
A Log-Barrier Method for Mesh Quality Improvement

Shankar P. Sastry¹, Suzanne M. Shontz², and Stephen A. Vavasis³

¹ Department of Computer Science and Engineering, The Pennsylvania State University, University Park, PA, U.S.A., sps210@cse.psu.edu

² Department of Computer Science and Engineering, The Pennsylvania State University, University Park, PA, U.S.A., shontz@cse.psu.edu

³ Department of Combinatorics and Optimization, University of Waterloo, Waterloo, ON, Canada, vavasis@math.uwaterloo.ca

Summary. The presence of a few poor-quality mesh elements can negatively affect the stability and efficiency of a finite element solver and the accuracy of the associated partial differential equation solution. We propose a mesh quality improvement method that improves the quality of the worst elements. Mesh quality improvement of the worst elements can be formulated as a nonsmooth unconstrained optimization problem, which can be reformulated as a smooth constrained optimization problem. Our technique solves the latter problem using a log-barrier interior point method and uses the gradient of the objective function to efficiently converge to a stationary point. The technique can be used with convex or nonconvex quality metrics. The method uses a logarithmic barrier function and performs global mesh quality improvement. Our method usually yields better quality meshes than existing methods for improvement of the worst quality elements, such as the active set, pattern search, and multidirectional search mesh quality improvement methods.

Keywords: mesh quality improvement, interior point method, log-barrier method

1 Introduction

High-quality meshes are necessary for the stability [1] and efficiency of finite element (FE) solvers and for the accuracy [2, 3, 4] of the solution of the associated partial differential equations (PDEs). Poor quality elements affect the conditioning of the linear system that arises from the PDE and the mesh [1]. Therefore, mesh quality improvement algorithms are used to improve the quality of mesh elements when the initial qualities of the elements obtained from mesh generators are poor or after mesh warping [5, 6, 7]. There are numerous geometric mesh quality improvement algorithms which are effective in improving the average mesh quality [8, 9, 10, 11, 12]. However, even a few poor quality elements can negatively affect the entire finite element analysis

due to the resulting ill-conditioned matrices that hinder the efficiency and the accuracy of the finite element solvers [13]. Therefore, we focus on the development of an algorithm to improve the worst quality mesh element. Our algorithm solves a smooth constrained optimization problem to improve the worst quality elements and can be classified as an interior point method [14].

Other algorithms have been developed for improvement of the quality of the worst element. Freitag and Plassmann developed an active set algorithm [15] for mesh quality improvement. However, their algorithm requires the the objective function specifying the mesh quality metrics to be convex. Park and Shontz developed two derivative-free optimization algorithms, namely the pattern search and multidirectional search methods [16], for mesh quality improvement. These algorithms do not use the gradient to optimize the mesh quality; thus, we expect their rate of convergence to be slow. They are described in more detail in Section 2.

Quality improvement of worst quality mesh elements is a nonsmooth unconstrained optimization problem. In Section 3, we describe the problem statement and show its reformulation into a constrained optimization problem. We solve this unconstrained optimization problem using an interior point method.

We develop a log-barrier interior point method [14] that seeks to improve the quality of the worst element in a mesh. Our method overcomes the disadvantages posed by the other algorithms presented above by employing a logarithmic barrier term, which is a function of the quality of the worst element. Though derived from classical optimization theory, the log-barrier method in our context has the following natural interpretation. On each iteration, the gradient of the log-barrier function points in a direction that is a weighted combination of the directions that improve each individual element. The weights are selected automatically in such a way that elements with the worst qualities have the highest weights (see (6) below). Therefore, the method globally updates vertex positions but concentrates on the improvement of the worst elements. In addition, the method has the built-in feature that the line search for maximizing the objective function will automatically prevent inversion, since the objective function (weighted element qualities) tends to minus infinity as inversion is approached. Our interior point method solves the primal formulation of the constrained optimization problem and can be used on both convex and nonconvex objective functions. The algorithm is presented in Section 4, and its characteristics are discussed in Section 5.

We run numerical experiments to assess the efficiency of our algorithm by comparing it against existing algorithms such as the active set, pattern search, and multidirectional search methods. Numerical experiments and their results are discussed in Section 6. We give our concluding remarks and indicate future research directions in Section 7.

2 Related Work

2.1 Derivative-Free Algorithms for Mesh Quality Improvement

Park and Shontz developed pattern search (PS) and multidirectional search (MDS) mesh quality improvement algorithms [16] to improve the worst quality mesh element. Their derivative-free algorithms do not compute the gradient of the objective function, but instead use function evaluations to move the vertices. These algorithms also employ local mesh quality improvement, as pattern search techniques are not efficient on large problems. The following objective function is maximized in order to improve the worst element mesh quality:

$$f(x) = \min_{1 \leq i \leq m} q_i(x), \quad (1)$$

where m is the number of elements, and $q_i(x)$ is a mesh quality metric.

Pattern Search (PS) Method. The PS method moves a mesh vertex in one of a pre-defined (usually orthogonal) set of directions. The direction and distance by which each vertex is moved is determined by evaluating objective function at the pattern points. A backtracking line search is used to untangle an element if vertex movement deems this necessary.

