Postprocessing Operations in CAD

The ensemble of activities by which the engineering significance of a mathematical field solution is evaluated is generally termed *postprocessing*. It includes the derivation of specific numerical results as well as their graphical presentation. A postprocessor is a major part of any design system since it allows relevant data to be extracted from the solution and presented in a way that has meaning to the user. Postprocessing should desirably be an interactive process, allowing the designer to query the solution. This chapter outlines the major requirements of postprocessing, while the following one describes a particular postprocessor structure by way of illustration.

Postprocessing is the activity of converting mathematical solutions into engineering results. This chapter examines the operations required in a selection of postprocessing tasks, in an attempt to exhibit unity in the processing requirements that underlie the great diversity of applicational needs. Thus the discussion here will begin by examining simple but illustrative design problems, and will then generalize to broader issues applicable not only to the examples treated but also to other cases.

Inductance Calculations

The calculation of terminal inductance values is probably the most common single requirement in magnetic device design. Despite its frequent occurrence, the determination of inductance is fraught with subtle difficulties not always evident at first glance. Examining a simple magnetic-core reactor and considering the methods available for calculating its inductance is therefore both indicative of general methods and useful in its own right.

A Simple Inductor

A simple, two-dimensional inductor may be formed by winding a coil around a highly permeable core, as shown in Fig. 1. It is assumed, for the

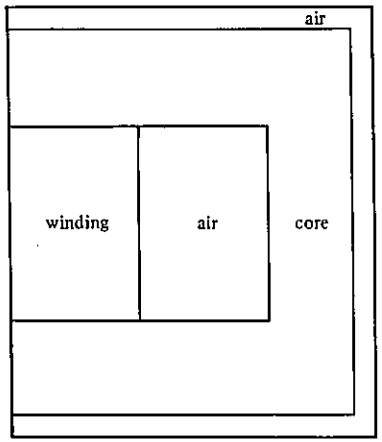

Figure 1. A simple inductor, assumed to extend infinitely in the direction normal to the paper.

sake of simplicity, that the device extends sufficiently far into the paper to allow purely two- dimensional analysis. Not only is the finite element analysis of the device made easy by this assumption, but basic conventional design rules may be applied for comparison when the inductor is simplified in this way.

Since this device is to furnish a prescribed value of terminal inductance, that is, to store energy in the magnetic field in its core, certain further simplifying assumptions are possible in the analysis. In addition to taking the field problem to be two-dimensional, two main simplifications will be made: the core iron will be assumed lossless, and external leakage flux, i.e., leakage flux not in the window area of the core will be neglected.

Assuming iron loss due to eddy currents to be absent permits finding the magnetic field by means of one or more static field solutions. The only loss accounted for will then be the ohmic loss in the winding conductor itself. This loss cannot be calculated from a magnetic field solution, but must be computed from the known wire resistance. The inductor will therefore be represented, to an approximation adequate for most purposes, by the equivalent circuit shown in Fig. 2.

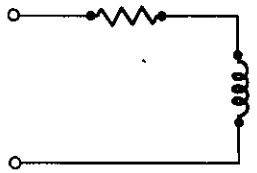

Figure 2. An equivalent circuit of the simple inductor shown in Fig. 1, assuming lossless iron.

The principal part of all the flux which links the current-carrying winding will very likely be contained in the iron core, while some small portion will close through the air space within the window opening of the core. Only a very minor amount will take a path outside the core itself. Hence, little error will be incurred in not modelling the air space and winding outside the core; it suffices to take the outer edge of the core perimeter to be a flux line and to model only the iron, air, and winding inside this flux-line boundary. Furthermore, the symmetry of the core and winding make it necessary for the core centerline to be a line of symmetry, and hence the separatrix flux line which separates the flux lines of clockwise circulation from those which close in a counterclockwise sense. It is therefore necessary to model only one half, say the right half, of the inductor. A finite element model of this half-problem appears in Fig. 3. While moderately crude, this discretization probably suffices to produce a magnetic field solution adequate for inductance computations. Any lingering doubts about the amount of external leakage flux may of course be dispelled by employing a similar model, but with some of the exterior air space explicitly included.

Physical realization of such a device is invariably subject to many requirements other than electromagnetic, such as weight limitations or the cost of material used. These constraints not infrequently will be handled

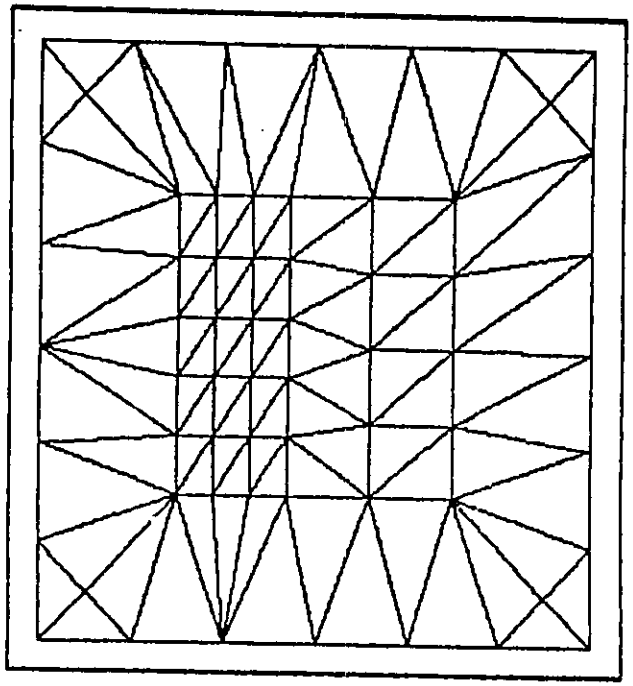

Figure 3. Finite element model used for analysis of the simple inductor.

by other CAD or CAE systems. Only its electrical design will be considered here.

Definitions of Inductance

The computation of inductance is a simple affair in principle: numerical values taken from a field solution are substituted into an appropriate expression that gives the inductance in terms of the magnetic field, and the result is evaluated. Unfortunately, there is no single definition of just what constitutes inductance for a nonlinear inductor. Indeed, there are several accepted definitions, of which two may be considered fundamental:

1. The number of flux linkages of the winding, divided by the current in the winding.
2. The energy stored in the inductor, divided by one-half the current squared.

All other accepted definitions make reference to either the flux linkages or the stored energy; they may therefore be considered variants of the above. Both definitions give identical results for linear inductors. In the nonlinear case, however, they do not. Which definition should be used, if indeed either, then depends very much on the uses to which the result will be put—or, what amounts to the same thing, the kind of measurement which the analysis is intended to mimic. Because periodic excitations (not necessarily sinusoidal) are of frequent interest, the second definition is often further refined in one direction or another, for example, to refer to the *average* stored energy and the *fundamental component* of current.

There are two experimental ways of measuring inductance, which parallel the two analytic definitions. One seeks to determine flux linkages more or less directly, the other measures time-averaged energy storage.

Direct measurement of flux linkages in an inductor is readily accomplished by a method widely used since the invention of the galvanometer and generally called the *ballistic fluxmeter* method. The principle is to open- circuit the inductor winding suddenly and to integrate the resulting terminal voltage over time so as to determine the total flux linkages that existed prior to open-circuiting. If $n\phi$ is the number of flux linkages, Faraday's law prescribes an open-circuit terminal voltage

$$e = -\frac{d(n\phi)}{dt}, \tag{1}$$

whose time integral, as measured by an integrator circuit, is equal to the initial value of flux linkages:

$$n\phi(0) = -\int_{-\infty}^{0} e(t)\,dt. \tag{2}$$

In practice, the current is usually reversed rather than merely disconnected. The number of flux linkages is thereby doubled, reducing the error of measurement slightly and simplifying the experimental procedure; but more importantly, hysteretic effects are substantially reduced or even eliminated. This method is directly related to the first, flux linkage based, definition of inductance. Its noteworthy aspect is that the measured value is not directly dependent on the shape of the material *B-H* curve and is therefore not related to the amount of energy stored; various different *B-H* curves could yield the same value of flux linkages. In other words, the measured result is truly the total of flux linkages; no approximations are involved. Inductance is then determined as the ratio

$$L = \frac{n\phi(0)}{i(0)}. \tag{3}$$

A second, very common, technique for measuring inductance is to find the root-mean-square values of voltage and current when a time-sinusoidal excitation is applied to the device. The inductance value is then found as the measured reactance divided by angular frequency,

$$L = \frac{X}{\omega} = \frac{E_{rms}}{\omega I_{rms}}. \tag{4}$$

This measurement is closely related to the second or *energy* definition of inductance, since the reactance is in principle the amount of energy stored per unit of coil current, averaged over the ac cycle. Let it be supposed for the moment that a source of sinusoidal current is connected to the terminals of the inductor, Fig. 2,

$$i(t) = I \sin \omega t. \tag{5}$$

The corresponding flux linkages will reflect the shape of the *B-H* curve of the core material, so the flux linkages, and hence the terminal voltage, will not be sinusoidal but must contain all odd harmonics as well:

$$e(t) = \sum_{k=1}^{\infty} E_{2k-1} \cos [(2k-1)\omega t + \theta_{2k-1}]. \tag{6}$$

When root-mean-square measurements are made, the measured voltage will include a fundamental term and an additional contribution from every harmonic term. That is to say,

$$E_{rms} = \sqrt{\frac{1}{2} \sum_{k=1}^{\infty} E_{2k-1}^2} \tag{7}$$

so that equation (4) reads

$$L = \frac{1}{2\omega I_{rms}} \sqrt{\sum_{k=1}^{\infty} E_{2k-1}^2}. \tag{8}$$

The laboratory measurements implied by (8) are easily carried out, so that (8), as well as a similar equation based on sinusoidal applied voltage and a nonsinusoidal resulting current, are commonly used as inductance definitions. To calculate inductance by simulating this measurement is feasible, but computationally a bit lengthy. Strictly speaking, it will be necessary to solve a static field problem for every one of a set of current values, from zero up to the peak value; then to find the corresponding flux linkages, and to differentiate these so as to obtain the voltage waveform. From the voltage waveform, the root-mean-square value can then be computed. Such a computation is only very rarely carried out, not only because it is complicated but also because the inductance value as defined by (8) is not ideally suited to many applications. Indeed, quite a few practical needs are better satisfied by a similar definition, but referring to the fundamental components only,

$$L = \frac{E_1^2}{2\omega I_{rms}}. \tag{9}$$

This definition is particularly useful if the winding under consideration is one of several, and it is known in advance that any reasonable interconnection of the windings will result in a great deal of harmonic cancellation. Such interconnections are often encountered in electric machines.

It is possible to make measurements which seek to determine the stored energy corresponding to a specific instantaneous current value, so as to apply directly the energy-based analytic definition of inductance. The principle is exactly opposite to that of the ballistic fluxmeter experiment: the inductor coil is suddenly connected to a pure resistance, and the instantaneous power in the resistor is integrated over time:

$$W = \int_0^{\infty} \frac{\partial}{\partial t} [n\phi(t)]\, i(t)\, dt. \tag{10}$$

As the energy stored in the inductor is gradually dissipated in the resistor, the current $i(t)$ falls, and the flux linkage $n\phi(t)$ falls with it in accordance with the shape of the *B-H* curve of the core material. The inductance value thus measured corresponds truly to the second definition given above, and faithfully reflects the saturation characteristics of the core material. Such measurements, however, are rarely made because they are experimentally difficult to carry out.

Inductance from Flux Linkage

When a CAD system is used to solve the magnetic field in the inductor described above, a static solution is normally produced. Static solutions can be closely related to the first, flux-linkage-based, definition of inductance, and the corresponding calculations are quite easily carried out.

Let the magnetic field be determined in the simple inductor described above. To make the matter concrete, the core is taken to have external

dimensions 0.8 m by 0.4 m, so that the half modelled in Fig. 3 measures 0.4 m by 0.4 m; the window size is 0.2 m by 0.2 m. The coil occupies a space 0.08 m by 0.2 m. The *B-H* curve for the model is as shown in Fig. 4. It may be noted that the curve is carried up to quite high saturation levels; the current increments along the H axis are 10000 ampere-meters between ticks, while the values of B extend up to 2.48 tesla. A solution, showing flux lines in the core and in the window space of the inductor, appears in Fig. 5. The coil excitation in this case is 500 ampere-turns. With the mean flux path length in the magnetic core of the order of one meter, it is clear from Fig. 4 that this solution is essentially a linear one; the magnetic material is working well within the linear region of its saturation curve throughout. The material permeability is correspondingly high and, as might be expected, practically all the magnetic flux is confined to the iron core.

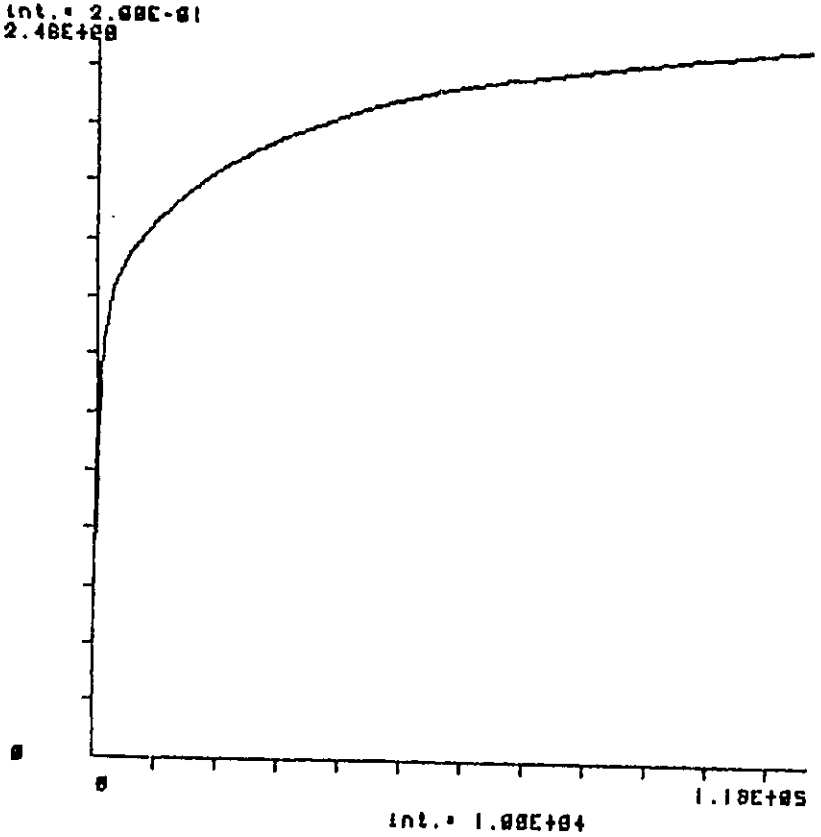

Figure 4. Magnetization curve used in inductance calculations for the simple inductor.

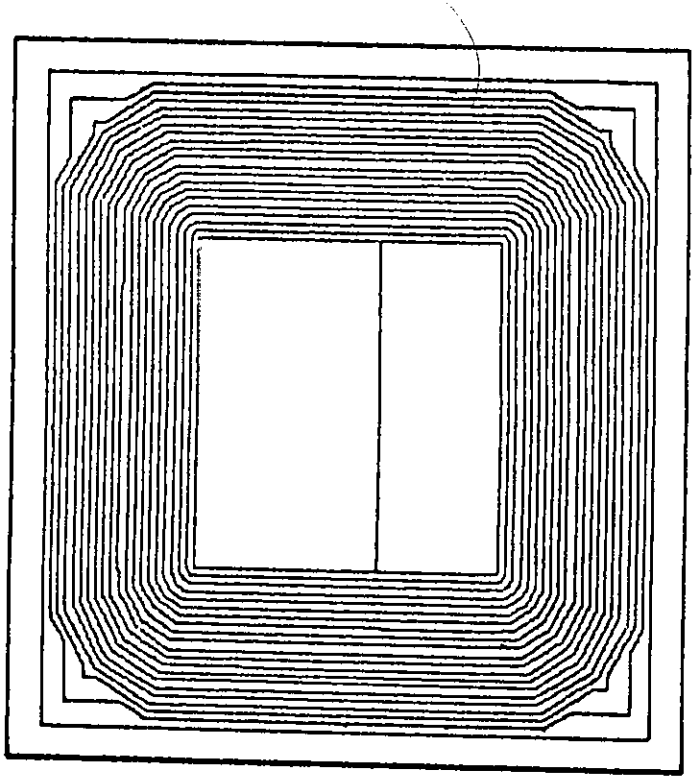

Figure 5. Flux distribution in the simple inductor at low saturation level (coil excitation 500 ampere-turns).

Inductance calculations based on flux linkage counts are very simple to carry out if the flux is principally confined to a clearly defined iron member or path, as it is in Fig. 5. In effect, it suffices to pretend that the winding can be replaced by a single filamentary coil and to calculate the flux spanned by that coil. In the case shown here, it does not even very much matter where the filamentary coil is placed, provided only that one side of it lies somewhere within the core window, while the other side is placed at a corresponding point in the other window. The precise position is unimportant, for so long as the coil is threaded by the iron core, it will link practically the same amount of flux.

