Asymptotics for general connexions at infinity

Carlos Simpson
CNRS, Laboratoire J. A. Dieudonné, Université de Nice-Sophia Antipolis
September 24th, 2003: preliminary notes for my talk in Toulouse on Sept. 26th

Introduction

This is a preliminary set of notes on the asymptotic behavior of the monodromy of connexions near a general point at ∞ in the space M_{DR} of connexions on a compact Riemann surface X. We will consider a path of connexions of the form $(E, \nabla + t\theta)$ which approaches the boundary divisor transversally at the point on the boundary of M_{DR} corresponding to a general Higgs bundle (E, θ). This is very similar to the situation considered in [1], and we import the vast majority of our techniques directly from there. However, our present situation is slightly more general with regard to the form of the family of connexions. We reduce to a case similar to that which was treated in [1], of a family of connexions of the form $d+B+tA$, but the matrix B may have poles. This difficulty of poles in the matrix B is the new phenomenon which is treated here. However, we are not able to get results as good as the precise asymptotic expansions of [1]. Our result of Theorem 2 (cf p. 14) is that if $m(t)$ denotes the family of monodromy or transport matrices for a given path, then the Laplace transform $f(\zeta)$ of m has an analytic continuation with locally finite singularities over the complex plane. The singularities are what determine the asymptotic behavior of $m(t)$. What we don't know is the behavior of $f(\zeta)$ near the singularities; the main question left open is whether f has polynomial growth at the singularities, and if so, to what extent the generalized Laurent series can be calculated from the individual terms in our integral expression for f.

Even though he doesn't appear in the references of [1], the ideas of J.-P. Ramis indirectly had a profound influence on that work (and hence on the present note). This can be traced to at least two inputs as follows:
(1) I had previously followed G. Laumon's course about ℓ-adic Fourier transform, which was partly inspired by the corresponding notions in complex function theory, a subject in which Ramis (and others in the subject of "resurgence") had a great influence; and
(2) at the time of writing [1] I was following N. Katz's course about expo-

nential sums, where again much of the inspiration came from Ramis's work (which Katz mentionned very often) on irregular singularities.

The compactified moduli space of connexions

Let X be a smooth projective curve over the complex numbers $\mathbf{C}$. Fix r. Recall that we have a moduli space M_{DR} of rank r vector bundles with integrable connexion on X. This has a compactification $M_{DR} \subset \overline{M}_{DR}$ constructed as follows. There is a moduli space $M_{Hod} \rightarrow \mathbf{A}^1$ for vector bundles with λ-connexion, $\lambda \in \mathbf{A}^1$. The fiber over $\lambda = 0$ is the moduli space M_{Dol} for semistable Higgs bundles of degree zero. This has a subvariety M_{Dol}^{nil} parametrizing the Higgs bundles (E, θ) such that θ is nilpotent as an Ω^1_X-valued endomorphism of E. Let M^*_{Dol} denote the complement of M^{nil}_{Dol} in M_{Dol} and let M^*_{Hod} denote the complement of M^{nil}_{Dol} in M_{Hod}. Then the algebraic group $\mathbf{G}_m$ acts on M_{Hod} preserving all of the above subvarieties, and

$$\overline{M}_{DR} := M^*_{Hod}/\mathbf{G}_m.$$

The complement of M_{Dol} in M_{Hod} (which is also the complement of M^*_{Dol} in M^*_{Hod}) is isomorphic to $M_{DR} \times \mathbf{G}_m$ and this gives the embedding $M_{DR} \hookrightarrow \overline{M}_{DR}$. The complementary divisor is given by

$$P_{DR} = M^*_{Dol}/\mathbf{G}_m.$$

In conclusion, this means that the points at ∞ in $\overline{M}_{DR}$ correspond to equivalence classes of semistable, non-nilpotent Higgs bundles (E, θ) under the equivalence relation

$$(E, \theta) \cong (E, u\theta)$$

for any $u \in \mathbf{G}_m$.

Recall that the moduli space M_{Dol} is an irreducible algebraic variety, so P_{DR} is also irreducible. The general point therefore corresponds to a "general" Higgs bundle (E, θ). For a general point, the spectral curve of θ is an irreducible curve with ramified map to X, such that the ramification points are all of the simplest type.

Curves going to infinity

The moduli spaces considered above are coarse only; however, in an etale neighborhood of the generic point they are fine moduli spaces, and also smooth. At a general point of P_{DR}, both $\overline{M}_{DR}$ and the divisor P_{DR} are smooth. Thus we can look for a curve cutting P_{DR} transversally at a general point. Such a curve may be obtained by taking the projection of a curve in M_{Hod} cutting M_{Dol} at a general point. In turn, this amounts to giving a family (E_c, ∇_c) where ∇_c is a $\lambda(c)$-connexion, parametrized by $c \in C$ for some curve C. Also in an etale neighborhood of the point $\lambda = 0$, the function $\lambda(c)$ should be etale. Note also that (E_0, ∇_0) should be a general semistable Higgs bundle of degree zero.

The easiest way to obtain such a curve is as follows: let (E, θ) be a general Higgs bundle, stable of degree zero. The bundle E is stable as a vector bundle (since stability is an open condition and it certainly holds on the subset of Higgs bundles with $\theta = 0$, so it holds at general points). In particular E supports a connexion ∇ and we can set

$$\nabla_\lambda := \lambda \nabla + \theta$$

for $\lambda \in \mathbf{A}^1$. Here the parameter is λ itself. The subset $\mathbf{G}_m \subset \mathbf{A}^1$ corresponds to points which are mapped into M_{DR}, and indeed the vector bundle with connexion corresponding to the above λ-connexion is

$$(E, \nabla + t\theta) \quad , \; t = \lambda^{-1}.$$

The map actually extends to a map from $\mathbf{A}^1$ into M_{DR} for the other coordinate chart $\mathbf{A}^1$ providing a neighborhood at ∞ in $\mathbf{P}^1$. In conclusion, the family of connexions $\{(E, \nabla + t\theta)\}$ corresponds to a morphism

$$\mathbf{P}^1 \to \overline{M}_{DR}$$

sending $t \in \mathbf{A}^1$ into M_{DR}, sending the point $t = \infty$ to a general point in the divisor P_{DR}, and the curve is transverse to the divisor at that point.

We will investigate the asymptotic behavior of the monodromy representations of the connexions $(E, \nabla + t\theta)$ as $t \to \infty$. Recall that the *Betti moduli space* M_B is the moduli space for representations of $\pi_1(X)$ up to conjugation, and we have an analytic isomorphism $M_{DR}^{\rm an} \cong M_B^{\rm an}$ sending a connexion to its monodromy representation. We will look at the asymptotics of the resulting analytic curve $\mathbf{A}^1 \to M_B$.

In order to set things up it will be useful to fix a basepoint $p \in X$ and a trivialization $\tau : E_p \cong \mathbf{C}^r$. Then for any $\gamma \in \pi_1(X, x)$ we obtain the monodromy matrix

$$\rho(E, \nabla + t\theta, \tau, \gamma) \in GL(r, \mathbf{C}).$$

Of course the monodromy matrices don't directly give functions on the moduli space M_B of representations, because they depend on the choice of trivialization τ. However, one has the Procesi coordinates which are certain polynomials in the monodromy matrices (for several γ at once) which are invariant under change of trivialization and give an embedding of the Betti moduli space M_B into an affine space. We will concentrate on looking at the asymptotic behavior of the monodromy matrices, but the resulting asymptotic information will also hold for any polynomials, and in particular for the Procesi coordinates. This will give asymptotic information about the image curve in M_B.

In any case, the functions we shall consider, be they the matrix coefficients of the monodromy ρ or some other polynomials in these, will be entire functions $m(t)$ on the complex line $t \in \mathbf{C}$. We will be looking to characterize their asymptotic properties.

The method we will use is the same as the method already used in [1] to treat exactly this question, for a more special class of curves going to infinity in M_{DR}. In that book was treated the case of families of connexions $(E, \nabla + t\theta)$ where

$$E = \mathcal{O}_X^r, \quad \nabla = d + B, \quad \theta = A$$

with A and B being $r \times r$ matrices of one-forms on X such that A is diagonal and B contains only zeros on the diagonal. In [1], a fairly precise description of the asymptotic behavior of the monodromy was obtained. It was also indicated how one should be able to reduce to this case in general; we shall explain that below. The only problem is that in the course of this reduction, one obtains the special situation but with B being a matrix of one-forms which has some poles on X. In this case the exact method used in [1] breaks down.

The purpose of the present paper is to try to remedy this situation as far as possible. We change very slightly the method (essentially by taking the more canonical gradient flows of the functions g_{ij} rather than the flows defined in Chapter 3 of [1], and also stopping the flows before arriving at the poles of B). However, we don't obtain the full results of [1], namely we can show an analytic continuation result for the Laplace transform of $m(t)$ (this Laplace transform is explained in more detail below), however we don't get good bounds or information about the singularities of the Laplace transform other than that they are locally finite sets of points. In particular we obtain information about the growth rate of $m(t)$ but not asymptotic expansions.

Even in order to obtain the analytic continuation, a much more detailed examination of the dynamics generated by the general method of [1] is necessary. This is the main body of the present paper (see Theorem 6). For the remainder of the technique we mostly refer to [1].

Thus while we treat a much more general type of curve going to infinity than was treated in [1], we obtain a weaker set of results for these curves. This leaves open the difficult question of what kinds of singularities the Laplace transforms have, and thus what type of asymptotic expansion we can get for $m(t)$. This question will be discussed at the end.

