Control of Stepping Motors

A Tutorial

by Douglas W. Jones
THE UNIVERSITY OF IOWA Department of Computer Science

This material expands on material originally posted to the rec.railroad newsgroup in 1990. Significant parts of this material have been republished as sections 5.2.10, 10.8, 10.9 and 10.10 of the *Handbook of Small Electric Motors* edited by W. H. Yeadon and A. W. Yeadon, McGraw-Hill, 2001.

Introduction

Stepping motors can be viewed as electric motors without commutators. Typically, all windings in the motor are part of the stator, and the rotor is either a permanent magnet or, in the case of variable reluctance motors, a toothed block of some magnetically soft material. All of the commutation must be handled externally by the motor controller, and typically, the motors and controllers are designed so that the motor may be held in any fixed position as well as being rotated one way or the other. Most steppers, as they are also known, can be stepped at audio frequencies, allowing them to spin quite quickly, and with an appropriate controller, they may be started and stopped "on a dime" at controlled orientations.

For some applications, there is a choice between using servomotors and stepping motors. Both types of motors offer similar opportunities for precise positioning, but they differ in a number of ways. Servomotors require analog feedback control systems of some type. Typically, this involves a potentiometer to provide feedback about the rotor position, and some mix of circuitry to drive a current through the motor inversely proportional to the difference between the desired position and the current position.

In making a choice between steppers and servos, a number of issues must be considered; which of these will matter depends on the application. For example, the repeatability of positioning done with a stepping motor depends on the geometry of the motor rotor, while the repeatability of positioning done with a servomotor generally depends on the stability of the potentiometer and other analog components in the feedback circuit.

Stepping motors can be used in simple open-loop control systems; these are generally adequate for systems that operate at low accelerations with static loads, but closed loop control may be essential for high accelerations, particularly if they involve variable loads. If a stepper in an open-loop control system is overtorqued, all knowledge of rotor position is lost and the system must be reinitialized; servomotors are not subject to this problem.

Stepping motors are known in German as *Schrittmotoren*, in French as *moteurs pas à pas*, and in Spanish as *motor paso paso*.

1. Stepping Motor Types

Stepping motors come in two varieties, *permanent magnet* and *variable reluctance* (there are also *hybrid* motors, which are indistinguishable from permanent magnet motors from the controller's point of view). Lacking a label on the motor, you can generally tell the two apart by feel when no power is applied. Permanent magnet motors tend to "cog" as you twist the rotor with your fingers, while variable reluctance motors almost spin freely (although they may cog slightly because of residual magnetization in the rotor). You can also distinguish between the two varieties with an ohmmeter. Variable reluctance motors usually have three (sometimes four) windings, with a common return, while permanent magnet motors usually have two independent windings, with or without center taps. Center-tapped windings are used in unipolar permanent magnet motors.

Stepping motors come in a wide range of angular resolution. The coarsest motors typically turn 90 degrees per step, while high resolution permanent magnet motors are commonly able to handle 1.8 or even 0.72 degrees per step. With an appropriate controller, most permanent magnet and hybrid motors can be run in half-steps, and some controllers can handle smaller fractional steps or microsteps.

For both permanent magnet and variable reluctance stepping motors, if just one winding of the motor is energised, the rotor (under no load) will snap to a fixed angle and then hold that angle until the torque exceeds the holding torque of the motor, at which point, the rotor will turn, trying to hold at each successive equilibrium point.

Variable Reluctance Motors

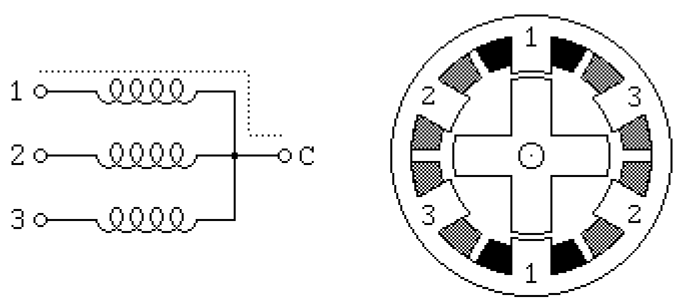

Figure 1.1

If your motor has three windings, typically connected as shown in the schematic diagram in Figure 1.1, with one terminal common to all windings, it is most likely a variable reluctance stepping motor. In use, the common wire typically goes to the positive supply and the windings are energized in sequence.

The cross section shown in Figure 1.1 is of 30 degree per step variable reluctance motor. The rotor in this motor has 4 teeth and the stator has 6 poles, with each winding wrapped around two opposite poles. With winding number 1 energised, the rotor teeth marked X are attracted to this winding's poles. If the current through winding 1 is turned off and winding 2 is turned on, the rotor will rotate 30 degrees clockwise so that the poles marked Y line up with the poles marked 2.

To rotate this motor continuously, we just apply power to the 3 windings in sequence. Assuming positive logic, where a 1 means turning on the current through a motor winding, the following control sequence will spin the motor illustrated in Figure 1.1 clockwise 24 steps or 2 revolutions:

```
  Winding 1 100100100100100100100100
  Winding 2 010010010010010010010010
  Winding 3 001001001001001001001001
                time --->
```

The section of this tutorial on [Mid-Level Control] provides details on methods for generating such sequences of control signals, while the section on [Control Circuits] discusses the power switching circuitry needed to drive the motor windings from such control sequences.

There are also variable reluctance stepping motors with 4 and 5 windings, requiring 5 or 6 wires. The principle for driving these motors is the same as that for the three winding variety, but it becomes important to work out the correct order to energise the windings to make the motor step nicely.

The motor geometry illustrated in Figure 1.1, giving 30 degrees per step, uses the fewest number of rotor teeth and stator poles that performs satisfactorily. Using more motor poles and more rotor teeth allows construction of motors with smaller step angle. Toothed faces on each pole and a correspondingly finely toothed rotor allows for step angles as small as a few degrees.

Unipolar Motors

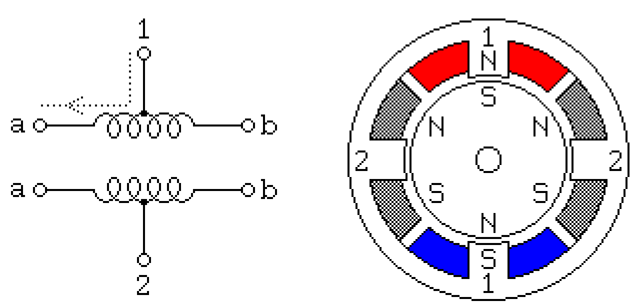

Figure 1.2

Unipolar stepping motors, both Permanent magnet and hybrid stepping motors with 5 or 6 wires are usually wired as shown in the schematic in Figure 1.2, with a center tap on each of two windings. In use, the center taps of the windings are typically wired to the positive supply, and the two ends of each winding are alternately grounded to reverse the direction of the field provided by that winding.

The motor cross section shown in Figure 1.2 is of a 30 degree per step permanent magnet or hybrid motor -- the difference between these two motor types is not relevant at this level of abstraction. Motor winding number 1 is distributed between the top and bottom stator pole, while motor winding number 2 is distributed between the left and right motor poles. The rotor is a permanent magnet with 6 poles, 3 south and 3 north, arranged around its circumfrence.

For higher angular resolutions, the rotor must have proportionally more poles. The 30 degree per step motor in the figure is one of the most common permanent magnet motor designs, although 15 and 7.5 degree per step motors are widely available. Permanent magnet motors with resolutions as good as 1.8 degrees per step are made, and hybrid motors are routinely built with 3.6 and 1.8 degrees per step, with resolutions as fine as 0.72 degrees per step available.

As shown in the figure, the current flowing from the center tap of winding 1 to terminal a causes the top stator pole to be a north pole while the bottom stator pole is a south pole. This attracts the rotor into the position shown. If the power to winding 1 is removed and winding 2 is energised, the rotor will turn 30 degrees, or one step.

To rotate the motor continuously, we just apply power to the two windings in sequence. Assuming positive logic, where a 1 means turning on the current through a motor winding, the following two control sequences will spin the motor illustrated in Figure 1.2 clockwise 24 steps or 4 revolutions:

```
Winding 1a 1000100010001000100010001
Winding 1b 0010001000100010001000100
Winding 2a 0100010001000100010001000
Winding 2b 0001000100010001000100010
            time --->

Winding 1a 1100110011001100110011001
Winding 1b 0011001100110011001100110
Winding 2a 0110011001100110011001100
Winding 2b 1001100110011001100110011
            time --->
```

Note that the two halves of each winding are never energized at the same time. Both sequences shown above will rotate a permanent magnet one step at a time. The top sequence only powers one winding at a time, as illustrated in the figure above; thus, it uses less power. The bottom sequence involves powering two windings at a time and generally produces a torque about 1.4 times greater than the top sequence while using twice as much power.

The section of this tutorial on Mid-Level Control provides details on methods for generating such sequences of control signals, while the section on Control Circuits discusses the power switching circuitry needed to drive the motor windings from such control sequences.

The step positions produced by the two sequences above are not the same; as a result, combining the two sequences allows half stepping, with the motor stopping alternately at the positions indicated by one or the other sequence. The combined sequence is as follows:

```
  Winding 1a 11000001110000011100000111
  Winding 1b 00011100000111000001110000
  Winding 2a 01110000011100000111000001
  Winding 2b 00000111000001110000011100
              time --->
```

Bipolar Motors

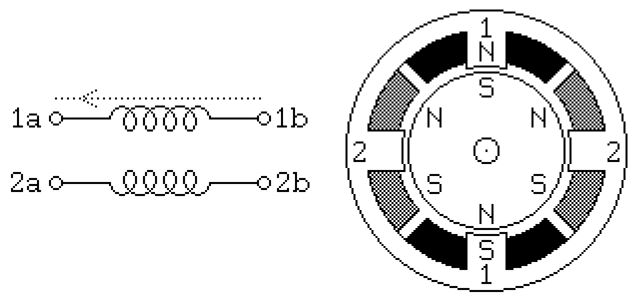

Figure 1.3

Bipolar permanent magnet and hybrid motors are constructed with exactly the same mechanism as is used on unipolar motors, but the two windings are wired more simply, with no center taps. Thus, the motor itself is simpler but the drive circuitry needed to reverse the polarity of each pair of motor poles is more complex. The schematic in Figure 1.3 shows how such a motor is wired, while the motor cross section shown here is exactly the same as the cross section shown in Figure 1.2.