Suppose next that a hypothetical single-turn filamentary coil is placed in the window space of the inductor as shown in Fig. 6(a). The flux ϕ linked by this turn can be expressed in terms of the flux density B and the area enclosed by the coil as

$$\phi = \int \mathbf{B} \cdot \mathbf{dS}. \tag{11}$$

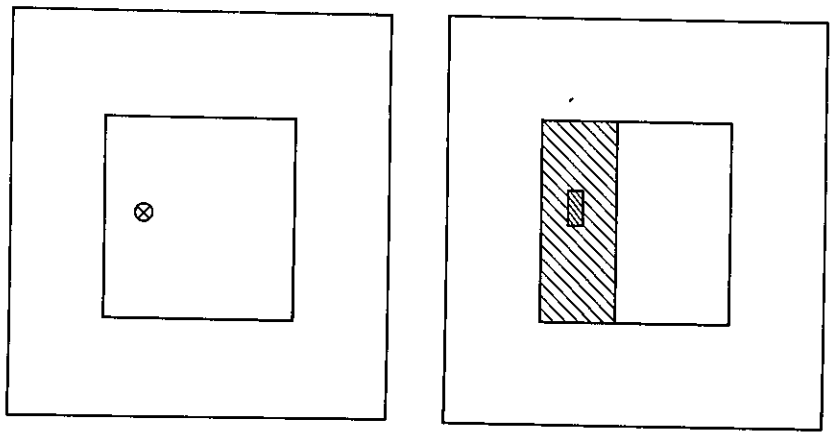

Figure 6. (a) A single-turn conductor or thin-sided coil in the inductor structure. (b) A slim-sided coil as part of a thick winding.

Substituting the vector potential **A** for the flux density **B**, in accordance with

$$\text{curl } \mathbf{A} = \mathbf{B}, \tag{12}$$

one obtains

$$\phi = \oint \mathbf{A} \cdot \mathbf{dl}, \tag{13}$$

where the integration is around the closed contour formed by the coil, i.e., following the wire of which the coil is made. Because the problem is essentially two-dimensional, the end contributions to the integral in (13) may be ignored. Along the two long sides of the coil, which lie within the window space of the core, the vector potential has constant values, say A_L on the left and A_R on the right. If the length of such a side is Z, then the integral in (13) takes the simple form

$$\phi = Z\,(A_R - A_L). \tag{14}$$

For a thin multi-turn coil, a similar development may be carried out as for the single turn above. If the coil is made of n turns, then integration must be carried out over the spiral contour traced by the wire. If the turns are tightly packed and occupy little space, the same argument may be applied to each one; the flux linkages $n\phi$ are therefore n times larger,

$$n\phi = n\,Z\,(A_R - A_L). \tag{15}$$

To keep the total exciting ampere-turns I of the n-turn coil the same as those of the single turn, the current must be reduced by a factor of n, to

$i = I/n$, because the same current is threaded through the core window n times. The inductance is therefore finally calculated as the number of flux linkages no divided by the coil current I/n:

$$L = \frac{n\phi}{I/n} = \frac{n^2\,Z\,(A_R - A_L)}{I}. \tag{16}$$

It will be noted that the inductance value is proportional to the square of the number n of turns in the winding. Computing the inductance of a multi-turn coil with thin sides is thereby reduced to determining the vector potentials A_R and A_L at its two sides and multiplying by $n^2\,Z$, mathematically straightforward operations.

The very simple flux linkage method of calculating inductance is not applicable to windings that fill an extended portion of space, because there is no particular point at which to measure the vector potential. In effect, thickening the coil sides requires forming some sort of average value of the vector potential over the space occupied by the winding. The mathematical development in this case is fortunately still quite simple, as may be seen from the following.

A thick winding may be regarded as a set of individual turns, or perhaps as a set of smaller, thin, multi-turn coils, all connected in series. The total flux linkage of the thick winding must then be the sum of the individual flux linkages of its component parts. In Fig. 6(b) the simple inductor appears again, with a particular small group of adjacent turns highlighted; if this group has n_k turns, say, then its flux linkages are

$$(n\phi)_k = n_k\,(A_{kR} - A_{kL}), \tag{17}$$

where it is assumed that the group is compact enough for the potential values at its left and right sides, A_{kL} and A_{kR}, to be substantially constant over the extent of the group. Summing over all the m elementary coils that make up the winding, the total flux linkage is

$$n\phi = \sum_{k=1}^{m} (n\phi)_k = \sum_{k=1}^{2m} n_k \frac{J}{|J|} A_k. \tag{18}$$

Here the last summation is carried out over coil *sides*, rather than coils, in the interest of generality. To ensure that oppositely directed sides enter into the summation with opposite signs, each term has attached to it the local direction of the current density **J**.

The flux linkage calculation as set out in equation (18) requires the winding to be divided into sections and their individual contributions to be summed. The inconvenience of explicit subdivision is avoided by using a more convenient formulation of the problem. To reformulate, it suffices to note that, since all turns in the winding are in series and therefore carry

the same current, the *signed* number of turns of a particular coil side, as in the rightmost member of (18), may be written

$$n_k \frac{J}{|J|} = \frac{\int_{S_k} J \, dS}{ni} n, \tag{19}$$

where the the surface of integration S_k is the cross-sectional area of the kth coil side. The numerator in equation (19) therefore represents the number of ampere-turns contributed by the kth elementary coil side, while the denominator equals the total number of ampere-turns for the whole coil. In simple words, (19) thus says that the number n_k of turns in the elementary coil side can be found by determining what fraction of total ampere-turns it contributes. Substituting (19) into (18), there results

$$n\phi = \frac{1}{i} \int_S A \, J \, dS, \tag{20}$$

where i represents, as previously, the current in the coil. The winding inductance is again calculated from its definition as the number of flux linkages per ampere:

$$L = \frac{Z}{i^2} \int_S A \, J \, dS, \tag{21}$$

where Z represents, as before, the length of the winding in the z direction. While this inductance expression is valid for two-dimensional cases as given here, it is readily generalized to cover situations where the coil sides are not straight, indeed to coil configurations not describable by two-dimensional approximations at all. Omitting the mathematical details, the result turns out to be

$$L = \frac{1}{i^2} \int_U \mathbf{A} \cdot \mathbf{J} \, dU, \tag{22}$$

where U is the volume occupied by the winding.

Although the number of turns n does not appear explicitly in equation (22), L is dependent on it through the values of J and A. Suppose, for example, that the number of turns is increased from n to N, and that the current i is correspondingly reduced to $(n/N)i$. The total number of ampere-turns remains unchanged in this process; so does the current density J. The vector potential A, however, depends only on J, and therefore remains unchanged as well. Thus, the numerator in (22) is not changed, while the denominator is altered by $(n/N)^2$. The inductance therefore increases by the factor $(N/n)^2$, in accordance with equation (22).

Stored Energy and Inductance

The inductance value of a reactor may be defined in terms of the magnetic energy W stored in it. This definition relies on the familiar energy expression

$$W = \frac{1}{2} L \, i^2, \tag{23}$$

where i is the terminal current of the inductor, as in the above. If the stored energy W can be determined, then the inductance value follows,

$$L = \frac{2W}{i^2}. \tag{24}$$

The stored energy contained within a magnetic device may be found by integrating the stored magnetic energy density w over the volume U of the device:

$$W = \int_U w \, dU. \tag{25}$$

But at any point in a magnetic material, the stored energy density is given by the area to the left of its magnetization characteristic:

$$w = \int_0^B \mathbf{H}(\mathbf{b}) \cdot \mathbf{db}, \tag{26}$$

where $\mathbf{b}$ is a dummy variable which follows the flux density along the B-H characteristic, which in general is a vector relationship, up to its final value B. The inductance of the inductor may therefore be calculated by using the relationship obtained by combining equations (24)–(26):

$$L = \frac{2}{i^2} \int_U \int_0^B \mathbf{H} \cdot \mathbf{db} \, dU \tag{27}$$

Evaluation of this quantity is fairly straightforward in well designed CAD systems, which have access to the material B-H curves. After all, the magnetization curve of every material must be known to the system if solution of the field problem is to be possible in the first place.

In several older magnetics CAD systems, the material magnetization characteristics are not actually available at postprocessing time; instead, the material reluctivity (inverse permeability, i.e., the value of H/B at the solution point) is known. In such circumstances, a rough approximation based on linear theory can sometimes lead to useful results. Linear or not, equation (26) may always be rewritten as

$$w = \int_0^B \nu \, \mathbf{b} \cdot \mathbf{db}. \tag{28}$$

For a magnetically linear material, the reluctivity may be moved across the integral sign, so that (28) becomes

$$w = \nu \int_0^B \mathbf{b} \cdot \mathbf{db} = \frac{B^2}{2\mu}. \tag{29}$$

Although this expression is not correct for nonlinear materials, it sometimes leads to results of adequate accuracy. That is to say,

$$L_{approx} = \frac{2}{i^2} \int_U \frac{B^2}{2\mu}\, dU \tag{30}$$

can furnish useful approximations where better ones are not available. Although the substitution of linear for nonlinear magnetization characteristics may seem a very crude approximation, it is frequently permissible because the use of *inductance* as a circuit parameter is often confined to near-linear cases anyway. If the approximation is very bad, then quite likely the use of inductance for whatever further purpose is likely to be a bad idea also!

An Example Calculation

The methods of inductance calculation discussed above may be illustrated by giving numerical examples. These relate to the inductor of Fig. 1, whose core material is characterized by the curve of Fig. 4.

A field solution obtained for very low excitations is shown, and briefly discussed, in Fig. 5. The flux density in that solution is everywhere fairly low, as the result of the low excitation value; the material permeability is high, and the flux is therefore almost entirely confined to the iron core. A different situation obtains when the excitation is raised to a much higher value. A second solution, computed for an excitation of 100 kiloampere-turns, is shown in Fig. 7. It is clear from Fig. 7 that most of the flux still resides in the iron core; but the leakage flux crossing the core window is considerably larger than in the earlier case. Even at this high saturation level, however, the leakage flux still represents only a modest fraction of the total flux, so that the same methods of inductance calculation may be applied, as is indeed done below.

Numerical values of flux densities at a few selected points are shown for both solutions in Table 1. The increase in leakage flux is immediately evident. Flux densities in the iron portions of the inductor rise by a factor of about 2.5 from one solution to the other; but the flux density in the winding space rises by 26 times, while that in the window air space increases to 1400 times its value at low saturation levels!

Table 1. Flux density values at selected points

Excitation (kA-turn)	Left limb T	Upper limb T	Right limb T	Lower limb T	Window mT	Winding mT
0.5	0.99	0.99	0.99	0.99	0.10	0.50
100	2.57	2.45	2.35	2.45	140.	13.0

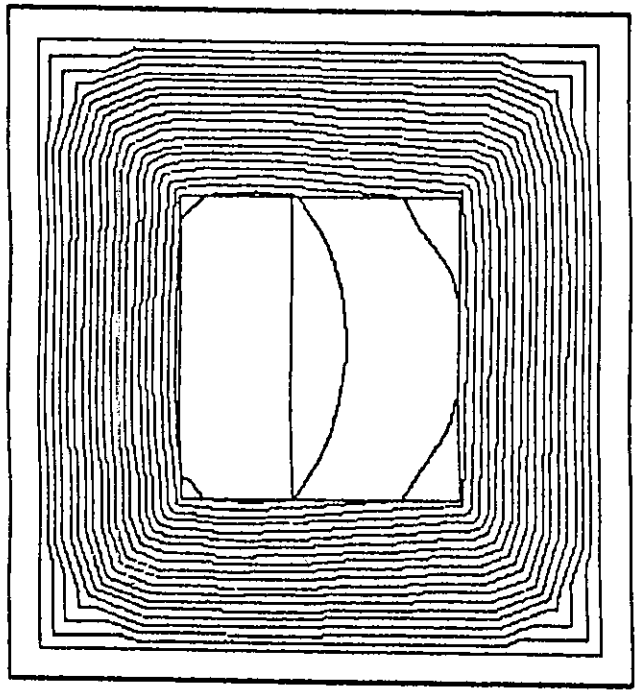

Figure 7. Flux plot for the simple inductor of Fig. 1, with a coil current to produce 100,000 ampere-turns.

Inductance values may be calculated from either solution, using the several techniques given above. Four values of inductance are given in Table 2 for either case: two computed using the flux linkage approach, two using the energy approach. The flux-linkage-based calculations include the approximate simple one,

$$L = \frac{n\phi}{I/n} = \frac{n^2 Z\,(A_r - A_L)}{I}, \tag{16}$$

as well as the more precise

$$L = \frac{\int_U \mathbf{A} \cdot \mathbf{J}\, dU}{i^2}. \tag{22}$$

The energy-based values are computed using first the full and correct stored energy expression,

$$L = \frac{2}{i^2} \int_U \int_0^B \mathbf{H} \cdot \mathbf{db}\, dU \tag{27}$$

and then also the quasi-linearized approximation

$$L = \frac{2}{i^2} \int \frac{B^2}{2\mu}\, dU \tag{30}$$

Results obtained by these four methods are shown in Table 2. They are expressed in microhenries, for an inductor with a winding assumed to contain exactly one massive turn; for realistic windings, the inductance values should be multiplied by n^2. For example, the correct inductance values for a 1000-turn winding are those shown in Table 2, multiplied by 10^6, that is, with Table 2 read as containing values in henries instead of microhenries. It is immediately evident that the two saturation levels are indeed extreme, their corresponding inductance values varying by about two orders of magnitude.

Table 2 is so arranged that the two most accurate values, those computed by using equations (21)–(22) and (27), occupy the two columns in the middle of the table. As can be seen from Table 2, the values are in very good agreement with each other for the unsaturated case.

The approximate energy-based calculation of equation (30) is seen to yield results in close agreement with those based on flux linkage totals. This phenomenon is perhaps not so surprising after all if one notes that the approximate energy calculation employs only the static permeability B/H, in other words, that the calculation relies heavily on one or a few points on the B-H curve; and so do the flux linkage methods. In essence, all three replace the B-H curve with a straight line that connects the origin with the current working point. The true nonlinear energy-based value, on the other hand, takes the curvature of the B-H characteristic fully into account.

Given that the different methods of inductance calculation can yield results as far apart as 0.974 and 2.556 for the same physical situation, which is correct? The answer, of course, is: both! It is important to keep in mind that the difference lies not so much in the computation methods

Table 2. Computed inductance values

Excitation (kA-turn)	L via flux linkages (microhenries)		L via stored energy (microhenries)	
kA-turn	(16)	(21)	(27)	(30)
0.5	198.96	198.96	181.50	198.96
100	2.556	2.556	0.974	2.556

as in the definitions of *inductance*. In fact, two different quantities are calculated and listed in Table 2, since the various definitions of what constitutes inductance coincide for the linear case—as indeed they do for the low-saturation case in Table 2—but are quite different for strongly saturated materials.

In practical design problems, difficulties are fortunately not as bad as they may seem from the above discussion. Most inductor designs have an air gap in the magnetic circuit to control the inductance under varying operating conditions. The inductor shown in Fig. 7 has been modified in Fig. 8 to include an air gap in one limb. The model also has to be changed to include some air space in and around the gap, in order to allow for the fringing flux which can now appear around the air gap. The gap introduced is 10 cm wide. The solution was recomputed for the two current levels described earlier and Fig. 8 shows the flux distribution for the saturated, 100 kA, case. The inductance values for the two current levels are shown in Table 3.

As can be seen, the variation in inductance between the saturated and unsaturated conditions is now considerably less, being of the order of 15% rather than the 80% without the air gap. The reason for the change is that

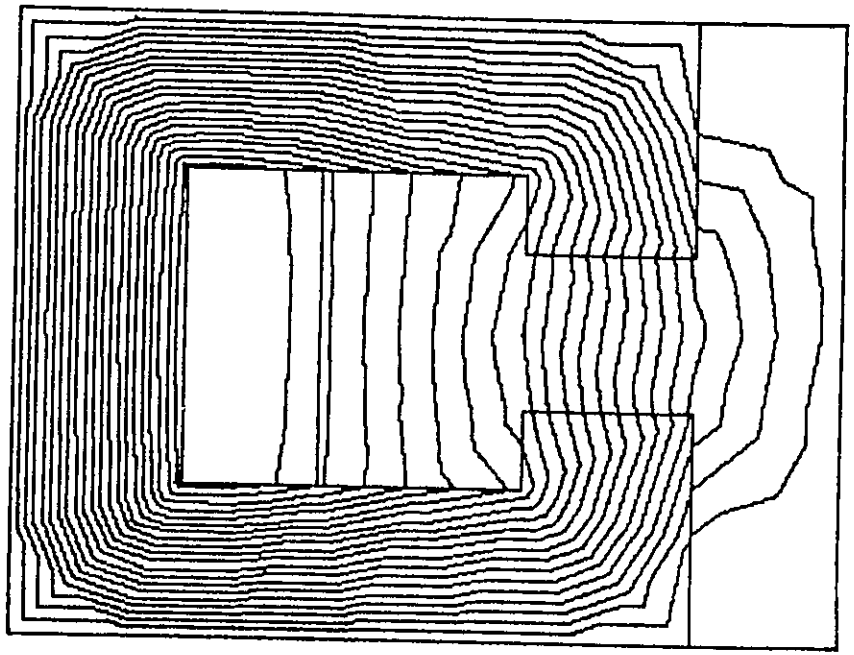

Figure 8. The inductor with air gap, subjected to an excitation of 100 kiloampere-turns.

Table 3. Computed inductance values with air gap

Excitation	L via flux linkages (microhenries)	L via stored energy (microhenries)	
kA-turn	(21)	(27)	(30)
0.5	2.750	2.750	2.750
100	2.177	1.823	2.177

the air gap dominates the problem in terms of where the energy is stored. The results detailed in Tables 1–3 were calculated on a finite element mesh having 76 nodes and 122 elements (for the problem in Fig. 8). The results of Table 3 were recomputed for a model having 543 nodes and 1039 elements but describing the same physical device. Considerable care was taken in the placement of the elements in the refined mesh in order to get a good solution for the flux distribution in the core and air gap of the inductor. The inductance values, corresponding to those in Table 3, for this highly refined model are shown in Table 4.

