

MOSFETs and catch diodes

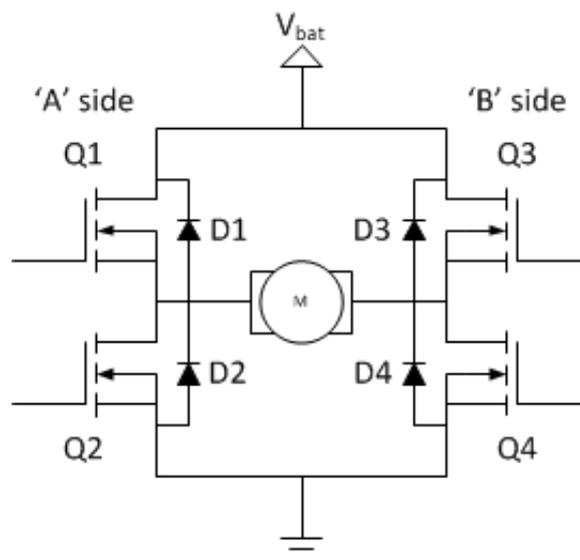
Author : Andras Tantos

Introduction

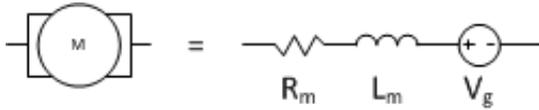
In the previous set of articles I've gone to excruciating details of each of the drive modes: the [sign-magnitude drive](#), the [lock anti-phase drive](#) and the [asynchronous sign-magnitude drive](#). To recap, the most important equations and properties of each of the drive modes are in the table below:

	Sign-magnitude drive	Lock anti-phase drive	Async. Sign-magnitude drive
Average motor current (I_{mot_avg})	$(V_{mot_avg} - V_g) / R_m$	$(V_{mot_avg} - V_g) / R_m$	$(V_{avg_conduct} - V_g) / R_m$
Average motor voltage (V_{mot_avg})	$V_{bat} * t_{on} / t_{cycle}$	$V_{bat} * (t_{on} - t_{off}) / t_{cycle}$	$(V_{bat} * t_{on} + V_g * t_{off_zero}) / t_{cycle}$
Maximum ripple current (I_{ripple_max})	$V_{bat} / L_m * t_{cycle} / 4$	$1/2 * V_{bat} / L_m * t_{cycle}$	$V_{bat} / L_m * t_{cycle} / 4$
Input capacitor (C_{input})	$1/64 * V_{bat}^2 / V_{max_bat_ripple}^2 * L_m * t_{cycle}$	$1/2 * I_{mot_avg} / V_{bat_ripple} * t_{cycle}$	$L_m / R_m * (I_{max} + (V_{bat} / R_m) * \ln(1 + I_{max} / (V_{bat} / R_m + I_{max}))) / V_{ripple_max}$
Regenerative braking	Yes	Yes	No
Dynamic braking	Yes	Yes	Yes

As usual, I include the bridge circuit schematic:



and our motor model as well as a reminder:



With all that math and theory behind us, we can come back to practical questions of how to go about actually designing an H-bridge. In this article I will go through the high level design decisions to make and the major component selection questions. I will leave some of the intricacies and the drive-circuit design for a later installment.

High level design parameters

The design of an H-bridge usually starts with making some high-level decisions. These are:

1. The maximum operating voltage (V_{bat}) of the bridge
2. The maximum (average) motor current that the bridge needs to handle
3. The drive mode of the bridge
4. The switching frequency of the bridge

The first two questions are relatively easy to answer when you have a particular application or motor to drive. The drive mode is a complex problem, but the above table should give you a rough idea of the major trade-offs between the options. We will come back to the nuances of the drive mode options when we talk about drive circuits. The switching frequency selection we haven't talked so much about yet so let's spend some time on it. There are several things to consider:

