

Harmonic Coordinates

Tony DeRose Mark Meyer
Pixar Technical Memo #06-02
Pixar Animation Studios

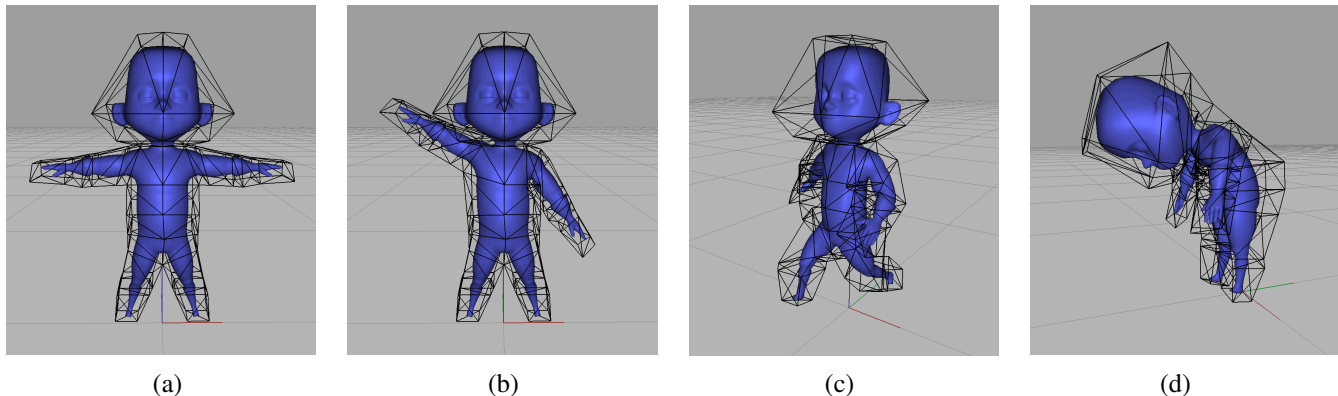


Figure 1: A character (shown in blue) being deformed by a cage (shown in black) using harmonic coordinates. (a) The character and cage at bind-time; (b) - (d) the deformed character corresponding to three different poses of the cage.

Abstract

Generalizations of barycentric coordinates in two and higher dimensions have been shown to have a number of applications in recent years, including finite element analysis, the definition of S-patches (n -sided generalizations of Bézier surfaces), free-form deformations, mesh parametrization, and interpolation. In this paper we present a new form of d dimensional generalized barycentric coordinates. The new coordinates are defined as solutions to Laplace's equation subject to carefully chosen boundary conditions. Since solutions to Laplace's equation are called harmonic functions, we call the new construction harmonic coordinates. We show that harmonic coordinates possess several properties that make them more attractive than mean value coordinates when used to define two and three dimensional deformations.

Keywords: Barycentric coordinates, mean value coordinates, free form deformations, rigging.

1 Introduction

Two dimensional barycentric coordinates are fundamental in a wide variety of applications, including Gouraud shading of triangles and the definition of triangular Bézier patches [Farin 2002]. Given a triangle with vertices T_1, T_2, T_3 , barycentric coordinates allow every point p in the plane of the triangle to be expressed uniquely as

$$p = \sum_{i=1,2,3} \beta_i(p) T_i \quad (1)$$

where the numbers $\beta_1(p), \beta_2(p), \beta_3(p)$ are the barycentric coordinates of p with respect to T_1, T_2, T_3 . They can be defined in many

ways, one of the simplest being as the unique linear functions satisfying the interpolation conditions:

$$\beta_i(T_j) = \delta_{i,j}, \quad i, j = 1, 2, 3. \quad (2)$$

Similarly, barycentric coordinates in three dimensions can be defined relative to a non-degenerate tetrahedron with vertices $T_1, T_2, T_3, T_4 \in \mathcal{R}^3$ as the unique linear functions that satisfy Equation 2 where the indices i and j run from 1 to 4 instead of from 1 to 3.

