ANNALES

DE

L'INSTITUT FOURIER

Simon DONALDSON

An elliptic boundary value problem for G_2-structures

AN ELLIPTIC BOUNDARY VALUE PROBLEM FOR G_2-STRUCTURES

by Simon DONALDSON

Dedicated to Jean-Pierre Demailly, for his 60th birthday

————

ABSTRACT. — We show that the G_2 holonomy equation on a seven-dimensional manifold with boundary, with prescribed 3-form on the boundary and modulo the action of diffeomorphisms, is elliptic. The main point is to set up a suitable linear elliptic theory. This result leads to a deformation theory, governed by a finite-dimensional obstruction space. We discuss conditions under which this obstruction space vanishes and as one application we establish the existence of certain G_2 cobordisms between two small deformations of a Calabi–Yau 3-fold.

RÉSUMÉ. — Nous montrons que l'équation d'holonomie G_2 sur une variété de dimension 7 à bord, avec 3-forme prescrite sur le bord et modulo l'action de difféomorphismes, est elliptique. Le point clé est de mettre en place une théorie linéaire elliptique adaptée. Avec ce résultat une théorie de la déformation est définie, gouvernée par un espace d'obstruction de dimension finie. Nous discutons les conditions pour lesquelles cet espace d'obstruction est trivial, et donnons une application en démontrant l'existence de certains G_2-cobordismes entre deux petites déformations d'une variété Calabi–Yau de dimension 3.

1. Introduction

The purpose of this paper is to develop a deformation theory for torsion-free G_2-structures on 7-manifolds with boundary. This extends the well-established theory for closed manifolds, going back to Bryant and Harvey (see [1, p. 561]) and further developed by Joyce [9, 10] and Hitchin [7, 8]. Recall that a torsion-free G_2-structure on an oriented 7-manifold M can be viewed as a closed 3-form ϕ which is "positive" (in a sense we recall below) at each point of M and which satisfies the nonlinear equation

$$(1.1) \qquad \qquad \mathrm{d} *_\phi \phi = 0,$$

————

where $*_\phi$ is the $*$-operator of the Riemannian metric g_ϕ defined by ϕ (which we also recall below). Begin with the standard case when M is a closed manifold and let c be a class in $H^3(M; \mathbf{R})$. Write $\mathcal{P}_c$ for the set of positive 3-forms representing c. Certainly a torsion-free structure ϕ defines a point in $\mathcal{P}_c$ with $c = [\phi]$. Conversely, if we have a c such that $\mathcal{P}_c$ is not empty then, as Hitchin observed, solutions of equation (1.1) in $\mathcal{P}_c$ correspond to critical points of the volume functional

$$(1.2) \qquad\qquad V(\phi) = \mathrm{Vol}(M, g_\phi)$$

on $\mathcal{P}_c$. In fact $\mathrm{d}(*_\phi\phi)$ can be regarded as the derivative dV of the volume functional on the infinite-dimensional space $\mathcal{P}_c$. The basic results of the standard theory can be summarised as follows.

(1) The derivative dV is a Fredholm section of the cotangent bundle of the quotient $\mathcal{Q}_c = \mathcal{P}_c/\mathcal{G}$ of $\mathcal{P}_c$ by the group $\mathcal{G}$ of diffeomorphisms of M isotopic to the identity. Thus the kernel of the Hessian of the volume functional on $\mathcal{Q}_c$ at a solution of (1.1) is finite-dimensional.

(2) In fact this kernel is always 0, which implies that if ϕ is a solution of (1.1) and if c' is sufficiently close to $c = [\phi]$ in $H^3(M)$ then there is a unique solution ϕ' in $\mathcal{Q}_{c'}$ close to ϕ. (Throughout this paper, cohomology is always taken with real coefficients.) In other words the "period map" $\phi \mapsto [\phi]$ defines a local homeomorphism from the moduli space of torsion-free G_2-structures to $H^3(M)$.

(3) In fact the Hessian of V on $\mathcal{Q}_c$ is negative-definite. A solution of (1.1) gives a strict local maximum for the volume functional on $\mathcal{Q}_c$.

Now we go on to the case of a compact, connected, oriented manifold M with non-empty boundary ∂M. If ρ is a closed 3-form on ∂M we define an *enhancement* of ρ to be an equivalence class of closed forms ϕ on M which restrict to ρ on the boundary, under the equivalence relation $\phi \sim \phi + \mathrm{d}\alpha$ for all 2-forms α which vanish on ∂M. So the set of enhancements is an affine space with tangent space $H^3(M, \partial M)$. There is an algebraic notion of a positive 3-form on ∂M. One definition is that these are exactly the forms which extend to positive forms on some neighbourhood of ∂M in M. Fix a closed positive form ρ on ∂M and enhancement $\hat{\rho}$. We write $\mathcal{P}_{\hat{\rho}}$ for the set of positive forms in the enhancement class (in general, $\mathcal{P}_{\hat{\rho}}$ could be the empty set) and $\mathcal{Q}_{\hat{\rho}}$ for the quotient by the identity component of the group of diffeomorphisms of M fixing the boundary pointwise. The boundary value problem, which was introduced in [5] and which we consider further here,

is to solve equation (1.1) for ϕ in $\mathcal{P}_{\hat{\rho}}$. Just as before, this is the Euler–Lagrange equation for the Hitchin volume functional, which descends to a functional on $\mathcal{Q}_{\hat{\rho}}$.

The author has only been able to extend the first of the three results from the standard theory above to this setting. That is (continuing the informal discussion, more precise technical statements are given later), we will show below (Proposition 4.5) that the derivative of the volume functional is a Fredholm section of the cotangent bundle of $\mathcal{Q}_{\hat{\rho}}$. This comes down to showing that our problem can be set up as an elliptic boundary value problem. The crucial linear result is Theorem 3.5. The kernel of the Hessian at a critical point ϕ is a finite dimensional vector space H_ϕ but this is not 0 in general. Similarly, we can show that the Hessian has finite index (i.e. a finite dimensional negative subspace) but we have not been able to show that the Hessian is semi-definite. We will discuss these questions at greater length in Section 5 below. In any event, we do know cases in which the space H_ϕ is zero and in such cases we get a straightforward deformation theory for our problem: for any enhanced boundary data sufficiently close to $\hat{\rho}$ there is a unique solution to the corresponding boundary value problem close to ϕ (Theorem 4.6). In Section 5 we give one application to the existence of "G_2-cobordisms" between closed 3-forms on a Calabi–Yau 3-fold (Theorem 5.8).

The authors's work is supported by the Simons Collaboration Grant "Special holonomy in Geometry, Analysis and Physics".

2. Review of standard theory

We begin with some purely algebraic statements.

- A 3-form $\phi \in \Lambda^3(V^*)$ on an oriented 7-dimensional real vector space V is called positive if the $\Lambda^7 V^*$-valued quadratic form on V

$$(2.1) \qquad v \mapsto i_v(\phi) \wedge i_v(\phi) \wedge \phi$$

  is positive definite. We fix a Euclidean structure g_ϕ in this conformal class by normalising so that $|\phi|^2 = 7$. Then, as in the Introduction, we have a 4-form $*_\phi \phi$ which we also write as $\Theta(\phi)$. So Θ is a smooth map from the space of positive 3-forms on V to $\Lambda^4 V^*$. The positive 3-forms on V form a single orbit under the action of $GL^+(V)$, so they are all equivalent. A convenient standard model for this paper is to take $V = \mathbf{R} \oplus \mathbf{C}^3 = \{(t, z_1, z_2, z_3)\}$ and

$$(2.2) \qquad \phi = \omega \wedge \mathrm{d}t + \mathrm{Im}(\mathrm{d}z_1 \mathrm{d}z_2 \mathrm{d}z_3),$$

  where ω is the standard symplectic form $\sum dx_a \wedge dy_a$ on $\mathbf{C}^3$.

- The stabiliser in $GL(V)$ of a positive 3-form is isomorphic to the exceptional Lie group G_2. Under the action of this group the forms decompose as

(2.3) $$\Lambda^2 = \Lambda^2_7 \oplus \Lambda^2_{14} \quad \Lambda^3 = \Lambda^3_1 \oplus \Lambda^3_7 \oplus \Lambda^3_{27}.$$

Here Λ^2_7 is the image of V under the map $v \mapsto i_v(\phi)$ and Λ^2_{14} is the orthogonal complement; Λ^3_7 is the image of V under the map $v \mapsto i_v(*_\phi\phi)$, Λ^3_1 is the span of ϕ and Λ^3_{27} is the orthogonal complement of their sum.

We have a quadratic form on Λ^2 defined by $\alpha \mapsto \alpha \wedge \alpha \wedge \phi$. The eigenspaces of this form, relative to the standard Euclidean structure, are $\Lambda^2_7, \Lambda^2_{14}$. For $\alpha_7 \in \Lambda^2_7$

(2.4) $$\alpha_7 \wedge \alpha_7 \wedge \phi = 2|\alpha_7|^2 \, \mathrm{vol},$$

and for $\alpha_{14} \in \Lambda^2_{14}$

(2.5) $$\alpha_{14} \wedge \alpha_{14} \wedge \phi = -|\alpha_{14}|^2 \, \mathrm{vol}\,.$$

- The volume form vol is a Λ^7-valued function on the open set of positive 3-forms. Its derivative is given by

(2.6) $$\mathrm{vol}(\phi + \delta\phi) = \mathrm{vol}(\phi) + \frac{1}{3}\delta\phi \wedge \Theta(\phi) + O(\delta\phi^2).$$

To identify the second derivative we write

$$\delta\phi = \delta_1\phi + \delta_7\phi + \delta_{27}\phi,$$

according to the decomposition (2.3). Then

$$\mathrm{vol}(\phi + \delta\phi) = \mathrm{vol}(\phi) + \frac{1}{3}\delta\phi \wedge (*_\phi\phi) + \frac{2}{3}q(\delta\phi)\,\mathrm{vol}(\phi) + O(\delta\phi^3),$$

where q is the quadratic form

(2.7) $$q(\delta\phi) = \frac{4}{3}|\delta_1\phi|^2 + |\delta_7\phi|^2 - |\delta_{27}|^2.$$

This formula also gives the derivative of the map Θ ([10, Proposition 10.3.5]):

(2.8) $$\Theta(\phi + \delta\phi) = \Theta(\phi) + \left(\frac{4}{3}*_\phi \delta_1\phi + *_\phi\delta_7\phi - *_\phi\delta_{27}\phi\right) + O(\delta\phi^2).$$

Now let M be an oriented 7-manifold and ϕ be a positive 3-form on M which defines a torsion-free G_2-structure, so both ϕ and $*_\phi\phi$ are closed forms. We can decompose the exterior derivative according to the decomposition of the forms

$$\Omega^2 = \Omega^2_7 \oplus \Omega^2_{14} \quad \Omega^3 = \Omega^3_1 \oplus \Omega^3_7 \oplus \Omega^3_{27}.$$

The resulting operators satisfy various identities, akin to the Kähler identities on Kähler manifolds. The following Proposition states the main identities we will need in this paper (there is a comprehensive treatment in [2]). Write $\chi : \Lambda^1 \to \Lambda_7^3$ for the bundle isomorphism $\chi(\eta) = *_\phi(\eta \wedge \phi)$. We also usually write $*$ for $*_\phi$ and d^* for the usual adjoint constructed using the metric g_ϕ.