- If the switching cycle time (t_{cycle} , the inverse of the switching frequency) is not quite a bit smaller than the electrical time-constant of the motor (L_m/R_m), the on-time and off-time current changes are no longer linear. This invalidates most of the calculations above in the table and in the previous articles so if you choose such a low frequency you're pretty much on your own. At the same time I see no reason to run at such a low frequency and I hope you won't either after reading through the rest of the items.
- The other factor to consider is that the ripple current (I_{ripple_max}) is directly proportional to t_{cycle} . So as your operating frequency increases, the ripple current decreases. Usually you want as low of a ripple current as possible because it stresses components, reduces efficiency (additional loss on various resistances on wires, connectors, switching elements etc.) and generates EMI noise.
- In some drive-modes the size of the required input capacitor depends on the operating frequency. The higher it is, the lower the input capacitance needs to be.
- The switching of the voltage generates an audible 'buzz' on the motor. To avoid this buzz you have to go up to ultrasonic ranges, basically above 20kHz. Of course if audible noise is not a concern in your application than this consideration doesn't apply to you.
- As we will discuss it later, the higher the switching frequency is, the higher the switching losses on the bridge are, and at some point they start to be a significant source of heat. This will limit your ability to increase the switching frequency arbitrarily high.
- As the switching frequency increases, you will want to turn the switching elements on-

and off- faster to minimize the above mentioned switching loss. That makes the drive-circuit design harder as well as make the circuit more 'noisy', emitting more EMI radiation.

Overall, today it seems that 20-40kHz switching frequency is a good compromise between these requirements for most designs. This is somewhat of a moving target however as component technology improves and of course depends on the application as well.

Component selection

Once you more or less have decided on the above design parameters, you can start looking for components. In a bridge, there are eight components of interest: the four switching elements and the four catch diodes. You of course will have to design the drive circuit as well that involves further component selection, but I'll discuss those issues later.

Switching elements – MOSFETs

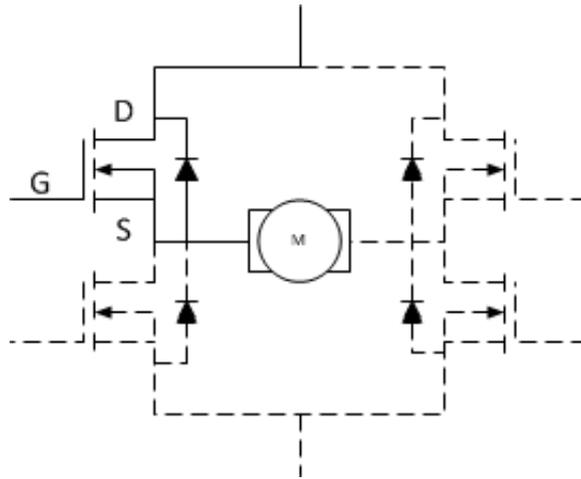
One of the key decision to make for an H-bridge is the selection of the switching elements. There are many factors to be considered, the most important ones are the operating current, the operating voltage and the switching frequency. For really high-voltage applications (several hundred volts) IGBTs are becoming popular and there are still some bridges out there using BJTs, however the vast majority of the modern designs use MOSFETs, so for the rest of the document I will assume MOSFET switching elements.

MOSFETs, when operated as switches, have two states: on and off. In the 'on' state they more or less behave like a small resistor, and their resistance is denoted by r_{dson} . Obviously the higher this value is, the higher the losses are on the MOSFET. While efficiency is not a big concern for most H-bridge designs, heat is. Since the loss on the MOSFET is converted to heat that has to be dissipated, the lower r_{dson} is the better. Another factor to consider is that r_{dson} is temperature-dependent and increases with temperature. Datasheets usually brag about r_{dson} at 25°C, but that hardly can be considered as normal operating condition. So always look for r_{dson} over the full temperature range to make sure you're operating within safe limits.

P- VERSUS N-CHANNEL

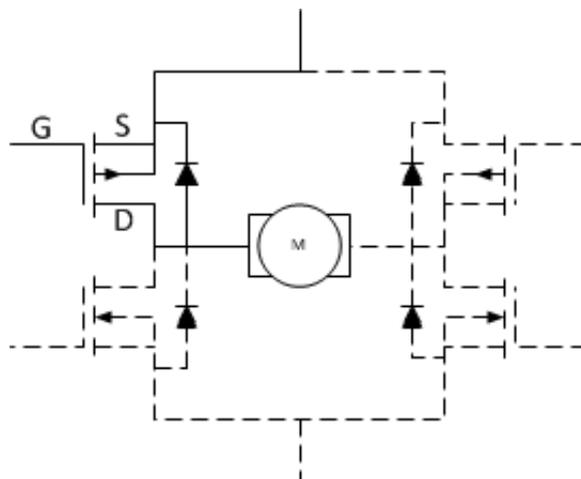
Right after deciding on MOSFET technology to use, the next question to answer is to use 'N'-channel or 'P'-channel MOSFETs. 'N' channel MOSFETs have a much lower r_{dson} so it would appear, that N-channel devices are desirable for their lower losses. For the low-side switches (Q2 and Q4) they are the obvious choice in deed. For the high-side switches (Q1 and Q3) the picture however is more complicated.

