Parameterization of triangular meshes

Michael S. Floater

November 10, 2009

Triangular meshes are often used to represent surfaces, at least initially, one reason being that meshes are relatively easy to generate from point cloud data. However, we often want a smoother surface representation, and hence the need arises to fit a smooth parametric surface through the vertices of the mesh. This requires making a suitable parameterization. Parameterizations are also useful for texture mapping and other processes in computer graphics. In this lecture we review some simple parameterization methods, including convex combination maps and discrete harmonic maps.

1 Parameterization of polygons

A standard approach to fitting a smooth parametric curve $\mathbf{c}(t)$ through a sequence of points $\mathbf{x}_i = (x_i, y_i, z_i) \in \mathbb{R}^3$, $i = 0, 1, \ldots, n$, is to start by choosing a *parameterization*, a corresponding increasing sequence of parameter values $t_0 < t_1 < \cdots < t_n$. Then by finding smooth functions $x, y, z : [t_0, t_n] \to \mathbb{R}$ for which $x(t_i) = x_i$, $y(t_i) = y_i$, $z(t_i) = z_i$, an interpolatory curve $\mathbf{c}(t) = (x(t), y(t), z(t))$ results, i.e., a curve $\mathbf{c} : [t_0, t_n] \to \mathbb{R}^3$ such that $\mathbf{c}(t_i) = \mathbf{x}_i$ for each i.

In order to discuss different parameterization methods, we will use the example of C^2 cubic spline interpolation. There is a unique C^2 cubic spline curve $\boldsymbol{\sigma} : [t_0, t_n] \to \mathbb{R}^d$ such that

$$\boldsymbol{\sigma}(t_i) = \mathbf{x}_i, \qquad i = 0, 1, \ldots, n,$$

and

$$\boldsymbol{\sigma}'(t_i) = \mathbf{m}_i, \qquad i = 0, n, \tag{1}$$

for some chosen vectors $\mathbf{m}_0$, $\mathbf{m}_n$. By 'cubic spline curve' we understand that $\boldsymbol{\sigma}$ is a cubic polynomial curve in each interval $[t_i, t_{i+1}]$, and that $\boldsymbol{\sigma}$ has C^2 continuity at the break points $t_1, \ldots, t_{n-1}$. The end conditions (1) are sometimes known as 'clamped' end conditions, and this kind of spline interpolation is sometimes called 'complete' spline interpolation.

The first thing to note about parameterization is that applying a linear transformation to $t_0, \ldots, t_n$ does not change the resulting spline curve. Provided we treat the end conditions (1) correctly, we simply get a reparameterization of $\boldsymbol{\sigma}$. To be precise, if we define

$$\hat{t}_i := \lambda t_i + \mu, \qquad i = 0, 1, \ldots, n,$$

for $\lambda > 0$ and we define the curve $\hat{\boldsymbol{\sigma}} : [\hat{t}_0, \hat{t}_n] \to \mathbb{R}^d$ by

$$\hat{\boldsymbol{\sigma}}(t) := \boldsymbol{\sigma}((t - \mu)/\lambda),$$

then $\hat{\boldsymbol{\sigma}}$ is clearly a C^2 cubic spline curve over the partition $\hat{t}_0 < \hat{t}_1 < \cdots < \hat{t}_n$, satisfying the conditions

$$\hat{\boldsymbol{\sigma}}(\hat{t}_i) = \mathbf{x}_i, \qquad i = 0, 1, \ldots, n,$$

and

$$\hat{\boldsymbol{\sigma}}'(\hat{t}_i) = \mathbf{m}_i/\lambda, \qquad i = 0, n.$$

For this reason, we may as well set $t_0 = 0$ and we typically specify a parameterization by recursively setting

$$t_{i+1} := t_i + d_i$$

for some chosen interval lengths $d_0, d_1, \ldots, d_{n-1} > 0$. Multiplying the interval lengths d_i by a common factor λ will not change the intrinsic geometry of the spline curve $\boldsymbol{\sigma}$ as long as we divide the vectors $\mathbf{m}_0$ and $\mathbf{m}_n$ by the same factor λ.