Multidirectional Search (MDS) Method. The MDS method uses search directions given by a simplex, i.e., a triangle (2D) or a tetrahedron (3D). The simplex is expanded, contracted, and/or reflected in order to determine the optimal position for a vertex. A backtracking line search is used to untangle an element, if needed, in a similar manner as for the PS method.

2.2 Active Set Method for Mesh Quality Improvement

Freitag and Plassmann developed an active set mesh quality improvement algorithm [15], which maximizes the quality of triangular or tetrahedral mesh elements. To guarantee convergence, the relevant objective function formed by quality metrics should possess convex level sets. Examples of such quality metrics include the minimum of the sine of the angles of the triangle in 2D and the aspect ratio quality metric in 3D for individual submeshes. This method employs a local quality improvement technique, where individual submeshes are optimized by moving one vertex at a time.

For each submesh, the objective function is defined as the quality of the worst element. The objective function described above in (1) is maximized in order to improve quality of the worst element in the mesh. They define *active value* as the current value of the objective function due to the vertex placement, x . They define *active set*, denoted by A , as a set of those functions that result in the active value. The nonsmooth optimization problem

of improving the quality of the worst element is solved using the steepest descent algorithm [14]. Here, the gradient and hence the descent direction is obtained by considering all possible convex combinations of active set gradients, $\nabla_x (q_i(x)) = g_i(x)$, and choosing the one that solves

$$\begin{aligned} \min_x \bar{g}^T \bar{g}, \text{ where } \bar{g} &= \sum_{i \in A} \beta_i g_i(x), \\ \text{s.t. } \sum_{i \in A} \beta_i &= 1, \beta_i \geq 0. \end{aligned}$$

A backtracking line search technique is used to determine the points at which the active set changes, and the vertex is moved to the point that results in the best quality improvement.

3 Mesh Quality Improvement Problem Formulation

In this section, we mathematically formulate our mesh quality improvement problem. Our problem involves the description of the quality metrics used to measure the mesh quality and the objective function that is optimized in order to improve the mesh quality. We use two quality metrics, i.e., one smooth and one nonsmooth, to demonstrate the effectiveness of our algorithm.

Aspect Ratio Quality Metric. We define the quality of tetrahedral mesh element i as

$$q_i(x) = \left(\frac{l_1^2 + l_2^2 + \dots + l_6^2}{6} \right)^{\frac{3}{2}} / \left(\text{vol} \times \frac{12}{\sqrt{2}} \right),$$

where x are the vertex positions, vol is the unsigned volume of the i^{th} tetrahedron, and l_j is the length of side j of the tetrahedron [17]. The range of this quality metric is from 1 to ∞ , where 1 is the quality of a regular tetrahedron. The quality tends to infinity as the tetrahedron becomes degenerate.

Nonsmooth Aspect Ratio Quality Metric. We define the quality of the i^{th} tetrahedral mesh element as

$$q_i(x) = \frac{\sqrt{2} l_{\max}^3}{12 \text{vol}},$$

where l_{\max} is the length of the longest edge of the tetrahedron. The range of this quality metric is from 1 to ∞ , where 1 is the quality of an equilateral tetrahedron. The quality tends to infinity as the tetrahedron becomes degenerate.

Objective Function. The problem of improving the worst quality element can be expressed as

$$\min_x \left(\max_{i \in [1, m]} q_i(x) \right).$$

However, our algorithm is designed for maximization of a quality metric. Thus, the objective function is modified by taking the reciprocal of the quality metric as follows:

$$\max_x \left(\min_{i \in [1, m]} \frac{1}{q_i(x)} \right).$$

This can be reformulated as a constrained optimization problem

$$\max_x t \text{ subject to } t < \frac{1}{q_i(x)}, \forall i \in [1, m]. \quad (2)$$

Note that all of these formulations improve the quality of the worst element, as minimizing a function is equivalent to maximizing its reciprocal.

4 A Log-Barrier Method for Mesh Quality Improvement

In this section, we describe our log-barrier method for mesh quality improvement. We develop our algorithm based on interior point methods, which can be used to solve constrained optimization problems [14]. Our method uses a logarithmic barrier term, which emphasizes the improvement of poor quality elements, and solves the constrained optimization problem using the gradient of the objective function.

4.1 Interior Point Methods

Interior point methods are a class of methods used to solve constrained optimization problems. For a constrained optimization problem, the objective function, $f(x)$, is maximized while respecting the constraint, $g(x) < 0$. When interior point methods are employed, the constraint is added as logarithmic barrier term to the objective function. The new objective function, $F(x, \mu)$, is given by:

$$F(x, \mu) = f(x) + \mu \log(-g(x)),$$

where $\mu > 0$. As we will describe in Section 5.1, μ is chosen such that it enables the satisfaction of the Karush Kuhn Tucker (KKT) conditions (i.e., first order, necessary conditions) [14]. The modified objective function is iteratively maximized using an unconstrained optimization algorithm. On every iteration, μ is reduced so that the barrier term is eventually negligible, and the original objective function, $f(x)$, is maximized subject to the constraint. Pseudocode for a typical interior point method is presented in Algorithm 1.

