Michigan Math. J. 55 (2007)

Invariant Metrics with Nonnegative Curvature on SO(4) and Other Lie Groups

Jack Huizenga & Kristopher Tapp

1. Introduction

The starting point for constructing all known examples of compact manifolds with positive (or even quasi-positive) curvature is the fact that bi-invariant metrics on compact Lie groups are nonnegatively curved. In order to generalize this fundamental starting point, we address the following problem: Given a compact Lie group G, classify the left-invariant metrics on G that have nonnegative curvature. New examples could potentially, via familiar quotient constructions, lead to new examples of quasi-positively curved spaces. On the other hand, proofs that there are no new examples would serve as further evidence that the known constructions are rigid and canonical.

The first two cases, $G = \mathrm{SO}(3)$ and $U(2)$, were completely solved in [1]. For $G = U(2)$, all such metrics lie in the closure of those coming from Cheeger's method, which is essentially the only known construction of nonnegatively curved left-invariant metrics. These classifications made use of techniques that work only in low dimensions. For higher-dimensional groups, more tools are necessary to approach the problem effectively. One important new tool is the following, which implies in particular that the nonnegatively curved metrics form a path-connected subset within the space of all left-invariant metrics.

Theorem 1.1. *If h is a left-invariant metric with nonnegative curvature on a compact Lie group G, then the unique inverse-linear path from any fixed bi-invariant metric $h(0)$ to $h(1) = h$ is through nonnegatively curved metrics.*

Here, a path of inner products on $\mathfrak{g} = T_eG$ (or the induced path of left-invariant metrics) is called *inverse-linear* if the inverses of the associated path of symmetric matrices form a straight line. So to classify the left-invariant metrics on G with nonnegative curvature, we can first classify the directions $h'(0)$ in which one can move away from a fixed bi-invariant metric $h(0)$ such that the inverse-linear path $h(t)$ appears (up to derivative information at $t = 0$) to remain nonnegatively curved. Then, for each candidate direction, we must check how far the nonnegative curvature is maintained along that path.

Received September 11, 2006. Revision received January 23, 2007.
Supported in part by NSF Grant nos. DMS-0303326 and DMS-0353634.

This is the approach we use for general G. In the case $G = \mathrm{SO}(4)$, our results provide strong evidence that all left-invariant metrics lie in the closure of those coming from Cheeger's method; that is, there do not seem to be any new examples. One of our stronger results toward the classification for SO(4) is the following.

THEOREM 1.2. *If h is a left-invariant metric with nonnegative curvature on* SO(4) *and if the matrix of h has an eigenvector in one of the simple factors of* $\mathrm{so}(4) = \mathrm{so}(3) \oplus \mathrm{so}(3)$, *then h is a known example of a metric of nonnegative curvature.*

The known examples come from Cheeger's method via an action of T^2 or S^3, as explained in Section 7. Those from a T^2 action have a singular eigenvector, as in Theorem 1.2.

The authors are pleased to thank Burkhard Wilking, Craig Sutton, Emily Proctor, Zachary Madden, Nela Vukmirovic, Angela Doyle, Min Kim, and the referee for numerous helpful discussions and comments on this work.

2. Cheeger's Method

In this section, we review Cheeger's method for altering a nonnegatively curved metric via a group of isometries and then use it to prove Theorem 1.1.

Let (M, h_0) be a nonnegatively curved manifold on which a compact Lie group G acts by isometries. Let h_R be a right-invariant metric on G with nonnegative curvature (often chosen to be bi-invariant). Observe that G acts on $M \times G$ as $g \star (p, a) = (g \star p, ag^{-1})$. The orbit space is diffeomorphic to M via the map $[p, g] \mapsto g \star p$. Consider the 1-parameter family of induced nonnegatively curved Riemannian submersion metrics, h_t, on this orbit space:

$$(M, h_t) = (M \times (G, (1/t)h_R))/G.$$

This family extends smoothly at $t = 0$ to the original metric h_0 on M. To describe the metric variation at a fixed $p \in M$, let $\{v_1, \dots, v_k\} \subset T_pM$ denote the values at p of the Killing fields on M associated to an h_R-orthonormal basis $\{e_1, \dots, e_k\}$ of the Lie algebra $\mathfrak{g}$ of G. Cheeger's formula in [2] implies that the path of matrices $A^t_{ij} = h_t(v_i, v_j)$ evolves according to

$$A^t = A^0(I + tA^0)^{-1}. \tag{2.1}$$

Several authors have derived curvature variation formulas, although they usually assume that h_R is bi-invariant; see [5; 6; 8; 9]. For this, it is useful to consider the bijection $\Phi_t \colon T_pM \to T_pM$, which describes h_t in terms of h_0 in the sense that, for all $X, Y \in T_pM$,

$$h_t(X, Y) = h_0(\Phi_t(X), Y).$$

This family of inner products on T_pM is *inverse-linear*. This means that the path $t \mapsto \Phi_t^{-1}$ is linear, so $\Phi_t = (I - t\Psi)^{-1}$ for some endomorphism $\Psi \colon T_pM \to T_pM$.

Cheeger mentioned that h_t has no more zero-curvature planes than h_0. A precise formulation of this comment, found for example in [6], is as follows.

LEMMA 2.1. *If the plane* $\sigma = \mathrm{span}\{X, Y\}$ *has positive curvature with respect to* h_0, *then the plane* $\Phi_t^{-1}(\sigma) = \mathrm{span}\{\Phi_t^{-1}(X), \Phi_t^{-1}(Y)\}$ *has positive curvature with respect to* h_t.

Hence the most natural variational approach is to differentiate the curvature with respect to h_t of the plane $\Phi_t^{-1}(\sigma)$; this was systematically studied in [5]. In the next section, we will borrow and generalize this idea.

Proof of Theorem 1.1. Let h be a left-invariant metric with nonnegative curvature on the compact Lie group G. Let h_0 be a fixed bi-invariant metric on G. Consider the family h_t of nonnegatively curved metrics on G defined by

$$(G, h_t) = ((G, h_0) \times (G, (1/t)h))/G,$$

where G acts diagonally on the right of both factors. In order for this action to be isometric, h must be reconsidered as a right-invariant metric on G, which is no problem because the left- and right-invariant metrics determined by an inner product on $\mathfrak{g}$ are isometric via the inversion map. Notice that each h_t is a left-invariant metric on G.

Let $\{E_1, \dots, E_k\}$ be an h_0-orthonormal basis of $\mathfrak{g}$ that diagonalizes h. Let $\{\lambda_1, \dots, \lambda_k\}$ be the corresponding eigenvalues of h, so that $\{e_i = E_i/\sqrt{\lambda_i}\}$ is an h-orthonormal basis of $\mathfrak{g}$. In equation (2.1), $v_i = e_i$ and $A^0 = \mathrm{diag}(1/\lambda_i)$, so $A^t = \mathrm{diag}(1/(\lambda_i + t))$. Thus, in the basis $\{E_i\}$, the matrix for Φ_t is

$$\Phi_t = \mathrm{diag}(1 + (1/\lambda_i)t)^{-1}.$$

Therefore, $\Phi_t = (I - t\Psi)^{-1}$, where $\Psi = \mathrm{diag}(-1/\lambda_i)$. We see that, as previously mentioned, the path is inverse-linear.

There is no value of t for which $h_t = h$. Instead we will show that the path h_t (for $t \in [0, \infty)$) visits scalings of all of the metrics along the unique inverse-linear path $\tilde{h}_s$ between $\tilde{h}_0 = h_0$ and $\tilde{h}_1 = h$. Let $\tilde{\Phi}_s$ determine this path, so that $\tilde{h}_s(X, Y) = h_0(\tilde{\Phi}_s X, Y)$ for all $X, Y \in \mathfrak{g}$. We have that $\tilde{\Phi}_s = (I - s\tilde{\Psi})^{-1}$, where $\tilde{\Psi}$ with respect to the basis $\{E_i\}$ is given by

$$\tilde{\Psi} = I - \tilde{\Phi}_1^{-1} = \mathrm{diag}(1 - 1/\lambda_i).$$

It is easy to see that the paths $\tilde{\Phi}_s$ (for $s \in [0, 1)$) and Φ_t (for $t \in [0, \infty)$) visit the same family of metrics up to scaling. More precisely, $c \cdot \tilde{\Phi}_s = \Phi_t$ when $t = s/(1 - s)$ and $c = 1 - s$. □

The method of the proof can be used to connect any two nonnegatively curved left-invariant metrics h_1 and h_2 on G through a path of nonnegatively curved metrics. The resultant path of inner products on $\mathfrak{g}$ is inverse-linear, but this is largely irrelevant to the question at hand because the path is not through left-invariant metrics.

3. Curvature Variation of 0-Planes

In this section and the next, we derive a curvature variation formula for an inverse-linear path of left-invariant metrics beginning at a bi-invariant metric.

Let G be a compact Lie group. Let h_t be an inverse-linear path of left-invariant metrics on G beginning at a bi-invariant metric h_0. The value of h_t at e is determined in terms of h_0 by some self-adjoint $\Phi_t\colon \mathfrak{g} \to \mathfrak{g}$ defined so that, for all $X, Y \in \mathfrak{g}$,

$$h(X,Y) = h_0(\Phi_t(X), Y).$$

Recall that "inverse-linear" means that

$$\Phi_t = (I - t\Psi)^{-1}$$

for some endomorphism $\Psi\colon \mathfrak{g} \to \mathfrak{g}$. Observe that $\Psi = \frac{d}{dt}\big|_{t=0}\Phi_t$ and hence Ψ is h_0-self-adjoint. For fixed $X, Y \in \mathfrak{g}$, define $\kappa(t)$ to be the unnormalized sectional curvature of $\{\Phi_t^{-1}X, \Phi_t^{-1}Y\}$ with respect to the metric h_t. The domain of $\kappa(t)$ is the open interval of t values for which Φ_t represents a nondegenerate metric; this interval depends on the eigenvalues of Ψ.

