Quaternion Algebraic Geometry

Dominic Widdows

June 5, 2006

Abstract

QUATERNION ALGEBRAIC GEOMETRY

Dominic Widdows

St Anne's College, Oxford

Thesis submitted Hilary Term, 2000, in support of application to supplicate for the degree of D.Phil.

This thesis is a collection of results about hypercomplex and quaternionic manifolds, focussing on two main areas. These are exterior forms and double complexes, and the 'algebraic geometry' of hypercomplex manifolds. The latter area is strongly influenced by techniques from quaternionic algebra.

A new double complex on quaternionic manifolds is presented, a quaternionic version of the Dolbeault complex on a complex manifold. It arises from the decomposition of real-valued exterior forms on a quaternionic manifold M into irreducible representations of Sp(1). This decomposition gives a double complex of differential forms and operators as a result of the Clebsch-Gordon formula $V_r \otimes V_1 \cong V_{r+1} \oplus V_{r-1}$ for Sp(1)-representations. The properties of the double complex are investigated, and it is established that it is elliptic in most places.

Joyce has created a new theory of quaternionic algebra [J1] by defining a quaternionic tensor product for A$\mathbb{H}$-modules ($\mathbb{H}$-modules equipped with a special real subspace). The theory can be described using sheaves over $\mathbb{C}P^1$, an interpretation due to Quillen [Q]. A$\mathbb{H}$-modules and their quaternionic tensor products are classified. Stable A$\mathbb{H}$-modules are described using Sp(1)-representations.

This theory is especially useful for describing hypercomplex manifolds and forming close analogies with complex geometry. Joyce has defined and investigated q-holomorphic functions on hypercomplex manifolds. There is also a q-holomorphic cotangent space which again arises as a result of the Clebsch-Gordon formula. A$\mathbb{H}$-module bundles are defined and their q-holomorphic sections explored.

Quaternion-valued differential forms on hypercomplex manifolds are of special interest. Their decomposition is finer than that of real forms, giving a second double complex with special advantages. The cohomology of these complexes leads to new invariants of compact quaternionic and hypercomplex manifolds.

Quaternion-valued vector fields are also studied, and lead to the definition of quaternionic Lie algebras. The investigation of finite-dimensional quaternionic Lie algebras allows the calculation of some simple quaternionic cohomology groups.

Acknowledgements

I would like to express gratitude, appreciation and respect for my supervisor Dominic Joyce. Without his inspiration as a mathematician this thesis would not have been conceived; without his patience and friendship it would certainly not have been completed.

Dedication

To the memory of Dr Peter Rowe, 1938-1998, who taught me relativity as an undergraduate in Durham. His determination to teach concepts before examination techniques introduced me to differential geometry.

Contents

Introduction

This aim of this thesis is to describe and develop various aspects of quaternionic algebra and geometry. The approach is based upon two pillars, namely the differential geometry of quaternionic manifolds and Joyce's recent theory of quaternionic algebra. Contributions are made to both fields of study, enabling these strands to be woven together in describing the algebraic geometry of hypercomplex manifolds.

A recurrent theme throughout will be representations of the group Sp(1) of unit quaternions. Our contribution to the theory of quaternionic manifolds relies on decomposing the Sp(1)-action on exterior forms, whilst the main new insight in quaternionic algebra is that the most important building blocks of Joyce's theory are best described and manipulated as Sp(1)-representations. The importance of Sp(1)-representations to both areas is chiefly responsible for the successful synthesis of methods in the work on hypercomplex manifolds.

Another frequent source of motivation is the behaviour of the complex numbers. Many situations in complex algebra and geometry have interesting quaternionic analogues. Aspects of complex geometry can often be described using the group U(1) of unit complex numbers; replacing this with the group Sp(1) can lead directly to quaternionic versions. The decomposition of exterior forms on quaternionic manifolds is precisely such an example, as is all the work on q-holomorphic functions and forms on hypercomplex manifolds. On the other hand, Joyce's quaternionic algebra is such a rich theory precisely because real subspaces of quaternionic vector spaces behave so differently from real subspaces of complex vector spaces.

Much of the original work presented is enticingly simple — indeed, I have often felt both surprised and privileged that it has not been carried out before. One of the main explanations for this is the relative unpopularity suffered by the quaternions in the 20^{th} century. This situation has left various aspects of quaternionic behaviour unexplored. To help understand the reasons for this omission and the consequent opportunities for development, the first chapter is devoted to a survey of the history of the quaternions and their applications. Background material also includes an introduction to the group Sp(1) and its representations. The irreducible representations on complex vector spaces and their tensor products are described, as are real and quaternionic representations.

Chapter 2 is about quaternionic structures in differential geometry. The approach is based on the work of Salamon [S3]. Taking complex manifolds as a model, hypercomplex manifolds (those possessing a torsion-free $\mathrm{GL}(n,\mathbb{H})$-structure) are defined, followed by the broader class of quaternionic manifolds (those possessing a torsion-free $\mathrm{Sp}(1)\mathrm{GL}(n,\mathbb{H})$-structure). After reviewing the Dolbeault complex, we consider the decomposition of differential forms on quaternionic manifolds, including an important elliptic complex discovered by Salamon upon which the integrability of the quaternionic structure depends.

This complex is in fact the top row of a hitherto undiscovered *double* complex on quaternionic manifolds, which is the subject of Chapter 3. As an Sp(1)-representation, the cotangent space of a quaternionic manifold M^{4n} takes the form $T^*M \cong 2nV_1$, where V_1 is the basic representation of Sp(1) on $\mathbb{C}^2$. Decomposition of the induced Sp(1)-representation on $\Lambda^k T^*M$ is a simple process achieved by considering weights. That this decomposition gives rise to a double complex results from the Clebsch-Gordon formula $V_r \otimes T^*M \cong V_r \otimes 2nV_1 \cong 2n(V_{r+1} \oplus V_{r-1})$. The new double complex is shown to be elliptic everywhere except along its bottom row, consisting of the basic representations V_1 and the trivial representations V_0. This double complex presents us with new quaternionic cohomology groups.

In the fourth chapter (which is partly a summary of the work of Joyce [J1] and Quillen [Q]) we move to our other major area of interest, the theory of quaternionic algebra. The building blocks of this theory are $\mathbb{H}$-modules equipped with a special real subspace. Such an object is called an A$\mathbb{H}$-module. Joyce has discovered a canonical tensor product operation for A$\mathbb{H}$-modules which is both associative and commutative. Using ideas from Quillen's work, we classify A$\mathbb{H}$-modules up to isomorphism. Dual A$\mathbb{H}$-modules are defined and shown to have interesting properties. Particularly well-behaved is the category of stable A$\mathbb{H}$-modules. In Chapter 5 it is shown that all stable A$\mathbb{H}$-modules and their duals are conveniently described using Sp(1)-representations.

The resulting theory is ideally adapted for describing hypercomplex geometry, a process begun in Chapter 6. A hypercomplex manifold M has a triple of global anticommuting complex structures which can be identified with the imaginary quaternions. This identification enables tensors on hypercomplex manifolds to be treated using the techniques of quaternionic algebra. Joyce has already used such an approach to define and investigate q-holomorphic functions on hypercomplex manifolds, which are seen as the quaternionic analogue of holomorphic functions. There is a natural product map on the A$\mathbb{H}$-module of q-holomorphic functions, which gives the q-holomorphic functions an algebraic structure which Joyce calls an H-algebra.

Using the Sp(1)-version of quaternionic algebra, we define a natural splitting of the quaternionic cotangent space $\mathbb{H} \otimes T^*M \cong A \oplus B$, and show that q-holomorphic functions are precisely those whose differentials take values in $A \subset \mathbb{H} \otimes T^*M$. The bundle A is hence defined to be the q-holomorphic cotangent space of M. These spaces are examples of A$\mathbb{H}$-module bundles or A$\mathbb{H}$-bundles, which we discuss. Several parallels with complex geometry arise. There are q-holomorphic A$\mathbb{H}$-bundles with q-holomorphic sections. Q-holomorphic sections are described using the quaternionic tensor product and the q-holomorphic cotangent space, and seen to form an H-algebra module over the q-holomorphic functions.

In the final chapter, such methods are applied to quaternion-valued tensors on hypercomplex manifolds. The double complex of Chapter 3 is revisited and adapted to quaternion-valued differential forms. The global complex structures give an extra decomposition which generalises the splitting $\mathbb{H} \otimes T^*M \cong A \oplus B$, further refining the double complex. The quaternion-valued double complex has advantages over the real-valued version, being elliptic in more places. The top row of the quaternion-valued double complex is particularly well-adapted to quaternionic algebra, which presents close parallels with the Dolbeault complex and motivates the definition of q-holomorphic k-forms.

Quaternion-valued vector fields are also interesting. The quaternionic tangent space splits as $\mathbb{H} \otimes TM \cong \widehat{A} \oplus \widehat{B}$ in the same way as the cotangent space. Vector fields taking values in $\widehat{A}$ are closed under the quaternionic tensor product and Lie bracket, a result which depends upon the integrability of the hypercomplex structure. This is the quaternionic analogue of the statement that on a complex manifold, the $(1,0)$ vector fields are closed under the Lie bracket. The vector fields in question therefore form a quaternionic Lie algebra, a new concept which we introduce. Interesting finite-dimensional quaternionic Lie algebras are used to calculate some quaternionic cohomology groups on Lie groups with left-invariant hypercomplex structures.

Chapter 1

The Quaternions and the Group Sp(1)

1.1 The Quaternions

The quaternions $\mathbb{H}$ are a four-dimensional real algebra generated by the identity element 1 and the symbols i_1, i_2 and i_3, so $\mathbb{H} = \{r_0 + r_1 i_1 + r_2 i_2 + r_3 i_3 : r_0, \ldots, r_3 \in \mathbb{R}\}$. Quaternions are added together component by component, and quaternion multiplication is given by the *quaternion relations*

$$i_1 i_2 = -i_2 i_1 = i_3, \quad i_2 i_3 = -i_3 i_2 = i_1, \quad i_3 i_1 = -i_1 i_3 = i_2, \quad i_1^2 = i_2^2 = i_3^2 = -1 \tag{1.1}$$

and the distributive law. The quaternion algebra is not commutative, though it does obey the associative law. The quaternions are a division algebra (an algebra with the property that $ab = 0$ implies that $a = 0$ or $b = 0$).

- Define the imaginary quaternions $\mathbb{I} = \langle i_1, i_2, i_3 \rangle$. The symbol $\mathbb{I}$ is not standard, but we will use it throughout.
- Define the conjugate $\bar{q}$ of $q = q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3$ by $\bar{q} = q_0 - q_1 i_1 - q_2 i_2 - q_3 i_3$. Then $\overline{(pq)} = \bar{q}\bar{p}$ for all $p, q \in \mathbb{H}$.
- Define the real and imaginary parts of q by $\mathrm{Re}(q) = q_0 \in \mathbb{R}$ and $\mathrm{Im}(q) = q_1 i_1 + q_2 i_2 + q_3 i_3 \in \mathbb{I}$. As with complex numbers, $\bar{q} = \mathrm{Re}(q) - \mathrm{Im}(q)$.
- We regard the real numbers $\mathbb{R}$ as a subfield of $\mathbb{H}$, and the quaternions as a direct sum $\mathbb{H} \cong \mathbb{R} \oplus \mathbb{I}$.
- Let $q = q_1 i_1 + q_2 i_2 + q_3 i_3 \in \mathbb{I}$. Then $q^2 = -1$ if and only if $q_1^2 + q_2^2 + q_3^2 = 1$, so the set of 'quaternionic square-roots of minus-one' is naturally isomorphic to the 2-sphere S^2. We shall often identify these sets, writing '$q \in S^2$' as a shorthand for '$q \in \mathbb{H} : q^2 = -1$'.
- If $q \in S^2$ then $\langle 1, q \rangle$ is a subfield of $\mathbb{H}$ isomorphic to $\mathbb{C}$. We shall call this subfield $\mathbb{C}_q$.

1.1.1 A History of the Quaternions

The quaternions were discovered by the Irish mathematician and physicist, William Rowan Hamilton (1805-1865), [1] whose contributions to mechanics are well-known and widely used. By 1835 Hamilton had helped to win acceptance for the system of complex numbers by showing that calculations with complex numbers are equivalent to calculations with ordered pairs of real numbers, governed by certain rules. At the time, complex numbers were being applied very effectively to problems in the plane $\mathbb{R}^2$. To Hamilton, the next logical step was to seek a similar *3-dimensional* number system which would revolutionise calculations in $\mathbb{R}^3$. For years, he struggled with this problem. In a touching letter to his son [H1], [2] dated shortly before his death in 1865, Hamilton writes:

> Every morning, on my coming down to breakfast, your brother and yourself used to ask me: "Well, Papa, can you multiply triplets?" Whereto I was always obliged to reply, with a sad shake of the head, "No, I can only add and subtract them".

For several years, Hamilton tried to manipulate the three symbols 1, i and j into an algebra. He finally realised that the secret was to introduce a *fourth* dimension. On 16th October, 1843, whilst walking with his wife, he had a flash of inspiration. In the same letter, he writes:

> An electric circuit seemed to close, and a spark flashed forth, the herald (as I foresaw immediately) of many long years to come of definitely directed thought and work ... I pulled out on the spot a pocket-book, which still exists, and made an entry there and then. Nor could I resist the impulse — unphilosophical as it may have been — to cut with a knife on the stone of Brougham Bridge, as we passed it, the fundamental formula with the symbols i, j, k:
>
> $$i^2 = j^2 = k^2 = ijk = -1,$$
>
> which contains the solution of the problem, but of course, as an inscription, has long since mouldered away.

Substituting i_1, i_2, i_3 for i, j, k, this gives the quaternion relations (1.1).

Rarely do we possess such a clear account of the genesis of a piece of mathematics. Most mathematical theories are invented gradually, and only after years of development can they be presented in a lecture course as a definitive set of axioms and results. The quaternions, on the other hand, "started into life, or light, full grown, on the 16th of Ocober, 1843..." [3] "Less than an hour elapsed" before Hamilton obtained leave of the Council of the Royal Irish Academy to read a paper on quaternions. The next day, Hamilton wrote a detailed letter to his friend and fellow mathematician John T. Graves

[1] Letters suggest that both Euler and Gauss were aware of the quaternion relations (1.1), though neither of them published the disovery [EKR, p. 192].

[2] Copies of Hamilton's most significant letters and papers concerning quaternions are currently available on the internet at `www.maths.tcd.ie/pub/HistMath/People/Hamilton`

[3] Letter to Professor P.G.Tait, an excerpt of which can be found on the same website as [H1].

[H2], giving us a clear account of the train of research which led him to his breakthrough. The discovery was published within a month on the 13th of November [H3].

The timing of the discovery amplified its impact upon Hamilton and his followers. The only other algebras known in 1843 were the real and complex numbers, both of which can be regarded as subalgebras of the quaternions. (It was not until 1858 that Cayley introduced matrices, and showed that the quaternion algebra could be realised as a subalgebra of the complex-valued 2×2 matrices.) As a result, Hamilton became the figurehead of a school of 'quaternionists', whose fervour for the new numbers far exceeded their usefulness. Hamilton believed his discovery to be of similar importance to that of the infinitesimal calculus, and devoted the rest of his career exclusively to its study. Echos of this zeal could still be heard this century; for example, while Eamon de Valera was President of Ireland (from 1959 to 1973), he would attend mathematical meetings whenever their title contained the word 'quaternions'!

Such excesses were bound to provoke a reaction, especially as it became clear that the quaternions are just one example of a number of possible algebras. Lord Kelvin, the famous Scottish physicist, once remarked that "Quaternions came from Hamilton after all his really good work had been done; and though beautifully ingenious, have been an unmixed evil to those who have touched them in any way" [EKR, p.193]. A belief that quaternions are somehow obsolete is often tacitly accepted to this day.

This is far from the case. The quaternions remain the simplest algebra after the real and complex numbers. Indeed, the real numbers $\mathbb{R}$, the complex numbers $\mathbb{C}$ and the quaternions $\mathbb{H}$ are the *only* associative division algebras, as was proved by Georg Frobenius in 1878: and amongst these the quaternions are the most general. The discovery of the quaternions provided enormous stimulation to algebraic research and it is thought that the term 'associative' was coined by Hamilton himself [H3, p.5] to describe quaternionic behaviour. Investigation into the nature of and constraints imposed by algebraic properties such as associativity and commutativity was greatly accelerated by the discovery of the quaternions.

The quaternions themselves have been used in various areas of mathematics. Most recently, quaternions have enjoyed prominence in computer science, because they are the simplest algebraic tool for describing rotations in three and four dimensions. Certainly, the numbers have fallen short of the early expectations of the quaternionists. However, quaternions do shine a light on certain areas of mathematics, and those who become familiar with them soon come to appreciate an intricacy and beauty which is all their own.

1.1.2 Quaternions and Matrices

In this section we will make use of the older notation $i = i_1$, $j = i_2$, $k = i_3$. This makes it easy to interpret i as the standard complex 'square root of -1' and j as a 'structure map' on the complex vector space $\mathbb{C}^2$.

It is well-known that the quaternions can be written as real or complex matrices, because there are isomorphisms from $\mathbb{H}$ into subalgebras of $\mathrm{Mat}(4, \mathbb{R})$ and $\mathrm{Mat}(2, \mathbb{C})$. The former of these is given by the mapping

$$q_0 + q_1 i + q_2 j + q_3 k \mapsto \begin{pmatrix} q_0 & q_1 & q_2 & q_3 \\ -q_1 & q_0 & -q_3 & q_2 \\ -q_2 & q_3 & q_0 & -q_1 \\ -q_3 & -q_2 & q_1 & q_0 \end{pmatrix}.$$

More commonly used is the mapping into $\text{Mat}(2, \mathbb{C})$. We can write every quaternion as a pair of complex numbers, using the equation

$$q_0 + q_1 i + q_2 j + q_3 k = (q_0 + q_1 i) + (q_2 + q_3 i)j. \tag{1.2}$$

In this way we obtain the expression $q = \alpha + \beta j \in \mathbb{H} \cong \mathbb{C}^2$. The map $j : \alpha + \beta j \mapsto -\bar{\beta} + \bar{\alpha} j$ is a conjugate-linear involution of $\mathbb{C}^2$ with $j^2 = -1$. This identification $\mathbb{H} \cong \mathbb{C}^2$ is not uniquely determined: each $q \in S^2$ determines a similar isomorphism.

Having written this down, it is easy to form the map

$$\iota : \mathbb{H} \to \mathcal{H} \subset \text{Mat}(2, \mathbb{C}) \qquad \alpha + \beta j \mapsto \begin{pmatrix} \alpha & \beta \\ -\bar{\beta} & \bar{\alpha} \end{pmatrix}. \tag{1.3}$$

The quaternion algebra can thus be realised as a real subalgebra of $\text{Mat}(2, \mathbb{C})$, using the identifications

$$1 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad i_1 = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \quad i_2 = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \quad i_3 = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}. \tag{1.4}$$

Note that the squared norm $q\bar{q}$ of a quaternion q is the same as the determinant of the matrix $\iota(q) \in \mathcal{H}$. The isomorphism ι gives an easy way to deduce that $\mathbb{H}$ is an associative division algebra; the inverse of any nonzero matrix $A \in \mathcal{H}$ is also in $\mathcal{H}$, and the only matrix in $\mathcal{H}$ whose determinant is zero is the zero matrix.

1.1.3 Simple Applications of the Quaternions

There are a number of ways in which quaternions can be used to express mathematical ideas. In many cases, a quaternionic description prefigures more modern descriptions. We will outline two main areas – vector analysis and Euclidean geometry. An excellent and readable account of most of the following can be found in Chapter 7 of [EKR].

Every quaternion can be uniquely written as the sum of its real and imaginary parts. If we identify the imaginary quaternions $\mathbb{I}$ with the real vector space $\mathbb{R}^3$, we can consider each quaternion $q = q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3$ as the sum of a scalar part q_0 and a vectorial part $(q_1, q_2, q_3) \in \mathbb{R}^3$ (indeed, it is in this context that the term 'vector' first appears [H4]). If we multiply together two imaginary quaternions p, $q \in \mathbb{I}$, we obtain a quaternionic version of the *scalar product* and *vector product* on $\mathbb{R}^3$, as follows:

$$pq = -p \cdot q + p \wedge q \in \mathbb{R} \oplus \mathbb{I} \cong \mathbb{H}. \tag{1.5}$$

Surprising as it seems nowadays, it was not until the 1880's that Josiah Willard Gibbs (1839-1903), a professor at Yale University, argued that the scalar and vector products had their own meaning, independent from their quaternionic origins. It was he who introduced the familiar notation $p \cdot q$ and $p \wedge q$ — before his time these were written '$-Spq$' and 'Vpq', indicating the 'scalar' and 'vector' parts of the quaternionic product $pq \in \mathbb{H}$.

Another crucial part of vector analysis which originated with Hamilton and the quaternions is the 'nabla' operator ∇ (so-called by Hamilton because the symbol ∇ is similar in shape to the Hebrew musical instrument of that name). The familiar *gradient operator* acting on a real differentiable function $f(x_1, x_2, x_3) : \mathbb{R}^3 \to \mathbb{R}$ was first written as

$$\nabla f := \frac{\partial f}{\partial x_1} i_1 + \frac{\partial f}{\partial x_2} i_2 + \frac{\partial f}{\partial x_3} i_3. \tag{1.6}$$

Hamilton went on to consider applying the operator ∇ to a 'differentiable quaternion field' $F(x_1, x_2, x_3) = f_1(x_1, x_2, x_3) i_1 + f_2(x_1, x_2, x_3) i_2 + f_3(x_1, x_2, x_3) i_3$, obtaining the equation

$$\nabla F = -\left(\frac{\partial f_1}{\partial x_1} + \frac{\partial f_2}{\partial x_2} + \frac{\partial f_3}{\partial x_3}\right) + \left(\frac{\partial f_3}{\partial x_2} - \frac{\partial f_2}{\partial x_3}\right) i_1 + \left(\frac{\partial f_1}{\partial x_3} - \frac{\partial f_3}{\partial x_1}\right) i_2 + \left(\frac{\partial f_2}{\partial x_1} - \frac{\partial f_1}{\partial x_2}\right) i_3,$$

which we recognise in modern terminology as

$$\nabla F = -\operatorname{div} F + \operatorname{curl} F.$$

Applying the ∇ operator to Equation (1.6) leads to the well-known *Laplacian operator* on $\mathbb{R}^3$:

$$\nabla^2 f = -\left(\frac{\partial^2 f}{\partial x_1^2} + \frac{\partial^2 f}{\partial x_2^2} + \frac{\partial^2 f}{\partial x_3^2}\right) =: \Delta f.$$

Having obtained the standard scalar product on $\mathbb{R}^3$, we can obtain the Euclidean metric on $\mathbb{R}^4$ in a similar fashion by relaxing the restriction in Equation (1.5) that p and q should be imaginary. For any $p,\, q \in \mathbb{H}$, we have

$$\operatorname{Re}(p\bar{q}) = p_0 q_0 + p_1 q_1 + p_2 q_2 + p_3 q_3 \in \mathbb{R},$$

and so we define the *canonical scalar product* on $\mathbb{H}$ by

$$\langle p, q \rangle = \operatorname{Re}(p\bar{q}). \tag{1.7}$$

We define the norm of a quaternion $q \in \mathbb{H}$ in the obvious way, putting

$$|q| = \sqrt{q\bar{q}}. \tag{1.8}$$

The norm function is multiplicative, *i.e.* $|pq| = |p||q|$ for $p,\, q \in \mathbb{H}$. As with complex numbers, we can combine the norm function with conjugation to obtain the inverse of $q \in \mathbb{H} \setminus \{0\}$ (it is easily verified that $q^{-1} = \bar{q}/|q|^2$ is the unique quaternion such that $qq^{-1} = q^{-1}q = 1$).

Another quaternionic formula, similar to Equation (1.7), is the identity

$$\operatorname{Re}(pq) = p_0 q_0 - p_1 q_1 - p_2 q_2 - p_3 q_3 \in \mathbb{R}. \tag{1.9}$$

If we regard p and q as four-vectors in spacetime with the 'time-axis' identified with $\mathbb{R} \subset \mathbb{H}$ and the 'spatial part' identified with $\mathbb{I}$, this is identical to the Lorentz metric of special relativity.

Let a and b be quaternions of unit norm, and consider the involution $f_{a,b} : \mathbb{H} \to \mathbb{H}$ given by $f_{a,b}(q) = aqb$. Then $|f_{a,b}(q)| = |q|$ and the function $f_{a,b} : \mathbb{R}^4 \to \mathbb{R}^4$ is a rotation. Clearly, $f_{-a,-b} = f_{a,b}$, so each of these choices for the pair a, b gives the same rotation. In 1855, Cayley showed that all rotations of $\mathbb{R}^4$ can be written in this fashion and that the two possibilities given above are the only two which give the same rotation. Also, all reflections can be obtained by the involutions $\bar{f}_{a,b}(q) = a\bar{q}b$.

One of the beauties of this system is that having obtained all rotations of $\mathbb{R}^4$, we obtain the rotations of $\mathbb{R}^3$ simply by putting $a = b^{-1}$, giving the inner automorphism of $\mathbb{H}$, $q \mapsto aqa^{-1}$. (By Cayley's theorem, all real-algebra automorphisms of $\mathbb{H}$ take this form.) This fixes the real line $\mathbb{R}$ and rotates the imaginary quaternions $\mathbb{I}$. Identifying $\mathbb{I}$ with $\mathbb{R}^3$, the map $q \mapsto aqa^{-1}$ is a rotation of $\mathbb{R}^3$. According to Cayley, within a year of the discovery of the quaternions Hamilton was aware that all rotations of $\mathbb{R}^3$ can be expressed in this fashion.

These discoveries provided much insight into the classical groups SO(3) and SO(4), and helped to develop our knowledge of transformation groups in general. There are interesting questions which arise. Why do the unit quaternions turn out to be so important? In view of the quaternionic version of the Lorentz metric in Equation (1.9), can we use quaternions to write Lorentz tranformations as elegantly? Are there other spaces which might lend themselves to quaternionic treatment? These questions are best addressed using the theory of Lie groups, which the pioneering work of Hamilton and Cayley helped to develop.

1.2 The Lie Group Sp(1) and its Representations

The unit quaternions form a subgroup of $\mathbb{H}$ under multiplication, which we call Sp(1). Its importance to the quaternions is equivalent to that of the circle group U(1) of unit complex numbers in complex analysis. As a manifold Sp(1) is the 3-sphere S^3. The multiplication and inverse maps are smooth, so Sp(1) is a compact Lie group. The isomorphism $\iota : \mathbb{H} \to \mathcal{H} \subset \mathrm{Mat}(2, \mathbb{C})$ of Equation (1.4) maps Sp(1) isomorphically onto SU(2).

The Lie algebra $\mathfrak{sp}(1)$ of Sp(1) is generated by the elements I, J and K, the Lie bracket being given by the relations $[I, J] = 2K$, $[J, K] = 2I$ and $[K, I] = 2J$. We can write these using the matrices of Equation (1.4):

$$I = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \quad J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \quad K = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}.$$

The complexification of $\mathfrak{sp}(1)$ is the Lie algebra $\mathfrak{sl}(2, \mathbb{C})$. This is generated over $\mathbb{C}$ by the elements

$$H = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad X = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \quad Y = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \tag{1.10}$$

and the relations

$$[H, X] = 2X, \quad [H, Y] = -2Y \quad \text{and} \quad [X, Y] = H. \tag{1.11}$$

Hence the equations

$$I = iH, \quad J = X - Y \quad \text{and} \quad K = i(X+Y) \tag{1.12}$$

give one possible identification $\mathfrak{sp}(1) \otimes_{\mathbb{R}} \mathbb{C} \cong \mathfrak{sl}(2, \mathbb{C})$.

For any $Q \in \mathfrak{sp}(1)$, the normaliser $N(Q)$ of Q is defined to be

$$N(Q) = \{P \in \mathfrak{sp}(1) : [P, Q] = 0\} = \langle Q \rangle,$$

which is a Cartan subalgebra of $\mathfrak{sp}(1)$. Identifying a unit vector $Q = aI+bJ+cK \in \mathfrak{sp}(1)$ with the corresponding imaginary quaternion $q = ai_2 + bi_2 + ci_3 \in S^2$, the exponential map $\exp : \mathfrak{sp}(1) \to \mathrm{Sp}(1)$ maps $\langle Q \rangle$ to the unit circle in $\mathbb{C}_q$, which we will call $\mathrm{U}(1)_q$.

In the previous section it was shown that rotations in three and four dimensions can be written in terms of unit quaternions. This is because there is a commutative diagram of Lie group homomorphisms

$$\begin{array}{ccc} \mathrm{Sp}(1) \cong \mathrm{Spin}(3) & \hookrightarrow & \mathrm{Sp}(1) \times \mathrm{Sp}(1) \cong \mathrm{Spin}(4) \\ \downarrow & & \downarrow \\ \mathrm{SO}(3) & \hookrightarrow & \mathrm{SO}(4). \end{array} \tag{1.13}$$

The horizontal arrows are inclusions, and the vertical arrows are $2:1$ coverings with kernels $\{1, -1\}$ and $\{(1,1), (-1,-1)\}$ respectively. The applications of these homomorphisms in Riemannian geometry are described by Salamon in [S1].

From this, we can see clearly why three- and four-dimensional Euclidean geometry fit so well in a quaternionic framework. We can also see why the same is not true for Lorentzian geometry. Here the important group is the Lorentz group SO(3,1). Whilst there is a double cover $\mathrm{SL}(2, \mathbb{C}) \to \mathrm{SO}_0(3,1)$, this does not restrict to a suitably interesting *real* homomorphism $\mathrm{Sp}(1) \to \mathrm{SO}_0(3,1)$. It is possible to write Lorentz transformations using quaternions (for a modern example see [dL]), but the author has found no way which is simple enough to be really effective. There have been attempts to use the *biquaternions* $\{p + iq : p + q \in \mathbb{H}\} \cong \mathbb{H} \otimes_{\mathbb{R}} \mathbb{C}$ to apply quaternions to special relativity [Sy], but the mental gymnastics involved in using four 'square roots of -1' are cumbersome: the equivalent description of 'spin transformations' using the matrices of the group SL(2,$\mathbb{C}$) are a more familiar and fertile ground.

1.2.1 The Representations of Sp(1)

A representation of a Lie group G on a vector space V is a Lie group homomorphism $\rho : G \to \mathrm{GL}(V)$. We will sometimes refer to V itself as a representation where the map ρ is understood. The differential $d\rho$ is a Lie algebra representation $d\rho : \mathfrak{g} \to \mathrm{End}(V)$.

The representations of $\mathrm{Sp}(1) \cong \mathrm{SU}(2)$ will play an important part in this thesis. Their theory is well-known and used in many situations. Good introductions to Sp(1)-representations include [BD, §2.5] (which describes the action of the group SU(2)) and [FH, Lecture 11] (which provides a description in terms of the action of the Lie algebra $\mathfrak{sl}(2, \mathbb{C})$). We recall those points which will be of particular importance.

Because it is a compact group, every representation of Sp(1) on a complex vector space V can be written as a direct sum of irreducible representations. The multiplicity of each irreducible in such a decomposition is uniquely determined. Moreover, there is a unique irreducible representation on $\mathbb{C}^n$ for every $n > 0$. This makes the representations of Sp(1) particularly easy to describe. Let V_1 be the basic representation of SL(2,$\mathbb{C}$) on $\mathbb{C}^2$ given by left-action of matrices upon column vectors. This coincides with the basic representation of Sp(1) by left-multiplication on $\mathbb{H} \cong \mathbb{C}^2$. The unique irreducible representation on $\mathbb{C}^{n+1}$ is then given by the n^{th} symmetric power of V_1, so we define

$$V_n = S^n(V_1).$$

The representation V_n is irreducible [BD, Proposition 5.1], and every irreducible representation of Sp(1) is of the form V_n for some nonnegative $n \in \mathbb{Z}$ [BD, Proposition 5.3].

Let $\mathbf{x}$ and $\mathbf{y}$ be a basis for $\mathbb{C}^2$, so that $V_1 = \langle \mathbf{x}, \mathbf{y} \rangle$. Then

$$V_n = S^n(V_1) = \langle \mathbf{x}^n, \mathbf{x}^{n-1}\mathbf{y}, \mathbf{x}^{n-2}\mathbf{y}^2, \ldots, \mathbf{x}^2\mathbf{y}^{n-2}, \mathbf{x}\mathbf{y}^{n-1}, \mathbf{y}^n \rangle.$$

The action of SL(2,$\mathbb{C}$) on V_n is given by the induced action on the space of homogeneous polynomials of degree n in the variables $\mathbf{x}$ and $\mathbf{y}$.

Each of the Lie group representations V_n is a representation of the Lie algebras $\mathfrak{sl}(2, \mathbb{C})$ and $\mathfrak{sp}(1)$. Another very important way to describe the structure of these representations is obtained by decomposing them into weight spaces (eigenspaces for the action of a Cartan subalgebra). In terms of H, X and Y, the $\mathfrak{sl}(2, \mathbb{C})$-action on V_1 is given by

$$\begin{array}{rclcrclcrcl} H(\mathbf{x}) & = & \mathbf{x} & & X(\mathbf{x}) & = & 0 & & Y(\mathbf{x}) & = & \mathbf{y} \\ H(\mathbf{y}) & = & -\mathbf{y} & & X(\mathbf{y}) & = & \mathbf{x} & & Y(\mathbf{y}) & = & 0. \end{array} \tag{1.14}$$

To obtain the induced action of $\mathfrak{sl}(2, \mathbb{C})$ on V_n we use the Leibniz rule [4] $A(a \cdot b) = A(a) \cdot b + a \cdot A(b)$. This gives

$$\begin{array}{rcl} H(\mathbf{x}^{n-k}\mathbf{y}^k) & = & (n-2k)(\mathbf{x}^{n-k}\mathbf{y}^k) \\ X(\mathbf{x}^{n-k}\mathbf{y}^k) & = & k(\mathbf{x}^{n-k+1}\mathbf{y}^{k-1}) \\ Y(\mathbf{x}^{n-k}\mathbf{y}^k) & = & (n-k)(\mathbf{x}^{n-k-1}\mathbf{y}^{k+1}). \end{array} \tag{1.15}$$

Each subspace $\langle \mathbf{x}^{n-k}\mathbf{y}^k \rangle \subset V_n$ is therefore a weight space of the representation V_n, and the weights are the integers

$$\{n, n-2, \ldots, n-2k, \ldots, 2-n, -n\}.$$

Thus V_n is also characterised by being the unique irreducible representation of $\mathfrak{sl}(2, \mathbb{C})$ with highest weight n. We can compute the action of $\mathfrak{sp}(1)$ on V_n by substituting I, J and K for H, X and Y using the relations of (1.12).

Another important operator which acts on an $\mathfrak{sp}(1)$-representation is the *Casimir operator* $C = I^2 + J^2 + K^2 = -(H^2 + 2XY + 2YX)$. It is easy to show that for any $\mathbf{x}^{n-k}\mathbf{y}^k \in V_n$,

$$C(\mathbf{x}^{n-k}\mathbf{y}^k) = -n(n+2)\mathbf{x}^{n-k}\mathbf{y}^k. \tag{1.16}$$

[4] This can be found in [FH, p. 110], which describes the action of Lie groups and Lie algebras on tensor products.

Each irreducible representation V_n is thus an eigenspace of the Casimir operator with eigenvalue $-n(n+2)$.

Let V_m and V_n be two Sp(1)-representations. Then their tensor product $V_m \otimes V_n$ is naturally an $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$-representation, [5] and also a representation of the diagonal Sp(1)-subgroup, the action of which is given by

$$g(u \otimes v) = g(u) \otimes g(v).$$

The irreducible decomposition of the diagonal Sp(1)-representation on $V_m \otimes V_n$ is given by the famous *Clebsch-Gordon formula*,

$$V_m \otimes V_n \cong V_{m+n} \oplus V_{m+n-2} \oplus \cdots \oplus V_{m-n+2} \oplus V_{m-n} \qquad \text{for } m \geq n. \quad (1.17)$$

This can be proved using characters [BD, Proposition 5.5] or weights [FH, Exercise 11.11].

Real and quaternionic representations

It is standard practice to work primarily with representations on complex vector spaces. Representations on real (and quaternionic) vector spaces are obtained using antilinear *structure maps*. A thorough guide to this process is in [BD, §2.6]. In the case of Sp(1)-representations, we define the structure map $\sigma_1 : V_1 \to V_1$ by

$$\sigma_1(z_1\mathbf{x} + z_2\mathbf{y}) = -\bar{z}_2\mathbf{x} + \bar{z}_1\mathbf{y} \qquad z_1, z_2 \in \mathbb{C}.$$

Then $\sigma_1^2 = -1$, and σ_1 coincides with the map j of Section 1.1.2. Let σ_n be the map which σ_1 induces on V_n, *i.e.*

$$\sigma_n(z_1\mathbf{x}^{n-k}\mathbf{y}^k) = (-1)^k \bar{z}_1\mathbf{x}^k\mathbf{y}^{n-k}. \quad (1.18)$$

If $n = 2m$ is even then $\sigma_{2m}^2 = 1$ and σ_{2m} is a *real structure* on V_{2m}. Let V_{2m}^σ be the set of fixed-points of σ_{2m}. Then $V_{2m}^\sigma \cong \mathbb{R}^{2m+1}$ is preserved by the action of Sp(1), and

$$V_{2m}^\sigma \otimes_{\mathbb{R}} \mathbb{C} \cong V_{2m}.$$

Thus V_{2m}^σ is a representation of Sp(1) on the real vector space $\mathbb{R}^{2m+1}$.

If on the other hand $n = 2m-1$ is odd, $\sigma_{2m-1}^2 = -1$. Then Sp(1) acts on the underlying *real* vector space $\mathbb{R}^{4m}$. This real vector space comes equipped with the complex structure i and the structure map σ_{2m-1}, in such a way that the subspace

$$\langle v,\ iv,\ \sigma_{2m-1}(v), i\sigma_{2m-1}(v)\rangle \cong \mathbb{R}^4$$

is isomorphic to the quaternions; thus $V_{2m-1} \cong \mathbb{H}^m$. This is why a complex antilinear map σ on a complex vector space V such that $\sigma^2 = -1$ is called a *quaternionic structure*.

---

[5] Since $\mathrm{Sp}(1) \times \mathrm{Sp}(1) \cong \mathrm{Spin}(4)$, we can construct all Spin(4) and hence all SO(4) representations in this fashion — see [S1, §3].

1.3 Difficulties with the Quaternions

Quaternions are far less predictable than their lower-dimensional cousins, due to the great complicating factor of non-commutativity. Many of the ideas which work beautifully for real or complex numbers are not suited to quaternions, and attempts to use them often result in lengthy and cumbersome mathematics — most of which dates from the 19^{th} century and is now almost forgotten. However, with modesty and care, quaternions can be used to recreate many of the structures over the real and complex numbers with which we are familiar. The purpose of this section is to outline some of the difficulties with a few examples, which highlight the need for caution: but also, it is hoped, point the way to some of the successes we will encounter.

1.3.1 The Fundamental Theorem of Algebra for Quaternions

The single biggest reason for using the complex numbers in preference to any other number field is the so-called 'Fundamental Theorem of Algebra' — every complex polynomial of degree n has precisely n zeros, counted with multiplicities. The real numbers are not so well-behaved: a real polynomial of degree n can have fewer than n real roots.

Quaternion behaviour is many degrees freer and less predictable. If we multiply together two 'linear factors', we obtain the following expression:

$$(a_1X + b_1)(a_2X + b_2) = a_1Xa_2X + a_1Xb_2 + b_1a_2X + b_1b_2.$$

It quickly becomes obvious that non-commutativity is going make any attempt to factorise a general polynomial extremely troublesome.

Moreover, there are many polynomials which display extreme behaviour. For example, the cubic $X^2i_1Xi_1 + i_1X^2i_1X - i_1Xi_1X^2 - Xi_1X^2i_1$ takes the value zero for all $X \in \mathbb{H}$. At the other extreme, since $i_1X - Xi_1 \in \mathbb{I}$ for all $X \in \mathbb{H}$, the equation $i_1X - Xi_1 + 1 = 0$ has no solutions at all!

In order to arrive at any kind of 'fundamental theorem of algebra', we need to restrict our attention considerably. We define a *monomial of degree n* to be an expression of the form

$$a_0Xa_1Xa_2 \cdots a_{n-1}Xa_n, \quad a_i \in \mathbb{H} \setminus \{0\}.$$

Then there is the following 'fundamental theorem of algebra for quaternions':

Theorem 1.3.1 *[EKR, p. 205] Let f be a polynomial over $\mathbb{H}$ of degree $n > 0$ of the form $m + g$, where m is a monomial of degree n and g is a polynomial of degree $< n$. Then the mapping $f : \mathbb{H} \to \mathbb{H}$ is surjective, and in particular f has zeros in $\mathbb{H}$.*

This is typical of the difficulties we encounter with quaternions. There *is* a quaternionic analogue of the theorems for real and complex numbers, but because of non-commutativity the quaternionic version is more complicated, less general and because of this less useful.

1.3.2 Calculus with Quaternions

The monumental successes of complex analysis make it natural to look for a similar theory for quaternions. Complex analysis can be described as the study of *holomorphic functions*. A function $f : \mathbb{C} \to \mathbb{C}$ is holomorphic if it has a well-defined complex derivative. One of the fundamental results in complex analysis is that every holomorphic function is analytic *i.e.* can be written as a convergent power series.

Sadly, neither of these definitions proves interesting when applied to quaternions — the former is too restrictive and the latter too general. For a function $f : \mathbb{H} \to \mathbb{H}$ to have a well-defined derivative

$$\frac{df}{dq} = \lim_{h\to 0}(f(q+h) - f(q))h^{-1},$$

it can be shown [Su, §3, Theorem 1] that f must take the form $f(q) = a + bq$ for some $a, b \in \mathbb{H}$, so the only functions which are quaternion-differentiable in this sense are affine linear.

In contrast to the complex case, the components of a quaternion can be written as quaternionic polynomials, *i.e.* for $q = q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3$,

$$q_0 = \frac{1}{4}(q - i_1 q i_1 - i_2 q i_2 - i_3 q i_3) \qquad q_1 = \frac{1}{4i_1}(q - i_1 q i_1 + i_2 q i_2 + i_3 q i_3) \quad \textit{etc.} \tag{1.19}$$

This takes us to the other extreme: the study of quaternionic power series is the same as the theory of real-analytic functions on $\mathbb{R}^4$. Hamilton and his followers were aware of this — it was in Hamilton's work on quaternions that some of the modern ideas in the theory of functions of several real variables first appeared. Despite the benefits of this work to mathematics as a whole, Hamilton never developed a successful 'quaternion calculus'.

It was not until the work of R. Fueter, in 1935, that a suitable definition of a 'regular quaternionic function' was found, using a quaternionic analogue of the Cauchy-Riemann equations. A *regular function* on $\mathbb{H}$ is defined to be a real-differentiable function $f : \mathbb{H} \to \mathbb{H}$ which satisfies the *Cauchy-Riemann-Fueter equations*:

$$\frac{\partial f}{\partial q_0} + \frac{\partial f}{\partial q_1} i_1 + \frac{\partial f}{\partial q_2} i_2 + \frac{\partial f}{\partial q_3} i_3 = 0. \tag{1.20}$$

The theory of quaternionic analysis which results is modestly successful, though it is little-known. The work of Fueter is described and extended in the papers of Deavours [D] and Sudbery [Su], which include quaternionic versions of Morera's theorem, Cauchy's theorem and the Cauchy integral formula.

As usual, non-commutativity causes immediate algebraic difficulties. If f and g are regular functions, it is easy to see that their sum $f + g$ must also be regular, as must the left scalar multiple qf for $q \in \mathbb{H}$; but their product fg, the composition $f \circ g$ and the right scalar multiple fq need not be. Hence regular functions form a left $\mathbb{H}$-module, but it is difficult to see how to describe any further algebraic structure.

1.3.3 Quaternion Linear Algebra

In the same way as we define real or complex vector spaces, we can define quaternionic vector spaces or $\mathbb{H}$-*modules* — a real vector space with an $\mathbb{H}$-action, which we shall call scalar multiplication. There is the added complication that we need to say whether this multiplication is on the left or the right. We will work with left $\mathbb{H}$-modules — this choice is arbitrary and has only notational effects on the resulting theory. A *left* $\mathbb{H}$*-module* is thus a real vector space U with an action of $\mathbb{H}$ on the left, which we write $(q,u) \mapsto q \cdot u$ or just qu, such that $p(qu) = (pq)u$ for all $p, q \in \mathbb{H}$ and $u \in U$. For example, $\mathbb{H}^n$ is an $\mathbb{H}$-module with the obvious left-multiplication.

Several of the familiar ideas which work for vector spaces over a commutative field work just as well for $\mathbb{H}$-modules. For example, we can define quaternion linear maps between $\mathbb{H}$-modules in the obvious way, and so we can define a dual $\mathbb{H}$-module $U^\times$ of quaternion linear maps $\alpha : U \to \mathbb{H}$.

However, if we try to define quaternion *bi*linear maps we run into trouble. If A, B and C are vector spaces over the commutative field $\mathbb{F}$, then an $\mathbb{F}$-bilinear map $\mu : A \times B \to C$ satisfies $\mu(f_1 a, b) = f_1\mu(a, b)$ and $\mu(a, f_2 b) = f_2\mu(a, b)$ for all $a \in A, b \in B, f_1, f_2 \in \mathbb{F}$. This is equivalent to having $\mu(f_1 a, f_2 b) = f_1 f_2 \mu(a, b)$.

Now suppose that U, V and W are (left) $\mathbb{H}$-modules. Let $q_1, q_2 \in \mathbb{H}$ and let $u \in U$, $v \in V$. If we try to define a bilinear map $\mu : U \times V \to W$ as above, then we need both $\mu(q_1 u, q_2 v) = q_1\mu(u, q_2 v) = q_2 q_1 \mu(u, v)$ and $\mu(q_1 u, q_2 v) = q_2\mu(q_1 u, v) = q_1 q_2 \mu(u, v)$. Since $q_1 q_2 \neq q_2 q_1$ in general, this cannot work.

Similar difficulties arise if we try to define a tensor product over the quaternions. If A and B are vector spaces over the commutative field $\mathbb{F}$, their tensor product over $\mathbb{F}$ is defined by

$$A \otimes_{\mathbb{F}} B = \frac{F(A, B)}{R(A, B)}, \tag{1.21}$$

where $F(A, B)$ is the vector space freely generated (over $\mathbb{F}$) by all elements $(a, b) \in A \times B$ and $R(A, B)$ is the subspace of $F(A, B)$ generated by all elements of the form

$$(a_1 + a_2, b) - (a_1, b) - (a_2, b) \qquad (a, b_1 + b_2) - (a, b_1) - (a, b_2)$$

$$f(a, b) - (f(a), b) \qquad \text{and} \qquad f(a, b) - (a, f(b)),$$

for $a, a_j \in A$, $b, b_j \in B$ and $f \in \mathbb{F}$. Not surprisingly, this process does not work for quaternions. Let U and V be left $\mathbb{H}$-modules. Then if we define the ideal $R(U, V)$ in the same way as above, we discover that $R(U, V)$ is equal to the whole of $F(U, V)$, so $U \otimes_{\mathbb{F}} V = \{0\}$.

One way around both of these difficulties is to demand that our $\mathbb{H}$-modules should also have an $\mathbb{H}$-action on the right. We can now define an $\mathbb{H}$-bilinear map to be one which satisfies $\mu(q_1 u, v q_2) = q_1 \mu(u, v) q_2$, or alternatively $\mu(q_1 u, q_2 v) = q_1 \mu(u, v)\bar{q}_2$. Similarly, for our tensor product we can define a generator $uq \otimes v - u \otimes qv$ (replacing $f(u) \otimes v - u \otimes f(v)$) for R. In this (more restricted) case we do obtain a well-defined 'quaternionic tensor product' $U \otimes_{\mathbb{H}} V = F(U, V)/R(U, V)$ which inherits a left $\mathbb{H}$-action from U and a right $\mathbb{H}$-action from V.

The drawback with this system is that it does not really provide any new insights. If we insist on having a left and a right $\mathbb{H}$-action, we restrict ourselves to talking about $\mathbb{H}$-modules of the form $\mathbb{H}^n = \mathbb{H} \otimes_{\mathbb{R}} \mathbb{R}^n$. In this case, we do get the useful relation $\mathbb{H}^m \otimes_{\mathbb{H}} \mathbb{H}^n \cong \mathbb{H}^{mn}$, but this is only saying that $\mathbb{H} \otimes_{\mathbb{R}} \mathbb{R}^m \otimes_{\mathbb{R}} \mathbb{R}^n \cong \mathbb{H} \otimes_{\mathbb{R}} \mathbb{R}^{mn}$. In this context, our 'quaternionic tensor product' is merely a real tensor product in a quaternionic setting.

1.3.4 Summary

By now we have become familiar with some of the more elementary ups and downs of the quaternions. In the 20^{th} century they have often been viewed as a sort of mathematical Cinderella, more recent techniques being used to describe phenomena which were first thought to be profoundly quaternionic. However, we have seen that it is possible to produce quaternionic analogues of some of the most basic algebraic and analytic ideas of real and complex numbers, often with interesting and useful results. In the next chapter we will continue to explore this process, turning our attention to quaternionic structures in differential geometry.

Chapter 2

Quaternionic Differential Geometry

In this chapter we review some of the ways in which quaternions are used to define geometric stuctures on differentiable manifolds. In the same way that real and complex manifolds are modelled locally by the vector spaces $\mathbb{R}^n$ and $\mathbb{C}^n$ respectively, there are manifolds which can be modelled locally by the $\mathbb{H}$-module $\mathbb{H}^n$. These models work by defining tensors whose action on the tangent spaces to a manifold is the same as the action of the quaternions on an $\mathbb{H}$-module. There are two important classes of manifolds which we shall consider: those which are called 'quaternionic manifolds', and a more restricted class called 'hypercomplex manifolds'.

In this chapter we describe these important geometric structures. We also review the decomposition of exterior forms on complex manifolds, and examine some of the parallels of this theory which have already been found in quaternionic geometry. Many of the algebraic and geometric foundations of the material in this chapter are collected in (or can be inferred from) Fujiki's comprehensive article [F].

2.1 Complex, Hypercomplex and Quaternionic Manifolds

Complex Manifolds

A complex manifold is a $2n$-dimensional real manifold M which admits an atlas of complex charts, all of whose transition functions are holomorphic maps from $\mathbb{C}^n$ to itself. As we saw in Section 1.3, the simplest notions of a 'quaternion-differentiable map from $\mathbb{H}$ to itself' are either very restrictive or too general, and the 'regular functions' of Fueter and Sudbery are not necessarily closed under composition. This makes the notion of 'quaternion-differentiable transition functions' less interesting that one might hope.

An equivalent way to define a complex manifold is by the existence of a special tensor called a complex structure. A complex structure on a real vector space V is an automorphism $I : V \to V$ such that $I^2 = -\operatorname{id}_V$. (It follows that $\dim V$ is even.) The complex structure I gives an isomorphism $V \cong \mathbb{C}^n$, since the operation 'multiplication by i' defines a standard complex structure on $\mathbb{C}^n$. An *almost complex structure* on a $2n$-dimensional real manifold M is a smooth tensor $I \in C^\infty(\operatorname{End}(TM))$ such that I is a complex structure on each of the fibres T_xM.

Now, if M is a complex manifold, each tangent space T_xM is isomorphic to $\mathbb{C}^n$, so taking I to be the standard complex structure on each T_xM defines an almost complex structure on M. An almost complex stucture I which arises in this way is called a complex structure on M, in which case I is said to be integrable. The famous *Newlander-Nirenberg theorem* states that an almost complex structure I is integrable if and only if the Nijenhuis tensor of I

$$N_I(X,Y) = [X,Y] + I[IX,Y] + I[X,IY] - [IX,IY]$$

vanishes for all $X, Y \in C^\infty(TM)$, for all $x \in M$. The Nijenhuis tensor N_I measures the $(0,1)$-component of the Lie bracket of two vector fields of type $(1,0)$. [1] On a complex manifold the Lie bracket of two $(1,0)$-vector fields must also be of type $(1,0)$.

Thus if I is an almost complex structure on M and $N_I \equiv 0$, then M is a complex manifold in the sense of the first definition given above. We can talk about the complex manifold (M,I) if we wish to make the extra geometric structure more explicit — especially as the manifold M might admit more than one complex structure.

Hypercomplex Manifolds

This way of defining a complex manifold adapts itself well to the quaternions. A *hypercomplex structure* on a real vector space V is a triple (I,J,K) of complex structures on V satisfying the equation $IJ = K$. (It follows that $\dim V$ is divisible by 4.) If we identify I, J and K with left-multiplication by i_1, i_2 and i_3, a hypercomplex structure gives an isomorphism $V \cong \mathbb{H}^n$. Equivalently, a hypercomplex structure is defined by a pair of complex structures I and J with $IJ = -JI$. It is easy to see that if (I,J,K) is a hypercomplex structure on V, then each element of the set $\{aI+bJ+cK : a^2+b^2+c^2 = 1\} \cong S^2$ is also a complex structure. We arrive at the following quaternionic version of a complex manifold:

Definition 2.1.1 An *almost hypercomplex structure* on a $4n$-dimensional manifold M is a triple (I,J,K) of almost complex structures on M which satisfy the relation $IJ = K$.

If all of the complex structures are integrable then (I,J,K) is called a *hypercomplex structure* on M, and M is a *hypercomplex manifold*.

A hypercomplex structure on M defines an isomorphism $T_xM \cong \mathbb{H}^n$ at each point $x \in M$. As on a vector space, a hypercomplex structure on a manifold M defines a 2-sphere S^2 of complex structures on M.

Some choices are inherent in this standard definition. A hypercomplex structure as defined above gives TM the structure of a left $\mathbb{H}$-module, since $IJ = K$. This induces a *right* $\mathbb{H}$-module structure on T^*M, using the standard definition $\langle I(\xi), X\rangle = \langle \xi, I(X)\rangle$ *etc.*, for all $\xi \in T^*M, X \in TM$, since on T^*M we now have

$$\langle \xi, IJ(X)\rangle = \langle I(\xi), J(X)\rangle = \langle JI(\xi), X\rangle.$$

In this thesis we will make more use of the hypercomplex structure on T^*M than that on TM. Because of this we will usually define our hypercomplex structures so that $IJ = K$ on T^*M rather than TM. This has only notational effect on the theory, but it does pay to

[1] Tensors of type (p,q) are defined in the next section.

be aware of how it affects other standard notations. For example, for us the hypercomplex structure acts trivially on the anti-self-dual 2-forms $\omega_1^- = dx_0 \wedge dx_1 - dx_2 \wedge dx_3$ *etc.* This encourages us to think of a connection whose curvature is acted upon trivially by the hypercomplex structure as anti-self-dual rather than self-dual, whereas some authors use the opposite convention. Such things are largely a matter of taste — we are choosing to follow the notation of Joyce [J1] for whom regular functions form a left $\mathbb{H}$-module, whereas Sudbery's regular functions form a right $\mathbb{H}$-module.

One important difference between complex and hypercomplex geometry is the existence of a special connection. A complex manifold generally admits many torsion-free connections which preserve the complex structure. By contrast, on a hypercomplex manifold there is a unique torsion-free connection ∇ such that

$$\nabla I = \nabla J = \nabla K = 0.$$

This was proved by Obata in 1956, and the connection ∇ is called the *Obata connection*.

Complex and hypercomplex manifolds can be decribed succinctly in terms of G-structures on manifolds. Let P be the principal frame bundle of M, *i.e.* the $\mathrm{GL}(n,\mathbb{R})$-bundle whose fibre over $x \in M$ is the group of isomorphisms $T_xM \cong \mathbb{R}^{4n}$. Let G be a Lie subgroup of $\mathrm{GL}(n,\mathbb{R})$. A *$G$-structure* Q on M is a principal subbundle of P with structure group G.

Suppose M^{2n} has an almost complex structure. The group of automorphisms of T_xM preserving such a structure is isomorphic to $\mathrm{GL}(n,\mathbb{C})$. Thus an almost complex structure I and a $\mathrm{GL}(n,\mathbb{C})$-structure Q on M contain the same information. The bundle Q admits a torsion-free connection if and only if there is a torsion-free linear connection ∇ on M with $\nabla I = 0$, in which case it is easy to show that I is integrable. Thus a complex manifold is precisely a real manifold M^{2n} with a $\mathrm{GL}(n,\mathbb{C})$-structure Q admitting a torsion-free connection (in which case Q itself is said to be 'integrable').

If M^{4n} has an almost hypercomplex structure then the group of automorphisms preserving this structure is isomorphic to $\mathrm{GL}(n,\mathbb{H})$. Following the same line of argument, a hypercomplex manifold is seen to be a real manifold M with an integrable $\mathrm{GL}(n,\mathbb{H})$-structure Q. The uniqueness of any torsion-free connection on Q follows from analysing the Lie algebra $\mathfrak{gl}(n,\mathbb{H})$. This process is described in [S3, §6].

