9x.19 PWM for DC Motors

Those of us who grew up with electric train sets remember well the difficulty in controlling their speed: the train is stationary; you gradually advance the control knob or level, but the train stubbornly sits there, then a moment later it rushes off. In those days the same recalcitrance afflicted sewing machines, with their pedal-operated speed control. It was a wonder anyone could control those things. But now you can get tools with speed controls that actually work – battery-powered drills, for example, where the speed is well-controlled as you advance the trigger, in spite of large variations of mechanical load as you drill your way through tough material.

9x.19.1 The Myth: PWM as Secret Sauce

These latter tools use permanent-magnet (PM) dc motors, with pulse-width modulation of the battery's dc voltage. They deliver plenty of torque at low speed, in contrast to the sewing machines and train sets of yesteryear. So it's easy to conclude (incorrectly, as we'll see) that it's the magic of PWM that delivers the desirable feature of well-controlled speed with adequate torque. To elaborate on this thinking, you might imagine that the pulsed current delivers high peak torque, with the motor gliding between pulses, and thus that PWM is more effective than the alternative of applying a steady dc voltage equal to the average PWM voltage (i.e., peak voltage times duty-cycle). This has caused considerable confusion, with statements like this: "Controlling the speed of something like a DC motor by varying the voltage (even in an efficient way) is a poor technique. It limits the motor to low torque at lower speeds because of limited power. That is why it is more common to see PWM (Pulse-Width Modulation) used to control motors (and lights/LEDs as well). PWM allows motors to have more torque at low speeds."

Such thinking would suggest that PWM is superior to applying a variable dc voltage, and that the experience with trains and sewing machines provides confirmation.

A. An Experiment

With characteristic arrogance, we decided to get to the bottom of this. First step, echoing eevblog's Dave Jones ("Don't turn it on, *Take it apart!*"), was a look at the innards of a typical battery-powered drill (Fig. 9x.55). Not much there – in particular there is no tachometer or other rotational-speed feedback, just PWM to a 2-terminal dc motor. Figure 9x.56 shows the voltage waveforms at increasing degrees of trigger depression.

Now for the fun: we intercepted the motor leads and applied dc from an external power supply. We did this to test the proposition that PWM is "better" than steady dc in terms of torque at low speed. We used a big drill bit to bore into hunks of plywood (Fig. 9x.57) at low speed, comparing the drill's muscle with both PWM and dc drive. And we found... (drumroll)... that there was no discernible difference!

B. Toy Trains and Sewing Machines

But this disagrees with our childhood "knowledge." What's going on? It turns out that those controllers of yesteryear did not provide a variable (and stiff) voltage source – instead they used a fixed voltage (either dc or ac) in series with a variable impedance (a *rheostat*, fancy name for a wirewound potentiometer that can dissipate at least a few watts). If you think about it, this is a *terrible* way to control a motor: Imagine you've set the rheostat so the motor's turning slowly, and then you increase the mechanical load on the motor; the latter then slows down, presenting a lower