The simplest choice is the *uniform parameterization* defined by $d_i = 1$, where the values t_i are uniformly spaced. But as early as 1967, Ahlberg, Nilson, and Walsh proposed using the *chordal parameterization* in which

$$d_i := \|\mathbf{x}_{i+1} - \mathbf{x}_i\|, \tag{2}$$

and $\|\cdot\|$ denotes the Euclidean norm in $\mathbb{R}^d$. The motivation behind this is that the distance between two points on a curve is a reasonable approximation to the length of the associated curve segment. Thus the hope is that the 'speed' $\|\boldsymbol{\sigma}'(t)\|$ of the spline curve might be close to unity at all $t \in [t_0, t_n]$. In practice, the chordal parameterization often does a better job of avoiding the cusps and self-intersections that sometimes occur with the uniform parameterization when the distances between points varies a lot.

Later, it was realized that the uniform and chordal parameterizations are the special cases $\mu = 0$ and $\mu = 1$ of the more general parameterization

$$d_i := \|\mathbf{x}_{i+1} - \mathbf{x}_i\|^\mu,$$

with $0 \leq \mu \leq 1$ acting as a kind of blending parameter. The choice $\mu = 1/2$ was termed by Lee the *centripetal parameterization*. He suggested that it typically leads to a spline curve that looks smoother than the uniform and chordal ones.

Consider now how we might generalize some of these methods to triangular meshes. The following theorem provides some ideas.

Theorem 1 *Suppose* $t_0 < t_1 < \cdots < t_n$ *and* $L_0, L_1, \ldots, L_{n-1} > 0$. *The following three statements are equivalent:*

1. *there is some* $\mu > 0$ *such that* $t_{i+1} = t_i + \mu L_i$ *for* $i = 0, 1, \ldots, n-1$,

2. $t_1, \ldots, t_{n-1}$ *minimize*

$$F(t_1, \ldots, t_{n-1}) = \sum_{i=0}^{n-1} \frac{(t_{i+1} - t_i)^2}{L_i},$$

3.

$$t_i = \left(\frac{L_i}{L_{i-1} + L_i} \right) t_{i-1} + \left(\frac{L_{i-1}}{L_{i-1} + L_i} \right) t_{i+1}, \qquad i = 1, \ldots, n-1.$$

2 Parameterization of triangular meshes

We now describe a general method for constructing a parameterization of triangular mesh in $\mathbb{R}^3$. We denote by $\mathcal{S}$ the set of triangles in the mesh and V its vertices and E its edges. We let $\Omega_{\mathcal{S}} \subset \mathbb{R}^3$ be the union of the triangles in $\mathcal{S}$. Then we define a parameterization of $\mathcal{S}$ as a continuous piecewise

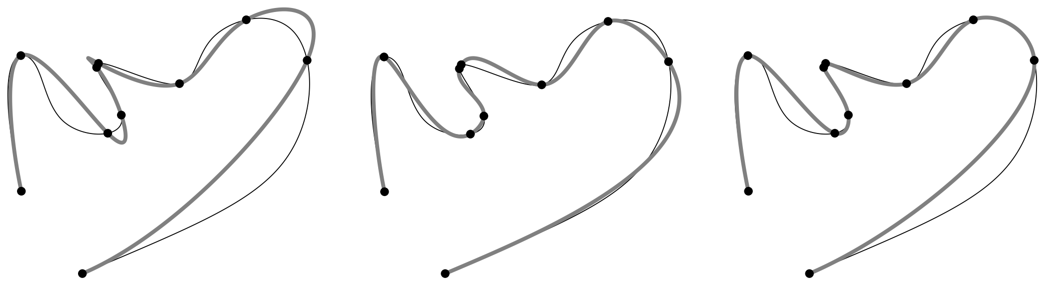

(a) uniform (b) chordal (c) centripetal

Figure 1: Choice of parameterization for cubic spline interpolation.

linear mapping $\psi : \Omega_{\mathcal{S}} \to \mathbb{R}^2$. Then ψ maps each vertex, edge, and triangle of $\mathcal{S}$ to a corresponding vertex, edge, and triangle in $\mathbb{R}^2$. Such a mapping is completely determined by the points $\psi(\mathbf{v})$, $\mathbf{v} \in V$. Let V_I denote the interior vertices of $\mathcal{S}$ and V_B the boundary ones. The boundary vertices of $\mathcal{S}$ form a polygon $\partial\mathcal{S}$ in $\mathbb{R}^3$ which we call the *boundary polygon* of $\mathcal{S}$. Two distinct vertices $\mathbf{v}$ and $\mathbf{w}$ in $\mathcal{S}$ are *neighbours* if they are the end points of some edge in $\mathcal{S}$. For each $\mathbf{v} \in V$, let

$$N_{\mathbf{v}} = \{\mathbf{w} \in V : [\mathbf{w}, \mathbf{v}] \in E\},$$

the set of neighbours of $\mathbf{v}$, where $E = E(\mathcal{S})$ is the set of edges in $\mathcal{S}$.