For N-channel devices to work, their drain needs to be at a higher potential than their source (otherwise their body-diode would open). So, when operated on the high-side, their source is connected to the motor terminal, and their drain is to the power supply:



This means that their source terminal potential can be anywhere between ground and V_{bat} . In order to turn the FET properly on, their gate should be (depending on construction) 3-12V higher than their source. At the same time, MOSFETs are very sensitive to the maximum gate-source voltage allowed on them. This is specified in the datasheet, but normally you will destroy your FET if you ever put more than 20V or so between its gate and source. This means that in order to properly and safely turn on a high-side N-MOSFET you need a varying gate-voltage that is potentially higher than V_{bat} . We'll come back to this later, for now it suffice to say, that this requirement complicates driver design quite a bit.

P-channel devices don't have this problem: their source is connected to V_{bat} , while their drain is on the motor terminal.



They need a negative 3-12V on their gate compared to their source to turn them on, so the the gate voltage is clearly always below V_{bat} and can be kept within safe limits without monitoring the motor terminal voltage. Not only P-channel devices usually have a higher r_{dson} , they are also slower to turn on and off. This will aggravate dynamic loss problems. and potentially complicate shoot-through protection.

Overall the usual trade-off is that for undemanding applications P-MOSFETs are selected for high-side operation as r_{dson} is not going to be a big problem, losses are manageable and they require much simpler drive circuitry. For high-current applications N-channel devices are a better compromise as P-FETs with comparable r_{dson} are either not available or extremely expensive making it reasonable to spend some extra money on the drive circuit.

Package selection and thermal management

Once you settle on the channel-type you can start looking at device datasheets. The goal is to arrive at is the maximum allowed r_{dson} for the device. It will depend on the maximum average motor current and the available cooling. Because MOSFETs – when they're on – can be basically thought of as resistors, the dissipated heat on them is going to be:

$$P = r_{dson} * I_{avg}^2$$

Now, the average current used here is the average current through the FET, not necessarily through the motor, but in most applications it's a good conservative approach to use the maximum allowed (average) motor current for this exercise. From this, the on-resistance for the FETs need to be:

$$r_{dson\ max} / I_{mot_avg_max}^2$$

If you want to figure out how much power a FET can dissipate, you have to start by looking at its packaging. As a guideline, the bigger the package is, the more heat it can dissipate. Common packages include SO-8, D-PAQ and D2-PAQ for surface mount and TO-92, TO-220 and TO-3 for through-hole mounting. There are of course other exotic packages and new ones are introduced almost every day.

The main characteristics that you're looking for is the 'thermal resistance' of the package, usually denoted as R_{θ} . The way to use it is this: the temperature difference between the two 'things' the thermal resistance is measured between will be $\Delta T = R_{\theta} * P$, where P is the dissipated power transferred between the two 'things'.

The datasheets also specify how hot the chip (or die) can become before it gets damaged.

The most valuable parameter consequently is the junction-to-ambient thermal resistance, that tells you how much hotter the chip is than the air that surrounds the package. There are some complications in determining that number as it depends on a lot of design-parameters so in many cases the datasheet will only contain another parameter, the junction-to-case thermal resistance. That tells you how much hotter the chip inside the package is than the outside of the case at a certain power dissipation but it leaves it up to you to figure out how hot the case can become.

When working with surface-mount packages, it is important to note that thermal characteristics are highly dependent on the actual PCB layout. Take the traditional SO-8 package as an example (one typical FET datasheet is here for the [FDS8447](#)) You'll see that with just the

minimum amount of copper on the board the package has $\sim 125^{\circ}\text{C}/\text{W}$ thermal resistance. With a 6.5cm^2 (1in^2) copper area, the same number is less than half of this, $\sim 50^{\circ}\text{C}/\text{W}$. Let's say now that the device itself can operate up to 150°C (this is called the maximum allowed junction temperature). If you can keep the surrounding temperature (called ambient temperature) below 50°C , you can dissipate somewhere between 0.8W ($(150-50)/125$) and 2W ($(150-50)/50$) of power on the FET depending on your PCB layout.

