

3ph Duo: 2 x 1kW Brushless Motor Controller w/ Field-Oriented Control

Design Notes [DRAFT]

Shane Colton
<scolton@mit.edu>

Controller v2.1
Design Notes Rev0
Revision Date: 01/12/2010

Notes:
TK: To come in a later revision.

Table of Contents

1	Design of the Controller.....	7
1.1	Modular, Optically-Isolated Half-Bridge.....	7
1.1.1	MOSFETs	8
1.1.2	Optocouplers	10
1.1.3	Drive Signal Inverter.....	11
1.1.4	DC/DC Converter (High-Side Supply).....	11
1.1.5	Full Half-Bridge.....	12
1.1.6	Controlling the Half-Bridge.....	15
1.2	Circuit Board Layout and Mechanical Design.....	16
1.2.1	Power/Signal Isolation	16
1.2.2	Power Supplies.....	17
1.2.3	Mechanical Constraints and Design.....	19
2	Field-Oriented Control Strategy	23
2.1	Modeling Field-Oriented Control	23
2.2	Synchronous Current Regulator.....	28
2.3	Modified Synchronous Current Regulator.....	29
3	Testing the Controller	33
3.1	Establishing a Baseline: q-Axis Control Only	33
3.2	Modified Synchronous Current Regulator: One Motor	35
3.3	Modified Synchronous Current Regulator: Two Motors.....	36
4	Brushless DC vs. Brushless AC.....	38
4.1	Six-Step Commutation.....	40
4.2	Sinusoidal Commutation.....	41
4.3	Comparison of Pure BLDC to Pure Sinusoidal Drive	43
4.4	Mixed Motor and Drive: Sinusoidal Drive of a Trapezoidal Motor.....	53
5	Appendices.....	59
5.1	MOSFET Selection and Analysis	59
5.2	Hall Effect Sensor Fault Tolerance.....	69
5.3	Full Controller Schematic.....	72
5.4	Table of Parameters for Test Motors	77
	References.....	78

Table of Figures

Figure 1: The high-level schematic of one power inverter, with a single half-bridge module highlighted.	7
Figure 2: The legend applied to electrical schematics of each module.	8
Figure 3: The MOSFET module, detailing which passives are included on the gate. This format, showing the high-level symbol and its equivalent full circuit schematic, with terminals labeled, will be used for all of the modules.	8
Figure 4: The MOSFET module as defined for multiple MOSFETs in parallel.	9
Figure 5: The module definition of the gate drive optocoupler.	10
Figure 6: The module definition of the drive signal inverter.	11
Figure 7: The module definition of the isolated high-side supply.	12
Figure 8: The high-level schematic of one half bridge. If you've randomly found yourself at this figure, see the module definitions above for expansion to the full schematic.	13
Figure 9: The shoot-through protection at work: One MOSFET gate turns off at the start of the transition(a), while the other MOSFET does not turn on until the end of the transition (b). The capacitor sets the duration of the transition, in this case a relatively large 8 μ s.	14
Figure 10: The high-level schematic of one half bridge, with isolated low-side power supply. If you've randomly found yourself at this figure, see the module definitions above for expansion to the full schematic.	15
Figure 11: Power and signal sections of the controller are physically as well as electrically isolated from each other.	17
Figure 12: The controller power supplies, an excerpt from the full schematic given in Appendix 4.3.	18
Figure 13: Mechanical design of the v1.0 controller, for a single motor.	19
Figure 14: The v1.0 controller was for a single motor, so a stack of two was required on the scooter. The stacked controllers took up foot space.	20
Figure 15: A size comparison of the v1.0 and v2.0 controllers, modeled in the scooter deck.	20
Figure 16: The v2.1 controller fits under a transparent cover flush with the scooter deck, so it does not occupy foot space. It is heat sinked directly to the aluminum chassis.	21
Figure 17: A close-up view of the power connections to the v2.0 controller, which were directly soldered instead of reliant on screw terminals.	22

Figure 18: The three controllers side-by-side: A stack of two v1.0 controllers without fans (right), a single v2.0 controller (middle), and a single v2.1 controller (left).....	22
Figure 19: The simple motor model used to illustrate the d - q coordinate system.....	23
Figure 20: Three current shown as vector quantities, their direction being defined by the coil they flow through (left). The resultant current, a vector sum of the three individual currents, is entirely on the d -axis (right).....	24
Figure 21: The optimum current vector for producing torque in a non-salient permanent magnet motor is on the q -axis, leading the magnets by 90° electrical.....	25
Figure 22: An illustration showing the magnetic flux, back EMF, and current vectors for optimal torque.	26
Figure 23: Current lag created by the motor inductance can reduce the torque and/or efficiency.	27
Figure 24: Phase advance adds extra lead to the drive voltage to bring the current back onto the q -axis.....	27
Figure 25: The block diagram for a synchronous current regulator, showing how the Park and inverse Park transforms can create a current controller acting in the rotor d - q frame.....	28
Figure 26: The modified synchronous current regulator used in this controller. It replaces the inverse Park transform with a simpler sine wave generator with phase advance.....	29
Figure 27: In the step response with no d -axis control, q -axis current still reaches 20A but d -axis current is non-zero. The total current is 23.3A.....	31
Figure 28: Comparison of the standard and synchronous current regulator step response to a 20A torque command with fixed motor speed. The total steady-state current is 20A in each case.	32
Figure 29: The B.W.D. Scooter, test vehicle for this controller, as two integrated 500W hub motors in its wheels.	33
Figure 30: A baseline controller with q -axis control only. The phase advance angle is always zero, i.e. the coil timing is fixed to the Hall effect sensors.....	34
Figure 31: Results of the baseline test with q -axis control only. As expected, some current lags behind onto the d -axis.....	35
Figure 32: The same motor and load as the baseline, now controlled with the modified synchronous current regulator.....	36
Figure 33: Field-oriented control of two motors simultaneously. d -axis current stays near zero at all speeds.....	37

Figure 34: The back EMF waveforms of two fictitious motors, one trapezoidal and one sinusoidal. If they are 14-pole motors, these waveforms both correspond to a speed of 635RPM.	39
Figure 35: An example of sinusoidal PWM at 15kHz showing the actual voltage applied to the motor terminal (red) and the average voltage (blue). For clarity, only a quarter-period is shown.	42
Figure 36: The back EMF and drive waveforms of the two ideal cases being considered: pure BLDC (a) and pure sinusoidal (b).....	43
Figure 37: The drive current (red) resulting from six-step commutation with motor inductance taken into account.	47
Figure 38: The instantaneous power converted through the back EMF.	48
Figure 39: The drive current (red) resulting from sinusoidal commutation with motor inductance taken into account.	49
Figure 40: The phasor diagram method of deriving the drive current in the sinusoidal motor with inductance taken into account.	49
Figure 41: The power output of the pure sinusoidal motor/drive, with inductance factored in, is much lower than the ideal case but still constant.....	50
Figure 42: A geometric solution to the optimum phase advance angle for this drive voltage amplitude. Black dimensions are constraints, gray dimensions are driven.	51
Figure 43: The 22° phase advance brings current (red) and back EMF (green) back in phase, though the drive current is still attenuated and power is lower than the ideal case with zero inductance. Here the drive voltage is also shown (blue) to illustrate the phase advance.	52
Figure 44: The 10V trapezoid is better approximated by the 12.16V sine wave, its fundamental component.....	54
Figure 45: The simulated power output (green) and power dissipated (blue) of 15.5V sinusoidal drive applied to 10V trapezoidal back EMF <i>with no inductance</i>	55
Figure 46: The simulated response of a trapezoidal motor to sinusoidal drive with no phase advance. The drive current (red) lags as expected and the power output (green) is much lower than the ideal zero-inductance case.....	55
Figure 47: The simulated response of a trapezoidal motor to sinusoidal drive with 15° phase advance. The drive current (red) is now in phase with the back EMF power output (green) is higher, but has more ripple.	56
Figure 48: The electrical symbol for an N-channel power MOSFET.	59

Figure 49: A typical set of power MOSFET I-V curves, showing that above 8.0V gate voltage, the output characteristic is that of a constant resistance, in this case about 3mΩ. Source: <http://www.irf.com/product-info/datasheets/data/irfb3077pbf.pdf>..... 60

Figure 50: A typical I-V curve for the body diode of a MOSFET, in this case the IRFB3077PbF. Source: <http://www.irf.com/product-info/datasheets/data/irfb3077pbf.pdf> 62

Figure 51: The effective internal capacitance from the gate to the source causes finite turn-on and turn-off time..... 63

Figure 52: An excerpt from the IXYS GWM 100-01X1 MOSFET module datasheet. Source: http://download.siliconexpert.com/pdfs/2008/12/03/semi_b/1/ixy/mosfet%20modules/gwm100-01x1.pdf 65

Figure 53: Two IXYS GWM 100-01X1 modules installed on the bottom side of the scooter controller..... 66

Figure 54: A worst-case sensor fault transient would instantaneously apply the full phase voltage 180° out of phase with the back EMF. This transient would cause the most stress on the DC bus capacitor..... 69

Figure 55: A simulation of the DC bus voltage and fault current for a worst-case sensor fault on the two scooter motors..... 70

1 Design of the Controller

1.1 Modular, Optically-Isolated Half-Bridge

The half-bridge is the fundamental unit of each power inverter used in the controller design. It is called a half-bridge because two of these modules could make a full H-bridge, which can be used to control reversible brushed DC motors. The term “phase leg” is also often used, since each half-bridge corresponds to a single phase of the brushless motor. Figure 1 shows a high level schematic of one inverter, highlighting an individual half-bridge module. Each inverter consists of three half-bridge modules and a DC bus capacitor. The purpose and sizing of the DC bus capacitor is discussed in detail in section TK. Although they share a common positive and negative DC voltage, the three half-bridge modules are otherwise isolated from each other.

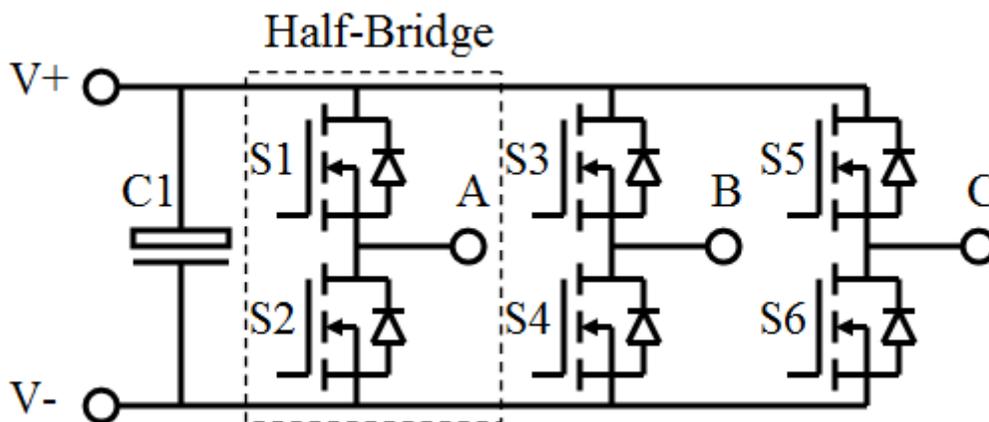


Figure 1: The high-level schematic of one power inverter, with a single half-bridge module highlighted.

The half-bridge module itself is composed of four smaller modules, each of which will be discussed in detail here before recombining them to look at control of the entire half-bridge module. *This is only one design for a modular half bridge and is not meant to represent standard practice.* It is the author’s preferred method and has some advantages that will be discussed.

Each half-bridge module consists of the following sub-modules:

1. MOSFETs (2). These are the high-speed power switches in the inverter. Alternative designs may use IGBTs or other transistors, but the focus here will be on MOSFETs.
2. Gate Drive Optocouplers (2). These are specially-design optocouplers that produce an output suitable for driving MOSFET gates. The input is an LED, and the input and output are electrically isolated. The LED light conducts the signal.
3. Isolated DC/DC Converters (1 or 2). These produce isolated power supplies for the gate drivers. This is necessary for the high-side driver and can also be used for the low-side.
4. Logic-level inverter (1). This is a useful IC for allowing both MOSFETs to be driven by a single logic signal.

As will be seen later, some of these modules may be combined into single-package ICs, even across half-bridges or inverters. However, this doesn’t affect the design and it is possible to create the equivalent circuit out of individual components.

Each of the modules of a single half-bridge will now be discussed in more detail. A high-level symbol and the associated full schematic for each module will be presented in figures. The following key applies to these Figures:

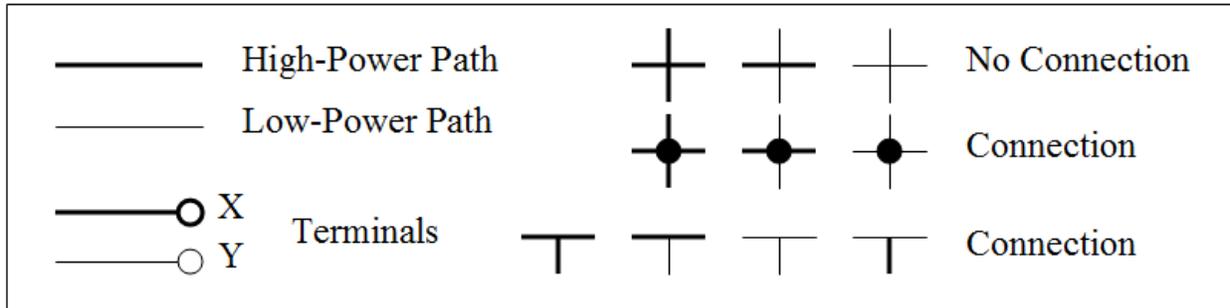


Figure 2: The legend applied to electrical schematics of each module.

1.1.1 MOSFETs

A detailed discussion about the selection and thermal analysis of individual MOSFETs is presented in Appendix 5.1. This section will look at the MOSFET and associated passive components used to create the power-switching part of the half bridge. It may consist of a single MOSFET or a group of MOSFETs grouped in parallel for high current capacity. First, the single MOSFET case is considered:

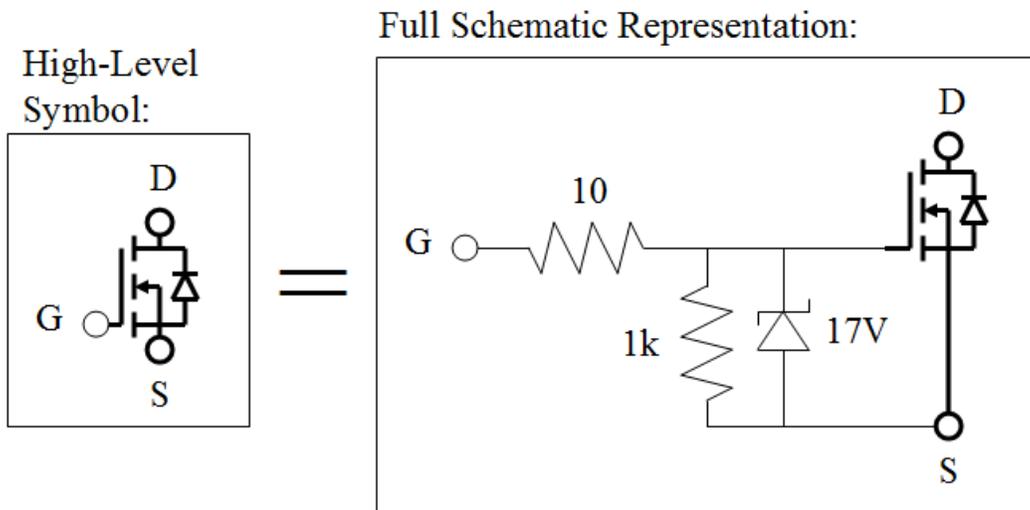


Figure 3: The MOSFET module, detailing which passives are included on the gate. This format, showing the high-level symbol and its equivalent full circuit schematic, with terminals labeled, will be used for all of the modules.

The MOSFET module as defined here includes the N-channel MOSFET and three passive components. The 10Ω resistor, called the “gate resistor,” limits the current that is sourced or sunk when the gate driver switches the gate on or off. Since the gate driver sources 15V, the maximum current draw is easily found to be 1.5A. The gate driver must be able to handle this amount of current for brief periods of time during switching.

The 1kΩ resistor is called a “pull-down” resistor. Strictly speaking, it is not necessary during proper operation of the gate driver. However, if the gate driver fails or becomes high-impedance for any reason, the 1kΩ resistor will allow the gate a path to discharge, turning off the MOSFET. Thus, it is a fail-safe feature. The value of 1kΩ is not critically important; any resistor that can discharge the gate in adequate time but not draw much additional current from the gate driver during normal operation would work.

The 17V Zener diode is a special type of diode called a “transient voltage suppressor,” or TVS. It is designed to break down when reverse-biased to greater than 17V, which is why it faces from source to gate. When it breaks down, it holds the voltage across its terminals to 17V, absorbing any excess energy and dissipating it as heat. It is also not necessary under normal operating conditions, when the gate voltage should never exceed 15V. However, in the event of a gate driver failure that produces an overvoltage, the TVS diode can protect the MOSFET gate itself. Conversely, if the MOSFET fails, the TVS diode may protect the gate driver from overvoltage. Thus, it is also a feature designed for added protection.

Next, the case of multiple MOSFETs in parallel is considered. This is often done for increasing the current capacity of a motor controller. MOSFETs are particularly good at parallel operation due to a positive temperature coefficient, meaning their effective resistance increases with temperature. This automatic feedback mechanism allows them to share current evenly. However, there are some concerns with paralleling MOSFETs that are addressed more thoroughly in [TKrefs]. Only the basics will be discussed here. The schematic representation for paralleled MOSFETs looks like this:

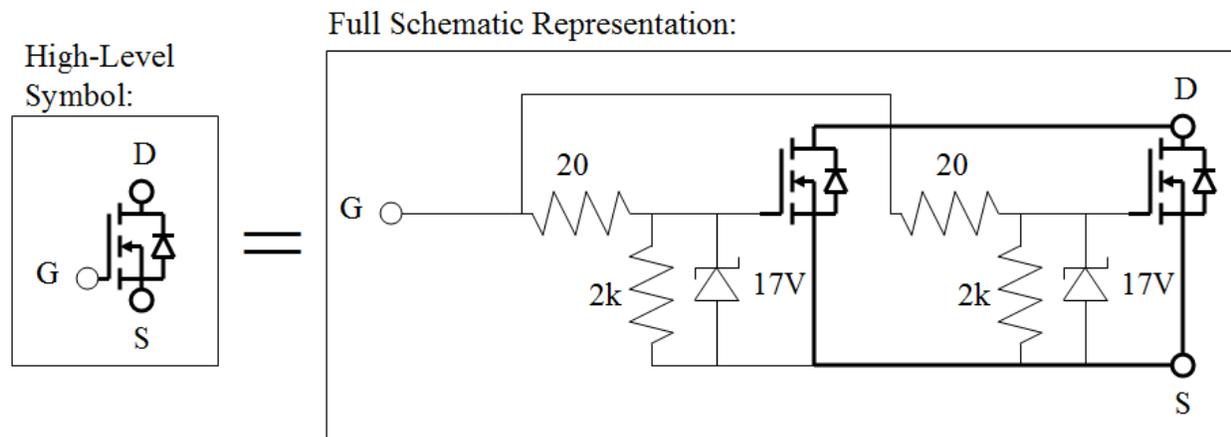


Figure 4: The MOSFET module as defined for multiple MOSFETs in parallel.

There are many ways to parallel MOSFETs. In the method depicted by Figure 4, each individual MOSFET has its own gate resistor and pull-down resistor. The resistor values are chosen so that the gate drive sees the same effective impedance of approximately 10Ω. However, since there are two MOSFETs, the turn-on will take roughly twice as long. Similarly, with four MOSFETs and gate resistors of 40Ω, turn-on would take four times as long. Although paralleling MOSFETs decreases the total on-resistance, it will also increase switching losses unless a more powerful gate driver is used.

Giving each MOSFET its own gate resistor is important for the prevention of ringing or other transients caused when MOSFETs are paralleled. From the gate driver to either gate, the effective impedance is 20Ω . However, the impedance from one gate to the other is 40Ω . Each gate is more likely to follow the gate driver than the other gate. If a single 10Ω resistor were used ahead of both gates, then the impedance from gate to gate would be very low. This could lead to unwanted transients.

Giving each MOSFET its own pull-down resistor and TVS diode is more a matter of choice, since both of those components are fail-safes. A single TVS diode and pull-down resistor could be used to protect entire parallel grouping. However, since these components are relatively small and inexpensive, there is not much harm in using one for each MOSFET.

1.1.2 Optocouplers

The optocouplers are special ICs that combine the function of a normal optocoupler with that of a low-side MOSFET gate driver. The purpose of the optocoupler is to electrically isolate one part of a circuit from another, in this case the sensitive signal lines from the noisy power inverter. It achieves this isolation by using light to carry the signal instead of copper traces. An LED shines on a photo sensor, all inside the IC, to convey the signal across an isolation barrier. A special class of optocoupler designed for driving MOSFET or IGBT gates uses the LED input signal to control a push-pull stage that is capable of both sourcing and sinking relatively high current bursts. There are a few examples of this, but the one used here is made by Avago Technologies, part number HCPL-3120 [2]. The module is simply defined as follows:

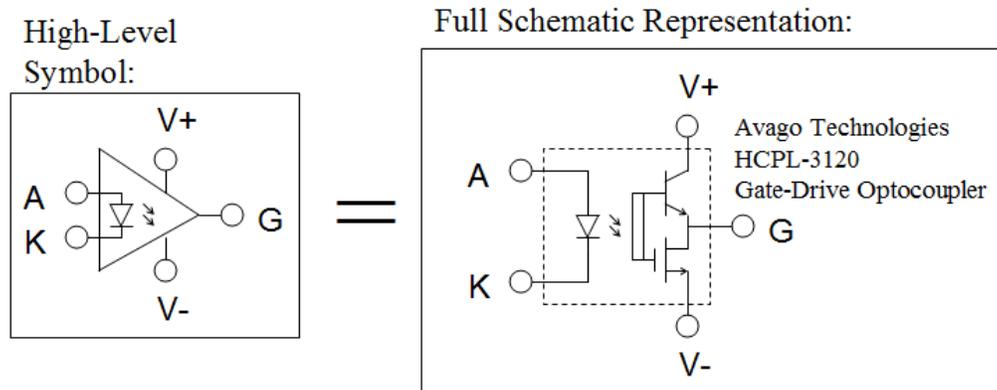


Figure 5: The module definition of the gate drive optocoupler.

The input to the module is considered to be a regular LED. The output is a gate drive stage that requires a positive and negative supply. The gate drive output is either connected to $V+$ or $V-$ based on the state of the LED. It can't take an intermediate or high-impedance state. This IC has a maximum output of $2.0A$ (sourcing or sinking). It also has a low-voltage cut-off at $11V$. No passives are included in the definition of this module, since they will be lumped into other modules or into the full half-bridge circuit assembly.

1.1.3 Drive Signal Inverter

A logic-level inverter is used to invert the high-side drive signal. This is common in half-bridge drives that use synchronous rectification, where one or the other switch is always on. However here it will be used in an unconventional way that also allows for both switches to be turned off. For now, only the module definition is presented:

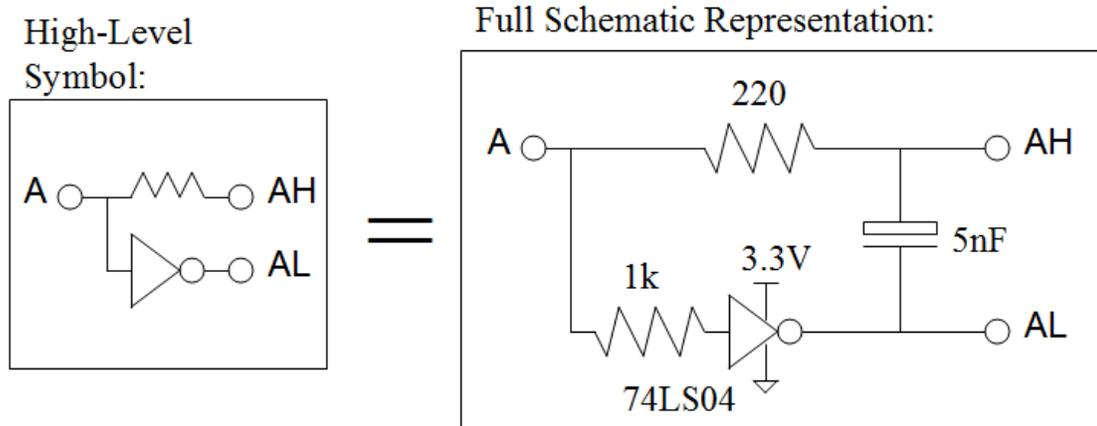


Figure 6: The module definition of the drive signal inverter.