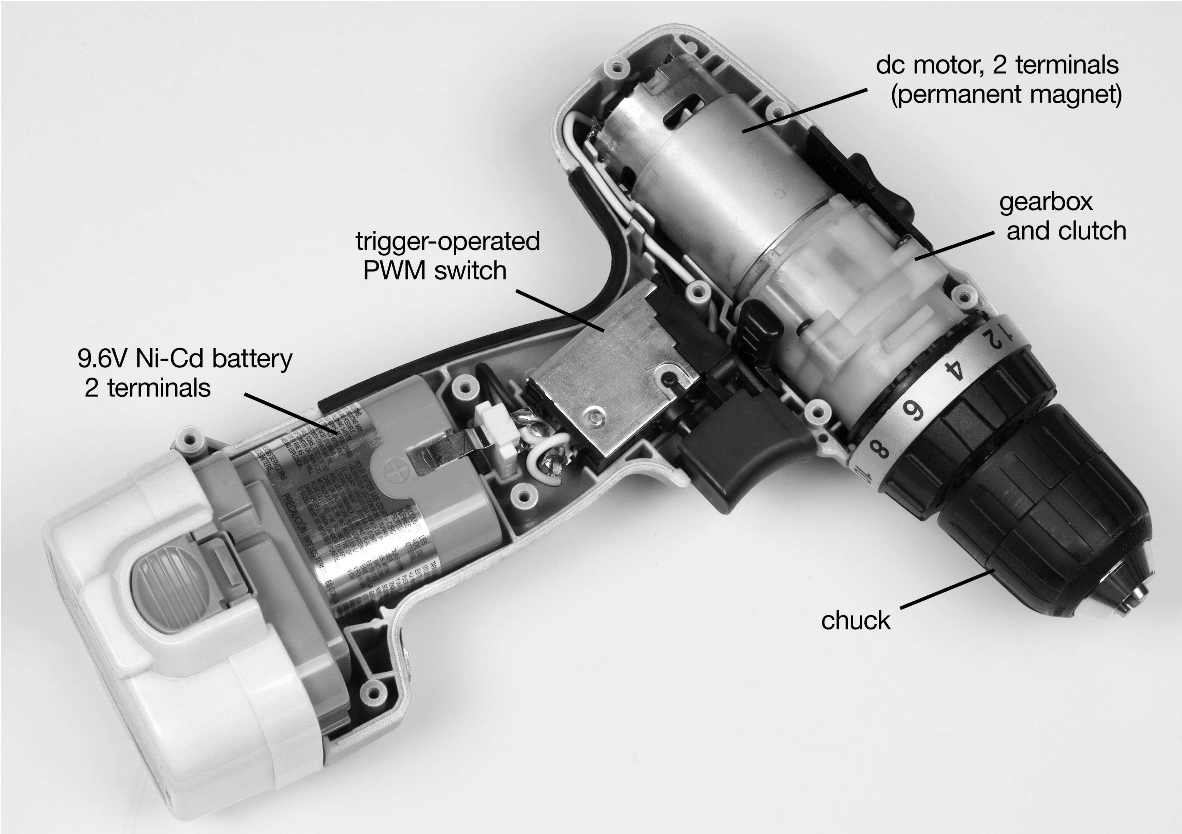

Figure 9x.55: *Half-clamshell view of a variable-speed electric drill (Makita 6260D). The shiny module behind the trigger generates full-voltage PWM to the permanent-magnet dc gearhead motor.*

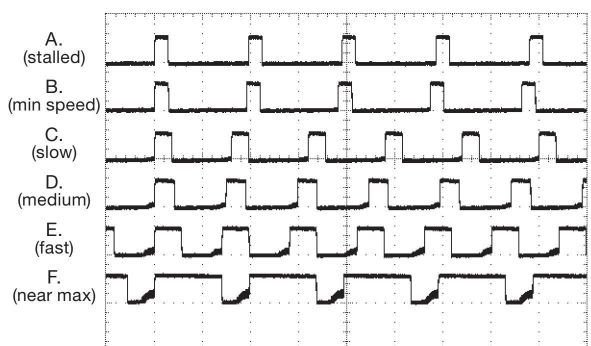

Figure 9x.56: *PWM voltage waveforms applied to the dc motor in Fig. 9x.55 as the trigger is advanced. When it is fully squeezed the waveform transitions to steady 10 V dc. Horizontal: 40 μs/div; Vertical: 20 V/div.*

impedance (because of lower "back-EMF"), which causes its voltage (and therefore rotational speed) to drop. Similarly, if the motor is turning at some speed under load and you decrease the load, the motor will turn faster, whereupon the increased back-EMF causes the voltage (and therefore motor speed) to rise still further.

