THE P-VERSION OF THE FINITE ELEMENT METHOD IN MSC/PROBE

John E. Schiermeier

The MacNeal-Schwendler Corporation
St. Louis Office
1600 S. Brentwood Blvd., Suite 840
St. Louis, Missouri 63144

1. INTRODUCTION

Structural analysis can generally be divided into two areas: global analysis, where general behavior, such as the load paths, of a structure is important; and local analysis, where the specific response, such as the detailed stress state, of a single component is important. The former is suited well to the h-version, which is used in MSC/NASTRAN, while the latter can take advantage of the p-version, which is incorporated in MSC/PROBE.

In this paper, the h-version and the p-version will be described. The convergence rates of the different extension processes will be discussed, and the quality control procedures for the p-version in MSC/PROBE will be presented. Then a sample problem for detailed stress analysis in MSC/PROBE will be analyzed, and the quality control procedures will be applied to verify that the solution is good.

2. P-VERSION OF THE FINITE ELEMENT METHOD

2.1 h-Version and p-Version

The finite element method provides approximations to analytical solutions in many fields, such as elasticity and heat transfer. The domain on which the solution is to be approximated is divided into subdomains, called finite elements. On these regions, the unknown quantity, such as displacement or temperature, is represented by a linear combination of simple polynomial functions, known as elemental shape functions. The quality of the approximation is determined in large part by the choice of the subdomains and the choice of the elemental basis functions. These two choices correspond to the two versions of the finite element method, the h-version and the p-version, respectively.

In the h-version of the finite element method, the number of shape functions is fixed for each element. The shape functions are usually polynomials of low degree, normally 1 (linear) or 2 (quadratic). These correspond to the linear and isoparametric elements in traditional finite element programs. The error of approximation is controlled by mesh refinement. Since the characteristic dimension of the elements is often denoted by h, this is called the h-version.

In the p-version of the finite element method, the number of shape functions for each element can be increased to include polynomials of higher degree. In this case, the error of approximation is controlled not only by mesh refinement, but also by increasing the polynomial degree of elements. Since the polynomial degree of the element is denoted by p, this is called the p-version.

The difference between the h-version and the p-version is represented in Figure 1.

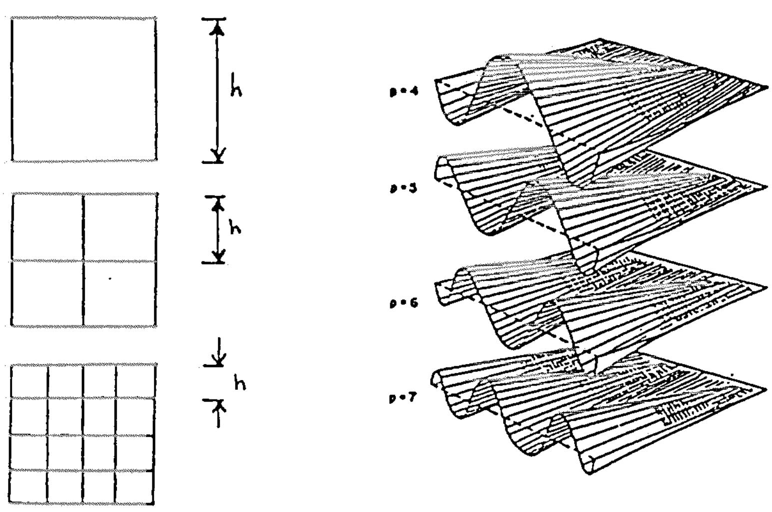

Figure 1: h-Version and p-Version

The hierarchic shape functions for the p-version quadrilateral are shown in Figure 2, where the polynomial degree is indicated on the left margin. The vertex modes for p=1 are non-zero at one vertex and decrease to zero at the other three. The side modes starting at p=2 are non-zero on one side and decrease to zero on the other three. The internal modes starting at p=4 are non-zero only on the interior of the element. The side and internal modes are not associated with physical nodes; they are functions that describe the variations between the vertices of the element.