The drive circuitry for such a motor requires an *H-bridge* control circuit for each winding; these are discussed in more detail in the section on Control Circuits. Briefly, an H-bridge allows the polarity of the power applied to each end of each winding to be controlled independently. The control sequences for single stepping such a motor are shown below, using + and - symbols to indicate the polarity of the power applied to each motor terminal:

```
  Terminal 1a +---+---+---+---  ++--++--++--++--
  Terminal 1b --+---+---+---+-  --++--++--++--++
  Terminal 2a -+---+---+---+--  -++--++--++--++-
  Terminal 2b ---+---+---+---+  +--++--++--++--+
              time --->
```

Note that these sequences are identical to those for a unipolar permanent magnet motor, at an abstract level, and that above the level of the H-bridge power switching electronics, the control systems for the two types of motor can be identical.

Note that many full H-bridge driver chips have one control input to enable the output and another to control the direction. Given two such bridge chips, one per winding, the following control sequences will spin the motor identically to the control sequences given above:

```
Enable    1 1010101010101010  1111111111111111
Direction 1 1x0x1x0x1x0x1x0x  1100110011001100
Enable    2 0101010101010101  1111111111111111
Direction 2 x1x0x1x0x1x0x1x0  0110011001100110
            time --->
```

To distinguish a bipolar permanent magnet motor from other 4 wire motors, measure the resistances between the different terminals. It is worth noting that some permanent magnet stepping motors have 4 independent windings, organized as two sets of two. Within each set, if the two windings are wired in series, the result can be used as a high voltage bipolar motor. If they are wired in parallel, the result can be used as a low voltage bipolar motor. If they are wired in series with a center tap, the result can be used as a low voltage unipolar motor.

2. Stepping Motor Physics

In any presentation covering the quantitative physics of a class of systems, it is important to beware of the units of measurement used! In this presentation of stepping motor physics, we will assume standard physical units:

	English	CGS	MKS
MASS	slug	gram	kilogram
FORCE	pound	dyne	newton
DISTANCE	foot	centimeter	meter
TIME	second	second	second
ANGLE	radian	radian	radian

A force of one pound will accelerate a mass of one slug at one foot per second squared. The same relationship holds between the force, mass, time and distance units of the other measurement systems. Most people prefer to measure angles in degrees, and the common engineering practice of specifying mass in pounds or force in kilograms will not yield correct results in the formulas given here! Care must be taken to convert such irregular units to one of the standard systems outlined above before applying the formulas given here!

Statics

For a motor that turns S radians per step, the plot of torque versus angular position for the rotor relative to some initial equilibrium position will generally approximate a sinusoid. The actual shape of the curve depends on the pole geometry of both rotor and stator, and neither this curve nor the geometry information is given in the motor data sheets I've seen! For permanent magnet and hybrid motors, the actual curve usually looks sinusoidal, but looks can be misleading. For variable reluctance motors, the curve rarely even looks sinusoidal; trapezoidal and even assymetrical sawtooth curves are not uncommon.

For a three-winding variable reluctance or permanent magnet motors with S radians per step, the period of the torque versus position curve will be 3S; for a 5-phase permanent magnet motor, the period will be 5S. For a two-winding permanent magnet or hybrid motor, the most common type, the period will be 4S, as illustrated in Figure 2.1:

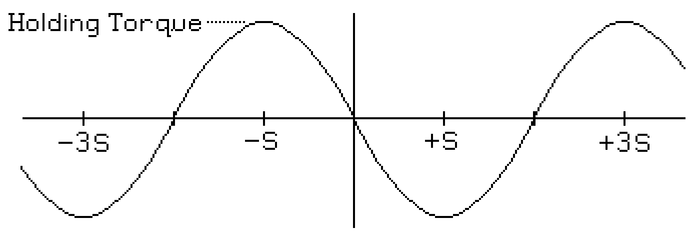

Figure 2.1

Again, for an ideal 2 winding permanent magnet motor, this can be mathematically expressed as:

$T = -h \sin(((\pi/2) / S) \Theta)$

where:

T -- torque
h -- holding torque
S -- step angle, in radians
Θ= shaft angle, in radians

But remember, subtle departures from the ideal sinusoid described here are very common.

The *single-winding holding torque* of a stepping motor is the peak value of the torque versus position curve when the maximum allowed current is flowing through one motor winding. If you attempt to apply a torque greater than this to the motor rotor while maintaining power to one winding, it will rotate freely.

It is sometimes useful to distinguish between the *electrical shaft angle* and the *mechanical shaft angle*. In the mechanical frame of reference, 2π radians is defined as one full revolution. In the electrical frame of reference, a revolution is defined as one period of the torque versus shaft angle curve. Throughout this tutorial, Θ refers to the mechanical shaft angle, and $((\pi/2)/S)\Theta$ gives the electrical angle for a motor with 4 steps per cycle of the torque curve.

Assuming that the torque versus angular position curve is a good approximation of a sinusoid, as long as the torque remains below the holding torque of the motor, the rotor will remain within 1/4 period of the equilibrium position. For a two-winding permanent magnet or hybrid motor, this means the rotor will remain within one step of the equilibrium position.

With no power to any of the motor windings, the torque does not always fall to zero! In variable reluctance stepping motors, residual magnetization in the magnetic circuits of the motor may lead to a small residual torque, and in permanent magnet and hybrid stepping motors, the combination of pole geometry and the permanently magnetized rotor may lead to significant torque with no applied power.

The residual torque in a permanent magnet or hybrid stepping motor is frequently referred to as the *cogging torque* or *detent torque* of the motor because a naive observer will frequently guess that there is a detent mechanism of some kind inside the motor. The most common motor designs yield a detent torque that varies sinusoidally with rotor angle, with an equilibrium position at every step and an amplitude of roughly 10% of the rated holding torque of the motor, but a quick survey of motors from one manufacturer (Phytron) shows values as high as 23% for one very small motor to a low of 2.6% for one mid-sized motor.

Half-Stepping and Microstepping

So long as no part of the magnetic circuit saturates, powering two motor windings simultaneously will produce a torque versus position curve that is the sum of the torque versus position curves for the two motor windings taken in isolation. For a two-winding permanent magnet or hybrid motor, the two curves will be S radians out of phase, and if the currents in the two windings are equal, the peaks and valleys of the sum will be displaced S/2 radians from the peaks of the original curves, as shown in Figure 2.2:

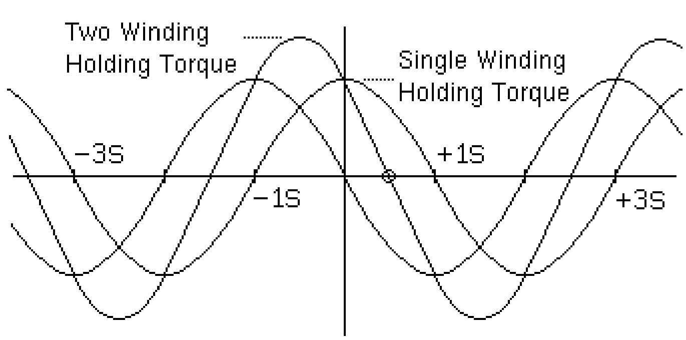

Figure 2.2

This is the basis of *half-stepping*. The *two-winding holding torque* is the peak of the composite torque curve when two windings are carrying their maximum rated current.

For common two-winding permanent magnet or hybrid stepping motors, the two-winding holding torque will be:

$h_2 = 2^{0.5}\ h_1$

where:

h_1 -- single-winding holding torque
h_2 -- two-winding holding torque

This assumes that no part of the magnetic circuit is saturated and that the torque versus position curve for each winding is an ideal sinusoid.

Most permanent-magnet and variable-reluctance stepping motor data sheets quote the two-winding holding torque and not the single-winding figure; in part, this is because it is larger, and in part, it is because the most common full-step controllers always apply power to two windings at once.

If any part of the motor's magnetic circuits is saturated, the two torque curves will not add linearly. As a result, the composite torque will be less than the sum of the component torques and the equilibrium position of the composite may not be exactly S/2 radians from the equilibria of the original.

Microstepping allows even smaller steps by using different currents through the two motor windings, as shown in Figure 2.3:

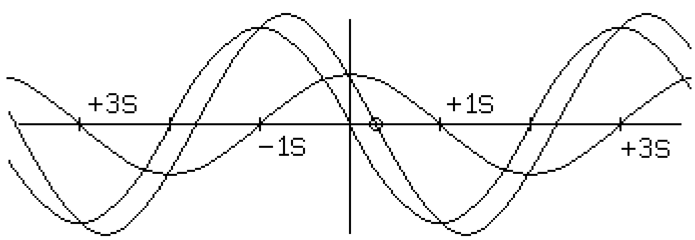

Figure 2.3

For a two-winding variable reluctance or permanent magnet motor, assuming nonsaturating magnetic circuits, and assuming perfectly sinusoidal torque versus position curves for each motor winding, the following formula gives the key characteristics of the composite torque curve:

$h = (a^2 + b^2)^{0.5}$
$x = (S / (\pi/2))\ arctan(b / a)$

where:

a -- torque applied by winding with equilibrium at 0 radians.
b -- torque applied by winding with equilibrium at S radians.
h -- holding torque of composite.

x -- equilibrium position, in radians.
S -- step angle, in radians.

In the absence of saturation, the torques *a* and *b* are directly proportional to the currents through the corresponding windings. It is quite common to work with normalized currents and torques, so that the single-winding holding torque or the maximum current allowed in one motor winding is 1.0.

Friction and the Dead Zone

The torque versus position curve shown in Figure 2.1 does not take into account the torque the motor must exert to overcome friction! Note that frictional forces may be divided into two large categories, static or sliding friction, which requires a constant torque to overcome, regardless of velocity, and dynamic friction or viscous drag, which offers a resistance that varies with velocity. Here, we are concerned with the impact of static friction. Suppose the torque needed to overcome the static friction on the driven system is 1/2 the peak torque of the motor, as illustrated in Figure 2.4.