Algorithm 1 Interior Point Method

Input: $f(x)$, $g(x)$, tol.
Output: $\max_x f(x)$ such that $g(x) < 0$.
Initialize $\mu_k > 0$ and x_0 such that $g(x_0) < 0$.
while $\mu_k \geq \text{tol}$ **do**
 Maximize $f(x) + \mu_k \log(-g(x))$ using any gradient-based optimization algorithm.
 Decrease μ_k towards 0.
end while

4.2 The Log-Barrier Method for Mesh Quality Improvement

We seek to improve the worst mesh element quality. A quantity, t , is defined, which is a function of the worst quality element such that $t < \min_i \frac{1}{q_i(x)}$, where $q_i(x)$ is the quality of the i^{th} element. When t is maximized, the quality of the worst element is improved. The expression $t - \frac{1}{q_i(x)} < 0$ is used as a constraint. We iteratively maximize

$$t + \mu_k \sum_{i=1}^m \log\left(\frac{1}{q_i(x)} - t\right) \quad (3)$$

to improve the quality of the worst element in the mesh. After every iteration, μ is brought closer to 0. We modify t such that $\frac{d}{dt}(F(x, t, \mu)) \approx 0$. For a fixed μ and x , the objective function is strictly concave in t . Therefore, setting its derivative to 0 corresponds to globally maximizing the objective w.r.t. t . The log-barrier method for mesh quality improvement is shown in Algorithm 2.

Algorithm 2 Interior Point Method for Mesh Quality Improvement

Initialize μ_k and $t < 1/q_i(x)$ for all $i \in [1, m]$ where $q_i(\cdot)$ is the quality metric function.
while not converged **do**
 Maximize $F(x, t, \mu_k) = t + \mu_k \sum_{i=1}^m \log\left(\frac{1}{q_i(x)} - t\right)$, where t and μ are held constant, using the nonlinear conjugate gradient method.
 Decrease μ_k towards 0.
 Update t to a new value such that $\frac{d}{dt}(F(x, t, \mu)) \approx 0$.
end while

Log Barrier Term for Nonsmooth Quality Metrics. In our paper, the nonsmooth aspect ratio quality metric is defined using the longest edge of a tetrahedron. Our method handles the metric using each of the edges in the tetrahedron to compute six qualities for the tetrahedron and uses them as additive terms in the log barrier function. Since each individual term (for

each edge in the tetrahedron) is smooth, the resulting log barrier function is also smooth.

5 Characteristics of the Log-Barrier Method for Mesh Quality Improvement

In this section, we show that the set of first-order necessary conditions, i.e., the KKT conditions [14], are satisfied for a solution of our constrained optimization problem. This fact is well known in the optimization community; we are including it here for the sake of completeness and also to provide more detail about the properties of the objective function and the algorithm. In addition, we examine the monotonicity of our algorithm.

5.1 Satisfaction of the KKT Conditions

For a solution, x^* , of a constrained optimization problem, the gradient of the Lagrangian vanishes at the solution, i.e., $\nabla_x L(x^*, t^*, \lambda^*) = 0$. For (2), the Lagrangian is given by

$$L(x, t, \lambda) = t + \sum_{i=1}^m \lambda_i \left(\frac{1}{q_i(x)} - t \right).$$

Hence, its gradient is given by

$$\nabla_x L(x, t, \lambda) = \sum_{i=1}^m \lambda_i \nabla_x \left(\frac{1}{q_i(x)} \right). \quad (4)$$

The nonlinear conjugate gradient (CG) step in the log barrier method computes x such that the gradient of the objective function given by (3), i.e., $\nabla_x F(x, t, \mu_k)$, vanishes. Thus,

$$\nabla_x F(x, t, \mu_k) = \mu_k \nabla_x \sum_{i=1}^m \log \left(\frac{1}{q_i(x)} - t \right) \quad (5)$$

$$= \mu_k \sum_{i=1}^m \frac{1}{\left(\frac{1}{q_i(x)} - t \right)} \nabla_x \left(\frac{1}{q_i(x)} \right). \quad (6)$$

From Equations (4) and (6), we see that, if λ_i satisfies

$$\lambda_i^* = \frac{\mu_k^*}{\frac{1}{q_i(x^*)} - t^*}, \quad (7)$$

then the solution obtained by our method satisfies the stationarity requirement of the KKT conditions. The stationarity conditions are satisfied, as

$$\nabla_x L(x^*, t^*, \lambda^*) = \sum_{i=1}^m \lambda_i^* \nabla_x \left(\frac{1}{q_i(x^*)} \right) = 0.$$

Primal feasibility is also satisfied, since

$$\frac{1}{q_i(x^*)} - t^* \geq 0.$$

Dual feasibility is satisfied if

$$\lambda_i^* \geq 0.$$

From (7), and since $\mu_k > 0$ and $\frac{1}{q_i(x^*)} - t^* > 0$ at the solution, we have

$$\lambda_i^* \geq 0.$$

The complementarity condition requires that

$$\lambda_i^* \left(\frac{1}{q_i(x^*)} - t^* \right) = 0.$$

Substituting for λ_i^* , we see that

$$\lambda_i^* \left(\frac{1}{q_i(x^*)} - t^* \right) = \mu_k.$$

The log-barrier method drives μ_k to 0 as $k \rightarrow \infty$. Thus, the complementarity condition is also satisfied. Therefore, our log-barrier method converges to stationary points. Our implementation explicitly checks that the line search exploration moves the vertices in an ascent direction.