The inverter itself is part of a standard logic IC, the 74LS04. This IC actually contains six inverters, so it can be shared between up to six half-bridges. Here, only one is needed. It is shown powered by 3.3V, the signal voltage used in this controller, but would work equally well with 5V. The 1k Ω resistor limits the current draw from the input pin. The 220 Ω resistor limits the LED current sent to the optocouplers. It is sized the same way as any typical LED current-limiting resistor. The value of 220 Ω should be well-sized for either 3.3V or 5V signals.

The 5nF capacitor across the output lines is a passive protection against shoot-through, the condition that occurs when both switches in a half bridge are temporarily on. This can occur during a transition; since it takes some time to charge and discharge the MOSFET gates, they can momentarily be in a state where both are on or partially on. Current passes through both MOSFETs, directly short-circuiting the DC bus. At best, this is inefficient. At worst, this can lead to controller destruction if both MOSFETs stay on for any significant length of time (milliseconds, even). The 5nF capacitor creates a slight delay between turn-off and turn-on, ensuring that this condition cannot occur. It only works because of the way the two outputs are connected to the high- and low-side optocouplers, so a full discussion is deferred until the full circuit is presented. However, it is a completely passive solution to the shoot-through problem: no additional hardware or software delay is needed.

1.1.4 DC/DC Converter (High-Side Supply)

The last module of the half-bridge is a special isolated DC/DC converter for the high-side supply. While there are many ways to do high-side drive, and this is a relatively expensive one, the advantage it offers in simplicity and modularity make it appealing. The purpose of the DC/DC converter is to create a power supply that is +15V with respect to the high-side MOSFET source. This source may in turn be at almost any voltage between the DC bus rails,

including the positive DV voltage. As such, the high side supply may need to create a voltage up to 15V higher than the positive DC rail. For this reason, an isolated supply with a 1:1 conversion is chosen. The specific IC used here is made by Texas Instruments, part number DCP021515 [3]. The module definition is presented here:

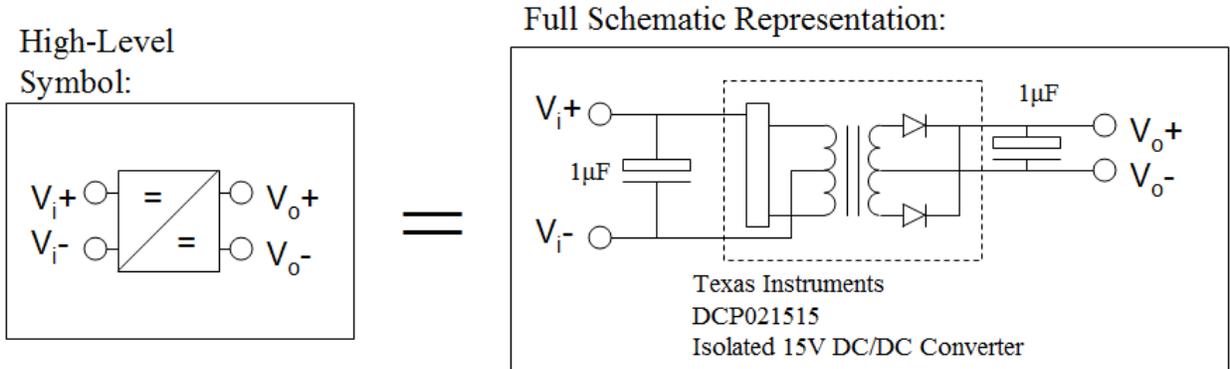


Figure 7: The module definition of the isolated high-side supply.

The DCP021515 is a 15V to 15V isolated, unregulated converter. It operates by transforming the 15V DC input to AC, sending it through a transformer, then rectifying it back to DC. The magnetic coupling in the transformer allows the input and output to be electrically isolated. It is unregulated, so the output may not be exactly 15V. In fact, is typically higher until loaded. The maximum gate voltage tolerable by the MOSFET is 20V, but its 17V TVS diode will protect it against overvoltage. The 1µF capacitors are recommended in the DCP021515 datasheet for smoothing the input and output. The maximum continuous output of is 2W, which is well below the gate drive requirements for this controller.

There is no reason, other than cost and board space, why the isolated supply cannot also be used on the low-side gate drive. In this case, the negative DC bus rail and the negative gate drive supply voltage would be electrically isolated, offering extra noise immunity. The author has successfully tried both configurations, and they will both be shown in the full circuit diagram below. The single isolated supply version is used in the controller presented here for lower cost and component count.

1.1.5 Full Half-Bridge

The four modules discussed above are combined to form a single half-bridge. The high-level symbols defined in the figures above are used to show the full half-bridge in one uncluttered image. For the full electrical schematic, see Appendix 5.3. The high-level schematic is shown in Figure 8.

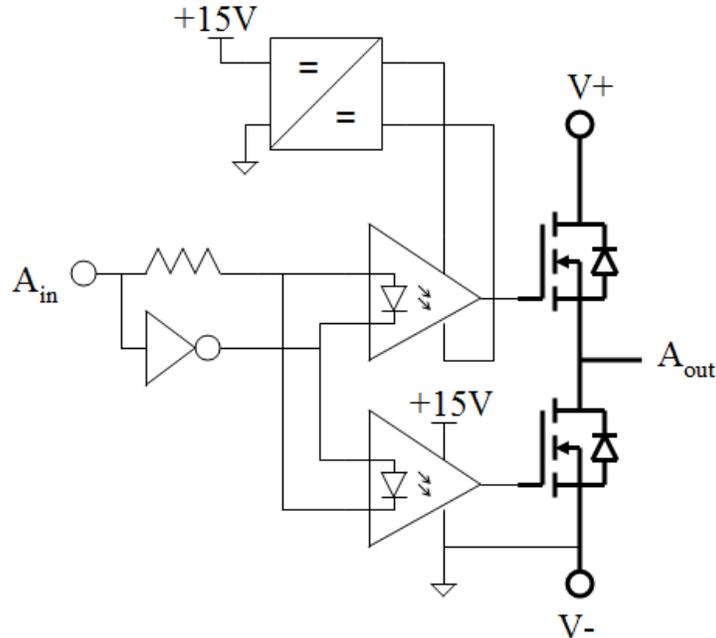


Figure 8: The high-level schematic of one half bridge. If you've randomly found yourself at this figure, see the module definitions above for expansion to the full schematic.

Together, this makes an isolated half-bridge with passive shoot-through protection. The V+ and V- terminals are intended to be connected to the DC bus. A_{out} connects to a single phase of a motor or other load, and A_{in} is the isolated input signal that controls A_{out}. A single +15V external supply is required to power the gate drivers. This can be derived from the DC bus, and in this case its ground is tied to the negative DC voltage.

Importantly, the inverter module outputs are connected to the two optocoupler LEDs in reverse-parallel. This ensures that both LEDs cannot be on at the same time, protecting against long-duration shoot-through. Combined with the output capacitor of the inverter module, short-duration shoot through is also prevented. In order to turn off one LED and turn on the other, the voltage across the inverter module outputs must go from 1.5V to -1.5V or vice versa. In between these two voltages, neither LED is on. The capacitor ensures that the time it takes to change over is longer than the time it takes one MOSFET to turn off. An example of what this signal looks like is shown in the Figure 9. The shoot-through delay can be directly measured or estimated using the formula:

$$\Delta t \approx RC \left(\frac{2V_f}{V_{sig}} \right),$$

where R is the LED current limiting resistor, C is the delay capacitor, V_f is the forward voltage of the optocoupler LED, and V_{sig} is the signal voltage being used. Using the values in this controller, the delay time is estimated to be:

$$\Delta t \approx (220\Omega)(5nF) \left[\frac{2(1.5V)}{3.3V} \right] = 1\mu s .$$

This is a conservative value that is much longer than the gate turn-on / turn-off time. It is still only a small fraction of the switching period.



Figure 9: The shoot-through protection at work: One MOSFET gate turns off at the start of the transition(a), while the other MOSFET does not turn on until the end of the transition (b). The capacitor sets the duration of the transition, in this case a relatively large $8\mu s$.

An alternative configuration, with a low-side DC/DC converter, is shown in Figure 10. This configuration differs only by the addition of a second isolated DC/DC converter, supplying the low-side gate drive optocoupler. While this is more expensive and has a higher component count, it offers the significant advantage of isolating the gate drive power supply entirely from the primary power. (That is, V- is no longer tied to the low-side gate driver.) In fact, this configuration creates three completely isolated power blocks: primary power for the inverter, 15V power for the gate drive, and 3.3V or 5V signal power. In this case, these three blocks need not share a ground.

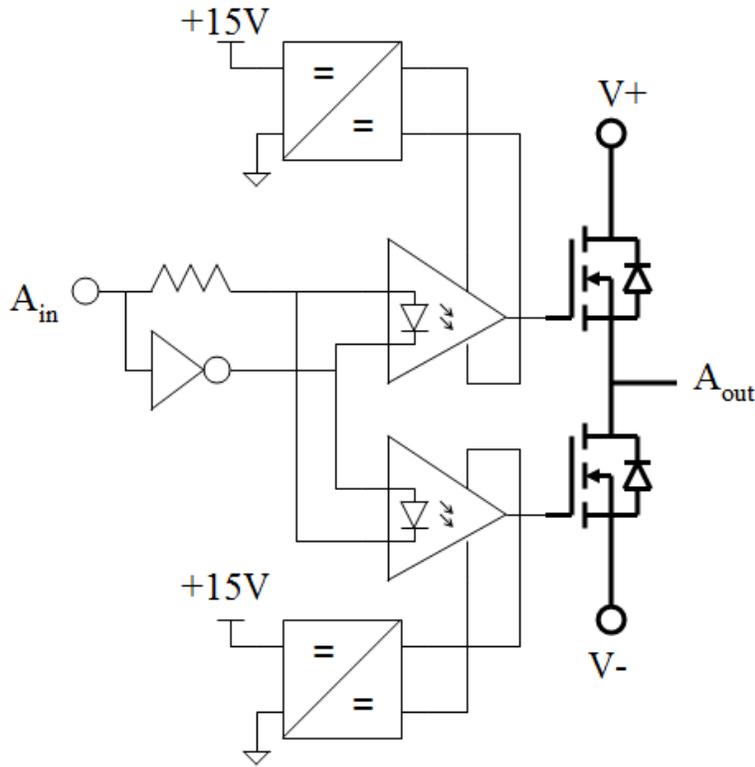


Figure 10: The high-level schematic of one half bridge, with isolated low-side power supply. If you've randomly found yourself at this figure, see the module definitions above for expansion to the full schematic.

The controller presented here uses the single isolated supply configuration (Figure 8). This choice was made for simplicity, low component count, smaller board size, and lower cost. From this point forward, the discussion will refer to the single isolated supply configuration and the high-level schematic of Figure 8.

1.1.6 Controlling the Half-Bridge

Now that the half-bridge module has been defined, a quick look at its control states is presented. The half bridge output really only has three desirable states: high, low, and freewheeling. These are the possibilities of having one or none of the MOSFETs conducting. If neither MOSFET is conducting, the output is freewheeling. (The term floating is avoided because diode conduction can still occur.) Obviously, the state in which both MOSFETs are on is neither desirable nor possible with the gate drive operating properly.

There is, however, only one input to the half bridge as defined in Figure 8: the input A_{in} . If A_{in} is driven by a 3.3V signal, the high side MOSFET will be on. If driven to 0V, the low-side MOSFET will be on. To achieve the third state, A_{in} must be left floating. Without the ability to source or sink current, neither LED will be able to turn on. This single input tri-stating operation, desirable in any half-bridge controller, comes automatically thanks to the need for LED drive current in this configuration. Most microcontrollers have the ability to drive a pin high, low, or leave it floating (input state), so this half bridge can be entirely controlled by a single microcontroller pin. Table 1 lists the three states of the half-bridge.

Table 1: The three half-bridge states.

Input A_{in}	Output A_{out}	Description
3.3V	V+	High-side MOSFET is on.
0V	V-	Low-side MOSFET is on.
Float	Freewheel	Neither MOSFET is on. Output may voltage undetermined. Diode conduction may still occur!

1.2 Circuit Board Layout and Mechanical Design

The inverter (Figure 1) is composed of three half bridge modules and a bypass capacitor. In this controller, two independent inverters are used to control two motors simultaneously. However, the inverters share the same circuit board and are driven by the same microcontroller. Additionally, support circuitry is required for reading Hall effect sensor signals, throttle commands. Robust power supplies that can operate in the noisy environment are also a necessary.

Unlike signal-level circuit design, much of the performance of a power electronics circuit depends on the *physical* layout of the circuit, in this case how components are distributed and how traces are routed on the PCB. This is integrally tied to the mechanical requirements of the hardware: space constraints, heat transfer, noise and vibration. For that reason, this section combines the mechanical and electrical layout design considerations.

1.2.1 Power/Signal Isolation

As much as possible, the power and signal paths in the controller should be physically and electrically isolated. Many circuit failures occur because power paths create noise on signal lines. Figure 11 shows a top view of the controller circuit board, divided into sections that illustrate physical power and signal isolation. The sections drawn continue on to the back side of the circuit board. The inverter and gate drive sections (two each, for two motors) are composed of the modular half bridge design detailed in Section 1.1. Gate drive is considered a “power” section even though it is relatively low current because it interacts directly with the MOSFETs at voltages higher than 15V.

The inverter sections include two bypass capacitors. Importantly, these two capacitors are placed *as close as possible* to the inverter power inputs; the traces connecting them are wide and short. This limits parasitic inductance and allows the capacitor to supply instantaneous bursts of energy to the inverter. It also helps isolates the rest of the circuit from high-frequency switching noise of the inverter. The importance of the bypass capacitor is further emphasized in [4].

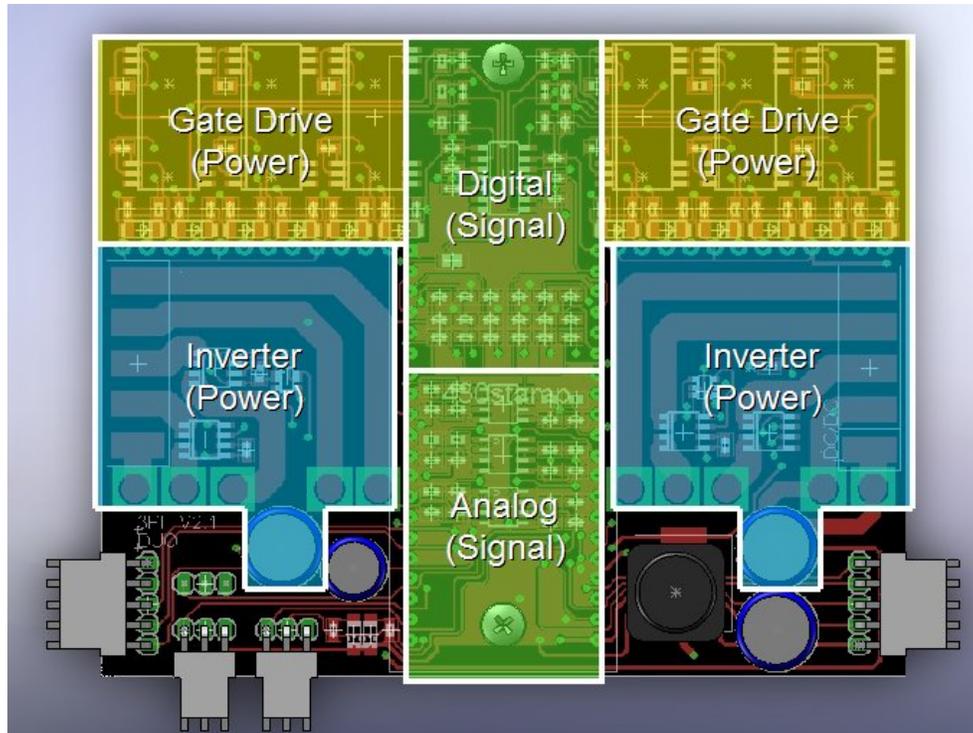


Figure 11: Power and signal sections of the controller are physically as well as electrically isolated from each other.

Physically isolating the power and signal sections is not sufficient; they must also be electrically isolated. Most importantly, the grounds must be well-controlled. One way to think about this: current flowing through the motors should never have to cross through a signal trace or ground plane on its way back out to the negative battery terminal. Simply putting a ground plane across the entire board is almost guaranteed to violate this condition. Additionally, if the power and signal electronics are to share a ground, the ground should be connected by a small trace *in only one place*. This is something not captured by an electrical schematic; it must be enforced during physical board layout.

One piece of isolation already discussed is the optically-coupled gate driver (Section 1.1.2). These optocouplers allow signals to be passed from the protected signal section of the board to the noisy power section by using light as an intermediary.

1.2.2 Power Supplies

An important interface between power and signal sections is at the power supplies. The power for microcontrollers, op-amps, inverters, and all signal-level devices is derived from the primary battery through several stages of power conversion. This is, in many cases, the point where interference from power electronics crosses over to the signal side. To prevent this, careful attention to the power supply layout is required. Figure 12 shows the power supply schematic used in this controller, part of the full schematic given in Appendix 5.3.

POWER SUPPLIES:

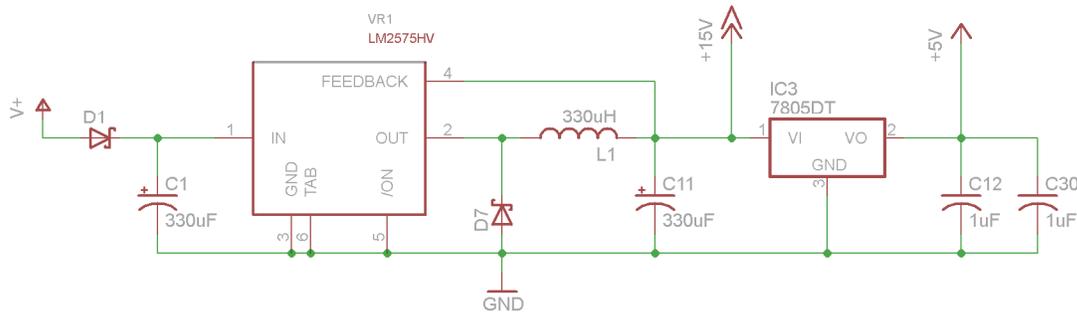


Figure 12: The controller power supplies, an excerpt from the full schematic given in Appendix 5.3.

The first power supply is an efficient switching regulator, that converts battery voltage (24-48V) to 15V. The regulator used is made by National Semiconductor, part number LM2575HV [5]. The “HV” stands for high voltage. This version can take up to a 63V input. The non-HV version can take a maximum input of 45V. Because it is a switching regulator, the LM2575HV requires an external diode (D7) and inductor (L1) to operate. It also requires a capacitor (C11) to smooth the output.

The switching regulator has its own bypass capacitor, C1, placed as close as possible to its inputs. This capacitor isolates the power supply from noise created by the inverters. Power from the battery, labeled V_+ , passes through a Schottky diode, D1, before even reaching the power supply input capacitor. This prevents current from flowing back from C1 to the positive DC bus, further isolating the power supplies from sudden increases in load on the inverters.

The 15V supply is used for powering the gate drivers. Thus, it is technically still on the “power” side of the circuit, susceptible to noise from the inverters. It can be further isolated from inverter noise by using the fully isolated half bridge layout, shown in Figure 10. However, in this controller the 15V supply directly powers the low-side gate drivers, as in Figure 8. Thus, the 15V supply is still susceptible to inverter noise.

The output of the 15V supply is used to create the signal-level power supply (5V). A common linear regulator (7805, made by several manufacturers) is used to do this step-down. Linear regulators produce a clean output (no switching), which is good for signal electronics. However, they dissipate a lot more heat than switching regulators. The full voltage drop at operating current is dissipated. Thus, at just 0.1A of current, the 7805 regulator here would be dissipating 1W. It is only required to power signal-level components, so the current draw is small, but heat dissipation is still an important consideration. The surface mount component is directly soldered to the ground plane, which can help with heat sinking.

The ground connection between the 15V power supply and the 7805 is *the only place* where signal and power grounds connect. This single-point connection is very important, as discussed in Section 1.2.1. At the output of the 7805 linear regulator are two 1uF ceramic capacitors for additional smoothing. These values could certainly be increased. All components connected to

the signal ground should also have their own bypass capacitors. The microcontroller board used in this controller has additional power supplies (3.3V) and smoothing capacitors.

1.2.3 Mechanical Constraints and Design

While these mechanical constraints are specific to one application of the controller (electric scooter), they provide an interesting insight into the design considerations of circuit layout that go beyond simple electrical connections. Space constraints, heat transfer, and vibration all played important roles in the design of the controller physical layout. To illustrate this, the v1.0 and v2.0/v2.1 controllers are compared.

Figure 13 shows the v1.0 mechanical design. It controlled one motor using discrete MOSFETs attached to aluminum heat sinks as the inverter stage. The heat sinks doubled as bus bars, with holes for wire connecting screws. (This heat sink/bus bar design is similar to that used in many commercial electric vehicle controllers.) Because the heat sinks were relatively thin, cooling was improved with a fan blowing directly on the inverter components.

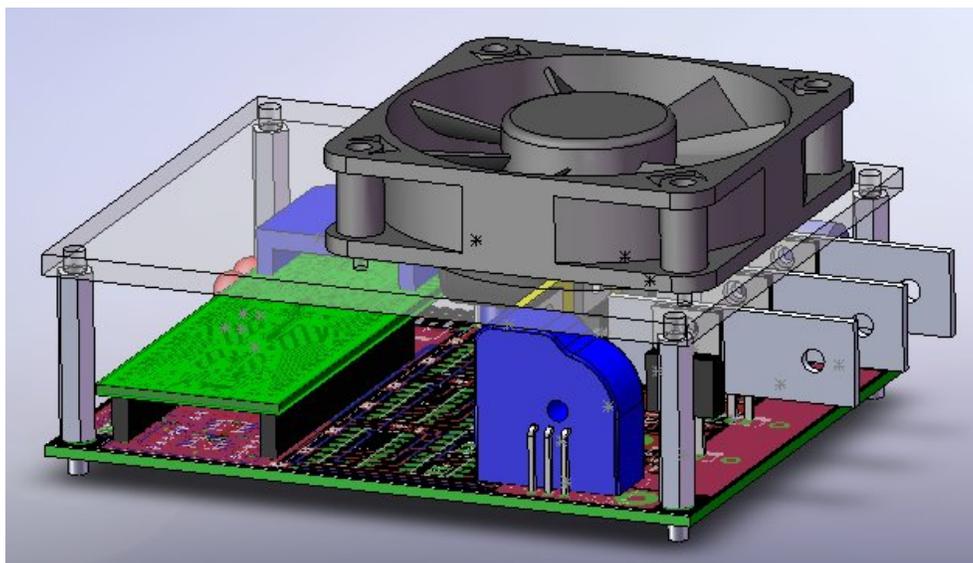


Figure 13: Mechanical design of the v1.0 controller, for a single motor.

While this design could likely have driven motors of 2kW or more, the power was unnecessary for the scooter and because it could only control one motor, twice this volume was needed. The controllers could be stacked to conserve space, but then the bottom fan would have reduced air flow. Figure 14 shows the physical space occupied by the v1.0 controller. In addition to taking up foot space on the scooter, it was visually unappealing.



Figure 14: The v1.0 controller was for a single motor, so a stack of two was required on the scooter. The stacked controllers took up foot space.

There were other mechanical problems with the v1.0 controller. It relied heavily on fasteners, for connecting MOSFETs to heat sinks, wires to bus bars, fans to covers, and for stacking controllers. These fasteners were prone to loosening due to mechanical vibration. The fan was also prone to vibration damage. Blades on both fans were broken during testing. Through-hole electrical components also tended to be loosened from sockets by vibration.

The v2.0 and v2.1 controller integrated two motors into a single board, saving a lot of space and allowing the controller to fit within the volume of the scooter deck where it would not occupy foot space. Figure 15 shows the v2.0 controller next to the stack of two v1.0 controllers. The v2.1 controller reduced the size even further by eliminating two large current sensors (blue) in favor of more compact surface mount current sensors. Figure 16 shows the v2.1 controller mounted inside the scooter deck.