The results for the fine mesh give inductance values which vary by about 6% on average from those calculated with the coarse discretization. If a similar test is performed without the air gap present, the variation is about 1%. These results highlight the fact that the discretization chosen for a particular problem depends heavily on the results which are to be obtained–the refined mesh took approximately 20 times longer to solve than the coarse system and yet the inductance values, based on global energy storage calculations, vary by less than 6%. However, if the flux distribution is considered and the degree of saturation at various parts of the circuit is of prime importance, then the refined mesh produces considerably more accurate results because the coarse system does not allow any flux redistribution around the magnetic circuit.

A Transformer Design Problem

In the design of small transformers, it is important to determine the conventional transformer equivalent circuit parameters. Not only are they significant as evaluation criteria, but they often appear as key points in

Table 4. Computed inductance values with air gap–refined mesh

Excitation	L via flux linkages (microhenries)	L via stored energy (microhenries)	
kA-turn	(21)	(27)	(30)
0.5	2.938	2.938	2.938
100	2.237	1.820	2.358

customer specifications. These parameters are experimentally determinable, and analytically calculable from field solutions in a manner which resembles that used for the simple inductor example above. However, since the transformer is a multi-winding device, both the questions that need to be asked and their answers tend to be a little more complicated.

The Conventional Equivalent Circuit

Many small transformers are constructed by placing a pair of windings around the center leg of a magnetic core, as indicated in Fig. 9(a). The windings are placed over each other, usually with the primary winding

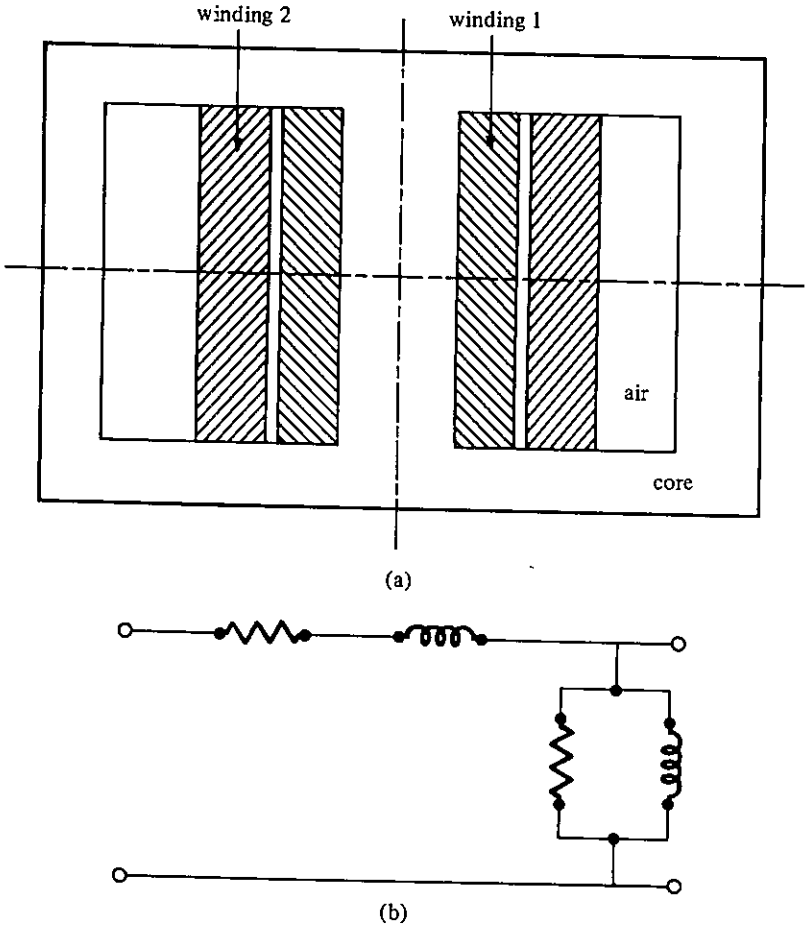

Figure 9. (a) A small two-winding transformer, with (b) its conventional equivalent circuit representation.

nearer the core center. The core is a built-up stack of E- and I-shaped stamped steel laminations, laid with the E and I pieces alternating in orientation, so that a closed but gapless magnetic structure is achieved. This structure resembles the simple inductor of Fig. 1, except for two details: the transformer comprises two windings while the inductor has one, and the inductor frequently (though not always) incorporates one or more air gaps in its magnetic path.

The transformer is structurally similar to a simple inductor, and the two devices are analyzed by techniques similar in principle. The assumption is once again made that very little leakage flux escapes the core, practically all the flux being confined to either the core itself and to the window space which contains the windings. End leakage will be ignored in a first analysis; some compensation for this approximation is possible by giving the air in the window space a relative permeability a little higher than unity. In a two-dimensional representation of the transformer, there are two symmetry planes, as indicated in Fig. 9(a); only one-quarter of the transformer need therefore be given an explicit representation for analysis.

Like many other electromagnetic devices, transformers are very frequently used as component parts in large, complex systems. They are commonly described by equivalent circuits, so that circuit descriptions of whole systems can be composed out of the equivalent circuits of their various component devices. Several equivalent circuits are available for transformers, all of roughly similar complexity. The most common equivalent circuit representation comprises two parts: a perfect (lossless and leakage-free) transformer with the correct turns ratio and the T-circuit shown in Fig. 9(b). The latter contains the circuit parameters which set the real device apart from an ideal transformer.

One reason why the T-shaped transformer equivalent circuit of Fig. 9(b) is in widespread use is that its components correspond fairly closely to the various physical phenomena in the transformer. The reactance in the vertical leg is essentially similar to the single reactance in Fig. 2: it accounts for the energy stored in the iron core. However, in the inductor the stored energy associated with leakage flux could be lumped in with the energy stored in the core; in the transformer it cannot, because there do exist measurements affected by the one and not the other. The series resistances in the horizontal legs of the T are the same wire resistances as in the inductor equivalent circuit of Fig. 2 above; the difference is that there are two windings and hence two resistances, while the inductor has only one. The shunt resistance in the vertical leg is a reasonably good representation of the power loss in the iron core. Identifying equivalent circuit elements with physical phenomena in this fashion is not strictly accurate; but the error involved is minor in most cases, because the numerical values of the vertical and horizontal branch circuit components in Fig. 9(b) differ by at least one, and occasionally as much as three, orders of magnitude.

The ideal transformer shown in Fig. 9(b) is usually omitted in analysis of systems containing the transformer. It is usual among transformer engineers to refer all quantities to the source-side winding, and to speak of the equivalent circuit as if it did not contain an ideal transformer or, what is the same thing, as if the turns ratio of the transformer were exactly unity. The voltages and currents that appear in the equivalent networks must then be scaled up and down, respectively, by the turns ratio to bring them into accord with the actual values. The minor inconvenience of voltage, current, and impedance referral, however, is richly compensated by the simplification that results when writing and solving circuit equations. This convention will be adhered to in what follows.

Short-Circuit Parameters

In the laboratory, the leakage reactance and winding resistance of the transformer are measured by means of the so-called *short-circuit* test. In this test, the secondary (load) terminals of the transformer are short-circuited, and the primary is fed with a low-voltage source so as to draw the rated current. Viewed in terms of the equivalent circuit of Fig. 9(b), the test arrangement amounts to short-circuiting the right-hand terminal pair. Very little current then flows in the shunt (magnetizing) branch of the circuit, since the impedance of this circuit branch is normally very high; typically, the shunt branch current might amount to 1–5% of the total. Primary-side measurements of voltage, current, and power therefore relate to the leakage reactance and winding resistance only; indeed, they provide a reasonable way of measuring these quantities.

Simulation of the classical short-circuit test is possible using a CAD system, but a better idea is to substitute another test, which might be called the *bucking test*; it is more accurate but much more difficult to perform in the laboratory. In the classical short-circuit test, a sinusoidal current source is connected to the left-hand terminals of Fig. 9(b) and the right-hand terminals are short-circuited. In the simulated test, two current sources are employed, one connected to the left and one to the right terminal pair. The two sources are identical, so that the current flowing into the left terminals is exactly equal to the current flowing out at the right. The magnetizing (vertical) branch of the equivalent circuit must then carry exactly zero current in the bucking test, in contrast to the conventional short-circuit test in which its current is approximately zero. For determining the horizontal branch parameters of Fig. 9(b), the bucking test is clearly superior because it eliminates the magnetizing branch completely. It is easy to implement in a CAD system, where it is only necessary to prescribe equal currents; in the laboratory, on the other hand, it is perfectly feasible in principle but very hard to carry out in practice. In the transformer core and window, such a test produces the flux distribution shown in Fig. 10. It will be seen that all the flux lines link at least part of one winding or the other. There are no flux lines which close in the iron

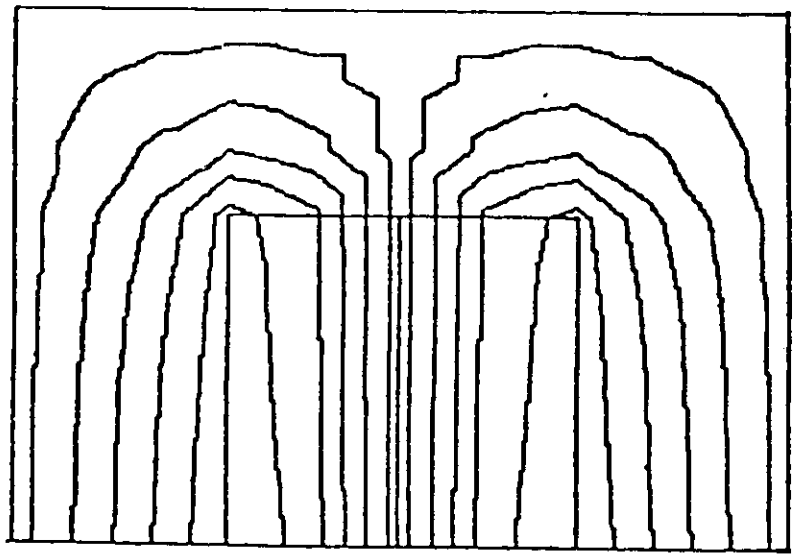

Figure 10. Flux plot for transformer on short circuit.

alone, showing that the two winding currents are exactly equal, with no net magnetomotive force applied to the core. To put the matter another way, there is zero net flux linkage of the secondary and primary windings.

The leakage inductances of the transformer can be determined by calculating the total stored energy associated with the solution of Fig. 9(b), using the technique of equation (27). Alternatively, the flux linkage technique of equation (21) may be applied; there should be no difference in the results, since almost all the stored energy will reside in the air. The flux density in the iron core is invariably very low at short circuit in a well designed transformer, and the energy stored in the iron core is therefore tiny compared to the energy stored in the leakage field in and around the windings. In fact, the short-circuit test simulation is very frequently performed on the assumption of infinite iron permeability, so that the energy stored in the core is not taken into account at all.

When short-circuit tests are performed in the laboratory, only the combined reactance value of the two windings (referred to the primary) is determined. It is conventional to apportion half the leakage inductance to the primary, half to the secondary winding. This apportionment may seem arbitrary but is often unavoidable. The single measurement made in the short-circuit test can only determine a single reactance, the combined total of primary and secondary leakages; separating them requires at least one further laboratory experiment. But the CAD system user, unlike the laboratory experimenter, obtains a full field solution from the (simulated) experiment, not merely one or two terminal values. He is therefore able to carry out additional simulated measurements by performing further

mathematical manipulations on the fields; furthermore, he is free to devise simulated measuring instruments which would be very difficult to realize in the laboratory. This flexibility permits separation of primary from secondary leakage inductances in cases where the windings are not symmetrically disposed and where a simple half-and-half split may not be appropriate. It is only necessary to write the total stored energy in the form

$$W = \frac{1}{2}\int_{S_p} A\,J\,dS + \frac{1}{2}\int_{S_s} A\,J\,dS, \tag{31}$$

where S_p is the cross-sectional area of the primary winding, S_s that of the secondary winding. In accordance with equation (20), this energy may be written in terms of the primary and secondary flux linkages $(n\phi)_p$ and $(n\phi)_s$ as

$$W = \frac{1}{2}\left[(n\phi)_p + (n\phi)_s\right] i. \tag{32}$$

Since no mutual flux linkages are shared between primary and secondary in the bucking test, it is proper to rewrite the energy W in terms of the primary and secondary leakage inductances L_p and L_s as

$$W = \frac{1}{2}\left[L_p i + L_s i\right] \tag{33}$$

showing that the individual leakage inductances are separately computable. The primary leakage inductance is thus, by (22),

$$L_p = \frac{1}{i^2}\int_{S_p} A\,J\,dS \tag{34}$$

and a similar expression holds for the secondary. The leakage inductances are separable, in other words, by measuring the primary and secondary flux linkages separately, a task not particularly difficult in the simulated bucking test but impossible in laboratory practice.

It should be noted that the winding resistances in the equivalent circuit cannot be calculated directly from the magnetic field distribution. As with the simple inductor, a value can be deduced from the length of wire used and its conductivity.

Magnetizing Inductance

The magnetizing reactance and core loss resistance of transformers are determined in the laboratory by performing an *open-circuit* test. In this test the secondary terminals are left open-circuited and the primary is excited at rated voltage. Measurements of the primary voltage, current, and power, together with the results from the short-circuit test, allow the

calculation of the magnetizing reactance and the core loss resistance. Since the magnetizing impedance of a well designed transformer is very considerably greater than the shunt impedances, the results of the open-circuit test are quite often used to calculate the magnetizing impedance directly, ignoring any correction for the leakage and winding resistance.

The open-circuit test is not only easy to perform in the laboratory, because the transformer is run at rated voltage without load, but it is also easy to model analytically. Leaving the secondary winding open-circuited simply requires that it is ignored in setting up the problem! In fact this test is identical to the simulation, already treated, of a simple inductor, since only one winding is taken into account in the modelling. With the primary winding on the left, the transformer of Fig. 9(a) exhibits the open-circuit flux distribution shown in Fig. 11.

Since the open-circuit test is performed at rated voltage, not at rated current, its simulation involves a few difficulties. First, most present-day CAD systems work best with prescribed currents rather than voltages. Secondly, the *magnetizing reactance* is ill defined, since most transformers are operated at sufficiently high saturation levels to cause the magnetizing current to contain a significant proportion of harmonics. With sinusoidal applied voltage, but nonsinusoidal current, the inductance—and hence the reactance—of the primary winding can be defined in several ways. This problem does not arise in the leakage calculations associated with the short-circuit test, since much of the magnetically stored energy is stored in air. The contrary is true for the open-circuit test: almost all the magnetic energy is stored in the iron core, which is after all nonlinear.

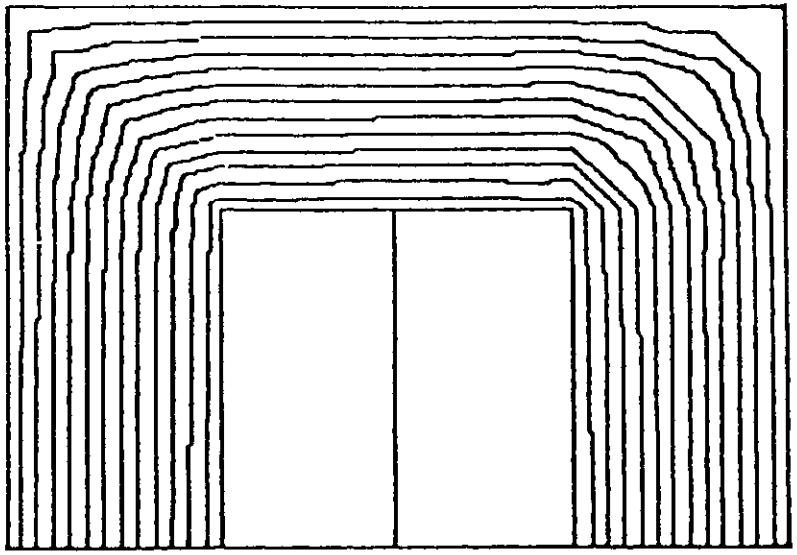

Figure 11. Flux plot for transformer on open circuit.

Like many other simulations, the open-circuit test can be carried out in a careful and fairly precise way, or it can be done quickly but less accurately. The slow and precise way proceeds as follows. The range of instantaneous magnetizing currents, up to the largest peak current likely to be encountered, is determined by any reasonable means (which may include an educated guess). A set C of N current values, which may number about ten, is chosen so as to span the range,

$$C = \{i_k,\ k = 1, \ldots, N\}. \tag{35}$$

The open-circuit field problem is solved for each one of these values, thereby producing a set A of vector potential distributions $A_k(x,y)$,

$$A = \{A_k(x,y),\ k = 1, \ldots, N\}. \tag{36}$$

The primary flux linkages are next computed for each solution, using equation (20),

$$(n\phi)_k = \frac{1}{i_k} \int_{S_p} A\ J\ dS. \tag{37}$$

Since every flux linkage value $(n\phi)_k$ in this set corresponds to a specific value i_k, the two sets taken together describe a function which assigns a specific flux linkage value to each current value, and vice versa:

$$(n\phi)_k = n\phi(i_k). \tag{38}$$

The resulting curve of flux linkages against current will roughly resemble a B-H curve. However, it will not be proportional to the B-H curve of the core material, because the flux distribution in the core actually changes as the material saturates.