PROPOSITION 2.1.

(1) The component $\mathrm{d}_1 : \Omega_{14}^2 \to \Omega_1^3$ is identically zero.

(2) The component $\mathrm{d}_7 : \Omega_{14}^2 \to \Omega_7^3$ is equal to the composite $\frac{1}{4}\chi \circ \mathrm{d}^*$ where $\mathrm{d}^* : \Omega_{14}^2 \to \Omega^1$, and the component $\mathrm{d}_7 : \Omega_7^2 \to \Omega_7^3$ is equal to the composite $-\frac{1}{2}\chi \circ \mathrm{d}^*$ where $\mathrm{d}^* : \Omega_7^2 \to \Omega^1$.

(3) For $\mathrm{d}_7 : \Omega^1 \to \Omega_7^2$ and $\mathrm{d}_{14} : \Omega^1 \to \Omega_{14}^2$ we have

$$\mathrm{d}^*\mathrm{d}_{14} = 2\mathrm{d}^*\mathrm{d}_7 = \frac{2}{3}\mathrm{d}^*\mathrm{d}$$

on Ω^1.

Proof. — For the first item, it suffices to prove that for a compactly supported $\alpha \in \Omega_{14}^2$ and function f the L^2-inner product $\langle \mathrm{d}\alpha, f\phi \rangle$ is zero. This inner product is

$$\int_M \mathrm{d}\alpha \wedge f * \phi = -\int_M \alpha \wedge \mathrm{d}f \wedge *\phi,$$

(using $\mathrm{d} * \phi = 0$) which vanishes since $df \wedge *\phi$ lies in Ω_7^5. For the second item we consider first an $\alpha \in \Omega_{14}^2$ as above and the inner product $\langle \mathrm{d}^*\alpha, \eta \rangle$ for a 1-form η. By definition this is $\langle \alpha, \mathrm{d}\eta \rangle$ and by (2.5) the latter can be expressed as

$$-\int_M \alpha \wedge \mathrm{d}\eta \wedge \phi.$$

By Stokes' Theorem (using, this time, $\mathrm{d}\phi = 0$) this is

$$\int_M \mathrm{d}\alpha \wedge \eta \wedge \phi = \langle \mathrm{d}\alpha, *(\eta \wedge \phi) \rangle = \langle \mathrm{d}_7\alpha, \chi(\eta) \rangle.$$

One computes readily that for any η we have

$$|\chi(\eta)|^2 = 4|\eta|^2,$$

and it follows that $\mathrm{d}_7\alpha = \frac{1}{4}\chi \circ \mathrm{d}^*\alpha$. The argument for the second part of the second item (for $\alpha \in \Omega_7^2$) is the same using (2.4).

For the third item: the equality $\mathrm{d}^*\mathrm{d}_{14} = 2\mathrm{d}^*\mathrm{d}_7$ follows from the second item and the fact that the component of d^2 from Ω^1 to Ω_7^3 is zero. The equality $\mathrm{d}^*\mathrm{d}_{14} = \frac{2}{3}\mathrm{d}^*\mathrm{d}$ follows in turn because $\mathrm{d}^*\mathrm{d} = \mathrm{d}^*\mathrm{d}_7 + \mathrm{d}^*\mathrm{d}_{14}$.   □

The variation of the volume functional (1.2) with respect to compactly supported variations of ϕ makes sense, even if M is not compact. Suppose for the moment that ϕ is any closed positive 3-form on M and that α is a 2-form with compact support. The pointwise formula (2.6) and integration by parts give

$$\mathrm{Vol}(\phi + \mathrm{d}\alpha) = \mathrm{Vol}(\phi) - \frac{1}{3}\int_M \alpha \wedge \mathrm{d}\Theta(\phi) + O(\alpha^2),$$

which shows that the torsion-free condition $\mathrm{d}\Theta(\phi) = 0$ is the Euler–Lagrange equation associated to the volume functional for exact variations. For any such ϕ we have

(2.9) $$*_\phi \, \mathrm{d}\Theta(\phi) \in \Omega^2_{14}(M).$$

This follows by direct calculation or, more conceptually, from the diffeomorphism invariance of the volume functional (see [4, Lemma 1]).

Now go back to assuming that ϕ defines a torsion-free G_2-structure, i.e. $\mathrm{d}\Theta(\phi) = 0$. For α in $\Omega^2(M)$ we define

(2.10) $$W(\alpha) = *_\phi \mathrm{d}\Theta(\phi + \mathrm{d}\alpha)$$

so the equation $W(\alpha) = 0$ is the torsion-free equation, for such variations. For any α we have $\mathrm{d}^* W(\alpha) = 0$ and $W(\alpha)$ takes values in the sub-bundle

$$*_\phi \Lambda^5_{14,\phi+\mathrm{d}\alpha} \subset \Lambda^2,$$

in an obvious notation. Let L be the linearisation of the nonlinear operator W at $\alpha = 0$, i.e. $W(\alpha) = L(\alpha) + O(\alpha^2)$. By (2.8) this linearised operator is given by the formula

(2.11) $$L(\alpha) = \frac{4}{3}\mathrm{d}^*\mathrm{d}_1\alpha + \mathrm{d}^*\mathrm{d}_7\alpha - \mathrm{d}^*\mathrm{d}_{27}\alpha.$$

PROPOSITION 2.2. — *The linear operator L vanishes on Ω^2_7 and takes values in Ω^2_{14}. For $\alpha = \alpha_7 + \alpha_{14}$ we have*

$$L(\alpha) = \mathrm{d}^*\mathrm{d}_7\alpha_{14} - \mathrm{d}^*\mathrm{d}_{27}\alpha_{14} = -\Delta\alpha_{14} + \frac{3}{2}\mathrm{d}_{14}\mathrm{d}^*\alpha_{14}$$

Proof. — The fact that L vanishes on Ω^2_7 follows from diffeomorphism invariance (or by direct calculation). Similarly, the fact that L takes values in Ω^2_{14} is a consequence of the fact above that $W(\alpha)$ is a section of $*_\phi\Lambda^5_{7,\phi+\mathrm{d}\alpha}$ (or can be shown by direct calculation). The formulae for $L(\alpha)$ follow from items (1) and (3) in Proposition 2.1.   □

There is a similar discussion for the Hessian of the volume functional. For α of compact support $\mathrm{Vol}(\phi + \mathrm{d}\alpha) = \mathrm{Vol}(\phi) + \frac{2}{3}Q(\alpha)$ where

(2.12) $$Q(\alpha) = \frac{4}{3}\|\mathrm{d}_1\alpha\|^2 + \|\mathrm{d}_7\alpha\|^2 - \|\mathrm{d}_{27}\alpha\|^2.$$

This can also be expressed, for $\alpha = \alpha_7 + \alpha_{14}$, as

$$(2.13) \qquad Q(\alpha) = \|d_7\alpha_{14}\|^2 - \|d_{27}\alpha_{14}\|^2 = \langle L\alpha_{14}, \alpha_{14}\rangle.$$

But we should emphasise that (2.12) is the "primary" formula (derived pointwise on M) and the passage to the expressions in (2.13) involves an application of Stokes' Theorem.

We will now sketch a treatment of the standard results for closed manifolds mentioned in the Introduction. In this sketch we will just give a formal treatment, ignoring the analytical aspects, but these will be taken up in a more general setting in Section 4. To avoid unimportant complications we suppose here that $H^2(M) = 0$.

With $c = [\phi] \in H^3(M)$, the tangent space of $\mathcal{P}_c$ at ϕ is

$$T\mathcal{P}_c = \mathrm{Im}(\mathrm{d} : \Omega^2 \to \Omega^3).$$

The infinitesimal action of the group $\mathcal{G}$ of diffeomorphisms of M is by the Lie derivative. But, since ϕ is closed, for a vector field v we have

$$\mathcal{L}_v\phi = \mathrm{d}(i_v(\phi))$$

and the 2-forms $i_v(\phi)$ are exactly Ω_7^2. So the tangent space of $\mathcal{Q}_c$ at ϕ is

$$T\mathcal{Q}_c = \frac{\mathrm{Im}(\mathrm{d} : \Omega^2(M) \to \Omega^3(M))}{\mathrm{d}\Omega_7^2}.$$

Let $\pi : \Omega_{14}^2 \to T\mathcal{Q}_c$ be the map induced by exterior derivative. This is obviously surjective and the kernel consists of those $\alpha_{14} \in \Omega_{14}^2$ such that there is an $\alpha_7 \in \Omega_7^2$ with $\mathrm{d}\alpha_7 = \mathrm{d}\alpha_{14}$. Under our assumption that $H^2(M) = 0$ this means that $\alpha_{14} - \alpha_7 = \mathrm{d}\eta$ for some $\eta \in \Omega^1$, so $\alpha_{14} = \mathrm{d}_{14}\eta$. Conversely, if $\alpha_{14} = \mathrm{d}_{14}\eta$ we can define $\alpha_7 = -\mathrm{d}_7\eta$. So we see that the kernel of π is the image of $\mathrm{d}_{14} : \Omega^1 \to \Omega_{14}^2$ and

$$T\mathcal{Q}_c = \Omega_{14}^2 / \mathrm{Im}\,\mathrm{d}_{14}.$$

By standard elliptic theory this can be identified with the kernel of the adjoint:

$$T\mathcal{Q}_c = \ker \mathrm{d}^* : \Omega_{14}^2 \to \Omega^1.$$

Now by items (1) and (2) of Proposition 2.1, for $\alpha \in \ker \mathrm{d}^* \subset \Omega_{14}^2$ the only component of $\mathrm{d}\alpha$ is $\mathrm{d}_{27}\alpha \in \Omega_{27}^3$. It follows that for such α the linearised operator $L(\alpha)$ is $-\Delta\alpha$. In other words, after taking account of the diffeomorphism group action in this way, the linearised operator is

$$-\Delta : \left(\ker \mathrm{d}^* \subset \Omega_{14}^2\right) \to \left(\ker \mathrm{d}^* \subset \Omega_{14}^2\right),$$

which is invertible. Similarly, with this representation of $T\mathcal{Q}_c$ the Hessian is

$$Q(\alpha) = -\|\mathrm{d}\alpha\|^2,$$

which is negative definite.

