Techniques for Modeling Torque Transfer between Concentric Cylindrical Components

Richard Tejeda
InForm Product Development, Inc.

Abstract

Finite elements that use a piecewise linear approximation of geometry are perfectly adequate for modeling cylindrical components such as shafts and hubs in many applications. However, linear elements present a challenge to the assessment of contact interfaces between curved surfaces, namely that faceted surfaces have peaks and valleys that can interlock with each other. A good example is when torque is transferred between a shaft and a hub via a key, collar, pin, or some other means. In this case, it can be difficult or impossible to control how the applied torque is shared between the interlocking mesh and the intended torque transfer device. If the goal of the analysis is to determine the strength of the actual torque transfer features (e.g., a keyway or spline), then it is critical to apply the correct load to them by eliminating or at least minimizing mesh interlocking. This paper discusses various strategies for circumventing the mesh interlocking problem.

Introduction

When finite elements that use a piecewise linear approximation of geometry are applied to a surface with curvature, the result is a faceted surface. Deviations from the true surface decrease when smaller elements are used, but the high and low spots can never be eliminated from the mesh. Faceted surfaces are often of little consequence, but they can make it difficult to accurately assess the contact interactions between a shaft and the bore of a mating component such as a gear, cam, or pulley. The problem is that the inevitable peaks and valleys on the meshed surfaces of the two components will interlock with each other and transfer loads in a non-physical manner. Even worse, the degree of interlocking can vary between different meshes, even if the nominal element size is the same.

There are a number of ways to deal with the mesh interlocking problem, but none of them work in every situation. This paper will describe some ways to approach this issue in LS-DYNA® MPP 971 R5.1.1 for the case of an 8-inch diameter steel hub mounted on a 4-inch diameter steel shaft. Despite the very specific nature of this example, the techniques described in this paper are applicable to many other instances of curved surface contact.

Journal Bearing Test Case

Before considering torque transfer between components, it is instructive to examine a frictionless line-on-line fit between the shaft and hub shown in Figure 1. In this example, the shaft is fixed in space and the force described in Figures 1 and 2 is applied to the hub. A smooth force profile was used in order to minimize artificial dynamic effects. The direction and location of the force follows the rotation of the hub. The hub will spin freely on the shaft under these circumstances, in essence acting as a journal bearing. Obviously no torque will be transferred, but this example problem provides an ideal test case for assessing the degree of mesh interlocking for various contact definitions.

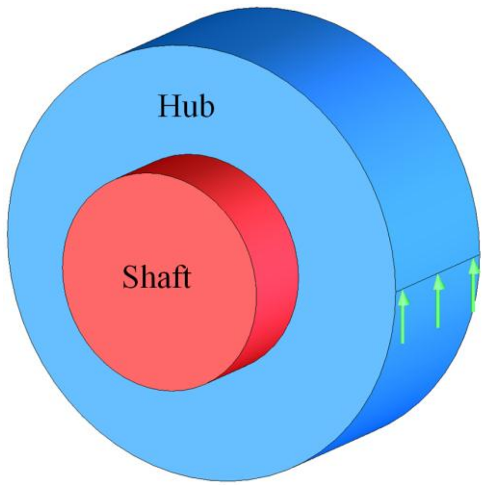

Figure 1. Example problem of an 8-inch diameter hub mounted on a 4-inch diameter shaft with a tangentially applied load.

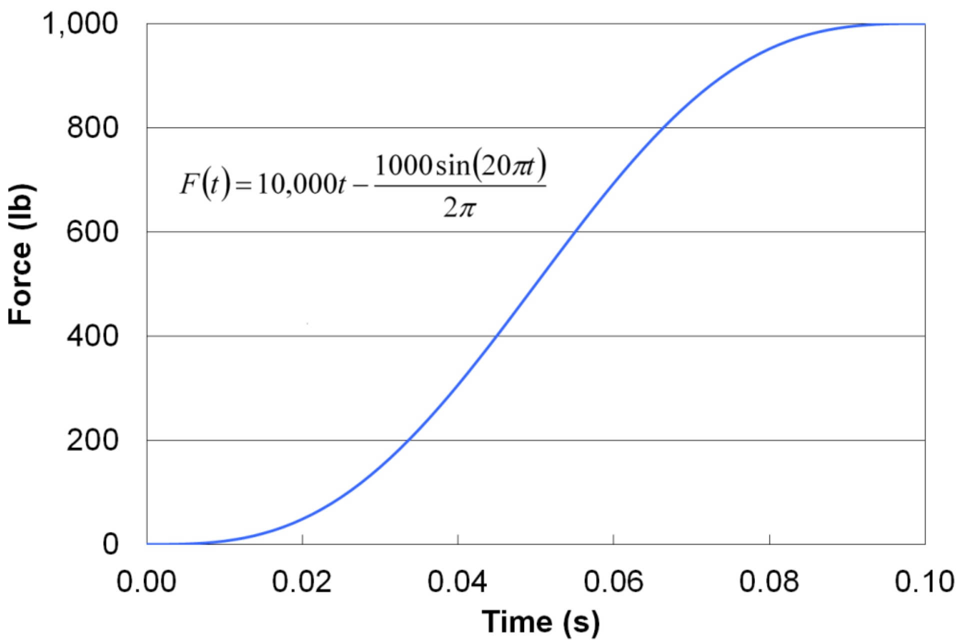

Figure 2. The magnitude of the applied load as a function of time.

Finite element solutions using the rather coarse mesh of constant-stress hexahedra shown in Figure 3 were obtained using an idealized rigid joint (*CONSTRAINED_JOINT), a geometric contact entity (*CONTACT_ENTITY), and a variety of penalty-based contacts (*CONTACT). For comparison, an analytical solution was also obtained by a straightforward application of Newton's second law. The speed of the outside diameter of the journal bearing was recorded for each solution method and presented in Figures 4, 5, and 6 for the purposes of comparison.

