Overdetermined systems, conformal differential geometry, and the BGG complex

Andreas Čap

University of Vienna
Faculty of Mathematics

IMA, July 2006

- Based on joint work with T. Branson, M. Eastwood, and A.R. Gover
- procedure for rewriting certain overdetermined systems in first order closed form
- works for symbols of geometric origin associated to various geometric structures
- we will restrict to the version for Riemannian manifolds
- comes from a method for constructing conformally invariant differential operators, which generalizes to parabolic geometries

Structure

- basic Riemannian geometry is closely related to representation theory of $O(n)$
- standard representation corresponds to the (co)tangent bundle
- use representation theory to organize symmetries

strategy

- Embed $O(n)$ into a larger group $G \cong O(n+1,1)$ and analyze representations of G from the point of view of this subgroup.
- In the example, we will deal with the standard representation of G.

Consider $\mathbb{V} = \mathbb{R}^{n+2}$ with the inner product

$$\left\langle \begin{pmatrix} x_0 \\ \vdots \\ x_{n+1} \end{pmatrix}, \begin{pmatrix} y_0 \\ \vdots \\ y_{n+1} \end{pmatrix} \right\rangle := x_0 y_{n+1} + x_{n+1} y_0 + \sum_{i=1}^{n} x_i y_i$$

- Basis vectors $e_1, \dots, e_n$ span a standard Euclidean $\mathbb{R}^n$
- The remaining two basis vectors span an $\mathbb{R}^2$ with a $(1,1)$–metric and light-cone–coordinates
- $\langle \ , \ \rangle$ has signature $(n+1,1)$ and hence $G := O(\mathbb{V})$ is isomorphic to $O(n+1,1)$.
- Mapping A to $\begin{pmatrix} 1 & 0 & 0 \\ 0 & A & 0 \\ 0 & 0 & 1 \end{pmatrix}$ defines an inclusion $O(n) \hookrightarrow G$.

As a representation of $O(n)$, we have

$$\mathbb{V} = \mathbb{V}_0 \oplus \mathbb{V}_1 \oplus \mathbb{V}_2 \cong \mathbb{R} \oplus \mathbb{R}^n \oplus \mathbb{R}$$

with $\mathbb{V}_0$ spanned by e_{n+1} and $\mathbb{V}_2$ spanned by e_0.

Notation: column vectors with $\mathbb{V}_2$ on top.

- For $k = 0, 1, 2$ consider $\Lambda^k \mathbb{R}^n \otimes \mathbb{V}$.
- This splits as $\Lambda^k \mathbb{R}^n \otimes \mathbb{V} = \oplus_i (\Lambda^k \mathbb{R}^n \otimes \mathbb{V}_i)$ but $\Lambda^k \mathbb{R}^n \otimes \mathbb{V}_1$ admits a finer decomposition.
- As an example, for $k = 1$ we get $\mathbb{R}^n \otimes \mathbb{R}^n = \mathbb{R} \oplus S^2_0 \mathbb{R}^n \oplus \Lambda^2 \mathbb{R}^n$.

Doing the decompositions for $k = 0, 1, 2$ we notice coincidences:

$$\begin{array}{ccc} \mathbb{R} & \mathbb{R}^n & \Lambda^2\mathbb{R}^n \\ \mathbb{R}^n & \mathbb{R} \oplus S^2_0\mathbb{R}^n \oplus \Lambda^2\mathbb{R}^n & \mathbb{R}^n \oplus W_2 \oplus \Lambda^3\mathbb{R}^n \\ \mathbb{R} & \mathbb{R}^n & \Lambda^2\mathbb{R}^n \end{array}$$

Assigning homogeneity $k + i$ to elements of $\Lambda^k\mathbb{R}^n \otimes \mathbb{V}_i$ we identify components of the same homogeneity.

Use these identifications to define $\partial : \mathbb{V} \to \mathbb{R}^n \otimes \mathbb{V}$ as well as $\delta^* : \Lambda^k\mathbb{R}^n \otimes \mathbb{V} \to \Lambda^{k-1}\mathbb{R}^n \otimes \mathbb{V}$ for $k = 0, 1$ such that $\delta^* \circ \delta^* = 0$. Explicit formulae in terms of bundles below.

Let (M, g) be a Riemannian manifold of dimension n.

- $\mathbb{V}$ gives rise to a vector bundle $V = V_0 \oplus V_1 \oplus V_2 \to M$.
- $\Lambda^k \mathbb{R}^n \otimes \mathbb{V}$ corresponds to $\Lambda^k T^*M \otimes V$.
- sections of $\Lambda^k T^*M \otimes V$ are triples consisting of two k–forms and one T^*M–valued k–form. We use subscripts $0, 1, 2$ to indicate components.