3. The boundary value problem–linear theory

This section represents the heart of this paper, in which we set up a linear elliptic boundary value problem. We suppose, as in the Introduction, that M is a compact, connected, oriented 7-manifold with non-empty boundary and that the 3-form ϕ defines a torsion-free G_2-structure on M. We will give a representation of the tangent space of $\mathcal{Q}_{\hat{\rho}}$ at ϕ for the enhanced boundary value $\hat{\rho}$ determined by ϕ and study the linearisation of the torsion-free equation in this representation.

As a preliminary, note that there are two different notions of "restriction to the boundary" of a p-form σ on M. One is that the pull back under the inclusion map, an element of $\Omega^p(\partial M)$, vanishes. We will denote this by the usual notation $\sigma|_{\partial M} = 0$. For the other, stronger, notion we mean that the restriction vanishes regarded as a section of the bundle $\Lambda^p T^* M|_{\partial M}$. We will denote this by the notation $\sigma\|_{\partial M} = 0$.

To begin we define a vector space

$$(3.1) \qquad T\mathcal{Q} = \frac{\{d\gamma : \gamma \in \Omega^2(M), \gamma|_{\partial M} = 0\}}{\{d\beta : \beta \in \Omega_7^2(M), \beta\|_{\partial M} = 0\}}.$$

The definition of our space $\mathcal{P}_{\hat{\rho}}$ and the identification of the vector fields on M with Λ_7^2 suggests that this should represent the tangent space of the infinite dimensional manifold $\mathcal{Q}_{\hat{\rho}}$ but we postpone any precise treatment of this for the present and just take (3.1) as a definition. Similarly we do not at this stage consider any topology on the vector space. In the same vein, we define a vector space $H_\phi \subset T\mathcal{Q}$ to be

$$(3.2) \qquad H_\phi = \frac{\{d\gamma : \gamma \in \Omega^2(M), L(\gamma) = 0, \gamma|_{\partial M} = 0\}}{\{d\beta : \beta \in \Omega_7^2(M), \beta\|_{\partial M} = 0\}}.$$

This is the space of solutions of the linearised equations modulo infinitesimal diffeomorphisms. Note that on the right hand side of (3.2) the denominator is a subspace of the numerator since L vanishes on Ω_7^2.

We now discuss the linear algebra of the decomposition $\Lambda_7^2 \oplus \Lambda_{14}^2$ on the boundary. We have a 3-form $\rho = \phi|_{\partial M}$ and a 2-form $\omega \in \Omega^2(\partial M)$ given by $\omega = i_\nu(\phi)$ where ν is a unit outward-pointing normal. At a point p on ∂M the situation corresponds to the model (2.2) on $\mathbf{R} \oplus \mathbf{C}^3$. There is a complex structure on the tangent space of ∂M at p; the 2-form ω is a

positive $(1,1)$ form and ρ is the real part of a complex volume form. In terms of the splitting $TM = T\partial M \oplus \mathbf{R}\nu$ at p the form ϕ is

$$\phi = \omega \wedge \nu^* + \rho,$$

where ν^* is the 1-form dual to ν. We define a bundle map

$$\chi_6 : \Lambda^1(\partial M) \to \Lambda^2(\partial M)$$

by

$$\chi_6(i_v\omega) = i_v(\rho),$$

for $v \in T\partial M$. We have a decomposition

(3.3) $$\Lambda^2(\partial M) = \Lambda_6^{2,\partial} \oplus \Lambda_8^{2,\partial} \oplus \Lambda_1^{2,\partial}$$

where, in terms of the complex structure, the summand $\Lambda_6^{2,\partial}$ consists of the real parts of forms of type $(2,0)$, the summand $\Lambda_8^{2,\partial}$ consists of the real $(1,1)$ forms orthogonal to ω and $\Lambda_1^{2,\partial}$ is the 1-dimensional space spanned by ω. Then χ_6 is a bundle isomorphism from $\Lambda^1(\partial M)$ to $\Lambda_6^{2,\partial}$.

LEMMA 3.1. — *At a boundary point:*
  (1) $\Lambda_7^2 = \Lambda_1^{2,\partial} \oplus \{a \wedge \nu^* + \chi_6(a) : a \in \Lambda^1(\partial M)\}$,
  (2) $\Lambda_{14}^2 = \Lambda_8^{2,\partial} \oplus \{2a \wedge \nu^* - \chi_6(a) : a \in \Lambda^1(\partial M)\}$.

This is straightforward to check, from the definitions. For a form $\alpha \in \Omega_{14}^2(M)$ we write $\alpha\|_{\partial,8}$ for the section of $\Lambda_8^{2,\partial}$ over ∂M defined by the decomposition in the second item of Lemma 3.1. Note that for a 2-form α in either of the spaces $\Omega_7^2, \Omega_{14}^2$ the two notions $\alpha|_{\partial M} = 0, \alpha\|_{\partial M} = 0$ are equivalent. Now define a vector space

(3.4) $$\hat{A} = \{\alpha \in \Omega_{14}^2(M) : \mathrm{d}^*\alpha = 0, \alpha\|_{\partial,8} = 0\}.$$

We define a linear map $F_{\hat{A}} : \hat{A} \to TQ$ as follows. It is clear from Lemma 3.1 that if $\alpha_{14} \in \Omega_{14}^2$ satisfies $\alpha_{14}\|_{\partial,8} = 0$ we can find a form $\beta_7 \in \Omega_7^2$ such that $(\alpha_{14} + \beta_7)|_{\partial M} = 0$. For $\alpha_{14} \in \hat{A}$ we define $F_{\hat{A}}(\alpha_{14})$ to be the equivalence class of $\mathrm{d}(\alpha_{14} + \beta_7)$ in TQ. The definition of TQ means that this is well-defined, independent of the choice of β_7.

We digress here to review some standard Hodge Theory for manifolds with boundary. For any p we consider the Laplace operator $\Delta : \Omega^p(M) \to \Omega^p(M)$ and the equation with boundary conditions

(3.5) $$\Delta\mu = \rho, \quad \mu|_{\partial M} = 0, \quad \mathrm{d}^*\mu|_{\partial M} = 0.$$

PROPOSITION 3.2.
  • *If $\mathrm{d}^*\rho = 0$ then any solution μ of (3.5) satisfies $\mathrm{d}^*\mu = 0$. In fact this holds without assuming that $\mu|_{\partial M} = 0$.*

- If $\rho = 0$ then a solution satisfies $\mathrm{d}\mu = 0$. The space of such solutions

$$\mathcal{H}^p = \{\mu \in \Omega^p : \mathrm{d}\mu = 0, \ \mathrm{d}^*\mu = 0, \ \mu|_{\partial M} = 0\},$$

  represents the relative cohomology group $H^p(M, \partial M)$.
- There is a solution of (3.5) if and only if ρ is L^2-orthogonal to $\mathcal{H}^p$. In this case we define $G\rho$ to be the unique solution μ orthogonal to $\mathcal{H}^p$.
- If $\rho = \mathrm{d}^*\sigma$ for some $\sigma \in \Omega^{p+1}$ then ρ is orthogonal to $\mathcal{H}^p$.

Apart from the third item, the proofs are straightforward variants of the usual theory for closed manifolds, checking boundary terms. The third item is an application of elliptic boundary value theory. These results go back to Spencer and Duff, Morrey and Friedrichs. A standard modern reference is Section 2.4 in [11].

With this theory at hand we can return to the space $\hat{A}$.

LEMMA 3.3. — For $\alpha \in \mathcal{H}^2$ the component $\pi_{14}(\alpha) \in \Omega_{14}^2$ lies in $\hat{A}$ and the induced map $\pi_{14} : \mathcal{H}^2 \to \hat{A}$ is injective.

Proof. — First suppose that $\alpha = \alpha_7 + \alpha_{14}$ lies in $\mathcal{H}^2$. Since $\mathrm{d}^*\alpha$ and $\mathrm{d}_7\alpha$ both vanish it follows from the second item of Proposition 2.1 that $\mathrm{d}^*\alpha_{14} = 0$. Since α vanishes on the boundary it follows that $\alpha_{14}\|_{\partial,8} = 0$ and thus $\alpha_{14} \in \hat{A}$. Suppose that $\alpha_{14} = 0$. Then $\alpha_7|_{\partial M} = 0$ and, as we noted above, this implies that $\alpha_7\|_{\partial M} = 0$. The 2-form α_7 is harmonic and by the general theory (see [10, Section 3.5.2] for example) the Bochner formula on Ω_7^2 is the same as that on Ω^1. Thus we have $\nabla^*\nabla\alpha_7 = 0$ (since the Bochner formula on Ω^1 involves the Ricci curvature, which vanishes in our case). Now integration by parts, using the boundary condition $\alpha_7\|_{\partial M} = 0$, shows that $\nabla\alpha_7 = 0$, and since α_7 vanishes on the boundary it must be zero everywhere. This shows that π_{14} induces an injection from $\mathcal{H}^2$ to $\hat{A}$. $\square$

Define $A \subset \hat{A}$ to be the orthogonal complement of $\pi_{14}\mathcal{H}^2$.

PROPOSITION 3.4. — The map $F_{\hat{A}} : \hat{A} \to T\mathcal{Q}$ is surjective with kernel $\pi_{14}(\mathcal{H}^2)$. Hence there is an induced isomorphism $F_A : A \to T\mathcal{Q}$.