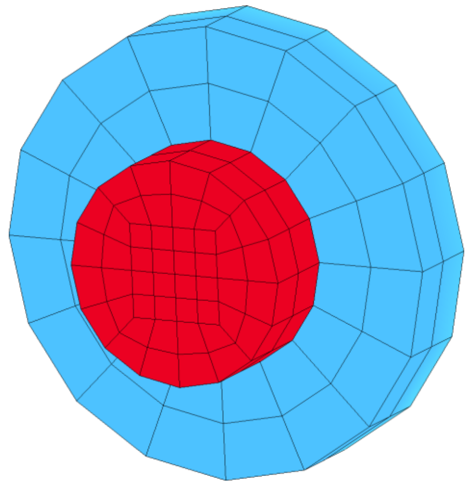

Figure 3. Mapped hexahedral mesh used in the journal bearing example problem.

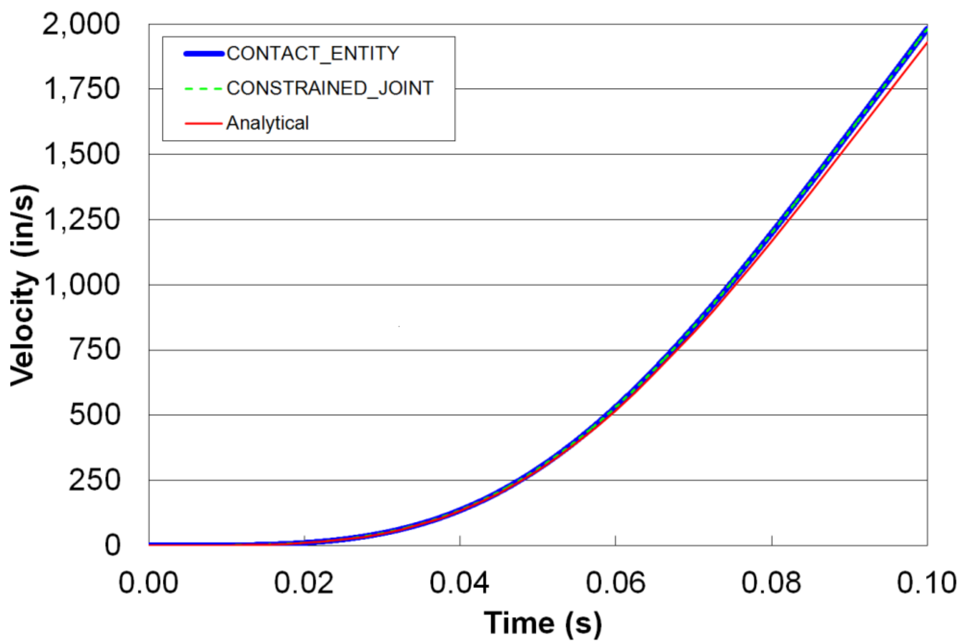

Figure 4. Velocity of the outside diameter of the journal bearing using *CONTACT_ENTITY and *CONSTRAINED_JOINT.

Not surprisingly, the results shown in Figure 4 indicate that the idealized joint and the geometric contact entity provide a nearly perfect match to the analytical solution. The slight discrepancy evident in Figure 4 is due to the fact that the relatively coarse mesh underestimates the inertia of the journal bearing by about 3%. Since both the idealized joint and the geometric contact entity involve rigid components, they are only a viable choice for situations in which the deformation of the shaft can be neglected. Note that the shaft and hub are assumed to be rigid bodies when idealized joints or contact entities are used. In all other cases, the shaft and hub are deformable bodies.

The results from the penalty-based contacts (see Figures 5 and 6) are perhaps of more interest because they are more widely applicable. As might be expected, SURFACE_TO_SURFACE contact severely restricts rotation of the journal bearing about the shaft and should be avoided whenever relative rotation of close fitting concentric cylindrical surfaces is of interest. The one way contact interfaces (ONE_WAY_SURFACE_TO_SURFACE and NODES_TO_SURFACE) provide an excellent match to the analytical results initially, but high frequency oscillations become objectionable as the journal bearing's speed increases. These oscillations may be of little consequence because in most torque transfer situations large-scale motion between the shaft and the hub is not present and would in fact represent failure of the joint. Note that the shaft is used as the master contact surface in both of the one way treatments.

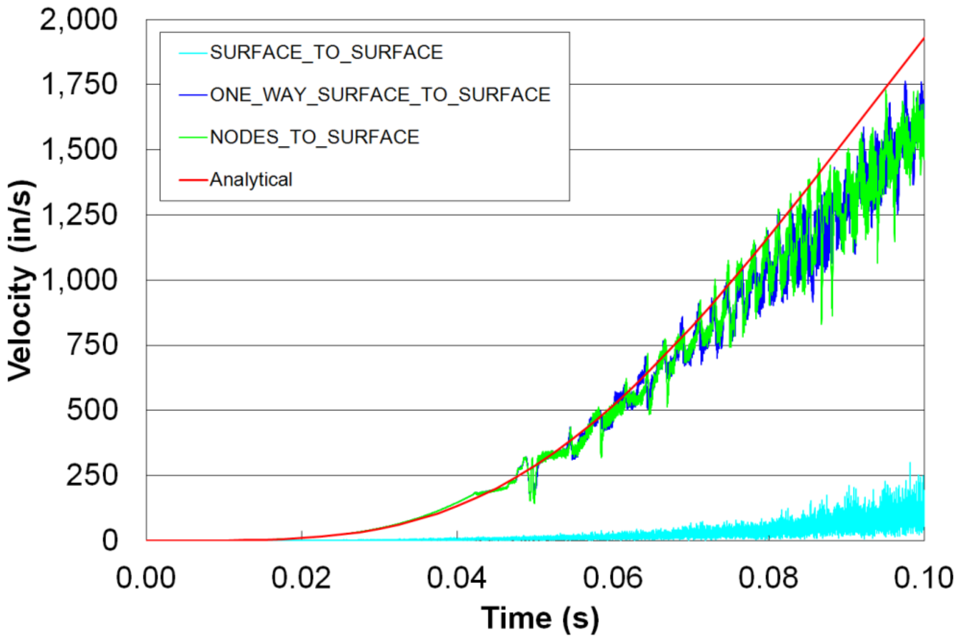

Figure 5. Velocity of the outside diameter of the journal bearing using three different penalty-based contacts.

