The Finite Element Method:
A Four-Article Series

The following four-article series was published in a newsletter of the *American Society of Mechanical Engineers (ASME)*. It serves as an introduction to the recent analysis discipline known as the **finite element method**. The author is an engineering consultant and expert witness specializing in finite element analysis.

Finite Element Analysis

1. ***Introduction***

2. ***Pre-processing***

3. ***Solution***

4. ***Post-processing***

FINITE ELEMENT ANALYSIS: Introduction
by Steve Roensch, Roensch & Associates

...First in a four-part series

Finite element analysis (FEA) is a fairly recent discipline crossing the boundaries of mathematics, physics, engineering and computer science. The method has wide application and enjoys extensive utilization in the structural, thermal and fluid analysis areas. The finite element method is comprised of three major phases: (1) **pre-processing**, in which the analyst develops a finite element mesh to divide the subject geometry into subdomains for mathematical analysis, and applies material properties and boundary conditions, (2) **solution**, during which the program derives the governing matrix equations from the model and solves for the primary quantities, and (3) **post-processing**, in which the analyst checks the validity of the solution, examines the values of primary quantities (such as displacements and stresses), and derives and examines additional quantities (such as specialized stresses and error indicators).

The advantages of FEA are numerous and important. A new design concept may be modeled to determine its real world behavior under various load environments, and may therefore be refined prior to the creation of drawings, when few dollars have been committed and changes are inexpensive. Once a detailed CAD model has been developed, FEA can analyze the design in detail, saving time and money by reducing the number of prototypes required. An existing product which is experiencing a field problem, or is simply being improved, can be analyzed to speed an engineering change and reduce its cost. In addition, FEA can be performed on increasingly affordable computer workstations and personal computers, and professional assistance is available.

It is also important to recognize the limitations of FEA. Commercial software packages and the required hardware, which have seen substantial price reductions, still require a significant investment. The method can reduce product testing, but cannot totally replace it. Probably most important, an inexperienced user can deliver incorrect answers, upon which expensive decisions will be based. FEA is a demanding tool, in that the analyst must be proficient not only in elasticity or fluids, but also in mathematics, computer science, and especially the finite element method itself.

Which FEA package to use is a subject that cannot possibly be covered in this short discussion, and the choice involves personal preferences as well as package functionality. Where to run the package depends on the type of analyses being performed. A typical finite element solution requires a fast, modern disk subsystem for acceptable performance. Memory requirements are of course dependent on the code, but in the interest of performance, the more the better, with a representative range measured in gigabytes per user. Processing power is the final link in the performance chain, with clock speed, cache, pipelining and multi-processing all contributing to the bottom line. These analyses can run for hours on the fastest systems, so computing power is of the essence.

One aspect often overlooked when entering the finite element area is education. Without adequate training on the finite element method and the specific FEA package, a new user will not be productive in a reasonable amount of time, and may in fact fail miserably. Expect to dedicate one to two weeks up front, and another one to two weeks over the first year, to either classroom or self-help education. It is also important that the user have a basic understanding of the computer's operating system.

Next month's article will go into detail on the pre-processing phase of the finite element method.

FINITE ELEMENT ANALYSIS: Pre-processing

by Steve Roensch, Roensch & Associates

...Second in a four-part series

As discussed last month, finite element analysis is comprised of pre-processing, solution and post-processing phases. The goals of pre-processing are to develop an appropriate finite element mesh, assign suitable material properties, and apply boundary conditions in the form of restraints and loads.

The finite element mesh subdivides the geometry into **elements**, upon which are found **nodes**. The nodes, which are really just point locations in space, are generally located at the element corners and perhaps near each midside. For a two-dimensional (2D) analysis, or a three-dimensional (3D) thin shell analysis, the elements are essentially 2D, but may be "warped" slightly to conform to a 3D surface. An example is the thin shell linear quadrilateral; *thin shell* implies essentially classical shell theory, *linear* defines the interpolation of mathematical quantities across the element, and *quadrilateral* describes the geometry. For a 3D solid analysis, the elements have physical thickness in all three dimensions. Common examples include solid linear brick and solid parabolic tetrahedral elements. In addition, there are many special elements, such as axisymmetric elements for situations in which the geometry, material and boundary conditions are all symmetric about an axis.

The model's degrees of freedom (dof) are assigned at the nodes. Solid elements generally have three translational dof per node. Rotations are accomplished through translations of groups of nodes relative to other nodes. Thin shell elements, on the other hand, have six dof per node: three translations and three rotations. The addition of rotational dof allows for evaluation of quantities through the shell, such as bending stresses due to rotation of one node relative to another. Thus, for structures in which classical thin shell theory is a valid approximation, carrying extra dof at each node bypasses the necessity of modeling the physical thickness. The assignment of nodal dof also depends on the class of analysis. For a thermal analysis, for example, only one temperature dof exists at each node.

