EQUALITY CONSTRAINTS, RIEMANNIAN MANIFOLDS AND DIRECT SEARCH METHODS*

David W. Dreisigmeyer†
Los Alamos National Laboratory
Los Alamos, NM 87545

December 15, 2006

Abstract

We present a general procedure for handling equality constraints in optimization problems that is of particular use in direct search methods. The central idea is to treat the equality constraints as implicitly defining a Riemannian manifold. Then the function and inequality constraints can be pulled-back to the tangent spaces of this manifold. One can then deal with the resulting inequality constrained optimization problem using any method of one's choosing. An advantage of this procedure is the implicit reduction in dimensionality of the original problem to that of the manifold. Additionally, under some restrictions, convergence results for the method used to solve the inequality constrained optimization problem can be carried over directly to our procedure.

Key words: Direct search methods, Riemannian manifolds, equality constraints, geodesics

AMS subject classifications: 90C56, 53C21

1 Introduction

The problem we will look at is

$$\min_{\mathbf{x}\in\mathbb{R}^n} \quad f(\mathbf{x}) \tag{1.1a}$$

$$\text{subject to :} \quad \mathbf{g}(\mathbf{x}) = \mathbf{0} \tag{1.1b}$$

$$\mathbf{h}(\mathbf{x}) \leq \mathbf{0}, \tag{1.1c}$$

where $\mathbf{g} : \mathbb{R}^n \to \mathbb{R}^m$. Our approach to dealing with (1.1) will be to employ some techniques from differential geometry to efficiently ensure that (1.1b) is always satisfied.

*LA-UR-06-7406
†email: dreisigm@lanl.gov

The way we will do this is by treating $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ as implicitly defining a manifold and, require that we always remain on this manifold as the algorithm proceeds. Since this effectively removes the equality constraints from further consideration, we then only need to concern ourselves with solving an inequality constrained optimization problem.

An outline of the general procedure we will use is the following:

1. Let the equality constraints $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ implicitly define a manifold $\mathcal{M}$.
2. Find an initial feasible point $\mathbf{y} \in \mathcal{M}$.
3. Pullback the inequality constraints $\mathbf{h}(\mathbf{x}) \leq \mathbf{0}$ and the function $f(\mathbf{x})$ from $\mathcal{M}$ to the tangent spaces of $\mathcal{M}$.
4. Use any desired procedure to solve the resulting reduced dimensional problem.

The specific reduced dimensional problem we will have to solve is

$$\min_{\mathbf{w} \in \mathbb{R}^{(n-m)}} \quad \widehat{f}(\mathbf{w}) \tag{1.2a}$$

$$\text{subject to :} \quad \widehat{\mathbf{h}}(\mathbf{w}) \leq \mathbf{0}, \tag{1.2b}$$

where that hat denotes the functions f and $\mathbf{h}$ after they are pulled-back to the tangent spaces and, $\mathbf{w}$ is a tangent vector. There are certain advantages to this procedure. First, we have an implicit reduction in the dimensionality of the problem from n to $(n-m)$, the dimension of the manifold. Secondly, we can choose any method we want to solve (1.2). With some qualifications, the convergence results associated with the method chosen to solve (1.2) carry over *without modification* to (1.1).

Our procedure would seem to be particular useful when employed in conjunction with the filter methods in [1, 8], the MADS algorithms in [2] or, the frame methods in [4, 5, 11]. These procedures can effectively deal with inequality constraints and, thus, are viable solution techniques for (1.2). By treating the equality constraints as a manifold, which conveniently allows us to always satisfy them, we can then easily extend the above algorithms to (1.1). Additionally, it may be difficult to use any derivative information about $\widehat{f}(\mathbf{w})$ and $\widehat{\mathbf{h}}(\mathbf{w})$ after the pullback operation. So direct search methods may be the only methods that can be effectively employed to solve (1.2).

Our paper is organized as follows. In section 2 we outline our general procedure and, relax some of the assumptions in section 3. An example is given in section 4. Finally, section 5 is a discussion of our results.