The *CONTACT_..._SMOOTH algorithms that have been implemented in the MPP version of LS-DYNA are of particular interest in this situation. Smooth contacts seem ideally suited to the problem at hand because the master surface defined by the mesh is replaced with a smooth curve-fitted surface [1]. Indeed, Figure 6 shows that NODES_TO_SURFACE_SMOOTH and ONE_WAY_SURFACE_TO_SURFACE_SMOOTH perform admirably despite the coarse mesh that was employed. On the other hand, SURFACE_TO_SURFACE_SMOOTH essentially locks the journal bearing to the shaft and, just like its non-smooth counterpart, should be avoided in the case of contact between concentric cylinders.

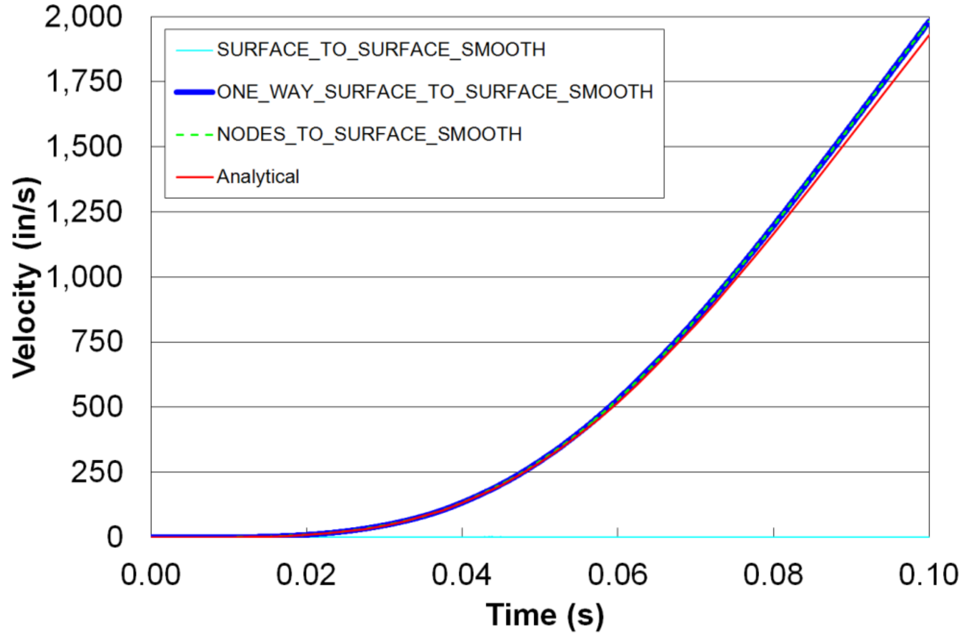

Figure 6. Velocity of the outside diameter of the journal bearing using three different penalty-based smooth contacts.

At this point, there is a temptation to completely abandon non-smooth contacts. However, if the mapped hexahedral mesh shown in Figure 3 is replaced with the automatically generated tetrahedral mesh of Figure 7, then severe interlocking occurs (see Figure 8). This troubling result appears to be due to the non-uniform distribution of nodes on the surface since there are no geometric discontinuities present on the cylindrical surfaces of this simple test case. For most applications, features such as keyways and holes are usually present and will almost certainly complicate the curve-fitted surface definition. The fact that the user can neither visualize nor control the shape of the master surface makes it difficult to recommend smooth contact as the first choice. The beauty of the LS-DYNA implementation is that once a model with a non-smooth contact is generated, changing to smooth contact requires only that _SMOOTH be added to the *CONTACT card. Therefore, it is recommended that the majority of model development take place using a non-smooth contact. Smooth contact can be easily implemented once a model is judged to be performing satisfactorily, but the user should be especially alert for anomalies in the model's contact behavior.

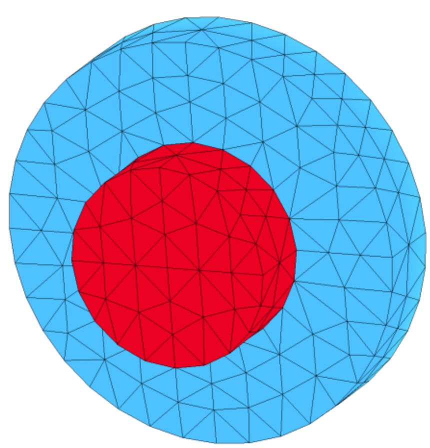

Figure 7. Tetrahedral mesh used in the journal bearing example problem.