Two important decisions here are inspired by properties of Cheeger's method: (1) restricting to inverse-linear paths, and (2) "twisting" the plane whose curvature we are tracking. Even though we are considering general paths not necessarily arising from Cheeger's method, Theorem 1.1 and several results to follow indicate that these decisions provide the correct approach.

If $Z_1, Z_2 \in \mathfrak{g}$, we write $\langle Z_1, Z_2\rangle = h_0(Z_1, Z_2)$, $|Z_1|^2 = h_0(Z_1, Z_1)$, and $|Z_1|^2_{h_t} = h_t(Z_1, Z_1) = \langle \Phi_t Z_1, Z_1\rangle$. We first describe $\kappa(t)$ in the important special case where $[X, Y] = 0$, so that $\kappa(0) = \frac{1}{4}|[X,Y]|^2 = 0$. In other words, we first study the variation of curvature for an initially zero-curvature plane.

PROPOSITION 3.1. *If* $[X,Y] = 0$, *then* $\kappa(0) = 0$, $\kappa'(0) = 0$, $\kappa''(0) = 0$, *and*

$$\begin{aligned}\tfrac{1}{6}\kappa'''(0) = {} & \langle [X, \Psi Y] + [\Psi X, Y], [\Psi X, \Psi Y]\rangle + \langle [\Psi X, X], \Psi[\Psi Y, Y]\rangle \\ & - \langle [X, \Psi Y], \Psi[X, \Psi Y]\rangle - \langle [X, \Psi Y], \Psi[\Psi X, Y]\rangle \\ & - \langle [\Psi X, Y], \Psi[\Psi X, Y]\rangle.\end{aligned}$$

Moreover, for all t in the domain of κ,

$$\kappa(t) = t^3 \cdot \tfrac{1}{6}\kappa'''(0) - t^4 \cdot \tfrac{3}{4}|[\Psi X, \Psi Y] - \Psi([\Psi X, Y] + [X, \Psi Y])|^2_{h_t}.$$

We will prove this proposition in the next section as a special case of a more general formula that does not assume that X and Y commute.

In the Taylor series of $\kappa(t)$ at 0, the first nonvanishing derivative is the third, after which the remaining tail sums to a nonpositive term involving the norm with respect to h_t of the vector

$$D = [\Psi X, \Psi Y] - \Psi([\Psi X, Y] + [X, \Psi Y]).$$

In light of our formula for $\kappa(t)$, we can assert the following definition.

DEFINITION 3.2. We call Ψ (or the variation Φ_t) *infinitesimally nonnegative* if the following equivalent conditions hold.

(1) For all $X, Y \in \mathfrak{g}$, there exists an $\varepsilon > 0$ such that $\kappa(t) \geq 0$ for $t \in [0, \varepsilon)$.
(2) For all commuting pairs $X, Y \in \mathfrak{g}$, $\kappa'''(0) \geq 0$, and $\kappa'''(0) = 0$ implies that $D = 0$.

If in the first condition a single choice of $\varepsilon > 0$ works for all pairs X, Y, then Φ_t has nonnegative curvature for $t \in [0, \varepsilon)$. In this case, we call the variation *locally nonnegative*. We do not know if infinitesimally nonnegative implies locally nonnegative. In any case, the infinitesimally nonnegative Ψ are the candidate directions; the best available derivative information predicts that the paths in these directions are through nonnegatively curved metrics.

It is significant that the tail of the power series for $\kappa(t)$ is nonpositive. In addition to demonstrating the equivalence of the two parts of Definition 3.2, this nonpositivity property immediately implies the following weak version of Theorem 1.1: *If h_t is nonnegatively curved for some $t > 0$, then Ψ is infinitesimally nonnegative.* This is the only version of Theorem 1.1 we will need throughout the rest of the paper. It states that one will locate all nonnegatively curved metrics by searching only along the infinitesimally nonnegative paths.

If one omits the plane twisting and instead defines $\kappa(t)$ as the unnormalized sectional curvature of $\{X, Y\}$, then $\kappa(0) = 0$ implies that $\kappa'(0) = 0$ and that $\kappa''(0) = |[X, \Psi Y] + [\Psi X, Y]|^2$. This statement is true without the assumption of an inverse-linear path so long as $\Psi = \frac{d}{dt}\big|_{t=0}\Phi_t$. It is interesting that $\kappa''(0) \geq 0$, but this means that the untwisted setup provides little help in deciding which variations remain nonnegatively curved. We will stick with the twisted version for the remainder of the paper.

Example 3.3. Suppose $H \subset G$ is a Lie subgroup with Lie algebra $\mathfrak{h} \subset \mathfrak{g}$. For $A \in \mathfrak{g}$, let $A^{\mathfrak{h}}$ and $A^{\mathfrak{p}}$ denote the projections of A onto and orthogonal to $\mathfrak{h}$ with respect to h_0. The variation $\Phi_t(A) = \frac{1}{1+t}A^{\mathfrak{h}} + A^{\mathfrak{p}}$ is inverse-linear and has nonnegative curvature for $t > 0$. In this variation, vectors tangent to H are gradually shrunk. The parameterization looks natural when re-described as a family of submersion metrics: $(G, h_t) = ((G, h_0) \times (H, (1/t)h_0))/H$. The $t = 0$ derivative is $\Psi A = -A^{\mathfrak{h}}$. Proposition 3.1 yields:

$$\tfrac{1}{6}\kappa'''(0) = |[X^{\mathfrak{h}}, Y^{\mathfrak{h}}]|^2. \tag{3.1}$$

Equation 3.1 (together with Lemma 2.1 and the nonpositivity of the tail of the power series for $\kappa(t)$) re-proves Eschenburg's formula from [3], which states that, with respect to the metric h_t (for fixed $t > 0$)), the plane spanned by $\Phi_t^{-1}(X)$ and $\Phi_t^{-1}(Y)$ has zero curvature if and only if $[X, Y] = 0$ and $[X^{\mathfrak{h}}, Y^{\mathfrak{h}}] = 0$.

The full domain of this variation is $(-1, \infty)$. As t decreases from 0 toward -1, vectors tangent to H are enlarged. Considering negative values of t for this variation is equivalent to considering positive values of t for the variation in the opposite direction, $-\Psi$. For this opposite variation, $\frac{1}{6}\kappa'''(0) = -|[X^{\mathfrak{h}}, Y^{\mathfrak{h}}]|^2$. Therefore, expanding $\mathfrak{h}$ immediately creates some negative curvature unless $[X^{\mathfrak{h}}, Y^{\mathfrak{h}}] = 0$ whenever $[X, Y] = 0$. If $\mathfrak{h}$ is abelian then $\kappa'''(0) = 0$ for all commuting X, Y, which suggests that enlarging an abelian subalgebra might preserve nonnegative curvature. Indeed, it is proven in [4] that enlarging an abelian subalgebra as far as $\frac{4}{3}$ always preserves nonnegative curvature. In Section 6, we will study this variation in greater depth to determine which subalgebras can be enlarged without losing nonnegative curvature.

Notice that, for $a > 0$, Ψ and $a\Psi$ generate different parameterizations of the same family of metrics. A slightly less obvious equivalence involves adding a multiple of the identity to Ψ.

PROPOSITION 3.4. *If Ψ is infinitesimally nonnegative, then so is $\tilde{\Psi} = \Psi + a \cdot I$ for any $a > 0$.*

This proposition gives the correct equivalence modulo which one should classify the infinitesimally nonnegative endomorphisms Ψ.

Proof of Proposition 3.4. The endomorphisms Ψ and $\tilde{\Psi}$ yield the same values for $\kappa'''(0)$ and D in Proposition 3.1. To verify this, it is convenient to use Equation 4.4.

An alternative proof is to observe that the inverse-linear paths $\Phi(t) = (I - t\Psi)^{-1}$ and $\tilde{\Phi}(s) = (I - s\tilde{\Psi})^{-1}$ visit the same family of metrics modulo scalings and reparameterizations. More precisely, $c \cdot \Phi(t) = \tilde{\Phi}(s)$ provided $c = 1 - s \cdot a$ and $t = s/(1 - s \cdot a)$. Note that this idea was used previously in the proof of Theorem 1.1.

4. Curvature Variation of General Planes

In this section we state and prove a generalization of Proposition 3.1 that does not assume X and Y commute. We use this result to prove the proposition.

Certain elements of $\mathfrak{g}$ will appear frequently in what follows, so to simplify the exposition we introduce the Lie algebra elements

$$\begin{aligned} A &= [\Psi X, Y] + [X, \Psi Y], \\ B &= [\Psi X, \Psi Y], \\ C &= [\Psi X, Y] + [\Psi Y, X], \\ D &= \Psi^2[X, Y] - \Psi A + B. \end{aligned}$$

The definition of D given here coincides with the definition in Section 3 when X and Y commute.

THEOREM 4.1. *For any t in the domain of κ,*

$$\kappa(t) = \alpha + \beta t + \gamma t^2 + \delta t^3 - \tfrac{3}{4} t^4 \cdot |D|^2_{h_t}, \tag{4.1}$$

where

$$\begin{aligned} \alpha &= \tfrac{1}{4}|[X, Y]|^2, \\ \beta &= -\tfrac{3}{4}\langle \Psi[X, Y], [X, Y] \rangle, \\ \gamma &= -\tfrac{3}{4}|\Psi[X, Y]|^2 + \tfrac{3}{2}\langle \Psi[X, Y], A \rangle - \tfrac{1}{2}\langle [X, Y], B \rangle \\ &\quad - \tfrac{1}{4}|A|^2 + \tfrac{1}{4}|C|^2 - \langle [\Psi X, X], [\Psi Y, Y] \rangle, \\ \delta &= -\tfrac{3}{4}\langle \Psi^3[X, Y], [X, Y] \rangle + \tfrac{3}{2}\langle \Psi^2[X, Y], A \rangle - \tfrac{3}{2}\langle \Psi[X, Y], B \rangle \\ &\quad - \tfrac{3}{4}\langle \Psi A, A \rangle - \tfrac{1}{4}\langle \Psi C, C \rangle + \langle \Psi[\Psi X, X], [\Psi Y, Y] \rangle + \langle A, B \rangle. \end{aligned}$$

There are two steps to the proof of this theorem: first we prove that equation (4.1) holds for all sufficiently small t, and then we show that each side of the equation is analytic. This allows us to invoke the well-known identity theorem that, if $f, g\colon I \to \mathbb{R}$ are analytic on an open interval I and if f and g agree on a subinterval of I, then $f = g$. We therefore conclude that (4.1) holds for all t. To accomplish the first step, we calculate the Taylor series of $\kappa(t)$ at $t = 0$. This calculation will also serve as the foundation for our analyticity arguments.