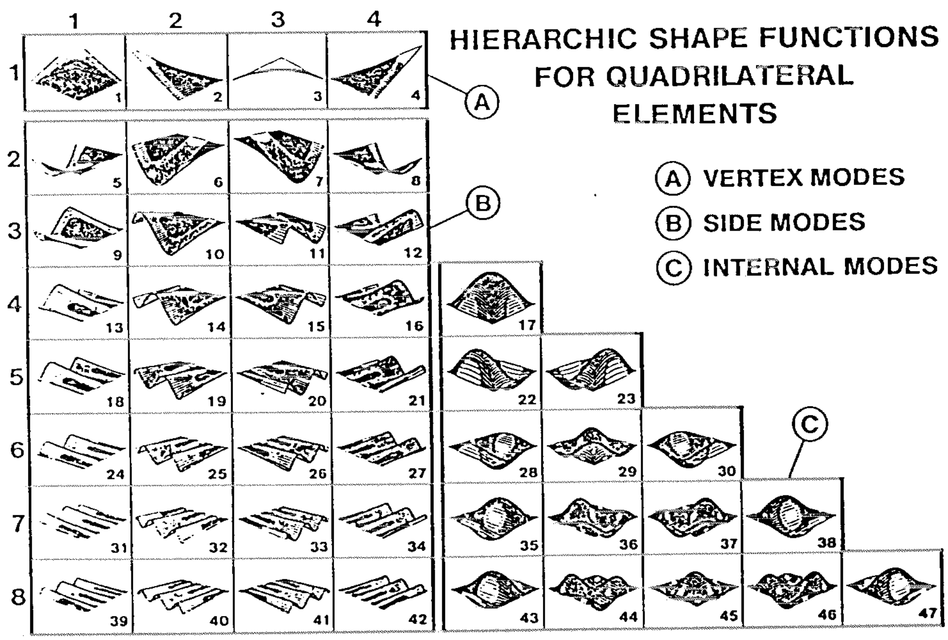

Figure 2: Hierarchic Shape Functions for the Quadrilateral

2.2 Discretization

In order to perform a finite element solution, which is an approximation to the exact solution, the problem must be discretized. This involves designing the finite element mesh and selecting the polynomial degree. That discretization must be designed so that engineering decisions based on the approximation are essentially the same as those that would be based on the exact solution, if the exact solution were known. As a result there must be reliable methods for determining how close the approximate solution is to the exact solution. This is done by performing extensions.

2.3 Extension Processes

Extensions involve increasing the number of degrees of freedom in an orderly fashion, so as to approach the exact solution as the number of degrees of freedom becomes infinite. Extension processes are important, because both the accuracy of the approximation and the control of the error are based on extensions. There are three basic types of extensions, corresponding to the h-version, p-version, and a combination of the two.

In the h-extension, the polynomial degree of the elements is fixed, while the number of elements is increased. This creates a dilemma for the user; uniform mesh refinement can be very wasteful, while localized refinement can be very cumbersome. In a production setting, the efforts of creating and managing such meshes can strongly discourage such refinement.

In the p-extension, a typically simple finite element mesh is fixed, while the polynomial degree of the elements is increased. This requires no additional work on the part of the user, and so it can be performed in a production setting.

In the hp-extension, selective refinement is performed in areas of singularities, where the stresses approach infinity, and the polynomial degree of the elements is increased simultaneously.

The global quality of the approximation is typically measured by the relative error in energy norm, which is related to the root-mean-square error in stresses, and so the performance of the extensions can be measured by the rate of convergence of the relative error with respect to the number of degrees of freedom.

It has been shown theoretically that convergence to the exact solution occurs for all extension processes, but with the differences occurring in the rates of convergence. These theoretical rates of convergence are shown in Figure 3.

The h-extension with uniform mesh refinement is the least efficient extension process, and for problems with singularities the rate is dependent on the polynomial level and strength of the singularity.

The p-extension with an ungraded mesh for problems with singularities has a convergence rate twice as great. Typically neither of the latter two extension processes would be used in a real problem.

The h-extension with optimally graded meshes, where optimal is defined as each element having the same error, has a convergence rate dependent only on the p-level. For smooth problems, the optimal mesh is a uniform mesh, but otherwise the optimal mesh is difficult to create in practice.

The p-extension with strongly-graded meshes around the singularities, which is how meshes for MSC/PROBE should be designed, consists of two parts: first is the exponential convergence rate, which approaches the theoretically optimal convergence rate; second is the algebraic rate, which is parallel to the p-extension for an ungraded mesh. More graded elements would postpone the transition, and the goal is to use enough grading to get the desired accuracy before the transition.

The hp-extension has a completely exponential convergence rate and is the most efficient. This is also the same rate as p-extension for smooth problems.