2 The general procedure

First let us look at the manifolds we will be working with. Useful reference material for this section is [9, 10, 12]. The Morse-Sard theorem tells us that with probability 1,

(1.1b) gives us an implicitly defined Riemannian manifold $\mathcal{M}$ of dimension $(n-m)$. Using this theorem, we will always assume that the Jacobian $\nabla\mathbf{g}(\mathbf{x})$ is full rank. This is simply the requirement that our set of points defined by (1.1b) is actually a manifold. A manifold $\mathcal{M}$ is, for us, a set of points that locally 'looks' like $\mathbb{R}^{(n-m)}$. We call $\mathcal{M}$ Riemannian if we have a sense of distances on $\mathcal{M}$. As an example, let $\mathcal{M}$ be implicitly defined by the function $g : \mathbb{R}^3 \to \mathbb{R}^1$ given by

$$\begin{aligned} g(\mathbf{x}) &= \mathbf{x}^T\mathbf{x} - 1 \\ &= 0. \end{aligned} \tag{2.1}$$

So our $\mathcal{M}$ is the unit sphere $\mathcal{S}^2$ embedded in $\mathbb{R}^3$ which has dimension $(3-1) = 2$. Locally $\mathcal{S}^2$ 'looks' like $\mathbb{R}^2$. That is, we can find, for a small enough region $\mathcal{U} \subset \mathcal{S}^2$, an invertible and differentiable mapping from $\mathcal{U}$ into $\mathbb{R}^2$. We also have a sense of distances on $\mathcal{S}^2$ that are inherited from it's embedding in $\mathbb{R}^3$. Given two points $\mathbf{p}, \mathbf{q} \in \mathcal{S}^2$, we take the distance between $\mathbf{p}$ and $\mathbf{q}$ as the length of the shortest path on $\mathcal{S}^2$ that connects $\mathbf{p}$ to $\mathbf{q}$. This shortest path is called a geodesic.

Now, take a point $\mathbf{y} \in \mathcal{M}$. The tangent space to $\mathcal{M}$ at $\mathbf{y}$, denoted $T_{\mathbf{y}}\mathcal{M}$, is the collection of all of the tangent vectors to $\mathcal{M}$ at $\mathbf{y}$. Since all of our manifolds are embedded in $\mathbb{R}^n$, we can concretely imagine this tangent space as being a copy of $\mathbb{R}^{(n-m)}$ translated to $\mathbf{y}$. Specifically, we have that $T_{\mathbf{y}}\mathcal{M}$ is given by $\mathrm{null}[\nabla\mathbf{g}(\mathbf{y})]$ translated to the base point $\mathbf{y}$. Given a tangent vector $\mathbf{w} \in T_{\mathbf{y}}\mathcal{M}$, we can find a geodesic path $\mathbf{x}(\tau) \in \mathcal{M}$ such that $\mathbf{x}(0) = \mathbf{y}$ and $\dot{\mathbf{x}}(0) = \mathbf{w}$. The specific equation we need to solve is [7]

$$\ddot{x}^k = \dot{\mathbf{x}}^T \left[\sum_{i=1}^{m} \left[-\mathbf{g}_{x^k}^T \left[\nabla\mathbf{g}(\nabla\mathbf{g})^T \right]^{-1} \right]^i \nabla^2 g^i \right] \dot{\mathbf{x}}, \tag{2.2}$$

where x^k is the kth component of $\mathbf{x}$, g^i is the ith component of $\mathbf{g}(\mathbf{x})$ and, $\mathbf{g}_{x^k}$ is the partial derivative of $\mathbf{g}(\mathbf{x})$ with respect to x^k. Equation (2.2) is the key to our procedure. Notice that we require that $\mathbf{g}(\mathbf{x})$ be C^2 in order to use (2.2). Also, if $\mathbf{g}(\mathbf{x})$ is C^2 then so is $\mathcal{M}$. Equation (2.2) can be practically implemented in conjunction with direct search methods [6]. We will not concern ourselves with the actual numerical solution of (2.2) in this paper.