Genericity results for the Higgs field

Before beginning to look more closely at the monodromy representations, we will consider some properties of general points (E, θ) on P_{DR}.

Suppose (E, θ) is a Higgs bundle. Suppose $P \in X$ and $v \in T_P X$; then we obtain the fiber E_P which is a vector space of rank r, with an endomorphism $\theta_P(v) \in End(E_P)$. We say that P is *singular* if $\theta_P(v)$ has an eigenvalue (i.e. zero of the characteristic polynomial) of multiplicity ≥ 2. We say that a singular point P is *generic* if there is exactly one eigenvalue of multiplicity ≥ 2; and if in a neighborhood of P with coordinate z and vector field v nonzero near P, the two eigenvalues $\alpha^{\pm}(z)$ of θ which come together at P, may be expressed as

$$\alpha^{\pm}(v) = \pm a\sqrt{z} + \ldots$$

The condition that all singular points are generic is a Zariski open condition on the moduli space of Higgs bundles.

Suppose P is a generic singular point. We obtain a set of $r-2$ distinct elements of $T_P^* X$ defined by the values of the multiplicity-one eigenvalues of θ at P. Call this set Σ_P. We say that P is *non-parallel* if the $0 \notin \Sigma_P$ and if $\Sigma_P \cup \{0\}$, viewed as a subset of the real two-dimensional space $T_P^* X$, doesn't have any colinear triples, nor any quadruples of points defining two parallel lines.

In terms of a coordinate z at P we can write the elements of Σ_P as

$$\alpha_i(P) = a_i dz$$

with a_i being distinct complex numbers. Then P is non-parallel if and only if the origin plus the set of values of $a_i \in \mathbf{C} \cong \mathbf{R}^2$ doesn't have any colinear triples or parallel quadruples. In turn this is equivalent to saying that the angular coordinates of the complex numbers a_i and $a_i - a_j$ are distinct. If we think of the first two eigenforms $\alpha_{1,2} = \alpha^{\pm}$ as defining forms which start with higher (fractional) powers of z so that their "leading terms" are zero then the non-parallel condition is just to say that the $\alpha_i - \alpha_j$ have distinct angular coordinates except for the double values forced by the fact that there are two whose leading terms are zero. This is the version of the condition we actually use.

The condition of being non-parallel is a real Zariski open condition. In particular, the condition that all singular points be generic and non-colinear, holds in the complement of a closed real algebraic subset of the moduli space. Therefore, if there is one such point then the set of such points is a dense real Zariski open subset.

To show that there is one point (E,θ) such that all of the singular points are generic and non-parallel, we can restrict to the case where $E = \mathcal{O}^{\oplus r}$ is a trivial bundle. In this case, θ corresponds to a matrix of holomorphic one-forms on X. We will consider a matrix of the form $A + \lambda B$ with A diagonal having entries α_i, and B is off-diagonal with λ small. The singular points are perturbations of the points where $\alpha_i(P) = \alpha_j(P)$. A simple calculation with a 2×2 matrix shows that the singularities are generic in this case. In order to obtain the non-colinear condition, it suffices to have that for a point P where $\alpha_i(P) = \alpha_j(P)$, the subset of other $\alpha_k(P)$ is non-parallel.

For a general choice of the α_i, this is the case. Indeed, the values of the $\alpha_k(P)$ move independently of each other when we move the choice of A (in other words, when A is general then the vector of $r-2$ values $\alpha_k(P)$ is general in $(\mathbf{R}^2)^{r-2}$), and a general set of $r-2$ points in $\mathbf{R}^2$ satisfies the non-parallel condition.

Pullback to a ramified covering and gauge transformations

Fix a general Higgs bundle (E,θ) on X. By taking a Galois completion of the spectral curve of θ and making a further two-fold ramified covering if necessary, we can obtain a ramified Galois covering

$$\varphi : Y \to X$$

such that the pullback Higgs field φ^* has a full set of eigen-one-forms defined on Y; and such that the ramification powers over singular points of θ are divisible by 4.

We have one-forms $\alpha_1, \ldots, \alpha_r$ and line sub-bundles

$$L_1, \ldots, L_r \subset \varphi^* E$$

such that at a general point of Y we have

$$\psi : \varphi^* E \cong L_1 \oplus \ldots \oplus L_r$$

with $\varphi^*\theta$ represented by the diagonal matrix with entries α_i. Note that $\varphi^*\theta$ preserves L_i (acting there by multiplication by α_i) globally on Y. However, the isomorphism ψ will only be meromorphic, and also the L_i are of degree < 0. Choose modifications L_i' of L_i (see Lemma 1 below) such that L_i' is of degree zero, and set

$$E' := L_1' \oplus \ldots \oplus L_r'.$$

Let θ' denote the diagonal Higgs field with entries α_i on E'. Let ∇' be a diagonal flat connexion on E'. We have a meromorphic map

$$\psi : E \to E',$$

and

$$\psi \circ \varphi^*\theta \circ \psi^{-1} = \theta'.$$

Suppose now that ∇ was a connexion on E, giving a connexion $\varphi^*\nabla$ on φ^*E. We can write

$$\psi \circ \varphi^*\nabla \circ \psi^{-1} = \nabla' + \beta$$

with β a matrix of meromorphic one-forms on Y.

A monodromy matrix of $(E, \nabla + t\theta)$ may be recovered at least as a transport matrix for the pullback bundle on Y. Indeed if γ is a loop in X based at p then it lifts to a path going from one lift p' of p to another lift p''. Thus it suffices to look at the problem of the asymptotics for transport matrices for the pullback family

$$\{(\varphi^*E, \varphi^*\nabla + t\varphi^*\theta)\}.$$

We may assume that p is not one of the singular points of θ, so p' and p'' will not be singular points of $\varphi^*\theta$. Then the transport matrices for this family are conjugate (by a conjugation which is constant in t) to the transport matrices for the family

$$\{(E', \nabla' + \beta + t\theta')\}.$$

Lemma 1 *In the above situation, the modifications L_i' of L_i may be chosen so that the diagonal entries of β are holomorphic. Furthermore the poles of the remaining entries of β are restricted to the points lying over singular points in X for the original Higgs field θ (the "turning points").*

Proof: Note first that, by definition, away from the singular points of θ the eigen-one-forms are distint so the eigenvectors form a basis for E, in other words the direct sum decomposition ψ is an isomorphism at these points. Thus ψ only has poles over the singular points of θ (hence the same for β).

Look now in a neighborhood of a point $P' \in Y$, lying over a singular point $P \in X$. Let z' denote a local coordinate at P' on Y, with z a local coordinate at P on X and with

$$z = (z')^m.$$

Our assumption on Y was that m is divisible by 4. In fact we may as well assume that $m = 4$ since raising to a further power doesn't modify the argument. Thus we can write

$$z' = z^{1/4}.$$

There are two eigenforms of θ which come together at P. Suppose that their lifts are α_1 and α_2. Then near P' we can write

$$\varphi^* E = U \oplus L_3 \oplus \ldots \oplus L_r$$

where U is the rank two subbundle of $\varphi^* E$ corresponding to eigenvalues α_1 and α_2. The direct sum decomposition is holomorphic at P' because, in X, the other eigenvalues of θ were distinct at P and different from the two singular ones (of course after the pullback all of the eigenforms have a value of zero at P' but the decomposition still holds nonetheless).

Now we use a little bit more detailed information about spectral curves for Higgs bundles: the general (E, θ) is obtained as the direct image of a line bundle on the spectral curve. This means that locally over P there is a two-fold branched covering with coordinate $u = z^{1/2}$ such that the rank 2 subbundle of E corresponding to the singular values looks like the direct image of the trivial bundle on the covering, and the 2×2 piece of θ looks like the action of multiplication by $udz = 2u^2 du$. The direct image, considered as a module over the series in z, is just the series in u; thus one can obtain a basis by looking at the odd and even powers of u: the basis vectors are $e_1 = 1$ and $e_2 = u$. In these terms we have

$$\theta e_1 = e_2 dz; \quad \theta e_2 = z e_1 dz.$$

Thus the 2×2 singular part of θ has matrix

$$\begin{pmatrix} 0 & z \\ 1 & 0 \end{pmatrix} dz.$$

Pulling back now to the covering Y which is locally 4-fold, we have a basis for U in which

$$\varphi^*\theta|_U = \begin{pmatrix} 0 & (z')^7 \\ (z')^3 & 0 \end{pmatrix} dz'.$$

On the other hand, since up until now our decomposition is holomorphic, the pullback connexion $\varphi^*\nabla$ may be written (in terms of our basis for U plus trivializations of the L_i for $i \geq 3$) as $d + B'$ where B' is a holomorphic matrix of one-forms. Since the basis can be pulled back from downstairs, we can even say that B' consists of one-forms pulled back from X.

To choose the modifications L_i' (for $i = 1, 2$) locally at P' we have to find a meromorphic change of basis for the bundle U, which diagonalizes $\varphi^*\theta|_U$. The eigenforms of the matrix are $\pm(z')^5 dz'$ and we can choose eigenvectors

$$e_\pm := \begin{pmatrix} z' \\ \pm(z')^{-1} \end{pmatrix}.$$

Note by calculation that

$$(\varphi^*\theta|_U)e_\pm = (\pm(z')^5 dz')e_\pm.$$

Choose the line bundles L_1' and L_2' to be spanned by the meromorphic sections e_+ and e_- of U. These are indeed eigen-subbundles for $\varphi^*\theta$. We just have to calculate the connexion $\varphi^*\nabla$ on the bundle $U' = L_1' \oplus L_2'$. This bundle is the same as the modification of U given by the meromorphic basis $z'e_1, (z')^{-1}e_2$.