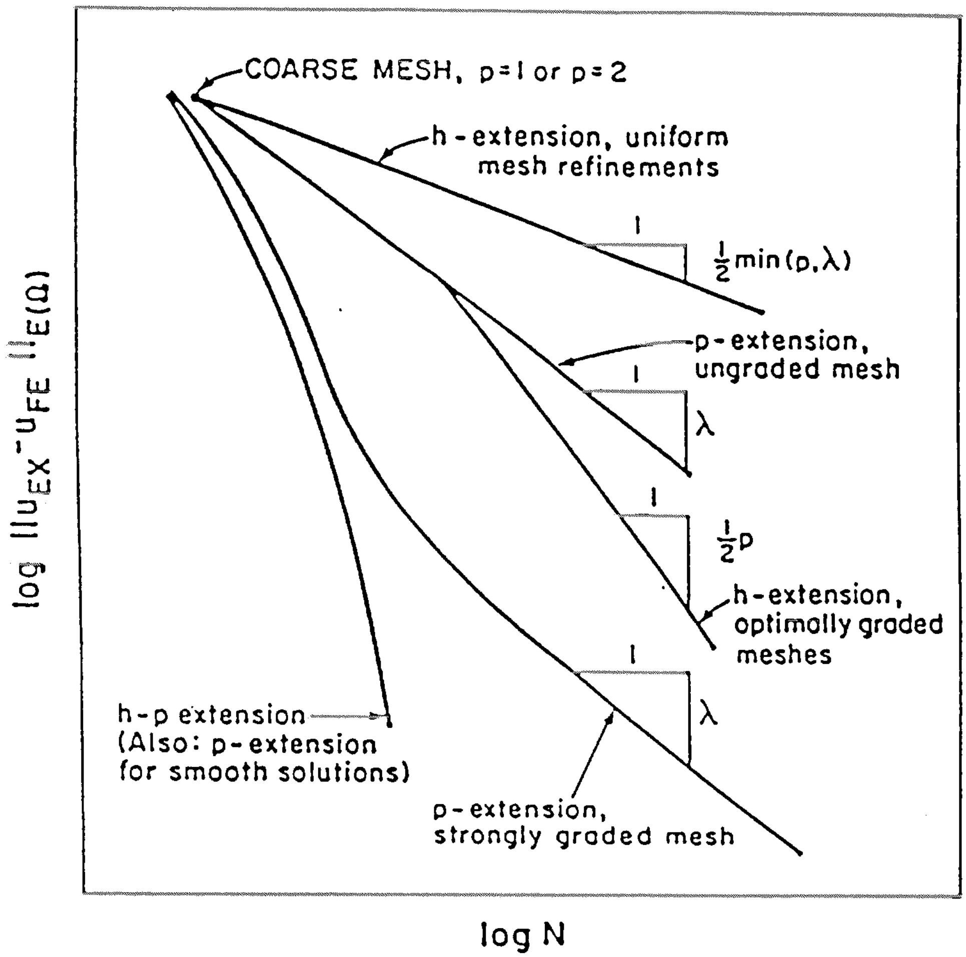

Figure 3: Theoretical Rates of Convergence

An example of the convergence rates for a problem with a singularity, the L-shaped domain, is shown in Figure 4. The singularity has strength lambda of 0.544. The h-extension with uniform mesh refinement and the p-extension with an ungraded mesh are shown, with the convergence rate of the p-extension twice that of the h-extension. The p-extension with a strongly graded mesh of two rings of elements around the singularity is shown, and it can be seen that the convergence is exponential through p=4 and then algebraic from p=6 to 8. If a 1% error in energy norm is desired, it would require approximately 10^3 degrees of freedom for the p-extension with graded mesh, 10^5 for the p-extension with ungraded mesh, and 10^7 for the h-extension with uniform refinement.

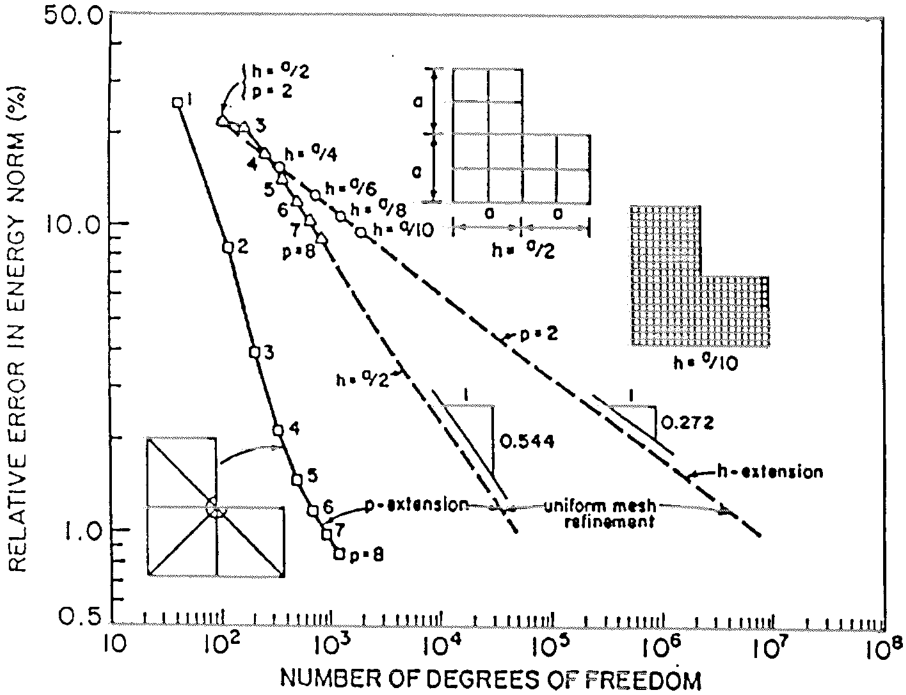

Figure 4: Rates of Convergence for L-Shaped Domain

2.4 Quality Control Procedures in the p-Version

After the p-extension has been performed, the quality and reliability of the final solution must be assessed. This involves three quality control procedures:

1. relative error in energy norm,
2. equilibrium of elements and inter-element continuity,
3. continuity and convergence of point functionals.