Figure 9x.57: *Victimized plywood, following drill-torque tests.*

C. Another Experiment

To compare the merits of voltage-drive with what we might call rheostat-drive, we rigged up two experiments, each of which let us see the ramp-up of a dc motor when driven either with a fixed low voltage (i.e., less than the motor's rated voltage), or with a series resistor from a higher voltage, the resistor chosen to produce the same lower motor voltage once things had stabilized. The ramp-up from a cold start is a simple proxy for torque, and it's easier to measure.

1. dc fan with tachometer. For this experiment we used a 120 mm 12 V brushless cooling fan (Delta AFB1212SH); it has a logic-level tachometer output that generates two pulses per revolution. We ran it at half voltage (6 V), using either a 6 V dc supply, or a 12 V supply with a 24 Ω series resistor (the value that produced 6 V across the fan after settling to its final speed). Figure 9x.58 shows the results (both fan voltage and tachometer pulses) when ramping up from a cold start. It's clear that a series resistor causes considerably reduced torque (slower ramp-up of both voltage and speed), as just explained in §9x.19.1B.

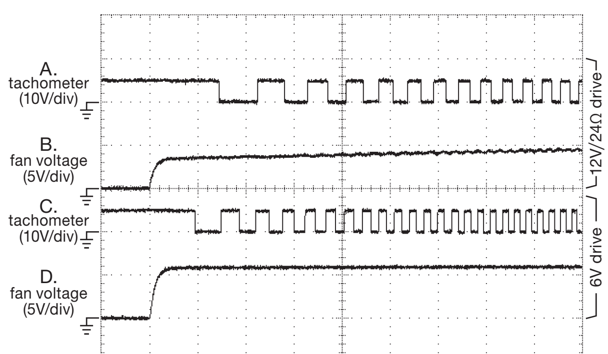

Figure 9x.58: *Tachometer output (A, C) and fan voltage (B, D) for a 12 V dc fan, as measured from a cold start for two configurations that run the fan at half voltage (6 V): The top pair used a 12 V supply with 24 Ω series resistor; the bottom pair used a 6 V supply and no resistor. Horizontal: 100 ms/div.*

2. dc motor-generator. For this experiment we used a pair of permanent-magnet dc motors[30] (marked Pittman 3140-0665) that we found in our vast collection of weird flea-market junk. They specify a strange operating voltage of 19.1 V (bringing to mind the famous outburst of the legendary I.I.Rabi, "who ordered *that*?"[31]) We used some heat-shrink tubing to couple the shafts, with the passive motor acting as a dc generator whose unloaded output voltage is proportional to rotational speed. Figure 9x.59 shows the results (motor drive voltage, and speed-proportional generator output voltage), demonstrating once again that the worst way to control motor speed is with a rheostat. Of perhaps greater interest, if you like myth debunking, is the demonstration that the thing behaves exactly the same with PWM (20 V amplitude, 50% duty cycle, 5 kHz) as with dc drive (10 V).

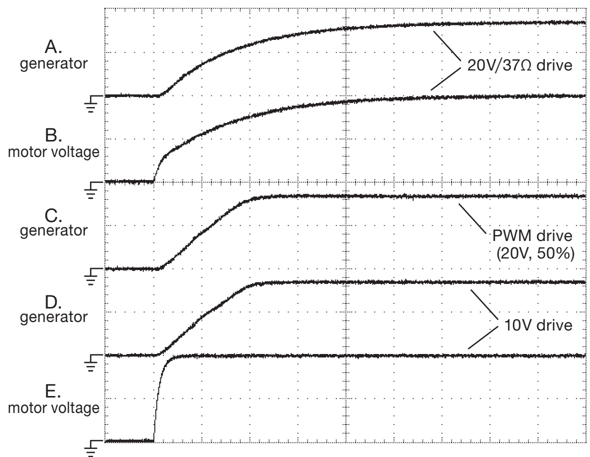

Figure 9x.59: *This time we tricked a pair of PM dc motors into acting as a motor-generator, with the passive motor's generated dc output serving as a measure of shaft rotation rate. As in Fig. 9x.58, the top pair (A,B) shows startup to half-voltage (10 V) with a series resistor, whereas the bottom pair (D,E) is powered with a 10 V voltage source. Trace C demonstrates the equivalence of PWM (versus dc) drive. Horizontal: 100 ms/div; Vertical: 5 V/div.*

