Getting Bodies into Contact - the Despair and Joy

Richard J. Tyrrell

Advanced Analysis Ltd

Abstract: This paper presents some tips learnt within the author's company for obtaining converged solutions in Abaqus contact simulations. Topics covered range from basic rules to follow in the setting up of problem definitions, through to more devious methods to obtain convergence in complex scenarios.

Keywords: Contact, Convergence, Hints, Tips, Standard, Explicit

1. Introduction

Many analysis problems involve the solution of the behaviour of parts being brought into (and out of) contact. The Abaqus software is well known for its ability to solve such problems.

However, many analysis practitioners are familiar with difficulties in obtaining converged solutions to such problems. The author's experience is that the resolution of such problems often involves many hours of frustration, with a disproportionate and unpredictable amount of time required to achieve satisfactory results.

This paper thus presents a summary of various methods in use by the author's company to resolve contact problems. It should be regarded as a start point for resolution of convergence problems, rather than as a substitute for experimentation and / or diligent study of Abaqus documentation.

Many of the methods presented are known to experienced Abaqus users: indeed, many of the approaches are discussed in the bowels of the Abaqus documentation. Some of the methods presented, however, represent innovative tricks which have proved useful, (and which perhaps should remain undocumented). All of the methods described here are with reference to the specific definition of contact pairs, as these represent the methods used for most work by the author. (The two alternate approaches are general contact, which tends to result in excessively large problems, and use of contact elements, which show no benefits over the use of contact pairs for most analysis work using solid models).

The author assumes no responsibility for the success or otherwise of the methodologies presented here: some work much of the time, some work in occasional instances, and some only appear useful during full moons when there is an "r" in the month. However, non convergence of analyses usually leads to desperation in methods to be attempted: here are some suggestions for alternate methods.

2. What's Happening?

The first step in the resolution of any problem of non-convergence is to identify what is happening, i.e. the reasons for failure to converge. Before getting involved in any details, it is essential to inspect the output from the "pre" checking of the model: all too often something is buried in there which provides a pointer to relatively obvious mistakes. It is useful to have the setting "split_dat = ON" in the environment file to get a separate output for the pre-processing messages.

The message ("**`xxx.msg`**") file is key to the identification of what is happening. The main things that should be looked for in this file are:

- Are magnitudes of Penetration Error and Force Error coming down, iteration by iteration?
- How do they compare with the Time Avg. Force and Largest Increment of Displacement?
- Are particular nodes or contact pairs consistently identified as the worse-case locations, i.e. the reason for none or slow convergence?
- Are there warnings of singularities or negative eigenvalues?

The action to take on finding problems with any of the above are listed in the sections below. If need be, the information may be enhanced by the use of the **`*PRINT, CONTACT`** option. However, this produces a large volume of data, which may be difficult to wade through.

3. Some Basic Rules

Before worrying about more sophisticated solution options, the basics of contact problem set up should be checked. These are well documented in the Abaqus manuals, but some of the main guidelines are worth illustrating.

3.1 Element types

With the default of NODE TO SURFACE contacts, the standard high order tetrahedral elements (C3D10) often give problems during model solution for complex problems with curved contact surfaces. The modified elements (C3D10M) have additional mid face nodes, and generally solve much better, albeit at a slightly higher computational cost. Even when a solution is found with the C3D10 elements, it may have implications for later analysis steps. Consider the simple test job shown in Figure 1. This comprises a lower cuboid block, with a rounded face upper part that has a pressure applied on its upper face. The "APPROACH" command is used (see later Section 5.1). Both elements types (C3D10 and C3D10M) achieve solutions for simple analyses such as these. The two views on the right side of this Figure show the resultant contact pressure distribution: that using the C3D10M elements (rightmost view) is a smoother and more credible result. (Note these analyses were performed using the SURFACE TO SURFACE contact definition – see Section 3.2)

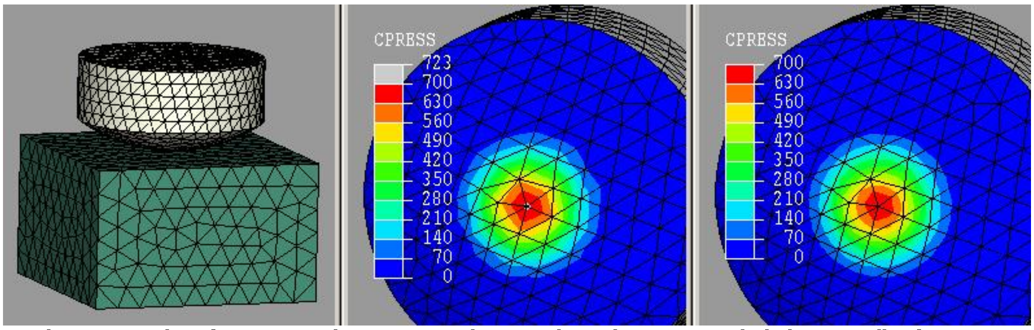

Figure 1: Simple Test Using C3D10 (Centre) and C3D10M (Right Hand) Elements

This improvement in solution may have a knock-on effect in subsequent analysis steps. In a more complex analysis of a connecting rod (Figure 2), the analysis using standard C3D10 elements fails to solve step 2, where there is some further movement of the contact between the big end shell and the rod forging. The solution using C3D10M elements is marginally more expensive both in terms of number of equations, and number of increments required … but a solution is obtained for the second step. The run histories (as evidenced by the .sta files) appear in Tables 1a and 1b.