The first step of the method is to choose any points $\psi(\mathbf{v}) \in \mathbb{R}^2$, for $\mathbf{v} \in V_B$, such that the boundary polygon $\partial\mathcal{S}$ of $\mathcal{S}$ is mapped into a simple polygon $\psi(\partial\mathcal{S})$ in the plane. In the second step, for $\mathbf{v} \in V_I$, we choose a set of strictly positive values $\lambda_{\mathbf{vw}}$, for $\mathbf{w} \in N_{\mathbf{v}}$, such that

$$\sum_{\mathbf{w} \in N_{\mathbf{v}}} \lambda_{\mathbf{vw}} = 1. \tag{3}$$

Then we let the points $\psi(\mathbf{v})$ in $\mathbb{R}^2$, for $\mathbf{v} \in V_I$, be the unique solutions of the linear system of equations

$$\psi(\mathbf{v}) = \sum_{\mathbf{w} \in N_{\mathbf{v}}} \lambda_{\mathbf{vw}} \psi(\mathbf{w}), \qquad \mathbf{v} \in V_I. \tag{4}$$

Since these equations force each point $\psi(\mathbf{v})$ to be a convex combination of its neighbouring points $\psi(\mathbf{w})$, we call ψ a *convex combination mapping*. Fig. 2 shows an example of a triangular mesh in $\mathbb{R}^3$. Fig. 3 shows a convex

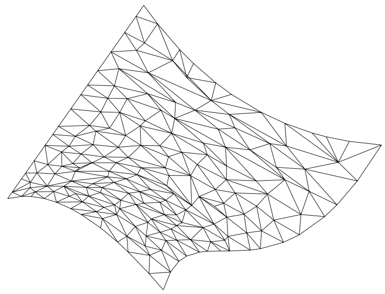

Figure 2: Triangular mesh in $\mathbb{R}^3$

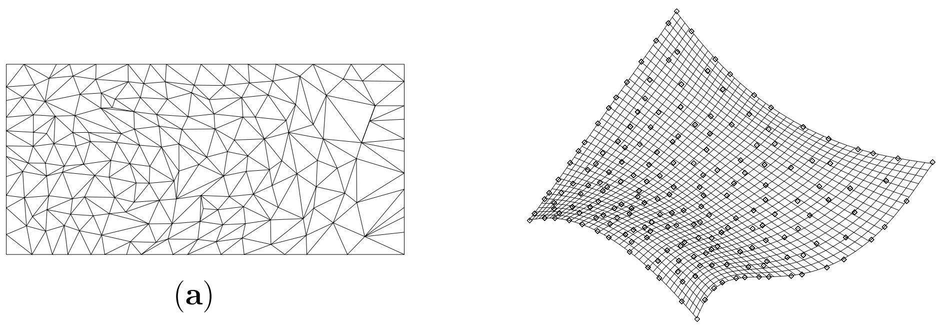

(**a**) (**b**)

Figure 3: Parameterization (triangular mesh in $\mathbb{R}^2$) (**a**) and resulting tensor-product spline approximation (**b**)

combination mapping of $\mathcal{S}$ into a planar mesh $\mathcal{T}$, whose boundary was chosen to be a rectangle. Fig. 3 also shows a tensor-product spline approximation (in fact a least square approximation) to the vertices of $\mathcal{S}$ based on their parameter points, the vertices of $\mathcal{T}$.