The maps ∂ and δ^*

Using abstract indices, we define $\partial : V \to T^*M \otimes V$, and $\delta^* : \Lambda^k T^*M \otimes V \to \Lambda^{k-1} T^*M \otimes V$ for $k = 0, 1$ by

$$\partial \begin{pmatrix} h \\ \phi_j \\ f \end{pmatrix} = \begin{pmatrix} 0 \\ hg_{ij} \\ -\phi_i \end{pmatrix}, \delta^* \begin{pmatrix} h_j \\ \phi_{jk} \\ f_j \end{pmatrix} = \begin{pmatrix} \frac{1}{n}\phi_j^j \\ -f_i \\ 0 \end{pmatrix}, \delta^* \begin{pmatrix} h_{ij} \\ \phi_{ijk} \\ f_{ij} \end{pmatrix} = \begin{pmatrix} \frac{-1}{n-1}\phi_{ik}{}^k \\ \frac{1}{2} f_{ij} \\ 0 \end{pmatrix}$$

Let ∇ be the component–wise Levi–Civita connection on V. Note that while ∂ and δ^* preserve homogeneities, ∇ raises homogeneity by one.

The basic system

- Define $\tilde{\nabla}$ on V by $\tilde{\nabla}\Sigma := \nabla\Sigma + \partial\Sigma$
- Choose a bundle map $A : V_0 \oplus V_1 \to S^2_0 T^*M$, and view it as $A : V \to T^*M \otimes V$.
- Consider the first order system

$$\tilde{\nabla}\Sigma + A(\Sigma) = \delta^*\psi \qquad \text{for some } \psi \in \Omega^2(M, V)$$

The core of the method is to equivalently rewrite this in two ways, once as a higher order system on $\Sigma_0 \in \Gamma(V_0)$ and once as a first order closed system.

First rewrite in terms of the V_0–component Σ_0. By construction A has values in $\ker(\delta^*)$ and $\delta^* \circ \delta^* = 0$, so $\tilde{\nabla}\Sigma + A(\Sigma) = \delta^*\psi$ implies $\delta^*\tilde{\nabla}\Sigma = 0$.

The operator $\delta^*\tilde{\nabla}$ on $\Gamma(V)$ is given by

$$\begin{pmatrix} h \\ \phi_j \\ f \end{pmatrix} \overset{\tilde{\nabla}_i}{\mapsto} \begin{pmatrix} \nabla_i h \\ \nabla_i \phi_j + h g_{ij} \\ \nabla_i f - \phi_i \end{pmatrix} \overset{\delta^*}{\mapsto} \begin{pmatrix} \frac{1}{n}\nabla^j \phi_j + h \\ -\nabla_j f + \phi_j \\ 0 \end{pmatrix}$$

It is evident, how to solve this:

- Choose $f \in \Gamma(V_0)$ arbitrarily
- put $\phi_i = \nabla_i f$, i.e. $\phi = df$
- put $h = -\frac{1}{n}\nabla^i \nabla_i f = -\frac{1}{n}\Delta f$, where Δ is the Laplacian

Proposition (splitting operator in degree zero)

Given $f \in \Gamma(V_0)$ there is a unique $\Sigma \in \Gamma(V)$ such that $\Sigma_0 = f$ and $\delta^*\tilde{\nabla}\Sigma = 0$. Mapping f to Σ defines a linear second order differential operator $L : \Gamma(V_0) \to \Gamma(V)$, which is explicitly given by

$$L(f) = \begin{pmatrix} -\frac{1}{n}\Delta f \\ \nabla_i f \\ f \end{pmatrix} = \sum_{\ell=0}^{2} (-1)^\ell (\delta^*\nabla)^\ell \begin{pmatrix} 0 \\ 0 \\ f \end{pmatrix}$$

Observe that the components of $L(f)$ in V_0 and V_1 are f and ∇f, respectively. Hence

For any $A : V_0 \oplus V_1 \to S^2_0 T^*M$, the map $f \mapsto A(L(f))$ is a first order operator $\Gamma(V_0) \to S^2_0 T^*M \subset T^*M \otimes V$, and all such operators can be obtained in that way.

We have seen that $\tilde{\nabla}\Sigma + A(\Sigma) = \delta^*\psi$ implies $\Sigma = L(f)$, where $f = \Sigma_0$. Now

$$\tilde{\nabla} L(f) = \tilde{\nabla}_i \begin{pmatrix} -\frac{1}{n}\Delta f \\ \nabla_j f \\ f \end{pmatrix} = \begin{pmatrix} -\frac{1}{n}\nabla_i \Delta f \\ \nabla_i \nabla_j f - \frac{1}{n} g_{ij} \Delta f \\ 0 \end{pmatrix}$$

$\nabla_i \nabla_j f - \frac{1}{n} g_{ij} f = \nabla_{(i}\nabla_{j)_0} f$, the tracefree part of $\nabla^2 f$. Adding $A(L(f))$ corresponds to adding $D(f)$ to the middle component, where $D : \Gamma(V_0) \to \Gamma(S^2_0 T^* M)$ is a first order operator.