Through-hole packages of course also perform very differently with or without heat-sinks. A TO-220 package for example has a $\sim 60^{\circ}\text{C}/\text{W}$ thermal resistance without a heat-sink (here's a typical TO-220 packaged FET datasheet, the [FDP55N06](#)). With the same temperature limits as before, this device can dissipate 1.67W of power. With a heat-sink, like this: <http://www.aavidthermalloy.com/products/standard/7023b-mtg> you will easily be able to dissipate 9 watts.

Once you figure out how much power you can dissipate on each FET, you can calculate the minimum allowed r_{dson} . As an example, taking the 1.67W power dissipation limit and assuming you want to be able to work with an average 10A of current over the transistor, you get a maximum r_{dson} of $16.7\text{m}\Omega$. That happens to be a quite nice fit for the [IRF1010Z](#) MOSFET, which has an r_{dson} of $7.5\text{m}\Omega$ at 25°C , and about twice of that at 150°C . Of course if you can provide better cooling with a heat-sink, or a fan (or both) than you can handle much higher average currents, but it is a losing game: if you can dissipate let's say four times as much power (6.4W), you can only handle twice the current (20A) with the same r_{dson} .

Turn-on and -off times

If you look at the previous power dissipation equation, you'll see that lowering r_{dson} is a more promising approach to increasing the current-delivery capability of the bridge without blowing your heat-budget: for example if you want to deliver twice the current you only need to cut r_{dson} in half. There's a catch though: the lower r_{dson} gets, the bigger the MOSFET becomes. The bigger the physical device is, the bigger its gate will be. The gate forms a capacitor towards the source and the drain. Since MOSFETs are voltage driven devices, their gate-source voltage has to be in a certain range (usually above $5\text{--}10\text{V}$) to be fully turned on, and in another range (less than a volt or so) to be turned off. So, the on and off transient has to charge and discharge these parasitic capacitors. If you have only limited current available to drive the gate (and you always do), the higher the gate-capacitance is, the longer it takes to charge or discharge it. Why is it important?

MOSFETs have a low r_{dson} when they are fully on, and they conduct almost no current when they're completely off. In both cases the dissipated power is relatively low. However when they transition between these two states, there will be a short period where r_{dson} is relatively high, but not high enough to stop significant current from flowing through the device. In these transitional periods both the voltage drop on the device (due to r_{dson}) and the current through it are significant, resulting in high power-dissipation. Naturally you would like to keep this transition time as low as possible from this perspective (we'll talk later about reasons why you don't want it to be too fast either), so a higher gate-capacitance will not be desirable. With a given gate-drive strength the gate capacitance limits the speed by which the element can be

turned on and off, and thus poses an operating frequency limit.

With that said, switching loss is usually not that big of an issue in a modern bridge for operating frequencies below let's say 40kHz, but becomes significant as frequency increases. After a certain point it is the main contributor to the dissipated heat.

Catch diodes

Catch diodes (D1..D4) are often overlooked or just briefly mentioned in most H-bridge descriptions. If you read the introductory article, you'll understand why: in the two most common drive modes they almost never conduct current for any significant amount of time, and their only purpose is to provide a path for the current to flow in the short transitions between the on-time and the off-time.

In the asynchronous drive mode however, the off-time current flows through the catch diode(s), so – especially for that mode – we'll have to pay more attention to these components.

Whenever a diode is conducting current, there will be a relatively constant voltage drop on it. This is called forward voltage drop and denoted as V_F . It is in the 500..1000mV range for most components. This voltage drop, combined with the current through the diode will produce some heat dissipation. The actual heat dissipation depends on the average current flowing through the diode and the percentage of the time the diode is open. The average motor current on the diode is the following (we're only talking about the asynchronous sign-magnitude drive mode here):

$$I_{\text{mot_avg}} = (V_{\text{avg_conduct}} - V_g) / R_m$$

and it flows through one of the diodes for $t_{\text{off_conduct}}$ amount of time, so the dissipated power is:

$$P_{\text{diode}} = (V_{\text{avg_conduct}} - V_g) / R_m * V_F * (t_{\text{off_conduct}}/t_{\text{cycle}})$$