Let us take a closer look at the linear system. We must show that it has a unique solution. To this end, note that it can be rewritten in the form

$$\psi(\mathbf{v}) - \sum_{\mathbf{w} \in N_{\mathbf{v}} \cap V_I} \lambda_{\mathbf{vw}} \psi(\mathbf{w}) = \sum_{\mathbf{w} \in N_{\mathbf{v}} \cap V_B} \lambda_{\mathbf{vw}} \psi(\mathbf{w}), \qquad \mathbf{v} \in V_I. \tag{5}$$

This can be written as the matrix equation

$$A\mathbf{x} = \mathbf{b},$$

where $\mathbf{x} = (\psi(\mathbf{w}))_{\mathbf{w}\in V_I}$ is the column vector of unknowns in some arbitrary ordering, $\mathbf{b}$ is the column vector whose elements are the right hand sides of (5), and the matrix $A = (a_{\mathbf{vw}})_{\mathbf{v},\mathbf{w}\in V_I}$ has dimension $n \times n$, with $n = |V_I|$, and elements

$$a_{\mathbf{vw}} = \begin{cases} 1, & \mathbf{w} = \mathbf{v}, \\ -\lambda_{\mathbf{vw}}, & \mathbf{w} \in N_{\mathbf{v}}, \\ 0, & \text{otherwise.} \end{cases}$$

The existence and uniqueness of the solution to (4) follows from the structure of the matrix A, namely that its off diagonal elements are either zero or negative and each row of A is diagonally dominant. Moreover every row corresponding to a vertex $\mathbf{v} \in V_I$ which has at least one neighbour in V_B is strictly diagonally dominant and every interior vertex can be connected to the boundary by a path of vertices. A standard result in linear algebra shows then that A is non-singular (in fact A is a so-called M-matrix, and such matrices frequently occur in numerical approximations to elliptic partial differential equations.

An interesting question is whether ψ is one-to-one. It will not be one-to-one in general but the following result gives a sufficient condition. The proof is beyond the scope of this course. We say that an interior edge of $\mathcal{S}$ is a *dividing edge* of $\mathcal{S}$ if both its end points are boundary vertices of $\mathcal{S}$.

Theorem 2 *If $\psi(\partial\Omega)$ is convex and no dividing edge $[v, w]$ is mapped by ψ into $\partial\Omega$ then ψ is injective.*

3 Choosing the Weights

A simple choice of weights $\lambda_{\mathbf{vw}}$ is to take them to be uniform, i.e., constant for each vertex $\mathbf{v}$, so that $\lambda_{\mathbf{vw}} = 1/d(\mathbf{v})$, $\mathbf{w} \in N_{\mathbf{v}}$ where $d(\mathbf{v})$ is the degree $|N_{\mathbf{v}}|$ of $\mathbf{v}$. Then every interior vertex $\psi(\mathbf{v})$ of the solution to the linear system will be the barycentre of its neighbours. However, numerical examples show that this uniform parameterization usually leads to poor spline surfaces when used for approximation. Look for example at Figures 4 and 5. Figure 4 shows a mesh $\mathcal{S}$ and Figure 5 shows a uniform parameterization $\mathcal{T}$ of $\mathcal{S}$ together with a Clough-Tocher interpolant (a C^1 piecewise cubic interpolant) to $\mathcal{S}$ over the mesh $\mathcal{T}$. Clearly the iso-curves are badly behaved.

One reason for the bad behaviour of the surface approximation is that the weights $\lambda_{\mathbf{vw}}$ are independent of the geometry of the vertices $\mathbf{v}$ of $\mathcal{S}$. In

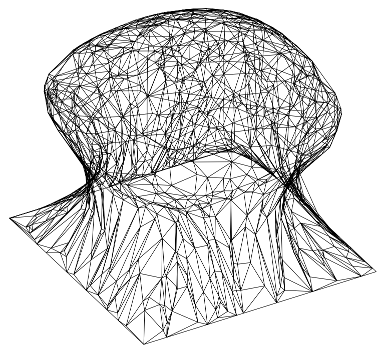

Figure 4: Triangular mesh in $\mathbb{R}^3$

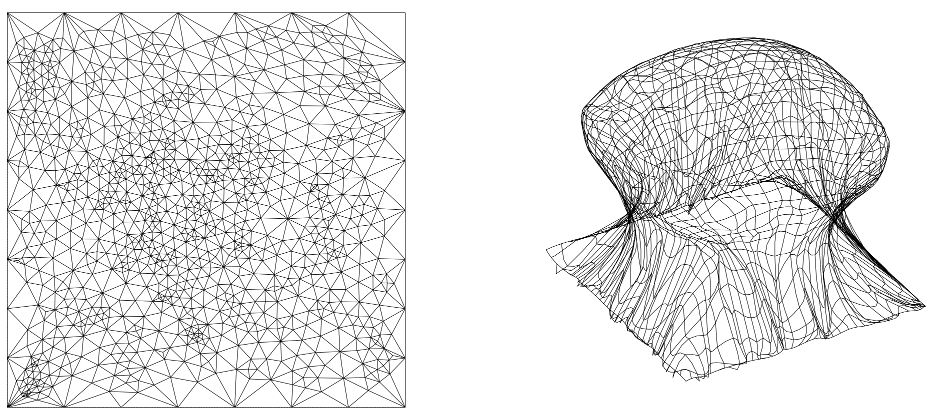

(**a**) (**b**)

