

INCREASING TAU3P ABORT-TIME VIA MESH QUALITY IMPROVEMENT

Nate Folwell^{1*}

Patrick Knupp^{2†}

Michael Brewer²

¹*Stanford Linear Accelerator Center, Stanford CA U.S.A. folwell@slac.stanford.edu*

²*Sandia National Laboratories, Albuquerque, NM U.S.A {mbrewer,pknupp}@sandia.gov*

ABSTRACT

The Tau3P time-domain computational electromagnetics code is used at The Stanford Linear Accelerator Center for a wide variety of accelerator design tasks. The code uses the Discrete Surface Integral method to solve Maxwell's equations on primal and dual hexahedral meshes. In this method run times are highly sensitive to mesh quality, particularly to minimum edge-length, element skew, and lack of smoothness. In this study we investigate methods for increasing Tau3P run times via mesh quality optimization. It is found that abort-time can be significantly increased via optimization of the mesh condition number or by Laplacian smoothing.

Keywords: mesh quality, mesh optimization, mesh smoothing, computational electromagnetics

1. INTRODUCTION

The Stanford Linear Accelerator Center (SLAC) is designing particle accelerators, in part, by performing computer simulations using a code called Tau3P to solve Maxwell's Equations in the time domain. The discretization algorithm used in Tau3P is the Discrete Surface Integral (DSI) scheme [1]. The algorithm has been proven to be conditionally stable for strictly orthogonal structured grids. The particle accelerator hardware consists of waveguides and other parts forming complex 3D geometrical configurations that are impossible to mesh with strictly orthogonal structured grids. Presently SLAC is using the CUBIT mesh generation code [2] to create unstructured hexahedral and tetrahedral meshes which accurately represent the geometry but may lack the high quality of orthogonal mapped meshes. No DSI stability condition has been proven for these more general meshes. Empirically, some simulations using these meshes exhibit unstable

behavior, suggesting that the DSI scheme is not even conditionally stable for the unstructured grids.

To work around this issue, a filtering technique [3] is used in Tau3P which permits useful calculations to be performed by delaying the onset of instability. With this filter, most Tau3P calculations are able to complete to the desired problem end-time. In some runs, however, the numerical solution grows exponentially, swamping the problem before the desired end-time is reached. To avoid computing meaningless solutions, Tau3P has a mechanism by which runs can be aborted before the end-time is reached. A comparison is made between the magnitude of the electric field vector at selected points and a value determined automatically given the problem. If the electric field vector exceeds the supplied value then the results are most likely non-physical, due to the unstable nature of the discretization algorithm. If this situation occurs, the run is aborted. We define the *abort-time* of a run to be the time (if any) at which the run is aborted.

It has long been observed that the abort-time encountered in many Tau3P simulations is sensitive to properties of the mesh. Much time has been spent tinkering with and regenerating meshes so that a given

*The first authors' work was funded by the DOE SciDAC Accelerator Project

†The second and third authors' work was funded by the DOE SciDAC TSTT Center under contract DE-AC04-94AL85000

simulation will not abort before the desired problem end-time is reached. To reduce the time needed to create an acceptable mesh, it would be helpful to know precisely which mesh properties play a crucial role in determining abort-time. An empirical investigation to identify the strongest correlations between the abort-time and various mesh properties [4] was conducted. Having identified the crucial mesh properties, we seek in this study to improve mesh quality with respect to the crucial properties, resulting in calculations that have a higher probability of running to completion.

2. MESH QUALITY METRICS

In the previous report [4], a series of meshes was studied to determine which quality metrics are most important in affecting abort-time. The meshes were created on a pillbox geometry (see Figure 1) that is representative of a simplified accelerator structure. Four cases (PBU, PBF, PPU, and PPF) were considered, resulting from combinations of beam/pulse excitation and unfiltered/filtered. To limit the number of cases to a manageable number, the present study focuses on just one of the four cases. The PPU case was selected because results are easiest to interpret from a physical viewpoint. The study showed that the most important metrics were:

- Minimum primal edge size (MPES).¹ This is the minimum edge length in the mesh, including boundary edges. In general, this metric can range from zero to infinity. For the pillbox geometry, the approximate acceptable range of this metric is MPES greater than 0.00040 meters (edge-lengths smaller than that tend to result in relatively small abort-times).
- Maximum primal condition number (MPCN). This is the maximum quadrilateral element condition number in the mesh. The quadrilateral condition number is defined as the maximum condition number at any of the four vertices of the quadrilateral. The condition number at a vertex is the sum of the squares of the two adjacent edge lengths, divided by twice the area. This metric ranges from one to infinity. The approximate acceptable range is between 1.0 and 2.0, with 1.0 being ideal. MPCN larger than 2.0 tends to cause relatively small abort-times.