PROPOSITION 4.2. *The Taylor series of $\kappa(t)$ at 0 is given by*

$$\kappa(t) = \alpha + \beta t + \gamma t^2 + \delta t^3 - \frac{3}{4}\sum_{n=4}^{\infty} t^n \langle \Psi^{n-4} D, D\rangle,$$

with convergence for $|t| < \|\Psi\|^{-1}$, where $\|\Psi\| = \sup_{|X|=1}|\Psi X|$ is the operator norm of Ψ.

Proof. In [7], Püttmann shows that the unnormalized sectional curvature of vectors $Z_1, Z_2 \in \mathfrak{g}$ with respect to a left-invariant metric h whose matrix with respect to h_0 is Φ is given by

$$\begin{aligned} k_h(Z_1, Z_2) = {}& \tfrac{1}{2}\langle [\Phi Z_1, Z_2] + [Z_1, \Phi Z_2], [Z_1, Z_2]\rangle - \tfrac{3}{4}|[Z_1, Z_2]|_h^2 \\ &+ \tfrac{1}{4}\langle [Z_1, \Phi Z_2] + [Z_2, \Phi Z_1], \Phi^{-1}([Z_1, \Phi Z_2] + [Z_2, \Phi Z_1])\rangle \\ &- \langle [Z_1, \Phi Z_1], \Phi^{-1}[Z_2, \Phi Z_2]\rangle. \end{aligned} \tag{4.2}$$

It follows that

$$\begin{aligned} \kappa(t) &= k_{h_t}(\Phi_t^{-1}X, \Phi_t^{-1}Y) \\ &= \tfrac{1}{2}\langle [X, \Phi_t^{-1}Y] + [\Phi_t^{-1}X, Y], [\Phi_t^{-1}X, \Phi_t^{-1}Y]\rangle \\ &\quad - \tfrac{3}{4}\langle \Phi_t[\Phi_t^{-1}X, \Phi_t^{-1}Y], [\Phi_t^{-1}X, \Phi_t^{-1}Y]\rangle \\ &\quad + \tfrac{1}{4}\langle [\Phi_t^{-1}X, Y] + [\Phi_t^{-1}Y, X], \Phi_t^{-1}([\Phi_t^{-1}X, Y] + [\Phi_t^{-1}Y, X])\rangle \\ &\quad - \langle [\Phi_t^{-1}X, X], \Phi_t^{-1}[\Phi_t^{-1}Y, Y]\rangle \\ &= I_1 - I_2 + I_3 - I_4. \end{aligned}$$

Using the expression $\Phi_t^{-1} = I - t\Psi$, we can easily simplify I_1, I_3, and I_4:

$$\begin{aligned} I_1 &= |[X, Y]|^2 - \frac{3t}{2}\langle [X, Y], A\rangle + t^2\left(\langle [X, Y], B\rangle + \frac{1}{2}|A|^2\right) - \frac{t^3}{2}\langle A, B\rangle, \\ I_3 &= \frac{t^2}{4}|C|^2 - \frac{t^3}{4}\langle C, \Psi C\rangle, \\ I_4 &= t^2\langle [\Psi X, X], [\Psi Y, Y]\rangle - t^3\langle [\Psi X, X], \Psi[\Psi Y, Y]\rangle. \end{aligned}$$

To calculate I_2, observe that if $|t| < \|\Psi\|^{-1}$ then

$$\Phi_t = \sum_{n=0}^{\infty} t^n \Psi^n,$$

with convergence in the space of endomorphisms of $\mathfrak{g}$ with the operator norm. From this formula we calculate

$$\begin{aligned}\frac{4}{3}I_2 &= \langle \Phi_t([X,Y]-tA+t^2B), [X,Y]-tA+t^2B\rangle \\ &= \sum_{n=0}^{\infty} t^n \langle \Psi^n[X,Y]-t\Psi^n A+t^2\Psi^n B, [X,Y]-tA+t^2B\rangle \\ &= \sum_{n=0}^{\infty} t^n \big(\langle \Psi^n[X,Y],[X,Y]\rangle - 2t\langle \Psi^n[X,Y],A\rangle \\ &\qquad\quad + t^2(\langle \Psi^n A, A\rangle + 2\langle \Psi^n[X,Y],B\rangle) \\ &\qquad\quad - 2t^3\langle \Psi^n A, B\rangle + t^4\langle \Psi^n B, B\rangle\big) \\ &= |[X,Y]|^2 + t(\langle \Psi[X,Y],[X,Y]\rangle - 2\langle [X,Y],A\rangle) \\ &\quad + t^2(\langle \Psi^2[X,Y],[X,Y]\rangle - 2\langle \Psi[X,Y],A\rangle + |A|^2 + 2\langle [X,Y],B\rangle) \\ &\quad + t^3(\langle \Psi^3[X,Y],[X,Y]\rangle - 2\langle \Psi^2[X,Y],A\rangle + \langle \Psi A, A\rangle \\ &\qquad\quad + 2\langle \Psi[X,Y],B\rangle - 2\langle A,B\rangle) \\ &\quad + \sum_{n=4}^{\infty} t^n \langle \Psi^{n-4}D, D\rangle.\end{aligned}$$

Combining the different terms proves the result. $\square$

The power series of $\kappa(t)$ would have been much messier if we were considering the unnormalized sectional curvature of X and Y with respect to h_t instead of the unnormalized sectional curvature of $\Phi_t^{-1}X$ and $\Phi_t^{-1}Y$. The value of twisting is apparent even at a purely computational level.

When $|t| < \|\Psi\|^{-1}$, we observe that

$$-\frac{3}{4}\sum_{n=4}^{\infty} t^n \langle \Psi^{n-4}D, D\rangle = -\frac{3}{4}t^4\langle \Phi_t D, D\rangle = -\frac{3}{4}t^4 \cdot |D|^2_{h_t}.$$

This proves that equation (4.1) holds for small t. Hence, to complete the proof of Theorem 4.1, we need only prove that $\kappa(t)$ and $|D|^2_{h_t}$ are analytic.

LEMMA 4.3. *The function $\kappa(t)$ is analytic on its domain of definition.*

Proof. Assume that t_0 is such that Φ_{t_0} corresponds to a metric on G. We show that κ is locally a power series at t_0. By (4.2) (Püttmann's formula), we clearly need only prove that $|[\Phi_t^{-1}X, \Phi_t^{-1}Y]|^2_{h_t}$ can be expressed as a power series near t_0. Because Ψ is h_0-self-adjoint, it can be diagonalized; say, $\Psi = \mathrm{diag}(a_1, \dots, a_d)$. We then have

$$\begin{aligned}\Phi_t &= \mathrm{diag}\left(\frac{1}{1-a_1t}, \dots, \frac{1}{1-a_dt}\right) \\ &= \mathrm{diag}\left(\frac{1}{1-a_it_0}\sum_{n=0}^{\infty}\left(\frac{a_i}{1-a_it_0}\right)^n (t-t_0)^n\right) \\ &= \Phi_{t_0}\sum_{n=0}^{\infty} \Phi_{t_0}^n \Psi^n (t-t_0)^n, \qquad (4.3)\end{aligned}$$

with convergence whenever $|t - t_0|$ is sufficiently small. We can use this expression for Φ_t together with the identity $\Phi_t^{-1} = I - t_0\Psi - (t - t_0)\Psi$ in order to expand $|[\Phi_t^{-1}X, \Phi_t^{-1}Y]|_{h_t}^2$ as a power series, as in the proof of Proposition 4.2. $\square$

Analyticity of $|D|_{h_t}^2$ also follows from equation (4.3), which completes the proof of Theorem 4.1.

Proof of Proposition 3.1. Assume X and Y commute. It is easy to see that $\alpha = \beta = 0$ and that δ equals 6 times the stated formula for $\kappa'''(0)$. All that remains to be shown is $\gamma = 0$. But the bi-invariance of h_0 and the Jacobi identity give the identity

$$\begin{aligned}
&\langle[\Psi X, Y], [X, \Psi Y]\rangle \\
&\quad = -\langle \Psi X, [[X, \Psi Y], Y]\rangle = \langle \Psi X, [[\Psi Y, Y], X] + [[Y, X], \Psi Y]\rangle \\
&\quad = \langle \Psi X, [[\Psi Y, Y], X]\rangle = -\langle[\Psi X, X], [\Psi Y, Y]\rangle, \qquad (4.4)
\end{aligned}$$

from which $\gamma = 0$ follows easily. $\square$

5. A General Rigidity Result

The next lemma is our primary tool for deriving rigidity statements about infinitesimally nonnegative variations. It plays an important role in Section 7, where we give a partial classification of the infinitesimally nonnegative endomorphisms of so(4).