Quaternionic Manifolds and the Structure Group $\mathrm{GL}(\mathbf{1},\mathbb{H})\mathrm{GL}(\mathbf{n},\mathbb{H})$

Not all of the manifolds which we wish to describe as 'quaternionic' admit hypercomplex structures. For example, the quaternionic projective line $\mathbb{H}P^1$ is diffeomorphic to the 4-sphere S^4. It is well-known that S^4 does not even admit a global almost complex structure; so $\mathbb{H}P^1$ can certainly not be hypercomplex, despite behaving extremely like the quaternions locally.

The reason (and the solution) for this difficulty is that $\mathrm{GL}(n,\mathbb{H})$ is not the largest subgroup of $\mathrm{GL}(4n,\mathbb{R})$ preserving a quaternionic structure. If we think of $\mathrm{GL}(n,\mathbb{H})$ as acting on $\mathbb{H}^n$ by right-multiplication by $n \times n$ quaternionic matrices, then the action of $\mathrm{GL}(n,\mathbb{H})$ commutes with that of the left $\mathbb{H}$-action of the group $\mathrm{GL}(1,\mathbb{H})$. Thus the group of symmetries of $\mathbb{H}^n$ is the product $\mathrm{GL}(1,\mathbb{H}) \times_{\mathbb{R}^*} \mathrm{GL}(n,\mathbb{H})$, which we write $\mathrm{GL}(1,\mathbb{H})\mathrm{GL}(n,\mathbb{H})$. We can multiply on the right by any real multiple of the identity (since $\mathrm{GL}(1,\mathbb{H})$ and $\mathrm{GL}(n,\mathbb{H})$ share the same centre $\mathbb{R}^*$), so without loss of generality we can reduce the first factor to $\mathrm{Sp}(1)$. Thus $\mathrm{GL}(1,\mathbb{H})\mathrm{GL}(n,\mathbb{H})$ is the same as $\mathrm{Sp}(1)\mathrm{GL}(n,\mathbb{H}) = \mathrm{Sp}(1) \times_{\mathbb{Z}_2} \mathrm{GL}(n,\mathbb{H})$.

Definition 2.1.2 [S3, 1.1] A *quaternionic manifold* is a $4n$-dimensional real manifold M ($n \geq 2$) with an $\mathrm{Sp}(1)\mathrm{GL}(n,\mathbb{H})$-structure Q admitting a torsion-free connection.

When $n = 1$ the situation is different, since $\mathrm{Sp}(1)\times\mathrm{Sp}(1) \cong \mathrm{SO}(4)$. In four dimensions we make the special definition that a quaternionic manifold is a self-dual conformal manifold.

In terms of tensors, quaternionic manifolds are a generalisation of hypercomplex manifolds in the following way. Each tangent space T_xM still admits a hypercomplex structure giving an isomorphism $T_xM \cong \mathbb{H}^n$, but this isomorphism does not necessarily arise from globally defined complex structures on M. There is still an S^2-bundle of local almost-complex structures which satisfy $IJ = K$, but it is free to 'rotate'. For a comprehensive study of quaternionic manifolds see [S3] and Chapter 9 of [S4].

Riemannian Manifolds in Complex and Quaternionic Geometry

Suppose the $\mathrm{GL}(n,\mathbb{C})$-structure Q on a complex manifold M admits a further reduction to an integrable $\mathrm{U}(n)$-structure Q'. Then M admits a Riemannian metric g with $g(IX, IY) = g(X, Y)$ for all $X, Y \in T_xM$ for all $x \in M$. We also define the differentiable 2-form $\omega(X, Y) = g(IX, Y)$. If such a metric arises from an integrable $\mathrm{U}(n)$-structure Q' then ω will be a closed 2-form, and M is a symplectic manifold — so M has compatible complex and symplectic structures. In this case M is called a *Kähler manifold*; an integrable $\mathrm{U}(n)$-structure is called a Kähler structure; the metric g is called a Kähler metric and the symplectic form ω is called a Kähler form.

The quaternionic analogue of the compact group $\mathrm{U}(n)$ is the group $\mathrm{Sp}(n)$. A hypercomplex manifold whose $\mathrm{GL}(n,\mathbb{H})$-structure Q reduces to an integrable $\mathrm{Sp}(n)$-structure Q' admits a metric g which is simultaneously Kähler for each of the complex structures I, J and K. Such a manifold is called *hyperkähler* . Using each of the complex structures, we define three independent symplectic forms ω_I, ω_J and ω_K. Then the complex 2-form $\omega_J + i\omega_K$ is holomorphic with respect to the complex structure I, and a hyperkähler manifold has compatible hypercomplex and complex-symplectic structures. Hyperkähler manifolds are studied in [HKLR], which gives a quotient construction for hyperkähler manifolds.

Similarly, if a quaternionic manifold has a metric compatible with the torsion-free $\mathrm{Sp}(1)\mathrm{GL}(n,\mathbb{H})$-structure, then the $\mathrm{Sp}(1)\mathrm{GL}(n,\mathbb{H})$-structure Q reduces to an $\mathrm{Sp}(1)\mathrm{Sp}(n)$-structure Q' and M is said to be *quaternionic Kähler*. The group $\mathrm{Sp}(1)\mathrm{Sp}(n)$ is a maximal proper subgroup of $\mathrm{SO}(4n)$ except when $n = 1$, where as we know $\mathrm{Sp}(1)\mathrm{Sp}(1) = \mathrm{SO}(4)$. In four dimensions a manifold is said to be quaternionic Kähler if and only if it is self-dual and Einstein. Quaternionic Kähler manifolds are the subject of [S2].

2.2 Differential Forms on Complex Manifolds

This section consists of background material in complex geometry, especially ideas which encourage the development of interesting quaternionic versions. More information on this and other aspects of complex geometry can be found in [W] and [GH, §0.2].

Let (M, I) be a complex manifold. Then I gives TM and T^*M the structure of a $\mathrm{U}(1)$-representation and the complexification $\mathbb{C}\otimes_{\mathbb{R}} TM$ splits into two weight spaces with weights $\pm i$. The same is true of $\mathbb{C}\otimes_{\mathbb{R}} T^*M$. There are various ways of writing these weight spaces; fairly standard is the notation $\mathbb{C}\otimes_{\mathbb{R}} T^*M = T^*_{1,0}M \oplus T^*_{0,1}M$, where these summands are the $+i$ and $-i$ eigenspaces of I respectively. However, for our purposes it will be more useful to adopt the notation of [S3], writing

$$\mathbb{C}\otimes_{\mathbb{R}} T^*M = \Lambda^{1,0}M \oplus \Lambda^{0,1}M, \tag{2.1}$$

so $\Lambda^{1,0}M = T^*_{1,0}M$ and $\Lambda^{0,1}M = T^*_{0,1}M$. A holomorphic function $f \in C^\infty(M, \mathbb{C})$ is a smooth function whose derivative takes values only in $\Lambda^{1,0}M$ for all $m \in M$, *i.e.* f is holomorphic if and only if $df \in C^\infty(M, \Lambda^{1,0}M)$. A closely linked statement is that $\Lambda^{1,0}M$ is a holomorphic vector bundle. Thus $\Lambda^{1,0}M$ is called the holomorphic cotangent space and $\Lambda^{0,1}M$ is called the antiholomorphic cotangent space of M. If we reverse the sense of the complex structure (*i.e.* if we swap I for $-I$) then we reverse the roles of the holomorphic and antiholomorphic spaces, which is why a function which is holomorphic with respect to I is antiholomorphic with respect to $-I$.

From standard multilinear algebra, the decomposition (2.1) gives rise to a decomposition of exterior forms of all powers

$$\mathbb{C}\otimes_{\mathbb{R}} \Lambda^k T^*M = \bigoplus_{p=0}^{k} \Lambda^p(T^*_{1,0}M) \otimes \Lambda^{k-p}(T^*_{0,1}M).$$

Define the bundle

$$\Lambda^{p,q}M = \Lambda^p T^*_{1,0}M \otimes \Lambda^q T^*_{0,1}M. \tag{2.2}$$

With this notation, Equation (2.1) is an example of the more general decomposition into types

$$\mathbb{C}\otimes_{\mathbb{R}} \Lambda^k T^*M = \bigoplus_{p+q=k} \Lambda^{p,q}M. \tag{2.3}$$

A smooth section of the bundle $\Lambda^{p,q}M$ is called a differential form of type (p, q) or just a (p, q)-form. We write $\Omega^{p,q}(M)$ for the set of (p, q)-forms on M, so

$$\Omega^{p,q}(M) = C^\infty(M, \Lambda^{p,q}M) \qquad \text{and} \qquad \Omega^k(M) = \bigoplus_{p+q=k} \Omega^{p,q}(M).$$

Define two first-order differential operators,

$$\begin{array}{l} \partial : \Omega^{p,q}(M) \to \Omega^{p+1,q}(M) \\ \partial = \pi^{p+1,q} \circ d \end{array} \quad \text{and} \quad \begin{array}{l} \overline{\partial} : \Omega^{p,q}(M) \to \Omega^{p,q+1}(M) \\ \overline{\partial} = \pi^{p,q+1} \circ d, \end{array} \tag{2.4}$$

where $\pi^{p,q}$ denotes the natural projection from $\mathbb{C}\otimes\Lambda^k M$ onto $\Lambda^{p,q}M$. The operator $\overline{\partial}$ is called the Dolbeault operator.

These definitions rely only on the fact that I is an almost complex structure. A crucial fact is that if I is integrable, these are the *only* two components represented in

Figure 2.1: The Dolbeault Complex

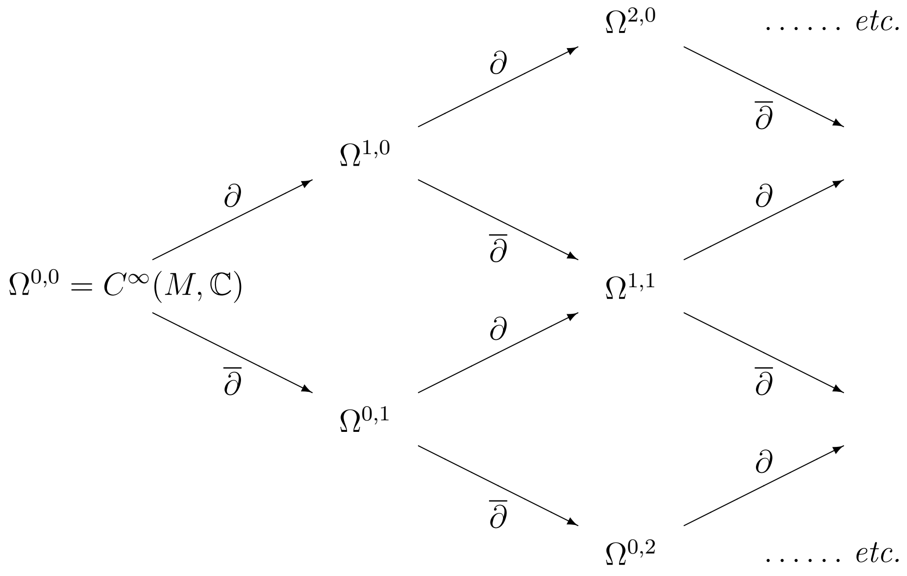

the exterior differential d, *i.e.* $d = \partial + \overline{\partial}$ [WW, Proposition 2.2.2, p.105]. An immediate consequence of this is that on a complex manifold M,

$$\partial^2 = \partial\overline{\partial} + \overline{\partial}\partial = \overline{\partial}^2 = 0. \tag{2.5}$$

This gives rise to the *Dolbeault complex.* Writing $\overline{\partial}^{p,q}$ for the particular map $\overline{\partial} : \Omega^{p,q} \longrightarrow \Omega^{p,q+1}$, we define the Dolbeault cohomology groups

$$H^{p,q}_{\overline{\partial}} = \frac{\mathrm{Ker}(\overline{\partial}^{p,q})}{\mathrm{Im}(\overline{\partial}^{p,q-1})}.$$

A function $f \in C^\infty(M,\mathbb{C})$ is holomorphic if and only if $\overline{\partial} f = 0$ and for this reason $\overline{\partial}$ is sometimes called the Cauchy-Riemann operator. Similarly, a $(p, 0)$-form α is said to be holomorphic if and only if $\overline{\partial}\alpha = 0$.

A useful way to think of the Dolbeault complex is as a decomposition of $\mathbb{C} \otimes \Lambda^k T^*M$ into types of U(1)-representation. The representations of U(1) on complex vector spaces are particularly easy to understand. Since U(1) is abelian, the irreducible representations all are one-dimensional. They are parametrised by the integers, taking the form

$$\varrho_n : \mathrm{U}(1) \to \mathrm{GL}(1,\mathbb{C}) = \mathbb{C}^* \qquad \varrho_n : e^{i\theta} \to e^{ni\theta}$$

for some $n \in \mathbb{Z}$. Representations of the Lie algebra $\mathfrak{u}(1)$ are then of the form $d\varrho_n : z \to nz$. The Dolbeault complex begins with the representation of the complex structure $\langle I \rangle \cong \mathfrak{u}(1)$ on $\mathbb{C} \otimes T^*M$. This induces a representation on $\mathbb{C} \otimes \Lambda^\bullet T^*M$. It is easy to see that if $\omega \in \Lambda^{p,q}M$, $I(\omega) = i(p - q)\omega$. In other words, $\Lambda^{p,q}M$ is the bundle of U(1)-representations of the type ϱ_{p-q}.

2.3 Differential Forms on Quaternionic Manifolds

Any G-structure on a manifold M induces a representation of G on the exterior algebra of M. Fujiki's account of this [F, §2] explains many quaternionic analogues of complex and Kähler geometry.

The decomposition of differential forms on quaternionic Kähler manifolds began by considering the *fundamental 4-form*

$$\Omega = \omega_I \wedge \omega_I + \omega_J \wedge \omega_J + \omega_K \wedge \omega_K,$$

where ω_I, ω_J and ω_K are the local Kähler forms associated to local almost complex structures I, J and K with $IJ = K$. The fundamental 4-form is globally defined and invariant under the induced action of Sp(1)Sp(n) on $\Lambda^4 T^*M$. Kraines [Kr] and Bonan [Bon] used the fundamental 4-form to decompose the space $\Lambda^k T^*M$ in a similar way to the Lefschetz decomposition of differential forms on a Kähler manifold [GH, p. 122]. A differential k-form μ is said to be *effective* if $\Omega \wedge * \mu = 0$, where $* : \Lambda^k T^*M \to \Lambda^{4n-k} T^*M$ is the Hodge star. This leads to the following theorem:

Theorem 2.3.1 *[Kr, Theorem 3.5][Bon, Theorem 2]*
For $k \leq 2n+2$, every every k-form ϕ admits a unique decomposition

$$\phi = \sum_{0 \leq j \leq k/4} \Omega^j \wedge \mu_{k-4j},$$

where the μ_{k-4j} are effective $(k-4j)$-forms.

Bonan further refines this decomposition for quaternion-valued forms, using exterior multiplication by the globally defined quaternionic 2-form $\Psi = i_1\omega_I + i_2\omega_J + i_3\omega_K$. Note that $\Psi \wedge \Psi = -2\Omega$.

Another way to consider the decomposition of forms on a quaternionic manifold is as representations of the group Sp(1)GL$(n,\mathbb{H})$. We express the Sp(1)GL$(n,\mathbb{H})$-representation on $\mathbb{H}^n$ by writing

$$\mathbb{H}^n \cong V_1 \otimes E, \tag{2.6}$$

where V_1 is the basic representation of Sp(1) on $\mathbb{C}^2$ and E is the basic representation of GL$(n,\mathbb{H})$ on $\mathbb{C}^{2n}$. (This uses the standard convention of working with complex representations, which in the presence of suitable structure maps can be thought of as complexified real representations. In this case, the structure map is the tensor product of the quaternionic structures on V_1 and E.)

This representation also describes the (co)tangent bundle of a quaternionic manifold in the following way. Let P be a principal G-bundle over the differentiable manifold M and let W be a G-module. We define the *associated bundle*

$$\mathbf{W} = P \times_G W = \frac{P \times W}{G},$$

where $g \in G$ acts on $(p,w) \in P \times W$ by $(p,w)\cdot g = (f \cdot g, g^{-1} \cdot w)$. Then $\mathbf{W}$ is a vector bundle over M with fibre W. We will usually just write W for $\mathbf{W}$, relying on context to distinguish between the bundle and the representation. Following Salamon [S3, §1], if M^{4n} is a quaternionic manifold with $\mathrm{Sp}(1)\mathrm{GL}(n, \mathbb{H})$-structure Q, then the cotangent bundle is a vector bundle associated to the principal bundle Q and the representation $V_1 \otimes E$, so that we write

$$(T^*M)^{\mathbb{C}} \cong V_1 \otimes E \tag{2.7}$$

(though we will usually omit the complexification sign). This induces an $\mathrm{Sp}(1)\mathrm{GL}(n, \mathbb{H})$-action on the bundle of exterior k-forms $\Lambda^k T^*M$,

$$\Lambda^k T^*M \cong \Lambda^k(V_1 \otimes E) \cong \bigoplus_{j=0}^{[k/2]} S^{k-2j}(V_1) \otimes L_j^k \cong \bigoplus_{j=0}^{[k/2]} V_{k-2j} \otimes L_j^k, \tag{2.8}$$

where L_j^k is an irreducible representation of $\mathrm{GL}(n, \mathbb{H})$. This decomposition is given by Salamon [S3, §4], along with more details concerning the nature of the $\mathrm{GL}(n, \mathbb{H})$ representations L_j^k. If M is quaternion Kähler, $\Lambda^k T^*M$ can be further decomposed into representations of the compact group $\mathrm{Sp}(1)\mathrm{Sp}(n)$. This refinement is performed in detail by Swann [Sw], and used to demonstrate significant results.

If we symmetrise completely on V_1 in Equation (2.8) to obtain V_k, we must antisymmetrise completely on E. Salamon therefore defines the irreducible subspace

$$A^k \cong V_k \otimes \Lambda^k E. \tag{2.9}$$

The bundle A^k can be described using the decomposition into types for the local almost complex structures on M as follows [S3, Proposition 4.2]: [2]

$$A^k = \sum_{I \in S^2} \Lambda_I^{k,0} M. \tag{2.10}$$

Letting p denote the natural projection $p : \Lambda^k T^*M \to A^k$ and setting $D = p \circ d$, we define a sequence of differential operators

$$0 \longrightarrow C^\infty(A^0) \xrightarrow{D=d} C^\infty(A^1 = T^*M) \xrightarrow{D} C^\infty(A^2) \xrightarrow{D} \dots \xrightarrow{D} C^\infty(A^{2n}) \longrightarrow 0. \tag{2.11}$$

This is accomplished using only the fact that M has an $\mathrm{Sp}(1)\mathrm{GL}(n, \mathbb{H})$-structure; such a manifold is called 'almost quaternionic'. The following theorem of Salamon relates the integrability of such a structure with the sequence of operators in (2.11):

Theorem 2.3.2 *[S3, Theorem 4.1] An almost quaternionic manifold is quaternionic if and only if (2.11) is a complex.*

This theorem is analogous to the familiar result in complex geometry that an almost complex structure on a manifold is integrable if and only if $\overline{\partial}^2 = 0$.

[2] This is because every Sp(1)-representation V_n is generated by its highest weight spaces taken with respect to all the different linear combinations of I, J and K. We will use such descriptions in detail in later chapters, and find that they play a prominent role in quaternionic algebra.

Chapter 3

A Double Complex on Quaternionic Manifolds

Until now we have been discussing known material. In this chapter we present a discovery which as far as the author can tell is new — a double complex of exterior forms on quaternionic manifolds. We argue that this is the best quaternionic analogue of the Dolbeault complex. The 'top row' of this double complex is exactly the complex (2.11) discovered by Simon Salamon, which plays a similar role to that of the $(k, 0)$-forms on a complex manifold.

The new double complex is obtained by simplifying the Bonan decomposition of Equation (2.8). Instead of using the more complicated structure of $\Lambda^k T^*M$ as an $\mathrm{Sp}(1)\mathrm{GL}(n, \mathbb{H})$-module, we consider only the action of the Sp(1)-factor and decompose $\Lambda^k T^*M$ into irreducible Sp(1)-representations — a fairly easy process achieved by considering weights. The resulting decomposition gives rise to a double complex through the Clebsch-Gordon formula, in particular the isomorphism $V_r \otimes V_1 \cong V_{r+1} \oplus V_{r-1}$. This encourages us to define two 'quaternionic Dolbeault' operators $\mathcal{D}$ and $\overline{\mathcal{D}}$, and leads to new cohomology groups on quaternionic manifolds.

The main geometric difference between this double complex and the Dolbeault complex is that whilst the Dolbeault complex is a diamond, our double complex forms an isosceles triangle, as if the diamond is 'folded in half'. This is more like the decomposition of *real*-valued forms on complex manifolds.

Determining where our double complex is elliptic has been far more difficult than for the de Rham or Dolbeault complexes. The ellipticity properties of our double complex are more similar to those of a real-valued version of the Dolbeault complex. Because of this similarity, we shall begin with a discussion of real-valued forms on complex manifolds.

3.1 Real forms on Complex Manifolds

Let M be a complex manifold and let $\omega \in \Lambda^{p,q} = \Lambda^{p,q}M$. Then $\overline{\omega} \in \Lambda^{q,p}$, and so $\omega + \overline{\omega}$ is a real-valued exterior form in $(\Lambda^{p,q} \oplus \Lambda^{q,p})_{\mathbb{R}}$, where the subscript $\mathbb{R}$ denotes the fact that we are talking about real forms. The space $(\Lambda^{p,q} \oplus \Lambda^{q,p})_{\mathbb{R}}$ is a real vector bundle associated to the principal $\mathrm{GL}(n, \mathbb{C})$-bundle defined by the complex structure. This gives a decomposition of real-valued exterior forms, [1]

$$\Lambda^k_{\mathbb{R}} T^*M = \bigoplus_{\substack{p+q=k \\ p>q}} (\Lambda^{p,q} \oplus \Lambda^{q,p})_{\mathbb{R}} \oplus \Lambda_{\mathbb{R}}^{\frac{k}{2},\frac{k}{2}}. \tag{3.1}$$

The condition $p > q$ ensures that we have no repetition. The bundle $\Lambda_{\mathbb{R}}^{\frac{k}{2},\frac{k}{2}}$ only appears when k is even. It is its own conjugate and so naturally a real vector bundle associated to the trivial representation of U(1).

Let $\omega + \overline{\omega} \in (\Omega^{p,q} \oplus \Omega^{q,p})_{\mathbb{R}}$. Then $\partial\omega + \overline{\partial}\overline{\omega} \in (\Omega^{p+1,q} \oplus \Omega^{q,p+1})_{\mathbb{R}}$. Call this operator $\partial \oplus \overline{\partial}$. Then

$$\begin{aligned} d(\omega + \overline{\omega}) &= (\partial \oplus \overline{\partial})(\omega + \overline{\omega}) + (\overline{\partial} \oplus \partial)(\omega + \overline{\omega}) \\ &\in (\Omega^{p+1,q} \oplus \Omega^{q,p+1})_{\mathbb{R}} \oplus (\Omega^{p,q+1} \oplus \Omega^{q+1,p})_{\mathbb{R}}. \end{aligned}$$

When $p = q$, $\omega = \overline{\omega}$, so $\partial \oplus \overline{\partial} = \overline{\partial} \oplus \partial = d$, and there is only one differential operator acting on $\Omega^{p,p}$. This gives the following double complex (where the real subscripts are omitted for convenience).

Figure 3.1: The Real Dolbeault Complex

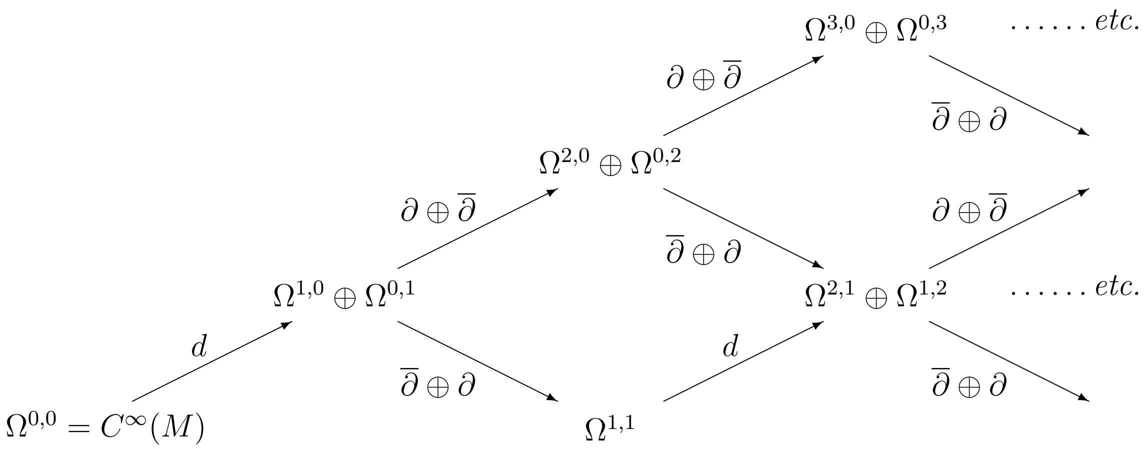

Thus there is a double complex of real forms on a complex manifold, obtained by decomposing $\Lambda^k_x T^*M$ into subrepresentations of the action of $\mathfrak{u}(1) = \langle I \rangle$, induced from the action on $T^*_x M$. This double complex is less well-behaved than its complex-valued counterpart; in particular, it is not elliptic everywhere. We shall now show what this means, and why it is significant.

Ellipticity

Whether or not a differential complex on a manifold is 'elliptic' is an important question, with striking topological, algebraic and physical consequences. Examples of elliptic

[1]This decomposition is also given by Reyes-Carión [R, §3.1], who calls the bundle $(\Lambda^{p,q} \oplus \Lambda^{q,p})_{\mathbb{R}}$ $[[\Lambda^{p,q}]]$.

complexes include the de Rham and Dolbeault complexes. A thorough description of elliptic operators and elliptic complexes can be found in [W, Chapter 5].

Let E and F be vector bundles over M, and let $\Phi : C^\infty(E) \to C^\infty(F)$ be a differential operator. (We will be working with first-order operators, and so will only describe this case.) At every point $x \in M$ and for every nonzero $\xi \in T^*_x M$, we define a linear map $\sigma_\Phi(x,\xi) : E_x \to F_x$ called the principal symbol of Φ, as follows. Let $\epsilon \in C^\infty(E)$ with $\epsilon(x) = e$, and let $f \in C^\infty(M)$ with $f(x) = 0$, $df(x) = \xi$. Then

$$\sigma_\Phi(x,\xi)e = \Phi(f\epsilon)|_x.$$

In coordinates, σ is often found by replacing the operator $\frac{\partial}{\partial x^j}$ with exterior multiplication by a cotangent vector ξ^j dual to $\frac{\partial}{\partial x^j}$. For example, for the exterior differential $d : \Omega^k(M) \to \Omega^{k+1}(M)$ we have $\sigma_d(x,\xi)\omega = \omega \wedge \xi$. The operator Φ is said to be *elliptic* at $x \in M$ if the symbol $\sigma_\Phi(x,\xi) : E_x \to F_x$ is an isomorphism for all nonzero $\xi \in T^*_x M$.

A complex $0 \xrightarrow{\Phi_0} C^\infty(E_0) \xrightarrow{\Phi_1} C^\infty(E_1) \xrightarrow{\Phi_2} C^\infty(E_2) \xrightarrow{\Phi_3} \dots \xrightarrow{\Phi_n} C^\infty(E_n) \xrightarrow{\Phi_{n+1}} 0$ is said to be *elliptic at* E_i if the symbol sequence $E_{i-1} \xrightarrow{\sigma_{\Phi_i}} E_i \xrightarrow{\sigma_{\Phi_{i+1}}} E_{i+1}$ is exact for all $\xi \in T^*_x M$ and for all $x \in M$. The link between these two forms of ellipticity is as follows. If we have a metric on each E_i then we can define a formal adjoint $\Phi^*_i : E_i \to E_{i-1}$. Linear algebra reveals that the complex is elliptic at E_i if and only if the Laplacian $\Phi^*_i\Phi_i + \Phi_{i-1}\Phi^*_{i-1}$ is an elliptic operator.

One important implication of this is that an elliptic complex on a compact manifold has finite-dimensional cohomology groups [W, Theorem 5.2, p. 147]. Whether a complex yields interesting cohomological information is in this way directly related to whether or not the complex is elliptic. The following Proposition answers this question for the real Dolbeault complex.

Proposition 3.1.1 *For $p > 0$, the upward complex*

$$0 \longrightarrow \Omega^{p,p} \longrightarrow \Omega^{p+1,p} \oplus \Omega^{p,p+1} \longrightarrow \Omega^{p+2,p} \oplus \Omega^{p,p+2} \longrightarrow \dots$$

is elliptic everywhere except at the first two spaces $\Omega^{p,p}$ and $\Omega^{p+1,p} \oplus \Omega^{p,p+1}$.

For $p = 0$, the 'leading edge' complex

$$0 \longrightarrow \Omega^{0,0} \longrightarrow \Omega^{1,0} \oplus \Omega^{0,1} \longrightarrow \Omega^{2,0} \oplus \Omega^{0,2} \longrightarrow \dots$$

is elliptic everywhere except at $\Omega^{1,0} \oplus \Omega^{0,1} = \Omega^1(M)$.

Proof. When $p > q + 1$, we have short sequences of the form

$$\begin{array}{ccccc} \Omega^{p-1,q} & \xrightarrow{\partial} & \Omega^{p,q} & \xrightarrow{\partial} & \Omega^{p+1,q} \\ \bigoplus & & \bigoplus & & \bigoplus \\ \Omega^{q,p-1} & \xrightarrow{\bar\partial} & \Omega^{q,p} & \xrightarrow{\bar\partial} & \Omega^{q,p+1}. \end{array} \tag{3.2}$$

Each such sequence is (a real subspace of) the direct sum of two elliptic sequences, and so is elliptic. Thus we have ellipticity at $\Omega^{p,q} \oplus \Omega^{q,p}$ whenever $p \geq q + 2$.

This leaves us to consider the case when $p = q$, and the sequence

$$0 \longrightarrow \Omega^{p,p} \begin{array}{c} \overset{\partial}{\nearrow} \\ \\ \underset{\bar\partial}{\searrow} \end{array} \begin{array}{ccccc} \Omega^{p+1,p} & \xrightarrow{\partial} & \Omega^{p+2,2} & \longrightarrow & \ldots etc. \\ \bigoplus & & \bigoplus & & \\ \Omega^{p,p+1} & \xrightarrow{\bar\partial} & \Omega^{2,p+2} & \longrightarrow & \ldots etc. \end{array} \tag{3.3}$$

This fails to be elliptic. An easy and instructive way to see this is to consider the simplest 4-dimensional example $M = \mathbb{C}^2$.

Let e^0, $e^1 = I(e^0)$, e^2 and $e^3 = I(e^2)$ form a basis for $T^*_x\mathbb{C}^2 \cong \mathbb{C}^2$, and let $e^{ab\dots}$ denote $e^a \wedge e^b \wedge \dots etc.$ Then $I(e^{01}) = e^{00} - e^{11} = 0$, so $e^{01} \in \Lambda^{1,1}$. The map from $\Lambda^{1,1}$ to $\Lambda^{2,1} \oplus \Lambda^{1,2}$ is just the exterior differential d. Since $\sigma_d(x, e^0)(e^{01}) = e^{01} \wedge e^0 = 0$ the symbol map $\sigma_d : \Lambda^{1,1} \to \Lambda^{2,1} \oplus \Lambda^{1,2}$ is not injective, so the symbol sequence is not exact at $\Lambda^{1,1}$.

Consider also $e^{123} \in \Lambda^{2,1} \oplus \Lambda^{1,2}$. Then $\sigma_{\partial\oplus\bar\partial}(x, e^0)(e^{123}) = 0$, since there is no bundle $\Lambda^{3,1} \oplus \Lambda^{1,3}$. But e^{123} has no e^0-factor, so is not the image under $\sigma_d(x, e^0)$ of any form $\alpha \in \Lambda^{1,1}$. Thus the symbol sequence fails to be exact at $\Lambda^{2,1} \oplus \Lambda^{1,2}$.

It is a simple matter to extend these counterexamples to higher dimensions and higher exterior powers. For $k = 0$, the situation is different. It is easy to show that the complex

$$0 \longrightarrow C^\infty(M) \xrightarrow{d} \Omega^{1,0} \oplus \Omega^{0,1} \xrightarrow{\partial\oplus\bar\partial} \Omega^{2,0} \oplus \Omega^{0,2} \longrightarrow \dots etc.$$

is elliptic everywhere except at $(\Omega^{1,0} \oplus \Omega^{0,1})$. ∎

This last sequence is given particular attention by Reyes-Carión [R, Lemma 2]. He shows that, when M is Kähler, ellipticity can be regained by adding the space $\langle\omega\rangle$ to the bundle $\Lambda^{2,0} \oplus \Lambda^{0,2}$, where ω is the real Kähler $(1,1)$-form.

The Real Dolbeault complex is thus elliptic except at the bottom of the isosceles triangle of spaces. Here the projection from $d(\Omega^{p,p})$ to $\Omega^{p+1,p} \oplus \Omega^{p,p+1}$ is the identity, and arguments based upon non-trivial projection maps no longer apply. We shall see that this situation is closely akin to that of differential forms on quaternionic manifolds, and that techniques motivated by this example yield similar results.

3.2 Construction of the Double Complex

Let M^{4n} be a quaternionic manifold. Then $T^*_xM \cong V_1 \otimes E$ as an $\mathrm{Sp}(1)\mathrm{GL}(n, \mathbb{H})$-representation for all $x \in M$. Suppose we consider just the action of the Sp(1)-factor. Then the (complexified) cotangent space effectively takes the form $V_1 \otimes \mathbb{C}^{2n} \cong 2nV_1$. Thus the Sp(1)-action on Λ^kT^*M is given by the representation $\Lambda^k(2nV_1)$.

To work out the irreducible decomposition of this representation we compute the weight space decomposition of $\Lambda^k(2nV_1)$ from that of $2nV_1$. [2] With respect to the action of a particular subgroup $\mathrm{U}(1)_q \subset \mathrm{Sp}(1)$, the representation $2nV_1$ has weights $+1$ and -1, each occuring with multiplicity $2n$. The weights of $\Lambda^k(2nV_1)$ are the k-wise distinct sums of these. Each weight r in $\Lambda^k(2nV_1)$ must therefore be a sum of p occurences of the weight '+1' and $p-r$ occurences of the weight '−1', where $2p-r = k$

[2] This process for calculating the weights of tensor, symmetric and exterior powers is a standard technique in representation theory — see for example [FH, §11.2].

and $0 \leq p \leq k$ (from which it follows immediately that $-k \leq r \leq k$ and $r \equiv k \bmod 2$). The number of ways to choose the p '+1' weights is $\binom{2n}{p}$, and the number of ways to choose the $(p-r)$ '−1' weights is $\binom{2n}{p-r}$, so the multiplicity of the weight r in the representation $\Lambda^k(2nV_1)$ is

$$\text{Mult}(r) = \binom{2n}{\frac{k+r}{2}}\binom{2n}{\frac{k-r}{2}}.$$

For $r \geq 0$, consider the difference $\text{Mult}(r) - \text{Mult}(r+2)$. This is the number of weight spaces of weight r which do not have any corresponding weight space of weight $r+2$. Each such weight space must therefore be the highest weight space in an irreducible subrepresentation $V_r \subseteq \Lambda^k T^*M$, from which it follows that the number of irreducibles V_r in $\Lambda^k(2nV_1)$ is equal to $\text{Mult}(r) - \text{Mult}(r+2)$. This demonstrates the following Proposition:

Proposition 3.2.1 *Let M^{4n} be a hypercomplex manifold. The decomposition into irreducibles of the induced representation of* $\text{Sp}(1)$ *on $\Lambda^k T^*M$ is*

$$\Lambda^k T^*M \cong \bigoplus_{r=0}^{k} \left[\binom{2n}{\frac{k+r}{2}}\binom{2n}{\frac{k-r}{2}} - \binom{2n}{\frac{k+r+2}{2}}\binom{2n}{\frac{k-r-2}{2}} \right] V_r,$$

where $r \equiv k \bmod 2$.

We will not always write the condition $r \equiv k \bmod 2$, assuming that $\binom{p}{q} = 0$ if $q \notin \mathbb{Z}$.

Definition 3.2.2 Let M^{4n} be a quaternionic manifold. Define the coefficient

$$\epsilon^n_{k,r} = \binom{2n}{\frac{k+r}{2}}\binom{2n}{\frac{k-r}{2}} - \binom{2n}{\frac{k+r+2}{2}}\binom{2n}{\frac{k-r-2}{2}},$$

and let $E_{k,r}$ be the vector bundle associated to the $\text{Sp}(1)$-representation $\epsilon^n_{k,r} V_r$. With this notation Proposition 3.2.1 is written

$$\Lambda^k T^*M \cong \bigoplus_{r=0}^{k} \epsilon^n_{k,r} V_r = \bigoplus_{r=0}^{k} E_{k,r}.$$

Our most important result is that this decomposition gives rise to a double complex of differential forms and operators on a quaternionic manifold.

Theorem 3.2.3 *The exterior derivative d maps $C^\infty(M, E_{k,r})$ to $C^\infty(M, E_{k+1,r+1} \oplus E_{k+1,r-1})$.*

Proof. Let ∇ be a torsion-free linear connection on M preserving the quaternionic structure. Then $\nabla : C^\infty(M, E_{k,r}) \to C^\infty(M, E_{k,r} \otimes T^*M)$. As $\text{Sp}(1)$-representations, this is

$$\nabla : C^\infty(M, \epsilon^n_{k,r} V_r) \to C^\infty(M, \epsilon^n_{k,r} V_r \otimes 2nV_1).$$

Using the Clebsch-Gordon formula we have $\epsilon^n_{k,r} V_r \otimes 2nV_1 \cong 2n\epsilon^n_{k,r}(V_{r+1} \oplus V_{r-1})$. Thus the image of $E_{k,r}$ under ∇ is contained in the V_{r+1} and V_{r-1} summands of $\Lambda^k T^*M \otimes T^*M$. Since ∇ is torsion-free, $d = \wedge \circ \nabla$, so d maps (sections of) $E_{k,r}$ to the V_{r+1} and V_{r-1} summands of $\Lambda^{k+1}T^*M$. Thus $d : C^\infty(M, E_{k,r}) \to C^\infty(M, E_{k+1,r+1} \oplus E_{k+1,r-1})$. ∎

This allows us to split the exterior differential d into two 'quaternionic Dolbeault operators'.

Definition 3.2.4 Let $\pi_{k,r}$ be the natural projection from $\Lambda^k T^*M$ onto $E_{k,r}$. Define the operators

$$\begin{array}{lcl} \mathcal{D} : C^\infty(E_{k,r}) \to C^\infty(E_{k+1,r+1}) & & \overline{\mathcal{D}} : C^\infty(E_{k,r}) \to C^\infty(E_{k+1,r-1}) \\ \mathcal{D} = \pi_{k+1,r+1} \circ d & \text{and} & \overline{\mathcal{D}} = \pi_{k+1,r-1} \circ d. \end{array} \tag{3.4}$$

Theorem 3.2.3 is equivalent to the following:

Proposition 3.2.5 *On a quaternionic manifold M, we have $d = \mathcal{D} + \overline{\mathcal{D}}$, and so*

$$\mathcal{D}^2 = \mathcal{D}\overline{\mathcal{D}} + \overline{\mathcal{D}}\mathcal{D} = \overline{\mathcal{D}}^2 = 0.$$

Proof. The first equation is equivalent to Theorem 3.2.3. The rest follows immediately from decomposing the equation $d^2 = 0$. ∎

Figure 3.2: The Double Complex

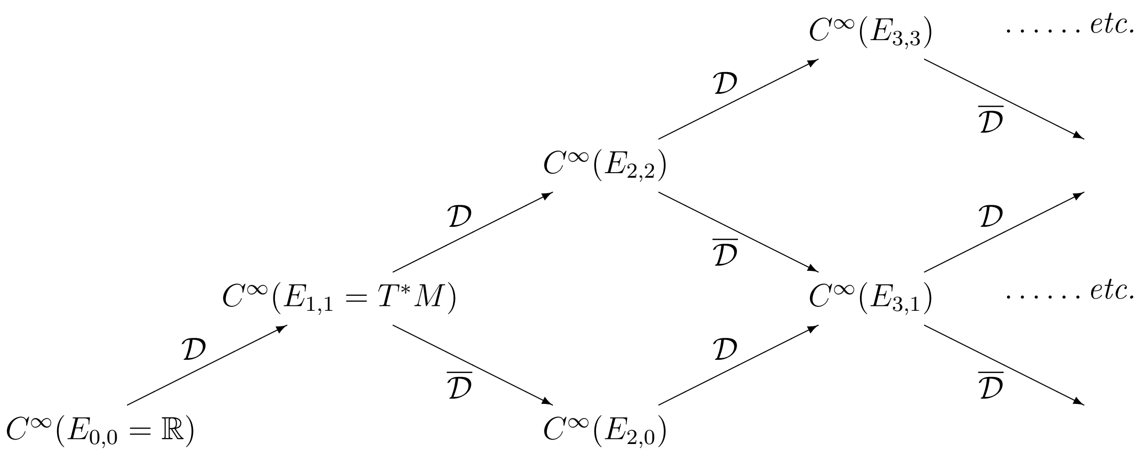

Here is our quaternionic analogue of the Dolbeault complex. There are strong similarities between this and the Real Dolbeault complex (Figure 3.1). Again, instead of a diamond as in the Dolbeault complex, the quaternionic version only extends upwards to form an isosceles triangle. This is essentially because there is one irreducible U(1)-representation for each integer, whereas there is one irreducible Sp(1)-representation only for each nonnegative integer. Note that this is a decomposition of real as well as complex differential forms; the operators $\mathcal{D}$ and $\overline{\mathcal{D}}$ map real forms to real forms.

By definition, the bundle $E_{k,k}$ is the bundle A^k of (2.9) — they are both the subbundle of $\Lambda^k T^*M$ which includes all Sp(1)-representations of the form V_k. Thus the leading edge of the double complex

$$0 \longrightarrow C^\infty(E_{0,0}) \xrightarrow{\mathcal{D}} C^\infty(E_{1,1}) \xrightarrow{\mathcal{D}} C^\infty(E_{2,2}) \xrightarrow{\mathcal{D}} \ldots \xrightarrow{\mathcal{D}} C^\infty(E_{2n,2n}) \xrightarrow{\mathcal{D}} 0$$

is precisely the complex (2.11) discovered by Salamon.

Example 3.2.6 Four Dimensions

This double complex is already very well-known and understood in four dimensions. Here there is a splitting only in the middle dimension, $\Lambda^2 T^*M \cong V_2 \oplus 3V_0$. Let I, J and K be local almost complex structures at $x \in M$, and let $e^0 \in T_x^*M$. Let $e^1 = I(e^0)$, $e^2 = J(e^0)$ and $e^3 = K(e^0)$. In this way we obtain a basis $\{e^0, \ldots, e^3\}$ for $T_x^*M \cong \mathbb{H}$. Using the notation $e^{ijk\cdots} = e^i \wedge e^j \wedge e^k \wedge \ldots etc.$, define the 2-forms

$$\omega_1^\pm = e^{01} \pm e^{23}, \quad \omega_2^\pm = e^{02} \pm e^{31}, \quad \omega_3^\pm = e^{03} \pm e^{12}. \tag{3.5}$$

Then I, J and K all act trivially [3] on the ω_j^-, so $E_{2,0} = \langle \omega_1^-, \omega_2^-, \omega_3^- \rangle$. The action of $\mathfrak{sp}(1)$ on the ω_j^+ is given by the multiplication table

$$\begin{array}{rclrclrcl} I(\omega_1^+) & = & 0 & I(\omega_2^+) & = & 2\omega_3^+ & I(\omega_3^+) & = & -2\omega_3^+ \\ J(\omega_1^+) & = & -2\omega_3^+ & J(\omega_2^+) & = & 0 & J(\omega_3^+) & = & 2\omega_1^+ \\ K(\omega_1^+) & = & 2\omega_2^+ & K(\omega_2^+) & = & -2\omega_1^+ & K(\omega_3^+) & = & 0. \end{array} \tag{3.6}$$

These are the relations of the irreducible $\mathfrak{sp}(1)$-representation V_2, and we see that $E_{2,2} = \langle \omega_1^+, \omega_2^+, \omega_3^+ \rangle$.

These bundles will be familiar to most readers: $E_{2,2}$ is the bundle of *self-dual* 2-forms Λ_+^2 and $E_{2,0}$ is the bundle of *anti-self-dual* 2-forms Λ_-^2. The celebrated splitting $\Lambda^2 T^*M \cong \Lambda_+^2 \oplus \Lambda_-^2$ is an invariant of the conformal class of any Riemannian 4-manifold, and $I^2 + J^2 + K^2 = -4(* + 1)$, where $* : \Lambda^k T^*M \to \Lambda^{4-k} T^*M$ is the Hodge star map.

This also serves to illustrate why in four dimensions we make the definition that a quaternionic manifold is a self-dual conformal manifold. The relationship between quaternionic, almost complex and Riemannian structures in four dimensions is described in more detail in [S4, Chapter 7], a classic reference being [AHS].

Because there is no suitable quaternionic version of holomorphic coordinates, there is no 'nice' co-ordinate expression for a typical section of $C^\infty(E_{k,r})$. In order to determine the decomposition of a differential form, the simplest way the author has found is to use the Casimir operator $\mathcal{C} = I^2 + J^2 + K^2$. Consider a k-form α. Then $\alpha \in E_{k,r}$ if and only if $(I^2 + J^2 + K^2)(\alpha) = -r(r+2)\alpha$. This mechanism also allows us to work out expressions for $\mathcal{D}$ and $\overline{\mathcal{D}}$ acting on α.

Lemma 3.2.7 *Let $\alpha \in C^\infty(E_{k,r})$. Then*

$$\mathcal{D}\alpha = -\frac{1}{4}\left((r-1) + \frac{1}{r+1}(I^2 + J^2 + K^2)\right) d\alpha$$

[3] We are assuming throughout that uppercase operators like I, J and K are acting as elements of a Lie algebra, not a Lie group. This makes no difference on T^*M but is important on the exterior powers $\Lambda^k T^*M$. In particular, for $k \neq 1$ we no longer expect $I^2 = J^2 = K^2 = -1$, and by 'act trivially' we mean 'annihilate'.

and

$$\overline{\mathcal{D}}\alpha = \frac{1}{4}\left((r+3) + \frac{1}{r+1}(I^2+J^2+K^2)\right) d\alpha.$$

Proof. We have $d\alpha = \mathcal{D}\alpha + \overline{\mathcal{D}}\alpha$, where $\mathcal{D}\alpha \in E_{k+1,r+1}$ and $\overline{\mathcal{D}}\alpha \in E_{k+1,r-1}$. Applying the Casimir operator gives

$$(I^2+J^2+K^2)(d\alpha) = -(r+1)(r+3)\mathcal{D}\alpha - (r+1)(r-1)\overline{\mathcal{D}}\alpha.$$

Rearranging these equations gives $\mathcal{D}\alpha$ and $\overline{\mathcal{D}}\alpha$. ∎

Writing $\mathcal{D}_{k,r}$ for the particular map $\mathcal{D} : C^\infty(E_{k,r}) \to C^\infty(E_{k+1,r+1})$, we define the quaternionic cohomology groups

$$H^{k,r}_{\mathcal{D}}(M) = \frac{\mathrm{Ker}(\mathcal{D}_{k,r})}{\mathrm{Im}(\mathcal{D}_{k-1,r-1})}. \tag{3.7}$$

3.3 Ellipticity and the Double Complex

In this section we shall determine where our double complex is elliptic and where it is not. It turns out that we have ellipticity everywhere except on the bottom two rows of the complex. This is exactly like the real Dolbeault complex of Figure 3.1, and though it is more difficult to prove for the quaternionic version, the guiding principles which determine where the two double complexes are elliptic are very much the same in both cases.

Here is the main result of this section:

Theorem 3.3.1 *For $2k \geq 4$, the complex*

$$0 \to E_{2k,0} \overset{\mathcal{D}}{\to} E_{2k+1,1} \overset{\mathcal{D}}{\to} E_{2k+2,2} \overset{\mathcal{D}}{\to} \ldots \overset{\mathcal{D}}{\to} E_{2n+k,2n-k} \overset{\mathcal{D}}{\to} 0$$

is elliptic everywhere except at $E_{2k,0}$ and $E_{2k+1,1}$, where it is not elliptic.

For $k = 1$ the complex is elliptic everywhere except at $E_{3,1}$, where it is not elliptic.

For $k = 0$ the complex is elliptic everywhere.

The rest of this section provides a proof of this Theorem. Note the strong similarity between this Theorem and Proposition 3.1.1, the analogous result for the Real Dolbeault complex. Again, it is the isosceles triangle as opposed to diamond shape which causes ellipticity to fail for the bottom row, because $d = \mathcal{D}$ on $E_{2k,0}$ and the projection from $d(C^\infty(E_{2k,0}))$ to $C^\infty(E_{2k+1,1})$ is the identity.

On a complex manifold M^{2n} with holomorphic coordinates z^j, the exterior forms $dz^{a_1} \wedge \ldots \wedge dz^{a_p} \wedge d\bar{z}^{b_1} \wedge \ldots \wedge d\bar{z}^{b_q}$ span $\Lambda^{p,q}$. This allows us to decompose any form $\omega \in \Lambda^{p,q}$, making it much easier to write down the kernels and images of maps which involve exterior multiplication. On a quaternionic manifold M^{4n} there is unfortunately no easy way to write down a local frame for the bundle $E_{k,r}$, because there is no quaternionic version of 'holomorphic coordinates'. However, we can decompose $E_{k,r}$ just enough to enable us to prove Theorem 3.3.1.

A principal observation is that since ellipticity is a local property, we can work on $\mathbb{H}^n$ without loss of generality. Secondly, since $\mathrm{GL}(n,\mathbb{H})$ acts transitively on $\mathbb{H}^n \setminus \{0\}$, if the symbol sequence $\ldots \xrightarrow{\sigma_{e^0}} E_{k,r} \xrightarrow{\sigma_{e^0}} E_{k+1,r+1} \xrightarrow{\sigma_{e^0}} \ldots$ is exact for any nonzero $e^0 \in T^*\mathbb{H}^n$ then it is exact for all nonzero $\xi \in T^*\mathbb{H}^n$. To prove Theorem 3.3.1, we choose one such e^0 and analyse the spaces $E_{k,r}$ accordingly.

3.3.1 Decomposition of the Spaces $E_{k,r}$

Let $e^0 \in T^*_x\mathbb{H}^n \cong \mathbb{H}^n$ and let (I, J, K) be the standard hypercomplex structure on $\mathbb{H}^n$. As in Example 3.2.6, define $e^1 = I(e^0)$, $e^2 = J(e^0)$ and $e^3 = K(e^0)$, so that $\langle e^0, \ldots, e^3\rangle \cong \mathbb{H}$. In this way we single out a particular copy of $\mathbb{H}$ which we call $\mathbb{H}_0$, obtaining a (nonnatural) splitting $T^*_x\mathbb{H}^n \cong \mathbb{H}^{n-1} \oplus \mathbb{H}_0$ which is preserved by action of the hypercomplex structure. This induces the decomposition $\Lambda^k\mathbb{H}^n \cong \bigoplus_{l=0}^4 \Lambda^{k-l}\mathbb{H}^{n-1} \otimes \Lambda^l\mathbb{H}_0$, which decomposes each $E_{k,r} \subset \Lambda^k\mathbb{H}^n$ according to how many differentials in the $\mathbb{H}_0$-direction are present.

Definition 3.3.2 Define the space $E^l_{k,r}$ to be the subspace of $E_{k,r}$ consisting of exterior forms with precisely l differentials in the $\mathbb{H}_0$-direction, *i.e.*

$$E^l_{k,r} \equiv E_{k,r} \cap (\Lambda^{k-l}\mathbb{H}^{n-1} \otimes \Lambda^l\mathbb{H}_0).$$

Then $E^l_{k,r}$ is preserved by the induced action of the hypercomplex structure on $\Lambda^k\mathbb{H}^n$. Thus we obtain an invariant decomposition $E_{k,r} = E^0_{k,r} \oplus E^1_{k,r} \oplus E^2_{k,r} \oplus E^3_{k,r} \oplus E^4_{k,r}$. Note that we can identify $E^0_{k,r}$ on $\mathbb{H}^n$ with $E_{k,r}$ on $\mathbb{H}^{n-1}$.

(Throughout the rest of this section, juxtaposition of exterior forms will denote exterior multiplication, for example αe^{ij} means $\alpha \wedge e^{ij}$.)

We can decompose these summands still further. Consider, for example, the bundle $E^1_{k,r}$. An exterior form $\alpha \in E^1_{k,r}$ is of the form $\alpha_0 e^0 + \alpha_1 e^1 + \alpha_2 e^2 + \alpha_3 e^3$, where $\alpha_j \in \Lambda^{k-1}\mathbb{H}^{n-1}$. Thus α is an element of $\Lambda^{k-1}\mathbb{H}^{n-1} \otimes 2V_1$, since $\mathbb{H}_0 \cong 2V_1$ as an $\mathfrak{sp}(1)$-representation. Since α is in a copy of the representation V_r, it follows from the isomorphism $V_r \otimes V_1 \cong V_{r+1} \oplus V_{r-1}$ that the α_j must be in a combination of V_{r+1} and V_{r-1} representations, *i.e.* $\alpha_j \in E^0_{k-1,r+1} \oplus E^0_{k-1,r-1}$. We write

$$\alpha = \alpha^+ + \alpha^- = (\alpha_0^+ + \alpha_0^-)e^0 + (\alpha_1^+ + \alpha_1^-)e^1 + (\alpha_2^+ + \alpha_2^-)e^2 + (\alpha_3^+ + \alpha_3^-)e^3,$$

where $\alpha_j^+ \in E^0_{k-1,r+1}$ and $\alpha_j^- \in E^0_{k-1,r-1}$.

The following Lemma allows us to consider α^+ and α^- separately.

Lemma 3.3.3 *If $\alpha = \alpha^+ + \alpha^- \in E^1_{k,r}$ then both α^+ and α^- must be in $E^1_{k,r}$.*

Proof. In terms of representations, the situation is of the form

$$(pV_{r+1} \oplus qV_{r-1}) \otimes 2V_1 \cong 2p(V_{r+2} \oplus V_r) \oplus 2q(V_r \oplus V_{r-2}),$$

where $\alpha^+ \in pV_{r+1}$ and $\alpha^- \in qV_{r-1}$. For α to be in the representation $2(p+q)V_r$, its components in the representations $2pV_{r+2}$ and $2qV_{r-2}$ must both vanish separately. The component in $2pV_{r+2}$ comes entirely from α^+, so for this to vanish we must have

$\alpha^+ \in 2(p+q)V_r$ independently of α^-. Similarly, for the component in $2qV_{r-2}$ to vanish, we must have $\alpha^- \in 2(p+q)V_r$. ∎

Thus we decompose the space $E^1_{k,r}$ into two summands, one coming from $E^0_{k-1,r-1} \otimes 2V_1$ and one from $E^0_{k-1,r+1} \otimes 2V_1$. We extend this decomposition to the cases $l = 0, 2, 3, 4$, defining the following notation.

Definition 3.3.4 Define the bundle $E^{l,m}_{k,r}$ to be the subbundle of $E^l_{k,r}$ arising from V_m-type representations in $\Lambda^{k-l}\mathbb{H}^{n-1}$. In other words,

$$E^{l,m}_{k,r} \equiv (E^0_{k-l,m} \otimes \Lambda^l \mathbb{H}_0) \cap E_{k,r}.$$

To recapitulate: for the space $E^{l,m}_{k,r}$, the bottom left index k refers to the exterior power of the form $\alpha \in \Lambda^k\mathbb{H}^n$; the bottom right index r refers to the irreducible Sp(1)-representation in which α lies; the top left index l refers to the number of differentials in the $\mathbb{H}_0$-direction and the top right index m refers to the irreducible Sp(1)-representation of the contributions from $\Lambda^{k-a}\mathbb{H}^{n-1}$ *before* wedging with forms in the $\mathbb{H}_0$-direction. This may appear slightly fiddly: it becomes rather simpler when we consider the specific splittings which Definition 3.3.4 allows us to write down.