¹MPES is not a scaled metric, but an absolute number. Because the comparisons we make are always on the same model, scaling is not an issue for the results presented in this paper. We recognize that it could be a problem if we were to try to make comparisons across geometric models. We have devised a number of scaled metrics to supplant MPES, but our studies to date have not found them as useful.




Figure 1: A typical hexahedral mesh of the pillbox geometry.

- Maximum primal smoothness metric (MPSM). Let A_k be the area of the k -th quadrilateral element in the cross-section of the mesh. Let element j share an edge with element k . Let

$$s_{jk} = \left| \frac{A_k - A_j}{A_k + A_j} \right|$$

$$t_k = \max_j s_{jk}$$

$$MPSM = \max_k(t_k) / \bar{t}$$

where \bar{t} is the average value of t_k over the entire mesh. This metric is effective in detecting jumps in area between adjacent quadrilaterals. If element k is on the boundary, it is not included in the calculation. This metric is only meaningful provided the elements areas are all positive. This metric ranges from zero to infinity, with an approximate acceptable range between zero and 5.0. When MPSM exceeds 5.0, abort-times tend to be relatively small.

- Maximum primal scaled Jacobian (MPSJ). This is the maximum sine of the angles in the mesh. The metric ranges from -1.0 to +1.0, with an approximate acceptable range of 0.40 to 1.0. When MPSJ is less than 0.40, abort-times tend to be relatively small.

These metrics reflect some of the basic properties of geometric mesh quality: size, smoothness, angle, and aspect ratio. Many other metrics which measure these generic properties are also correlated with abort-time and thus could have been used in this study.

The previous study [4] established the importance of these metrics by generating a set of 25 pillbox




Figure 2: Scatter plots for MPCN-2D metric

meshes and determining an abort-time for each of the meshes. For each of the 25 meshes one may calculate various mesh quality metric values. Each metric value can, in turn, be paired with the abort-time for the given mesh to create a scatter plot. If a given metric is a poor predictor of abort-time, the scatter plot will show no pattern. On the other hand, a strong predictor of abort-time will show a distinctive 'L' pattern. Figure 2 shows the scatter plot obtained for the Maximum Primal Condition Number metric. This metric exhibits the 'L' pattern: if the mesh has good quality (small condition number values) then the abort-time is independent of the metric (horizontal branch of the 'L'). In that case, the abort-time is determined in conjunction with other mesh quality metrics. If the mesh has poor quality (large condition numbers), the points in the scatter plot lie on the vertical branch of the 'L'. This indicates that the abort-time is relatively small and that this is likely due to mesh elements having poor shape quality (large condition number). If a metric exhibits a strong 'L' pattern in its scatter plot, one can define an *acceptable range* for the metric, namely, the set of values for which abort-time is independent of the metric. In the case of the MPCN scatter plot, for example, the acceptable range is approximately between 1 and 2.

The present situation for Tau3P meshing is iterative:

- Generate a mesh for a given geometry,
- Perform a Tau3P run,
- If the run aborts (prematurely), generate a new mesh and return to previous step.

- If the run does not abort prematurely, halt.

Such a procedure is clearly expensive in terms of the number of man-hours that potentially may be devoted to a given meshing problem because both the meshing step and the Tau3P calculation are within the iterative loop. Smoothing is one way of producing another mesh for the iterative procedure. Our results will show that smoothing often produces a mesh which increases the abort-time, thus increasing the likelihood that the iteration will halt. However, to further reduce the time spent in meshing, the following modified iteration is proposed:

- Generate a mesh for a given geometry,
- Calculate the quality metrics (MPES, etc),
- If any metric lies outside the acceptable range, improve the mesh in various ways until the metrics are acceptable,
- Perform a Tau3P run, halt the procedure

In this approach, the iteration loop does not include Tau3P, so the iteration proceeds much faster than the previous iterative scheme. If one's best mesh still results in a pre-maturely aborted Tau3P run, one should probably not continue to search for better meshes because one has already used the best mesh.

The goal of the present and future investigations is to find ways to improve Tau3P abort-time on the PPU (Pulse Pillbox Unfiltered) and, later, the other cases by modifying the meshes to improve the various mesh quality metrics. In doing so, it is believed that a methodology will be developed that can be applied to more realistic meshing problems that commonly arise in the modeling and simulation of accelerators. Our tools for improving the abort-time are smoothing (node repositioning), edge biasing, interval assignment, remeshing of certain regions, and inserting prisms. In this paper we confine our activities to smoothing.