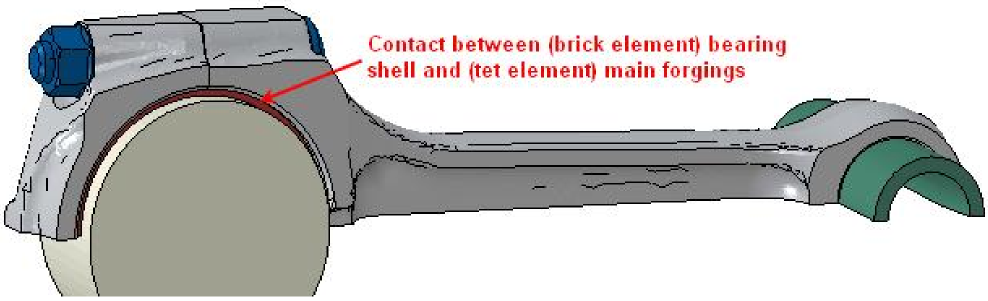

Figure 2: Connecting Rod Model

Table 1a: Connecting rod .sta file, C3D10 elements

STEP	INC	ATT	SEVERE DISCON ITERS	EQUIL ITERS	TOTAL ITERS	TOTAL TIME/ FREQ	STEP TIME/LPF	INC OF TIME/LPF
1	1	1	4	1	5	0.340	0.340	0.3400
1	2	1	0	3	3	0.680	0.680	0.3400
1	3	1	0	3	3	1.00	1.00	0.3200
2	1	1U	8	0	8	1.00	0.000	0.5100
			………	Etc	……			
2	1	5U	8	0	8	1.00	0.000	0.001992

THE ANALYSIS HAS NOT BEEN COMPLETED

Table 1b: Connecting rod .sta file, C3D10M elements

```
SUMMARY OF JOB INFORMATION:
 STEP  INC ATT SEVERE EQUIL TOTAL  TOTAL      STEP       INC OF
               DISCON ITERS ITERS  TIME/    TIME/LPF    TIME/LPF
               ITERS               FREQ
   1     1   1     7     1     8  0.340      0.340      0.3400
   1     2   1     0     3     3  0.680      0.680      0.3400
   1     3   1     0     3     3  1.00       1.00       0.3200
   2     1   1     2     1     3  1.51       0.510      0.5100
   2     2   1     0     3     3  2.00       1.00       0.4900
THE ANALYSIS HAS COMPLETED SUCCESSFULLY
```

For large models, the C3D10M elements have the disadvantage that they cause some increase in the model size, due to the use of the additional midface nodes. In the past, before the introduction of the SURFACE TO SURFACE method, a simple code was developed to modify only those elements that appear in contact definitions from C3D10 to C3D10M types as we only want these midface nodes on the elements actually involved in contact faces. The user manual (and the running Abaqus software) pronounce warnings due to the inherent incompatibility of the two elements types (the volume inside a mesh remains at C3D10 type elements): however, no difference was observed in the results of any analyses with these mixed element formulations.

It may be noted that the C3D10I elements appear to behave exactly the same as the baseline, C3D10 ones. Indeed, to date no advantage has been noted in the use of these elements, and they have the disadvantages of being more fussy in terms of element shape checks. Thus, they are not in use by the author.

3.2 Contact Types

The default surface definition (as defined on the *CONTACT PAIR instruction) is to use NODE TO SURFACE definitions. As a general rule, the SURFACE TO SURFACE algorithm is more reliable. For example, the simple test illustrated in Figure 1 failed to converge when using the NODE TO SURFACE algorithm, for either the C3D10 or the C3D10M elements.

Likewise, using the same connecting rod model as described in Section 3.1, but changing the important contact definition to a SURFACE TO SURFACE type, enables both the C3D10 and the C3D10M models to run. Interestingly, these ran slightly quicker than when using the NODE TO SURFACE definition, probably because a more stable solution was found.

In some instances (particularly where contacts exhibit “chatter”) it is beneficial to define the contact as a “symmetric” pair. In this, the slave and master surfaces are defined twice, with the second occurrence reversed, such that both sides of the contact act as the master, and vice-versa. Care must be taken with local results interpretation if this method is used.

Nearly all work performed by the author is based on geometrically-defined (element based) surfaces. Node based surfaces are only used in exceptional circumstances, as they are more prone to numerical instabilities. Because of this, and the reasons above, the SURFACE TO SURFACE method is used for all analyses by the author’s company. It is used for all the subsequent test cases described in this paper.

Whenever possible, small sliding definitions are used as they are less computationally expensive. These apply in particular to joints where there is minimal relative movement between parts, such as across bolted joints or at interference fit components. However, such joints rarely lead to the convergence problems which form the subject of this paper: hence, all analyses referred to here use finite sliding contact definitions.