A manifold is called geodesically complete if all of $T_{\mathbf{y}}\mathcal{M}$ can be mapped into $\mathcal{M}$ for every $\mathbf{y} \in \mathcal{M}$. The mapping is denoted by $\mathrm{Exp}_{\mathbf{y}} : T_{\mathbf{y}}\mathcal{M} \to \mathcal{M}$. We will, for now, assume that our manifolds are geodesically complete. This implies that $\mathbf{x}(\tau)$ given by (2.2) will be defined for all $\tau \in \mathbf{R}$. Geodesic completeness also implies that any two points on our manifold can be connected by a geodesic. An example of a geodesically complete manifold is $\mathcal{S}^2$ defined by (2.1).

Now, given our $\mathcal{M}$ implicitly defined by (1.1b), the inequality constraints in (1.1c) will confine us to allowed regions of $\mathcal{M}$. We will take (1.1c) as implicitly defining k full dimensional disjoint regions $V_i \subset \mathbb{R}^n$, $i = 1, \ldots, k$, with $V_i \bigcap \mathcal{M} \neq \emptyset$. Let $\mathcal{U}_i = V_i \bigcap \mathcal{M}$, where we will assume that the $\mathcal{U}_i$ have dimension $(n-m)$ everywhere. Then the $\mathcal{U}_i$ are the feasible sets to (1.1). Let us start at an initial point $\mathbf{y} \in \mathcal{U}_j$.

Then it is a simple matter to apply, e.g., the pattern search method in [1] to $T_{\mathbf{y}}\mathcal{M}$. In particular, the inequality constraints in (1.1c) will become 'black-box' constraints on the allowed tangent vectors $\mathbf{w} \in T_{\mathbf{y}}\mathcal{M}$. It is these tangent vectors that, when used as initial conditions in (2.2), will result in a point $\mathbf{x}(1) \in \mathcal{U}_j$. Note two things. First, if we start in $\mathcal{U}_j$, we will always remain around $\mathcal{U}_j$, assuming the $\mathcal{U}_i$ are sufficiently separated. Depending on the algorithm used to solve (1.2), it may be possible to jump from $\mathcal{U}_j$ to $\mathcal{U}_l$ if they lie close enough together on $\mathcal{M}$. This will not typically be the case and, under the assumption of the geodesic completeness of $\mathcal{M}$, will not cause any problems even if it does occur. Secondly, given a tangent vector $\mathbf{w}$, we evaluate the corresponding geodesic at $\tau = 1$. Then the length that we travelled from $\mathbf{x}(0) = \mathbf{y}$ to $\mathbf{x}(1)$ will be the same as the Euclidean length of $\mathbf{w}$.

Some general comments are in order. Why should the inequality constraints in (1.1c) be taken as 'black-box' constraints on $T_{\mathbf{y}}\mathcal{M}$? The reason is that it would often be very difficult to restate the constraints that define $\mathcal{U}_j$ directly in terms of constraints on $T_{\mathbf{y}}\mathcal{M}$. Instead, we will solve (2.2) and find what the constraint value is at our new point on the manifold. We will then assign this value, along with the function evaluation, to the corresponding tangent vector. Also, the assumption of geodesic completeness tells us that we can reach any point in $\mathcal{U}_j$ from any initial point $\mathbf{y} \in \mathcal{U}_j$. This is really what we need for the algorithm, so the assumption that $\mathcal{M}$ is geodesically complete can be relaxed. We examine this more in section 3. Finally, all of the results concerning the method chosen to solve (1.2) carry over to our extension without modification. This is because (2.2) allows us to implicitly restate the problem of minimizing $f(\mathbf{x})$ in (1.1a) over $\mathcal{U}_j \subset \mathcal{M}$ as one of minimizing $\widehat{f}(\mathbf{w})$ over $T_{\mathbf{y}}\mathcal{M} \simeq \mathbb{R}^{(n-m)}$, subject to our 'black-box' constraints.