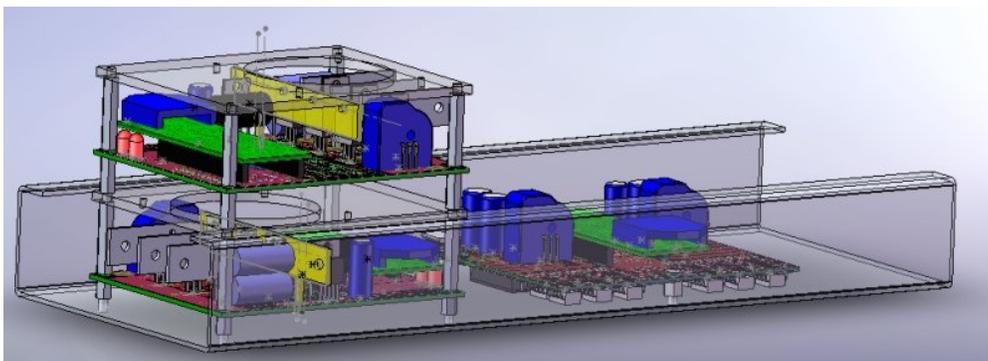


Figure 15: A size comparison of the v1.0 and v2.0 controllers, modeled in the scooter deck.

To achieve these space savings, many components were changed from through-hole to surface mount. In particular, switching to a surface mount MOSFET module that integrates all six MOSFETs into a single package was a key space-saving change. The power capability of this MOSFET package is less than the six discrete MOSFETs in the v1.0 controller, but still sufficient for the scooter motors. The MOSFET package also has a large heat sink plate that is in contact (with thermal paste) with the aluminum scooter chassis. The heat sink plates are shown in Figure 53 This eliminates the heat sink/bus bar and fan. For a detailed discussion of MOSFET selection and thermal analysis, see Appendix 5.1.



Figure 16: The v2.1 controller fits under a transparent cover flush with the scooter deck, so it does not occupy foot space. It is heat sunk directly to the aluminum chassis.

Other measures were taken to minimize the affect of mechanical vibration on the controller. Surface mount components were used wherever possible in place of sockets and through hole components. Power connections are directly soldered to the circuit board, rather than attached to bus bars with screw terminals, which also saves space. A close up of the power connections is shown in Figure 17.

The progress from v1.0 to v2.1 represents a good deal of mechanical modification, making the latest version much more compact and tolerant to vibration. Some power was sacrificed to fit two motor drives on a single board, but with dual outputs and good passive heat transfer to the scooter chassis, the v2.1 controller still has a respectable power density. Figure 18 shows the three versions next to each other as a final size comparison.

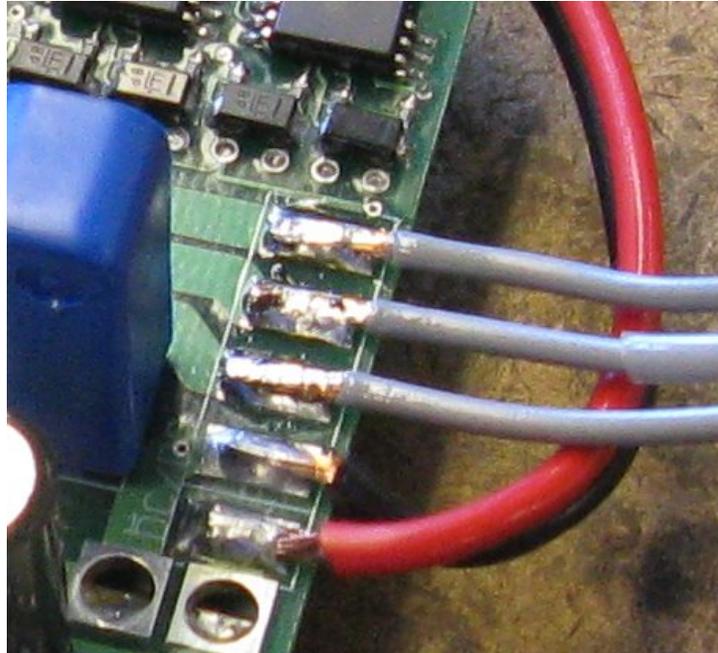


Figure 17: A close-up view of the power connections to the v2.0 controller, which were directly soldered instead of reliant on screw terminals.

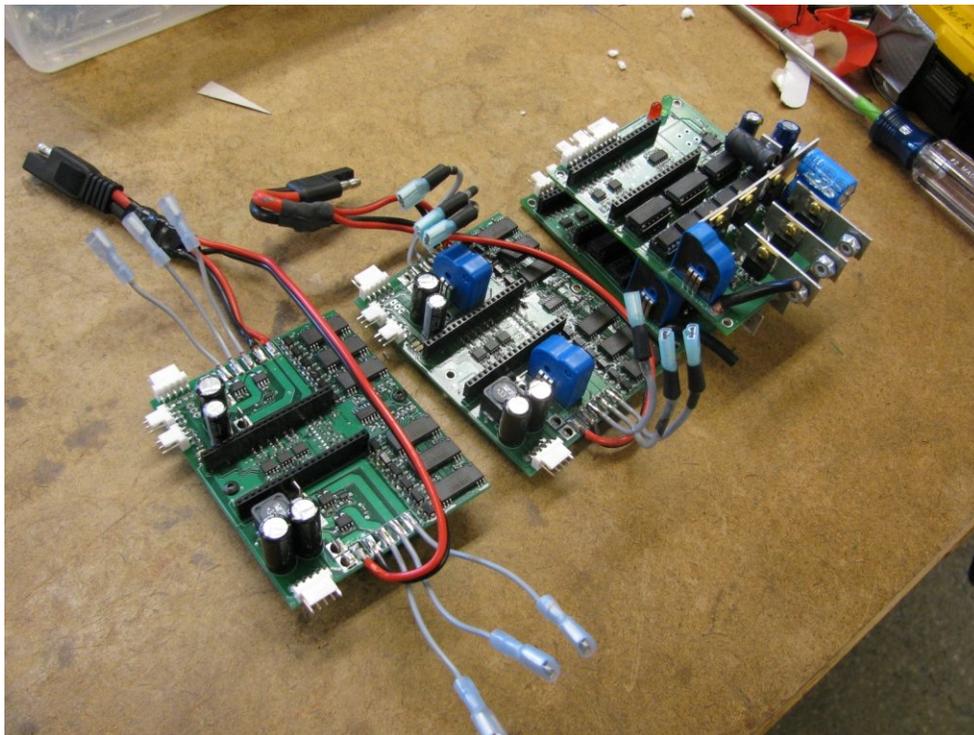


Figure 18: The three controllers side-by-side: A stack of two v1.0 controllers without fans (right), a single v2.0 controller (middle), and a single v2.1 controller (left).

2 Field-Oriented Control Strategy

The controller presented here employs field-oriented control, an advanced method of motor control that offers several advantages over traditional BLDC control. Field-oriented control uses sinusoidal commutation, which results in quieter motor operation and less vibration than BLDC control. It relies on accurate rotor position estimates and current sensing to determine both the magnitude and the relative phase of current in the motor windings. Using the current vector as a feedback element, field-oriented control adjusts the timing of sinusoidal commutation to achieve optimal torque even in the presence of current lag due to winding inductance.

2.1 Modeling Field-Oriented Control

A simple motor model with two poles and a three-phase winding shows the coordinate system used in field-oriented control:

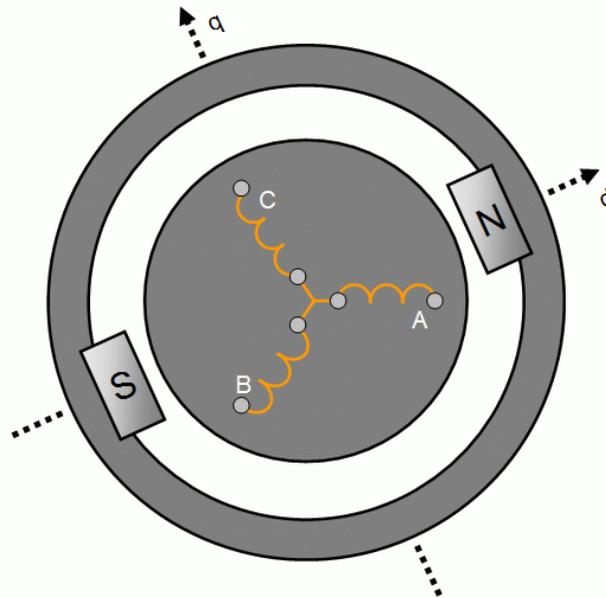


Figure 19: The simple motor model used to illustrate the d - q coordinate system.

This motor has an outside rotor with two magnets, N and S, evenly spaced at 180° intervals. On the inside stator, three winding coils, A, B, and C, are evenly spaced at 120° intervals and connected at a single point (wye-connected). Since the problem is two-dimensional, a set of orthogonal axes is useful. In motor control, the axes are called d (direct) and q (quadrature) instead of x and y . In permanent magnet motors, the d -axis is defined to be on the magnetic poles of the rotor. The q -axis is defined to be in between magnetic poles, specifically one quarter cycle ahead of the d -axis in the direction of rotation. (Thus, in Figure 19, a counterclockwise rotation is implied.) The axes are fixed to the rotor, rotating as seen from the stator.

In a motor with more than two poles, the d -axis and q -axis are not physically orthogonal. The d -axis is always on magnetic poles and the q -axis is always in between magnetic poles. In a four pole motor, for example, they are physically separated by 45° . However, they are still considered to be 90 electrical degrees apart. From a control standpoint, the only difference between a two-pole motor and a four-pole motor is the number of electrical cycles per mechanical revolution (one and two, respectively). The number of poles, then, is like a gear ratio. However, the angles used in field-oriented control are the electrical angles. This makes it possible to use the simple two-pole motor, in which the electrical and physical coordinate systems are the same, as a general model for field-oriented control.

One basic premise of field-oriented control is that motor parameters are vector quantities. That is, they have both a magnitude and a direction. The direction of certain quantities is obvious: magnetic flux from the permanent magnets is always aligned with the d -axis, for example. Quantities that can be attributed to coil A, B, or C take the direction defined by that coil. Figure 20 show an example of this.

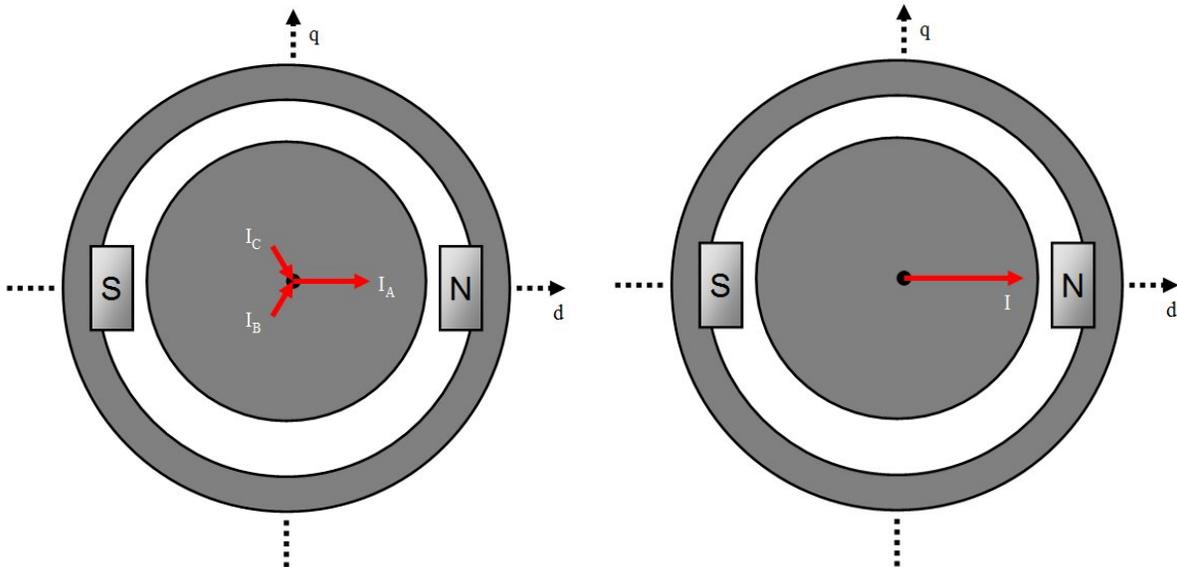


Figure 20: Three current shown as vector quantities, their direction being defined by the coil they flow through (left). The resultant current, a vector sum of the three individual currents, is entirely on the d -axis (right).

In this example, current is sent into coils B and C and comes out of coil A. The magnitude of each current, indicated by the length of the red arrows, is such that the sum of currents into the motor is zero. This must be true since the three coils are connected at a single point. By adding the three current vectors, the resultant current can be found. In this case, the resultant current is entirely on the d -axis. Note that this is only true when the rotor is in this exact position. With the rotor in another position, the same three currents could produce a different vector in the d - q plane.

Torque is produced by interaction of stator current with the rotor magnetic field. The stator current turns the stator coils and steel into an electromagnet, with its poles oriented in the direction of the current vector. The permanent magnets will try to align themselves with the stator-generated poles. Thus, the d -axis will try to align with stator current. In Figure 20, this is

already the case, and the rotor will resist moving from its position. In order to produce positive torque, stator current must always lead the rotor magnets, pulling them along. The optimum lead angle for producing torque in a non-salient permanent magnet motor is 90° *electrical*, i.e. current entirely on the q -axis, as shown in Figure 21.

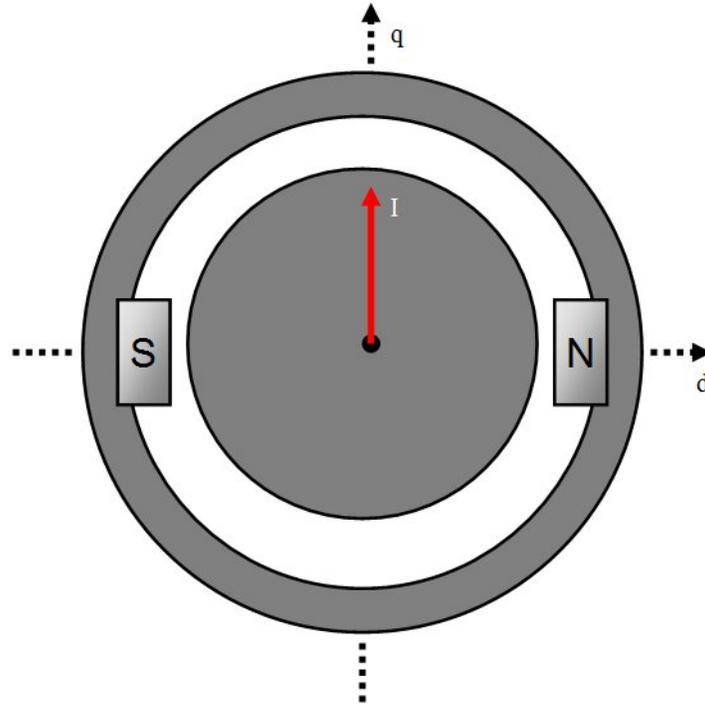


Figure 21: The optimum current vector for producing torque in a non-salient permanent magnet motor is on the q -axis, leading the magnets by 90° electrical.

To produce this current vector at this rotor position using only the three coil directions available, current should be sent into coil B and come out of coil C. Coil A, which is on the d -axis at this rotor position, will have zero current. Part of the job of field-oriented control is to perform these trigonometric projections in real time, creating an arbitrary current vector in the d - q plane.

Another way to prove that the optimum current vector for this type of motor is on the q -axis is to consider back EMF. Back EMF is the voltage that the motor would produce at its three terminals if there was no load. This voltage comes about due to the changing magnetic flux in the three coils as the permanent magnets pass by. Its mathematical definition is:

$$E = \frac{d\lambda}{dt},$$

where λ is the flux linkage of the winding, the amount of magnetic flux passing through the coil times the number of turns in the coil. We already stated that the magnetic flux is a vector quantity and its direction is aligned with the d -axis. As seen from a point on the stator, flux is assumed to be a sinusoidally-varying quantity in time, an assumption scrutinized in Section 3

and in [1]. If this is true, then the time derivative of flux will *lead* the flux itself by 90° electrical. Thus, the back EMF vector will be on the q -axis:

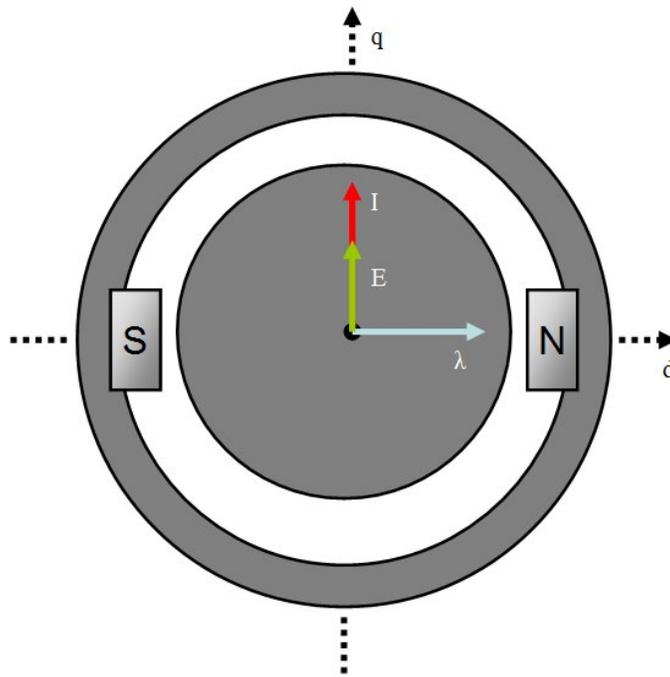


Figure 22: An illustration showing the magnetic flux, back EMF, and current vectors for optimal torque.

Physically, this means that if we measured the voltage on all three motor coils while spinning the rotor, then do the vector sum, the resultant back EMF vector would fall between magnets, on the q -axis. It still rotates with the rotor, but its phase is such that it is always 90° electrical ahead of the magnet axis.

From Figure 22, the optimum placement of the current vector is more obvious. Power converted by the motor is current times back EMF, just as it would be in a brushed DC motor. However, since the motor quantities are now vectors, it is the dot product that matters. For a given back EMF and current, aligning the two vectors produces the most power (and torque). Thus the current should be driven into the back EMF, on the q -axis, for most efficient torque production. As the rotor spins, the controller must continuously place the current vector on the “new” q -axis.

An interesting case to consider is that of fixed coil timing, where the voltage applied to each coil is directly linked to the position of the rotor by some constant mapping. A simple version of this is six-step commutation, where three Hall-effect sensors trigger different switch states that set a coil high, low, or off. A slightly more complex version is sinusoidal commutation, where rotor position is directly linked to a sine table that sets voltages on each phase. In this case, a resultant voltage vector can be defined. An obvious choice might be to set the motor timing such that the voltage vector is on the q -axis. The problem with this choice occurs under load. Because the motor coils are also inductors, they resist rapid changes in current. Thus, the current vector will actually lag behind the voltage vector by some amount:

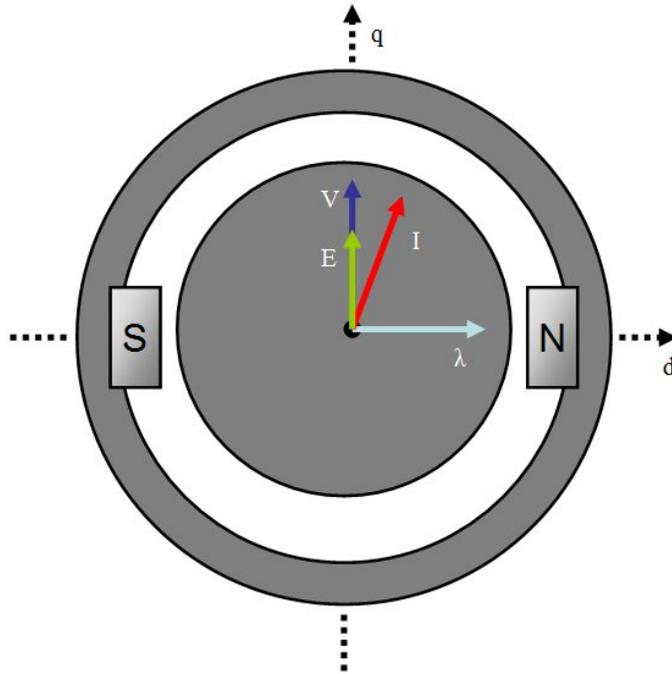


Figure 23: Current lag created by the motor inductance can reduce the torque and/or efficiency.

Current now has a non-zero d -axis component. The current vector is no longer in phase with the back EMF vector, meaning torque production is decreased. To recover the lost torque, the magnitude of the current can be increased, but this will result in higher I^2R losses and lower efficiency. If, however, the voltage applied to the motor windings is advanced by some amount, current will be brought back onto the q -axis:

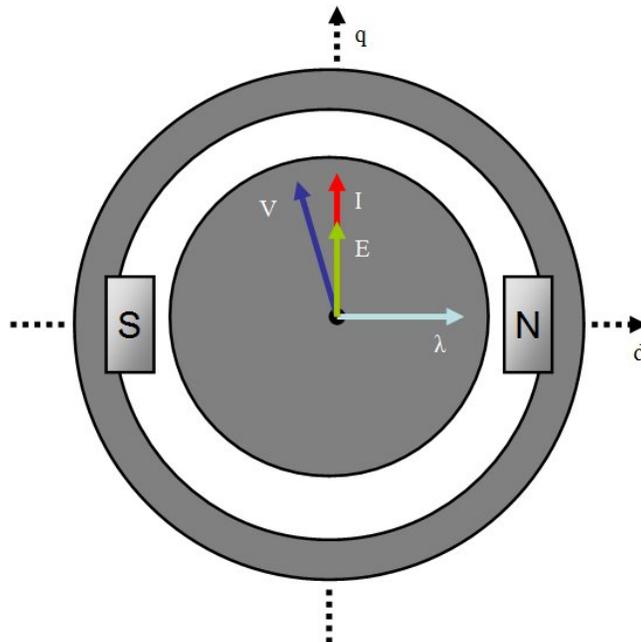


Figure 24: Phase advance adds extra lead to the drive voltage to bring the current back onto the q -axis.

The angle by which the voltage should be advanced to return current to the q -axis is dependent on motor inductance, resistance, and speed. Thus, it is different for every motor and difficult to predict. Still, controllers with dead reckoning schemes for advancing timing can achieve sufficient results, especially if operating conditions are not very variable. Such controllers are said to have “phase advance” or “timing advance,” but may not be using true field-oriented control.

Field-oriented control goes a step further than dead-reckoned phase advance. The current vector is derived in real time based on phase current measurements. Usually, the current is projected into its q -axis and d -axis component. The matrix operation that encapsulates these projections is called the Park transform. The two components are then used in a feedback control system to set the voltage vector in exactly the right location to achieve optimal torque (q -axis only current). There are many different ways to achieve this feedback control. The controller presented here uses a slight modification of a method called a “synchronous current regulator.”

2.2 Synchronous Current Regulator

A standard synchronous current regulator block diagram is shown in Figure 25. The defining feature that separates a synchronous current regulator from any other type of current control is the presence of the Park and inverse Park transforms. These use the instantaneous rotor position to project currents onto the d - q plane. Control action is taken in this frame of reference, then commands are projected back onto the A, B, and C axes as PWM commands.