5.2 Monotonicity

In our algorithm, the optimization method maximizes the objective function given by (3),

$$F(x, t, \mu_k) = t + \mu_k \sum_{i=1}^m \log \left(\frac{1}{q_i(x)} - t \right),$$

on every iteration. Because t and $\mu_k > 0$ are constants for a given iteration, the maximization of the objective function is equivalent to maximization of the sum of the logarithmic terms. This is equivalent to maximizing the product of the terms (without taking their logarithms).

For simplicity of the analysis, let us now examine the monotonicity of our algorithm when employed on a patch having only two elements. If we plot the qualities of the two elements on the X and Y axes, we obtain hyperbolic contours representing the objective function as shown in Fig. 1. In Fig. 1, P represents a patch with near-equal qualities of the elements, and Q represents a patch with unequal element qualities. The symbols a , b , c , and d represent

the paths the patches can take in order to maximize the objective function. Ideally, P should take path b , and Q should take path d so that the qualities of both the elements improve. In many cases, this is not possible, as improving the quality of one of the elements decreases the quality of the other. In the near-equal case, if P takes path a , the quality of the worst element decreases. Thus, we see that our algorithm may not monotonically increase the quality of the worst element in the mesh. For the unequal case, path c also improves the quality of the worst element.

Fig. 2 shows how the nonsmooth objective function for maximizing the worst quality element in Fig. 2(a) is converted to a smooth objective function in Fig. 2(b). In Fig. 2(a), the nonsmooth aspect ratio is plotted for a patch with a free vertex in the square formed by the diagonal from $(0, 0)$ to $(1, 1)$. Other vertices are on the perimeter of the square at $(0, 0)$, $(0, 0.5)$, $(0, 1)$, $(0.5, 1)$, $(1, 1)$, $(1, 0.5)$, $(1, 0)$, and $(0.5, 0)$. Note that the contours are nonsmooth when plotting the worst quality element in the patch. For illustration purposes, we chose the point $(0.1, 0.1)$, where the function is nonsmooth and set t in the log-barrier objective function as some quantity less than the worst element quality in the patch with the free vertex at $(0.1, 0.1)$. When the contours of the objective function are plotted, we see in Fig. 2(b) that they are smooth. Our algorithm moves the free vertex at $(0.1, 0.1)$ to a point close to $(0.5, 0.5)$.




Fig. 1. Illustration of possible nonmonotonicity in the convergence of our algorithm. The X and Y axes represent the qualities of two elements in a patch, and the hyperbolic contours represent the sum of the two qualities on a logarithmic scale (which is maximized in an iteration). P and Q are possible locations of qualities of the patch. The symbols a , b , c , and d are the possible paths our algorithm can take. Although the objective function is maximized, notice that the quality of the worst element may not improve in all cases.

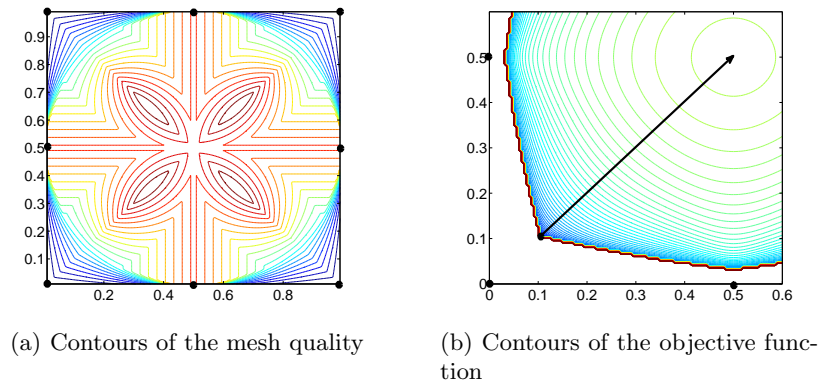


Fig. 2. Contours of the worst element quality and log-barrier objective function for a patch with eight vertices on the perimeter. (a) Contours of the nonsmooth aspect ratio quality metric. They are nonsmooth at points where two worst quality elements are present. (b) Contours of the log-barrier objective function with the free vertex at $(0.1, 0.1)$. Notice how smooth the contours are.

6 Numerical Experiments

In this section, we describe the numerical experiments that we designed to evaluate the performance of our method. We implemented our algorithm, the PS, and the MDS methods in the Mesquite Mesh Quality Improvement Toolkit Version 2.0.0 [18]. The active set method was already implemented in Mesquite. For each of the methods, the movement of the surface vertices was enabled. For the star meshes, the gradient of the objective function for the boundary vertices was projected onto the respective planes. For the sphere meshes, if the boundary vertices moved away from the surface, they were snapped back onto the surface.

Star and sphere (Fig. 3) meshes of various sizes were constructed using CUBIT [19]. In order to test our algorithm on challenging meshes, 50% of the vertices in the original meshes were randomly perturbed. The following three experiments were performed.

