2}_{loc}(\mathbb{R}^2,N)$ as $i\to\infty$.

In fact, in principle, there may be many bubbles developing near each singular point (cf. [31]) i.e. we might be able to take different sequences a_i and λ_i which lead to convergence to a different bubble. (Different meaning not just the old bubble distorted a bit.) It follows from Theorems 2.3.1 and 2.4.2 that

$$\lim_{t\to\infty}E(u(t))-E(u_\infty)\geq E(\omega), \tag{2.4.2}$$

where ω is the bubble extracted at infinite time, and similarly from Theorems 2.3.1 and 2.4.3 that

$$\lim_{t \uparrow T} E(u(t)) - E(u(T)) \geq E(\omega), \tag{2.4.3}$$

for ω the bubble extracted at finite time. However, if we capture *all* the bubbles at each singular point and replace the right-hand side with the sum of all their energies, then it is a result of Ding-Tian [6] that we would have equality in (2.4.2) and (2.4.3). For more details, and the most refined result (based on the work of Struwe [29], Ding-Tian [6], Qing-Tian [18] and Lin-Wang [16]) see [32].

Note that the behaviour of the flow at the singularity can still be quite bad. For example, the flow map $u(T)$ at time T might be discontinuous, despite being in $W^{1,2}(M,N)$, and it is sometimes possible to change the sequence of times $t_i \uparrow T$ and end up with a different set of bubbles, even with different images. See [32] for details of these constructions.

Both types of singularity – finite and infinite time – can actually occur in practice.

An example of Chang-Ding [3] says that we can find a degree zero map from S^2 to S^2 that flows smoothly for all time, and develops two bubbles at infinite time (one degree 1 and one degree -1) with the map u_∞ from Theorem 2.3.1 being constant.

[draw a picture!!!!!]

Meanwhile, an example of Chang-Ding-Ye [4] says that we can find a degree one map from S^2 to S^2 that develops a singularity in finite time, when two bubbles are created (one degree 1 and one degree -1) leaving the flow to continue smoothly forever afterwards and converge to the identity map.

[draw a picture!!!!!]

We can now return to Theorem 2.3.2 to consider what the flow does with a degree 1 map from T^2 to S^2. What can happen (and will definitely happen for maps with energy close to their minimum possible value in their homotopy class) is that the flow generates one bubble as a singularity, leaving a map homotopic to a constant map. What is unclear is whether that happens in finite time, leaving a degree zero map that then flows on and converges to a constant map, or whether that happens at infinite time, leaving a constant map u_∞. This is unknown to this day.

2.5 No bubbling for nonpositively curved targets

A significant consequence of the singularity analysis of the previous section is that if we are dealing with a target manifold (N, G) for which there do not exist any nonconstant harmonic maps $S^2 \mapsto (N, G)$, then we can be sure that there there are never any singularities in the flow, even at infinite time, i.e. both S_1 and S_2 are empty. One such situation is:

Lemma 2.5.1. *Every harmonic map $S^2 \to (N, G)$ into a Riemannian manifold of nonpositive sectional curvature is a constant map.*

Proof. By lifting to the universal cover, we may assume that (N, G) is simply connected. The squared distance function $d^2(x_0, \cdot)$ to any fixed point $x_0 \in M$ is a strictly convex function because of the nonpositive curvature [1], and any harmonic map from any closed Riemannian manifold into any Riemannian manifold supporting a convex function is necessarily constant because the composition of the map and the convex function must be subharmonic [10]. □

Thus we recover the original Eells-Sampson theorem (see Theorem 2.2.1 and Remark 2.2.2) in the 2D case as a toy instance of this argument.

2.6 Finding branched minimal immersions with the harmonic map flow

We saw at the end of the previous chapter/lecture that to find branched minimal spheres, we only have to find nonconstant harmonic maps from S^2. We will now give an example of how the harmonic map flow can achieve this.

Theorem 2.6.1. *Suppose the (closed) target N satisfies $\pi_2(N) \neq \{0\}$. Then for any target metric G, there exists a branched minimal immersion $S^2 \mapsto (N, G)$.*

Proof. Choose any homotopically nontrivial smooth map $u_0 : S^2 \to N$. By Struwe's Theorem 2.3.1, we can run the harmonic map flow starting at u_0. Either the flow exists for all time and converges in C^0 to our desired branched minimal immersion u_∞, or there exists a singularity at finite and/or infinite time. But if there is a singularity, then we must have a bubble ω according to Lemmata 2.4.2 and 2.4.3, which would itself give the desired branched minimal immersion. □

On the other hand, if $\pi_2(N) = \{0\}$...