Developing the mesh is usually the most time-consuming task in FEA. In the past, node locations were keyed in manually to approximate the geometry. The more modern approach is to develop the mesh directly on the CAD geometry, which will be (1) **wireframe**, with points and curves representing edges, (2) **surfaced**, with surfaces defining boundaries, or (3) **solid**, defining where the material is. Solid geometry is preferred, but often a surfacing package can create a complex blend that a solids package will not handle. As far as geometric detail, an underlying rule of FEA is to "model what is there", and yet simplifying assumptions simply must be applied to avoid huge models. Analyst experience is of the essence.

The geometry is meshed with a mapping algorithm or an automatic free-meshing algorithm. The first maps a rectangular grid onto a geometric region, which must therefore have the correct number of sides. Mapped meshes can use the accurate and cheap solid linear brick 3D element, but can be very time-consuming, if not impossible, to apply to complex geometries. Free-meshing automatically subdivides meshing regions into elements, with the advantages of fast meshing, easy mesh-size transitioning (for a denser mesh in regions of large gradient), and adaptive capabilities. Disadvantages include generation of huge models, generation of distorted elements, and, in 3D, the use of the rather expensive solid parabolic tetrahedral element. It is always important to check elemental distortion prior to solution. A badly distorted element will cause a matrix singularity, killing the solution. A less distorted element may solve, but can deliver very poor answers. Acceptable levels of distortion are dependent upon the solver being used.

Material properties required vary with the type of solution. A linear statics analysis, for example, will require an elastic modulus, Poisson's ratio and perhaps a density for each material. Thermal properties are required for a thermal analysis. Examples of restraints are declaring a nodal

translation or temperature. Loads include forces, pressures and heat flux. It is preferable to apply boundary conditions to the CAD geometry, with the FEA package transferring them to the underlying model, to allow for simpler application of adaptive and optimization algorithms. It is worth noting that the largest error in the entire process is often in the boundary conditions. Running multiple cases as a sensitivity analysis may be required.

Next month's article will discuss the solution phase of the finite element method.

FINITE ELEMENT ANALYSIS: Solution
by Steve Roensch, Roensch & Associates

...Third in a four-part series

While the pre-processing and post-processing phases of the finite element method are interactive and time-consuming for the analyst, the solution is often a batch process, and is demanding of computer resource. The governing equations are assembled into matrix form and are solved numerically. The assembly process depends not only on the type of analysis (e.g. static or dynamic), but also on the model's element types and properties, material properties and boundary conditions.

In the case of a linear static structural analysis, the assembled equation is of the form **Kd = r**, where **K** is the system stiffness matrix, **d** is the nodal degree of freedom (dof) displacement vector, and **r** is the applied nodal load vector. To appreciate this equation, one must begin with the underlying elasticity theory. The strain-displacement relation may be introduced into the stress-strain relation to express stress in terms of displacement. Under the assumption of compatibility, the differential equations of equilibrium in concert with the boundary conditions then determine a unique displacement field solution, which in turn determines the strain and stress fields. The chances of directly solving these equations are slim to none for anything but the most trivial geometries, hence the need for approximate numerical techniques presents itself.

A finite element mesh is actually a displacement-nodal displacement relation, which, through the element interpolation scheme, determines the displacement anywhere in an element given the values of its nodal dof. Introducing this relation into the strain-displacement relation, we may express strain in terms of the nodal displacement, element interpolation scheme and differential operator matrix. Recalling that the expression for the potential energy of an elastic body includes an integral for strain energy stored (dependent upon the strain field) and integrals for work done by external forces (dependent upon the displacement field), we can therefore express system potential energy in terms of nodal displacement.

Applying the principle of minimum potential energy, we may set the partial derivative of potential energy with respect to the nodal dof vector to zero, resulting in: a summation of element stiffness integrals, multiplied by the nodal displacement vector, equals a summation of load integrals. Each stiffness integral results in an element stiffness matrix, which sum to produce the system stiffness matrix, and the summation of load integrals yields the applied load vector, resulting in **Kd = r**. In practice, integration rules are applied to elements, loads appear in the **r** vector, and nodal dof boundary conditions may appear in the **d** vector or may be partitioned out of the equation.

Solution methods for finite element matrix equations are plentiful. In the case of the linear static **Kd = r**, inverting **K** is computationally expensive and numerically unstable. A better technique is Cholesky factorization, a form of Gauss elimination, and a minor variation on the "LDU" factorization theme. The **K** matrix may be efficiently factored into **LDU**, where **L** is lower triangular, **D** is diagonal, and **U** is upper triangular, resulting in **LDUd = r**. Since **L** and **D** are easily inverted, and **U** is upper triangular, **d** may be determined by back-substitution. Another popular approach is the wavefront method, which assembles and reduces the equations at the same time. Some of the best modern solution methods employ sparse matrix techniques. Because node-to-node stiffnesses are non-zero only for nearby node pairs, the stiffness matrix has a large number of zero entries. This can be exploited to reduce solution time and storage by a factor of 10 or more. Improved solution methods are continually being developed. The key point is that the analyst must understand the solution technique being applied.