6.1 Effect of Parameters on Algorithmic Performance

For our first experiment, the following set of parameters were modified to determine their effect on the performance of the mesh quality improvement. Three variants of the nonlinear conjugate gradient method, i.e., the Fletcher-Reeves, Polak-Ribière, and Hestenes-Stiefel variants were used to improve the mesh quality in the inner loop. Two, four, and eight CG iterations per outer iteration were used in each of the experiments. The parameter μ was reduced to 90%, 60%, and 30% of its value after every outer iteration. We used the




Fig. 3. Unperturbed coarse meshes on the two domains representative of actual meshes considered in this paper.

largest star mesh with approximately 1.012 million elements. We maximized the reciprocal of the aspect ratio quality metric for each of the subexperiments. The subexperiments were carried out for all combinations of these parameters. The numerical experiments were run until the quality of the worst element did not improve for five successive iterations.

Hestenes-Stiefel Conjugate Gradient Method. For this subexperiment, the Hestenes-Stiefel CG method was used for mesh quality improvement. The results from all the parameter combinations are shown in Fig. 4. Fig. 4(a) shows the plot of mesh quality verses time when μ was reduced to 30% of its value in every outer iteration for a set of CG iterations per outer iteration. It can be clearly seen that eight CG iterations for every outer iteration give the best mesh quality improvement. The results indicate that eight and four iterations of Hestenes-Stiefel CG give the best performance for the 60% and 90% cases, respectively.

Fletcher-Reeves and Polak-Ribière Conjugate Gradient Methods. We repeated the above subexperiment using the Fletcher-Reeves and Polak-Ribière CG methods. The results for two iterations of Polak-Ribière CG and four iterations of Fletcher-Reeves CG are show in Figs. 5(a), (b), and (c). For all the numerical experiments we conducted, it was seen that four Fletcher-Reeves CG iterations per outer iteration returned the best performance Polak-Ribière CG method gave best performance when two CG iterations were carried out for every outer iteration.

Parameters for Best Performance of Log-Barrier Method. Fig. 5 shows a summary of the best results obtained for the above subexperiments. Figs. 5(a), (b), and (c) show the best performance of each of the CG variants for μ being reduced to 30%, 60%, and 90% of its value after every outer iteration, respectively. Fig. 5(d) summarizes the best performing results from Figs. 5(a), (b), and (c). For each of the cases, the Hestenes-Stiefel CG variant gives the best performance. When μ is reduced to 90% of its value after four

iterations of the Hestenes-Stiefel CG algorithm, maximum quality improvement is seen. Thus, we use this set of parameter values in our subsequent experiments, described in Sections 6.3 and 6.4, in which we compare the performance of our methods with existing algorithms for worst element mesh quality improvement including the active set, PS, and MDS methods.

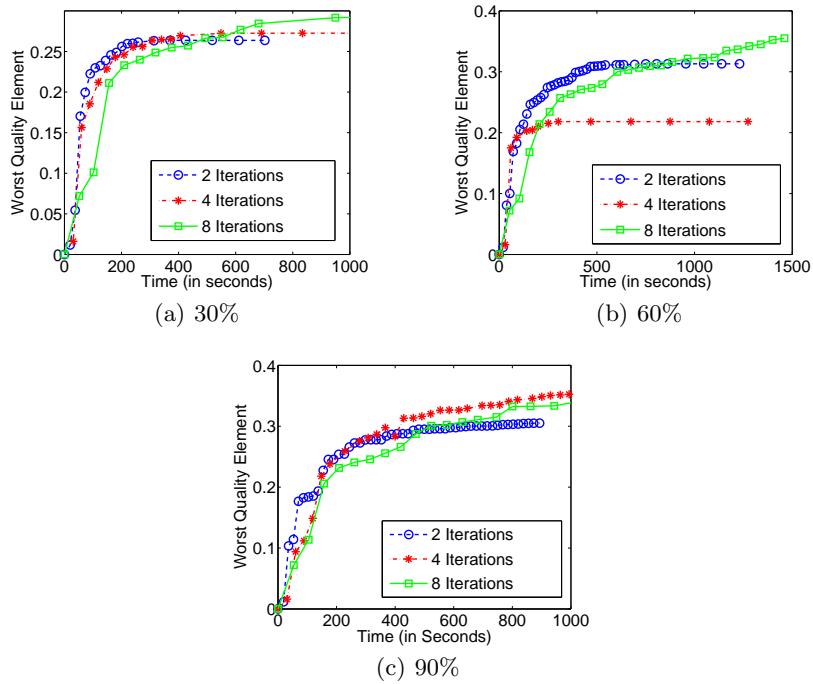


Fig. 4. Results obtained by using the Hestenes-Stiefel CG on the star mesh. The reciprocal of the aspect ratio is maximized using the log-barrier method. The percentages refer to the factor to which μ is reduced after every outer iteration. The number of iterations refer to the number of CG iterations per outer iteration.

6.2 Scalability

For this experiment, the reciprocal of the aspect ratio metric was maximized to improve the quality of the perturbed meshes using all the methods described in Section 2. Two inner iterations were carried out for every outer iteration each the method. The log-barrier method employed the Hestenes-Stiefel conjugate gradient algorithm [14] in the inner loop. After every outer iteration, μ was reduced to 90% of its value in the previous iteration. In order to accurately estimate the time per iteration, our implementation was executed for 50 iterations on each star mesh.