An element $\begin{pmatrix} h_i \\ \phi_{ij} \\ f_i \end{pmatrix} \in \Omega^1(M, V)$ is of the form $\delta^*\psi$ for some $\psi \in \Omega^2(M, V)$ iff $f_i = 0$ and ϕ_{ij} is skew. Since our middle component is symmetric by construction it has to vanish and we get:

Proposition

For any first order operator $D : \Gamma(V_0) \to \Gamma(S^2_0 T^*M)$ there is a bundle map $A : V \to T^*M \otimes V$ such that $f \mapsto L(f)$ and $\Sigma \mapsto \Sigma_0$ induce inverse bijections between solutions of $\nabla_{(i}\nabla_{j)_0} f + D(f) = 0$ and of $\tilde{\nabla}\Sigma + A(\Sigma) = \delta^*\psi$ for some $\psi \in \Omega^2(M, V)$.

We also see that if $\Sigma = \begin{pmatrix} h \\ \phi_i \\ f \end{pmatrix}$ satisfies $\tilde{\nabla}\Sigma + A(\Sigma) = \delta^*\psi$ for some $\psi \in \Omega^2(M, V)$, then we actually have

$$\begin{cases} \nabla_i h + \tau_i = 0 \\ \nabla_i \phi_j + h g_{ij} + A_{ij}(f, \phi) = 0 \\ \nabla_i f - \phi_i = 0 \end{cases}$$

for some one–form τ_i. To rewrite in closed form, it remains to compute τ_i.

Let $d^{\tilde{\nabla}} : \Omega^1(M, V) \to \Omega^2(M, V)$ be the covariant exterior derivative associated to $\tilde{\nabla}$. Explicitly

$$d^{\tilde{\nabla}}\alpha(\xi, \eta) = \tilde{\nabla}_\xi(\alpha(\eta)) - \tilde{\nabla}_\eta(\alpha(\xi)) - \alpha([\xi, \eta]).$$

- $d^{\tilde{\nabla}}\tilde{\nabla}\Sigma = R \bullet \Sigma$, the action of the Riemann curvature on Σ.
- $d^{\tilde{\nabla}}(A(\Sigma))$ is concentrated in the middle component, and depends only on Σ_0, Σ_1 and their first derivatives.
- Elements which are concentrated in the top component are reproduced by $\delta^* d^{\tilde{\nabla}}$.

Hence τ_i equals the top component of $-\delta^*(R \bullet \Sigma + d^{\tilde{\nabla}}(A(\Sigma)))$, and inserting the equations for $\nabla\Sigma_0$ and $\nabla\Sigma_1$ we already have, we obtain an expression in terms of the values of Σ.

Theorem (main result for the example)

For any first order operator $D:\Gamma(V_0)\to\Gamma(S^2_0T^*M)$ there is a bundle map $C:V\to T^*M\otimes V$ such that $f\mapsto L(f)$ and $\Sigma\mapsto\Sigma_0$ induce inverse isomorphisms between solutions of $\nabla_{(i}\nabla_{j)_0}f+D(f)=0$ and of $\tilde\nabla\Sigma+C(\Sigma)=0$.

Consequences

- Since any solution of $\tilde\nabla\Sigma+C(\Sigma)=0$ is determined by its value in one point, we conclude that any solution of $\nabla_{(i}\nabla_{j)_0}f+D(f)=0$ is determined by the values of f, df and Δf in one point.
- If D is linear then C can be chosen to be linear. The space of solutions has dimension $\leq n+2$, and equality is only possible if the connection $\tilde\nabla+C$ is flat.

There are interesting applications of this result. Given a Riemannian metric g on M and a nowhere vanishing function f, we can rescale the metric conformally to $\hat{g}$ by multiplying by an appropriate power of f and compute how such a change affects the components of the Riemann curvature.

- Weyl curvature remains unchanged
- Change of scalar curvature is governed by the Yamabe operator (conformal Laplacian)
- For $\hat{g} = \frac{1}{f^2}g$, the change of tracefree part of Ricci is governed by an operator with principal part $\nabla_{(i}\nabla_{j)_0}f$

Rescalings of g to *Einstein metrics* (i.e. metrics with Ricci proportional to the metric) are in bijective correspondence with nowhere vanishing solutions of a system of the form

$$\nabla_{(i}\nabla_{j)_0}f + A_{ij}f = 0.$$

Here we replace the standard representation $\mathbb{V}$ of $G \cong O(n+1,1)$ by an arbitrary finite dimensional irreducible representation. (This works for spinor representations using $Spin(n) \hookrightarrow Spin(n+1,1)$.)