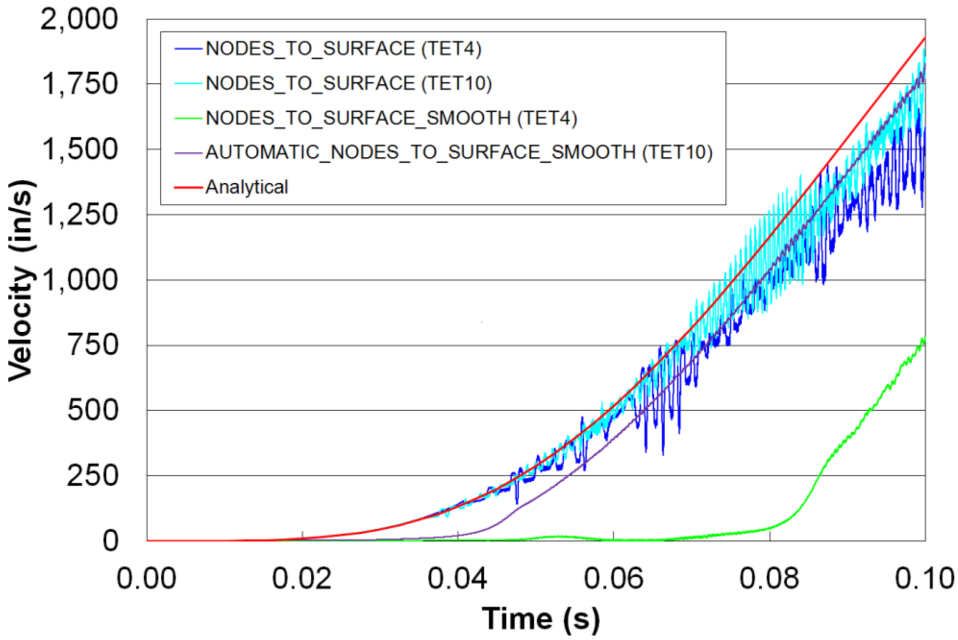

Figure 8. Velocity of the outside diameter of the journal bearing for an automatically generated tetrahedral mesh.

Although non-smooth contacts seem to be less sensitive to the node distribution on the surface than smooth contacts, non-smooth contacts have their own mesh requirements. For example, if the shaft is rotated with respect to the journal bearing as shown in Figure 9, then the meshes will interlock. This difficulty can be overcome by ensuring that there is a compatible mesh at the contact interface. In other words, master and slave nodes should be coincident in the contact region.

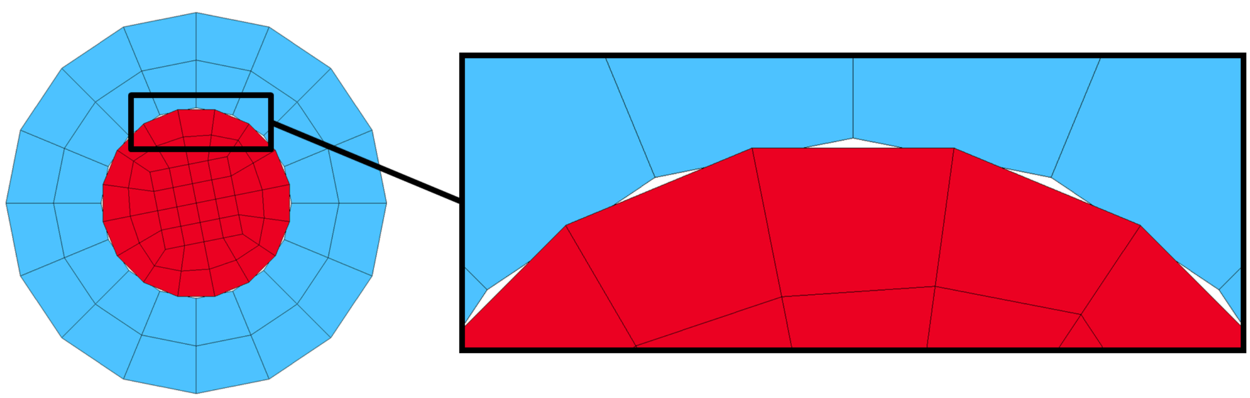

Figure 9. Incompatible mesh on the journal bearing and shaft can cause locking with NODE_TO_SURFACE contact.

The results presented thus far have used the default settings on all of the *CONTACT and *CONTROL cards with the exception of the viscous damping coefficient (VDC), which is set to 20 percent of critical. It is possible to make improvements by modifying the default values, but these modifications can also do more harm than good. For example, decreasing the penalty stiffness factor (SLSFAC on the *CONTROL_CONTACT card) can smooth out some of the objectionable oscillations evident in NODES_TO_SURFACE contact and can even allow some relative motion of the journal bearing when using SURFACE_TO_SURFACE contact (see Figure 10). However, these "improvements" come at the expense of allowing a greater amount of penetration between the contact surfaces, which is of questionable value. A similar effect can be achieved by artificially increasing the diameter of the bore or decreasing the diameter of the shaft by a small amount. Increasing the penalty stiffness tends to increase resistance to journal bearing movement so that approach is generally not recommended.

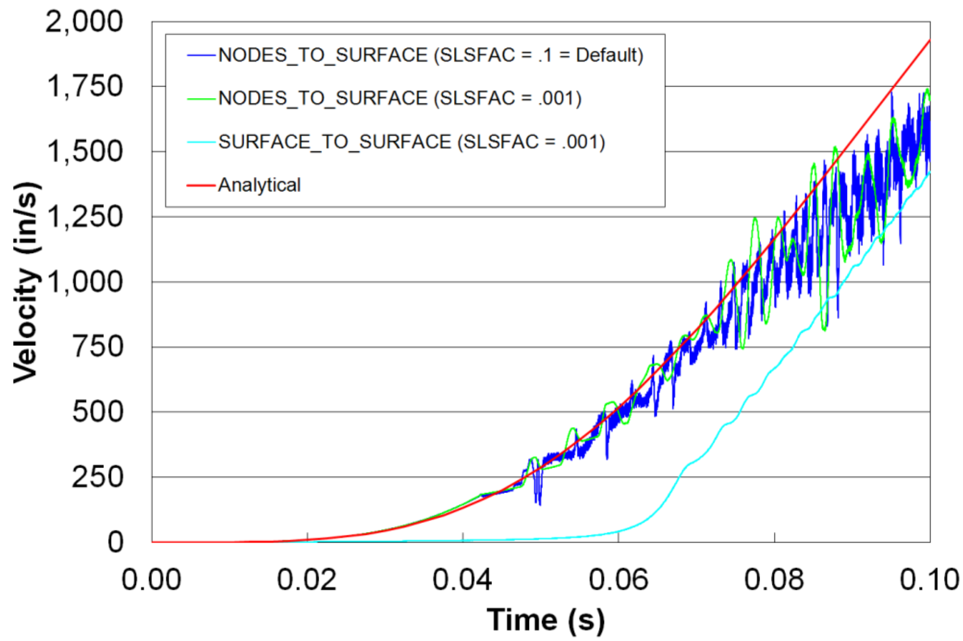

Figure 10. Velocity of the outside diameter of the journal bearing with various penalty stiffness scaling factors.