When an open-circuit test is carried out in the laboratory, the primary winding is excited with a sinusoidal voltage, say

$$e(t) = E \cos \omega t. \tag{39}$$

By Faraday's law, the primary flux linkages must therefore vary in a sinusoidal fashion,

$$n\phi = -(1/\omega)\ E \sin \omega t, \tag{40}$$

so that the sequence of time instants t_k may be determined at which the tabulated values of flux linkage are reached during a quarter cycle:

$$t_k = -(1/\omega) \arcsin [-\omega n\phi(i_k)/E]. \tag{41}$$

But this relationship creates a set T of time values,

$$T = \{t_k,\ k = 1, \ldots, N\}, \tag{42}$$

which correspond to the current values C of equation (35). If these current values are plotted against time, the magnetizing current waveform is obtained. It may be written as a Fourier series, as detailed further below. Because the total current waveform has quarter-cycle symmetry and begins at $i(0) = 0$, the series contains odd sine terms only:

$$i(t) = \sum_{m=1}^{\infty} I_m \sin (2m-1)\omega t. \tag{43}$$

The magnetizing inductance is then calculated by taking the fundamental components only,

$$L_{mag} = \frac{E}{\omega I_1}. \tag{44}$$

The fundamental-frequency component definition of inductance is probably the best to choose in this application, because a main use of small transformer equivalent circuits is in the calculation of fundamental-frequency currents in systems containing transformers. Getting the fundamental components right is therefore likely to be important, more so than, say, finding the stored energies with high accuracy. The same cannot be said, however, of applications involving transients. Most transient phenomena involve transfer of flux linkages or stored energies and therefore should use inductances based on total quantities, rather than specific Fourier series components.

A quick but less accurate estimate of the magnetizing reactance can be obtained by means of a single static solution. The primary winding current is set at the root-mean-square value expected, and its flux linkages are computed from the field solution. A value of terminal voltage is then obtained, on the assumption that the value obtained for flux linkages is the root-mean-square value also. If the resulting voltage value is within a few percent of the rated voltage, the magnetizing reactance value is probably accurate to a few percent also.

It may be worth noting that no attempt has been made in either of the above methods to allow for the primary leakage reactance. Since both calculations actually deal with *terminal* voltage, the reactance values obtained represent not the magnetizing reactance, but rather the sum of magnetizing reactance and leakage reactance. To obtain a value for the magnetizing reactance alone, it suffices to subtract the leakage reactance value obtained from the bucking test simulation described above.

Harmonic Analysis

The harmonic content of magnetizing current is a quantity of considerable interest to transformer designers as well as to the system engineer using transformers as circuit components. If hysteresis is absent, as has been

assumed in the above, then the magnetizing current will have quarter-cycle symmetry; that is,

$$i(\omega t + \pi) = -i(\omega t). \tag{45}$$

The Fourier series of $i(\omega t)$ can then contain no odd terms; and

$$i(\omega t + \pi/2) = -i(-\omega t) \tag{46}$$

so there can be no cosine terms. The Fourier series of $i(\omega t)$ therefore has the form already alluded to,

$$i(t) = \sum_{m=1}^{\infty} I_m \sin (2m-1)\omega t. \tag{43}$$

In any real situation, the true current $i(t)$ cannot be known, since an infinite amount of data would be needed to determine all Fourier series components. Instead of (43), the practical analyst must remain content with the approximate current $i_M(t)$, which contains only M terms of the series,

$$i_M(t) = \sum_{m=1}^{M} I_m \sin (2m-1)\omega t. \tag{47}$$

The number M is of course limited to at most N, the number of data points contained in the sets of current and time values, equations (35) and (42). This restriction is not particularly severe, for there is rarely much interest in harmonics beyond the fifth or (in some three-phase systems) eleventh. That is to say, M is limited to 3, or 6 at most, while present CAD systems make it easy to compute solutions for, say, $N = 10$.

Since the number N of data points available is frequently larger than the number M of Fourier series terms required, it is best to compute the limited number of terms by a least-squares approximation. To do so, the series (47) is formally evaluated at each of the N time instants t_k,

$$i_M(t_k) = \sum_{m=1}^{M} I_m \sin (2m-1)\omega t_k. \tag{48}$$

The coefficients I_m are chosen so as to make the squared difference between left- and right-hand sides of (48), taken at all the N time instants, as small as possible. That is to say, the I_m are chosen so as to make first derivatives vanish:

$$\frac{\partial}{\partial I_n} \sum_{k=1}^{N} [i(t_k) - i_M(t_k)]^2 = 0. \tag{49}$$

Substituting (48) into (49) and differentiating, a set of simultaneous equations is obtained. For every value of m,

$$\sum_{n=1}^{M}\sum_{k=1}^{N}\sin(2n-1)\omega t_k\ \sin(2m-1)\omega t_k\ I_n$$

$$=\sum_{k=1}^{N}\sin(2m-1)\omega t_k\ i(t_k) \qquad (50)$$

must hold. These equations are easily recast in matrix form for numerical solution, noting that the inner summation (with index k) on the left-hand side collapses the trigonometric functions into a symmetric square matrix with indices n,m.

The above discussion assumes that hysteresis is absent. Indeed, no other assumption can at present be fruitfully made, for no general-purpose CAD systems now available can treat hysteretic behavior in solving field problems. The analysis given here is most accurate at high saturation levels, when the hysteresis loop width is small compared to the excursion in magnetic field.

Loss Estimation

The core loss resistance to be included in the transformer equivalent circuit may be estimated roughly from a single static field solution. The estimating technique is simple: the local power loss density p (usually given on a per unit mass basis) is weighted by the mass density $m(x,y)$ of the stacked laminations and integrated over the volume of magnetic material so as to obtain the total loss P. Thus

$$P = Z\int_{S_{iron}} p\ m(x,y)\ dS, \qquad (51)$$

where Z represents the stacking depth of the core iron, and S_{iron} is the area occupied by iron in the transverse (x-y) plane. The loss density $p(x,y)$ of course depends on the magnetic events at the point (x,y). It is usually expressed as a function of peak flux density,

$$p = p(B_p). \qquad (52)$$

The function $p(B_p)$ must be obtained either from the iron supplier or by direct experimental measurement. Most suppliers of magnetic materials provide core loss curves for sheet and strip stock, in watts per kilogram or pound as a function of the peak flux density. By referring to mass rather than volume, any required stacking factors are automatically taken care of by the mass density $m(x,y)$ in equation (51). Curves are usually available for sinusoidal flux density variation, at 50 Hz and 60 Hz for the thicker materials and 400 Hz for the thinner grades. Data for other frequencies or working conditions are furnished only infrequently.

Since the loss curves refer to peak flux density, the loss is estimated from a field solution obtained for the open-circuit test with the primary

winding carrying peak current. This solution is identical to one obtained for determining the magnetizing reactance. The potential values in it are differentiated, to yield the corresponding flux density B_k,

$$\mathbf{B}_k = \text{curl}\ (\mathbf{1}_z A_k). \qquad (53)$$

These represent peak values, since they are derived from the peak values of potential. The loss density curve supplied by the manufacturer is then used to determine the loss density everywhere in the core, as in (52), and the total loss is obtained by integration over the whole core, as in equation (51).

The core loss resistance R to be included in the transformer equivalent circuit may be calculated from the estimated core loss P. By simple circuit theory, the value of R is

$$R = \frac{E_{rms}^2}{P}. \qquad (54)$$

If voltage drop in the primary resistance and leakage inductance is neglected—a reasonable approximation for the open-circuit test—then E_{rms} is the root-mean-square voltage at the primary terminals. Because the transformer may be assumed to be excited by a sinusoidal voltage, (54) may be rewritten in terms of the peak voltage E_p,

$$R = \frac{E_p^2}{2P}. \qquad (55)$$

But the peak terminal voltage is related to the primary flux linkages $n\phi$ by equation (40). Hence the value for core loss resistance is finally obtained as

$$R = \frac{\omega(n\phi)_p^2}{2P}. \qquad (56)$$

It should be noted that the loss as estimated using the standard core loss curves refers only to the loss incurred by alternating flux, not rotational flux density. Rotating fluxes arise in many devices through the redistribution of flux which arises from saturation. The distribution of fluxes in a transformer core, or indeed in any magnetic device, is clearly not the same at high and low saturation levels. As the exciting current is raised, regions of high flux density saturate first, crowding flux into areas of lower density. At high saturation levels, flux densities therefore tend to be much more uniform than at low total flux values, and the local direction of flux is not the same in the two cases. As a transformer on open-circuit test traverses the alternating current cycle, the flux density vector **B** at most points of the core iron not only varies in magnitude but also changes in direction as flux redistribution takes place. Hence the vector **B** rotates in

direction as its magnitude alternates; its tip generally traces an elliptic path. Iron loss curves, however, make no allowance for rotating components in the flux density. Estimates of loss are therefore usually less accurate than estimates of the corresponding magnetizing reactance.

Mutual Inductances

Many phenomena in electromagnetic devices are conveniently described in terms of mutual inductance values, and many designer-years are spent calculating values of mutual inductances. Ways of extracting mutual inductances from field solutions therefore merit more than passing mention.

Mutual Inductances and Energy

Since there are at least two reasonable bases on which to define self-inductance—stored energy and flux linkages—it should not be surprising that mutual inductances can be similarly defined.

One note of caution is perhaps in order: the computation of inductances, like the computation of almost anything else in a CAD system, should be undertaken only as a *simulation* of experimental measurements. A corollary is that no attempt should ever be made to calculate quantities which cannot be defined in terms of a physically feasible experiment. While this principle may seem obvious, it is very frequently violated by classical engineering electromagnetics. For example, the concept of *internal inductance* of a conductor is enshrined in textbooks and has been so for a hundred years. It is well known and calculable for a round wire, for which analytic solutions of the skin-effect problem exist. It is not calculable for other, more complicated shapes—not because of shortcomings in the available processes of calculation, but because the classical definition of internal inductance itself makes no sense for other conductor shapes. (It assumes that one flux line coincides with the conductor surface, partitioning all other flux lines into those *internal* to the conductor and those *external* to it. Such a separatrix flux line can only exist if the conductor is round.) Many of the classically useful concepts and quantities of electric machine engineering in particular are imprecise because their definitions contain built-in presuppositions about geometric shapes: examples are end-turn inductance, slot leakage inductance, and zig-zag leakage. If there is any doubt at all about some well established parameter, the one sure test is to invent a physical experiment for measuring the relevant quantity. In effect, the invention of such an experiment redefines the quantity in experimental terms valid for all cases. Even if the measurements are impractical to carry out, their simulation may be quite straightforward.

Defining mutual inductance in terms of stored energy is a secure procedure, for the stored energy is a quantity easily measured, at least in principle, by an integrating wattmeter. The mutual inductance in a two-winding system may be determined by a difference measurement involving two experiments, one with the two winding currents oriented to have their fluxes adding, the other with the fluxes bucking each other, as in Fig. 12. It is of course not necessary to know at the outset which is which; it is only essential to conduct two experiments, with connections to one winding reversed relative to their orientation in the other experiment, i.e., to have the winding currents i_1 and i_2 in one case, i_1 and $-i_2$ in the other. Altering the order of experiments will make the mutual inductance have either a positive or a negative sign.

Let W_1 and W_2 be the stored energies in the two experimental cases. By simple circuit theory, they are related to the self-inductances L_{11}, L_{22} and the mutual inductances M_{12}, M_{21} through the relationships

$$W_1 = \frac{1}{2}\,(L_{11}\,i_1^2 + M_{12}\,i_1\,i_2 + M_{21}\,i_2\,i_1 + L_{22}\,i_2^2) \tag{57}$$

and

$$W_2 = \frac{1}{2}\,(L_{11}\,i_1^2 - M_{12}\,i_1\,i_2 - M_{21}\,i_2\,i_1 + L_{22}\,i_2^2). \tag{58}$$

By subtracting, the difference in stored energies is obtained. There immediately results

$$M_{12} + M_{21} = \frac{W_1 - W_2}{i_1\,i_2}, \tag{59}$$

and since mutual inductances are reciprocal, $M_{12} = M_{21} = M$,

$$M = \frac{W_1 - W_2}{2\,i_1\,i_2}. \tag{60}$$

The procedure as given above requires measurement or computation of stored energy, followed by subtraction to find the energy difference. If the energies W_1 and W_2 are not very different, the result can be considerably in error. This *subtraction of elephants* difficulty arises no matter whether the data are obtained by direct physical experimentation or by computer simulation: if the difference looks like small mice, the result is of doubtful accuracy. To be more precise: as many significant figures will be lost in the subtraction as there are similar leading digits in W_1 and W_2. It is always well to inspect results, including intermediate results, from time to time as calculations proceed and to bear this potential difficulty in mind.

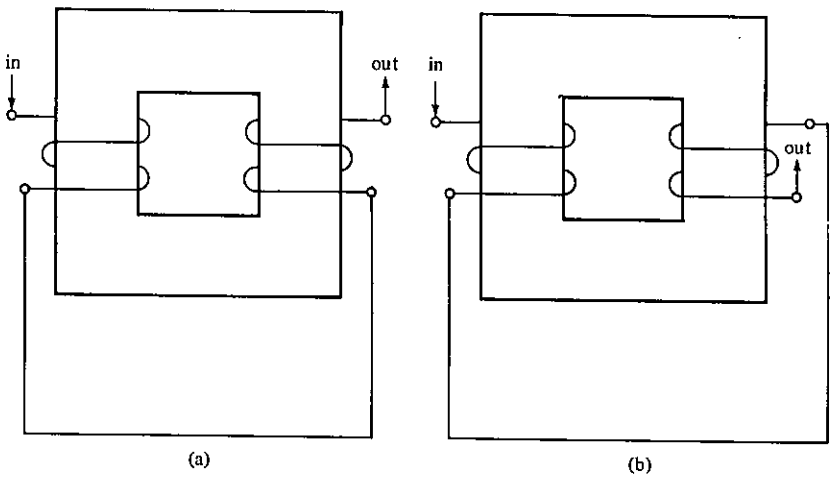

Figure 12. Two windings connected together for purposes of mutual inductance determination: (a) series aiding, (b) bucking.

The stored energy calculations required in the above simulation may be carried out by exactly the same techniques as followed in the simple inductor case. The appropriate equations to use are generally (25) and (26), but circumstances may allow simplifications arising from magnetic linearity or geometric shape.

Mutual Inductance and Flux Linkage

Like self-inductances, mutual inductances may be defined in terms of flux linkages. Such a definition leads to quite direct and simple calculations, possibly easier than the energy-based approach. In linear problems, the two definitions again coincide, so the results obtained are identical to within the error inherent in numerical computation; the choice of method is therefore a matter of convenience. In nonlinear cases, however, flux linkages and stored energy must lead to different results because they reflect two different definitions of inductance.

In terms of flux linkages, the mutual inductance M_{12} is defined as *the flux linkages of winding 1 caused by the current in winding 2*, divided by the current causing them:

$$M_{12} = \frac{(n\phi)_{12}}{i^2}, \tag{61}$$

As already discussed in connection with equation (20), flux linkages are best determined with reference to currents, so that any nonuniformity in the winding density is automatically introduced as a weighting factor. To

determine the flux linkages in winding 1, it is therefore best to imagine that a very small current i_1 is made to flow in that winding, one small enough to cause negligible change in vector potential. This small current will nevertheless cause a current density J_1 to exist in winding 1. Generalizing on equation (20) slightly, the required flux linkage may then be calculated as

$$(n\phi)_{12} = \int A \left[\frac{J_1}{i^2}\right] dS. \tag{62}$$

Combining (61) and (62), the mutual inductance is thus

$$M_{12} = \frac{1}{i_2} \int A \left[\frac{J_1}{i_1}\right] dS. \tag{63}$$

The region of integration here may be the cross-sectional area of winding 1, or any larger area. A larger area is perfectly acceptable, since the current density J_1 will vanish in any portion not actually occupied by winding 1. No extra contribution will accrue to the integral even if the area is too large. As a practical matter, the area is frequently chosen to include the entire problem region, thereby reducing work in defining the region of integration.

When the calculation of (63) is actually carried out in a CAD system, there is no need to keep i_1 small, provided the vector potential A as used in the calculation is the vector potential that results from solving the field problem with winding 2 excited, but without current in winding 1. In other words, there is no need to risk numerical error by choosing i_1 to be very tiny; any convenient value will do.

Force Calculations

Many of the uses to which magnetic devices are put may be classified as electromechanical; that is, they are used to convert energy between electrical and mechanical forms. Indeed, mechanical force production is the major reason for the existence of devices such as actuators and electric motors. Consequently, the end product of a magnetic field analysis may well be the evaluation of the mechanical force produced by the device and its variation with changes in excitation or position. The principles involved are introduced and discussed here with reference to illustrative examples.

A Magnetic Bearing

Permanent magnet bearings are employed in some watt-hour meters. Such meters are in essence small electric motors designed to measure electricity

consumption, which must have very consistent performance in order to meet legal and commercial needs. The usual type of meter includes a rotating aluminum disk, driven by eddy currents caused by two excitation coils—one coil intended to measure instantaneous current, the other instantaneous voltage. Integration over time is performed mechanically, by having the disk actuate a rotating counter. For satisfactory operation the rotating disk mechanism should exhibit little friction, and the bearing characteristics should show negligible change with diurnal and seasonal weather variations as well as over the lifetime of the device, which is likely to be measured in decades.

To meet stability and consistency requirements, one type of watt-hour meter incorporates a permanent magnet suspension system with an auto-compensating magnetic shunt to allow for the temperature variation of permanent magnets. Such a configuration is represented in Fig. 13. A compensating shunt (flux diverter) bypasses part of the magnetic flux of the permanent magnet. With an appropriate choice of materials and dimensions, the amount of flux diverted can be made to vary with temperature in the sense opposite to the permanent magnet material. For example, an increase in total flux, resulting from temperature variations, is compensated for by an increase in the fraction of flux diverted. In this way, the suspension height can be kept practically constant, and the magnetic bearing can produce a watt-hour meter with precision that rivals mechanical suspension, but without the wear that results from mechanical contact.

For the magnetic bearing of Fig. 13, the field at nominal design temperature is shown in Fig. 14. The analysis shown was carried out using explicit current-carrying coils rather than intrinsic permanent magnet models, and the coil sides (which actually are empty air space) are clearly visible in Fig. 14. The flux diverter accounts for very little leakage at the design temperature, which corresponds to summertime operation.