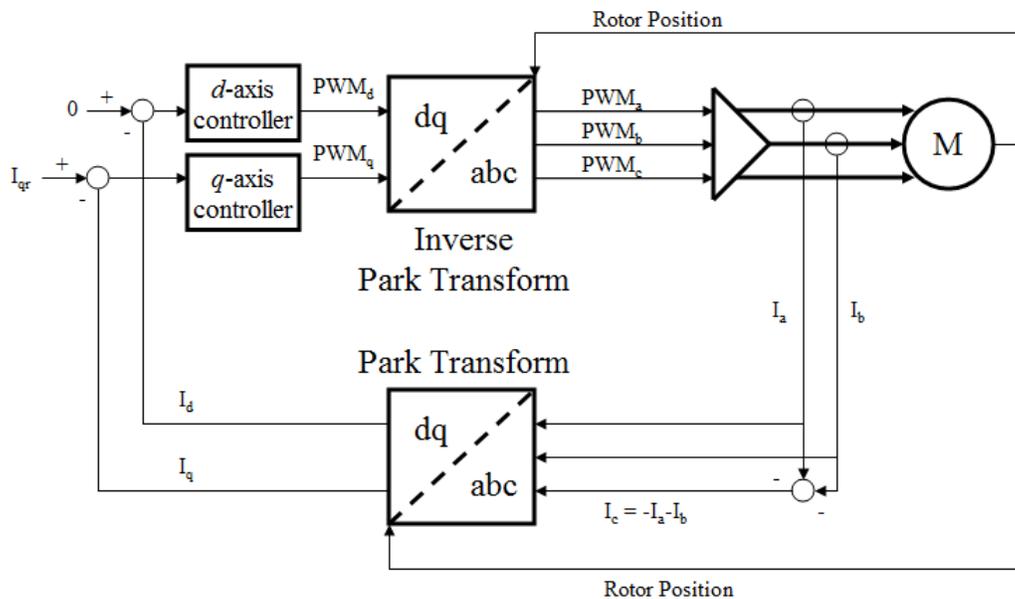


Figure 25: The block diagram for a synchronous current regulator, showing how the Park and inverse Park transforms can create a current controller acting in the rotor d - q frame.

The control itself is decoupled into d -axis control and q -axis control. The goal of the d -axis controller in this case is simply to keep the d -axis current at zero, so the d -axis reference current is always zero. The q -axis controller is the torque controller, and I_{qr} is the reference current

relating to the torque demand. The output of these two controllers are d - and q -axis PWM command, which are re-mapped onto A, B, and C and sent to the motor.

The synchronous current regulator has many advantages over other current control methods. It is simple to understand, particularly in the context of torque control. It also has DC steady-state reference. Even though the rotor and three-phase currents are AC, the d -axis and q -axis references are constant values in steady-state operation. Thus, a simple proportional-integral (PI) controller can achieve zero steady-state error [1]. Other control methods that attempt to feedback-control the phase currents are chasing an AC reference, making it more difficult to achieve zero steady-state error.

One disadvantage of the synchronous current regulator is the computation involved in the Park and inverse Park transforms. Although conceptually simple, these matrix transformations involve trigonometric functions. However, processing power is no longer a limiting factor and this method of current control is easily implemented in embedded systems.

2.3 Modified Synchronous Current Regulator

The controller presented here uses a modified synchronous current regulator designed to minimize processing power requirements. It does so by separating the control effort into two loops: a fast loop for generating PWMs and a slow loop for executing current control. The modified synchronous current regulator is shown in Figure 26.

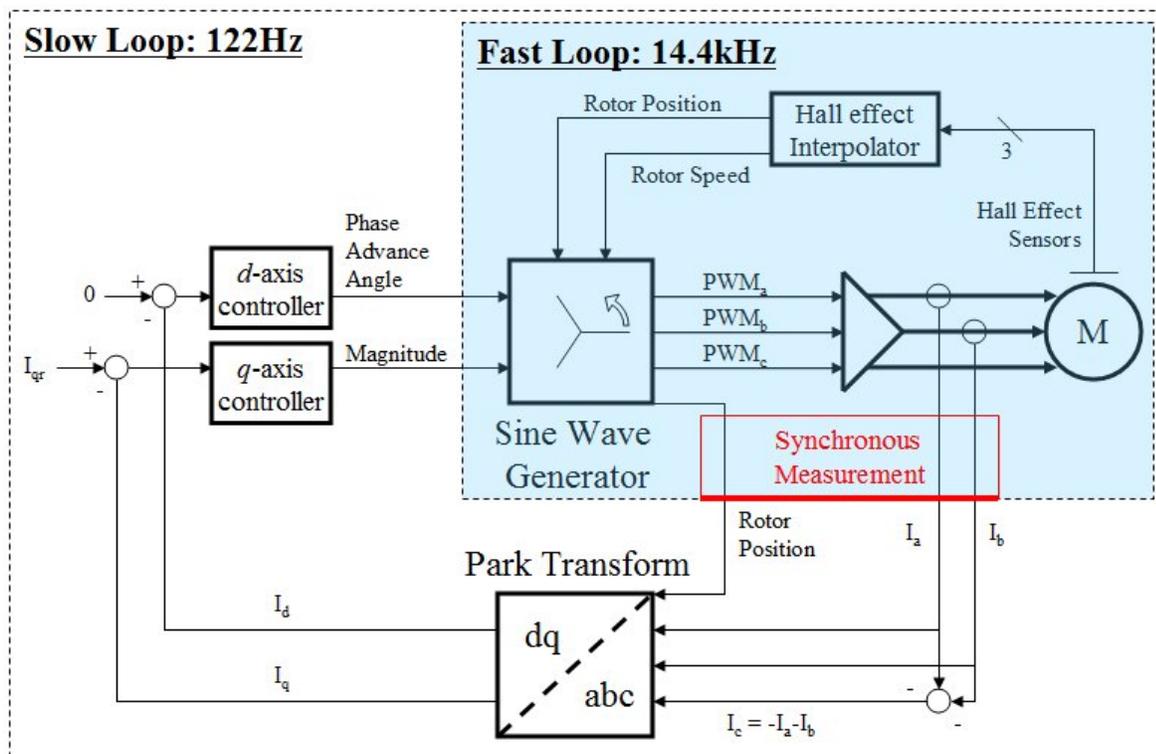


Figure 26: The modified synchronous current regulator used in this controller. It replaces the inverse Park transform with a simpler sine wave generator with phase advance.

In the fast loop, a sine wave generator replaces the inverse Park transform. The sine wave generator is based on a look-up table, and the index in the table is driven by rotor position. However, the controller is designed to work with motors that only have Hall effect sensor feedback. (These are motors that would otherwise be controlled by six-step commutation.) The Hall effect sensors produce an absolute position reference, but not with enough resolution to generate smooth sine waves. An interpolation routine fills in the gaps in information, feeding an estimated rotor speed to the sine wave generator that it uses to produce a smooth position estimate. When a new absolute position comes in, it overrides the estimate.

The fast loop executes at 14.4kHz, the PWM frequency. Thus, it produces new commands every time the PWM is reset. This is the maximum possible update rate and resolution for sine wave generation; anything higher would not be translated into commands. The sine wave generator can produce sine waves with a time resolution of about 69 μ s. It involves no floating point calculation.

The slow loop runs at 122Hz, much slower than the PWM frequency, but still faster than the mechanical time constants of the system. It contains the Park transform, as well as the d - and q -axis control loops. These are all floating-point operations, which are processor-intensive. However, the slow loop has much longer (8.2ms) to execute.

Information is passed from the fast loop to the slow loop to make a synchronous current measurement. When the slow loop measures the phase A and B currents, it locks in the instantaneous rotor position estimate and uses this value for the Park transformation. Even if the Park transformation takes many cycles to complete and the rotor moves during the processing, it is using the locked-in value of rotor position. This ensures that the current measurement is truly synchronous.

The inputs to the modified synchronous current regulator are unchanged: d -axis reference current is always zero and q -axis reference current is the torque command. Similarly, the two axes have their own independent control blocks. The fundamental difference is the output of these control blocks. Whereas before, a d -axis and q -axis command were established, the outputs now are a magnitude and phase angle for the sine wave generator. The magnitude is a simple scaling operation applied equally to all three PWM outputs. The phase is just a shift in the look-up table. These are processor-friendly operations that replace the inverse Park transformation.

It's important to note that this modification does actually change the functionality of the synchronous current regulator. Now, torque control is executed solely by adjusting the magnitude of the sine waves while phase control is handled independently. Assuming the loops are stable, the steady-state operating point is the same. This is entirely driven by the references. What will be different is the transient response; how the commands adjust themselves.

To illustrate the difference (and similarity) between the standard and modified synchronous current regulator, a simple simulation runs each on a typical motor. The conditions of the simulation are listed in Table 2. Additionally, the motor speed is held constant and there is no processing delay. A step input in the torque command, I_{qr} , to 20A is given, and the step response of the system is plotted.

Table 2: Simulation parameters for the simple fixed-speed step input analysis of the standard and modified synchronous current regulators.

Parameter	Value	Notes
Phase Resistance	0.167Ω	rear scooter motor
Synchronous Reactance	0.100Ω	similar to rear scooter motor at 500RPM
<i>q</i> -axis Gain	0.7863V/(As)	Volts per Amp second (error integrator)
<i>d</i> -axis Gain	0.7863V/(As)	Volts per Amp second (error integrator) synchronous current regulator
<i>d</i> -axis Gain	8.578°/(As)	degrees per amp second (error integrator) <i>modified</i> synchronous current regulator

First, the simulation is run with fixed timing, voltage placed directly on the *q*-axis as in Figure 23. The *q*-axis controller is a simple proportional control followed by an integrator. (It is not a P.I. controller, though.) The overall gain of the controller is 0.7863 Volts per amp-second. As expected the current lags into the *d*-axis:

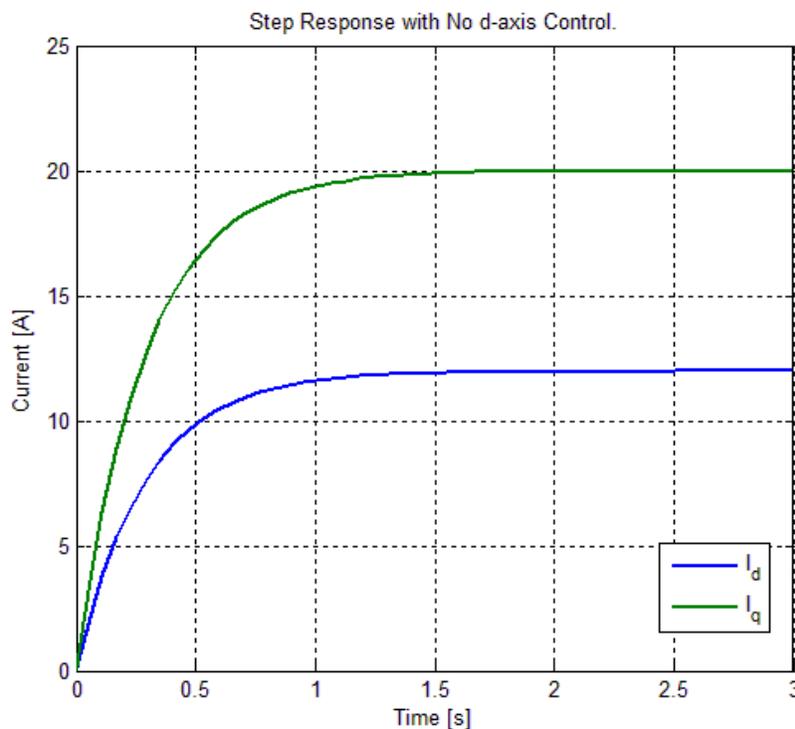


Figure 27: In the step response with no *d*-axis control, *q*-axis current still reaches 20A but *d*-axis current is non-zero. The total current is 23.3A.

Next, the *d*-axis controller is added in. In the case of the standard synchronous current regulator, it is identical to the *q*-axis controller and gets the same gain. In the case of the synchronous current regulator, the *q*-axis controller sets the magnitude and the *d*-axis controller sets the phase of the sine wave generator, as in Figure 26. The *q*-axis gain is left at 0.7863 Volts per Amp-second. The *d*-axis controller, though, now has units of degrees per Amp-second, and there is no

immediately obvious basis for setting this value. For this simulation, 8.578 degrees per Amp-second is used. The phase advance angle is limited to 90°, so a d -axis current of 10A would cause the controller to traverse the full range in just over one second. Figure 28 shows the step response of both the standard and the modified synchronous current regulator under.

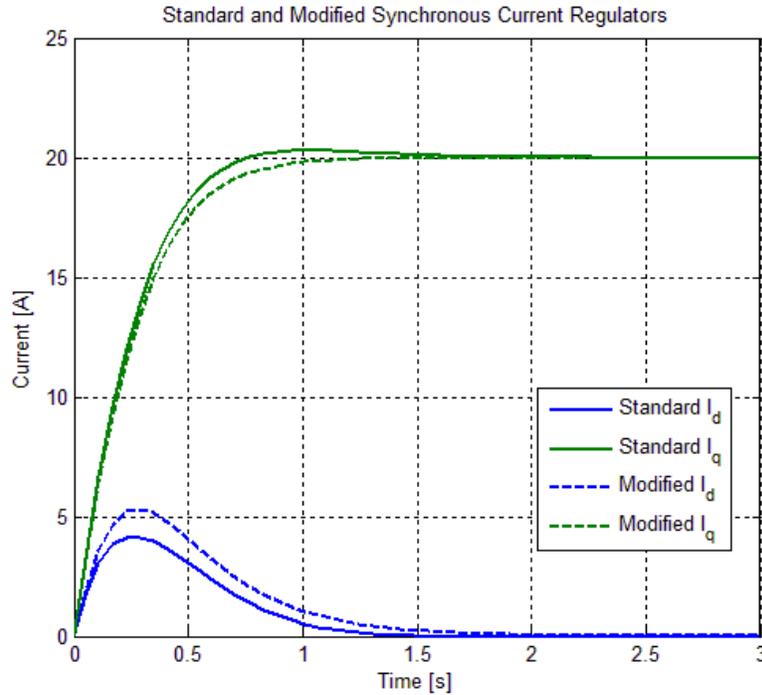


Figure 28: Comparison of the standard and synchronous current regulator step response to a 20A torque command with fixed motor speed. The total steady-state current is 20A in each case.

Both versions of the synchronous current regulator are able to drive the d -axis current to zero by advancing the phase of the voltage. The steady-state current in each case is 20A, exactly on the q -axis. The only difference is the transient response, and with these gains the difference is minor. The modified regulator shows a slightly longer settling time and more d -axis transient current, but less q -axis overshoot. However, choosing different gains for either regulator can change the response significantly. Most importantly, both the standard and the modified synchronous current regulators can both achieve the desired result of keeping current on the q -axis.

3 Testing the Controller

The controller has been designed for the B.W.D Scooter, a prototype electric kick scooter with custom hub motors in both wheels. The scooter, shown in Figure 29, was built as part of the Edgerton Center Summer Engineering Workshop in 2009. The two 500W outrunner motors have trapezoidal back EMF and adjustable-timing Hall effect sensors, so they are designed for BLDC commutation. For a detailed analysis of the differences between trapezoidal and sinusoidal motors, and how field-oriented control applies to each, see Section 4 and [1].

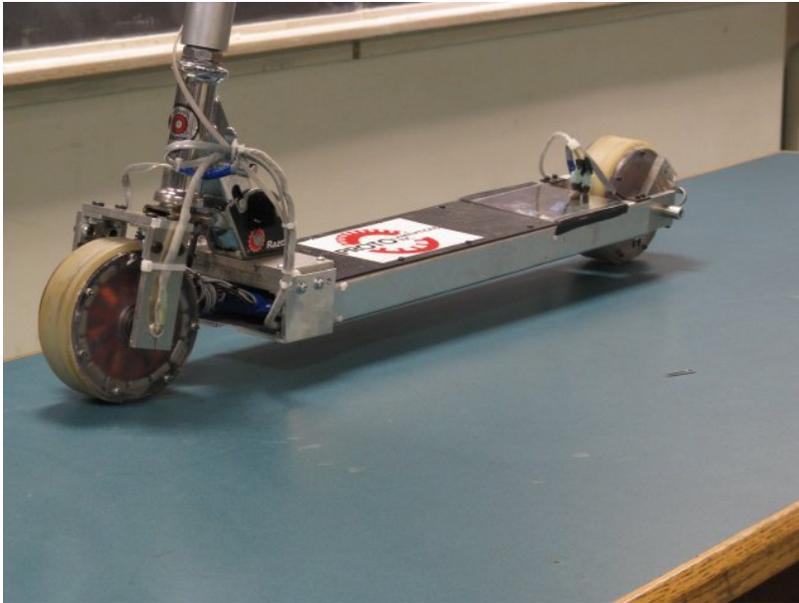


Figure 29: The B.W.D. Scooter, test vehicle for this controller, as two integrated 500W hub motors in its wheels.

The two motors tested differ only in the number of turns per phase of their windings. The front motor has 60 turns per phase, while the rear motor has 90 turns per phase. This gives the rear motor a higher torque, but lower top speed. In all other ways they are identical. Appendix 5.4 gives more detailed specifications for each motor. The test power supply is a 33V, 4.4Ah lithium iron phosphate battery pack.

While all three controllers (v1.0, v2.0, and v2.1) were tested using the B.W.D. Scooter, only v2.1 implements field-oriented control as described in Section 2. The others use standard BLDC six-step commutation as described in Section 4.1. Six-step commutation is simple to implement and test, and is commonly used with trapezoidal motors such as these to good effect. However, the purpose of the v2.1 controller is to test the modified synchronous current regulator described in Section 2.3, and that will be the focus of this section.

3.1 Establishing a Baseline: q -Axis Control Only

Before testing the field-oriented control scheme, a baseline operating point was established for comparison. Comparing the sinusoidally-commutated field oriented controller to standard six-step BLDC control is not easy or particularly useful, since the operating characteristics are so different (see Section 4). Therefore, the baseline used is a sinusoidally-commutated control scheme with no opportunity for phase advance. It may still, however, measure and control q -axis

current to maintain torque control. This baseline control scheme is depicted in Figure 30, a slight modification of Figure 26.

In this controller, the d-axis controller is eliminated and the phase advance angle is fixed at zero. This means the commutation, though still sinusoidal, is fixed to the Hall effect sensors. The exact timing is set by rotating the Hall effect sensors at no load until the motor is spinning at its slowest stable operating speed. This corresponds to a voltage peak on the q -axis, as in Figure 23. And as in Figure 23, current lag is expected when the motor is loaded at speed. Figure 31 shows the results of this baseline test.

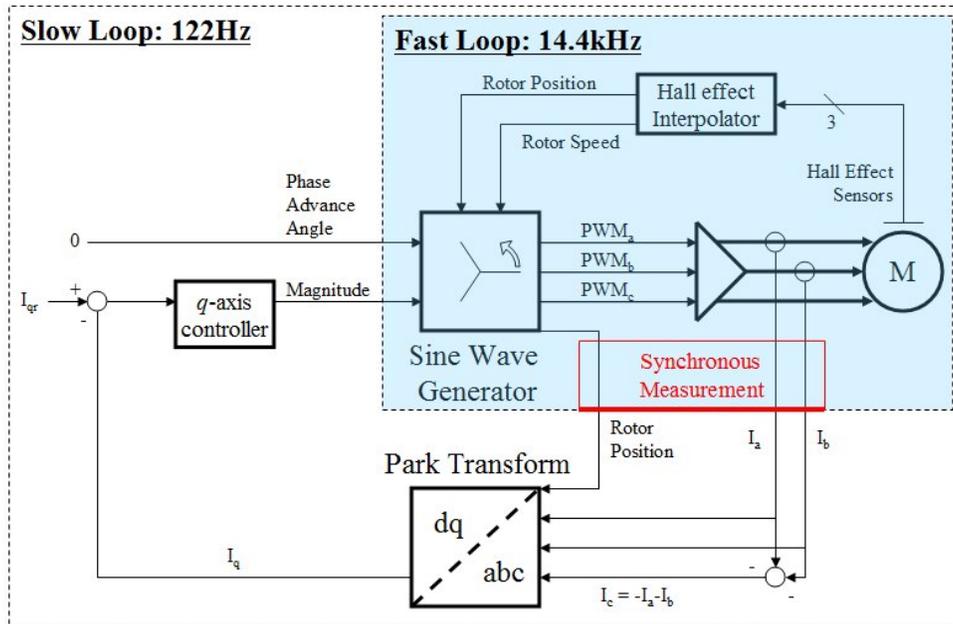


Figure 30: A baseline controller with q-axis control only. The phase advance angle is always zero, i.e. the coil timing is fixed to the Hall effect sensors.

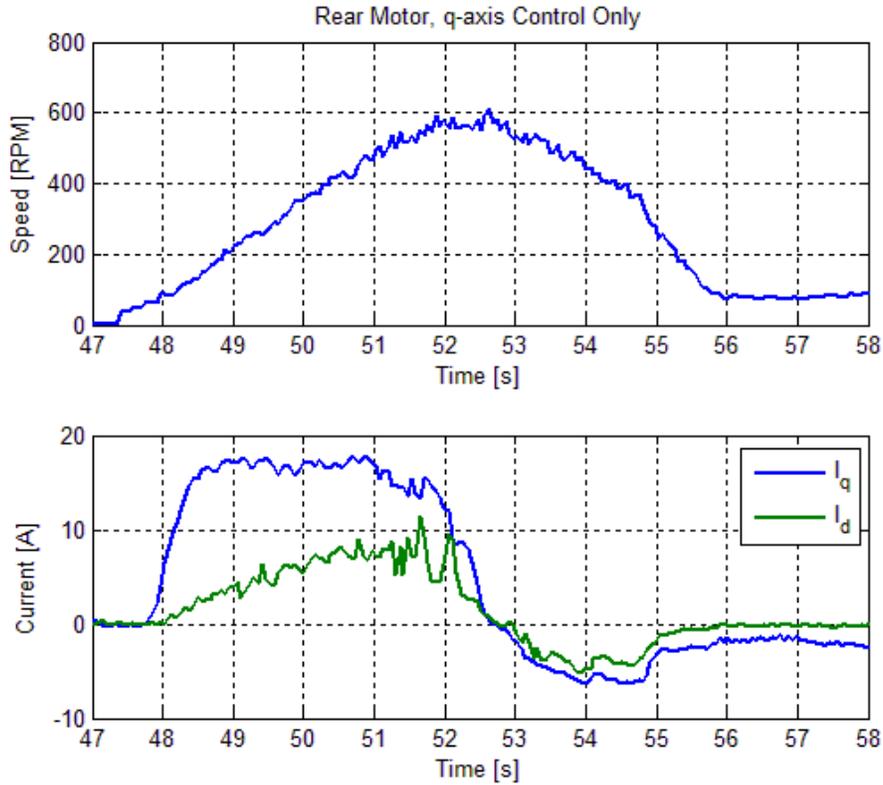


Figure 31: Results of the baseline test with q -axis control only. As expected, some current lags behind onto the d -axis.

As predicted, current begins to lag behind voltage and onto the d -axis. While the q -axis controller still does its job, maintaining the proper q -axis current for the requested torque, the magnitude of the *total* current vector is increased due to the d -axis component. The total current magnitude is what determines dissipation in the winding resistance, so the presence of d -axis current yields more dissipation for the same amount of torque, or conversely less torque for the same amount of dissipation. In other words, the motor efficiency is lower.

In this baseline test, there is a region of constant q -axis current (17A) from 49 to 51 seconds. During this time, the speed increases from 220 to 500rpm (280rpm increase). The d -axis current increases from 4A to 7A and the total current magnitude increases from 17.5A to 18.4A. As speed increases further, the ratio of d -axis to q -axis current increases. This ratio is the tangent of the angle by which current lags the q -axis.

3.2 Modified Synchronous Current Regulator: One Motor

Now for comparison, the same motor and load are controlled using the modified synchronous current regulator described in Section 2.3 and shown in Figure 23. The result is shown in Figure 32.

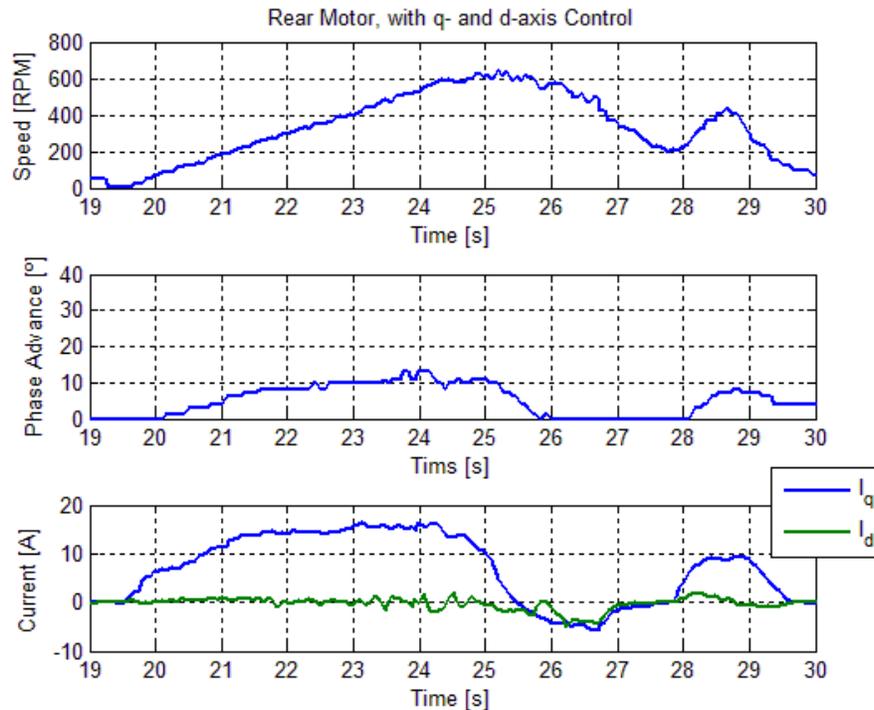


Figure 32: The same motor and load as the baseline, now controlled with the modified synchronous current regulator.