As described in Ju *et. al.* [Ju et al. 2005], most of the uses of barycentric coordinates stem from their use in the construction of interpolating functions. Gouraud shading is a familiar example where colors c_1, c_2, c_3 assigned to the vertices of a triangle are interpolated across the triangle according to

$$c(p) = \sum_{i=1,2,3} \beta_i(p) c_i \quad (3)$$

A similar formula can be used to define a deformation of two-space. Namely, let T'_1, T'_2 and T'_3 denote new positions for the vertices of the original triangle. The deformed position of point p can then be defined as

$$p' = \sum_{i=1,2,3} \beta_i(p) T'_i \quad (4)$$

Consider now the problem of defining such two dimensional coordinates (and corresponding interpolants) relative to polygons with more than three vertices. The interpolation conditions are still important, but it is no longer sufficient to require the coordinate functions to be linear – there are too many interpolation conditions to satisfy with linear functions. An interesting and important question then is how to generalize barycentric coordinates to arbitrary closed polygons in the plane and arbitrary closed polyhedra in space.

The problem of generalizing barycentric coordinates is surprisingly rich, and has received considerable attention in recent years ([Wachpress 1975], [R.Sibson 1981], [Loop and DeRose 1989], [Warren 1996], [Meyer et al. 2002], [Floater 2003], [Floater et al. 2005], [Ju et al. 2005]). Ju *et. al.* provide a good overview in [Ju et al. 2005]. Since not all properties of barycentric coordinates can

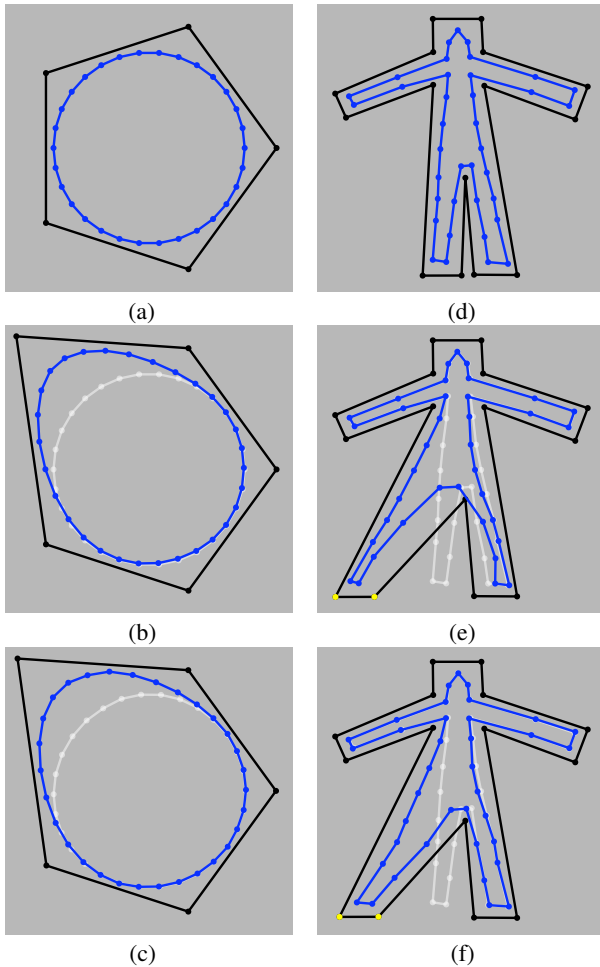


Figure 2: Two dimensional generalized barycentric coordinates used to define deformations of two different objects (shown in blue) using cages (shown in black). The top row shows the cages and objects at “bind” time. The second row shows modified cages and the corresponding deformed objects using mean value coordinates. The last row shows modified cages and deformed objects using harmonic coordinates. In the last two rows, the original undeformed object is shown in white. The two methods perform similarly for convex shapes. In the bipedal case, harmonic coordinates perform better in that the motion of cage points in the left leg does not influence points in the right leg.

be retained in the generalization, the richness results from the many different ways the properties can be relaxed.