Theorem 2.6.2. *If $\pi_2(N) = \{0\}$, then we can* homotop *an arbitrary map u_0 (from an arbitrary closed surface) into a harmonic map.*

Proof. Again, we apply Struwe's Theorem 2.3.1. The claim is simply that the asymptotic harmonic map u_∞ is homotopic to u_0. For example, it is true that any flow with isolated point singularities does not change its homotopy class as we pass through the singular time.

EXPLAIN and DRAW PICTURE! □

This latter theorem does not help us directly find branched minimal immersions: If the domain is S^2, then we will always be able to homotop u_0 to a constant map by hypothesis, and this is harmonic but of no help. On the other hand, if the domain is not S^2 then the harmonic map we found in Theorem 2.6.2 need not be weakly conformal, and therefore not necessarily a branched minimal immersion. The next chapter will attempt to rectify that.

Chapter 3

How might we find higher genus branched minimal immersions?

3.1 Trying to flow map and domain metric

We saw in the last section how the harmonic map flow can find harmonic maps from surfaces, and how when the domain is S^2 we can thus find branched minimal immersions in many cases.

Of course, for genus 1 or higher surfaces, we can still find the harmonic maps, but they will not represent branched minimal immersions in general, since they need not be weakly conformal.

Instead, we return to Section 1.5 where we discussed the first variation of energy. Lemma 1.5.2 tells us that if we have a nonconstant critical point of E with respect to variations of both u and g, then we have a branched minimal immersion as desired. Our proposal is, therefore, to do a gradient flow with respect to both the map u and the metric g.

By the first variation formula of Lemma 1.5.1, it would thus appear that the gradient flow should be

$$\frac{\partial u}{\partial t} = \tau_g(u); \qquad \frac{\partial g}{\partial t} = \frac{1}{4} Re(\Phi(u, g)), \tag{3.1.1}$$

Unfortunately, these PDE have no obvious smoothing effect for the metric. This means that the flow cannot be expected to reduce to 'complexity' of the initial data, and will be very hard to analyse.

It will turn out that these difficulties can be overcome by exploiting some symmetries of the problem. The first one we have already remarked in the first section:

Lemma 3.1.1. *The energy functional for maps from surfaces is invariant under conformal deformation of the domain metric:*

$$E(u, g) = E(u, fg) \qquad \text{for all smooth functions } f : M \to (0, \infty).$$

As we saw, the proof is simply because $|du|^2_{fg} = f^{-1}|du|^2_g$, while $d\mu_{fg} = f d\mu_g$.

Thus both the harmonic map energy, and the notion of harmonic will be independent of the conformal factor of the domain metric. Of course, this means that the energy, and the notion of harmonic only depends on the conformal structure of the domain (although the flow will also depend on the conformal factor). However, instead of taking this viewpoint, we persist with metrics, but choose them to be very special, as we describe in the next section.

3.2 Geometrising surfaces

Lemma 3.2.1. *Given a closed Riemannian surface* (M, g)*, there exists a smooth positive function* $f : M \to (0, \infty)$ *such that* (M, fg) *has constant Gauss curvature* 1*,* 0 *or* -1*.*

If the surface is topologically S^2*, then we can make the Gauss curvature identically* 1*. If it is* T^2*, then we can make the curvature* 0*, and the total area* 1*. If it is a higher genus surface then we can make the Gauss curvature* -1*.*

For each of the oriented surfaces in the lemma, we denote by $\mathcal{M}_c$ the space of all metrics with constant curvature -1, 0 or 1, and with unit area in the middle case. (This middle case is the only case in which the area is not fixed by the Gauss-Bonnet theorem.)

Note that for genus $\gamma \geq 2$, the conformal hyperbolic metric is unique. But it's not for $\gamma = 0$.

Proof sketch. To prove this, first scale the metric homothetically (i.e. replace g with cg for some $c > 0$) so that the *average* Gauss curvature K_0 is either 1, 0 or -1. In the case that $K_0 = 0$, we scale so that the area is 1. We can then simply run the normalised Ricci flow

$$\frac{\partial g}{\partial t} = (K_0 - K)\, g.$$

It turns out that this flow will smoothly and conformally deform an arbitrary surface to one of constant curvature K_0 as $t \to \infty$, without changing the area, or indeed (by the Gauss-Bonnet theorem) the average curvature. The Gauss-Bonnet theorem tells us which surfaces will have which curvature: 1, 0 or -1. See [34] for more on this topic. □

The idea then is to try to do a gradient flow for E with respect to the map u, and the metric g constrained to be one of the constant curvature metrics from Lemma 3.2.1. The space of constant curvature metrics is infinite dimensional, in some sense, but we will see that it behaves like a finite dimensional space, which will make the equation for g under gradient flow on E look more like an ODE.