LEMMA 5.1. *Assume that Ψ is infinitesimally nonnegative. Let $\mathfrak{p}_0$ be the eigenspace of Ψ corresponding to the smallest eigenvalue. If $X \in \mathfrak{p}_0$, $Y \in \mathfrak{g}$, and $[X, Y] = 0$, then $[X, \Psi Y] \in \mathfrak{p}_0$.*

Proof. Proposition 3.1 applied to X and Y yields

$$\tfrac{1}{6}\kappa'''(0) = a_0|[X, \Psi Y]|^2 - \langle[X, \Psi Y], \Psi[X, \Psi Y]\rangle,$$

where a_0 is the smallest eigenvalue. This is negative unless $[X, \Psi Y] \in \mathfrak{p}_0$. $\square$

The next proposition is a global version of Lemma 5.1. The argument used in its proof serves as the prototype for how we transform rigidity statements about infinitesimally nonnegative endomorphisms into rigidity statements about nonnegatively curved metrics.

PROPOSITION 5.2. *Assume that Φ is the matrix of a nonnegatively curved metric h. Let $\mathfrak{p}_0$ be the eigenspace of Φ corresponding to the smallest eigenvalue. If $X \in \mathfrak{p}_0$, $Y \in \mathfrak{g}$, and $[X, Y] = 0$, then $[X, \Phi^{-1}Y] \in \mathfrak{p}_0$.*

Proof. Let $\Psi = I - \Phi^{-1}$, so that $\Phi_t = (I - t\Psi)^{-1}$ is the unique inverse-linear path from h_0 to $h_1 = h$. By Theorem 1.1, Ψ must be infinitesimally nonnegative. Observe that Ψ and Φ have the same smallest eigenspace $\mathfrak{p}_0$. Proposition 5.1 now yields

$$[X, \Psi Y] = [X, (I - \Phi^{-1})Y] = -[X, \Phi^{-1}Y] \in \mathfrak{p}_0. \qquad \square$$

We note that this result can also be derived directly from Püttmann's formula.

6. Enlarging Subalgebras

Here we continue the discussion on enlarging subalgebras begun in Example 3.3. Let $H \subset G$ be a Lie subgroup of the Lie group G with Lie algebra $\mathfrak{h} \subset \mathfrak{g}$. For $Z \in \mathfrak{g}$, denote by $Z^{\mathfrak{h}}$ and $Z^{\mathfrak{p}}$ the projections of Z onto $\mathfrak{h}$ and its h_0-orthogonal complement $\mathfrak{p}$. Let $\Psi(Z) = Z^{\mathfrak{h}}$, so $\Phi_t = (I - t\Psi)^{-1}$ is the inverse-linear variation that gradually expands vectors in $\mathfrak{h}$ as t increases from 0. If $\mathfrak{h}$ is abelian, it is easy to use the formulas for the coefficients of the power series of $\kappa(t)$ in tandem with the analyticity of κ to prove

$$\kappa(t) = \frac{1}{4}|[X,Y]|^2 - \frac{3}{4}|[X,Y]^{\mathfrak{h}}|^2 \cdot \frac{t}{1-t} \quad (-\infty < t < 1). \tag{6.1}$$

From this formula we can show that enlarging $\mathfrak{h}$ by a factor of up to $\frac{4}{3}$ always preserves nonnegative curvature, a result that first appeared in [4]. In fact, the particularly nice form of $\kappa(t)$ allows us to prove a stronger statement as follows.

THEOREM 6.1. *Scaling the abelian subalgebra $\mathfrak{h} \subset \mathfrak{g}$ preserves nonnegative curvature if and only if no vector in $[\mathfrak{g},\mathfrak{g}]$ has the square of its norm expanded by more than $\frac{4}{3}$.*

Proof. By equation (6.1), the metric h_t is nonnegatively curved if and only if

$$|Z^{\mathfrak{h}}|^2 \cdot \frac{t}{1-t} \leq \frac{1}{3}|Z|^2 \tag{6.2}$$

holds for all $Z \in [\mathfrak{g},\mathfrak{g}]$. Since

$$|Z|^2_{h_t} = \langle \Phi_t Z, Z\rangle = \left\langle Z + \frac{t}{1-t}Z^{\mathfrak{h}}, Z\right\rangle = |Z|^2 + |Z^{\mathfrak{h}}|^2 \cdot \frac{t}{1-t},$$

it follows that inequality (6.2) is equivalent to requiring $|Z|^2_{h_t} \leq \frac{4}{3}|Z|^2$ to hold for all $Z \in [\mathfrak{g},\mathfrak{g}]$. □

If $[\mathfrak{g},\mathfrak{g}] \cap \mathfrak{h} \neq \{0\}$, then Theorem 6.1 means that $\mathfrak{h}$ can be scaled up by a factor of no more than $\frac{4}{3}$. At the other extreme, if $[\mathfrak{g},\mathfrak{g}] \perp \mathfrak{h}$ then $\mathfrak{h}$ can be expanded up by an arbitrary amount. This was already known, since if $\mathfrak{h}$ is orthogonal to $[\mathfrak{g},\mathfrak{g}]$ then $\mathfrak{h}$ is contained in the center of $\mathfrak{g}$. This rescaling then stays within the family of bi-invariant metrics on $\mathfrak{g}$.

When $\mathfrak{h}$ is not abelian, things are not quite so simple. In this case the power series simplifies to

$$\begin{aligned}\kappa(t) = \frac{1}{4}|[X,Y]|^2 - \frac{3}{4}|[X,Y]^{\mathfrak{h}}|^2 t + \frac{3}{4}|B|^2 t^2 \\ - \frac{1}{4}|B|^2 t^3 - \frac{3}{4}|[X^{\mathfrak{p}},Y^{\mathfrak{p}}]^{\mathfrak{h}}|^2 \cdot \frac{t^2}{1-t}.\end{aligned}$$

We can use this formula to classify exactly which subalgebras of $\mathfrak{g}$ can be enlarged a small amount while maintaining nonnegative curvature.

THEOREM 6.2. *Expanding the subalgebra $\mathfrak{h} \subset \mathfrak{g}$ by a small amount preserves nonnegative curvature if and only if there exists a constant c such that $|[X^{\mathfrak{h}}, Y^{\mathfrak{h}}]| \leq c \cdot |[X,Y]|$ for all $X, Y \in \mathfrak{g}$.*

We omit the lengthy but easy proof for the reason that we do not know if there are any interesting examples of subalgebras for which the latter condition holds. It clearly holds when $\mathfrak{h}$ is either abelian or an ideal of $\mathfrak{g}$ (or the sum of an ideal and an orthogonal abelian subalgebra), but it is already known that such subalgebras can be enlarged while maintaining nonnegative curvature.

7. Known Metrics on SO(4) with Nonnegative Curvature

Each known example of a left-invariant metric h with nonnegative curvature on $G = \mathrm{SO}(4)$ comes from Cheeger's construction. In this section, we catalog each known example in terms of the eigenvalue and eigenvector structure of the map Φ representing it with respect to a fixed bi-invariant metric h_0, meaning that $h(A,B) = h_0(\Phi A, B)$.

7.1. Product Metrics

The Lie algebra $\mathfrak{g} = \mathrm{so}(4)$ is a product $\mathfrak{g} = \mathfrak{g}_1 \oplus \mathfrak{g}_2$, where each factor is isomorphic to so(3). The two factors are h_0-orthogonal. If they are h-orthogonal, then h is a product metric on SO(4)'s double cover $S^3 \times S^3$. The classification of product metrics with nonnegative curvature reduces to the classification of left-invariant metrics with nonnegative curvature on SO(3), which was solved in [1]. Observe that, for any product metric, $\mathfrak{g}$ decomposes into three 2-dimensional Φ-invariant abelian subalgebras obtained by pairing eigenvectors from the two factors.

As for infinitesimal examples, if Ψ is a product (meaning that $\Psi(\mathfrak{g}_1) \subset \mathfrak{g}_1$ or equivalently that $\Psi(\mathfrak{g}_2) \subset \mathfrak{g}_2$) then the inverse-linear path $\Phi_t = (I - t\Psi)^{-1}$ it generates is through product metrics, which have nonnegative curvature for small t.

7.2. Torus Actions

Let $\{A_1, A_2, A_3\}$ and $\{B_1, B_2, B_3\}$ be h_0-orthonormal bases of $\mathfrak{g}_1$ and $\mathfrak{g}_2$, respectively. After scaling $\mathfrak{g}_1$ and $\mathfrak{g}_2$ by factors c and d (respectively), enlarging the abelian subalgebra $\tau = \mathrm{span}\{A_3, B_1\}$ by $\frac{4}{3}$, and then further altering the metric on τ via the remaining T^2-action on G, one obtains a nonnegatively curved metric h with matrix Φ of the form

$$\begin{pmatrix} c & 0 & 0 & 0 & 0 & 0 \\ 0 & c & 0 & 0 & 0 & 0 \\ 0 & 0 & a_1 & a_3 & 0 & 0 \\ 0 & 0 & a_3 & a_2 & 0 & 0 \\ 0 & 0 & 0 & 0 & d & 0 \\ 0 & 0 & 0 & 0 & 0 & d \end{pmatrix} \tag{7.1}$$

with respect to the basis $\{A_1, A_2, A_3, B_1, B_2, B_3\}$. In the final alteration, any right-invariant (and hence bi-invariant and flat) metric on T^2 can be used. The only restriction on Φ—which derives from this final alteration shrinking only vectors—is that the norm on τ determined by the matrix $\left(\begin{smallmatrix} a_1 & a_3 \\ a_3 & a_2 \end{smallmatrix}\right)$ is strictly bounded above by the norm determined by $\left(\begin{smallmatrix}(4/3)c & 0\\ 0 & (4/3)d\end{smallmatrix}\right)$. The limit points of such a metric are also nonnegatively curved. That is, we must consider the closure of the known examples, which transforms the strict inequality into a nonstrict one.