3.3 Element Sizes

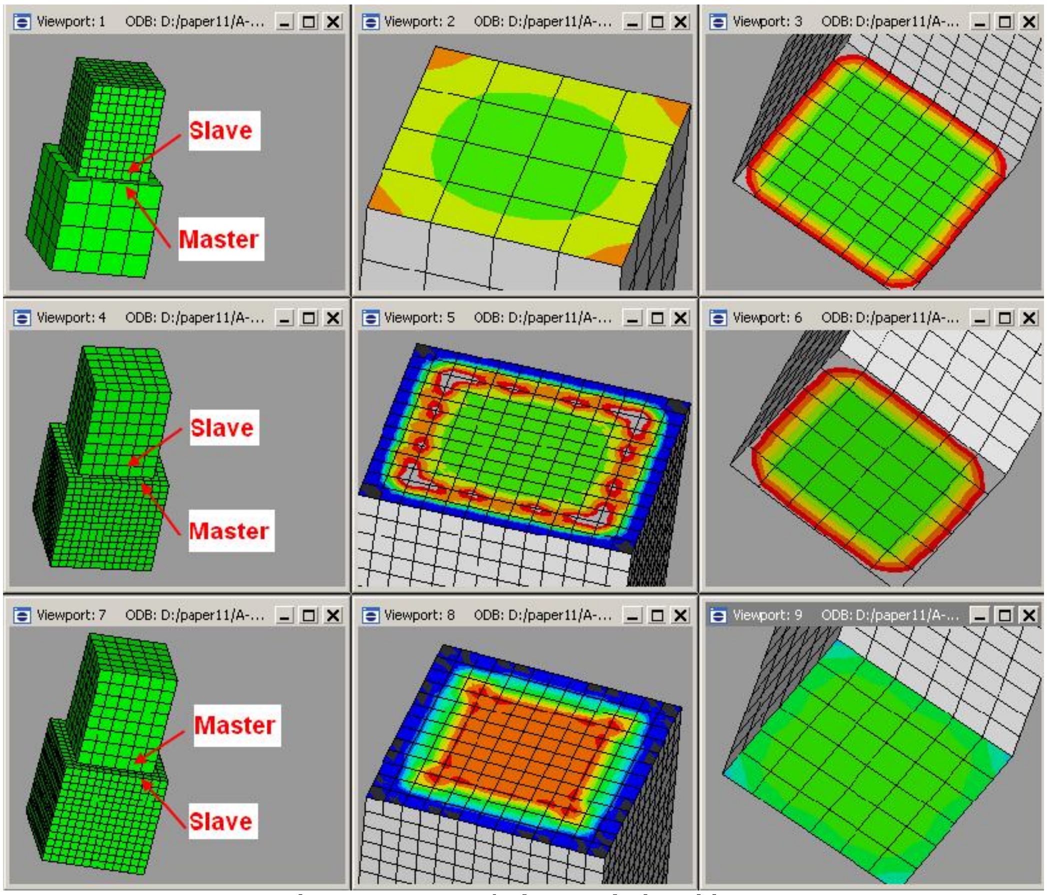

Figure 3: Master / Slave Relationships

The general rules are:

- Elements on slave surfaces should be smaller than those on master surfaces
- Whenever possible, the slave surface should be fully enveloped by the Master surface.

However, these rules may be broken in the case of relatively simple contacts. The examples shown in Figure 3 are for the same problem of an upper block being pushed down onto a larger lower block: the top row follows the above rules, the middle row has the master elements smaller than the slave ones (slave is still the upper block), and the lower row has the slave and master reversed, i.e. the slave surface extends beyond the master. All analyses solved easily. The differences in presented CPRESS results are worth noting: all are to the same scale. These examples illustrate some of the dangers of setting up problems incorrectly, and also that care must be used in interpreting results due to nodal results averaging in the generation of these views within Abaqus/Viewer.

3.4 Surface Smoothing

The solution of contacts involving curved surfaces often benefits from using the various smoothing options available: these are well documented and will not be expanded upon here.

3.5 Contact Softening

The default hard contact formulation may lead to local chattering and non-convergence of solution. A simple modification is to change the contact to an exponentially softened one. A typical start point specification for metallic components in mm-kkg-secs units appears:

```
*Surface Behavior, pressure-overclosure=EXPONENTIAL
0.01, 10.
```

The use of tabular softening specifications provides a powerful tool for resolution of some large displacement analyses: see Section 10 for an example.

4. Avoidance

One of the best ways to overcome contact convergence problems is to avoid the need for a contact altogether. Interference fit parts often fall into the category of possible avoidance. Figure 4 shows the 3rd Invariant stress results for the hot assembly load case of a main bearing shell fitted into an engine crankcase. The left side view shows the conventional analysis results, the right side view shows the same parts as analysed using a local increase in the bearing shell temperature and orthotropic material properties to give a different expansion coefficient in the axial direction than the radial one. In this case, the bearing shell and housing joint is modelled either as a *TIE or as a *CONTACT PAIR, TIED. Table 2 presents the definitions used here for the orthotropic material expansions: the axial expansion is set the same as that of the housing (including temperature dependence), so avoiding artificial stresses at the junction in the axial (Z) direction when the differential temperature is applied to the bearing shell. Note the change in the other (X,Y) direction values used for cold assembly, to those used for the hot assembly (when all these parts of the crankcase are at around 120^{o}C. The amount of additional temperature to specify for the assembly load is readily calculated, based on an assumed differential expansion rate, the required interference fit and assembly temperature. The values for differential expansion and additional temperature at the hot assembled condition are equally readily calculated by hand.