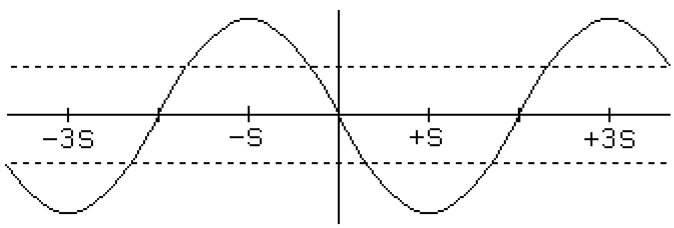

Figure 2.4

The dotted lines in Figure 2.4 show the torque needed to overcome friction; only that part of the torque curve outside the dotted lines is available to move the rotor. The curve showing the available torque as a function of shaft angle is the difference between these curves, as shown in Figure 2.5:

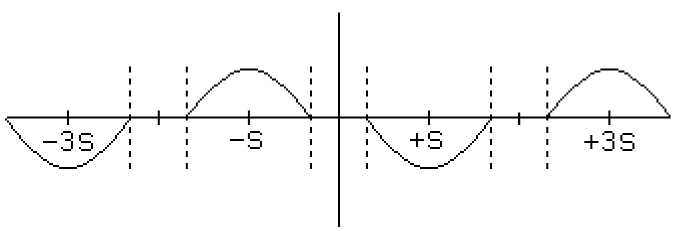

Figure 2.5

Note that the consequences of static friction are twofold. First, the total torque available to move the load is reduced, and second, there is a *dead zone* about each of the equilibria of the ideal motor. If the motor rotor is positioned anywhere within the dead zone for the current equilibrium position, the frictional torque will exceed the torque applied by the motor windings, and the rotor will not move. Assuming an ideal sinusoidal torque versus position curve in the absence of friction, the angular width of these dead zones will be:

$d = 2 (S / (\pi/2)) \, arcsin(f / h) = (S / (\pi/4)) \, arcsin(f / h)$

where:

d -- width of dead zone, in radians
S -- step angle, in radians
f -- torque needed to overcome static friction
h -- holding torque

The important thing to note about the dead zone is that it limits the ultimate positioning accuracy! For the example, where the static friction is 1/2 the peak torque, a 90° per step motor will have dead-zones 60° wide! That means that successive steps may be as large as 150° and as small as 30°, depending on where in the dead zone the rotor stops after each step!

The presence of a dead zone has a significant impact on the utility of microstepping! If the dead zone is x° wide, then microstepping with a step size smaller than x° may not move the rotor at all. Thus, for systems intended to use high resolution microstepping, it is very important to minimize static friction.

Dynamics

Each time you step the motor, you electronically move the equilibrium position S radians. This moves the entire curve illustrated in Figure 2.1 a distance of S radians, as shown in Figure 2.6:

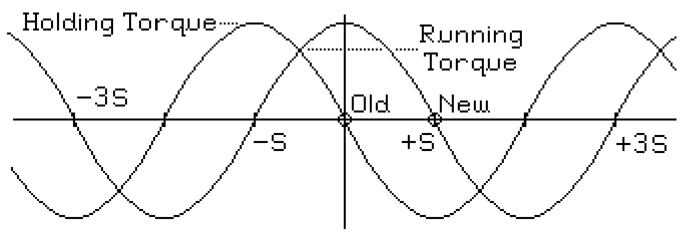

Figure 2.6

The first thing to note about the process of taking one step is that the maximum available torque is at a minimum when the rotor is halfway from one step to the next. This minimum determines the *running torque*, the maximum torque the motor can drive as it steps slowly forward. For common two-winding permanent magnet motors with ideal sinusoidal torque versus position curves and holding torque h, this will be $h/(2^{0.5})$. If the motor is stepped by powering two windings at a time, the running torque of an ideal two-winding permanent magnet motor will be the same as the single-winding holding torque.

It shoud be noted that at higher stepping speeds, the running torque is sometimes defined as the *pull-out torque*. That is, it is the maximum frictional torque the motor can overcome on a rotating load before the load is pulled out of step by the friction. Some motor data sheets define a second torque figure, the *pull-in torque*. This is the maximum frictional torque that the motor can overcome to accelerate a stopped load to synchronous speed. The pull-in torques documented on stepping motor data sheets are of questionable value because the pull-in torque depends on the moment of inertia of the load used when they were measured, and few motor data sheets document this!

In practice, there is always some friction, so after the equilibrium position moves one step, the rotor is likely to oscillate briefly about the new equilibrium position. The resulting trajectory may resemble the one shown in Figure 2.7:

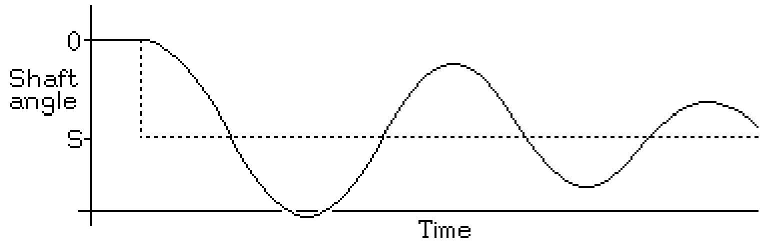

Figure 2.7

Here, the trajectory of the equilibrium position is shown as a dotted line, while the solid curve shows the trajectory of the motor rotor.

Resonance

The resonant frequency of the motor rotor depends on the amplitude of the oscillation; but as the amplitude decreases, the resonant frequency rises to a well-defined small-amplitude frequency. This frequency depends on the step angle and on the ratio of the holding torque to the moment of inertia of the rotor. Either a higher torque or a lower moment will increase the frequency!

Formally, the small-amplitude resonance can be computed as follows: First, recall Newton's law for angular acceleration:

$T = \mu A$

where:

T -- torque applied to rotor
µ -- moment of inertia of rotor and load
A -- angular acceleration, in radians per second per second

We assume that, for small amplitudes, the torque on the rotor can be approximated as a linear function of the displacement from the equilibrium position. Therefore, Hooke's law applies:

$T = -k \Theta$

where:

k -- the "spring constant" of the system, in torque units per radian
Θ -- angular position of rotor, in radians

We can equate the two formulas for the torque to get:

$\mu A = -k \Theta$

Note that acceleration is the second derivitive of position with respect to time:

$A = d^2\Theta/dt^2$

so we can rewrite this the above in differential equation form:

$d^2\Theta/dt^2 = -(k/\mu)\ \Theta$

To solve this, recall that, for:

$f(\ t\ ) = a\ sin\ bt$

The derivitives are:

$df(\ t\ )/dt = ab\ cos\ bt$
$d^2f(\ t\ )/dt^2 = -ab^2\ sin\ bt = -b^2\ f(t)$

Note that, throughout this discussion, we assumed that the rotor is resonating. Therefore, it has an equation of motion something like:

$\Theta = a\ sin\ (2\pi\ f\ t)$
a = angular amplitude of resonance
f = resonant frequency

This is an admissable solution to the above differential equation if we agree that:

$b = 2\pi\ f$
$b^2 = k/\mu$

Solving for the resonant frequency f as a function of k and μ, we get:

$f = (\ k/\mu\ )^{0.5} / 2\pi$

It is crucial to note that it is the moment of inertia of the rotor plus any coupled load that matters. The moment of the rotor, in isolation, is irrelevant! Some motor data sheets include information on resonance, but if any load is coupled to the rotor, the resonant frequency will change!

In practice, this oscillation can cause significant problems when the stepping rate is anywhere near a resonant frequency of the system; the result frequently appears as random and uncontrollable motion.

Resonance and the Ideal Motor

Up to this point, we have dealt only with the small-angle spring constant k for the system. This can be measured experimentally, but if the motor's torque versus position curve is sinusoidal, it is also a simple function of the motor's holding torque. Recall that:

$T = -h\ sin(\ ((\pi/2)/S)\ \Theta)$

The small angle spring constant k is the negative derivitive of T at the origin.

$k = -dT / d\Theta = -(-h((\pi/2)/S)\cos(0)) = (\pi/2)(h / S)$

Substituting this into the formula for frequency, we get:

$f = ((\pi/2)(h / S) / \mu)^{0.5} / 2\pi = (h / (8\pi \mu S))^{0.5}$

Given that the holding torque and resonant frequency of the system are easily measured, the easiest way to determine the moment of inertia of the moving parts in a system driven by a stepping motor is indirectly from the above relationship!

$\mu = h / (8\pi f^2 S)$

For practical purposes, it is usually not the torque or the moment of inertia that matters, but rather, the maximum sustainable acceleration that matters! Conveniently, this is a simple function of the resonant frequency! Starting with the Newton's law for angular acceleration:

$A = T / \mu$

We can substitute the above formula for the moment of inertia as a function of resonant frequency, and then substitute the maximum sustainable running torque as a function of the holding torque to get:

$A = (h / (2^{0.5})) / (h / (8\pi f^2 S)) = 8\pi S f^2 / (2^{0.5})$

Measuring acceleration in steps per second squared instead of in radians per second squared, this simplifies to:

$A_{steps} = A / S = 8\pi f^2 / (2^{0.5})$

Thus, for an ideal motor with a sinusoidal torque versus rotor position function, the maximum acceleration in steps per second squared is a trivial function of the resonant frequency of the motor and rigidly coupled load!

For a two-winding permanent-magnet or variable-reluctance motor, with an ideal sinusoidal torque-versus-position characteristic, the two-winding holding torque is a simple function of the single-winding holding torque:

$h_2 = 2^{0.5} h_1$

where:

h_1 -- single-winding holding torque
h_2 -- two-winding holding torque

Substituting this into the formula for resonant frequency, we can find the ratios of the resonant frequencies in these two operating modes:

$f_1 = (h_1 / ...)^{0.5}$
$f_2 = (h_2 / ...)^{0.5} = (2^{0.5} h_1 / ...)^{0.5} = 2^{0.25} (h_1 / ...)^{0.5} = 2^{0.25} f_1 = 1.189... f_1$

This relationship only holds if the torque provided by the motor does not vary appreciably as the stepping rate varies between these two frequencies.

In general, as will be discussed later, the available torque will tend to remain relatively constant up until some cutoff stepping rate, and then it will fall. Therefore, this relationship only holds if the resonant frequencies are below this cutoff stepping rate. At stepping rates above the cutoff rate, the two frequencies will be closer to each other!

Living with Resonance

If a rigidly mounted stepping motor is rigidly coupled to a frictionless load and then stepped at a frequency near the resonant frequency, energy will be pumped into the resonant system, and the result of this is that the motor will literally lose control. There are three basic ways to deal with this problem:

Controlling resonance in the mechanism

Use of elastomeric motor mounts or elastomeric couplings between motor and load can drain energy out of the resonant system, preventing energy from accumulating to the extent that it allows the motor rotor to escape from control.

Or, viscous damping can be used. Here, the damping will not only draw energy out of the resonant modes of the system, but it will also subtract from the total torque available at higher speeds. Magnetic eddy current damping is equivalent to viscous damping for our purposes.

Figure 2.8 illustrates the use of elastomeric couplings and viscous damping in two typical stepping motor applications, one using a lead screw to drive a load, and the other using a tendon drive:

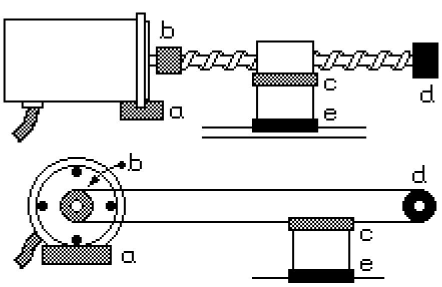

Figure 2.8

In Figure 2.8, elastomeric moter mounts are shown at a and elastomeric couplings between the motor and load are shown at b and c. The end bearing for the lead screw or tendon, at d, offers an opportunity for viscous damping, as do the ways on which the load slides, at e. Even the friction found in sealed ballbearings or teflon on steel ways can provide enough damping to prevent resonance problems.