The first quality control procedure is the relative error in energy norm, which is based on an estimated global strain energy computed from theoretical considerations. The estimate uses a sequence of three solutions which must be hierarchic, so that the lower-order solutions are subsets of the higher-order solutions. Since the polynomial shape functions in MSC/PROBE were developed to be hierarchic, this condition is met. The global energy is the only quantity for which theoretical estimates exist, and provides a global measure of the solution quality. After the solution is verified to be good in a global sense, more localized checking is then performed.

The second quality control procedure is the equilibrium of elements and inter-element continuity. This involves explicitly integrating the stresses over the element boundaries to calculate the resultant forces. The forces for each side of an element may then be summed to get the resultant force on the element, which should be zero. Similarly, the forces between two adjacent elements may be compared to check whether Newton's action/reaction principle is satisfied. If these element-level tests, performed on elements in the area of interest, are satisfied, then even more localized checking is performed.

The third quality control procedure is the continuity and convergence of point functionals, such as stresses and strains. This is the most localized and rigorous test, since the previous tests involved integrations over larger areas that could hide local problems. This test involves comparing the continuity of point functionals on boundaries between adjacent elements in the region of interest. It also involves examining the convergence of the functionals at any particular point as the polynomial level is increasing. This can not only give a value of the function, but also an estimate of the bounds on the accuracy.

Once all three quality control procedures have passed, the finite element approximation may be accepted. Then the quantities of interest, such as the location and magnitude of the maximum stress, may be calculated from the solution and used in making engineering decisions.

3. EXAMPLE PROBLEM: SPLICING FIXTURE

3.1 Description

The example problem chosen to illustrate the usage of MSC/PROBE with the p-version is a splicing fixture used to join structural components, as shown in Figure 5 with symmetry planes indicated. There is a force from a bolt applied through a washer to the top of the hole, and the top rim of the fixture is constrained in the vertical direction.

The goal of this analysis is to determine the location and magnitude of the maximum principal stress in the model. The p-version in MSC/PROBE is used in the detailed stress analysis, whereas MSC/NASTRAN could be used for a structural analysis of the fixture with the components that it joins. The resulting loads or displacements from the structural analysis could also be used as boundary conditions for the detailed stress analysis.