Lemma 3.3.5 *Let $E^{l,m}_{k,r}$ be as above. We have the following decompositions:*

$$E^0_{k,r} = E^{0,r}_{k,r} \qquad E^1_{k,r} = E^{1,r+1}_{k,r} \oplus E^{1,r-1}_{k,r} \qquad E^2_{k,r} = E^{2,r+2}_{k,r} \oplus E^{2,r}_{k,r} \oplus E^{2,r-2}_{k,r}$$

$$E^3_{k,r} = E^{3,r+1}_{k,r} \oplus E^{3,r-1}_{k,r} \qquad \textit{and} \qquad E^4_{k,r} = E^{4,r}_{k,r}.$$

Proof. The first isomorphism is trivial, as is the last (since the hypercomplex structure acts trivially on $\Lambda^4\mathbb{H}_0$). The second isomorphism is Lemma 3.3.3, and the fourth follows in exactly the same way since $\Lambda^3\mathbb{H}_0 \cong 2V_1$ also. The middle isomorphism follows a similar argument. ∎

Recall the self-dual forms and anti-self-dual forms in Example 3.2.6. The bundle $E^{2,r}_{k,r}$ splits according to whether its contribution from $\Lambda^2\mathbb{H}_0$ is self-dual or anti-self-dual. We will call these summands $E^{2,r+}_{k,r}$ and $E^{2,r-}_{k,r}$ respectively, so $E^{2,r}_{k,r} = E^{2,r+}_{k,r} \oplus E^{2,r-}_{k,r}$.

3.3.2 Lie in conditions

We have analysed the bundle $E_{k,r}$ into a number of different subbundles. We now determine when a particular exterior form lies in one of these subbundles. Consider a form $\alpha = \alpha_1 e^{s_1 \ldots s_a} + \alpha_2 e^{t_1 \ldots t_a} + \ldots$ *etc.* where $\alpha_j \in E^0_{k-a,b}$. For α to lie in one of the spaces $E^{a,b}_{k,r}$ the α_j will usually have to satisfy some simultaneous equations. Since these are the conditions for a form to lie in a particular Lie algebra representation, we will refer to such equations as 'Lie in conditions'.

To begin with, we mention three trivial Lie in conditions. Let $\alpha \in E^0_{k,r}$. That $\alpha \in E^{0,r}_{k,r}$ is obvious, as is $\alpha e^{0123} \in E^{4,r}_{k,r}$, since wedging with e^{0123} has no effect on the $\mathfrak{sp}(1)$-action.

Likewise, the $\mathfrak{sp}(1)$-action on the anti-self-dual 2-forms $\omega_1^- = e^{01} - e^{23}$, $\omega_2^- = e^{02} - e^{31}$ and $\omega_3^- = e^{03} - e^{12}$ is trivial, so $\alpha\omega_j^- \in E^{2,r-}_{k,r}$ for all $j = 1, 2, 3$.

This leaves the following three situations: those arising from taking exterior products with 1-forms, 3-forms and the self-dual 2-forms ω_j^+. As usual when we want to know which representation an exterior form is in, we apply the Casimir operator.

The cases $l = 1$ and $l = 3$

Let $\alpha_j \in E^0_{k,r}$. Then $\alpha = \alpha_0 e^0 + \alpha_1 e^1 + \alpha_2 e^2 + \alpha_3 e^3 \in E^{1,r}_{k+1,r+1} \oplus E^{1,r}_{k+1,r-1}$, and α is entirely in $E^{1,r}_{k+1,r+1}$ if and only if $(I^2 + J^2 + K^2)\alpha = -(r+1)(r+3)\alpha$.

By the usual (Leibniz) rule for a Lie algebra action on a tensor product, we have that $I^2(\alpha_j e^j) = I^2(\alpha_j)e^j + 2I(\alpha_j)I(e^j) + \alpha_j I^2(e^j)$, *etc.* Thus

$$\begin{aligned}(I^2 + J^2 + K^2)\alpha &= \sum_{j=0}^{3} \Big[(I^2 + J^2 + K^2)(\alpha_j)e^j + \alpha_j(I^2 + J^2 + K^2)(e^j) \; + \\ &\qquad + 2\Big(I(\alpha_j)I(e^j) + J(\alpha_j)J(e^j) + K(\alpha_j)K(e^j)\Big)\Big] \\ &= -r(r+2)\alpha - 3\alpha + 2\sum_{j=0}^{3}\Big(I(\alpha_j)I(e^j) + J(\alpha_j)J(e^j) + K(\alpha_j)K(e^j)\Big) \\ &= (-r^2 - 2r - 3)\alpha + 2\,\Big(I(\alpha_0)e^1 - I(\alpha_1)e^0 + I(\alpha_2)e^3 - I(\alpha_3)e^2 + \\ &\qquad +J(\alpha_0)e^2 - J(\alpha_1)e^3 - J(\alpha_2)e^0 + J(\alpha_3)e^1 + \\ &\qquad +K(\alpha_0)e^3 + K(\alpha_1)e^2 - K(\alpha_2)e^1 - K(\alpha_3)e^0\Big).\end{aligned} \tag{3.8}$$

For $\alpha \in E^{1,r}_{k+1,r+1}$ we need this to be equal to $-(r+1)(r+3)\alpha$, which is the case if and only if

$$\begin{aligned}-r\alpha &= I(\alpha_0)e^1 - I(\alpha_1)e^0 + I(\alpha_2)e^3 - I(\alpha_3)e^2 + J(\alpha_0)e^2 - J(\alpha_1)e^3 - J(\alpha_2)e^0 + J(\alpha_3)e^1 + \\ &\quad + K(\alpha_0)e^3 + K(\alpha_1)e^2 - K(\alpha_2)e^1 - K(\alpha_3)e^0.\end{aligned}$$

Since the α_j have no e^j-factors and the action of I, J and K preserves this property, this equation can only be satisfied if it holds for each of the e^j-components separately. We conclude that $\alpha \in E^{1,r}_{k+1,r+1}$ if and only if α_0, α_1, α_2 and α_3 satisfy the following Lie in conditions: [4]

$$\begin{aligned} r\alpha_0 - I(\alpha_1) - J(\alpha_2) - K(\alpha_3) &= 0 \\ r\alpha_1 + I(\alpha_0) + J(\alpha_3) - K(\alpha_2) &= 0 \\ r\alpha_2 - I(\alpha_3) + J(\alpha_0) + K(\alpha_1) &= 0 \\ r\alpha_3 + I(\alpha_2) - J(\alpha_1) + K(\alpha_0) &= 0. \end{aligned} \tag{3.9}$$

Suppose instead that $\alpha \in E^{1,r}_{k+1,r-1}$. Then $(I^2+J^2+K^2)\alpha = -(r-1)(r+1)\alpha$. Putting this alternative into Equation (3.8) gives the result that $\alpha \in E^{1,r}_{k+1,r-1}$ if and only if

[4] Our interest in these conditions arises from a consideration of exterior forms, but the equations describe $\mathfrak{sp}(1)$-representations in general: they are the conditions that $\alpha \in V_r \otimes V_1$ must satisfy to be in the V_{r+1} subspace of $V_{r+1} \oplus V_{r-1} \cong V_r \otimes V_1$. The other Lie in conditions have similar interpretations.

$$
\begin{array}{rcl}
(r+2)\alpha_0 + I(\alpha_1) + J(\alpha_2) + K(\alpha_3) & = & 0 \\
(r+2)\alpha_1 - I(\alpha_0) - J(\alpha_3) + K(\alpha_2) & = & 0 \\
(r+2)\alpha_2 + I(\alpha_3) - J(\alpha_0) - K(\alpha_1) & = & 0 \\
(r+2)\alpha_3 - I(\alpha_2) + J(\alpha_1) - K(\alpha_0) & = & 0.
\end{array} \tag{3.10}
$$

Consider now $\alpha = \alpha_0 e^{123} + \alpha_1 e^{032} + \alpha_2 e^{013} + \alpha_3 e^{021} \in E^{3,r}_{k+3,r+1} \oplus E^{3,r}_{k+3,r-1}$. Since $\Lambda^3\mathbb{H}_0 \cong \mathbb{H}_0$, the Lie in conditions are exactly the same: for α to be in $E^{3,r}_{k+3,r+1}$ we need the α_j to satisfy Equations (3.9), and for α to be in $E^{3,r}_{k+3,r-1}$ we need the α_j to satisfy Equations (3.10).

The case $l = 2$

We have already noted that wedging a form $\beta \in E^0_{k,r}$ with an anti-self-dual 2-form ω_j^- has no effect on the $\mathfrak{sp}(1)$-action, so $\beta\omega_j^- \in E^{2,r-}_{k+2,r}$. Thus we only have to consider the effect of wedging with the self-dual 2-forms $\langle \omega_1^+, \omega_2^+, \omega_3^+ \rangle \cong V_2 \subset \Lambda^2\mathbb{H}_0$. By the Clebsch-Gordon formula, the decomposition takes the form $V_r \otimes V_2 \cong V_{r+2} \oplus V_r \oplus V_{r-2}$. Thus for $\beta = \beta_1\omega_1^+ + \beta_2\omega_2^+ + \beta_3\omega_3^+$ we want to establish the Lie in conditions for β to be in $E^{2,r}_{k+2,r+2}$, $E^{2,r+}_{k+2,r}$ and $E^{2,r}_{k+2,r-2}$.

We calculate these Lie in conditions in a similar fashion to the previous cases, by considering the action of the Casimir operator $I^2+J^2+K^2$ on β and using the multiplication table (3.6). The following Lie in conditions are then easy to deduce:

$$
\beta \in E^{2,r}_{k+2,r+2} \iff \left\{ \begin{array}{rcl}
(r+4)\beta_1 & = & J(\beta_3) - K(\beta_2) \\
(r+4)\beta_2 & = & K(\beta_1) - I(\beta_3) \\
(r+4)\beta_3 & = & I(\beta_2) - J(\beta_1).
\end{array} \right. \tag{3.11}
$$

$$
\beta \in E^{2,r+}_{k+2,r} \iff \left\{ \begin{array}{rcl}
2\beta_1 & = & J(\beta_3) - K(\beta_2) \\
2\beta_2 & = & K(\beta_1) - I(\beta_3) \\
2\beta_3 & = & I(\beta_2) - J(\beta_1).
\end{array} \right. \tag{3.12}
$$

$$
\beta \in E^{2,r}_{k+2,r-2} \iff \left\{ \begin{array}{rcl}
(2-r)\beta_1 & = & J(\beta_3) - K(\beta_2) \\
(2-r)\beta_2 & = & K(\beta_1) - I(\beta_3) \\
(2-r)\beta_3 & = & I(\beta_2) - J(\beta_1).
\end{array} \right. \tag{3.13}
$$

Equation 3.12 is particularly interesting. Since this equation singles out the V_r-representation in the direct sum $V_{r+2} \oplus V_r \oplus V_{r+2}$, it must have $\dim V_r = r+1$ linearly independent solutions. Let $\beta_0 \in V_r$ and let $\beta_1 = I(\beta_0)$, $\beta_2 = J(\beta_0)$, $\beta_3 = K(\beta_0)$. Using the Lie algebra relations $2I = [J,K] = JK - KJ$, it is easy to see that β_1, β_2 and β_3 satisfy Equation 3.12. Moreover, there are $r+1$ linearly independent solutions of this form (for $r \neq 0$). We conclude that *all* the solutions of Equation (3.12) take the form $\beta_1 = I(\beta_0)$, $\beta_2 = J(\beta_0)$, $\beta_3 = K(\beta_0)$.

3.3.3 The Symbol Sequence and Proof of Theorem 3.3.1

We now describe the principal symbol of $\mathcal{D}$, and examine its behaviour in the context of the decompositions of Definition 3.3.2 and Lemma 3.3.5. This leads to a proof of Theorem 3.3.1. First we obtain the principal symbol from the formula for $\mathcal{D}$ in Lemma 3.2.7 by replacing $d\alpha$ with αe^0.

Proposition 3.3.6 *Let* $x \in \mathbb{H}^n$, $e^0 \in T_x^*\mathbb{H}^n$ *and* $\alpha \in E_{k,r}$. *The principal symbol mapping* $\sigma_{\mathcal{D}}(x, e^0) : E_{k,r} \to E_{k+1,r+1}$ *is given by*

$$\sigma_{\mathcal{D}}(x, e^0)(\alpha) = \frac{1}{2(r+1)}\left((r+2)\alpha e^0 - I(\alpha)e^1 - J(\alpha)e^2 - K(\alpha)e^3\right).$$

Proof. Replacing $d\alpha$ with αe^0 in the formula for $\mathcal{D}$ obtained in Lemma 3.2.7, we have

$$\begin{aligned}
\sigma_{\mathcal{D}}(x, e^0)(\alpha) &= -\frac{1}{4}\Big((r-1) + \frac{1}{r+1}(I^2 + J^2 + K^2)\Big)\alpha e^0 \\
&= \frac{-1}{4(r+1)}\Big[\,(\,(r-1)(r+1) - r(r+2) - 3\,)\,\alpha e^0 \\
&\qquad +2\left(I(\alpha)e^1 + J(\alpha)e^2 + K(\alpha)e^3\right)\Big] \\
&= \frac{1}{2(r+1)}\Big((r+2)\alpha e^0 - I(\alpha)e^1 - J(\alpha)e^2 - K(\alpha)e^3\Big),
\end{aligned}$$

as required. ∎

Corollary 3.3.7 *The principal symbol* $\sigma_{\mathcal{D}}(x, e^0)$ *maps the space* $E_{k,r}^{l,m}$ *to the space* $E_{k+1,r+1}^{l+1,m}$.

Proof. We already know that $\sigma_{\mathcal{D}} : E_{k,r} \to E_{k+1,r+1}$, by definition. Using Lemma 3.3.6, we see that $\sigma_{\mathcal{D}}(x, e^0)$ increases the number of differentials in the $\mathbb{H}_0$-direction by one, so the index l increases by one. The only action in the other directions is the $\mathfrak{sp}(1)$-action, which preserves the irreducible decomposition of the contribution from $\Lambda^{k-a}\mathbb{H}^{n-1}$, so the index m remains the same. ∎

(To save space we shall use σ as an abbreviation for $\sigma_{\mathcal{D}}(x, e^0)$ for the rest of this section.)

The point of all this work on decomposition now becomes apparent. Since $\sigma : E_{k,r}^l \to E_{k+1,r+1}^{l+1}$, we can reduce the (somewhat indefinite) symbol sequence

$$\ldots \xrightarrow{\sigma} E_{k-1,r-1} \xrightarrow{\sigma} E_{k,r} \xrightarrow{\sigma} E_{k+1,r+1} \xrightarrow{\sigma} \ldots \textit{etc.}$$

to the 5-space sequence

$$0 \xrightarrow{\sigma} E_{k-2,r-2}^0 \xrightarrow{\sigma} E_{k-1,r-1}^1 \xrightarrow{\sigma} E_{k,r}^2 \xrightarrow{\sigma} E_{k+1,r+1}^3 \xrightarrow{\sigma} E_{k+2,r+2}^4 \xrightarrow{\sigma} 0. \qquad (3.14)$$

Using Lemma 3.3.5 as well, we can analyse this sequence still further according to the different (top right) m-indices, obtaining three short sequences (for $k \geq 2$, $k \equiv r \bmod 2$)

$$\begin{array}{ccccccccccc}
 & & & & 0 & \to & E_{k,r}^{2,r+2} & \to & E_{k+1,r+1}^{3,r+2} & \to & E_{k+2,r+2}^{4,r+2} & \to 0 \\
 & & & & & & \oplus & & \oplus & & & \\
 & & 0 & \to & E_{k-1,r-1}^{1,r} & \to & E_{k,r}^{2,r} & \to & E_{k+1,r+1}^{3,r} & \to & 0 & \\
 & & & & \oplus & & \oplus & & & & & \\
0 & \to & E_{k-2,r-2}^{0,r-2} & \to & E_{k-1,r-1}^{1,r-2} & \to & E_{k,r}^{2,r-2} & \to & 0. & & &
\end{array} \qquad (3.15)$$

This reduces the problem of determining where the operator $\mathcal{D}$ is elliptic to the problem of determining when these three sequences are exact.

For a sequence $0 \to A \to B \to C \to 0$ to be exact, it is necessary that $\dim A - \dim B + \dim C = 0$. Given this condition, if the sequence is exact at any two out of A, B and C it is exact at the third. We shall show that for $r \neq 0$ this dimension sum does equal zero.

Lemma 3.3.8 *For $r > 0$, each of the sequences in (3.15) satisfies the dimension condition above, i.e. the alternating sum of the dimensions vanishes.*

Proof. Let $r > 0$. We calculate the dimensions of the spaces $E^{l,m}_{k,r}$ for $l = 0, \dots, 4$. Recall the notation $E_{k,r} = \epsilon^n_{k,r} V_r$ from Definition 3.2.2. It is clear that $\dim E^0_{k,r} = (r+1)\epsilon^{n-1}_{k,r}$, since $E^0_{k,r}$ on $\mathbb{H}^n$ is simply $E_{k,r}$ on $\mathbb{H}^{n-1}$. Thus $\dim E^{0,r-2}_{k-2,r-2} = (r-1)\epsilon^{n-1}_{k-2,r-2}$ and $\dim E^{4,r+2}_{k+2,r+2} = (r+3)\epsilon^{n-1}_{k-2,r+2}$.

The cases $a = 1$ and $a = 3$ are easy to work out since they are of the form $E^0_{k,r} \otimes 2V_1$. For $a = 1$, we have $\dim E^{1,r-2}_{k-1,r-1} = 2r\epsilon^{n-1}_{k-2,r-2}$ and $\dim E^{1,r}_{k-1,r-1} = 2r\epsilon^{n-1}_{k-2,r}$. For $a = 3$, $\dim E^{3,r}_{k+1,r+1} = 2(r+2)\epsilon^{n-1}_{k-2,r}$ and $\dim E^{3,r+2}_{k+1,r+1} = 2(r+2)\epsilon^{n-1}_{k-2,r+2}$.

The case $a = 2$ is slightly more complicated, as we have to take into account exterior products with the self-dual 2-forms V_2 and anti-self-dual 2-forms $3V_0$ in $\Lambda^2\mathbb{H}_0$. The spaces $E^{2,r+2}_{k,r}$ and $E^{2,r-2}_{k,r}$ receive contributions only from the self-dual part V_2, from which we infer that $\dim E^{2,r+2}_{k,r} = (r+1)\epsilon^{n-1}_{k-2,r+2}$ and $\dim E^{2,r-2}_{k,r} = (r+1)\epsilon^{n-1}_{k-2,r-2}$. Finally, the space $E^{2,r+}_{k,r}$ has dimension $(r+1)\epsilon^{n-1}_{k-2,r}$ and the space $E^{2,r-}_{k,r}$ has dimension $3(r+1)\epsilon^{n-1}_{k-2,r}$, giving $E^{2,r}_{k,r}$ a total dimension of $4(r+1)\epsilon^{n-1}_{k-2,r}$.

It is now a simple matter to verify that for the top sequence of (3.15)

$$\epsilon^{n-1}_{k-2,r+2}(r + 1 - 2(r+2) + r + 3) = 0,$$

for the middle sequence

$$\epsilon^{n-1}_{k-2,r}(2r - 4(r+1) + 2(r+2)) = 0,$$

and for the bottom sequence

$$\epsilon^{n-1}_{k-2,r-2}(r - 1 - 2r + r + 1) = 0.$$

∎

The case $r = 0$ is different. Here the bottom sequence of (3.15) disappears altogether, the top sequence still being exact. Exactness is lost in the middle sequence. Since the isomorphism $\epsilon^{n-1}_{k-2,0}V_0 \otimes V_2 \cong \epsilon^{n-1}_{k-2,0}V_2$ gives no trivial V_0-representations, there is no space $E^{2,0+}_{k,0}$. Thus $E^{2,0}_{k,0}$ is 'too small' — we are left with a sequence

$$0 \longrightarrow 3\epsilon^{n-1}_{k-2,0}V_0 \longrightarrow 2\epsilon^{n-1}_{k-2,0}V_1 \longrightarrow 0,$$

which cannot be exact. (As there is no space $E^2_{0,0}$, this problem does not arise for the leading edge $0 \to E_{0,0} \to E_{1,1} \to \dots$ *etc.*)

We are finally in a position to prove Theorem 3.3.1, which now follows from:

Proposition 3.3.9 *When $r \neq 0$, the three sequences of (3.15) are exact.*

Proof. Consider first the top sequence $0 \xrightarrow{\sigma} E^{2,r+2}_{k,r} \xrightarrow{\sigma} E^{3,r+2}_{k+1,r+1} \xrightarrow{\sigma} E^{4,r+2}_{k+2,r+2} \longrightarrow 0$. The Clebsch-Gordon formula shows that there are no spaces $E^{3,r+2}_{k+1,r-1}$ or $E^{4,r+2}_{k+2,r}$. Thus $\overline{\mathcal{D}} = 0$ on $E^{2,r+2}_{k,r}$ and $E^{3,r+2}_{k+1,r+1}$, so $\mathcal{D} = d$ for the top sequence. It is easy to check using the relevant Lie in conditions that the map $\wedge e^0 : E^{2,r+2}_{k,r} \to E^{3,r+2}_{k+1,r+1}$ is injective and the map $\wedge e^0 : E^{3,r+2}_{k+1,r+1} \to E^{4,r+2}_{k+2,r+2}$ is surjective.

To show exactness at $E^1_{k-1,r-1}$, consider $\alpha = \alpha_0 e^0 + \alpha_1 e^1 + \alpha_2 e^2 + \alpha_3 e^3 \in E^1_{k-1,r-1}$. A calculation using Proposition 3.3.6 shows that

$$\begin{aligned}\sigma(\alpha) \;=\; & \frac{1}{2r}\big((r\alpha_1 + I(\alpha_0))e^{10} + (r\alpha_2 + J(\alpha_0))e^{20} + (r\alpha_3 + K(\alpha_0))e^{30} + \\ & + (2\alpha_1 - J(\alpha_3) + K(\alpha_2))e^{32} + (2\alpha_2 - K(\alpha_1) + I(\alpha_3))e^{13} + (2\alpha_3 - I(\alpha_2) + K(\alpha_1))e^{21}\big).\end{aligned} \tag{3.16}$$

Since the α_i have no e^j-components, $\sigma(\alpha) = 0$ if and only if all these components vanish. This occurs if and only if $\alpha_1 = -\frac{1}{r}I(\alpha_0)$, $\alpha_2 = -\frac{1}{r}J(\alpha_0)$, $\alpha_3 = -\frac{1}{r}K(\alpha_0)$ (since as remarked in Section 3.3.1 these equations also guarantee that $2\alpha_1 - J(\alpha_3) + K(\alpha_2) = 0$ *etc.*), in which case it is clear that

$$\sigma(\alpha) = 0 \Longleftrightarrow \alpha = \sigma\left(\tfrac{2(r+1)}{r}\alpha_0\right).$$

This shows that the sequence $E^0_{k-2,r-2} \to E^1_{k-1,r-1} \to E^2_{k,r}$ is exact. Restricting to $E^{1,r}_{k-1,r-1}$ and $E^{1,r-2}_{k-1,r-1}$, we see that exactness holds at these spaces in the middle and bottom sequences respectively of (3.15).

Consider $\alpha \in E^0_{k-2,r-2}$. Then

$$\sigma(\alpha) = \frac{1}{2(r-1)}\left(r\alpha e^0 - I(\alpha)e^1 - J(\alpha)e^2 - K(\alpha)e^3\right).$$

Since these are linearly independent, $\sigma(\alpha) = 0$ if and only if $\alpha = 0$, and $\sigma : E^{0,r-2}_{k-2,r-2} \to E^{1,r-2}_{k-1,r-1}$ is injective. Hence the bottom sequence $0 \longrightarrow E^{0,r-2}_{k-2,r-2} \xrightarrow{\sigma} E^{1,r-2}_{k-1,r-1} \xrightarrow{\sigma} E^{2,r-2}_{k,r} \longrightarrow 0$ is exact.

Finally, we show that the middle sequence $0 \longrightarrow E^{1,r}_{k-1,r-1} \xrightarrow{\sigma} E^{2,r}_{k,r} \xrightarrow{\sigma} E^{3,r}_{k+1,r+1} \longrightarrow 0$ is exact at $E^{2,r}_{k,r}$, which is now sufficient to show that the sequence is exact.

Let $\beta = \beta_1\omega_1^+ + \beta_2\omega_2^+ + \beta_3\omega_3^+ \in E^{2,r+}_{k,r}$. Recall the Lie in condition (3.12) that β must take the form $\beta = \frac{1}{r}\left(I(\beta_0)\omega_1^+ + J(\beta_0)\omega_2^+ + K(\beta_0)\omega_3^+\right)$ for some $\beta_0 \in E^0_{k-2,r}$. (The $\frac{1}{r}$-factor makes no difference here and is useful for cancellations.) Thus a general element of $E^{2,r}_{k,r}$ is of the form

$$\beta + \gamma = \frac{1}{r}\left(I(\beta_0)\omega_1^+ + J(\beta_0)\omega_2^+ + K(\beta_0)\omega_3^+\right) + \gamma_1\omega_1^- + \gamma_2\omega_2^- + \gamma_3\omega_3^-,$$

for $\beta_0, \gamma_j \in E^{0,r}_{k-2,r}$. A similar calculation to that of (3.16) shows that

$$\sigma(\beta+\gamma) = 0 \Longleftrightarrow \left\{\begin{array}{rcl} (r+2)\beta_0 + I(\gamma_1) + J(\gamma_2) + K(\gamma_3) & = & 0 \\ (r+2)\gamma_1 - I(\beta_0) - J(\gamma_3) + K(\gamma_2) & = & 0 \\ (r+2)\gamma_2 + I(\gamma_3) - J(\beta_0) - K(\gamma_1) & = & 0 \\ (r+2)\gamma_3 - I(\gamma_2) + J(\gamma_1) - K(\beta_0) & = & 0. \end{array}\right.$$

But this is exactly the Lie in condition (3.10) which we need for $\beta_0 e^0 + \gamma_1 e^1 + \gamma_2 e^2 + \gamma_3 e^3$ to be in $E^{1,r}_{k-1,r-1}$, in which case we have

$$\beta + \gamma = \sigma\left(2(\beta_0 e^0 + \gamma_1 e^1 + \gamma_2 e^2 + \gamma_3 e^3)\right).$$

This demonstrates exactness at $E^{2,r}_{k,r}$ and so the middle sequence is exact. ∎

As a counterexample for the case $r = 0$ and $k \geq 4$, consider $\alpha \in E^0_{k-4,0}$. Then $\alpha e^{0123} \in E^{4,0}_{k,0}$ and $\sigma(\alpha e^{0123}) = 0$, so $\sigma : E_{k,0} \to E_{k+1,1}$ is not injective, which is exactly the same as saying that the symbol sequence is not exact at $E_{k,0}$. It is easy to see that this counterexample does not arise when $k = 0$ or 2, and to show that the maps $\sigma : E_{0,0} \to E_{1,1}$ and $\sigma : E_{2,0} \to E_{3,1}$ *are* injective.

As a counterexample for the case $r = 1$ and $k \geq 2$, consider $\alpha \in E^0_{k-2,0}$. Then $\alpha e^{123} \in E^{3,0}_{k+1,1}$ and $\alpha e^{123} \wedge e^0 \in E^4_{k+2,0}$. Thus $\sigma(\alpha e^{123}) = 0$. Since αe^{123} has no e^0-components at all it is clear that $\alpha e^{123} \neq \sigma(\beta)$ for any $\beta \in E^2_{k,0}$. Thus the symbol sequence fails to be exact at $E_{k+1,1}$. Again, it is easy to see that this counterexample does not arise when $k = 0$, and to show that the sequence $E_{0,0} \xrightarrow{\sigma} E_{1,1} \xrightarrow{\sigma} E_{2,2}$ *is* exact at $E_{1,1}$.

This concludes our proof of Theorem 3.3.1.

3.4 Quaternion-valued forms on Hypercomplex Manifolds

Let M be a hypercomplex manifold. Then M has a triple (I, J, K) of complex structures which we can identify globally with the imaginary quaternions. Thus we have globally defined operators which generate the $\mathfrak{sp}(1)$-action on $\Lambda^k T^*M$.

Consider also the quaternions themselves. Equation (2.6) describes the $\mathrm{Sp}(1)\mathrm{GL}(n,\mathbb{H})$-representation on $\mathbb{H}^n$ as $V_1 \otimes E$. In the case $n = 1$ this reduces to the representation

$$\mathbb{H} \cong V_1 \otimes V_1, \tag{3.17}$$

where we can interpret the left-hand copy of V_1 as the left-action $(p,q) \mapsto pq$, and the right-hand copy of V_1 as the right-action $(p,q) \mapsto qp^{-1}$, for $q \in \mathbb{H}$ and $p \in \mathrm{Sp}(1)$.

We can now use our globally defined hypercomplex structure to combine the $\mathrm{Sp}(1)$-actions on $\mathbb{H}$ and $\Lambda^k T^*M$. This motivates a thorough investigation of quaternion-valued forms on hypercomplex manifolds. Consider, for example, quaternion-valued exterior forms in the bundle $E_{k,r} = \epsilon^n_{k,r} V_r$. The $\mathrm{Sp}(1)$-action on these forms is described by the representation

$$\mathbb{H} \otimes E_{k,r} \cong V_1 \otimes V_1 \otimes \epsilon^n_{k,r} V_r.$$

Leaving the left $\mathbb{H}$-action untouched, we consider the effect of the right $\mathbb{H}$-action and the hypercomplex structure simultaneously. This amounts to applying the operators

$$\mathcal{I} : \alpha \to I(\alpha) - \alpha i_1, \qquad \mathcal{J} : \alpha \to J(\alpha) - \alpha i_2 \qquad \text{and} \qquad \mathcal{K} : \alpha \to K(\alpha) - \alpha i_3$$

to $\alpha \in \mathbb{H} \otimes E_{k,r}$. Under this diagonal action the tensor product $V_1 \otimes \epsilon^n_{k,r} V_r$ splits, giving the representation

$$\mathbb{H} \otimes E_{k,r} \cong V_1 \otimes \epsilon^n_{k,r}(V_{r+1} \oplus V_{r-1}). \tag{3.18}$$

Each of these summands inherits the structure of a left $\mathbb{H}$-module from the left-action V_1, which is not affected by our splitting.

This situation mirrors our discussion of real and complex forms on complex manifolds. There is a decomposition of real-valued forms, which is taken further when we consider complex-valued forms. In the same way, considering quaternion-valued forms on a hypercomplex manifold allows us to take our decomposition further.

This point of view turns out to be very fruitful. It will, over the next few chapters, lead to quaternionic analogues of holomorphic functions and k-forms, the holomorphic tangent and cotangent spaces, and complex Lie groups and Lie algebras.

The algebraic foundation for this geometry lies in considering objects like our left $\mathbb{H}$-modules in Equation (3.18). Each of the summands $V_1 \otimes \epsilon^n_{k,r} V_{r\pm 1}$ is a left $\mathbb{H}$-module which arises as a submodule of $\mathbb{H} \otimes E_{k,r} \cong (r+1)\epsilon^n_{k,r}\mathbb{H}^n$. Thus each summand is an $\mathbb{H}$-linear submodule of $\mathbb{H}^n$. In the next chapter we will introduce a new algebraic theory which is based upon such objects.

Chapter 4

Developments in Quaternionic Algebra

This chapter describes an algebraic theory which will be central to our description of hypercomplex geometry. The theory is that of my supervisor, Dominic Joyce, and is presented in [J1]. The basic objects of study are $\mathbb{H}$-submodules U of $\mathbb{H} \otimes \mathbb{R}^n$. Joyce shows that the inclusion $\iota_U : U \to \mathbb{H}^n$ is determined up to isomorphism by the $\mathbb{H}$-module structure of U and the choice of a real vector subspace $U' \subset U$ satisfying a certain condition. The pair (U, U') is an *augmented* $\mathbb{H}$*-module*, or A$\mathbb{H}$-module.

The most important discovery in [J1] is a canonical tensor product for A$\mathbb{H}$-modules with interesting properties. (Recall from Section 1.3 that the most obvious definitions of a tensor product over the quaternions are not especially fruitful.) For two A$\mathbb{H}$-modules $U \subset \mathbb{H}^m$ and $V \subset \mathbb{H}^n$, we can define a unique A$\mathbb{H}$-module $U \otimes_{\mathbb{H}} V \subset \mathbb{H}^{mn}$. The operation '$\otimes_{\mathbb{H}}$' will be called the quaternionic tensor product. It has similar properties to the tensor product over a commuting field; for example it is both associative and commutative. This allows us to develop the algebra of A$\mathbb{H}$-modules as a parallel to that of vector spaces over $\mathbb{R}$ or $\mathbb{C}$. This analogy is particularly strong for certain well-behaved A$\mathbb{H}$-modules which will be called stable A$\mathbb{H}$-modules.

There are other algebraic operations which are equivalent to Joyce's quaternionic tensor product. A sheaf-theoretic point of view is presented by Quillen [Q], in which he discovers a contravariant equivalence of tensor categories between A$\mathbb{H}$-modules and regular sheaves on a real form of $\mathbb{C}P^1$. This allows us to classify all A$\mathbb{H}$-modules and determine their tensor products. In the next chapter, we will see that the most important classes of A$\mathbb{H}$-modules are conveniently described and manipulated using Sp(1)-representations.

4.1 The Quaternionic Algebra of Joyce

The following is a summary of parts of Joyce's theory of quaternionic algebra. The interested reader should consult [J1] for more details and proofs.

4.1.1 A$\mathbb{H}$-Modules

We begin by defining $\mathbb{H}$-modules and their dual spaces. A (left) $\mathbb{H}$-module is a real vector space U with an action of $\mathbb{H}$ on the left which we write as $(q, u) \mapsto q \cdot u$ or qu, such that $p(q(u)) = (pq)(u)$ for $p, q \in \mathbb{H}$ and $u \in U$. For our purposes, all $\mathbb{H}$-modules will be left $\mathbb{H}$-modules. By $\dim U$ we will always mean the dimension of U as a *real* vector space, even if U is an $\mathbb{H}$-module.

We write U^* for the dual vector space of U. If U is an $\mathbb{H}$-module we also define the *dual* $\mathbb{H}$-*module* $U^\times$ of linear maps $\alpha : U \to \mathbb{H}$ that satisfy $\alpha(qu) = q\alpha(u)$ for all $q \in \mathbb{H}$ and $u \in U$. If $q \in \mathbb{H}$ and $\alpha \in U^\times$, define $q \cdot \alpha$ by $(q \cdot \alpha)(u) = \alpha(u)\overline{q}$ for $u \in U$. Then $q \cdot \alpha \in U^\times$, and $U^\times$ is a (left) $\mathbb{H}$-module. Dual $\mathbb{H}$-modules behave just like dual vector spaces.

Definition 4.1.1 [J1, Definition 2.2] Let U be an $\mathbb{H}$-module. Let U' be a real vector subspace of U. Define a real vector subspace $U^\dagger$ of $U^\times$ by

$$U^\dagger = \{\alpha \in U^\times : \alpha(u) \in \mathbb{I} \text{ for all } u \in U'\}. \tag{4.1}$$

Conversely, $U^\dagger$ determines U' (at least for finite-dimensional U) by

$$U' = \{u \in U : \alpha(u) \in \mathbb{I} \text{ for all } \alpha \in U^\dagger\}. \tag{4.2}$$

An *augmented* $\mathbb{H}$-*module*, or A$\mathbb{H}$-*module*, is a pair (U, U') such that if $u \in U$ and $\alpha(u) = 0$ for all $\alpha \in U^\dagger$, then $u = 0$. We consider $\mathbb{H}$ to be an A$\mathbb{H}$-module, with $\mathbb{H}' = \mathbb{I}$. A$\mathbb{H}$-modules should be thought of as the quaternionic analogues of real vector spaces.

Usually we will refer to U itself as an A$\mathbb{H}$-module, assuming that U' is also given. If we consider only the real part $\mathrm{Re}(\alpha(u))$ (for $u \in U$ and $\alpha \in U^\times$), we can interpret $U^\times$ as the dual of U as a real vector space, and then $U^\dagger$ is the annihilator of U'. Thus if U is finite-dimensional, $\dim U' + \dim U^\dagger = \dim U = \dim U^\times$ and an isomorphism $U \cong U^\times$ determines an isomorphism $U/U' \cong U^\dagger$.

Let U be an A$\mathbb{H}$-module and let $u, v \in U$ such that $\alpha(u) = \alpha(v)$ for all $\alpha \in U^\dagger$. Then since α is a linear map we have $\alpha(u - v) = 0$ for all $\alpha \in U^\dagger$ and it follows from Definition 4.1.1 that $u = v$. Thus U is an A$\mathbb{H}$-module if and only if each $u \in U$ is uniquely determined by the values of $\alpha(u)$ for $\alpha \in U^\dagger$. In effect, the definition of an A$\mathbb{H}$-module demands that U' should not be *too large*. Definition 4.1.1 demands for any $u \in U$, its $\mathbb{H}$-linear span $\mathbb{H} \cdot u$ should not be entirely contained in U', so $\dim(U' \cap \mathbb{H} \cdot u) \leq 3$.

If V is an A$\mathbb{H}$-module, we say that U *is an* A$\mathbb{H}$-*submodule of* V if U is an $\mathbb{H}$-submodule of V and $U' = U \cap V'$. As $U^\dagger$ is the restriction of $V^\dagger$ to U, if $\alpha(u) = 0$ for all $\alpha \in U^\dagger$ then $u = 0$, so U is an A$\mathbb{H}$-module. Define W to be the quotient $\mathbb{H}$-module V/U and define W' to be the real subspace $(V' + U)/U$ of W. We would like to define (W, W') to be the *quotient* A$\mathbb{H}$-*module* V/U. However, there is a catch: W may not be an A$\mathbb{H}$-module, as the condition in Definition 4.1.1 may not be satisfied.

Example 4.1.2 [J1, Definition 6.1] Let $Y \subset \mathbb{H}^3$ be the set $Y = \{(q_1, q_2, q_3) : q_1 i_1 + q_2 i_2 + q_3 i_3 = 0\}$. Then $Y \cong \mathbb{H}^2$ is a left $\mathbb{H}$-module. Define a real subspace $Y' = Y \cap \mathbb{I}^3$; so $Y' = \{(q_1, q_2, q_3) : q_j \in \mathbb{I} \text{ and } q_1 i_1 + q_2 i_2 + q_3 i_3 = 0\}$. Then $\dim Y = 8$ and $\dim Y' = 5$.

Let $\nu : Y \to \mathbb{H}$, $\nu(q_1, q_2, q_3) = i_1 q_1 + i_2 q_2 + i_3 q_3$. Then $\mathrm{im}(\nu) = \mathbb{I}$ and $\ker(\nu) = Y'$. Since $Y/Y' \cong (Y^\dagger)^*$, ν induces an isomorphism $(Y^\dagger)^* \cong \mathbb{I} \cong V_2$.

Here is the natural concept of linear map between A$\mathbb{H}$-modules:

Definition 4.1.3 Let U, V be A$\mathbb{H}$-modules and let $\phi : U \to V$ be $\mathbb{H}$-linear. We say that ϕ is an A$\mathbb{H}$-*morphism* if $\phi(U') \subset V'$. If ϕ is also an isomorphism of $\mathbb{H}$-modules we say ϕ is an A$\mathbb{H}$-*isomorphism*.

One obvious question is whether we can classify A$\mathbb{H}$-modules up to A$\mathbb{H}$-isomorphism. Real and complex vector spaces are classified by dimension, but clearly this is not true for A$\mathbb{H}$-modules, as there are several choices of U' for each $\mathbb{H}$-module $U \cong \mathbb{H}^n$. We will return to this question in some detail later.

We note the following points:

- If $\phi : U \to V$ and $\psi : V \to W$ are A$\mathbb{H}$-morphisms, then $\psi \circ \phi : U \to W$ is an A$\mathbb{H}$-morphism.

- Let U, V be A$\mathbb{H}$-modules and $\phi : U \to V$ an A$\mathbb{H}$-morphism. Define an $\mathbb{H}$-linear map $\phi^\times : V^\times \to U^\times$ by $\phi^\times(\beta)(u) = \beta(\phi(u))$ for $\beta \in V^\times$ and $u \in U$. Then $\phi(U') \subset V'$ implies that $\phi^\times(V^\dagger) \subset U^\dagger$.

- Let U be an A$\mathbb{H}$-module. Then $\mathbb{H}\otimes(U^\dagger)^*$ is an $\mathbb{H}$-module, with $\mathbb{H}$-action $p\cdot(q\otimes x) = (pq) \otimes x$. Define a map $\iota_U : U \to \mathbb{H} \otimes (U^\dagger)^*$ by $\iota_U(u) \cdot \alpha = \alpha(u)$, for $u \in U$ and $\alpha \in U^\dagger$. Then ι_U is $\mathbb{H}$-linear, so that $\iota_U(U)$ is an $\mathbb{H}$-submodule of $\mathbb{H} \otimes (U^\dagger)^*$.

  Suppose $u \in \ker \iota_U$. Then $\alpha(u) = 0$ for all $\alpha \in U^\dagger$, so that $u = 0$ as U is an A$\mathbb{H}$-module. Thus ι_U is injective, and $\iota_U(U) \cong U$.

- From Equation (4.2), it follows that $\iota_U(U') = \iota_U(U) \cap (\mathbb{I} \otimes (U^\dagger)^*)$. Thus the A$\mathbb{H}$-module (U, U') is determined by the $\mathbb{H}$-submodule $\iota_U(U)$.

This shows as promised that every A$\mathbb{H}$-module is isomorphic to a (left) submodule of $(\mathbb{H} \otimes \mathbb{R}^n, \mathbb{I} \otimes \mathbb{R}^n)$ for $n = \dim(U^\dagger)^*$. Example 4.1.2 shows how ths works for the A$\mathbb{H}$-module Y. As an abstract A$\mathbb{H}$-module Y is isomorphic to $\mathbb{H}^2$ and $(Y^\dagger)^* \cong V_2$. One of the easiest and most symmetrical ways to obtain Y is as an 8-dimensional subspace of $\mathbb{H}^3 = \mathbb{H} \otimes V_2$. We will find this version of events very useful in many situations, as we shall see immediately.

4.1.2 The Quaternionic Tensor Product

Let U and V be A$\mathbb{H}$-modules. Then they can be regarded as subspaces of $\mathbb{H} \otimes (U^\dagger)^*$ and $\mathbb{H} \otimes (V^\dagger)^*$ respectively. Since the $\mathbb{H}$-action on both of these is the same, we can paste these A$\mathbb{H}$-modules together to get a product A$\mathbb{H}$-module. Here is the key idea of the theory:

Definition 4.1.4 [J1, Definition 4.2] Let U, V be A$\mathbb{H}$-modules. Then $\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$ is an $\mathbb{H}$-module, with $\mathbb{H}$-action $p \cdot (q \otimes x \otimes y) = (pq) \otimes x \otimes y$. Exchanging the factors of $\mathbb{H}$ and $(U^\dagger)^*$, we may regard $(U^\dagger)^* \otimes \iota_V(V)$ as a subspace of $\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$. Thus $\iota_U(U) \otimes (V^\dagger)^*$ and $(U^\dagger)^* \otimes \iota_V(V)$ are A$\mathbb{H}$-submodules of $\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$. We define their intersection to be the *quaternionic tensor product of U and V*,

$$U \otimes_{\mathbb{H}} V = (\iota_U(U) \otimes (V^\dagger)^*) \cap ((U^\dagger)^* \otimes \iota_V(V)) \subset \mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*. \quad (4.3)$$

The vector subspace $(U \otimes_{\mathbb{H}} V)'$ is then given by $(U \otimes_{\mathbb{H}} V)' = (U \otimes_{\mathbb{H}} V) \cap (\mathbb{I} \otimes (U^\dagger)^* \otimes (V^\dagger)^*)$ and with this definition $U \otimes_{\mathbb{H}} V$ is an A$\mathbb{H}$-module. The operation $\otimes_{\mathbb{H}}$ will be called the *quaternionic tensor product.*

A few words of explanation may be useful at this point. At the end of Chapter 1 we saw that it is possible to define a sort of tensor product $\mathbb{H}^m \otimes_{\mathbb{H}} \mathbb{H}^n \cong \mathbb{H}^{mn}$, but that this is not really any different from taking tensor products over $\mathbb{R}$. The theory of A$\mathbb{H}$-modules and the quaternionic tensor product is a way of taking the quaternionic behaviour into account. How U' behaves in relation to the $\mathbb{H}$-action determines a particular subspace $\iota_U(U)$ of $\mathbb{H} \otimes \mathbb{R}^n$. The quaternionic tensor product is the natural way of combining these choices for A$\mathbb{H}$-modules $U \subseteq \mathbb{H} \otimes \mathbb{R}^m$ and $V \subseteq \mathbb{H} \otimes \mathbb{R}^n$ into an A$\mathbb{H}$-module $U \otimes_{\mathbb{H}} V \subseteq \mathbb{H} \otimes \mathbb{R}^{mn}$.

We also define the tensor product of two A$\mathbb{H}$-morphisms:

Definition 4.1.5 Let U, V, W, X be A$\mathbb{H}$-modules, and let $\phi : U \to W$ and $\psi : V \to X$ be A$\mathbb{H}$-morphisms. Then $\phi^\times(W^\dagger) \subset U^\dagger$ and $\psi^\times(X^\dagger) \subset V^\dagger$. Taking the duals gives maps $(\phi^\times)^* : (U^\dagger)^* \to (W^\dagger)^*$ and $(\psi^\times)^* : (V^\dagger)^* \to (X^\dagger)^*$. Combining these, we have a map

$$\mathrm{id} \otimes (\phi^\times)^* \otimes (\psi^\times)^* : \mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^* \to \mathbb{H} \otimes (W^\dagger)^* \otimes (X^\dagger)^*. \tag{4.4}$$

Define $\phi \otimes_{\mathbb{H}} \psi : U \otimes_{\mathbb{H}} V \to W \otimes_{\mathbb{H}} X$ to be the restriction of $\mathrm{id} \otimes (\phi^\times)^* \otimes (\psi^\times)^*$ to $U \otimes_{\mathbb{H}} V$. Then $\phi \otimes_{\mathbb{H}} \psi$ is an A$\mathbb{H}$-morphism from $U \otimes_{\mathbb{H}} V$ to $W \otimes_{\mathbb{H}} X$. This is the *quaternionic tensor product of* ϕ *and* ψ.

It can be proved [J1, Lemma 4.3] that there are canonical A$\mathbb{H}$-isomorphisms

$$\mathbb{H} \otimes_{\mathbb{H}} U \cong U, \quad U \otimes_{\mathbb{H}} V \cong V \otimes_{\mathbb{H}} U \quad \text{and} \quad (U \otimes_{\mathbb{H}} V) \otimes_{\mathbb{H}} W \cong U \otimes_{\mathbb{H}} (V \otimes_{\mathbb{H}} W). \tag{4.5}$$

This tells us that $\otimes_{\mathbb{H}}$ is commutative and associative, and that the A$\mathbb{H}$-module $\mathbb{H}$ acts as an identity element for $\otimes_{\mathbb{H}}$. Since $\otimes_{\mathbb{H}}$ is commutative and associative we can define symmetric and antisymmetric products of A$\mathbb{H}$-modules:

Definition 4.1.6 [J1, 4.4] Let U be an A$\mathbb{H}$-module. Write $\bigotimes_{\mathbb{H}}^k U$ for the product $U \otimes_{\mathbb{H}} \cdots \otimes_{\mathbb{H}} U$ of k copies of U, with $\bigotimes_{\mathbb{H}}^0 U = \mathbb{H}$. Then the k^{th} symmetric group S_k acts on $\bigotimes_{\mathbb{H}}^k U$ by permutation of the U factors in the obvious way. Define $S_{\mathbb{H}}^k U$ and $\Lambda_{\mathbb{H}}^k U$ to be the A$\mathbb{H}$-submodules of $\bigotimes_{\mathbb{H}}^k U$ which are symmetric and antisymmetric respectively under the action of S_k.

Much of the algebra that works over $\mathbb{R}$ or $\mathbb{C}$ can be adapted to work over $\mathbb{H}$, using A$\mathbb{H}$-modules and the quaternionic tensor product instead of vector spaces and the real or complex tensor product. There are, however, many situations where the quaternionic tensor product behaves differently from the standard real or complex tensor product. For example, the dimension of $U \otimes_{\mathbb{H}} V$ can behave strangely. It can vary discontinuously under smooth variations of U' or V', and it is possible to have $U \otimes_{\mathbb{H}} V = \{0\}$ when both U and V are non-zero. If ϕ and ψ are both injective A$\mathbb{H}$-morphisms, it is possible to prove [J1, Lemma 7.4] that $\phi \otimes_{\mathbb{H}} \psi$ is also injective. However, if ϕ and ψ are both surjective then $\phi \otimes_{\mathbb{H}} \psi$ is not necessarily surjective.

Given $u \in U$ and $v \in V$ it is not possible in general to define an element $u \otimes_{\mathbb{H}} v \in U \otimes_{\mathbb{H}} V$. However, we do have the following special case:

Lemma 4.1.7 *[J1, 4.6] Let U, V be A$\mathbb{H}$-modules, and let $u \in U$ and $v \in V$ be nonzero. Suppose that $\alpha(u)\beta(v) = \beta(v)\alpha(u) \in \mathbb{H}$ for every $\alpha \in U^\dagger$ and $\beta \in V^\dagger$. Define an* *element* $u \otimes_{\mathbb{H}} v$ *of* $\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$ *by* $(u \otimes_{\mathbb{H}} v) \cdot (\alpha \otimes \beta) = \alpha(u)\beta(v) \in \mathbb{H}$*. Then* $u \otimes_{\mathbb{H}} v$ *is a nonzero element of* $U \otimes_{\mathbb{H}} V$*.*

It is easy to visualise how this Lemma 'works'. If $\alpha(u)\beta(v) = \beta(v)\alpha(u) \in \mathbb{H}$ for every $\alpha \in U^\dagger$ and $\beta \in V^\dagger$, then $\alpha(u)$ and $\beta(v)$ must both be in some commutative subfield $\mathbb{C}_q \subset \mathbb{H}$. This is the same as saying that

$$\iota_U(u) \in \mathbb{C}_q \otimes (U^\dagger)^* \qquad \text{and} \qquad \iota_V(v) \in \mathbb{C}_q \otimes (V^\dagger)^*,$$

and the element $u \otimes_{\mathbb{H}} v$ is just the complex tensor product $\iota_U(u) \otimes_{\mathbb{C}_q} \iota_V(v) \in \mathbb{C}_q \otimes (U^\dagger)^* \otimes (V^\dagger)^*$. So Lemma 4.1.7 tells us that on complex subfields of $\mathbb{H}$, the quaternionic tensor product is the same as the complex tensor product.

4.1.3 Stable and Semistable A$\mathbb{H}$-Modules

In this section we define two special sorts of A$\mathbb{H}$-modules, which we shall call semistable and stable. These A$\mathbb{H}$-modules behave particularly well, and we can exploit their 'nice' properties to cement further the analogy between real and quaternionic algebra.

Definition 4.1.8 [J1, §8] Let U be a finite-dimensional A$\mathbb{H}$-module. We say that U is *semistable* if it is generated over $\mathbb{H}$ by the subspaces $U' \cap qU'$ for $q \in S^2$.

We can describe semistable A$\mathbb{H}$-modules by the following property:

Lemma 4.1.9 *Suppose that* U *is semistable, with* $\dim U = 4j$ *and* $\dim U' = 2j + r$*, for integers* j, r*. Then* $U' + qU' = U$ *for generic* $q \in S^2$*. Thus* $r \geq 0$*.*

The next logical step is to require this property for *all* $q \in S^2$, motivating the following definition:

Definition 4.1.10 Let U be a finite-dimensional A$\mathbb{H}$-module. We say that U is *stable* if $U = U' + qU'$ for all $q \in S^2$.

In effect, our definitions of stable and semistable A$\mathbb{H}$-modules act as a balance to Definition 4.1.1 by demanding that U' should not be too small. Many of the properties of semistable and stable A$\mathbb{H}$-modules can be characterised by exploring the properties of a particularly important type of A$\mathbb{H}$-module:

Definition 4.1.11 Let $q \in \mathbb{I} \setminus \{0\}$. Define an A$\mathbb{H}$-module X_q by $X_q = \mathbb{H}$, $X'_q = \{p \in \mathbb{H} : pq = -qp\}$. In other words, X'_q is the subspace of $\mathbb{H}$ which is perpendicular to $\mathbb{C}_q$ with respect to the standard scalar product.

We quote the following results, mainly taken from [J1, §8]:

- X_q is semistable, but not stable.
- $X_q = X_{\lambda q}$ for all $\lambda \in \mathbb{R} \setminus \{0\}$, but for $p \neq \lambda q$, X_p and X_q are not A$\mathbb{H}$-isomorphic to one another. There is thus a distinct A$\mathbb{H}$-module X_q given by each pair of antipodal points $\{q, -q\}$ for $q \in S^2$.

- There is a canonical AH-isomorphism $X_q \otimes_{\mathbb{H}} X_q \cong X_q$, but if $p \neq \lambda q$ then $X_p \otimes_{\mathbb{H}} X_q = \{0\}$.
- Let $\chi_q : X_q \to \mathbb{H}$ be the identity map on $\mathbb{H}$. Then χ_q and $\mathrm{id} \otimes_{\mathbb{H}} \chi_q : U \otimes_{\mathbb{H}} X_q \to U \otimes_{\mathbb{H}} \mathbb{H} \cong U$ are injective AH-morphisms.
- There is an isomorphism $(U \otimes_{\mathbb{H}} X_q)' \cong U' \cap qU' \cong \mathbb{C}_q^n$. It follows that $U \otimes_{\mathbb{H}} X_q \cong nX_q$.
- Therefore if $q \in S^2$ and U is an AH-module with $\dim U = 4j$ and $\dim U' = 2j + r$, then $U \otimes_{\mathbb{H}} X_q \cong nX_q$ with $n \geq r$.
- If U is semistable then $U \otimes_{\mathbb{H}} X_q \cong rX_q$ for generic $q \in S^2$, by Lemma 4.1.9.
- An AH-module U is stable if and only if $U \otimes_{\mathbb{H}} X_q \cong rX_q$ for all $q \in S^2$.
- It follows that if U and V are stable AH-modules then $U \otimes_{\mathbb{H}} V \otimes_{\mathbb{H}} X_q \cong U \otimes_{\mathbb{H}} (sX_q) \cong rsX_q$ for all $q \in S^2$, so using the associativity of the quaternionic tensor we infer that $U \otimes_{\mathbb{H}} V$ is a stable AH-module.

The AH-module X_q is an important bridge from quaternionic to complex algebra. In effect, $X_q = \mathbb{C}_q \oplus \mathbb{C}_q^{\perp}$, and the operation '$\otimes_{\mathbb{H}} X_q$' converts an AH-module U into copies of X_q, so it turns the AH-module structure on U which is quaternionic information into a set of what are effectively complex vector spaces.

It is clear that all stable AH-modules are semistable. There is a sense in which the X_q's are the 'only' class of AH-modules which are semistable but not stable, due to the following Proposition:

Proposition 4.1.12 *[J1, 8.8]: Let V be a finite-dimensional* AH*-module. Then V is semistable if and only if $V \cong U \oplus (\bigoplus_{i=1}^{l} X_{q_i})$, where U is stable and $q_i \in S^2$.*

Joyce also shows that generic AH-modules with appropriate dimensions are stable or semistable:

Lemma 4.1.13 *[J1, 8.9] Let j, r be integers with $0 \leq r \leq j$. Let $U = \mathbb{H}^j$ and let U' be a real vector subspace of U with $\dim U' = 2j + r$. For generic subspaces U', (U, U') is a semistable* AH*-module. If $r > 0$ then for generic subspaces U', (U, U') is a stable* AH*-module.*

The benefits of working with stable and semistable AH-modules become increasingly apparent as one becomes more familiar with the theory. For now, we will quote the following theorems:

Theorem 4.1.14 *[J1, 9.1] Let U and V be stable* AH*-modules with*

$$\dim U = 4j, \quad \dim U' = 2j + r, \quad \dim V = 4k \quad \text{and} \quad \dim V' = 2k + s. \tag{4.6}$$

Then $U \otimes_{\mathbb{H}} V$ is a stable AH*-module with $\dim(U \otimes_{\mathbb{H}} V) = 4l$ and $\dim(U \otimes_{\mathbb{H}} V)' = 2l + t$, where $l = js + rk - rs$ and $t = rs$.*

Proof.(Sketch of formula for total dimension) The proof works along the following lines. If $\dim U = 4j$ and $\dim U' = 2j + r$, then $\dim(U^\dagger)^* = 2j - r$ and similarly $\dim(V^\dagger)^* = 2k - s$. So $\dim(\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*) = 4(2j - r)(2k - s)$.