Controlling resonance in the low-level drive circuitry

A resonating motor rotor will induce an alternating current voltage in the motor windings. If some motor winding is not currently being driven, shorting this winding will impose a drag on the motor rotor that is exactly equivalent to using a magnetic eddy current damper.

If some motor winding is currently being driven, the AC voltage induced by the resonance will tend to modulate the current through the winding. Clamping the motor current with an external inductor will counteract the resonance. Schemes based on this idea are incorporated into some of the drive circuits illustrated in later sections of this tutorial.

Controlling resonance in the high-level control system

The high level control system can avoid driving the motor at known resonant frequencies, accelerating and decelerating through these frequencies and never attempting sustained rotation at these speeds.

Recall that the resonant frequency of a motor in half-stepped mode will vary by up to 20% from one half-step to the next. As a result, half-stepping pumps energy into the resonant system less efficiently than full stepping. Furthermore, when operating near these resonant frequencies, the motor control system may preferentially use only the two-winding half steps when operating near the single-winding resonant frequency, and only the single-winding half steps when operating near the two-winding resonant frequency. Figure 2.9 illustrates this:

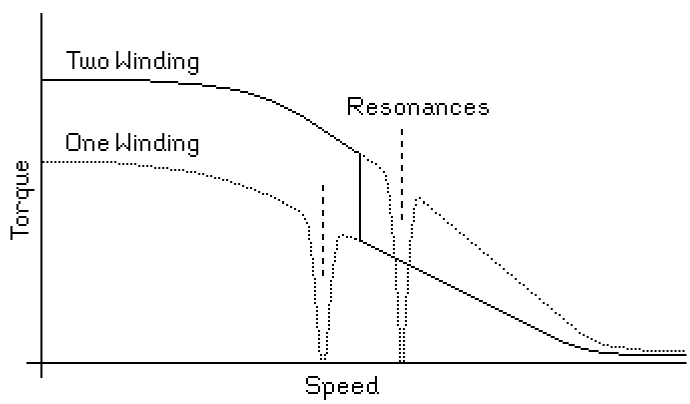

Figure 2.9

The darkened curve in Figure 2.9 shows the operating torque achieved by a simple control scheme that delivers useful torque over a wide range of speeds despite the fact that the available torque drops to zero at each resonance in the system. This solution is particularly effective if the resonant frequencies are sharply defined and well separated. This will be the case in minimally damped systems operating well below the cutoff speed defined in the next section.

Torque versus Speed

An important consideration in designing high-speed stepping motor controllers is the effect of the inductance of the motor windings. As with the torque versus angular position information, this is frequently poorly documented in motor data sheets, and indeed, for variable reluctance stepping motors, it is not a constant! The inductance of the motor winding determines the rise and fall time of the current through the windings. While we might hope for a square-wave plot of current versus time, the inductance forces an exponential, as illustrated in Figure 2.10:

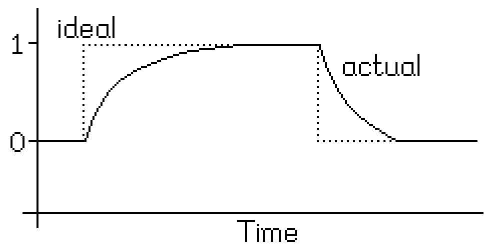

Figure 2.10

The details of the current-versus-time function through each winding depend as much on the drive circuitry as they do on the motor itself! It is quite common for the time constants of these exponentials to differ. The rise time is determined by the drive voltage and drive circuitry, while the fall time depends on the circuitry used to dissipate the stored energy in the motor winding.

At low stepping rates, the rise and fall times of the current through the motor windings has little effect on the motor's performance, but at higher speeds, the effect of the inductance of the motor windings is to reduce the available torque, as shown in Figure 2.11:

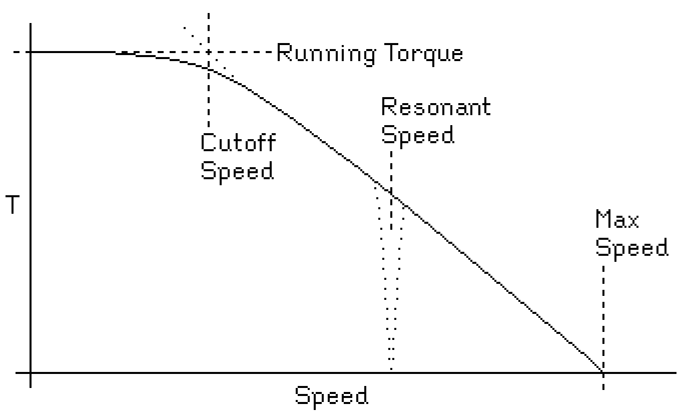

Figure 2.11

The motor's *maximum speed* is defined as the speed at which the available torque falls to zero. Measuring maximum speed can be difficult when there are resonance problems, because these cause the torque to drop to zero prematurely. The *cutoff speed* is the speed above which the torque begins to fall. When the motor is operating below its cutoff speed, the rise and fall times of the current through the motor windings occupy an insignificant fraction of each step, while at the cutoff speed, the step duration is comparable to the sum of the rise and fall times. Note that a sharp cutoff is rare, and therefore, statements of a motor's cutoff speed are, of necessity, approximate.

The details of the torque versus speed relationship depend on the details of the rise and fall times in the motor windings, and these depend on the motor control system as well as the motor. Therefore, the cutoff speed and maximum speed for any particular motor depend, in part, on the control system! The torque versus speed curves published in motor data sheets occasionally come with documentation of the motor controller used to obtain that curve, but this is far from universal practice!

Similarly, the resonant speed depends on the moment of inertia of the entire rotating system, not just the motor rotor, and the extent to which the torque drops at resonance depends on the presence of mechanical damping and on the nature of the control system. Some published torque versus speed curves show very clear resonances without documenting the moment of inertia of the hardware that may have been attached to the motor shaft in order to make torque measurements.

The torque versus speed curve shown in Figure 2.11 is typical of the simplest of control systems. More complex control systems sometimes introduce electronic resonances that act to increase the available torque above the motor's low-speed torque. A common result of this is a peak in the available torque near the cutoff speed.

Electromagnetic Issues

In a permanent magnet or hybrid stepping motor, the magnetic field of the motor rotor changes with changes in shaft angle. The result of this is that turning the motor rotor induces an AC voltage in each motor winding. This is referred to as the *counter EMF* because the voltage induced in each motor winding is always in phase with and counter to the ideal waveform required to turn the motor in the same direction. Both the frequency and amplitude of the counter EMF increase with rotor speed, and therefore, counter EMF contributes to the decline in torque with increased stepping rate.

Variable reluctance stepping motors also induce counter EMF! This is because, as the stator winding pulls a tooth of the rotor towards its equilibrium position, the reluctance of the magnetic circuit declines. This decline increases the inductance of the stator winding, and this change in inductance demands a decrease in the current through the winding in order to conserve energy. This decrease is evidenced as a counter EMF.

The reactance (inductance and resistance) of the motor windings limits the current flowing through them. Thus, by ohms law, increasing the voltage will increase the current, and therefore increase the available torque. The increased voltage also serves to overcome the counter EMF induced in the motor windings, but the voltage cannot be increased arbitrarily! Thermal, magnetic and electronic considerations all serve to limit the useful torque that a motor can produce.

The heat given off by the motor windings is due to both simple resistive losses, eddy current losses, and hysteresis losses. If this heat is not conducted away from the motor adequately, the motor windings will overheat. The simplest failure this can cause is insulation breakdown, but it can also heat a permanent magnet rotor to above its curie temperature, the temperature at which permanent magnets lose their magnetization. This is a particular risk with many modern high strength magnetic alloys.

Even if the motor is attached to an adequate heat sink, increased drive voltage will not necessarily lead to increased torque. Most motors are designed so that, with the rated current flowing through the windings, the magnetic circuits of the motor are near saturation. Increased current will not lead to an appreciably increased magnetic field in such a motor!

Given a drive system that limits the current through each motor winding to the rated maximum for that winding, but uses high voltages to achieve a higher cutoff torque and higher torques above cutoff, there are other limits that come into play. At high speeds, the motor windings must, of necessity, carry high frequency AC signals. This leads to eddy current losses in the magnetic circuits of the motor, and it leads to skin effect losses in the motor windings.

Motors designed for very high speed running should, therefore, have magnetic structures using very thin laminations or even nonconductive ferrite materials, and they should have small gauge wire in their windings to minimize skin effect losses. Common high torque motors have large-gauge motor windings and coarse core laminations, and at high speeds, such motors can easily overheat and should therefore be derated accordingly for high speed running!

It is also worth noting that the best way to demagnetize something is to expose it to a high frequency-high amplitude magnetic field. Running the control system to spin the rotor at high speed when the rotor is actually stalled, or spinning the rotor at high speed against a control system trying to hold the rotor in a fixed position will both expose the rotor to a high amplitude high-frequency field. If such operating conditions are common, particularly if the motor is run near the curie temperature of the permanent magnets, demagnetization is a serious risk and the field strengths (and expected torques) should be reduced accordingly!

3. Basic Stepping Motor Control Circuits

This section of the stepper tutorial deals with the basic final stage drive circuitry for stepping motors. This circuitry is centered on a single issue, switching the current in each motor winding on and off, and controlling its direction. The circuitry discussed in this section is connected directly to the motor windings and the motor power supply, and this circuitry is controlled by a digital system that determines when the switches are turned on or off.

This section covers all types of motors, from the elementary circuitry needed to control a variable reluctance motor, to the H-bridge circuitry needed to control a bipolar permanent magnet motor. Each class of drive circuit is illustrated with practical examples, but these examples are not intended as an exhaustive catalog of the commercially available control circuits, nor is the information given here intended to substitute for the information found on the manufacturer's component data sheets for the parts mentioned.

This section only covers the most elementary control circuitry for each class of motor. All of these circuits assume that the motor power supply provides a drive voltage no greater than the motor's rated voltage, and this significantly limits motor performance. The next section, on current limited drive circuitry, covers practical high-performance drive circuits.