Observe that $\mathfrak{g}$ decomposes into three 2-dimensional Φ-invariant abelian subalgebras: one equals τ, and the other two are obtained by pairing vectors in $\mathfrak{g}_1$ with vectors in $\mathfrak{g}_2$. Observe also that any endomorphism Ψ with the matrix form of (7.1) will generate an inverse-linear variation $\Phi_t = (I - t\Psi)^{-1}$. These metrics will be nonnegatively curved for some interval $t \in [0, \varepsilon)$. The parameters $\{c, d, a_1, a_2, a_3\}$ defining Ψ are unrestricted, although they do determine ε.

7.3. S^3-Actions

Let $\tilde{h}$ denote the bi-invariant metric on $S^3 \times S^3$ obtained from h_0 by rescaling $\mathfrak{g}_1$ and $\mathfrak{g}_2$ by the factors a and b, respectively. Let g_R denote a right-invariant metric with nonnegative curvature on S^3 with eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\}$ and eigenvectors $\{e_1, e_2, e_3\}$. Define a metric h by

$$(S^3 \times S^3, h) = ((S^3 \times S^3, \tilde{h}) \times (S^3, g_R))/S^3,$$

where S^3 acts diagonally. Consider the basis

$$\mathfrak{g} = \mathfrak{g}_1 \oplus \mathfrak{g}_2 = \operatorname{span}\{A_1, A_2, A_3\} \oplus \operatorname{span}\{B_1, B_2, B_3\},$$

where $A_i = (e_i, 0)$ and $B_i = (0, e_i)$. Let $V_i = \operatorname{span}\{A_i, B_i\}$, which for each i is a 2-dimensional abelian subalgebra of $\mathfrak{g}$. Notice that the three V_i are mutually orthogonal with respect to h_0, $\tilde{h}$, and h. It therefore suffices to describe h in terms of h_0 separately on each V_i.

Thus, the matrix representing $\tilde{h}$ in terms of h_0 on V_i in the basis $\{A_i, B_i\}$ is $M_i = \left(\begin{smallmatrix}a & 0\\ 0 & b\end{smallmatrix}\right)$, and the matrix representing h in terms of $\tilde{h}$ in the basis $\{A_i + B_i, bA_i - aB_i\}$ is $N_i = \left(\begin{smallmatrix}t_i & 0\\ 0 & 1\end{smallmatrix}\right)$, where $t_i = \lambda_i/(1+\lambda_i)$. If we let $T = \left(\begin{smallmatrix}1 & b\\ 1 & -a\end{smallmatrix}\right)$ be the change of basis matrix, then the matrix we seek that represents h in terms of h_0 on V_i in the basis $\{A_i, B_i\}$ is

$$\Phi_i = M_i(TN_iT^{-1}) = \frac{1}{a+b}\begin{pmatrix} a(b+at_i) & ab(t_i-1)\\ ab(t_i-1) & b(a+bt_i)\end{pmatrix}. \tag{7.2}$$

In summary, $\mathfrak{g}$ decomposes into the three Φ-invariant 2-dimensional abelian subalgebras $\{V_1, V_2, V_3\}$. However, with only the five parameters $\{a, b, t_1, t_2, t_3\}$ under our control and with restrictions on the values of t, we do not attain the full 9-parameter family of metrics for which the subalgebras $\{V_1, V_2, V_3\}$ are Φ-invariant.

Infinitesimal examples have the form $\Psi := I - \Phi^{-1}$ with Φ in the form of equation (7.2). A calculation shows that all such matrices have the form $\Psi = \operatorname{diag}(\Psi_1, \Psi_2, \Psi_3)$, where

$$\Psi_i = \begin{pmatrix}\alpha & 0\\ 0 & \beta\end{pmatrix} - \frac{1}{2\lambda_i}\begin{pmatrix}1 & 1\\ 1 & 1\end{pmatrix}. \tag{7.3}$$

The parameters α, β are free, but the parameters $\{\lambda_1, \lambda_2, \lambda_3\}$ are restricted to be eigenvalues of a nonnegatively curved metric on SO(3).

8. Infinitesimal Rigidity for SO(4)

In this section, we assume that $G = \mathrm{SO}(4)$ and that $\Psi\colon \mathfrak{g} \to \mathfrak{g}$ is infinitesimally nonnegative, and we prove rigidity results for Ψ. In the next section, we translate these infinitesimal rigidity results into global theorems.

Recall that $\mathfrak{g} = \mathrm{so}(4) = \mathfrak{g}_1 \oplus \mathfrak{g}_2$ is a product; $X \in \mathfrak{g}$ is called *regular* if it has nonzero projections onto both $\mathfrak{g}_1$ and $\mathfrak{g}_2$ and is called *singular* otherwise. We give G the most natural bi-invariant metric h_0, so that any orthonormal bases of the factors $\mathfrak{g}_1$ and $\mathfrak{g}_2$ behave like the quaternions $\{\mathbf{i},\mathbf{j},\mathbf{k}\}$ with respect to their Lie bracket structure. We will show in Section 10 that no loss of essential information results from restricting ourselves to working only with this bi-invariant metric.

The previous section classified the known possibilities of Ψ into three types originating from (1) products, (2) torus actions, and (3) S^3-actions. In the first two cases, Ψ has a nonzero singular eigenvector; in the third case, it does not.

THEOREM 8.1. *If Ψ has a nonzero singular eigenvector, then either Ψ is a product or Ψ has the form of* (7.1). *In either case, h_t is a family of known examples with nonnegative curvature for sufficiently small t.*

If Ψ has no nonzero singular eigenvectors, we hypothesize that Ψ is a known example coming from an S^3-action. A first step in this direction is to locate three Ψ-invariant abelian subalgebras. The following theorem falls just short of this goal.

THEOREM 8.2. *There are orthonormal bases $\{A_1, A_2, A_3\}$ and $\{B_1, B_2, B_3\}$ of the two factors of $\mathfrak{g} = \mathfrak{g}_1 \oplus \mathfrak{g}_2$ such that, with respect to the basis $\{A_1, B_1, A_2, B_2, A_3, B_3\}$, Ψ has the form*

$$\Psi = \begin{pmatrix} a_1 & a_3 & 0 & 0 & 0 & 0 \\ a_3 & a_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & b_1 & b_3 & \lambda & 0 \\ 0 & 0 & b_3 & b_2 & 0 & \mu \\ 0 & 0 & \lambda & 0 & c_1 & c_3 \\ 0 & 0 & 0 & \mu & c_3 & c_2 \end{pmatrix}.$$

We conjecture that $\lambda = \mu = 0$, which means that $\mathfrak{g}$ decomposes into three orthogonal Ψ-invariant abelian subalgebras (as it should). Even granting this conjecture, there remains the work of reducing this 9-parameter family to the 5-parameter family of known examples from equation (7.3). This appears to be a computationally difficult problem.

The remainder of this section is devoted to proving Theorems 8.1 and 8.2. We begin with a weak version of Theorem 8.1. Recall that $\mathfrak{p}_0$ denotes the eigenspace corresponding to the smallest eigenvalue, a_0, of Ψ.

LEMMA 8.3. *If $\mathfrak{p}_0$ contains a nonzero singular vector, then either Ψ is a product or Ψ has the form of* (7.1).

Proof. Without loss of generality, assume there exists a nonzero vector $X_1 \in \mathfrak{g}_1 \cap \mathfrak{p}_0$. Assume that Ψ is not a product, so there exists a $\hat{Y} \in \mathfrak{g}_2$ such that $\Psi\hat{Y}$ has a nonzero projection, X_2, onto $\mathfrak{g}_1$. Observe that X_1 and X_2 are orthogonal because

$$\langle X_1, X_2 \rangle = \langle X_1, \Psi\hat{Y} \rangle = \langle \Psi X_1, \hat{Y} \rangle = a_0 \langle X_1, \hat{Y} \rangle = 0.$$

Let $X_3 = [X_1, \Psi\hat{Y}] \in \mathfrak{g}_1$, which (by Lemma 5.1) lies in $\mathfrak{p}_0$ and so $\operatorname{span}\{X_1, X_3\} \subset \mathfrak{p}_0$. Let Y_2 be the projection of ΨX_2 onto $\mathfrak{g}_2$, which is a nonzero vector by the self-adjoint property of Ψ. Complete $\{Y_2\}$ to an orthogonal basis $\{Y_1, Y_2, Y_3\}$ of $\mathfrak{g}_2$, ordered so that their bracket structure is like that of $\{\mathbf{i}, \mathbf{j}, \mathbf{k}\}$. Notice that $\Psi(\operatorname{span}\{Y_1, Y_3\}) \subset \mathfrak{g}_2$ (again by the self-adjoint property of Ψ). In summary, after scaling all the vectors to unit length, we have an orthonormal basis

$$\mathfrak{g} = \mathfrak{g}_1 \oplus \mathfrak{g}_2 = \operatorname{span}\{X_1, X_2, X_3\} \oplus \operatorname{span}\{Y_1, Y_2, Y_3\}$$

with $\operatorname{span}\{X_1, X_3\} \subset \mathfrak{p}_0$, and $\Psi X_2 = cY_2 + \lambda X_2$ (for some $c, \lambda \in \mathbb{R}$ with $c \neq 0$) and $\Psi(\operatorname{span}\{Y_1, Y_3\}) \subset \mathfrak{g}_2$.