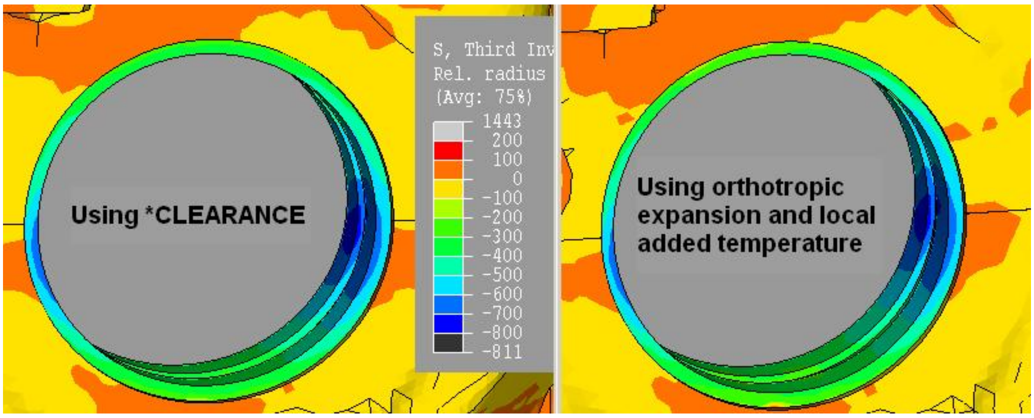

Figure 4: Use of Orthotropic Expansion Coefficients

Table 2: Example Orthotropic Expansion Coefficients

```
*MATERIAL, NAME=MMAINS
                  --- etc ---
*EXPANSION, TYPE=ORTHO
      4E-05,      4E-05,  2.03E-05,    20.
      4E-05,      4E-05,  2.03E-05,    60.
 4.718E-05, 4.718E-05,  2.03E-05,    80.
 4.718E-05, 4.718E-05,  2.21E-05,   120.
 4.718E-05, 4.718E-05,  2.21E-05,   200.
*SPECIFIC HEAT
20.,
```

Inventive use of the wide range of facilities available in Abaqus can yield very significant advantages in analysis run times – e.g. this methodology for a typical crankcase assembly, when applied to all bearing inserts, plus valve seats and valve guides, generally yields a reduction in the time for the solution of the assembly load case of the order of 70%.

5. The Use of Controls

The Abaqus software provides the user with a great deal of control over the way the software tests for convergence. It is time well spent to become familiar with the various levels of control possible. These methods are well described in the Abaqus documentation: thus, they will not be covered in detail here.

5.1 Simple Controls

In approximate order of desperation in use, the author's company tends to invoke:

```
*CONTACT CONTROLS, AUTOMATIC TOLERANCES
```

For jobs with initial gaps, particularly if gross motions are indicated in the message file:

```
*CONTACT CONTROLS, APPROACH, MASTER= xxx, SLAVE = xxx
```

Alternately, light springs may be attached to components with initial rigid body motion, and then removed using ***MODEL CHANGE, REMOVE** once the contact has settled. The **APPROACH** method tends to result in a more rapid solution, but may produce incorrect results prior to closure of any contact – e.g. see Section 10.

For jobs when a low number of nodes appear regularly in the Penetration or Force error messages:

```
*CONTACT CONTROLS, MAXCHP = xxx, PERRMX= xxx, UERRMX = xxx
```

When all else fails:

```
*CONTACT CONTROLS, STABILIZE, { MASTER=xxx, SLAVE = xxx}
```

And if it seems to fail in multiple places:

```
*STATIC, STABILIZE
```

The STABILIZE options should be used with extreme caution, and should always be followed by a repeat of the step but without the STABILIZE option. This is because the automatic stabilisation tends to introduce significant errors.

5.2 Advanced Controls

The advanced CONTROLS options tend to influence the convergence of the main model, rather than the convergence of the contacts themselves. However, one of the main reasons for failure to converge in analyses with initial rigid body motions possible is that the force levels throughout the assembly are low. As convergence is judged in terms of a percentage residual compared against the time average force, it is useful to artificially set the time average force using:

```
*CONTROLS, PARAMETERS = FIELD, FIELD = DISPLACEMENT
 , , , <value>
```

, where the **<value>** to be used may be estimated from results from a prior TIED analysis.

For many problems of contact, however, achieving convergence of the contact constraints themselves is only half the problem. In these cases it may be beneficial to adjust the ratio of convergence criteria – Abaqus documentation includes the comment that the defaults are likely to be too stringent, and experience indicates that a relaxation of most factors by an order of 10x does not significantly affect the results obtained.