3.3 The space of constant curvature metrics $\mathcal{M}_c$

We would like to understand the 'tangent space' $T\mathcal{M}_c$ to the space $\mathcal{M}_c$ of constant curvature metrics on a closed orientable surface as in Lemma 3.2.1. This is because if we want g to flow within $\mathcal{M}_c$ we had better make sure that $\frac{\partial g}{\partial t}$ lies within $T\mathcal{M}_c$.

Imagine a one parameter family $g(s)$ of such constant curvature metrics. The fact that we are insisting the curvature remains constant imposes a constraint on $h := \frac{\partial g}{\partial s}$ and we would like to work out what that is.

First of all we note a decomposition one can make for an arbitrary change of metric, on domain manifolds of any dimension.

Lemma 3.3.1. *Suppose (M, g) is a closed Riemannian manifold/surface, and $h \in \Gamma(\mathrm{Sym}^2 T^*M)$. Then there exists a vector field X on M and a section h_0 of $\mathrm{Sym}^2 T^*M$ that is divergence free, i.e. $\delta h_0 \equiv 0$, giving the L^2-orthogonal decomposition*

$$h = h_0 + \mathcal{L}_X g.$$

(Note that divergence is defined by $\delta := -\mathrm{tr}\nabla$ for us.)

This is analogous to the decomposition of a vector field H into the sum of a divergence-free vector field H_0 and the gradient of a function f. To prove that we would solve the PDE

$$\Delta f = \mathrm{div} H$$

and set $H_0 = H - \nabla f$. The proof below, taken from [35], is essentially the same.

Proof. Recall that the formal adjoint of $\delta : \Gamma(\mathrm{Sym}^2 T^*M) \to \Gamma(T^*M)$ is the operator $\delta^* : \Gamma(T^*M) \to \Gamma(\mathrm{Sym}^2 T^*M)$ given by $\omega \mapsto \frac{1}{2}\mathcal{L}_{\omega^\#} g$. We claim that we can find a 1-form ω solving the elliptic PDE

$$\delta\delta^*\omega = \delta h. \tag{3.3.1}$$

Here, $\delta\delta^*$ is a self-adjoint elliptic operator, a type of Laplacian. Therefore by the Fredholm alternative we will have a solution ω of (3.3.1) precisely when the right-hand side δh is orthogonal to the kernel of $\delta\delta^*$.

But if $\sigma \in \ker(\delta\delta^*)$ then

$$\|\delta^*\sigma\|^2 = \langle \delta\delta^*\sigma, \sigma\rangle = 0,$$

so $\delta^*\sigma = 0$ and thus

$$\langle \delta h, \sigma\rangle = \langle h, \delta^*\sigma\rangle = 0,$$

i.e. δh is automatically orthogonal to the kernel of $\delta\delta^*$, and we can always solve (3.3.1).

It remains to set $X = \omega^\#$ and $h_0 = h - \mathcal{L}_X g = h - \delta^*\omega$. □

Now let's restrict M to be two dimensional, g to be of constant curvature, and h to be a direction we could move g that preserves this constant curvature to first order.

Lemma 3.3.2. *Suppose $g(s)$ is a one parameter family of metrics in $\mathcal{M}_c$ on the surface M. Set $h = \frac{\partial g}{\partial s} \in \Gamma(\mathrm{Sym}^2 T^*M)$ at $s = 0$. Then there exists a vector field X on M and a section h_0 of $\mathrm{Sym}^2 T^*M$ that is divergence free, i.e. $\delta h_0 \equiv 0$, and also trace free, i.e. $\mathrm{tr} h_0 = 0$, such that*

$$h = h_0 + \mathcal{L}_X g.$$

Tensors such as h_0 are often called *transverse traceless*. The condition $\mathrm{tr}h_0 \equiv 0$ is saying that when we move a metric in the direction h_0, the volume measure does not change, because in general

$$\frac{\partial}{\partial t}d\mu = \frac{1}{2}(\mathrm{tr}h_0)d\mu. \tag{3.3.2}$$

See for example [33, Proposition 2.3.12]