9x.19.2 Wrapup: PWM versus DC for Motor Drive

Evidently the poor performance of those rheostat-controlled sewing machines and toy trains was a result of using a variable impedance in series with the motor's power supply, nicely seen in the data above. In those heady years (of the 1950's) a rheostat was far less expensive than a variable transformer ("Variac"), and before the era of silicon power semiconductors it was difficult to implement a high-current variable dc supply.

It's no longer difficult or expensive to make the latter; so, looping back to the confusion about PWM versus dc drive, and given the prevalence of PWM motor control, are we to conclude that PWM provides superior control of PM dc motors?

The answer is no: PWM is simply a convenient way to achieve the benefits of variable dc drive, without having to build a dc-dc converter, with its inductive energy storage, capacitive smoothing components, and feedback regulation. In effect, with PWM the motor's inductance and mechanical inertia[32] substitutes for the *LC* components in a dc-dc converter. You can think of PWM+motor

[30] These included a nice quadrature optical encoder on the back end; evidently these were used to drive the pen assembly in HP 7470A *x-y* plotters.
[31] Referring to the just-discovered muon interloper, not to some motor's odd choice of voltage.
[32] Typical switching frequencies are upward of a few kHz (to put them above audibility), well above the response of the motor.

Figure 9x.60: *A vintage sewing-machine pedal-operated rheostat-type control. This thing gets* hot*, so the power resistor is kept safely inside a well-ventilated cage.*

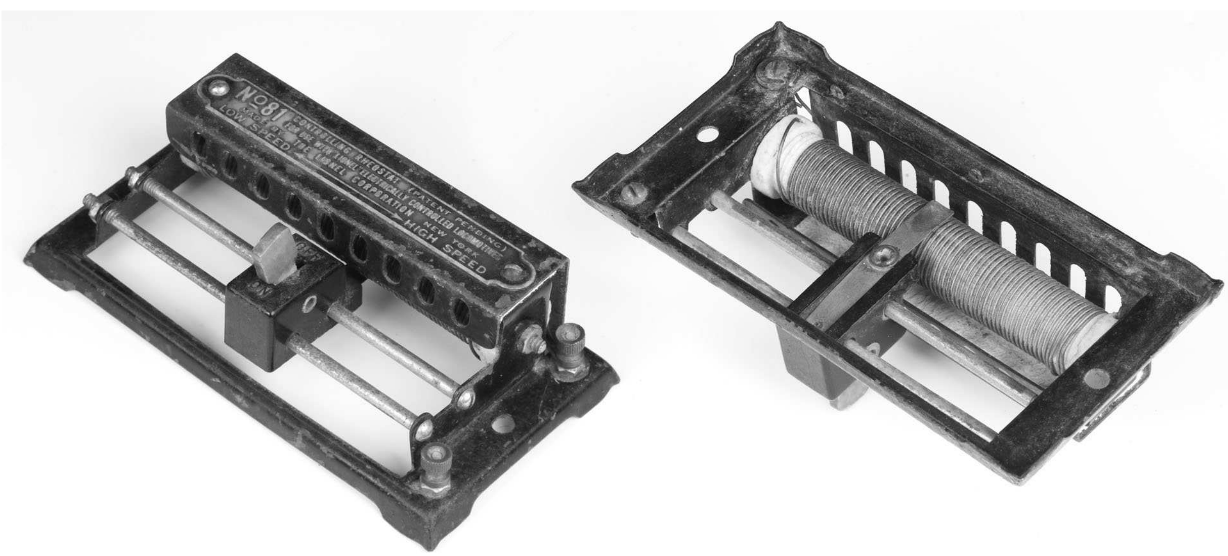

Figure 9x.61: *Vintage Lionel model 81 train-set rheostat. The handle moves a sliding contact along the wirewound power resistor inside the ventilated enclosure.*