Dynamic analysis for too many analysts means normal modes. Knowledge of the natural frequencies and mode shapes of a design may be enough in the case of a single-frequency vibration of an existing product or prototype, with FEA being used to investigate the effects of mass, stiffness and damping modifications. When investigating a future product, or an existing

design with multiple modes excited, forced response modeling should be used to apply the expected transient or frequency environment to estimate the displacement and even dynamic stress at each time step.

This discussion has assumed h-code elements, for which the order of the interpolation polynomials is fixed. Another technique, p-code, increases the order iteratively until convergence, with error estimates available after one analysis. Finally, the boundary element method places elements only along the geometrical boundary. These techniques have limitations, but expect to see more of them in the near future.

Next month's article will discuss the post-processing phase of the finite element method.

FINITE ELEMENT ANALYSIS: Post-processing
by Steve Roensch, Roensch & Associates

...Last in a four-part series

After a finite element model has been prepared and checked, boundary conditions have been applied, and the model has been solved, it is time to investigate the results of the analysis. This activity is known as the post-processing phase of the finite element method.

Post-processing begins with a thorough check for problems that may have occurred during solution. Most solvers provide a log file, which should be searched for warnings or errors, and which will also provide a quantitative measure of how well-behaved the numerical procedures were during solution. Next, reaction loads at restrained nodes should be summed and examined as a "sanity check". Reaction loads that do not closely balance the applied load resultant for a linear static analysis should cast doubt on the validity of other results. Error norms such as strain energy density and stress deviation among adjacent elements might be looked at next, but for h-code analyses these quantities are best used to target subsequent adaptive remeshing.

Once the solution is verified to be free of numerical problems, the quantities of interest may be examined. Many display options are available, the choice of which depends on the mathematical form of the quantity as well as its physical meaning. For example, the displacement of a solid linear brick element's node is a 3-component spatial vector, and the model's overall displacement is often displayed by superposing the deformed shape over the undeformed shape. Dynamic viewing and animation capabilities aid greatly in obtaining an understanding of the deformation pattern. Stresses, being tensor quantities, currently lack a good single visualization technique, and thus derived stress quantities are extracted and displayed. Principal stress vectors may be displayed as color-coded arrows, indicating both direction and magnitude. The magnitude of principal stresses or of a scalar failure stress such as the Von Mises stress may be displayed on the model as colored bands. When this type of display is treated as a 3D object subjected to light sources, the resulting image is known as a *shaded image stress plot.* Displacement magnitude may also be displayed by colored bands, but this can lead to misinterpretation as a stress plot.

An area of post-processing that is rapidly gaining popularity is that of adaptive remeshing. Error norms such as strain energy density are used to remesh the model, placing a denser mesh in regions needing improvement and a coarser mesh in areas of overkill. Adaptivity requires an associative link between the model and the underlying CAD geometry, and works best if boundary conditions may be applied directly to the geometry, as well. Adaptive remeshing is a recent demonstration of the iterative nature of h-code analysis.

Optimization is another area enjoying recent advancement. Based on the values of various results, the model is modified automatically in an attempt to satisfy certain performance criteria and is solved again. The process iterates until some convergence criterion is met. In its scalar form, optimization modifies beam cross-sectional properties, thin shell thicknesses and/or material properties in an attempt to meet maximum stress constraints, maximum deflection constraints, and/or vibrational frequency constraints. Shape optimization is more complex, with the actual 3D model boundaries being modified. This is best accomplished by using the driving dimensions as optimization parameters, but mesh quality at each iteration can be a concern.

Another direction clearly visible in the finite element field is the integration of FEA packages with so-called "mechanism" packages, which analyze motion and forces of large-displacement multi-body systems. A long-term goal would be real-time computation and display of displacements and stresses in a multi-body system undergoing large displacement motion, with frictional effects and fluid flow taken into account when necessary. It is difficult to estimate the increase in computing power necessary to accomplish this feat, but 2 or 3 orders of magnitude is probably close. Algorithms to integrate these fields of analysis may be expected to follow the computing power increases.

In summary, the finite element method is a relatively recent discipline that has quickly become a mature method, especially for structural and thermal analysis. The costs of applying this technology to everyday design tasks have been dropping, while the capabilities delivered by the method expand constantly. With education in the technique and in the commercial software packages becoming more and more available, the question has moved from "Why apply FEA?" to "Why not?". The method is fully capable of delivering higher quality products in a shorter design cycle with a reduced chance of field failure, provided it is applied by a capable analyst. It is also a valid indication of thorough design practices, should an unexpected litigation crop up. The time is now for industry to make greater use of this and other analysis techniques.