The mesh in Figure 3 was intentionally constructed with large elements in order to clearly demonstrate the mesh interlocking problem. The refined hexahedral and tetrahedral meshes shown in Figure 11 are more typical and do in fact improve the results of the simulation, but mesh interlocking problems are still clearly present in the refined mesh results (see Figure 12).

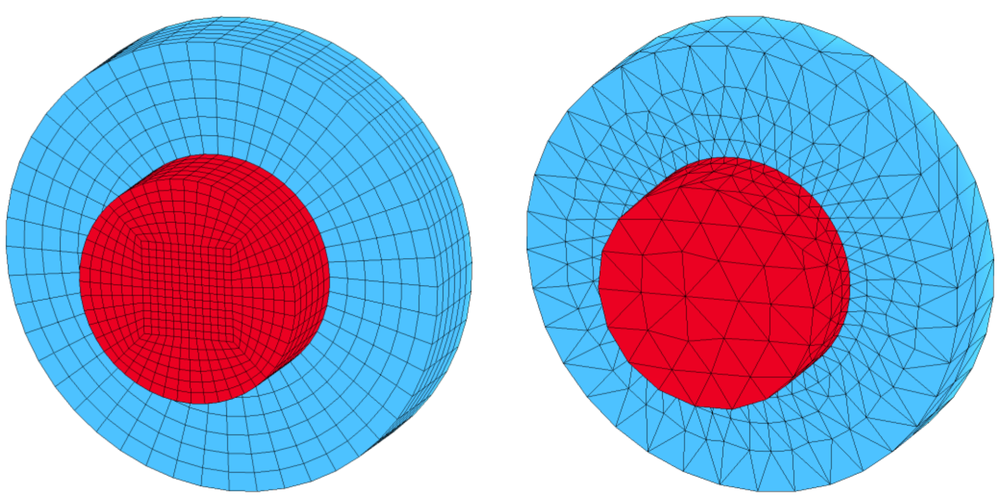

(a) (b)

Figure 11. Refinements of (a) the hexahedral mesh and (b) the tetrahedral mesh.

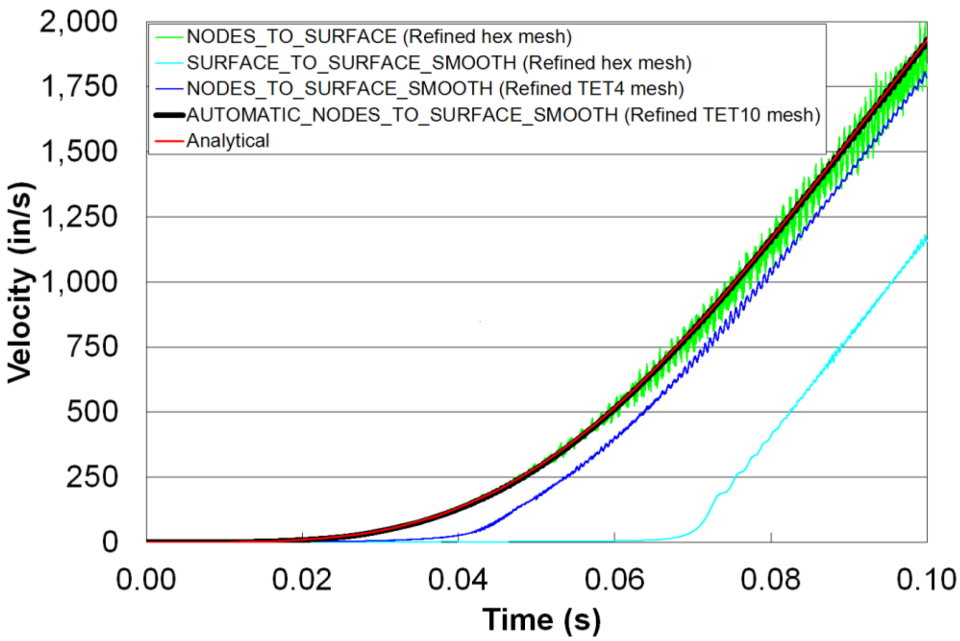

Figure 12. Velocity of the outside diameter of the journal bearing using various penalty-based contacts and refined element size.

The results for the refined TET10 mesh in Figure 12 indicate that smooth contact can function almost perfectly with enough mesh refinement. However, the results are dependent on the nature of the surface mesh and the computational cost may be prohibitive in many cases.

Interference Fit Test Case

A common way to transfer torque between a shaft and a hub is through an interference fit. The journal bearing test case that was described in the previous section can be easily modified to address torque transfer through an interference fit. The main task is to ensure that the proper frictional torque exists at the interface between the shaft and the hub.