We are particularly interested in using generalized barycentric coordinates for character deformation, as shown in Figure 1 and as described in Ju *et. al* [Ju *et al.* 2005]. In this application, an object to be deformed is positioned relative to a closed shape that we’ll call a cage. Examples are shown in Figures 1 and Figure 2. The object is then “bound” to the cage by computing generalized barycentric coordinates $g_i(p)$ of each object point p relative to the cage vertices C_i . As the cage vertices are moved to new locations C'_i , the deformed points p' are computed from

$$p' = \sum_i g_i(p) C'_i \quad (5)$$

Of the various generalized barycentric coordinate formulations available, mean value coordinates [Floater 2003; Floater *et al.* 2005; Ju *et al.* 2005] are particularly useful in this application because:

- The cage that controls the deformation can be any simple closed polygon in two dimensions, and any simple closed triangular mesh in three dimensions.
- The coordinates are smooth, so the deformation is smooth.
- The coordinates reproduce linear functions, so the object doesn’t “pop” when it is bound. That is, the coordinates are such that setting C'_i to C_i in Equation 5 results in p' reducing to p .

A second example motivated by the articulation of bipedal characters is shown in the second column of Figure 2. Notice how the modified cage points on the leg on the left in Figure 2(e) influence the position of object points in the leg on the right. This occurs because mean value coordinates are based on Euclidean (straight-line) distances between points of the cage and points of the object. Although the influence is noticeable in still images, the movement of points in the right leg when the left leg cage points is particularly striking in interactive use, as demonstrated in the accompanying video. The behavior of 3D mean value coordinates is similar, and is highly undesirable for the articulation of characters in feature film production.

What is needed for character articulation is a form of generalized barycentric coordinates that adds the following properties to those enjoyed by mean value coordinates:

- Non-negativity. This implies that object points move in the same direction as cage points; negative coordinates mean they can move in the opposite direction. Negativity of mean value coordinates is responsible for the right leg points in Figure 2(e) moving in the opposite direction from the left leg cage points. Negativity is also responsible for the collapsing of the left leg points in Figure 2(e).
- Interior locality. Informally, the coordinates should fall off as a function of the distance between cage points and object points *as measured within the cage*.

In this paper we show that such coordinates can be produced as solutions to Laplace’s equation with appropriately chosen boundary conditions. Since solutions to Laplace’s equation are generically referred to as harmonic functions, we therefore call these coordinates *harmonic coordinates*, and the deformations they generate *harmonic deformations*.¹

1.1 Previous work

Laplace’s equation, harmonic functions, and harmonic maps have often been mentioned in previous constructions of barycentric coordinates in two dimensions. For instance, the “cotangent weights” of [Pinkhall and Polthier 1993] and [Meyer *et al.* 2002] can be derived from piecewise linear discretizations of Laplace’s equation. Similarly, Floater’s construction of mean value coordinates was motivated by the mean value theorem for harmonic functions. It is somewhat surprising to us that direct solution of Laplace’s equation has never been used to create generalized barycentric coordinates, but that seems to be the case.

Another connection between Laplace’s equation and mean value coordinates comes from the motivation given in [Ju *et al.* 2005]. They derive mean value coordinates starting with an interpolant they call the mean value interpolant. The mean value interpolant to a function f defined on a closed boundary works as follows. To

¹Since each component of a harmonic deformation is a harmonic function, many texts refer to such deformations as harmonic maps. We prefer the term harmonic deformation because of the context in which they’re used in this paper.