Note first that the matrix B' of one-forms pulled back from X consists of one-forms which have zeros at least like $(z')^3 dz'$. Thus B' transported to U' is still a matrix of holomorphic one-forms so it doesn't affect our lemma. In particular we just have to consider the transport to U' of the connexion d_U constant with respect to the basis (e_1, e_2) on the bundle U.

Calculate

$$d_U(a_+e_+ + a_-e_-) = d_U \begin{pmatrix} (a_+ + a_-)z' \\ (a_+ - a_-)(z')^{-1} \end{pmatrix}$$

$$= \begin{pmatrix} (da_+ + da_-)z' \\ (da_+ - da_-)(z')^{-1} \end{pmatrix} + \begin{pmatrix} (a_+ + a_-)(d\log z')z' \\ -(a_+ - a_-)(d\log z')(z')^{-1} \end{pmatrix}$$

and with the notation $d_{U'}$ for the constant connexion on the bundle U' with respect to its basis $e_\pm$, this is equal to

$$= d_{U'}(a_+e_+ + a_-e_-) + a_+(d\log z')e_- + a_-(d\log z')e_+.$$

We conclude that the connexion matrix β is, up to a holomorphic piece, just the 2×2 matrix

$$\begin{pmatrix} 0 & (z')^{-1} \\ (z')^{-1} & 0 \end{pmatrix} dz'.$$

In particular the diagonal terms of β are holomorphic, as desired for Lemma 1.
///

Remarks:
(i) The above proof gives further information: the only terms with poles in the matrix β are the off-diagonal terms corresponding to the two eigenvalues which came together originally downstairs in X; and these terms have exactly logarithmic (i.e. first-order) poles with residue 1. This information might be useful in trying to improve the current results in order to obtain precise expansions at the singularities of the Laplace transform of the monodromy.
(ii) The fact that we had to go to a covering whose ramification power is divisible by 4 rather than just 2 (as would be sufficient just for diagonalizing θ) is somewhat mysterious; it probably indicates that we (or some of us at least) don't fully understand what is going on here.

Let β^{diag} denote the matrix of diagonal entries of β. Let $Z = \widetilde{Y}$ be the universal covering. Over Z we can use the diagonal connexion $\nabla' + \beta^{\mathrm{diag}}$ to trivialize

$$E'|_Z \cong \mathcal{O}_Z^r.$$

With respect to this trivialization , our family now has the form of a family of connexions

$$\{(\mathcal{O}_Z^r, d + B + tA)\}$$

where A (corresponding to the pullback of θ' to Z) is the diagonal matrix whose entries are the pullbacks of the α_i; and where B is a matrix whose diagonal entries are zero, and whose off-diagonal entries are meromorphic with poles at the points lying over singular points for θ.

We can now apply the method developped in [1] to this family of connexions. Note that it is important to know that the diagonal entries of A come from forms on the compact Riemann surface Y; on the other hand the fact that B is only defined over the universal covering Z is not a problem. The next two sections will constitute a brief discussion of how the method of [1] works; however the reader is refered back there for the full details.

Laplace transform of the monodromy operators

We now look at a family of connexions of the form $d+B+tA$ on the trivial bundle $\mathcal{O}^r$ on the universal covering Z of the ramified cover Y, where A is a diagonal matrix with one-forms α_i along the diagonal, and B is a matrix of meromorphic one-forms with zeros on the diagonal. We assume that the poles of B are at points P which come from the original singular points of the Higgs field θ on X. We make no assumption about the order of poles, in spite of the additional information given by Remark (i) after the proof of Lemma 1 above.

Assume that p and q are two points in Z, not on the singular points. These could correspond to two lifts of the base point as explained above. Choose a path γ from p to q not passing through the singular points. We obtain the *transport matrix* $m(t)$ for continuing solutions of the ordinary differential equation $(d+B+tA)f=0$ from p to q along the path γ. Note that $m(t)$ is a holomorphic $r\times r$-matrix-valued function defined for all $t\in \mathbf{C}$.

Recall from [1] that after a gauge transformation and an expansion as a sum of iterated integrals, we obtain a formula for the transport matrix. One way of thinking of this formula is to look at the transport for the connexion $d+sB+tA$ and expand in a Taylor series in s about the point $s=0$, then evaluate at $s=1$. The terms in the expansion are the higher derivatives in s, at $s=0$, which are functions of t. A concrete derivation of the formula is given in [1]. It says

$$m(t)=\sum_I \int_{\eta_I} b_I e^{tg_I}$$

where:

—the sum is taken over multi-indices of the form $I=(i_0,i_1,\ldots,i_k)$ where

we note $k = |I|$;
—for a multi-index I we denote by Z_I the product of $k = |I|$ factors $Z \times \ldots \times Z$;
—in Z_I we have a cycle

$$\eta_I := \{(\gamma(t_1), \ldots, \gamma(t_k))\}$$

for $0 \leq t_1 \leq \ldots \leq t_k \leq 1$ where γ is viewed as a path parametrized by $t \in [0, 1]$;
—the cycle η_I should be thought of as representing a class in a relative homology group of Z_I relative to the simplex formed by points where $z_i = z_{i+1}$ or at the ends $z_1 = p$ or $z_k = q$;
—the matrix B leads to a (now meromorphic) matrix-valued k-form b_I on Z_I defined as follows: if the entries of B are denoted $b_{ij}(z)dz$ then

$$b_I = b_{i_k i_{k-1}}(z_k)dz_k \wedge \ldots \wedge b_{i_1 i_0}(z_1)dz_1 \mathbf{e}_{i_k i_0}$$

where $\mathbf{e}_{i_k i_0}$ denotes the elementary matrix with zeros everywhere except for a 1 in the $i_k i_0$ place; —and finally g_I is a holomorphic function $Z_I \to \mathbf{C}$ defined by integrating the one-forms α_i as follows:

$$g_I(z_1, \ldots, z_k) = \int_p^{z_1} \alpha_{i_1} + \ldots + \int z_k{}^q \alpha_{i_k}.$$

The fact that b_I is meromorphic rather than holomorphic is the only difference between our present situation and the situation of [1]. Note that because our path γ misses the singular points and thus the poles of B, the cycle η_I is supported away from the poles of b_I. We will be applying essentially the same technique of moving the cycle of integration η, but we need to do additional work to make sure it stays away from the poles of b_I.

It is useful to have the formula

$$g_I(z_1, \ldots, z_k) = g_{i_0 i_1}(z_1) + \ldots + g_{i_{k-1} i_k}(z_k) + \int_p^q \alpha_{i_k},$$

where

$$g_{ij}(z) := \int_p^z \alpha_i - \alpha_j.$$

Our formula for m gives a preliminary bound of the form

$$|m(t)| \leq Ce^a |t|.$$

Indeed, along the path γ the one-forms b_{ij} are bounded, so

$$|b_I| \leq C^k$$

on η_I; also we have a bound $|g_I(z)| \leq a$ for $z \in \eta_I$, uniform in I; and finally the cycle of integration η_I has size $(k!)^{-1}$. Putting these together gives the bound for $m(t)$.

Recall now that the *Laplace transform* of a function $m(t)$ which satisfies a bound such as the above, is by definition the integral

$$f(\zeta) := \int_0^\infty m(t)e^{-\zeta t}dt$$

where $\zeta \in \mathbf{C}$ with $|\zeta| > a$ and the path of integration is taken in a suitably chosen direction so that the integrand is rapidly decreasing at infinity. In our case since $m(t)$ is a matrix, $f(\zeta)$ is also a matrix. We can recover $m(t)$ by the inverse transform

$$m(t) = \frac{1}{2\pi i}\oint f(\zeta)e^{\zeta t}d\zeta$$

with the integral being taken over a loop going around once counterclockwise in the region $|\zeta| > a$.

The singularities of $f(\zeta)$ are directly related to the asymptotic behavior of $m(t)$. This is a classical subject which we don't discuss in the present version. One might just note that by the inverse transform, there exist functions $m(t)$ satisfying the preliminary bound $|m(t)| \leq Ce^a|t|$ but such that the Laplace transforms $f(\zeta)$ have arbitrarily bad singularities in the region $|\zeta| \leq a$. Thus getting any nontrivial restrictions on the singularities of f amounts to a restriction on which types of functions $m(t)$ can occur.

In our case, the expansion formula for $m(t)$ leads to a similar formula for the Laplace tranform:

$$f(\zeta) = \sum_I \int_{\eta_I} \frac{b_I}{g_I - \zeta}.$$

In particular this formula (plus the easy bounds for b_I and the size of η_I) automatically gives an analytic continuation of $f(\zeta)$ to the complement of the union

$$g(\eta) := \bigcup_I g_I(\eta_I) \subset \mathbf{C}.$$

A naive approach would be to try to move the path γ so as to move the union of images $g(\eta)$ and analytically continue f to a larger region in this way. The 3×3 example at the end of [1] shows that this approach cannot be optimal. In fact, we should instead move each cycle of integration η_I individually. Unfortunately this has to be done with great care in order to maintain control of the sizes of the individual terms so that the infinite sum over I still converges.

We finish the section by stating our main theorem. Recall that a function such as $f(\zeta)$ defined on $|\zeta| > a$ is said to have an *analytic continuation with locally finite branching* if for every $M > 0$ there is a finite set of points $S_M \subset \mathbf{C}$ such that if σ is any piecewise linear path in $\mathbf{C} - S_M$ starting at a point where $|\zeta| > a$ and such that the length of σ is $\leq M$, then $f(\zeta)$ can be analytically continued along σ.