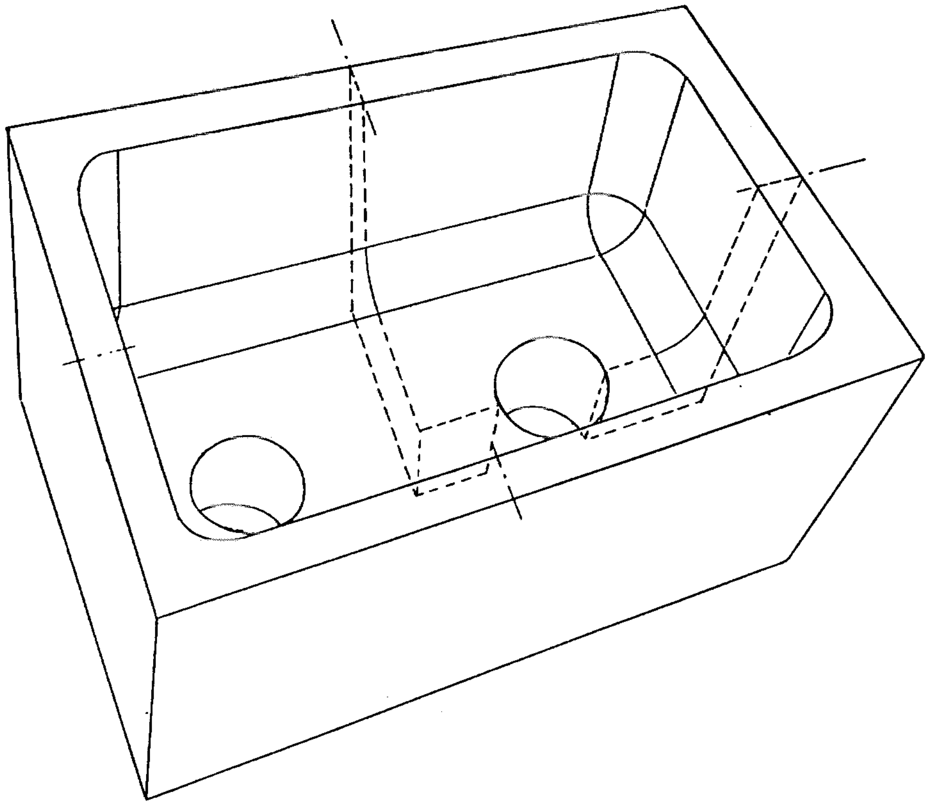

Figure 5: Splicing Fixture

3.2 Finite Element Mesh

The finite element mesh for the splicing fixture takes advantage of several of the characteristics of the p-version. Larger elements are acceptable, since each element has sufficient degrees of freedom to model the behavior on the interior of the element. The hole and fillets are modelled with exact circles instead of multiple linear or quadratic segments, so that fewer elements are necessary to accurately model the geometry. This is especially important, considering that the stress concentration is expected to be in one of these areas, and discontinuous tangents between the elements would create singularities that might affect the results. After the geometric features have been modelled, a minimum number of elements may be used to fill in the rest of the model. The p-version elements are robust enough so that large aspect ratios or large distortion generally does not degrade the performance.

In the p-version, loads and constraints are applied over finite areas, since a point load or a point constraint creates a singularity in the theory of elasticity. In addition, the loads and constraints should be modelled realistically for detailed stress analysis.

The mesh for the splicing fixture has 18 elements and is shown in Figure 6. Only one quarter of the fixture is modelled, because of symmetry, and so symmetry conditions are applied to the front and left faces. The top faces are constrained in the vertical direction, and the bolt load through the washer is applied to the tops of the three elements around the hole.

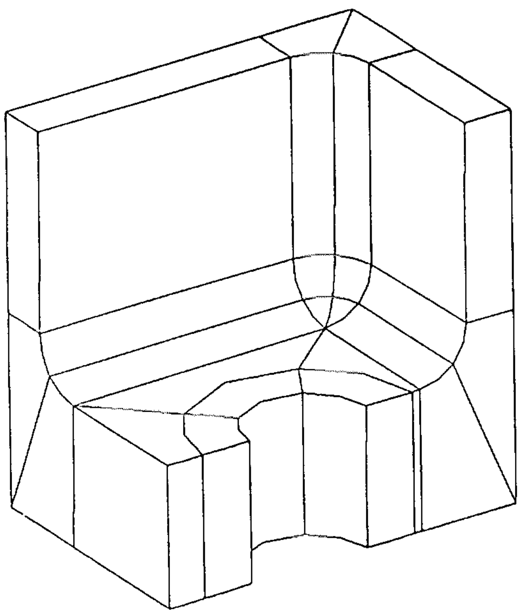

Figure 6: Finite Element Mesh

MSC/PROBE performs checks on the aspect ratios and distortions of elements to provide warnings to the user. In this case, there were warnings that the top two elements at the intersection of the three fillets had moderately high distortion indices. If these are in a region where the solution is smooth, as they are expected to be, this is not a problem. Otherwise the user can modify the mesh in that area.

If there were singularities in the model, those would be isolated by layers of refinement in a geometric progression. It is usually not necessary to model for the stress concentrations, since the higher-order polynomials can accurately represent the stress gradients.

3.3 Finite Element Solution

After the topology, material, load, and constraint data has been entered, a low-order analysis may be run to verify the model. At p=1, this is inexpensive to perform and enables the user to see if the model behaves as expected for the given boundary conditions.

The deformed shape of the fixture is shown in Figure 7. As expected, the hole is being pulled downward by the applied load, and the walls are being pulled inward tangent to the applied constraints.

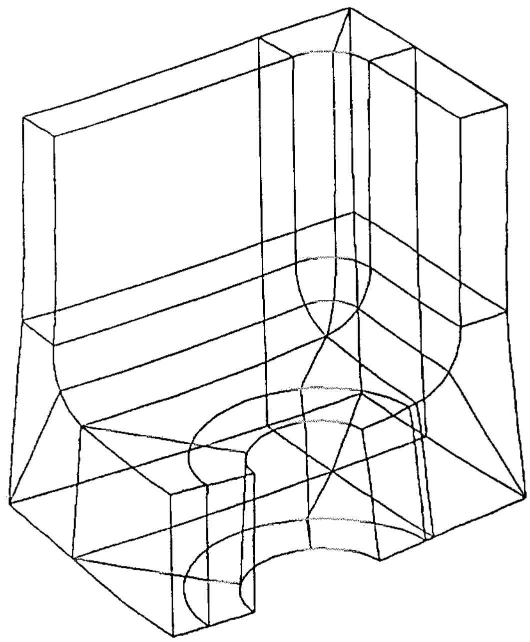

Figure 7: Deformed Shape of Splicing Fixture

After the model has been verified, the user can automatically run a progression of p-levels without changing the model. This can be submitted to batch mode or a compute server. MSC/PROBE can run up to p=8, but that is generally not necessary. The higher p-levels can be run efficiently by choosing to save the stiffness matrices from the lower p-levels, since the lower-order matrices are embedded in the higher-order matrices. This is a result of the hierarchic nature of the polynomial levels in MSC/PROBE.