Consider the case of a 4-inch diameter shaft and an 8-inch diameter hub with 0.002 inches of diametral interference. The interface pressure is given by the following formula [2]:

$$p = \frac{\delta E}{4R_{shaft}}\left[1-\left(\frac{R_{shaft}}{R_{hub}}\right)^2\right]$$

p = interface pressure
δ = diametral interference
E = Elastic modulus
R_{shaft} = outside radius of the shaft
R_{hub} = outside radius of the hub

For steel components, $E = 30\text{x}10^6$ psi, so p = 5,625 psi. With a hub that is 1.5 inches wide and an assumed friction coefficient of 0.01, this interface pressure will generate a frictional torque of 2,120 in·lb according to [2]:

$$T = 2\pi\mu p R_{shaft}^2 L$$

T = Frictional torque at the interface

μ = Coefficient of friction
L = Length of the shaft/hub interface

Since the applied force is acting at the hub radius of 4 inches, the hub will slip when the load reaches 530 lb.

The corresponding finite element solution can be found by first initializing the stress in the shaft and hub. *CONTACT_NODES_TO_SURFACE_INTERFERENCE provides a convenient method for developing the preload. The shaft and hub are meshed with the initial interference included and the contact algorithm finds the equilibrium position of the interface during a stress initialization solution phase [1]. In this paper, explicit dynamic relaxation is used, but presumably any other means (implicit dynamic relaxation, transient explicit solution with mass damping, etc) would suffice. Some experimentation is required to determine an appropriate penalty stiffness scale factor as well as suitable convergence criteria for dynamic relaxation. This task is greatly simplified by the fact that the analytical solution is readily available in this case. Figure 13 shows the stress distribution at the end of the dynamic relaxation phase.

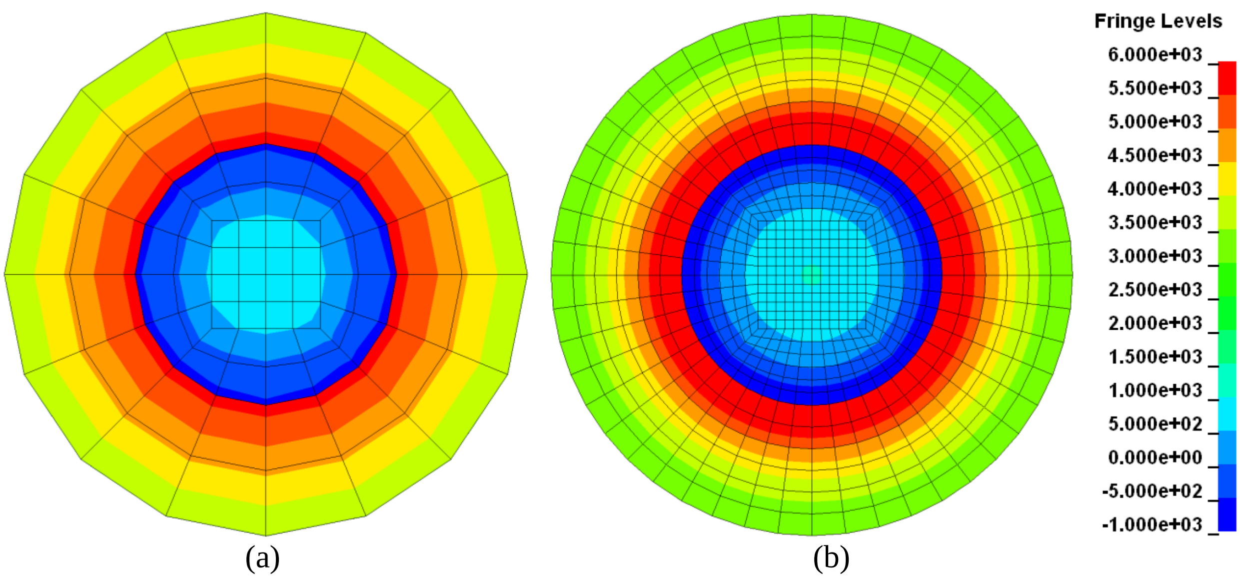

Figure 13. Contours of first principal stress in (a) the coarse mesh and (b) the refined mesh at the end of the dynamic relaxation phase using NODES_TO_SURFACE_INTERFERENCE.

Unfortunately, the smooth embodiments of the interference contacts have not yet been implemented in LS-DYNA, so a different method of stress initialization has to be employed. The classic approach to this problem is to induce the preload via thermal expansion of the shaft and/or thermal contraction of the hub. NODES_TO_SURFACE_SMOOTH in conjunction with a thermal load produces an initial stress distribution very similar to the one generated by NODES_TO_SURFACE_INTERFERENCE (see Figure 14). Results from the torque application phase of the solution are presented in Figure 15 for both smooth and non-smooth contact.