Note that it is important to realize that values set using the *CONTROLS options remain in force for subsequent analysis steps. Thus, it is necessary to invoke the reset command in subsequent steps, to avoid unexpected results:

```
*CONTROLS, RESET
```

5.3 Time Incrementation

The default changes in attempted time increments tend to lead to rather confusing histories. Although modification of the defaults has only a secondary effect on the likelihood of a problem solving (or not), the author typically modifies the defaults such that increment times are simply halved or doubled on success or failure of prior increments. This change is often used in conjunction with an upper limit on the increment size (set at the *STATIC command) to enable easy visualization of the contact behaviour throughout the step:

```
*STATIC
8.0E-3, 1., 1e-06, 8.0E-03
```

And

```
*CONTROLS, PARAMETERS = TIME INCREMENTATION
** Defaults:
** 4, 8, 9, 16, 10, 4, 12, 5, 6, 3, 50, 50, 6
** 0.25, 0.5, 0.75, 0.85, 0.25, 0.25, 1.5, 0.75
** 0.8, 1.5, 1.25, 2.0, 0.95, 0.1, 1.0, 0.95
** Changes
 , ,
0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 2.0, 1.0
 , 2.0
**
```

6. Enforced Displacement

Consider the test job illustrated in Section 3.3, but this time with an initial gap between the components. If the loading comprises a force, then initial movement of the upper block is unrestrained, and will lead to infinite magnitude rigid body motion, no matter how small the applied load or initial increment size.

This problem may be overcome in many cases by the use of the APPROACH instruction. However, even with this some problems may arise.

One method commonly adopted to overcome this is to restrain the moving part with weak springs (to ground, or to adjacent parts), and then to apply light loads in a preliminary step to bring the parts into contact. This method is commonly used, but may still lead to difficult convergence, and also involves the setting up of multiple springs, which must then be removed (deactivated) in subsequent steps. Experience is needed in the specification of a suitable spring stiffness.

It is standard practice within the author's company to solve such initial rigid-body motion problems by invoking a first step with displacement loading, rather than force loading. However, the very nature of contact problems is that the absolute magnitude of displacement of relative parts to obtain a settled contact condition is unknown prior to the analysis. Thus, enforced displacement of the rear face of the upper block in this example would be likely to either result in excessive forces at the contact, or in the contact not quite being achieved at the end of that first analysis step.

Thus, the enforced displacement is generally imposed via a spring. Thus, for example, if we expect a movement of approximately 1.0mm prior to initial contact, we might specify a weak SPRING2 element attached to the part, and move the free end of that spring 1.1mm. This will prevent any rigid body motions, and give a contact which is lightly settled prior to subsequent steps. The weak spring may be deactivated, or it's free-end boundary condition simply removed, in subsequent steps.