Variable Reluctance Motors

Typical controllers for variable reluctance stepping motors are variations on the outline shown in Figure 3.1:

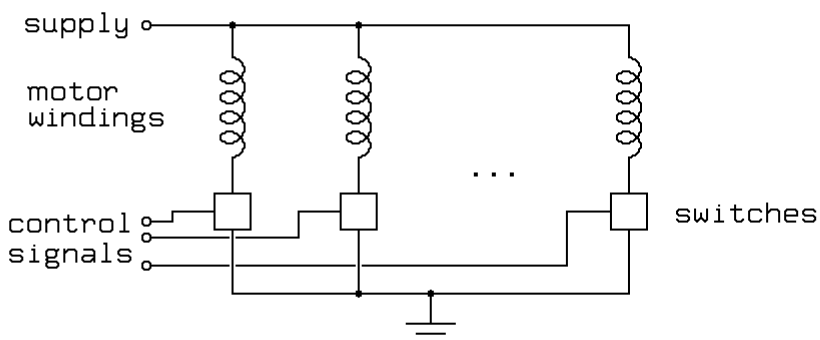

Figure 3.1

In Figure 3.1, boxes are used to represent switches; a control unit, not shown, is responsible for providing the control signals to open and close the switches at the appropriate times in order to spin the motors. In many cases, the control unit will be a computer or programmable interface controller, with software directly generating the outputs needed to control the switches, but in other cases, additional control circuitry is introduced, sometimes gratuitously!

Motor windings, solenoids and similar devices are all inductive loads. As such, the current through the motor winding cannot be turned on or off instantaneously without involving infinite voltages! When the switch controlling a motor winding is closed, allowing current to flow, the result of this is a slow rise in current. When the switch controlling a motor winding is opened, the result of this is a voltage spike that can seriously damage the switch unless care is taken to deal with it appropriately.

There are two basic ways of dealing with this voltage spike. One is to bridge the motor winding with a diode, and the other is to bridge the motor winding with a capacitor. Figure 3.2 illustrates both approaches:

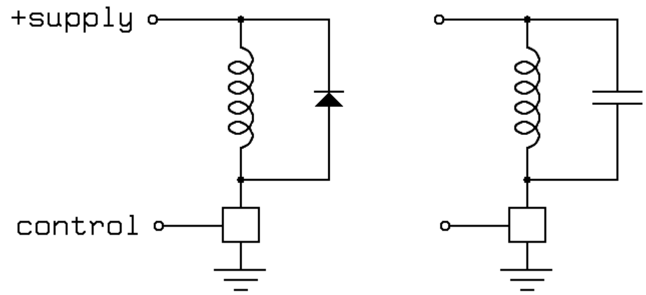

Figure 3.2

The diode shown in Figure 3.2 must be able to conduct the full current through the motor winding, but it will only conduct briefly each time the switch is turned off, as the current through the winding decays. If relatively slow diodes such as the common 1N400X family are used together with a fast switch, it may be necessary to add a small capacitor in parallel with the diode.

The capacitor shown in Figure 3.2 poses more complex design problems! When the switch is closed, the capacitor will discharge through the switch to ground, and the switch must be able to handle this brief spike of discharge current. A resistor in series with the capacitor or in series with the power supply will limit this current. When the switch is opened, the stored energy in the motor winding will charge the capacitor up to a voltage significantly above the supply voltage, and the switch must be able to tolerate this voltage. To solve for the size of the capacitor, we equate the two formulas for the stored energy in a resonant circuit:

$P = C V^2 / 2$
$P = L I^2 / 2$

where:

P -- stored energy, in watt seconds or coulomb volts
C -- capacity, in farads
V -- voltage across capacitor
L -- inductance of motor winding, in henrys
I -- current through motor winding

Solving for the minimum size of capacitor required to prevent overvoltage on the switch is fairly easy:

$C > L I^2 / (V_b - V_s)^2$

where:

V_b -- the breakdown voltage of the switch
V_s -- the supply voltage

Variable reluctance motors have variable inductance that depends on the shaft angle. Therefore, worst-case design must be used to select the capacitor. Furthermore, motor inductances are frequently poorly documented, if at all.

The capacitor and motor winding, in combination, form a resonant circuit. If the control system drives the motor at frequencies near the resonant frequency of this circuit, the motor current through the motor windings, and therefore, the torque exerted by the motor, will be quite different from the steady state torque at the nominal operating voltage! The resonant frequency is:

$$f = 1 / (2\pi (L C)^{0.5})$$

Again, the electrical resonant frequency for a variable reluctance motor will depend on shaft angle! When a variable reluctance motors is operated with the exciting pulses near resonance, the oscillating current in the motor winding will lead to a magnetic field that goes to zero at twice the resonant frequency, and this can severely reduce the available torque!

Unipolar Permanent Magnet and Hybrid Motors

Typical controllers for unipolar stepping motors are variations on the outline shown in Figure 3.3:

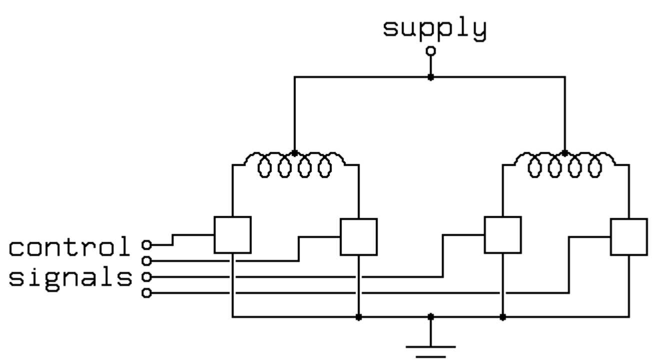

Figure 3.3

In Figure 3.3, as in Figure 3.1, boxes are used to represent switches; a control unit, not shown, is responsible for providing the control signals to open and close the switches at the appropriate times in order to spin the motors. The control unit is commonly a computer or programmable interface controller, with software directly generating the outputs needed to control the switches.

As with drive circuitry for variable reluctance motors, we must deal with the inductive kick produced when each of these switches is turned off. Again, we may shunt the inductive kick using diodes, but now, 4 diodes are required, as shown in Figure 3.4:

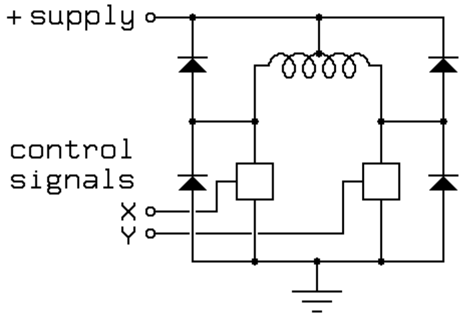

Figure 3.4

The extra diodes are required because the motor winding is not two independent inductors, it is a single center-tapped inductor with the center tap at a fixed voltage. This acts as an autotransformer! When one end of the motor winding is pulled down, the other end will fly up, and visa versa. When a switch opens, the inductive kickback will drive that end of the motor winding to the positive supply, where it is clamped by the diode. The opposite end will fly downward, and if it was not floating at the supply voltage at the time, it will fall below ground, reversing the voltage across the switch at that end. Some switches are immune to such reversals, but others can be seriously damaged.

A capacitor may also be used to limit the kickback voltage, as shown in Figure 3.5:

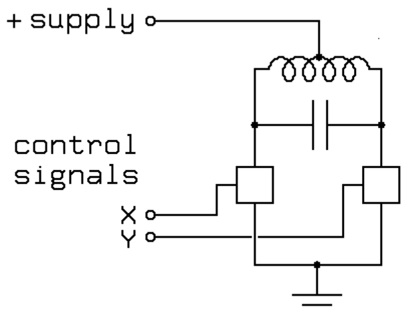

Figure 3.5

The rules for sizing the capacitor shown in Figure 3.5 are the same as the rules for sizing the capacitor shown in Figure 3.2, but the effect of resonance is quite different! With a permanent magnet motor, if the capacitor is driven at or near the resonant frequency, the torque will increase to as much as twice the low-speed torque! The resulting torque versus speed curve may be quite complex, as illustrated in Figure 3.6.

Figure 3.6 shows a peak in the available torque at the electrical resonant frequency, and a valley at the mechanical resonant frequency. If the electrical resonant frequency is placed appropriately above what would have been the cutoff speed for the motor using a diode-based driver, the effect can be a considerable increase in the effective cutoff speed.

The mechanical resonant frequency depends on the torque, so if the mechanical resonant frequency is anywhere near the electrical resonance, it will be shifted by the electrical resonance! Furthermore, the width of the mechanical resonance depends on the local slope of the torque versus speed curve; if the torque drops with speed, the mechanical resonance will be sharper, while if the torque climbs with speed, it will be broader or even split into multiple resonant frequencies.

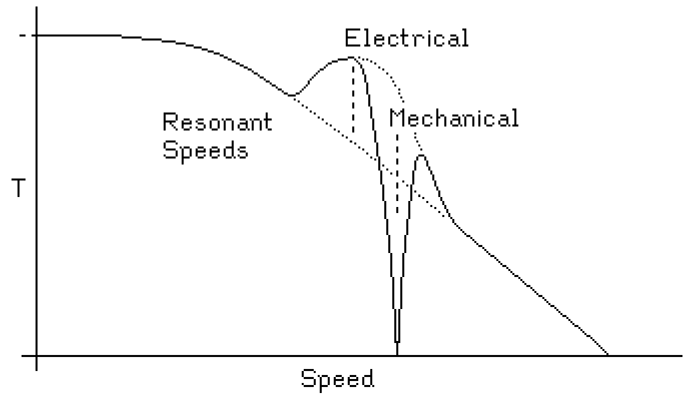

Figure 3.6

Bipolar Motors and H-Bridges

Things are more complex for bipolar permanent magnet stepping motors because these have no center taps on their windings. Therefore, to reverse the direction of the field produced by a motor winding, we need to reverse the current through the winding. We could use a double-pole double throw switch to do this electromechanically; the electronic equivalent of such a switch is called an H-bridge and is outlined in Figure 3.9:

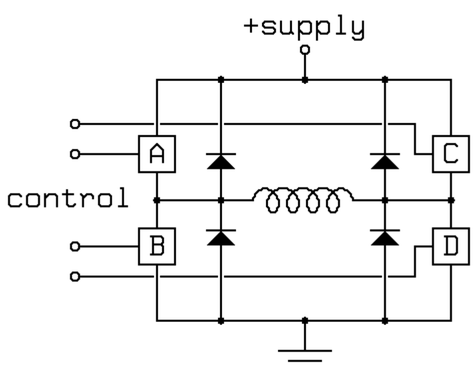

Figure 3.9

As with the unipolar drive circuits discussed previously, the switches used in the H-bridge must be protected from the voltage spikes caused by turning the power off in a motor winding. This is usually done with diodes, as shown in Figure 3.9.

It is worth noting that H-bridges are applicable not only to the control of bipolar stepping motors, but also to the control of DC motors, push-pull solenoids (those with permanent magnet plungers) and many other applications.