The Lie algebra $\mathfrak{g}$ of G

For our choice of inner product, this has the form

$$\mathfrak{g} = \left\{ \begin{pmatrix} a & Z & 0 \\ X & A & -Z^t \\ 0 & -X^t & -a \end{pmatrix} : \begin{matrix} A \in \mathfrak{o}(n), a \in \mathbb{R}, \\ X \in \mathbb{R}^n, Z \in \mathbb{R}^{n*} \end{matrix} \right\}.$$

- The central block formed by A corresponds to $O(n) \subset G$
- The element $E := \left(\begin{smallmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{smallmatrix}\right)$ is called the *grading element*
- Note that $\mathbb{V} = \mathbb{V}_0 \oplus \mathbb{V}_1 \oplus \mathbb{V}_2$ comes from eigenspaces for E.

The $|1|$–grading of $\mathfrak{g}$

- $\mathrm{ad}(E) = [E, \ ]$ is diagonalizable on $\mathfrak{g}$ with eigenvalues $-1, 0, 1$
- The decomposition $\mathfrak{g} = \mathfrak{g}_{-1} \oplus \mathfrak{g}_0 \oplus \mathfrak{g}_1$ into eigenspaces has the property that $[\mathfrak{g}_i, \mathfrak{g}_j] \subset \mathfrak{g}_{i+j}$ (with $\mathfrak{g}_\ell = \{0\}$ for $\ell \neq 0, \pm 1$).
- The adjoint action of $\mathfrak{o}(n) \subset \mathfrak{g}_0$ preserves each $\mathfrak{g}_i$, and $\mathfrak{g}_{\pm 1}$ is isomorphic to $\mathbb{R}^n$.

representations of $\mathfrak{g}$

Let $\mathbb{W}$ be a finite dimensional irreducible representation of $\mathfrak{g}$

- E acts diagonalizably on $\mathbb{W}$ with eigenvalues $j_0, j_0 + 1, \dots, j_0 + N$ for some $j_0 \in \mathbb{R}$ and $N \in \mathbb{N}$.
- denoting eigenspaces by $\mathbb{W}_j$ for $j = 0, \dots, N$ we have $\mathfrak{g}_i \cdot \mathbb{W}_j \subset \mathbb{W}_{i+j}$.

Via highest weights, irreducible representations of $\mathfrak{g}$ correspond to pairs $(\mathbb{W}_0, r)$, where $\mathbb{W}_0$ is an irreducible representation of $\mathfrak{o}(n)$ and $r \geq 1 \in \mathbb{N}$. If $\mathbb{W}$ corresponds to $(\mathbb{W}_0, r)$ then $\mathbb{W}_0$ is isomorphic to the lowest E–eigenspace of $\mathbb{W}$, and N can be computed from $\mathbb{W}_0$ and r. The standard representation $\mathbb{V}$ corresponds to $(\mathbb{R}, 2)$.

The map ∂^*

Let $\mathbb{W} = \mathbb{W}_0 \oplus \cdots \oplus \mathbb{W}_N$ be an irreducible representation of $\mathfrak{g}$. Define $\partial^* : \Lambda^k \mathfrak{g}_1 \otimes \mathbb{W} \to \Lambda^{k-1} \mathfrak{g}_1 \otimes \mathbb{W}$ by

$$\partial^*(Z_1 \wedge \cdots \wedge Z_k \otimes w) := \sum_{i=1}^{k} (-1)^i Z_1 \wedge \cdots \wedge \widehat{Z_i} \wedge \cdots \wedge Z_k \otimes Z_i \cdot w.$$

This preserves homogeneity and since $\mathfrak{g}_1 \subset \mathfrak{g}$ is commutative, we get $Z_i \cdot Z_j \cdot w = Z_j \cdot Z_i \cdot w$, which implies $\partial^* \circ \partial^* = 0$.

The map ∂

$\mathfrak{g}_1 \cong (\mathfrak{g}_{-1})^*$ via the Killing form. Viewing elements $\Lambda^k \mathfrak{g}_1 \otimes \mathbb{W}$ as k–linear alternating maps on $\mathfrak{g}_{-1}$, we define $\partial : \Lambda^k \mathfrak{g}_1 \otimes \mathbb{W} \to \Lambda^{k+1} \mathfrak{g}_1 \otimes \mathbb{W}$ by

$$\partial\alpha(X_0, \dots, X_k) := \sum_{i=0}^k (-1)^i X_i \cdot \alpha(X_0, \dots, \widehat{X_i}, \dots, X_k).$$

This preserves homogeneity and $\partial \circ \partial = 0$.