as comprising a buck converter, equivalent to a dc voltage equal to the switched voltage times the switching duty cycle.[33] So all that's needed is a PWM switching signal to drive a MOSFET, with a catch diode (or MOSFET) to complete the motor drive. Easy peasy.

Looking slightly deeper, there are some disadvantages to PWM (compared with dc), namely the high switching frequency (and its harmonics) generates additional losses in the magnetic materials, and there are dynamic switching losses associated with switch and parasitic capacitances ("hard switching"). Overall, however, PWM is easy, and efficient enough to make it the technique of choice in dc motor control.[34]

One might add, as an afterthought, that if your goal is to tightly control a motor's rpm over

[33] And, as with a buck converter, the current out of the motor's return lead is dc, with little ripple.

[34] And in some cases the pulsating torque caused by PWM can be beneficial, particularly at low speeds where it can overcome "stiction" and "cogging."

varying mechanical load, you might consider actively closing a PID loop, adjusting PWM duty cycle with tachometer feedback.[35] Active control beats the pants off passive control; the latter depends on motor current changing according to the difference between applied voltage (via PWM, or whatever) and the speed-proportional back-EMF, a control mechanism that is compromised by non-zero resistive losses in the motor windings.[36]

9x.19.3 Afterward: DC Motor Model

We've been a bit casual about what's actually going on inside the dc motor, having arrogantly asserted that there's a "back-EMF" that's proportional to rotation speed, and that the torque depends on the current. It's worth looking into this a bit more. The treatment that follows[37] was suggested by Steven Leeb, to whom we are indebted.

Figure 9x.62 is the basic model for a permanent-magnet dc motor (with brushes or equivalent commutation). *Mechanically* it's a machine whose shaft is turning at an angular velocity ω radians/s, while providing a torque into some mechanical load of τ newton-meters (abbreviated N·m). *Electrically* it looks like a voltage source ("back-EMF") V_{bemf}, proportional to angular velocity, in series with the motor's winding resistance R_{m} and inductance L_{m}. The torque is strictly proportional to current, $\tau = k_{\text{i}} I$ (Lorenz force on a current-carrying wire), and the back-EMF is strictly proportional to rotational velocity, $V_{\text{bemf}} = k_{\text{v}}\omega$ (Faraday's law of induction, with the motor acting as a generator).

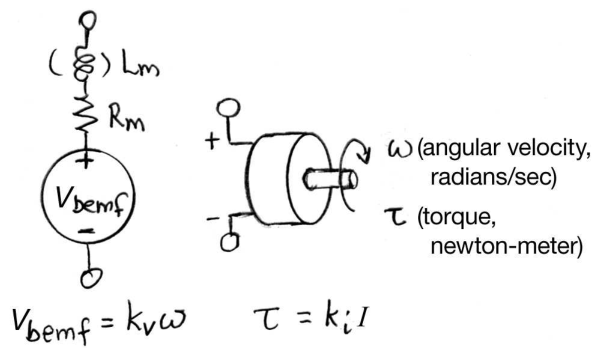

Figure 9x.62: *Simplified electrical and mechanical model of a permanent-magnet dc motor. R_{m} and L_{m} represent the resistance and inductance of the motor windings. The motor generates a torque proportional to the current flowing through the windings; and their motion through the magnetic field generates a "back-EMF" proportional to rotational speed. The proportionality constants are equal: $k_{\text{v}} = k_{\text{i}}$.*

Now here's something totally cool: the constants k_{v} and k_{i} are equal! This is easy to see, by equating the electrical power converted to motion, $P_{\text{elec}} = V_{\text{bemf}} I = k_{\text{v}}\omega I$, to the mechanical power delivered $P_{\text{mech}} = \tau\omega = k_{\text{i}} I\omega$.[38]

[35] Jim Roberge made a nice video in 1985 demonstrating just such a system – go to YouTube.com and enter this string: uHtKGf4AymM. We are indebted to Prof. Steven Leeb at MIT for this link, as well as for teachings on the secrets of dc motors.