Proof. — Consider any $\gamma = \gamma_{14} + \gamma_7 \in \Omega^2(M)$ with $\gamma|_{\partial M} = 0$. We apply Proposition 2.2 with $\rho = \mathrm{d}^*\gamma_{14}$, so we find an $\eta = G(\mathrm{d}^*\gamma_{14})$ with $\eta|_{\partial M} = 0$ and $\mathrm{d}^*\mathrm{d}\eta = \mathrm{d}^*\gamma_{14}$. Now, by the third item of Proposition 2.1, we have $\mathrm{d}^*\mathrm{d}_{14}\eta = \frac{2}{3}\mathrm{d}^*\mathrm{d}\eta = \frac{2}{3}\mathrm{d}^*\gamma_{14}$. This means that $\alpha_{14} = \gamma_{14} - \frac{3}{2}\mathrm{d}_{14}\eta$ satisfies $\mathrm{d}^*\alpha_{14} = 0$. Also, since η vanishes on the boundary so does $\mathrm{d}\eta$, and this means that $\alpha_{14}\|_{\partial,8} = 0$. Thus α_{14} lies in $\hat{A}$. Going back to the definition

of $F_{\hat{A}}$: the form

$$\Gamma = \alpha_{14} + \gamma_7 - \frac{3}{2}\mathrm{d}_7\eta = \gamma - \frac{3}{2}\mathrm{d}\eta$$

vanishes on the boundary, so $F_{\hat{A}}(\alpha_{14})$ is the equivalence class of $\mathrm{d}\Gamma$ in $T\mathcal{Q}$. But $\mathrm{d}^2\eta = 0$ so $\mathrm{d}\Gamma = \mathrm{d}\gamma$. This shows that $F_{\hat{A}}$ is surjective.

In the other direction, suppose that $F_{\hat{A}}(\alpha_{14}) = 0$, for some $\alpha_{14} \in \hat{A}$. This means that we can choose an $\alpha_7 \in \Omega_7^2$ such that $\alpha = \alpha_7 + \alpha_{14}$ restricts to zero on the boundary and $\mathrm{d}\alpha = 0$. As in the proof of Lemma 3.3, the condition $\mathrm{d}^*\alpha_{14} = 0$ implies that $\mathrm{d}^*\alpha = 0$, so α lies in $\mathcal{H}^2$ and α_{14} is in $\pi_{14}(\mathcal{H}^2)$. □

We can now set-up the linear boundary value problem which is the main point of this paper.

THEOREM 3.5. — *For ρ in $\Omega_{14}^2(M)$ the equation $\Delta\alpha = \rho$ for $\alpha \in \Omega_{14}^2$, with boundary conditions*

$$\alpha\|_{\partial,8} = 0, \quad \mathrm{d}^*\alpha|_{\partial M} = 0,$$

is a self-adjoint elliptic boundary value problem. Moreover if $\mathrm{d}^*\rho = 0$ then a solution α satisfies $\mathrm{d}^*\alpha = 0$.

The statement that this is a self-adjoint elliptic boundary value problem has the following standard consequences. Define

$$(3.6) \qquad \tilde{H} = \{\alpha \in \Omega_{14}^2 : \Delta\alpha = 0, \ \alpha\|_{\partial,8} = 0, \ \mathrm{d}^*\alpha|_{\partial M} = 0\}.$$

Then

    (1) $\tilde{H}$ is finite dimensional;
    (2) a solution to the boundary value problem in Theorem 3.5 exists if and only if ρ is L^2-orthogonal to $\tilde{H}$;
    (3) in such a case we have elliptic estimates

$$\|\alpha\|_{L_k^2} \leqslant C_k \|\rho\|_{L_{k-2}^2}.$$

The proof of Theorem 3.5 extends across the next few pages, including Lemmas 3.6 and 3.7.

There is a standard definition of an elliptic boundary value problem (see [12, Chapter 5] for example) but the general theory is somewhat complicated and we do not need much of it here. We take as known the theory of the Dirichlet problem for the Laplace operator on Ω_{14}^2 and, for simplicity, we assume initially that the only solution $\alpha \in \Omega_{14}^2$ of $\Delta\alpha = 0$ with $\alpha\|_{\partial M} = 0$ is $\alpha = 0$. Then, by the standard theory, for any $\rho \in \Omega_{14}^2$ and section θ of the restriction of Λ_{14}^2 to ∂M there is a unique solution $\alpha = \alpha(\rho, \theta)$ of the equation $\Delta\alpha = \rho$ with $\alpha\|_{\partial M} = \theta$. Now consider a 1-form a on ∂M

and set $\theta(a) = 2a \wedge \nu^* - \chi_6(a)$ as in Lemma 3.1. The boundary value problem in Theorem 3.5 then becomes an equation for a:

$$d^*\alpha(\rho, \theta(a))|_{\partial M} = 0.$$

Write $d^*\alpha(\rho, 0)|_{\partial M} = -\sigma$ and let

$$P : \Omega^1(\partial M) \to \Omega^1(\partial M)$$

be the operator which maps a to $d^*\alpha(\theta(a), 0)|_{\partial M}$. Then the equation to solve for a is

(3.7) $P(a) = \sigma.$

The claim that the boundary value problem is elliptic is equivalent to the claim that P is an elliptic pseudo-differential operator of order 1 on ∂M. When this holds the equation (3.7) can be solved for a provided that σ lies in a subspace of finite codimension and the solution a is unique up to a finite dimensional kernel. These facts imply the corresponding statements about the boundary value problem. The simplifying assumption on the solubility of the Dirichlet problem is unnecessary since the whole discussion can take place modulo finite dimensional subspaces. From another point of view we can run the arguments above in a model "flat" case (as in the proof of Lemma 3.6 below) and use the solution there to construct a parametrix for our boundary value problem.

The ellipticity of the operator P is a condition on the symbol and this symbol can be described as follows.

LEMMA 3.6. — *Define an operator* $\tilde{P} : \Omega^1(\partial M) \to \Omega^1(\partial M)$ *by*

$$\tilde{P}(a) = 2\Delta^{1/2}a - d^*_{\partial M}\chi_6(a).$$

Then P and $\tilde{P}$ have the same symbol.

Proof. — To see this we can consider a situation where the geometry is locally flat, so we can work in $\mathbf{C}^3 \times (-\infty, 0]$ with boundary $\mathbf{C}^3 = \mathbf{C}^3 \times \{0\}$ and co-ordinate t in the $\mathbf{R}$ factor. We write a 2-form as

$$\alpha = 2a_t \wedge dt - \Psi_t$$

where a_t, Ψ_t are respectively t-dependent 1-forms and 2-forms on $\mathbf{C}^3$. Then

$$d^*\alpha = 2\frac{da_t}{dt} - d^*_6\Psi_t + 2(d^*_6 a_t)dt,$$

where d^*_6 denotes the d^* operator on $\mathbf{C}^3$. In our situation we have $\alpha \in \Omega^2_{14}$ so $\Psi_t = \chi_6 a_t + \Theta_t$, where Θ_t takes values in $\Lambda^2_{\partial,8}$ and $\Theta_0 = 0$. Thus the

restriction of $d^*\alpha$ to the boundary is given by

$$2\frac{\mathrm{d}a_t}{\mathrm{d}t}\Big|_{t=0} - \mathrm{d}_6^*(\chi_6(a_0)),$$

where $a(t)$ is the harmonic extension of the given 1-form a on the boundary. Thus $\frac{\mathrm{d}a}{\mathrm{d}t}\big|_{t=0}$ is obtained from a_0 by applying the "Dirichlet-to-Neumann" operator, whose symbol is the same as that of $\Delta^{1/2}$, and this gives the statement in the Lemma. $\qquad\square$

The symbol of $\Delta^{1/2}$ at a cotangent vector ξ on ∂M is multiplication by $|\xi|$. Thus the ellipticity of our boundary value problem follows from the following statement.

LEMMA 3.7. — If ξ is a unit cotangent vector on ∂M the symbol of $a \mapsto \mathrm{d}_{\partial M}^* \chi_6(a)$ at ξ has eigenvalues $0, \pm 1$

Proof. — To prove this Lemma we can compute in the flat model $\mathbf{C}^3$ as in (2.2), with standard co-ordinates $z_i = x_i + \sqrt{-1}y_i$, with $\omega = \mathrm{d}x_1\mathrm{d}y_1 + \mathrm{d}x_2\mathrm{d}y_2 + \mathrm{d}x_3\mathrm{d}y_3$ and with

$$\rho = -\mathrm{d}y_1\mathrm{d}y_2\mathrm{d}y_3 + \mathrm{d}y_1\mathrm{d}x_2\mathrm{d}x_3 + \mathrm{d}x_1\mathrm{d}y_2\mathrm{d}x_3 + \mathrm{d}x_1\mathrm{d}x_2\mathrm{d}y_3.$$

One finds from the definition that if $a = \sum_{i=1}^3 \lambda_i\mathrm{d}y_i + \mu_i\mathrm{d}x_i$ then

$$\chi_6(a) = \sum_{\mathrm{cyclic}} \mu_i(\mathrm{d}y_j\mathrm{d}y_k - \mathrm{d}x_j\mathrm{d}x_k) + \lambda_i(\mathrm{d}y_j\mathrm{d}x_k + \mathrm{d}x_j\mathrm{d}y_k),$$

where the notation means that (ijk) run over cyclic permutations of (123). Since the symmetry group $SU(3)$ acts transitively on the unit sphere it suffices to check any given unit co-tangent vector ξ, so we take $\xi = dx_1$. In other words we have to pull out the $\partial_1 = \frac{\partial}{\partial x_1}$ term in $d^*\chi_6(a)$. This is

$$(\partial_1\mu_3)\mathrm{d}x_2 - (\partial_1\mu_2)\mathrm{d}x_3 + (\partial_1\lambda_2)\mathrm{d}y_3 - (\partial_1\lambda_3)\mathrm{d}y_2.$$

To get the symbol we replace the derivative ∂_1 by multiplication by $\sqrt{-1}$. We see that the symbol at dx_1 is the linear map Σ with

$$\Sigma(\mathrm{d}x_1) = 0 \qquad \Sigma(\mathrm{d}x_2) = -\sqrt{-1}\mathrm{d}x_3 \qquad \Sigma(\mathrm{d}x_3) = \sqrt{-1}\mathrm{d}x_2$$

$$\Sigma(\mathrm{d}y_1) = 0 \qquad \Sigma(\mathrm{d}y_2) = \sqrt{-1}\mathrm{d}y_3 \qquad \Sigma(\mathrm{d}y_3) = -\sqrt{-1}\mathrm{d}y_2.$$

This linear map Σ has eigenvalues $0, 1, -1$ (each with multiplicity two). $\qquad\square$