Theorem 2 *Suppose $m(t)$ is the transport matrix from p to q for a family of connexions on the trivial bundle $\mathcal{O}^r_Z$ of the form $\{d + B + tA\}$. Suppose that A is diagonal with one-forms α_i, coming from the pullback of a general Higgs field θ over the original curve X, and suppose that B is a meromorphic matrix of one-forms with poles only at points lying over the singular points of θ. Let $f(\zeta)$ denote the Laplace transform of $m(t)$. Then f has an analytic continuation with locally finite branching.*

The remainder of these notes is devoted to explaining the proof.

Analytic continuation of the Laplace transform—the method

We now recall the basic method of [1] for moving the cycles η_I in order to analytically continue $f(\zeta)$. We refer there for most details and concentrate here just on stating what the end result is. Still we need a minimum amount of notation. We work with *pro-chains* which are formal sums of the form $\eta = \sum_I \eta_I$ of chains on the Z_I (and now we are forced to assume that the support of η_I avoids points $(z_1, \ldots, P, \ldots, z_k)$ where P is any singular point, because of the poles of B). We have a boundary operator denoted $\partial + A$ where ∂ is the usual boundary operator on each η_I individually, and A (different from the matrix of one-forms considered above) is a signed sum of face maps corresponding to the inclusions $Z_{I'} \to Z_I$ obtained when some $z_i = z_{i+1}$. Our original pro-chain of integration in the integral expansion satisfies $(\partial + A)\eta = 0$. We can write the expansion formula as an integral over the pro-chain $\eta = \sum_I \eta_I$ as

$$f(\zeta) = \int_\eta \frac{b}{g - \zeta}$$

where b is the union of forms b_I on Z_I and g is the union of functions g_I.

In general if we add to η a boundary term of the form $(\partial + A)\kappa$ then the integral doesn't change:

$$\int_{\eta+(\partial+A)\kappa} \frac{b}{g-\zeta} = \int_{\eta} \frac{b}{g-\zeta}.$$

With this in mind, the goal is to progressively modify η by addition of boundary terms, so as to obtain the analytic continuation.

Suppose now in general that η is a pro-chain obtained by a sequence of such modifications, allowing us to analytically continue $f(\zeta)$ to a neighborhood of a point $\zeta_0 \in \mathbf{C}$. In particular the image $g(\eta)$ doesn't meet a disc around ζ_0. Fix a line segment S going from ζ_0 to another point ζ_1; we would like to continue f in a neighborhood of S. By making a rotation in the complex plane (which can be seen as a rotation of the original Higgs field) we may without loss of generality assume that the segment S is parallel to the real axis and the real part of z_1 is smaller than the real part of ζ_0. Let u be a cut-off function for a neighborhood of S and write

$$\eta = \eta' + \eta'', \quad \eta' = g^*(u) \cdot \eta.$$

Let N_1 be the support of u, which is a neigborhood of S, and let N_2 be the support of du which is an oval going around S but not touching it. Let N_3 be the neighborhood of S where u is identically 1. Let D be a disc around ζ_0, such that $g(\eta)$ misses D, and which we may assume has radius bigger than the width of N_1. Then

$$g(\eta') \subset N_1 - (N_1 \cap D),$$

$$g(\eta'') \subset \mathbf{C} - (N_3 \cup D),$$

and $(\partial A)\eta' = -(\partial A)\eta''$ with

$$g((\partial A)\eta') \subset N_2 - (N_2 \cap D).$$

In particular the integral over the chain η'' is already analytically continued in a neighborhood N_3 of S, so we can ignore this piece. The piece η' needs to be moved; it corresponds to what was denoted η in Chapter 4 of [1]. After choosing flows (our choice of flows will be slightly different, see below) the process from that chapter applied to η' gives pro-chains denoted $F\tau$, $F\psi$

and $FK\varphi$ with formulae (using the operators F, K and H from Chapter 4, cf the start of Chapter 5)

$$F\tau = \sum_{r,s} F(-KA)^r H(AK)^s \eta',$$

$$F\psi = \sum_r F(-KA)^r K(\partial + A)\eta',$$

$$FK\varphi = \sum_r FK(AK)^r \eta'.$$

The operator K corresponds to applying the flows in the various coordinates. This has the effect of decreasing the real part $\Re g$. In fact in our case we will use flows along vector fields which are multiples of ∇g_{ij}, in particular the flows respect strictly the imaginary part of g. This differs from the case of [1] and means we can avoid discussion of "angular sectors" such as on pages 52-53. Thus, in our case, when we apply a flow to a point, the new point has the same value of $\Im g$, and the real part $\Re g$ is decreased.

The operator F is related to the use of buffers; we refer to [1] for that discussion and heretofore ignore it. The operator A is the boundary operator discussed above; and the operator H is just the result of doing the flows K after unit time. In particular, we expect the operator H to significantly decrease the value of the real part of g. Thus, we hope that $g(F\tau)$ will go beyond the segment S and be supported outside of N_3 for example. On the other hand, given that the boundary term $(\partial + A)\eta'$ is supported in the U-shaped region N_2, we already have that $g(F\psi)$ is supported away from N_3. Finally, the chain $FK\varphi$ is the one which serves as homotopy between the other ones (cf [1] Lemma 4.4):

$$\eta' + (\partial + A)FK\varphi = F\tau - F\psi.$$

In particular

$$f(\zeta) = \int_{F\tau - F\psi + \eta''} \frac{b}{g - \zeta}.$$

A major part of the proof of the analytic continuation is to bound the sizes of all of the above chains, justifying in particular the change of integrals in the preceding line which amounts to reorganizing an infinite sum. This is done exactly as in [1] and we won't get into that here.

With these bounds among other things, and because of the discussion of the supports $g(F\psi)$ and $g(\eta'')$ above, the term $\int_{F\psi+\eta''} \frac{b}{g-\zeta}$ is already analytically continued in the neighborhood $\zeta \in N_3$.

Another thing we don't discuss any further is why everything which interests us is going to take place inside a fixed relatively compact subset in Z. This again is discussed adequately in [1] and there is no change in our present situation.

In order to complete the proof of the analytic continuation for Theorem 2, we just have to choose our flowing vector fields in such a way that everything stays a bounded distance away from the singular points which are poles of B, and then verify that the support $g(F\tau)$ indeed goes away from the segment S. Of course this will depend on the choice of segment S because it shouldn't cross the actual singularities of f!

Description of cells using trees

As was used in [1], the chains defined above can be expressed as sums of cells; the cells are cubical in nature, and are parametrized by trees. We are most interested in the chain $F\tau$ although what we say also applies to the other ones. These chains are unions of cells which have the form of a family of cubes parametrized by points in one of the original cells η'_I. We call these things just cubes. The points in the cubes which occur are parametrized by "trees" by which we mean binary planar trees T sandwiched between a top horizontal line and a bottom horizontal line, and provided with additional information as follows. For each top vertex of the tree (i.e. where an edge meets the top horizontal line) we should specify a point $z \in Z$. For each region in the complement of the tree (between the top and bottom horizontal lines) we should specify an index. In this way each edge of the tree is provided with left and right indices which will be denoted i_e and j_e below. This indexation also leads to an indexation of the segments between vertices along the top and bottom lines. We obtain in this way multi-indices $I^{\rm top}$ and $I^{\rm bot}$, and the collection of points $(z_1, \ldots, z_k)$ attached to the top vertices gives a point $z^{\rm top} \in Z_{I^{\rm top}}$. This point should be in the support of η'. (Also note that in reality the chain η' is a union of cells possiblly multiplied by cutoff functions; the contribution of the cube corresponding to the tree is of course multiplied by the corresponding factor in the expression of η'.) Finally, to each edge e of the tree is assigned a "length" $s(e) \in [0, 1]$.

Henceforth by the notation *tree* T we shall mean a tree provided with all of the above information.

The cubes in question which enter into the chain $F\tau$ (or the other ones) are obtained by specifying which of the edges are to have variable lengths s; some of the other edges are specified as always having length 0 or length 1. A simple count shows that in order to get a cube of the right dimension, half of the edge lengths should be variable and half should be fixed. The chain $F\tau$ is a sum of various cubes coming from cubes of trees of this form.

We have to explain how a tree T leads to a point in Z_I. Specifically it leads to a point in $Z_{I^{\mathrm{bot}}(T)}$. Suppose we have made a choice of vector fields V_{ij} for each pair of indices i, j. Then a *flowing map* $\Phi : T \rightarrow Z$ is a map from the topological realization of the tree, into Z, satisfying the following properties:
—if v is a top vertex which is assigned a point z in the information contained in T, then $\Phi(v) = z$;
—if e is an edge with left and right indices i_e and j_e and with initial vertex v and terminal vertex v', then $\Phi(e)$ is the flow curve for flowing along the vector field $V_{i_e j_e}$ from $\Phi(v)$ to $\Phi(v')$, where the flow is done for time $s = s(e)$. This determines $\Phi(v')$ as a function of $\Phi(v)$ and the information in the tree. Thus by recursion we determine the $\Phi(v)$ for all vertices, as well as the paths $\Phi(e)$ for the edges e (the map Φ on the edges is only well determined up to reparametrization because we don't fix a parametrization of the edges; the length s is abstract, since it is convenient to picture even edges assigned $s = 0$ as being actual edges).