Proof. We already know from the previous lemma that $h = h_0 + \mathcal{L}_X g$ for some divergence free h_0 and vector field X. We need merely to work out what extra constraint we have if we ask that the derivative of the curvature as we move g in the direction h is zero, that is

$$DK_g[h] \equiv 0,$$

where K is the Gauss curvature. It is clear that

$$DK_g[\mathcal{L}_X g] = X(K) \equiv 0,$$

since K is constant. Therefore we just require that $DK_g[h_0] \equiv 0$. The Gauss curvature is half the scalar curvature R, and there is a general formula for the derivative of the scalar curvature as we deform the metric (see e.g. [33, Proposition 2.3.9]) in any dimension, given by

$$DR_g[h_0] = -\Delta_g(\mathrm{tr}h_0) + \delta\delta h_0 - \langle \mathrm{Ric}, h_0 \rangle,$$

which implies here, because $R = 2K$, $DK_g[h_0] = 0$, $\delta h_0 \equiv 0$ and $\mathrm{Ric} = Kg$, that

$$0 = 2DK_g[h_0] = -\Delta_g(\mathrm{tr}h_0) - K\mathrm{tr}h_0.$$

By multiplying this equation by $\mathrm{tr}h_0$ and integrating, we find that

$$0 = \int \left(|\nabla(\mathrm{tr}h_0)|^2 - K(\mathrm{tr}h_0)^2\right).$$

Now if $K \equiv -1$, then this immediately implies that $\mathrm{tr}h_0 \equiv 0$ as desired. If $K \equiv 0$, then it implies that $\mathrm{tr}h_0$ is constant, and since the total area is assumed to be fixed, we find by (3.3.2) that $\mathrm{tr}h_0 \equiv 0$ again. On the other hand, if $K \equiv 1$, then we are on the round sphere where we know the first two eigenvalues of the Laplacian to be 0 and 2, whereas we have proved that $-\Delta(\mathrm{tr}h_0) = \mathrm{tr}h_0$, from which we can deduce that $\mathrm{tr}h_0 \equiv 0$, otherwise 1 would also be an eigenvalue. □

The *transverse traceless* condition from the previous lemma can be described in terms of holomorphic quadratic differentials:

Lemma 3.3.3. *On a closed orientable surface (M, g), the space of tensors $h \in \Gamma(\mathrm{Sym}^2 T^*M)$ that satisfy $\delta h = 0$ and $\mathrm{tr}h = 0$ is precisely the space*

$$\{Re(\Psi) \ : \ \Psi \textit{ is a holomorphic quadratic differential}\},$$

which in turn can be identified with

$$\mathcal{H} := \{\Psi \ : \ \Psi \textit{ is a holomorphic quadratic differential}\},$$

via the bijection $\Psi \mapsto Re(\Psi)$.

Note as before that this final identification is a bijection: If the real part of Ψ is zero, then Ψ is zero. We restrict the proof to the first assertion of the lemma.

Proof. If h is any tensor that is divergence free and traceless, decompose it as in (1.4.1). Now, $\mathrm{tr} dz^2 = 0$, so the traceless condition implies that $a \equiv 0$, i.e. h is the real part of a quadratic differential. Let's write

$$h = Re(\psi dz^2).$$

It remains to exploit the fact that h is divergence free. But by (1.3.4), we have

$$0 = \delta h = \delta Re(\psi dz^2) = Re(\delta(\psi dz^2)) = Re(-\frac{2\psi_{\bar{z}}}{\rho^2} dz) = -\frac{\psi_{\bar{z}}}{\rho^2} dz - \frac{\bar{\psi}_z}{\rho^2} d\bar{z}.$$

We deduce that $\psi_{\bar{z}} = 0$, i.e. ψdz^2 is holomorphic as desired.

Clearly the argument works the other way too: given the real part of a holomorphic quadratic differential, it is trace free, and divergence free. □

3.4 The space of holomorphic quadratic differentials - preliminaries

Now we have described the tangent space to $\mathcal{M}_c$ in terms of holomorphic quadratic differentials as

$$T\mathcal{M}_c = Re(\mathcal{H}) \oplus \{\mathcal{L}_X g\},$$

we are going to need some information about the entire space $\mathcal{H}$ of such differentials, and we start by working out its dimension, which crucially is finite.

If the domain is S^2, then we already saw in Lemma 1.5.4 that there are no holomorphic quadratic differentials other than the zero differential, i.e. $\dim_{\mathbb{C}} \mathcal{H} = 0$.