Next we establish the self-adjoint property. Let us denote the boundary conditions in the statement of Theorem 3.5 by (BC). Recall that in general the adjoint boundary conditions $(BC)^*$ are defined by saying that $\beta \in \Omega_{14}^2$ satisfies $(BC)^*$ if and only if we have an equality of L^2 inner products $\langle\Delta\alpha, \beta\rangle = \langle\alpha, \Delta\beta\rangle$ for all α satisfying (BC). We want to show that the

boundary conditions $(BC)^*$ are the same as (BC). To do this it suffices, by a simple dimension-counting argument, to show that if α, β both satisfy (BC) then $\langle \Delta\alpha, \beta \rangle = \langle \alpha, \Delta\beta \rangle$. We can see this by an indirect argument using the operator L given by (13). First we claim that if $\alpha, \beta \in \Omega_{14}^2$ satisfy $\alpha\|_{\partial,8} = \beta\|_{\partial,8} = 0$ then $\langle L\alpha, \beta \rangle = \langle \alpha, L\beta \rangle$. Indeed choose $\alpha', \beta' \in \Omega_7^2$ such that $(\alpha + \alpha')|_{\partial M}, (\beta + \beta')|_{\partial M}$ vanish. Then

$$\langle L\alpha, \beta \rangle = \langle L(\alpha + \alpha'), \beta + \beta' \rangle,$$

since L vanishes on Ω_7^2 and maps to Ω_{14}^2. Let $\langle \cdot, \cdot \rangle_Q$ be the symmetric bilinear form associated to the quadratic form Q on Ω_M^2

$$(3.8) \qquad \langle \gamma_1, \gamma_2 \rangle_Q = \frac{4}{3} \langle d_1\gamma_1, d_1\gamma_2 \rangle + \langle d_7\gamma_1, d_7\gamma_2 \rangle - \langle d_{27}\gamma_1, d_{27}\gamma_2 \rangle.$$

Then *if γ_1, γ_2 vanish on ∂M* we have

$$(3.9) \qquad \langle \gamma_1 \gamma_2 \rangle_Q = \langle L\gamma_1, \gamma_2 \rangle = \langle \gamma_1 L\gamma_2 \rangle.$$

We apply this to $\alpha + \alpha', \beta + \beta'$, which vanish on the boundary by construction, so we have

$$\langle L(\alpha + \alpha'), \beta + \beta' \rangle = \langle \alpha + \alpha', \beta + \beta' \rangle_Q$$

which is symmetric in α, β. This we have shown that $\langle L\alpha, \beta \rangle$ is symmetric in α, β. Now, using Proposition 2.1 we can write $\Delta = L + \frac{3}{2}d_{14}d^*$ on Ω_{14}^2. If α, β satisfy the other part of (BC) that is, if $d^*\alpha|_{\partial M} = d^*\beta|_{\partial M} = 0$, then

$$\langle dd^*\alpha, \beta \rangle = \langle d^*\alpha, d^*\beta \rangle = \langle \alpha, dd^*\beta \rangle$$

since the relevant boundary terms vanish. This completes the proof of self-adjointness.

The last statement in Theorem 3.5 (that $d^*\rho = 0$ implies $d^*\alpha = 0$) is a particular case of the first item in Proposition 2.2. This completes the proof of Theorem 3.5.

The operator L and the symmetric form $\langle \cdot, \cdot \rangle_Q$ are related by a boundary term. For $\alpha \in \Omega_{14}^2$ with $\alpha\|_{\partial,8} = 0$ we define a 1-form on ∂M:

$$\alpha\|_{\partial,6} = \chi_6^{-1}(\alpha|_{\partial M}).$$

Recall also that we have a 2-form ω on ∂M given by the contraction of ϕ with the normal vector.

PROPOSITION 3.8. — If α_{14}, β_{14} are in Ω_{14}^2 with $\alpha_{14}\|_{\partial,8}, \beta_{14}\|_{\partial,8} = 0$ then

$$\langle L\alpha_{14}, \beta_{14} \rangle = \langle \alpha_{14}, \beta_{14} \rangle_Q + \langle \alpha_{14}\|_{\partial,6}, \beta_{14}\|_{\partial,6} \rangle_\partial$$

where, for 1-forms a, b on ∂M,

$$\langle a, b \rangle_{\partial} = -\int_{\partial M} \chi_6(a) \wedge \mathrm{d}(b \wedge \omega).$$

It follows from this proposition that, given $\rho \in \Omega_{14}^2$ with $\mathrm{d}^*\rho = 0$, our linear boundary value problem is the Euler–Lagrange equation associated to the functional on A given by

(3.10) $$- \|\mathrm{d}\alpha\|^2 + \langle \alpha\|_{\partial,6}, \alpha\|_{\partial,6} \rangle_{\partial} - \langle \rho, \alpha \rangle.$$

Proof. — We give a derivation of Proposition 3.8 although we will not use the result, so this is a digression from our main path. First, one can check that $\langle \cdot, \cdot \rangle_{\partial}$ is symmetric, so by polarisation it suffices to prove the formula for $\beta_{14} = \alpha_{14}$. As before, choose α_7 so that $\alpha = \alpha_{14} + \alpha_7$ vanishes on the boundary. It follows then by integration by parts that

$$\langle \alpha, \alpha_7 \rangle_Q = \langle \alpha, L\alpha_7 \rangle = 0$$

and

$$\langle \alpha, \alpha_{14} \rangle_Q = \langle \alpha, L\alpha_{14} \rangle = \langle \alpha_{14}, L\alpha_{14} \rangle.$$

Thus

$$\langle \alpha_{14}, \alpha_7 \rangle_Q = -\langle \alpha_7, \alpha_7 \rangle_Q$$

and

$$\langle \alpha_{14}, \alpha_{14} \rangle_Q = \langle \alpha_{14}, L\alpha_{14} \rangle + \langle \alpha_7, \alpha_7 \rangle_Q.$$

Let v be the vector field on M such that $\alpha_7 = i_v(\phi)$ and let $S : \Lambda^3 \to \Lambda^3$ be the bundle map equal to $\frac{4}{3}, +1, -1$ on the factors $\Lambda_1^3, \Lambda_7^3, \Lambda_{27}^3$ respectively. Thus by (2.8) the first variation of $\Theta(\phi)$ for a variation $\delta\phi$ in ϕ is $*S(\delta\phi)$. Take $\delta\phi = \mathrm{d}\alpha_7 = \mathcal{L}_v\phi$, so that by diffeomorphism invariance of the constructions $*S(\mathrm{d}\alpha_7) = \mathcal{L}_v(*\phi) = \mathrm{d}i_v(*\phi)$. Now

$$\langle \alpha_7, \alpha_7 \rangle_Q = \langle \mathrm{d}\alpha_7, S(\mathrm{d}\alpha_7) \rangle = \int_M \mathrm{d}\alpha_7 \wedge \mathrm{d}(i_v(*\phi)),$$

and we can write this as a boundary integral

$$\int_{\partial M} \alpha_7 \wedge \mathrm{d}(i_v * \phi).$$

On the boundary the assumption that $(\alpha_7 + \alpha_{14})|_{\partial M} = 0$ implies that v is tangential to ∂M. Then

$$i_v(*\phi) = \frac{1}{2}i_v(\omega^2) = -a \wedge \omega,$$

where $a = -i_v(\omega)$. Thus, on the boundary $\alpha_7 = -a \wedge \nu^* - \chi_6(a)$ and $\alpha_{14}|_{\partial M} = \chi_6(a)$, so $a = \alpha_{14}\|_{\partial,6}$ and the boundary term is the integral of $-\chi_6(a) \wedge \mathrm{d}(a \wedge \omega)$ as required. □

Remarks 3.9.

- The ellipticity of our boundary value problem depends crucially on the factor 2 appearing in the decomposition of Λ_{14}^2 in Lemma 3.1; more precisely that this factor is not ± 1. By contrast there is a similar-looking boundary value problem for forms in Ω_7^2 which is *not* elliptic. For this we consider the equation $\Delta \gamma = 0$ for $\gamma \in \Omega_7^2$ with boundary condition, in the obvious notation, $\mathrm{d}^* \gamma|_{\partial M} = 0$ and $\gamma\|_{\partial,1} = 0$. Under the identification between Ω_7^2 and Ω^1, the solutions correspond to 1-forms η on M with $\mathrm{d}_7 \eta = 0, \mathrm{d}^* \eta = 0$ and with $i_\nu(\eta) = 0$ on ∂M. This is the gauge-fixed, abelian, "G_2-instanton" equation and the space of solutions is infinite-dimensional.

- By standard theory, there is a complete set of eigenfunctions associated to our problem i.e. solutions of $-\Delta \alpha = \lambda \alpha$ satisfying (BC). The spectrum is discrete and bounded above so there are only finitely many positive eigenvalues. This 1-sided boundedness can be seen by considering the 1-parameter family of product metrics on $M \times S^1$ with the length of the S^1-factor equal to κ, lying in the interval $[2\pi, 4\pi]$ say. We consider sections of the bundle $\pi^*(\Lambda_{14}^2)$ lifted by the projection $\pi : M \times S^1 \to M$. There is an obvious way to lift the boundary conditions (BC) to $M \times S^1$ and we consider the operator

$$-\Delta_{M \times S^1} = -\Delta_M + \left(\frac{\mathrm{d}}{\mathrm{d}\theta} \right)^2 ,$$

with these lifted boundary conditions. The same discussion as before shows that this is an elliptic boundary value problem. The crucial fact is that the eigenvalues of the symbol Σ in the proof of Lemma 3.6 have modulus *less than* 2. If the spectrum of our original problem is not bounded above there are eigenfunctions α_i with eigenvalues $\lambda_i \to \infty$. For all large i we can choose parameter values κ_i such that $\sqrt{\lambda_i} = 0 \mod \kappa_i$ Then the sections

$$\tilde{\alpha}_i = \alpha_i \cos(\sqrt{\lambda_i} \theta)$$

satisfy $\Delta_{M \times S^1} \tilde{\alpha}_i = 0$, and this plainly contradicts the elliptic estimate on $M \times S^1$, which holds uniformly for parameter values $\kappa \in [2\pi, 4\pi]$.

4. The nonlinear problem

4.1. Gauge fixing

So far in this paper the connection between the linear theory and the deformations of G_2-structures has only been made at a formal level. We now correct this and develop the full nonlinear theory. There are two aspects to the nonlinearity, the first involving the action of the diffeomorphism group and the second involving the nonlinear nature of the torsion-free condition. For the first, happily, we are able to refer to the careful treatment of Fine, Lotay and Singer in [6]. This treats a 4-dimensional problem, but the proofs, extending the well-known results of Ebin for closed manifolds, go over immediately to our situation.