[36] Going into this a bit further, the voltage difference ($\Delta V = V_{\text{PWM}} - V_{\text{back-emf}}$) appears across the winding's total impedance ($R + jX$), but what creates torque is the *current* through the winding. So, for example, if inductance dominates the winding's impedance, then the torque will increase only according to the time integral of the speed error. The differential term in a PID loop could do wonders here.

[37] In which we ignore the transient effects due to the (often ignorable) motor inductance. Strictly speaking, the motor's (electrical, energy-storing) inductance and its (mechanical, energy-storing) moment of inertia form a second-order system, with some natural frequency and damping factor. But the electrical time scale is ordinarily far faster than the mechanical timescale, justifying the simplified (electrically quasi-static) first-order treatment below. Of course, for the steady-state behavior one can always ignore inductance (and inertia as well).

[38] Note that it's the motor's internal voltage V_{bemf} that figures into this proof. Because of the motor's resistance

Now to the business of dc (or PWM) drive versus rheostat drive. Looking first at the limiting cases, for a perfect *current source* drive (i.e., of unlimited voltage compliance) the motor simply produces constant torque. With no mechanical load it will accelerate without limit, and the terminal voltage will track the angular speed. With a mechanical load whose torque is proportional to speed (say $\tau{=}\beta\omega$, a viscous drag) the motor will accelerate until the load's torque matches the motor's current-determined torque: $\omega_{\text{final}}{=}kI/\beta$. And if the "coefficient of viscosity" β changes, so will the equilibrium speed. It's rubbery – it behaves like the hard-to-control sewing machine.

At the other limit (a *voltage source*), the motor's steady-state velocity is simply $\omega_{\text{final}}{=}V/k$, the speed at which the back-EMF equals the applied voltage. For a lossless motor ($R_{\text{m}}{=}0$) that speed is independent of load torque – any attempt to lower the speed by increasing the load causes a mismatch of drive voltage and back-EMF, and is thus met with an enormous current (the voltage mismatch across $R_{\text{m}}{=}0$) that maintains the speed.[39]

For the realistic case of finite (but small) R_{m}, the acceleration to final speed is finite; and loading the motor with mechanical torque causes some degree of slowing, because the current required to produce the torque creates some voltage drop across R_{m}, thus leaving somewhat less voltage to balance the back-EMF. But the motor fights the slowdown: the reduced back-EMF puts more voltage across R_{m}, thus more current (and more torque). The lower R_{m}, the more effective the motor is in running at a load-independent constant speed.

For the in-between case (driving the motor with a rheostat in series with a dc voltage), the situation is also in-between. Rather than hand-waving, let's put this more quantitatively. Figure 9x.63 shows the arrangement, where we've ignored the (small) winding impedance; i.e., an ideal lossless motor. For a given supply voltage V_{s} and rheostat setting R_{s}, the motor current (for rotation speed ω) is just

$$I = \frac{V_{\text{s}} - V_{\text{bemf}}}{R_{\text{s}}} = \frac{V_{\text{s}} - k\omega}{R_{\text{s}}}, \tag{9x.1}$$

producing a motor torque

$$\tau = kI = \frac{k}{R_{\text{s}}}(V_{\text{s}} - k\omega) = \frac{kV_{\text{s}}}{R_{\text{s}}} - \frac{k^2\omega}{R_{\text{s}}}. \tag{9x.2}$$