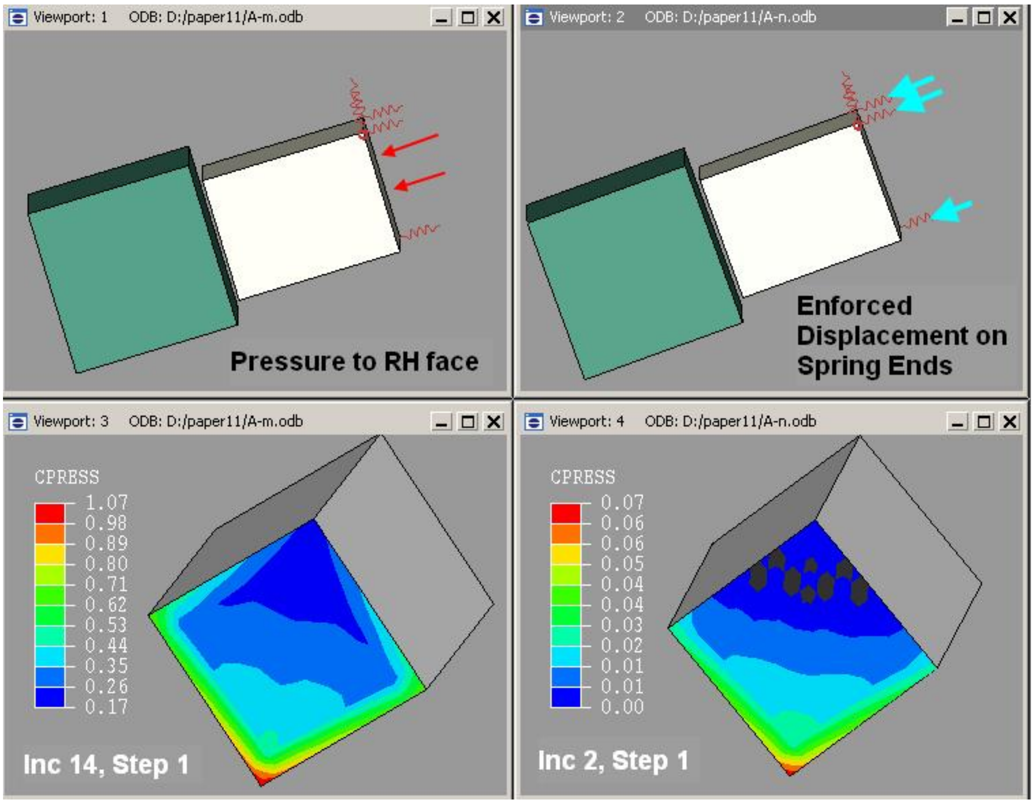

Figure 5: Use of enforced displacement for settlement

In the example shown in Figure 5, the left side views show the contact pressures obtained when using light springs to restrain rigid body motion of the upper part, with a light pressure applied to the upper face of the moving upper block. The right side views show the same problem, but with the initial contact obtained by enforced displacement of the retention spring in the closing (right-to-left in these views) direction. The results are essentially the same (contact pressures are low, with distribution caused by the retaining spring arrangement): the difference is that the analysis with springs to ground plus pressure load took 14 increments to solve the initial step, whereas the one with enforced displacement took only 2.

7. Penetration Errors – A Simple Dodge

Very often, the convergence of a contact will be shown in the message file to fail due to the penetration error being too large, "compared to displacement increment". This may occur both with problems with rigid body motions, and also with generally small movements throughout the whole model. The use of the APPROACH option as above usually overcomes these problems, but has some detrimental effects on interim results. The use of modification of the displacement convergence criterion may also introduce subtle problems.

A simple dodge is to introduce some part in the model that has artificially high movement. Typically, this may be achieved by generating a low-stiffness spring type SPRING2, with one node attached to a fully restrained point on the model, and the other end moved using a *BOUNDARY instruction. As the one end is fixed to an already-restrained point in the main model, this spring has no effect on the results obtained, but does provide a large value for "LARGEST INCREMENT OF DISPLACEMENT", so easing the criteria for penetration.

8. Force Residual – A Simple Dodge

In a similar manner to the introduction of a large displacement to ease penetration errors, the CONTROL option on time average force may be avoided by simply generating multiple small elements undergoing load, such that the average stress is the model is raised.

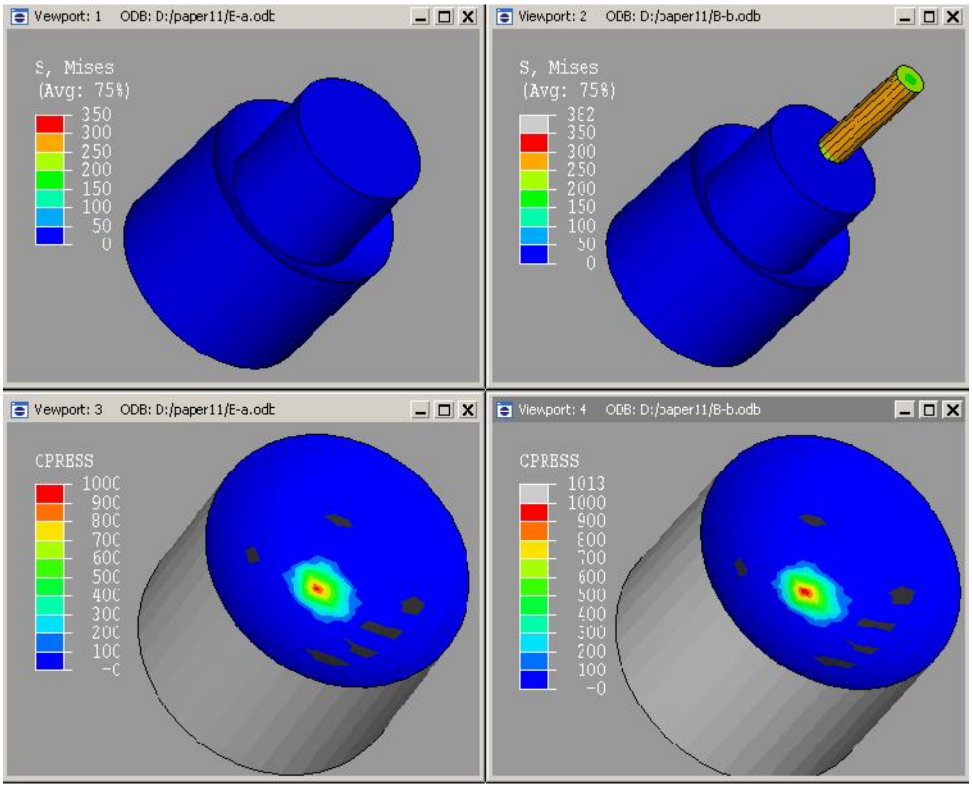

Figure 6: Use of artificial high stress

Consider the test jobs shown in Figure 6. This model starts with an initial clearance between the smaller plunger and the lower socket. For the analysis shown in the left-side views, load was applied to the rear face of the plunger via a point connected to the rear (uppermost) face using a distributing coupling. The analysis failed to completely solve: after 163 increments, it halted at a step time of 95% of complete (note that the peak contact stress is marginally lower than that for the right hand model).

The model shown in the right hand views was exactly the same in terms of the lower socket, the plunger, and the contact definitions. However, the load was applied via an additional column of small elements, which were in turn connected by a distributing coupling to the rear of the plunger. This model had high stress in the additional small elements, and solved completely in 3 increments.