The standard approach to constructing a slice for the action of the diffeomorphism group on some space of tensors is to consider, at a tensor τ, the variations $\delta\tau$ which are L^2 orthogonal to the Lie derivatives $L_v\tau$, for all vector fields τ. Under the identification of tangent vectors with Λ_7^2 the Lie derivatives of the closed 3-form ϕ are the image of d : $\Omega_7^2 \to \Omega^3$. So in our case the standard slice for the diffeomorphism group action is given by variations $\delta\phi$ with $\pi_7(\mathrm{d}^*(\delta\phi)) = 0$, where π_7 is the projection to Λ_7^2. As usual, it is convenient to work with Banach spaces and following [6] we will use Sobolev spaces, although Hölder spaces would work just as well. We fix some suitably large s, say $s = 5$, and consider the set of maps from M to M which are equal to the identity on the boundary and with $s+1$ derivatives in L^2. Such maps are C^1 and it makes sense to consider diffeomorphisms of this class. These diffeomorphisms form a topological group and we define $\mathcal{G}^{s+1}$ to be the identity component of these L_{s+1}^2 diffeomorphisms. The group $\mathcal{G}^{s+1}$ acts on the space of L_s^2 3-forms on M. With our choice of the Sobolev index s these forms are also C^1. In particular the notion of positive 3-form makes sense. For $\delta > 0$ let

$$S_\delta = \{\phi + \chi : \|\chi\|_{L_s^2} < \delta,\ \pi_7 \mathrm{d}^*\chi = 0\}.$$

PROPOSITION 4.1. — There are constants $\epsilon, \delta > 0$ such that for every 3-form $\tilde{\phi}$ with $\|\tilde{\phi} - \phi\|_{L_s^2} < \epsilon$ there is a unique diffeomorphism $f \in \mathcal{G}^{s+1}$ such that $f^*(\tilde{\phi})$ lies in S_δ.

The statement is modelled on Theorem 2.1 of [6] and the proof is essentially the same so we do not go into it in detail here. However we do want to recall the linear result which underpins the proof.

LEMMA 4.2. — *For any $\sigma \in \Omega_7^2$ there is a unique $\gamma = \Gamma(\sigma) \in \Omega_7^2$ with $\gamma\|_{\partial M} = 0$ and $\pi_7 d^* d\gamma = \sigma$. The map Γ extends to a bounded map from L_{s-1}^2 to L_{s+1}^2.*

Proof. — It is easy to check that $\pi_7 d^* d$ is a self-adjoint elliptic operator on Ω_7^2. Thus by the standard theory it suffices to show uniqueness, in other words that for $\gamma \in \Omega_7^2$ if $\pi_7 d^* d\gamma = 0$ and $\gamma\|_{\partial M} = 0$ then $\gamma = 0$. Integrating by parts, the conditions imply that $d\gamma = 0$. Using the identification of Λ_7^2 with tangent vectors, γ corresponds to a vector field v, vanishing on the boundary of M, with $L_v(\phi) = 0$. This implies that v is a Killing field for the metric g_ϕ and it is a simple fact from Riemannian geometry that vanishing on the boundary forces v to vanish everywhere. $\qquad\square$

The set of closed 3-forms in a given enhancement class is preserved by the diffeomorphism group $\mathcal{G}$. Thus we immediately get from Proposition 3.8 a model for a neighbourhood in $\mathcal{Q}_{\hat\rho}$. Define

$$T = \{d\alpha : \alpha \in \Omega^2(M), \ \alpha|_{\partial M} = 0, \ \pi_7 d^* d\alpha = 0\},$$

and let T_s be the L_s^2 completion of T. Let $\mathcal{P}_{\hat\rho}^s$ be the L_s^2 version of $\mathcal{P}_{\hat\rho}$, in an obvious sense, and $\mathcal{Q}_{\hat\rho}^s$ be the quotient by $\mathcal{G}^{s+1}$ Then Proposition 3.8 implies that for suitable small $\delta > 0$ the map $\chi \mapsto \phi + \chi$ induces a homeomorphism from the ball $\{\chi \in T_s : \|\chi\|_{L_s^2} < \delta\}$ to a neighbourhood of $[\phi] \in \mathcal{Q}_{\hat\rho}^s$. Slightly more generally, if a 3-form ϕ_1 is sufficiently close to ϕ in L_s^2 norm and if $\hat\rho_1$ is the corresponding enhancement class, the map $\chi \mapsto \phi_1 + \chi$ induces a homeomorphism from this same ball in T_s to a neighbourhood of $[\phi_1]$ in $\mathcal{Q}_{\hat\rho_1}^s$.

It follows immediately from Lemma 4.2 that the natural map from T to $T\mathcal{Q}$ is an isomorphism. Define a vector space

$$B = \{\alpha = \alpha_7 + \alpha_{14} \in \Omega^2(M) : \alpha_{14} \in A, \ \alpha|_{\partial M} = 0, \ \pi_7 d^* d\alpha = 0\}.$$

The map $\alpha \mapsto \alpha_{14}$ induces a map $p : B \to A$ and Lemma 4.2 implies that this is an isomorphism. To spell this out, it is clear from the decomposition of the forms on the boundary that we can choose a smooth bundle map $\tau : \Lambda_{14}^2 \to \Lambda_7^2$, supported in a neighbourhood of the boundary, such that for all $\beta \in \Omega_{14}^2$ we have $\beta + \tau(\beta)|_{\partial M} = 0$. Now for $\alpha_{14} \in A$ define

(4.1) $\qquad F(\alpha_{14}) = \alpha_{14} + \tau(\alpha_{14}) - \Gamma(\pi_7 d^* d(\alpha_{14} + \tau(\alpha_{14})))$.

Then F maps to B and is the inverse to p. It is clear from the formula that this extends to an isomorphism (of topological vector spaces) $F : A_{s+1} \to B_{s+1}$ where B_{s+1} is the L_{s+1}^2 completion of B. Now the exterior derivative induces a map from B to T and it follows from Proposition 3.2 and the above-noted isomorphism of T and $T\mathcal{Q}$ that this is an isomorphism and

also for the corresponding Sobolev versions. Putting all this together we have the following result.

PROPOSITION 4.3. — There is a $\delta' > 0$ such that, if ϕ_1 is sufficiently close to ϕ in L_s^2, the map $\alpha \mapsto \phi_1 + \mathrm{d}F(\alpha)$ induces a homeomorphism from the ball in A_{s+1}: $\{\alpha \in A_{s+1} : \|\alpha\|_{L_{s+1}^2} < \delta'\}$ to a neighbourhood of $[\phi_1]$ in $\mathcal{Q}_{\hat{\rho}_1}^s$ where $\hat{\rho}_1$ is the enhanced boundary value determined by ϕ_1.

4.2. A Fredholm equation

We want to represent the solutions of our boundary value problem, for given data $\hat{\rho}$, as the zeros of a Fredholm map. Recall that for $\alpha \in \Omega^2(M)$ we defined $W(\alpha) = *_\phi \mathrm{d}\Theta(\phi + \mathrm{d}\alpha)$. In our description from Proposition 4.3 of a neighbourhood in $\mathcal{Q}_{\hat{\rho}}^s$ the torsion-free equation is $W(F(\alpha)) = 0$, for small $\alpha \in A_{s+1}$. The complication is that for general α the term $W(F(\alpha))$ lies in the subspace $*_\phi \Omega_{14,\tilde{\phi}}^5$ where $\tilde{\phi} = \phi + \mathrm{d}F(\alpha)$ and this space also depends on $\tilde{\phi}$. We use a projection construction to get around this, which essentially amounts to constructing a local trivialisation of the cotangent bundle of $\mathcal{Q}_{\hat{\rho}}^s$.

For any α as above, write $\sigma = W(F(\alpha)) = *_\phi \mathrm{d}\tilde{\phi}$. Then $\mathrm{d}^*\sigma = 0$ and σ is a section of the bundle $*_\phi \Lambda_{14,\tilde{\phi}}^5$. Let $\sigma = \sigma_7 + \sigma_{14}$. By our Hodge theory result, Proposition 2.2, there is a unique $\eta \in \Omega^1$ orthogonal to $\mathcal{H}^1$ solving the equation $\mathrm{d}^*\mathrm{d}\eta = \mathrm{d}^*\sigma_7$ with $\eta|_{\partial M} = 0$. Now recall (as in the proof of Proposition 3.2) that $\mathrm{d}^*\mathrm{d} = \frac{3}{2}\mathrm{d}^*\mathrm{d}_{14}$ on Ω^1. Thus

$$\mathrm{d}^*\sigma_{14} = -\mathrm{d}^*\sigma_7 = -\frac{3}{2}\mathrm{d}^*\mathrm{d}_{14}\eta.$$

Set $\tilde{\sigma}_{14} = \sigma_{14} + \frac{3}{2}\mathrm{d}_{14}\eta$, so $\tilde{\sigma}_{14}$ lies in Ω_{14}^2 and satisfies $\mathrm{d}^*\tilde{\sigma}_{14} = 0$.

LEMMA 4.4. — There is a $\delta'' > 0$ such that if $\|\alpha\|_{L_{s+1}^2} < \delta''$ then $\sigma = 0$ if and only if $\tilde{\sigma}_{14} = 0$.

Proof. — We know that σ is a section of the bundle $*_\phi \Lambda_{14,\tilde{\phi}}^5$. When $\tilde{\phi} = \phi$ this is exactly the bundle Λ_{14}^2. If $\tilde{\phi}$ is close to ϕ in C^0 we can use the standard graph construction. There is a bundle map

$$S_{\tilde{\phi}} : \Lambda_{14}^2 \to \Lambda_7^2$$

such that elements of $*_\phi \Lambda_{14,\tilde{\phi}}^5$ are of the form $\tau_{14} + S_{\tilde{\phi}}\tau_{14}$. If α is small in L_{s+1}^2 then $S_{\tilde{\phi}}$ is small in L_s^2. In the preceding discussion, we have $\sigma_7 = S_{\tilde{\phi}}\sigma_{14}$ so

$$\|\sigma_7\|_{L_k^2} \leqslant \epsilon \|\sigma_{14}\|_{L_k^2},$$

for $k \leqslant s-1$, where we can make ϵ as small as we please by taking α small in L^2_{s+1}. On the other hand, the elliptic estimates for the boundary value problem give

$$\|\mathrm{d}_{14}\eta\|_{L^2_k} \leqslant C\|\sigma_7\|_{L^2_k}.$$

The Lemma follows by choosing δ'' such that $\frac{3}{2}C\epsilon < 1$.   □

Write $\tilde{\sigma}_{14} = \Pi_{\tilde{\phi}}(\sigma)$. The conclusion of Proposition 4.3 and Lemma 4.4 is that the solutions of the torsion free equation in a neighbourhood of $[\phi]$ in $\mathcal{Q}_{\hat{\rho}}$ correspond to the zeros of a map $\mathcal{F}$ defined by

$$(4.2) \qquad\qquad \mathcal{F}(\alpha) = \Pi_{\tilde{\phi}}\left(W(\phi + \mathrm{d}F\alpha)\right).$$

This map $\mathcal{F}$ takes values in the space $A' = \{\alpha \in \Omega^2_{14} : \mathrm{d}^*\alpha = 0\}$. Write A'_{s-1} for the L^2_{s-1} completion.