3. SMOOTHING STRATEGY

A typical 3D pillbox mesh (See Figure 1) consists of hexahedral elements. The hexahedral mesh is generated by first creating quadrilateral meshes on the beam cross-section (a quarter circle) and on the collar cross-section (a quarter annulus). These meshes are 'swept' along the beam axis to create hexahedra.

We limited our attention in this study to sixteen 2D beam cross-section meshes. The previous study showed that these cross-sections cover the full range of Tau3P abort-times. QC2 is our 'best' mesh in that,

without smoothing, it resulted in the largest abort-time. QC3 is our 'worst' mesh, having a singularity on the boundary.² The abort-time for QC3 is several orders of magnitude smaller than the other meshes. Four representative cross-section meshes are shown in Figure 3.

Our approach to smoothing of the pillbox mesh has potentially two stages: (1) smooth the beam cross-section mesh consisting of quadrilateral elements (nodes on the bounding curves of the cross-section are held fixed), and (2) smooth the hexahedral mesh (holding nodes on the bounding surfaces fixed).

We do not smooth the collar cross-section mesh (see Figure 4) since it would be difficult to improve the quality obtained on that surface (the meshes are smooth and nearly orthogonal). For the same reason, we do not attempt to improve the quality of any of the linking surface meshes. Finally, since the 3D mesh is a simple translation of the 2D cross-section meshes, there is little to be gained by smoothing the hexahedral mesh as in Step 2 above. For more general geometries, the choice of which surfaces to smooth would have to be considered on a case by case basis.³ This issue will be explored in future work.

In the original study there were sixteen quadrilateral meshes generated on the beam (quarter-circle) cross-section. For each such cross-section, there are two meshes of interest in this study which we denote by the following types:

- Type UU. Neither the quadrilateral nor the hexahedral meshes are smoothed,
- Type SU. The quadrilateral mesh on the beam cross-section is smoothed, then swept to create the hexahedral mesh. The hexahedral mesh is not smoothed.

4. SMOOTHING THE BEAM CROSS-SECTION

In this section we compare Type UU meshes to Type SU meshes to explore potential benefits of 2D mesh smoothing to improve quality (and thus abort-time). To create the smoothed SU meshes, we first try condition number smoothing [5], as implemented in the CUBIT code. This smoother minimizes the following

²That does not mean that there is an element with zero area. Rather the underlying continuum map has a singularity. No element in this mesh has zero area.

³On more complex geometries one may still wish to smooth some or all of the surface meshes, thus 2D metrics remain relevant. On the other hand, 3D metrics are clearly needed when one needs to smooth the volume.




Figure 3: Meshes in the cross-section of QC2, QC3, QC5 and QC11C.




Figure 4: Typical collar mesh.

objective function by changing vertex positions:

$$F = \sum_{\varepsilon_j} Q(\varepsilon_j) \quad (1)$$

where

$$Q(\varepsilon_j) = \sum_{i=1}^4 \kappa_{ij}^2 \quad (2)$$

and

$$\kappa_{ij} = \frac{L_{ij}^2 + L_{i+1,j}^2}{2A_{ij}} \quad (3)$$

This smoother is designed to improve element shape (angles and aspect ratio) while avoiding inverted elements. A relatively tight tolerance was used during the optimization to ensure the SU meshes were close to the minimum of F .

A visual inspection of the SU meshes reveals little difference between the unsmoothed and smoothed meshes, so the latter are not shown to save space. Differences are more apparent in the metric values. The metrics in the original study, MPCN, MPSM, and MPSJ, were calculated using only the quadrilateral beam meshes. That is, they are 2D metrics. To be more precise in this study, we re-label these metrics MPCN-2D, MPSM-2D, and MPSJ-2D. The MPES metric in the previous study was not a 2D metric as it used information from the hexahedral mesh. For consistency we've introduced a new 2D metric, MPES-2D, which is computed using only the quadrilateral beam mesh. An additional metric, average primal condition number, APCN-2D, was added to monitor changes in the average values of the metrics to better understand the effect of condition number smoothing on abort-time.

We give values of MPES-2D, MPCN-2D, MPSM-2D, MPSJ-2D, and APCN-2D in Tables 1, 2, 3, 4, and

5 for both the initial unsmoothed (UU) meshes and for the cross-section smoothed (SU) meshes, along with the percent increase or decrease in abort-time. Table 6 gives the Tau3P abort-time for both the unsmoothed and smoothed meshes. Thirteen of the sixteen cross-sections smoothed resulted in an increased abort-time, a success rate of 81%! The average increase in abort-time was 402% and in several cases the abort-time increased an order of magnitude via smoothing. Abort-time decreased due to smoothing on three cross-sections, one, as much as 89.2%. We thus see that smoothing is potentially of considerable use in improving Tau3P abort-times but cannot be applied routinely.