Immediately, the difference is clear. With the d -axis controller running, d -axis current is held near zero during the entire course of acceleration. This is accomplished by advancing the phase of the voltage, as in Figure 24, to accommodate for current lag. The exact amount of phase advance is controlled in real time to keep the d -axis current at zero. In this case, it varies from 0° to 13° .

Though the torque command is different, there is a region of fairly constant q -axis current in this test as well between 22 and 24 seconds. The average current is about 15A in this window. The speed increases from 300rpm to 540rpm (an increase of 240rpm) during this period of time. The acceleration difference is proportional to the difference in q -axis current (15A vs. 17A), however in this case there is no d -axis current. As a result, the magnitude of the *total* current vector is the same as the q -axis current. Only torque-producing current contributes to dissipation. This is clearly a more efficient operating point.

3.3 Modified Synchronous Current Regulator: Two Motors

The controller can execute field-oriented control on both motors simultaneously. It does not contribute much new information to the theoretical discussion, since the motors are independent. Therefore, a simple confirmation that the d -axis current controller works is shown in Figure 33.

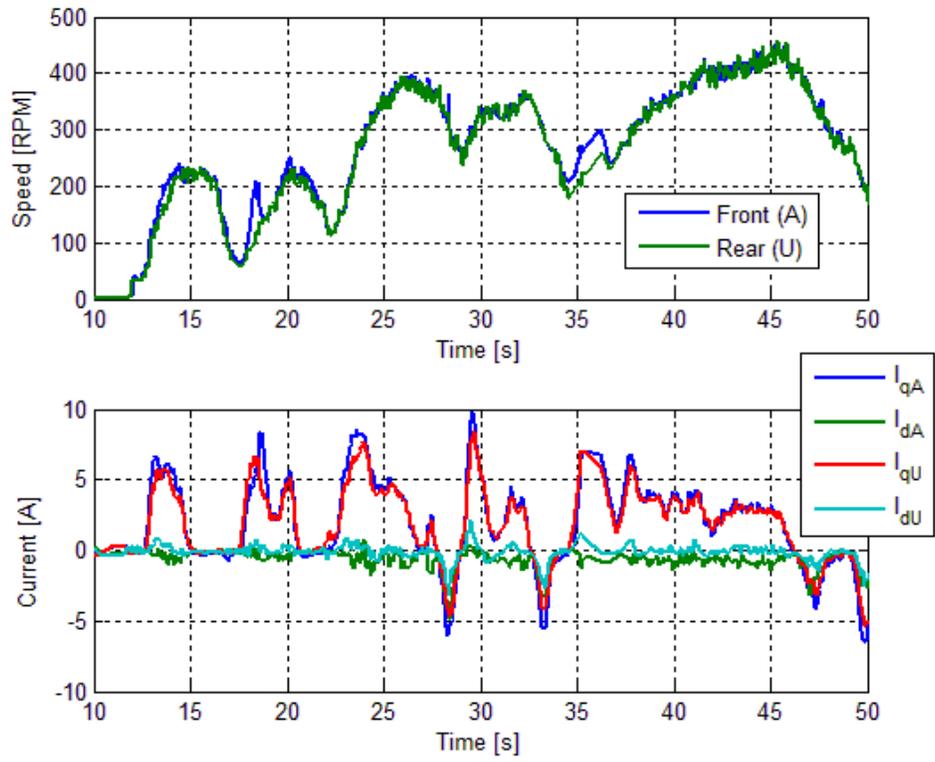


Figure 33: Field-oriented control of two motors simultaneously. d -axis current stays near zero at all speeds.

As expected, d -axis current is held near zero over the entire range of tested speeds and current loads.

4 Brushless DC vs. Brushless AC

This section tackles the differences (and similarities) between brushless DC motors and brushless AC motors, most often called Permanent Magnet AC (PMAc) motors or Permanent Magnet Synchronous Motors (PMSM). The analysis itself is not complicated, but sorting out a consistent and unambiguous definition of the problem is where the challenge lies. This is in large part due to the fact that both the motor and the control strategy are involved in the definition. The motor, specifically the shape of its back EMF waveform (trapezoidal or sinusoidal), is only part of the equation and must be matched with a drive strategy to form a complete definition of the problem. One of the most thorough approaches to this challenge is contained in the S.M. thesis of James Mevey [1]. To directly quote Mevey:

It is the author's opinion that the difference between trap and sine [brushless motors] is surrounded by more misunderstanding and confusion than any other subject in the field of brushless motor control.

To start tackling the problem, brushless motors themselves can be broken into two different types: sinusoidal and trapezoidal. These, it should be noted, are just the extremes of a large spectrum of possible real motors. However, these two extremes will be used to bound the analysis. The classification is based on the shape of the back EMF waveform of the motor, the voltage it produces at its terminals as a function of rotor position *with no load*. The back EMF is also a function of angular velocity. Specifically, the period of the back EMF waveform is one electrical period of the motor. This is implicit in defining the back EMF waveform shape as a function of rotor position. What isn't yet implicit is that the amplitude of the back EMF waveform is directly proportional to the angular velocity. This comes from the rate component in the definition of voltage as the rate of change of magnetic flux through the windings. Angular velocity, therefore, only scales the waveform linearly. The normalized shape does not change based on speed.

The sinusoidal and trapezoidal back EMF waveforms for two motors *at a given speed* are shown in Figure 34. These are line-to-neutral back EMF waveforms for a wye-connected motor. Both motors are fictitious, but the amplitude of the trapezoidal back EMF waveform is chosen to match the given motor constant for the rear scooter motor if it were a pure trapezoidal motor with a pure BLDC drive (described below). The sinusoidal back EMF waveform is simply given the same peak value for easy analysis. This is not meant to imply equal power or equal potential to produce torque.

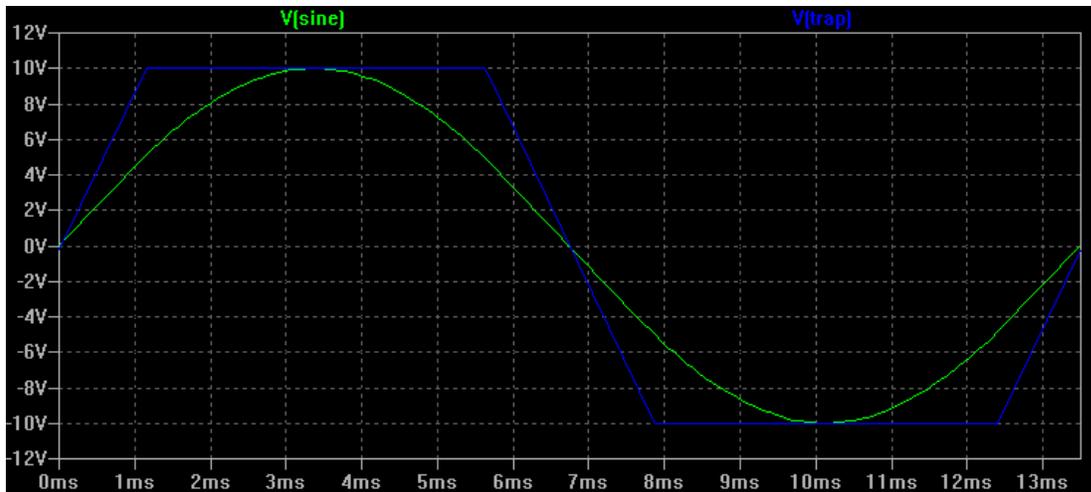


Figure 34: The back EMF waveforms of two fictitious motors, one trapezoidal and one sinusoidal. If they are 14-pole motors, these waveforms both correspond to a speed of 635RPM.

Both waveforms are shown at an electrical frequency of 74.1Hz, corresponding to a mechanical speed of 635RPM in a 14-pole motor. The trapezoidal motor's line-to-neutral back EMF has a flat top that is 120° electrical wide. The inverse motor constant, expressed as the RPM/V, is based on the line-to-line voltage peak of the trapezoidal motor. The line-to-line peak is simply twice the line-to-neutral peak, or in this case 20V. That puts the inverse motor constant (kV) at 32RPM/V, the same as the rear scooter motor. The sinusoidal motor is simply chosen to have the same line-to-neutral peak value.

The shape of the no-load back EMF waveform is entirely determined by the motor. It is a function of the geometry of the motor, including the magnets, windings, and stator and rotor steel. It is in no way influenced by the drive method. Either type of motor can be driven by any controller without influencing its back EMF waveform. Trapezoidal vs. sinusoidal is thus a completely independent classification for the motor itself, and does not by itself make a complete definition of brushless DC and brushless AC. Real motors mostly fall somewhere in between the two.

Motors that have a back EMF that looks more trapezoidal are often small and inexpensive. Concentrated windings, discrete magnets with no skew, and thinner steel that saturates are all features that make the motor both easier to make and more likely to have a trapezoidal back EMF. Motors for hobby and remote control vehicles often take this form and they are often referred to as brushless DC motors.

Motors that have a back EMF that looks more sinusoidal are often larger, more expensive, and designed for low-speed motion control. Overlapped windings, skewed magnets, and thick steel that operates well below saturation are features that make the motor more expensive to make, but give it a more sinusoidal back EMF. These motors are found most in motion control (servomotors) and vehicles. They are sometimes called permanent magnet AC (PMAC) motors or permanent magnet synchronous motors (PMSM).

The other important classification used is of the drive method. The controller, in this case controlling a three-phase inverter, can use many different methods to decide how to set the six switches at any given point. Included in these are six-step commutation and sinusoidal commutation, the two methods that will be considered in this analysis.

4.1 Six-Step Commutation

Six-step commutation is very common in controllers that use three Hall effect sensors to detect rotor position. Inexpensive sensorless controllers also use six-step commutation. It is the easiest form of brushless motor control to implement, and takes the form of a simple state machine that can be implemented in either hardware or software.

The electrical period of the motor (specifically, of the back EMF, which is fixed to the rotor position and velocity) is divided into six segments of equal length, 60° electrical each. Each of these segments corresponds to a different Hall effect sensor state, which in turn sets a different inverter switch state. The full set of states is listed in Table 3. The switch designations in this table refer to the inverter layout shown in Figure 1.

Table 3: The states in six-step commutation.

Electrical Angle	Hall Effect State	S1	S2	S3	S4	S5	S6
0°-60°	{0,0,1}	PWM	~PWM	OFF	ON	OFF	OFF
60°-120°	{0,1,1}	OFF	OFF	OFF	ON	PWM	~PWM
120°-180°	{0,1,0}	OFF	ON	OFF	OFF	PWM	~PWM
180°-240°	{1,1,0}	OFF	ON	PWM	~PWM	OFF	OFF
240°-300°	{1,0,0}	OFF	OFF	PWM	~PWM	OFF	ON
300°-360°	{1,0,1}	PWM	~PWM	OFF	OFF	OFF	ON

Table 3 contains a tremendous amount of information, maybe more than can be processed in one quick look. It is easier to see the patterns as the states progress. First, notice that as the rotor progresses in either direction through electrical angle, each individual Hall effect sensor goes through a repeating sequence: {0, 0, 0, 1, 1, 1, 0, 0, 0, 1, 1, 1, ...}. In fact, the sensor is just reporting the “polarity” of the nearest magnet, 1 or 0 could be N or S. Each N magnet and each S magnet spans half an electrical period, 180°, or three states. By offsetting the sensors by 120° electrical, this repeating sequence is also offset. This results in three Hall effect sensor values which encode six states. There are actually eight total states, but {1,1,1} and {0,0,0} are undefined.

The degrees of freedom represented by the state table are also confusing to consider. Assuming the sensor positions are fixed, there are still three sensor wires and three power wires which might be wired in any permutation to the controller. At most, this represents 3!×3! or 36 total permutations. With good color coding and documentation, this shouldn’t be a problem, but often it comes down to trial and error. Knowing how many permutations need to be tested is useful in this case.

Luckily many of these permutations are redundant due to the 120° rotational symmetry of the three-phase set. In fact, there are only 12 that would result in a unique operation, and only two that result in proper operation: one forwards and one reverse. The other states may cause oscillation, if the progression through the state table is backwards, or poor performance, if the states are shifted. Usually the correct states are obvious, once found. The trial-and-error process can be carried out either in software or physically by swapping wires.

If the sensors themselves can be moved, as is the case with adjustable-timing motors, there is one extra degree of freedom that can be used to slightly shift all the states by a few degrees. Usually, the sensors are set so that the phase being PWMed is at the peak of its back EMF waveform. Physically, the coils of this phase will be between two magnets, and it will try to attract one of the magnets into alignment with it while repelling the other. There are reasons to “advance” the timing of the sensors by rotating them slightly from this point. The benefits are similar to the benefits of using field-oriented control, although in the case of physically moving sensors, the motor is optimized ahead of time for particular operating point instead of being dynamically adjusted. For the purpose of this analysis, the sensors will be considered fixed at the point where the back EMF peak is centered in the 60° segment of the phase being PWMed.

4.2 Sinusoidal Commutation

Sinusoidal commutation is the more straightforward to understand. In this drive method, the half-bridges are used with synchronous rectification, so that at any given time three switches are closed and three are open. With the constraint that the high and low switches of any given leg can never be closed at the same time, this means that each leg has either a high or a low switch closed at any given time. The three legs, though, might operate independently of each other, out of sync or even at different frequencies. This is uncommon, but not technically impossible. Each half bridge is given a single control input: the duty cycle. In this drive method, the duty cycle is defined as the ratio of on-time of the high switch to total time. The average voltage applied to the motor terminal is the duty cycle multiplied by the total DC voltage. An example of how this might look with a 15kHz PWM frequency and a 74.1Hz electrical frequency is shown in Figure 35.

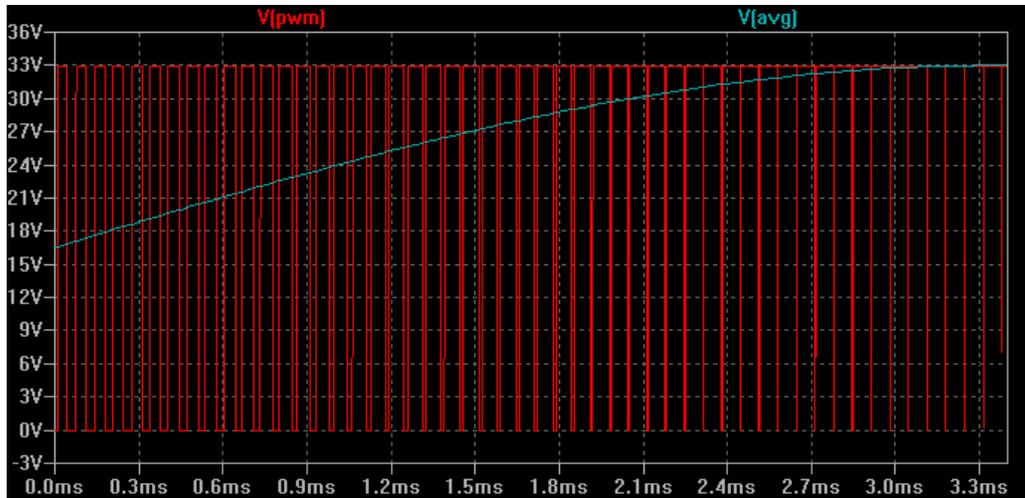


Figure 35: An example of sinusoidal PWM at 15kHz showing the actual voltage applied to the motor terminal (red) and the average voltage (blue). For clarity, only a quarter-period is shown.

At 50% duty cycle on any leg of the inverter, the average voltage applied to the motor terminal is half the DC voltage, 16.5V in this case. As the duty cycle increases, the average voltage increases. By modulating the duty cycle to be sinusoidal, the average voltage applied is sinusoidal as well. It should be clear from this that the exact frequency used does not matter, as long as it is significantly higher than the electrical frequency of the motor.

The amplitude and offset of the sine wave are controllable parameters within the limits of the DC bus voltage. In Figure TK, the average applied voltage has amplitude of 16.5V and an offset of 16.5V. With a 33V DC bus, this is the only possible offset, for reasons that should be obvious. However, imagine a sine wave with amplitude of 10V. One way to achieve this would be to vary the duty cycle sinusoidally between 19.7% and 80.3% (centered at 50%). This would produce the 10V sine wave centered at 16.5V. Another option is to vary the duty cycle between 0% and 60.6%. This would give the same amplitude, but an offset of 10V. If all three phases have the same offset, the motor is not affected. (Only the voltage difference between phases is seen by the motor.) This is useful in the practical implementation of the PWM, as discussed in Section TK.

If the half bridges are used with synchronous rectification, the duty cycle and thus the average voltage applied to each motor terminal could be controlled to follow *any* trajectory, though sinusoidal is the only one considered here. For analysis, the average voltage waveform is considered rather than the high-frequency PWM that produces it. With synchronous rectification, *the voltage on each motor terminal is determined entirely by the controller*. This is an important and significant difference between the analysis of sinusoidal and six-step commutation.

4.3 Comparison of Pure BLDC to Pure Sinusoidal Drive

In the following section, two fictitious motor/drive combinations are considered to illustrate the difference between pure BLDC control with a trapezoidal motor and pure sinusoidal control with a sinusoidal motor. These two cases will serve as benchmarks for comparison of non-ideal and mixed motor/drive combinations. The motor in each case is similar to the rear scooter motor, the properties of which are given in Table 15, with the only major difference being that motor inductance is neglected. Table 4 here explicitly defines the two cases being considered and Figure 36 shows the ideal back EMF and drive waveforms. The back EMF waveforms are identical to those shown in Figure 34.

Table 4: The conditions considered for the ideal BLDC and sinusoidal motor/drive combinations.

Pure BLDC	Pure Sinusoidal
<ul style="list-style-type: none"> • $R_a = 0.167\Omega$ • $L_a = 0\text{mH}$ • Line-to-neutral back EMF is trapezoidal with 120° flat top, peak is 10V at 635RPM. • Drive is $\pm 20\text{A}$ (constant) during the flat portions of the back EMF, zero during the sloped portions. 	<ul style="list-style-type: none"> • $R_a = 0.167\Omega$ • $L_a = 0\text{mH}$ • Line-to-neutral back EMF is sinusoidal, peak is 10V at 635RPM. • Drive is a $\pm 20\text{A}$ peak sinusoid, <i>in phase</i> with the back EMF.

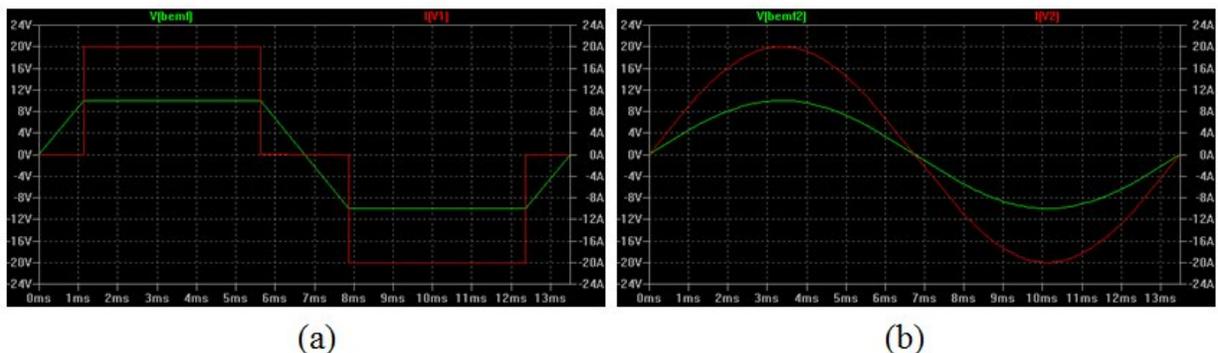


Figure 36: The back EMF and drive waveforms of the two ideal cases being considered: pure BLDC (a) and pure sinusoidal (b).

In these two cases, the effects of motor inductance are ignored. This is most troublesome in the case of pure BLDC control, where current is instantaneously turned on and off. It will be important to revisit this assumption later. For now, the ideal case can be used to make a useful comparison between these two cases.

In the case of pure BLDC, power is only being converted during the square pulses of current, and during these pulses the voltage and current are constant. This makes analysis very easy. Assuming the three phases are staggered according to the state machine presented in Table 3, there will always be one phase with positive voltage and current, one with negative voltage and current, and one with zero current and a sloping voltage. The power converted through the back EMF, then, is *constant*. Transitions are instantaneous, so that when one phase turns off another immediately turns on. If the speed of the motor is held constant by the load, this means that the

torque is also constant. This is an important and commonly misunderstood point: *in ideal BLDC drive with a trapezoidal motor, torque is constant*. The power, torque, dissipated power, and applied voltage can be found as follows:

$$\begin{aligned}
 P &= 2EI = (2)(10V)(20A) = 400W \\
 T &= \frac{P}{\Omega} = \frac{400W}{635RPM} = \frac{400W}{66.5rad/s} = 6.0Nm \\
 P_r &= 2I^2R_a = (2)(20A)^2(0.167\Omega) = 134W \\
 V_{DC} &= 2(E + IR_a) = 2[10V + (20A)(0.167\Omega)] = 26.7V
 \end{aligned}$$

where P is the mechanical power output, T is the torque output, P_r is the power dissipated in the winding resistance, and V_{DC} is the minimum DC bus voltage that would be required to create this operating point. The efficiency of the ideal BLDC drive at this operating point is 75%.

At first glance, the sinusoidal drive seems more difficult to analyze. However, the analysis is greatly aided by the following trigonometric identity:

$$\sin^2(\theta) + \sin^2\left(\theta - \frac{2\pi}{3}\right) + \sin^2\left(\theta + \frac{2\pi}{3}\right) = \frac{3}{2}.$$

This means that, as expected, torque and power are also constant with the ideal 3-phase sinusoidal drive. Applying this identity, the power, torque, and dissipated power of the ideal sinusoidal drive can also be easily calculated using the amplitudes of the current and back EMF values for E and I :

$$\begin{aligned}
 P &= \frac{3}{2}EI = \frac{3}{2}(10V)(20A) = 300W \\
 T &= \frac{P}{\Omega} = \frac{300W}{635RPM} = \frac{300W}{66.5rad/s} = 4.5Nm \\
 P_r &= \frac{3}{2}I^2R_a = \frac{3}{2}(20A)^2(0.167\Omega) = 100W
 \end{aligned}$$

In the pure sinusoidal case, the motor puts out less power and torque, but also dissipates less power in the windings resistance. The efficiency is still 75%. To achieve the same power level as the pure BLDC motor, some efficiency would have to be sacrificed due to the I^2 term. No conclusion about the relative power-handling capabilities of the two types of motor/drive is being made, though, because the back EMF waveforms here are chosen arbitrarily, not based on any real or fair comparison of two similarly-sized motors.

To calculate the minimum DC bus voltage required to create this operating point in the sinusoidal motor/drive, the line-to-line voltage amplitude is needed:

$$V_{DC} = V_{ll} = \sqrt{3}V_{ln} = \sqrt{3}V_a$$

$$V_{DC} = \sqrt{3}(E + IR_a) = \sqrt{3}[10V + (20A)(0.167\Omega)] = 23.1V$$

This result assumes that the controller has freedom to manipulate the zero-sequence voltage. That is, the “neutral” point in line-to-neutral can be moved. More commonly, though, the neutral point is fixed, for example at half the DC bus voltage. Fixing the neutral voltage would result in a minimum DC bus voltage of 26.7V, the same as in the pure BLDC case. Because this is the more common practice, and because the torque generated by the sinusoidal motor/drive at this operating point is less than the BLDC motor/drive, it is an often-stated conclusion that sinusoidal drive needs a higher bus voltage to produce the same torque at a given speed [TKref]. This needs to be weighed against many other factors that start showing up when the analysis moves beyond these two ideal cases.