Lemma (B. Kostant)

∂ and ∂^* are adjoint with respect to an inner product of Lie theoretic origin, and we get an algebraic Hodge decomposition

$$\Lambda^k \mathfrak{g}_1 \otimes \mathbb{W} = \mathrm{im}(\partial) \oplus (\ker(\partial) \cap \ker(\partial^*)) \oplus \mathrm{im}(\partial^*).$$

$\mathrm{im}(\partial) \oplus (\ker(\partial) \cap \ker(\partial^*)) = \ker(\partial)$ and likewise for ∂^*.

Kostant's version of the Bott–Borel–Weil theorem describes $\ker(\partial) \cap \ker(\partial^*)$ as a representation of $\mathfrak{g}_0$ (and hence of $\mathfrak{o}(n)$).

The Cartan product

Let $\mathbb{E}$ and $\mathbb{F}$ be two finite dimensional irreducible representations of a semisimple Lie algebra. Then

- There is a unique irreducible component $\mathbb{E} \circledcirc \mathbb{F} \subset \mathbb{E} \otimes \mathbb{F}$ whose highest weight is the sum of the highest weights of $\mathbb{E}$ and $\mathbb{F}$.
- There is a unique (up to multiples) nonzero equivariant map $\mathbb{E} \otimes \mathbb{F} \to \mathbb{E} \circledcirc \mathbb{F}$.

Both the space $\mathbb{E} \circledcirc \mathbb{F}$ and the map onto it is called the *Cartan product*. It represents the "main part" of the tensor product.

Theorem (very special case of Kostant's version of BBW)

Let $\mathbb{W} = \mathbb{W}_0 \oplus \cdots \oplus \mathbb{W}_N$ be the irreducible representation corresponding to $(\mathbb{W}_0, r)$. Then

- In degree zero, we have $\ker(\partial) = \mathbb{W}_0$ and $\operatorname{im}(\partial^*) = \mathbb{W}_1 \oplus \cdots \oplus \mathbb{W}_N$.
- In degree one, $\ker(\partial) \cap \ker(\partial^*) \cong S^r_0\mathfrak{g}_1 \circledcirc \mathbb{W}_0$, it is contained in $\mathfrak{g}_1 \otimes \mathbb{W}_{r-1}$ and the only irreducible component of $\Lambda^*\mathfrak{g}_1 \otimes \mathbb{W}$ of this isomorphism type.

In each degree ∂ induces an isomorphisms

$$\Lambda^k\mathfrak{g}_1 \otimes \mathbb{W} \supset \operatorname{im}(\partial^*) \to \operatorname{im}(\partial) \subset \Lambda^{k+1}\mathfrak{g}_1 \otimes \mathbb{W},$$

while ∂^* induces an isomorphism in the opposite direction. Replace ∂^* by δ^* defined by $\delta^*|_{\ker(\partial^*)} = 0$ and $\delta^*|_{\operatorname{im}(\partial)} = \partial^{-1}$. This has the same kernel and image as ∂^*.

For $i > 0$ we have $\partial : \mathbb{W}_i \to \mathfrak{g}_1 \otimes \mathbb{W}_{i-1}$. Acting with ∂ on the second factor we move to $\otimes^2 \mathfrak{g}_1 \otimes \mathbb{W}_{i-2}$, and iterating we get

$$\phi_i := (\mathrm{id} \otimes \cdots \otimes \partial) \circ \ldots \circ (\mathrm{id} \otimes \partial) \circ \partial : \mathbb{W}_i \to \otimes^i \mathfrak{g}_1 \otimes \mathbb{W}_0.$$

Interpreted as a multilinear map on $\mathfrak{g}_{-1}$, $\phi_i(w)$ is simply given by $(X_1, \ldots, X_i) \mapsto X_1 \cdots X_i \cdot w$, so ϕ_i has values in $S^i \mathfrak{g}_1 \otimes \mathbb{W}_0$. Together with Kostant's result simple direct arguments show

Proposition

Suppose that $\mathbb{W}$ corresponds to $(\mathbb{W}_0, r)$ and let $\mathbb{K} \subset S^r \mathfrak{g}_1 \otimes \mathbb{W}_0$ be the kernel of the Cartan product. The for each i, the map ϕ_i defines an isomorphism from $\mathbb{W}_i$ onto its image, which is given by

$$\mathrm{im}(\phi_i) = \begin{cases} S^i \mathfrak{g}_1 \otimes \mathbb{W}_0 & i < r \\ S^i \mathfrak{g}_1 \otimes \mathbb{W}_0 \cap S^{i-r} \mathfrak{g}_1 \otimes \mathbb{K} & i \geq r \end{cases}$$