Although it is necessary to obtain lifting force values as a final analytic result, preliminary examination of possible designs may require no more than simple visual inspection of field distributions. Often an extremely

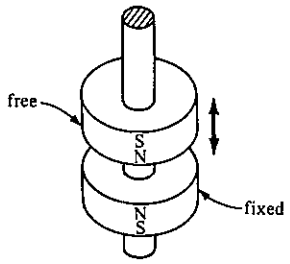

Figure 13. A simple vertical-shaft magnetic bearing used in a watt-hour meter, based on permanent magnets.

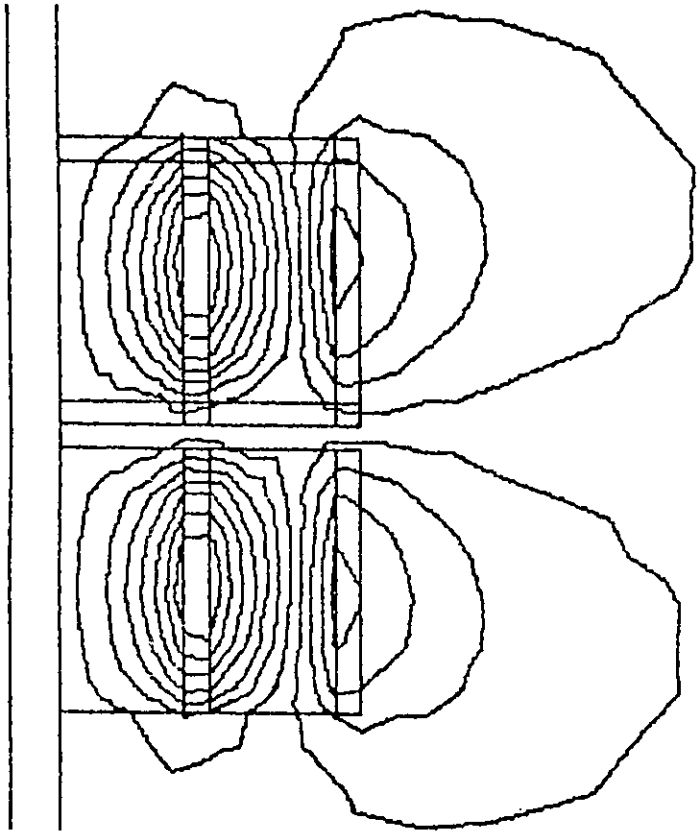

Figure 14. Magnetic field (flux lines) of the permanent magnet bearing shown in Fig. 13.

useful indication of the behavior of a magnetic system can be obtained in this qualitative way. In Fig. 14 the flux lines are crowded together in the air gap between the two magnets and are directed horizontally. This strong tangential component of field in the air gap suggests that there is a substantial repulsion force between the two parts of the bearing. The magnitude of the force can be gauged very roughly by observing both the density of the lines and how horizontal they appear to be.

The actual value of the force in a magnetic system may be determined in several ways. Since there is in reality only one force, all the ways should produce the same result, provided there is no computational error. However, because of the numerical approximations which have been made in the solution of the field equations and the distribution of the force itself, the several methods may not give identical answers. In fact, the most appropriate method to use may well depend on the device being analyzed. Some of the more common approaches and their pitfalls are outlined below.

The Method of Virtual Work

The force exerted on a body may be evaluated by determining the work done when it is slightly displaced from its rest position. In the absence of frictional losses, this work must equal the change in the energy stored in the entire electromechanical system or device. Thus, if the stored energy is evaluated for the device in two positions separated by a small displacement, then the difference in energy divided by the distance will give a value for the force. Suppose, for example, that the stored energy is W_1 in position 1 and W_2 in position 2, the two positions being separated by a displacement x_{12}. The force required is then given by

$$F = \frac{W_2 - W_1}{x_{12}}. \tag{64}$$

This approach to force calculation is widely employed in mechanics, where it is known as the *method of virtual work*. It evidently relies on the assumption that the force does not change significantly during the motion and thus is valid for small displacements only. In the completely general case where displacements may take place in various directions, the force is a vector quantity, given by

$$\mathbf{F} = \text{grad } W(x), \tag{65}$$

where $W(x)$ is the stored energy, viewed as a position function of the vector displacement x. Equation (64) may be regarded as a special case of (65), for one-dimensional movements.

The lifting force for the magnetic bearing described here was computed by the method of virtual work. Since lift force is the desired quantity, the upper magnet of the bearing was displaced so as to halve the vertical air gap between magnets, from 2.0 mm to 1.0 mm. (By way of comparison, the bearing magnet diameter is 18 mm.) The force is likely to change more or less linearly with distance, in view of the large bearing magnet diameter, so that the value obtained is likely to represent a good estimate for the actual force at an air gap of 1.5 mm. Values for the stored energies in the two positions and the force are as follows:

Stored energy with 2 mm gap	=	2.9476 mJ
Stored energy with 1 mm gap	=	2.8892 mJ
Change in stored energy	=	0.0584 mJ

and the force follows directly as

$$\text{Force} = \frac{0.0584 \text{ mJ}}{1.000 \text{ mm}} = 0.0584 \text{ N}.$$

It may be noted that the energy difference in this case is about 2% of the stored energy itself. In other words, nearly two leading significant figures

are alike in the two energies, and nearly two significant figures will therefore be lost to roundoff error in the subtraction. If the energies themselves can be relied on to about five figures, then three-figure accuracy of the force is all that can be hoped for.

The figures above are fairly typical for a calculation based on the virtual work method, although in this case a comparatively large displacement can be made. In systems where only a very small displacement is possible and the forces are likely to be small, the change in stored energy may be a fraction of a percent of the total energy. The difficulty of *subtraction of elephants* then arises once again, just as it did with equations (59)–(60). There are two ways it can be alleviated: wisdom in solution of the field problems and wisdom in differentiation.

When forces (or mutual inductances!) are computed by subtracting energies, accuracy can be considerably enhanced if the error inherent in the two energies themselves is similar. If, say, W_1 as computed contains some error e_1, as $W_1 + e_1$, and W_2 is similarly computed as $W_2 + e_2$, then obviously

$$(W_2 + e_2) - (W_1 + e_1) = (W_2 - W_1) + (e_2 - e_1). \tag{66}$$

The error in the energy difference will certainly be no larger than the error in either energy; indeed it will generally be smaller, if it can be guaranteed that both e_2 and e_1 have the same sign. Such a guarantee is easily furnished if the field problems underlying W_1 and W_2 are both solved using energy-minimizing finite element methods. Fortunately, such methods are by far the most popular ones in present-day CAD systems. The best results will clearly be obtained if e_2 and e_1 are not only of the same sign but of similar magnitude. Their magnitudes are very often mainly affected by the discretization error that arises from the finite element mesh. If W_1 and W_2 are solutions computed on the same finite element mesh, or very similar meshes, errors are likely to be similar and the energy difference therefore much more accurate than the simple rules of thumb would indicate.

The results obtainable by virtual work calculations can sometimes be improved by employing not the simple equation (64), but its more general version (65). Instead of computing energy at the two ends of the displacement x_{12}, a set of several stored energies can be found at several points along the displacement vector. If a smooth curve is fitted to these several values, and differentiation performed according to (65), considerable error smoothing can be achieved at the price of increased computation.

Maxwell Stresses

The second most common approach to determining electromechanical forces is that known as the *Maxwell stress tensor* method. In contrast to the virtual work technique, which employs a volume integral to determine the stored energy, the Maxwell stress approach computes local stress at all

points of a bounding surface, then sums the local stresses by means of a surface integral to find the net force.

The Maxwell stress tensor method may be derived from the elementary force density expression which relates the force density vector **f** to the flux density **B** and the current density **J** by

$$\mathbf{f} = \mathbf{J} \times \mathbf{B}. \tag{67}$$

This expression is derived in theoretical electromagnetics from the fundamental force relationship between two moving charges, and represents the magnetic portion of the *Lorentz force*.

The expression given in equation (67) describes a force density vector, which possesses components in each of the coordinate directions and has dimensions of force per unit volume. Thus the force on a body in a particular direction may be found by integrating the appropriate component of the force vector over the entire volume. It is common, however, to reduce the volume integral described above to a surface integral in order to simplify the overall force calculation. If a substitution is made for **J** in equation (67) by using Maxwell's equations, the force expression becomes

$$\mathbf{f} = \nu\, \mathbf{B} \times \operatorname{curl} \mathbf{B}. \tag{68}$$

The left-hand side of this equation may be expanded into three components. The x directed component is typical and has the following form:

$$f_x = \nu B_z \frac{\partial B_x}{\partial z} - \nu B_z \frac{\partial B_z}{\partial x} - \nu B_y \frac{\partial B_y}{\partial x} + \nu B_y \frac{\partial B_x}{\partial y}. \tag{69}$$

If a term $\nu B_x(\partial B_x/\partial x)$ is simultaneously added to and subtracted from equation (69), and the identity

$$\frac{\partial}{\partial x}(B_x)^2 = 2B_x \frac{\partial B_x}{\partial x} \tag{70}$$

is used, then the force component becomes

$$fx = \nu \left[\frac{1}{2} \frac{\partial}{\partial x}(B_x)^2 + B_z \frac{\partial B_x}{\partial z} + B_y \frac{\partial B_x}{\partial y} - \frac{1}{2} \frac{\partial}{\partial x}(B_x^2 + B_y^2 + B_z^2) \right]. \tag{71}$$

Some further manipulation gives

$$f_x = \left[\frac{\partial}{\partial x}\left(B_x^2 - \frac{1}{2} \mid B \mid^2\right) + \frac{\partial}{\partial y}(B_x B_y) + \frac{\partial}{\partial z}(B_x B_z) - B_x \operatorname{div} \mathbf{B} \right]. \tag{72}$$

Since div **B** = 0, the last term disappears and the remaining expression may be recognized as the divergence of a vector $\mathbf{f}_x$, whose components are

$$f_{xx} = \nu\left[B_x^2 - \frac{1}{2} \mid B \mid^2\right], \tag{73}$$

$$f_{xy} = \nu B_x B_y,$$

$$f_{xz} = \nu B_x B_z. \tag{73}$$

A similar development holds for each of the other force components ($\mathbf{f}_y$ and $\mathbf{f}_z$). Thus the force expression, equation (67), may be written as the divergence of some *tensor* **T**:

$$\mathbf{f} = \int_U \mathbf{J} \times \mathbf{B}\, dU = \frac{1}{\mu} \int_U \operatorname{div} \mathbf{T}\, dU. \tag{74}$$

Making use of the divergence theorem, this volume integral may be reduced to a surface integral. Then the force becomes

$$\mathbf{f} = \frac{1}{\mu} \oint_S \mathbf{T} \cdot \mathbf{dS}, \tag{75}$$

where the surface vector **dS** is taken as the outward normal on S.

The tensor **T** defined above has the dimensions of stress and is commonly known as the *second Maxwell stress tensor* or the *magnetic stress tensor*. A more complete derivation of this tensor may be obtained by considering both the magnetic and electric components of the Lorentz force. In this case, an electric stress tensor may be defined in addition to the magnetic one described here. The electric stress tensor has essentially the same form as its magnetic counterpart with **B** replaced by **E** and reluctivity replaced by permittivity. This form is useful for force calculations in electrostatic fields. Furthermore, when both electric and magnetic forces are considered together a link between the two fields appears which may be recognized as the *Poynting vector* representing power flow within the volume.

The complete magnetic stress tensor **T**, written out in full, has the following form:

$$\mathbf{T} = \begin{bmatrix} \left(B_x^2 - \frac{1}{2} \mid B \mid^2\right) & B_x B_y & B_x B_z \\ B_y B_x & \left(B_y^2 - \frac{1}{2} \mid B \mid^2\right) & B_y B_z \\ B_z B_x & B_z B_y & \left(B_z^2 - \frac{1}{2} \mid B \mid^2\right) \end{bmatrix}. \tag{76}$$

Equation (76) provides the local values of all components of magnetic stress along each of the coordinate axes. The rigid body force acting on an object is obtained by integrating these components over its bounding surface, as in equation (75).

It should be remembered that each stress component is in fact a vector and thus the dot product with the normal should be applied to each term of the appropriate horizontal row in turn. In two dimensions the 3×3 matrix reduces to 2×2, i.e. the top left hand corner of the tensor, and the surface integral becomes a contour integral. Thus, if the surface is parallel to the x-coordinate axis in a two- dimensional system, the x-directed component of force is given by B_xB_y and the y directed force is $(B_x^2 - |B|^2/2)$. In the case of two blocks of iron facing each other, such as might be encountered in an electromagnet where all the flux is directed between the poles, there is no sideways force because B_x is zero. The attractive force between the two poles is then given by $B_y^2/(2\mu_0)$, which may be recognized as the conventional expression for the force between two magnetic poles.

To determine the forces on a rigid body, the surface of integration S should of course encompass the body. In principle, S should be the surface of the body itself. In computational practice, it is often found convenient to place this surface in the air region surrounding the machine part or other object on which the force is to be found. In effect, a "piece of air" is therewith attached to the rigid body, as if it were a solid. Since the air carries no currents and has no magnetic properties different from free space, there is no harm in this hypothetical attachment.

Although the component values of equation (76) have units of stress, they do not necessarily give correct local stress values. The mathematical reason is that the divergence theorem, on which their derivation is based, has a meaning only for complete closed surface integrals. Physically, local stresses may be statically indeterminate and therefore not available at all from such a global calculation. However, their closed surface integral is guaranteed to have a physical meaning and to represent total force correctly. Of course, the surface may be closed only in a restricted sense. For example, in a periodic structure the surface need only embrace one period, the remainder of the structure being taken care of implicitly.

The expressions given above may be rewritten in terms of the normal and tangential components of flux density at each point on the surface. In two-dimensional problems the calculation of force requires the determination of the normal and tangential values of flux density at each point along the contour. Difficulties may be encountered as the result of numerical cancellation (the *subtraction of elephants* problem once again!) when Maxwell stress calculations are applied to finite element solutions. First of all, flux density components are obtained from potential solutions by differentiation, a process which commonly emphasizes errors. If the element discretization of the air region is relatively coarse, the error in evaluating the pointwise stresses may be considerable. Of course, this problem may be overcome by increasing the number of elements. Secondly, if the contour of integration actually passes through a node connecting several elements, the flux density is multi-valued, and no current CAD system incorporates any rational way of determining which value of flux density to choose, or whether to apply some form of smoothing or

averaging. Finally, and probably most significantly, the numerical round-off problem already alluded to remains omnipresent.

If the tangential force is to be determined in a device such as an electric motor, as for the calculation of torque, then the tangential component of stress may not necessarily be directed continuously in the direction of rotation—any particular pole of the rotor body may be attracted both toward the stator magnetic pole opposite it and to the one behind it. The relative strengths of these two forces depend on the angle between the rotor and stator magnetic axes. However, the stress distribution evaluates both forces and the resultant is correctly the difference between them. As with the virtual work method, the useful force may be small compared to the individual, positively and negatively directed, force components. In some situations, substantial numerical errors may result.

In practical applications of the Maxwell stress approach to force calculation, it is advisable to evaluate the force using several contours and then to average the results. The deviation between results obtained for different contours will often serve to indicate the likely accuracy level achieved, even though strict upper and lower bounds are not available. If these cautionary notes convey the impression that the Maxwell stress approach has drawbacks, it is worth noting that it also has the advantage of being computationally cheap; it requires just one field solution, while the virtual work approach demands a minimum of two.

Current–Force Interactions

The force density of equation (67), used as the starting point for the Maxwell stress evaluation, may of course be employed directly in the calculation of electromechanical forces. Since (67) gives a body force density, the net force on a conductor, viewed as a rigid body, is

$$\mathbf{F} = \int_U \mathbf{J} \times \mathbf{B}\, dU. \tag{77}$$

This expression effectively involves integration only over the current-carrying regions and thus may involve little calculation. Its advantage is also its disadvantage: it cannot be employed to find the forces on magnetic objects which do not carry any current, and it is therefore of limited use. For example, in computing the torque of an electric machine, (77) is not useful, because it will not produce the forces tending to move the rotor, only the force exerted by the conductors on the slot walls. In general, this is a feeble force compared to the magnetic surface forces that move the rotor.

A specialized form of (77) has been used for over a century by electric machine designers. If all the current is confined to a filamentary conductor of length l, placed in a field of uniform flux density B, the general form of force in (77) reduces to

$$F = B\, i\, l, \tag{78}$$

where i is the conductor current. This expression will correctly give the forces, and hence the machine torque, if (and only if) the rotor is perfectly smooth and cylindrical; for only in that case are there no other forces acting on the rotor. It is curious to observe that removal of simplifying assumptions, as in introducing rotor slotting in this case, can sometimes actually make results worse!

Rate of Change of Inductance

Still another approach to computing electromechanical forces is based on finding the rate of change of inductance with displacement. This approach is closely related to the principle of virtual work and is particularly useful for devices whose desired end product is motion. The motion may be continuous, as in rotating or linear induction machines; or it may be transient, as in stepping motors or loudspeakers. The latter class includes a host of devices intended to supply kinetic energy on a pulsed basis, as for example in an electromagnetic hammer.

The following discussion is based on a ballistic linear actuator, the hammer of a door chime, as an example of a motional device in whose analysis forces are best calculated using the rate of change of inductance. The designer's aim here is to maximize the kinetic energy possessed by the hammer at the moment of striking its target, and conversely, to position the target so as to have it struck at the moment when the hammer possesses maximum kinetic energy. Fig. 15 illustrates the device in question, an axisymmetric solenoidal coil and a moving plunger. The plunger is attached to a nonmagnetic hammer, which is not shown in the drawing.