Non-Ideality: Motor Inductance

The first place where the simple analysis carried out for the pure BLDC and pure sinusoidal motor/drive combinations breaks down is with the introduction of motor winding inductance. This inductance makes it practically impossible to achieve a current waveform as in Figure 36a, with its instantaneous steps. It also introduces an electrical low-pass filter when combined with the winding resistance. If no countermeasures are taken by the controller, this will cause the drive current to lag the back EMF, reducing torque and efficiency. The effect of motor inductance on the motor/drive combination is complex enough to merit simulation, particularly in the case of the BLDC drive. Using simulation and some first-principles analysis, the effect is analyzed here and compared to the ideal results described above.

The ideal BLDC drive (Figure 36) with motor inductance is impossible to achieve. The energy stored in the inductance of the phase to be turned off must go somewhere. If the current is literally forced to zero, such as by opening a mechanical switch, this would result in a high voltage that could create an arc. If a semiconductor switch with no freewheeling diodes was used, the semiconductor itself would be forced above its breakdown voltage and absorb all the excess energy in itself. If it doesn't fail immediately, repetition of this will cause heat build up which can cause it to fail within seconds.

The three-phase inverter (Figure 1) invariably has switches with freewheeling diodes. This gives current a place to go even when both switches are off. If positive current was flowing through the phase winding, current will be pulled through the low side diode from the negative DC voltage. If negative current was flowing, current will be pushed through the high side diode into the positive DC voltage. Some power is dissipated in the diode drop of approximately 1V, and this may have a significant effect on the controller heating as discussed in Appendix 5.1. For the most part, though, the magnetically stored energy is dissipated in the winding resistance and the current decays to zero.

To assess whether this decay time is even worth considering, a rough estimate can be made by using the linearized constitutive equation for an inductor:

$$\frac{\Delta I}{\Delta t} = \frac{V}{L} \Rightarrow \Delta t = (\Delta I) \left(\frac{L}{V} \right).$$

Here, the inductance, L , is the phase inductance L_a . For this example, we will use the phase inductance of the rear scooter motor, given as 0.5mH in Table 15. The voltage to use is the voltage across the inductor after the switches are turned off. If the neutral point stays at half the total applied voltage, this would be 20V for the fictional operating point being used in these analyses. This is a flimsy assumption, but the result will tell us that further investigation is necessary anyway. The change in current is 20A and the change in time is to be calculated. That gives:

$$\Delta t = (\Delta I) \left(\frac{L}{V} \right) = (20A) \left(\frac{0.5mH}{20V} \right) = 0.5ms .$$

This is a significant fraction of the commutation period of 13.5ms, so further investigation is warranted with this motor and at this operating point. (If this value was 10 μ s, for example, or if the commutation period were much longer, the results of the ideal BLDC analysis would be accurate enough to stand.) A summary of this quick analysis gives a good way of judging the relative accuracy of the ideal BLDC analysis based on the motor and the operating point. The following criteria all support the accuracy of the ideal BLDC analysis:

1. Trapezoidal line-to-neutral back EMF with a 120° flat top.
2. Low winding inductance.
3. Low current.
4. High voltage.
5. Long commutation period (a function of both motor speed and pole count).

To go any deeper into the analysis, though, simulation is necessary. A simulation carried out in SPICE can predict all the relevant operating parameters of the motor based on the back EMF waveform and drive logic. To simulate BLDC drive with inductance, a SPICE model of the inverter is constructed using (ideal) switches with freewheeling diodes. This ideal inverter is used to drive the ideal trapezoidal back EMF waveform depicted in Figure 36a according to the state machine in Table 3. Importantly, the timing of the states is fixed and defined so that the phase being PWMed is exactly at its peak, i.e. no phase advance is used yet. *The majority of BLDC controllers operate in this way.* Table 5 summarizes the simulation parameters.

Table 5: BLDC with inductance simulation parameters.

Summary of BLDC w/ Inductance Simulation Parameters
<ul style="list-style-type: none"> • Back EMF as defined in Figure 36a. This is a fictitious waveform, but it is similar to the actual waveform of the rear scooter motor operating at 635RPM. • $R_a = 0.167\Omega$, $L_a = 0.5mH$. These are the measured parameters of the rear scooter motor as defined in Table 15. • DC bus voltage of 26.7V, which would produce a steady-state current of 20A. • Ideal inverter with freewheeling diodes. (No switch on-resistance or turn on/off delay.)

- Drive logic as defined in Table 3.
- Drive states are exactly in phase with back EMF, so that 0° in Table 3 corresponds to the start of the flat peak of phase A's back EMF. No phase advance.

The simulation results, shown in Figure 37, reveal a drive current that is far from a simple low-pass filtered version of the drive current in Figure 36a. There is a significant lag induced by the inductance, shifting the current out of phase. However, the shape of the current is completely distorted. A jagged transition even occurs in the center of the on period, where before the current was constant. This is due to the diode conduction of another phase at that transition, which moves the neutral point voltage.

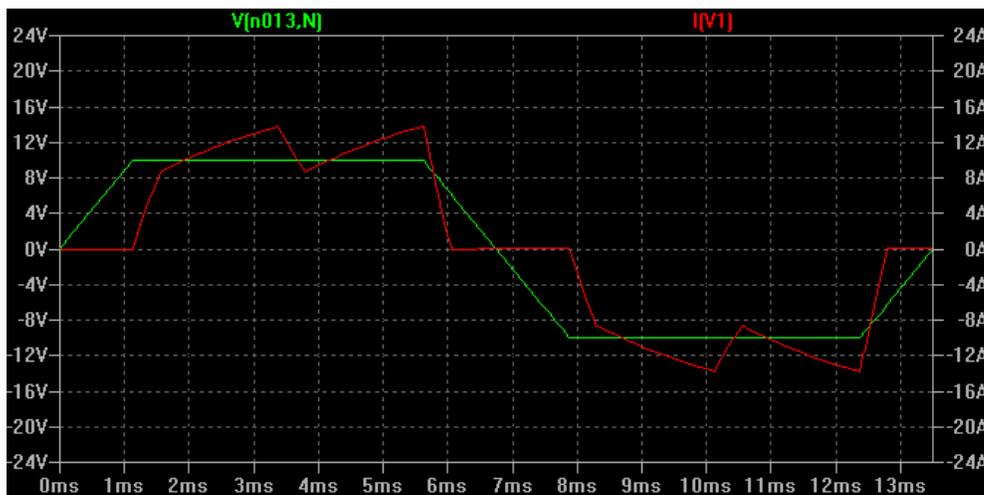


Figure 37: The drive current (red) resulting from six-step commutation with motor inductance taken into account.

The power (and torque) out of this drive is significantly lower and contains a large ripple. Figure 38 shows the instantaneous power converted through the back EMF of all three phases. The average power output is 227W. To summarize: The drive logic and DC bus voltage are exactly the same as in the ideal analysis done above. By simply adding in the inductance of the motor, the average power output is reduced from 400W to 227W and a large torque ripple is introduced at six times the commutation frequency.

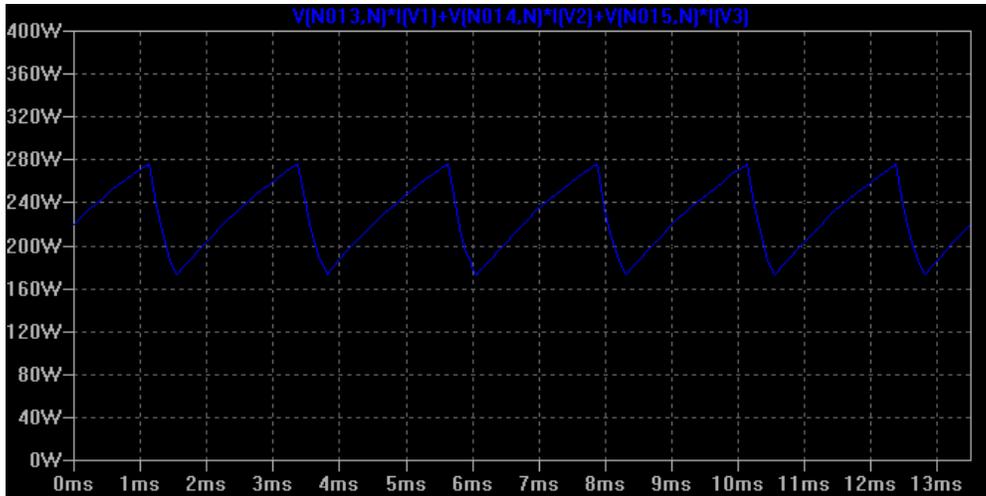


Figure 38: The instantaneous power converted through the back EMF.

To be fair, the power lost is not going to heat: the average dissipation in this case is only 43W. The motor is simply able to convert less power now at this operating point. To get more power out of it, the DC bus voltage, would have to be increased from the value of 26.7V used in the ideal case. In a real controller, this corresponds to increasing the duty cycle of the PWM. But this means that the achievable operating area in the torque/speed plane of the motor will be reduced for a given maximum DC bus voltage.

Many inexpensive BLDC motor controllers use phase advance (also called timing advance) to address the problem of reduced torque due to the effects of motor inductance [TKrefs]. Phase advance involves shifting the starting point of the commutation so that the motor current and back EMF remain as close as possible to in phase. This increases the power output (and the dissipation), but does not reduce torque ripple. To implement timing advance in increments of anything other than 60° electrical, some form of estimator must be used to predict an intermediate angle based on Hall effect sensor transitions. This same type of time-based interpolation can be used for sinusoidal control, as discussed in Section TK.

For comparison, the motor inductance is also introduced into a pure sinusoidal simulation, using the same back EMF (Figure 36b) and a sinusoidal drive voltage. In this case, the synchronous inductance is used. This is the effective inductance that each phase sees during balanced sinusoidal operation, and is only applicable to pure sinusoidal motors and drive. A detailed derivation of the synchronous inductance can be found in [1], Appendix B. For now, it is simply set to $1.5L_a$, where L_a is the same phase inductance used in the BLDC case. The phase angle of the drive signal is left at zero, i.e. in phase with the back EMF. If the motor inductance were zero, as is the case in the ideal analysis, the current and back EMF would be in phase. The applied voltage is set to be a sine wave with amplitude of 13.35V, such that if the motor inductance was zero and everything was in phase, the current amplitude would be 20A. The results of this simulation are shown in Figure 39.

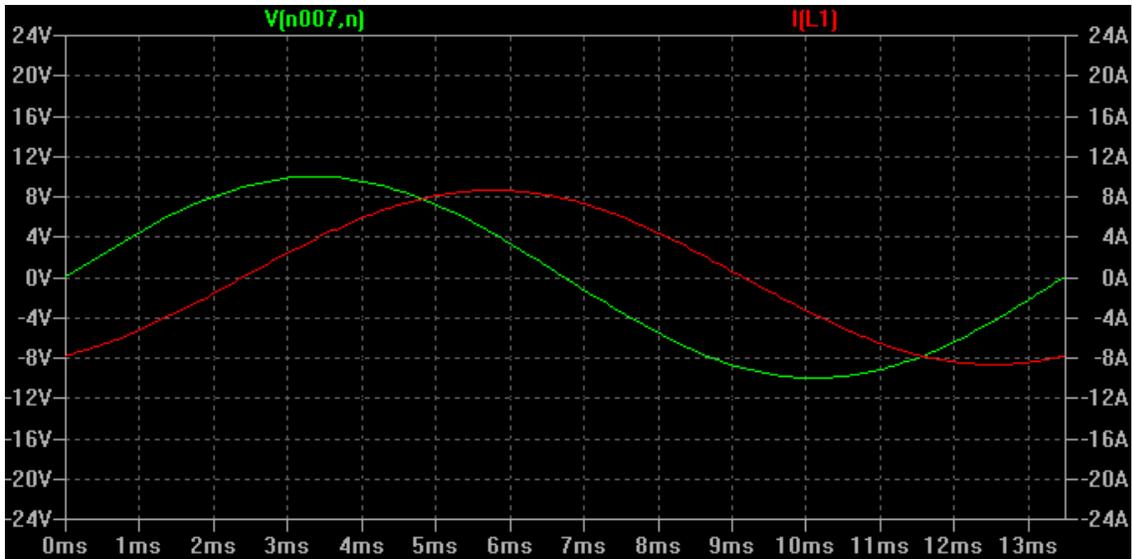


Figure 39: The drive current (red) resulting from sinusoidal commutation with motor inductance taken into account.

As expected, the current is both shifted (lagging) and attenuated, but still sinusoidal. This is one of the nice aspects of pure sinusoidal motors and drives: the waveforms pass through without changing shape or frequency. Simple trigonometry or analysis in complex quantities can be used to solve the motor's operating parameters without simulation:

$$|I| = \frac{|V - E|}{|Z|} = \frac{|V - E|}{|R_a + j\omega L_s|} = \frac{13.35V - 10V}{\sqrt{(0.167\Omega)^2 + [(466 \text{ rad/s})(1.5)(0.5mH)]^2}} = 8.65A$$

$$\angle I = \arctan\left(\frac{|j\omega L_s|}{|R_a|}\right) = \arctan\left(\frac{(466 \text{ rad/s})(1.5)(0.5mH)}{0.167\Omega}\right) = 64^\circ$$

Figure 40 shows the same information derived geometrically. This is the phasor diagram, and won't be elaborated on other than to illustrate that the current vector (shown as the voltage across the resistor) is lagging the drive and back EMF vectors, which are both in phase. The lag is a substantial 64° at this operating point.

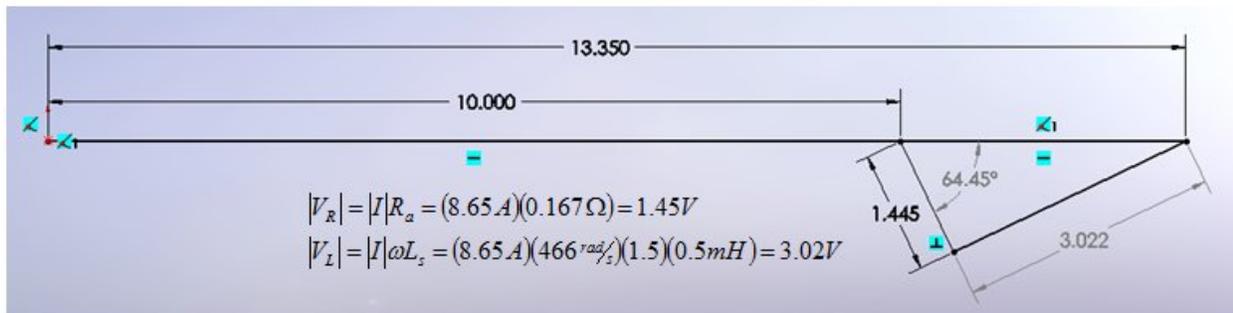


Figure 40: The phasor diagram method of deriving the drive current in the sinusoidal motor with inductance taken into account.

Likewise, the power and torque can be predicted without any simulation. Power converted is related to the dot product of the current and back EMF vectors shown above. However, since these are represented as peak amplitudes here, a factor of 0.5 is introduced to account for the integral of $\sin^2(\theta)$. In other words, the values used to calculate power are the RMS values. Also remembering that there are three phases contributing, this gives:

$$P = \frac{3}{2} |I||E| \cos(64^\circ) = \frac{3}{2} (8.65A)(10V) \cos(64^\circ) = 57W.$$

Importantly, this is a constant power; there is no ripple even with the inductance and associated phase shift. This is confirmed by simulation as shown in Figure 41.

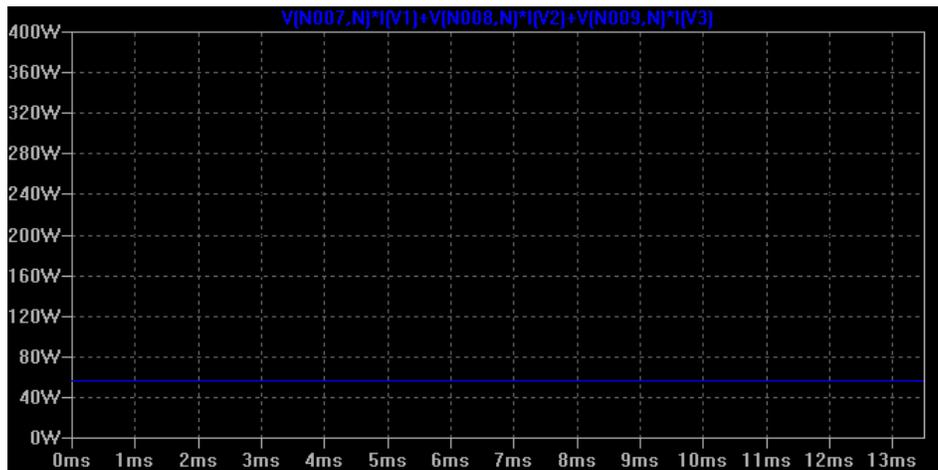


Figure 41: The power output of the pure sinusoidal motor/drive, with inductance factored in, is much lower than the ideal case but still constant.

Again, the power difference between this and the ideal sinusoidal drive with zero inductance (which was 300W) is not lost as heat. The dissipation here is a constant 19W, easily found by $3I^2R_a$ with the RMS value for current I . With the same drive voltage and timing, the motor is simply not able to convert as much power because of the phase shift created by the motor inductance. In this case, the affect of the inductance on the sinusoidal motor is even greater than it is on the pure BLDC motor: power is reduced by 81% compared to the ideal case with zero inductance!

One way to increase the power output here is simply to increase the drive voltage amplitude without changing the phase angle. This is justified since the motor dissipation is also greatly reduced. As a baseline, at this operating point the dissipation in the pure sinusoidal case was 100W and the power output was 300W, for an efficiency of 75%. With inductance, what would happen if the drive voltage amplitude were increased until the dissipation reaches 100W? That's easy to predict, since it does not affect the phase angle. The current amplitude would return to 20A, and the power would be:

$$P = \frac{3}{2} |I||E| \cos(64^\circ) = \frac{3}{2} (20A)(10V) \cos(64^\circ) = 132W.$$

This is still far short of the 300W achieved if the motor inductance was zero. With 100W dissipation, the efficiency is now a mere 57%. Additionally, the drive voltage amplitude is now much larger. It can be found by adding the three vector components of voltage. As a shortcut, the similarity of the triangle formed by the three vectors in Figure 40 can be used, since the phase angle does not change. This gives a simple proportion:

$$V = 10V + (3.35V) \frac{20A}{8.65A} = 17.75V.$$

This result can also be confirmed by simulation or geometrically. The DC bus voltage required to produce this drive voltage amplitude is 35.5V. If this is higher than the maximum DC voltage, usually the battery voltage of the system, then it could be possible that this operating point can no longer be reached. (In the case of the scooter, the battery is 33V.)

Another way to increase the power converted is to use phase advance to shift the current and back EMF back in phase. This can be done without increasing the drive voltage amplitude. This is similar to phase advance in the BLDC case, but easier to analyze due to the preservation of the sinusoidal waveform. A geometric approach is taken here in Figure 42.

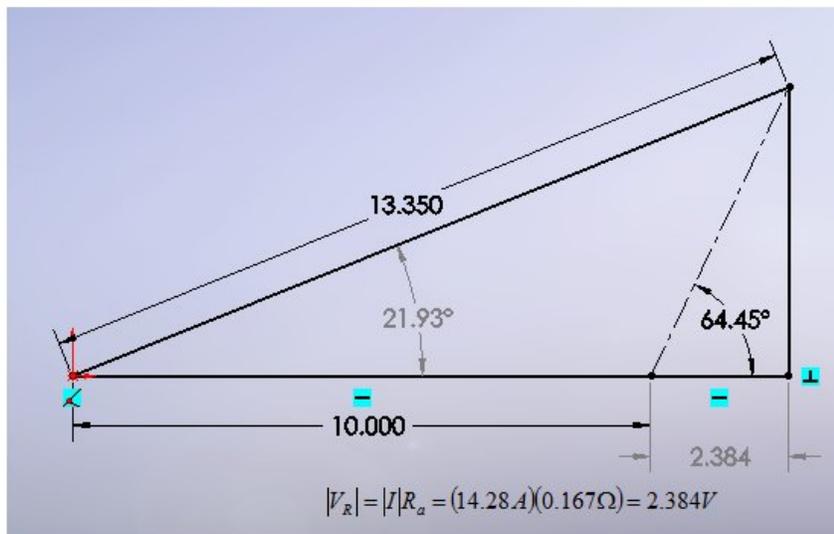


Figure 42: A geometric solution to the optimum phase advance angle for this drive voltage amplitude. Black dimensions are constraints, gray dimensions are driven.

Here, a geometric solver is used to swing the drive voltage (13.35V) out ahead of the back EMF (10V) by some unknown angle. The resistive and reactive impedances are constrained to the same ratio as before by holding the angle of the right triangle they form constant. The result is a phase advance angle of 22° to bring the current and back EMF in phase. The resulting current is 14.28A, as calculated from the voltage drop across the resistor, which is what is actually being solved. This would give a total power of 214W and a dissipation of 51W. Adding a 22° phase advance to the simulation confirms this geometric shortcut. The simulated current with this phase advance is shown in Figure 43.

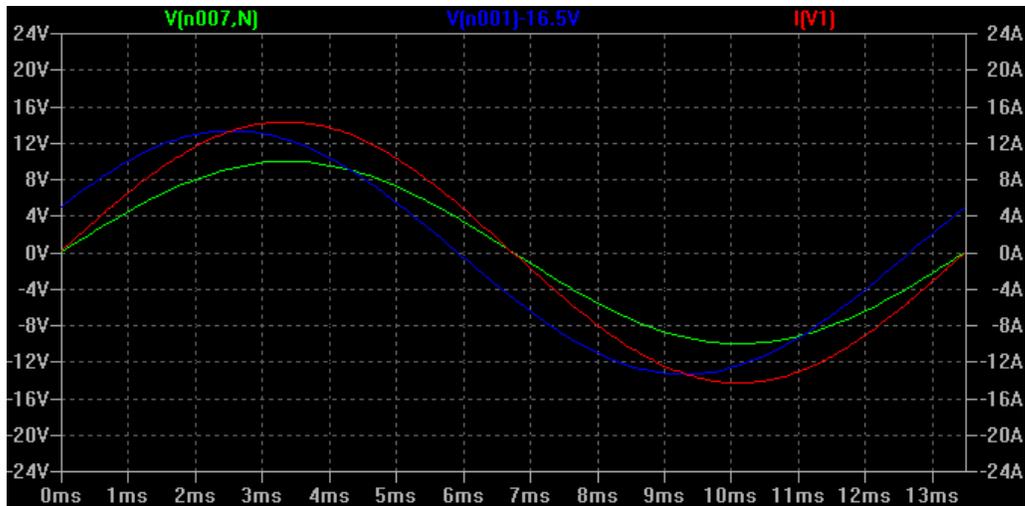


Figure 43: The 22° phase advance brings current (red) and back EMF (green) back in phase, though the drive current is still attenuated and power is lower than the ideal case with zero inductance. Here the drive voltage is also shown (blue) to illustrate the phase advance.

Phase advance with no increase in drive voltage amplitude nets more power for less dissipation than increasing the drive voltage amplitude with no phase advance. It also does not increase the minimum DC bus voltage required. It should, therefore, be the first choice method of getting the most power out. Even with an ideal phase advance, though, the power is still lower than it would be with zero inductance. Thus, a combination of phase advance and increasing the drive voltage amplitude might be required to meet power demand. How to effectively use phase advance and drive voltage amplitude to control power (torque) is a main topic of this report and is discussed in Section 0.

4.4 Mixed Motor and Drive: Sinusoidal Drive of a Trapezoidal Motor

So far, only BLDC drive with a trapezoidal motor and sinusoidal drive with a sinusoidal motor have been considered in this analysis. In this section, the consequences of mixing sinusoidal drive with a trapezoidal motor are explored. This fits into the goal of applying advanced control techniques to low-cost motor system, which tend to be trapezoidal. To be effective, these techniques should prove to be significantly better than simple six-step BLDC control with fixed timing.

Even so-called trapezoidal motors are not pure trapezoids with 120° flat tops. They may have rounded corners, uneven slopes, or other harmonics resulting from their geometry and windings. Any factions that “round off” the trapezoid, bringing it closer to a sinusoid, may improve the motor’s operation with sinusoidal drive. No real conclusion can be made about the differences except on a case-by-case basis. But as an interesting extreme case, what would happen if a pure sinusoidal drive was used on a pure trapezoidal motor?