Let $A = \iota_U(U) \otimes (V^\dagger)^*$ and $B = (U^\dagger)^* \otimes \iota_V(V)$, so that $U \otimes_{\mathbb{H}} V = A \cap B$. Then $\dim A = 4j(2k - s)$ and $\dim B = 4k(2j - r)$, and $\dim(A + B) = \dim A + \dim B - \dim(A \cap B)$. Now, if $r, s \geq 0$ then $\dim A + \dim B \geq \dim(\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*)$, and so if the subspaces A and B are suitably transverse in $\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$ we expect that $A + B = \mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$, in which case

$$\begin{aligned} \dim(U \otimes_{\mathbb{H}} V) &= \dim A + \dim B - \dim(\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*) \\ &= 4(js + rk - rs). \end{aligned}$$

The rest of Joyce's proof consists of showing that if U and V are stable then this intersection *is* transverse and finding $(U \otimes_{\mathbb{H}} V)'$, from which it is easy to see that $U \otimes_{\mathbb{H}} V$ is stable. ∎

In fact, these dimension formulae still hold if V is only semistable. Theorem 4.1.14 and Proposition 4.1.12 combine to give the following:

Corollary 4.1.15 *[J1, 9.3] Let U, V be semistable* AH*-modules. Then $U \otimes_{\mathbb{H}} V$ is semistable.*

Thus both stable and semistable AH-modules form subcategories of the tensor category of AH-modules, closed under direct and tensor products. Let U be a stable AH-module, with $\dim U = 4j$ and $\dim U' = 2j + r$. We define r to be the *virtual dimension* of U. Then Proposition 4.1.12 shows that the virtual dimension of $U \otimes_{\mathbb{H}} V$ is the product of the virtual dimensions of U and V. We end this section by quoting the following result:

Proposition 4.1.16 *[J1, 9.6] Let U be a stable* AH*-module, with $\dim U = 4j$ and $\dim U' = 2j + r$. Let n be a positive integer. Then $S^n_{\mathbb{H}} U$ and $\Lambda^n_{\mathbb{H}} U$ are stable* AH*-modules, with $\dim(S^n_{\mathbb{H}} U) = 4k$, $\dim(S^n_{\mathbb{H}} U)' = 2k + s$, $\dim(\Lambda^n_{\mathbb{H}} U) = 4l$ and $\dim(\Lambda^n_{\mathbb{H}} U)' = 2l + t$, where*

$$k = (j-r)\binom{r+n-1}{n-1} + \binom{r+n-1}{n}, \quad s = \binom{r+n-1}{n}, \quad l = (j-r)\binom{r-1}{n-1} + \binom{r}{n}, \quad t = \binom{r}{n}.$$

4.2 Duality in Quaternionic Algebra

The objects which are naturally dual to AH-modules are called SH-modules. Under certain circumstances an SH-module can also be regarded as an AH-module. In this case we obtain interesting algebraic results which use dual AH-modules to tell us about AH-morphisms between AH-modules.

4.2.1 S$\mathbb{H}$-modules

In this section we will describe the class of objects which are dual to A$\mathbb{H}$-modules. These will be called *strengthened* $\mathbb{H}$*-modules*, or S$\mathbb{H}$-modules. S$\mathbb{H}$-modules are introduced by Quillen in [Q], as a link between sheaves and A$\mathbb{H}$-modules. Being very much part of quaternionic algebra rather than sheaf theory, we discuss them in this section.

Let (U, U') be an A$\mathbb{H}$-module. Then $u \in U$ is completely determined by the values of $\alpha(u)$ for $\alpha \in U^\dagger$, and if u is determined then so is $\beta(u)$ for all $\beta \in U^\times$. So for all $\beta \in U$, $\beta(u)$ is determined by the action of $U^\dagger$ on u. Since the only other structure present is the $\mathbb{H}$-action, each $\beta \in U^\times$ must be an $\mathbb{H}$-linear combination of elements of $U^\dagger$, so $U^\times$ is generated over $\mathbb{H}$ by $U^\dagger$. The converse is clearly true as well — if every $\beta \in U^\times$ is of the form $q \cdot \alpha$ for some $\alpha \in U^\dagger$, then (U, U') is an A$\mathbb{H}$-module by Definition 4.1.1. The natural dual to an A$\mathbb{H}$-module is thus an $\mathbb{H}$-module equipped with a generating real subspace.

Definition 4.2.1 Let Q be a left $\mathbb{H}$-module and $Q^\dagger$ a real linear subspace of Q. We say that the pair $(Q, Q^\dagger)$ is a *strengthened* $\mathbb{H}$*-module* or S$\mathbb{H}$*-module* if Q is generated over $\mathbb{H}$ by $Q^\dagger$.

If $U = (U, U')$ is an A$\mathbb{H}$-module then $(U^\times, U^\dagger)$ is an S$\mathbb{H}$-module, the S$\mathbb{H}$*-module corresponding to* U. Just as we sometimes write U for the A$\mathbb{H}$-module (U, U'), we will often write $U^\times$ for the S$\mathbb{H}$-module $(U^\times, U^\dagger)$.

A further link between these two ideas is provided by choosing a (hyperhermitian) metric on the A$\mathbb{H}$-module (U, U'), giving an $\mathbb{H}$-module isomorphism $U \cong U^\times$. This in turn identifies $U^\dagger$ with $(U^\dagger)^*$, and so realises $(U^\dagger)^*$ as a subspace of U which is perpendicular to U'. We see that choosing a metric gives us a decomposition $U \cong U' \oplus (U^\dagger)^*$, where $(U, (U^\dagger)^*)$ is an S$\mathbb{H}$-module. Every A$\mathbb{H}$-module can thus be regarded as an S$\mathbb{H}$-module — the point of view depends on whether we think of U' or $(U')^\perp$ as the 'special' subspace. The simplest example is that of the quaternions themselves: we can regard them as the A$\mathbb{H}$-module $(\mathbb{H}, \mathbb{I})$ or the S$\mathbb{H}$-module $(\mathbb{H}, \mathbb{R})$, and these definitions are exactly equivalent.

We define morphisms and quaternionic tensor products for S$\mathbb{H}$-modules, by taking the definitions from their corresponding A$\mathbb{H}$-modules: the whole theory works in exactly the same way. For example, if (U, U') and (V, V') are A$\mathbb{H}$-modules and $\phi : U \to V$ is an A$\mathbb{H}$-morphism, then $(U^\times, U^\dagger)$ and $(V^\times, V^\dagger)$ are S$\mathbb{H}$-modules and the dual $\mathbb{H}$-morphism $\phi^\times : V^\times \to U^\times$ satisfies $\phi^\times(V^\dagger) \subseteq U^\dagger$, which makes $\phi^\times$ an S$\mathbb{H}$*-morphism.* We write $U^\times \otimes^{\mathbb{H}} V^\times$ for the quaternionic tensor product of two S$\mathbb{H}$-modules; in other words we define

$$U^\times \otimes^{\mathbb{H}} V^\times = (U \otimes_{\mathbb{H}} V)^\times. \tag{4.7}$$

Another useful example is given by stable and semistable A$\mathbb{H}$-modules. An A$\mathbb{H}$-module is stable (respectively semistable) if and only if it satisfies the identity $U = U' + qU'$ for all (respectively for generic) $q \in \mathbb{I}$. But

$$U' + qU' = U \iff (U')^\perp \cap q(U')^\perp = \{0\},$$

where $(U')^\perp$ is the subspace orthogonal to U' with respect to a (hyperhermitian) metric on U. Since there is an S$\mathbb{H}$-isomorphism $(U, (U')^\perp) \cong (U^\times, U^\dagger)$, we see that

$$U' + qU' = U \iff U^\dagger \cap qU^\dagger = \{0\}.$$

Definition 4.2.2 An $\mathbb{SH}$-module is called *stable* (respectively *semistable*) if and only if it has the property that $U^\dagger \cap qU^\dagger = \{0\}$ for all (respectively for generic) $q \in \mathbb{I}$.

This is sometimes easier to demonstrate than the property for the corresponding $\mathbb{AH}$-modules. We shall work happily with either $\mathbb{AH}$-modules or $\mathbb{SH}$-modules according to the needs of each situation, since their theories are interchangeable.

4.2.2 Dual $\mathbb{AH}$-modules

In this section we take a new step and ask what happens if we consider $(U^\times, U^\dagger)$ as an $\mathbb{AH}$-module — the *dual* $\mathbb{AH}$*-module* of (U, U'). There are immediate attractions to this approach. If V is a vector space over the commutative field $\mathbb{F}$ then we define the dual space V^* to be the space of $\mathbb{F}$-linear maps $\phi : V \to \mathbb{F}$. In the same way, if U is an $\mathbb{H}$-module we define $U^\times$ to be the space of $\mathbb{H}$-linear maps $\phi : U \to \mathbb{H}$. Once we also have a real subspace $U' \subset U$ we define $U^\dagger$ to be the set of maps

$$U^\dagger = \{\alpha \in U^\times : \alpha(u) \in \mathbb{I} \text{ for all } u \in U'\}.$$

But an $\mathbb{H}$-linear map $\alpha : U \to \mathbb{H}$ such that $\alpha(U') \subseteq \mathbb{I}$ is precisely an $\mathbb{AH}$-morphism from U into $\mathbb{H}$. Thus the space $(U^\times, U^\dagger)$ consists of two sets of maps $\phi : U \to \mathbb{H}$, namely the $\mathbb{H}$-linear maps and $\mathbb{AH}$-morphisms respectively. This suggests that defining $(U^\times, U^\dagger)$ to be the dual $\mathbb{AH}$-module of (U, U') could be a good quaternionic analogue of the concept of a dual vector space in real or complex algebra.

There is an obvious possible catch: $(U^\times, U^\dagger)$ might not even *be* an $\mathbb{AH}$-module! Thus if we are to talk about dual $\mathbb{AH}$-modules, we need to discern which $\mathbb{AH}$-modules have well-defined duals. We have in fact already done this in Section 4.2.1.

Lemma 4.2.3 *The* $\mathbb{AH}$*-module* (U, U') *has a well defined dual* $\mathbb{AH}$*-module* $(U^\times, U^\dagger)$ *if and only if* (U, U') *is also an* $\mathbb{SH}$*-module.*

Proof. In Section 4.2.1 we showed that (U, U') is an $\mathbb{AH}$-module if and only if $(U^\times, U^\dagger)$ is an $\mathbb{SH}$-module. We simply reverse this argument: $(U^\times, U^\dagger)$ is an $\mathbb{AH}$-module if and only if the $\mathbb{AH}$-module (U, U') is also an $\mathbb{SH}$-module. ∎

Definition 4.2.4 Let U be an $\mathbb{H}$-module and U' a real subspace of U. We say that the pair (U, U') is a *strengthened augmented* $\mathbb{H}$*-module*, or $\mathbb{SAH}$*-module*, if (U, U') is both an $\mathbb{AH}$-module and an $\mathbb{SH}$-module.

A comprehensive way to sum this up is to say that (U, U') is an $\mathbb{AH}$-module if and only if it has no submodule isomorphic to $(\mathbb{H}, \mathbb{H})$, an $\mathbb{SH}$-module if and only if it has no submodule isomorphic to $(\mathbb{H}, \{0\})$, and an $\mathbb{SAH}$-module if and only if it has no submodule isomorphic to either $(\mathbb{H}, \mathbb{H})$ or $(\mathbb{H}, \{0\})$. Unless otherwise stated, when we refer to properties of an $\mathbb{SAH}$-module such as stability, we mean this in terms of the structure of U as an $\mathbb{AH}$-module.

Example 4.2.5 Any semistable $\mathbb{AH}$-module U has the property that $U = U' + qU'$ for generic $q \in \mathbb{I}$, so U is also an $\mathbb{SH}$-module and so an $\mathbb{SAH}$-module.

Definition 4.2.6 Let $U = (U, U')$ be an A$\mathbb{H}$-module which is also an SA$\mathbb{H}$-module. Then we define $U^\times = (U^\times, U^\dagger)$ to be the *dual* A$\mathbb{H}$-*module* of U.

With this definition, it is clear that $U^\times$ is also an SA$\mathbb{H}$-module. For finite dimensional U, there are canonical isomorphisms $U \cong (U^\times)^\times$ and $(U^\times)^\dagger \cong U'$.

Definition 4.2.7 Let (U, U') be an A$\mathbb{H}$-module. Then U is called *antistable* if and only if $U' \cap qU' = \{0\}$ for all $q \in \mathbb{I}$.

It is clear that antistable A$\mathbb{H}$-modules are almost always dual to stable A$\mathbb{H}$-modules. The only irreducible exception is the A$\mathbb{H}$-module $(\mathbb{H}, \{0\})$, which we regard as an antistable A$\mathbb{H}$-module in spite of the fact that its dual $(\mathbb{H}, \mathbb{H})$ is not an A$\mathbb{H}$-module.

4.2.3 Spaces of A$\mathbb{H}$-morphisms and Duality

The space $U^\dagger$ is, as we have remarked, the space of A$\mathbb{H}$-morphisms $\phi : U \to \mathbb{H}$. This fact is an example of a more general result which makes the theory of dual spaces in quaternionic algebra particularly useful. Let A and B be (free, finite dimensional) modules over the commutative ring R, and let $\mathrm{Hom}_R(A, B)$ denote the space of R-linear maps $\phi : A \to B$. It is well-known that there is a canonical isomorphism $\mathrm{Hom}_R(A, B) \cong A^* \otimes_R B$.

There is an analogous result in quaternionic algebra. We start with the following definition:

Definition 4.2.8 Let U and V be A$\mathbb{H}$-modules. Then $\mathrm{Hom}_{\mathrm{A}\mathbb{H}}(U, V)$ denotes the space of A$\mathbb{H}$-morphisms from U into V.

Here is the main result of this section:

Theorem 4.2.9 *Let U be an* SA$\mathbb{H}$-*module and V be an* A$\mathbb{H}$-*module. Then there is a canonical isomorphism*

$$\mathrm{Hom}_{\mathrm{A}\mathbb{H}}(U, V) \cong (U^\times \otimes_{\mathbb{H}} V)'.$$

Proof. Let $\phi \in (U^\times \otimes_{\mathbb{H}} V)'$. Then

$$\phi \in (\iota_{U^\times}(U^\times) \otimes (V^\dagger)^*) \cap ((U')^* \otimes \iota_V(V)) \cap (\mathbb{I} \otimes (U')^* \otimes (V^\dagger)^*),$$

or equivalently

$$\phi \in (\iota_{U^\times}(U^\dagger) \otimes (V^\dagger)^*) \cap ((U')^* \otimes \iota_V(V')).$$

Consider $\phi \in \iota_{U^\times}(U^\dagger) \otimes (V^\dagger)^*$. The mapping $\iota_{U^\times}$ identifies $U^\dagger$ with $\iota_{U^\times}(U^\dagger)$. Using this and the canonical isomorphism of real vector spaces $\mathrm{Hom}(A, B) \cong A^* \otimes B$ we see that ϕ is exactly equivalent to an real linear map

$$\Phi : (U^\dagger)^* \to (V^\dagger)^*,$$

which in turn is equivalent to an $\mathbb{H}$-linear map

$$\Phi_{\mathbb{H}} : \mathbb{H} \otimes (U^\dagger)^* \to \mathbb{H} \otimes (V^\dagger)^*,$$

which by definition is an AH-morphism.

Every AH-morphism $\psi : U \to V$ is equivalent to an AH-morphism $\Psi : \mathbb{H} \otimes (U^\dagger)^* \to \mathbb{H} \otimes (V^\dagger)^*$ with the property that $\Psi : \iota_U(U) \to \iota_V(V)$. Using the fact that $\phi \in (U')^* \otimes \iota_V(V')$ and the natural identification $U' \cong \iota_U(U')$ we see that

$$\Phi_{\mathbb{H}} : \iota_U(U') \to \iota_V(V').$$

Since U is an SAH-module, $\iota_U(U)$ is generated over $\mathbb{H}$ by $\iota_U(U')$, from which it follows that $\Phi_{\mathbb{H}} : \iota_U(U) \to \iota_V(V)$. Thus $\Phi_{\mathbb{H}}$ is equivalent to an AH-morphism from U to V. Reversing these steps, we can construct an element of $(U^\times \otimes_{\mathbb{H}} V)'$ from each AH-morphism from U to V. ∎

The quaternionic tensor product can be used in this way to tell us about spaces of AH-morphisms, which is another piece of evidence suggesting that Joyce's definition of the quaternionic tensor product is the right one for AH-modules.

4.3 Real Subspaces of Complex Vector Spaces

We have seen that choosing different real subspaces U' of an $\mathbb{H}$-module $U \cong \mathbb{H}^n$ gives rise to different algebraic properties. As always with the quaternions, it is useful to compare this situation with that of the complex numbers. As an instructive example (and a bit of light relief!) we shall give a classification of real linear subspaces of complex vector spaces up to complex linear isomorphism. In other words, we classify the possible orbits of a subspace $\mathbb{R}^k$ of $\mathbb{C}^n$ under the action of $\mathrm{GL}(n, \mathbb{C})$. This is not difficult, though as far as the author can tell both the problem and its solution are original.

Theorem 4.3.1 *Let $U' \cong \mathbb{R}^k$ be a real linear subspace of $\mathbb{C}^n$. Then we can choose a complex basis $\{e^j : j = 1, \dots, n\}$ of $\mathbb{C}^n$ such that*

$$U' = \langle e^1, \dots, e^p, ie^1, \dots, ie^q \rangle_{\mathbb{R}},$$

where $q \leq p \leq n$, $p + q = k$.

Proof. Both $U' + iU'$ and $U' \cap iU'$ are complex subspaces of $\mathbb{C}^n$. Let $U' + iU' \cong \mathbb{C}^p$ and let $U' \cap iU' \cong \mathbb{C}^q$. Then $\dim_{\mathbb{R}}(U') = p + q$.

Choose a complex basis $\{e^1, \dots, e^q\}$ for $U' \cap iU'$. Then $\{e^1, ie^1, \dots, e^q, ie^q\}$ is a real basis for $U' \cap iU'$. Extend this to a real basis $\{e^1, ie^1, \dots, e^q, ie^q, e^{q+1}, \dots, e^p\}$ for U'. Then $\{e^1, \dots, e^p\}$ spans $U' + iU' \cong \mathbb{C}^p$ (over $\mathbb{C}$), and so $\{e^1, \dots, e^p\}$ is linearly independent over $\mathbb{C}$. The result follows. ∎

It is easy to see from this theorem that choosing a real subspace of $\mathbb{C}^n$ is always compatible with the complex structure, in the sense that each basis vector of the real subspace U' can be chosen to lie in one and only one copy of $\mathbb{C}$. The pair $(U, U') \cong (\mathbb{C}^n, \mathbb{R}^k)$ can always be completely reduced to a direct sum of copies of $\mathbb{C}$, each of which contains 0, 1 or 2 basis vectors for U'.

The situation is very different for $\mathbb{H}$-modules. For example, consider the AH-module (U, U') with

$$U = \mathbb{H}^2 \qquad \text{and} \qquad U' = \langle (1,0)(0,1), (i_1, i_2) \rangle.$$

There is no way to decompose U into two separate copies of $\mathbb{H}$, the first of which contains 2 basis vectors for U' and the second of which contains the remaining basis vector.

This motivates the following definition:

Definition 4.3.2 An A$\mathbb{H}$-module (U, U') is *irreducible* if and only if it cannot be written as a direct sum of two non-trivial A$\mathbb{H}$-modules, *i.e.* there are no two non-trivial A$\mathbb{H}$-modules (U_1, U_1'), (U_2, U_2') such that $(U, U') = (U_1 \oplus U_2, U_1' \oplus U_2')$.

The classification of irreducible A$\mathbb{H}$-modules up to A$\mathbb{H}$-isomorphism is a much more difficult problem than its analogue for complex vector spaces. It will be addressed in the next section using a class of algebraic objects called K-modules.

4.4 K-modules

A K-module is an algebraic object based on a pair of complex vector spaces. Real and quaternionic vector spaces, as so often, occur as complex vector spaces with particular structure maps.

Definition 4.4.1 [Q, 4.1] A *K-module* is a pair (W, V) of (finite dimensional) complex vector spaces together with a linear map $e : W \to H \otimes V$, where $H \cong \mathbb{C}^2$ is the basic representation of $\mathrm{GL}(2, \mathbb{C})$.

The reason why K-modules are important to quaternionic algebra is that in the presence of suitable structure maps, some K-modules are equivalent to A$\mathbb{H}$-modules. This allows us to use the classification of irreducible K-modules, a problem with a known solution, to write down all irreducible A$\mathbb{H}$-modules very explicitly. This link between K-modules and A$\mathbb{H}$-modules was discovered by Quillen [Q]. Quillen is more interested in an interpretation of quaternionic algebra in terms of sheaves over the Riemann sphere, a powerful theory which we will review in the next section. We follow a slightly different approach from that in Quillen's paper to obtain a more immediate link between K-modules and A$\mathbb{H}$-modules.

A K-module $e : W \to H \otimes V$ is called indecomposable if it cannot be written as the direct sum of two non-trivial K-modules. A K-module morphism is a map $\phi : (W_1 \stackrel{e_1}{\to} H \otimes V_1) \longrightarrow (W_2 \stackrel{e_2}{\to} H \otimes V_2)$ which respects the K-module structure. A K-module can equivalently be defined as a pair of linear maps $e_1, e_2 : W \to V$. In this guise, a K-module is a representation of the *Kronecker quiver*. We recover the first definition by setting $e(w) = h_1 \otimes e_1(w) + h_2 \otimes e_2(w)$ where $\{h_1, h_2\}$ is a basis for H. Representations of the Kronecker quiver are discussed in Benson's book [Ben, §4.3]. The important result is the following classification theorem of Kronecker (which we have summarised slightly), which shows that every indecomposable K-module is isomorphic to one of three basic types.

Theorem 4.4.2 *[Ben, p. 101] Let $e_1, e_2 : W \to V$ be a pair of linear maps constituting an indecomposable K-module. Then one of the following holds:*

(i) *The vector spaces W and V have the same dimension. In this case, if* $\det e_1 \neq 0$

then e_1 and e_2 can be written in the form

$$e_1 \mapsto \mathrm{id} \qquad e_2 \mapsto \begin{pmatrix} \alpha & & & 0 \\ 1 & \alpha & & \\ & \ddots & \ddots & \\ 0 & & 1 & \alpha \end{pmatrix}.$$

If $\det e_1 = 0$ *a modification is necessary which in some sense corresponds to the 'rational canonical form at infinity'.*

(ii) *The dimension of W is one larger than the dimension of V, and bases may be chosen so that e_1 and e_2 are represented by the matrices*

$$e_1 = \begin{pmatrix} 1 & & 0 & 0 \\ & \ddots & & \vdots \\ 0 & & 1 & 0 \end{pmatrix} \qquad e_2 = \begin{pmatrix} 0 & 1 & & 0 \\ \vdots & & \ddots & \\ 0 & 0 & & 1 \end{pmatrix}.$$

(iii) *The dimension of W is one smaller than the dimension of V, and bases may be chosen so that e_1 and e_2 are represented by the transposes of the above matrices.*

Definition 4.4.3 Define $\mathcal{X}_1^{n,\alpha}$ to be the indecomposable K-module of type (i) with $\dim V = n$ and α on the leading diagonal of e_2. (We write $\mathcal{X}_1^{n,\infty}$ for the case $\det e_1 = 0$). Define $\mathcal{X}_2^n$ to be the irreducible K-module of type (ii) with $\dim V = n$. Define $\mathcal{X}_3^n$ to be the irreducible K-module of type (iii) with $\dim V = n$.

Let σ_H be the standard quaternionic structure map on H defined by $\sigma_H(z_1 h_1 + z_2 h_2) = -\bar{z}_2 h_1 + \bar{z}_1 h_2$. This gives $H \otimes V$ the structure of a complex $\mathbb{H}$-module. Our aim is for the K-module $e : W \to H \otimes V$ to define a real subspace of a real $\mathbb{H}$-module. This is accomplished by compatible structure maps on W and V.

Definition 4.4.4 [Q, 11.2] An SK-module is a K-module $e : W \to H \otimes V$ equipped with antilinear operators σ_W and σ_V of squares 1 and -1 respectively, such that $e \cdot \sigma_W = (\sigma_H \otimes \sigma_V) \cdot e$.

Suppose $W \to H \otimes V$ is an SK-module. Then σ_W is a real structure on W, and its set of fixed points is the real vector space W^σ. The map $\sigma_H \otimes \sigma_V$ is also a real structure, and in the same way we define the real vector space $(H \otimes V)^\sigma$. Then $(H \otimes V)^\sigma$ is a real $\mathbb{H}$-module, whose $\mathbb{H}$-module structure is inherited from that on H, and $e(W^\sigma)$ is a real subspace of this $\mathbb{H}$-module. In many circumstances an SK-module is therefore equivalent to an A$\mathbb{H}$-module.

Example 4.4.5 Consider the K-module $\mathcal{X}_2^2$ which has $\dim W = 3$ and $\dim V = 2$. Let $\{w_1, w_2, w_3\}$ be a basis for W and let $\{v_1, v_2\}$ be a basis for V. We have

$$e(w_1) = h_1 \otimes v_1, \qquad e(w_2) = h_1 \otimes v_2 + h_2 \otimes v_1, \qquad e(w_3) = h_2 \otimes v_2.$$

Let σ_V be the standard quaternionic structure on V, so $\sigma_V(v_1) = v_2$ and $\sigma_V(v_2) = -v_1$. There is a compatible real structure σ_W on W given by $\sigma_W(w_1) = w_3$, $\sigma(w_3)_W = w_1$

and $\sigma_W(w_2) = -w_2$, so that with these structure maps $\mathcal{X}_2^2$ is an SK-module. We have real vector spaces

$$W^\sigma = \langle w_1 + w_3, iw_2, i(w_1 - w_3) \rangle$$

and

$$(H \otimes V)^\sigma = \left\langle \begin{array}{ll} h_1 \otimes v_2 - h_2 \otimes v_1 & i(h_1 \otimes v_2 + h_2 \otimes v_1) \\ h_1 \otimes v_1 + h_2 \otimes v_2 & i(h_1 \otimes v_1 - h_2 \otimes v_2) \end{array} \right\rangle .$$

An $\mathbb{H}$-module isomorphism $(H \otimes V)^\sigma \cong \mathbb{H}$ is obtained by mapping these basis vectors to 1, i_1, i_2 and i_3 respectively. Under this isomorphism the real subspace $e(W^\sigma)$ is mapped to the imaginary quaternions $\mathbb{I}$. This demonstrates explicitly that the SK-module $\mathcal{X}_2^2$ is equivalent to the A$\mathbb{H}$-module $(\mathbb{H}, \mathbb{I}) = \mathbb{H}$.

Every real subspace U' of a quaternionic vector space U can be obtained in this fashion, so a classification of SK-modules gives a classification of A$\mathbb{H}$-modules.

Corollary 4.4.6 *Every indecomposable SK-module is isomorphic to one of the following:*

(A) *The direct sum of a pair of K-modules of type (i) of the form* $\mathcal{X}_1^{n,\alpha} \oplus \mathcal{X}_1^{n,-\bar{\alpha}^{-1}}$.

(B) *An indecomposable K-module of type (ii) or (iii) with* $\dim V$ *even; in other words* $\mathcal{X}_2^{2m}$ *or* $\mathcal{X}_3^{2m}$.

(C) *An indecomposable K-module of type (ii) or (iii) with* $\dim V$ *odd, tensored with the basic representation* H *equipped with its standard structure map* σ_H*; in other words* $\mathcal{X}_2^{2m+1} \otimes H$ *or* $\mathcal{X}_3^{2m+1} \otimes H$.

Proof. This follows from Theorem 4.4.2 and explicit calculations using standard structure maps on the vector spaces V. ∎

It remains to check which pairs (U, U') arising in this fashion are A$\mathbb{H}$-modules. As noted earlier, a pair (U, U') fails to be an A$\mathbb{H}$-module if and only if it has a subspace of the form $(\mathbb{H}, \mathbb{H})$. This pair is given by the SK-module $H \otimes \mathcal{X}_3^1$. The 'annihilating K-module' $\mathcal{X}_3^0$ also fails to give an A$\mathbb{H}$-module. Any SK-module containing neither of these indecomposables is equivalent to an A$\mathbb{H}$-module.

Here are some important facts about irreducible A$\mathbb{H}$-modules which can now be deduced:

- SK-modules of type (ii) correspond to stable A$\mathbb{H}$-modules and SK-modules of type (iii) correspond to antistable A$\mathbb{H}$-modules.
- Indecomposable SK-modules of type (i) of the form $\mathcal{X}_1^{1,\alpha} \oplus \mathcal{X}_1^{1,-\bar{\alpha}^{-1}}$ correspond to the semistable A$\mathbb{H}$-modules X_q.
- The irreducible stable A$\mathbb{H}$-module corresponding to $\mathcal{X}_2^{2m}$ has $\dim U = 4m$ and $\dim U' = 2m + 1$. Thus $U \cong \mathbb{H}^m$ and the virtual dimension of U is 1.
- The irreducible stable A$\mathbb{H}$-module of the form $H \otimes \mathcal{X}_2^{2m+1}$ has $\dim U = 4(2m+1)$ and $\dim U' = 4(m+1)$. Thus $U \cong \mathbb{H}^{2m+1}$ and the virtual dimension of U is 2.
- This shows that there is an irreducible stable A$\mathbb{H}$-module with virtual dimension 1 in every dimension and an irreducible stable A$\mathbb{H}$-module with virtual dimension 2 in every odd dimension.

- The isomorphism class of an irreducible stable $\mathbb{AH}$-module is thus uniquely determined by the dimension and virtual dimension of U.

4.5 The Sheaf-Theoretic approach of Quillen

Much of Joyce's quaternionic algebra can be described using (coherent) sheaves over the complex projective line $\mathbb{C}P^1$. This interpretation is due to Daniel Quillen [Q]. Quillen's paper works by recognising that certain exact sequences of sheaf cohomology groups are K-modules. Thus in the presence of certain structure maps, we obtain $\mathbb{SH}$-modules. (Quillen deals primarily with $\mathbb{SH}$-modules rather than $\mathbb{AH}$-modules.) Quillen uses slightly different structure maps from those we used in the previous section to obtain SK-modules, but the resulting theory is exactly the same.

The most interesting new result in this section is that the equivalence between sheaves and $\mathbb{SH}$-modules respects tensor products, enabling us to calculate the quaternionic tensor product of two $\mathbb{SH}$-modules from knowing the tensor products of the corresponding sheaves. Thus by the end of this section we will have succeeded in classifying all $\mathbb{AH}$-modules and their tensor products.

4.5.1 Sheaves on the Riemann Sphere

We describe the algebraic geometry of (coherent) sheaves over $\mathbb{C}P^1$. [1] Quillen demonstrates that every coherent sheaf over $\mathbb{C}P^1$ is the direct sum of a holomorphic vector bundle and a torsion sheaf (one whose support is finite). [2] These summands factorise very neatly — every torsion sheaf is the sum of indecomposable sheaves supported at a single point, and every holomorphic vector bundle is a sum of holomorphic line bundles. We will describe the vector bundles first.

Every holomorphic line bundle over $\mathbb{C}P^n$ is a tensor power of the hyperplane section bundle L [GH, p. 145]. In the case $n = 1$, we use the open cover of $\mathbb{C}P^1 = \mathbb{C} \cup \{\infty\}$ consisting of the two open sets $U_0 = \mathbb{C}P^1 \backslash \{\infty\}$ and $U_1 = \mathbb{C}P^1 \backslash \{0\}$. A holomorphic line bundle over $\mathbb{C}P^1$ is determined by a holomorphic transition function $\psi : U_0 \cap U_1 \to \mathbb{C}^*$, so $\psi : \mathbb{C}^* \to \mathbb{C}^*$. Two transition functions ψ and ψ' determine the same line bundle if and only if there exist non-vanishing holomorphic functions $f, g : \mathbb{C}^* \to \mathbb{C}^*$ such that

$$\psi' = \frac{f}{g}\psi.$$

Two functions are equivalent under this relation if and only if they have the same winding number, so each line bundle on $\mathbb{C}P^1$ is given by one of the transition functions $g(z) = z^n$, $n \in \mathbb{Z}$. The line bundle given by the transition function z^n is in fact $\bigotimes^n L$. Following standard notation, we write $\mathcal{O}(n)$ for the sheaf of its holomorphic sections. Thus $\mathcal{O} =$

[1] Background material can be found in [GH], which introduces sheaves and their cohomology [pp. 34-49], coherent sheaves [pp. 678-704], holomorphic vector bundles [pp. 66-71] and holomorphic line bundles [pp. 132-139]. A more thorough exposition of the differential geometry of holomorphic vector bundles, including many of the properties of sheaves used in Quillen's paper, can be found in Kobayashi's book [K].

[2] This uses the convention of identifying a holomorphic vector bundle with its sheaf of holomorphic sections.

$\mathcal{O}(0)$ is the structure sheaf of $\mathbb{C}P^1$. It is easy to see (by multiplying the transition functions together) that $L^n \otimes L^m \cong L^{n+m}$, or in sheaf-theoretic terms $\mathcal{O}(n) \otimes_{\mathcal{O}} \mathcal{O}(m) \cong \mathcal{O}(n+m)$.

Every holomorphic vector bundle over $\mathbb{C}P^1$ can be written as a direct sum of these line bundles, the summands being unique up to order. [3] Thus every holomorphic vector bundle E is a sum of irreducible line bundles, and can be written $E = \bigoplus_{-\infty}^{\infty} a_n L^n$, where the multiplicities a_n are unique (though the decomposition itself may not be).

This leaves us to consider sheaves which are supported at a finite set of points, which are called torsion sheaves. Quillen [Q, §2] demonstrates that every coherent sheaf over $\mathbb{C}P^1$ is the sum of a vector bundle and a torsion sheaf. Torsion sheaves themselves split into sheaves supported at one point only. Let $z \in \mathbb{C}P^1$ be such a point, and let $\mathcal{O}_z$ be the ring of germs of holomorphic functions at z. Define m_z to be the unique maximal ideal of $\mathcal{O}_z$ consisting of germs of functions whose first derivative vanishes at z. Every torsion sheaf splits into sheaves of the form $\mathcal{O}_z/(m_z)^n$, which we write $\mathcal{O}/m_z^n$ by extending m_z by $\mathcal{O}$ on the complement of z. We have the following Theorem:

Theorem 4.5.1 *[Q, 2.3] Any coherent sheaf over $\mathbb{C}P^1$ splits with unique multiplicities into indecomposable sheaves of the form $\mathcal{O}(n)$ for $n \in \mathbb{Z}$ and $\mathcal{O}/m_z^n$ for $n \geq 1$ and $z \in \mathbb{C}P^1$.*

Cohomology Groups and Exact Sequences

There are various ways to calculate the cohomology groups of these sheaves. [4] The method Quillen outlines uses the properties of exact sequences of sheaves. It is from the maps in these sequences that we obtain K-modules and thence $\mathbb{SH}$-modules.

Let $H \cong \mathbb{C}^2$ be the basic representation of $\mathrm{GL}(2, \mathbb{C})$. Then $\mathbb{C}P^1$ can be identified with the set of quotient lines of H and there is a basic exact sequence

$$0 \to \Lambda^2 H \otimes \mathcal{O}(-1) \to H \otimes \mathcal{O} \to \mathcal{O}(1) \to 0. \tag{4.8}$$

Tensoring (over $\mathcal{O}$) with the sheaf F and choosing an identification $\Lambda^2 H \cong \mathbb{C}$ yields the exact sequence

$$0 \to F(-1) \to H \otimes F \to F(1) \to 0, \tag{4.9}$$

where $F(n)$ denotes the sheaf $F \otimes_{\mathcal{O}} \mathcal{O}(n)$.

Example 4.5.2 If we put $F = \mathcal{O}(n), n \geq 0$, we have the exact sequence

$$0 \to \mathcal{O}(n-1) \to H \otimes \mathcal{O}(n) \to \mathcal{O}(n+1) \to 0. \tag{4.10}$$

We know that $H^0(\mathcal{O}) \cong \mathbb{C}$ (global holomorphic functions on $\mathbb{C}P(1)$) and that $H^0(\mathcal{O}(-1)) = 0$. By the exact sequences of (4.10) and induction, it follows that $H^0(\mathcal{O}(n)) \cong S^n(H)$ for $n \geq 0$ and zero otherwise, and there is an exact sequence of cohomology groups given by

$$0 \to S^{n-1}H \to H \otimes S^n H \to S^{n+1}H \to 0. \tag{4.11}$$

[3] This follows from the *Harder-Narasimhan filtration* of a holomorphic vector bundle E over any Riemann surface M [K, p. 137]

[4] For example [W, p.11], where $H^0(\mathcal{O}(n))$ is shown to be isomorphic to the (complex) vector space of homogeneous polynomials of degree n in 2 variables.

Torsion sheaves can be dealt with in a similar fashion, using a resolution involving the sheaf cohomology groups $H^0(F)$. It is easy to see that $H^0(\mathcal{O}/m_z^n) \cong \mathbb{C}^n$. We call the sheaves $\mathcal{O}(n)$ where $n \geq 0$, torsion sheaves, and sums thereof *regular sheaves.* For all regular sheaves F, the first cohomology group $H^1(F)$ is zero.

This leaves the sheaves $\mathcal{O}(n)$ where $n < 0$ and sums thereof, which we call *negative vector bundles.* These must be treated slightly differently, using the first cohomology groups $H^1(F)$. Since $H^0(\mathcal{O}(n)) = 0$ for $n < 0$, we obtain an exact sequence

$$0 \to H^1(\mathcal{O}(n-1)) \to H \otimes H^1(\mathcal{O}(n)) \to H^1(\mathcal{O}(n+1)) \to 0, \tag{4.12}$$

from which it follows that $H^1(\mathcal{O}(n)) \cong S^{-n-2}H$ for $n \leq -2$ and zero otherwise.

Sheaves and K-modules

Consider the regular sheaf $\mathcal{O}(n)$ for $n \geq 0$. The exact sequence (4.10) gives rise to an injection of cohomology groups $H^0(\mathcal{O}(n-1)) \to H^0(H \otimes \mathcal{O}(n))$ which takes the form $e : S^{n-1}H \to H \otimes S^nH$. This is clearly a K-module, the irreducible K-module $\mathcal{X}_3^{n+1}$. Similarly for the torsion sheaves $\mathcal{O}/m_x^n$ there is an injection $H^0(\mathcal{O}/m_x^n(-1)) \to H^0(H \otimes \mathcal{O}/m_x^n)$. This gives a K-module of type (i) and $\dim W = n$. Thus we obtain a K-module from each regular sheaf F which we call ξ^+F.

Let $\mathcal{O}(n)$, $n < 0$ be a negative vector bundle. The exact sequence (4.12) gives a map $H^1(\mathcal{O}(n-1)) \to H \otimes H^1(\mathcal{O}(n))$ which is equivalent to the indecomposable K-module $\mathcal{X}_3^{-n-1}$. Thus for any negative vector bundle G we obtain a K-module which we call ξ^-G.

Comparing the classifications of indecomposable sheaves and K-modules (Theorems 4.4.2 and 4.5.1), it is clear that these categories are equivalent. Quillen proves this in detail [Q, §§4,5] and uses the tensor product of sheaves $F \otimes_{\mathcal{O}} G$ to construct an equivalent tensor product operation for K-modules [Q, §6].

The rest of the programme begins to take shape. Some sheaves will correspond to SK-modules, from which we obtain S$\mathbb{H}$-modules. Quillen formulates this slightly differently from our treatment in Section 4.4. Let $W \to H \otimes V$ be a K-module. Instead of a real structure on W and a quaternionic structure on V, Quillen uses K-modules with a *quaternionic* structure σ_W on W and a *real* structure σ_V on V, such that the map e intertwines σ_W and $\sigma_H \otimes \sigma_V$. In this situation, W and $H \otimes V$ are $\mathbb{H}$-modules and e is an $\mathbb{H}$-linear map. He calls this structure a *σK-module.* A σK-module is not itself an S$\mathbb{H}$-module, but the inclusion of the real subspace V^σ in the cokernel $(H \otimes V)/W$ is an S$\mathbb{H}$-module if the K-module has no submodule for which the map e is surjective, in which case the K-module is called *reduced.*

There is a parallel description in terms of sheaves. Let $\sigma : z \mapsto \bar{z}^{-1}$ be the antipodal map on the Riemann sphere. This induces a map of sheaves $\sigma^* : F \mapsto \sigma^*(F)$ which we call the *σ-transform*, and allows us to define a 'σ-invariant sheaf' or just 'σ-sheaf'. For example, a torsion sheaf F is a σ-sheaf if it is supported at a finite set of points which is preserved by the antipodal map σ — so it must consist of sheaves of the form $\mathcal{O}/m_z^n \oplus \mathcal{O}/m_{\sigma(z)}^n$, where σ^* interchanges the two summands. Quillen investigates σ-sheaves thoroughly, and discovers that:

Proposition 4.5.3 *[Q, 10.7] Any σ-sheaf splits with unique multiplicities into the following irreducible σ-sheaves:*
(1) $\mathcal{O}/(m_z m_{\sigma(z)})^n$ for any pair $\{z, \sigma(z)\}$ of antipodal points and $n \geq 1$.
(2) $\mathcal{O}(2m)$ for $m \in \mathbb{Z}$.
(3) $\mathcal{O}(2m+1) \otimes H$ for $m \in \mathbb{Z}$.

The formal similarity between this result and Corollary 4.4.6 is clear. He also proves that:

Proposition 4.5.4 *[Q, 12.6] The categories of reduced σK-modules and $\mathrm{S}\mathbb{H}$-modules are equivalent.*

We can also use σ-sheaves to obtain SK-modules which lead directly to $\mathrm{S}\mathbb{H}$-modules.

Example 4.5.5 Let F be a regular σ-sheaf. Then we have the exact sequence

$$0 \to H^0(F) \to H \otimes H^0(F(1)) \to H^0(F(2)) \to 0, \tag{4.13}$$

and $H^0(F) \to H \otimes H^0(F(1))$ is an SK-module $e : W \to H \otimes V$. Taking real subspaces gives an $\mathrm{S}\mathbb{H}$-module which we call $\eta^+(F)$.

Example 4.5.6 Let G be a negative σ-vector bundle with no summand $H \otimes \mathcal{O}(-1)$ or $\mathcal{O}(-2)$. Then we have the exact sequence

$$0 \to H^1(G) \to H \otimes H^1(G(1)) \to H^1(G(2)) \to 0, \tag{4.14}$$

Just as in the previous example, this sequence gives an SK-module

$$H^1(G) \to H \otimes H^1(G(1)). \tag{4.15}$$

We call this $\mathrm{S}\mathbb{H}$-module $\eta^-(G)$. We also define $\eta^-(\mathcal{O}(-2)) = \eta^-(H \otimes \mathcal{O}(-1)) = 0$.

If we have a σ-sheaf $A = F+G$, with F and G as above, then $H^0(G) = H^0(G(1)) = 0$ and $H^1(F) = H^1(F(1)) = 0$; so $\eta^+(G) = \eta^-(F) = 0$. In theory, we could combine the functors η^+ and η^- into a single functor $\eta = \eta^+ + \eta^-$, since $\eta(A) = \eta^+(F) + \eta^-(G)$ as required.

4.5.2 Sheaves and the Quaternionic Tensor Product

We have seen how the tensor product of sheaves encourages us to define a 'reduced tensor product' operation for K-modules. It turns out that this tensor product agrees remarkably with the quaternionic tensor product for $\mathrm{S}\mathbb{H}$-modules. This is a considerable bonus from Quillen's theory — the correspondences between σ-sheaves, SK-modules and $\mathrm{S}\mathbb{H}$-modules allow us to compute tensor products in each category. The main theorem is as follows:

Theorem 4.5.7 *[Q, 7.1]* [5]

[5] Quillen proves this theorem for the tensor product of K-modules — the version given here is obtained by performing the simple translation into $\mathrm{S}\mathbb{H}$-modules.

Let F_i *be regular* σ*-sheaves and* G_i *be negative* σ*-vector bundles. Then*

$$\eta^+F_1 \otimes^{\mathbb{H}} \eta^+F_2 = \eta^+(F_1 \otimes_{\mathcal{O}} F_2),$$

$$\eta^+F_1 \otimes^{\mathbb{H}} \eta^-G_1 = \eta^-(F_1 \otimes_{\mathcal{O}} G_1)$$

and

$$\eta^-G_1 \otimes^{\mathbb{H}} \eta^-G_2 = \{0\}.$$

Since we will usually work with A$\mathbb{H}$-modules, we will often find ourselves using this theorem for the corresponding A$\mathbb{H}$-modules, which of course take the form $(\eta^+F)^\times$ and $(\eta^-G)^\times$.

Thus if F is a torsion σ-sheaf and G is a negative σ vector bundle, $(\eta^+F)^\times \otimes_{\mathbb{H}} (\eta^-G)^\times = \{0\}$. If $F = \mathcal{O}(m)$ and $G = \mathcal{O}(n)$ (possibly tensored with H if m or n is odd) with $m \geq 0$ and $n < -2$ then

$$(\eta^+F)^\times \otimes_{\mathbb{H}} (\eta^-G)^\times = \begin{cases} 0 & m+n \geq -2 \\ \eta^-(F \otimes_{\mathcal{O}} G)^\times & m+n \leq -3 \end{cases}$$

since $\eta^-(\mathcal{O}(k)) = \{0\}$ for $k \geq -2$.

The sheaf-theoretic approach is thus a very powerful tool for describing quaternionic algebra. For example, it is possible to obtain the dimension theorems of Section 4.1.3 by translating known results about the degree and rank of the tensor product of two sheaves.

Chapter 5

Quaternionic Algebra and Sp(1)-representations

The representations of the group Sp(1) occur in so many different situations, from Kähler geometry to particle physics, that they are by far the most ubiquitous Lie group representations in modern mathematical literature. Given this versatility, it is no surprise that these representations are a powerful tool in quaternionic algebra, especially since Sp(1) is just the group of unit quaternions and the Lie algebra $\mathfrak{sp}(1)$ can be identified with $\mathbb{I}$.

In this chapter, we will see how stable and antistable $\mathrm{A}\mathbb{H}$-modules can be handled using Sp(1)-representations. We encountered the germ of this idea in Section 3.4, where we came across the space $\mathbb{H} \otimes E_{k,r}$ and its splitting into two $\mathbb{H}$-submodules. Because the Riemann sphere $\mathbb{C}P^1$ can be described as the homogeneous space Sp(1)/U(1), the holomorphic sections of line bundles over $\mathbb{C}P^1$ are naturally Sp(1)-representations. The structure maps necessary to define an SK-module $e : W \to H \otimes V$ of type (ii) or (iii) arise from Sp(1)-representations on W, H and V. Not only do Sp(1)-representations underlie all of these phenomena — they also make the theory of quaternionic algebra very easy to predict and manipulate.

This point of view turns out to have fruitful applications in hypercomplex geometry, towards which our exposition is deliberately geared. We use representations to explain the structure of stable $\mathrm{A}\mathbb{H}$-modules and their tensor products. We also investigate the role of semistable $\mathrm{A}\mathbb{H}$-modules and their interaction with the Sp(1)-representation structure of stable $\mathrm{A}\mathbb{H}$-modules.

5.1 Stable $\mathrm{A}\mathbb{H}$-modules and Sp(1)-representations

5.1.1 Sp(1)-representations on the quaternions

As a motivating example, we review the case of the quaternions themselves, viewed as the stable $\mathrm{A}\mathbb{H}$-module $(\mathbb{H}, \mathbb{I})$. Recall the description of the quaternions as a tensor product of two Sp(1)-representations $\mathbb{H} \cong V_1 \otimes V_1$, given in Equation (3.17), where the left hand copy of V_1 gives the left $\mathbb{H}$-action, and the right-hand copy gives the right $\mathbb{H}$-action. [1] In other words, we think of $\mathbb{H}$ as an $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$-representation by defining

$$(p,q) : r \mapsto prq^{-1} \qquad r \in \mathbb{H}, \quad p,q \in \mathrm{Sp}(1) \subset \mathbb{H}.$$

Consider now the action of the diagonal Sp(1)-subgroup $\{(q,q) : q \in \mathrm{Sp}(1)\} \subset \mathrm{Sp}(1) \times \mathrm{Sp}(1)$ on $V_1 \otimes V_1$. The Clebsch-Gordon formula gives the splitting $V_1 \otimes V_1 \cong V_2 \oplus V_0$ (equivalent to the standard isomorphism $V \otimes V \cong S^2 V \oplus \Lambda^2 V$). Each of these summands inherits a real structure from the real structure on $V_1 \otimes V_1$ so we obtain the splitting

$$V_1 \otimes V_1 \cong V_2 \oplus V_0 \tag{5.1}$$

into real subspaces of dimensions three and one respectively, just as we would expect. This is the same as taking the action by conjugation $r \mapsto qrq^{-1}$, which as we know preserves the splitting $\mathbb{H} \cong \mathbb{I} \oplus \mathbb{R}$. For the quaternions, the A$\mathbb{H}$-module structure $\mathbb{H}' \cong \mathbb{I}$ and $(\mathbb{H}^\dagger)^* \cong \mathbb{R}$ is a concept which arises naturally when we take *both* the Sp(1) actions into account. It is this account of the A$\mathbb{H}$-module $\mathbb{H}$ which we will generalise to all stable A$\mathbb{H}$-modules.

This description of the quaternions is very similar to that of Example 4.4.5, where $\mathbb{H} \cong \mathcal{X}_2^2$. In terms of Sp(1)-actions, the basic vector space H we used so much in the previous chapter is simply a copy of the basic representation V_1.

5.1.2 Notation for Several Sp(1)-representations

It will be a sound investment at this point to introduce some notation to help us keep track of the structure of representations when we have several copies of Sp(1) acting on a vector space. We have already encountered the action of $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$ on $V_1 \otimes V_1$. Here we have two copies of Sp(1) acting, so there is already the possibility of ambiguity concerning which Sp(1) is acting on which V_1. We remove this ambiguity by writing upper-case superscripts with the groups and the representations, to make it clear which group is acting on which vector space.

For left $\mathbb{H}$-modules there will always be a left $\mathbb{H}$-action to consider. We will denote this by V_1^L, and the copy of Sp(1) which acts on this factor by $\mathrm{Sp}(1)^L$. Other copies of Sp(1) and other representations will be labelled with the letters M, N *etc.* So we would write the above example as

$$\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M \quad \text{acting on} \quad V_1^L \otimes V_1^M.$$

When we decompose such a representation using the Clebsch-Gordon formula, we are decomposing the action of the diagonal subgroup $\{(q,q)\} \subset \mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$. We will call this subgroup $\mathrm{Sp}(1)^{LM}$, thus stating explicitly of which two groups this is the diagonal subgroup. Similarly, we can combine superscripts for the representations to write

$$V_1^L \otimes V_1^M \cong V_2^{LM} \oplus V_0^{LM}.$$

[1] We are talking about the representation $V_1 \otimes V_1$ as a *real* representation on $\mathbb{R}^4$, implicitly using the induced map $\sigma_1 \otimes \sigma_1$ as a real structure on $V_1 \otimes V_1$. For more details, refer to Section 1.2.1.

This book-keeping comes into its own when we come to consider tensor products of many Sp(1)-representations. For example, if we have three copies of Sp(1) acting, we write this as

$$\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M \times \mathrm{Sp}(1)^N \quad \text{acting on} \quad V_1^L \otimes V_j^M \otimes V_k^N.$$

In this situation there are various diagonal actions we could be interested in, and we can join the superscripts as above to indicate exactly which one we are considering. For example, supposing we want to restrict to the diagonal subgroup in the first two copies of Sp(1), *i.e.* $\{(q,q)\} \times \mathrm{Sp}(1)^N$. We denote this subgroup $\mathrm{Sp}(1)^{LM} \times \mathrm{Sp}(1)^N$. We combine superscripts for the representations in the same way, so that we now have

$$\mathrm{Sp}(1)^{LM} \times \mathrm{Sp}(1)^N \quad \text{acting on} \quad (V_{j+1}^{LM} \oplus V_{j-1}^{LM}) \otimes V_k^N.$$

If, however, we considered the diagonal subgroup of the first and last copies of Sp(1), we would write this as

$$\mathrm{Sp}(1)^{LN} \times \mathrm{Sp}(1)^M \quad \text{acting on} \quad (V_{k+1}^{LN} \oplus V_{k-1}^{LN}) \otimes V_j^M.$$

This provides an unambiguous and (it is hoped) easy way to understand tensor products of several representations and their decompositions into irreducibles under different diagonal actions.

5.1.3 Irreducible stable $A\mathbb{H}$-modules

We have described the quaternions as an $\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$-representation. This allows us to interpret the primed part $\mathbb{H}' \cong \mathbb{I}$ as a representation of the diagonal subgroup of $\mathrm{Sp}(1)^{LM} \subset \mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$. In this section we will demonstrate how this idea can be adapted to describe more general stable $A\mathbb{H}$-modules. The basic idea is exactly the same — a stable $A\mathbb{H}$-module is a (real) $\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$-representation $V_1^L \otimes W^M$, where $W^M = \bigoplus_1^k a_j V_j^M$. The left $\mathbb{H}$-action is given by the action of the *left* subgroup $\mathrm{Sp}(1)^L \subset \mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$. The primed part is then a representation of the *diagonal* subgroup $\mathrm{Sp}(1)^{LM} \subset \mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$.

Let (U, U') be an irreducible stable (or antistable) $A\mathbb{H}$-module. Let M be the group of $A\mathbb{H}$-automorphisms of U whose (real) determinant is equal to 1. Using Theorem 4.2.9, we see that $M = \pm 1$ if the virtual dimension U is 1, and $M \cong \mathrm{Sp}(1)$ if the virtual dimension of U is 2.

Consider now the more general group G of real linear isomorphisms $\phi : U \to U$ such that:

- $\phi(U') = U'$,
- The (real) determinant of ϕ is 1,
- There exists some $q \in \mathrm{Sp}(1)$ such that $\phi(pu) = (qpq^{-1})\phi(u)$ for all $p \in \mathbb{H}, u \in U$.

Then G is a compact Lie group of which M is a normal subgroup. Because of this the Lie algebra of G splits into two orthogonal ideals

$$\mathfrak{g} = \mathfrak{sp}(1) \oplus \mathfrak{m}$$

and the exponential map determines a homomorphism $\rho : \mathrm{Sp}(1) \to G$. Since the elements of G map U' to itself, ρ is a representation of Sp(1) on U'.

Since every stable A$\mathbb{H}$-module is a sum of such irreducibles, there is an action of Sp(1) on U' for all stable A$\mathbb{H}$-modules U, and the irreducible decomposition of U as an A$\mathbb{H}$-module determines the irreducible decomposition of the Sp(1)-action on U', and vice versa. As we shall see, there is a unique irreducible stable A$\mathbb{H}$-module for each irreducible Sp(1)-representation.

Let U be a stable A$\mathbb{H}$-module and suppose that U' is preserved by some diagonal action of Sp(1). This diagonal action will be the result of the left $\mathbb{H}$-action on V_1 and some other Sp(1)-action on U. Thus our A$\mathbb{H}$-modules will follow the basic form

$$U = V_1^L \otimes V_m^M \tag{5.2}$$

as a representation of the group $\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$. The primed part of U is then the V_{n+1}^{LM} summand in the decomposition

$$V_1^L \otimes V_n^M \cong V_{n+1}^{LM} \oplus V_{n-1}^{LM}.$$

For example, recall the splitting

$$\mathbb{H} \otimes E_{k,r} \cong V_1 \otimes \varepsilon_{k,r}^n (V_{r+1} \oplus V_{r-1})$$

from Section 3.4. We see that each of these summands is an A$\mathbb{H}$-submodule of $\mathbb{H} \otimes E_{k,r}$. The larger (left-hand) submodule is stable; the smaller one is antistable. Both stable and antistable A$\mathbb{H}$-modules arise as $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$-representations.

This is not necessarily the case for A$\mathbb{H}$-modules which are neither stable nor antistable. For example, there is no representation of Sp(1) on X_q which couples with the left $\mathbb{H}$-action to give a representation of Sp(1) on X_q'.

A$\mathbb{H}$-modules of the form $\mathbf{V_1 \otimes V_{2m-1}}$

Consider an even-dimensional irreducible Sp(1)-representation V_{2m-1}. Because $\sigma = \sigma_1 \otimes \sigma_{2m-1}$ is a real structure, there is a real representation $V_1^L \otimes V_{2m-1}^M \cong \mathbb{R}^{4m}$ with a left $\mathbb{H}$-action defined by $q : a \otimes b \mapsto (qa) \otimes b$. Under the diagonal $\mathrm{Sp}(1)^{LM}$-action $q : a \otimes b \mapsto (qa) \otimes (qb)$, we have the splitting

$$V_1^L \otimes V_{2m-1}^M \cong V_{2m}^{LM} \oplus V_{2m-2}^{LM}, \tag{5.3}$$

which is a splitting of real vector spaces (technically we could write $(V_1^L \otimes V_{2m-1}^M)^\sigma \cong (V_{2m}^{LM})^\sigma \oplus (V_{2m-2}^{LM})^\sigma$).