Our general method is given by Procedure 2.1. Note that in Step 4, it may be rather difficult to use any method other than a direct search one. The methods in [1, 2, 5] seem particularly well suited to be used in conjunction with Procedure 2.1. The reason for this is that we will not normally have explicit formulas for $\widehat{f}(\mathbf{w})$ or $\widehat{\mathbf{h}}(\mathbf{w})$. So typically the availability of derivative information will, at best, be limited and/or expensive to calculate. This situation can be mitigated if we can reasonably find or estimate $\nabla_{\mathbf{w}}\mathbf{x}(\tau)|_{\tau=1}$, where $\mathbf{x}(\tau)$ is the geodesic found via (2.2) with initial condition $\mathbf{x}(0) = \mathbf{y}$ and $\dot{\mathbf{x}}(0) = \mathbf{w}$. Then we can find or estimate, e.g.,

$$\begin{aligned} \nabla_{\mathbf{w}}\widehat{f}(\mathbf{w}) &= \nabla_{\mathbf{w}} f(\mathbf{x}(\tau))\,|_{\tau=1} \\ &= [\nabla_{\mathbf{w}}\mathbf{x}(\tau)]^T \nabla_{\mathbf{x}} f(\mathbf{x}(\tau))\,|_{\tau=1}\ . \end{aligned} \tag{2.3}$$

For the direct search methods, one could employ the simplex gradient method in [3] for estimating the gradient.

There are three key assumptions we have made in this section:

1. $\mathcal{M}$ is geodesically complete;
2. $\mathcal{M}$ is C^2, and;
3. $\mathcal{M}$ is connected, which is implied by 1.

Procedure 2.1 ***The General Method for Problem (1.1)***

1. *Let the equality constraints* $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ *implicitly define a manifold* $\mathcal{M}$*. Assume that* $\mathcal{M}$ *is geodesically complete.*

2. *Find an initial feasible point* $\mathbf{y} \in \mathcal{M}$*.*

3. *Use the geodesic formula (2.2) to pullback the inequality constraints* $\mathbf{h}(\mathbf{x}) \leq \mathbf{0}$ *and, the function* $f(\mathbf{x})$ *from* $\mathcal{M}$ *to* $T_{\mathbf{y}}\mathcal{M}$*. That is, given any tangent vector* $\mathbf{w} \in T_{\mathbf{y}}\mathcal{M}$ *that results in a geodesic* $\mathbf{x}(\tau)$ *such that* $\mathbf{x}(1) = \mathbf{z}$*, associate with* $\mathbf{w}$ *the function values* $\mathbf{h}(\mathbf{z})$ *and* $f(\mathbf{z})$*. Call these pulled-back functions* $\widehat{\mathbf{h}}(\mathbf{w})$ *and* $\widehat{f}(\mathbf{w})$*.*

4. *Use any desired procedure to solve the resulting reduced dimensional problem given by (1.2).*

5. *Let* $\mathbf{w}^*$ *be the solution to (1.2) found in Step 4. If* $\mathbf{x}^*(\tau)$ *is the geodesic corresponding to* $\mathbf{w}^*$ *via (2.2) and,* $\mathbf{z}^* = \mathbf{x}^*(1)$*, then* $\mathbf{z}^*$ *is the solution to (1.1).*

We look at relaxing 1 in section 3. We will not relax 2 since this is required for (2.2). Assumption 3 will also not be dealt with in this paper. So, if $\mathcal{M}$ consists of, e.g., two pieces and, our initial point $\mathbf{y}$ is on one of the pieces, we assume we remain on the initial piece.

3 Relaxing the geodesic completeness assumption

Now we look at relaxing our assumption that $\mathcal{M}$ is geodesically complete. What we will do is allow ourselves to move around on $\mathcal{M}$ initially. All we need is that eventually our procedure will enter and remain in a neighborhood $\mathcal{U} \subset \mathcal{M}$ such that there is a point $\mathbf{y} \in \mathcal{U}$ that can be connected via a geodesic to any other point in $\mathcal{U}$. Under this assumption of 'local geodesic completeness', we can use this $\mathbf{y}$ as our initial feasible point in Procedure 2.1. This allows us to again carry over the convergence results of the method used to solve (1.2) to our modified algorithm. We call $\mathcal{U}$ a star-shaped set at $\mathbf{y}$ [13]. That is, any $\mathbf{x} \in \mathcal{U}$ can be connected to $\mathbf{y}$ by a geodesic that lies completely in $\mathcal{U}$.