Now, when the average current is high, the bridge is in continuous current mode, so we can simplify these equations a little bit, and get:

$$P_{\text{diode}} = (V_{\text{bat}} * t_{\text{on}}/t_{\text{cycle}} - V_g)/R_m * V_F * (t_{\text{cycle}} - t_{\text{on}})/t_{\text{cycle}}$$

This is a quadratic equation over t_{on} , and reaches its maximum when $t_{\text{on}} = t_{\text{cycle}}/2$. Finally, since we don't know V_g , the generator voltage, the conservative approach to take is to assume V_g to be 0. (In fact the absolute most conservative is to assume V_g is $-V_{\text{bat}}$, but that's a really extreme condition. I will stay with $V_g = 0$ for now, but you can do the math for the other case if you're extra cautious) With that we get:

$$P_{\text{diode}} = 1/4 * V_{\text{bat}} * V_F / R_m$$

As an example, if we assume R_m is 1?, V_{bat} is 20V and V_F is only 500mV, we get 2.5W heat dissipated on the diode. You can see that the heat dissipation on the diode (again only for

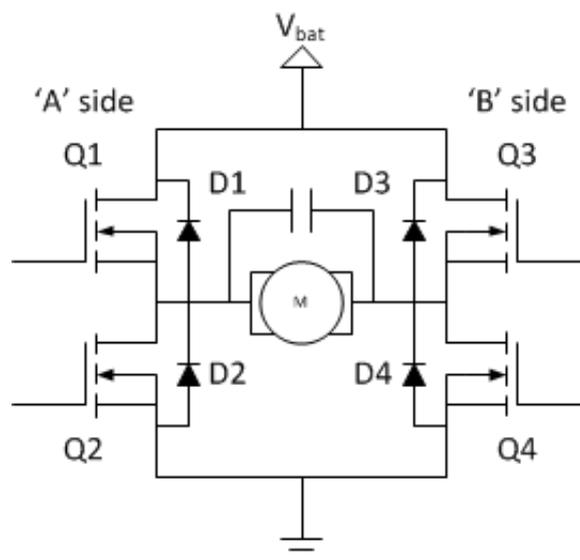
asynchronous sign-magnitude drive) can be significantly higher than that on the MOSFETs.

One important feature of MOSFET transistors is that they contain an intrinsic (unavoidable, built-in) diode between their drain and source. This diode acts as a catch diode in an H-bridge configuration, and most MOSFET datasheets specify the parameters of this diode. In many bridge designs it is possible to use this built-in diode of the transistors and not provide external ones, but the specification of this diode needs to meet the design requirements obviously. For bipolar transistors there's no such intrinsic diode so external diodes always have to be provided.

If you decide to use external diodes with MOSFETs, make sure that those diodes have a V_F lower than that of the body-diodes of the transistors. Otherwise, the diodes inside the FETs will open first and will divert the current from your external diodes.

A nice advantage of using the internal diodes is that – being on the same die, inside the same package – the cooling and heat-sinks that you provide for the FETs will automatically work for the diodes as well. Of course you have to make sure that the cooling is in fact adequate for the diodes, not just the FETs but in many cases this approach results in a simpler mechanical construction.

Diodes – mostly when they're off – have a small capacitance between their leads. This capacitance will have to be discharged before the device can turn on and leads to a delay in response to a sudden change in the voltage. This capacitance depends on many factors, but in general grows with the surface area of the P-N transition, that is the current carrying capability of the device. In short, the beefier the device, the slower it is. When the bridge turns off, the motor current will need a way to continue flowing. The motor will forward-bias a diode (or diodes) in the bridge to create a route for that current, however the turn-on delay of the diodes will create a problem. Without mitigation the motor voltage can rise to dangerous levels and damage the FETs. To bridge this interval where neither the switches nor the diodes conduct a capacitor has to be connected to the terminals of the motor:



Some motors contain this capacitor already – they are needed for other reasons as well – but many require an external one. This capacitor will conduct the current until the diodes open, but the terminal voltage of the motor will still rapidly increase. It is important to select diodes with a short turn-on delay, and this is the reason that Schottky-type diodes are preferred in this role.

Where to go from here

In this article we've went through the high-level design decisions that need to be made for an H-Bridge design and the various concerns we have to deal with when selecting the major components: the switching elements and the catch diodes. With that foundation, in the next installment of the series we can go through the various drive circuit options, that is, how to generate the gate voltages for the MOSFETs in the various drive modes.