For a given choice of vector fields V_{ij} and of information attached to the tree T, the flowing map exists and is unique. This determines a point given by the values z at the bottom vertices,

$$z^{\mathrm{bot}} \in Z_{I^{\mathrm{bot}}}.$$

The tree T contributes this point to the chain in question.

See [1], pages 54-55 for a discussion of how the chains such as $F\tau$ correspond to sums of cubes coming from trees. The only information we need to retain is that for the trees entering into the chain $F\tau$, there is always (up to reparametrization of the planar embedding) a horizontal line which cuts the tree along a sequence of edges, such that all of these edges are assigned the fixed value $s = 1$. This corresponds to the operator H in the formula for $F\tau$ written above.

We finish this section by pointing out the relationship between $g(z^{\mathrm{top}})$ and $g(z^{\mathrm{bot}})$. This is the key point in our discussion, because z^{top} is the input point coming from the chain η' and z^{bot} is the output point which goes into the resulting chain $F\tau$. We want to prove that the real part of $g(z^{\rm bot})$ can be moved down past the end of the segment S.

If e is an edge of T and $s' \in [0, s(e)]$ then we can define the tree T' obtained by *pruning* T at (e, s'). This is obtained by cutting off everything below e and sending the bottom vertex of e to the line at the bottom. The indices associated to regions in the complement follow accordingly. Finally we set $s(e) := s'$ in the new tree T'.

Suppose for the same edge e we also pick $s'' \in [s', s(e)]$. Then we obtain a different pruning denoted T'' (which has almost all the same information except for the length of the edge e). Let v' (resp. v'') denote the bottom vertices corresponding to e in the trees T' (resp. T''). Let Φ' (resp. Φ'') denote the flowing map for T' (resp. T''). These coincide and coincide with Φ on the parts of the trees that are in common (the unpruned parts). We have

$$g(z^{\rm bot}(T'')) = g(z^{\rm bot}(T')) + \int_{\Phi'(v')}^{\Phi''(v'')} dg_{i_e j_e}.$$

Note that the segment of $\Phi(e)$ going from $\Phi'(v')$ to $\Phi''(v'')$ is a flow curve for the vector field $V_{i_e j_e}$, and it flows for time $s'' - s'$.

We isolate the following condition on our vector fields:

Condition (R): the vector fields $V_{i_e j_e}$ are negative multiples of the gradient vector fields for the real functions $\Re g_{i_e j_e}$.

With this condition we get that the integral of $dg_{i_e j_e}$ along a flow curve for $V_{i_e j_e}$ is a negative real number, so this gives

$$g(z^{\rm bot}(T'')) - g(z^{\rm bot}(T')) \in \mathbf{R}_{\leq 0}.$$

If we prune at an edge e with $s' = s(e)$ then it amounts to cutting off the tree at the lower vertex of e. If furthermore all of the length vectors assigned to edges below e are 0, then $g(z^{\rm bot}(T')) = g(z^{\rm bot}(T))$.

By recurrence we obtain the following description:

$$g(z^{\rm bot}(T)) = g(z^{\rm top}(T)) + \sum_e \int_{\Phi(e)} dg_{i_e j_e},$$

and in particular

$$g(z^{\rm bot}(T)) - g(z^{\rm top}(T)) \in \mathbf{R}_{\leq 0}.$$

This explains more precisely something which was said about the supports in the previous section.

There is also another way to prune a tree: if e is an edge such that $i_e = j_e$ then we can cut off e and all of the edges below it, and consolidate the two edges above and to the side of e into one edge. The only difficulty here is that the consolidated edge might have total length > 1 but this doesn't affect the remainder of our argument (since at this point we can ignore questions about the sizes of the cells). Let T' denote the pruned tree obtained in this way. We again have

$$g(z^{\mathrm{bot}}(T)) - g(z^{\mathrm{bot}}(T')) \in \mathbf{R}_{\leq 0}.$$

In general we will be trying to show for the trees which arise in $F\tau$, that the real part of $g(z^{\mathrm{bot}}(T))$ is small enough. If we can show it for T' then it follows also for T. In this way we can reduce for the remainder of the argument, to the case where $i_e \neq j_e$ for all edges of T.

Another small remark is that technically speaking we should give a separate discussion for the edges which were called "side edges" in [1]. There is nothing new to say here so we leave it to the reader.

Calculations of gradient flows

We express the gradient of the real part of a holomorphic function, as a vector field in a usual coordinate and in logarithmic coordinates. This is of course trivial but we do the calculation just to get the formula right. Suppose z is a coordinate in a coordinate patch on X. The metric on X may be expressed by the real-valued positive function

$$h(z) := \frac{|dz|^2}{2}.$$

Write $z = x + iy$. Note that dx and dy are perpendicular and have the same length, so

$$h(z) = |dx|^2.$$

The real tangent space has orthogonal basis

$$\{\frac{\partial}{\partial x}, \frac{\partial}{\partial y}\}$$

and the formula

$$1 = |\frac{\partial}{\partial x} \cdot dx| = h^{1/2} |\frac{\partial}{\partial x}|$$

yields

$$|\frac{\partial}{\partial x}| = h^{-1/2}.$$

In particular an orthonormal basis for the real tangent space is given by

$$\{h^{1/2}\frac{\partial}{\partial x}, h^{1/2}\frac{\partial}{\partial y}\}.$$

Thus we have the formula, for any function a:

$$\nabla a = h\frac{\partial a}{\partial x}\frac{\partial}{\partial x} + h\frac{\partial a}{\partial y}\frac{\partial}{\partial y}.$$

Now suppose $g = a + ib$ is a holomorphic function (with a, b real), and pose $f(z) := \frac{\partial g}{\partial z}$ so that $dg = f(z)dz$. Write $f(z) = u + iv$ with u, v real, and expand:

$$(u + iv)(dx + idy) = \frac{\partial a}{\partial x}dx + \frac{\partial a}{\partial y}dy + i\frac{\partial b}{\partial x}dx + i\frac{\partial b}{\partial y}dy.$$

Comparing both sides we get

$$u = \frac{\partial a}{\partial x}, \quad v = -\frac{\partial a}{\partial y}.$$

Note that $a = \Re g$ is the real part of g, so finally we have the formula

$$\nabla \Re g = h(z)\left((\Re\frac{\partial g}{\partial z})\frac{\partial}{\partial x} - (\Im\frac{\partial g}{\partial z})\frac{\partial}{\partial y}\right).$$

Suppose now that w is a local coordinate at a point P, and consider

$$g = a_m w^m.$$

Let $z = -i\log w$ so $w = e^{iz}$, and writing $z = x + iy$ we have $w = e^{ix-y}$. Then

$$g(z) = a_m e^{imz}; \quad \frac{\partial g}{\partial z} = mia_m e^{imz}.$$

If we write $mia_m = e^{r+is}$ then

$$\frac{\partial g}{\partial z} = e^{r-my+i(s+mx)},$$

so

$$\nabla \Re g = h(z) e^{r-my} \left(\cos(s+mx) \frac{\partial}{\partial x}, \sin(s+mx) \frac{\partial}{\partial y} \right).$$

The *asymptotes* are the values $x = B$ where $\cos(s+mx) = 0$. At these points, the gradient flow vector field is vertical (going either up or down, depending on the sign of $\sin(s+mx)$). If the flow goes up, then it stays on the vertical line until $y = \infty$.

Note that the gradient of $\Re g$ is perpendicular to the level curves of $\Re g$, so it is parallel to the level curves of $\Im g$. Which is to say that the level curves of $\Im g$ are the flow lines. This gives an idea of the dynamics of the flow. We have

$$\Im g = \Im(im^{-1} e^{(r-my)+i(s+mx)}) = -m^{-1} e^{r-my} \sin(s+mx).$$

Thus a curve $\Im g = C$ is given by

$$e^{-my} = \frac{-mC}{e^r \sin(s+mx)}$$

or (noting that the sign of C must be chosen so that the right hand side is positive)

$$y = m^{-1} r \log|\sin(s+mx)| - m^{-1} \log|mC|.$$

In particular the level curves are all vertical translates of the same curve; this curve $y = m^{-1} r \log|\sin(s+mx)|$ has vertical asymptotes at the points where $\sin(s+mx) = 0$. Note however that at the asymptotes, we get $y \to -\infty$; whereas our coordinate patch corresponds to a region $y > y_0$. Thus, every gradient flow except for the inbound (i.e. upward) flows directly on the asymptotes, eventually turns around and exits the coordinate patch. This of course corresponds to what the classical picture looks like in terms of the original coordinate w.

Also we can calculate the second derivative (which depends only on x and not on which level curve we are on, since they are all vertical translates). Consider for example points where $\sin(s+mx) > 0$. There

$$\frac{dy}{dx} = \frac{r\cos(s+mx)}{\sin(s+mx)}$$

and

$$\frac{d^2 y}{dx^2} = \frac{-rm}{\sin^2(s+mx)}$$

In particular note that we have a uniform bound everywhere:

$$\frac{d^2y}{dx^2} \leq -\gamma,$$

here with $\gamma = rm$.

Suppose now more generally that g is a holomorphic function with Taylor expansion

$$g = a_m w^m + a_{m+1} w^{m+1} + \dots .$$

Then we will get

$$h(z)^{-1} e^{my-r} \nabla \Re g = \left(\cos(s + mx) \frac{\partial}{\partial x}, \sin(s + mx) \frac{\partial}{\partial y} \right) + O(e^- y).$$

In particular, the direction of the gradient flow for g is determined, up to an error term in $O(e^{-y})$, by the vector $(\cos(s + mx), \sin(s + mx))$.