To understand better what smoothing has wrought, we present in Figure 5 plots of percent change in the five metrics MPCN, MPSM, APCN, MPES and MPSJ verses percent change in abort-time. Each line segment corresponds to a single run such as QC1. The origin represents the UU position, while the other end of each line segment is the SU position.

Looking first at the plot for APCN, one sees that this metric value always decreases. This is because the condition number objective function that we minimize serves as an upper bound to the average condition number given by APCN. Next, consider the plot for MPCN. The south-east (SE) quadrant of this plot consists of those runs for which MPCN decreased (i.e., the quality improved) while the actual abort-time increased. Runs lying in the SE-quadrant thus have abort-time predictions which are consistent with abort-times achieved. Eight of the sixteen runs lie in the SE quadrant. Similarly, the north-west (NW) quadrant consists of runs for which MPCN increased (i.e., the quality decreased) while the actual abort-time decreased. Runs in the NW-quadrant have abort-time predictions consistent with actual abort-times. Only one run (QC3) lies in the NW-quadrant. Runs in the NE and SW-quadrants have abort-time predictions which are inconsistent with the actual abort-times achieved. There were two runs in the SW-quadrant and five runs in the NE-quadrant. Presumably, runs in the SW and NE-quadrants appear inconsistent because other metrics are determining the actual abort-time achieved. Plots for MPES and MPSJ show similar results. QC10tbias stands out as highly unusual in that while MPCN decreased modestly, abort-time decreased considerably.

QC10tbias How can we explain the behavior of QC10tbias (see Figure 6 for the UU and SU meshes)? As the figure shows, the mesh consists of a near-polar coordinate system in the outer regions of the beam, combined with a triangle-primitive (three block-structured submeshes) near the beam center.

Mesh	MPES-2D		
	Type UU	Type SU	% Change
QC1	0.00090	0.00092	+2.2
QC2	0.00107	0.00105	-1.9
QC2b	0.00102	0.00103	+1.0
QC3	0.00175	0.00143	-18.3
QC5	0.00107	0.00118	+10.3
QC5b	0.00109	0.00109	+0.0
QC6	0.00070	0.00091	+30.0
QC7	0.00143	0.00101	-29.4
QC10t	0.00039	0.00037	-5.1
QC10p	0.00039	0.00047	+20.5
QC10c	0.00055	0.00067	+21.8
QC11	0.00079	0.00082	+3.8
QC11C	0.00060	0.00060	+0.0
QC12	0.00040	0.00086	+115
QC10tbias	0.00039	0.00032	-18.0
QC11bias	0.00079	0.00087	+3.8

Table 1: MPES-2D Metric for cross-section condition number smoothed meshes.

The UU and SU meshes are nearly indistinguishable in the figure. Overlaying them directly on top of one another shows that the only visible difference in the SU mesh shows smoothing has pulled in the azimuthal lines somewhat closer to the beam center. Other possibly significant changes can be detected by looking at the metrics. First, the worst angle in the mesh occurs at the three-valent node. This angle remains virtually unchanged by smoothing and cannot account for the change in abort-time. The MPES metric shows that the minimum mesh edge-length was small in the UU mesh and made smaller by smoothing, this change might partly account for the decrease in abort-time. The MPSM metrics shows the SU mesh is smoother than the UU mesh (in terms of area transitions). This should have helped increase the abort-time, but evidently, this was insufficient to overcome other, adverse changes in the mesh. The most likely explanation for the decrease in abort-time is thus the decrease in MPES, especially since this occurred near the beam center, a critical region in which the solution is developed. In terms of developing a practical procedure for increasing abort-time, it would be helpful if one could know a priori that smoothing this mesh via condition number would decrease the abort-time, but a great deal of experience would be needed to have foreseen this result.

QC11bias We also examine QC11bias to determine why smoothing had a slightly adverse effect on the abort-time. Figure 7 shows that the main change in the smoothed mesh is the improvement in MPSJ which occurs for one of the quadrilaterals in the paved portion of the mesh. MPSM decreased 6.1% but this

Mesh	MPCN-2D		
	Type UU	Type SU	% Change
QC1	1.20	1.24	+3.3
QC2	1.38	1.28	-7.9
QC2b	1.38	1.21	-12.3
QC3	2.29	2.55	+11.4
QC5	1.30	1.25	-4.6
QC5b	1.31	1.27	-3.1
QC6	1.88	1.32	-29.8
QC7	1.43	1.26	-11.9
QC10t	1.53	1.55	+1.3
QC10p	3.49	3.49	+0.0
QC10c	1.41	1.46	+3.6
QC11	1.37	1.41	+2.9
QC11C	1.72	1.73	+0.6
QC12	1.64	1.49	-9.2
QC10tbias	1.53	1.52	-0.6
QC11bias	1.81	1.79	-1.0

Table 2: MPCN-2D Metrics for cross-section condition number smoothed meshes.

does not seem sufficient to account for the decrease in abort-time. Again, it is next to impossible to determine a priori from the metrics that condition number smoothing would decrease the abort-time in this example since the UU mesh has no obvious defect. Probably what would have worked better in this case is to add a few intervals along the biased curves to reduce the aspect ratio and to improve the paved mesh quality.

