Question:

How do the different Mechanica convergence methods work?

Answer:

Answer provided by Rich King, Mechanica R&D

MPA algorithm

Global RMS stress is really just an energy norm. In Szabo's book, p86 [1] it's shown the integral over the model of strain energy is equivalent to global RMS stress. So we use the difference in global strain energy from 2 p-passes as an error norm and refer to it as "global RMS stress index".

In MPA, we have local and global error norms. We end the P-loop when all global error norms selected by the user have been met (e.g. disp and rms stress, or measures).
If the global norms are not met, we use the local norms to adapt:
-loop over all edges in the model. If the displacement along the edge has changed too much since the last p-pass, the p-order of the edge must be increased.
-loop over all elements in the model. If the strain energy in the element has changed too much since the last p-pass, the p-order of the all the edges of the element must be increased.

Example:

```
        Konvergenz wird geprüft                    (18:13:08)
            Nicht konvergierte Elemente:         32
            Nicht konvergierte Kanten:            31
            Messgrößenkonvergenz:          3.4%
            Lokaler Index für Verschiebung/Energie:    100.0%
            Index für Globale RMS-Spannung:      3.4%

         Schätzungen des RMS-Spannungsfehlers:

         Lastsatz          Spannungsfehler % max. Hauptspannung
         ----------------  ------------  -----------------
         Last_C0_6WK        1.53e+01       0.4% of  3.53e+03

         Ressourcenprüfung                                    (18:13:10)
   Rechendauer     (Sek.):     458.35
   CPU-Zeit         (Sek.):     246.36
   Speicherbelegung      (KB):     580853
   ArbVerz-Plattenbelegung (KB):     124804

   Analyse konvergierte innerhalb 5% auf
   Messgrößen.
```

Here, the displacement/energy index is still showing an error of 100%, which is why 31 edges and 32 elements are not converged. But the global norm, which in this case is the measures, has converged within 5%, so the p-loop terminates.

SPA algorithm

SPA uses the results of a single p-pass, with all the edges in the model initially set at 3. The stresses are smoothed using a global technique described in [2]. The smoothed stresses (also called "superconverged") are compared with the "raw" or directly computed stresses, at various points in each element, to calculate an error. If the error exceeds a tolerance (default 8%), the size of the error is used to compute an estimate of how many functions are needed in the element. There's a discussion of how to relate number of functions needed to error in [2]. From the number of functions needed in the element, we can determine the polynomial order each edge must be increased. So for an element with really high error, the p-order of its edges could jump from 3 to a high value, as much as 9.

WF5 local control

In WF5 we've added more control to SPA:

-by picking advanced control, you can set the 8% default to a different value

-you can also pick geometric entities (points, edges, surfaces, or components), to set a local value. The local stress target will then be used instead of the global target for all elements that are in the components or touching the points, edges, or surfaces

-you can define excluded elements which are ignored in the SPA calculation, and can specify a p-order to limit them too

References

[1] Szabo, B, and Babusca, I, Finite Element Analysis, Wiley, 1991 (the classic work on the P method)

[2] Zienkiewicz, O, and Zhu, J, “A Simple Error Estimator and Adaptive Procedure for Practical Engineering Analysis”, Int. J. Numer. Methods Engr. 24 (1987), 337-357.