Applying Proposition 3.1 to the vectors X_2 and Y_1 gives

$$\begin{aligned}
\kappa'''(0) &= 6\langle [\Psi X_2, Y_1], [\Psi X_2, \Psi Y_1] \rangle - 6\langle [\Psi X_2, Y_1], \Psi[\Psi X_2, Y_1] \rangle \\
&= 6\langle [cY_2, Y_1], [cY_2, \Psi Y_1] \rangle - 6\langle [cY_2, Y_1], \Psi[cY_2, Y_1] \rangle \\
&= -6c^2 \langle Y_3, [Y_2, \Psi Y_1] \rangle - 6c^2 \langle Y_3, \Psi Y_3 \rangle \geq 0.
\end{aligned}$$

Observe that

$$\begin{aligned}
\langle Y_3, [Y_2, \Psi Y_1] \rangle &= \langle Y_3, [Y_2, \text{projection of } \Psi Y_1 \text{ onto } Y_1] \rangle \\
&= \langle Y_3, [Y_2, \langle \Psi Y_1, Y_1 \rangle Y_1] \rangle \\
&= -\langle \Psi Y_1, Y_1 \rangle,
\end{aligned}$$

from which we conclude

$$\langle Y_1, \Psi Y_1 \rangle \geq \langle Y_3, \Psi Y_3 \rangle.$$

Similarly, applying Proposition 3.1 to the vectors X_2 and Y_3 yields the reverse inequality, so

$$\langle Y_1, \Psi Y_1 \rangle = \langle Y_3, \Psi Y_3 \rangle.$$

Replacing Y_1 and Y_3 with any other orthonormal basis of $\operatorname{span}\{Y_1, Y_3\}$ yields the same conclusion. In other words, for any angle θ, if we set $a = \cos(\theta)$ and $b = \sin(\theta)$ then

$$\langle aY_1 + bY_3, \Psi(aY_1 + bY_3) \rangle = \langle bY_1 - aY_3, \Psi(bY_1 - aY_3) \rangle.$$

This implies that $\langle Y_1, \Psi Y_3 \rangle = \langle \Psi Y_1, Y_3 \rangle = 0$. The linear map from $\operatorname{span}\{Y_1, Y_3\}$ to $\mathbb{R}$ sending $Y \mapsto \langle \Psi Y, Y_2 \rangle$ has a nonzero vector in its kernel. Assume without loss of generality that Y_1 is in this kernel, and note that Y_1 is an eigenvector of Ψ.

In the ordered basis $\{X_1, X_2, X_3, Y_1, Y_2, Y_3\}$, we thus far have

$$\Psi = \begin{pmatrix} a_0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \lambda & 0 & 0 & c & 0 \\ 0 & 0 & a_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \beta & 0 & 0 \\ 0 & c & 0 & 0 & \gamma & s \\ 0 & 0 & 0 & 0 & s & \beta \end{pmatrix}.$$

Applying our $\kappa'''(0)$ formula to $X = X_2$ and to $Y = aY_2 + bY_3$ gives

$$\kappa'''(0) = 6bc^2(as + b\beta) - 6b^2c^2\beta = 6bc^2as.$$

Since $\kappa'''(0) \geq 0$ for all choices of $\{a, b\}$ and since $c \neq 0$, it follows that $s = 0$. After re-ordering the basis, Ψ has the form of (7.1). $\square$

THEOREM 8.4. *The eigenspace $\mathfrak{p}_0$ contains a nonzero vector that belongs to a Ψ-invariant 2-dimensional abelian subalgebra of $\mathfrak{g}$.*

Proof. If $\mathfrak{p}_0$ contains a nonzero singular vector then the conclusion follows easily from Lemma 8.3, so we assume that this is not the case. When $A = (A_1, A_2) \in \mathfrak{g} = \mathfrak{g}_1 \oplus \mathfrak{g}_2$ is regular, let

$$\bar{A} = \left(\frac{|A_2|}{|A_1|}A_1, -\frac{|A_1|}{|A_2|}A_2\right),$$

which commutes with A, is orthogonal to A, and has the same norm as A.

The proof is indirect. We assume for each $A \in \mathfrak{p}_0$ that $\operatorname{span}\{A, \bar{A}\}$ is not Ψ-invariant, and we then derive a contradiction.

Let $A \in \mathfrak{p}_0$ be of unit length. Since Ψ is self-adjoint, we know that $\Psi\bar{A}$ is orthogonal to A. Observe that $\bar{A}$ is not an eigenvector of Ψ; if it were, then $\operatorname{span}\{A, \bar{A}\}$ would be an invariant abelian subalgebra. Therefore, $[A, \Psi\bar{A}]$ is nonzero. Let B be the unit-length vector in the direction of $[A, \Psi\bar{A}]$. By Lemma 5.1, $B \in \mathfrak{p}_0$. Notice that B is orthogonal to A and to $\bar{A}$.

So far we know that $\dim(\mathfrak{p}_0) \geq 2$. Clearly $\dim(\mathfrak{p}_0) \leq 3$, because it contains no nonzero singular vectors and hence intersects $\mathfrak{g}_1$ and $\mathfrak{g}_2$ trivially. We wish to prove $\dim(\mathfrak{p}_0) = 2$. Suppose to the contrary that $\dim(\mathfrak{p}_0) = 3$. Consider the map from $\mathfrak{p}_0$ to $\mathfrak{p}_0$ defined as

$$Z \mapsto [Z, \Psi\bar{Z}].$$

By the foregoing arguments, this map sends each unit-length $Z \in \mathfrak{p}_0$ to a nonzero vector in $\mathfrak{p}_0$ that is orthogonal to Z. This map therefore induces a smooth non-vanishing vector field on the unit 2-sphere in $\mathfrak{p}_0$, which is a contradiction. Thus, $\dim(\mathfrak{p}_0) = 2$. Here A and B play symmetric roles in that $[B, \Psi\bar{B}]$ is parallel to A (because it lies in $\mathfrak{p}_0$ and is perpendicular to B), and A is orthogonal to B and to $\bar{B}$.

Choose unit-length vectors $C_1 \in \mathfrak{g}_1$ and $C_2 \in \mathfrak{g}_2$ such that $\{A, \bar{A}, B, \bar{B}, C_1, C_2\}$ is an orthonormal basis of $\mathfrak{g}$. For $i = 1, 2$, the $\mathfrak{g}_i$-components of $\{A, B, C_i\}$ form an orthogonal basis of $\mathfrak{g}_i$. The C_i can be chosen so that these orthogonal bases are oriented and so, after normalizing, they act like $\{\mathbf{i},\mathbf{j},\mathbf{k}\}$ with respect to their Lie bracket structure. For purposes of calculating Lie brackets in this basis, we lose no generality in assuming that, for some $a, b \in (0,1)$,

$$\begin{aligned} A &= (a\mathbf{i}, \sqrt{1-a^2}\mathbf{i}), \quad & B &= (b\mathbf{j}, \sqrt{1-b^2}\mathbf{j}), \quad & C_1 &= (\mathbf{k}, 0); \\ \bar{A} &= (\sqrt{1-a^2}\mathbf{i}, -a\mathbf{i}), \quad & \bar{B} &= (\sqrt{1-b^2}\mathbf{j}, -b\mathbf{j}), \quad & C_2 &= (0, \mathbf{k}). \end{aligned} \tag{8.1}$$

Notice that $\langle \Psi\bar{A}, \bar{B}\rangle = \langle \Psi\bar{B}, \bar{A}\rangle = 0$, because if $\Psi\bar{A}$ had a nonzero $\bar{B}$-component then $[A, \Psi\bar{A}]$ would have nonzero C_1- and C_2-components.

In the basis $\{A, \bar{A}, B, \bar{B}, C_1, C_2\}$, Ψ has the form

$$\Psi = \begin{pmatrix} a_0 & 0 & 0 & 0 & 0 & 0 \\ 0 & p & 0 & 0 & \alpha_1 & \alpha_2 \\ 0 & 0 & a_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & q & \beta_1 & \beta_2 \\ 0 & \alpha_1 & 0 & \beta_1 & f_1 & f_2 \\ 0 & \alpha_2 & 0 & \beta_2 & f_2 & f_3 \end{pmatrix}. \tag{8.2}$$

There are a few obvious restrictions among the variables determining Ψ. For example, since $[A, \Psi\bar{A}]$ is parallel to B and since $[B, \Psi\bar{B}]$ is parallel to A, we have

$$\frac{\alpha_1}{\alpha_2} = \frac{\beta_2}{\beta_1} = \frac{b\sqrt{1-a^2}}{a\sqrt{1-b^2}} \tag{8.3}$$

and obtain

$$\Psi = \begin{pmatrix} a_0 & 0 & 0 & 0 & 0 & 0 \\ 0 & p & 0 & 0 & \alpha & \alpha\cdot s \\ 0 & 0 & a_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & q & \beta\cdot s & \beta \\ 0 & \alpha & 0 & \beta\cdot s & f_1 & f_2 \\ 0 & \alpha\cdot s & 0 & \beta & f_2 & f_3 \end{pmatrix}, \tag{8.4}$$

where $s = a\sqrt{1-b^2}/b\sqrt{1-a^2} > 0$ and $\alpha, \beta \neq 0$.