Passing to Riemannian n–manifolds, consider

- the bundle $W = W_0 \oplus \cdots \oplus W_N$ induced by $\mathbb{W}$
- the bundle maps ∂ and δ^* on the bundles $\Lambda^k T^*M \otimes W$ and $\phi_i : W_i \to S^i T^*M \otimes W_0$ induced by the respective maps.
- The subbundle $H_1 := \ker(\delta^*) \cap \ker(\partial) \subset T^*M \otimes W_{r-1}$, which is isomorphic to $S^r_0 T^*M \circledcirc W_0$

The basic system

- Define $\tilde{\nabla}$ on W by $\tilde{\nabla}\Sigma := \nabla\Sigma + \partial\Sigma$.
- Choose a bundle map $A : W_0 \oplus \cdots \oplus W_{r-1} \to H_1$ and view it as $A : W \to T^*M \otimes W$.
- Consider the first order system

$$\tilde{\nabla}\Sigma + A(\Sigma) = \delta^*\psi \qquad \text{for some } \psi \in \Omega^2(M, W)$$

For $f \in \Gamma(W_0)$ we now define $L(f) := \sum_{i=0}^{N}(-1)^i(\delta^*\nabla)^i f$.

Proposition

(1) We have $L(f)_0 = f$ and $\delta^*\tilde{\nabla}L(f) = 0$, and these two properties characterize $L(f)$.
(2) Mapping f to $(L(f)_0, \dots, L(f)_{r-1})$ induces an isomorphism from the $(r-1)$st jet prolongation $J^{r-1}W_0$ to $W_0 \oplus \dots \oplus W_{r-1}$.

Sketch of proof (1) The first property is evident. By definition, we have $(\tilde{\nabla}\Sigma)_i = \nabla\Sigma_i + \partial(\Sigma_{i+1})$. By construction $\delta^*\partial$ is the identity on $W_1 \oplus \dots \oplus W_N$, so $\delta^*\tilde{\nabla}\Sigma = 0$ is equivalent to $\Sigma_{i+1} = -\delta^*\nabla\Sigma_i$ for all $i \geq 0$.
(2) On shows that, up to lower order terms, $L(f)_i$ is obtained by applying ϕ_i^{-1} to the symmetrization of $\nabla^i f$. Then (2) follows directly from from the fact that ϕ_i is an isomorphism for $i < r$.

Define $D^{\mathbb{W}}(f)$ to be the H_1–component of $\tilde{\nabla}L(f) \in \Omega^1(M, W)$. The above observations easily imply that

- $D^{\mathbb{W}}$ is nonzero, of order r, and hence its principal symbol is a multiple of the Cartan product
- the ith classical prolongation of this symbol is isomorphic to the bundle W_{r+i}

Proposition

Let $D : \Gamma(W_0) \to \Gamma(H_1)$ be a differential operator of order $< r$. Then there is a bundle map $A : W \to T^*M \otimes W$ as above such that $f \mapsto L(f)$ and $\Sigma \mapsto \Sigma_0$ induce inverse bijections between solutions of $D^{\mathbb{W}}(f) + D(f) = 0$ and of $\tilde{\nabla}\Sigma + A(\Sigma) = \partial^*\psi$ for some $\psi \in \Omega^2(M, W)$.

Proof.

- Choose A in such a way that $D(f) = A(L(f))$. Then $D^{\mathbb{W}}(f) + D(f)$ is the H_1–component of $\tilde{\nabla}L(f) + A(L(f))$.
- $\delta^*(\tilde{\nabla}\Sigma + A(\Sigma)) = 0$ is equivalent to $\partial^*\tilde{\nabla}\Sigma = 0$ and hence to $\Sigma = L(\Sigma_0)$.
- A section of $\ker(\delta^*)$ has values in the subbundle $\operatorname{im}(\delta^*)$ iff its H_1 component is trivial.

To rewrite in closed form, let $d^{\tilde{\nabla}}$ be the covariant exterior derivative.

- on $\Gamma(\mathrm{im}(\partial^*)) \subset \Omega^1(M, W)$ the operator $\delta^* d^{\tilde{\nabla}}$ reproduces the lowest nonzero component
- for a solution Σ of the basic system, one may therefore compute the lowest nonzero component of $\tilde{\nabla}\Sigma + A(\Sigma)$ from $\delta^* d^{\tilde{\nabla}}(\tilde{\nabla}\Sigma + A(\Sigma))$.
- $d^{\tilde{\nabla}}\tilde{\nabla}\Sigma = R \bullet \Sigma$, and $d^{\tilde{\nabla}}A(\Sigma)$ depends only on $\Sigma_0, \dots, \Sigma_{r-1}$ and their first derivatives.
- This procedure can be iterated and the order of the individual components can be controlled well enough to show that inserting know equations for lower components and their derivatives into higher components leads to

Theorem

For any differential operator $D : \Gamma(W_0) \to \Gamma(H_1)$ of order $r-1$, there is a bundle map $C : V \to T^*M \otimes V$ such that $f \mapsto L(f)$ and $\Sigma \mapsto \Sigma_0$ induce inverse isomorphisms between the sets of solutions of $D^{\mathbb{W}}(f) + D(f) = 0$ and of $\tilde{\nabla}\Sigma + C(\Sigma) = 0$.