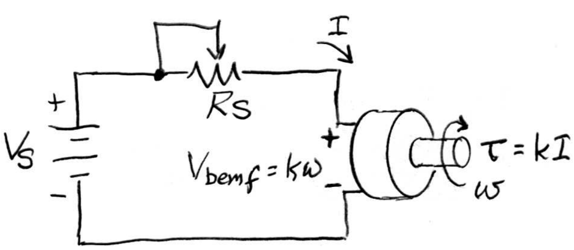

Figure 9x.63: *A fixed dc supply powers an ideal dc motor, with (undesirable) speed control through a rheostat (variable resistor).*

Equation 9x.2 lets us plot a family of curves of torque versus angular speed (Fig. 9x.64A). These show that the unloaded speed is just $\omega{=}V_{\text{s}}/k$, the speed at which the back-EMF plus resistive drop equals the supply voltage. Since there's no load, it's insensitive to series resistance (though in reality the motor itself has some bearing and air friction). These curves also show that the starting torque ($\omega{=}0$) increases inversely with series resistance.

R_{m}, the voltage you supply to run the motor (call it V_{s}) is greater than V_{bemf}, unless you have an "ideal motor"

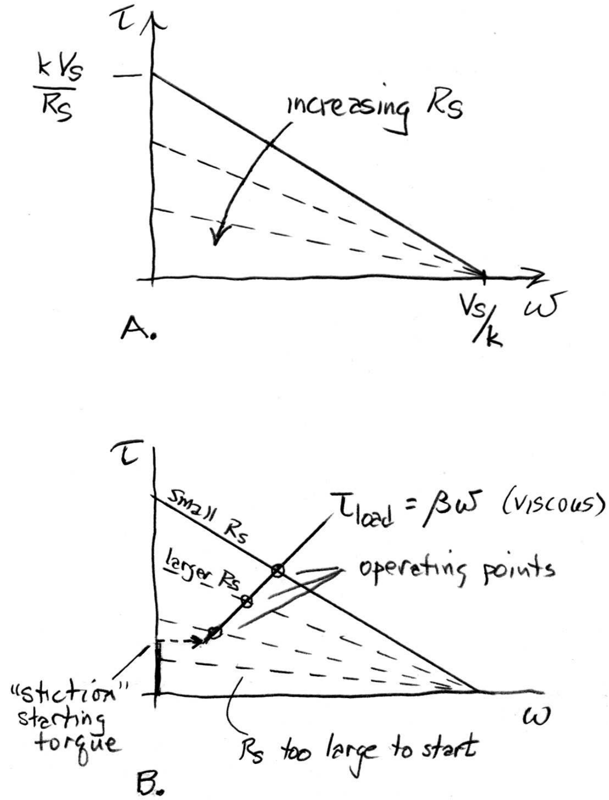

Figure 9x.64: *A. Operating characteristic (load line) for the circuit of Fig. 9x.63, according to eq'n 9x.2. B. Determining operating points for a linear load cursed with some stiction.*

From these operating curves (analogous to *load lines*, see Appendix F in the main volume) it's easy to figure out how a loaded motor will behave, by overlaying a curve representing the mechanical load's torque versus angular speed (Fig. 9x.64B). Here we've drawn a viscous load ($\tau \propto \omega$) that is sticky at rest and requires a minimum torque to overcome its "stiction" (static friction). For the motor to start with that load you must have $R_s < kV_s/\tau_{\text{stiction}}$: a graphical representation of the annoying "deadband" affliction of those rheostat controllers of yesteryear.

the bottom line: for good control of motor speed in situations of varying mechanical loading, power your dc motors from a low-impedance voltage source, adjusting the voltage to set the rotation speed. Equivalently, you can use pulse-width modulation to control motor speed, starting from a voltage source that runs the motor at maximum rated speed for 100% duty cycle. And for the ultimate control, use tachometer feedback.