Using Lemma 5.1, we can now prove that $s = 1$ and consequently $a = b$. Indeed, for every $Z \in \operatorname{span}\{A, B\}$, we have $[Z, \Psi\bar{Z}] \in \operatorname{span}\{A, B\}$. In particular, let $Z_t = (\cos t)A + (\sin t)B$, so

$$\begin{aligned} \bar{Z}_t = \big(&f(t)(a\cos(t)\mathbf{i} + b\sin(t)\mathbf{j}), \\ &-(1/f(t))(\sqrt{1-a^2}\cos(t)\mathbf{i} + \sqrt{1-b^2}\sin(t)\mathbf{j})\big), \end{aligned}$$

where

$$f(t) = \sqrt{\frac{(1-a^2)\cos^2(t) + (1-b^2)\sin^2(t)}{a^2\cos^2(t) + b^2\sin^2(t)}}.$$

We will use that the following vector lies in $\operatorname{span}\{A, B\}$:

$$
\begin{aligned}
Q &= \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}[Z_t, \Psi\bar{Z}_t] \\
&= [B, \Psi\bar{A}] + [A, \Psi(f'(0)a\mathbf{i} + f(0)b\mathbf{j}, -g'(0)\sqrt{1-a^2}\mathbf{i} - g(0)\sqrt{1-b^2}\mathbf{j})] \\
&= [B, \Psi\bar{A}] + [A, \Psi(f(0)b\mathbf{j}, -g(0)\sqrt{1-b^2}\mathbf{j})] \\
&= [B, \Psi\bar{A}] + \left[A, \Psi\left(\frac{b\sqrt{1-a^2}}{a}\mathbf{j}, -\frac{a\sqrt{1-b^2}}{\sqrt{1-a^2}}\mathbf{j}\right)\right] \\
&= [B, \Psi\bar{A}] + [A, \Psi(\sqrt{1-b^2}\cdot s^{-1}\mathbf{j}, -b\cdot s\mathbf{j})].
\end{aligned}
$$

In particular, Q is perpendicular to $\bar{A}$ and so

$$
\begin{aligned}
0 = \langle Q, \bar{A}\rangle &= \langle [B, \Psi\bar{A}], \bar{A}\rangle + \langle [A, \Psi(\sqrt{1-b^2}\cdot s^{-1}\mathbf{j}, -b\cdot s\mathbf{j})], \bar{A}\rangle \\
&= \langle [B, \Psi\bar{A}], \bar{A}\rangle = -\langle \Psi\bar{A}, [B, \bar{A}]\rangle \\
&= -\langle p\bar{A} + (\alpha\mathbf{k}, \alpha s\mathbf{k}), [(b\mathbf{j}, \sqrt{1-b^2}\mathbf{j}), (\sqrt{1-a^2}\mathbf{i}, -a\mathbf{i})]\rangle \\
&= -\langle p\bar{A} + (\alpha\mathbf{k}, \alpha s\mathbf{k}), (-b\sqrt{1-a^2}\mathbf{k}, a\sqrt{1-b^2}\mathbf{k})\rangle \\
&= \alpha b\sqrt{1-a^2} - s\alpha a\sqrt{1-b^2},
\end{aligned}
$$

which implies $s = b\sqrt{1-a^2}/a\sqrt{1-b^2} = s^{-1}$. It follows that $s = 1$ and hence $a = b$. That the orthogonal projection of Q onto $\operatorname{span}\{C_1, C_2\}$ is zero is now equivalent to

$$
p(-b\sqrt{1-a^2}\mathbf{k}, a\sqrt{1-b^2}\mathbf{k}) + q(a\sqrt{1-b^2}\mathbf{k}, -b\sqrt{1-a^2}\mathbf{k}) = 0. \tag{8.5}
$$

Since $a = b$, this implies that $q = p$. Therefore,

$$
\Psi = \begin{pmatrix} a_0 & 0 & 0 & 0 & 0 & 0 \\ 0 & p & 0 & 0 & \alpha & \alpha \\ 0 & 0 & a_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & p & \beta & \beta \\ 0 & \alpha & 0 & \beta & f_1 & f_2 \\ 0 & \alpha & 0 & \beta & f_2 & f_3 \end{pmatrix}. \tag{8.6}
$$

Since $a = b$, it is easy to see that $[A, \bar{B}] + [B, \bar{A}] = 0$. This implies that $V_1 = \beta\bar{A} - \alpha\bar{B}$ commutes with $V_2 = \beta A - \alpha B$. Since $V_2 \in \mathfrak{p}_0$ and since V_1 is an eigenvector of Ψ (with eigenvalue p), it follows that $\operatorname{span}\{V_1, V_2\}$ is a Ψ-invariant 2-dimensional abelian subalgebra of $\mathfrak{g}$ containing a nonzero vector in $\mathfrak{p}_0$. This is a contradiction. $\square$

Proof of Theorem 8.2. By Theorem 8.4, there exists a Ψ-invariant abelian subalgebra of $\mathfrak{g}$ that is spanned by some $A_1 \in \mathfrak{g}_1$ and some $B_1 \in \mathfrak{g}_2$. Let V_1 denote the orthogonal complement of A_1 in $\mathfrak{g}_1$ and let V_2 denote the orthogonal complement of B_1 in $\mathfrak{g}_2$.

Let $\pi_1\colon \mathfrak{g} \to \mathfrak{g}_1$ and $\pi_2\colon \mathfrak{g} \to \mathfrak{g}_2$ denote the projections. Define $T_1\colon V_1 \to V_2$ as $T_1 = \pi_2 \circ \Psi|_{V_1}$ and define $T_2\colon V_2 \to V_1$ as $T_2 = \pi_1 \circ \Psi|_{V_2}$. Observe that for all $A \in V_1$ and $B \in V_2$,

$$\langle T_1 A, B\rangle = \langle \Psi A, B\rangle = \langle A, \Psi B\rangle = \langle A, T_2 B\rangle.$$

Let S^1 denote the circle of unit-length vectors in V_1, and let $R\colon S^1 \to S^1$ denote a 90° rotation. Define $F\colon S^1 \to \mathbb{R}$ by $F(A) = \langle T_1(A), T_1(R(A))\rangle$. For all $A \in S^1$,

$$F(R(A)) = \langle T_1(R(A)), T_1(-A)\rangle = -F(A).$$

This implies that there exists an $A_2 \in S^1$ such that $F(A_2) = 0$. Let $A_3 = R(A_2)$. First suppose T_1 (and hence also T_2) is nonsingular. Define $B_2 = T_1(A_2)/|T_1(A_2)|$ and $B_3 = T_1(A_3)/|T_1(A_3)|$. The equality $F(A_2) = 0$ immediately implies that B_2 and B_3 are orthogonal and that $T_2(B_2) \parallel A_2$ and $T_2(B_3) \parallel A_3$. Thus, the basis $\{A_1, A_2, A_3, B_1, B_2, B_3\}$ satisfies the conclusion of the theorem.

If T_1 (and hence also T_2) is singular, then the arbitrary orthonormal bases $\{A_2, A_3\}$ of V_1 and $\{B_2, B_3\}$ of V_2 work provided $A_2 \in \ker(T_1)$ and $B_2 \in \ker(T_2)$. □

Our final proof in this section is due to Nela Vukmirovic and Zachary Madden.

Proof of Theorem 8.1. Choose bases $\{A_1, A_2, A_3\}$ of $\mathfrak{g}_1$ and $\{B_1, B_2, B_3\}$ of $\mathfrak{g}_2$ so that Ψ has the matrix form of Theorem 8.2. With respect to the ordering $\{A_3, A_2, A_1, B_1, B_2, B_3\}$, Ψ then has the form

$$\Psi = \begin{pmatrix} c_1 & \lambda & 0 & 0 & 0 & c_3 \\ \lambda & b_1 & 0 & 0 & b_3 & 0 \\ 0 & 0 & a_1 & a_3 & 0 & 0 \\ 0 & 0 & a_3 & a_2 & 0 & 0 \\ 0 & b_3 & 0 & 0 & b_2 & \mu \\ c_3 & 0 & 0 & 0 & \mu & c_2 \end{pmatrix}.$$

If $a_3 = 0$ then the result follows from Lemma 8.3, so we can assume $a_3 \neq 0$. To complete the proof, we show that $c_1 = b_1$, $b_2 = c_2$, and $\lambda = \mu = b_3 = c_3 = 0$, which puts Ψ into the form of (7.1). The hypothesis that Ψ has a nonzero singular eigenvector implies that $b_3 = 0$ or $c_3 = 0$. Without loss of generality, assume $b_3 = 0$. Henceforth, the value $\kappa'''(0)$ with respect to the commuting pair $X = \alpha_1 A_1 + \alpha_2 A_2 + \alpha_3 A_3$ and $Y = \beta_1 B_1 + \beta_2 B_2 + \beta_3 B_3$ will be denoted by $[\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3]$. These 6-tuples are easily expanded using Maple or Mathematica.

First,

$$[0, \pm 1, 1, 1, 0, 0] = c_3^2(a_2 - b_2) \pm 4a_3^2\lambda \geq 0.$$

However, since $[0,0,1,0,1,0] + [0,0,1,0,0,1] = c_3^2(b_2 - a_2) \geq 0$, we deduce that $\lambda = 0$ and consequently $c_3^2(b_2 - a_2) = 0$. Similarly,

$$[1, 0, 0, 0, \pm 1, 1] = c_3^2(a_1 - b_1) \pm 4a_3^2\mu \geq 0.$$

But $[1,0,0,0,1,0] + [1,0,0,0,0,1] = c_3^2(b_1 - a_1) \geq 0$, so it follows that $\mu = 0$ and $c_3^2(b_1 - a_1) = 0$.

Furthermore, the inequalities $[0,1,0,1,0,0] \geq 0$ and $[0,0,1,1,0,0] \geq 0$ give (respectively) the plus and minus versions of the inequality $\pm a_3^2(b_1 - c_1) \geq 0$. Analogously, after examining $[1,0,0,0,1,0]$ and $[1,0,0,0,0,1]$ we conclude that $\pm a_3^2(b_2 - c_2) \geq 0$. Because a_3 is nonzero, we obtain $b_1 = c_1$ and $b_2 = c_2$.

All that remains to be shown is that $c_3 = 0$. If $c_3 \neq 0$, then $a_1 = b_1$ and $a_2 = b_2$. By considering $[1,1,1,1,1,1]$, $[1,1,1,-1,1,1]$, $[1,1,1,1,-1,1]$, and $[1,1,1,1,1,-1]$, we deduce that $\pm a_3^2 c_3 \geq 0$, which implies $c_3 = 0$. Thus, Ψ has the form of (7.1). □

9. Global Rigidity for SO(4)

The previous section partially classified the infinitesimally nonnegative endomorphisms for $G = \mathrm{SO}(4)$. We now translate these infinitesimal results into a partial classification of the nonnegatively curved left-invariant metrics on SO(4).