Proposition 5.1.1 *The pair $(V_1^L \otimes V_{2m-1}^M, V_{2m}^{LM})$ forms a stable* A$\mathbb{H}$*-module (U, U') with $U \cong \mathbb{H}^m$ and $U' \cong \mathbb{R}^{2m+1}$.*

Proof. Consider the maximal stable submodule W of U. Since W is an $\mathbb{H}$-submodule it must be invariant under the left $\mathbb{H}$-action. Also, W must depend solely on the $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$-representation structure: in particular W' must be preserved by the diagonal action. So W' must be an $\mathrm{Sp}(1)^{LM}$-invariant subspace of $U' = V_{2m}$ and by Schur's Lemma [FH, p.7] $W' = V_{2m}$ or $W' = \{0\}$. Since $\dim U' > \frac{1}{2} \dim U$ we must have $W \neq \{0\}$. Hence $(W, W') = (U, U')$ and thus U is stable. ∎

Lemma 5.1.2 *The* $A\mathbb{H}$*-module* $U = V_1 \otimes V_{2m-1}$ *is the irreducible stable* $A\mathbb{H}$*-module corresponding to the SK-module* $\mathcal{X}_2^{2m}$*.*

Proof. This follows from Proposition 5.1.1 and the remarks in Section 4.4. The virtual dimension of U is 1, from which it follows that U must be irreducible. The isomorphism with $\mathcal{X}_2^{2m}$ follows because irreducible stable $A\mathbb{H}$-modules are uniquely determined by their dimensions. ∎

It is instructive to describe the splitting of $V_1 \otimes V_{2m-1}$ thoroughly in terms of basis vectors. To make statements less cumbersome, we let $n = 2m - 1$ throughout, bearing in mind that n is odd.

Let $V_1 = \langle \mathbf{x}, \mathbf{y} \rangle$ and $V_n = \langle \mathbf{a}^n, \mathbf{a}^{n-1}\mathbf{b}, \ldots, \mathbf{a}\mathbf{b}^{n-1}, \mathbf{b}^n \rangle$ be Sp(1)-representations. The actions of $\mathfrak{sl}(2, \mathbb{C})$ on V_1 and V_n are given by Equations (1.14) and (1.16) of Section (1.2.1). We want to understand the actions on the tensor product

$$V_1 \otimes V_n = \left\langle \begin{array}{c} \mathbf{x} \otimes \mathbf{a}^n, \mathbf{x} \otimes \mathbf{a}^{n-1}\mathbf{b}, \ldots, \mathbf{x} \otimes \mathbf{a}\mathbf{b}^{n-1}, \mathbf{x} \otimes \mathbf{b}^n \\ \mathbf{y} \otimes \mathbf{a}^n, \mathbf{y} \otimes \mathbf{a}^{n-1}\mathbf{b}, \ldots, \mathbf{y} \otimes \mathbf{a}\mathbf{b}^{n-1}, \mathbf{y} \otimes \mathbf{b}^n \end{array} \right\rangle.$$

In particular, we would like to find out how the left $\mathbb{H}$-action interacts with the splitting $V_1 \otimes V_n \cong V_{n+1} \oplus V_{n-1}$.

Consider the structure of $\mathcal{X}_2^{2m}$ as a K-module $e : W \to H \otimes V$, where $H \cong V_1$ and $V \cong V_n$. Using these isomorphisms and Theorem 4.4.2, the image $e(W)$ is spanned by the vectors

$$\{\mathbf{x}\otimes\mathbf{a}^n, \mathbf{x}\otimes\mathbf{a}^{n-1}\mathbf{b}+\mathbf{y}\otimes\mathbf{a}^n, \ldots, \mathbf{x}\otimes\mathbf{a}^{n-k}\mathbf{b}^k+\mathbf{y}\otimes\mathbf{a}^{n-k+1}\mathbf{b}^{k-1}, \ldots, \mathbf{x}\otimes\mathbf{b}^n+\mathbf{y}\otimes\mathbf{a}\mathbf{b}^{n-1}, \mathbf{y}\otimes\mathbf{b}^n\}.$$

It is easier to observe the $\mathfrak{sp}(1)$ action on the complementary subspace which we can identify as the $U^\dagger$-part of an $S\mathbb{H}$-module. Given a suitable choice of metric, the perpendicular subspace to $e(W)$ is spanned by vectors of the form $\mathbf{x}\otimes\mathbf{a}^{n-k}\mathbf{b}^k - \mathbf{y}\otimes\mathbf{a}^{n-k+1}\mathbf{b}^{k-1}$. Calculating the action of the Casimir operator $C = H^2 + 2XY + 2YX$ reveals that

$$C(\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k - \mathbf{y} \otimes \mathbf{a}^{n-k+1}\mathbf{b}^{k-1}) = (n+1)(n-1)(\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k - \mathbf{y} \otimes \mathbf{a}^{n-k+1}\mathbf{b}^{k-1}).$$

This shows that $\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k - \mathbf{y} \otimes \mathbf{a}^{n-k+1}\mathbf{b}^{k-1} \in V_{n-1}$, giving the result that

$$V_{n-1} = \mathrm{Span}\{\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k - \mathbf{y} \otimes \mathbf{a}^{n-k+1}\mathbf{b}^{k-1} : 1 \le k \le n\}. \tag{5.4}$$

The subspaces $e(W)$ and $e(W)^\perp$ of the SK-module $H \otimes V$ are thus equivalent to the subspaces V_{n+1} and V_{n-1} respectively in the splitting $V_1 \otimes V_n \cong V_{n+1} \oplus V_{n-1}$.

The SK-module structure maps σ_W and σ_V are exactly the standard real structure σ_{n-1} and the quaternionic structure σ_n introduced in Section 1.2.1. The $\mathbb{H}$-action is as usual determined by the action of $\mathfrak{sp}(1)$ on $\mathbf{x}$ and $\mathbf{y}$ using the correspondence

$$1 \longleftrightarrow \mathbf{x} \quad i_1 \longleftrightarrow i\mathbf{x} \quad i_2 \longleftrightarrow \mathbf{y} \quad i_3 \longleftrightarrow -i\mathbf{y}.$$

This formalism enables us to write out the structure of $V_1 \otimes V_n$ as an $\mathbb{H}$-module in the same fashion as in Example 4.4.5.

A$\mathbb{H}$-modules of the form $\mathbf{V_1 \otimes V_{2m}}$

We can also obtain a stable A$\mathbb{H}$-module from an odd-dimensional irreducible Sp(1)-representation $V_{2m} \cong \mathbb{C}^{2m+1}$. If we take the tensor product $V_1^L \otimes V_{2m}^M \cong \mathbb{C}^{4m+2}$ we obtain the splitting $V_1^L \otimes V_{2m}^M \cong V_{2m+1}^{LM} \oplus V_{2m-1}^{LM}$ and a left $\mathbb{H}$-action in the same way as above. However this does not restrict to an $\mathbb{H}$-action on any suitable *real* vector space U such that $U \otimes_{\mathbb{R}} \mathbb{C} = V_1 \otimes V_{2m}$ (this is obviously impossible since $\mathbb{R}^{4m+2} \ncong \mathbb{H}^k$ for any k). The reason for this is that the structure map $\sigma_1 \otimes \sigma_{2m}$ has square -1 instead of 1, and so $V_1^L \otimes V_{2m}^M$ is a quaternionic rather than a real representation of $\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$.

There are two ways round this difficulty. Firstly, we could simply take the underlying real vector space $\mathbb{R}^{8m+4} \cong V_1 \otimes V_{2m}$ to be an $\mathbb{H}$-module. Secondly, we can tensor with $H \cong \mathbb{C}^2$ equipped with its standard structure map. The vector space H is unaffected by the $\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$-action; thus we can think of $V_1^L \otimes V_{2m}^M \otimes H$ as a direct sum of two copies of $V_1^L \otimes V_{2m}^M$, which we write $2V_1^L \otimes V_{2m}^M$. This space comes equipped with a real structure $\sigma = \sigma_1 \otimes \sigma_{2m} \otimes \sigma_H$, and so we have a stable A$\mathbb{H}$-module

$$((2V_1^L \otimes V_{2m}^M)^\sigma, (2V_{2m+1}^{LM})^\sigma). \tag{5.5}$$

(As usual, once the correct structure maps have been specified, we will not usually mention the σ-superscript.) This approach is the equivalent of dealing with SK-modules of the form $\mathcal{X}_{k=2,3}^{2m+1} \otimes H$ and the σ-sheaves $\mathcal{O}(2m+1) \otimes H$.

Both these approaches give exactly the same A$\mathbb{H}$-module; both effectively leave the $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$-representation $V_1 \otimes V_{2m}$ untouched, whilst doubling the dimension of the real vector space we are considering so that it is divisible by four.

5.1.4 General Stable and Antistable A$\mathbb{H}$-modules

Definition 5.1.3 Let U_{2n} denote the A$\mathbb{H}$-module $(V_1^L \otimes V_{2n+1}^M, V_{2n+2}^{LM})$.

Let U_{2n-1} denote the A$\mathbb{H}$-module $(2V_1^L \otimes V_{2n}^M, 2V_{2n+1}^{LM})$.

The A$\mathbb{H}$-module U_{2n} corresponds to the SK-module $\mathcal{X}_2^{2n+2}$ and the σ-sheaf $\mathcal{O}(2n)$.

The A$\mathbb{H}$-module U_{2n-1} corresponds to the SK-module $\mathcal{X}_2^{2n+1} \otimes H$ and the σ-sheaf $\mathcal{O}(2n-1) \otimes H$.

By analogy with the quaternions themselves, we will refer to the action of the 'left' subgroup $\mathrm{Sp}(1)^L$ on V_1 as the *left $\mathbb{H}$-action*, and the action of the 'right' subgroup $\mathrm{Sp}(1)^M$ on V_n as the *$\mathrm{Sp}(1)^M$-action*. (This could also be taken to signify 'module' action.) We will often omit the superscripts L and M from expressions like $V_1^L \otimes V_n^M$ if the context leaves no ambiguity as to which group acts on what. Thus we write

$$U_n = aV_1 \otimes V_{n+1},$$

where $a = 1$ if n is even and $a = 2$ if n is odd.

This formulation allows us to see the relationship between stable A$\mathbb{H}$- and S$\mathbb{H}$-modules very explicitly. The A$\mathbb{H}$-module $U_n = aV_1 \otimes V_{n+1}$ splits as $a(V_{n+2} \oplus V_n)$. We can choose to regard this as a stable A$\mathbb{H}$-module by thinking of aV_{n+2} as the 'primed part', or as a stable S$\mathbb{H}$-module by regarding aV_n as the 'generating real subspace'.

The classification results of Sections 4.4 and 4.5 allow us to state the following theorem:

Theorem 5.1.4 *Every stable* $\mathrm{A}\mathbb{H}$*-module can be written as a direct sum of the irreducibles* U_n *with unique multiplicities.*

Consider the direct sum $U = \bigoplus_{j=0}^{n} a_j U_j$. We can write U more explicitly in terms of $\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^M$-representations, as the sum

$$U = V_1^L \otimes \left(\bigoplus_{j=1}^{m} V_{2j+1}^M \oplus 2 \bigoplus_{k=1}^{n} V_{2k}^M \right). \tag{5.6}$$

The left $\mathbb{H}$-action on V_1^L is common to all the irreducibles. The $\mathrm{Sp}(1)^M$-action can be much more complicated. However, we know that since Sp(1) is a compact group, *any* such representation can be written as a sum of irreducibles with unique multiplicities. Having done this, it is then easy to separate these representations to form separate $\mathrm{A}\mathbb{H}$-modules, provided that each odd-dimensional representation V_{2k}^M appears with even multiplicity. Thus the following is equivalent to Theorem 5.1.4:

Theorem 5.1.5 *To every stable* $\mathrm{A}\mathbb{H}$*-module* (U, U') *can be attached an* $\mathrm{Sp}(1)^M$*-action which intertwines with the left* $\mathbb{H}$*-action in such a way that the diagonal* $\mathrm{Sp}(1)^{LM}$*-action preserves* U'*.*

In the decomposition of Equation (5.6), each irreducible subrepresentation of the $\mathrm{Sp}(1)^M$-action contributes 1 to the virtual dimension of U. Thus the virtual dimension of $\bigoplus_{j=0}^{n} c_j U_j$ is equal to $\sum_{j \text{ even}} c_j + 2\sum_{j \text{ odd}} c_j$.

Antistable $\mathrm{A}\mathbb{H}$-modules

Let $U_n = aV_1 \otimes V_{n+1}$ be a stable $\mathrm{A}\mathbb{H}$-module. Then its dual $\mathrm{A}\mathbb{H}$-module $U_n^\times$ is an antistable $\mathrm{A}\mathbb{H}$-module. Just like stable $\mathrm{S}\mathbb{H}$-modules, antistable $\mathrm{A}\mathbb{H}$-modules are formed by taking the smaller summand in the splitting $aV_1 \otimes V_{n+1} = a(V_{n+2} \oplus V_n)$. The following Lemma then follows immediately from Definition 5.1.3.

Lemma 5.1.6 *The antistable* $\mathrm{A}\mathbb{H}$*-module* $U_{2n}^\times$ *takes the form* $(V_1^L \otimes V_{2n+1}^M, V_{2n}^{LM})$.

The antistable $\mathrm{A}\mathbb{H}$*-module* $U_{2n-1}^\times$ *takes the form* $(2V_1^L \otimes V_{2n}^M, 2V_{2n-1}^{LM})$.

The $\mathrm{A}\mathbb{H}$-module $U_{2n}^\times$ corresponds to the SK-module $\mathcal{X}_3^{2n+2}$ and the σ-sheaf $\mathcal{O}(-2n-4)$. The $\mathrm{A}\mathbb{H}$-module $U_{2n-1}^\times$ corresponds to the SK-module $\mathcal{X}_3^{2n+1} \otimes H$ and the σ-sheaf $\mathcal{O}(-2n-3) \otimes H$. There is a 'unique factorisation theorem' for antistable $\mathrm{A}\mathbb{H}$-modules which is exactly dual to Theorem 5.1.4.

5.1.5 Line Bundles over $\mathbb{C}\mathrm{P}^1$ and Sp(1)-representations

There is naturally a link between Sp(1)-representations and the cohomology groups of vector bundles over $\mathbb{C}P^1$. In Section 4.5.1 we demonstrated that $H^0(\mathcal{O}(n)) \cong S^n(H)$, where $H \cong \mathbb{C}^2$ is the basic representation of GL(2,$\mathbb{C}$). From the inclusion SL(2,$\mathbb{C}$) $\subset$ GL(2,$\mathbb{C}$), H is also the basic representation V_1 of SL(2,$\mathbb{C}$) and therefore Sp(1). The induced action of Sp(1) on $S^n(H)$ is by definition the irreducible representation V_n. Thus the cohomology groups of line bundles over $\mathbb{C}P^1$ are Sp(1)-representations; we have

$$H^0(\mathcal{O}(n)) \cong V_n \qquad \text{and} \qquad H^1(\mathcal{O}(-n)) \cong V_{n-2}. \tag{5.7}$$

The exact sequence (4.11) of Section 4.5.1 is thus the same as the exact sequence

$$0 \longrightarrow V_{n-1} \longrightarrow V_1 \otimes V_n \cong V_{n-1} \oplus V_{n+1} \longrightarrow V_{n+1} \longrightarrow 0. \tag{5.8}$$

The fact that the cohomology groups of line bundles over $\mathbb{C}P^1$ have the structure of irreducible Sp(1)-representations is already known in the context of the theory of homogeneous spaces. Let G be a compact Lie group and let T be a maximal toral subgroup. Then the homogeneous space G/T has a homogeneous complex structure. (This famous result is due to Borel.) The right action of T on G gives G the structure of a principal T-bundle over G/T. Let $\mathfrak{t}$ be the Lie algebra of T, so that $\mathfrak{t}$ is a Cartan subalgebra of $\mathfrak{g}$. For each dominant weight $\lambda \in \mathfrak{t}^*$ there is a one-dimensional representation $\mathbb{C}_\lambda$ of T. The holomorphic line bundle associated to the principal bundle G and the representation λ is then

$$\begin{aligned} L_\lambda &= G \times_T \mathbb{C}_\lambda \\ &= (G \times \mathbb{C}_\lambda)/\{(g,v) \sim (gt, t^{-1}v), t \in T\}. \end{aligned}$$

Since G acts on L_λ, the cohomology groups of L_λ are naturally representations of G. For more information see [FH, p. 382-393].

In the case of the group Sp(1), each maximal torus is isomorphic to U(1), and the homogeneous space $\mathrm{Sp}(1)/\mathrm{U}(1) \cong \mathbb{C}P^1$ is the Hopf fibration $S^1 \hookrightarrow S^3 \to S^2$. The line bundle $\mathrm{Sp}(1) \times_{\mathrm{U}(1)} \mathbb{C}_\lambda$ is then $L^{-\lambda}$, where L is the hyperplane section bundle of $\mathbb{C}P^1$.

5.2 Sp(1)-Representations and the Quaternionic Tensor Product

This section describes the quaternionic algebra of stable and antistable A$\mathbb{H}$-modules using the ideas of the previous section.

5.2.1 The inclusion map $\iota_U(U)$

We begin by discussing the map ι_U and its image. Let U_n be an irreducible stable A$\mathbb{H}$-module. Then

$$U_n = a(V_1^L \otimes V_{n+1}^M), \quad U_n' = aV_{n+2}^{LM} \quad \text{and} \quad U_n^\dagger \cong (U_n^\dagger)^* = aV_n^{LM}.$$

There is an injective map $\iota_{U_n} : U_n \to \mathbb{H} \otimes (U_n^\dagger)^*$. This map has a natural interpretation in terms of the Sp(1)-representations involved. Writing the quaternions as the stable A$\mathbb{H}$-module $V_1^L \otimes V_1^R$, we have

$$\mathbb{H} \otimes (U_n^\dagger)^* \cong V_1^L \otimes V_1^R \otimes aV_n^N.$$

This is exactly like the motivating example of $\mathbb{H} \otimes E_{k,r}$ in Section 3.4. Leaving the left-action untouched and taking the diagonal $\mathrm{Sp}(1)^{RN}$-action gives the isomorphism

$$\mathbb{H} \otimes (U_n^\dagger)^* \cong V_1^L \otimes a(V_{n+1}^{RN} \oplus V_{n-1}^{RN}) \tag{5.9}$$

as an $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$-representation. The A$\mathbb{H}$-submodule $\iota_{U_n}(U_n)$ is clearly the $V_1^L \otimes aV_{n+1}^{RN}$ subrepresentation of $\mathbb{H} \otimes (U_n^\dagger)^*$.

Antistable A$\mathbb{H}$-modules behave in a similar fashion. Consider the A$\mathbb{H}$-module $U_n^\times$, so that

$$U_n^\times = a(V_1^L \otimes V_{n+1}^M), \quad (U_n^\times)' = aV_n^{LM} \quad \text{and} \quad (U_n^\times)^\dagger \cong ((U_n^\times)^\dagger)^* = aV_{n+2}^{LM}.$$

There is a similar splitting

$$\mathbb{H} \otimes ((U_n^\times)^\dagger)^* \cong V_1^L \otimes V_1^R \otimes aV_{n+2}^N \cong V_1^L \otimes a(V_{n+3}^{RN} \oplus V_{n+1}^{RN}).$$

This time, the A$\mathbb{H}$-submodule $\iota_{U_n^\times}(U_n^\times)$ is the *smaller* A$\mathbb{H}$-submodule $V_1^L \otimes aV_{n+1}^{RN}$. Thus the splitting

$$\mathbb{H} \otimes aV_n^N \cong aV_1^L \otimes (V_{n+1}^{RN} \oplus V_{n-1}^{RN}) \tag{5.10}$$

splits $\mathbb{H} \otimes aV_n$ into the direct sum of a stable A$\mathbb{H}$-module isomorphic to U_n and an antistable A$\mathbb{H}$-module isomorphic to $U_{n-2}^\times$.

The subspaces $\iota_{U_n}(U_n')$ and $\iota_{U_n^\times}((U_n^\times)')$ have a similar interpretation. Treating the imaginary quaternions $\mathbb{I}$ as a copy of V_2^{LR}, we have

$$\mathbb{I} \otimes (U_n^\dagger)^* \cong V_2^{LR} \otimes aV_n^N \cong a(V_{n+2}^{LRN} \oplus V_n^{LRN} \oplus V_{n-2}^{LRN}).$$

and $\iota_{U_n}(U_n')$ is the aV_{n+2} subrepresentation of $\mathbb{I} \otimes (U_n^\dagger)^*$. In exactly the same way, for $U_n^\times$ we have

$$\mathbb{I} \otimes ((U_n^\times)^\dagger)^* \cong V_2^{LR} \otimes aV_{n+2}^N \cong a(V_{n+4}^{LRN} \oplus V_{n+2}^{LRN} \oplus V_n^{LRN}).$$

In this case, $\iota_{U_n^\times}((U_n^\times)')$ is the smallest subrepresentation aV_n^{LRN}. This also shows why we would not expect U' to be closed under the left $\mathbb{H}$-action — the group $\mathrm{Sp}(1)^L$ does not act upon it, since we do not have an intact copy of V_1^L.

It is worth noting that so far we have been able consistently to interpret stable A$\mathbb{H}$-modules and their subspaces as representations of highest weight in tensor products of $\mathrm{Sp}(1)$-representations, and antistable A$\mathbb{H}$-modules and their subspaces as representations of lowest weight.

5.2.2 Tensor products of stable A$\mathbb{H}$-modules

We shall now see how to use our description of stable A$\mathbb{H}$-modules to form the quaternionic tensor product. The results in this section can be obtained through Quillen's sheaf-theoretic version of A$\mathbb{H}$-modules by using Theorem 4.5.7. However, the author hopes that including a little more description will help the reader to get more of a feel for what is going on.

Let $U_m = aV_1 \otimes V_{m+1}$, $U_n = bV_1 \otimes V_{n+1}$ be stable A$\mathbb{H}$-modules. By Definition 4.1.4,

$$U_m \otimes_{\mathbb{H}} U_n = (\iota_{U_m}(U_m) \otimes (U_n^\dagger)^*) \cap ((U_m^\dagger)^* \otimes \iota_{U_n}(U_n)) \subset \mathbb{H} \otimes (U_m^\dagger)^* \otimes (U_n^\dagger)^*.$$

In terms of Sp(1)-representations,

$$\mathbb{H} \otimes (U_m^\dagger)^* \otimes (U_n^\dagger)^* \cong V_1^L \otimes V_1^R \otimes aV_m^P \otimes bV_n^Q. \tag{5.11}$$

Using Equation (5.9), we write $\iota_{U_m}(U_m) \cong aV_1^L \otimes V_{m+1}^{RP} \subset V_1^L \otimes a(V_{m+1}^{RP} \oplus V_{m-1}^{RP}) \cong \mathbb{H} \otimes (U_m^\dagger)^*$. Tensoring this expression with $(U_n^\dagger)^* \cong bV_n^Q$ gives

$$\iota_{U_m}(U_m) \otimes (U_n^\dagger)^* \cong a(V_1^L \otimes V_{m+1}^{RP}) \otimes bV_n^Q. \tag{5.12}$$

In the same way, we form the isomorphism

$$(U_m^\dagger)^* \otimes \iota_{U_n}(U_n) \cong aV_m^P \otimes b(V_1^L \otimes V_{n+1}^{RQ}). \tag{5.13}$$

A rearrangement of the factors leaves us considering the spaces $abV_1^L \otimes V_{m+1}^{RP} \otimes V_n^Q$ and $abV_1^L \otimes V_m^P \otimes V_{n+1}^{RQ}$. We now have an $\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^{RP} \times \mathrm{Sp}(1)^Q$-representation and an $\mathrm{Sp}(1)^L \times \mathrm{Sp}(1)^P \times \mathrm{Sp}(1)^{RQ}$-representation. From these we want to obtain a single $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$-representation which leaves the left $\mathbb{H}$-action intact. The way to proceed is to leave the V_1^L-factor in each of these expressions alone and consider the representations of the diagonal subgroup $\mathrm{Sp}(1)^{RPQ}$. We examine the factors $V_{m+1}^{RP} \otimes V_n^Q$ and $V_m^P \otimes V_{n+1}^{RQ}$. To obtain a stable A$\mathbb{H}$-module, we want to reduce these two $\mathrm{Sp}(1) \times \mathrm{Sp}(1)$-representations to a single Sp(1)-representation. In so doing, we hope to find the intersection of these two spaces.

This is summed up in the following diagram:

$$\begin{array}{ccccc}
 & & \mathbb{H} \otimes (U_m^\dagger)^* \otimes (U_n^\dagger)^* \cong V_1^L \otimes V_1^R \otimes aV_m^P \otimes bV_n^Q & & \\
U_m \otimes (U_n^\dagger)^* & \nearrow & & \nwarrow & (U_m^\dagger)^* \otimes U_n \\
\downarrow \cong & & & & \downarrow \cong \\
\iota_{U_m}(U_m) \otimes (U_n^\dagger)^* \cong V_1^L \otimes aV_{m+1}^{RP} \otimes bV_n^Q & & & & (U_m^\dagger)^* \otimes \iota_{U_n}(U_n) \cong aV_m^P \otimes V_1^L \otimes bV_{n-1}^{RQ} \\
 & \nwarrow & & \nearrow & \\
 & & U_m \otimes_{\mathbb{H}} U_n \cong abV_1^L \otimes (\bigoplus V_?^{RPQ}) & &
\end{array}$$

The upward arrows here are inclusion maps. The argument goes in the opposite direction, as we restrict our attention to particular subspaces. If we go down the left hand side, we consider the diagonal action of the subgroup $\mathrm{Sp}(1)^{RP}$, and restrict to the higher weight subspace $\iota_{U_m}(U_m) \otimes (U_n^\dagger)^*$. We then consider the action of $\mathrm{Sp}(1)^{RPQ}$ on this. If on the other hand we go down the *right* hand side, we consider the diagonal action of the subgroup $\mathrm{Sp}(1)^{RQ}$, restrict to the higher weight subspace $(U_m^\dagger)^* \otimes \iota_{U_n}(U_n)$, and then consider the action of $\mathrm{Sp}(1)^{RPQ}$ on *this.* At each 'half-way stage' we are considering representations of diagonal subgroups of different pairs of groups, and in both cases we take the higher weight representation in a sum $V_{k+1} \oplus V_{k-1}$ and discard the V_{k-1} part. Since we do this for different diagonal subgroups we expect to be left with different subspaces.

Using the Clebsch-Gordon formula, we obtain the two decompositions

$$V_{m+1}^{RP} \otimes V_n^Q \cong \bigoplus_{j=0}^{\min\{m+1,n\}} V_{m+1+n-2j}^{RPQ} \qquad \text{and} \qquad V_m^P \otimes V_{n+1}^{RQ} \cong \bigoplus_{j=0}^{\min\{m,n+1\}} V_{m+n+1-2j}^{PRQ}.$$

These decompositions contain fairly similar summands, with differences arising as the index j approaches the region of $\min\{m, n\}$. However, just because we have two $\mathrm{Sp}(1)^{RPQ}$-representations of the same weight, we cannot say that they are automatically the same subspace of $\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$. We want to know which parts end up contributing to the final $\mathrm{Sp}(1)^{RPQ}$-representation *whichever path we take.* This will identify the subspace $U_m \otimes_{\mathbb{H}} U_n \subseteq \mathbb{H} \otimes (U_m^\dagger)^* \otimes (U_n^\dagger)^*$.

One thing that we can guarantee for any $m, n > 0$ is that the representation with highest weight will be the same in both cases — both expressions have leading summand V_{m+n+1}. We conjecture that this is the summand which we find in $U_m \otimes_{\mathbb{H}} U_n$. This would fit well with the observation that stable A$\mathbb{H}$-modules arise as representations of highest weight in decompositions of tensor products of $\mathrm{Sp}(1)$-representations.

We will show that this is in fact the case, using Joyce's dimension formulae for stable A$\mathbb{H}$-modules. Here is the main result of this section:

Theorem 5.2.1 *Let U_m, U_n be irreducible stable* A$\mathbb{H}$*-modules. If m or n is even then*

$$U_m \otimes_{\mathbb{H}} U_n \cong U_{m+n}.$$

If m and n are both odd then

$$U_m \otimes_{\mathbb{H}} U_n \cong 4U_{m+n}.$$

Proof. We have already noted that each irreducible representation of the $\mathrm{Sp}(1)^M$-action on a stable A$\mathbb{H}$-module U contributes 1 to the virtual dimension of U. Thus any stable A$\mathbb{H}$-module of virtual dimension k must be a sum of at least $k/2$ and at most k irreducibles, depending on whether the irreducibles are odd or even.

We will deal with the three possible cases in turn.

Case 1(m and n both even): Let $m = 2p$, $n = 2q$. Then

$$\dim U_m = 4(p+1) \quad \dim U_m' = 2p+3 \quad \dim U_n = 4(q+1) \quad \text{and} \quad \dim U_n' = 2q+3.$$

Using Theorem 4.1.14 we find that $\dim U_m \otimes_{\mathbb{H}} U_n = 4(p+q+1)$ and that the virtual dimension of $U_m \otimes_{\mathbb{H}} U_n$ is equal to 1. But any stable A$\mathbb{H}$-module whose virtual dimension is equal to 1 must be irreducible. The irreducible stable A$\mathbb{H}$-module whose dimension is $4(p+q+1)$ and whose virtual dimension is 1 is $V_1 \otimes V_{2(p+q)+1} = U_{m+n}$. Hence $U_m \otimes_{\mathbb{H}} U_n = U_{m+n}$.

Case 2(m even and n odd): Let $m = 2p$, $n = 2q-1$. Then

$$\dim U_m = 4(p+1) \quad \dim U'_m = 2p+3 \quad \dim U_n = 4(2q+1) \quad \text{and} \quad \dim U'_n = 4(q+1).$$

Using Theorem 4.1.14 we find that $\dim U_m \otimes_{\mathbb{H}} U_n = 4(2p+2q+1)$ and that the virtual dimension of $U_m \otimes_{\mathbb{H}} U_n$ is equal to 2. Thus $U_m \otimes_{\mathbb{H}} U_n$ must be either an even irreducible or a sum of two odd irreducibles.

Consider the space $\iota_{U_m}(U_m) \otimes (U_n^\dagger)^* \cong a(V_1^L \otimes V_{m+1}^{RP}) \otimes bV_n^Q$ of Equation (5.12). For $m = 2p$ and $n = 2q-1$ this becomes

$$2V_1^L \otimes V_{2p+1}^{RP} \otimes V_{2q-1}^Q \cong V_1^L \otimes 2(V_{2p+2q}^{RPQ} \oplus V_{2p+2q-2}^{RPQ} \oplus \ldots) \tag{5.14}$$

The virtual dimension of the tensor product $U_m \otimes_{\mathbb{H}} U_n$ must be equal to 2, so we cannot have more than 2 of the irreducibles of the $\mathrm{Sp}(1)^{RPQ}$-action. We also need a total dimension of $4(2p+2q+1)$. Examining Equation (5.14) we see that the only way this can occur is if $U_m \otimes_{\mathbb{H}} U_n \cong V_1 \otimes 2V_{2p+2q}$, as all the other irreducibles of the $\mathrm{Sp}(1)^{RPQ}$-action have smaller dimension. Hence $U_m \otimes_{\mathbb{H}} U_n \cong 2V_1 \otimes V_{2p+2q} = U_{m+n}$.

Case 3 (m and n both odd): The argument is very similar to that of Case 2.

Let $m = 2p-1$, $n = 2q-1$. Then

$$\dim U_m = 4(2p+1) \quad \dim U'_m = 4(p+1) \quad \dim U_n = 4(2q+1) \quad \text{and} \quad \dim U'_n = 4(q+1).$$

Using Theorem 4.1.14 we find that $\dim U_m \otimes_{\mathbb{H}} U_n = 16(p+q)$ and that the virtual dimension of $U_m \otimes_{\mathbb{H}} U_n$ is equal to 4.

Consider the space $\iota_{U_m}(U_m) \otimes (U_n^\dagger)^* \cong a(V_1^L \otimes V_{m+1}^{RP}) \otimes bV_n^Q$ of Equation (5.12). For $m = 2p-1$ and $n = 2q-1$ this becomes

$$4V_1^L \otimes V_{2p}^{RP} \otimes V_{2q-1}^Q \cong V_1^L \otimes 4(V_{2p+2q-1}^{RPQ} \oplus V_{2p+2q-3}^{RPQ} + \ldots). \tag{5.15}$$

The only way $U_m \otimes_{\mathbb{H}} U_n$ can have a virtual dimension of four and a total dimension of $16(p+q)$ is if $U_m \otimes_{\mathbb{H}} U_n \cong 4V_1 \otimes V_{2p+2q-1} = 4U_{m+n}$. ∎

Quaternionic tensor products of more general stable A$\mathbb{H}$-modules can be computed from this result by splitting into irreducibles and using the fact that the quaternionic tensor product is distributive for direct sums.

This result is parallel to Theorem 4.5.7 applied to non-negative vector bundles. For the canonical sheaves $\mathcal{O}(n)$ over $\mathbb{C}P^1$, $H^0(\mathcal{O}(n)) \cong V_n$. The isomorphism $\mathcal{O}(n) \otimes_{\mathcal{O}} \mathcal{O}(m) \cong \mathcal{O}(n+m)$ induces a map of cohomology groups $H^0(\mathcal{O}(m)) \otimes H^0(\mathcal{O}(n)) \to H^0(\mathcal{O}(m+n))$. In terms of Sp(1)-representations, this is a map

$$V_m^P \otimes V_n^Q \cong V_{m+n}^{PQ} \oplus V_{m+n-2}^{PQ} \oplus \ldots \to V_{m+n}.$$

The map in question is projection onto the irreducible of highest weight V_{n+m}. This is really what this whole section has been about — the idea that the behaviour of stable A$\mathbb{H}$-modules can be thoroughly and flexibly described by taking subrepresentations of highest weight in tensor products of Sp(1)-representations.

5.2.3 Tensor Products of Antistable A$\mathbb{H}$-modules

It is not difficult to extend Joyce's results for tensor products of stable A$\mathbb{H}$-modules to irreducible antistable A$\mathbb{H}$-modules — we can follow the same argument as in the proof of Theorem 4.1.14, since the generic properties of sums and intersections guaranteed by stability also hold if one or both of the A$\mathbb{H}$-modules is irreducible and antistable.

Let U and V be antistable A$\mathbb{H}$-modules with $\dim U = 4j$, $\dim U' = 2j-r$, $\dim V = 4k$ and $\dim V' = 2k - s$. Let $A = \iota_U(U) \otimes (V^\dagger)^*$ and let $B = (U^\dagger)^* \otimes \iota_V(V)$. Then $\dim A = 4j(2k+s)$ and $\dim B = 4k(2j+r)$, so $\dim(\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*) = 4(2j+r)(2k+s) > \dim A + \dim B$. Thus in generic situations we would expect $\dim A \cap B = \dim U \otimes_{\mathbb{H}} V = 0$.

Suppose instead that V is stable, so now $\dim V' = 2k+s$. A similar calculation shows that $\dim A + \dim B \geq \dim \mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$ if and only if $s(j+r) \geq kr$. In this case we might expect $\dim U \otimes_{\mathbb{H}} V = \dim A + \dim B - \dim(\mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*)$. For example, let $U = U_m^\times = (aV_1 \otimes V_{m+1})^\times$ and $V = U_n = bV_1 \otimes V_{n+1}$. Then we would expect that $\dim(U_m^\times \otimes_{\mathbb{H}} U_n) = 2ab(m-n+2)$.

As in Section 5.2, we can describe what is going on in terms of diagonal actions on tensor products of Sp(1)-representations. We will illustrate the case $U_m \otimes_{\mathbb{H}} U_n^\times$. Let $U_m = aV_1 \otimes V_{m+1}$ and $U_n^\times = (bV_1 \otimes V_{n+1})^\times$, so $(U_m^\dagger)^* \cong V_m$ and $((U_n^\times)^\dagger)^* \cong V_{n+2}$. This gives rise to the standard descriptions

$$\mathbb{H} \otimes_{\mathbb{R}} (U_m^\dagger)^* \cong aV_1^L \otimes V_1^R \otimes V_m^P \cong aV_1^L \otimes (V_{m+1}^{RP} \oplus V_{m-1}^{RP})$$

and

$$\mathbb{H} \otimes_{\mathbb{R}} ((U_n^\times)^\dagger)^* \cong bV_1^L \otimes V_1^R \otimes V_{n+2}^Q \cong bV_1^L \otimes (V_{n+3}^{RQ} \oplus V_{n+1}^{RQ}).$$

The only difference between this and equation (5.9) is that we have an antistable A$\mathbb{H}$-module involved, and thus to obtain $\iota_{U_m^\times}(U_m^\times)$, we take the *smaller* summand of $V_{n+3} \oplus V_{n+1}$. We use the Clebsch-Gordon formula to describe

$$\iota_{U_m}(U_m) \otimes_{\mathbb{R}} ((U_n^\times)^\dagger)^* \cong abV_1^L \otimes V_{m+1}^{RP} \otimes V_{n+2}^Q \cong abV_1^L \otimes \left(\bigoplus_{j=0}^{\min\{m+1,n+2\}} V_{m+n+3-2j}^{RPQ} \right)$$

and

$$(U_m^\dagger)^* \otimes_{\mathbb{R}} \iota_{U_n^\times}(U_n^\times) \cong abV_1^L \otimes V_m^P \otimes V_{n+1}^{RQ} \cong abV_1^L \otimes \left(\bigoplus_{j=0}^{\min\{m,n+1\}} V_{m+n+1-2j}^{RPQ} \right).$$

As with stable A$\mathbb{H}$-modules, our task is to find which of these summands is in the intersection $U_m \otimes_{\mathbb{H}} U_n^\times = \iota_{U_m}(U_m) \otimes_{\mathbb{R}} ((U_n^\times)^\dagger)^* \cap (U_m^\dagger)^* \otimes_{\mathbb{R}} \iota_{U_n^\times}(U_n^\times)$. This time since $\min\{m+1, n+2\} = \min\{m, n+1\}$, we can guarantee that the summand of *smallest* weight will appear in both expressions. From our dimensional arguments, we only expect a non-zero intersection if $m < n+2$, in which case $\min\{m, n+1\} = m$, and we would predict that $U_m \otimes_{\mathbb{H}} U_n^\times$ contains the summand V_{n-m+1}. Because its virtual dimension is not positive, $U_m \otimes_{\mathbb{H}} U_n^\times$ cannot be stable. This suggests that

$$U_m \otimes_{\mathbb{H}} U_n^\times \cong (abV_1 \otimes V_{n-m+1})^\times = \begin{cases} \{0\} & \text{if } m \geq n+2 \\ U_{n-m}^\times & n \text{ or } m \text{ even and } m < n+2 \\ 4U_{n-m}^\times & n \text{ and } m \text{ both odd and } m < n+2. \end{cases} \tag{5.16}$$

Rather than try to emulate Joyce's (difficult) proof of Theorem 4.1.14, we will confirm these conjectures by appealing to Quillen's powerful results.

Proposition 5.2.2 *Let U_m be an irreducible stable* $\mathrm{A}\mathbb{H}$*-module and let $U_n^\times$ be an irreducible antistable* $\mathrm{A}\mathbb{H}$*-module.*

If $m \geq n+2$ then $U_m \otimes_{\mathbb{H}} U_n^\times = \{0\}$.

If $m < n+2$ and m or n is even then $U_m \otimes_{\mathbb{H}} U_n^\times \cong U_{n-m}^\times$.

If $m < n+2$ and m and n are both odd then $U_m \otimes_{\mathbb{H}} U_n^\times \cong 4U_{n-m}^\times$.

Let $U_m^\times$ and $U_n^\times$ be antistable irreducible $\mathrm{A}\mathbb{H}$*-modules. Then $U_m^\times \otimes_{\mathbb{H}} U_n^\times = \{0\}$.*

Proof. This follows from Theorem 4.5.7 (due to Quillen), using the correspondences $U_m = \eta^+(a\mathcal{O}(m))$ and $U_n^\times = \eta^-(a\mathcal{O}(-n-4))$, where as usual $a = 1$ or 2 depending on whether m, n are even or odd. ∎

We can use this result about tensor products of antistable $\mathrm{A}\mathbb{H}$-modules to tell us about $\mathrm{A}\mathbb{H}$-morphisms between *stable* $\mathrm{A}\mathbb{H}$-modules, using the isomorphism

$$\mathrm{Hom}_{\mathrm{A}\mathbb{H}}(U_m, U_n) \cong (U_m^\times \otimes_{\mathbb{H}} U_n)'$$

of Theorem 4.2.9.

Proposition 5.2.3 *Let U_m and U_n be stable* $\mathrm{A}\mathbb{H}$*-modules. Then*

$$\mathrm{Hom}_{\mathrm{A}\mathbb{H}}(U_m, U_n) \cong \begin{cases} (aU_{m-n}^\times)' = aV_{m-n} & n \leq m \\ \{0\} & n > m \end{cases}$$

where $a = 4$ if m and n are both odd and $a = 1$ otherwise.

Proof. This follows immediately by combining Theorems 4.2.9 and 5.2.2. ∎

In particular, since $U_0 = \mathbb{H}$, we see that there are always $\mathrm{A}\mathbb{H}$-morphisms from U_n into $\mathbb{H}$ (and indeed, this is a defining property for $\mathrm{A}\mathbb{H}$-modules), but *never* $\mathrm{A}\mathbb{H}$-morphisms from $\mathbb{H}$ into U_n unless $n = 0$.

Similarly, we can now see that there are always $\mathrm{A}\mathbb{H}$-morphisms from antistable $\mathrm{A}\mathbb{H}$-modules into stable $\mathrm{A}\mathbb{H}$-modules, but never $\mathrm{A}\mathbb{H}$-morphisms from stable $\mathrm{A}\mathbb{H}$-modules into antistable $\mathrm{A}\mathbb{H}$-modules.

5.3 Semistable $\mathrm{A}\mathbb{H}$-modules and Sp(1)-representations

In this section we shall consider how the $\mathrm{A}\mathbb{H}$-module X_q fits into the picture of stable $\mathrm{A}\mathbb{H}$-modules and Sp(1)-representations. Recall from Section 4.1.3 that for $q \in S^2$, $X_q = \mathbb{H}$ and $X_q' = \{p \in \mathbb{H} : pq = -qp\}$ so that $X_q^\dagger \cong (X_q^\dagger)^* \cong \mathbb{C}_q$.

If we consider the dual $\mathrm{A}\mathbb{H}$-module $X_q^\times \cong (\mathbb{H}, \mathbb{C}_q)$ we see that the left-multiplication $L_q : \mathbb{H} \to \mathbb{H}$ defined by $L_q(p) = q \cdot p$ gives an $\mathrm{A}\mathbb{H}$-isomorphism $X_q \cong X_q^\times$. This suggests that Theorem 4.2.9 might be particularly interesting in the case of X_q. For any $\mathrm{A}\mathbb{H}$-module U, there is a canonical isomorphism

$$\mathrm{Hom}_{\mathrm{A}\mathbb{H}}(X_q, U) \cong (X_q^\times \otimes_{\mathbb{H}} U)'$$

and so

$$\mathrm{Hom}_{\mathrm{A}\mathbb{H}}(X_q^\times, U) \cong \mathrm{Hom}_{\mathrm{A}\mathbb{H}}(X_q, U) \cong (X_q \otimes_{\mathbb{H}} U)' \cong (X_q^\times \otimes_{\mathbb{H}} U)'.$$

Let $\phi : X_q^\times \to U$ be an A$\mathbb{H}$-morphism. Then we need $\phi(1) = u \in U'$ and $\phi(q) \in U'$. Since ϕ is $\mathbb{H}$-linear, $\phi(q) = qu$, and so $u \in U' \cap qU'$. It follows that

$$\mathrm{Hom}_{\mathrm{A}\mathbb{H}}(X_q, U) \cong (X_q \otimes_{\mathbb{H}} U)' \cong U' \cap qU'. \tag{5.17}$$

As noted by Joyce (see the summary in Section 4.1.3), the second of these isomorphisms is given by the map $(\mathrm{id}_U \otimes_{\mathbb{H}} \chi_q) : (U \otimes_{\mathbb{H}} X_q)' \to (U \otimes_{\mathbb{H}} \mathbb{H})' \cong U'$.

Though X_q' is not itself an Sp(1)-representation, the subspaces $\mathbb{C}_q$ and $X_q' = \mathbb{C}_q^\perp$ are acted on by the Cartan subgroup $\mathrm{U}(1)_q \subset \mathrm{Sp}(1)$. As we shall see, taking the tensor product of a stable A$\mathbb{H}$-module with the A$\mathbb{H}$-module X_q serves to restrict attention from information about Sp(1)-representations to information concerning representations of the group $\mathrm{U}(1)_q$ and its Lie algebra $\mathfrak{u}(1)_q$.

Let $Q = aI_1 + bI_2 + cI_3 \in \mathfrak{sp}(1)$ with $a^2 + b^2 + c^2 = 1$ and let $q = ai_1 + bi_2 + ci_3 \in S^2$. Any irreducible representation of $\mathfrak{sp}(1)$ is also a representation of the subalgebra $\mathfrak{u}(1)_q = \langle Q \rangle$. The analysis of V_n as a $\mathrm{U}(1)_q$-representation is already familiar: it is the decomposition of V_n into weight spaces of the operator $Q : V_n \to V_n$. If we choose an identification $\mathfrak{sp}(1) \otimes_{\mathbb{R}} \mathbb{C} = \mathfrak{sl}(2, \mathbb{C})$ such that $Q = iH$, this is exactly the same as the decomposition of V_n into eigenspaces of H with weights $\{-n, -n+2, \ldots, n-2, n\}$. This decomposition gives important information about the action of q on the A$\mathbb{H}$-module $U_n = aV_1 \otimes V_{n+1}$.

This is exactly what we have done in the explicit calculations of Section 5.1.3 for the case $q = i_1$. We define a basis $\langle \mathbf{x}, \mathbf{y} \rangle$ for the space V_1 in such a way that $Q(\mathbf{x}) = i\mathbf{x}$ and $Q(\mathbf{y}) = -i\mathbf{y}$. This also gives the left action of $q \in \mathrm{Sp}(1)$ on V_1, since the actions of q and Q coincide for this representation (see Section 1.2.1). This gives the left action of $q \in \mathbb{H}$ on $U_n = aV_1 \otimes V_{n+1}$.

The goal of this discussion is to describe the A$\mathbb{H}$-module $U_n \otimes_{\mathbb{H}} X_q$. We do this with the aid of the following lemma. For ease of notation we work with the A$\mathbb{H}$-module U_{n-1}.

Lemma 5.3.1 *The subspace $U_{n-1}' \cap qU_{n-1}'$ is given by the sum of the weight spaces of Q with highest and lowest possible weights.*

Proof. Recall from Section 5.1.3 that

$$V_1 \otimes V_n = \left\langle \begin{array}{c} \mathbf{x} \otimes \mathbf{a}^n, \mathbf{x} \otimes \mathbf{a}^{n-1}\mathbf{b}, \ldots, \mathbf{x} \otimes \mathbf{a}\mathbf{b}^{n-1}, \mathbf{x} \otimes \mathbf{b}^n \\ \mathbf{y} \otimes \mathbf{a}^n, \mathbf{y} \otimes \mathbf{a}^{n-1}\mathbf{b}, \ldots, \mathbf{y} \otimes \mathbf{a}\mathbf{b}^{n-1}, \mathbf{y} \otimes \mathbf{b}^n \end{array} \right\rangle,$$

and that the space $(U_{n-1})' \cong V_{n+1}$ is spanned by the vectors

$$\{\mathbf{x}\otimes\mathbf{a}^n, \mathbf{x}\otimes\mathbf{a}^{n-1}\mathbf{b}+\mathbf{y}\otimes\mathbf{a}^n, \ldots, \mathbf{x}\otimes\mathbf{a}^{n-k}\mathbf{b}^k+\mathbf{y}\otimes\mathbf{a}^{n-k+1}\mathbf{b}^{k-1}, \ldots, \mathbf{x}\otimes\mathbf{b}^n+\mathbf{y}\otimes\mathbf{a}\mathbf{b}^{n-1}, \mathbf{y}\otimes\mathbf{b}^n\}.$$

Define the basis vectors $w_k \equiv \mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k + \mathbf{y} \otimes \mathbf{a}^{n-k+1}\mathbf{b}^{k-1}$, including $w_0 = \mathbf{x} \otimes \mathbf{a}^n$ and $w_{n+1} = \mathbf{y} \otimes \mathbf{b}^n$.

The action of $H \in \mathfrak{sl}(2, \mathbb{C})$ is given by $H(\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k) = (n - 2k + 1)\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k$ and $H(\mathbf{y} \otimes \mathbf{a}^{n-k}\mathbf{b}^k) = (n - 2k - 1)\mathbf{y} \otimes \mathbf{a}^{n-k}\mathbf{b}^k$. Thus each basis vector w_k is a weight vector with weight $(n - 2k + 1)$. The two extreme vectors $\mathbf{x} \otimes \mathbf{a}^n$ and $\mathbf{y} \otimes \mathbf{b}^n$ are also weight vectors with weights $n+1$ and $-n-1$ respectively. Note that all these weights are different.

The left $\mathbb{H}$-action of q on $V_1 \otimes V_n$ is given by

$$q(\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k) = i\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k \qquad q(\mathbf{y} \otimes \mathbf{a}^{n-k}\mathbf{b}^k) = -i\mathbf{y} \otimes \mathbf{a}^{n-k}\mathbf{b}^k.$$

Left multiplication by q therefore preserves the weight space decomposition with respect to Q of a vector $w \in V_1 \otimes V_n$.

On the basis vectors $w_k = \mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k + \mathbf{y} \otimes \mathbf{a}^{n-k+1}\mathbf{b}^{k-1}$ we have

$$q(\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k + \mathbf{y} \otimes \mathbf{a}^{n-k+1}\mathbf{b}^{k-1}) = i\mathbf{x} \otimes \mathbf{a}^{n-k}\mathbf{b}^k - i\mathbf{y} \otimes \mathbf{a}^{n-k+1}\mathbf{b}^{k-1}.$$

Since each vector w_k has a different weight from all the others we have $q(w_k) \in V_{n+1}$ if and only if $q(w_k) \in \langle w_k \rangle$, and

$$q(\sum \lambda_k w_k) \in V_{n+1} \iff \lambda_k = 0 \text{ or } q(w_k) \in V_{n+1} \qquad \text{for all } k.$$

It is evident that

$$q(\mathbf{x} \otimes \mathbf{a}^n) = i\mathbf{x} \otimes \mathbf{a}^n \in V_{n+1} \quad \text{and} \quad q(\mathbf{y} \otimes \mathbf{a}^n) = -i\mathbf{y} \otimes \mathbf{a}^n \in V_{n+1},$$

but for all the other basis vectors w_k,

$$q(w_k) \notin \langle w_k \rangle \quad \text{and so} \quad q(w_k) \notin V_{n+1}.$$

Hence

$$V_{n+1} \cap qV_{n+1} = \langle \mathbf{x} \otimes \mathbf{a}^n, \mathbf{y} \otimes \mathbf{b}^n \rangle,$$

the weight spaces with highest and lowest weight. Taking the σ-invariant subspace $\langle \mathbf{x} \otimes \mathbf{a}^n - \mathbf{y} \otimes \mathbf{b}^n \rangle$ if U_{n-1} is an even A$\mathbb{H}$-module yields the desired result for the A$\mathbb{H}$-module U_{n-1}. ∎

Since the A$\mathbb{H}$-submodule $(\mathrm{id} \otimes_{\mathbb{H}} \chi_q)(U_{n-1} \otimes_{\mathbb{H}} X_q) \subset U_{n-1}$ is the subspace generated over $\mathbb{H}$ by $U' \cap qU'$, this demonstrates the main result of this section which describes $U_n \otimes_{\mathbb{H}} X_q$ as follows:

Theorem 5.3.2 *Let $U_n = aV_1 \otimes V_{n+1}$ be an irreducible stable* A$\mathbb{H}$*-module, whose eigenspace decomposition with respect to $Q \in \mathfrak{sp}(1)$ takes the form*

$$V_1 \otimes V_{n+1} = \left\langle \begin{array}{c} \mathbf{x} \otimes \mathbf{a}^{n+1}, \mathbf{x} \otimes \mathbf{a}^n\mathbf{b}, \ldots, \mathbf{x} \otimes \mathbf{a}\mathbf{b}^n, \mathbf{x} \otimes \mathbf{b}^{n+1} \\ \mathbf{y} \otimes \mathbf{a}^{n+1}, \mathbf{y} \otimes \mathbf{a}^n\mathbf{b}, \ldots, \mathbf{y} \otimes \mathbf{a}\mathbf{b}^n, \mathbf{y} \otimes \mathbf{b}^{n+1} \end{array} \right\rangle.$$

If U_n is an even A$\mathbb{H}$*-module then $U_n \otimes_{\mathbb{H}} X_q \cong X_q$ and*

$$(\mathrm{id} \otimes_{\mathbb{H}} \chi_q)(U_n \otimes_{\mathbb{H}} X_q) = \mathbb{H} \cdot \langle \mathbf{x} \otimes \mathbf{a}^{n+1} - \mathbf{y} \otimes \mathbf{b}^{n+1} \rangle_{\mathbb{R}},$$

where the $\mathbb{H}$-action is induced by the $\mathbb{H}$-action on V_1.

If U_n is an odd A$\mathbb{H}$*-module then $U_n \otimes_{\mathbb{H}} X_q \cong 2X_q$ and*

$$(\mathrm{id} \otimes_{\mathbb{H}} \chi_q)(U_n \otimes_{\mathbb{H}} X_q) = \mathbb{H} \cdot \langle \mathbf{x} \otimes \mathbf{a}^{n+1}, \mathbf{y} \otimes \mathbf{b}^{n+1} \rangle_{\mathbb{R}} = V_1 \otimes \langle \mathbf{a}^{n+1}, \mathbf{b}^{n+1} \rangle.$$

Thus the quaternionic tensor product $U_n \otimes_{\mathbb{H}} X_q$ picks out the representations of extreme weight in the decomposition of U_n into weight spaces of Q.

5.4 Examples and Summary of $A\mathbb{H}$-modules

By now, the reader should be familiar with the ideas of quaternionic algebra. In this section we shall briefly sum up this information, giving explicit constructions of the $A\mathbb{H}$-modules which will occur most frequently in the following chapter, in the forms in which they occur most naturally.

Example 5.4.1 Recall the $A\mathbb{H}$-module $Y = \{(q_1, q_2, q_3) : q_1 i_1 + q_2 i_2 + q_3 i_3 = 0\}$ of Example 4.1.2. Since Y is stable and has virtual dimension 1, it follows that Y is irreducible and is isomorphic to $U_2 = V_1 \otimes V_3$. A calculation shows that $(Y^\dagger)^* \cong V_2$ and that the equation $q_1 i_1 + q_2 i_2 + q_3 i_3 = 0$ is precisely the condition for (q_1, q_2, q_3) to lie in the subspace $V_1^L \otimes V_3^{RM}$ of $\mathbb{H} \otimes V_2^M \cong V_1^L \otimes V_1^R \otimes V_2^M \cong V_1^L \otimes (V_3^{RM} \oplus V_1^{RM})$.

Consider the $A\mathbb{H}$-modules $\bigotimes_{\mathbb{H}}^k Y$, $S_{\mathbb{H}}^k Y$ and $\Lambda_{\mathbb{H}}^k Y$. Using the dimension formulae of Theorem 4.1.14 and Proposition 4.1.16, we discover that $\Lambda_{\mathbb{H}}^n Y = \{0\}$ and that $\bigotimes_{\mathbb{H}}^n Y = S_{\mathbb{H}}^n Y$ with $\dim(S_{\mathbb{H}}^n Y) = 4(n+1)$ and $\dim(S_{\mathbb{H}}^n Y)' = 2n+3$. From this we deduce that $S_{\mathbb{H}}^n Y \cong U_{2n}$, and that all the even irreducible stable $A\mathbb{H}$-modules can be realised as tensor powers of the $A\mathbb{H}$-module Y.

Example 5.4.2 [J1, Example 10.1] Let $Z \subset \mathbb{H} \otimes \mathbb{R}^4$ be the set

$$Z = \{(q_0, q_1, q_2, q_3) : q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3 = 0\}.$$

Then $Z \cong \mathbb{H}^3$ is a left $\mathbb{H}$-module. Define a real subspace $Z' = \{(q_0, q_1, q_2, q_3) : q_j \in \mathbb{I} \text{ and } q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3 = 0\}$. Then $\dim Z = 12$, $\dim Z' = 8$ and Z is a stable $A\mathbb{H}$-module.

In fact, Z is isomorphic to the first odd irreducible $A\mathbb{H}$-module $U_1 \cong 2V_1 \otimes V_2$. We have $(Z^\dagger)^* \cong \mathbb{R}^4 \cong 2V_1$. The equation $q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3 = 0$ is the condition for (q_0, q_1, q_2, q_3) to lie in the subspace $2V_1^L \otimes V_2^{RM}$ of $\mathbb{H} \otimes 2V_1^M \cong V_1^L \otimes V_1^R \otimes 2V_1^M \cong 2V_1^L \otimes (V_2^{RM} \oplus V_0^{RM})$.

Example 5.4.3 We can obtain the rest of the odd $A\mathbb{H}$-modules as tensor products of those above. From Theorem 5.2.1 we know that

$$U_{2n+1} = U_{2n} \otimes_{\mathbb{H}} U_1.$$

This combined with the previous examples shows that

$$U_{2n+1} \cong S_{\mathbb{H}}^n Y \otimes_{\mathbb{H}} Z.$$

Thus we have obtained all the irreducible stable $A\mathbb{H}$-modules in terms of previously known $A\mathbb{H}$-modules. We can also use these constructions to write down formulae giving all those irreducible antistable $A\mathbb{H}$-modules which are dual to stable $A\mathbb{H}$-modules.