Figure 3: Mean value vs harmonic interpolation. (a) The straight-line paths corresponding to mean value interpolation. (b) The Brownian paths corresponding to harmonic interpolation.

compute a value for each interior point p , consider each point x on the boundary. Multiply $f(x)$ by the reciprocal distance from x to p , then average over all x (see Figure 3(a)). This definition makes it clear that mean value coordinates involve straight-line distances irrespective of the visibility of x from p . An alternative interpolant that respects visibility is to average not over all straight-line paths, but rather to average over all *Brownian paths* leaving p , where the value assigned to each path is the value of f at the point the path first hits the boundary (see Figure 3(b)). Although this definition at first seems intractable to compute, it is a famous result from stochastic processes (c.f. [Port and Stone 1978], [Bass 1995]) that the interpolant thus produced (in any dimension) in fact satisfies Laplace’s equation subject to the boundary conditions given by f .²

2 Theory

In this section we formalize the discussion of Section 1. Let C be a closed (not necessarily convex) volume in d dimensions with a piecewise linear boundary. Geometers call such shapes polytopes, but because we have specific uses in mind, we refer to these shapes instead as cages. In two dimensions, a cage is a region of the plane bounded by a closed polygon (such as the one shown in Figure 2), and in three dimensions a cage is a closed region of space bounded by planar (though not necessarily triangular) faces. For each of the vertices C_i of the cage, we seek a function $h_i(p)$ defined on C subject to the following conditions:

1. Interpolation: $h_i(C_j) = \delta_{i,j}$.
2. Affine-invariance: $\sum_i h_i(p) = 1$ for all $p \in C$.
3. Strict generalization of barycentric coordinates: when C is a simplex, $h_i(p)$ is the barycentric coordinate of p with respect to C_i .
4. Smoothness: The functions $h_i(p)$ are at least C^1 smooth.
5. Non-negativity: $h_i(p) \geq 0$, for all $p \in C$.
6. Linear reproduction: Given an arbitrary function $f(p)$, the coordinate functions can be used to define an interpolant $H[f](p)$ according to:

$$H[f](p) = \sum_i h_i(p) f(C_i) \quad (6)$$

Following Ju *et. al* [Ju *et al.* 2005], we require $H[f](p)$ to be exact for linear functions. As shown by Ju *et. al.*, taking $f(p) = p$ means that

$$p = \sum_i h_i(p) C_i \quad (7)$$

which is the “non-popping” condition mentioned in Section 1.

²We thank [name omitted for review purposes] for pointing out this connection to us.

7. Interior locality: We quantify the notion of interior locality introduced above as follows: interior locality holds, if, in addition to non-negativity, the coordinate functions have no interior extrema.

Mean value coordinates possess all but two of these properties: namely, non-negativity and interior locality. We claim that coordinate functions satisfying all seven properties can be obtained as solutions to Laplace’s equation

$$\nabla^2 h_i(p) = 0, p \in \text{Int}(C) \quad (8)$$

if the boundary conditions are carefully chosen.

To gain some insight into how the boundary conditions are determined, we consider first the construction of harmonic coordinates in two dimensions. It will then be clear how the construction generalizes to d dimensions. For reasons that will soon become apparent, the appropriate boundary conditions for $h_i(p)$ in two dimensions are as follows. Let ∂p denote a point on the boundary ∂C of C , then

$$h_i(\partial p) = \phi_i(\partial p), \text{ for all } \partial p \in \partial C \quad (9)$$

where $\phi_i(\partial p)$ is the (univariate) piecewise linear function such that $\phi_i(C_j) = \delta_{i,j}$. For example, if C is the cage shown in Figure 4(a), then $\phi_i(\partial p)$ is the piecewise linear function defined on the edges e_1, \dots, e_{15} such that $\phi_i(C_j) = \delta_{i,j}$, for $i, j = 1, \dots, 15$.

We now show that functions satisfying Equation 8 subject to Equation 9 possess the properties enumerated above. It turns out that the linear reproduction property subsumes several other conditions, so for purposes of proof we verify the conditions in a different order than the one presented above.