9. Tracking Thickness

Inspection of the contact opening ("COPEN") is often a useful guide as to the performance of a contact during solution. It may also comprise a main output from analysis – e.g. a recent analysis of an electric motor required the air gap between stator and rotor magnets to be determined under a range of loading conditions.

By default, the opening is only monitored (and hence used in contact calculations and also available for output) when the gap between master and slave is less than the distance at which contact forces may be transmitted. Experience has shown that some problems with large relative motions between contact faces fail to converge because surface facets that should come into the range of contact during an increment are "missed". This may lead to excessive forces being transmitted at the artificial edge of the part of the surface that does lie in the monitored region. Figure 7 shows the CPRESS and COPEN values on a slave surface (from the example presented in Section 10), for an increment just before the analysis halted due to lack of convergence. Surface facets just beyond the erroneous contact pressure were just outside the thickness being tracked, and hence did not contribute to the contact solution.

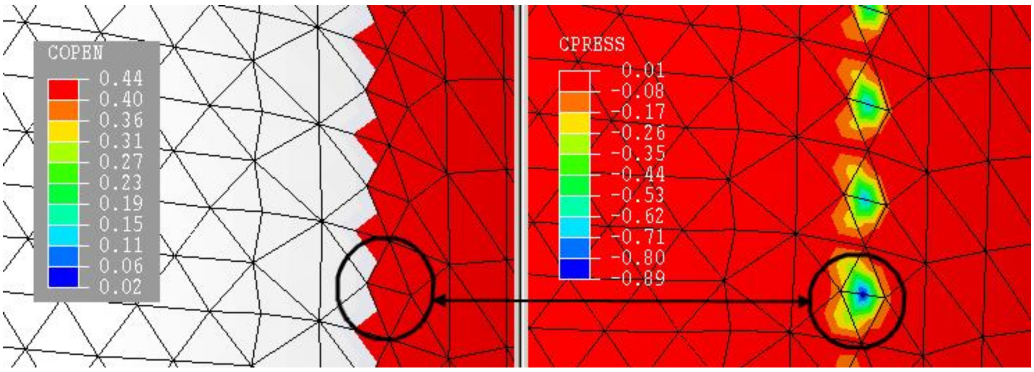

Figure 7: Facets outside of default contact tracking distance

The use of the **TRACKING THICKNESS** specification on the ***SURFACE INTERACTION** definition avoids these problems, albeit at some increase in computational cost. This option is also useful for those cases where a gap needs to be monitored (via output of the COPEN variable), even if that gap is rather large, such as for the electric motor air gap mentioned above.

10. All Together

An example is shown in Figure 8 of an ostensibly simple problem, but which required many of the above methods to obtain a converged solution. The problem comprised prediction of the seal between a 6mm diameter pipe end nipple and socket when the nut was done up, for various combinations of friction and manufacturing tolerance.

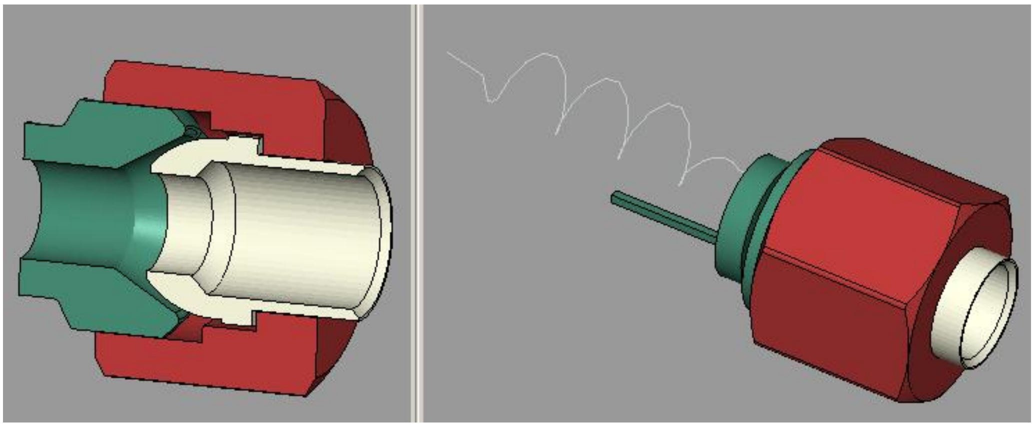

Figure 8: Analysis of pipe connector

The base of the socket (shown green) was restrained from motion.

The nut to socket threaded joint was represented by a pair of cylindrical surfaces. These surfaces were also connected via distributing couplings to points between which a connector element was specified to model the thread friction. (Attempts were also made to use the ***CLEARANCE, ……, BOLT** method, but these proved difficult to implement with clear understanding of the friction effects). The relation between rotation of the nut and axial movement of the nut was achieved using a ***EQUATION** between the relevant rotational and translatory degrees of freedom of a single node coupled to the nut flanges.

The pipe end was unconstrained at the start of the analysis. Other joints were friction-controlled contacts.

At the start of the analysis, there were axial-direction gaps of approximately 0.3mm between the rear face of the nipple and the nut, and between the nipple and socket sealing faces.