Example 5.4.4 Consider the $A\mathbb{H}$-module (U, U') where

$$U = \mathbb{H}^2 \quad \text{and} \quad U' = \langle (1,0), (i_1, i_2), (0,1), (0, i_1) \rangle.$$

Then $\{(0, q) : q \in \mathbb{H}\}$ is an $A\mathbb{H}$-submodule of U isomorphic to X_{i_1}. However, there is no complementary $A\mathbb{H}$-submodule V such that $U \cong X_{i_1} \oplus V$, and indeed U is irreducible.

It is also clear that U is *not* semistable — nor is it the dual $A\mathbb{H}$-module of any semistable $A\mathbb{H}$-module. We call such an $A\mathbb{H}$-module *irregular*.

Irregular $A\mathbb{H}$-modules correspond to type (i) SK-modules of the form $\mathcal{X}_1^{n,\alpha} \oplus \mathcal{X}_1^{n,-\bar{\alpha}^{-1}}$ and torsion sheaves of the form $\mathcal{O}/m_x^n$, where $n \geq 2$. In the present case which singles out the subfield $\mathbb{C}_{i_1}$, we have

$$U \cong \eta^+(\mathcal{O}/m^2_{\{0,\infty\}}).$$

For a general formula linking pairs of antipodal points $\{z, \sigma(z)\}$ in $\mathbb{C}P^1$ and complex subfields of $\mathbb{H}$ see [Q, §14].

These are the only $A\mathbb{H}$-modules which we do not describe in detail, for two reasons. Firstly, they are badly behaved compared to semistable and antistable $A\mathbb{H}$-modules. Secondly, as far as the author is aware, they do not arise naturally in geometrical situations in the way that the other $A\mathbb{H}$-modules do.

Example 5.4.5 There is one remaining irreducible to consider — the $A\mathbb{H}$-module $(\mathbb{H}, \{0\})$. This trivially satisfies Definition 4.1.1, and so is an $A\mathbb{H}$-module. Any $A\mathbb{H}$-module whose primed part is zero is a direct sum of copies of $(\mathbb{H}, \{0\})$. It is easy to see that for any irreducible $A\mathbb{H}$-module U, $U \otimes_{\mathbb{H}} (\mathbb{H}, \{0\}) = \{0\}$ unless $U = \mathbb{H}$, in which case we have $\mathbb{H} \otimes_{\mathbb{H}} (\mathbb{H}, \{0\}) = (\mathbb{H}, \{0\})$.

Though badly behaved, the $A\mathbb{H}$-module $(\mathbb{H}, \{0\})$ *does* arise naturally in quaternionic algebra — for example, $Z \otimes_{\mathbb{H}} \mathbb{H}^\times = (\mathbb{H}, \{0\})$. This suggests the notation

$$(\mathbb{H}, \{0\}) = U^\times_{-1},$$

in which case this result agrees with Theorem 5.2.2. Such notation is consistent with the sheaf description of antistable $A\mathbb{H}$-modules, since we have

$$U^\times_{-1} = \eta^-(2\mathcal{O}(-3))$$

as expected. Thus we interpret $(\mathbb{H}, \{0\})$ as the 'antistable part' of $2V_1 \otimes V_0$. The dual space $(\mathbb{H}, \mathbb{H})$ is of course not an $A\mathbb{H}$-module, so there is no stable $A\mathbb{H}$-module U_{-1} which is dual to $U^\times_{-1}$. In spite of this we still regard $U^\times_{-1}$ as 'antistable', because treating it as an exception every time would be cumbersome.

We end this chapter with a diagram (overleaf) summarising much of our theory.

Figure 5.1: Irreducible A$\mathbb{H}$-modules and the ratio of their dimensions.

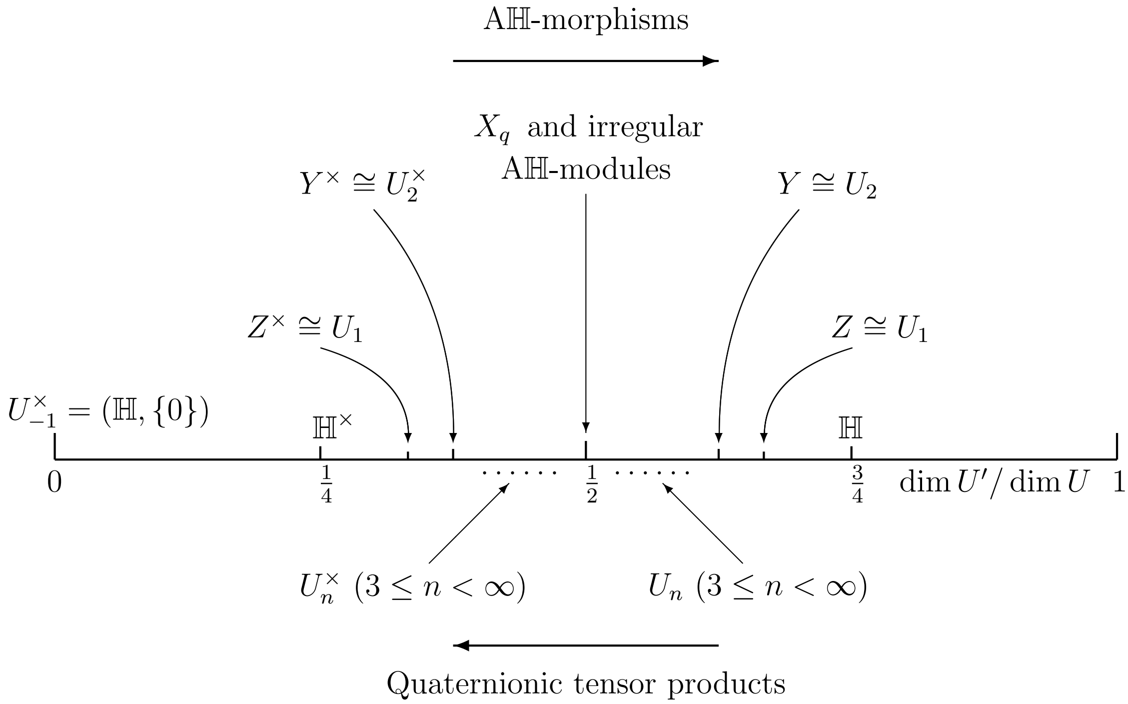

This describes the A$\mathbb{H}$-modules we have met so far — all the finite-dimensional irreducible A$\mathbb{H}$-modules. The quaternionic tensor product of two A$\mathbb{H}$-modules is always to the left of both of them in Figure 5.1. On the other hand, there are A$\mathbb{H}$-morphisms from an A$\mathbb{H}$-module into itself and any A$\mathbb{H}$-modules *to its right* — and never from an A$\mathbb{H}$-module to any A$\mathbb{H}$-module to its left. These statements are closely linked by Theorem 4.2.9. If U and V are irreducible stable or antistable A$\mathbb{H}$-modules, this demonstrates that that there will always be A$\mathbb{H}$-morphisms from $U \otimes_{\mathbb{H}} V$ into U and V, but *never* any A$\mathbb{H}$-morphisms from U or V into $U \otimes_{\mathbb{H}} V$, unless U or V is equal to $\mathbb{H}$.

Chapter 6

Hypercomplex Manifolds

This chapter uses the quaternionic algebra developed in Chapters 4 and 5 to describe hypercomplex manifolds. It is in three parts (the first of which is a summary of Joyce's work, the other two being original). The first part (Section 6.1) summarises Joyce's theory of q-holomorphic functions. A q-holomorphic function on a hypercomplex manifold M is a smooth function $f : M \to \mathbb{H}$ which satisfies a quaternionic version of the Cauchy-Riemann equations. We let $\mathcal{P}_M$ denote the A$\mathbb{H}$-module of q-holomorphic functions on M. The q-holomorphic functions on the hypercomplex manifold $\mathbb{H}$ are precisely the regular functions of Fueter and Sudbery. Because of the noncommutativity of the quaternions, the product of two q-holomorphic functions is not in general q-holomorphic. Nonetheless, q-holomorphic functions possess a rich algebraic structure. Joyce attempts to capture this using the concept of an *H-algebra*, a quaternionic version of a commutative algebra over the real or complex numbers.

In the second part (Section 6.2), we use the Sp(1)-representation $T^*M \cong 2nV_1$ defined by the hypercomplex structure to obtain a natural splitting of the quaternionic cotangent space of a hypercomplex manifold M, which we write $\mathbb{H} \otimes T^*M \cong A \oplus B$. This is precisely a version of the splitting $\mathbb{H} \otimes aV_k \cong aV_1 \otimes (V_{k+1} \oplus V_{k-1})$ used in the previous chapter to construct and describe all stable and antistable A$\mathbb{H}$-modules, and the Sp(1)-version of quaternionic algebra gives a complete description of the geometric situation. A function f is q-holomorphic if and only if its differential df takes values in the subspace $A \subset \mathbb{H} \otimes T^*M$. This is very similar to the situation in complex geometry where the differential of a holomorphic function takes values in the holomorphic cotangent space $\Lambda^{1,0}$. It follows that the subbundle A should be thought of as the *q-holomorphic cotangent space* of M, its complement B being the *q-antiholomorphic cotangent space.*

The third part (Sections 6.3 and 6.4) is about A$\mathbb{H}$-bundles, which are smooth vector bundles whose fibres are A$\mathbb{H}$-modules. They can be defined over any smooth manifold M, but are most interesting when M is hypercomplex. An A$\mathbb{H}$-bundle E is said to be q-holomorphic if it is simultaneously holomorphic with respect to the whole 2-sphere of complex structures on M. This condition is met if and only if E carries an anti-self-dual connection which is compatible with the structure of E as an A$\mathbb{H}$-bundle. We generalise the theory of q-holomorphic functions to that of *q-holomorphic sections*, so that a q-holomorphic function is precisely a q-holomorphic section of the trivial bundle $M \times \mathbb{H}$. We investigate the algebraic structure of q-holomorphic sections in some detail, showing that the q-holomorphic sections of a q-holomorphic vector bundle E form an

H-algebra module over $\mathcal{P}_M$.

6.1 Q-holomorphic Functions and H-algebras

6.1.1 Quaternion-valued functions

This section is mainly a summary of [J1, §§3,5], to which the reader is referred for more detail and proofs of the important results. Let M be a smooth manifold and let $C^\infty(M,\mathbb{H})$ be the vector space of smooth quaternion-valued functions on M. An $\mathbb{H}$-action on $C^\infty(M,\mathbb{H})$ is defined by setting $(q \cdot f)(m) = q(f(m))$ for all $m \in M$, $f \in C^\infty(M,\mathbb{H})$. Thus $C^\infty(M,\mathbb{H})$ is a left $\mathbb{H}$-module. What is less immediately obvious is that $C^\infty(M,\mathbb{H})$ is an A$\mathbb{H}$-module.

Lemma 6.1.1 *Define a linear subspace*

$$C^\infty(M,\mathbb{H})' = \{f \in C^\infty(M,\mathbb{H}) : f(m) \in \mathbb{I} \quad \text{for all } m \in M\} = C^\infty(M,\mathbb{I}).$$

With these definitions, $(C^\infty(M,\mathbb{H}), C^\infty(M,\mathbb{H})')$ *is an* A$\mathbb{H}$*-module.*

Proof. We use the fact that M can be embedded in $C^\infty(M,\mathbb{H})^\dagger$ in the following way. For each $m \in M$, define the 'evaluation map'

$$\theta_m : C^\infty(M,\mathbb{H}) \to \mathbb{H} \qquad \text{by} \qquad \theta_m(f) = f(m).$$

Then $\theta_m(q \cdot f) = q \cdot \theta_m(f)$, so $\theta_m \in C^\infty(M,\mathbb{H})^\times$. Also, if $f \in C^\infty(M,\mathbb{I})$ then $\theta_m(f) \in \mathbb{I}$ for all $m \in M$, so $\theta_m \in C^\infty(M,\mathbb{H})^\dagger$.

Suppose that $f \in C^\infty(M,\mathbb{H})$, and $\alpha(f) = 0$ for all $\alpha \in C^\infty(M,\mathbb{H})^\dagger$. Since $\theta_m \in C^\infty(M,\mathbb{H})^\dagger$, $f(m) = 0$ for all $m \in M$; so $f \equiv 0$. Thus $C^\infty(M,\mathbb{H})$ is an A$\mathbb{H}$-module, by Definition 4.1.1. ∎

This technique of linking a point $m \in M$ to an element of $C^\infty(M,\mathbb{H})^\dagger$ is extremely useful. Joyce has used this process to reconstruct hypercomplex manifolds from their H-algebras. Note that we have assumed no geometric structure on M other than that of a smooth manifold.

Complex and quaternionic functions

For each $q \in S^2$ let $\iota_q : \mathbb{C} \to \mathbb{H}$ be the inclusion obtained by identifying $i \in \mathbb{C}$ with $q \in S^2$, so that $\iota_q(a+ib) = a+qb$. [1] For any real vector space E, the inclusion ι_q extends to a map $\iota_q : \mathbb{C}\otimes E \to \mathbb{H}\otimes E$ given by $\iota_q(e_0 + ie_1) = e_0 + qe_1$ (for $e_0, e_1 \in E$). Note that the images of ι_q and ι_{-q} are the same, but the two maps are not identical: since $\iota_q(a+ib) = a+qb$ and $\iota_{-q}(a+ib) = a-qb$, we see that $\iota_q(z) = \iota_{-q}(\bar{z})$.

Let $f = a+ib$ be a complex-valued function on M. Then for every $q \in S^2$, the function $\iota_q(f) = a+qb$ is a quaternion-valued function on M. In this way we obtain quaternion-valued functions from complex ones, and we shall soon see that this construction can be used to obtain q-holomorphic functions from holomorphic ones.

[1] The map ι_q will be distinguished from the inclusion map ι_U of A$\mathbb{H}$-modules by the use of lower-case rather than upper-case subscripts.

6.1.2 Q-holomorphic functions

In this section we will define q-holomorphic functions, the hypercomplex version of holomorphic functions, and familiarise ourselves with some of their basic properties. Consider the complex manifold (M, I). A complex-valued function $f = f_0 + idf_1 \in C^\infty(M, \mathbb{C})$ is holomorphic (with respect to I) if and only if f satisfies the *Cauchy-Riemann equations*

$$df_0 + Idf_1 = 0. \tag{6.1}$$

Let (M, I_1, I_2, I_3) be a hypercomplex manifold. Here is the definition of a q-holomorphic function on M.

Definition 6.1.2 Let $f \in C^\infty(M, \mathbb{H})$ be a smooth $\mathbb{H}$-valued function on M. Then $f = f_0 + f_1 i_1 + f_2 i_2 + f_3 i_3$ with $f_j \in C^\infty(M)$. The function f is said to be *q-holomorphic* if and only if it satisfies the *Cauchy-Riemann-Fueter equations*

$$Df \equiv df_0 + I_1 df_1 + I_2 df_2 + I_3 df_3 = 0.$$

The set of all q-holomorphic functions on M is called $\mathcal{P}_M$.

The term q-holomorphic is short for *quaternion-holomorphic*, and it is intended to indicate that a q-holomorphic function on a hypercomplex manifold is the appropriate quaternionic analogue of a holomorphic function on a complex manifold. If a function f is q-holomorphic then the function $q \cdot f$ is also q-holomorphic for all $q \in \mathbb{H}$. Thus the set of q-holomorphic functions $\mathcal{P}_M$ forms a left $\mathbb{H}$-submodule of $C^\infty(M, \mathbb{H})$. We adopt the obvious definition for $\mathcal{P}'_M$, namely

$$\mathcal{P}'_M = \{f \in \mathcal{P}_M : f(m) \in \mathbb{I} \text{ for all } m \in M\}.$$

So $\mathcal{P}'_M = \mathcal{P}_M \cap C^\infty(M, \mathbb{H})'$, and $\mathcal{P}_M$ is an A$\mathbb{H}$-submodule of $C^\infty(M, \mathbb{H})$. Thus the q-holomorphic functions $\mathcal{P}_M$ on a hypercomplex manifold M form an A$\mathbb{H}$-module.

The product of two q-holomorphic functions is not in general q-holomorphic — we can observe this simply by noting that all constant functions are trivially q-holomorphic, but $\mathcal{P}_M$ is not even closed under right-multiplication by quaternions. Thus q-holomorphic functions do not form an algebra in the same sense that the holomorphic functions on a complex manifold do.

We show that there are many interesting q-holomorphic functions on a hypercomplex manifold M, by observing that every complex-valued function on M which is holomorphic with respect to any complex structure $Q \in S^2$ gives rise to a q-holomorphic function.

Lemma 6.1.3 *Let* $q = a_1 i_1 + a_2 i_2 + a_3 i_3 \in S^2$ *and let* $Q = a_1 I_1 + a_2 I_2 + a_3 I_3$ *be the corresponding complex structure on* M*. Let* $f = x + iy \in C^\infty(M, \mathbb{C})$ *be holomorphic with respect to* Q*. Then* $\iota_q(f) = x + qy$ *is a q-holomorphic function.*

Proof. The proof is a simple substitution. If $f = x + iy$ is holomorphic with respect to Q then it satisfies the Cauchy-Riemann equations $dx + Q(dy) = 0$, in which case

$$0 = dx + (a_1 I_1 + a_2 I_2 + a_3 I_3)(dy) = D(x + a_1 y i_1 + a_2 y i_2 + a_3 y i_3) = D(x + qy).$$

So the $\mathbb{H}$-valued function $\iota_q(f) = x + qy$ is q-holomorphic. ∎

If $f = x + iy \in C^\infty(M, \mathbb{C})$ is holomorphic with respect to Q, then its conjugate $\bar{f} = x - iy$ is holomorphic with respect to $-Q$. Does this mean that both $x + qy$ and $x - qy$ are q-holomorphic functions? Closer inspection shows that this is not the case — we see that

$$\iota_q(x + iy) = \iota_{-q}(x - iy) = x + qy.$$

So the holomorphic functions with respect to Q and $-Q$ are mapped to *the same* functions when we consider their images under $\iota_{\pm q}$ in $C^\infty(M, \mathbb{H})$. We shall see that this is a consequence of (indeed is equivalent to) the fact that the A$\mathbb{H}$-modules X_q and X_{-q} are the same.

Now, if the functions $f, g \in C^\infty(M, \mathbb{C}_q)$ are both holomorphic with respect to Q, then so is $fg \in C^\infty(M, \mathbb{C}_q)$. By Lemma 6.1.3, $\iota_q(f)$, $\iota_q(g)$, and $\iota_q(fg) = \iota_q(f)\iota_q(g)$ are all q-holomorphic. In this special case where the q-holomorphic functions f and g both take values in a commuting subfield of $\mathbb{H}$, their product is also q-holomorphic. This is reminiscent of the situation described in Lemma 4.1.7, where two elements $u \in U$ and $v \in V$ such that $\iota_U(u) \in \mathbb{C}_q \otimes (U^\dagger)^*$ and $\iota_V(v) \in \mathbb{C}_q \otimes (V^\dagger)^*$ define an element $u \otimes_{\mathbb{H}} v \in U \otimes_{\mathbb{H}} V$. In this case we have two well-defined 'products' of f and g — their quaternionic tensor product $f \otimes_{\mathbb{H}} g \in \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_M$ and their product as $\mathbb{C}_q$-valued functions $fg = gf \in \mathcal{P}_M$. This algebraic situation is described by the theory of H-algebras.

6.1.3 H-algebras and Q-holomorphic functions

An *H-algebra* is a quaternionic version of an algebra over a commutative field. Let $\mathbb{F}$ be a commutative field (usually the real or complex numbers). An $\mathbb{F}$-algebra is a vector space A over $\mathbb{F}$, equipped with an $\mathbb{F}$-bilinear *multiplication map* $\mu : A \times A \to A$ which has certain algebraic properties. For example if $\mu(a, b) = \mu(b, a)$ for all $a, b \in A$ then μ is said to be commutative. As we have already seen in Section 1.3, this formulation is of no great use to us when seeking a quaternionic analogue because the non-commutativity of the quaternions makes the notion of an $\mathbb{H}$-bilinear map untenable.

The axioms for an algebra over $\mathbb{F}$ can alternatively be written in terms of tensor products. The main feature of tensor algebra is that a bilinear map on the cartesian product $A \times B$ translates into a linear map on the tensor product $A \otimes_{\mathbb{F}} B$. So our bilinear multiplication map $\mu : A \times A \to A$ becomes an $\mathbb{F}$-linear map $\mu : A \otimes_{\mathbb{F}} A \to A$. The commutative axiom $\mu(a, b) = \mu(b, a)$ becomes $\mu(a \otimes b) = \mu(b \otimes a)$, so $\mu(a \otimes b - b \otimes a) = \mu(a \wedge b) = 0$. Hence we obtain a 'tensor algebra version' of the commutative axiom, saying that $\mu : A \otimes_{\mathbb{F}} A \to A$ is commutative if and only if $\Lambda^2 A \subseteq \ker \mu$.

Once we have reformulated our axioms in terms of tensor products and linear maps we can translate them into quaternionic algebra, replacing 'vector space', 'linear map' and 'tensor product' with 'A$\mathbb{H}$-module', 'A$\mathbb{H}$-morphism' and 'quaternionic tensor product'. This is precisely what Joyce does, producing the following definition [J1, §5]:

Axiom H. (*i*) $\mathcal{P}$ is an A$\mathbb{H}$-module.

(*ii*) There is an A$\mathbb{H}$-morphism $\mu_{\mathcal{P}} : \mathcal{P} \otimes_{\mathbb{H}} \mathcal{P} \to \mathcal{P}$, called the *multiplication map*.

(*iii*) $\Lambda^2_{\mathbb{H}}\mathcal{P} \subset \ker \mu_{\mathcal{P}}$. Thus $\mu_{\mathcal{P}}$ is *commutative.*

(*iv*) The A$\mathbb{H}$-morphisms $\mu_{\mathcal{P}} : \mathcal{P}\otimes_{\mathbb{H}}\mathcal{P} \to \mathcal{P}$ and $\mathrm{id} : \mathcal{P} \to \mathcal{P}$ combine to give A$\mathbb{H}$-morphisms $\mu_{\mathcal{P}}\otimes_{\mathbb{H}}\mathrm{id}$ and $\mathrm{id}\otimes_{\mathbb{H}}\mu_{\mathcal{P}} : \mathcal{P}\otimes_{\mathbb{H}}\mathcal{P}\otimes_{\mathbb{H}}\mathcal{P} \to \mathcal{P}\otimes_{\mathbb{H}}\mathcal{P}$. Composing with $\mu_{\mathcal{P}}$ gives A$\mathbb{H}$-morphisms $\mu_{\mathcal{P}} \circ (\mu_{\mathcal{P}}\otimes_{\mathbb{H}}\mathrm{id})$ and $\mu_{\mathcal{P}} \circ (\mathrm{id}\otimes_{\mathbb{H}}\mu_{\mathcal{P}}) : \mathcal{P}\otimes_{\mathbb{H}}\mathcal{P}\otimes_{\mathbb{H}}\mathcal{P} \to \mathcal{P}$. Then $\mu_{\mathcal{P}} \circ (\mu_{\mathcal{P}}\otimes_{\mathbb{H}}\mathrm{id}) = \mu_{\mathcal{P}} \circ (\mathrm{id}\otimes_{\mathbb{H}}\mu_{\mathcal{P}})$. This is *associativity of multiplication.*

(*v*) An element $1 \in A$ called the *identity* is given, with $1 \notin \mathcal{P}'$ and $\mathbb{I}\cdot 1 \subseteq \mathcal{P}'$.

(*vi*) Part (*v*) implies that if $\alpha \in \mathcal{P}^\dagger$ then $\alpha(1) \in \mathbb{R}$. Thus for each $a \in \mathcal{P}$, $1\otimes_{\mathbb{H}}a$ and $a\otimes_{\mathbb{H}}1 \in \mathcal{P}\otimes_{\mathbb{H}}\mathcal{P}$ by Lemma 4.1.7. Then $\mu_{\mathcal{P}}(1\otimes_{\mathbb{H}}a) = \mu_{\mathcal{P}}(a\otimes_{\mathbb{H}}1) = a$ for each $a \in \mathcal{P}$. Thus 1 *is a multiplicative identity.*

Definition 6.1.4 $\mathcal{P}$ is an *H-algebra* if $\mathcal{P}$ satisfies Axiom H.

Here H-algebra stands for *Hamilton algebra*. We also define morphisms of H-algebras:

Definition 6.1.5 Let $\mathcal{P}, \mathcal{Q}$ be H-algebras, and let $\phi : \mathcal{P} \to \mathcal{Q}$ be an A$\mathbb{H}$-morphism. Write $1_{\mathcal{P}}, 1_{\mathcal{Q}}$ for the identities in $\mathcal{P}, \mathcal{Q}$ respectively. We say that ϕ is an *H-algebra morphism* if $\phi(1_{\mathcal{P}}) = 1_{\mathcal{Q}}$ and $\mu_{\mathcal{Q}} \circ (\phi\otimes_{\mathbb{H}}\phi) = \phi \circ \mu_{\mathcal{P}}$ as A$\mathbb{H}$-morphisms $\mathcal{P}\otimes_{\mathbb{H}}\mathcal{P} \to \mathcal{Q}$.

Let M be a hypercomplex manifold. Joyce has constructed a multiplication map μ_M with respect to which the q-holomorphic functions $\mathcal{P}_M$ form an H-algebra, and such that for $f, g \in C^\infty(M, \mathbb{C}_q)$ holomorphic with respect to Q we have $\mu(f\otimes_{\mathbb{H}}g) = fg$. In this way the the H-algebra $\mathcal{P}_M$ includes the algebras of holomorphic functions with respect to all the different complex structures on M.

Let M, N and (therefore) $M \times N$ be hypercomplex manifolds. We define an A$\mathbb{H}$-morphism $\phi : \mathcal{P}_M\otimes_{\mathbb{H}}\mathcal{P}_N \to \mathcal{P}_{M\times N}$. Let U and V be A$\mathbb{H}$-modules. Let $x \in U^\dagger \otimes V^\dagger$ and $y \in U\otimes_{\mathbb{H}}V$. Then $y \cdot x \in \mathbb{H}$, where '$\cdot$' contracts together the factors of $U^\dagger \otimes V^\dagger$ with those of $(U^\dagger)^* \otimes (V^\dagger)^*$ in $U\otimes_{\mathbb{H}}V \subset \mathbb{H} \otimes (U^\dagger)^* \otimes (V^\dagger)^*$. Define a linear map

$$\lambda_{UV} : U^\dagger \otimes V^\dagger \to (U\otimes_{\mathbb{H}}V)^\dagger \quad \text{by setting} \quad \lambda_{UV}(x)(y) = y \cdot x, \tag{6.2}$$

for all $x \in U^\dagger \otimes V^\dagger$ and $y \in U\otimes_{\mathbb{H}}V$. [2]

Consider again the maps $\theta_m \in \mathcal{P}_M^\dagger$ of Lemma 6.1.1, which allow us to interpret points of $m \in M$ as A$\mathbb{H}$-morphisms $\theta_m : \mathcal{P}_M \to \mathbb{H}$. In the same way, we can associate to each point $(m, n) \in M \times N$ the map $\theta_m \otimes \theta_n \in \mathcal{P}_M^\dagger \otimes \mathcal{P}_N^\dagger$. Then $\lambda_{\mathcal{P}_M,\mathcal{P}_N}(\theta_m \otimes \theta_n) \in (\mathcal{P}_M\otimes_{\mathbb{H}}\mathcal{P}_N)^\dagger$.

Definition 6.1.6 We define a map $\phi : \mathcal{P}_M\otimes_{\mathbb{H}}\mathcal{P}_N \to C^\infty(M \times N, \mathbb{H})$ as follows. Let $f \in \mathcal{P}_M\otimes_{\mathbb{H}}\mathcal{P}_N$. Then $\lambda_{\mathcal{P}_M,\mathcal{P}_N}(\theta_m \otimes \theta_n) \cdot f \in \mathbb{H}$ and we define $\phi(f) : M \times N \to \mathbb{H}$ by setting

$$\phi(f)(m, n) = \lambda_{\mathcal{P}_M,\mathcal{P}_N}(\theta_m \otimes \theta_n) \cdot f.$$

For infinite-dimensional vector spaces U and V we define $U \otimes V$ to consist of finite sums of elements $u \otimes v$. Thus $\phi(f)$ is a sum of finitely many smooth functions, and so is smooth. It is easy to see that $\phi(f)$ is q-holomorphic, since $D_{M\times N} =$

[2]See [J1, Definition 4.2], where Joyce defines this map and uses it to prove that $U\otimes_{\mathbb{H}}V$ is an A$\mathbb{H}$-module.

$D_M(\phi(f)) + D_N(\phi(f)) = 0$, *i.e.* we can evaluate the derivatives in the M and N directions separately. It is clear that ϕ is $\mathbb{H}$-linear and that $\phi(\mathcal{P}'_M \otimes_{\mathbb{H}} \mathcal{P}'_N) \subseteq \mathcal{P}'_{M\times N}$. Thus we have a canonical A$\mathbb{H}$-morphism $\phi : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_N \to \mathcal{P}_{M\times N}$.

The second (and easier) step is to consider the case $M = N$, so $\phi : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_M \to \mathcal{P}_{M\times M}$. The diagonal submanifold $M_{\text{diag}} = \{(m, m) : m \in M\} \subset M \times M$ is a submanifold of $M \times M$ isomorphic to M as a hypercomplex manifold, and each q-holomorphic function on $M \times M$ restricts to a q-holomorphic function on M_{diag}. Let ρ be the restriction map $\rho : \mathcal{P}_{M\times M} \to \mathcal{P}_{M_{\text{diag}}}$; then ρ is an A$\mathbb{H}$-morphism. Thus we can define an A$\mathbb{H}$-morphism

$$\mu_M = \rho \circ \phi : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_M \to \mathcal{P}_M. \tag{6.3}$$

Here is the key theorem of this section:

Theorem 6.1.7 *[J1, 5.5]. Let M be a hypercomplex manifold, so that $\mathcal{P}_M$ is an* A$\mathbb{H}$-*module. Let $1 \in \mathcal{P}_M$ be the constant function on M with value 1, and let $\mu_M : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_M \to \mathcal{P}_M$ be the* A$\mathbb{H}$-*morphism $\mu_M = \rho \circ \phi$ of Equation (6.3). With these definitions, $\mathcal{P}_M$ is an H-algebra.*

This statement is also true for $C^\infty(M, \mathbb{H})$. An A$\mathbb{H}$-morphism $\phi : C^\infty(M, \mathbb{H}) \otimes_{\mathbb{H}} C^\infty(N, \mathbb{H}) \to C^\infty(M \times N, \mathbb{H})$ can be defined just as in Definition 6.1.6, and ρ is still the obvious restriction. The q-holomorphic functions $\mathcal{P}_M$ form an *H-subalgebra* of $C^\infty(M, \mathbb{H})$, and the inclusion map $\text{id} : \mathcal{P}_M \to \mathbb{C}^\infty(M, \mathbb{H})$ is an H-algebra morphism.

Example 6.1.8 Q-holomorphic functions on $\mathbb{H}$ *[J1, §10].*

Consider the manifold $\mathbb{H}$ with coordinates (x_0, x_1, x_2, x_3) representing the quaternion $x_0 + x_1 i_1 + x_2 i_2 + x_3 i_3$. Then $\mathbb{H}$ is naturally a complex manifold with hypercomplex structure (I_1, I_2, I_3) given by

$$I_1 dx_2 = dx_3, \quad I_2 dx_3 = dx_1, \quad I_3 dx_1 = dx_2 \quad \text{and} \quad I_j dx_0 = dx_j, \quad j = 1, 2, 3. \tag{6.4}$$

Consider the linear $\mathbb{H}$-valued functions on $\mathbb{H}$. Using the Cauchy-Riemann-Fueter Equation (1.20) with the hypercomplex structure in (6.4), we find that $f = q_0 x_0 + q_1 x_1 + q_2 x_2 + q_3 x_3$ is q-holomorphic if and only if $q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3 = 0$. Thus the linear q-holomorphic functions on $\mathbb{H}$ form an A$\mathbb{H}$-submodule of the set of all linear functions, and this submodule is isomorphic to the A$\mathbb{H}$-module $Z \cong U_1$ of Example 5.4.2.

Consider the A$\mathbb{H}$-module $U^{(k)}$ of homogeneous q-holomorphic polynomials of degree k. The spaces $U^{(k)}$ are important in quaternionic analysis and are studied by Sudbery [Su, §6]. Joyce uses the dimension formulae of Proposition 4.1.16 to prove that $S^k_{\mathbb{H}} Z \cong U^{(k)}$. We can easily show that

$$S^k_{\mathbb{H}} Z \cong (k+1) V_1 \otimes V_{k+1} = \begin{cases} (k+1) U_k & k \text{ even} \\ \frac{1}{2}(k+1) U_k & k \text{ odd.} \end{cases}$$

This constructs not only the spaces $U^{(k)}$, but also their structure as $\text{Sp}(1)$-representations, which is crucial to Sudbery's approach.

The space of all q-holomorphic polynomials on $\mathbb{H}$ is therefore given by the sum $\bigoplus_{k=0}^\infty S^k_{\mathbb{H}} Z$, which is naturally an H-algebra. Joyce [J1, Example 5.1] calls this F^Z, the *free* H-*algebra generated by* Z, and this idea enables us to write down the multiplication map $\mu_{\mathbb{H}}$. The full H-algebra $\mathcal{P}_{\mathbb{H}}$ is constructed by adding in convergent power series.

6.2 The Quaternionic Cotangent Space

Let M be a complex manifold and recall the splitting $\mathbb{C} \otimes_{\mathbb{R}} T^*M = \Lambda^{1,0}M \oplus \Lambda^{0,1}M$ of Equation (2.1). A function $f \in C^\infty(M, \mathbb{C})$ is holomorphic if $df \in \Lambda^{1,0}M$. The Cauchy-Riemann operator is defined by $\overline{\partial} = \pi^{0,1} \circ d$, and f is holomorphic if and only if $\overline{\partial} f = 0$.

This section presents the quaternionic analogue of this description. Let M be a hypercomplex manifold. We show that there is a natural splitting of the quaternionic cotangent space $\mathbb{H} \otimes T^*M \cong A \oplus B$. The bundle A is then the *q-holomorphic cotangent space* of M. A function $f \in C^\infty(M, \mathbb{H})$ is q-holomorphic if and only if $df \in A$, and the Cauchy-Riemann-Fueter operator can be written $\bar{\delta} = \pi^B \circ d$, where π^B is the natural projection to the subspace $B \subset \mathbb{H} \otimes T^*M$. The ease with which quaternionic algebra presents such close parallels with complex geometry could scarcely be more rewarding.

Theorem 6.2.1 *Let M^{4n} be a hypercomplex manifold and let $\mathbb{H} \otimes T^*M$ be the quaternionic cotangent bundle of M. The hypercomplex structure determines a natural splitting*

$$\mathbb{H} \otimes T^*M \cong A \oplus B,$$

where $\dim A = 12n$ and $\dim B = 4n$.

Proof. Recall from Section 3.2 that $T^*M \cong 2nV_1$ as an Sp(1)-representation. Call this representation $2nV_1^G$ (since it is the action defined by the geometric structure of M). Following the standard arguments of Chapter 5, we have a splitting

$$\begin{aligned} \mathbb{H} \otimes T^*M &\cong V_1^L \otimes V_1^R \otimes 2nV_1^G \\ &\cong 2nV_1^L \otimes (V_2^{RG} \oplus V_0^{RG}) \\ &\cong A \oplus B, \end{aligned}$$

where $A = 2nV_1^L \otimes V_2^{RG}$ and $B = 2nV_1^L \otimes V_0^{RG}$. ∎

This is a situation with which we are by now very familiar. Using the theory of Chapter 5, we immediately deduce that

- The (fibres of the) subspaces A and B are A$\mathbb{H}$-modules with $A \cong nU_1$ and $B \cong nU_{-1}^{\times}$.
- Dim $A = 12n$ and $\dim A' = 8n$, where $A' = A \cap \mathbb{I} \otimes T^*M$. A is a stable A$\mathbb{H}$-module.
- Dim $B = 4n$ and $\dim B' = 0$, where $B' = B \cap \mathbb{I} \otimes T^*M$. B is an antistable A$\mathbb{H}$-module.
- The mapping $A \oplus B \hookrightarrow \mathbb{H} \otimes T^*M$ is an injective A$\mathbb{H}$-morphism which is an isomorphism of the total spaces. It is not an A$\mathbb{H}$-isomorphism because $\dim(A' \oplus B')$ is smaller than $\dim(\mathbb{I} \otimes T^*M)$.

The importance of this splitting lies partly in the following result:

Theorem 6.2.2 *A function* $f \in C^\infty(M, \mathbb{H})$ *is q-holomorphic if and only if* $df \in C^\infty(A)$.

The proof of this theorem will be in two parts. Firstly, we investigate the projection operators π^A and π^B from $\mathbb{H} \otimes T^*M$ onto A and B. We then show that the Cauchy-Riemann-Fueter operator can be written as $\pi^B \circ d$.

To work out projection maps to A and B we use the Casimir operator for the diagonal $\mathrm{Sp}(1)^{RG}$-action, obtained by coupling right-multiplication on $\mathbb{H}$ with the action of the hypercomplex structure. The diagonal Lie algebra action is given by

$$\mathcal{I}(\omega) = I_1(\omega) - \omega i_1 \qquad \mathcal{J}(\omega) = I_2(\omega) - \omega i_2 \qquad \mathcal{K}(\omega) = I_3(\omega) - \omega i_3. \tag{6.5}$$

Then $C(\omega) = (\mathcal{I}^2 + \mathcal{J}^2 + \mathcal{K}^2)(\omega) = -6\omega - 2\chi(\omega)$, where $\chi(\omega) = I_1(\omega)i_1 + I_2(\omega)i_2 + I_3(\omega)i_3$. The operator χ satisfies the equation $\chi^2 = 3 - 2\chi$, so has eigenvalues $+1$ and -3. If ω is in the -3 eigenspace then we have $C(\omega) = 0$, so $\omega \in V_0^{RG}$. If ω is in the $+1$ eigenspace then we have $C(\omega) = -8$, so $\omega \in V_2^{RG}$. Thus the subspaces A and B are the eigenspaces of the operator χ. This enables us to write down expressions for π^A and π^B such that $\pi^A + \pi^B = 1$.

Lemma 6.2.3 *Projection maps* $\pi^A : \mathbb{H} \otimes T^*M \to A$ *and* $\pi^B : \mathbb{H} \otimes T^*M \to B$ *are given by*

$$\pi^A(\omega) = \frac{1}{4}(3\omega + I_1(\omega)i_1 + I_2(\omega)i_2 + I_3(\omega)i_3) \qquad \pi^B(\omega) = \frac{1}{4}(\omega - I_1(\omega)i_1 - I_2(\omega)i_2 - I_3(\omega)i_3).$$

These maps depend only on the structure of M *as a hypercomplex manifold.*

Just as we defined the Dolbeault operators in Equation (2.4), we define new differential operators δ and $\bar{\delta}$:

$$\begin{array}{l} \delta : \mathbb{C}^\infty(M, \mathbb{H}) \to C^\infty(M, A) \\ \delta = \pi^A \circ d \end{array} \quad \text{and} \quad \begin{array}{l} \bar{\delta} : \mathbb{C}^\infty(M, \mathbb{H}) \to C^\infty(M, B) \\ \bar{\delta} = \pi^B \circ d. \end{array} \tag{6.6}$$

Then for a function $f \in C^\infty(M, \mathbb{H})$, $df = \delta f + \bar{\delta} f$.

Proposition 6.2.4 *A function* $f \in C^\infty(M, \mathbb{H})$ *is q-holomorphic if and only if* $\bar{\delta} f = 0$.

Proof. Let $f = f_0 + f_1 i_1 + f_2 i_2 + f_3 i_3$ be a q-holomorphic function on M, so $Df = df_0 + I_1 df_1 + I_2 df_2 + I_3 df_3 = 0$. This is exactly the real part of the equation

$$4\bar{\delta} f \equiv df - I_1 df i_1 - I_2 df i_2 - I_3 df i_3 = 0. \tag{6.7}$$

Moreover, the three imaginary parts are the equations $I_j(Df) = 0$, which are satisfied if and only if $Df = 0$. Thus the real equation $Df = 0$ is exactly the same as the quaternionic equation $\bar{\delta} f = 0$. ∎

It now follows that a function f is q-holomorphic if and only if $df = \delta f \in C^\infty(M, A)$. This proves Theorem 6.2.2 and motivates the following definition:

Definition 6.2.5 Let M be a hypercomplex manifold and let $\mathbb{H} \otimes T^*M \cong A \oplus B$, as in Theorem 6.2.1. Then A is the *q-holomorphic cotangent space* of M and B is the *q-antiholomorphic cotangent space* of M.

Here are some practical methods for writing the spaces A and B. Let $\omega = 1 \otimes \omega_0 + i_1 \otimes \omega_1 + i_2 \otimes \omega_2 + i_3 \otimes \omega_3$. It is straightforward to verify that

$$\omega \in A \Longleftrightarrow \omega_0 + I_1(\omega_1) + I_2(\omega_2) + I_3(\omega_3) = 0,$$

and

$$\omega \in B \Longleftrightarrow \omega_0 = I_1(\omega_1) = I_2(\omega_2) = I_3(\omega_3).$$

Let M^{4n} be a hypercomplex manifold and let $\{e^a : a = 0, \ldots, 4n-1\}$ be a basis for T_x^*M such that $I_b(e^{4a}) = e^{4a+b}$. In other words, we choose a particular isomorphism $T_x^*M \cong \mathbb{H}^n$ such that the subspace $\langle e^{4a}, e^{4a+1}, e^{4a+2}, e^{4a+3} \rangle$ is a copy of $\mathbb{H}$ with its standard hypercomplex structure. With respect to this basis the fibres of A and B at the point x are given by

$$A_x = \left\{ \sum q_a e^a : q_{4b} + q_{4b+1} i_1 + q_{4b+2} i_2 + q_{4b+3} i_3 = 0 \quad \text{for all } b = 0, \ldots, n-1 \right\} \tag{6.8}$$

and

$$B_x = \left\{ \sum q_a e^a : q_{4b} = -q_{4b+1} i_1 = -q_{4b+2} i_2 = -q_{4b+3} i_3 \quad \text{for all } b = 0, \ldots, n-1 \right\}. \tag{6.9}$$

Equation (6.8) is a higher dimensional version of the equation $q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3 = 0$ which gives the q-holomorphic cotangent space on $\mathbb{H}$ (Example 6.1.8). This gives an A$\mathbb{H}$-isomorphism between A_x and the A$\mathbb{H}$-module $nZ \cong nU_1$.

Let us pause for a moment to reflect on what has happened. On a complex manifold, we have a splitting of $\mathbb{C} \otimes T^*M$ into two spaces of equal dimension, which are conjugate to one another. On a hypercomplex manifold, we have a splitting into two subspaces, but they are not equal in size. Instead, one of them has three times the dimension of the other, because the subspaces are defined by the equation $V_1 \otimes V_1 \cong V_2 \oplus V_0$. This type of splitting into spaces of dimension $3n$ and n is typical of the quaternions, echoing the fundamental isomorphism $\mathbb{H} \cong \mathbb{I} \oplus \mathbb{R}$.

6.2.1 The Holomorphic and Q-holomorphic cotangent spaces

Let M^{4n} be a hypercomplex manifold with hypercomplex structure (I_1, I_2, I_3). Let $q = a_1 i_1 + a_2 i_2 + a_3 i_3 \in S^2$ and let $Q = a_1 I_1 + a_2 I_2 + a_3 I_3$ be the corresponding complex structure on M. We have already seen (Lemma 6.1.3) that if f is a complex-valued function on M which is holomorphic with respect to the complex structure Q, then $\iota_q(f)$ is a q-holomorphic function. This is a correspondence not only of holomorphic and q-holomorphic functions, but also of the holomorphic and q-holomorphic cotangent spaces.

Each complex structure Q on M defines a holomorphic cotangent space $\Lambda^{1,0}_Q$. Using the embedding ι_q we can consider $\mathbb{C}\otimes T^*M \cong \mathbb{C}_q\otimes T^*M$ as a subspace of $\mathbb{H}\otimes T^*M$. The relationship between these subspaces of $\mathbb{H}\otimes T^*M$ is clarified in the following Lemma:

Lemma 6.2.6 *Let $A \subset \mathbb{H}\otimes T^*M$ be the q-holomorphic cotangent space of a hypercomplex manifold M and let $\Lambda^{1,0}_Q \subset \mathbb{C}\otimes T^*M$ be the holomorphic cotangent space with respect to the complex structure Q. Then*

$$\mathbb{H}\cdot \iota_q(\Lambda^{1,0}_Q) = (\mathrm{id}_A \otimes_{\mathbb{H}} \chi_q)(A\otimes_{\mathbb{H}} X_q).$$

Proof. We illustrate the case $\dim M = 4$, the general result holding in exactly the same way but involving more coordinates. For any point $x\in M$ we choose a basis $\{e^0, e^1, e^2, e^3\}$ for $T^*_x M$ such that the hypercomplex structure at x is given by the standard relations (6.4). Without loss of generality take $q = i_1$ and consider the complex manifold (M, I_1). Then

$$\Lambda^{1,0}_{I_1} = \langle e^0 + ie^1, e^2 + ie^3\rangle_{\mathbb{C}}.$$

Since $\iota_{i_1}(a+ib) = a + i_1 b$, we see that

$$\mathbb{H}\cdot \iota_q(\Lambda^{1,0}_{I_1}) = \mathbb{H}\cdot\langle e^0 + i_1e^1, e^2 + i_1e^3\rangle. \tag{6.10}$$

Recall from Section 5.3 that $(\mathrm{id}_U \otimes_{\mathbb{H}} \chi_q)(U\otimes_{\mathbb{H}} X_q)' = U' \cap qU'$ for any A$\mathbb{H}$-module U. Using Equation (6.8), we have $A = \{(q_0e^0 + \ldots + q_3e^3) : q_0 + q_1i_1 + q_2i_2 + q_3i_3 = 0\}$ and $A' = \mathbb{I}\otimes T^*M$. It follows that

$$\begin{aligned} A'\cap i_1A' &= \{q_0e^0 + \ldots + q_3e^3 \in A : q_j \in \langle i_2, i_3\rangle\} \\ &= \langle i_2e^0 - i_3e^1, i_3e^0 + i_2e^1, i_2e^2 - i_3e^3, i_3e^2 + i_2e^3\rangle, \end{aligned} \tag{6.11}$$

and the image $(\mathrm{id}_A \otimes_{\mathbb{H}} \chi_{i_1})(A\otimes_{\mathbb{H}} X_{i_1}) \cong 2X_{i_1}$ is generated over $\mathbb{H}$ by this subspace. Observe that $-i_2\cdot(i_2e^0 - i_3e^1) = e^0 + i_1e^1$ and that $-i_2\cdot(i_2e^2 - i_3e^3) = e^2 + i_1e^3$. Therefore $\mathbb{H}\cdot(A'\cap i_1A')$ is exactly the same as the subspace $\mathbb{H}\cdot\iota_q(\Lambda^{1,0}_{I_1})$ of Equation (6.10). ∎

This gives us a description of the q-holomorphic and q-antiholomorphic cotangent spaces of M in terms of complex geometry:

Corollary 6.2.7 *Let M^{4n} be a hypercomplex manifold, with q-holomorphic cotangent space A and q-antiholomorphic cotangent space B. Then*

$$A = \sum_{Q\in S^2} \mathbb{H}\cdot\iota_q(\Lambda^{1,0}_Q)$$

and

$$B = \bigcap_{Q\in S^2} \mathbb{H}\cdot\iota_q(\Lambda^{0,1}_Q).$$

Thus the q-holomorphic cotangent space is generated over $\mathbb{H}$ by the different holomorphic cotangent spaces.

Proof. Use the fact that A is a stable A$\mathbb{H}$-module; thus A is generated over $\mathbb{H}$ by the subspaces $A'\cap qA'$ for $q\in\mathbb{I}$. The proof now follows immediately from Lemma 6.2.6. ∎

6.3 A$\mathbb{H}$-bundles

An A$\mathbb{H}$-bundle is the natural quaternionic analogue of a real or complex vector bundle: a set of A$\mathbb{H}$-modules E parametrised smoothly by a base manifold M.

Definition 6.3.1 Let M be a differentiable manifold and (W, W') a fixed A$\mathbb{H}$-module. A *smooth* A$\mathbb{H}$*-module bundle with fibre* W or simply A$\mathbb{H}$*-bundle* is a family $\{(E_x, E_x')\}_{x \in M}$ of A$\mathbb{H}$-modules parametrised by M, together with a differentiable manifold structure on $E = \bigcup_{x \in M} E_x$ such that

- The projection map $\pi : E \to M$, $\pi : E_x \mapsto x$ is C^∞.
- For every $x_0 \in M$, there exists an open set $U \subseteq M$ containing x_0 and a diffeomorphism
$$\phi_U : \pi^{-1}(U) \to U \times W$$
such that $\phi(E_x) \mapsto \{x\} \times W$ is an A$\mathbb{H}$-isomorphism for each $x \in U$.

A section of an A$\mathbb{H}$-bundle is a smooth map $\alpha : M \to E$ such that $\alpha(x) \in E_x$ for all $x \in M$, and the left $\mathbb{H}$-module of smooth sections of E is denoted by $C^\infty(M, E)$. Define $C^\infty(M, E)' = C^\infty(M, E')$. With this definition, it is easy to adapt Lemma 6.1.1 to show that if E is an A$\mathbb{H}$-bundle then $C^\infty(M, E)$ is an A$\mathbb{H}$-module. We can multiply sections on the left by quaternion-valued functions, and $C^\infty(M, E)$ is a left module over the ring $C^\infty(M, \mathbb{H})$.

The standard real or complex algebraic operations which give new bundles from old can be transfered *en masse* to quaternionic algebra and A$\mathbb{H}$-bundles. Let (E, π_1, M) and (F, π_2, N) be A$\mathbb{H}$-bundles. We define a bundle $(E \otimes_{\mathbb{H}} F, M \times N)$ by setting $(E \otimes_{\mathbb{H}} F)_{(m,n)} = E_m \otimes_{\mathbb{H}} F_n$. We can define a bundle $E \oplus F$ in the same way.

In particular, let E and F be A$\mathbb{H}$-bundles over M. By restricting to the diagonal submanifold of $M \times M$ we define a bundle $(E \otimes_{\mathbb{H}} F, M)$ by setting $(E \otimes_{\mathbb{H}} F)_m = E_m \otimes_{\mathbb{H}} F_m$. In the same way $(E \oplus F, M)$, $\Lambda^k_{\mathbb{H}} E$ and $S^k_{\mathbb{H}}(E)$ are A$\mathbb{H}$-bundles, as is $E^\times$ if the fibres of E are SA$\mathbb{H}$-modules. We will always make clear whether an A$\mathbb{H}$-bundle '$E \otimes_{\mathbb{H}} F$' refers to $(E \otimes_{\mathbb{H}} F, M \times M)$ or the diagonal restriction $(E \otimes_{\mathbb{H}} F, M)$.

The q-holomorphic and q-antiholomorphic cotangent bundles A and B are both A$\mathbb{H}$-bundles, and are A$\mathbb{H}$-subbundles of $\mathbb{H} \otimes T^*M$. In fact, every A$\mathbb{H}$-bundle can be regarded as an A$\mathbb{H}$-subbundle of some bundle $\mathbb{H} \otimes V$, as is easy to demonstrate:

Lemma 6.3.2 *Let* $E = (E, \pi, M)$ *be an* A$\mathbb{H}$*-bundle. Then* E *is isomorphic to an* A$\mathbb{H}$*-subbundle of* $\mathbb{H} \otimes V$*, for some real vector bundle* (V, π_1, M)*.*

Proof. This is an extension of the fact that any A$\mathbb{H}$-module W is isomorphic to $\iota_W(W) \subseteq \mathbb{H} \otimes (W^\dagger)^*$. Consider the map $\iota_{E_x}(E_x) : E_x \to \mathbb{H} \otimes (E_x^\dagger)^*$ for $x \in M$. Let $V_x = (E_x^\dagger)^*$. Then $V = \bigcup_{x \in M} V_x$ is a real vector bundle, and $\mathbb{H} \otimes V$ is an A$\mathbb{H}$-bundle. Thus E is isomorphic to the A$\mathbb{H}$-subbundle

$$\bigcup_{x \in M} \iota_{E_x}(E_x) \subseteq \mathbb{H} \otimes V,$$

and the lemma is proved. ∎

Abusing notation slightly, if E is an A$\mathbb{H}$-bundle we will often write $\iota_E : E \to \mathbb{H}\otimes(E^\dagger)^*$ for the set of inclusion maps $\bigcup_{x\in M}(\iota_{E_x} : E_x \to \mathbb{H}\otimes(E_x^\dagger)^*)$.

We could broaden our definition of an A$\mathbb{H}$-bundle, demanding only that E should be a real vector bundle possessing a left $\mathbb{H}$-action and a real subbundle E' such that (E_x, E'_x) is an A$\mathbb{H}$-module for all $x \in M$. The following example shows why this definition would be unwieldy:

Example 6.3.3 Let $M = \mathbb{R}$, and define a real vector bundle $E = \mathbb{R}\times\mathbb{H}$. Define a real subspace $E' \subset E$ by

$$E'_x = \langle i_1, i_2 + xi_3\rangle.$$

It is easy to check that $(E_x, E'_x) = X_{i_3 - xi_2}$ as an A$\mathbb{H}$-module. Thus E is *not* an A$\mathbb{H}$-bundle, since the fibres E_x and E_y are not A$\mathbb{H}$-isomorphic for $x \neq y$. There is good reason for excluding E from being an A$\mathbb{H}$-bundle. For example,

$$(E\otimes_{\mathbb{H}} X_{i_3})_x = \begin{cases} X_{i_3} & \text{if } x = 0 \\ 0 & \text{otherwise.} \end{cases}$$

So despite the fact that E is a smooth vector bundle, $E\otimes_{\mathbb{H}} X_q$ is not; if we considered this broader class of objects to be 'A$\mathbb{H}$-bundles' then they would not be closed under the tensor product.

Let V be a real vector bundle over the manifold M with fibre $\mathbb{R}^k$, and let P^V be the bundle of frames associated with V. Then P^V is a principal bundle with fibre $\mathrm{GL}(k,\mathbb{R})$, and $V = P^V \times_{\mathrm{GL}(k,\mathbb{R})} \mathbb{R}^k$. In the same way, let E be an A$\mathbb{H}$-bundle with fibre W and let $G = \mathrm{Aut}_{\mathrm{A}\mathbb{H}}(W)$. We can define a principal G-bundle P^E associated with E so that $E = P^E \times_G W$. As we know, G will be a subgroup of $\mathrm{GL}(k,\mathbb{H})$. In most cases, G will be significantly smaller than $\mathrm{GL}(k,\mathbb{H})$. For example, if $W = U_{2j}$ for some j then by Theorem 4.2.9 we have $\mathrm{Hom}_{\mathrm{A}\mathbb{H}}(U_{2j}, U_{2j}) \cong \mathbb{R}$, so $\mathrm{Aut}_{\mathrm{A}\mathbb{H}}(U_{2j}) = \mathbb{R}^* = \mathbb{R}\setminus\{0\}$ and P^E is just a principal $\mathbb{R}^*$-bundle. Principal bundles associated with A$\mathbb{H}$-bundles tend to be much smaller than those associated with real or complex vector bundles.

6.3.1 Connections on A$\mathbb{H}$-bundles

Definition 6.3.4 Let E be an A$\mathbb{H}$-bundle over M. An A$\mathbb{H}$-*connection* on E is an A$\mathbb{H}$-morphism

$$\nabla_E : C^\infty(M,E) \longrightarrow C^\infty(M, E\otimes T^*M) = \Omega^1(M,E)$$

which satisfies the rule

$$\nabla_E(f\cdot\alpha) = df\otimes\alpha + f\cdot\nabla_E\alpha \qquad \text{for all } \alpha\in C^\infty(M,E),\ f\in C^\infty(M,\mathbb{H}).$$

We will sometimes omit the subscript and just write '∇' when the A$\mathbb{H}$-bundle E is clearly implied. Let ∇ be an A$\mathbb{H}$-connection on E. Then the map $\mathrm{id}\otimes(\nabla_E^\dagger)^*$ is an A$\mathbb{H}$-connection on $\mathbb{H}\otimes(E^\dagger)^*$ which preserves $\iota_E(E)$. Thus every A$\mathbb{H}$-connection ∇ on E is equivalent to a connection $(\nabla^\dagger)^*$ on the real vector bundle $(E^\dagger)^*$. We will not distinguish between ∇ and $(\nabla^\dagger)^*$, but will write ∇ for both. This equivalence allows us to construct A$\mathbb{H}$-connections on tensor products, in very much the same way that Joyce constructs the A$\mathbb{H}$-morphism $\phi\otimes_{\mathbb{H}}\psi$ from A$\mathbb{H}$-morphisms ϕ and ψ [J1, Definition 4.4].