With 4 switches, the basic H-bridge offers 16 possible operating modes, 7 of which short out the power supply! The following operating modes are of interest:

Forward mode, switches A and D closed.
Reverse mode, switches B and C closed.

These are the usual operating modes, allowing current to flow from the supply, through the motor winding and onward to ground. Figure 3.10 illustrates forward mode:

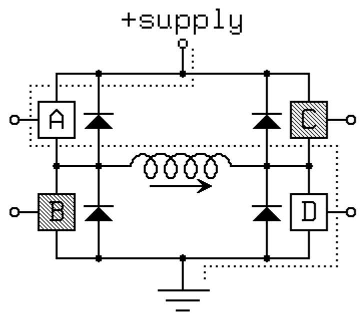

Figure 3.10

Fast decay mode or *coasting mode*, all switches open.
Any current flowing through the motor winding will be working against the full supply voltage, plus two diode drops, so current will decay quickly. This mode provides little or no dynamic braking effect on the motor rotor, so the rotor will coast freely if all motor windings are powered in this mode. Figure 3.11 illustrates the current flow immediately after switching from forward running mode to fast decay mode.

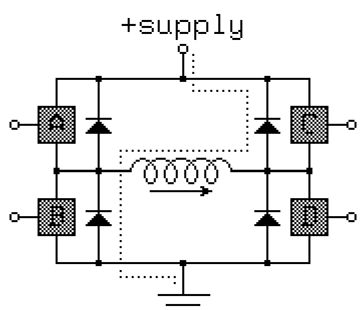

Figure 3.11

Slow decay modes or *dynamic braking modes*.
In these modes, current may recirculate through the motor winding with minimum resistance. As a result, if current is flowing in a motor winding when one of these modes is entered, the current will decay slowly, and if the motor rotor is turning, it will induce a current that will act as a brake on the rotor. Figure 3.12 illustrates one of the many useful slow-decay modes, with switch D closed; if the motor winding has recently been in forward running mode, the state of switch B may be either open or closed:

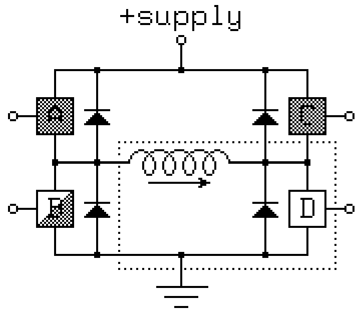

Figure 3.12

Most H-bridges are designed so that the logic necessary to prevent a short circuit is included at a very low level in the design. Figure 3.13 illustrates what is probably the best arrangement:

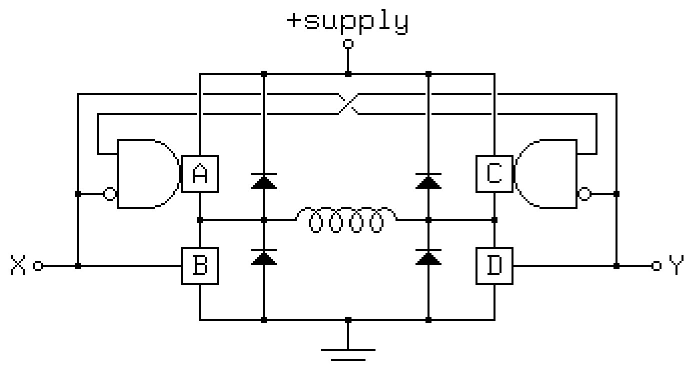

Figure 3.13

Here, the following operating modes are available:

XY	ABCD	Mode
00	0000	fast decay
01	1001	**forward**
10	0110	**reverse**
11	0101	slow decay

The advantage of this arrangement is that all of the useful operating modes are preserved, and they are encoded with a minimum number of bits; the latter is important when using a microcontroller or computer system to drive the H-bridge because many such systems have only limited numbers of bits available for parallel output. Sadly, few of the integrated H-bridge chips on the market have such a simple control scheme.

4. Current Limiting for Stepping Motors

Small stepping motors, such as those used for head positioning on floppy disk drives, are usually driven at a low DC voltage, and the current through the motor windings is usually limited by the internal resistance of the winding. High torque motors, on the other hand, are frequently built with very low resistance windings; when driven by any reasonable supply voltage, these motors typically require external current limiting circuitry.

There is good reason to run a stepping motor at a supply voltage above that needed to push the maximum rated current through the motor windings. Running a motor at higher voltages leads to a faster rise in the current through the windings when they are turned on, and this, in turn, leads to a higher cutoff speed for the motor and higher torques at speeds above the cutoff.

Microstepping, where the control system positions the motor rotor between half steps, also requires external current limiting circuitry. For example, to position the rotor 1/4 of the way from one step to another, it might be necessary to run one motor winding at full current while the other is run at approximately 1/3 of that current.

The remainder of this section discusses various circuits for limiting the current through the windings of a stepping motor, starting with simple resistive limiters and moving up to choppers and other switching regulators. Most of these current limiters are appropriate for many other applications, including limiting the current through conventional DC motors and other inductive loads.

Resistive Current Limiters

The easiest to understand current limiter is a series resistor. Most motor manufacturers recommended this approach in their literature up until the early 1980's, and most motor data sheets still give performance curves for motors driven by such circuits. The typical circuits used to control the current through one winding of a permanent magnet or hybrid motor are shown in Figure 4.1.

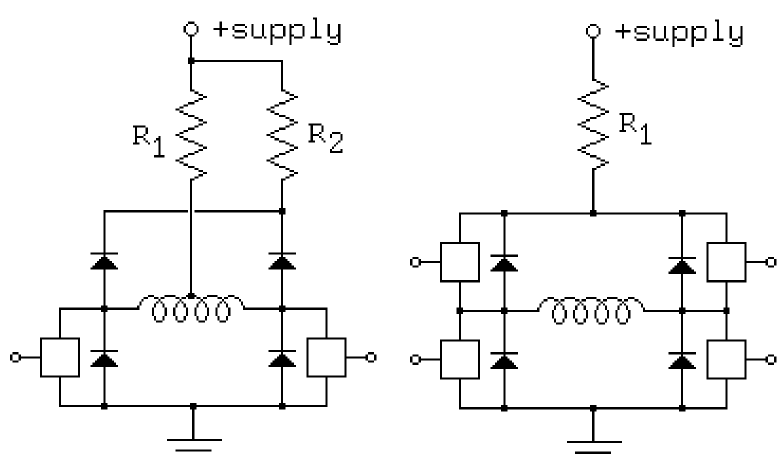

Figure 4.1

R_1 in this figure limits the current through the motor winding. Given a rated current of I and a motor winding with a resistance R_w, Ohm's law sets the maximum supply voltage as $I(R_w+R_1)$. Given that the inductance of the motor motor winding is L_w, the time constant for the motor winding will be $L_w/(R_w+R_1)$. Figure 4.2 illustrates the effect of increasing the resistance and the operating voltage on the rise and fall times of the current through one winding of a stepping motor.

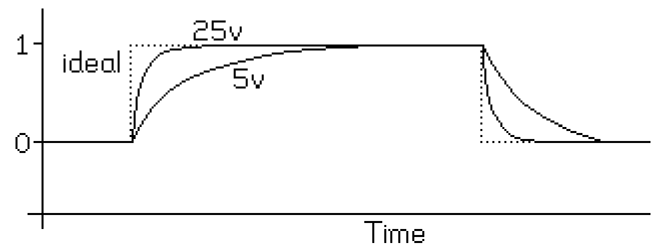

Figure 4.2

R_2 is shown only in the unipolar example in Figure 4.1 because it is particularly useful there. For a bipolar H-bridge drive, when all switches are turned off, current flows from ground to the motor supply through R_1, so the current through the motor winding will decay quite quickly. In the unipolar case, R_2 is necessary to equal this performance.

Note: When the switches in the H-bridge circuit shown in Figure 4.1 are opened, the direction of current flow through R_1 will reverse almost instantaneously! If R_1 has any inductance, for example, if it is wire-wound, it must either be bypassed with a capacitor to handle the voltage kick caused by this current reversal, or R_2 must be added to the H-bridge.

Given the rated maximum current through each winding and the supply voltage, the resistance and wattage of R_1 is easy to compute. R_2 if it is included, poses more interesting problems. The resistance of R_2 depends on the maximum voltage the switches can handle. For example, if the supply voltage is 24 volts, and the switches are rated at 75 volts, the drop across R_2 can be as much as 51 volts without harming the transistors. Given an operating current of 1.5 amps, R_2 can be a 34 ohm resistor. Note that an interesting alternative is to use a zener diode in place of R_2.

Figuring the peak average power R_2 must dissipate is a wonderful exercise in dynamics; the inductance of the motor windings is frequently undocumented and may vary with the rotor position. The power dissapated in R_2 also depends on the control system. The worst case occurs when the control system chops the power to one winding at a high enough frequency that the current through the motor winding is effectively constant; the maximum power is then a function of the duty cycle of the chopper and the ratios of the resistances in the circuit during the on and off phases of the chopper. Under normal operating conditions, the peak power dissipation will be significantly lower.

Linear Current Limiters

A pair of high wattage power resistors can cost more than a pair of power transistors plus a heat sink, particularly if forced air cooling is available. Furthermore, a transistorized constant current source, as shown in Figure 4.3, will give faster rise times through the motor windings than the current limiting resistor shown in Figure 4.1. This is because a current source will deliver the full supply voltage across the motor winding until the current reaches the rated current; only then will the current source drop the voltage.

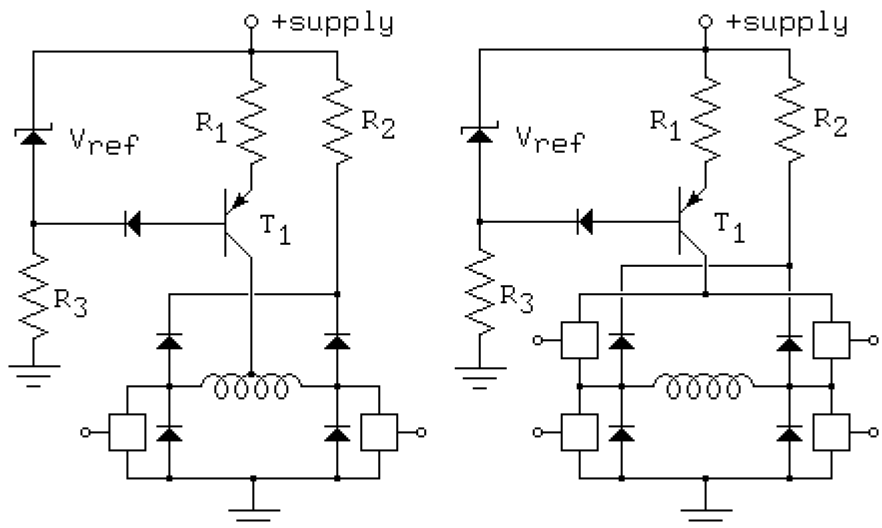

Figure 4.3

In Figure 4.3, a transistorized current source (T_1 plus R_1) has been substituted for the current limiting resistor R_1 used in the examples in Figure 4.1. The regulated voltage supplied to the base of T_1 serves to regulate the voltage across the sense resistor R_1, and this, in turn, maintains a constant current through R_1 so long as any current is allowed to flow through the motor winding.