Assume $G = \mathrm{SO}(4)$. Let Φ be the matrix for a nonnegatively curved left-invariant metric h on G. The variation $\Phi_t = (I - t\Psi)^{-1}$ satisfies $\Phi_1 = \Phi$ as long as we choose $\Psi = I - \Phi^{-1}$. By Theorem 1.1, this variation is through nonnegatively curved metrics and so Ψ is infinitesimally nonnegative. We will apply restrictions on Ψ from Section 8 in order to prove rigidity theorems about Φ.

First, we prove a global analogue of Theorem 8.1. This theorem implies Theorem 1.2.

THEOREM 9.1. *If Φ has a singular eigenvector, then either h is a product metric or h comes from a torus action. In either case, h is a known example of a metric of nonnegative curvature.*

Proof. Since Φ has a singular eigenvector, so does Ψ. According to Theorem 8.1, either Ψ is a product or Ψ can be written in the form of (7.1). If Ψ is a product then Φ is a product, which means that h is a product metric. If instead Ψ has the form of (7.1), then so does Φ.

Assume Φ has the form of (7.1); we must prove that Φ satisfies the $\frac{4}{3}$-restriction shared by all known examples. Permuting some basis vectors if necessary, we may assume that A_1, A_2, A_3 and B_1, B_2, B_3 behave like the quaternions $\mathbf{i}, \mathbf{j}, \mathbf{k}$ with respect to their Lie bracket structure. Denote by $\tilde{h}$ the metric on τ corresponding to the matrix

$$\begin{pmatrix} \frac{4}{3} \cdot c & 0 \\ 0 & \frac{4}{3} \cdot d \end{pmatrix}.$$

We must prove that

$$|\alpha A_3 + \beta B_1|_h^2 \leq |\alpha A_3 + \beta B_1|_{\tilde{h}}^2$$

holds for all $\alpha, \beta \in \mathbb{R}$.

Consider the unnormalized sectional curvature of the vectors $\alpha A_1 + \beta B_2$ and $A_2 + B_3$ with respect to h. We have

$$[\Phi(\alpha A_1 + \beta B_2), A_2 + B_3] = \alpha c A_3 + \beta d B_1,$$

$$[\alpha A_1 + \beta B_2, \Phi(A_2 + B_3)] = \alpha c A_3 + \beta d B_1,$$

$$[\alpha A_1 + \beta B_2, A_2 + B_3] = \alpha A_3 + \beta B_1;$$

thus, by Püttmann's formula,

$$\begin{aligned} k_h(\alpha A_1 + \beta B_2, A_2 + B_3) &= \langle \alpha c A_3 + \beta d B_1, \alpha A_3 + \beta B_1 \rangle - \tfrac{3}{4}|\alpha A_3 + \beta B_1|_h^2 \\ &= \tfrac{3}{4}(|\alpha A_3 + \beta B_1|_{\tilde{h}}^2 - |\alpha A_3 + \beta B_1|_h^2). \end{aligned}$$

Since h is nonnegatively curved, this proves the required inequality. □

Similarly, we obtain a global version of Theorem 8.2 as follows.

THEOREM 9.2. *There are orthonormal bases* $\{A_1, A_2, A_3\}$ *and* $\{B_1, B_2, B_3\}$ *of the two factors of* $\mathfrak{g} = \mathfrak{g}_1 \oplus \mathfrak{g}_2$ *such that, with respect to the basis* $\{A_1, B_1, A_2, B_2, A_3, B_3\}$, Φ *has the form*

$$\Phi = \begin{pmatrix} a_1 & a_3 & 0 & 0 & 0 & 0 \\ a_3 & a_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & b_1 & b_3 & \lambda & 0 \\ 0 & 0 & b_3 & b_2 & 0 & \mu \\ 0 & 0 & \lambda & 0 & c_1 & c_3 \\ 0 & 0 & 0 & \mu & c_3 & c_2 \end{pmatrix}.$$

In particular, $\mathfrak{g}$ *has a* 2*-dimensional* Φ*-invariant abelian subalgebra.*

Proof. By Theorem 8.4, $\mathfrak{g}$ has a 2-dimensional Ψ-invariant abelian subalgebra. This subalgebra is also Φ-invariant. The result follows by mimicking the proof of Theorem 8.2. □

10. Changing the Initial Bi-invariant Metric

Let h_0 be a fixed bi-invariant metric, and consider a second bi-invariant metric h_1. If h is a nonnegatively curved left-invariant metric then, according to Theorem 1.1, the unique inverse-linear paths from h_0 to h and from h_1 to h are through nonnegatively curved metrics. We can view this as saying that the inverse-linear path from h_0 to h is through nonnegatively curved metrics if and only if the inverse-linear path from h_1 to h is.

In light of this result, it is natural to ask whether the inverse-linear path from h_0 to h is infinitesimally nonnegative if and only if the inverse-linear path from h_1 to h is. The main result of this section is an affirmative answer, which shows that the concept of "infinitesimally nonnegative" is independent of the starting bi-invariant metric. This means that, when classifying the infinitesimally nonnegative endomorphisms of $\mathfrak{g}$ with respect to a bi-invariant metric, the choice of bi-invariant metric is essentially irrelevant.

THEOREM 10.1. *The inverse-linear path from* h_0 *to* h *is infinitesimally nonnegative if and only if the inverse-linear path from* h_1 *to* h *is.*

For the proof of this theorem, let M be the matrix of h_1 with respect to h_0, let Φ be the matrix of h with respect to h_0, let Θ be the matrix of h with respect to h_1, and put $\Psi = I - \Phi^{-1}$ and $\Upsilon = I - \Theta^{-1}$. Theorem 10.1 is a consequence of the following result.

PROPOSITION 10.2. *For any commuting vectors X and Y in $\mathfrak{g}$,*

$$D^{\Upsilon}_{X,Y} = D^{\Psi}_{MX,MY} \quad \text{and} \quad \delta^{\Upsilon,h_1}_{X,Y} = \delta^{\Psi,h_0}_{MX,MY};$$

here, for instance, $\delta^{\Upsilon,h_1}_{MX,MY}$ denotes the coefficient δ in the power series of the function $\kappa(t)$ defined with respect to the endomorphism Ψ, the bi-invariant metric h_0, and the commuting pair of vectors MX, MY. Hence Ψ is infinitesimally nonnegative if and only if Υ is.

Proof. Write

$$\mathfrak{g} = \mathfrak{g}_1 \oplus \cdots \oplus \mathfrak{g}_r \oplus Z(\mathfrak{g}),$$

where the $\mathfrak{g}_i$ are simple subalgebras and $Z(\mathfrak{g})$ is the center of $\mathfrak{g}$. The simple subalgebras have unique bi-invariant metrics up to a scalar multiple, any choice of inner product on $Z(\mathfrak{g})$ is bi-invariant, and all bi-invariant metrics on $\mathfrak{g}$ arise as product metrics from this decomposition. We can diagonalize M with respect to a basis respecting the decomposition, and M will have a single eigenvalue corresponding to each simple factor $\mathfrak{g}_i$ and arbitrary eigenvalues on basis vectors in $Z(\mathfrak{g})$. This allows us to factor $M = M_1 \cdots M_s$, where each M_i scales an ideal of $\mathfrak{g}$ and leaves its orthogonal complement fixed. By induction, it suffices to prove the preceding formulas for $M = M_1$, where M acts on $\mathfrak{g}$ by $Z \mapsto \lambda Z^{\mathfrak{h}} + Z^{\mathfrak{k}}$ for some $\lambda > 0$ and where $\mathfrak{h}, \mathfrak{k}$ are ideals of $\mathfrak{g}$ with $\mathfrak{g} = \mathfrak{h} \oplus \mathfrak{k}$. This special case follows from a long straightforward calculation using the definitions of D and δ. □

We conjecture that the formulas of this proposition are a special case of a formula relating $\kappa^{\Upsilon,h_1}_{X,Y}(t)$ to $\kappa^{\Psi,h_0}_{MX,MY}(t)$. For instance, in the special case where $M = \lambda I$ is a scalar multiple of the identity, the formula

$$\left(\frac{\lambda}{1-(1-\lambda)t}\right)^3 \cdot \kappa^{\Upsilon,h_1}_{X,Y}(t) = \kappa^{\Psi,h_0}_{MX,MY}\left(\frac{\lambda t}{1-(1-\lambda)t}\right) \quad (0 \le t \le 1)$$

holds, even when X and Y do not commute, and can be demonstrated using the techniques of Section 4.

References

[1] N. Brown, R. Finck, M. Spencer, K. Tapp, and Z. Wu, *Invariant metrics with nonnegative curvature on compact Lie groups,* Canad. Math. Bull. 50 (2007), 24–34.

[2] J. Cheeger, *Some examples of manifolds of nonnegative curvature,* J. Differential Geom. 8 (1972), 623–628.

[3] J.-H. Eschenburg, *Inhomogeneous spaces of positive curvature,* Differential Geom. Appl. 2 (1992), 123–132.

[4] K. Grove and W. Ziller, *Curvature and symmetry of Milnor spheres,* Ann. of Math. (2) 152 (2000), 331–367.

[5] M. Müter, *Krümmungserhöhende deformationen mittels Gruppenaktionen,* dissertation, Westfalische Wilhelms-Univ. Munster, 1987.

[6] P. Petersen and F. Wilhelm, *Examples of Riemannian manifolds with positive curvature almost everywhere,* Geom. Topol. 3 (1999), 331–367.
[7] T. Püttmann, *Optimal pinching constants of odd dimensional homogeneous spaces,* Ph.D. thesis, Ruhr-Universität, 1991.
[8] M. Strake, *Curvature increasing metric variations,* Math. Ann. 276 (1987), 633–641.
[9] B. Wilking, *Manifolds with positive sectional curvature almost everywhere,* Invent. Math. 148 (2002), 117–141.

J. Huizenga
Department of Mathematics
Harvard University
Cambridge, MA 02138

huizenga@math.harvard.edu

K. Tapp
Department of Mathematics
and Computer Science
Suffolk University
Boston, MA 02114

ktapp@mcs.suffolk.edu