The asymptotes are no longer vertical curves, but they remain in bands $x \in B_{ij,a}$. Also we can choose A in the definition of steepness, so that at non-steep parts of the level curves we still have a bound

$$\frac{d^2y}{dx^2} \leq -\gamma.$$

Choice of the vector fields V_{ij}

The only thing left to be determined in order to fix our procedure for moving the cycle of integration is to choose the vector fields V_{ij}. We will fix discs $D_\epsilon(P)$ of radius ϵ around the singular points P and let δ denote a function which is equal to a constant value λ outside the $D_\epsilon(P)$ and which is equal to zero in the neighborhoods $D_{\epsilon/2}(P)$ of the singular points P. Use a metric h, say pulled back from a smooth metric on Y, to define the gradient of a function; and set

$$V_{ij} := -\delta \nabla \Re g_{ij}.$$

The first thing to notice is that flowing along the vector fields V_{ij} can never get us to the interior of the neighborhoods $D_{\epsilon/2}(P)$. Any flowing map Φ defined using our vector fields, will map T into the complement of the $D_{\epsilon/2}(P)$ provided the top points $z_i^{\rm top}$ are not in $D_{\epsilon/2}(P)$. Thus everything we do will take place in a region of Z where the matrix B is defined and holomorphic, and even has a uniform bound (however the bound depends on the choice of ϵ).

Furthermore note that flowing along V_{ij} is the same as flowing along $-\lambda\nabla\Re g_{ij}$ unless the flow curve enters a disc of the form $D_\epsilon(P)$. In particular, once we have chosen ϵ, we can choose the constant value λ large enough so that the following property holds: if $\Phi(e)$ is a flow curve for V_{ij} which never enters into any $D_\epsilon(P)$, and which flows for time $s \geq 1$, then

$$\int_{\Phi(e)} dg_{ij} \leq \zeta_1 - \zeta_0.$$

Recall that ζ_0, ζ_1 were the endpoints of the segment S parallel to the real axis, along which we wanted to analytically continue; and $\zeta_1 - \zeta_0$ is a negative real number.

Denote by ν the width of the neighborhood N_3 of S. It is smaller than the diameter of the disc D around ζ_0. Note that because of our hypothesis on the support of η', for any tree which enters into $F\tau$ we have

$$\Re g(z^{\rm top}(T)) \leq \Re\zeta_0 - \nu.$$

Thus if T contains an edge e assigned a length $s \geq 1$, and such that the flow curve $\Phi(e)$ doesn't go into any $D_\epsilon(P)$, then

$$\Re g(z^{\rm bot}(T)) \leq \Re\zeta_1 - \nu$$

in other words $g(z^{\rm bot}(T))$ is outside of the neighborhood N_3.

¿From the above discussion, the only case which can pose a problem is when all of the edges of length ≥ 1 flow into neighborhoods $D_\epsilon(P)$. Recall that for trees entering into $F\tau$, there is a whole horizontal line of edges of length ≥ 1. We can therefore prune T to get a tree T' such that all of the bottom vertices v satisfy $\Phi(v) \in D_\epsilon(P(v))$. We consider this case in the next section.

Results on the dynamics of our flowing maps

Let P denote one of our singular points. We will consider a system of discs centered at P:

$$D_\epsilon(P) \subset D_\xi(P) \subset D_u(P) \subset D_w(P).$$

We will first fix u and w so that certain things are true in a coordinate system for $D_w(P)$ (and say $u = w/2$). Then once u and w are fixed we will let $\epsilon \to 0$. Finally $\xi > \epsilon$ will be a function of ϵ with $\xi \to 0$ when $\epsilon \to 0$.

Our first lemma bounds the number of outgoing subtrees.

Lemma 3 *If T is a tree with one top edge e, and if $\Phi : T \to X$ is a flowing map such that the images of all bottom vertices are contained in some $D_\epsilon(P_i)$, and if $\Phi(e)$ exits from $D_u(P)$ then T contains a strand σ such that $\Phi(\sigma)$ exits from $D_w(P)$ also.*

Our next lemma gives a normal form for any subtree which stays entirely within $D_u(P)$.

Lemma 4 *If T is a tree with one edge e at the top, and if Φ is a flowing map from T into $D_u(P) \subset X$ such that all of the bottom vertices are mapped into $D_\epsilon(P)$, then the curve $\Phi(e)$ passes into $D_\xi(P)$, and flows along a vector field $V_{i_e j_e}$ in an ingoing sector near an ingoing curve $G_{i_e j_e}$.*

The last of our preliminary lemmas bounds the number of subtrees having the previous normal form.

Lemma 5 *There is a number K such that if T is a tree consisting of one edge strand κ plus a number of sub-trees coming out of κ, and if Φ is a flowing map from T into $D_u(P)$ with the property that all the sub-trees coming out of κ are covered by Lemma 4, then there are $\leq K$ of these sub-trees.*

For the proofs of these lemmas, we will use a logarithmic coordinate system for $D_w(P)$. If z_D denotes the coordinate in the disc then we introduce $z_L = -i \log z_D$ and write $z_L = x + iy$ $z_D = e^{ix-y}$.

The disc $D_w(P)$ is given by $y > y_0$.

The vector fields $V_{ij} = \nabla g_{ij}$ are approximately equal (up to a term smaller by a factor of $O(e^{-y})$) to the standard vector fields $V'_{ij} = \nabla g'_{ij}$ where g'_{ij} is the leading term in the Taylor expansion for g_{ij} at P.

Because of this, we obtain the following facts. The asymptotic directions (which are close to vertical lines) occur in bands of the form $x \in B_{ij,a}$ where $B_{ij,a} \subset \mathbf{R}$ are intervals which can be made as small as we like by modifying y_0. These intervals are disjoint, except for the asymptotes of the pairs $\{V_{ik}, V_{jk}\}$ or $\{V_{ki}, V_{kj}\}$, where i, j are the two indices attached to P, and k is any index different from these two. In those cases the pairs share the same values $B_{ij,a}$ and the same bands. We say that a vector field V_{ij} is *attached* to an interval B if $B = B_{ij,a}$. The only intervals with more than one vector field attached to them are those described above.

It is worth mentionning why we have this disjointness property. This is because of the non-parallel condition on the eigenforms of θ at the singular points. The non-parallel condition implies that the bands, which are the solutions of $s + mx = 0$ modulo π, are distinct, because the values of s (which are the angular coordinates of the constants attached to the leading terms of g_{ij} as explained in the preceding section) are different exactly by the non-parallel condition. Notice that the exponents m are the same for all of the values ij except the two attached to the singular point; for those which are attached the value m' is bigger. The non-parallel condition gives disjointness for all of the bands except the ones corresponding to the attached indices ij and ji. For those guys, note that if we make a general rotation of everything, those asymptotic solutions of $s + m'x$ move differently than the solutions of $s + mx$, so those bands are disjoint from all the other ones. The general rotation of everything corresponds to a condition that the line segments in the complex plane along which we analytically continue, might be constrained to not be parallel to a certain finite number of directions. This doesn't hurt our ability to analytically continue the function.

We can fix a number $A > 0$ with the following properties: outside of an asymptotic band for V_{ij} or V_{ji}, the slope of the vector V_{ij} satisfies

$$\left|\frac{dy}{dx}(V_{ij})\right| \leq A.$$

Inside an asymptotic band B, only the vector fields V_{ij} which are attached to B can have slope bigger than A or less than $-A$.

Suppose now that $(x(t), y(t))$ is a flow along one of the vector fields V_{ij}. We say that the path is *steep* if

$$\left|\frac{dy}{dx}\right| > A,$$

and we say that it is *not steep* otherwise. We say that the path is *ingoing* if $\frac{dy}{dx} > 0$ and outgoing otherwise. Note that with our logarithmic coordinate system, outgoing is downward and ingoing is upward. The coordinate patch (i.e. choice of y_0) and the choice of A can be made so that all of the paths satisfy the following property:

—once the path is steep and outgoing, it remains steep and outgoing for the remainder of the time of definition, and ends up leaving the region $y > y_0$. This is true even though the vector field is not exactly equal to the standard model but only close to it.

On the other hand, the direction, i.e. the sign of $\frac{dx}{dt}$ remains the same throughout the interval where the path is not steep. Call this sign $(-1)^m$. In particular we can think of the path as being parametrized by x.

We have a bound, in the region where the path is not steep:

$$\frac{d^2y}{dx^2} \leq -\gamma$$

with $\gamma > 0$ a positive constant.

In particular, once the path is outgoing it remains outgoing for the remainder of its period of definition. This is because of the second derivative when it is not steep, and the fact that when it becomes steep and outgoing then it stays that way.