In summary, the important points illustrated in this section are, first, that condition number smoothing is often successful in significantly increasing abort-time. Second, our model of the relationship between abort-time and certain quality metrics is generally but not completely effective in predicting the effect of a particular smoothing operation on abort-time. This is seen by the preponderance of runs that fall in the SE and NW-quadrants of our compass plots. It is not known if our model could be made more effective by including additional metrics or if some other effect, such as relation of the mesh to the solution, is clouding the picture.

5. LAPLACIAN SMOOTHING

To understand better whether or not the choice of condition number smoothing is critical to increasing abort-time, we smoothed the same sixteen cross sections using Cubit’s Laplacian smoother. We believe this is a fair comparison because the QC meshes have good topology on a convex domain. Thus Laplacian smoothing is not expected to invert any of the cells of the mesh. In addition, it will tend to make the mesh




Figure 5: Effect of Condition Number Smoothing on Abort-Time




Figure 6: The UU and SU meshes for QC10tbias

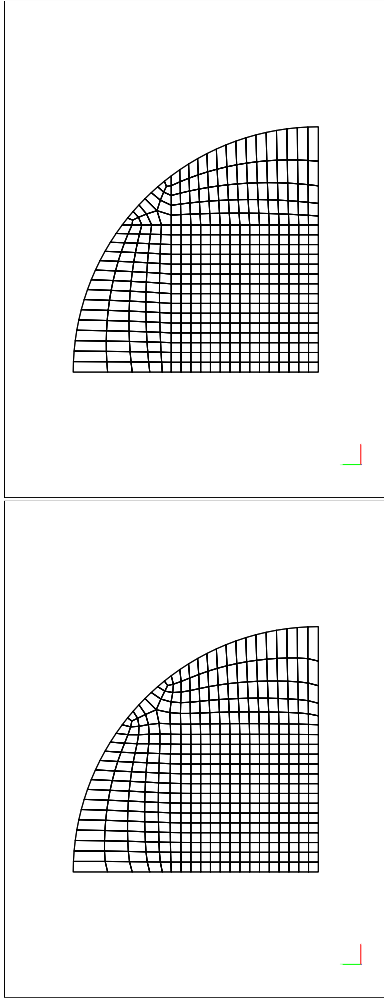


Figure 7: The UU and SU meshes for QC11bias

Mesh	MPSM-2D		
	Type UU	Type SU	% Change
QC1	5.81	6.41	+10.3
QC2	4.35	4.17	-4.1
QC2b	4.05	3.84	-5.2
QC3	3.29	8.42	+156
QC5	5.40	4.30	-20.4
QC5b	5.43	5.33	-1.8
QC6	6.07	6.42	+5.8
QC7	6.38	3.49	-45.3
QC10t	9.23	4.25	-53.9
QC10p	10.2	5.05	-50.5
QC10c	8.27	4.67	-43.5
QC11	5.79	6.19	+6.9
QC11C	8.05	10.0	+24.2
QC12	7.47	6.15	-17.7
QC10tbias	5.80	4.83	-16.6
QC11bias	6.11	6.21	-6.1

Table 3: MPSM-2D Metrics for cross-section condition number smoothed meshes.

lines smoother. The basic question is does smoothing of the mesh lines increase abort-time as effectively as condition number smoothing?

We give values of MPES-2D, MPCN-2D, MPSM-2D, MPSJ-2D, and APCN-2D in Tables 7, 8, 9, 10, and 11 for both the initial unsmoothed (UU) meshes and for the cross-section Laplacian-smoothed (SU) meshes, along with the percent increase or decrease in abort-time. Table 12 gives the Tau3P abort-time for both the unsmoothed and smoothed meshes. Thirteen of the sixteen cross-sections smoothed resulted in an increased abort-time, the same success rate as was achieved by condition number smoothing. The average increase in abort-time was 333% and in several cases the abort-time increased an order of magnitude via smoothing. Abort-time decreased due to smoothing on three cross-sections, one, as much as 61.7%. In six of sixteen cross-sections, Laplacian smoothing was more effective than condition number in increasing abort-time. We thus see that Laplacian smoothing is potentially of considerable use in improving Tau3P abort-times (at least when it does not invert the mesh). As illustrated by Figure 9, meshes which do not visually appear to have good quality can, in fact, increase the abort-time (probably in this case because more mesh nodes are concentrated in the critical area of the beam).