Figure 5: Uniform parameterization (**a**) and Clough-Tocher interpolant (**b**)

practice it better to choose weights for which the Euclidean distance

$$||\mathbf{v} - \sum_{\mathbf{w} \in N_\mathbf{v}} \lambda_{\mathbf{vw}} \mathbf{w}|| \tag{6}$$

between $\mathbf{v}$ and $\sum_{\mathbf{w} \in N_\mathbf{v}} \lambda_{\mathbf{vw}} \mathbf{w}$ is as small as possible and in particular, when $\mathbf{v}$ and its neighbours lie in a plane, we should have

$$\mathbf{v} = \sum_{\mathbf{w} \in N_\mathbf{v}} \lambda_{\mathbf{vw}} \mathbf{w}. \tag{7}$$

This latter condition implies linear precision: if the whole mesh $\mathcal{S}$ lies in a plane and $\psi|_{\partial\mathcal{S}}$ is an affine mapping then the whole convex combination mapping ψ is an affine mapping.

A choice of weights that achieves this is the mean value weights:

$$\lambda_{\mathbf{vw}} = w_{\mathbf{vw}} / \sum_{\mathbf{u} \in N_\mathbf{v}} w_{\mathbf{vu}},$$

where

$$w_{\mathbf{vw}} = \frac{\tan(\alpha/2) + \tan(\beta/2)}{||\mathbf{w} - \mathbf{v}||},$$

and α and β are the angles at the vertex $\mathbf{v}$ of the two triangles adjacent to the edge $[\mathbf{v}, \mathbf{w}]$. Figure 6 shows the result of interpolating $\mathcal{S}$ of Figure 4 with a Clough-Tocher interpolant over the mean value parameterization of $\mathcal{S}$. The surface approximation is clearly better than that of Figure 5, using uniform parameterization.

The effect of choosing different boundary polygons $\psi(\partial\mathcal{S})$ is shown in Figures 7 and 8. In Fig. 8(**a**) the parameter points of the boundary vertices of the mesh in 7 were distributed by chord length on a rectangle and in in Fig. 8(**b**) on a circle.

4 Weighted Least Squares

A special case of convex combination parameterizations arises from minimizing so-called spring energy. First, choose as before a convex polygon $\psi(\partial\mathcal{S})$. Secondly, for each interior edge $[\mathbf{v}, \mathbf{w}]$ in $\mathcal{S}$ choose some value $\mu_{\mathbf{vw}} = \mu_{\mathbf{wv}} > 0$. Then let the points $\psi(\mathbf{v})$, $\mathbf{v} \in V_I$, minimize the function

$$F = \sum_{[\mathbf{v},\mathbf{w}] \in E} \mu_{\mathbf{vw}} ||\psi(\mathbf{v}) - \psi(\mathbf{w})||^2. \tag{8}$$

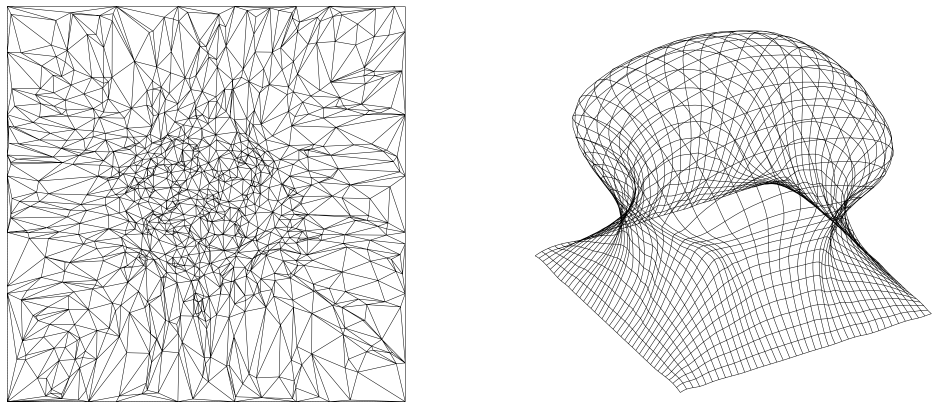

(**a**) (**b**)