This modified algorithm allows us to significantly extend Procedure 2.1. The basic idea is this. We can employ any method we wish to find a feasible point $\mathbf{y}$ that will give us the star-shaped set property we desire. Once we find such a $\mathbf{y}$, we use this in Procedure 2.1 as our initial point. Now we can use any, potentially different, method we desire to solve (1.2). Since the convergence results for the second method we choose carry over, we now have convergence results for our new algorithm. As long as we eventually find such a $\mathbf{y}$ in a finite number of steps, we are guaranteed this carry over of convergence results. This may require, e.g., restricting ourselves to a finite number

of search steps in a direct search method.

Let us give an example. Assume we think $\mathbf{y}_1$ is a good choice to use in Procedure 2.1. However, as we attempt to solve (1.2), suppose we find that $\mathbf{y}_1$ is not sufficient and, instead, we need to switch to $\mathbf{y}_2$. Now, if it turns out that $\mathbf{y}_2$ does give us the required star-shaped set property, we then have our convergence guarantees. Some optimization methods over Riemannian manifolds are covered in [6, 7, 13].

Various other relaxations are also possible. For instance, we do not necessarily have to require that our base point $\mathbf{y}$ satisfy the inequality constraints in (1.1c). In this case we only need that the feasible region $\mathcal{U}_j$ is reachable via geodesics through $\mathbf{y}$ defined by tangent vectors in $T_{\mathbf{y}}\mathcal{M}$. This may be especially useful for geodesically complete manifolds. Also, if a frame based method as in [4, 11] is used to solve (1.2), it may be possible to relax the above requirement of finding a fixed base point $\mathbf{y}$ in a finite number of steps and still carry over the convergence results. Such a method is presented in [6] for problems without any inequality constraints. Finally, we assume that (2.2) always had a solution at $\tau = 1$. If this is not the case, one would need to refine the tangent vectors used as initial conditions in (2.2). The reason we refine the tangent vectors is because their lengths correspond to the length travelled along the geodesic from $\tau = 0$ to $\tau = 1$.

4 An example

Here we'll look at a relatively simple but illuminating example. Let us consider minimizing some function $f(\mathbf{x})$ over the upper half of $\mathcal{S}^2$. So our optimization problem is

$$\min_{\mathbf{x}\in\mathbb{R}^3} \quad f(\mathbf{x}) \tag{4.1a}$$

$$\text{subject to :} \quad \mathbf{x}^T\mathbf{x} = 1 \tag{4.1b}$$

$$z \geq 0. \tag{4.1c}$$

Then (2.2) is given by

$$\ddot{x}^k = -x^k\dot{\mathbf{x}}^T\dot{\mathbf{x}}. \tag{4.2}$$

Now, if we take $\mathbf{y}$ to be, e.g., the north pole of $\mathcal{S}^2$, we can solve (4.2) with the initial conditions $\mathbf{x}(0) = \mathbf{y}$ and $\dot{\mathbf{x}}(0) = \mathbf{w} \in T_{\mathbf{y}}\mathcal{S}^2$. This is all we need to use Procedure 2.1.

The example here is simple enough that we can write down explicitly the problem (1.2) that Procedure 2.1 is solving implicitly. The first thing is to switch to spherical coordinates, so that (4.1) becomes

$$\min_{\theta,\phi} \quad \widetilde{f}(\theta,\phi) \tag{4.3a}$$

$$\text{subject to :} \quad \phi \leq \frac{\pi}{2}. \tag{4.3b}$$

Now we can rewrite every $\mathbf{w} \in T_{\mathbf{y}}\mathcal{S}^2$ in polar coordinates as

$$\mathbf{w} \quad = \quad \|\mathbf{w}\| \left[\mathbf{w}_1 \cos(\theta) + \mathbf{w}_2 \sin(\theta)\right], \tag{4.4}$$