Consequences • Any solution of $D^{\mathbb{W}}(f) + D(f) = 0$ is uniquely determined by the value of $L(f)$ and hence by the N–jet of f in a single point

• If D is linear, then the dimension of the space of solutions of $D^{\mathbb{W}}(f) + D(f) = 0$ is $\leq \dim(\mathbb{W})$. Equality is only possible, if the linear connection $\Sigma \mapsto \tilde{\nabla}\Sigma + C(\Sigma)$ on the bundle W is flat.

• Since N and $\dim(\mathbb{W})$ can be easily computed from $(\mathbb{W}_0, r)$, these bounds are available without going through the procedure. They both turn out to be sharp.

A *conformal structure* on a smooth n–manifold M is an equivalence class $[g]$ of Riemannian metric obtained from each other by multiplication by a positive smooth function. We will assume $n \geq 3$ throughout.

Conformal interpretation of $O(n+1,1)$

- Take $G = O(\mathbb{V})$ as before and let $P \subset G$ be the stabilizer of the isotropic line generated by e_0.
- Then G/P is the space of all isotropic lines in $\mathbb{V}$, and it is easy to see that $G/P \cong S^n$.
- In terms of the grading $\mathfrak{g} = \mathfrak{g}_{-1} \oplus \mathfrak{g}_0 \oplus \mathfrak{g}_1$, the Lie algebra of P is $\mathfrak{p} = \mathfrak{g}_0 \oplus \mathfrak{g}_1$.
- The inner product on $\mathbb{V}$ induces a conformal structure on G/P, so G acts by conformal isometries. It turns out that all conformal isometries of S^n are obtained in this way.

Conformal interpretation of $O(n+1,1)$

- Let $P_+ \subset P$ be the closed normal subgroup of elements which fix $o = eP \in G/P$ to first order. Then P_+ has Lie algebra $\mathfrak{g}_1$ and $\exp : \mathfrak{g}_1 \to P_+$ is a diffeomorphism.
- $T_oG/P \cong \mathfrak{g}/\mathfrak{p} \cong \mathfrak{g}_{-1}$ and mapping $g \in P$ to $T_o\ell_g$ induces an isomorphism $P/P_+ \cong CO(\mathfrak{g}_{-1})$.
- G/P_+ is a principal bundle over G/P with structure group $P/P_+ \cong CO(n)$, which can be naturally identified with the conformal frame bundle of S^n.

Theorem (E. Cartan ∼1925)

Let $(M, [g])$ be a conformal manifold. Then the conformal frame bundle of M can be canonically extended to a principal P–bundle $p : \mathcal{G} \to M$, which can be endowed with a canonical Cartan connection $\omega \in \Omega^1(\mathcal{G}, \mathfrak{g})$.

conformally invariant differential operators

- idea: operators intrinsic to a conformal structure
- formally: operators on Riemannian manifolds, which can be written terms of g, ∇ and R, and remain unchanged if g is replaced by a conformally equivalent metric
- the direct approach suggested by this definition works well for low order but quickly gets out of hand
- more conceptual approaches are based on the canonical Cartan connection

The machinery of BGG sequences is one of these conceptual approaches. To set up the stage, we have to rethink the ingredients of our prolongation procedure from the point of view of the group P.

- For $g \in P$, the grading of $\mathfrak{g}$ is not $\mathrm{Ad}(g)$–invariant, but $\mathfrak{p} = \mathfrak{g}_0 \oplus \mathfrak{g}_1$ and $\mathfrak{g}_1$ are.
- If $\mathbb{W}$ is an irreducible representation of G, with decomposition into E–eigenspaces $\mathbb{W} = \mathbb{W}_0 \oplus \cdots \oplus \mathbb{W}_N$, then each of the subspaces $\mathbb{W}^i := \mathbb{W}_i \oplus \cdots \oplus \mathbb{W}_N$ is P–invariant.
- $\partial^* : \Lambda^k \mathfrak{g}_1 \otimes \mathbb{W} \to \Lambda^{k-1} \mathfrak{g}_1 \otimes \mathbb{W}$ is P–equivariant.
- $\partial : \Lambda^k \mathfrak{g}_1 \otimes \mathbb{W} \to \Lambda^{k+1} \mathfrak{g}_1 \otimes \mathbb{W}$ is *not* P–equivariant. (While P naturally acts on $\mathfrak{g}_{-1}$ via the identification with $\mathfrak{g}/\mathfrak{p}$, the action of $\mathfrak{g}_{-1}$ on $\mathbb{W}$ has no interpretation in this picture.)