We now note the *additive relation* for the vector fields at vertices of a tree. Suppose we are in the situation of a flowing map $\Phi : T \rightarrow X$. At any vertex v of T with edges noted e_1, e_2, e_3 (say e_1 ingoing and e_2, e_3 outgoing), we have three indices i, j, k such that

$$i_{e_1} = i_{e_2} = i;\ j_{e_2} = i_{e_3} = j;\ j_{e_1} = j_{e_3} = k.$$

Then the three vector fields for the edges are respectively V_{ik}, V_{ij}, V_{jk} and we have the relation

$$V_{ik}(Phi(v)) = V_{ij}(Phi(v)) + V_{jk}(Phi(v)).$$

Proof of Lemma 3. The disc $D_u(P)$ will be determined by $y > y_1$ for some y_1 fixed as a function of y_0 (and in fact one could take $y_1 = y_0 + 1$ for example). A consequence of the additive relation is that if $V_{ik}(\Phi(v))$ is outgoing (i.e. $\frac{dy}{dt} \leq 0$ along this vector) then one of the other two $V_{ij}(Phi(v))$ or $V_{jk}(Phi(v))$ will also be outgoing. As we have noted above, if the flow along any edge is outgoing at some point then it is outgoing for all further points. In particular if at any point in the tree the flow is outgoing then we can choose a strand going down to the bottom, along which the flow is always outgoing. If there is an edge which crosses out of $D_u(P)$, at the crossing point it has $\frac{dy}{dt} \leq 0$, so we get a strand which maintains $\frac{dy}{dt} \leq 0$ as long as it stays inside $D_w(P)$. In particular the strand cannot go back to $D_\epsilon(P)$ so it must exit from $D_w(P)$ (here using the hypothesis that any strand must end in some $D_\epsilon(P_i)$). This completes the proof of Lemma 3.
///

Now we come to the proofs of Lemmas 4 and 5. Fix notations $L := \log \epsilon$ and $L_1 := \log \xi$. Thus we will let $L \to -\infty$ and we have to specify L_1 as a function of L such that $L_1 \to -\infty$ too. Our discs $D_\epsilon(P)$ and $D_\xi(P)$ respectively become the regions $y < L$ and $y < L_1$. We will specify L_1 as a function of L so as to make the proofs of Lemmas 4 and 5 work.

In both lemmas, we lift the maps Φ into maps into the coordinate chart for the logarithmic coordinates.

Proof of Lemma 4. At any point where the flow is not steep, the second derivative is bounded above by $-\gamma$. In particular the flow becomes outgoing before it becomes steep again. Furthermore, if v is a vertex with indices i, j, k as above, such that the vector field $V_{ik}(\Phi(v))$ is not steep but is ingoing, then the additive relation insures that one of the other two flows $V_{ij}(Phi(v))$ or $V_{jk}(Phi(v))$ has slope less than or equal to the slope of $V_{ik}(\Phi(v))$. In reality, we draw a line through the first vector, and note that one of the two other vectors has to lie below or on the line. Note that this gives two cases: either the new vector changes direction (i.e. the sign $(-1)^m$ changes) and the new vector is in fact outgoing; or else the direction stays the same and the slope decreases. Thus if t_0 is any point in T where the flow is ingoing but not steep, then we can choose a strand σ below t_0 with the property that at the end of the strand the flow becomes outgoing; and along the strand the direction stays the same and the second derivative satisfies

$$\frac{d}{dx}((-1)^m \frac{dy}{dx}) \leq -\gamma$$

in a distributional sense. Then (noting by $x(t), y(t)$) the coordinates of the image point $\Phi(t)$ for $t \in \sigma$) we have

$$y(t) \leq y(t_0) + A(-1)^m(x(t) - x(t_0)) - \frac{\gamma}{2}(x(t) - x(t_0))^2$$

for any $t \geq t_0$. In particular there is a number N such that

$$y(t) \leq y(t_0) + N$$

further along the strand. We will choose $L_1 = L - N$.

Recall now that in the hypotheses of the lemma, we suppose that all strands in the tree remain inside $D_u(P)$ and also finish in $D_\epsilon(P)$. However, we construct above a strand which eventually becomes outgoing; therefore the strand must enter the region corresponding to $D_\epsilon(P)$ before it becomes outgoing (and notice also that it could simply stop inside this region before becoming outgoing, a case not mentionned above). In particular, if there is any point t_0 corresponding to a non-steep ingoing flow, or of course to any sort of outgoing flow, then we have to have $y(t_0) + N > L$ or $y(t_0) > L_1$.

Now we can complete the proof of the lemma. If v is any vertex, such that the incoming edge is steep and ingoing, then one of the two outgoing edges has to be either non-steep and ingoing, or outgoing. This is verified from the fact that at most two different vector fields can be attached as ingoing asymptotic vector fields for the same band B. From what was said above, the bottom vertex of the first edge e of the tree must satisfy $y(\Phi(v)) > L_1$, in other words the first edge continues all the way until $D_\xi(P)$. Also the part of the edge e which is outside of $D_\xi(P)$ must be contained in an ingoing asymptotic band for its vector field V_{ij} and the flow is steep at all points of $\Phi(e)$ which are outside of $D_\xi(P)$. This completes the proof of Lemma 4. ///

Proof of Lemma 5. Consider a vertex v along κ where a subtree in the normal form of Lemma 4 comes off. Use the same notation as previously for the edjes and indices adjoining v. For the sake of simplicity we assume that κ corresponds to the two leftmost edges e_1 and e_2 at v. The upper edge of the subtree is thus e_3 with indices jk.

Note from the proof of 4 that V_{jk} is ingoing and steep at $\Phi(v)$.

As a first case, note that if the anterior edge e_1 of κ has V_{ik} which is outgoing and steep, then the subsequent edge e_2 of κ is also outgoing and steep. In particular at any point where κ becomes outgoing and steep, it remains that way and in fact will leave the region $y > y_1$ before it goes into any other band B. By looking at the possible combinatorics of the indices one sees, even in the case of two vector fields sharing the same band, that there can be no further normal-form vertices on κ.

In view of the previous paragraph we may restrict our attention to the places where κ is either not steep, or else steep but ingoing. However, if it is steep but ingoing then again at most one vertex with a normal-form subtree can correspond to the current band; thus at some point κ leaves this band and must become non-steep. On the other hand, once κ is non-steep, it doesn't change to become steep and ingoing. It doesn't do this in the middle of an edge, because of the second derivative condition. It doesn't do it at a vertex because the edge e_3 which comes off is steep and ingoing, and a V_{ik} which is not steep couldn't be the sum of two steep and ingoing vectors.

¿From the two previous paragraphs then, we may (at the price of at most two extra normal-form subtrees) restrict our attention to the region where κ is non-steep. Now one sees again from the additive relation that if V_{ik} and V_{ij} are non-steep, whereas V_{jk} is steep and ingoing, then the directions of V_{ik} and V_{ij} must be the same. Indeed, if not then we would have

$$V_{jk} = V_{ik} + (-V_{ij})$$

which would be a sum of two vectors in the same non-steep quadrant, so V_{jk} in a steep quadrant would be impossible.

Since the sign $(-1)^m$ of $\frac{dx}{dt}$ doesn't change, we can use x to parametrize κ. Furthermore the derivative $(-1)^m\frac{dy}{dx}$ is decreasing along κ (note that at any vertices where a subtree in normal form comes off, the remaining outgoing edge of κ has a bigger derivative than the ingoing edge, because of the additive relation).

In other words, the second derivative is distributionally less than the constant $-\gamma$, so at some time t with $|x(t)-x_0| \leq 2A/\gamma$ we get to $(-1)^m\frac{dy}{dx} \leq -A$, i.e. κ becomes steep and outgoing. We get that the non-steep part of the path κ is parametrized by an interval in the x-coordinate, of length $\leq 2A/\gamma$. There is a bound K so that such an interval can cross (or go near) at most $K-2$ asymptotic bands. A band is attached to at most two pairs of indices, but only one of these can lead correspond to a normal-form subtree. Thus (counting the two we may have missed above) the number of normal-form subtrees attached to γ is $\leq K$. This completes the proof of Lemma 5.
///

We now come to the statement of our theorem. Fix u, w as above, and let $L_1 := L + N$ be the function determined by the above proofs. For any ϵ put $L := \log \epsilon$ and set $\xi := e^{L_1} = e^N \epsilon$. Note that $\xi \to 0$ as $\epsilon \to 0$.

Theorem 6 *There is a bound K depending on $u, w, \mathbf{B}$ but independent of ϵ, ξ with the following properties. Suppose T is a tree and $\Phi : T \to X$ is a flowing map such that the top vertices are outside of any $D_w(P_i)$ and such that the bottom vertices are each mapped into some $D_\epsilon(P_i)$. Suppose furthermore that $g(\Phi) \geq g(\text{top}) - \mathbf{B}$. Then we can cut T into a tree T' onto which are attached subtrees, such that Φ maps the bottom vertices of T' into various $D_\xi(P_i)$ and such that the number of bottom vertices of T' is bounded by K.*

Proof: Among the subtrees that we strip off are any ones starting with edges e for which $i_e = j_e$. In particular we may assume from the start that T has no such edges.

Next group the bottom vertices into series connected by intervals where the bounding loop of the interval is mapped into $D_u(P)$. There is a bound K_1 for the number of such series, because any loop which goes out of $D_u(P)$ has to contribute at least a certain fixed amount to $g(\mathrm{top}) - g(\Phi)$. Next we can look at a specific series. It is the set of bottom vertices of a subtree T_1 obtained by taking the union of all of the loops joining the bottom vertices together. Note in particular that $\Phi(T_1) \subset D_u(P)$. Let κ denote the boundary path of T_1. Note that the subtree T_1 doesn't necessarily include all strands emanating from all of its vertices. However, if v is a vertex on κ corresponding to an adjoining edge e not in κ, then either e goes into the interior of the region bounded by κ, in which case e starts a subtree mapped into $D_u(P)$ and such that all bottom edges go into $D_\epsilon(P)$; or else it goes out of the region bounded by κ in which case e is not a part of the tree T_1. In the former case, the normal form of Lemma 4 applies to the subtree starting at e. In the latter case, the subtree starting at e could be in normal form or not. However, if e is an edge going out of κ such that the subtree starting at e is not in the normal form of Lemma 4, then this subtree contains at least one strand which goes out of $D_u(P)$. By Lemma 3 it also contains a strand which goes out of $D_w(P)$ and there is a global bound K_2 on the number of such edges e. If we cut κ at vertices v where such edges e go out, then it is cut into $\leq K_2$ strands κ' and each little strand has only vertices corresponding to normal-form subtrees. Finally, by the bound of Lemma 5 there are no more than K_3 such vertices on each little strand κ'. Each of these normal-form subtrees can be cut at the point where it goes into $D_\xi(P)$, and there is only one such point for each subtree. Thus if we trim off the tree T_1 at all of the points where the strands enter $D_\xi(P)$, there are at most K_2K_3 bottom vertices. Finally, since there were at most K_1 subtrees T_1 corresponding to series of bottom vertices, we can trim off T to a tree T' where there are at most $K_1K_2K_3$ bottom vertices, all going inside some $D_\xi(P_i)$. This proves the theorem.
///

Proof of Theorem 2

By a *singular point* we mean a point $y = (y_1, \ldots, y_k) \in Z_I$ such that the y_n are singular points of the functions $g_{i_n i_{n+1}}$. Note that (in our special situation coming originally from a general Higgs bundle (E, θ)) the singular points are all points denoted P above, lying over singularities of the original Higgs field

θ.