- Interpolation: by construction $h_i(C_j) = \phi_i(C_j) = \delta_{i,j}$.
- Smoothness: Away from the boundary harmonic coordinates are solutions to Laplace’s equation, and hence they are C^∞ .
- Non-negativity: harmonic functions achieve their extrema at their boundaries. Since boundary values are restricted to $[0, 1]$, interior values are also restricted to $[0, 1]$.
- Linear reproduction: Let $f(p)$ be an arbitrary linear function. We need to show that $H[f](p) = f(p)$, where $H[f](p)$ is defined as in Equation 6. We begin by establishing that $H[f](p) = f(p)$ everywhere on the boundary of C . If ∂p is a point on the boundary of C , then by construction

$$H[f](\partial p) = \sum_i h_i(\partial p) f(C_i) = \sum_i \phi_i(\partial p) f(C_i) \quad (10)$$

The functions $\phi_i(\partial p)$ are the univariate linear B-spline basis functions (commonly known as the “hat function” basis), which are capable of reproducing all linear functions on ∂C (in fact, they reproduce all piecewise linear functions on ∂C).

Next we extend the result to the interior of C . Note that since $f(p)$ is linear, all second derivatives vanish, and in particular $\nabla^2 f(p) = 0$; thus $f(p)$ satisfies Laplace’s equation on the interior of C . $H[f](p)$ also satisfies Laplace’s equation on the interior, because for interior points $\cdot p$:

$$\begin{aligned} \nabla^2 H[f](\cdot p) &= \nabla^2 \sum_i h_i(\cdot p) f(C_i) \\ &= \sum_i f(C_i) \nabla^2 h_i(\cdot p) \\ &= \sum_i f(C_i) 0 \\ &= 0 \end{aligned}$$

Since $f(p)$ and $H[f](p)$ agree on their boundaries and are both solutions to the same differential equation, by uniqueness of solutions to PDEs, they must be the same function.




Figure 4: A comparison of coordinate functions for a concave cage. (a) A 2D cage with vertices C_1, \dots, C_{15} ; (b) the value of the mean value coordinate for C_2 (yellow indicates positive values, green indicates negative values); (c) the value of the harmonic coordinate for C_2 (red denotes the exterior of the cage where the function is undefined). To accentuate values near zero, intensities of yellow and green are proportional to the square root of the coordinate function value. The significant influence of the position of C_2 on object points in the leg on the right is indicated by the presence of green in the right leg of (b). The corresponding influence in (c) is essentially zero.

- Affine invariance: The function $f(p) = 1$ is linear, so affine invariance follows immediately from the linear reproduction property.
- Strict generalization of barycentric coordinates. If the cage C consists of a single triangle harmonic coordinates reduce to barycentric coordinates. Let $\beta_j(p)$ denote the barycentric coordinates of p with respect to the triangle. To establish that $h_j(p) = \beta_j(p)$, note that $\beta_j(p)$ is a linear function, so we can use the linear reproduction property above by taking $f(p) = \beta_j(p)$:

$$\begin{aligned}
 \beta_j(p) &= H[\beta_j](p) \\
 &= \sum_i h_i(p) \beta_j(C_i) \\
 &= \sum_i h_i(p) \delta_{i,j} \\
 &= h_j(p)
 \end{aligned}$$

- Interior locality: follows from non-negativity and the fact that harmonic functions possess no interior extrema.

To generalize from two to d dimensions, we first back up and consider harmonic coordinates in one dimension. In one dimension a cage is a line segment bounded by two vertices C_0 and C_1 , and Laplace's equation reduces to

$$\frac{d^2 h_i(p)}{dp^2} = 0. \quad (11)$$

Thus, $h_i(p)$ is a linear function, and the proper (zero dimensional) boundary conditions come from the interpolation property: $h_i(C_j) = \delta_{i,j}$.

With this insight, we can rephrase the two dimensional construction as: to construct two dimensional harmonic coordinates, start with the interpolation conditions $h_i(C_j) = \delta_{i,j}$. This determines the coordinates on the 0-dimensional facets (the vertices) of C . Next, extend the coordinates to the 1-dimensional facets (the edges) of C using the one dimensional version of Laplace's equation. Finally, extend them to the two dimensional facets (the interior) of C using the two dimensional version of Laplace's equation.