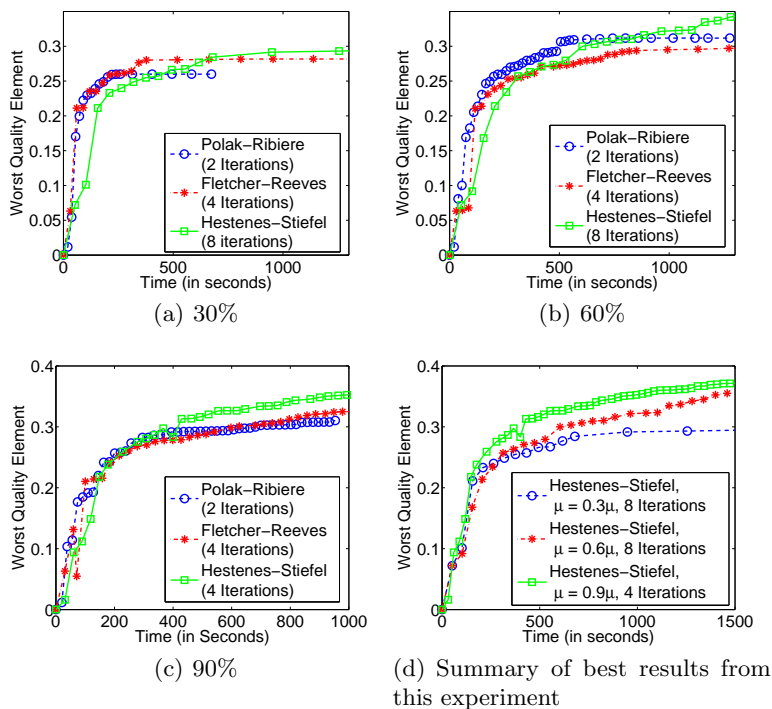


Fig. 5. Results obtained by using variants of CG on the star mesh. The reciprocal of the aspect ratio is maximized using the log barrier method. The percentages refer to the factor to which μ is reduced after every outer iteration. The number of iterations refer to the number of CG iterations per outer iteration.

The results from the scalability experiment are shown in Fig. 6, where the time per outer iteration scales linearly (for each method) with the number of elements in the mesh, as the time required to compute the gradient and move the vertices is directly proportional to the mesh size. We determined the order of convergence as a function of the problem size for our method. The order of convergence, α , is given by $T = k * m^\alpha$, where T is the time to convergence, m is the number of mesh elements, and $k > 0$ is a constant. In order to determine α , a least squares fit was computed by taking the logarithm of both sides. The value of α was found to be 0.9946.

6.3 Comparison with Existing Mesh Quality Improvement Methods for the Aspect Ratio Quality Metric

For this experiment, we used the largest meshes for each domain containing approximately 1.012 million and 1.014 million elements in the star and the

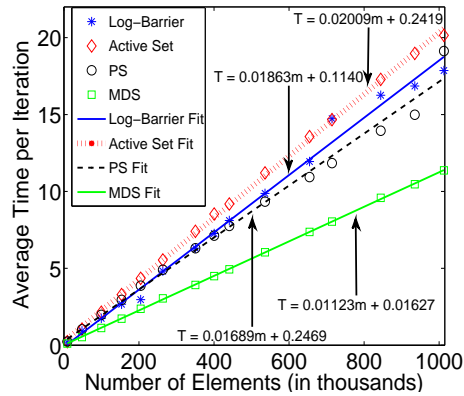


Fig. 6. Scalability experiment results. The data and the linear regression fits are shown. In the equations, T refers to average time per iteration, and m refers to number of elements (in thousands). The reciprocal of the aspect ratio quality metric was maximized. The average time per iteration for the first 50 outer iterations of all the methods were used to compute the least squares fit. For the log-barrier method, two Hestenes-Stiefel CG iterations were carried out per outer iteration.

sphere mesh, respectively. The reciprocal of the aspect ratio quality metric was maximized by the four algorithms previously discussed: the log-barrier, active set, PS, and MDS methods. Several experiments were conducted to find the values of the various parameters in each of the algorithms resulting in the best performance (measured as the worst element quality at the time by which our algorithm converged). The numerical experiments were carried out until the worst element quality remained the same for five iterations. We present results for the best performance for each of the algorithms.

The results for the mesh quality improvement on the star mesh are shown in Fig. 7(a). Our method improved the mesh quality by the greatest amount when compared to the other methods. It was followed by the MDS, PS, and then the active set method. A closer inspection of the plot reveals that, in the initial iterations, the active set method was the fastest method to significantly improve the worst quality. The method was followed by the MDS and PS methods. During the initial phase, our method was slower than the rest.

The slow convergence of the other three methods, despite their initial performance, may have been caused by the slow propagation of unequal patches due to their use of local mesh quality improvement. The active set method was able to quickly improve the worst quality by a significant amount within two iterations but became stagnant afterward. The active set method moves every vertex to the optimal location with respect to the patch.

The optimal vertex locations are approximately determined in each iteration for the PS and MDS methods. Thus, unequal patches are present

throughout the mesh. This enables steady improvement of the worst element quality in MDS and PS. In MDS, we noticed a behavior similar to the active set method where the worst element quality was constant for seven iterations and then the method converged to a mesh with a slightly better quality.