One way to approach the problem is with Fourier series analysis. The trapezoidal back EMF can be broken down into an infinite series of pure sinusoidal components:

$$E = A \frac{4}{\pi} \sum_{n=odd} a_n \cos(n\theta).$$

Here cosines are used instead of sines, but this is just a matter of shifting. Also, only odd components are required to create the back EMF, since it has symmetry. The overall magnitude is adjusted through the scaling factor A . For the 120° trapezoidal back EMF shown in Figure 36a, with a peak of 10V, the scaling factor and first four non-zero coefficients of the Fourier series are given in Table 6.

Table 6: The scaling factor and first four non-zero coefficients of the Fourier series expansion of a 120° trapezoid with a peak of 10V.

Coefficient	Exact Value	Decimal Value
A	10V	10.0000V
a_1	$3/(\pi)$	0.9549
a_3	$-2/(3\pi)$	-0.2122
a_5	$3/(25\pi)$	0.0382
a_7	$-2/(27\pi)$	-0.0236

One way to approach the problem of applying sinusoidal drive to this trapezoidal back EMF is to look at only the fundamental. In this case, the fundamental component is:

$$E_1 = A \frac{4}{\pi} a_1 \cos(\theta) = 10V \frac{4}{\pi} \frac{3}{\pi} \cos(\theta) = 12.16V \cos(\theta).$$

This overall fundamental amplitude of 12.16V shows why the two back EMFs being compared in Figure 36 don’t necessarily represent a fair comparison. The fundamental component of the trapezoidal back EMF with a 10V peak is larger than the sinusoidal back

EMF with an amplitude of 10V. To some degree, this might be interpreted as showing that the trapezoidal back EMF with the same peak has inherently more capability to convert power. For this reason, a direct application of the 13.35V amplitude sinusoidal drive used in the pure sinusoidal simulation to the pure trapezoidal back EMF would be an unfair comparison and would result in low power conversion. Instead, the fundamental component of the trapezoidal back EMF will be used for analysis. To further illustrate this, the 10V trapezoid, 10V sine wave, and 12.16V sine wave are shown on the same scale in Figure 44. Clearly, the 12.16V sine wave is a better overall approximation of the trapezoid than the 10V sine wave.

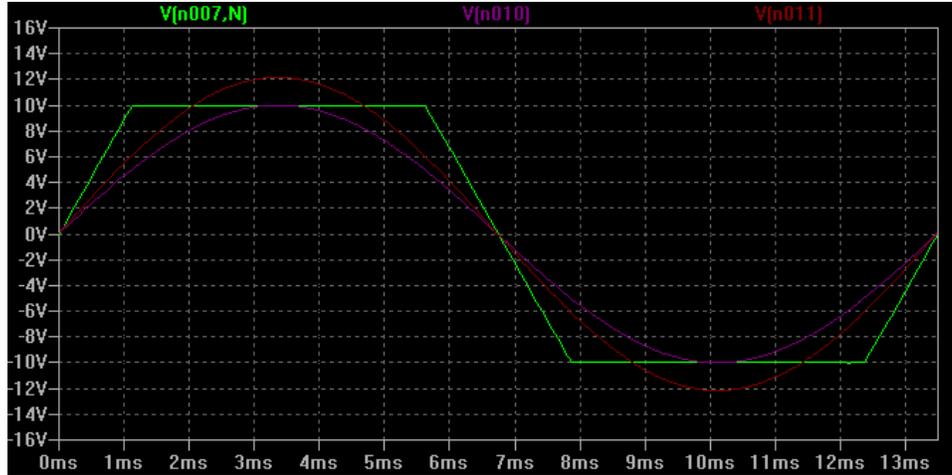


Figure 44: The 10V trapezoid is better approximated by the 12.16V sine wave, its fundamental component.

Using the fundamental component of the trapezoid as a basis for analysis, the drive voltage magnitude required to produce a 20A current, neglecting inductance, is simply:

$$|V| = |E_1| + |I|R_a = 12.16V + (20A)(0.167\Omega) = 15.5V .$$

And the power converted by the fundamental component of the back EMF would be:

$$P = \frac{3}{2}|I||E_1| = \frac{3}{2}(20A)(12.16V) = 365W .$$

The power converted by the higher-order harmonic components of the back EMF can also be calculated. They can be positive or negative and will contribute their frequency content to the overall power output, creating ripple. Rather than carry this analysis out term-by-term, the result is simulated by applying a 15.5V sinusoidal drive to the trapezoidal back EMF *with inductance set to zero*. The power output and power dissipation are shown in Figure 45. Average power output is 361W and average power dissipation is 103W, giving an efficiency of 78%. There is considerable ripple, though, at six times the commutation frequency.

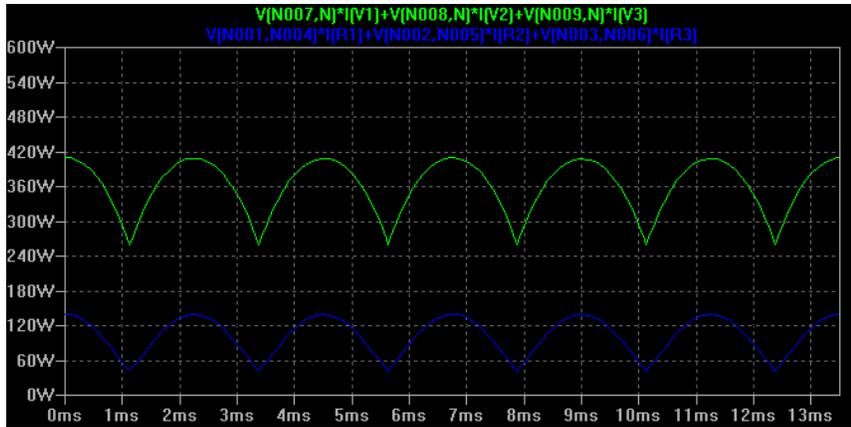


Figure 45: The simulated power output (green) and power dissipated (blue) of 15.5V sinusoidal drive applied to 10V trapezoidal back EMF *with no inductance*.

The most interesting case, though, is with inductance also factored in. Using the 15.5V drive voltage amplitude and keeping it in phase with the back EMF fundamental, the phase lag and drive current attenuation show up again. In this simulation, the phase inductance, L_a , is used rather than the synchronous inductance. They differ only by a factor of 1.5, but in this case the motor does not have sinusoidal-distributed windings and the mutual inductance of the windings is small, so the usual definition of synchronous inductance [1] does not apply. The result of this simulation is shown in Figure 46.

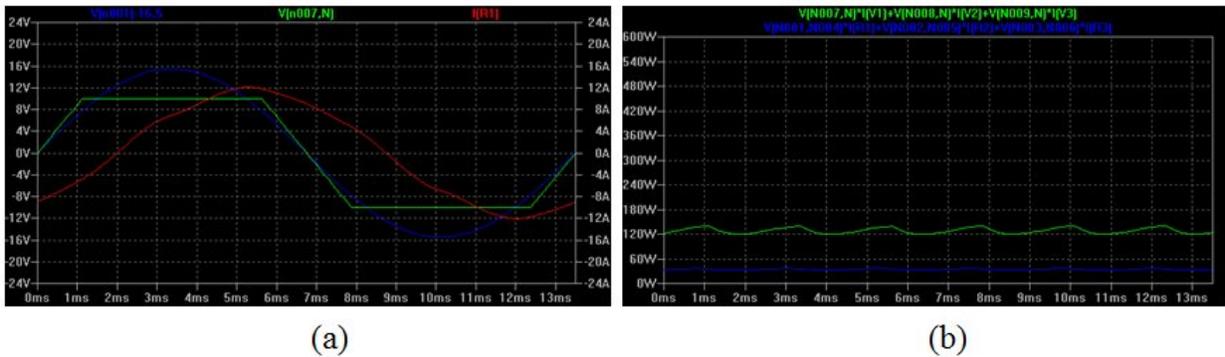
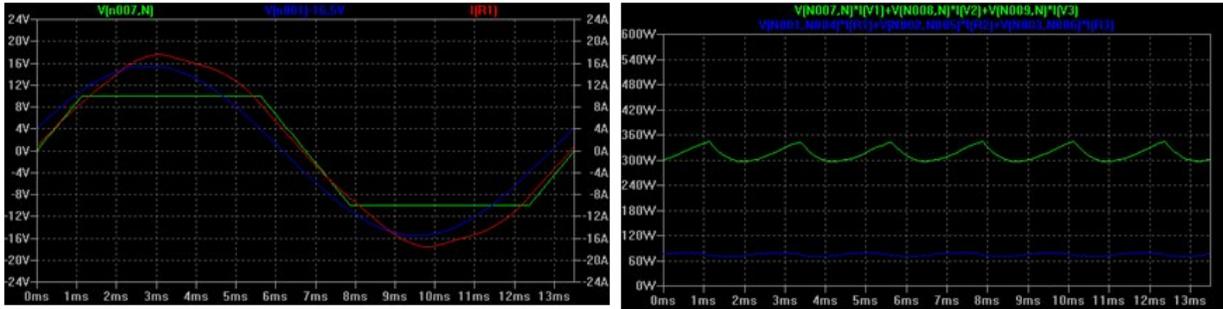


Figure 46: The simulated response of a trapezoidal motor to sinusoidal drive with no phase advance. The drive current (red) lags as expected and the power output (green) is much lower than the ideal zero-inductance case.

The average power output is only 128W, but the ripple is relatively small. Not surprisingly, the affect of the motor inductance is somewhere between that of pure BLDC and that of pure sinusoidal. That is, the attenuation is larger than that of pure BLDC, but the ripple is smaller.

The last drive method to make before summarizing everything is sinusoidal drive of a trapezoidal motor *with phase advance*. To estimate the optimum amount of phase advance, the geometric derivation of Figure 42 is used with the new values of drive voltage amplitude and motor inductance. Also, the 12.16V fundamental of the back EMF is used instead of the 10V sinusoid. The predicted result is an optimum phase advance of 15° and a resulting current of 16.9A. Figure 47 shows the results of the same simulation with this 15° phase advance. The average power output is now 315W, but with considerably more ripple.



(a)

(b)

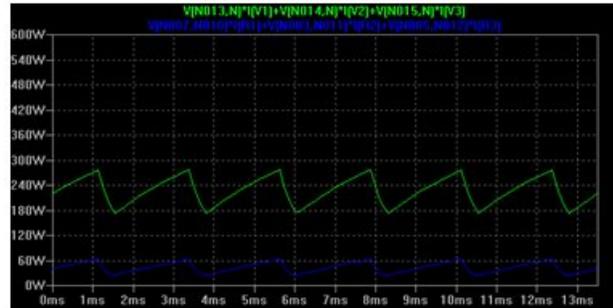
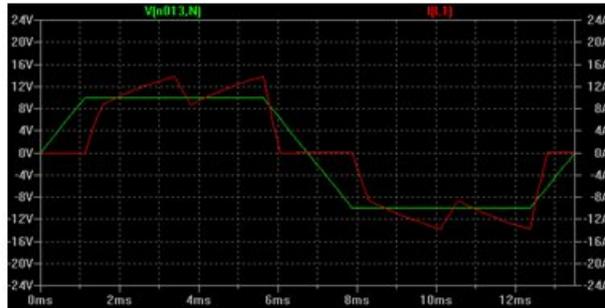
Figure 47: The simulated response of a trapezoidal motor to sinusoidal drive with 15° phase advance. The drive current (red) is now in phase with the back EMF power output (green) is higher, but has more ripple.

The concept of phase advance still works with a trapezoidal motor. The current waveform, although not quite a perfect sinusoid, can be shifted back into phase with the back EMF for more power. In fact, this can be more effective than phase advancing the BLDC drive. To summarize the comparison, the four simulated results are compared in Table 7. All four cases drive the same trapezoidal back EMF and assume the same values for motor inductance and resistance.

Table 7: A summary of the four cases considered for driving this trapezoidal motor.

Case 1: BLDC Drive, No Phase Advance

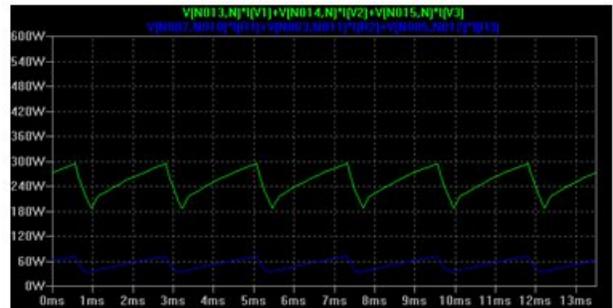
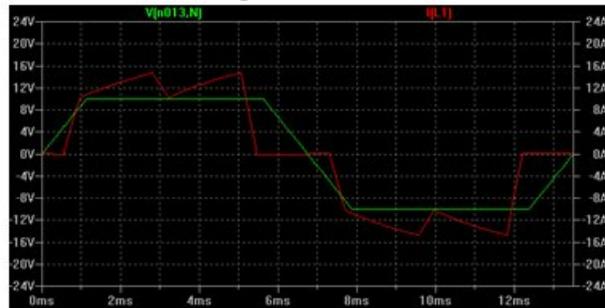
- Drive: Six-step as defined in Table 3.
- DC Bus Voltage: 26.7V



- Average Power Out: 227W | Power Ripple: 102W | Power Dissipated: 43W

Case 2: BLDC Drive, 15° Phase Advance

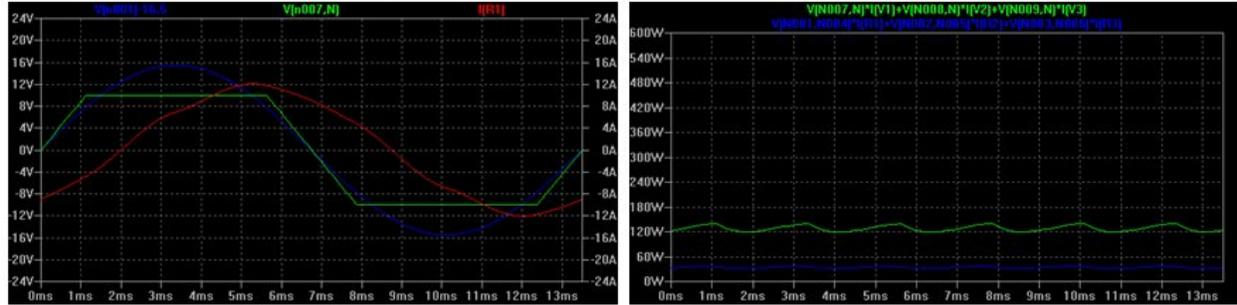
- Drive: Six-step as defined in Table 3, advanced 15°.
- DC Bus Voltage: 26.7V



- Average Power Out: 250W | Power Ripple: 107W | Power Dissipated: 53W

Case 3: Sinusoidal Drive, No Phase Advance

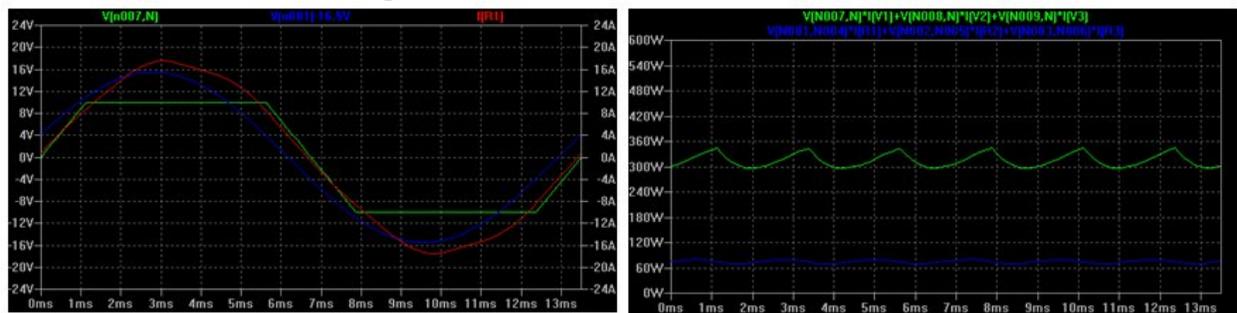
- Drive: Sinusoidal, in phase with back EMF fundamental.
- Drive Voltage Amplitude: 15.5V
- Minimum DC Bus Voltage: 31V



- Average Power Out: 128W | Power Ripple: 20W | Power Dissipated: 34W

Case 4: Sinusoidal Drive, 15° Phase Advance

- Drive: Sinusoidal, advanced 15° with respect to back EMF fundamental.
- Drive Voltage Amplitude: 15.5V
- Minimum DC Bus Voltage: 31V



- Average Power Out: 315W | Power Ripple: 48W | Power Dissipated: 75W

Even with a purely trapezoidal motor, sinusoidal drive offers some advantage in this case. Though it requires the use of phase advance and a higher DC bus voltage, the result is a higher achievable power with considerably less ripple as compared to pure BLDC control with or without phase advance. As the actual back EMF deviates from the pure trapezoidal form towards a more sinusoidal form, the benefits of sinusoidal drive will become even greater. Effectively controlling the phase advance is a key component of sinusoidal control with any type of motor, and it is this control that is studied further in later sections.

5 Appendices

5.1 MOSFET Selection and Analysis

MOSFETs are the preferred transistors for low-voltage motor controllers and are in many ways the closest thing to an ideal, voltage-controlled switch that power semiconductors have to offer. MOSFETs differ from other types of transistors in important ways that make them well-suited for low voltage applications. Specifically, they do not have a saturation voltage but are instead modeled as having a constant resistance (linear current vs. voltage) when fully on. This may seem particularly non-ideal, but in low-voltage applications where the saturation voltage of a BJT or IGBT would be a significant portion of the total system voltage, the linear I-V curve of a MOSFET is preferable. Many modern MOSFETs have extremely low resistance values, and they can be easily paralleled for more current capability since they have a positive temperature coefficient. As a result, very high power densities can be achieved in small controller packages.

Just as there are two different types of bipolar transistors (NPN, PNP), there are two different types of MOSFET (N-channel, P-channel). There are also differences between power MOSFETs and signal MOSFETs. The vast majority of motor controllers use only n-channel power MOSFETs, so only that type is described here. From this point forward, the use of “MOSFET” implies an N-channel power MOSFET. The electrical symbol for a MOSFET is shown in Figure 48. The majority of MOSFETs have three accessible terminals (or pins) a gate, a drain, and a source. The presence of an anti-parallel “body diode” is indicated in the symbol; this diode is part of the MOSFET and will *always* conduct current from source to drain if the voltage of the source is higher than that of the drain.

N-Channel Power MOSFET

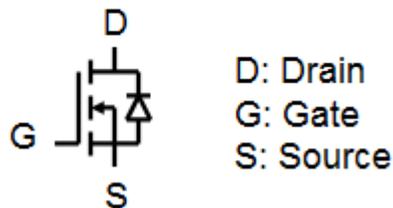


Figure 48: The electrical symbol for an N-channel power MOSFET.

As with any power transistor, the function of the MOSFET is to convert a low-power signal into an amplified high-power output. The low-power signal is applied to the gate, in the form of a voltage. This voltage, called the gate voltage, is measured *with respect to the source*. This becomes important for driving MOSFETs which do not have their source connected to the lowest system voltage (see Section 1.1.2). One way to think about a transistor is as a variable resistance controlled by the low-power signal. In the case of a MOSFET, the resistance from the drain to the source is controlled by the voltage from the gate to the source. When the gate voltage is below a certain “threshold” value, or if it is negative, the resistance from the drain to the source is very high and very little current flows. Above that threshold voltage, the resistance from drain to source will decrease as the gate voltage increases. There is a lower limit to the resistance,

which for a power MOSFET is typically achieved when the gate voltage is at or above 8-10V. This can be seen in a typical set of power MOSFET I-V curve, such as in Figure 49. This curve set is for the International Rectifier IRFB3077PbF MOSFET.

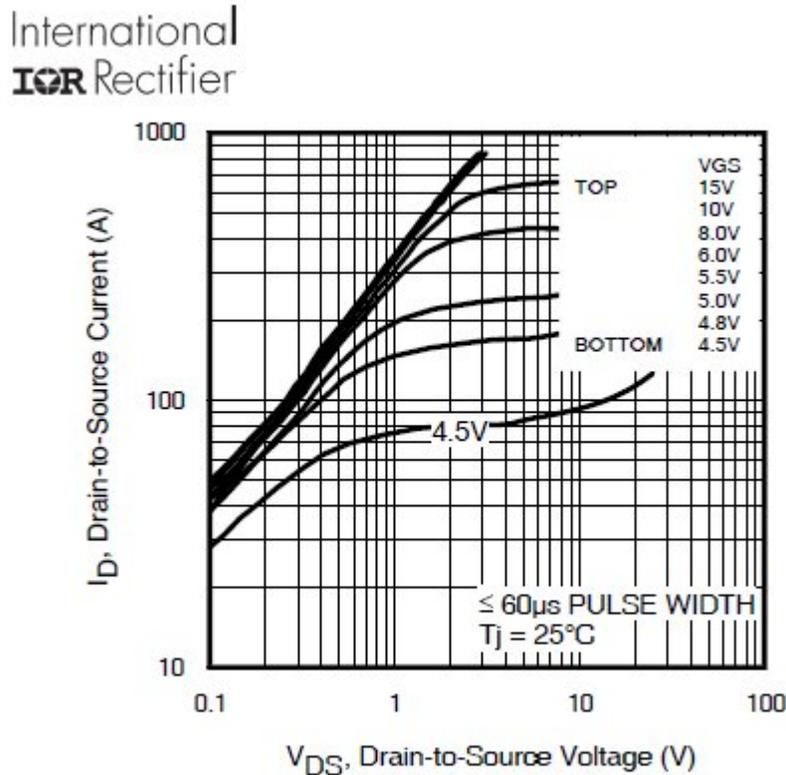


Figure 49: A typical set of power MOSFET I-V curves, showing that above 8.0V gate voltage, the output characteristic is that of a constant resistance, in this case about 3m Ω . Source: <http://www.irf.com/product-info/datasheets/data/irfb3077pbf.pdf>

MOSFETs are most often used at the two extremes of the I-V curves. When the gate voltage is zero (or negative), the resistance is so high that the MOSFET can effectively be treated as an open-circuit, keeping in mind that the body diode may still conduct from source to drain. When the gate voltage is above 8-10V, it can be considered fully on and treated as a constant resistance. MOSFET ratings always include the value of this resistance, often labeled R_{DS} or $R_{DS(on)}$. The MOSFET depicted in Figure 49 has a $R_{DS(on)}$ rating of 3.3m Ω .

To the extent which the MOSFETs in a motor controller are kept in one of these two states at all times (a bad assumption, which will be looked at below), a first-order sizing of the MOSFETs can be very easy. As an example, a one could ask if the IRFB3077PbF would be a good MOSFET to use for a 1.2kW (50A @ 24V) brushed DC motor controller. Table 8 lists some other relevant specifications for this MOSFET. The first thing to check is that the MOSFET voltage rating is a good deal higher than the system voltage, which in this case it is. With an on-state resistance of 3.3m Ω , the power dissipated in the MOSFET at 50A would be:

$$P_{dis} = I^2 R_{DS(on)} = (50A)^2 (0.0033\Omega) = 8.25W.$$

Table 8: Some specifications for the IRFB3077PbF MOSFET. Source: <http://www.irf.com/product-info/datasheets/data/irfb3077pbf.pdf>

Specification	Value
V_{DSS}	75V
$R_{DS(on) \text{ max.}}$	3.3m Ω
$R_{\theta JA}$	62 °C/W
$R_{\theta JS}$	0.9°C/W
$T_J \text{ max.}$	175°C

The first thing to note is the apparent efficiency of the MOSFET. Only 8.25W are dissipated from a 1.2kW system, yielding an efficiency of 99.3%. The voltage drop across the MOSFET is only 0.165V at 50A. Compared to that of an IGBT or BJT, the MOSFET voltage drop is much lower in this case. Only at much higher currents, when the IR drop becomes more than a typical IGBT saturation voltage, would an IGBT be more efficient. This isn't the whole story, since the MOSFET has other loss mechanisms which will be discussed below, but in a well-designed controller this can account for the majority of the loss. Often, the losses in the bus capacitors are on par with the losses in the MOSFETs themselves. Even so, a motor controller's efficiency can be in the upper 90% range.

By multiplying the dissipated power by the thermal resistance of the MOSFET, the temperature rise at the semiconductor junction can be found. If that temperature rise above ambient is well under the maximum junction temperature, T_J in Table 8, the MOSFET will work. MOSFET datasheets usually give multiple thermal resistance values. One, here labeled $R_{\theta JA}$, may denote the junction-to-ambient thermal resistance, if there is no heat sinking or forced convection. In this case, with 8.25W dissipation, the MOSFET cannot be operated with free-air convection only. However, the thermal resistance to a greased, well-cooled heat sink, here labeled $R_{\theta JS}$, is much lower. Under these conditions, this MOSFET will easily handle 50A. These numbers, especially the relative difference between free-air convection and heat sinking, are representative of many MOSFET packages. The importance of heat sinking and heat removal from the sink itself should be very clear from this. This and the heating of bus capacitors represent the two biggest limitations on continuous current capacity of a motor controller.