In the case that the domain is T^2, we can lift any quadratic differential to its universal cover, which is $\mathbb{C}$, and write it globally as ψdz^2 where ψ is holomorphic and periodic. Liouville's theorem again tells us that ψ must be a constant (complex number). On the other hand, we obtain a holomorphic quadratic differential for any choice of this complex number. Thus the space of holomorphic differentials is one (complex) dimensional, i.e. $\dim_{\mathbb{C}} \mathcal{H} = 1$.

For domains (M, g) of genus $\gamma \geq 2$, an argument using the Riemann-Roch theorem shows that the complex dimension of $\mathcal{H}$ is $3(\gamma - 1)$. (See, for example, [12, Corollary 5.4.2].)

Later we will see a geometric explanation for this dimension: It must coincide with the dimension of the space of hyperbolic metrics (modulo isometries) essentially by Lemmata 3.3.2 and 3.3.3, and we will sketch a geometric picture of how any hyperbolic metric can be deformed, in terms of 'pairs of pants decompositions', in precisely $6(\gamma - 1)$ directions.

3.5 Definition of Teichmüller harmonic map flow

We now return to our first attempt at a flow, (3.1.1).

We will revise the equations by asking that g lies in the space $\mathcal{M}_c$ of constant curvature metrics, and remains there. Clearly then we must force $\frac{\partial g}{\partial t}$ to lie in the 'tangent space' of $\mathcal{M}_c$, and that consists of tensors that are the real parts of holomorphic quadratic differentials.

The simplest way of achieving that is simply to project the Hopf differential onto the space of holomorphic quadratic differentials before taking its real part. We will see later how this is motivated by taking a gradient flow on E in a space of maps and metrics somewhat related to Teichmüller space.

This leads us to the so-called Teichmüller harmonic map flow, that we introduced with Melanie Rupflin, which is defined (up to an arbitrary coupling constant $\eta > 0$ that will be motivated later) by

$$\frac{\partial u}{\partial t} = \tau_g(u); \qquad \frac{\partial g}{\partial t} = \frac{\eta^2}{4} Re(P_g(\Phi(u,g))), \tag{3.5.1}$$

where P_g is the projection from the entire space of quadratic differentials to the holomorphic ones.

By plugging these equations into the first variation formula from Lemma 1.5.1, we see that for a smooth solution, the energy decreases according to:

$$\frac{dE}{dt} = -\int_M \left[|\tau_g(u)|^2 + \left(\frac{\eta}{4}\right)^2 |Re(P_g(\Phi(u,g)))|^2 \right]. \tag{3.5.2}$$

Intuitively, a key advantage of the flow equations (3.5.1) over the equations 3.1.1 is that after projecting the Hopf differential onto $\mathcal{H}$, using P_g, we have g moving in a finite dimensional space of directions, as we saw in Section 3.4, and this almost amounts to g satisfying an ODE. This will allow us to get existence of solutions, although first we inspect what the Teichmüller harmonic map flow does for domains of different genera.

3.5.1 Domain is a sphere

In this case, we saw that any holomorphic quadratic differential is identically zero. (Therefore when we project the Hopf differential on $\mathcal{H}$, we get zero.) Thus $\frac{\partial g}{\partial t} = 0$ and the Teichmüller harmonic map flow is just the harmonic map flow (which already does well at finding branched minimal immersions, as we saw before).

3.5.2 Domain is a torus

In this case, we saw that the holomorphic quadratic differentials are parallel. There is a one complex dimensional space of directions we can move. In fact, these 'horizontal' directions are integrable - i.e moving g so that $\frac{\partial g}{\partial t}$ is always the real part of a holomorphic quadratic differential will keep g within a two real dimensional subspace of the space of unit area metrics on the torus. This makes the analysis a lot easier.

With a bit of work, one can show the Teichmüller harmonic map flow is equivalent to a flow that was studied by Ding-Li-Liu [5] - see [20].

3.5.3 Domain has genus at least 2

This is the case that will concern us most from now. For each $g \in \mathcal{M}_c$ – i.e. each hyperbolic metric g in this case – we have a space of directions $Re(\mathcal{H})$ making up a distribution on $\mathcal{M}_c$. But as opposed to the torus case, *this distribution is not integrable.* In other words, given $g \in \mathcal{M}_c$, there does not exist a $6(\gamma - 1)$ real dimensional submanifold of $\mathcal{M}_c$ through g on which we remain as we move around $\mathcal{M}_c$ always moving in horizontal directions. Thus the equation for g is more complicated than a finite dimensional ODE, and this makes the PDE theory tricky.