In Figure 8, we use the same type of plots shown in Figure 5 to examine the relationship between percentage change in MPCN, MPSM, APCN, MPES and MPSJ and percentage change in abort-time. The plot of APCN shows that most of the time the average con-

Mesh	MPSJ-2D		
	Type UU	Type SU	% Change
QC1	0.836	0.809	-3.2
QC2	0.725	0.783	+8.0
QC2b	0.727	0.828	+13.9
QC3	0.039	0.039	+0.00
QC5	0.729	0.805	+10.4
QC5b	0.770	0.805	+4.6
QC6	0.706	0.760	+7.6
QC7	0.707	0.802	+13.4
QC10t	0.758	0.660	-12.9
QC10p	0.760	0.814	+7.1
QC10c	0.710	0.694	-2.3
QC11	0.737	0.708	-3.9
QC11C	0.679	0.608	-10.5
QC12	0.609	0.673	+10.5
QC10tbias	0.758	0.764	+0.8
QC11bias	0.737	0.760	+14.5

Table 4: MPSJ-2D Metrics for cross-section condition number smoothed meshes.

dition number decreases. The two runs, QC10p and QC11bias stand out as having an increase APCN. The increase in QC11bias is a modest +0.9% and can be explained as due to the biasing along the sides of the mesh. The increase in QC10p is fairly large compared to most other runs and could be attributed to a high amount of 'pulling' towards the beam. The effect of Laplacian smoothing on MPCN is seen to be mostly random. We see again a small number of runs in the SE and NW quadrants for MPCN. Also the most runs occur in the NE quadrant for MPCN. For MPSM we see runs in all four quadrants, but with most occurring in the NE quadrant which is what we would expect.

6. CONCLUSION

In this study we smoothed sixteen pillbox cross-section meshes. Two different smoothers were used: condition number and Laplacian. Both smoothers proved generally effective in increasing the abort-time compared to the unsmoothed meshes. Of the unsmoothed meshes, QC2 had the largest abort-time (1.10e-06). Of the smoothed meshes, QC5 with Laplacian smoothing had the largest abort-time (1.36e-06). Thus smoothing was able to improve on the overall best abort-time achieved with no smoothing by 23.6%. This result suggests that, on this simple model problem, one can obtain good quality meshes even without smoothing. In such cases, smoothing will result in only a small increase in abort-time. In spite of this result we remain confident that smoothing will be an important tool in increasing abort-time in general because there are much more difficult SLAC geometries to mesh than the one selected for this study. This will be especially true




Figure 8: Effect of Laplacian Smoothing on Abort-Time

Mesh	APCN-2D		
	Type UU	Type SU	% Change
QC1	1.08	1.06	-1.8
QC2	1.02	1.02	+0.0
QC2b	1.02	1.02	+0.0
QC3	1.21	1.20	-1.0
QC5	1.02	1.02	+0.0
QC5b	1.02	1.02	+0.0
QC6	1.03	1.02	-1.0
QC7	1.05	1.04	-1.0
QC10t	1.25	1.24	-0.8
QC10p	2.05	2.00	-2.4
QC10c	1.10	1.10	+0.0
QC11	1.07	1.07	+0.0
QC11C	1.06	1.04	-1.9
QC12	1.06	1.04	-1.9
QC10tbias	1.25	1.24	-0.8
QC11bias	1.08	1.07	-0.9

Table 5: APCN-2D Metric for cross-section condition number smoothed meshes.

when applied to 3D hex meshes. On difficult geometries, the initial, unsmoothed mesh is likely to be of relatively poor quality, perhaps more along the lines of some of the less effective templates used in this study. As this study shows, when the initial quality is relatively poor, smoothing can result in considerable gains in abort-time (300-400 percent on average and an order of magnitude in some cases).