The ballistic actuator is forced to act by applying a voltage V to the terminals of its solenoid. When the voltage is suddenly applied, current begins to flow in the coil, and a magnetic force appears to accelerate the plunger, thereby endowing it with kinetic energy. Neglecting the coil resistance, the instantaneous power input to the coil is, in accordance with Faraday's law as given in equation (1),

$$i\,V = i\,\frac{d(Li)}{dt}. \tag{79}$$

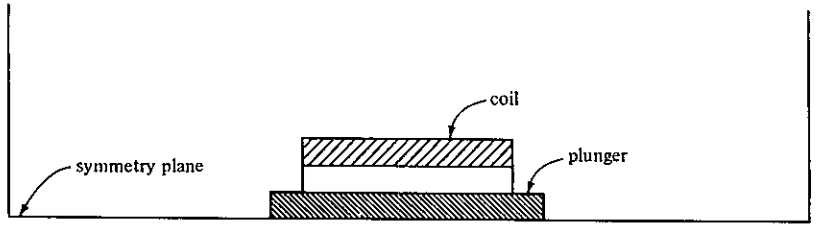

Figure 15. Ballistic actuator in its resting position. The plunger is centered within the coil structure.

This equation expresses a power balance: the input power must exactly equal the rate of increase of internal energy. Rewriting, (79) becomes

$$i\,V = i\,L\,\frac{di}{dt} + i^2\,\frac{dL}{dt}. \tag{80}$$

The second term on the right-hand side results entirely from motion and represents the rate of change of stored energy. The rate of change of magnetically stored energy due to motion alone is obtainable from equation (23) as

$$\frac{dW_m}{dt} = \frac{1}{2}\,i^2\,\frac{dL}{dt}. \tag{81}$$

The remaining energy must appear as the mechanical, i.e., kinetic, energy of the plunger. On the other hand, the kinetic energy is known to be expressible in terms of the mass m and the velocity v of the hammer and plunger as

$$W_k = \frac{1}{2}\,mv^2. \tag{82}$$

Equating the kinetic energy to itself, there is obtained

$$\frac{1}{2}\,i^2\,\frac{dL}{dt} = \frac{1}{2}\,mv^2. \tag{83}$$

Rewriting, using the chain rule of differentiation, there results

$$i^2\,\frac{dL}{dx}\,\frac{dx}{dt} = mv^2. \tag{84}$$

This equation can finally be solved for the hammer velocity. Noting that

$$v = \frac{dx}{dt}, \tag{85}$$

it immediately follows that v is given by

$$v = \frac{i^2}{m}\,\frac{dL}{dx}. \tag{86}$$

Thus the velocity of the plunger can be determined from the change of inductance as the plunger moves from one position to the next. The computation required is thus to find the coil inductance L as a function of position, using an energy-based definition of inductance, then differentiating (or taking finite divided differences) and substituting in (86). The force acting on the plunger at any point in its travel may be determined subsequently by any one of the force calculation methods discussed above. Since the mass of the plunger is presumed to be known, its acceleration may also be determined.

The flux distribution for the plunger in its equilibrium (symmetric) position is shown in Fig. 15; Fig. 16 shows the flux lines when the plunger is displaced to one side. The field plots clearly demonstrate that the force is one tending to restore the plunger to move to the centered position.

For a plunger 1.25 times the length of the coil and an excitation of 300 ampere-turns, typical results for the centered position of Fig. 15 and for the off-center position shown in Fig. 16 are as follows:

Position:	Centered	Off-center
Magnetic energy (mJ)	1.555	0.771
Plunger kinetic energy (mJ)	1.418	0.634
Plunger velocity (m/s)	0.476	0.319
Force on plunger (mN)	0.0	57.0

The plunger is considered to be magnetic but nonconducting throughout this discussion. If the plunger is made of conductive material then the motion through the magnetic field would induce eddy currents which, in turn, would alter the device performance.

Local Field Values

The determination of local phenomena—field components, power densities, and particle forces—is an important part of the CAD art. Extracting

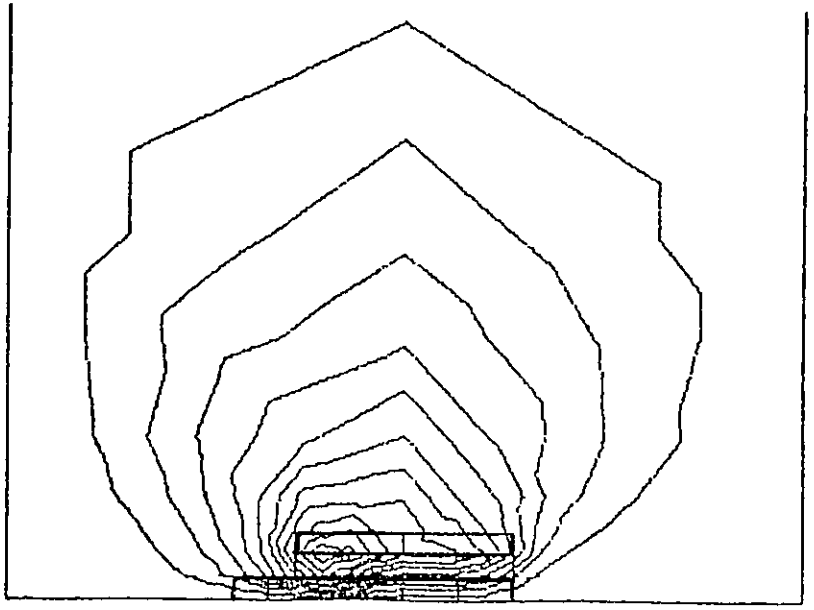

Figure 16. Ballistic actuator with plunger displaced to one side.

local results from solutions is therefore just as important as determination of inductances or forces. Of course, the extraction of local values requires somewhat different postprocessing operations.

Local phenomena rival global ones in importance because the electromagnetic device designer typically must satisfy two distinct forms of specifications: device performance and device feasibility. Both may involve either global quantities—stored energies, power levels, or total forces—or local values, such as flux densities, field uniformities, or particle trajectories. Devices used as *system components* interact with their environments through terminal parameters which reflect global performance specifications, although their feasibility limitations may be imposed by local internal phenomena. Thus a power reactor will be expected to have a particular inductance value, and an electric motor to furnish a specified torque; both designs will no doubt be constrained by local power loss densities or electric field gradients. In contrast to system components, *instruments* are often specified in terms of local behavior and limited by global considerations. An electron lens, for example, may have its performance prescribed in terms of aberration coefficients, which depend heavily on details of local field structure; its feasibility will be dependent on total power loss and material volume.

Local Field Components

Most present CAD systems employ the vector potential **A** to represent magnetic events, so that the determination of any field quantities must take this potential function as the point of departure. The most direct, and easily computed, quantity is of course the flux density, given in general by

$$\mathbf{B} = \text{curl } \mathbf{A}. \tag{12}$$

In two-dimensional cases where the vector potential is directed entirely into the plane of solution, only a single vector component appears; **A** seems to be a scalar quantity even though it is in reality a one-component vector. The general form (12) then assumes the special form

$$\mathbf{B} = \text{curl } (\mathbf{1}_z A). \tag{87}$$

$\mathbf{1}_z$ being the unit vector in the z direction. In terms of its Cartesian vector components, the flux density is thus

$$B_x = -\frac{\partial A_z}{\partial y}, \tag{88}$$

$$B_y = +\frac{\partial A_z}{\partial x}. \tag{89}$$

If required, the corresponding magnetic field values are found by first determining the local value of reluctivity, then multiplying. The reluctivity

is of course a function of the entire flux density, not merely one component,

$$\nu = \nu\ (B_x, B_y), \qquad (90)$$

so that to find the magnetic field **H**, determination of *both* components of **B** is necessary. This determination is generally done automatically by good CAD systems, and the user need not be explicitly aware of it, until the need arises to push the limits of the system. The field components are of course given by

$$H_x = -\nu \frac{\partial A_z}{\partial y} \qquad (91)$$

and

$$H_y = +\nu \frac{\partial A_z}{\partial x}. \qquad (92)$$

Comparable expressions are easily derived for components in other coordinate systems. In point of fact, most CAD work is conveniently done in either Cartesian or polar coordinates, as a matter of user preference. The finite element methods commonly used are substantially coordinate-independent, but users preferences tend to favor polar coordinates for the analysis of rotating devices, Cartesians for translational motion or symmetry.

Where a magnetic scalar potential function Ω is employed in solution, the vector quantity directly obtainable is the magnetic field **H**, not the flux density, for

$$\mathbf{H} = -\text{grad}\ \Omega \qquad (93)$$

whose components are clearly obtained by forming the partial derivatives

$$H_x = -\frac{\partial \Omega}{\partial x} \qquad (94)$$

and

$$H_y = -\frac{\partial \Omega}{\partial y}. \qquad (95)$$

To obtain the corresponding components of flux density **B**, H_x and H_y are multiplied by the local material permeability. In nonlinear materials, the permeability is a function of *all* components of **H**, as is evident from (90), so all components of **H** must be determined to find any component of **B**. While in principle the permeability may be a tensor quantity, few current CAD systems allow for tensor permeabilities to be taken into account automatically.

A Recording Head

An example of the need for local field values may be found in the examination of a magnetic recording head. The specification of a recording head must first and foremost be stated in terms of performance, that is, in terms of the quantity and quality of information it is actually capable of writing onto the recording medium. Secondary specifications may involve global values such as terminal impedance. The feasibility of a proposed design, however, is very much bound up with global values, such as the power loss in the device, and its associated temperature rise.

The primary point of interest in the analysis of a recording head is presumably determination of the fields and field gradients within the recording medium. The magnetization of the recording medium requires that the field gradient be high enough at the writing point, but it must be considerably lower elsewhere to avoid destroying information already recorded. Fig. 17 shows an outline drawing of a perpendicular recording head and the magnetic field produced by it, assuming the recording medium to be nonmagnetic. Ideally, the analysis should take care of both nonlinearity and hysteresis effects and might well include a reasonable model of the magnetic recording process itself. These requirements unfortunately surpass by a wide margin what currently available CAD systems are capable of doing. However, while awaiting the ultimate in software, the engineering analyst can still obtain a great deal of useful data from only a nonlinear analysis.

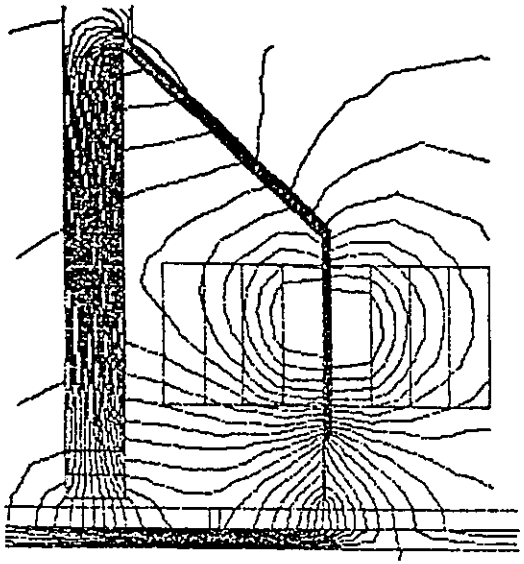

Figure 17. A perpendicular magnetic recording head. The recording medium is assumed to be nonmagnetic; it is backed by a magnetic material (bottom edge of drawing).

Similar requirements as to local field behavior arise in the design of large magnets for use in particle accelerators or nuclear magnetic resonance diagnostic machines. In these latter devices not only are the local field values important, but an even more crucial evaluative criterion is the *uniformity* of the field. It might be noted, however, that high field uniformity implies that the vector potential is required to vary linearly with position. This fact in turn implies that a very fine discretization of the air gap is not needed, at least for a first analysis!

Components in Local Coordinates

In many devices the local field values of importance are not the field or its components at a specific point, but rather along some track or contour. In the magnetic recording head example, the fields at specific points are certainly interesting; but most designers would consider a more global picture of greater interest and value than mere spot readings. Thus the first, essentially qualitative, representation of a solution desired by analysts is a flux plot, as in Fig. 17. But having examined the overall picture and found it satisfactory, the designer ordinarily wishes to move on to more quantitative data. In the present example, that probably means a curve such as Fig. 18, showing the perpendicular component of the magnetic field at the upper surface of the recording medium.

If the local values of field prove more or less satisfactory on examination, the designer may well attempt to modify the shape of the head, in order to improve the head performance as expressed by the curve in Fig. 18. The next calculation may very likely seek to determine the ampere-

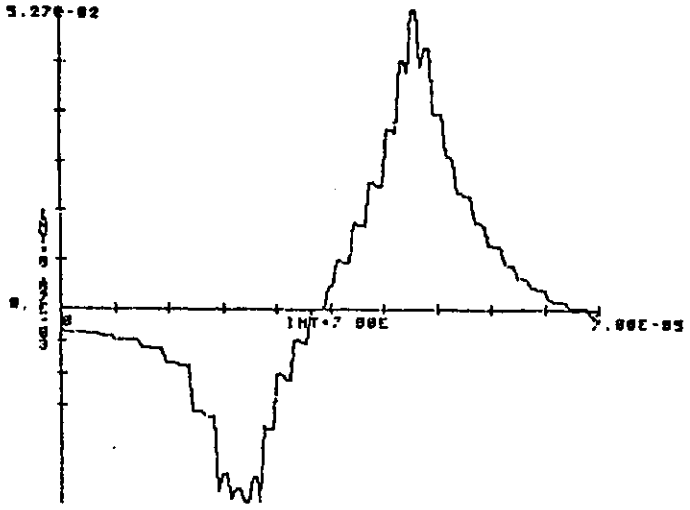

Figure 18. Vertical component of magnetic field under the recording head, at the top surface of the recording medium, plotted against longitudinal distance.

turn requirement of the iron member of the head itself. In other words, the designer may wish to evaluate the magnetomotive force M_{iron} given by

$$M_{iron} = \int_{iron} \mathbf{H} \cdot \mathbf{dl}. \tag{96}$$

The contour of integration must follow an inverted U-shaped path, more or less as in Fig. 17. This path will not in general coincide with coordinate directions, and it will be necessary to resolve the vector **H** into components tangential to the path (parallel to the direction of the line segment **dl**) and normal to it. Only the tangential component then enters the calculation in (96). This situation is similar to the component calculations encountered in the Maxwell stress expressions, equations (75) and (76), where the *normal* and *tangential* directions refer to the local orientation of the surface of integration or, in two-dimensional problems, the contour of integration.

In rotating electric machines, most contours of interest are either circular, or follow the complicated outline of some machine part. Probably the most commonly required field component plot in such equipment is the air gap flux density, for it has served generations of designers as a traditional working tool and has the immense advantage of familiarity. The radial component of flux density in particular gives both a qualitative and quantitative indication of the coupling across the machine air gap and may also be used in the calculation of the forces exerted on the shaft. The tangential component contributes to the calculation of the thrust forces.

Where field or potential values are required along a contour, for purposes of display or calculation, the Cartesian vector components can be resolved to yield the normal and tangential component values. Two reasonable ways of proceeding may be suggested: contours likely to be needed may be predefined in the CAD system, or the user may be given facilities for defining his own. Air-gap fluxes of machines, for example, are a very common requirement and justify inclusion of circular arcs as standard contours in general-purpose CAD systems; the same may be said of straight-line segments. Other, less common, shapes are probably better catered for by allowing them to be defined as needed.

Local field values to be displayed or further processed may be selected by location, or by value. Selection by location—choosing values at points or along contours—has already been discussed at some length. Selection by value typically implies searching. For example, the designer may inquire after the *highest value of flux density* and the place where it occurs; or, more broadly, he may wish to examine *all regions where the flux density exceeds 90% of its peak value*. Selection by value is often combined with selection by location, as in *highest flux density within the recording medium*, or *loss density above average value in the left half*.

To permit selection by location, interactive CAD systems may include facilities for the user to identify chosen sections of the problem model.

These sections may be two-dimensional portions (complete areas) within the model, such as might be needed in the calculation of the loss in a tooth; or one-dimensional portions, i.e. contours, along which the distribution of the field is required; or zero-dimensional pieces, i.e. points, at which the local value of the field (or its component in a given direction) might be needed. The geometric entities as well as the mathematical operations required are similar in principle, but different in details in each case.

To permit selection by value, control languages for postprocessing need to include suitable command verbs. This matter will be dealt with in some detail in connection with the structure of postprocessing programs themselves.

Determination of Terminal Voltage

Finding the motionally induced terminal voltage of rotating machines is a design requirement traditionally fulfilled by calculating the air-gap flux density. In CAD systems which solve for the vector potential, this approach is feasible, but usually bad. It does produce results, but does through tortuous computations what can be achieved simply and quickly, and with better accuracy, through a calculation based on flux linkages.

The induced motional voltage in a thin electric machine coil of n turns is easily calculated. Let the coil have one side placed at a particular point L in the $r-\theta$ plane, the other side at a point R. If the axial length of the machine is Z, then the flux linkages of this coil at any particular moment are given by

$$n\phi = nZ\,(A_R - A_L). \tag{15}$$

The coil voltage is, as usual, given by Faraday's law,

$$e = -\frac{d}{dt}(n\phi). \tag{1}$$

In a rotating machine, the points R and L will move along circular paths of fixed radii r_R and r_L, with their angular positions changing by an increment $d\theta$ in time dt. Thus, by the chain rule of differentiation,

$$e = -\frac{d\theta}{dt}\,\frac{d}{d\theta}\,(n\phi). \tag{97}$$

The first factor on the right will of course be recognized immediately as the rotational speed of the machine.

Where coils of substantial cross-sectional area, and possibly nonuniform winding density, are employed, the above expression may be generalized by employing the broader definition (20) of flux linkages. A small current

i_c is imagined to flow in the coil, causing a current density J_c to exist in it. Integrating over S, the entire area occupied by the coil, (97) then becomes

$$e = -\frac{d\theta}{dt}\,\frac{d}{d\theta}\,\frac{\int_S A\,J_c\,dS}{i_c}. \tag{98}$$

The computational procedure implied by this expression is simple: the flux linkages are evaluated for two positions of the rotor. They are subtracted and divided by the angular displacement so as to form an approximation to the rate of change of flux linkages with angle. Multiplication by the rotational speed then yields the generated voltage directly.