The results for the sphere mesh are shown in Fig. 7(b). As in the earlier case, our method was able to improve the mesh quality by the greatest amount. Here the PS method was very competitive and converged faster than our method, but to a lower optimal value. The MDS and active set methods also converged to a lower optimal value.

# Vertices	# Elements	Initial Worst Quality	Final Worst Quality
2,128	9,099	1.8544e-03	5.2991e-02
9,501	48,219	8.9960e-04	5.6172e-02
21,972	99,684	4.6232e-04	2.0447e-02
29,096	153,780	1.2875e-05	3.4908e-02
38,163	204,612	4.2422e-05	4.0148e-02
48,880	263,602	5.7753e-04	4.5985e-02
64,673	350,303	9.4858e-06	3.3304e-02
73,617	400,522	3.1302e-05	3.4584e-02
80,926	440,711	6.1325e-05	2.9847e-02
97,981	535,921	7.2535e-05	2.3580e-02
119,137	654,606	4.5933e-05	4.0360e-02
129,952	714,495	7.9804e-06	2.3341e-02
152,929	844,425	4.9885e-06	3.1510e-02
169,024	935,178	4.9250e-06	1.4332e-02
182,760	1,012,632	2.3352e-04	3.0506e-02

Table 1. Number of vertices and elements in the star meshes with their initial and final qualities after 50 outer iterations of quality improvement using the log-barrier method. The objective function that is formed from the reciprocal of the aspect ratio quality metric is maximized. The aspect ratio metric is a smooth, convex quality metric.

6.4 Comparison with Existing Mesh Quality Improvement Methods for the Nonsmooth Aspect Ratio Quality Metric

Through this experiment, we demonstrate that our algorithm is also efficient for mesh quality improvement when a nonsmooth or nonconvex quality metric is used to define the objective function. For this experiment, the objective function that was formed from the reciprocal of the nonsmooth aspect ratio quality metric was maximized by the log barrier, PS, and MDS methods. The active set method is designed to be used with a convex objective function, and yields a tangled mesh when used with a nonconvex quality metric. Thus, we have shown only the results for the other three methods in Fig. 8(a). The numerical experiments were carried out until the quality of the worst element did not improve for five iterations. It can be clearly seen in Fig. 8(a)

that our method yields a better quality mesh than the other methods. In Fig. 8(b), it can be seen that our log-barrier method converged to a mesh of somewhat lower quality than the PS method. We seek to determine the line search parameter values that can make our method converge to a mesh with better quality.

We also found that the PS and MDS methods were very sensitive to parameter changes, but our log-barrier method and the active set method are not as sensitive. We have shown the effectiveness of our method compared to other existing methods. Our method can also improve the quality of a mesh even when it is assessed using a nonsmooth, nonconvex metric.

7 Conclusions and Future Work

We have presented a log barrier interior point method for improving the worst quality elements in a finite element mesh. We reformulated the non-smooth problem of maximizing the quality of the worst element as a smooth constrained optimization problem, which is solved using a log barrier interior point method. Our method uses a log barrier function whose gradient places a greater emphasis on poor quality elements in the mesh and performs global mesh quality improvement. Our method usually yields a better quality mesh than other existing worst quality mesh improvement methods and scales roughly linearly with the problem size.

We have used the conjugate gradient method to perform mesh quality improvement in the inner loop. We plan to use Newton’s method for the objective function maximization instead of the nonlinear CG method because it uses second-order information which may lead to faster convergence. The constrained optimization problem can also be solved using primal-dual Newton-

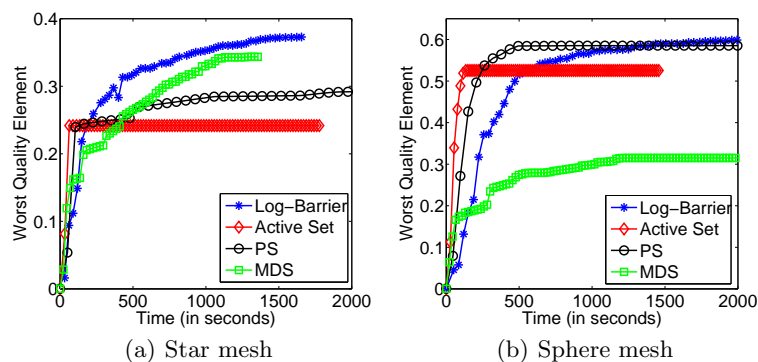


Fig. 7. Results from the experiment that compares the mesh quality improvement algorithms. The aspect ratio quality metric was improved in the meshes.