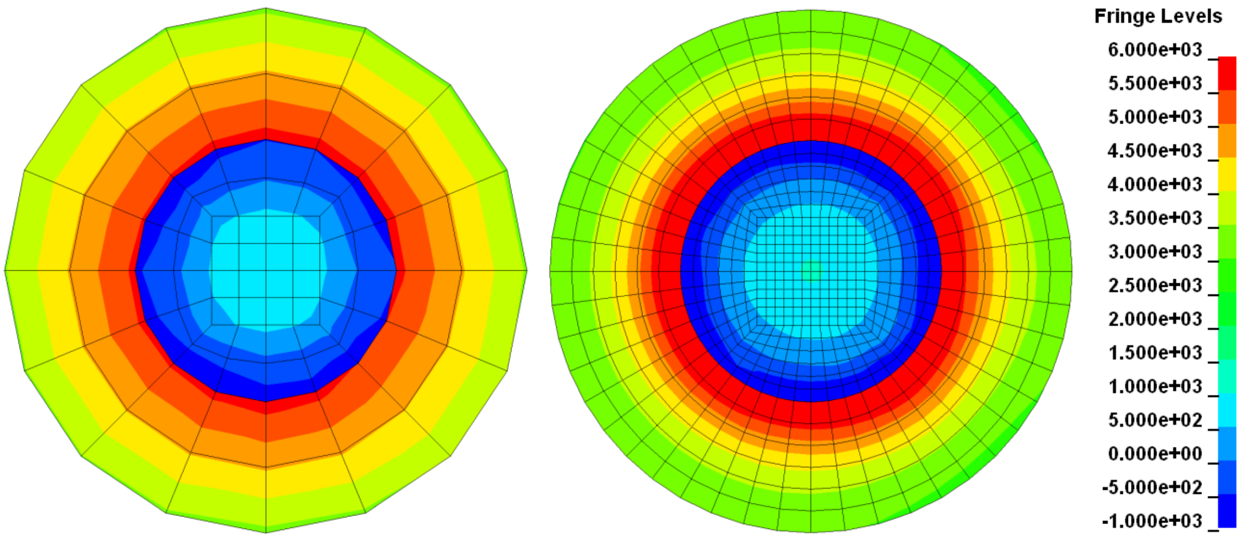

Figure 14. Contours of first principal stress in (a) the coarse mesh and (b) the refined mesh at the end of the dynamic relaxation phase using NODES_TO_SURFACE_SMOOTH in conjunction with a thermal load.

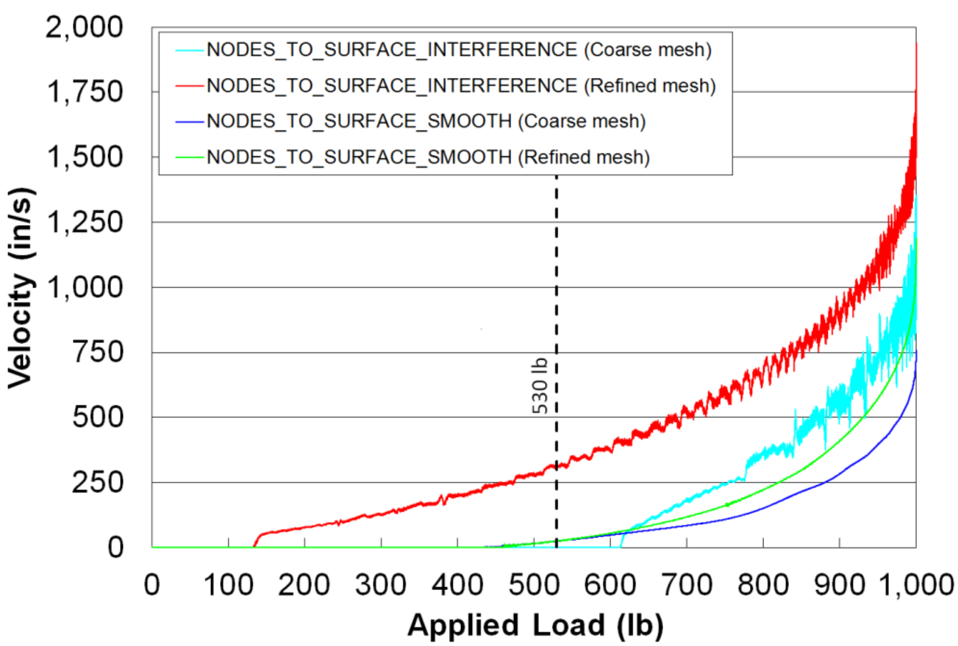

Figure 15. Velocity of the outside diameter of the press fit hub for various penalty-based contacts and element sizes.

Despite the fact that NODES_TO_SURFACE_INTERFERENCE does an excellent job of generating the correct interface pressure with a coarse mesh, the hub clings to the shaft about 15% longer than it should. On the other hand, NODES_TO_SURFACE_INTERFERENCE over-predicts the interface pressure substantially when the refined mesh is used, but the hub slips much too early. NODES_TO_SURFACE_SMOOTH cause the hub to slip about 15% too soon even though the interface pressure is under-predicted in the case of the coarse mesh and over-predicted in the case of the refined mesh. The interference fit test case was also attempted with the coarse tetrahedral meshes shown in the previous section, but these results are so obviously erroneous that they are not included here.

Keyway Test Case

Assembling and disassembling an interference fit can be so difficult (especially for large diameter shafts) that a key is often used to transmit torque between the shaft and hub. Consider the coarse and refined meshes of the keyed shaft shown in Figure 16. The shaft radius is half of the hub radius so the reaction force between the key and hub is expected to be twice as large as the force applied to the hub. For example, the hub/key reaction force ought to be 2,000 lb when the force applied to the hub is 1,000 lb (see Figures 1 and 2 for a complete definition of the applied load). As shown in Figure 17, the reaction force predicted by the finite element model varies rather dramatically depending on the contact definition and the degree of mesh refinement that is used.

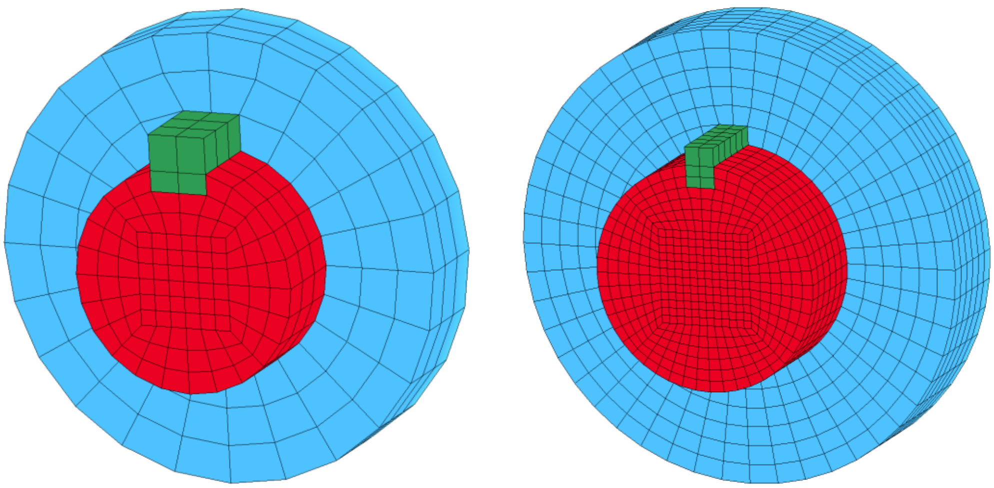

(a) (b)

Figure 16. Finite element models with (a) a coarse mesh and (b) a refined mesh for the keyway test case.