Analyses were performed in two steps: first to obtain light closure of the joints, then application of the tightening torque to the nut flanges.

Initial attempts at solution made use of the APPROACH option. However, these were observed to introduce incorrect results prior to closure of the various gaps, which had an adverse effect on the motion of the various parts. As the objective of the work was partly to identify just how frictional and tolerance effects changed the condition and position of the seal, this was not acceptable.

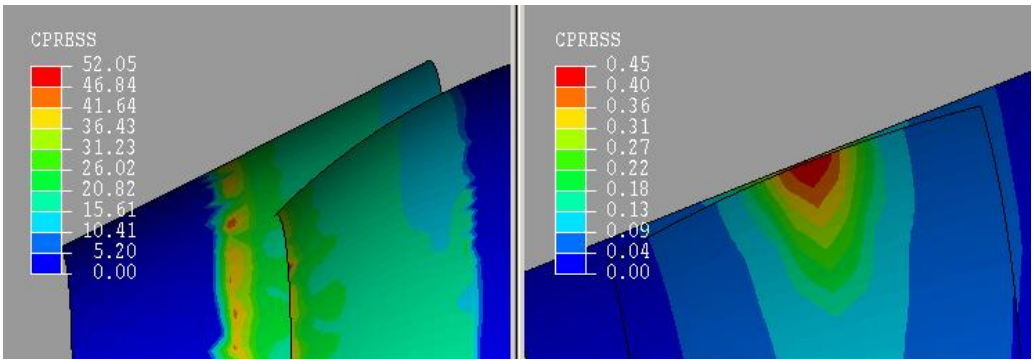

Figure 9: Contact pressures during initial settlement

Figure 9 presents the contact pressure distribution between the nipple end and the socket, just prior to “true” contact being achieved. The left view shows the results when the APPROACH option is used: it may be seen that there are high pressures and hence significant loads in the other parts, which would incorrectly affect their behaviour during the near zero load duration of the nut tightening prior to contact. The right hand view shows the same region, but for the final method used. This has negligible contact pressure values, and contact has just started to form between the (out-of round) nipple and socket.

In all, the successful analysis made use of:

- Material plasticity
- Controls on time incrementation only
- Enforced displacement of the nut via a spring during the initial, contact closing step. This as described in Section 6.
- Spring with artificial enforced displacement as Section 7 (seen in the right hand view of Figure 8)
- A cantilever attached to the (restrained) socket base, also visible in the right hand view of Figure 8, to which a lateral load was applied from the start of the step (using an amplitude definition to put all of the load on at the step start), to increase the time averaged force as described in Section 8. The relevant amplitude definition is simple, and appears:

```
*Amplitude, name=StepLoad
  0., 1.,    1., 1.
```

– A tabular definition of the pressure – closure relationship between the various parts. The form of this relationship is shown in Figure 10: note the pressure scale values! The presence of a very low pressure for large clearances mean that the contacts were active right from the start of the analysis, but with very light loads. This provided the required stability during the low-force initial closing stage of the analysis.

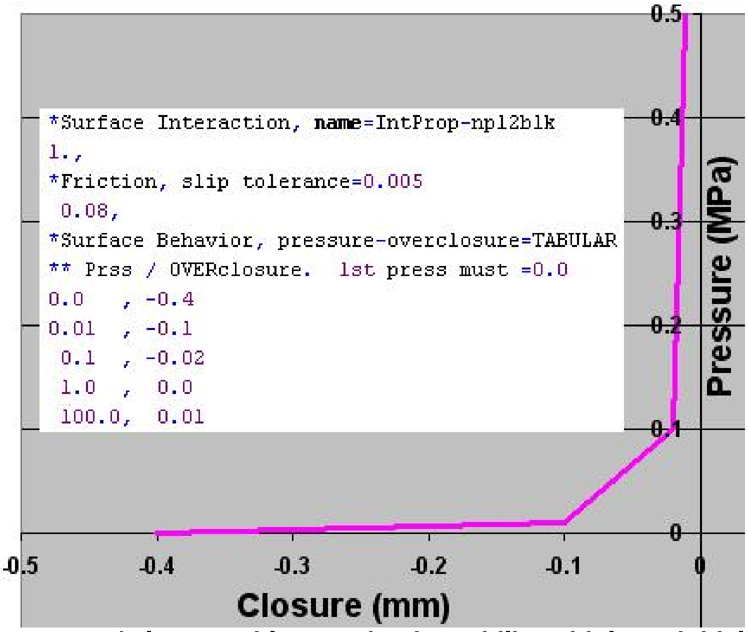

Figure 10: Pressure / closure table to maintain stability with large initial gap.

11. If All Else Fails

It should be noted that the Abaqus / Explicit code generally resolves contact problems far more easily than the Abaqus / Standard software. Much use is made by the author's company of the explicit facility for otherwise intractable problems. Details of its application and use are beyond the scope of this paper.

12. Summary

This paper presents a preliminary checklist of ideas which may help in the resolution of difficult contact – related analyses. It does not attempt to be either comprehensive, nor a substitute for revision or study of the various resolution methods described in Abaqus documentation.