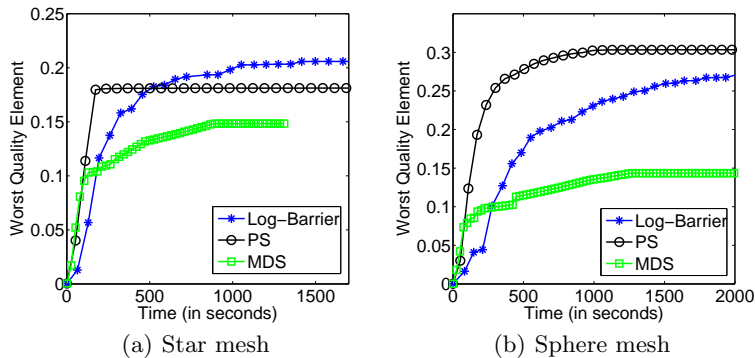


Fig. 8. Results from the experiment that compares the mesh quality improvement algorithms. The nonsmooth aspect ratio quality metric was improved in the meshes. The active set method is designed to be used with a convex objective function, and it yields a tangled mesh when used with a nonconvex quality metric, and hence its performance is not shown here.

based methods. An example of such a method includes the Mehrotra predictor-corrector method [20]. We plan to explore such methods to determine the most efficient solver for mesh quality improvement. We will also conduct research on other barrier functions that may make our method more efficient.

Our method can be naturally extended to handle mesh untangling by appropriately choosing the barrier term in the objective function. In this case, the barrier term should place the highest emphasis on inverted elements, medium emphasis on poor quality elements, and the least emphasis on good quality elements. Finally, by combining these approaches, our method may be used for simultaneous untangling and quality improvement.

Acknowledgments

The work of the first two authors was supported in part by NSF grants CNS-0720749 and NSF CAREER Award OCI-1054459. The work of the third author was supported in part by a Discovery grant from NSERC (Canada) and a grant from the U.S. Air Force Office of Scientific Research.

References

1. E. Fried, “Condition of finite element matrices generated from nonuniform meshes,” *AIAA Journal*, vol. 10, pp. 219–221, 1972.
2. I. Babuska and M. Suri, “The p and h-p versions of the finite element method, basic principles, and properties,” in *SIAM Review*, vol. 35, pp. 579–632, 1994.

3. M. Berzins, "Solution-based mesh quality for triangular and tetrahedral meshes," in *Proc. of the 6th International Meshing Roundtable*, pp. 427–436, Sandia National Laboratories, 1997.
4. M. Berzins, "Mesh quality - Geometry, error estimates, or both?," in *Proc. of the 7th International Meshing Roundtable*, pp. 229–237, Sandia National Laboratories, 1998.
5. S. Shontz and S. Vavasis, "Analysis of and workarounds for element reversal for a finite element-based algorithm for warping triangular and tetrahedral meshes," *BIT, Numerical Mathematics*, vol. 50, pp. 863–884, 2010.
6. S. Shontz and S. Vavasis, "A robust solution procedure for hyperelastic solids with large boundary deformation." *Engineering with Computers*, Published online at <http://www.springerlink.com>, DOI: 10.1007/s00366-011-0225-y., 2011.
7. P. Knupp, "Updating meshes on deforming domains: An application of the target-matrix paradigm," *Commun. Num. Meth. Engr.*, vol. 24, pp. 467–476, 2007.
8. P. Knupp, "Matrix norms and the condition number: A general framework to improve mesh quality via node-movement," in *Proc. of the 8th International Meshing Roundtable*, pp. 13–22, 1999.
9. P. Knupp and L. Freitag, "Tetrahedral mesh improvement via optimization of the element condition number," *Int. J. Numer. Meth. Eng.*, vol. 53, pp. 1377–1391, 2002.
10. N. Amenta, M. Bern, and D. Eppstein, "Optimal point placement for mesh smoothing," *In Proc. of the 8th ACM-SIAM Symposium on Discrete Algorithms*, pp. 528–537, 1997.
11. T. Munson, "Mesh shape-quality optimization using the inverse mean-ratio metric," *Mathematical Programming*, vol. 110, pp. 561–590, 2007.
12. A. Plaza, J. Suárez, M. Padrón, S. Falcón, and D. Amieiro, "Mesh quality improvement and other properties in the four-triangles longest-edge partition," *Comput. Aided Geom. D.*, vol. 21, no. 4, pp. 353 – 369, 2004.
13. J. Shewchuk, "What is a good linear element? Interpolation, conditioning, and quality measures," in *Proc. of the 11th International Meshing Roundtable*, pp. 115–126, Sandia National Laboratories, 2002.
14. J. Nocedal and S. Wright, *Numerical Optimization*. New York: Springer, 2nd ed., 2006.
15. L. Freitag and P. Plassmann, "Local optimization-based simplicial mesh untangling and improvement," *Int. J. Numer. Meth. Eng.*, vol. 49, pp. 109–125, 2000.
16. J. Park and S. Shontz, "Two derivative-free optimization algorithms for mesh quality improvement," in *Proc. of the 2010 International Conference on Computational Science*, vol. 1, pp. 387–396, 2010.
17. V. Parthasarathy, C. Graichen, and A. Hathaway, "A comparison of tetrahedron quality measures," *Finite Elem. Anal. Des.*, vol. 15, pp. 255–261, 1994.
18. M. Brewer, L. Freitag Diachin, P. Knupp, T. Laurent, and D. Melander, "The mesquite mesh quality improvement toolkit," in *Proc. of the 12th International Meshing Roundtable*, pp. 239–250, Sandia National Laboratories, 2003.
19. "CUBIT generation and mesh generation toolkit." <http://cubit.sandia.gov/>.
20. S. Mehrotra, "On the implementation of a primal-dual interior point method," *SIAM J. Optimiz.*, vol. 2, no. 4, pp. 575–601, 1992.