Recall that we say that a point $z \in Z_I$ is *beyond* p *at distance* $\leq M$ if there is a tree embedded in Z with one top vertex at p, with the bottom vertices at the y_i, and with total length $\leq M$ measured in the (Finsler) metric defined by the differentials dg_{ij} (see Chapter 6 of [1]). The path γ may be considered as a tree with one spine and lots of edges of length zero coming off at the various points. Thus if M_0 is the length of γ in the Finsler metric, the points of the original chain of integration η are beyond p at distance $\leq M_0$.

Define S_M to be the set of complex values of the form $g(y)$ where y are singular points which are beyond p at distance $\leq M_0+2M$. This is the subset which is to enter into the definition of analytic continuation with locally finite branching for $f(\zeta)$.

Suppose we have already analytically continued f along a piecewise linear path of length $\leq M_1$. Inductively we may assume that the points of η are beyond p at distance $\leq M_0 + 2M_1$. If we add a segment S then the total length is $\leq M$ where $M = M_1 + |S|$.

Fix our neighborhoods N_i and a number ν so that the neighborhood N_1 of the segment S stays at a distance $> \nu$ away from the points of S_M. Let K be the bound of Theorem 6. Choose ϵ small enough so that if $z \in D_\xi(P)$ then for any i

$$\left| \int_P^z \alpha_i \right| < \frac{\nu}{K}.$$

We show that all points of the chain $F\tau$ are sent (by g) outside of N_3. Suppose on the contrary that we had a point, corresponding to a tree T, such that $g(z^{\mathrm{bot}}(T)) \in N_3$.

By Theorem 6 there exists a pruning T' of T with $\leq K$ bottom vertices, such that for every bottom vertex v of T' we have $\Phi(v) \in D_\xi(P(v))$ for some singular point $P(v)$. In particular, the point $z^{\mathrm{bot}}(T')$ which is the vector of these $\Phi(v)$ is near to a point $y = (\ldots, P(v), \ldots)$. More precisely we obtain from $k \leq K$ and the bound above,

$$|g(y) - g(z^{\mathrm{bot}}(T'))| < \nu.$$

On the other hand, if $g(z^{\mathrm{bot}}(T'))$ were inside N_3 then the singular point y would occur below points of η at distance $\leq |S|$, and hence below points of p at distance $\leq M_0 + 2M_1 + |S| < M$, therefore $g(y)$ must be included in S_M. On the other hand, the point $g(z^{\mathrm{bot}}(T'))$ occurs on the real segment between $g(z^{\mathrm{bot}}(T))$ and some point of $g(\eta')$. This contradicts the assumption that the neighborhood N_1 stays away from S_M by distance at least ν. This shows that all points of $g(F\tau)$ are outside of N_3, and completes the proof that we can analytically continue $f(\zeta)$ along the segment S.

Finally in order to maintain the inductive hypothesis we note that, cutting everything off fairly close to the segment S we can insure that the points of the new cycle of integration $F\tau$ (and also $F\psi$) are beyond points of η at distance $\leq 2|S|$, hence they are beyond p at distance $\leq M_0 + 2M$ as required. This completes the proof of Theorem 2

Conclusion

We close with a few more general remarks about the consequences of Theorem 2. First recall that a product of functions $m_1(t)m_2(t)$ is transformed into a convolution of their Laplace transforms. Furthermore, if f_1 and f_2 have analytic continuations with locally finite branching, then the same is true for their convolution $f_1 * f_2$. This is proved in [1], proof of Lemma 11.1 (noting that the proof of locally finite branching for the convolutions uses only locally finite branching for the two functions). Thus it follows from Theorem 2 that if $P(t)$ is any polynomial in transport matrices for various paths, then the Laplace transform of $P(t)$ has an analytic continuation with locally finite branching.

The next remark is about whether the singularities of the Laplace transform disappear or not. In the course of the proof, we describe a fairly explicit family of sets of singular points S_M, however it is possible that $f(\zeta)$ has a holomorphic continuation across some of these points. One can note that it is not possible that *all* of the singularities disappear, unless $m(t)$ is identically zero, because of the inverse Laplace transform. There remains the possibility that $f(\zeta)$ could have a single singular point, say at $\zeta = 0$. This would correspond to the case where $m(t)$ has sub-exponential growth.

We say that m is *sub-exponential* if f has a single singularity at the origin, and *exponential* otherwise. The exponential case could involve only two singularities. Thus we say that m is *strictly exponential* if the convex hull of the singularities of f contains a nonempty interior. In the strictly exponential case we obtain (possibly after multiplying by a pure exponential to translate f so that the origin is in the interior) a bound of the form $|m(t)| \geq ce^{a|t|}$. This might not be the case in the non strictly exponential case. However, even in the non-strictly exponential case, for sectors covering all but two directions we have such a bound, and in particular there is a positive lower bound for the possible exponents a which can enter into bounds of the form $|m(t)| \leq Ce^{a|t|}$.

One can prove that the monodromy representation m of a general family of connexions is exponential (i.e. the Laplace transform has at least two singularities). This is done by specializing to the case treated in [1]. That case, of a family of connexions on the trivial bundle of the form $d + B + tA$ with A diagonal and B off-diagonal, everything being holmorphic on X, leads to asymptotic expansions whose coefficients can be calculated. One route is to note that for generic values of A and B, calculation of the coefficients gives nonzero coefficients at more than one singular point. Another route would be to note that if there were only one singularity for the monodromy matrices for this family, then the monodromy representation would actually have polynomial growth. Also if there were two singularities then for a sequence of values of $t \to \infty$ the monodromy matrices would have polynomial growth. Those possibilities are ruled out by specializing again to a direct sum of a 2×2 system and trivial systems, and noting that for 2×2 systems we have proven (in the paper [2]) that the monodromy representation always has growth at least $e^{t^{1/k}}$ for some integer k.

In any case by either of these two routes we can conclude that the Laplace transform for at least one monodromy matrix in the special system (and also for at least one Procesi coordinate of the corresponding point in M_B) has at least two singularities. Now a Hartogs argument shows that this must also be the case for the Laplace transform for any of a dense set of generic families of connexions $(E, \nabla + t\theta)$. This can be stated as follows.

Theorem 7 *For each family $(E, \nabla + t\theta)$ going to infinity at a generic Higgs bundle (E, θ), let ρ_t denote the family of monodromy representations, thought of as a point in M_B, and measure the size of ρ_t by the Procesi coordinates. Say that ρ_t is* exponential *if its Laplace transform has at least two singularities, which also means that for at least some sequences of points $t \to \infty$ we have a positive a with*

$$|\rho_t| \geq ce^{a|t|}.$$

The set of Higgs bundles (E, θ) for which, for at least one ∇, the monodromy representation is exponential, is dense in the usual topology of the divisor at infinity in $\overline{M}_{DR}$.

This result improves, at least for certain generic points at infinity approached from certain sectors, the bound given in [2].

Finally it is important to reiterate that, in spite of these consequences, the result of Theorem 2 is highly unsatisfactory in that it doesn't say anything about the behavior of the Laplace transform $f(\zeta)$ near the singularities. It doesn't even seem clear what the answer will be here: on the one hand one can imagine that an improvement of the present analysis, potentially based on Remark (i) following the proof of Lemma 1, might lead to a polynomial bound for the singularities. On the other hand, a crude look at the present argument yields no such bound, and it is also quite concievable that the poles in the matrix B lead unavoidably to more complicated singularities of $f(\zeta)$. This is undoubtedly true in the general case where B has poles of order > 1.

This problem also leads to the unsatisfactory statement of Theorem 7: if we could calculate exactly where the singularities were we could probably show that for generic values of (E, θ) the singularities would span a convex hull with nonempty interior, in other words that the monodromy families ρ_t would be strictly exponential. This would be a more significant improvement of the result of [2].

The result of Theorem 2 should be thought of as a weak form of "resurgence" for the monodromy function $m(t)$ and its Laplace transform. The problem of getting more precise information about this behaviour is probably most naturally attacked using new ideas and techniques such as have been developped by the school of J.-P. Ramis.

References

[1] C. Simpson. *Asymptotic Behavior of Monodromy*, **L. N. M. 1502**, Springer (1991).

[2] C. Simpson. A lower bound for the size of monodromy of systems of ordinary differential equations. *Algebraic Geometry and Analytic Geometry—ICM-90 Satellite Conference Proceedings*, Springer (1991), 198-230.