The extension to three and higher dimensions follows immediately: The harmonic coordinates $h_i(p)$ for a d dimensional cage C with vertices C_i , are the unique functions such that:

1. $h_i(C_j) = \delta_{i,j}$.

2. On every facet of dimension $k \leq d$, the k dimensional Laplace equation is satisfied.

To prove that d dimensional harmonic coordinates defined in this way possess the required properties, we can use induction on the facet dimension, starting with the 1-facets as the base case. The proofs given above for two dimensions are actually more general; they are valid in any dimension k assuming that linear reproduction is achieved on the $k-1$ facets. These proofs therefore serve as the inductive step.

Having defined coordinates in this way, by construction we have the following additional property that is shared by barycentric coordinates and Warren's [Warren 1996] construction:

- Dimension reduction: d dimensional harmonic coordinates, when restricted to a $k < d$ dimensional facet, reduce to k dimensional harmonic coordinates.

For example, a three dimensional cage bounded by triangular facets possesses harmonic coordinates that reduce to barycentric coordinates on the faces. Similarly, a dodecahedral cage will have 3D harmonic coordinates that reduce to 2D harmonic coordinates on its pentagonal faces.

3 Implementation

Our current implementation of harmonic coordinates is limited to two and three dimensions. In both cases we use a simple hierarchical finite difference solver, though in principle any solution method for Laplace's equation, such as a finite element method, could be used.

Now for some details. First we'll describe the non-hierarchical version of the solver. We'll then describe the extension to the hierarchical solver. For each vertex C_i of the cage, we approximate $h_i(p)$ over the interior of the cage as follows:

1. Allocate a regular grid of cells that is large enough to enclose the cage. We choose the grid to contain 2^s cells on a side. All two dimensional examples have been computed with $s = 6$; three dimensional examples use $s = 7$. Each grid cell contains a value, and a tag, where the tag is one of UNTYPED, BOUNDARY, INTERIOR, or EXTERIOR.

2. Initialize the grid by:

- (a) Tag all cells as UNTYPED.

- (b) Scan-convert boundary conditions into the grid, marking each scan converted cell with the BOUNDARY tag. In two dimensions, the function $\phi_i(p)$ as defined in Section 2 is scan-converted into the grid. In three dimensions, our implementation is currently restricted to triangular faces, meaning that the boundary values varying in a piecewise linear fashion. We therefore use a simple voxel-based triangle scan-converter in this stage.
 - (c) Starting with one of the corner cells, flood fill the exterior, marking each visited cell with the EXTERIOR tag. The flood fill recursion stops when BOUNDARY tags are reached. Since the boundary is closed, only the exterior cells are visited during this stage.
 - (d) Mark remaining UNTYPED cells as INTERIOR with harmonic coordinate value equal to 0.
3. Laplacian smooth: For each INTERIOR cell, replace the value of the cell with the average of the value of its neighbors. In 2D cells are considered to be 4-connected; in 3D they are considered to be 6-connected. This Laplacian smoothing step is performed iteratively until the termination criterion is reached. Our solver terminates when the average change to a cell drops below a specified threshold τ . All examples in this paper have used $\tau = 10^{-5}$.

The solver described above can be significantly accelerated by noting that Laplace’s equation produces very smoothly varying functions. By first solving the problem at a lower resolution, better starting points for the iteration can be obtained. The hierarchical solver exploits this observation by “pulling” the boundary conditions up to a coarser level, recursively solving there, “pushing” the coarse solution down to the finer level, then iterating the Laplacian smoothing step until convergence is reached.

The pulling step in two dimensions computes a coarse level grid of size $2^{s-1} \times 2^{s-1}$ from a fine level grid of size $2^s \times 2^s$. Each coarse level grid cell represents four “children” cells on the finer level. In three dimensions, each coarse level grid cell represents eight children cells on the finer level. In both cases, a coarse cell is tagged as a BOUNDARY if at least one child is tagged as a BOUNDARY; it is tagged as EXTERIOR if all children are EXTERIOR, and it is tagged as INTERIOR if all children are INTERIOR. The value of a coarse level BOUNDARY cell is the average of the finer level BOUNDARY cells. INTERIOR cells on the coarse level are initialized with a value of zero.