where we have aligned the coordinates vectors $\mathbf{w}_i$, $i \in \{1, 2\}$, in $T_{\mathbf{y}}\mathcal{S}^2$ so that we can use the same θ coordinate as we use on $\mathcal{S}^2$. Additionally, the constraint on ϕ becomes a constraint on the allowed $\mathbf{w}$ to be used as an initial condition in (4.2). Specifically, since the length travelled by the geodesic $\mathbf{x}(\tau)$ from $\tau = 0$ to $\tau = 1$ is given by $\|\mathbf{w}\|$, we have that $\|\mathbf{w}\| = \phi$. Then, (4.3) can be rewritten as

$$\min_{\mathbf{w} \in \mathbb{R}^2} \qquad \widehat{f}(\mathbf{w}) \tag{4.5a}$$

$$\text{subject to :} \qquad \|\mathbf{w}\| \leq \frac{\pi}{2}, \tag{4.5b}$$

which is the same form as in (1.2).

Using Procedure 2.1 to solve our initial problem in (4.1) means that we are implicitly solving (4.5). Any method that one would choose to solve (4.5) can also be used in Procedure 2.1 to find the solution to (4.1). Generally, any method that one would choose to solve (1.2) can also be used in Procedure 2.1 to find the solution to (1.1). However, for this example we were able to derive explicit formulas for $\widehat{f}(\mathbf{w})$ and our inequality constraint. So we could use a solution method that relies on derivative information. Typically this would not be the case and, direct search methods would be the algorithms of choice to solve (1.2).

5 Discussion

We've shown how to deal with equality constraints in optimization problems by treating them as implicitly defining a Riemannian manifold. By using this framework, the equality constraints can be removed from the original problem. We then have an implicitly defined optimization problem that only has inequality constraints. Then methods like those in [1, 2, 5, 6, 8] can be extended to optimization problems that have both inequality and equality constraints. Additionally, the dimensionality of the original problem is reduced to the dimensionality of the manifold. In particular, our procedure does not require the introduction of any slack variables as in [6].

6 Acknowledgments

The author would like to thank John E. Dennis, JR (Rice University) for very helpful discussions on previous versions of this paper.

References

[1] C. Audet and J. E. Dennis, JR. A pattern search filter method for nonlinear programming without derivatives. *SIAM Journal on Optimization*, 14:980–1010, 2004.

[2] C. Audet and J. E. Dennis, JR. Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization*, 17:188–217, 2006.

[3] D. M. Bortz and C. T. Kelley. The simplex gradient and noisy optimization problems. In *Computational Methods in Optimal Design and Control*, pages 77–90, 1998.

[4] I. D. Coope and C. J. Price. Frame based methods for unconstrained optimization. *Journal of Optimization Theory and Applications*, 107:261–274, 2000.

[5] J. E. Dennis, JR., C. J. Price, and I. D. Coope. Direct search methods for nonlinearly constrained optimization using filters and frames. *Optimization and Engineering*, 5:123–144, 2004.

[6] D. W. Dreisigmeyer. Direct search algorithms over Riemannian manifolds. *In preparation*.

[7] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20:303–353, 1998.

[8] R. Fletcher and S. Leyffer. Nonlinear programming without a penalty function. *Mathematical Programming, Series A*, 91:239–269, 2002.

[9] M. W. Hirsch. *Differential Topology*. Springer, 1976.

[10] P. J. Olver. *Applications of Lie Groups to Differential Equations*. Springer, 2nd edition, 1993.

[11] C. J. Price and I. D. Coope. Frames and grids in unconstrained and linearly constrained optimization: a nonsmooth approach. *SIAM Journal on Optimization*, 14:415–438, 2003.

[12] M. Spivak. *A Comprehensive Introduction to Differential Geometry*, volume 1. Publish or Perish, 3rd edition, 2005.

[13] C. Udriste. *Convex Functions and Optimization Methods on Riemannian Manifolds*. Kluwer, 1994.