3.6 Comments on Teichmüller space and the Weil-Petersson metric

Let us assume for the moment that given an initial map u_0 and initial metric $g_0 \in \mathcal{M}_c$, there exists a short-time solution to the Teichmüller harmonic map flow (3.5.1). While the flow exists, the metric g traces out a path in $\mathcal{M}_c$; for example for genus at least two, g moves within the space of hyperbolic metrics . It is natural to work within this space modulo diffeomorphisms: two metrics should be deemed to be the same if one is the pull-back of the other by a diffeomorphism. This space is called moduli space. In fact, a much more regular space is obtained if one considers the quotient

$$\mathcal{T}(M) := \mathcal{M}_c/\mathcal{D}_0,$$

where

$$\mathcal{D}_0 := \{\text{diffeomorphisms } M \mapsto M \text{ that are homotopic to the identity map}\},$$

and this space could be taken as a definition of Teichmüller space.

Teichmüller space comes with a natural metric, the so-called Weil-Petersson metric, which is basically the L^2 metric. Slightly more precisely, by Lemma 3.3.2, a tangent vector in $\mathcal{T}(M)$ can be viewed as a transverse-traceless tensor h_0, or a holomorphic quadratic differential (by also Lemma 3.3.3). Note that we can ignore the term $\mathcal{L}_X g$ since it is just keeping the metric the same, but modifying it by a diffeomorphism. We can take the L^2 norm of h_0 (modulo an ambiguous scaling constant) to give the norm of the tangent vector in $\mathcal{T}(M)$.

The $g(t)$ part of the Teichmüller harmonic map flow can be projected down onto Teichmüller space, and moves a finite distance in a finite time.

When working on the torus, this is perfect, because Teichmüller space is complete in that case. However, for surfaces of larger genus, Teichmüller space is *incomplete* and the flow can in theory move g all the way to the boundary in finite time. We will see in the next section what that looks like in some detail, but for now we satisfy ourselves with the assertion that as a hyperbolic metric is deformed to the boundary of Teichmüller space, the injectivity radius must decrease to zero.

3.7 Medium-time existence of solutions

The existence theory naturally extends the existence theory of Struwe for the harmonic map flow that we saw in Theorem 2.3.1. Indeed, in the case that $M = S^2$, the flows coincide.

However, for higher genus M, the fact that the metric is evolving creates some new technical problems. In the case that $M = T^2$, the existence theory was given by Ding-Li-Liu [5]. Roughly, there exists a global weak solution with bubbling.

In the case $genus(M) \geq 2$, the initial existence theory can only go up to any such point that the domain degenerates in the sense described in the last section.

We use the terminology $\mathcal{M}_{-1}$ for the space of hyperbolic metrics on M – i.e. the hyperbolic case of $\mathcal{M}_c$.

Theorem 3.7.1 ([21])**.** *For any initial data $(u_0, g_0) \in C^\infty(M, N) \times \mathcal{M}_{-1}$ there exists a (weak) solution (u, g) of (3.5.1) defined on a maximal interval $[0, T)$, $T \leq \infty$, satisfying the following properties*

(i) *The solution (u, g) is smooth away from at most finitely many singular times $T_i \in (0, T)$. Furthermore as $t \to T_i$ the metrics $g(t)$ converge smoothly to an element $g(T_i) \in \mathcal{M}_{-1}$ and the maps $u(t)$ converge to $u(T_i)$ weakly in H^1 and smoothly away from a finite set of points $S(T_i)$ in M.*

(ii) *The energy $t \mapsto E(u(t), g(t))$ is non-increasing.*

(iii) *The solution exists until the metrics degenerate in moduli space; i.e. if the maximal existence time $T < \infty$ then the infimum of the injectivity radius of $(M, g(t))$ over $t \in [0, T)$ is zero.*

Furthermore, this solution is uniquely determined by its initial data in the class of all weak solutions with non-increasing energy.

3.8 Teichmüller harmonic map flow as a gradient flow

We have justified the flow equations (3.5.1) as a modification of the naive first choice of equations (3.1.1). We would now like to show that the flow equations are the natural gradient flow equations for E, once we have determined the correct space in which to work.

The trick is to exploit another symmetry of the energy E, given by the action of diffeomorphisms of M. More precisely, we define

$$\mathcal{D}_0 := \{\text{diffeomorphisms } M \mapsto M \text{ that are homotopic to the identity map}\},$$

SORRY, HAVEN'T TYPED THIS YET!!!! Don't worry, you don't need it for the next sections.