The pushing step propagates values from coarse level cells to INTERIOR cells on the finer level. Specifically, all INTERIOR cells on the finer level receive the value of their parent cell on the coarse level.

4 Results

The behavior of harmonic deformations in two dimensions is illustrated in the accompanying video as well as in Figure 2. The behavior of three dimensional harmonic deformations is illustrated in Figure 1, where we have bound an object containing 8019 vertices to a cage containing 112 vertices. The binding time for this example was 262 seconds, using the hierarchical solver with a finest grid with 2^7 cells on a side, and a coarsest grid with 2^4 cells on a side. The termination tolerance τ was 10^{-5} . The corresponding bind time for mean value coordinates was 443 seconds using the algorithm as published in Figure 4 of [Ju et al. 2005].

Notice that the bind time for harmonic coordinates is faster than that for mean value coordinates in this case. The primary reason is that the harmonic coordinate solver computes an entire coordinate

| Subdivisions | Object vertices | MVC (in sec) | HC (in sec) |
|--------------|-----------------|--------------|-------------|
| 0 | 21 | 0.16 | 29 |
| 2 | 242 | 1.7 | 30 |
| 4 | 3842 | 24 | 30 |
| 5 | 15,362 | 113 | 30 |

Table 1: A comparison of the binding time of the mean value and harmonic coordinate solvers as the number of object points increases. These examples were generated by using an icosahedron as the cage, and a subdivided dodecahedron as the object. The “Subdivisions” column indicates the number of times the dodecahedron was subdivided. Note that the time required for harmonic coordinates is relatively insensitive to the number of object vertices.

function at a time. Once a coordinate function is computed it is very inexpensive to look up the value for each of the object points. The running time of the harmonic solver is therefore most strongly dependent on the number of cage vertices. The mean value solver, on the other hand, iterates over the entire cage for each of the object points. When the number of object points is small compared to the number of cage points, the mean value solver is faster. As the number of object points increases, the harmonic solver eventually outperforms the mean value solver. This trend is demonstrated in Table 1.

One potential disadvantage of harmonic coordinates compared to mean value coordinates is memory overhead. A straightforward implementation of mean value coordinates requires only a constant amount of additional memory for each face of the cage, whereas the memory requirements of our simple harmonic solver is dominated by the solver grid. In two dimensions the grids are typically small (roughly 40Kbytes in our examples), but in three dimensions the grids can become rather large (roughly 25Mbytes in our examples).

Harmonic coordinates computed as above are only numerical approximations, where the cell size and termination threshold determine the accuracy. The approximation error will, in general, cause each object point p to experience a residual $\Delta(p)$ when it is bound to the cage:

$$\Delta(p) = p - \sum_i h_i(p) C_i \quad (12)$$

Another source of error occurs when coordinates below a threshold are removed from the sum, a process we call *sparsification*. Residuals due to sparsification occur for both mean value coordinates and harmonic coordinates. In cases where the residuals are too large, either because of an inaccurate solve or because of overly aggressive sparsification, the residuals can be computed and stored at bind time on a per object point basis. They can then be added back at deformation time to improve the accuracy of the deformation with little run-time overhead. The examples used in this paper and the accompanying video were accurate enough that residuals were not used.

4.1 Extension to cell complexes

Harmonic coordinates as formulated thus far are defined relative to a cage consisting of a polytope, meaning the deformation is controlled entirely by boundary vertices of the cage. Once the behavior of the boundary is set, the behavior on the entire interior is completely determined. In many instances this is ideal. However, it is sometimes helpful to give artists additional control over interior details of the deformation. A simple example is shown in Figure 5 where an additional isolated vertex has been added to refine control of the deformation in the area of the head of the character.