In most practical cases the generated voltage can be found from a single field solution. It is not necessary to compute two solutions for two distinct positions of the rotor, for it usually suffices to leave the rotor and stator alone but to shift the position of the winding by one slot pitch. An exception to this general rule occurs when both rotor and stator are slotted, and the air gap is small. The generated voltage then includes a significant slot ripple component, which will be ignored if the rotor is effectively rotated by exactly one slot pitch. Recomputation of course involves extra work. There is consolation, however, in the fact that the conventional air-gap flux density method would not exhibit the slot ripple at all.

The traditional method of computing generated voltage examines the air-gap flux density, and in effect computes the generated voltage that would result if the coil were located in the air gap. The accuracy of this technique is lower than can be expected of direct flux linkage computation as above, and of course the amount of work is considerably larger. Unless there is some particular reason to the contrary, the air-gap flux method is therefore not to be recommended.

Fields in Axisymmetric Problems

Many problems encountered in magnetic devices have some form of axial symmetry. Solution in cylindrical coordinates actually produces true three-dimensional solutions in such cases. For example, the magnetic bearing and the ballistic actuator discussed above are actually axisymmetric devices.

When postprocessing operations are carried out on axisymmetric fields, care must be exercised on two counts. First, the common vector operators *grad, curl,* and *div* have forms which are different from those encountered in Cartesian coordinates. Secondly, most axisymmetric solver programs do not actually compute the vector potential, but a closely related potential-like function weighted by the radial coordinate r.

Differentiation and integration of vector quantities in cylindrical coordinates are not difficult operations, and their general forms can be found in

standard books on electromagnetics. The classical axisymmetric problem has purely solenoidal current densities, with the vector potential **A** purely azimuthal, that is, possessing only a single component A_θ. In this rather special situation, the curl operation is simplified, so that the flux density components are now derived as

$$B_r = -\frac{\partial A_\theta}{\partial z} \tag{99}$$

and

$$B_z = \frac{1}{r}\frac{\partial}{\partial r}(r\,A_\theta). \tag{100}$$

The azimuthal component of **B**, which would have to be derived purely from r and z directed components of **A**, does not exist.

Most CAD systems do not solve for the magnetic vector potential directly because the axisymmetric Poisson and Laplace equations include a singularity at $r = 0$. The singularity arises from a multiplier of the form r^{-1} and is removed by modifying the potential by a similar but opposite multiplier. In the MagNet system, for example, the modified potential

$$U = \frac{A_\theta}{r} \tag{101}$$

is used. The components of **B** for this case are obtained by direct substitution into equations (99) and (100), as

$$B_r = -\,r\,\frac{\partial U}{\partial z} \tag{102}$$

and

$$B_z = 2U + r\,\frac{\partial U}{\partial r}. \tag{103}$$

In these expressions, it is interesting to note that r appears as a multiplier; thus U does not have to vanish along the axis of rotational symmetry in order for A_θ to vanish. That is to say, the modified potential U can be a regular function at the axis, which is precisely the reason for introducing it in the first place. The flux density components are regular also; in fact, at the axis itself, B_z is proportional to U, as can be seen readily from equation (102).

Some older CAD systems solve directly for the potential **A**, ignoring the singularity whenever it can be ignored and dealing with it in an ad hoc fashion wherever it cannot. In some special-purpose programs which solve for **A** directly, the regions modelled are in fact simply forbidden from including the axis of symmetry. There also exist other quite useful modified potentials different from U of equation (101). While the several

ways of regularizing the differential equations all make solution feasible, it is important to be aware which modified potential is used before attempting to derive flux densities or any other quantities from it.

In axisymmetric problems, flux lines are not lines of constant A, but rather lines of constant rA. This fact easily follows from the general relationship

$$\phi = \oint \mathbf{A}\cdot\mathbf{dl}. \tag{13}$$

Consider a closed contour which proceeds radially outward from the axis of symmetry, then follows a circular arc at $r = constant$, $z = constant$ for an angular distance of θ_0, and finally returns to the starting point along a radial path. The flux ϕ_0 contained within this wedge-shaped contour may be evaluated by reference to (13). Since the vector potential **A** is everywhere purely azimuthal, no contribution accrues to the integral along the two radial sides, where **A** is orthogonal to the distance element **dl**. Along the circular arc, **A** and **dl** are exactly parallel, so the integral evaluates to exactly $rA\theta_0$. Thus, in this particular case,

$$\phi_0 = rA\;\theta_0. \tag{104}$$

If the circular arc is displaced to some other point in the r-z plane in such a way as to keep ϕ_0 constant, then the new point must lie on the same flux line as the old, for if it did not, moving from one point to the other must have changed the flux linked. Flux lines in an axisymmetric problem are therefore lines of constant rA, not lines of constant A, and it is necessary to keep track of which are plotted in any given situation. Furthermore, if the modified potential U of equation (101) is employed, it follows that

$$\phi_0 = r^2\;U\;\theta_0. \tag{105}$$

In other words, flux lines are lines of constant r^2U, and it is these which designers commonly wish to inspect. As a corollary, the requirement that the flux density **B** remain unchanged in the air gap of a magnet is equivalent to requiring that U remain constant. The modified potential U, in other words, is smoother than the vector potential **A** itself, hence computationally better behaved.

All the foregoing discussions for the calculation of terminal conditions, mechanical effects, and local values in the x-y plane remain valid in the axisymmetric case, provided due care is taken to use the correct forms of the differential operators, as well as the correct potential function.

Linear Time-Varying Problems

Many magnetic devices are initially analyzed on the assumption that the fields are essentially static, that is, no induced currents flow anywhere. For

some devices refined analysis is subsequently required to determine the effect of eddy currents which actually do occur. Still other devices cannot work at all if no induced currents are allowed to exist. Most CAD systems therefore permit time-varying as well as static fields to be analyzed.

The postprocessing operations applicable to static field solutions, as described this far, are for the most part equally valid for time-dependent problems. However, there are several new quantities that a designer might require, which only have value in time-varying cases. Some of these will be reviewed in the following.

Time-Harmonic Analysis

In general, computer software for time-varying problems can be classified into time-domain and frequency-domain programs. Time-domain programs work by generating a sequence of solutions, one for each of a series of time values; frequency-domain programs solve for sinusoidal excitations at one or more fixed frequencies. This distinction may be familiar from circuit theory, where a corresponding classification of programs is possible. The time-domain approach is capable of dealing with arbitrary excitations applied to nonlinear devices, but its very generality is also its disadvantage: it generates extremely large data files at the cost of great quantities of computer time. In essence, it produces a movie film of the field behavior, one frame per time instant; all the frames need to be computed and stored. The frequency-domain technique is very compact and cheap, for the volume of data to be stored is just double that of a static solution, while the computing time required is greater than that for static solutions, but greater only by a modest amount. It is unfortunately applicable only to linear problems, because it is based on the premise that *all* time-dependent phenomena are sinusoidal, a premise satisfied only by linear systems.

Most present CAD systems make provision for analysis of time-harmonic phenomena in linear materials, but they do not provide equally extensive facilities for the general time-varying case. This restriction is not nearly so draconian as might seem, for many design specifications and many concepts on which system evaluation is based rely on notions of impedance, phase delay, and amplitude, notions which are strictly valid only for sinusoidally time-varying phenomena. The present discussion will therefore be confined to frequency-domain analysis.

As discussed in the chapter *The Potential Equations of Magnetics*, sinusoidally varying fields are conveniently described by combining the magnetic vector potential **A** and (sometimes) the electric scalar potential V. The vector potential is determined by solving the complex diffusion equation

$$\nabla^2 \mathbf{A} + j\omega\mu g\mathbf{A} = -\mu\mathbf{J}_O \tag{106}$$

which is obtained from the general time-varying case through replacement of the time derivative by the phase quadrature operator $j\omega$. The vector

potential **A** is still assumed entirely z directed, but it is now a complex number rather than a pure real as was the case in static field problems. Here, as earlier, $\mathbf{J}_O$ represents the current density that would exist at extremely low frequencies, i.e., the current distribution that would obtain if the excitations all took place so slowly that no induced currents existed and all convection currents assumed a spatial distribution in accordance with material resistivities only. Briefly, though not totally accurately, $\mathbf{J}_O$ is sometimes referred to as the *dc distribution* of current density.

Of course, it must not be forgotten that the solution to Equation (106) really represents two solutions to two problems,

$$\nabla^2 \mathrm{Re}[\mathbf{A}] - \omega\mu g\ \mathrm{Im}[\mathbf{A}] = -\mu\ \mathrm{Re}[\mathbf{J}_O] \tag{107}$$

and

$$\nabla^2\ \mathrm{Im}[\mathbf{A}] + \omega\mu g\ \mathrm{Re}[\mathbf{A}] = -\mu\ \mathrm{Im}[\mathbf{J}_O]. \tag{108}$$

The point here is that Re[**A**], the solution of (107), represents the magnetic field at that moment in the a-c cycle when the real part of the driving current, Re[$\mathbf{J}_O$], reaches its peak and its imaginary part Im[$\mathbf{J}_O$] vanishes. This solution is altogether independent of the other solution Im[**A**], which represents the field at the moment when the imaginary part Im[$\mathbf{J}_O$] reaches its peak and Re[$\mathbf{J}_O$] vanishes. At all other times, the real time-varying vector potential $\mathbf{A}(t)$, which might alternatively have been found by using a time-domain solution, can be determined by superposing suitably weighted components. Its value for any time t is given by

$$\mathbf{A}(t) = \mathrm{Re}[\mathbf{A}]\cos\omega t + \mathrm{Im}[\mathbf{A}]\sin\omega t. \tag{109}$$

By choosing a sequence of time values t, say at every 10° of the ac cycle, and evaluating the potential **A**(t) at each time instant in the sequence, a series of "snapshots" of the field can be generated. The two solutions Re[**A**] and Im[**A**] suffice for this purpose, no other field solution is required. Often, experienced analysts take the sequence of time values to lie at 90° intervals; in other words, they plot only Re[**A**] and Im[**A**] and do not bother with intermediate values. This abbreviated approach is suitable for anyone who has acquired considerable familiarity with time-harmonic field problems and is therefore capable of performing the required interpolation ("how **A** gets from Re[**A**] to Im[**A**]") mentally. Most engineers are well advised to generate at least two or three intermediate plots as well, so as to gain a clearer appreciation of how the flux lines of one plot gradually merge and split in time to produce the time-quadrature picture.

As time goes on, magnetic flux lines in time-harmonic problems do not merely pulsate, they shift and twist about in space. The plotting of field patterns through time is therefore an important tool in learning how magnetic devices behave and how to make design improvements in them. As

suggested above, one good way of plotting the fields is to produce a sequence of time "snapshots", movie frames as it were, which can be viewed in succession. Such sequences of snapshots can even be combined into motion pictures. Although actually making a motion picture is a very demanding technical task best turned over to professionals, the motion picture can sometimes communicate in seconds things which cannot be got across in minutes or hours of explanation. It is therefore of potentially very great value in management presentations or sales work. The device designer is well advised to keep in mind, while examining field plots, that novel devices or even design changes will need to be explained at some time motion pictures can communicate some things very effectively!

Current Densities

A quantity of frequent interest in time-harmonic fields is the density of induced (eddy) currents, and the density of total currents that results. In classical terminology, *eddy currents* are generally understood to be currents induced in conductive material which would not otherwise carry any current; the term *skin effect* is used to denote current density distribution in material that carries both excitation current ($\mathbf{J}_O$ in the above) and eddy currents. Here the term *induced currents* is employed to denote currents that arise from time-varying magnetic fields in both situations.

The induced currents that result from time variations in the magnetic field can always be obtained by combining Ohm's law,

$$\mathbf{J} = g\,\mathbf{E}, \tag{110}$$

with the general rule for electric fields

$$\mathbf{E} = -\frac{\partial \mathbf{A}}{\partial t} - \text{grad } V. \tag{111}$$

Combining, and replacing the time derivative by the complex operator valid in the time-harmonic case,

$$\mathbf{J} = -j\omega g\,\mathbf{A} - g\,\text{grad } V. \tag{112}$$

Here $\mathbf{J}$ is the total current density, attributable to all the mechanisms that make current flow. However, the two terms on the right-hand side may be identified separately, as the dc current distribution

$$\mathbf{J}_O = -g\,\text{grad } V \tag{113}$$

and the *induced current density*

$$\mathbf{J}_i = -g\,\omega\mathbf{A}. \tag{114}$$

It should be noted that the general equation (112) relies only on the definition of $\mathbf{A}$ as a vector whose curl produces the flux density $\mathbf{B}$; it is therefore quite independent of the choice of gauge.

Considerable interest attaches to plotting the eddy current densities that occur in conductive parts of various magnetic devices. In the two-dimensional cases discussed here for the most part, $\mathbf{A}$ is z directed or (in axisymmetric problems) θ directed, so the induced current density $\mathbf{J}_i$ must therefore be z directed or θ directed also. Plotting is therefore not a fundamental problem—there is still only one component of $\mathbf{J}$ for a single-component $\mathbf{A}$—but it is slightly more complicated because the form of display is less well established conventionally. While it is perfectly possible to draw contours of equal magnitude of current density, such lines do not always convey the desired information. A better choice, in CAD systems which permit it, is to produce a *zone plot*. In such a plot, the space to be plotted is divided into a set of *zones*, so that one zone includes all that portion of space in which the current density lies between its maximum value and (say) 90% of maximum; another zone in which the current density lies between 90% and 80% of maximum, and so on down to zero. Plotting is then done by filling in the highest-density zone with a very bright color, the next zone with a less bright color, . . . , and the zero-density zone with a dark color. People generally find it easy to associate bright colors with high densities, dark colors with low ones, and can therefore interpret zone plots rather easily.

Equiphase plots of current density values can sometimes be quite illuminating. These will be considered more fully later, in connection with equiphase plots of other phasor quantities such as fields and the Poynting vector.

Power Flow and the Poynting Vector

The Poynting vector is a quantity generally associated with the flow of energy in electromagnetic devices and used to measure power transferred. It is established by considering the stored energy and the power dissipated in a device and relating them through the law of energy conservation.

Let W_m be the magnetically stored energy in a closed region of space U. If the region contains magnetically linear materials, as it must if time-harmonic analysis is to be useful, then this energy is given by

$$W_m = \frac{1}{2}\int_U \mathbf{B}\cdot\mathbf{H}\,dU, \tag{115}$$

while the electrically stored energy W_e in the same region is

$$W_e = \frac{1}{2}\int_U \mathbf{E}\cdot\mathbf{D}\,dU. \tag{116}$$

The power P_e dissipated in the region U is similarly given by

$$P_e = \int_U \mathbf{E} \cdot \mathbf{J}\, dU. \tag{117}$$

Now the law of conservation of energy clearly requires that the input power P into the region U must equal the dissipation plus the rate of increase of stored energy. That is,

$$P = P_e + \frac{\partial}{\partial t}\,(W_m + W_e). \tag{118}$$

Substituting the explicit expressions for P_e, W_m, and W_e from above, the input power turns out to be

$$P = \int_U \mathbf{E} \cdot \left[\frac{\partial \mathbf{D}}{\partial t} + \mathbf{J}\right] dU + \int_U \mathbf{H} \cdot \frac{\partial \mathbf{B}}{\partial t}\, dU. \tag{119}$$

But the factor in parentheses, and the time derivative of **B**, are readily recognized as right members of the Maxwell curl equations. Substituting, there finally results

$$P = \int_U \mathbf{E} \cdot \text{curl } \mathbf{H}\, dU - \int_U \mathbf{H} \cdot \text{curl } \mathbf{E}\, dU. \tag{120}$$

A well-known vector identity permits rewriting the latter in terms of the cross-product of the vectors **E** and **H** as

$$P = \int_U \text{div } (\mathbf{E} \times \mathbf{H})\, dU, \tag{121}$$

which may be converted, making use of the divergence theorem, into the closed surface integral over the bounding surface ∂U of the region U:

$$P = \oint_U (\mathbf{E} \times \mathbf{H}) \cdot \mathbf{dS}. \tag{122}$$

The vector in parentheses is usually known as *Poynting's vector* and denoted by **S**,

$$\mathbf{S} = \mathbf{E} \times \mathbf{H}. \tag{123}$$

The Poynting vector is conventionally thought to measure local power flow, much as W_e and W_m are held to measure local stored energy densities. It is often plotted and viewed as a quantity indicative of the performance of a device, especially in so far as portions of a device may

be concerned. If some part of a device hardly matters in power flow or energy transformation, should it be there at all?

The Poynting vector may be written in terms of the vector potential **A**, so as to facilitate its evaluation in CAD systems. Substituting from (12) and (111) for the magnetic and electric quantities respectively, equation (123) assumes the form

$$\mathbf{S} = -j(\omega/\mu)\, \mathbf{A} \times \text{curl } \mathbf{A} - (1/\mu)\, \text{grad } V \times \text{curl } \mathbf{A}, \tag{124}$$

which only involves the potentials normally available from field solutions.

The Poynting vector is generally a two-component quantity even in problems where only a single component of vector potential exists. It is difficult to plot in a graphic display, because there is no conventional form of representation for time-varying multicomponent vectors. One fashion in which the Poynting vector, and indeed other planar vectors such as **B** or **H**, can be displayed is in two independent plots, showing phase and magnitude.