Running a sequence of solutions with increasing polynomial order allows the accuracy of the model to be verified, and then the convergence of the quantity of interest can be examined.

3.4 Quality Control Procedures

3.4.1 Relative Error in Energy Norm

The first quality control procedure, as outlined in Section 2.4, is the relative error in energy norm, which is related to the root-mean-square error in stresses, and gives a global measure of the solution quality. MSC/PROBE calculates the extrapolated strain energy based on the highest three p-levels, and then computes the relative error in energy norm at each p-level and the convergence rate between each p-level. These values are given in Table 1, along with the number of unconstrained degrees of freedom and strain energy for each p-level, and the percent error in energy norm vs. number of degrees of freedom is shown in a log-log plot in Figure 8. Since the problem is smooth, the exponential rate of convergence from Figure 3 applies.

Table 1: Percent Error in Energy Norm

p	DOF	strain energy	conv. rate	% EEN
1	124	7.7672	0.00	54.83
2	420	10.3569	1.22	25.98
3	728	10.8607	2.03	14.89
4	1261	11.0170	1.84	8.99
5	2025	11.0747	2.17	5.38
6	3074	11.0934	2.08	3.48
7	4462	11.1008	2.14	2.33
8	6237	11.1039	2.14	1.63

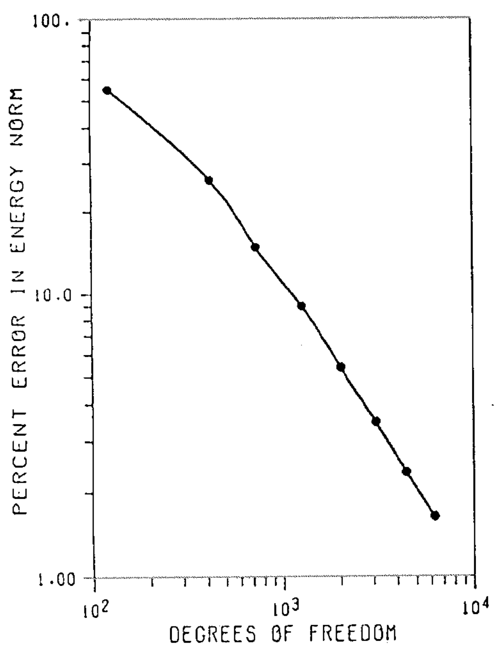

Figure 8: Percent Error in Energy Norm

The error in energy converges to 1.63% at p=8, which is very good. Note that the relative error in energy norm is the square root of the relative error in energy, so that the error in energy is much less. If there were singularities in the model that dominated the behavior, or other gross modelling errors, the effects would show up in the energy.

3.4.2 Equilibrium of Elements and Inter-Element Continuity

The second quality control procedure is the equilibrium of individual elements and inter-element continuity. The stresses on the element faces are explicitly calculated from the displacements and then integrated to find the resultant forces. The sum of the resultant forces on each element should be zero, so the error in equilibrium provides an indicator of where changes in the mesh might be necessary. In addition, comparing the force transfer between two adjacent elements to see if Newton's action/reaction principle is satisfied provides another indication.

These resultant forces are different from the nodal forces that are obtained by multiplying the element stiffness matrix and solution vector. The nodal forces must sum to zero or to the applied nodal force, since that is implied by the global matrix equation. An imbalance in the nodal forces is an indication of errors in the linear solver, not in the finite element approximation.

MSC/PROBE produces reports of the error in equilibrium of each element and of the difference in forces on each face. These values are normalized by characteristic sizes and forces and then ranked to indicate areas that may need mesh changes. The four elements with the highest error in equilibrium per unit volume are shown in Figure 9.

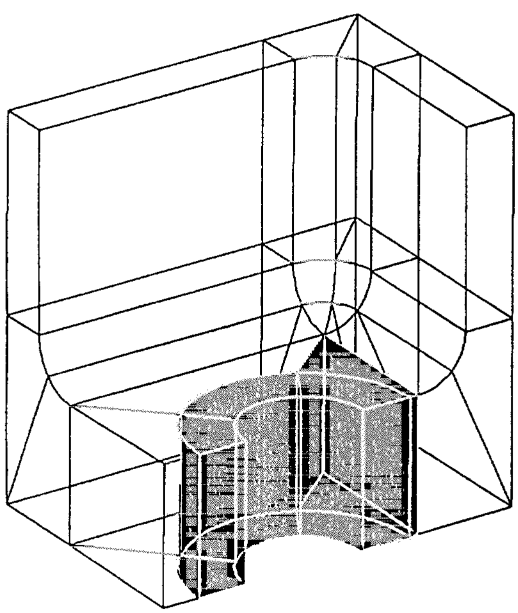

Figure 9: Errors in Equilibrium

These four elements occur in the area of the applied load, where stresses are expected to be the highest, and their errors in equilibrium are not significantly higher than the other elements. When normalized by the highest force on each element, the errors in equilibrium are all a very small percentage. This is an indication that the force transfer in the model is good.