It is also possible to extend the cage to include interior faces and edges; an example of including a collection of interior edges is




Figure 5: An example of interior control. (a) shows the cage and object at bind time, where an isolated interior vertex has been added to the cage; (b) shows the deformed object in response to movement of the interior vertex.

shown in Figure 6. As demonstrated in this figure, the interior controls need not form a manifold — it is sufficient for the interior of the cage to form what is known as a *linear cell complex*. Intuitively, a linear cell complex is a collection of “cells” (vertices, linear edges, and planar faces) with the property that the intersection of any two cells is either empty or is another cell in the collection. Harmonic coordinates are easily adapted to such cages by treating the interior facets in exactly the same way as the boundary.³

Since harmonic functions are guaranteed to be only continuous at interior boundary conditions, harmonic coordinates are only C^0 smooth across interior facets. In practice this means that if interior facets are used and smoothly deformed objects are desired, the interior facets should be placed so the object being deformed does not cross them.

5 Summary

We have provided a new and easy to implement construction for generalized barycentric coordinates as solutions to Laplace’s equation subject to carefully chosen boundary conditions — ones that correspond to lower dimensional solutions to Laplace’s equation. These harmonic coordinates improve on mean value coordinates in that they are guaranteed to be positive everywhere in the interior of the cage, and their influence falls off with distance as measured within the cage. Moreover, we have shown that the construction of harmonic coordinates can be carried out in any dimension, and we’ve show that the cage can be augmented with additional interior vertices, edges, and faces to provide more detailed control when necessary.

Unlike mean value coordinates, harmonic coordinates are defined only within the cage, and they do not possess a closed form expression. The memory requirements of harmonic coordinates are also considerably larger than mean value coordinates, especially in three dimensions. However, we’ve show that they can be efficiently approximated using a hierarchical solver, and we’ve shown that in the common case where the number of object points greatly exceeds the number of cage points, they are faster to compute than mean value coordinates. Once coordinates have been computed, the cost of evaluating harmonic deformations is identical to that of deformations based on mean value coordinates.

References

BASS, R. 1995. *Probabilistic Techniques in Analysis*. Springer-Verlag.

³Note to reviewers: we have not yet implemented harmonic coordinates for cell complexes in 3D, but we will have such examples prior to final publication.




Figure 6: An example of interior control using a linear cell complex. (a) An object deformed using a cage with no interior controls. (b) the bind-time situation for the same object where an interior cell complex has been added to the cage. (c) the deformed object resulting from cage with interior controls. Note that the influence of the modified cage point is much more local in (c).

FARIN, G. 2002. *Curves and surfaces for CAGD: a practical guide*, 5th ed. Morgan Kaufmann Publishers Inc.

FLOATER, M. S., KOS, G., AND REIMERS, M. 2005. Mean value coordinates in 3d. *Computer Aided Geometric Design* 22, 623–631.

FLOATER, M. 2003. Mean value coordinates. *Computer Aided Geometric Design* 20, 1, 19–27.

JU, T., SCHAEFER, S., AND WARREN, J. 2005. Mean value coordinates for closed triangular meshes. *ACM Trans. Graph.* 24, 3, 561–566.

LOOP, C. T., AND DEROSE, T. D. 1989. A multisided generalization of bézier surfaces. *ACM Trans. Graph.* 8, 3, 204–234.

MEYER, M., LEE, H., BARR, A., AND DESBRUN, M. 2002. Generalized barycentric coordinates for irregular polygons. *Journal of Graphics Tools* 7, 1, 13–22.

PINKHALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 15–36.

PORT, S. C., AND STONE, C. J. 1978. *Brownian Motion and Classical Potential Theory*. Academic Press.

R. SIBSON. 1981. A brief description of natural neighbor interpolation. In *Interpreting Multivariate Data*, V. Barnett, Ed. John Wiley, 21–36.

WACHPRESS, E. 1975. *A Rational Finite Element Basis*. Academic Press.

WARREN, J. 1996. Barycentric coordinates for convex polytopes. *Advances in Computational Mathematics* 6, 97–108.