PROPOSITION 4.5. — $\mathcal{F}$ extends to a smooth Fredholm map of index 0 from a neighbourhood of the origin in A_{s+1} to A'_{s-1}. The derivative at $\alpha = 0$ is $-\Delta : A_{s+1} \to A'_{s-1}$.

Proof. — The proof is straightforward, given Theorem 3.5. We compute the derivative formally:

$$W(F(\alpha)) = L(F(\alpha)) + O(\alpha^2),$$

(by definition of the linearised operator L);

$$L(F(\alpha)) = L(\alpha),$$

(since $F(\alpha)$ differs from α by a term in Ω^2_7 and L vanishes on Ω^2_7);

$$L(\alpha) = -\Delta\alpha,$$

(by the discussion in Section 3). Clearly the derivative of the projection term $\Pi_{\tilde{\phi}}$ at $\alpha = 0$ is the identity on Ω^2_{14}. Now one can check that these calculations are compatible with the Sobolev structures.

We know that Δ is self-adjoint on A so the index is zero and the cokernel can be identified with the kernel. This kernel is

$$\mathcal{H}_\phi = A \cap \tilde{H} = \{\alpha \in \Omega^2_{14} : \mathrm{d}^*\alpha = 0,\ \Delta\alpha = 0,\ \alpha\|_{\partial,8} = 0\},$$

and the isomorphism from A to $T\mathcal{Q}$ takes $\mathcal{H}_\phi$ to the space H_ϕ defined in (3.2).   □

This proposition achieves our goal of representing the solutions of the torsion-free equation in $\mathcal{Q}^s_{\hat{\rho}}$ as the zeros of a Fredholm map. By standard elliptic regularity any L^2_{s+1} solution is smooth.

To complete the discussion, we consider varying the boundary data $\hat{\rho}$. Let θ be a closed 3-form on M which is small in L^2_s, so that $\phi + \theta$ is a

closed positive 3-form on M which defines perturbed enhanced boundary data $\hat{\rho}(\theta)$. By Proposition 4.3, the map

$$\alpha \mapsto \phi + \theta + \mathrm{d}F(\alpha)$$

gives a homeomorphism from the δ''-ball in A^{s+1} to a neighbourhood of $[\phi + \theta]$ in $\mathcal{Q}_{\hat{\rho}(\theta)}$. In this neighbourhood the solutions of the torsion-free equation correspond to the zeros of a perturbed map

$$\mathcal{F}_\theta(\alpha) = \Pi_{\tilde{\phi}}\left(W(\phi + \theta + \mathrm{d}F\alpha)\right),$$

where $\tilde{\phi} = \phi + \theta + \mathrm{d}F(\alpha)$. This is a smooth map in the two variables $\theta \in \ker d \cap L_s^2, \alpha \in A_{s+1}$, taking values in A'_{s-1}. Thus we can apply standard theory to study the solutions of the torsion-free equation for nearby boundary data. In particular we have by the implicit function theorem:

THEOREM 4.6. — *If the vector space H_ϕ is zero then for smooth θ which are sufficiently small in L_s^2 there is a unique solution of the torsion-free equation in $\mathcal{Q}_{\hat{\rho}(\theta)}$ close to $\rho + \theta$.*

More generally, if H_ϕ is not zero the standard theory of Fredholm maps gives a finite-dimensional "Kuranishi model" for the solutions of the torsion-free equation.

This discussion of the local structure in $\mathcal{Q}_{\hat{\rho}}$ can be extended to the Hitchin functional. We define a quadratic form on A by

(4.3) $$q_\phi(\alpha) = -\langle \alpha, \Delta \alpha \rangle.$$

So the eigenvalues associated to our boundary value, which we discussed in Section 3, are the eigenvalues of the quadratic form q_ϕ relative to the L^2 form. Standard theory (as described in [3, Proposition 2.5], for example) gives a diffeomorphism from one neighbourhood of the origin in A_{s+1} to another which takes the Hitchin functional to a sum $f \circ \pi + \frac{2}{3}q_\phi$ where f is a real-valued function on $\mathcal{H}_\phi$ and $\pi : A_{s+1} \to \mathcal{H}_\phi$ is L^2 projection.

We emphasise that the crucial difference in this case of manifolds with boundary, compared with the closed case, is that the quadratic form q_ϕ is not manifestly negative definite due to the boundary term in Proposition 3.8.

5. Examples and questions

We have seen that the deformation theory for our boundary value problem is governed by a finite-dimensional vector space

$$H_\phi = \frac{\{\mathrm{d}\alpha : \alpha|_{\partial M} = 0, \, L(\alpha) = 0\}}{\{\mathrm{d}\alpha_7 : \alpha_7|_{\partial M} = 0\}}.$$

Now we define another vector space

$$(5.1) \qquad K_\phi = \frac{\{d\alpha_7 : \alpha_7 \in \Omega_7^2(M)\} \cap \{d\alpha : \alpha \in \Omega^2(M), \, \alpha|_{\partial M} = 0\}}{\{d\alpha_7 : \alpha_7|_{\partial M} = 0\}}.$$

There is an obvious map $E : K_\phi \to H_\phi$ arising from the fact that $L(\alpha_7)$ vanishes for any $\alpha_7 \in \Omega_7^2$ and it also obvious that this map E is injective.

Next recall that we have a decomposition of forms on the boundary

$$\Omega^2(\partial M) = \Omega_1^{2,\partial} \oplus \Omega_6^{2,\partial} \oplus \Omega_8^{2,\partial},$$

and write

$$\Omega_7^{2,\partial} = \Omega_1^{2,\partial} \oplus \Omega_6^{2,\partial}.$$

Thus the forms in $\Omega_7^{2,\partial}$ are exactly the restrictions to the boundary of forms in $\Omega_7^2(M)$. We define

$$V_\phi = \{\theta \in \Omega_7^{2,\partial} : d\theta = 0\},$$

and

$$W_\phi = \{\alpha_7 \in \Omega_7^2(M) : d\alpha_7 = 0\}.$$

So there is a restriction map

$$\iota : W_\phi \to V_\phi.$$

The space W_ϕ corresponds to the vector fields on M preserving ϕ (which are Killing fields for the metric) and, as we have noted in the proof of Lemma 4.4, the map ι is an injection.

The exact sequence of the pair $(M, \partial M)$ gives a co-boundary map from $H^2(\partial M)$ to $H^3(M, \partial M)$. Since an element of V_ϕ defines a class in $H^2(\partial M)$ we have a map, which we denote by

$$p : V_\phi \to H^3(M, \partial M).$$

Proposition 5.1. — *There is an isomorphism*

$$K_\phi \cong \ker p / (\ker p \cap \operatorname{Im} \iota).$$

Proof. — For simplicity, we just prove that if the right hand side is zero then so also is K_ϕ (which is what we will use). So suppose that we have a pair α, α_7 representing a class in K_ϕ (i.e. $d\alpha = d\alpha_7$ and $\alpha|_{\partial=0}$). Thus the restriction of α_7 to the boundary is a 2-form, θ say, in V_ϕ. Recall that in general the definition of the boundary map is that we extend θ to some 2-form Θ over M and take the cohomology class of $d\Theta$ in $H^3(M, \partial M)$. In our case we can take $\Theta = \alpha_7$ and the fact that $d\alpha_7 = d\alpha$ with $\alpha|_M = 0$ says exactly that $p(\theta) = 0$. So by assumption θ lies in the image of ι, say $\theta = \iota(\tilde\alpha_7)$. But now we can replace α_7 by $\alpha_7 - \tilde\alpha_7$, representing the same

class in K_ϕ. Thus we may as well suppose that α_7 restricts to zero on the boundary, which means that the class in K_ϕ is zero. $\qquad\square$

THEOREM 5.2. — *Suppose that M is a domain with smooth boundary in a closed manifold with torsion-free G_2 structure (M^+, ϕ^+) and that ϕ is the restriction of ϕ^+. Then $E : K_\phi \to H_\phi$ is an isomorphism.*

As we mentioned before, injectivity is trivial so we have to prove that the map is surjective. The proof is simply to extend deformations from M to M^+ and then apply the standard theory on the closed manifold. The only complication is that we have to work with forms that are not smooth.

Let $\alpha \in \Omega^2(M)$ represent a class in H_ϕ, so $\alpha|_{\partial M} = 0$ and $L(\alpha) = 0$. The first step is to show that we can suppose that α satisfies the stronger condition $\alpha\|_{\partial M} = 0$. Indeed suppose that in a normal product neighbourhood, with normal coordinate t,

$$\alpha = a_t \mathrm{d}t + b_t,$$

where a_t, b_t are t-dependent forms on ∂M. The hypothesis is that $b_0 = 0$. Let ϵ be the 1-form ta_0 in this product neighbourhood, extended smoothly over M. Then $\mathrm{d}\epsilon = t\mathrm{d}a_0 t - a_0 \mathrm{d}t$ and $\alpha' = \alpha + \mathrm{d}\epsilon$ satisfies $\alpha'\|_{\partial M} = 0$. Since α' represents the same class in H_ϕ we may as well suppose that $\alpha\|_{\partial M} = 0$.

Next let $\underline{\alpha}$ be the 2-form on M^+ equal to α on M and extended by zero over the complement. This extension is not smooth but it is Lipschitz, so $\underline{\alpha} \in L_1^p(M^+)$ for all p. We apply the standard theory, as sketched in Section 2, to $\underline{\alpha}$. Thus we solve the equation $\Delta\eta = \mathrm{d}^*\underline{\alpha}_{14}$. The Lipschitz condition means that $\mathrm{d}^*\underline{\alpha}_{14}$ has no distributional component and standard elliptic theory gives $\eta \in L_2^p(M^+)$ for all p. We find a harmonic form h_{14} on M^+ such that $\tilde{\alpha}_{14} = \underline{\alpha}_{14} - \mathrm{d}_{14}\eta + h_{14}$ is orthogonal to the harmonic space. Then by construction $\mathrm{d}^*\tilde{\alpha}_{14} = 0$. Now we bring in the hypothesis that $L(\alpha) = 0$. Taking the inner product with α this implies that $Q(\underline{\alpha}) = 0$ (since the relevant boundary term vanishes). The arguments for smooth forms all extend to L_1^p forms, and we deduce that $Q(\tilde{\alpha}_{14}) = 0$. But since $\mathrm{d}^*\tilde{\alpha}_{14} = 0$ we have $Q(\tilde{\alpha}_{14}) = \|\mathrm{d}\tilde{\alpha}_{14}\|^2$ so $\mathrm{d}\tilde{\alpha}_{14} = 0$ and hence $\tilde{\alpha}_{14}$ is harmonic. But this means that $\tilde{\alpha}_{14}$ vanishes, since it was chosen to be orthogonal to the harmonic forms.