Typically, R_1 will have as low a resistance as possible, in order to avoid the high cost of a power resistor. For example, if the forward voltage drops across the diode in series with the base T_1 and V_{BE} for T_1 are both 0.65 volts, and if a 3.3 volt zener diode is used for a reference, the voltage across R_1 will be maintained at about 2.0 volts, so if R_1 is 2 ohms, this circuit will limit the current to 1 amp, and R_1 must be able to handle 2 watts.

R_3 in Figure 4.3 must be sized in terms of the current gain of T_1 so that sufficient current flows through R_1 and R_3 to allow T_1 to conduct the full rated motor current.

The transistor T_1 used as a current regulator in Figure 4.3 is run in linear mode, and therfore, it must dissapate quite a bit of power. For example, if the motor windings have a resistance of 5 ohms and a rated current of 1 amp, and a 25 volt power supply is used, T_1 plus R_1 will dissapate, between them, 20 watts! The circuits discussed in the following sections avoid this waste of power while retaining the performance advantages of the circuit given here.

When an H-bridge bipolar drive is used with a resistive current limiter, as shown in Figure 4.1, the resistor R_2 was not needed because current could flow backwards through R_1. When a transistorized current limiter is used, current cannot flow backwards through T_1, so a separate current path back to the positive supply must be provided to handle the decaying current through the motor windings when the switches are opened. R_2 serves this purpose here, but a zener diode may be substituted to provide even faster turn-off.

The performance of a motor run with a current limited power supply is noticably better than the performance of the same motor run with a resistively limited supply, as illustrated in Figure 4.4:

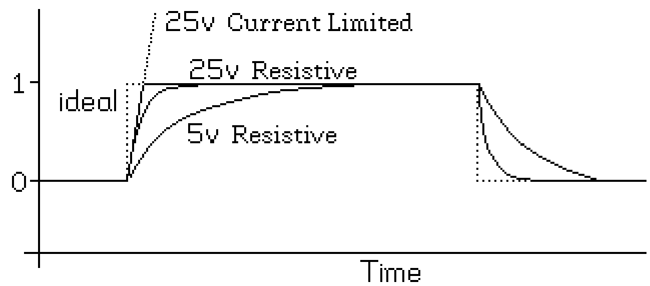

Figure 4.4

With either a current limited supply or a resistive current limiter, the initial rate of increase of the current through the inductive motor winding when the power is turned on depends only on the inductance of the winding and the supply voltage. As the current increases, the voltage drop across a resistive current limiter will increase, dropping the voltage applied to the motor winding, and therefore, dropping the rate of increase of the current through the winding. As a result, the current will only approach the rated current of the motor winding asymptotically

In contrast, with a pure current limiter, the current through the motor winding will increase almost linearly until the current limiter cuts in, allowing the current to reach the limit value quite quickly. In fact, the current rise is not linear; rather, the current rises asymptotically towards a limit established by the resistance of the motor winding and the resistance of the sense resistor in the current limiter. This maximum is usually well above the rated current for the motor winding.

Open Loop Current Limiters

Both the resistive and the linear transistorized current limiters discussed above automatically limit the current through the motor winding, but at a considerable cost, in terms of wasted heat. There are two schemes that eliminate this expense, although at some risk because of the lack of feeback about the current through the motor.

Use of a Voltage Boost

If you plot the voltage across the motor winding as a function of time, assuming the use of a transistorized current limiter such as is illustrated in Figure 4.3, and assuming a 1 amp 5 ohm motor winding, the result will be something like that illustrated in Figure 4.5:

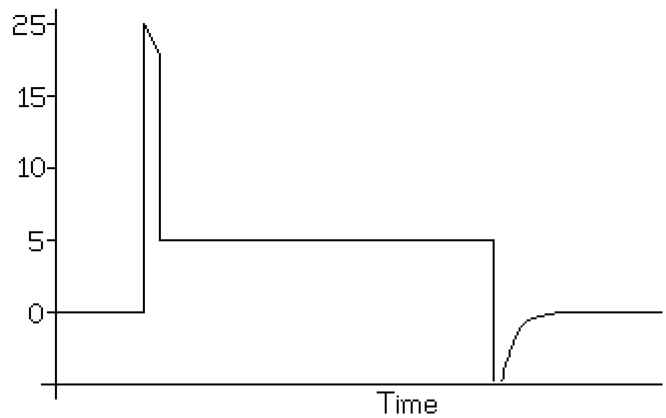

Figure 4.5

As long as the current is below the current limiter's set point, almost the full supply voltage is applied across the motor winding. Once the current reaches the set point, the voltage across the motor winding falls to that needed to sustain the current at the set point, and when the switches open, the voltage reverses briefly as current flows through the diode network and R_2.

An alternative way to get this voltage profile is to use a dual-voltage power supply, turning on the high voltage for as long as it takes to bring the current in the motor winding up to the rated current, and then turning off the high voltage and turning on the sustaining voltage. Some motor controllers do this directly, without monitoring the current through the motor windings. This provides excellent performance and minimizes power losses in the regulator, but it offers a dangerous temptation.

If the motor does not deliver enough torque, it is tempting to simply lengthen the high-voltage pulse at the time the motor winding is turned on. This will usually provide more torque, although saturation of the magnetic circuits frequently leads to less torque than might be expected, but the cost is high! The risk of burning out the motor is quite real, as is the risk of demagnitizing the motor rotor if it is turned against the imposed field while running hot. Therefore, if a dual-voltage supply is used, the temptation to raise the torque in this way should be avoided!

The problems with dual voltage supplies are particularly serious when the time intervals are under software control, because in this case, it is common for the software to be written by a programmer who is insufficiently aware of the physical and electrical characteristics of the control system.

Use of Pulse Width Modulation

Another alternative approach to controlling the current through the motor winding is to use a simple power supply controlled by *pulse width modulaton* (PWM) or by a *chopper*. During the time the current through the motor winding is increasing, the control system leaves the supply attached with a 100% duty cycle. Once the current is up to the full rated current, the control system changes the duty cycle to that required to maintain the current. Figure 4.6 illustrates this scheme:

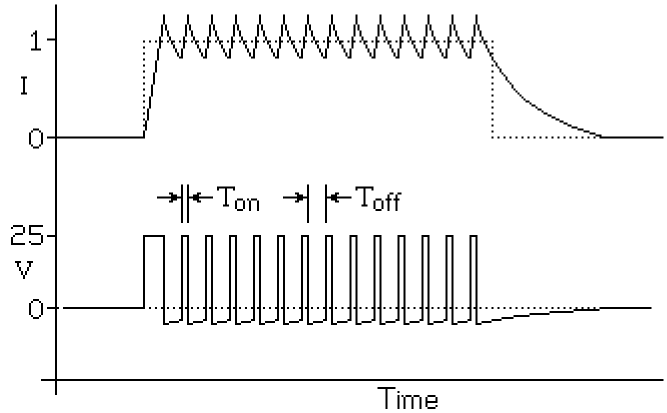

Figure 4.6

For any chopper or pulse width modulator, we can define the duty-cycle D as the fraction of each cycle that the switch is closed:

$D = T_{on} / (T_{on} + T_{off})$

where

T_{on} -- time the switch is closed during each cycle
T_{off} -- time the switch is open during each cycle

The voltage curve shown above indicates the full supply voltage being applied to the motor winding during the on-phase of every chopper cycle, while when the chopper is off, a negative voltage is shown. This is the result of the forward voltage drop in the diodes that are used to shunt the current when the switches turn off, plus the external resistance used to speed the decay of the current through the motor winding.

For large values of T_{on} or T_{off}, the exponential nature of the rise and fall of the current through the motor winding is significant, but for sufficiently small values, we can approximate these as linear. Assuming that the chopper is working to maintain a current of I and that the amplitude is small, we will approximate the rates of rise and fall in the current in terms of the voltage across the motor winding when the switch is closed and when it is open:

$V_{on} = V_{supply} - I(R_{winding} + R_{on})$
$V_{off} = V_{diode} + I(R_{winding} + R_{off})$

Here, we lump together all resistances in series with the winding and power supply in the on state as R_{on}, and we lump together all resistances in the current recirculation path when the switch(es) are open as R_{off}. The forward voltage drops of any diodes in the current recirculation path have been lumped as V_{diode}; if the off-state recirculation path runs from ground to the power supply (H-bridge fast decay mode), the supply voltage must also be included in V_{diode}. Forward voltage drops of any switches in the on-state and off-state paths should also be incorporated into these voltages.

To solve for the duty cycle, we first note that:

$dI/dt = V/L$

where

I -- current through the motor winding
V -- voltage across the winding
L -- inductance of the winding

We then substitute the specific voltages for each phase of operation:

$I_{ripple} / T_{off} = V_{off} / L$
$I_{ripple} / T_{on} = V_{on} / L$

where

I_{ripple} -- the peak to peak ripple in the current

Solving for T_{off} and T_{on} and then substituting these into the definition of the duty cycle of the chopper, we get:

$D = T_{on} / (T_{on} + T_{off}) = V_{off} / (V_{on} + V_{off})$

If the forward voltage drops in diodes and switches are negligable, and if the only significant resistance is that of the motor winding itself, this simplifies to:

$D = I\ R_{winding} / V_{supply} = V_{running} / V_{supply}$

This special case is particularly desirable because it delivers all of the power to the motor winding, with no losses in the regulation system, without regard for the difference between the supply voltage and the running voltage.

The AC ripple I_{ripple} superimposed on the running current by a chopper can be a source of minor problems; at high frequencies, it can be a source of RF emissions, and at audio frequencies, it can be a source of annoying noise. For example, with audio frequency chopping, most stepper controlled systems will "squeel", sometimes loudly, when the rotor is displaced from the equilibrium position. To find the ripple amplitude, first recall that:

$I_{ripple} / T_{off} = V_{off} / L$

Then solve for I_{ripple}:

$I_{ripple} = T_{off}\ V_{off} / L$

Thus, to reduce the ripple amplitude at any particular duty cycle, it is necessary to increase the chopper frequency. This cannot be done without limit because switching losses increase with frequency. Note that this change has no significant effect on AC losses; the decrease in such losses due to decreased amplitude in the ripple is generally offset by the effect of increasing frequency.