associated bundles

Via the canonical Cartan connection ω we get isomorphisms $TM \cong \mathcal{G} \times_P (\mathfrak{g}/\mathfrak{p})$ and $T^*M \cong \mathcal{G} \times_P \mathfrak{g}_1$. Any usual conformally natural bundle is associated to $\mathcal{G}$ via the quotient homomorphism $P \to P/P_+ \cong CO(n)$.

tractor bundles

Let $\mathbb{W}$ be an irreducible representation of G. By restriction, it is also a representation of P.

- the induced bundle $\mathcal{W} := \mathcal{G} \times_P \mathbb{W}$ is called a *tractor bundle*
- geometrically, these are unusual objects, since the action of conformal isometries needs second order information
- we get $\partial^* : \Lambda^k T^*M \otimes \mathcal{W} \to \Lambda^{k-1} T^*M \otimes \mathcal{W}$ and $\ker(\partial^*)$ and $\operatorname{im}(\partial^*)$ are natural subbundles.
- the subquotient $H_k := \ker(\partial^*)/\operatorname{im}(\partial^*)$ is a usual conformally natural bundle and explicitly computable using Kostant's BBW.

tractor connections

- the canonical Cartan connection ω induces linear connections on tractor bundles (but not on general associated vector bundles)
- the tractor connection $\nabla^{\mathcal{W}}$ automatically mixes differential and algebraic components; it lowest homogeneous component is tensorial and induced by ∂.

The BGG machinery (1)

Extend $\nabla^{\mathcal{W}}$ to the covariant exterior derivative $d^{\mathcal{W}}$.

- $\partial^* d^{\mathcal{W}}$ defines a conformally invariant operator on $\Omega^k(M, \mathcal{W})$
- restricted to $\Gamma(\text{im}(\partial^*))$, the lowest homogeneous component of this operator is tensorial and invertible
- this implies that $\partial^* d^{\mathcal{W}}|_{\Gamma(\text{im}(\partial^*))}$ is invertible, and the inverse Q is a differential operator.

The BGG machinery (2)

- let $\pi_H : \Gamma(\ker(\partial^*)) \to \Gamma(H_k)$ be the natural tensorial projection
- For $f \in \Gamma(H_k)$ choose $\phi \in \Gamma(\ker(\partial^*))$ such that $\pi_H(\phi) = f$ and consider $\phi - Q\partial^* d^{\mathcal{W}}\phi$.
- This depends only on f, so we obtain a differential operator $L : \Gamma(H_k) \to \Omega^k(M, \mathcal{W})$, such that $\pi_H(L(f)) = f$.
- $L(f)$ further satisfies $\partial^* L(f) = 0$ and $\partial^* d^{\mathcal{W}} L(f) = 0$ and is characterized by these properties
- The BGG operator $D^{\mathcal{W}} : \Gamma(H_k) \to \Gamma(H_{k+1})$ is then defined by $D^{\mathcal{W}}(f) := \pi_H d^{\mathcal{W}} L(f)$.

These BGG operators are conformally invariant by construction. To see that they are non–trivial, one has to look at locally conformally flat manifolds.

The locally conformally flat case

- $\nabla^{\mathcal{W}}$ has trivial curvature, so $(\Omega^*(M,\mathcal{W}), d^{\mathcal{W}})$ is a fine resolution of the sheaf of parallel sections of $\mathcal{W}$
- Using $d^{\mathcal{W}} \circ d^{\mathcal{W}} = 0$, one easily shows that $L \circ D^{\mathcal{W}} = d^{\mathcal{W}} \circ L$ and $(\Gamma(H_*), D^{\mathcal{W}})$ is also a complex.
- L induces an isomorphism in cohomology, so $(\Gamma(H_*), D^{\mathcal{W}})$ is also a fine resolution. In particular, all $D^{\mathcal{W}}$ are nontrivial.
- We already know that $H_0 = W_0$ and $H_1 = S^r_0 T^*M \circledcirc W_0$. By naturality, the principal symbol of $D^{\mathcal{W}} : \Gamma(H_0) \to \Gamma(H_1)$ must be the Cartan product.
- Since L identifies $\ker(D^{\mathcal{W}})$ with the space of parallel sections of the flat connection $\nabla^{\mathcal{W}}$, sharpness of the bounds in the prolongation procedure follows.