If one is going to smooth meshes to increase the abort-time, it is natural to ask what smoother should be used. As this paper shows, both condition number and Laplacian smoothing can be effective. Both smoothers increased the UU abort-time for a given topology around 80% of the time. Condition number smoothing increased the abort-time an average of 400% while Laplacian smoothing increased the abort-time an average of 333%. The abort-time achieved with condition number was better than the abort-time achieved with Laplacian smoothing in 60% of the cases. On the other hand, Laplacian smoothing gave the largest abort-time over all of the mesh topologies (QC5). In one case (QC11bias), neither condition number nor Laplacian smoothing was able to improve the UU abort-time. Often, one can use the metrics to predict which of the smoothed meshes will give the larger abort-time. For example, in QC12, the condition number-smoothed mesh clearly has better quality than the Laplacian mesh: MPCN is 1.49 vs. 2.21, MPSM is 6.15 vs. 8.15, and MPSJ is 0.673 vs. 0.453, respectively. Indeed, the abort-time for QC12 was almost six times longer for condition number than for Laplacian. In many cases, however, it is much less clear from looking at the metrics, which of

Mesh	UU Abort	SU Abort	% Change
QC1	2.30e-07	5.45e-07	+137.0
QC2	1.10e-06	1.21e-06	+10.0
QC2b	2.68e-07	9.26e-07	+245.5
QC3	2.62e-10	2.39e-10	-8.7
QC5	2.15e-07	9.57e-07	+345.1
QC5b	2.02e-07	8.30e-07	+310.9
QC6	2.96e-08	3.01e-07	+916.9
QC7	2.53e-08	3.16e-07	+1149
QC10t	2.00e-08	1.10e-07	+450.0
QC10p	1.48e-08	1.20e-07	+710.0
QC10c	2.71e-08	3.10e-07	+1044
QC11	8.93e-08	2.55e-07	+185.6
QC11C	2.84e-8	2.86e-07	+907.0
QC12	2.12e-07	4.90e-07	+131.1
QC10tbias	1.52e-07	1.64e-08	-89.2
QC11bias	8.81e-08	7.63e-08	-13.4

Table 6: PPU Abort-times with percentage change for cross-section condition number-smoothed meshes.

the two smoothing methods will give the larger abort-time. For example, the abort-time for QC10tbias was almost twenty times longer for Laplacian smoothing than for Condition number, yet a comparison of their metrics does not reveal anything particularly in favor of Laplacian (except possibly MPSM). Unless there is something obviously wrong with a given mesh, it does not appear possible to reliably predict the abort-time in advance using just metric values. This makes achievement of our modified iteration procedure (described in section 2) appear less attainable.

An obvious future line of inquiry would be to construct a hybrid objective function which trades off various aspects of mesh smoothness, edge-length, and element shape in order to simultaneously improve as many aspects of mesh quality as possible. This idea will be explored using the MESQUITE software [6].

To make this work of practical value to the SLAC meshing effort we will need to develop a methodology for when to smooth, what to smooth, and to incorporate other mesh quality improvement techniques such as prism insertion. The methodology will be tested on more difficult meshing problems that will require investigation into 3D smoothing and mesh quality.

References

- [1] Niel Madsen. Divergence Preserving Discrete Surface Integral Methods for Maxwell's Curl Equations Using Non-orthogonal Unstructured Grids. *Journal of Computational Physics* **119** (1995) 34-45.




Figure 9: The SU mesh for QC10tbias when Laplacian smoothed

- [2] Sandia National Labs. CUBIT Mesh Generation Toolkit. *webpage* <http://cubit.sandia.gov/> (2002).
- [3] *Full Wave Analysis of Microwave Components in the Time-Domain with Unstructured Grids*, Numerical Modeling Group, Stanford Linear Accelerator, Stanford CA.
- [4] N. Folwell and P. Knupp, *Sensitivity of Tau3P Abort-Time to Mesh Quality*, SAND2003-1132P, Sandia National Laboratories, Albuquerque, 2003.
- [5] P. Knupp, *Hexahedral and Tetrahedral Mesh Shape Optimization*, Intl. J. Numer. Methods in Engr., in press, 2003.
- [6] M. Brewer, L. Diachin, P. Knupp, T. Leurent, and D. Melander, *The Mesquite Mesh Quality Improvement Toolkit*, 12th International Meshing Roundtable, 2003.

Mesh	MPES-2D		
	Type UU	Type SU	% Change
QC1	0.00090	0.00097	+7.8
QC2	0.00107	0.00116	+8.4
QC2b	0.00102	0.00109	+6.9
QC3	0.00175	0.00157	-10.3
QC5	0.00107	0.00119	+11.2
QC5b	0.00109	0.00114	+4.6
QC6	0.00070	0.00092	+31.4
QC7	0.00143	0.00091	-36.4
QC10t	0.00039	0.00029	-25.6
QC10p	0.00039	0.00036	-7.7
QC10c	0.00055	0.00065	+18.2
QC11	0.00079	0.00079	+0.0
QC11C	0.00060	0.00082	+36.7
QC12	0.00040	0.00096	+140
QC10tbias	0.00039	0.00029	-25.6
QC11bias	0.00079	0.00079	+0.0

Table 7: MPES-2D Metric for cross-section Laplacian smoothed meshes.