The primary problem with use of a simple chopping or pulse-width modulation control scheme is that it is completely open loop. Design of good chopper based control systems requires knowledge of motor characteristics such as inductance that are frequently poorly documented, and as with dual-voltage supplies, when motor performance is marginal, it is very tempting to increase the duty-cycle without attention to the long-term effects of this on the motor. In the designs that follow, this weakness will be addressed by introducing feedback loops into the low level drive system to directly monitor the current and determine the duty cycle.

5. Microstepping of Stepping Motors

Microstepping serves two purposes. First, it allows a stepping motor to stop and hold a position between the full or half-step positions, second, it largely eliminates the jerky character of low speed stepping motor operation and the noise at intermediate speeds, and third, it reduces problems with resonance.

Although some microstepping controllers offer hundreds of intermediate positions between steps, it is worth noting that microstepping does not generally offer great precision, both because of linearity problems and because of the effects of static friction.

Sine Cosine Microstepping

Recall, from the discussion in Part 2 of this tutorial, on Stepping Motor Physics, that for an ideal two-winding variable reluctance or permanent magnet motor the torque versus shaft angle curve is determined by the following formulas:

$h = (a^2 + b^2)^{0.5}$
$x = (S / (\pi/2))\ arctan(b / a)$

where:

a -- torque applied by winding with equilibrium at angle 0.
b -- torque applied by winding with equilibrium at angle S.
h -- holding torque of composite.
x -- equilibrium position.
S -- step angle.

This formula is quite general, but it offers little in the way of guidance for how to select appropriate values of the current through the two windings of the motor. A common solution is to arrange the torques applied by the two windings so that their sum *h* has a constant magnitude equal to the single-winding holding torque. This is referred to as sine-cosine microstepping:

$a = h_1 \, sin(((\pi/2)/S)\Theta)$
$b = h_1 \, cos(((\pi/2)/S)\Theta)$

where:

h_1 -- single-winding holding torque
$((\pi/2)/S)\Theta$ -- the electrical shaft angle

Given that none of the magnetic circuits are saturated, the torque and the current are linearly related. As a result, to hold the motor rotor to angle Θ), we set the currents through the two windings as:

$I_a = I_{max} \, sin(((\pi/2)/S)\Theta)$
$I_b = I_{max} \, cos(((\pi/2)/S)\Theta)$

where:

I_a -- current through winding with equilibrium at angle 0.
I_b -- current through winding with equilibrium at angle S.
I_{max} -- maximum allowed current through any motor winding.

Keep in mind that these formulas apply to two-winding permanent magnet or hybrid stepping motors. Three pole or five pole motors have more complex behavior, and the magnetic fields in variable reluctance motors don't add following the simple rules that apply to the other motor types.

Limits of Microstepping

The utility of microstepping is limited by at least three consideraitons. First, if there is any static friction in the system, the angular precision achievable with microstepping will be limited. This effect was discussed in more detail in the discussion in Part 2 of this tutorial, on Stepping Motor Physics, in the discussion of friction and the dead zone.

Detent Effects

The second problem involves the non-sinusoidal character of the torque versus shaft-angle curves on real motors. Sometimes, this is attributed to the detent torque on permanent magnet and hybrid motors, but in fact, both detent torque and the shape of the torque versus angle curves are products of poorly understood aspects of motor geometry, specifically, the shapes of the teeth on the rotor and stator. These teeth are almost always rectangular, and I am aware of no detailed study of the impact of different tooth profiles on the shapes of these curves.

Most commercially available microstepping controllers provide a fair approximation of the sine-cosine drive current that would drive an ideal stepping motor to uniformly spaced steps. Ideal motors are rare, and when such a controller is used with a real motor, a plot of the actual motor position as a function of the expected position will generally look something like the plot shown in Figure 5.1.

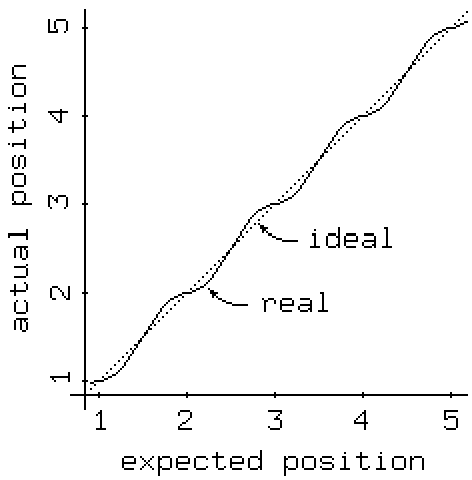

Figure 5.1

Note that the motor is at its expected position at every full step and at every half step, but that there is significant positioning error in the intermediate positions. The curve shown is the curve that would result from a perfect sin-cosine microstepping controller used with a motor that had a torque versus position curve that included a significant 4th harmonic component, usually attributed to the detent torque.

Quantization

The third problem arises because most applications of microstepping involve digital control systems, and thus, the current through each motor winding is quantized, controlled by a digital to analog converter. Furthermore, if typical PWM current limiting circuitry is used, the current through each motor winding is not held perfectly constant, but rather, oscillates around the current control circuit's set point. As a result, the best a typical microstepping controller can do is approximate the desired currents through each motor winding.

The effect of this quantization is easily seen if the available current through one motor winding is plotted on the X axis and the available current through the other motor winding is plotted on the Y axis. Figure 5.2 shows such a plot for a motor controller offering only 4 uniformly spaced current settings for each motor winding:

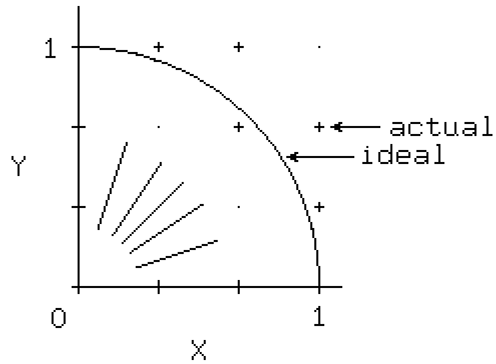

Figure 5.2

Of the 16 available combinations of currents through the motor windings, 6 combinations lead to roughly equally spaced microsteps. There is a clear tradeoff between minimizing the variation in torque and minimizing the error in motor position, and the best available motor positions are hardly uniformly spaced! Use of higher precision digital to analog conversion in the current control system reduces the severity of this problem, but it cannot eliminate it!

Plotting the actual rotor position of a motor using the microstep plan outlined in Figure 5.2 versus the expected position gives the curve shown in Figure 5.3:

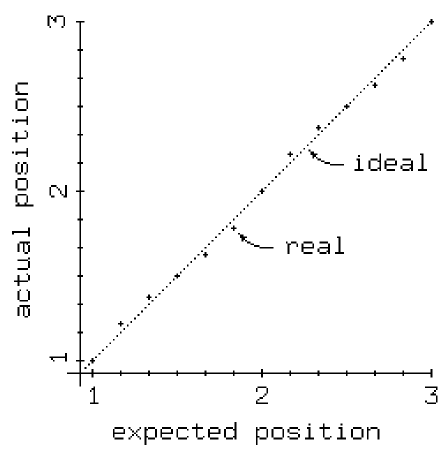

Figure 5.3

It is very common for the initial microsteps taken away from any full step position to be larger than the intended microstep size, and this tends to give the curve a staircase shape, with the downward steps aligned with the full step positions where only one motor winding carries current. The sign of the error at intermediate positions tends to fluctuate, but generally, the position errors are smallest between the full step positions, when both motor windings carry significant current.

Another way of looking at the available microsteps is to plot the equilibrium position on the horizontal axis, in fractions of a full-step, while plotting the torque at each available equilibrium position on the vertical axis. If we assume a 4-bit digital-to-analog converter, giving 16 current levels for each each motor winding, there are 256 equilibrium positions. Of these, 52 offer holding torques within 10% of the desired value, and only 33 are within 5%; these 33 points are shown in bold in Figure 5.4:

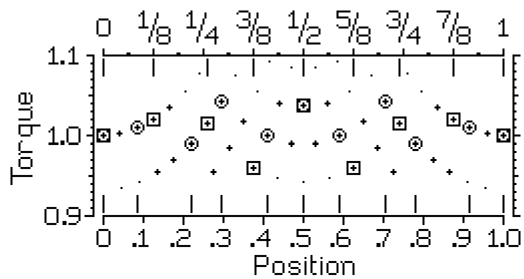

Figure 5.4

If torque variations are to be held within 10%, it is fairly easy to select 8 almost-uniformly spaced microsteps from among those shown in Figure 5.4; these are boxed in the figure. The maximum errors occur at the 1/4 step points; the maximum error is .008 full step or .06 microsteps. This error will be irrelevant if the dead-zone is wider than this.

If 10 microsteps are desired, the situation is worse. The best choices, still holding the maximum torque variation to 10%, gives a maximum position error of .026 full steps or .26 microsteps. Doubling the allowable variation in torque approximately halves the positioning error for the 10 microstep example, but does nothing to improve the 8 microstep example.

One option which some motor control system designers have explored involves the use of nonlinear digital to analog converters. This is an excellent solution for small numbers of microsteps, but building converters with essentially sinusoidal transfer functions is difficult if high precision is desired.

Typical Control Circuits

As typically used, a microstepping controller for one motor winding involves a current limited H-bridge or unipolar drive circuit, where the current is set by a reference voltage. The reference voltage is then determined by an analog-to-digital converter, as shown in Figure 5.5:.

Figure 5.5 assumes a current limited motor controller such as is shown in Figures 4.7, 4.8, 4.10 or 4.11. For all of these drivers, the state of the X and Y inputs determines the whether the motor winding is on or off and if on, the direction of the current through the winding. The V_0 through V_n inputs determine the reference voltage and this the current through the motor winding.

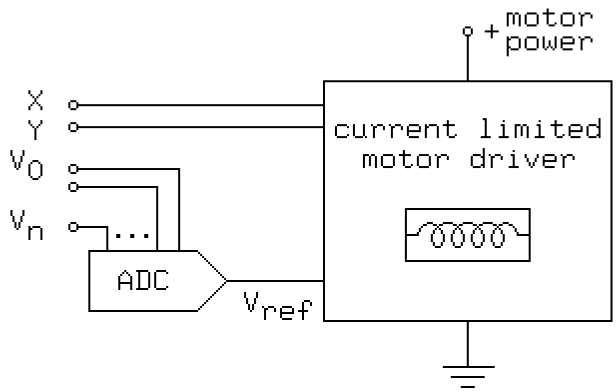

Figure 5.5