The report for the differences in forces gives similar information. In addition, it gives the resultant forces on all constrained faces and the difference between resultant and applied forces on all loaded faces.

3.4.3 Continuity and Convergence of Point Functionals

The third and most localized quality control procedure is the continuity and convergence of point functionals, such as the maximum stress. This is often the quantity of interest. The stresses are also more difficult to accurately obtain, since they are based on derivatives of the displacements. Since the continuity of stresses between elements is not enforced in MSC/PROBE, and typically the stresses are continuous in the physical part being modelled, the continuity provides a useful check.

The discontinuities of stresses in the whole model can be observed in a contour plot, since the post-processing model is defined to prevent averaging of values at the nodes. If the contours have noticeable breaks between the MSC/PROBE elements, the stresses are not continuous. Note that some post-processors do not always draw smooth contours, so that if there are breaks within the MSC/PROBE elements, it is from the post-processor. A contour plot of the maximum principal stress is shown in Figure 10. There are only some minor breaks in the contours, indicating good overall continuity.

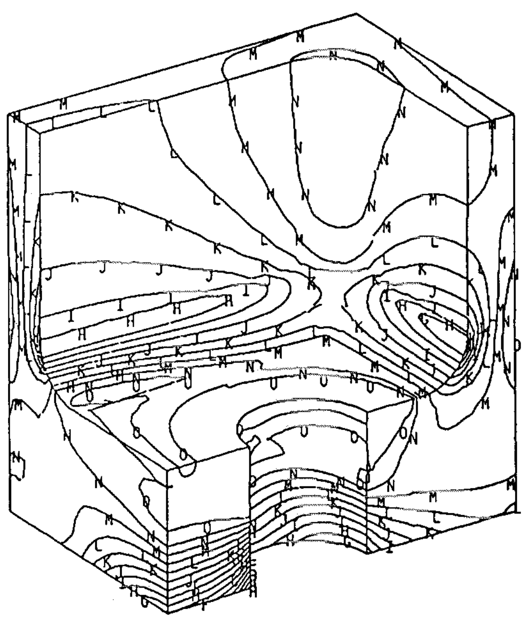

Figure 10: Maximum Principal Stress Contours
(A=0 to O=49000)

Points within the area of interest may be checked more closely. In this case, the highest stress region identified by the contours is at the bottom edge of the hole. The highest point is on the front plane of symmetry, and is not shared by multiple elements. The next point into the hole on the bottom edge is shared by two elements, and at p=8 the two maximum principal stresses are 32510 psi and 32680 psi, a difference of 0.52%. The next point to the left on the bottom edge is also shared by two elements, and at p=8 the two maximum principal stresses are 33120 psi and 33060 psi, a difference of 0.18%. The continuity of the stresses in this high-stress region indicates that the solution is good.

MSC/PROBE has a feature which allows the determination of the maximum stress within the entire model or a specified region. The method works by sampling the stresses to a user-specified refinement, and then sorting the results. The location and magnitude of the maximum stress should converge as the p-level increases. For this model, the maximum principal stress was located at the same point for p=3 to 8. This point was on the left side of the lower edge of the hole, as shown earlier. The convergence of the maximum principal stress at this point is given in Table 2 and shown in Figure 11. Since the principal stress is based on the other components of the stress, convergence is more difficult than for an individual component of stress.

Table 2: Convergence of Maximum Principal Stress

p	DOF	stress
1	124	31260.
2	420	46620.
3	728	49800.
4	1261	49990.
5	2025	52200.
6	3074	51030.
7	4462	51340.
8	6237	51260.