A. Series Resistance: Op-amp Analogy

Here's another way[40] to think about this business of series resistance. Look at Figure 9x.65. We apply a supply voltage V_s to the (non-ideal) motor, creating a current I that produces a torque $\tau = kI$.

(R_m=0). But $k_v = k_i$, regardless.

[39] And, in that unrealistic case of R_m=0, the motor draws infinite current initially, producing infinite acceleration until reaching final speed.

[40] Steven Leeb, again!

But the current is proportional to $V_\mathrm{s}-V_\mathrm{bemf}$ and inversely proportional to R_m: $I=(V_\mathrm{s}-k\omega)/R_\mathrm{m}$. That is, you can think of the motor as a differential amplifier that converts the difference voltage $\Delta V=V_\mathrm{s}-V_\mathrm{bemf}$ to torque, with gain $(\tau/\Delta V)$ equal to k/R_m.

Viewed this way, a good motor has high gain (k/R_m), just as a good op-amp has high open-loop gain G_OL. And it doesn't matter if the resistance comes from the motor or an external series resistance R_s, which simply adds to the motor's internal resistance for a final (even lower) gain of k/R_total.

Finally, closing the loop around the "amplifier" gives us the circuit of Figure 9x.66. The inner (mechanical) loop includes a proportional ("viscous") torque drag, plus rotor inertia (which integrates net torque to spin up the angular speed); the outer loop is the by-now familiar voltage-to-torque characteristic of the motor, including the effect of its internal resistance and any external series resistance that is foolishly added. Don't add the latter!

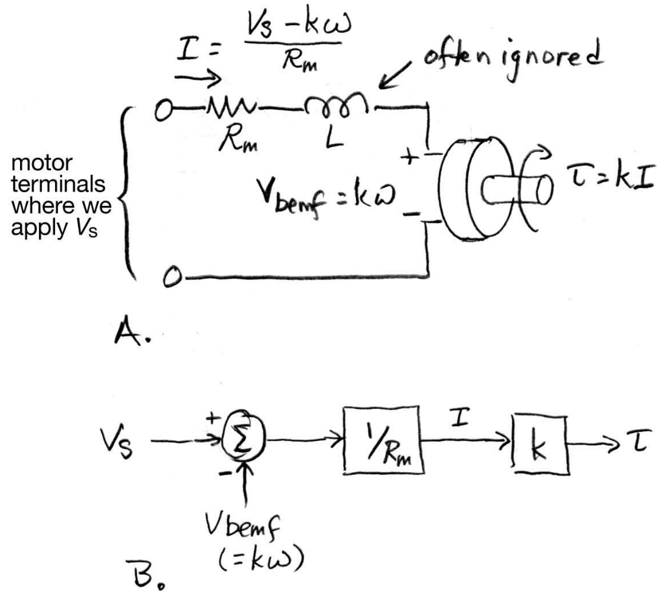

Figure 9x.65: *Motor model, viewed as a gain block that converts the voltage difference $V_\mathrm{s}-k\omega$ to a torque τ, with a "gain" of k/R_m.*

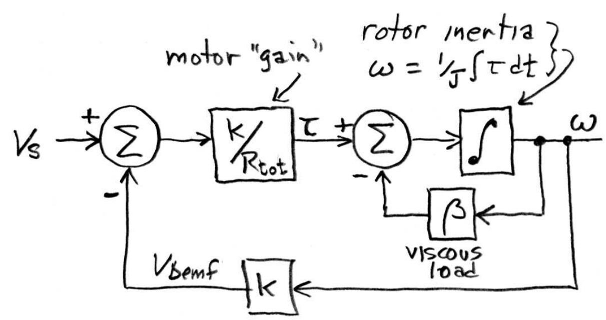

Figure 9x.66: *Closing the loop around the motor-as-amplifier. The inner loop is the mechanical system (viscous load torque proportional to speed, plus rotor inertia) that converts torque to angular speed; the outer loop takes it from there.*