Lemma 6.3.5 *Let* (E, M) *and* (F, N) *be* $\mathrm{A}\mathbb{H}$*-bundles and let* ∇_E *and* ∇_F *be* $\mathrm{A}\mathbb{H}$*-connections on* E *and* F*. Then* ∇_E *and* ∇_F *induce an* $\mathrm{A}\mathbb{H}$*-connection* $\nabla_{E \otimes_{\mathbb{H}} F}$ *on the* $\mathrm{A}\mathbb{H}$*-bundle* $(E \otimes_{\mathbb{H}} F, M \times N)$.

Proof. We regard ∇_E as a connection on $(E^\dagger)^*$. The bundle $((E^\dagger)^*, M)$ extends trivially to a bundle $((E^\dagger)^*, M \times N)$ over $M \times N$. Using the natural identification $T^*_{(m,n)}(M \times N) \cong T^*_m M \oplus T^*_n N$, define a differential operator $\nabla_{E,M}$ on $((E^\dagger)^*, M \times N)$, where $\nabla_{E,M}$ differentiates only in the M-directions of $M \times N$. In the same way define a connection $\nabla_{F,N}$ on $((F^\dagger)^*, M \times N)$ which differentiates only in the N-directions. Adding these gives a differential operator

$$\nabla_{E,M} \otimes \mathrm{id} + \mathrm{id} \otimes \nabla_{F,N} : C^\infty((E^\dagger)^* \otimes (F^\dagger)^*) \longrightarrow C^\infty((E^\dagger)^* \otimes (F^\dagger)^* \otimes T^*(M \times N)).$$

This operator is a connection on the real vector bundle $(E^\dagger)^* \otimes (F^\dagger)^*$, so the operator $\mathrm{id} \otimes (\nabla_{E,M} \otimes \mathrm{id} + \mathrm{id} \otimes \nabla_{F,N})$ on the bundle $(\mathbb{H} \otimes (E^\dagger)^* \otimes (F^\dagger)^*, M \times N)$ is an $\mathrm{A}\mathbb{H}$-connection.

Since ∇_E is an $\mathrm{A}\mathbb{H}$-connection, the operator $\mathrm{id} \otimes \nabla_{E,M} \otimes \mathrm{id}$ maps sections of $\iota_E(E) \otimes (F^\dagger)^*$ to sections of $\iota_E(E) \otimes (F^\dagger)^* \otimes T^*(M \times N)$, and acts trivially on the $(F^\dagger)^*$ factor; thus $\mathrm{id} \otimes \nabla_{E,M} \otimes \mathrm{id}$ maps sections of $E \otimes_{\mathbb{H}} F$ to sections of $E \otimes_{\mathbb{H}} F \otimes T^*(M \times N)$. The same is true for $\mathrm{id} \otimes \mathrm{id} \otimes \nabla_{F,N}$, and so the $\mathrm{A}\mathbb{H}$-connection $\mathrm{id} \otimes (\nabla_{E,M} \otimes \mathrm{id} + \mathrm{id} \otimes \nabla_{F,N})$ preserves the $\mathrm{A}\mathbb{H}$-subbundle $(E \otimes_{\mathbb{H}} F, M \times N)$, on which it is therefore an $\mathrm{A}\mathbb{H}$-connection. We call this connection $\nabla_{E \otimes_{\mathbb{H}} F}$. ∎

If E and F are $\mathrm{A}\mathbb{H}$-bundles over the same manifold M we can use Lemma 6.3.5 to define an $\mathrm{A}\mathbb{H}$-connection on the diagonal bundle $(M, E \otimes_{\mathbb{H}} F)$.

This process can be described in terms of principal bundles. [3] Suppose that $\phi \in \mathrm{Aut}_{\mathrm{A}\mathbb{H}}(W)$ for some $\mathrm{A}\mathbb{H}$-module W. Then ϕ induces a real linear isomorphism $(\phi^\times)^* : (W^\dagger)^* \to (W^\dagger)^*$. In this way the action of $G = \mathrm{Aut}_{\mathrm{A}\mathbb{H}}(W)$ on W gives rise to an action on $(W^\dagger)^*$. Let E be an $\mathrm{A}\mathbb{H}$-bundle with fibre W, and let P be the associated principal bundle with fibre G, so that $E = P \times_G W$. An $\mathrm{A}\mathbb{H}$-connection ∇_E on E gives rise to a connection D_P on the principal bundle P, and vice versa. Since connections on principal bundles always exist, we can use this to guarantee that every $\mathrm{A}\mathbb{H}$-bundle has an $\mathrm{A}\mathbb{H}$-connection.

Let F be another $\mathrm{A}\mathbb{H}$-bundle with fibre V, let $H = \mathrm{Aut}_{\mathrm{A}\mathbb{H}}(V)$ and let Q be the associated principal bundle with fibre H, so that $F = Q \times_H V$. An $\mathrm{A}\mathbb{H}$-connection ∇_F on F gives rise to a connection D_Q on the principal bundle Q. Consider the principal bundle $P \times Q$, a bundle over $M \times N$ with fibre $G \times H$. Since G acts on $(W^\dagger)^*$ and H acts on $(V^\dagger)^*$, there is an induced action of $G \times H$ on $W \otimes_{\mathbb{H}} V$. The $\mathrm{A}\mathbb{H}$-bundle $(E \otimes_{\mathbb{H}} F, M \times N)$ is the associated bundle $(P \times Q) \times_{G \times H} W \otimes_{\mathbb{H}} V$. Since the connections D_P and D_Q define a unique connection $D_P \oplus D_Q$ on the product bundle $P \times Q$, this gives rise to a connection on the associated bundle $E \otimes_{\mathbb{H}} F$. Since $G \times H$ acts as $\mathrm{A}\mathbb{H}$-automorphisms on $W \otimes_{\mathbb{H}} V$, this induced connection is an $\mathrm{A}\mathbb{H}$-connection, which we can identify with the $\mathrm{A}\mathbb{H}$-connection $\nabla_{E \otimes_{\mathbb{H}} F}$ of Lemma 6.3.5.

---

[3] A similar discussion for connections in complex vector bundles can be found in [K, §1.5].

6.4 Q-holomorphic $A\mathbb{H}$-bundles

Using complex geometry as a model, we define a *q-holomorphic* $A\mathbb{H}$*-bundle* over a hypercomplex manifold M to be one which is holomorphic with respect to each complex structure $Q \in S^2$. An $A\mathbb{H}$-bundle is q-holomorphic if and only if it admits a connection whose curvature takes values only in $E_{2,0} = \Lambda^2_-$. We define the q-holomorphic sections of a q-holomorphic vector bundle using a version of the Cauchy-Riemann-Fueter equations. Q-holomorphic sections have interesting algebraic properties, which generalise those of Joyce's q-holomorphic functions. We give an interpretation of q-holomorphic sections in terms of quaternionic algebra and the q-holomorphic cotangent space, which makes the step-by-step correspondence with complex geometry extremely clear.

Here is the main definition of this section:

Definition 6.4.1 Let E be an $A\mathbb{H}$-bundle over the hypercomplex manifold M, equipped with an $A\mathbb{H}$-connection ∇. Suppose that ∇ gives E the structure of a holomorphic vector bundle over the complex manifold (M, Q) for all $Q \in S^2$. Then (E, ∇) is a *q-holomorphic* $A\mathbb{H}$*-bundle.*

If we want to be really specific we can write (E, π, M, ∇) for our q-holomorphic bundle; on the other hand, if a connection is already specified, we can simply refer to the total space E as a q-holomorphic $A\mathbb{H}$-bundle. The existence of such a connection ensures that the different holomorphic structures are compatible. A lot can be learned about q-holomorphic $A\mathbb{H}$-bundles by studying the connection ∇.

6.4.1 Anti-self-dual connections and q-holomorphic $A\mathbb{H}$-bundles

Let E be a complex vector bundle over the complex manifold (M, I) and let ∇ be a connection on E. Just as we split the exterior differential $d = \partial + \overline{\partial}$, we can define $\partial^E = \pi^{1,0} \circ \nabla : C^\infty(E) \to C^\infty(E \otimes \Lambda^{1,0} M)$ and $\overline{\partial}^E = \pi^{0,1} \circ \nabla : C^\infty(E) \to C^\infty(E \otimes \Lambda^{0,1} M)$.

Proposition 6.4.2 *The connection ∇ gives E the structure of a holomorphic vector bundle if and only if $\overline{\partial}^E \circ \overline{\partial}^E = 0$, i.e. the $(0,2)$-component of the curvature R of ∇ vanishes. Conversely, if E is a holomorphic vector bundle then E admits such a connection.*

Proof. This is Propositions 3.5 and 3.7 of [K, p. 9]. ∎

Let E be a q-holomorphic $A\mathbb{H}$-bundle over the hypercomplex manifold M, so E has the structure of a holomorphic vector bundle over (M, Q) for every complex structure $Q \in S^2$. Let ∇ be a connection on E and let R be the curvature of ∇. Each $Q \in S^2$ defines a splitting $\nabla = \partial^E_Q + \overline{\partial}^E_Q$, and so a decomposition of the curvature tensor $R = R^{2,0}_Q + R^{1,1}_Q + R^{0,2}_Q$. By Proposition 6.4.2, ∇ defines a holomorphic structure on E with respect to Q if and only if $R^{0,2}_Q \equiv 0$. If we reverse the sense of Q then we reverse the decomposition of R, so $R^{0,2}_{-Q} = R^{2,0}_Q$. Thus ∇ gives E the structure of a holomorphic vector bundle with respect to every $Q \in S^2$ if and only if $R^{2,0}_Q = R^{0,2}_Q = 0$ for all $Q \in S^2$, *i.e.* the curvature of ∇ is of type $(1,1)$ with respect to all $Q \in S^2$.

A 2-form ω is of type $(1,1)$ with respect to every $Q \in S^2$ if and only if it is annihilated by the action of $\mathfrak{sp}(1)$ on $\Lambda^2 T^*M$, so $\omega \in E_{2,0} \subset \Lambda^2 T^*M$. By analogy with the 4-dimensional theory (see Example 3.2.6), we call $E_{2,2}$ the space of *self-dual* 2-forms Λ^2_+ and $E_{2,0}$ the space of *anti-self-dual* 2-forms Λ^2_-. A connection ∇ whose curvature R takes values only in $C^\infty(\mathrm{End}(E) \otimes E_{2,0})$ is therefore called an *anti-self-dual connection.*

Several authors have considered self-dual and anti-self-dual connections, particulary on quaternionic Kähler manifolds: for example Galicki and Poon [GP], Nitta [N] and Mamone Capria and Salamon [MS]. This is partly because such connections give minima of a Yang-Mills funtional on M. It is important to note that Mamone Capria and Salamon refer to the bundle $E_{2,0}$ as self-dual rather than anti-self-dual, and so refer to connections taking values in $E_{2,0}$ as self-dual connections. There is no unanimous convention in the literature: our choice is made because we are using the conventions that $I_j(e^0) = e^j$ not $-e^j$, and that the volume form e^{0123} gives a positive orientation of $\mathbb{H}$.

We have established the following result:

Theorem 6.4.3 *Let E be an* $\mathrm{A}\mathbb{H}$*-bundle over the hypercomplex manifold M equipped with an anti-self-dual* $\mathrm{A}\mathbb{H}$*-connection nabla. Then (E, ∇) is a q-holomorphic* $\mathrm{A}\mathbb{H}$*-bundle.*

Conversely, an $\mathrm{A}\mathbb{H}$*-bundle E equipped with an* $\mathrm{A}\mathbb{H}$*-connection ∇ is q-holomorphic if and only if ∇ is anti-self-dual.*

This is similar to the idea of *hyperholomorphic bundles* described by Verbitsky [V, §2]. Verbitsky considers the case where B is a Hermitian vector bundle over a hyperkähler manifold, so the fibres of such a bundle are not normally $\mathbb{H}$-modules. There is no reason why his definition cannot be extended to hypercomplex manifolds, since as we have seen, the Hermitian inner product is not necessary to force the connection to be anti-self-dual if it is to be compatible with the 2-sphere of holomorphic structures. If E is a q-holomorphic $\mathrm{A}\mathbb{H}$-bundle then it is easy to see that $\mathbb{H} \otimes (E^\dagger)^*$ is also q-holomorphic, and $(E^\dagger)^* \otimes \mathbb{C}$ is a hyperholomorphic bundle. Conversely, if B is a hyperholomorphic bundle with a real structure σ such that $B = B^\sigma \otimes \mathbb{C}$, then $\mathbb{H} \otimes B^\sigma$ is a q-holomorphic $\mathrm{A}\mathbb{H}$-bundle. Any $\mathrm{A}\mathbb{H}$-subbundle E of $\mathbb{H} \otimes B^\sigma$ which is preserved by the anti-self-dual connection ∇ will also be q-holomorphic.

The following Proposition gives further insight into the analogy between holomorphic and q-holomorphic bundles:

Proposition 6.4.4 *Let (E, π, M, ∇) be a q-holomorphic* $\mathrm{A}\mathbb{H}$*-bundle over the hypercomplex manifold M. Then E is itself a hypercomplex manifold.*

Proof. Consider the splitting of TE into horizontal and vertical subbundles $TE \cong E \oplus TM$, where the vertical subbundle E is naturally defined by the structure of E as an $\mathrm{A}\mathbb{H}$-bundle and the horizontal subbundle TM is defined by the connection.

Define a hypercomplex structure (I_1, I_2, I_3) on TE as follows. On the vertical subbundle isomorphic to E, the action of I_1, I_2 and I_3 is given by the fixed left $\mathbb{H}$-action of i_1, i_2 and i_3 on the fibres. On the horizontal subbundle isomorphic to TM, define the hypercomplex structure (I_1, I_2, I_3) to be the horizontal lift of the hypercomplex structure on M. This defines an almost-hypercomplex structure on E. The integrability of this structure is guaranteed by the fact that (E, I_j) is a holomorphic vector bundle for $j = 1, 2, 3$. ∎

Example 6.4.5 Let M be a hypercomplex manifold. Then the Obata connection ∇ on T^*M is anti-self-dual and defines an anti-self-dual A$\mathbb{H}$-connection on $\mathbb{H} \otimes T^*M$, which is thus a q-holomorphic A$\mathbb{H}$-bundle. Let $A \subset \mathbb{H} \otimes T^*M$ be the q-holomorphic cotangent space of M,

$$A = \{\omega_0 + i_1 \otimes \omega_1 + i_2 \otimes \omega_2 + i_3 \otimes \omega_3 : \omega_0 + I_1\omega_1 + I_2\omega_2 + I_3\omega_3 = 0, \quad \omega_j \in T^*M\}.$$

Let $a = a_0 + i_1 a_1 + i_2 a_2 + i_3 a_3 \in C^\infty(M, A)$. Then $\nabla a_0 + I_1(\nabla a_1) + I_2(\nabla a_2) + I_3(\nabla a_3) = 0$ (where I_1, I_2 and I_3 act on the A factor of $A \otimes T^*M$), since $\nabla I_j = 0$ for $j = 1, 2, 3$. It follows that $\nabla a \in C^\infty(A \otimes T^*M)$. Hence the Obata connection defines an A$\mathbb{H}$-connection on A. Since the Obata connection is anti-self-dual, A is a q-holomorphic A$\mathbb{H}$-bundle.

Example 6.4.6 Let E and F be q-holomorphic A$\mathbb{H}$-bundles. Then so are the various associated bundles $E \otimes_{\mathbb{H}} F$, $E \oplus F$, $\Lambda^k_{\mathbb{H}} E$ and so on, all with their induced A$\mathbb{H}$-connections.

6.4.2 Q-holomorphic sections

Q-holomorphic sections of A$\mathbb{H}$-bundles are defined using a version of the Cauchy-Riemann-Fueter equations. We cannot automatically rewrite Equation (6.7) to define a Cauchy-Riemann-Fueter operator on a general A$\mathbb{H}$-bundle E, because the fibres of E will not in general have a well-defined right $\mathbb{H}$-action. Instead, we use the inclusion map ι_E to manipulate sections of e in a more manageable form.

Definition 6.4.7 Let E be an A$\mathbb{H}$-bundle over the hypercomplex manifold M equipped with an A$\mathbb{H}$-connection ∇. Let $e \in C^\infty(M, E)$ so that $\iota_E(e) = e_0 + i_1 e_1 + i_2 e_2 + i_3 e_3$, $e_j \in C^\infty(M, (E^\dagger))^*$. Then e is a *q-holomorphic section* of E if and only if

$$\nabla e_0 + I_1(\nabla e_1) + I_2(\nabla e_2) + I_3(\nabla e_3) = 0, \tag{6.12}$$

where I_1, I_2 and I_3 act on the T^*M factor of $E \otimes T^*M$.

Let $\mathcal{P}(M, E) = \mathcal{P}(E)$ be the space of q-holomorphic sections of E. Let $\mathcal{P}(E)' = \mathcal{P}(E) \cap C^\infty(M, E')$ so that $\mathcal{P}(E)$ is an A$\mathbb{H}$-submodule of $C^\infty(M, E)$.

Q-holomorphic sections are the natural generalisation of q-holomorphic functions, which are precisely the q-holomorphic sections of the trivial bundle $M \times \mathbb{H}$ equipped with the flat connection. So $\mathcal{P}_M = \mathcal{P}(M \times \mathbb{H})$.

A general A$\mathbb{H}$-bundle might have no q-holomorphic sections. However, if the A$\mathbb{H}$-bundle (E, ∇) is q-holomorphic then E is holomorphic with respect to the complex structure Q, and admits holomorphic sections. It is easy to adapt Lemma 6.1.3 to show that the holomorphic sections of this holomorphic vector bundle give rise to q-holomorphic sections of E.

Not all q-holomorphic A$\mathbb{H}$-bundles have interesting q-holomorphic sections. For example, the q-antiholomorphic cotangent space $B \subset \mathbb{H} \otimes T^*M$ is closed under the action of the Obata connection, and so is technically a q-holomorphic A$\mathbb{H}$-bundle. Recall that

$$B = \{\omega - i_1 \otimes I_1(\omega) - i_2 \otimes I_2(\omega) - i_3 \otimes I_3(\omega) : \omega \in T^*M\}.$$

Let $b = b_0 - i_1 I_1 b_0 - i_2 I_2 b_0 - i_3 I_3 b_0 \in C^\infty(M, B)$, so b is q-holomorphic if and only if

$$\nabla b_0 - I_1 \nabla I_1 b_0 - I_2 \nabla I_2 b_0 - I_3 \nabla I_3 b_0 = 0.$$

Since $\nabla I_j = 0$ and $I_j^2 = -1$, this equation is satisfied if and only if $\nabla b = 0$, and the A$\mathbb{H}$-bundle B admits no non-trivial q-holomorphic sections. In fact, this follows from the fact that $B' = 0$, and such behaviour is predicted and described by the algebra of the quaternionic cotangent space.

Q-holomorphic sections and the quaternionic cotangent space

Let E be a q-holomorphic A$\mathbb{H}$-bundle over the hypercomplex manifold M. We can describe the q-holomorphic sections of E using the q-holomorphic cotangent space, just as we did for q-holomorphic functions in the previous chapter. The inclusion map ι_E and the A$\mathbb{H}$-connection ∇ (regarded as a connection on $(E^\dagger)^*$) define an A$\mathbb{H}$-morphism $\nabla \circ \iota_E : C^\infty(M, E) \to C^\infty(M, \mathbb{H} \otimes (E^\dagger)^* \otimes T^*M)$. After swapping the $\mathbb{H}$ and $(E^\dagger)^*$ factors, we have a copy of the quaternionic cotangent space $\mathbb{H} \otimes T^*M$, which we can split into the q-holomorphic and q-antiholomorphic spaces A and B. Thus we have a map $\nabla \circ \iota_E : C^\infty(M, E) \to C^\infty(M, (E^\dagger)^* \otimes (A \oplus B))$. We can now use the projections π^A and π^B of the previous chapter to split the action of $\nabla \circ \iota_E$ into two operators.

Definition 6.4.8 Let (E, ∇) be a q-holomorphic A$\mathbb{H}$-bundle. We define the pair of operators $\delta^E : C^\infty(E) \to C^\infty((E^\dagger)^* \otimes A)$ and $\bar{\delta}^E : C^\infty(E) \to C^\infty((E^\dagger)^* \otimes B)$ by

$$\delta^E = (\mathrm{id}_{(E^\dagger)^*} \otimes \pi^A) \circ \nabla \circ \iota_E$$

and

$$\bar{\delta}^E = (\mathrm{id}_{(E^\dagger)^*} \otimes \pi^B) \circ \nabla \circ \iota_E.$$

Hence we regard E as a subspace of $\mathbb{H} \otimes (E^\dagger)^*$, and then the action of ∇ on E splits into a q-holomorphic part δ^E and a q-antiholomorphic part $\bar{\delta}^E$ so that in effect $\nabla = \delta^E + \bar{\delta}^E$. This is an exact analogue of the complex case where a connection ∇ splits as $\nabla = \pi^{1,0} \circ \nabla + \pi^{0,1} \circ \nabla$. Let V be a holomorphic vector bundle with a connection ∇ compatible with the holomorphic structure, so that $\overline{\partial}^E = \overline{\partial}$. A section $v \in C^\infty(V)$ is holomorphic if and only if $\pi^{0,1} \circ \nabla(s) = 0$. In just the same way, the operator $e \mapsto \nabla e_0 + I_1(\nabla e_1) + I_2(\nabla e_2) + I_3(\nabla e_3)$ of Definition 6.4.7 is precisely the real part of $\bar{\delta}^E$, and a section $e \in C^\infty(E)$ is q-holomorphic if and only if $\bar{\delta}^E(e) = 0$.

This analogy goes further. In complex geometry, a section $v \in C^\infty(V)$ is holomorphic if and only if $\overline{\partial} v = 0$, which means that $\nabla v \in C^\infty(V \otimes_{\mathbb{C}} \Lambda^{1,0} M)$. Here is the quaternionic analogue of this statement:

Proposition 6.4.9 *Let* (E, π, M, ∇) *be a q-holomorphic* $\mathrm{A}\mathbb{H}$*-bundle, and let* $e \in C^\infty(E)$ *be a q-holomorphic section. Then*

$$\nabla \circ \iota_E(e) \in C^\infty(E \otimes_{\mathbb{H}} A),$$

where A *is the q-holomorphic cotangent space of* M*.*

Proof. Let e be q-holomorphic, so $\bar{\delta}^E(e) = 0$ and $\nabla \circ \iota_E(e) = \delta^E(e) \in C^\infty(\mathbb{H} \otimes (E^\dagger)^* \otimes T^*M)$. Using the identification $(A^\dagger)^* \cong T^*M$ we can regard $\delta^E(e)$ as a section of $\mathbb{H} \otimes (E^\dagger)^* \otimes (A^\dagger)^*$.

The induced connection ∇ on $(E^\dagger)^*$ preserves $\iota_E(E)$, so $\delta^E(e) \in C^\infty(\iota_E(E) \otimes (A^\dagger)^*)$. Clearly $\delta^E(e) \in C^\infty((E^\dagger)^* \otimes \iota_A(A))$, since δ^E is defined by projection to this subspace. Thus $\nabla \circ \iota_E(e) \in C^\infty(E \otimes_{\mathbb{H}} A)$, by Definition 4.1.4. ∎

In complex geometry, a holomorphic section v of a holomorphic vector bundle V is one whose covariant derivative ∇v takes values in the complex tensor product of V with the holomorphic cotangent space. In hypercomplex geometry, a q-holomorphic section e of a q-holomorphic vector bundle E is one whose covariant derivative ∇e takes values in the *quaternionic* tensor product of E with the *q-holomorphic* cotangent space.

6.4.3 Sections of Tensor Products

For real and complex vector bundles, there is a natural inclusion $C^\infty(M, E) \otimes C^\infty(N, F) \hookrightarrow C^\infty(M \times N, E \otimes F)$ given by $(e \otimes f)(m, n) = e(m) \otimes f(n)$. The quaternionic analogue of this map is more delicate. Let E and F be $\mathrm{A}\mathbb{H}$-bundles. We want to define an $\mathrm{A}\mathbb{H}$-morphism $\phi : C^\infty(M, E) \otimes_{\mathbb{H}} C^\infty(N, F) \to C^\infty(M \times N, E \otimes_{\mathbb{H}} F)$. The obvious difficulty is that for sections e and f of E and F respectively, there will not in general be a section $e \otimes_{\mathbb{H}} f$ of the $\mathrm{A}\mathbb{H}$-bundle $(E \otimes_{\mathbb{H}} F, M \times N)$. Instead we use a generalisation of Joyce's map $\phi : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_N \to \mathcal{P}_{M \times N}$ (Definition 6.1.6).

Consider first the fibres E_m and F_n for $m \in M$ and $n \in N$. Let $\alpha \in E_m^\dagger$ and $\beta \in F_n^\dagger$. Define an $\mathrm{A}\mathbb{H}$-morphism $\alpha_m : C^\infty(M, E) \to \mathbb{H}$ by $\alpha_m(e) = \alpha(e(m))$ for all $e \in C^\infty(M, E)$. Then $\alpha_m \in C^\infty(M, E)^\dagger$. Similarly, define the 'evaluation at $n \in N$', $\beta_n(f) = \beta(f(n))$. Then $\beta_n \in C^\infty(N, F)^\dagger$. The operators α_m and β_m are generalisations of Joyce's $\theta_m \in C^\infty(M, \mathbb{H})^\dagger$, introduced in Lemma 6.1.1. With $\mathbb{H}$-valued functions (sections of $M \times \mathbb{H}$), the evaluation map θ_m generates the whole of $\mathbb{H}_m^\dagger \cong \mathbb{R}$. For more general $\mathrm{A}\mathbb{H}$-bundles, we need to consider the combination of a point m at which to evaluate sections, *and* an $\mathrm{A}\mathbb{H}$-morphism $\alpha \in E_m^\dagger$, since we can no longer assume that the map $\alpha = \mathrm{id}$ generates the whole of $E_m^\dagger$.

For $\mathrm{A}\mathbb{H}$-modules U and V, recall the linear map $\lambda_{UV} : U^\dagger \otimes V^\dagger \to (U \otimes_{\mathbb{H}} V)^\dagger$ of Equation (6.2). Thus there are linear maps $\lambda_{E_m, F_n} : E_m^\dagger \otimes F_n^\dagger \to (E_m \otimes_{\mathbb{H}} F_n)^\dagger$ and $\lambda_{C^\infty(M,E), C^\infty(N,F)} : C^\infty(M, E)^\dagger \otimes C^\infty(N, F)^\dagger \to (C^\infty(M, E) \otimes_{\mathbb{H}} C^\infty(N, F))^\dagger$.

Definition 6.4.10 Let $\epsilon \in C^\infty(M, E) \otimes_{\mathbb{H}} C^\infty(N, F)$. Define $\phi_{E,F}(\epsilon)(m, n) \in E_m \otimes_{\mathbb{H}} F_n$ by the equation

$$\lambda_{E_m, F_n}(\alpha \otimes \beta) \cdot \phi_{E,F}(\epsilon)(m, n) = \lambda_{C^\infty(M,E), C^\infty(N,F)}(\alpha_m \otimes \beta_n) \cdot \epsilon$$

for all $m \in M$, $\alpha \in E_m^\dagger$ and for all $n \in N$, $\beta \in F_n^\dagger$.

Thus $\phi_{E,F}(\epsilon)$ defines a section of the AH-bundle $(E \otimes_{\mathbb{H}} F, M \times N)$. By the same arguments as in Definition 6.1.6, $\phi_{E,F}(\epsilon)$ is smooth because it is a finite sum of smooth sections, and so we have a linear map

$$\phi_{E,F} : C^\infty(M, E) \otimes_{\mathbb{H}} C^\infty(M, F) \to C^\infty(M \times N, E \otimes_{\mathbb{H}} F).$$

It is easy to see that $\phi_{E,F}$ is an injective A$\mathbb{H}$-morphism. We shall call $\phi_{E,F}$ the *section product map* for A$\mathbb{H}$-bundles.

If $e \in C^\infty(M, E)$ and $f \in C^\infty(N, F)$ satisfy the conditions of Lemma 4.1.7 (so if $\iota_E(e)$ and $\iota_F(f)$ both take values in $\mathbb{C}_q$ for some $q \in S^2$) then the section $\phi_{E,F}(e \otimes_{\mathbb{H}} f)$ is the complex product of the sections, so $\phi(e \otimes_{\mathbb{H}} f)(m, n) = e(m) \otimes_{\mathbb{H}} f(n) \cong e(m) \otimes_{\mathbb{C}_q} f(n)$. The section product map is the natural tool for relating tensor products of sections with sections of tensor products and allows us to treat sections of more complicated A$\mathbb{H}$-bundles using similar techniques to those used by Joyce for q-holomorphic functions.

It is well known that sums and tensor products of holomorphic sections are themselves holomorphic. The same is true for sums of q-holomorphic sections, and there is an analogous description for q-holomorphic sections of quaternionic tensor products, using the section product map $\phi_{E,F}$. This is a generalisation of the fact (Definition 6.1.6) that $\phi : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_N \to \mathcal{P}_{M \times N}$.

Theorem 6.4.11 *Let (E, π_1, M, ∇_E) and (F, π_2, N, ∇_F) be q-holomorphic* A$\mathbb{H}$*-bundles, and let $\mathcal{P}(M, E) \subset C^\infty(M, E)$ and $\mathcal{P}(N, F) \subset C^\infty(N, F)$ be the* A$\mathbb{H}$*-modules of their q-holomorphic sections. Then*

$$\phi_{E,F} : \mathcal{P}(M, E) \otimes_{\mathbb{H}} \mathcal{P}(N, F) \longrightarrow \mathcal{P}(M \times N, E \otimes_{\mathbb{H}} F),$$

where $E \otimes_{\mathbb{H}} F$ is equipped with the connection $\nabla_{E \otimes_{\mathbb{H}} F}$, so the section product map takes q-holomorphic sections to q-holomorphic sections.

Proof. Consider the anti-self-dual connection $\nabla_{E \otimes_{\mathbb{H}} F} = \mathrm{id} \otimes (\nabla_{E,M} \otimes \mathrm{id} + \mathrm{id} \otimes \nabla_{F,N})$. Then $\nabla_{E \otimes_{\mathbb{H}} F} = \delta^{E \otimes_{\mathbb{H}} F} + \bar{\delta}^{E \otimes_{\mathbb{H}} F}$. Using the natural splitting $T^*(M \times N) \cong T^*M \oplus T^*N$, we see that $\bar{\delta}^{E \otimes_{\mathbb{H}} F} = \bar{\delta}^{E,M} \otimes \mathrm{id} + \mathrm{id} \otimes \bar{\delta}^{F,N}$, where $\bar{\delta}^{E,M}$ differentiates sections of $\iota_E(E)$ in the M directions and then projects to the q-antiholomorphic cotangent space of $M \times N$, and similarly for $\bar{\delta}^{F,N}$.

Let $\epsilon \in \mathcal{P}(E) \otimes_{\mathbb{H}} \mathcal{P}(F)$, and let $\phi_{E,F}(\epsilon) \in C^\infty(M \times N, E \otimes_{\mathbb{H}} F)$, where $\phi_{E,F}$ is the section product map of Definition 6.4.10. Since $\epsilon \in \iota_{\mathcal{P}(E)}(\mathcal{P}(E)) \otimes (\mathcal{P}(F)^\dagger)^*$, it follows that $(\bar{\delta}^{E,M} \otimes \mathrm{id})(\phi_{E,F}(\epsilon)) = 0$. Similarly, $(\mathrm{id} \otimes \bar{\delta}^{F,N})(\phi_{E,F}(\epsilon)) = 0$. Hence $\bar{\delta}^{E \otimes_{\mathbb{H}} F}(\phi_{E,F}(\epsilon)) = 0$, and $\phi_{E,F}(\epsilon)$ is q-holomorphic. ∎

Just as for q-holomorphic functions, when $M = N$ we can restrict to the diagonal bundle $(E \otimes_{\mathbb{H}} F, M)$. The restriction map ρ clearly preserves q-holomorphic sections, and we obtain a natural product $\rho \circ \phi_{E,F} : \mathcal{P}(M, E) \otimes_{\mathbb{H}} \mathcal{P}(M, F) \to \mathcal{P}(M, E \otimes_{\mathbb{H}} F)$. In particular, let E be a q-holomorphic A$\mathbb{H}$-bundle over M, let $\mathcal{P}(E)$ be the A$\mathbb{H}$-module of q-holomorphic sections of E and let $\mathcal{P}_M = \mathcal{P}(\mathbb{H})$ be the H-algebra of q-holomorphic functions on M. Then by Theorem 6.4.11, we have a natural A$\mathbb{H}$-morphism

$$\rho \circ \phi_{\mathbb{H},E} : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}(E) \longrightarrow \mathcal{P}(\mathbb{H} \otimes_{\mathbb{H}} E) \cong \mathcal{P}(E). \tag{6.13}$$

We can descibe such an algebraic situation by saying that $\mathcal{P}(E)$ is an *H-algebra module* over $\mathcal{P}_M$. Here is the definition of an H-algebra module:

Axiom M. (*i*) $\mathcal{Q}$ is an A$\mathbb{H}$-module and $(\mathcal{P}, \mu_\mathcal{P})$ is an H-algebra.

(*ii*) There is an A$\mathbb{H}$-morphism $\mu_\mathcal{Q} : \mathcal{P} \otimes_\mathbb{H} \mathcal{Q} \to \mathcal{Q}$, called the *module multiplication map.*

(*iii*) The maps $\mu_\mathcal{P}$ and $\mu_\mathcal{Q}$ combine to give A$\mathbb{H}$-morphisms $\mu_\mathcal{P} \otimes_\mathbb{H} \mathrm{id}$ and $\mathrm{id} \otimes_\mathbb{H} \mu_\mathcal{Q} : \mathcal{P} \otimes_\mathbb{H} \mathcal{P} \otimes_\mathbb{H} \mathcal{Q} \to \mathcal{P} \otimes_\mathbb{H} \mathcal{Q}$. Composing with $\mu_\mathcal{Q}$ gives A$\mathbb{H}$-morphisms $\mu_\mathcal{Q} \circ (\mu_\mathcal{P} \otimes_\mathbb{H} \mathrm{id})$ and $\mu_\mathcal{Q} \circ (\mathrm{id} \otimes_\mathbb{H} \mu_\mathcal{Q}) : \mathcal{P} \otimes_\mathbb{H} \mathcal{P} \otimes_\mathbb{H} \mathcal{Q} \to \mathcal{Q}$. Then $\mu_\mathcal{Q} \circ (\mu_\mathcal{P} \otimes_\mathbb{H} \mathrm{id}) = \mu_\mathcal{Q} \circ (\mathrm{id} \otimes_\mathbb{H} \mu_\mathcal{Q})$. This is *associativity of module multiplication.*

(*iv*) For $u \in \mathcal{Q}$, $1 \otimes_\mathbb{H} u \in \mathcal{P} \otimes_\mathbb{H} \mathcal{Q}$ by Lemma 4.1.7. Then $\mu_\mathcal{Q}(1 \otimes_\mathbb{H} u) = u$ for all $u \in \mathcal{Q}$. Thus 1 acts as an identity on $\mathcal{Q}$.

Definition 6.4.12 $\mathcal{Q}$ is an *H-algebra module* if $\mathcal{Q}$ satisfies Axiom M.

The idea of an H-algebra module was suggested by Joyce, and the axioms follow a very similar pattern to those for an H-algebra (Definition 6.1.4). It is relatively easy to see that the q-holomorphic sections $\mathcal{P}(E)$ of a q-holomorphic A$\mathbb{H}$-bundle form an H-algebra module over the H-algebra $\mathcal{P}_M$, the module multiplication map being the section product map $\phi_{\mathbb{H},E}$ of Equation (6.13). The proof of this statement is obtained by adapting Joyce's proof that the q-holomorphic functions $\mathcal{P}_M$ form an H-algebra [J1, Theorem 5.5].

One possible application of this idea is to study anti-self-dual connections (instantons) on $\mathbb{H}$. Different connections will give rise to different H-algebra modules of q-holomorphic sections. This suggests that the theory of H-algebra modules over $\mathcal{P}_\mathbb{H}$ might lead to an algebraic description of instantons on $\mathbb{H}$.

Chapter 7

Quaternion Valued Forms and Vector Fields

This final chapter uses the algebra and geometry developed so far to describe quaternion-valued tensors on hypercomplex manifolds. On a hypercomplex manifold we have global complex structures, which allows us to adapt the real-valued double complex of Chapter 3 to decompose quaternion-valued forms. The splitting $\mathbb{H} \otimes T^*M \cong A \oplus B$ of the previous chapter is the first example of this type of decomposition. More generally, for all k and r we obtain a splitting of $\mathbb{H} \otimes E_{k,r}$ into two A$\mathbb{H}$-bundles. This decomposition gives rise to a double complex of quaternion-valued forms on hypercomplex manifolds, which has some advantages over the real-valued double complex. Not only the top row but the top two rows of the quaternion-valued double complex are elliptic, and in four dimensions the whole complex is elliptic. The top row of the quaternion-valued double complex is particularly well-behaved, and can be constructed using quaternionic algebra. This allows us to define q-holomorphic k-forms, and to describe their algebraic structure using ideas from the previous chapter.

A similar approach to quaternion-valued vector fields is also fruitful. Just as with the quaternionic cotangent space, there is a splitting of the quaternionic tangent space $\mathbb{H} \otimes TM$ using which we define a hypercomplex version of the '$(1,0)$ vector fields' on a complex manifold. These vector fields are closed under the Lie bracket (suitably adapted to the quaternionic situation). This encourages us to adapt the axioms for a Lie algebra to form a quaternionic version, in a similar way to that in which Joyce arrived at H-algebras. The Lie bracket on hypercomplex manifolds defines a natural operation on vector fields which satisfies these axioms.

In recent times, Spindel *et al.* [SSTP] and Joyce [J3] demonstrated the existence of invariant hypercomplex structures on certain compact Lie groups and their homogeneous spaces. This discovery presents us with lots of examples of finite-dimensional quaternionic Lie algebras. We use this idea to calculate the quaternionic cohomology groups of the group U(2), and suggest how these methods may be extended to higher-dimensional hypercomplex Lie groups.

7.1 The Quaternion-valued Double Complex

In Chapter 3, we saw how the real-valued exterior forms $\Lambda^k T^*M$ on a quaternionic manifold M are acted upon by the principal Sp(1)-bundle Q of local almost complex structures on M. Recall that the subbundle of $\Lambda^k T^*M$ consisting of V_r-type representations is denoted $E_{k,r} = \varepsilon^n_{k,r} V_r$.

Suppose in addition that M is hypercomplex, and so possesses global complex structures I_1, I_2 and $I_3 = I_1 I_2$ (in other words, Q is the trivial bundle $M \times \mathrm{Sp}(1)$). Instead of just real or complex forms (which we can think of as taking values in the trivial representation V_0), consider forms taking values in some Sp(1)-representation W, *i.e.* sections of the bundle $W \otimes \Lambda^k T^*M$. Since the Sp(1)-action on $\Lambda^k T^*M$ is now defined by *global* complex structures, we can take the diagonal action under which the subspace $W \otimes E_{k,r}$ splits according to the Clebsch-Gordon formula.

The situation in which we are particularly interested is that of forms taking values in the quaternions $\mathbb{H} = V_1^L \otimes V_1^R$. As suggested in Section 3.4, we obtain splittings by coupling the right Sp(1)-action on $\mathbb{H}$ with the Sp(1)-action on $\Lambda^k T^*M$. In symbols, this takes the form

$$\mathbb{H} \otimes E_{k,r} \cong V_1^L \otimes V_1^R \otimes \varepsilon^n_{k,r} V_r^G \cong \varepsilon^n_{k,r} V_1^L \otimes (V_{r+1}^{RG} \oplus V_{r-1}^{RG}). \tag{7.1}$$

Proposition 7.1.1 *Let M^{4n} be a hypercomplex manifold. The* $\mathrm{A}\mathbb{H}$*-bundle* $\mathbb{H} \otimes \Lambda^k T^*M$ *decomposes as*

$$\mathbb{H} \otimes \Lambda^k T^*M \cong V_1^L \otimes \left(\bigoplus_{r=0}^{k+1} (\varepsilon^n_{k,r+1} + \varepsilon^n_{k,r-1}) V_r^{RG} \right),$$

where $r \equiv k+1 \bmod 2$.

Proof. By Proposition 3.2.1, we have

$$\mathbb{H} \otimes \Lambda^k T^*M \cong V_1^L \otimes V_1^R \otimes \left(\bigoplus_{r=0}^{k} \varepsilon^n_{k,r} V_r^G \right),$$

where the 'Sp $(1)^G$-action' is the action of Sp(1) on $\Lambda^k T^*M$ induced by the hypercomplex structure. Taking the diagonal $\mathrm{Sp}(1)^{RG}$-action using the Clebsch-Gordon formula, this becomes

$$\mathbb{H} \otimes \Lambda^k T^*M \cong V_1^L \otimes \left(\bigoplus_{r=0}^{k} \varepsilon^n_{k,r} (V_{r+1}^{RG} \oplus V_{r-1}^{RG}) \right). \tag{7.2}$$

Collecting together the V_r representations yields the formula in the Proposition. ∎

Definition 7.1.2 Define $F_{k,r}$ to be the subspace $(\varepsilon^n_{k,r+2} + \varepsilon^n_{k,r}) V_1 \otimes V_{r+1} \subseteq \mathbb{H} \otimes \Lambda^k T^*M$.

The primed part of the space $F_{k,r}$ is obtained using the theory of Chapter 5. As in Equation 5.10, Equation 7.1 is a decomposition of $\mathbb{H} \otimes E_{k,r}$ into stable and antistable $\mathrm{A}\mathbb{H}$-modules. (This is a generalisation of the splitting $\mathbb{H} \otimes T^*M \cong A \oplus B$.) Defining $F'_{k,r} = F_{k,r} \cap (\mathbb{I} \otimes \Lambda^k T^*M)$, the space $F_{k,r}$ is an $\mathrm{A}\mathbb{H}$-subbundle of $\mathbb{H} \otimes (F^\dagger_{k,r})^* = \mathbb{H} \otimes$

$(E_{k,r+2} \oplus E_{k,r})$, and each (fibre of the) AH-bundle $F_{k,r}$ is the direct sum of stable and antistable components. The splitting

$$\mathbb{H} \otimes E_{k,r} \cong \varepsilon^n_{k,r} V^L_1 \otimes (V^{RG}_{r+1} \oplus V^{RG}_{r-1})$$

gives an $\mathbb{H}$-module isomorphism

$$\mathbb{H} \otimes E_{k,r} \cong \begin{cases} \varepsilon^n_{k,r} U_r \oplus \varepsilon^n_{k,r} U^\times_{r-2} & \text{for } r \text{ even} \\ \\ \frac{1}{2}\varepsilon^n_{k,r} U_r \oplus \frac{1}{2}\varepsilon^n_{k,r} U^\times_{r-2} & \text{for } r \text{ odd.} \end{cases} \tag{7.3}$$

As with the decomposition $\mathbb{H} \otimes T^*M \cong A \oplus B$, these are not AH-isomorphisms because some of the primed part is lost in the splitting. We give names to these spaces as follows (where as usual $a = 1$ if n is even and $a = 2$ if n is odd):

Definition 7.1.3 Define $F^\uparrow_{k,r}$ to be the AH-bundle $\frac{1}{a}\varepsilon^n_{k,r} U_r \subseteq \mathbb{H} \otimes E_{k,r}$. Define $F^\downarrow_{k,r}$ to be the AH-bundle $\frac{1}{a}\varepsilon^n_{k,r+2} U^\times_r \subseteq \mathbb{H} \otimes E_{k,r+2}$. Thus $F^\uparrow_{k,r}$ is the U_r-type subspace of $\mathbb{H} \otimes \Lambda^k T^*M$ and $F^\downarrow_{k,r}$ is the $U^\times_r$-type subspace of $\mathbb{H} \otimes \Lambda^k T^*M$, so that

$$F_{k,r} = F^\uparrow_{k,r} \oplus F^\downarrow_{k,r}.$$

The q-holomorphic cotangent space A is $F_{1,1} = F^\uparrow_{1,1}$ and the q-antiholomorphic cotangent space B is $F_{1,-1} = F^\downarrow_{1,-1}$.

With these definitions, we have $({F^\uparrow_{k,r}}^\dagger)^* = ({F^\downarrow_{k,r-2}}^\dagger)^* = E_{k,r}$, and $\mathbb{H} \otimes E_{k,r} \cong F^\uparrow_{k,r} \oplus F^\downarrow_{k,r-2}$. Just as with the splitting $\mathbb{H} \otimes T^*M \cong A \oplus B$, there is an injective AH-morphism

$$F^\uparrow_{k,r} \oplus F^\downarrow_{k,r-2} \hookrightarrow \mathbb{H} \otimes E_{k,r}$$

which is an $\mathbb{H}$-linear isomorphism of the total spaces but is not injective on the primed parts. A short calculation shows that

$$\begin{array}{lcl} \dim F^\uparrow_{k,r} = 2(r+2)\varepsilon^n_{k,r} & & \dim F^\downarrow_{k,r} = 2(r+2)\varepsilon^n_{k,r+2} \\ & \text{and} & \\ \dim {F^\uparrow_{k,r}}' = (r+3)\varepsilon^n_{k,r} & & \dim {F^\downarrow_{k,r}}' = (r+1)\varepsilon^n_{k,r+2}. \end{array}$$

Since we can consider the bundles $F^\uparrow_{k,r}$ and $F^\downarrow_{k,r}$ separately, it may appear strange to mix up the stable and antistable fibres in the single bundle $F_{k,r}$. However, it soon becomes clear that exterior differentiation does not necessarily map stable fibres to stable fibres or antistable fibres to antistable fibres, so in order to obtain a double complex it is necessary to amalgamate the stable and antistable contributions.

The definition of the operators δ and $\bar{\delta}$ of Section 6.2 can now be generalised to cover the whole of $\mathbb{H} \otimes \Lambda^\bullet T^*M$, giving rise to a double complex of quaternion-valued forms.

Definition 7.1.4 Let $\pi_{k,r}$ be the natural projection map $\pi_{k,r} : \mathbb{H} \otimes \Lambda^k T^*M \to F_{k,r}$. Define the differential operators $\delta, \bar{\delta} : \Omega^k(M, \mathbb{H}) \to \Omega^{k+1}(M, \mathbb{H})$ by

$$\begin{array}{lcl} \delta : C^\infty(F_{k,r}) \to C^\infty(F_{k+1,r+1}) & & \bar{\delta} : C^\infty(F_{k,r}) \to C^\infty(F_{k+1,r-1}) \\ & \text{and} & \\ \delta = \pi_{k+1,r+1} \circ d & & \bar{\delta} = \pi_{k+1,r-1} \circ d\,. \end{array}$$

Theorem 7.1.5 *The exterior derivative d maps $C^\infty(M, F_{k,r})$ to $C^\infty(M, F_{k+1,r+1} \oplus F_{k+1,r-1})$, so*

$$d = \delta + \bar{\delta}.$$

It follows that

$$\delta^2 = \delta\bar{\delta} + \bar{\delta}\delta = \bar{\delta}^2 = 0.$$

and the decomposition $\mathbb{H} \otimes \Lambda^k T^*M = \bigoplus_{r=0}^{k+1} F_{k,r}$ gives rise to a double complex.

Proof. The proof works in exactly the same way as that of Theorem 3.2.3. Let ∇ be the Obata connection on M, so $\nabla : C^\infty(M, F_{k,r}) \to C^\infty(M, F_{k,r} \otimes T^*M)$. Since

$$F_{k,r} \otimes T^*M = (\varepsilon^n_{k,r+2} + \varepsilon^n_{k,r})V_1 \otimes V_{r+1} \otimes 2nV_1,$$

it follows (from the Clebsch-Gordon splitting $V_{r+1} \otimes 2nV_1 \cong 2n(V_{r+2} \oplus V_r)$, followed by the antisymmetrisation $d = \wedge \circ \nabla$) that $d : C^\infty(M, F_{k,r}) \to C^\infty(M, F_{k+1,r+1} \oplus F_{k+1,r-1})$. The rest of the theorem follows automatically. ∎

Figure 7.1: The Quaternion-Valued Double Complex

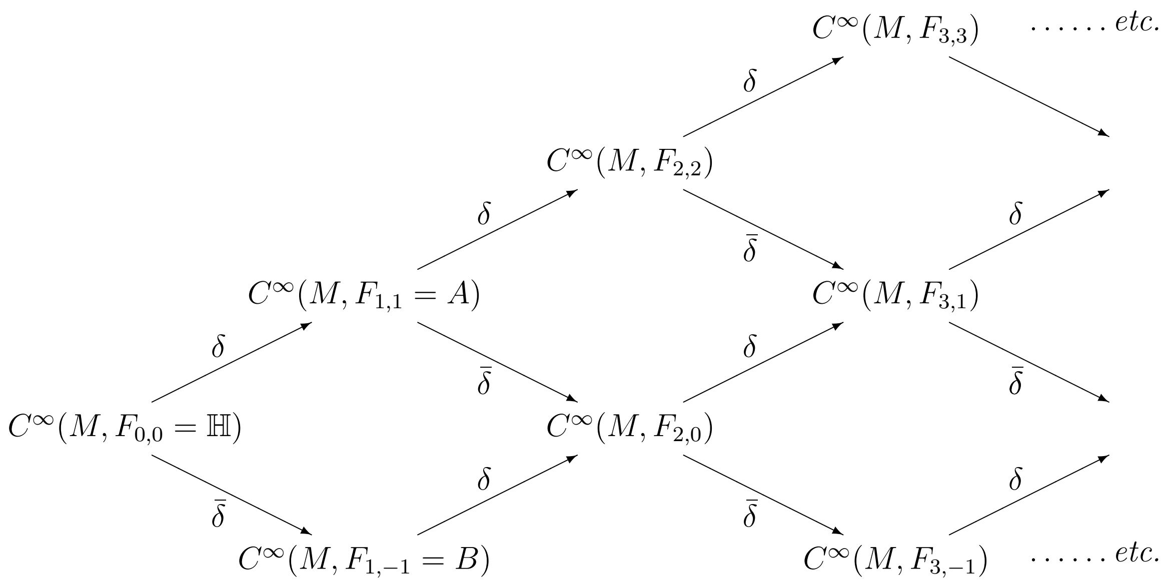

As already hinted, the operators δ and $\bar{\delta}$ do *not* preserve stable or antistable subspaces. For example, though we have $\bar{\delta} : F_{k,r} \to F_{k+1,r-1}$, it is not the case that $\bar{\delta} : F^{\uparrow}_{k,r} \to F^{\uparrow}_{k+1,r-1}$. This can be observed in the simplest of cases — a quaternion-valued function f is a section of $F_{0,0} = F^{\uparrow}_{0,0}$, and unless f is q-holomorphic $\bar{\delta} f$ is a nonzero section of $F^{\downarrow}_{1,-1}$.

Much of the theory from the real-valued double complex of Chapter 3 can be adapted to describe the quaternion-valued version as well. For example, the operators δ and $\bar{\delta}$ can be expressed in a similar fashion to $\mathcal{D}$ and $\overline{\mathcal{D}}$, using the Casimir element technique of Lemma 3.2.7. The Casimir operator in question is that of the diagonal Lie algebra action given by the operators

$$\mathcal{I}(\omega) = I_1(\omega) - \omega i_1 \qquad \mathcal{J}(\omega) = I_2(\omega) - \omega i_2 \qquad \mathcal{K}(\omega) = I_3(\omega) - \omega i_3.$$

of Equation (6.5). This leads to the following adaptation of Lemma 3.2.7:

Lemma 7.1.6 *Let* $\alpha \in C^\infty(F_{k,r})$. *Then*

$$\delta\alpha = -\frac{1}{4}\left(r + \frac{1}{r+2}(\mathcal{I}^2 + \mathcal{J}^2 + \mathcal{K}^2)\right) d\alpha$$

and

$$\bar{\delta}\alpha = \frac{1}{4}\left((r+4) + \frac{1}{r+2}(\mathcal{I}^2 + \mathcal{J}^2 + \mathcal{K}^2)\right) d\alpha.$$

The results on ellipticity in Section 3.3 can also be adapted to the new situation. We can infer that the operator δ is elliptic except at the bottom spaces $F_{2k-1,-1}$ and $F_{2k,0}$. Again, the operator δ *is* elliptic at some of these lowest-weight spaces for low exterior powers; in particular the leading edge of spaces $F_{k,k}$ forms a complex which is elliptic throughout. Closer examination also reveals that the 'second row' of spaces $F_{k,k-2}$ is also an elliptic complex with respect to δ.

Lemma 7.1.7 *The complex*

$$0 \longrightarrow C^\infty(F_{1,-1}) \xrightarrow{\delta} C^\infty(F_{2,0}) \xrightarrow{\delta} \ldots \xrightarrow{\delta} C^\infty(F_{2n+1,2n-1}) \xrightarrow{\delta} 0$$

is elliptic.

Proof. Ellipticity at $C^\infty(F_{k,k-2})$ for all $k \geq 3$ follows from a generalisation of the techniques used in the proof of Theorem 3.3.1. We need to show that the complex is elliptic at $C^\infty(F_{1,-1}) = C^\infty(B)$ and $C^\infty(F_{2,0})$. As in Section 3.3, it is enough to choose some $e^0 \in T^*M$ and show that the symbol sequence

$$0 \longrightarrow F_{1,-1} \xrightarrow{\sigma} F_{2,0} \xrightarrow{\sigma} F_{3,1} \xrightarrow{\sigma} \ldots \textit{etc.}$$

is exact, where $\sigma(\omega) = \sigma_\delta(\omega, e^0) = \pi_{k+1,r+1}(\omega e^0)$ for $\omega \in F_{k,r}$. (As usual ωe^0 means $\omega \wedge e^0$.)

Since $\delta = d$ on $F_{1,-1} = B$, we have $\sigma(\beta) = \beta e^0$ for all $\beta \in B$. It follows from the expression for B in Equation (6.9) that $\sigma : B \to F_{2,0}$ is injective, so the complex is elliptic at B.

Let $\omega \in F_{2,0}$. Then $\sigma(\omega) = 0$ if and only if $\omega e^0 \in F_{3,-1}$, which is the case if and only if $\mathcal{I}(\omega e^0) = \mathcal{J}(\omega e^0) = \mathcal{K}(\omega e^0) = 0$. The first of these equations is

$$I_1(\omega e^0) - \omega e^0 i_1 = I_1(\omega)e^0 + \omega e^1 - \omega e^0 i_1 = 0.$$

It follows by taking exterior product with e^0 that $\omega e^{10} = 0$. The same arguments for $\mathcal{J}$ and $\mathcal{K}$ show that

$$\sigma(\omega) = 0 \Longrightarrow \omega e^{10} = \omega e^{20} = \omega e^{30} = 0,$$

and so ωe^0 must be equal to zero and $\omega = \gamma e^0$ for some $\gamma \in \mathbb{H} \otimes T^*M$.

It remains to show that we can choose $\beta \in B$ such that $\omega = \beta e^0$. Suppose instead that $\omega = \alpha e^0$, with $\alpha \in A$. Since the complex is elliptic at A it follows that $\alpha e^0 \in F_{2,0}$ if and only if $\alpha = \sigma(q) = \frac{1}{4}q(3e^0 + e^1 i_1 + e^2 i_2 + e^3 i_3)$ for some $q \in \mathbb{H}$. But then $\alpha e^0 = \beta e^0$, where $\beta = \frac{1}{4}q(-e^0 + e^1 i_1 + e^2 i_2 + e^3 i_3) \in B$. Hence we can find $\beta \in B$ such that $\sigma(\omega) = 0$ implies that $\omega = \sigma(\beta)$ for all $\omega \in F_{2,0}$, so the complex is exact at $F_{2,0}$. This completes the proof. ∎

This is an unexpected bonus — on hypercomplex manifolds, the double complex of quaternion-valued forms has not just one but two rows which with respect to the operator δ are elliptic throughout, namely the top row $F_{k,k}$ and the second row $F_{k,k-2}$.