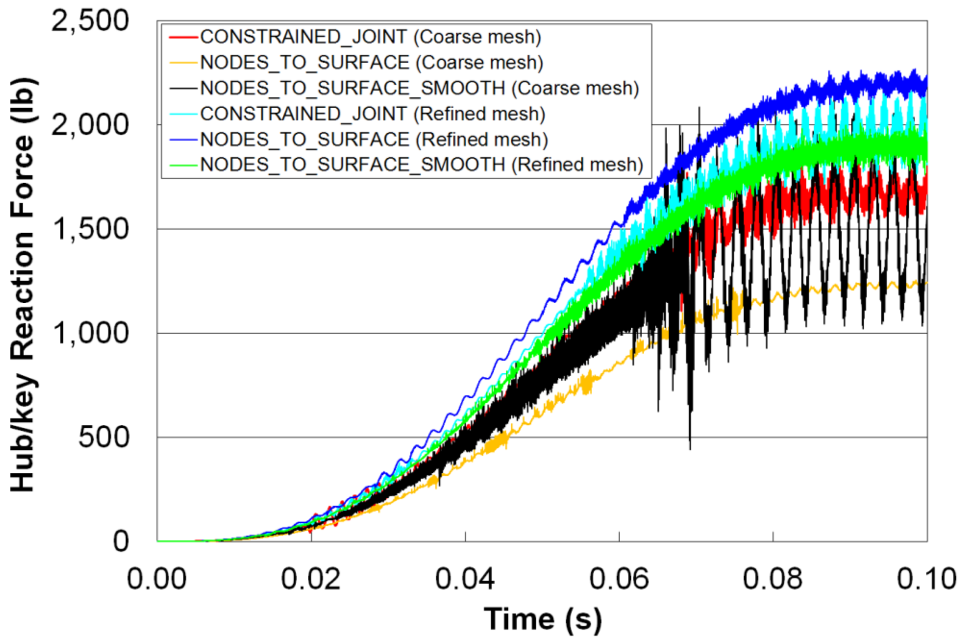

Figure 17. Reaction force on the key for various contact options.

Summary and Conclusions

The example problems presented in this paper demonstrate how mesh interlocking can make it difficult to correctly transfer torque between concentric cylindrical components in a finite element model. Fortunately, LS-DYNA provides a variety of methods for modeling contact interactions. The method that the user selects will depend on the specific requirements of the particular application, but the following guidelines and suggestions may be useful for problems that involve concentric cylinders:

- Idealized joints and geometric contact entities are not susceptible to mesh interlocking, but require the components to be rigid. Even if rigid components are not acceptable for a particular application, idealized joints and geometric contact entities may be useful during model development for determining the correct load path in an assembly.
- NODES_TO_SURFACE and ONE_WAY_SURFACE_TO_SURFACE are usually quite robust, especially when a relatively coarse mesh is employed. Either one is a good first choice when modeling concentric cylinders. SURFACE_TO_SURFACE contacts should be avoided in this application.
- Smooth contacts perform remarkably well in certain situations, but they can also be more sensitive to small variations in the mesh than their non-smooth counterparts. It is recommended that smooth contact only be implemented after a model has been fully debugged with a non-smooth contact.
- The default settings on the *CONTACT cards are usually adequate. Adjustments should only be required in situations such as contact between materials of dissimilar stiffness and/or mesh density.
- The inner cylinder should always be the master surface in the contact definition.
- The master and slave nodes should be coincident in the contact region.
- Increasing the penalty stiffness scale factor increases the likelihood of mesh interlocking.
- It is very difficult to control the point at which an interference fit will slip. For a simple interference fit, hand calculations are probably more reliable since the finite element solution may not be conservative. For more complicated situations such as load-sharing between an interference fit and a key, it may be wise to take the conservative approach of neglecting the torque carried by the interference fit.
- As with any finite element solution, a mesh refinement study is recommended to determine the appropriate element size for a particular problem.

This study also led to the “development wish-list” shown below.

- Implementation of *CONTACT…SMOOTH_INTERFERENCE.
- A way to visualize smooth contact surfaces in LS-PrePost (similar to the images in reference [3]).
- A way to modify the curve fitting parameters used by the smooth contacts.

These developments would make smooth contact more user friendly, but would not necessarily eliminate the mesh interlocking problem. In order achieve that lofty goal a fundamental shift away from traditional polynomial-based finite element analysis as proposed by Hughes et al. [4] is probably required. Implementing NURBS-based isogeometric analysis is, perhaps, not justified by the mesh interlocking problem alone. Nevertheless, the contact treatments described by Temizer et al. [5] seem promising and may be the ideal approach to curved surface contact.

References

[1] LS-DYNA Keyword User's Manual version 971, Livermore Software Technology Corporation, May 2007.

[2] M. Lindeburg, Mechanical Engineering Reference Manual, Eleventh Edition, Professional Publications, Inc.

[3] X. Zhu "Smooth Contact" *FEA Information*, pp. 36–47, February 2008.

[4] T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs "Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement," *Comput. Methods Appl. Mech. Engrg.*, 194:4135–4195.

[5] I. Temizer, P. Wriggers, and T.J.R. Hughes "Contact Treatment in Isogeometric Analysis with NURBS", ICES REPORT 10-30, The Institute for Computational Engineering and Sciences, The University of Texas at Austin, July 2010.