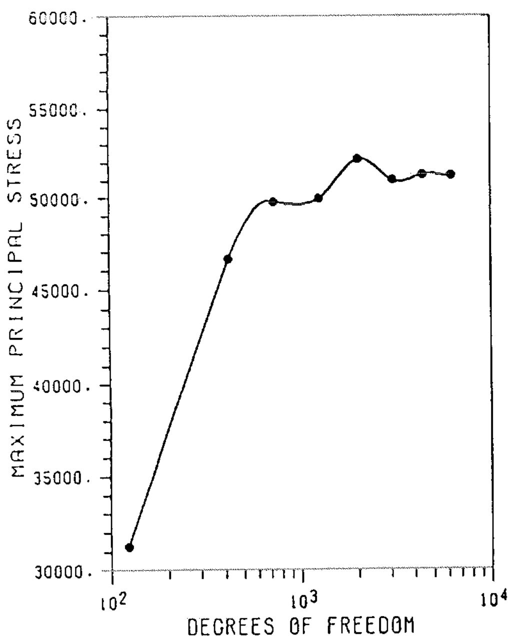

Figure 11: Convergence of Maximum Principal Stress

It can be seen that the maximum stress oscillates and converges strongly to at least two significant digits with a value of 51260 psi at p=8 and a variation of only 0.60% from p=6 to p=8. This is the real quantity of interest in the problem, and the very strong convergence of the maximum stress, in a region where there are strong stress gradients, shows the quality of the solution. With this convergence even at the lower p-levels, it was probably not necessary to run p=7 and 8.

3.5 Post-Processing

After the quality of the solution has been verified, the usual post-processing may be performed. This can consist of deformed shape or contour plots, more detailed stress information, or other data. For instance, the profile of maximum stress component, which in this case is the circumferential stress, up the left edge of the hole is shown in Figure 12.

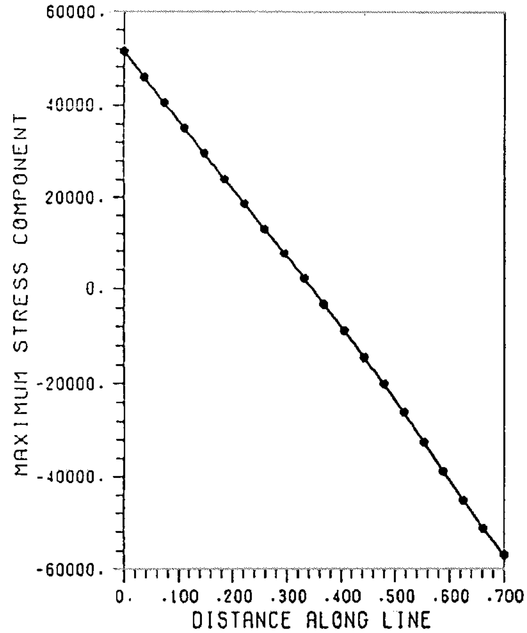

Figure 12: Maximum Stress Component through Hole

With the ability in the p-version to see convergence of important data, the analyst can more confidently make decisions regarding sizing, materials, strength, fatigue, or other characteristics of the component.

4. SUMMARY

The p-version of the finite element method and its implementation in MSC/PROBE were presented. The p-version differs from the h-version in that the h-version uses more elements to increase the number of degrees of freedom, whereas the p-version uses higher-order polynomials with the same number of elements. Extension processes were compared, and the hp-extension is the most efficient with an exponentional rate of convergence. The p-extension with a strongly graded mesh has the same exponential rate initially, and then changes to an algebraic rate. It is recommended that meshes for MSC/PROBE be designed with the strong grading to take advantage of the exponential rate of convergence.

Three quality control procedures for the p-version available in MSC/PROBE were presented. These are:

1. relative error in energy norm,
2. equilibrium of elements and inter-element continuity,
3. continuity and convergence of point functionals.

After the solution has been calculated, the quality control procedures are applied to assess the accuracy of the solution. With multiple solutions on the same mesh, convergence of the quantity of interest, such as maximum stress, may be observed.

A sample problem was presented to illustrate some of the p-version principles. This problem was a splicing fixture, and the goal of the analysis was to find the location and magnitude of the maximum principal stress. The design of the p-version mesh with 18 elements was shown, and the quality control procedures listed above were applied to check the accuracy of the solution. The maximum stress was shown to converge strongly to at least two significant digits with a value of 51260 psi at p=8 and a variation of only 0.60% from p=6 to p=8. This shows that it was probably not necessary to run p=7 and 8.

The client performed two traditional h-version analyses on the fixture: the first had 2567 elements with 5409 degrees of freedom, and resulted in a stress of 48100 psi; the second had 4752 elements with 13409 degrees of freedom and resulted in a stress of 53800 psi. With only these two results and no more refined models, it is difficult to extrapolate the actual stress. Often the second analysis would have never been performed in a production setting. With the p-version in MSC/PROBE, the analyst can run eight solutions, using the degrees of freedom in a more efficient way, and see good convergence of the results.

5. BIBLIOGRAPHY

1. MSC/PROBE Technical Overview. (This includes an extensive bibliography of the p-version of the finite element method.)

2. "Control of Error in Local Stress Analysis with MSC/PROBE, a New p-Version FEA Program," presented at the 1988 USAF Structural Integrity Program Conference, November, 1988.

3. MSC/PROBE Sample Problem: Splicing Fixture, 3D Stress Analysis.

4. MSC/PROBE Sample Problem: Automobile Crankshaft, 3D Stress Analysis.