Equal-magnitude plots of vector quantities are often best arranged as zone plots, while equiphase contours are best left as line contours. The reasons are almost solely those of ease in visual perception and interpretation. Energy transport (power flow) in a time-harmonic system is essentially a wave phenomenon: a packet of energy is sent off from its source and allowed to propagate toward its eventual sink, possibly attenuating through dissipation as it travels. In one period, any point in the region U executes one full cycle of time events. Conversely, the wave or energy packet propagates during one period precisely as far as it is necessary to go to find another space point where electrical events are 360° lagging in phase. In other words, energy travels in one period from one equiphase contour to the next one 360° removed. Plotting the equiphase contours thus makes evident the flow of energy through the device, and thereby exhibits to the experienced designer whether and how device performance might be improved.

Exploiting Superposition

In time-harmonic problems the assumption of magnetic linearity is essential, for in its absence the notions of sinusoidal phasor analysis become invalid. However, if linearity is assumed, then superposition may be used, permitting simplification of many problems.

Where a magnetically linear device has altogether N distinct excitation ports—windings, conductors, or other current-carrying entities—at most a total of N distinct field solutions is required at any one frequency to cover *all* possible cases. To choose a specific example, consider a transformer with N windings, in which eddy current losses may occur in the lamina-

tions or in other conductive materials. Now any possible excitation state of this transformer may be described by specifying all its winding currents, all of which jointly may be viewed as the driving current density $\mathbf{J}_O$ of equation (106). If $\mathbf{J}_{Ok}$ is the current density distribution caused by a *unit* current in winding k, then any set of winding currents $\{i_k, k = 1, \ldots, N\}$ will cause the excitation current distribution in the transformer to have the value

$$\mathbf{J}_O = \sum_{k=1}^{N} i_k \, \mathbf{J}_{Ok}(x,y). \tag{125}$$

Now a solution may be sought to equation (106) in the presence of only one of the winding currents. That is to say, all but one of the i_k may be set to zero, that one being given unity value. Equation (106) then assumes the form

$$\nabla^2 \mathbf{A}_k + j\omega\mu g \, \mathbf{A}_k = -\mu \, \mathbf{J}_{Ok}. \tag{126}$$

In the presence of many different currents i_k, the solutions $\mathbf{A}_k$, each valid for a single unit current, superpose. For the combined current $\mathbf{J}_O$ of (125), the relevant differential equation is thus (106), and its solution is immediately obtained as

$$\mathbf{A} = \sum_{k=1}^{N} i_k \, \mathbf{A}_k. \tag{127}$$

There is never any need to solve more than N field problems, for once the *N basis solutions* of (126) have been generated, any and all other solutions can be obtained as in (127). It should be noted that any eddy currents which may exist in conductive parts of the device are included in the solutions of (126), and therefore also of (127).

When solutions are constructed by superposition, it is worth noting that the solver phase of the CAD system may at times not be used to produce any actual solution at all. It may be employed solely as a generator of basis solutions, with the composite solutions, as in equation (127), created by a postprocessing program. Considerable computation time may be saved by this approach since the fastest available solvers operate in approximately $O(N^2)$ time whereas scaling and adding operations occur in $O(N)$ time.

Electrostatic Calculations

The operations and solution systems described earlier in this book have all referred to magnetic systems. However, many problems of an electrostatic nature are also described by either the Laplace or Poisson equation. Such electric field problems can be solved by using magnetostatic packages provided the user can determine the appropriate equivalences. To illustrate, a simple capacitor with mixed dielectrics may be considered.

A Parallel-Plate Capacitor

A parallel-plate capacitor, imagined to extend infinitely into the paper, is illustrated in Fig. 19. Initially, the model is defined with a dielectric of air and only one of the possible symmetries in the problem is used. The equation to be solved is a scalar Laplace equation in which the scalar potential is the voltage in the system. The electric field $\mathbf{E}$ is related to the electric scalar potential by

$$\mathbf{E} = -\operatorname{grad} V, \tag{128}$$

a special case of (111), with no time variations.

When an equivalence is drawn with magnetostatics problems, the scalar potential V in the electrostatic model is substituted for the vector potential $\mathbf{A}$ of magnetostatics. The magnetostatic Poisson equation for nonlinear materials may be written

$$\operatorname{curl} \nu \operatorname{curl} \mathbf{A} = \mathbf{J}, \tag{129}$$

whose mixed-dielectric electrostatic equivalent is

$$\operatorname{div} \epsilon \operatorname{grad} V = -\rho. \tag{130}$$

In a two-dimensional Cartesian coordinate system these two equations reduce to identical form, when written out in detail. Hence an equivalence may be drawn between the current and charge densities, the vector and scalar potentials, as well as the material properties (reluctivity and permittivity). The latter equivalence is obvious but must be treated with caution if modelling of material property curves is required, for most magnetic material modelling programs allow permeability, rather than reluctivity, to be specified by the user. That is, the user needs to specify an $\mathbf{E}$ against $\mathbf{D}$ curve.

Once the analogous quantities have been identified, solution of the electrostatic problem proceeds exactly as if the problem were magnetostatic.

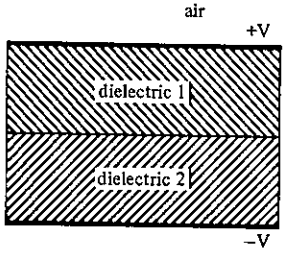

Figure 19. A simple parallel-plate capacitor.

For the capacitor of Fig. 19, an equipotential plot is shown in Fig. 20. Both dielectric slabs are assumed to be air in this case. The analogy with magnetics problems is immediately apparent on inspection of the solution. Although the finite element subdivision employed for the solution is not explicitly shown, it can be seen from Fig. 20 that the elements are generally quite large. An analogous solution, for the same geometric shapes but two different dielectrics, appears in Fig. 21. The lower portion of the dielectric space in this case has a relative permittivity of 20. The electric field can be seen to be considerably lower in the dielectric than it is in air.

The capacitance of multi-dielectric systems is probably the quantity of major interest, analogously to inductance in magnetics problems. It may be calculated by using substantially similar techniques, as will now be discussed in some detail.

Determination of Capacitance

The capacitance of an electrostatic device may be determined in either of two ways, which are analogous to the ways of computing inductance in

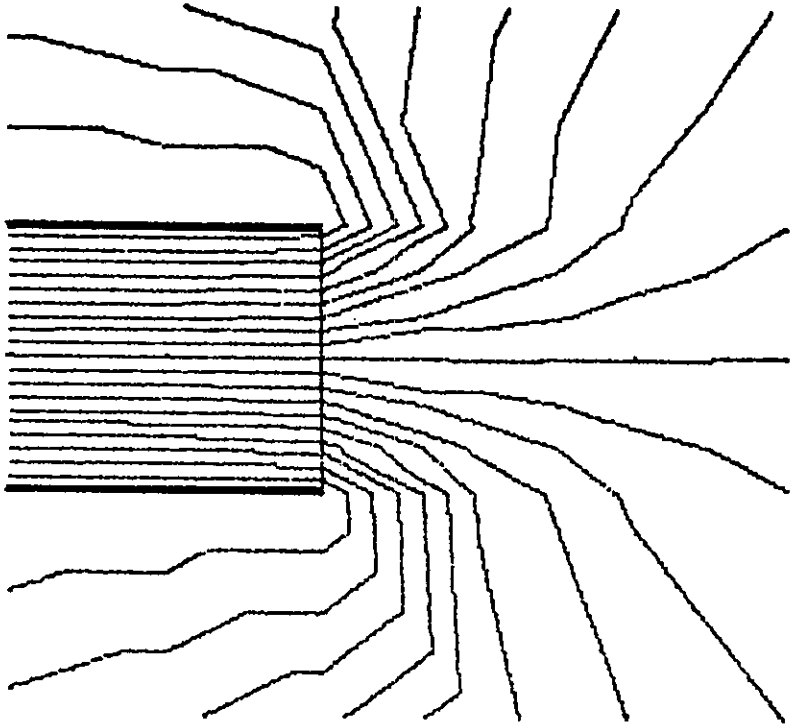

Figure 20. Electric field of the capacitor, with all dielectric material having the permittivity of air.

magnetics problems. One approach requires a calculation of the stored energy, the other a determination of the total electric flux. The capacitance C may be defined in two corresponding ways, either in terms of the electrically stored energy W,

$$W = \frac{1}{2} CV_0^2, \tag{131}$$

or as the ratio of charge to potential difference V_0 between the electrodes,

$$C = \frac{q}{V_0}. \tag{132}$$

The total electric flux is, by definition, equal to the total charge q. To determine it, one or the other electrode may be encased in a closed surface S, to which Gauss' law is then applied:

$$q = \oint_S \mathbf{D} \cdot \mathbf{dS}. \tag{133}$$

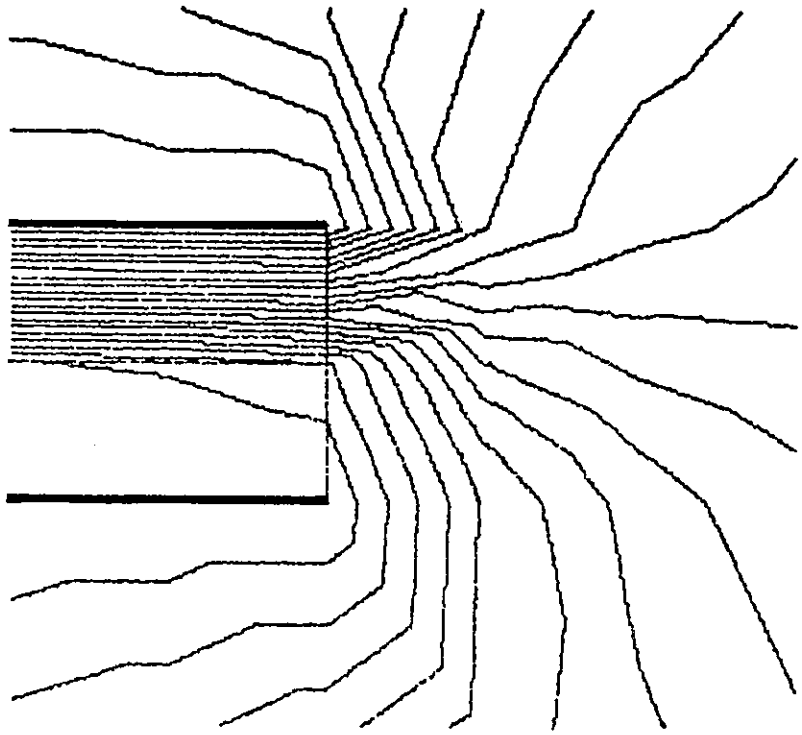

Figure 21. Electric field of the capacitor with mixed dielectrics.

The capacitance C is then found by combining (132) and (133) to yield

$$C = \frac{\oint_S \epsilon \text{ grad } V \cdot \mathbf{dS}}{V_0}, \quad (134)$$

where (128) has been employed to eliminate the flux density and to substitute the computed quantity V. It should be noted that any surface S which completely encloses the electrode is equivalent to any other in principle, since Gauss' law is equally applicable in all cases. In two dimensions, the surface integral of course reduces to a closed contour integral in the x-y plane.

Finding the capacitance via the energy route involves use of equation (116) to determine the stored electric energy, then equation (131) to determine C. Combining the two, there is obtained

$$C = \frac{\int_U E^2 \, dU}{V_0^2}. \quad (135)$$

Because finite element methods generally are energy-based methods, (135) usually provides the more accurate basis for calculation.

To illustrate, the capacitance was determined for a capacitor having a plate width of 0.1 m and a plate separation of 0.01 m. The gap is half filled with air and half filled with a dielectric of given relative permittivity, taken as either unity (to correspond with Fig. 20) or 20. Both an energy calculation and a contour integration were used. In the latter case, two contours were chosen, both the full width of the model; the first passed through the upper half of the air gap, the second through the lower half. Thus, for the dielectric model, the second contour passed through the dielectric layer. Values for all three calculations are given in Table 5. These results may, of course, be checked roughly by the usual simple formula for a parallel-plate capacitor, provided that the fringing at edges is neglected. This calculation gives a value of 88.5 pF for the air capacitance and 168.6 pF when the dielectric is present. The actual values, when fringing is considered, should of course be higher.

Characteristic Impedance and Phase Velocity

A two-dimensional problem of some interest is that of a parallel-strip transmission line similar to the mixed-dielectric capacitor described above.

Table 5. Capacitance (pF) of parallel-plate capacitor

Permittivity	Contour integral 1	Contour integral 2	Stored energy
1.0	99.34	98.81	110.94
20.0	181.95	194.34	196.12

If a quasi-TEM wave is assumed to propagate along such a line, the wave impedance and phase velocity may be determined directly from the two capacitance values computed above. The characteristic impedance of any transmission line is given by

$$Z_0 = \left[\frac{L}{C}\right]^{1/2}, \quad (136)$$

while its phase velocity, for quasi-TEM waves, is

$$v = \frac{1}{(LC)^{1/2}}. \quad (137)$$

When the dielectric material is air throughout, the phase velocity must be equal to the velocity c of light. If C_0 is the capacitance in that case, then

$$c = \frac{1}{(LC_0)^{1/2}} \quad (138)$$

and hence the line inductance is

$$L = \frac{1}{c^2 C_0}. \quad (139)$$

Since both dielectrics have the permeability of free space, the inductance value is unaffected by the presence of the dielectric. Thus, the characteristic impedance of (136) becomes

$$Z_0 = \frac{1}{(c^2 C_0 C)^{1/2}} \quad (140)$$

and the phase velocity

$$v = \frac{C_0}{C^{1/2} c}. \quad (141)$$

For the device described above the characteristic impedance is then 22.6 ohms.

Effects of Numerical Approximations

Finite element methods form the mathematical basis for nearly all current general-purpose CAD systems intended for magnetics problems. In postprocessing, the various mathematical operations—differentiation, integration, multiplication, and whatever else may be required—cannot be performed on the exact fields themselves, but must use the finite element approximations computed by solver programs instead. The precision of

results, and even their interpretation, may be affected by the numerical approximations involved, and some brief consideration of the nature of these approximations may not be out of place here.

Finite Element Approximations

Finite element methods invariably produce polynomial approximations to field solutions. Over each finite element, the potential field is computed as a polynomial expression in the coordinate quantities, whose degree is usually fairly low. The approximations are so constructed that the potential solution is continuous everywhere; but its derivatives are not continuous across the element edges. If the potential is imagined to describe a surface above the x-y plane, the surface is described by a polynomial over each of a set of patches. At the patch edges, the pieces of surface join so there are no "holes", but the joints are not necessarily smooth; there may be creases. Thus, if elements of order N are used, the potential is given by polynomials of order N within each element, but it is C^0 continuous globally. Arithmetic and analytic operations in postprocessing operate directly on the polynomials, and the operations may on occasion trespass on the limits of the possible.

One immediate result of piecewise continuity is that the flux density distribution along a specified contour through a problem model will not be smooth. The flux density plots shown at various points throughout this book make this phenomenon quite evident. It is of course perfectly possible to beautify the results by polynomial (or other) interpolative smoothing. In general, the plots shown in this book have not been subjected to any cosmetic improvement, in the belief that raggedness in results, should it occur, conveys information about the reliability of results, and it would be a mistake to destroy this additional information.

The postprocessing operations outlined in this chapter include a broad variety of mathematical operations. Not all of these are permissible on piecewise polynomial functions in all cases; and not all produce piecewise polynomial results from piecewise polynomial arguments. At least the following elementary mathematical operations are included in the list of requirements:

Arithmetic: addition and subtraction of potential or vector fields, multiplication of fields by scalars.
Arithmetic: multiplication and division of potential (scalar) field quantities.
Vector Operations: cross and dot products of vector fields.
Functional Operations: absolute values, square roots, trigonometric functions, and other general functions of potential fields.
Differential Calculus: gradients of potential fields; curls and divergences of vector fields.
Integral Calculus: line and surface integrals over defined geometric regions.

These will be briefly examined in the following, with a view particularly to the accuracy to be expected in practical cases.

Arithmetic Operations

The elementary arithmetic operations fortunately do not affect the nature of finite element approximations. When piecewise polynomial functions of degree N or lower are added, the result is always a field of piecewise polynomials of similar degree. Thus the elementary arithmetic operations—addition, subtraction, and multiplication by a scalar—may be performed without effect on the nature of numerical approximation. As always, however, a cautionary note must be sounded regarding subtraction: small differences of large numbers may exhibit round-off error accumulation to a considerable extent. Wherever possible, it is thus preferable to formulate problems in such a way as to avoid numerical formation of small differences.

Multiplication of fields invariably produces results which are piecewise polynomial, but not of the same degree as the arguments. For example, multiplication of a first-degree polynomial field by itself, thus creating, say $v^2(x,y)$ from $v(x,y)$, produces a piecewise polynomial field of degree 2. No existing CAD system can handle the growth in polynomial order that results from sequences of repeated multiplications, and the only practical approach is to approximate the high-order polynomials by others of lower order.

Division of piecewise polynomial fields can only yield piecewise polynomials in rare cases. In general, the result is a piecewise rational fraction, continuous but not necessarily bounded. Again, no CAD system is capable of handling such functions, and the usual technique is to approximate them by piecewise polynomials in a least-squares sense. It is clear that some of the fundamental properties of the functions being approximated might be lost in the process. For example, the quotient $u(x,y)/v(x,y)$ will be unbounded if $v(x,y)$ vanishes anywhere in the region of interest; but an approximation to the quotient cannot be unbounded if it is to be expressible in terms of polynomials. The division operation therefore requires particular caution.

Operations on vector quantities are not inherently different from operations on scalars. The inner product of two vector fields, after all, is expressible as the combination of products of scalars, the vector components; thus, any problems likely to arise with scalar fields are likely to recur when vectors are considered. However, it must be noted that troubles arising from the division operation are entirely absent, since division of vectors is undefined.

Differentiation and Integration

Finite element methods produce piecewise polynomial approximations continuous in function value, but not necessarily continuous in any