Figure 6: Mean value parameterization (**a**) and Clough-Tocher interpolant (**b**)

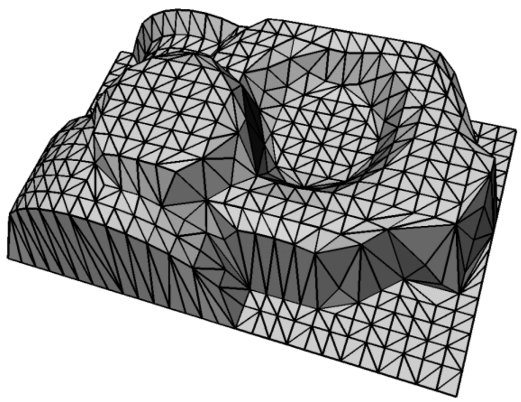

Figure 7: Triangular mesh in $\mathbb{R}^3$

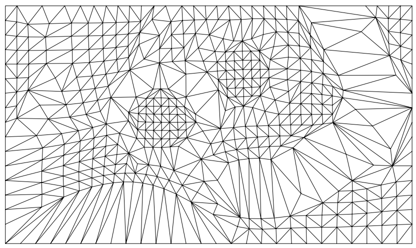

(a)

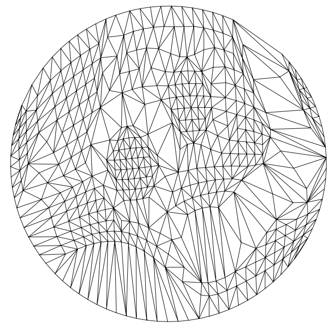

(b)

Figure 8: Mean value parameterizations of the mesh in Figure 7

The normal equations for (8) are

$$\psi(\mathbf{v}) = \frac{\sum_{\mathbf{w}\in N_{\mathbf{v}}} \mu_{\mathbf{vw}}\psi(\mathbf{w})}{\sum_{\mathbf{w}\in N_{\mathbf{v}}} \mu_{\mathbf{vw}}}, \qquad \mathbf{v} \in V_I$$

and so minimizing F is equivalent to solving (4) where

$$\lambda_{\mathbf{vw}} = \frac{\mu_{\mathbf{vw}}}{\sum_{\mathbf{w}\in N_{\mathbf{v}}} \mu_{\mathbf{vw}}}, \qquad \mathbf{v} \in V_I, \quad \mathbf{w} \in N_{\mathbf{v}}.$$

Notice that in general $\lambda_{\mathbf{vw}} \neq \lambda_{\mathbf{wv}}$ even though $\mu_{\mathbf{vw}} = \mu_{\mathbf{wv}}$. As an example one might choose $\mu_{\mathbf{vw}} = 1/||\mathbf{v} - \mathbf{w}||$, but this method does not have linear precision. Currently it is not known whether it is possible to find positive coefficients $\mu_{\mathbf{vw}}$ that yield linear precision. A choice of coefficients which does have linear precision comes from the finite element discretization of the harmonic map, but the coefficients are not in general positive, and so Theorem 2 is no longer applicable, and the mapping in some cases is not injective. However, in practice this pararmeterization often is injective and yields results similar to that of the mean value parameterization. The discrete harmonic map is based on the fact that harmonic maps minimize the *Dirichlet energy* which is defined as

$$E_D(\psi) = \int_{\Omega_S} \|\nabla\psi\|^2.$$

A calculation shows that $E_D(\psi)$ can be expressed in the form (8) with

$$\mu_{\mathbf{vw}} = \frac{1}{4}(\cot\delta + \cot\gamma),$$

where δ and γ are the angles opposite to $[\mathbf{v}, \mathbf{w}]$ in the two triangles adjacent to it. Since

$$\cot\delta + \cot\gamma = \frac{\sin(\delta+\gamma)}{\sin\delta\sin\gamma},$$

this means that $\mu_{\mathbf{vw}} \geq 0$ if and only if $\delta + \gamma \leq \pi$.