Example 7.1.8 Quaternion-valued forms in four dimensions

Figure 7.2: The Quaternion-Valued Double Complex in Four Dimensions

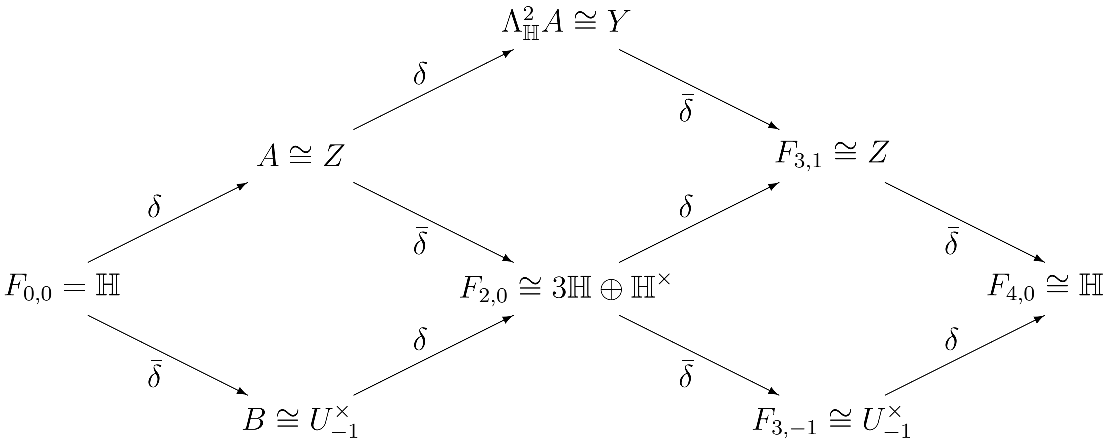

In four dimensions the situation is particularly friendly towards quaternion-valued forms, because the whole double complex is elliptic. To show this, choose a standard basis of 1-forms $\{e^0, \ldots, e^3\}$ for T^*M so that as usual $I_j(e^0) = e^j$. Explicitly, we have $\mathbb{H} \otimes T^*M \cong F_{1,1} \oplus F_{1,-1}$, where

$$F_{1,1} = A = \{q_0 e^0 + q_1 e^1 + q_2 e^2 + q_3 e^3 : q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3 = 0\},$$

$$F_{1,-1} = B = \{q_0 e^0 + q_1 e^1 + q_2 e^2 + q_3 e^3 : q_0 = q_1 i_1 = q_2 i_2 = q_3 i_3\}.$$

Next, $\mathbb{H} \otimes \Lambda^2 T^*M \cong F_{2,2} \oplus F_{2,0}$, where

$$F_{2,2} = \Lambda^2_{\mathbb{H}} A = \{q_1 e^{01+23} + q_2 e^{02+31} + q_3 e^{03+12} : q_1 i_1 + q_2 i_2 + q_3 i_3 = 0\},$$

$$\begin{aligned} F_{2,0} &= F^{\uparrow}_{2,0} \oplus F^{\downarrow}_{2,0} \\ &= \langle e^{01-23}, e^{02-31}, e^{03-12} \rangle_{\mathbb{H}} \oplus \{q_1 e^{01+23} + q_2 e^{02+31} + q_3 e^{03+12} : q_1 i_1 = q_2 i_2 = q_3 i_3\}. \end{aligned}$$

Lastly, $\mathbb{H} \otimes \Lambda^3 T^*M \cong F_{3,1} \oplus F_{3,-1}$, where

$$F_{3,1} = \{q_0 e^{123} + q_1 e^{032} + q_2 e^{013} + q_3 e^{021} : q_0 + q_1 i_1 + q_2 i_2 + q_3 i_3 = 0\},$$

$$F_{3,-1} = \{q_0 e^{123} + q_1 e^{032} + q_2 e^{013} + q_3 e^{021} : q_0 = q_1 i_1 = q_2 i_2 = q_3 i_3\}.$$

It follows that the symbol map σ_δ is an isomorphism between $F_{3,-1}$ and $F_{4,0}$, and so in four dimensions the entire quaternion-valued double complex is elliptic.

7.1.1 The Top Row $\Lambda^k_{\mathbb{H}}A$ and Q-holomorphic k-forms

In contrast with the the role of the Cauchy-Riemann operator $\overline{\partial}$ in the Dolbeault complex, the Cauchy-Riemann-Fueter operator $\bar{\delta}$ in the quaternion-valued double complex (Figure 7.1) does not begin a longer elliptic complex. [1] On the other hand, the operator δ on functions does extend to give the elliptic complex

$$0 \longrightarrow C^\infty(F_{0,0}) \xrightarrow{\delta} C^\infty(F_{1,1}) \xrightarrow{\delta} \dots \xrightarrow{\delta} C^\infty(F_{2n,2n}) \xrightarrow{\delta} 0.$$

This section discusses these top row spaces $F_{k,k}$ which are of particular interest.

In complex geometry the Hodge decomposition of forms results immediately from the isomorphism $\mathbb{C}\otimes T^*M \cong \Lambda^{1,0}\oplus\Lambda^{0,1}$ and the isomorphism in exterior algebra $\Lambda^k(U\oplus V) = \bigoplus_{p+q=k}\Lambda^pU\otimes\Lambda^qV$. With hypercomplex geometry we are not quite so lucky, since the splitting $A\oplus B \cong \mathbb{H}\otimes T^*M$ is not an A$\mathbb{H}$-isomorphism but rather an injective A$\mathbb{H}$-morphism which is not surjective on the primed parts. Despite the fact that we can identify $\mathbb{H}\otimes\Lambda^kT^*M$ with $\Lambda^k_{\mathbb{H}}(\mathbb{H}\otimes T^*M)$, the induced A$\mathbb{H}$-morphism

$$\iota_k : \Lambda^k_{\mathbb{H}}(A\oplus B) \hookrightarrow \mathbb{H}\otimes\Lambda^kT^*M$$

is therefore *not* an isomorphism for $k>1$. This is why we resort to our more complicated analysis of the Sp(1)-representation on $\mathbb{H}\otimes\Lambda^kT^*M$ to discover the quaternionic version of the Dolbeault complex. In spite of this, the inclusion map ι_k is still of special interest, because it describes explicitly the top row of the double complex.

Proposition 7.1.9 *Let A be the q-holomorphic cotangent space of a hypercomplex manifold M. The inclusion map $\iota_k : \Lambda^k_{\mathbb{H}}(A\oplus B) \hookrightarrow \mathbb{H}\otimes T^*M$ identifies $\Lambda^k_{\mathbb{H}}A$ with the highest space $F_{k,k}\subseteq \mathbb{H}\otimes\Lambda^kT^*M$.*

Proof. There is a natural identification $\Lambda^k_{\mathbb{H}}(A\oplus B) = \bigoplus_{p+q=k}(\Lambda^p_{\mathbb{H}}A\otimes_{\mathbb{H}}\Lambda^q_{\mathbb{H}}B)$, and so

$$\iota_k : \bigoplus_{p+q=k} \Lambda^p_{\mathbb{H}}A\otimes_{\mathbb{H}}\Lambda^q_{\mathbb{H}}B \hookrightarrow \mathbb{H}\otimes\Lambda^kT^*M.$$

Since $B' = \{0\}$, $B\otimes_{\mathbb{H}}B = \{0\}$ and so $\Lambda^k_{\mathbb{H}}B = \{0\}$ for $k>1$; also $A\otimes_{\mathbb{H}}B = \{0\}$. Thus $\Lambda^k_{\mathbb{H}}(A\oplus B) = \Lambda^k_{\mathbb{H}}A$ for $k>1$, and the map $\iota_k : \Lambda^k_{\mathbb{H}}(A\oplus B) \to \mathbb{H}\otimes\Lambda^kT^*M$ is none other than the normal inclusion map $\iota : \Lambda^k_{\mathbb{H}}A \to \mathbb{H}\otimes(\Lambda^k_{\mathbb{H}}A^\dagger)^*$. It follows that there is an A$\mathbb{H}$-submodule of $\mathbb{H}\otimes\Lambda^kT^*M$ which is isomorphic to $\Lambda^k_{\mathbb{H}}A$.

Now, $A\cong nU_1$, so by Theorem 5.2.1, $\Lambda^k_{\mathbb{H}}A \subseteq \bigotimes^k_{\mathbb{H}}A \cong mU_k$ for some m. (Recall that $U_k = aV_1\otimes V_{n+1}$ where $a=1$ for k even and $a=2$ for k odd.) By Proposition 4.1.16, $\dim\Lambda^k_{\mathbb{H}}(nU_1) = 2(k+2)\binom{2n}{k}$. It follows that

$$\Lambda^k_{\mathbb{H}}A \cong \frac{1}{a}\binom{2n}{k}U_k = \binom{2n}{k}V_1\otimes V_{k+1}. \tag{7.4}$$

Thus $\iota_k(\Lambda^k_{\mathbb{H}}A)$ must be a subspace of $\mathbb{H}\otimes\Lambda^kT^*M$ of this form.

[1] In 1991, Baston [Bas] succeeded in extending the Cauchy-Riemann-Fueter operator to a locally exact complex with a different construction involving second order operators.

From Definition 7.1.2, we see that $F_{k,k} = \binom{2n}{k} V_1 \otimes V_{k+1}$ as well. Since there are no other representations of this weight in $\mathbb{H} \otimes \Lambda^k T^*M$, it follows that

$$\iota_k(\Lambda^k_{\mathbb{H}} A) = F_{k,k},$$

so there is a natural isomorphism $\iota_k : \Lambda^k_{\mathbb{H}} A \cong F_{k,k}$. ∎

Just as we define the q-holomorphic cotangent space A as being a particular submodule of $\mathbb{H} \otimes T^*M$, so that $(A^\dagger)^* = T^*M$, we identify its exterior powers with the corresponding A$\mathbb{H}$-submodules of $\mathbb{H} \otimes \Lambda^k T^*M$. Thus we will omit to write the 'ι_k', writing $\Lambda^k_{\mathbb{H}} A = F_{k,k}$ and $(\Lambda^k_{\mathbb{H}} A^\dagger)^* = E_{k,k}$.

In Section 6.2.1 we showed that the q-holomorphic cotangent space A is generated over $\mathbb{H}$ by the various holomorphic cotangent spaces (Corollary 6.2.7). We generalise this result to higher exterior powers as follows:

Theorem 7.1.10 [2] *Let A be the q-holomorphic cotangent space of a hypercomplex manifold M^{4n}. Then*

$$(\mathrm{id} \otimes_{\mathbb{H}} \chi_q)(\Lambda^k_{\mathbb{H}} A \otimes_{\mathbb{H}} X_q) = \mathbb{H} \cdot \iota_q(\Lambda^{k,0}_Q).$$

It follows that

$$\Lambda^k_{\mathbb{H}} A = \sum_{Q \in S^2} \mathbb{H} \cdot \iota_q(\Lambda^{k,0}_Q),$$

i.e. the space $\Lambda^k_{\mathbb{H}} A$ is generated over $\mathbb{H}$ by the spaces of $(k,0)$-forms $\Lambda^{k,0}_Q$.

Proof. Since $\Lambda^k_{\mathbb{H}} A \cong \binom{2n}{k} V_1 \otimes V_{k+1}$, it follows that $\Lambda^k_{\mathbb{H}} A \otimes_{\mathbb{H}} X_q \cong \binom{2n}{k} X_q$. This is A$\mathbb{H}$-isomorphic to the submodule $(\mathrm{id} \otimes_{\mathbb{H}} \chi_q)(\Lambda^k_{\mathbb{H}} A \otimes_{\mathbb{H}} X_q) \subset \Lambda^k_{\mathbb{H}} A$, which by Theorem 5.3.2 is generated over $\mathbb{H}$ by the weight-spaces of Q with extremal weight.

Consider the weights of the action of $Q \in S^2 \subset \mathfrak{sp}(1)$. The highest-weight vectors in $\mathbb{C} \otimes \Lambda^k T^*M$ are the $(k,0)$-forms $\Lambda^{k,0}_Q$, which are mapped to $\mathbb{C}_q \otimes \Lambda^k T^*M$ by the map ι_q. Thus $\mathbb{H} \cdot \iota_q(\Lambda^{k,0}_Q) \subset (\mathrm{id} \otimes_{\mathbb{H}} \chi_q)(\Lambda^k_{\mathbb{H}} A \otimes_{\mathbb{H}} X_q)$, and since $\dim_{\mathbb{C}} \Lambda^{k,0}_Q = \binom{2n}{k}$ we see that

$$\mathbb{H} \cdot \iota_q(\Lambda^{k,0}_Q) = (\mathrm{id} \otimes_{\mathbb{H}} \chi_q)(\Lambda^k_{\mathbb{H}} A \otimes_{\mathbb{H}} X_q),$$

proving the first part of the Theorem.

Since $\Lambda^k_{\mathbb{H}} A$ is stable, it is generated by these subspaces. The result follows. ∎

Just as the q-holomorphic cotangent space A is our quaternionic analogue of the holomorphic cotangent space $T^*_{1,0}M$, the A$\mathbb{H}$-bundle $\Lambda^k_{\mathbb{H}} A$ is the quaternionic analogue of the bundle of $(k,0)$-forms $\Lambda^{k,0} = \Lambda^k_{\mathbb{C}} T^*_{1,0}M$. Both are formed in the same way using exterior algebra over their respective fields, and both form the 'top row' of their double complexes. This suggests a natural definition of a 'q-holomorphic k-form':

Definition 7.1.11 Let M be a hypercomplex manifold. A quaternion-valued k-form $\omega \in \Omega^k(M, \mathbb{H})$ is *q-holomorphic* if and only if $\omega \in C^\infty(M, \Lambda^k_{\mathbb{H}} A)$ and $\bar{\delta}\omega = 0$. The A$\mathbb{H}$-module of q-holomorphic k-forms on M will be written $\mathcal{P}^k_M$, so the H-algebra of q-holomorphic functions on M is $\mathcal{P}^0_M = \mathcal{P}_M$.

[2] This theorem is really a quaternionic version of (Salamon's) Equation 2.10, which is essentially the same result for complexified k-forms.

This gives rise to what we may call the *q-holomorphic de Rham complex*

$$0 \longrightarrow \mathcal{P}_M \xrightarrow{d=\delta} \mathcal{P}_M^1 \xrightarrow{d=\delta} \mathcal{P}_M^2 \xrightarrow{d=\delta} \dots \xrightarrow{d=\delta} \mathcal{P}_M^{2n-1} \xrightarrow{d=\delta} \mathcal{P}_M^{2n} \xrightarrow{d=\delta} 0. \qquad (7.5)$$

Just as (on a complex manifold) $\Lambda^{k,0}$ is a holomorphic vector bundle, it is easy to see that $(\Lambda^k_{\mathbb{H}}A, \nabla)$ is a q-holomorphic A$\mathbb{H}$-bundle where ∇ is (the connection induced by) the Obata connection on M. The q-holomorphic k-forms $\mathcal{P}_M^k$ are precisely the q-holomorphic sections $\mathcal{P}(\Lambda^k_{\mathbb{H}}A)$ of the A$\mathbb{H}$-bundle $(\Lambda^k_{\mathbb{H}}A, \nabla)$ as introduced in Definition 6.4.7, as the following Proposition demonstrates:

Proposition 7.1.12 *A quaternion-valued k-form $\omega \in C^\infty(M, \Lambda^k_{\mathbb{H}}A)$ satisfies the equation $\bar{\delta}\omega = 0$ if and only if $\nabla\omega \in C^\infty(\Lambda^k_{\mathbb{H}}A \otimes_{\mathbb{H}} A)$, i.e. ω is a q-holomorphic section of $(\Lambda^k_{\mathbb{H}}A, \nabla)$.*

Proof. The 'if' part is automatic, since $\nabla\omega \in C^\infty(\Lambda^k_{\mathbb{H}}A \otimes_{\mathbb{H}} A)$ implies that $d\omega = \wedge \circ \nabla\omega \in C^\infty(\Lambda^{k+1}_{\mathbb{H}}A)$ and so $\bar{\delta}\omega = 0$.

The reverse implication is nontrivial and depends upon analysing the Sp(1)-representation on $\Lambda^k_{\mathbb{H}}A \otimes T^*M$. Using Equation (7.4), we have

$$\begin{aligned} \Lambda^k_{\mathbb{H}}A \otimes T^*M &\cong \binom{2n}{k} V_1^L \otimes V_{k+1}^M \otimes 2nV_1^G \\ &\cong 2n\binom{2n}{k} V_1^L \otimes (V_{k+2}^{MG} \oplus V_k^{MG}). \end{aligned}$$

The bundle $\Lambda^k_{\mathbb{H}}A \otimes_{\mathbb{H}} A$ is precisely the higher-weight subspace $2n\binom{2n}{k}V_1^L \otimes V_{k+2}^{GH}$. The complementary subspace $2n\binom{2n}{k}V_1^L \otimes V_k^{GH}$ is revealed by Proposition 7.1.1 to be none other than the bundle $F_{k+1,k-1} \subset \mathbb{H} \otimes \Lambda^{k+1}T^*M$. The component of $\nabla\omega$ taking values in $F_{k+1,k-1}$ is of course $\bar{\delta}\omega$, the vanishing of which therefore guarantees that $\nabla\omega \in C^\infty(\Lambda^k_{\mathbb{H}}A \otimes_{\mathbb{H}} A)$. ∎

The q-holomorphic de Rham complex therefore inherits a a rich and interesting algebraic structure. We have already noted that $d : \mathcal{P}_M^k \to \mathcal{P}_M^{k+1}$. It follows from the theory of q-holomorphic sections that q-holomorphic forms are closed under the tensor product. Explicitly, let

$$\phi_{\Lambda^k_{\mathbb{H}}A, \Lambda^l_{\mathbb{H}}A} : C^\infty(M, \Lambda^k_{\mathbb{H}}A) \otimes_{\mathbb{H}} C^\infty(M, \Lambda^l_{\mathbb{H}}A) \longrightarrow C^\infty(M \times M, \Lambda^k_{\mathbb{H}}A \otimes_{\mathbb{H}} \Lambda^l_{\mathbb{H}}A)$$

be the section product map (Definition 6.4.10). Define $\phi_{k,l}$ to be the restriction to $C^\infty(M, \Lambda^{k+l}_{\mathbb{H}}A)$, brought about by the restriction $\rho : M \times M \to M$ to the diagonal submanifold $M_{\text{diag}} \subset M \times M$ followed by the skewing map $\wedge : C^\infty(M, \Lambda^k_{\mathbb{H}}A \otimes_{\mathbb{H}} \Lambda^l_{\mathbb{H}}A) \to C^\infty(M, \Lambda^{k+l}_{\mathbb{H}}A)$. It follows (from Theorem 6.4.11 and Proposition 7.1.12) that

$$\phi_{k,l} : \mathcal{P}_M^k \otimes_{\mathbb{H}} \mathcal{P}_M^l \longrightarrow \mathcal{P}_M^{k+l}.$$

Thus the H-algebra structure on the q-holomorphic functions $\mathcal{P}_M$ extends to one on the q-holomorphic k-forms $\mathcal{P}_M^\bullet$, and we say that $\mathcal{P}_M^\bullet$ forms a *differential graded H-algebra.* It is well known that on a real or complex manifold M one can use exterior products over $\mathbb{R}$ or $\mathbb{C}$ to give an algebraic structure to de Rham or Dolbeault cohomology, which is often called the *cohomology algebra* of M. It may be that the differential graded H-algebra structure on $\mathcal{P}_M^\bullet$ can be used to give a similar description of the quaternionic cohomology of a hypercomplex manifold.

7.2 Vector Fields and Quaternionic Lie Algebras

Quaternionic algebra can also be used to describe vector fields on a hypercomplex manifold M. Using similar ideas to those of Section 6.2, we define a splitting of the quaternionic tangent space $\mathbb{H} \otimes TM \cong \widehat{A} \oplus \widehat{B}$. Quaternion-valued vector fields which take values in the subspace $\widehat{A} \subset \mathbb{H} \otimes TM$ turn out to be a good hypercomplex analogue of the '$(1,0)$ vector fields' in complex geometry. In particular, an almost hypercomplex structure is integrable if and only if these vector fields are closed under a quaternionic version of the Lie bracket operator. This encourages us to treat these vector fields as a *quaternionic Lie algebra*, a new concept which we proceed to define and explore.

7.2.1 Vector Fields on Hypercomplex Manifolds

As so often, we take our inspiration from complex geometry. An almost complex structure I on a manifold M defines a splitting $\mathbb{C} \otimes TM \cong T^{1,0}M \oplus T^{0,1}M$, where $T^{1,0}M$ is the holomorphic and $T^{0,1}M$ the antiholomorphic tangent space of M. Let $\mathcal{V}$ be the set of smooth vector fields on M. The Lie bracket is a bilinear map from $\mathcal{V} \times \mathcal{V}$ to $\mathcal{V}$. Let $\mathcal{V}^{1,0} = C^\infty(M, T^{1,0}M)$ be the vector fields of type $(1,0)$ on M. The almost complex structure I is integrable if and only if the Lie bracket preserves $(1,0)$ vector fields, which is expressed by the inclusion

$$[\mathcal{V}^{1,0}, \mathcal{V}^{1,0}] \subseteq \mathcal{V}^{1,0}. \tag{7.6}$$

The obstruction to this equation is the Nijenhuis tensor N_I which measures the $(0,1)$ component of the Lie bracket of two $(1,0)$ vector fields.

We present a similar theorem for quaternionic vector fields on hypercomplex manifolds. Let M^{4n} be a hypercomplex manifold. We define a splitting of the quaternionic tangent space $\mathbb{H} \otimes TM$, which is roughly dual to the splitting $\mathbb{H} \otimes T^*M \cong A \oplus B$ of Section 6.2.

Definition 7.2.1 Let M^{4n} be a hypercomplex manifold so that $TM \cong 2nV_1$ as an Sp(1)-representation. The quaternionic tangent space $\mathbb{H} \otimes TM$ splits according to the equation

$$\mathbb{H} \otimes TM \cong V_1^L \otimes V_1^R \otimes 2nV_1^G \cong 2nV_1^L \otimes (V_2^{RG} \otimes V_0^{RG}).$$

Define the A$\mathbb{H}$-subbundles $\widehat{A} = 2nV_1^L \otimes V_2^{RG}$ and $\widehat{B} = 2nV_1^L \otimes V_0^{RG}$, using the natural definitions $\widehat{A}' = \widehat{A} \cap (\mathbb{I} \otimes TM)$ and $\widehat{B}' = \widehat{B} \cap (\mathbb{I} \otimes TM) = \{0\}$. Then $\widehat{A}$ is the *q-holomorphic tangent bundle* and $\widehat{B}$ is the *q-antiholomorphic tangent bundle* of M.

This definition is perfectly natural, though not quite ideal from the point of view of quaternionic algebra. We would like $\widehat{A}$ and $\widehat{B}$ to be dual to the cotangent spaces A and B. However, whilst there are $\mathbb{H}$-linear bundle isomorphisms $\widehat{A} \cong A^\times$ and $\widehat{B} \cong B^\times$, these spaces are not isomorphic as A$\mathbb{H}$-bundles, nor is there any fruitful way to alter the definitions to make them so.

The Lie bracket of vector fields is a bilinear map $[\ ,\ ] : \mathcal{V} \times \mathcal{V} \to \mathcal{V}$, and so defines a natural linear map $\lambda : \mathcal{V} \otimes \mathcal{V} \to \mathcal{V}$. As with H-algebras, we will find it much more meaningful to talk about linear maps on tensor products, and we can use this formulation to obtain a quaternionic analogue of Equation (7.6).

Definition 7.2.2 Let M be a hypercomplex manifold with q-holomorphic tangent space $\widehat{A} \subset \mathbb{H} \otimes TM$. Let $\mathcal{V} = C^\infty(M, TM)$ be the vector space of smooth real vector fields on M.

Define $\mathcal{V}_{\mathbb{H}} = \mathbb{H} \otimes \mathcal{V}$ to be the A$\mathbb{H}$-module of smooth quaternion-valued vector fields on M, so $\mathcal{V}_{\mathbb{H}} = C^\infty(M, \mathbb{H} \otimes TM)$. Define $\mathcal{V}_A = C^\infty(M, \widehat{A})$ to be the A$\mathbb{H}$-submodule of $\mathbb{H}$-valued vector fields on M taking values in $\widehat{A} \subset \mathbb{H} \otimes TM$. We will refer to elements of $\mathcal{V}_A$ as *A-type vector fields.*

The relationship between these vector fields and the map $\lambda : \mathcal{V} \otimes \mathcal{V} \to \mathcal{V}$ is particularly interesting. Using the canonical isomorphism $(\mathbb{H} \otimes TM) \otimes_{\mathbb{H}} (\mathbb{H} \otimes TM) \cong \mathbb{H} \otimes TM \otimes TM$, the Lie bracket defines an A$\mathbb{H}$-morphism (which we shall also call λ)

$$\lambda : \mathcal{V}_{\mathbb{H}} \otimes_{\mathbb{H}} \mathcal{V}_{\mathbb{H}} \to \mathcal{V}_{\mathbb{H}},$$

which is effectively a Lie bracket operation on quaternionic vector fields.

Here is the quaternionic version of Equation (7.6). The formulation and proof is similar in spirit to Theorem 6.4.11.

Theorem 7.2.3 *Let M be a hypercomplex manifold with q-holomorphic tangent space $\widehat{A}$, and let $\mathcal{V}_A$ denote the space of A-type vector fields on M. Then the Lie bracket λ on quaternionic vector fields preserves $\mathcal{V}_A$, so that*

$$\lambda : \mathcal{V}_A \otimes_{\mathbb{H}} \mathcal{V}_A \to \mathcal{V}_A.$$

Proof. To begin with, we use the section product map $\phi_{\widehat{A},\widehat{A}}$ (Definition 6.4.10) to regard elements of $\mathcal{V}_A \otimes_{\mathbb{H}} \mathcal{V}_A$ as sections in $C^\infty(\widehat{A} \otimes_{\mathbb{H}} \widehat{A}) \subset C^\infty(\mathbb{H} \otimes TM \otimes TM)$.

Let $v_j, w_j \in \mathcal{V}$ be vector fields such that $\sum_j q_j \otimes v_j \otimes w_j \in C^\infty(\widehat{A} \otimes_{\mathbb{H}} \widehat{A})$, which means that $\sum q_j \otimes v_j, \sum q_j \otimes w_j \in \mathcal{V}_A$. We want to find an expression for $\lambda(\sum q_j \otimes v_j \otimes w_j) = \sum q_j \otimes [v_j, w_j]$.

The Obata connection ∇ is an A$\mathbb{H}$-connection on $\widehat{A}$ (and in fact is anti-self-dual, so that $\widehat{A}$ is a q-holomorphic A$\mathbb{H}$-bundle). Thus $\nabla(\sum q_j \otimes w_j)$ is an element of $C^\infty(\widehat{A} \otimes T^*M)$. Contracting the T^*M-factor with $v_j \in TM$, it follows that

$$\sum q_j \otimes \nabla_{v_j} w_j \in \mathcal{V}_A,$$

and similarly

$$\sum q_j \otimes \nabla_{w_j} v_j \in \mathcal{V}_A.$$

Since ∇ is torsion-free, the Lie bracket $[v_j, w_j]$ is given by the difference $\nabla_{v_j} w_j - \nabla_{w_j} v_j$. It follows immediately that

$$\lambda\left(\sum q_j \otimes v_j \otimes w_j\right) = \sum q_j \otimes (\nabla_{v_j} w_j - \nabla_{w_j} v_j) \in \mathcal{V}_A,$$

proving the theorem. ∎

The reason why the hypercomplex structure must be integrable to obtain this result is that the integrability of I, J, and K ensures that the connection ∇ with $\nabla I = \nabla J = \nabla K = 0$ is torsion-free, and otherwise we would not have $[v, w] = \nabla_v w - \nabla_w v$.

Another way to understand this result is in terms of the decomposition of tensors with respect to different complex structures. Just as in Theorem 7.1.10, the bundle $\widehat{A}$ is generated over $\mathbb{H}$ by the tensors of type $(1,0)$ with respect to the different complex structures, and the bundle $\widehat{A} \otimes_{\mathbb{H}} \widehat{A}$ is generated by the tensors of type $(2,0)$. In other words, we have

$$\widehat{A} = \sum_{Q \in S^2} \mathbb{H} \cdot \iota_q(T_Q^{1,0} M) \qquad \text{and} \qquad \widehat{A} \otimes_{\mathbb{H}} \widehat{A} = \sum_{Q \in S^2} \mathbb{H} \cdot \iota_q(T_Q^{1,0} M \otimes T_Q^{1,0} M).$$

If every complex structure $Q \in S^2$ is integrable, we have

$$[\mathcal{V}_Q^{1,0}, \mathcal{V}_Q^{1,0}] \subseteq \mathcal{V}_Q^{1,0}$$

for all $Q \in S^2$, where $\mathcal{V}_Q^{1,0}$ denotes the vector fields which are of type $(1,0)$ with respect to the complex structure Q. Since the fibres of $\widehat{A}$ are stable, it follows that the Lie bracket must map sections of $\widehat{A} \otimes_{\mathbb{H}} \widehat{A}$ to sections of $\widehat{A}$.

7.2.2 Quaternionic Lie Algebras

The result of the previous section encourages us to think of the vector fields $\mathcal{V}_A$ as a *quaternionic Lie algebra* with respect to the Lie bracket map λ. We describe this idea as an abstract algebraic structure, and see how it fits in with some of Joyce's other algebraic structures over the quaternions.

As always, we do not talk about bilinear maps, but rather about linear maps on tensor products. A quaternionic Lie algebra will be an A$\mathbb{H}$-module A together with an A$\mathbb{H}$-morphism $\lambda : A \otimes_{\mathbb{H}} A \to A$, whose properties reflect those of a Lie bracket on a real or complex vector space: namely antisymmetry and the Jacobi identity. Here are the axioms for a quaternionic Lie algebra:

Axiom QL. *(i)* A is an A$\mathbb{H}$-module and there is an A$\mathbb{H}$-morphism $\lambda = \lambda_A : A \otimes_{\mathbb{H}} A \to A$ called the *Lie bracket*.

(ii) $S^2_{\mathbb{H}} A \subset \ker \lambda$. Thus λ is *antisymmetric*.

(iii) The composition $\lambda \circ (\mathrm{id} \otimes_{\mathbb{H}} \lambda)$ defines an A$\mathbb{H}$-morphism from $A \otimes_{\mathbb{H}} A \otimes_{\mathbb{H}} A$ to A such that $\Lambda^3_{\mathbb{H}} A \subset \ker(\lambda \circ (\mathrm{id} \otimes_{\mathbb{H}} \lambda))$. This is the *Jacobi identity* for λ.

Definition 7.2.4 The pair (A, λ) is a *quaternionic Lie algebra* if it satisfies Axiom QL.

We will often refer to A itself as a quaternionic Lie algebra when the map λ is understood. Axiom QL(iii) is probably the least familiar-looking of these axioms. By way of explanation, let (V, λ) be a real or complex Lie algebra. Then in terms of tensor products the Jacobi identity is

$$\lambda \circ (\mathrm{id} \otimes \lambda)(x \otimes y \otimes z + y \otimes z \otimes x + z \otimes x \otimes y) = 0.$$

Since we can identify $x \otimes y - y \otimes x$ with $x \wedge y$, we see that the Jacobi identity is equivalent to the condition that $\lambda \circ (\mathrm{id} \otimes \lambda)(x \wedge y \wedge z) = 0$, so that $\Lambda^3 V \subset \ker(\lambda \circ (\mathrm{id} \otimes \lambda))$.

Example 7.2.5 Let M be a hypercomplex manifold. The space of quaternion-valued vector fields $\mathcal{V}_{\mathbb{H}}$ is a quaternionic Lie algebra with respect to the Lie bracket operator $\lambda : \mathcal{V}_{\mathbb{H}} \otimes_{\mathbb{H}} \mathcal{V}_{\mathbb{H}} \to \mathcal{V}_{\mathbb{H}}$. Axiom QL follows from the corresponding identities satisfied by the Lie bracket on real vector fields.

Theorem 7.2.3 now shows that the A-type vector fields $\mathcal{V}_A$ form a quaternionic Lie subalgebra of $\mathcal{V}_{\mathbb{H}}$. This is the quaternionic version of the well-known fact that on a complex manifold, the $(1,0)$ vector fields form a complex Lie subalgebra of the complex vector fields.

Joyce has already considered the notion of a Lie algebra over the quaternions, but in a different fashion. In [J1, §6], he writes down axioms called Axiom L and Axiom P, which define quaternionic analogues of Lie algebras and Poisson algebras: but instead of a Lie bracket $\lambda : A \otimes_{\mathbb{H}} A \to A$, Joyce's quaternionic Lie bracket is a map $\xi : A \otimes_{\mathbb{H}} A \to A \otimes_{\mathbb{H}} Y$, where $Y \cong U_2$ is the A$\mathbb{H}$-module of Example 4.1.2.

The reason for this is that Joyce's main purpose is to describe the algebraic structure of q-holomorphic functions on hyperkähler manifolds. Since a hyperkähler manifold M has three independent symplectic forms, two functions f and g have three different Poisson brackets and the H-algebra $\mathcal{P}_M$ of q-holomorphic functions on M has three independent Poisson structures. Using the fact that $\mathbb{I} \cong \mathbb{R}^3$, Joyce describes the three Poisson brackets using a single map $\xi : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_M \to \mathcal{P}_M \otimes \mathbb{I}$. Recalling that $(Y^\dagger)^* \cong V_2 = \mathbb{I}$, we have a map $\xi : \mathcal{P}_M \otimes_{\mathbb{H}} \mathcal{P}_M \to \iota_{\mathcal{P}_M}(\mathcal{P}_M) \otimes (Y^\dagger)^*$ which is in fact an A$\mathbb{H}$-morphism whose image is contained in $\mathcal{P}_M \otimes_{\mathbb{H}} Y$. This is why, for Joyce, quaternionic Lie algebras and Poisson algebras are defined by an A$\mathbb{H}$-morphism $\xi : A \otimes_{\mathbb{H}} A \to A \otimes_{\mathbb{H}} Y$.

We can relate these two algebraic ideas very simply by choosing an A$\mathbb{H}$-morphism $\eta : Y \to \mathbb{H}$. The space of such A$\mathbb{H}$-morphisms is of course $Y^\dagger \cong V_2$. Then for every A$\mathbb{H}$-module A there is a map $\mathrm{id} \otimes_{\mathbb{H}} \eta : A \otimes_{\mathbb{H}} Y \to A \otimes_{\mathbb{H}} \mathbb{H} \cong A$. Suppose that $\xi : A \otimes_{\mathbb{H}} A \to A \otimes_{\mathbb{H}} Y$ satisfies Joyce's Axiom L. Define a Lie bracket $\lambda : A \otimes_{\mathbb{H}} A \to A$ by setting $\lambda = (\mathrm{id} \otimes_{\mathbb{H}} \eta) \circ \xi$. It follows immediately that the pair (A, λ) is a quaternionic Lie algebra in the sense of Axiom QL.

To go in the opposite direction, suppose that (A, λ) is a quaternionic Lie algebra, and consider the A$\mathbb{H}$-module $A \otimes_{\mathbb{H}} Y$. In order to define an HL-algebra we need to form a map from $(A \otimes_{\mathbb{H}} Y) \otimes_{\mathbb{H}} (A \otimes_{\mathbb{H}} Y)$ to $(A \otimes_{\mathbb{H}} Y) \otimes_{\mathbb{H}} Y$. Let τ be the natural isomorphism $\tau : A \otimes_{\mathbb{H}} Y \otimes_{\mathbb{H}} A \otimes_{\mathbb{H}} Y \to A \otimes_{\mathbb{H}} A \otimes_{\mathbb{H}} Y \otimes_{\mathbb{H}} Y$ which interchanges the second and third factors. Define an A$\mathbb{H}$-morphism

$$\xi = (\lambda \otimes_{\mathbb{H}} \mathrm{id} \otimes_{\mathbb{H}} \mathrm{id}) \circ \tau : (A \otimes_{\mathbb{H}} Y) \otimes_{\mathbb{H}} (A \otimes_{\mathbb{H}} Y) \to (A \otimes_{\mathbb{H}} Y) \otimes_{\mathbb{H}} Y.$$

Then ξ satisfies Joyce's Axiom L, and the pair $(A \otimes_{\mathbb{H}} Y, \xi)$ forms an HL-algebra. Precisely which HL-algebras may be obtained from quaternionic Lie algebras and vice versa using these constructions remains open to question.

7.3 Hypercomplex Lie groups

As a final example, we consider hypercomplex structures on compact Lie groups, and show how these give rise to finite-dimensional quaternionic Lie algebras. Since the early

1950s, the mathematical world has been aware that every compact Lie group of even dimension is a homogeneous complex manifold. [3] This was first announced by Samelson, whose proof is an extension of Borel's celebrated result that the quotient of a compact Lie group by its maximal torus is a homogeneous complex manifold. In 1988 and 1992 respectively, Spindel *et al.* [SSTP] and Joyce [J3] demonstrated independently that these results extend to hypercomplex geometry. Joyce's approach also gives hypercomplex structures on more general homogeneous spaces. It relies on being able to decompose the Lie algebra of a compact Lie group as follows:

Lemma 7.3.1 *[J3, Lemma 4.1] Let G be a compact Lie group, with Lie algebra $\mathfrak{g}$. Then $\mathfrak{g}$ can be decomposed as*

$$\mathfrak{g} = \mathfrak{b} \oplus \sum_{k=1}^{n} \mathfrak{d}_k \oplus \sum_{k=1}^{n} \mathfrak{f}_k, \tag{7.7}$$

where $\mathfrak{b}$ is abelian, $\mathfrak{d}_k$ is a subalgebra of $\mathfrak{g}$ isomorphic to $\mathfrak{su}(2)$, $\mathfrak{b}+\sum_k \mathfrak{d}_k$ contains the Lie algebra of a maximal torus of G, and $\mathfrak{f}_1, \ldots, \mathfrak{f}_n$ are (possibly empty) vector subspaces of $\mathfrak{g}$, such that for each $k = 1, 2, \ldots, n$, $\mathfrak{f}_k$ satisfies the following two conditions:

(i) $[\mathfrak{d}_l, \mathfrak{f}_k] = \{0\}$ *whenever $l < k$, and*

(ii) $\mathfrak{f}_k$ *is closed under the Lie bracket with $\mathfrak{d}_k$, and the Lie bracket action of $\mathfrak{d}_k$ on $\mathfrak{f}_k$ is isomorphic to the sum of m copies of the basic representation V_1 of $\mathfrak{su}(2)$ on $\mathbb{C}^2$, for some integer m.*

By adding an additional p (with $0 \leq p \leq \mathrm{Max}(3, \mathrm{rk}\, G)$) copies of the abelian Lie algebra $\mathfrak{u}(1)$ to $\mathfrak{b}$ if necessary, this decomposition allows us to define a hypercomplex structure on $p\mathfrak{u}(1) \oplus \mathfrak{g}$ as follows. For a 1-dimensional subspace $\mathfrak{b}_k$ of $p\mathfrak{u}(1) \oplus \mathfrak{b}$ there is an isomorphism $\mathfrak{b}_k \oplus \mathfrak{d}_k \cong \mathbb{R} \oplus \mathbb{I} = \mathbb{H}$. The action of $\mathfrak{d}_k$ on $\mathfrak{f}_k$ gives an isomorphism $\mathfrak{f}_k \cong \mathbb{H}^m$. This gives an isomorphism $p\mathfrak{u}(1) \oplus \mathfrak{g} \cong \mathbb{H}^l$, in other words a hypercomplex structure. Normally there are many choices to be made in such an isomorphism, which give rise to distinct hypercomplex structures. That such a hypercomplex structure on the vector space $p\mathfrak{u}(1) \oplus \mathfrak{g}$ defines an integrable hypercomplex structure on the manifold $\mathrm{U}(1)^p \times G$ follows from Samelson's original work on homogeneous complex manifolds. This leads to the following result:

Theorem 7.3.2 *[J3, Theorem 4.2] Let G be a compact Lie group. Then there exists an integer p with $0 \leq p \leq \mathrm{Max}(3, \mathrm{rk}\, G)$ such that $\mathrm{U}(1)^p \times G$ admits a left-invariant homogeneous hypercomplex structure.*

There is a strong link between these hypercomplex structures and the quaternionic Lie algebras of the previous section. Suppose that $\mathfrak{g}$ is any real Lie algebra. Then $\mathfrak{g}_{\mathbb{H}} \equiv \mathbb{H} \otimes \mathfrak{g}$ is a quaternionic Lie algebra because Axiom QL is obviously satisfied. The quaternionic Lie algebra structure of $\mathfrak{g}$ is much more interesting when $\mathfrak{g}$ admits a hypercomplex structure as described above. In this case we define the subspace

$$\mathfrak{g}_A = \{v_0 + i_1 v_1 + i_2 v_2 + i_3 v_3 : v_0 + I_1 v_1 + I_2 v_2 + I_3 v_3 = 0\} \subset \mathfrak{g}_{\mathbb{H}}$$

which we shall call the space of *A-type elements* of $\mathfrak{g}_{\mathbb{H}}$. Lie groups with integrable left-invariant hypercomplex structures then give rise to interesting quaternionic Lie algebras.

[3] Note that not all 'Lie groups possessing a complex structure' are complex Lie groups, because their multiplication and inverse maps might not be holomorphic.

Corollary 7.3.3 *Let G be a hypercomplex Lie group with hypercomplex structure (I_1, I_2, I_3), and let $\mathfrak{g}_A \subset \mathfrak{g}_{\mathbb{H}}$ be the subspace of A-type elements of $\mathfrak{g}_{\mathbb{H}}$. Then $\mathfrak{g}_A$ is a quaternionic Lie subalgebra of $\mathfrak{g}_{\mathbb{H}}$.*

Proof. By Theorem 7.2.3, the set of A-type vector fields on G is closed under the Lie bracket operator λ. Since λ also preserves the left-invariant vector fields $\mathfrak{g}$, it preserves $\mathfrak{g}_A$. ∎

In this way, Joyce's hypercomplex structures on compact Lie groups give rise to many interesting finite-dimensional quaternionic Lie algebras. We can also begin to calculate the quaternionic cohomology groups of these manifolds. On a Lie group G the exterior differential d on G-invariant k-forms can be expressed as a formal differential $d : \Lambda^k\mathfrak{g}^* \to \Lambda^{k+1}\mathfrak{g}^*$. The map d is induced by the dual of the Lie bracket, which is a linear map $\lambda : \mathfrak{g} \otimes \mathfrak{g} \to \mathfrak{g}$. Since λ is antisymmetric it is effectively a map $\lambda : \Lambda^2\mathfrak{g} \to \mathfrak{g}$, so its dual is the map $d = \lambda^* : \mathfrak{g}^* \to \Lambda^2\mathfrak{g}^*$. In practice, this means that

$$d\omega(u, v) = \omega([u, v]) \qquad\qquad \omega \in \mathfrak{g}^* \quad v, w \in \mathfrak{g}.$$

The map d extends to a unique antiderivation on $\Lambda^\bullet\mathfrak{g}^*$ in the usual fashion. Questions about the cohomology of G as a real, complex or hypercomplex manifold can then be rephrased in terms of the cohomology of the complex $(\Lambda^\bullet\mathfrak{g}^*, d)$. This is true at least for G-invariant forms, which accounts for all the de Rham cohomology since it is known that every de Rham cohomology class has a G-invariant representative. This is less obvious for Dolbeault and quaternionic cohomology groups, whose theory is possibly more subtle for this reason.

Example 7.3.4 Let $M = \mathrm{U}(2)$ be the unitary group in two dimensions, *i.e.* the subgroup of $\mathrm{GL}(2, \mathbb{C})$ preserving the standard hermitian metric on $\mathbb{C}^2$. It is well-known that U(2) is isomorphic to $\mathrm{U}(1) \times_{\mathbb{Z}_2} \mathrm{SU}(2)$, which is diffeomorphic to the Hopf surface $S^1 \times S^3$.

The Lie algebra $\mathfrak{u}(2)$ appears naturally in the form of Equation (7.7), thanks to the decomposition $\mathfrak{u}(2) = \mathfrak{u}(1) \oplus \mathfrak{su}(2)$. Let $\mathfrak{u}(1) = \langle e_0\rangle$ and $\mathfrak{su}(2) = \langle e_1, e_2, e_3\rangle$, so that

$$[e_0, e_j] = 0 \qquad \text{and} \qquad [e_i, e_j] = \varepsilon_{ijk}e_k$$

where $i, j, k \subset \{1, 2, 3\}$. Let $\{e^\alpha\}$ be the dual basis for $\mathfrak{g}^*$. It follows that $de^0 = 0$ and $de^i = e^j \wedge e^k$, where $\{i, j, k\}$ is an even permutation of $\{1, 2, 3\}$. Thus e^0 and e^{123} both generate de Rham cohomology classes, and we have $b^0(M) = b^1(M) = b^3(M) = b^4(M) = 1$, $b^2(M) = 0$. This complex is best described using the structure of $\mathfrak{u}(2)$ as a representation τ of $\mathfrak{su}(2) = \langle e_1, e_2, e_3\rangle$. This subgroup acts on $\mathfrak{u}(2)$ via the adjoint representation, so that $\tau_v(w) = [v, w]$ for $v \in \mathfrak{su}(2)$, $w \in u(2)$. Since $[v, e^0] = 0$, e^0 generates a copy of the trivial representation V_0 and $\langle e^1, e^2, e^3\rangle$ is just the adjoint representation of $\mathfrak{su}(2)$, which is V_2. Thus $T^*M \cong V_0 \oplus V_2$, and this induces a representation of $\mathfrak{su}(2)$ on $\Lambda^\bullet T^*M$ by the usual Leibniz rule $\tau_v(w_1 \wedge w_2) = [v, w_1] \wedge w_2 + w_1 \wedge [v, w_2]$. This gives the following decompositions:

$$\begin{array}{rcl} T^*M = & \langle e^0\rangle \oplus \langle e^1, e^2, e^3\rangle & \cong V_0 \oplus V_2 \\ \Lambda^2T^*M = & \langle e^{23}, e^{31}, e^{12}\rangle \oplus \langle e^{01}, e^{02}, e^{03}\rangle & \cong V_2 \oplus V_2 \\ \Lambda^3T^*M = & \langle e^{032}, e^{013}, e^{021}\rangle \oplus \langle e^{123}\rangle & \cong V_2 \oplus V_0 \\ \Lambda^4T^*M = & \langle e^{0123}\rangle & \cong V_0. \end{array} \tag{7.8}$$

The benefit of this approach is that the map d is $\mathfrak{su}(2)$-equivariant. Thus once we have shown that $de^1 = e^{23}$, it follows immediately that d gives an $\mathfrak{su}(2)$-equivariant isomorphism

$$d : \langle e^1, e^2, e^3 \rangle \to \langle e^{23}, e^{31}, e^{12} \rangle.$$

This allows us to read off cohomological information, including the quaternionic cohomology groups $H^{k,r}_{\mathcal{D}}(M)$ and $H^{k,r}_{\delta}(M)$. To do this we compare the decomposition of Equation (7.8) with the decomposition of $\Lambda^k T^*M$ induced by the hypercomplex structure. In other words, we have two representations of $\mathfrak{sp}(1)$ on T^*M. The first is the representation $T^*M \cong 2V_1$ defined by the hypercomplex structure, and the second is the representation $T^*M \cong V_0 \oplus V_2$ given by the adjoint action of the subalgebra $\mathfrak{su}(2) \subset \mathfrak{u}(2)$. The quaternionic cohomology of M is defined using the first action, but the d operator is best described using the second. By collating the two we obtain a complete picture of the situation.

In four dimensions, $E_{2,2}$ is the space of self-dual 2-forms $\langle e^{01+23}, e^{02+31}, e^{03+12} \rangle$, where $e^{ab+cd} = e^{ab} + e^{cd}$. The image under d of T^*M is $\langle e^{01}, e^{02}, e^{03} \rangle$, and the projection $\pi_{2,2}$ is an $\mathfrak{su}(2)$-equivariant surjection onto $E_{2,2}$. The sequence

$$0 \longrightarrow E_{0,0} \longrightarrow E_{1,1} \longrightarrow E_{2,2} \longrightarrow 0$$

is therefore exact at $E_{2,2}$, and we obtain the standard self-dual cohomology of $S^1 \times S^3$,

$$H^{0,0}_{\mathcal{D}}(M) = \mathbb{R} \qquad H^{1,1}_{\mathcal{D}}(M) = \langle e^0 \rangle \cong \mathbb{R} \qquad H^{2,2}_{\mathcal{D}}(M) = 0. \tag{7.9}$$

Moving to quaternion-valued forms, it is easy to show that non-zero δ-cohomology occurs only at $F_{0,0}$, $F_{1,1}$, $F_{3,-1}$ and $F_{4,0}$, the sequence $0 \to B \to F_{2,0} \to F_{3,1} \to 0$ being exact. Explicitly, we have

$$\begin{array}{rclrcl} H^{0,0}_{\delta}(M) & = & \mathbb{H} & H^{1,1}_{\delta}(M) & = & \langle 3e^0 + i_1 e^1 + i_2 e^2 + i_3 e^3 \rangle_{\mathbb{H}} \cong U^{\times}_{-1} \\ H^{4,4}_{\delta}(M) & = & \langle e^{0123} \rangle_{\mathbb{H}} \cong \mathbb{H} & H^{3,-1}_{\delta}(M) & = & \langle e^{123} - i_1 e^{032} - i_2 e^{013} - i_3 e^{021} \rangle_{\mathbb{H}} \cong U^{\times}_{-1}. \end{array} \tag{7.10}$$

It follows from the property of ellipticity that the cohomology sequences should be exact, as indeed they are. However, they are clearly not A$\mathbb{H}$-exact, because they are not exact on the primed parts. The quaternionic algebra of such phenonomena might be interesting and merit closer study.

It would be desirable to extend this work to higher-dimensional groups and homogeneous spaces. The group SU(3) provides an interesting case. A similar analysis to that above may provide the correct results, though because of the additional four dimensions this option is difficult and complicated. The principle would nonetheless be very much the same, and essentially involves comparing different $\mathfrak{sp}(1)$-actions, one defined by Joyce's hypercomplex structure and the other arising from the adjoint representation of a particular subalgebra $\mathfrak{su}(2) \subset \mathfrak{su}(3)$. For example, it is easy to demonstrate that $H^{1,1}_{\delta}(\mathrm{SU}(3)) \neq 0$. Quaternionic cohomology is in this case certainly not a subset of de Rham cohomology, since $b^1(\mathrm{SU}(3)) = 0$.

A logical precursor to developing the quaternionic cohomology theory of hypercomplex Lie groups would be to understand thoroughly the Dolbeault cohomology of homogeneous complex manifolds. Surprisingly, this theory seems to be lacking or at best extremely obscure. The theory for complex groups and homogeneous spaces *is* documented in Bott's important paper [Bot, §2], but the link between theory and results in this paper is quite opaque and certainly contains some mistakes. Pittie's paper [P] uses a simpler bicomplex than Bott's, conjecturing that the cohomology of these complexes is the same. Certain hypercomplex nilmanifolds discussed by Dotti and Fino [DF] might also provide fruitful rseults. This would be an interesting area for future research.

Bibliography

[AHS] M.F. Atiyah, N.J. Hitchin, I.M. Singer: *Self-duality in four-dimensional Riemannian geometry*, Proc R. Soc. Lond. A. 362 (1978), pp. 425-61.

[Bas] R. J. Baston: *Quaternionic complexes*, J. Geom. Phys. 8 (1992), pp. 29-52.

[Ben] D. Benson: *Representations and cohomology I: Basic representation theory of finite groups and associative algebras*, Cambridge University Press, Cambridge (1991).

[Bes] A.L. Besse: *Einstein Manifolds*, Springer-Verlag (1987).

[Bon] E. Bonan: *Sur l'algèbre extérieure d'une variété presque hermitienne quaternionique*, C. R. Acad. Sc. Paris 295 (1982), pp. 115-118 and 296 (1983), pp. 601-602.

[Bot] R. Bott: *Homogeneous Vector Bundles*, Ann. of Math. 66 no. 2 (1957), pp. 203-248.

[BD] T. Bröcker, T. tom Dieck: *Representations of Compact Lie Groups*, Springer-Verlag, Graduate Texts in Math. 98 (1985).

[D] C.A. Deavours: *The Quaternion Calculus*, Amer. Math. Monthly 80 (1973), pp. 995-1008.

[DF] I. G. Dotti, A. Fino: *Abelian Hypercomplex 8-dimensional nilmanifolds*, Ann. Global Anal. Geom. 18 (2000), pp. 47-59.

[EKR] Ebbinghaus *et al*: *Numbers*, Springer-Verlag, Graduate Texts in Math. 123 (1990) (in particular M. Koecher, R. Remmert: Chapter 7, *Hamilton's Quaternions*).

[F] A. Fujiki: *On the de Rham cohomology group of a compact Kähler symplectic manifold*, Adv. Stud. Pure Math. 10, 1987 (Algebraic Geometry, Sendai, 1985), pp. 105-165.

[FH] W. Fulton, J. Harris: *Representation Theory*, Springer-Verlag, Graduate Texts in Math. 129 (1991).

[GH] P. Griffiths, J. Harris: *Principles of Algebraic Geometry*, John Wiley and Sons (1978).

[GP] K. Galicki, Y. S. Poon: *Duality and Yang-Mills fields on quaternionic Kähler manifolds*, J. Math. Phys. 32 (1991), pp. 1263-1268.

[H1] Sir William Rowan Hamilton: Letter to the Rev. Archibald Hamilton, August 5th 1865. Hamilton's Math. Papers III p. 15.

[H2] Sir William Rowan Hamilton: Letter to John T.Graves, October 17th, 1843. London, Edinburgh and Dublin Philosophical Magazine and Journal of Science 25 (1844), pp. 189-95.

[H3] Sir William Rowan Hamilton: *On a new Species of Imaginary Quantities connected with a theory of Quaternions*, Proc. Royal Irish Acad. 2 (1843), pp. 424-434.

[H4] Sir William Rowan Hamilton: *On Quaternions*, Proc. Royal Irish Acad. 3 (1847), pp. 1-16.

[HKLR] N.J. Hitchin, A. Karlhede, U. Lindström, M. Roček: *Hyperkähler metrics and supersymmetry*, Commun. Math. Phys. 108 (1987), pp. 535-589.

[J1] D. Joyce: *Hypercomplex Algebraic Geometry*, Quart. J. Math. Oxford, 49 (1998), pp. 129-162.

[J2] D. Joyce: *The Hypercomplex quotient and the Quaternionic quotient*, Math. Ann. 290 (1991), pp. 323-340.

[J3] D. Joyce: *Compact Hypercomplex and Quaternionic manifolds*, J. Diff. Geom. 35 (1992), pp. 743-761.

[K] S. Kobayashi: *Differential Geometry of Complex Vector Bundles*, Princeton University Press (1987).

[Kr] V. Y. Kraines: *Topology of Quaternionic Manifolds*, Trans. Amer. Math. Soc. 122 (1966) pp. 357-367.

[MS] M. Mamone Capria, S.M. Salamon: *Yang-Mills fields on quaternionic spaces*, Nonlinearity 1 (1988), pp. 517-530.

[N] T. Nitta: *Vector bundles over Quaternionic Kähler manifolds*, Tohoku Math. J. 40 (1988), pp. 425-440.

[P] H. V. Pittie: *The nondegeneration of the Hodge–de Rham spectral sequence*, Bull. AMS 10 (1989), pp. 19-22.

[Q] D. Quillen: *Quaternionic Algebra and Sheaves on the Riemann Sphere*, Quart. J. Math. Oxford 49 (1998), pp. 163-198.

[R] R. Reyes-Carrión: A generalisation of the notion of instanton, Diff. Geom. and its Apps. 8 (1998), pp. 1-20.

[S1] S. Salamon: *Topics in four-dimensional Riemannian geometry*, Geometry seminar "Luigi Bianchi", Springer-Verlag lecture notes in math. 1022 (1982).

[S2] S. Salamon: *Quaternionic Kähler Manifolds*, Invent. Math. 67 (1982), pp. 143-171.

[S3] S. Salamon: *Differential Geometry of Quaternionic Manifolds*, Ann. Scient. Éc. Norm. Sup. Paris. series 4 no. 19 (1986), pp. 31-55.

[S4] S. Salamon: *Riemannian geometry and holonomy groups*, Pitman Res. Notes in Math. 201, Longman, Harlow (1989).

[SSTP] Ph. Spindel, A. Sevrin, W. Troost, A. van Proeyen: *Extended supersymmetric σ-models on group manifolds*, Nuclear Phys. B 308 (1988), pp. 662-698.

[dL] S. de Leo: *Quaternions and Special Relativity*, J. Math. Phys. 37.3 (1996), pp. 2955-2968.

[Su] A. Sudbery: *Quaternionic Analysis*, Math. Proc. Camb. Phil. Soc. 85 (1979), pp. 199-225.

[Sw] A. Swann: *Quaternion Kähler Geometry and the Fundamental 4-form*, Proc. Curvature Geom. workshop, University of Lancaster (1989), pp. 165-173.

[Sy] J.L. Synge: *Quaternions, Lorentz Transformations and the Conway-Dirac-Eddington Matrices*, Comm. Dublin Inst. for advanced studies, series A no. 21 (1972).

[V] M. Verbitsky: *Hyperholomorphic bundles over a Hyperkähler manifold*, J. Alg. Geom. 5 (1996), pp. 633-669.

[v.d.W] B.L. van der Waerden: *A History of Algebra*, Springer-Verlag (1985).

[WW] R.S. Ward, R.O. Wells: *Twistor Geometry and Field Theory*, Cambridge Monographs on Math. Phys. (1990).

[W] R.O Wells: *Differential Analysis on Complex Manifolds*, Springer-Verlag, Graduate Texts in Math. 65 (1980).