Then

$$\mathrm{d}\underline{\alpha} = \mathrm{d}(\underline{\alpha}_7 + \underline{\alpha}_{14} + \mathrm{d}_7\eta + \mathrm{d}_{14}\eta + h) = \mathrm{d}(\tilde{\alpha}_7)$$

with $\tilde{\alpha}_7 = \underline{\alpha}_7 + \mathrm{d}_7\eta$. If we know that $\tilde{\alpha}_7$ is smooth on the manifold-with-boundary M, then we have shown that the pair $\alpha, \tilde{\alpha}_7$ represents a class in K_ϕ mapping by E to the given class in H_ϕ, thus completing the proof of the theorem. This smoothness follows from the following Lemma.

LEMMA 5.3. — *Suppose $\tilde{\alpha}_7$ is an L_1^p section of Λ_7^2 over M^+ such that $\mathrm{d}\tilde{\alpha}_7$ is smooth up to the boundary on M. Then $\tilde{\alpha}_7$ is also smooth up to the boundary on M.*

Proof. — Let D be the operator $\mathrm{d} : \Omega_{\partial,7}^2 \to \Omega^3(\partial M)$. This is an overdetermined elliptic operator, so if θ lies in some Sobolev space on ∂M and $D\theta$ is smooth on ∂M then θ is also smooth. We apply this to the restriction of $\tilde{\alpha}_7$ to ∂M. This lies *a priori* in a Sobolev space $L_{1-1/p}^p$ but now we see that it is a smooth form on ∂M. The relation between d^* and d_7 on Ω_7^2 shows that $\mathrm{d}^*\tilde{\alpha}_7$ is also smooth up to the boundary on M. So the same is true of $\Delta\tilde{\alpha}_7$. Thus $\tilde{\alpha}_7$ solves an elliptic boundary value problem

$$\Delta\tilde{\alpha}_7 = \rho, \quad \tilde{\alpha}_7|_{\partial M} = \sigma, \quad \mathrm{d}^*\tilde{\alpha}_7|_{\partial M} = \tau,$$

with ρ smooth up to the boundary on M and σ, τ smooth on ∂M. It follows by elliptic regularity that $\tilde{\alpha}_7$ is smooth up to the boundary on M. $\square$

The example discussed in [5] is an annular region in $\mathbf{R}^7$ which can be embedded in a compact torus, so Theorem 5.2 applies. In that example K_ϕ is not zero, so the same is true of H_ϕ and the deformation problem is obstructed.

In a similar vein we have

PROPOSITION 5.4. — *For $(M, \phi) \subset (M^+, \phi^+)$ as in Theorem 5.2 the quadratic form q_ϕ is negative semi-definite.*

This is essentially the same (in slightly different language) as [5, Proposition 1].

These results raise the following questions.

QUESTION 5.5. — *Is it true that for all (M, ϕ) we have $H_\phi = K_\phi$?*

QUESTION 5.6. — *Is it true that for all (M, ϕ) the form q_ϕ is negative semi-definite?*

The author has spent some effort attempting to answer these questions, without success. By the same argument as in [5, Proposition 1] an equivalent form of Question 5.6 is to ask whether the inequality

$$\|\mathrm{d}_{27}\alpha_{14}\| \geqslant \|\mathrm{d}_7\alpha_{14}\|$$

holds for all *compactly supported* $\alpha_{14} \in \Omega_{14}^2$.

The equation $D\theta = 0$ (in the notation of Lemma 5.3) is highly overdetermined, so one expects that typically the space V_ϕ is 0, and hence also K_ϕ. If the answer to Question 5.5 above is affirmative it would follow that in most situations the space H_ϕ is zero i.e. that the same three facts for the

closed manifold theory reviewed in the introduction hold for the boundary value problem, in most situations.

Leaving aside this question aside: Theorem 5.2 and Proposition 5.1 can be used to supply examples where H_ϕ vanishes. We will just consider one class of examples here. Let N be a closed Calabi–Yau 3-fold i.e. a complex 3-fold with Kähler form ω and holomorphic 3-form Θ such that ω and $\rho = \mathrm{Im}(\Theta)$ are equivalent to the standard model (2.2) at each point. For $L > 0$ let M_L be the manifold with boundary $N \times [0, L]$ with 3-form $\phi = \omega \mathrm{d}t + \rho$.

LEMMA 5.7. — *For this ϕ on M_L we have $H_\phi = 0$.*

Proof. — We can embed M_L in a closed manifold $M \times S^1$ so Theorem 5.2 applies. Thus we have to identify the space V_ϕ which is the sum of two copies of the kernel $\ker D$ of the operator D on N. Taking cohomology we have a map $h : \ker D \to H^2(N)$ and it follows from the Hodge decomposition that the image of this is $V = \mathbf{R}[\omega] + H_{\mathbf{R}}^{2,0}$ where $H_{\mathbf{R}}^{2,0}$ denotes the real part of complex cohomology. We claim that h is injective, so that $\ker D$ is isomorphic to V. Suppose that θ lies in the kernel of h, so $\theta = \mathrm{d}\eta$ for some 1-form η on N. In other words the component of $\mathrm{d}\eta$ in $\Lambda_0^{1,1}$ vanishes. The Hodge–Riemann bilinear relations give that

$$\mathrm{d}\eta \wedge \mathrm{d}\eta \wedge \omega = \left(2|\mathrm{d}_1\eta|^2 + |\mathrm{d}_6\eta|^2\right) \mathrm{vol}_6,$$

(where d_1 denotes the component in $\mathbf{R}\omega$ and d_6 the component in $\Lambda_{\mathbf{R}}^{2,0}$). So we have, by Stokes theorem,

$$0 = \int_N \mathrm{d}\eta \wedge \mathrm{d}\eta \wedge \omega = 2\|\mathrm{d}_1\eta\|^2 + \|\mathrm{d}_6\eta\|^2.$$

Hence $\mathrm{d}\eta = 0$ and the claim is proved.

We now have $V_\phi = V \oplus V$ with one copy of V for each boundary component. In this case $H^3(M_L\partial M_L) = H^2(N)$ and the map $p : V_\phi \to H^3(M_L\partial M_L)$ is

$$p(\theta_1, \theta_2) = \theta_1 - \theta_2.$$

So the kernel of p is the diagonal copy of V in $V \oplus V$. On the other hand it is clear that the space W_ϕ is isomorphic to V and that ι maps on to the diagonal so we see from Proposition 5.1 that $K_\phi = 0$.   □

We can apply our main result to get an existence theorem for deformations of these product manifolds. To state this we need to pin down the choice of enhancement data. Recall that the space of enhancements is an affine space modelled on $H^3(M, \partial M)$ but with no canonical origin. In our case we have $H^3(M_L, \partial M_L) = H^2(N)$, as above. Fix 2-cycles σ_a in N

representing a basis for $H_2(N)$. Then for any 3-form ψ on M_L we define

$$I_a = \int_{\sigma_a \times [0,L]} \psi.$$

The collection of these integrals can be regarded as an element $I(\psi) \in H^2(N)$. This induces an identification between the enhancements of a given boundary form and $H^2(N)$. Clearly $I(\phi_L) = L[\omega]$.

THEOREM 5.8. — *For L and (N, ω, ρ) as above there is a neighbourhood U of ρ in the space of closed forms on N (in the C^∞ topology) and a neighbourhood U' of $L[\omega]$ in $H^2(N)$ such that if ρ_0, ρ_L are in U and define the same cohomology class in $H^3(N)$ and ν is in U' then there is a torsion free G_2-structure ϕ on M_L which restricts to ρ_0, ρ_L on the two boundary components and with $I(\phi) = \nu$.*

Of course there is also a uniqueness statement, for solutions close to ϕ_L.

BIBLIOGRAPHY

[1] R. L. BRYANT, "Metrics with exceptional holonomy", *Ann. Math.* (1987), p. 525-576.

[2] ———, "Some remarks on G_2-structures", in *Gökova Geometry and Topology (Gökova, 2004-2005*, International Press, 2005, p. 71-109.

[3] S. K. DONALDSON, *Floer homology groups in Yang-Mills theory*, Cambridge Tracts in Mathematics, vol. 147, Cambridge University Press, 2002, vii+236 pages.

[4] ———, "Adiabatic limits of co-associative Kovalev-Lefschetz fibrations", in *Algebra, geometry, and physics in the 21st century*, Progress in Mathematics, vol. 324, Birkhäuser, 2017, p. 1-29.

[5] ———, "Boundary value problems in dimensions seven, four and three related to exceptional holonomy", in *Geometry and physics. Vol. I*, Oxford University Press, 2018.

[6] J. FINE, J. D. LOTAY & M. SINGER, "The space of hyperkähler metrics on a 4-manifold with boundary", *Forum Math. Sigma* **5** (2017), Art. ID e6, 50 p.

[7] N. HITCHIN, "The geometry of three-forms in six dimensions", *J. Differ. Geom.* **55** (2000), no. 3, p. 547-576.

[8] ———, "Stable forms and special metrics", in *Global differential geometry: the mathematical legacy of Alfred Gray (Bilbao, 2000)*, Contemporary Mathematics, vol. 288, American Mathematical Society, 2000, p. 70-89.

[9] D. JOYCE, "Compact Riemannian manifolds with exceptional holonomy G_2. I", in *Essays on Einstein manifolds*, Surveys in Differential Geometry, vol. 6, International Press, 1996, p. 291-328.

[10] ———, *Compact manifolds with special holonomy*, Oxford Mathematical Monographs, Oxford University Press, 2000, xii+436 pages.

[11] G. SCHWARZ, *Hodge decomposition. A method for solving boundary value problems*, Lecture Notes in Mathematics, vol. 1607, Springer, 1995, 155 pages.

[12] M. E. TAYLOR, *Pseudodifferential operators*, Princeton Mathematical Series, vol. 34, Princeton University Press, 1981, xi+452 pages.

Simon DONALDSON
The Simons Center for Geometry and Physics
Stony Brook and Imperial College
London (UK)
s.donaldson@imperial.ac.uk