Mesh	MPCN-2D		
	Type UU	Type SU	% Change
QC1	1.20	1.40	+16.7
QC2	1.38	1.38	+0.0
QC2b	1.38	1.37	-0.7
QC3	2.29	2.55	+11.4
QC5	1.30	1.30	+0.0
QC5b	1.31	1.41	+7.6
QC6	1.88	1.21	-35.6
QC7	1.43	1.58	+10.5
QC10t	1.53	1.52	-0.6
QC10p	3.49	3.62	+3.7
QC10c	1.41	1.29	-8.5
QC11	1.37	1.90	+38.7
QC11C	1.72	3.23	+87.8
QC12	1.64	2.21	+34.8
QC10tbias	1.53	1.59	+3.9
QC11bias	1.81	2.15	+18.8

Table 8: MPCN-2D Metrics for cross-section Laplacian smoothed meshes.

Mesh	MPSM-2D		
	Type UU	Type SU	% Change
QC1	5.81	6.83	+17.6
QC2	4.35	3.98	-8.5
QC2b	4.05	3.69	-8.9
QC3	3.29	7.26	+121
QC5	5.40	4.24	-21.5
QC5b	5.43	5.34	-1.7
QC6	6.07	4.77	-21.4
QC7	6.38	5.05	-20.9
QC10t	9.23	2.80	-69.7
QC10p	10.2	4.62	-54.7
QC10c	8.27	3.53	-57.3
QC11	5.79	9.34	+61.3
QC11C	8.05	10.9	+35.4
QC12	7.47	8.15	+9.1
QC10tbias	5.80	2.78	-52.1
QC11bias	6.11	8.54	+39.8

Table 9: MPSM-2D Metrics for cross-section Laplacian smoothed meshes.

Mesh	APCN-2D		
	Type UU	Type SU	% Change
QC1	1.08	1.03	-4.6
QC2	1.02	1.02	+0.0
QC2b	1.02	1.02	+0.0
QC3	1.21	1.21	+0.0
QC5	1.02	1.02	+0.0
QC5b	1.02	1.02	+0.0
QC6	1.03	1.01	-1.9
QC7	1.05	1.03	-1.9
QC10t	1.25	1.24	-0.8
QC10p	2.05	2.20	+7.3
QC10c	1.10	1.06	-3.6
QC11	1.07	1.07	+0.0
QC11C	1.06	1.03	-2.8
QC12	1.06	1.04	-1.9
QC10tbias	1.25	1.25	+0.0
QC11bias	1.08	1.09	+0.9

Table 11: APCN-2D Metric for cross-section Laplacian smoothed meshes.

Mesh	MPSJ-2D		
	Type UU	Type SU	% Change
QC1	0.836	0.716	-14.4
QC2	0.725	0.723	-2.8
QC2b	0.727	0.731	+5.5
QC3	0.039	0.039	+0.0
QC5	0.729	0.772	+5.9
QC5b	0.770	0.712	-7.5
QC6	0.706	0.837	+18.6
QC7	0.707	0.670	-5.2
QC10t	0.758	0.660	-12.9
QC10p	0.760	0.591	-22.2
QC10c	0.710	0.804	+13.2
QC11	0.737	0.526	-28.6
QC11C	0.679	0.310	-54.3
QC12	0.609	0.453	-25.6
QC10tbias	0.758	0.631	-16.8
QC11bias	0.737	0.510	-30.8

Table 10: MPSJ-2D Metrics for cross-section Laplacian smoothed meshes.

Mesh	UU Abort	SU Abort	% Change
QC1	2.30e-07	6.60e-07	+187
QC2	1.10e-06	4.21e-07	-61.7
QC2b	2.68e-07	4.26e-07	+59.0
QC3	2.62e-10	2.62e-10	+0.0
QC5	2.15e-07	1.36e-06	+533
QC5b	2.02e-07	5.58e-07	+176
QC6	2.96e-08	2.14e-07	+623
QC7	2.53e-08	1.43e-07	+465
QC10t	2.00e-08	2.03e-07	+915
QC10p	1.48e-08	6.03e-08	+307
QC10c	2.71e-08	1.69e-07	+524
QC11	8.93e-08	1.32e-07	+47.8
QC11C	2.84e-8	4.60e-07	+1520
QC12	2.12e-07	8.36e-08	-60.6
QC10tbias	1.52e-07	3.07e-07	+102
QC11bias	8.81e-08	8.26e-08	-6.2

Table 12: PPU Abort-times with percentage change for cross-section Laplacian-smoothed meshes.