There are two other main loss mechanisms for a MOSFET: diode loss and switching loss. Diode loss is the most straightforward to calculate: it is the heat dissipated in the body diode when current is conducted through it. The characteristics of this diode are typically given as an I-V chart in the datasheet, such as is shown in Figure 50. The current in this case is the reverse current flowing through the diode. By multiplying the current by the diode voltage drop, the power dissipated by the diode can be found. In this case, the diode drop at 50A and 25°C would be 0.7V. The power dissipated by diode conduction is thus:

$$P_{dis} = IV_{SD(50A)} = (50A)(0.7V) = 35W.$$

This is much larger than the conductive dissipation of the MOSFET, and for controllers rely on the body diodes to conduct current under normal operation, this can be the most significant source of dissipation. In some brushed DC motor controllers, this is the case. In most brushless

and three-phase controller, though, diode conduction is avoided by switching the high and low side MOSFETs with complementary signals, such that one or the other is always on. This is called synchronous rectification, and is discussed in more detail in Section 1.1.3. Current may still flow from source to drain, but it flows through the switch element itself instead of the diode. Thus, the losses are the same as the conductive losses in the forward direction. The diode is still present as a back-up, allowing current to flow somewhere at all times to prevent voltage spikes.

IRFB3077PbF

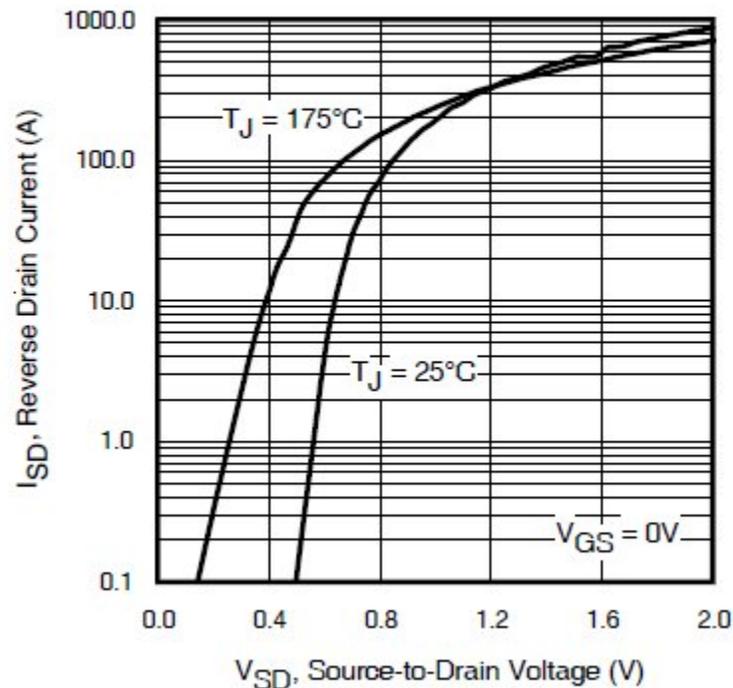


Figure 50: A typical I-V curve for the body diode of a MOSFET, in this case the IRFB3077PbF. Source: <http://www.irf.com/product-info/datasheets/data/irfb3077pbf.pdf>

In six-step square-wave drive, described in Section 4.1, a diode still conducts during the “off” period of the commutation cycle, dissipating the energy stored in the inductance of the phase being turned off. Most of the energy is dissipated by the windings of the motor, but some is lost in the MOSFET diode as well. Depending on the winding time constant, the fraction of time spent in diode conduction may actually be very low. A very conservative worst-case estimate can be made by assuming that, at any given time, one diode in the three-phase bridge is conducting the full motor current. Obviously this can’t be true, but it will give a high estimate of heat generation which can be used to size MOSFETs with considerable safety margin.

In sine-wave drive, three MOSFETs are on at any given time (synchronous rectification of each half-bridge) and diode conduction is rarely required. Thus, diode losses are minimal. Only during faults, brief dead times, and other short-duration events do the diodes conduct. However, it may still be appropriate to size the MOSFETs such that one continuous diode loss can be tolerated, for a large margin of safety.

Switching loss is the last and more difficult to analyze MOSFET loss mechanism. This is the loss that occurs as a result of the fact that the MOSFET cannot be turned on or off instantaneously. The MOSFET gate has a small but important capacitance, modeled as in Figure 51, that creates a low-pass filter with the external gate resistor, slowing down the turn-on and turn-off transitions. As it transitions between having a very large resistance while off to having its on-state resistance, it passes intermediate resistance values that dissipate much more power than either extreme. Because the resistance from drain to source is a nonlinear function of the gate voltage, it is hard to analytically estimate the switching losses with no point of reference. A brute-force approach to the analysis might be to integrate the instantaneous power loss along the path the MOSFET takes through an I-V curve as the gate is charged. Simulation is the most effective method for predicting switching losses, as tools like SPICE have extensive MOSFET models that accounts for the nonlinearities. Since the goal of this analysis is to size the MOSFET with a large safety margin, a simple worst-case method is presented instead.

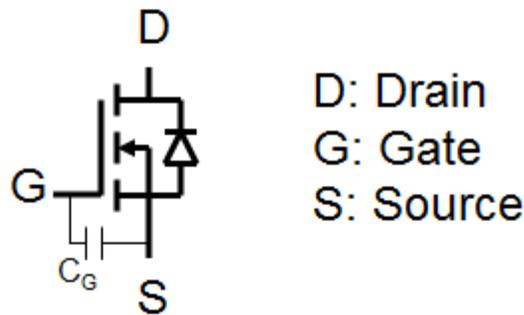


Figure 51: The effective internal capacitance from the gate to the source causes finite turn-on and turn-off time.

Though the path through the I-V curve during a switching event is unknown, the most extreme case would be that during a switching transient, the MOSFET has both the maximum voltage across it and the maximum current flowing through it at the same time. This means that for the entire duration of the turn-on or turn-off time, it takes on the worst possible resistance value: the one that dissipates the most power. The maximum voltage is simply the full DC voltage of the motor controller power supply. The maximum current could have many meanings. For extreme robustness, this could be the motor stall current. In most cases, it is the expected maximum operating current, which is much less than the stall current.

Another way of looking at it is that, in the worst-case scenario, *all* of the power is going into heating of the MOSFET during a switching transient. This should reveal the importance of a robust gate driver: If the MOSFET were somehow to get stuck in the half-on state, it would be disastrous. This is the most catastrophic, fastest, and in many cases most common failure mode of MOSFETs.

By keeping the switching transients very short compared to the time between them, the average power dissipation during switching can be kept low. The time of a switching transient is the time it takes the gate to get from 0V to the fully-on voltage (10V) or to get from 10V to 0V. Because gate drives have a limited current output, a gate resistor is used to keep the current applied in a

reasonable range. This, with the gate capacitance, determines a time constant that governs the voltage rise or fall on the gate. If a 15V gate drive is used, the time to get to 10V is approximately one time constant (63% of 15 is 9.45). A gross overestimate of the energy dissipated during one turn-on or turn-off event can be calculated, then as follows:

$$E_{dis} = I_{max} V_{max} R_G C_G,$$

where R_G is the gate resistor used in series with the gate driver and C_G is the effective gate capacitance, often listed on the MOSFET datasheet as the “input capacitance.” The power dissipated is this energy multiplied by the number of turn-on and turn-off events per second, or twice the switching frequency.

$$P_{dis} = 2f_{sw} E_{dis} = 2f_{sw} I_{max} V_{max} R_G C_G.$$

Using the maximum current and voltage of the example 1.2kW DC motor drive, as well as the published C_G of the IRFB3077PbF MOSFET (9.4nF), a 10Ω gate resistance (~1.5A gate driver current), and a 16kHz switching frequency, the power dissipated through switching is 3.61W. Even with this worst-case scenario analysis, the switching dissipation is less than the conductive dissipation and much less than the diode dissipation. In a well-designed motor controller, this should be the case. However, in many instances the time constant of the gate resistor and capacitance is too long or the switching frequency too high, resulting in large switching losses. For example, the well-known Open-Source Motor Controller (OSMC) schematic¹ uses a similar MOSFET and a 150Ω, which would increase switching losses by a factor of 15, all other parameters held constant.

The only thing intrinsic to the MOSFET itself that affects switching losses are the shape of the I-V curve as a function of gate voltage, and the gate capacitance. As far as the worst-case analysis goes, only the gate capacitance matters. The smaller the gate capacitance, the faster a MOSFET will turn on and off with a given gate driver, which will result in lower switching losses. Sometimes, the MOSFET datasheet will list a gate charge, in Coulombs, instead of a gate (input) capacitance. In this case, the gate voltage will usually be specified as well, and the gate capacitance can be calculated by:

$$C_G = \frac{Q_G}{V_{GS}}.$$

In other cases, a MOSFET datasheet may explicitly define a switching energy dissipation at a given set of operating conditions (a gate voltage and resistance, as well as the operating voltage and current across the switch, are the necessary parameters). In this case, the switching losses can be found by scaling this value appropriately, with the following considerations in mind:

1. Switching losses in a MOSFET will be the switching energy loss scaled by double the switching frequency, to account for turn-on and turn-off time.

¹ OSMC v3.22 Schematic: <http://www.robotpower.com/downloads/osmc3-22sch-clean.pdf>

2. Switching energy loss will scale proportionally with switched current.
3. Switching energy loss will scale proportionally with switched voltage (V_{DS}).
4. Switching energy loss will scale proportionally with the gate resistance used.

For example, Figure 52 is an excerpt from the IXYS GWM 100-01X1 MOSFET module used in the scooter controller. In this datasheet, Q_g is given instead of C_g . The gate capacitance is:

$$C_G = \frac{Q_G}{V_{GS}} = \frac{90nC}{10V} = 9nF.$$

So its gate capacitance is similar to that of the IRFB3077PbF, which is expected since they are of similar power and voltage ratings. With the same criteria used to evaluate the worst-case switching loss in the IRFB3077PbF, the worst-case loss in the 100-01X1 is calculated to be 3.46W. Now using the explicit energy loss value given in the datasheet, and scaling appropriately for the same operating conditions, the more accurate switching loss is calculated as:

$$P_{dis} = 2f_{sw}E_{dis} = 2f_{sw}E_{on} \left(\frac{24V}{48V} \right) \left(\frac{50A}{70A} \right) \left(\frac{10\Omega}{33\Omega} \right) = 1.39W.$$

Thus the worst-case estimate was about 2.5 times higher than the more accurate result calculated using the appropriately-scaled energy loss numbers from the datasheet. It should also be noted that the gate drive voltage used to calculate the energy loss in the datasheet is 10V, not 15V. Using a higher gate drive voltage will further reduce the switching losses. In general, the worst-case analysis should yield a result that is several times larger than the true dissipation. In either case, the switching losses are smaller than the conductive losses of the example 1.2kW DC motor driver.

Q_g	} $V_{GS} = 10V; V_{DS} = 65V; I_D = 90A$	90	nC
Q_{gs}		30	nC
Q_{gd}		30	nC
$t_{d(on)}$	} inductive load $V_{GS} = 10V; V_{DS} = 48V$ $I_D = 70A; R_G = 33\Omega;$ $T_J = 125^\circ C$	130	ns
t_r		95	ns
$t_{d(off)}$		290	ns
t_f		55	ns
E_{on}		0.4	mJ
E_{off}		0.4	mJ
E_{recoff}	0.007	mJ	

Figure 52: An excerpt from the IXYS GWM 100-01X1 MOSFET module datasheet. Source: http://download.siliconexpert.com/pdfs/2008/12/03/semi_b/1/ixy/mosfet%20modules/gwm100-01x1.pdf

The purpose of this analysis is to highlight the most important factors to consider when choosing a MOSFET. Often, there are tradeoffs between these factors. In choosing a MOSFET or MOSFET module for a motor controller these tradeoffs become important design decisions.

There are also other factors not related to the electrical specifications of the MOSFET. For example, form factor and ease of assembly. MOSFET modules (single packages with two or more MOSFETs in them) offer a significant advantage in this regard, since they can be attached to an isolated heat sink.

The first version of the scooter controller used six discrete IRFB3077PbF MOSFETs per motor. Whereas the IRFB3077PbF offers advantages over the GWM 100-01X1 in some performance categories (lower R_{DS} and R_{θ}), and is similar in others (C_G), the benefit of having a single package three-phase bridge made the IXYS module the better choice for the more compact second version of the scooter controller. Figure 53 shows a the two GWM 100-01X1 modules installed on the controller board, heat tab facing out for thermal pasting to a heat sink.

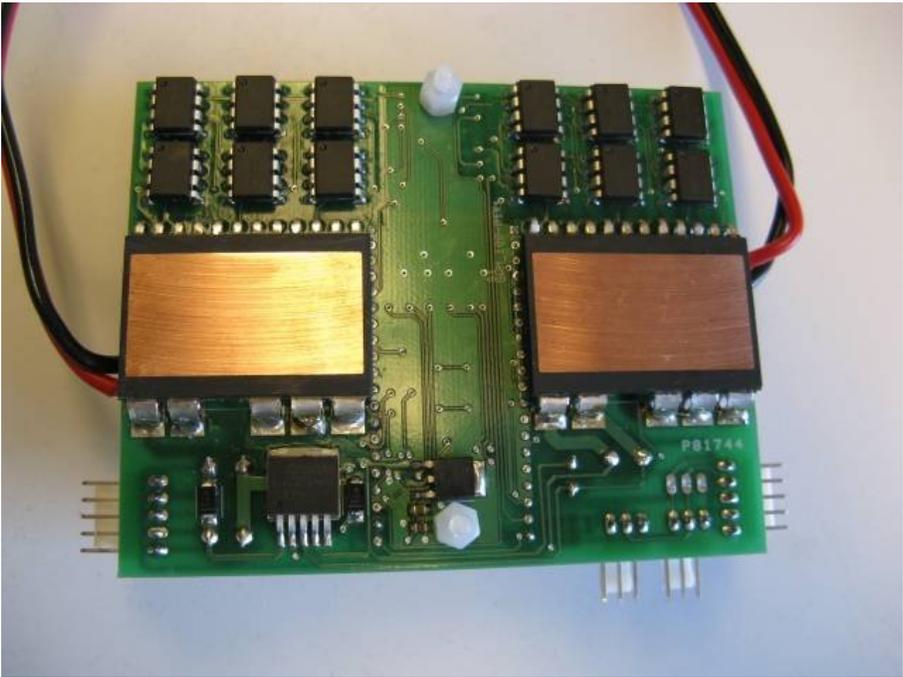


Figure 53: Two IXYS GWM 100-01X1 modules installed on the bottom side of the scooter controller.

The following analysis seeks to estimate the power dissipation and temperature rise of the MOSFET modules under the operating conditions of the scooter motors, as well as at a more extreme maximum operating point, representative of the maximum values imposed by components other than the motors. Table 9 defines the two operating points. Additionally, two drive conditions are considered: sine wave drive with synchronous rectification and exaggerated worst-case square wave drive with one diode always conducting the full motor current.

Table 9: Two operating points to be evaluated for MOSFET heat dissipation.

	Scooter	Maximum
Operating Voltage	33V	48V
Current	20A	30A
Heat Sink Temperature	30°C	75°C
Switching Frequency	15kHz	15kHz
Gate Resistance	10Ω	10Ω

First is an estimate of the power dissipation by conductive losses. In square wave drive, only two MOSFETs in the bridge are conducting current at any given time. In sine wave drive, three are conducting at any one point, but the RMS current is the relevant parameter. For the purpose of this analysis, the RMS current will be set to the value listed in Table 9. Note that the two drives (square wave versus sine wave) will have different torque-speed outputs! The relevant equation is:

$$P_{dis,on} = nI^2 R_{DS(on)},$$

where n is the number of MOSFET conducting, I is the current being considered, and $R_{DS(on)}$ is the on resistance of a single MOSFET. Since the $R_{DS(on)}$ value is a function of temperature, a higher value² than that specified in the datasheet is used to calculate the conductive dissipation at the maximum operating point. Table 10 shows the calculation of conductive dissipation for the four cases being considered.

Table 10: The conductive dissipation, $P_{dis,on}$, for the four cases being considered.

Scooter, Square-Wave Drive $P_{dis,on} = (2)(20A)^2(7.5m\Omega) = 6W$	Maximum, Square Wave Drive $P_{dis,on} = (2)(30A)^2(14m\Omega) = 25.2W$
Scooter, Sine Wave Drive $P_{dis,on} = (3)(20A)^2(7.5m\Omega) = 9W$	Maximum, Sine Wave Drive $P_{dis,on} = (3)(30A)^2(14m\Omega) = 37.8W$

Next, the diode dissipation is considered for all four cases. In both sine wave drive cases, the diode dissipation is negligible. In square wave drive, an exaggerated worst-case estimate can be made by assuming one diode (either the high side or low side) in the “off” phase leg is conducting the full motor current at any given time. The diode I-V characteristic at different junction temperatures³ is given in Figure 16 of the GWM 100-01X1 datasheet. The equation used to calculate diode dissipation is:

$$P_{dis,d} = IV_{SD},$$

where I the current being considered and V_{SD} is the diode voltage drop, given by the MOSFET data sheet for different values of diode current. Table 11 shows the calculation of diode dissipation for all four cases.

Table 11: The diode dissipation, $P_{dis,d}$, for the four cases being considered.

Scooter, Square Wave Drive $P_{dis,d} = (20A)(0.8V) = 16W$	Maximum, Square Wave Drive $P_{dis,d} = (30A)(0.82V) = 24.6W$
Scooter, Sine Wave Drive $P_{dis,d} = 0W$	Maximum, Sine Wave Drive $P_{dis,d} = 0W$

² This higher value can be found iteratively using the chart in Figure 5 of the GWM 100-01X1 datasheet. For simplicity, it is taken to be the value at $T_j=125^\circ\text{C}$ here.

³ For simplicity, the 25°C curve is used for all cases. Since the diode voltage drop decreases with temperature, this yields a conservative estimate for the higher-temperature case.

Finally, switching losses are considered. By appropriately scaling the given energy dissipation per switching transient to the operating conditions under consideration, an estimate of the switching loss power can be made:

$$P_{dis,sw} = 2nf_{sw}E_{dis} = (2n)(15kHz)(0.4mJ)\left(\frac{V}{48V}\right)\left(\frac{I}{70A}\right)\left(\frac{10\Omega}{33\Omega}\right),$$

where n is the number of MOSFETs being switched at any given time, V is the operating voltage, and I is the operating current. In square wave drive, two MOSFETs are switched at a time. In sine wave drive, all six are switched. Here, RMS current is not really appropriate, since the operating voltage is a peak quantity. However, the switching dissipation is already a roughly estimated quantity it will be seen that these losses represent a small fraction of the total losses in the cases being considered, making the true nature of V and I used here less important to MOSFET selection. For brevity, the full calculation is omitted and the results for all four cases are shown in Table 12.

Table 12: The switching dissipation, $P_{dis,sw}$, for the four cases being considered.

Scooter, Square Wave Drive $P_{dis,sw} = 1.43W$	Maximum, Square Wave Drive $P_{dis,sw} = 3.12W$
Scooter, Sine Wave Drive $P_{dis,sw} = 4.29W$	Maximum, Sine Wave Drive $P_{dis,sw} = 9.35W$

Adding the three components of dissipation together gives the total power dissipated, $P_{dis,total}$, in the MOSFET in these four cases. Multiplying by the MOSFET module's thermal resistivity, $1.1^{\circ}C/W$, gives the temperature rise of the junction above the heat sink temperature. Table 13 summarizes this.

Table 13: The total dissipation and junction temperature rise for the four cases considered.

Scooter, Square Wave Drive $P_{dis,total} = 23.4W$ $T_j = 30^{\circ}C + (23.4W)(1.1^{\circ}C/W) = 55.7^{\circ}C$	Maximum, Square Wave Drive $P_{dis,total} = 52.9W$ $T_j = 75^{\circ}C + (52.9W)(1.1^{\circ}C/W) = 133^{\circ}C$
Scooter, Sine Wave Drive $P_{dis,total} = 13.3W$ $T_j = 30^{\circ}C + (13.3W)(1.1^{\circ}C/W) = 44.6^{\circ}C$	Maximum, Sine Wave Drive $P_{dis,total} = 47.2W$ $T_j = 75^{\circ}C + (47.2W)(1.1^{\circ}C/W) = 127^{\circ}C$

The junction temperature calculations in all four cases are still within the maximum junction temperature operating range of the MOSFET. In the case of the scooter motor operating point, it is clear that the MOSFETs are not the bottleneck. The "maximum" operating point, with a sink temperature of $75^{\circ}C$, represents the upper limit of what the controller might be capable of. Other limiting factors, such as the current sensors, board traces, and DC bus capacitor, may limit operation further and are explored in separate sections. In all four cases, the MOSFET efficiencies are very high. This efficiency is not very meaningful, though, since there are still losses in the other components to consider.

5.2 Hall Effect Sensor Fault Tolerance

For a motor with Hall effect sensor-based commutation, the controller should be designed to tolerate sensor faults, for examples: (1) a missed signal, (2) an extra signal, (3) an incorrect signal state, or (4) a complete loss of signal. (1), (2), and (3) could be fairly common in normal operation. (4) would indicate a more severe problem, such as a disconnected sensor cable.

The sensor fault tolerance of the controller is a function of the operating conditions when the fault occurs, as well as the duration of the fault. Non-repetitive faults that last for only the time between two commutation interrupts would have little effect on the motor or power electronics, since little heat would be generated during that period of time. The assumption is that the sensor fault will be corrected by the next sensor signal. One case where this might not hold true would be if the fault puts extraordinary stress on the DC-link capacitor such that its voltage drops below the cutoff threshold for the auxiliary power supplies, which could cause a controller brown-out or reset.

The worst-case scenario sensor fault would occur at full speed and would put the applied voltage at the exact opposite polarity as the back EMF. This is full-speed “plug braking,” with no current control, and the transient would put the most severe stress on the bus capacitor. Figure 54 illustrates what this transient would look like electrically.

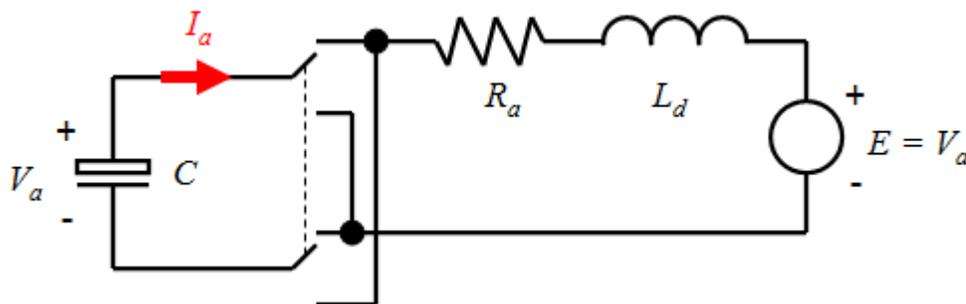


Figure 54: A worst-case sensor fault transient would instantaneously apply the full phase voltage 180° out of phase with the back EMF. This transient would cause the most stress on the DC bus capacitor.

The scenario of the fault is as follows: The motor is operating at full speed and no load, such that very little current is flowing. That is, I_a is zero until the time of the fault. The DC bus capacitor is fully charged to the DC voltage and the voltage applied to the motor is in phase with the back EMF. When the sensor fault occurs, it is as if the applied voltage is instantaneously switched to the reverse polarity, or put 180° out of phase with back EMF. This is illustrated in Figure 54 by a DPDT switch that is thrown at the time of the fault. One other assumption made is that only the bus capacitor supplies current during the fault, i.e. the battery cable inductance is very large. This is unrealistic, but will lead to a more conservative answer. Remembering that V_a is actually half the DC voltage, and that E remains constant at the initial value of V_a during the fault, the following second-order system describes the fault:

$$\frac{dV_{DC}}{dt} = -\frac{I_a}{C}$$

$$\frac{dI_a}{dt} = \frac{V_a + E - I_a R_a}{L_d} = \frac{0.5V_{DC} + E - I_a R_a}{L_d}$$

Simulating this system from $t=0$ using the operating parameters of the scooter motors and controller yields the voltage and current profiles in Figure 55.

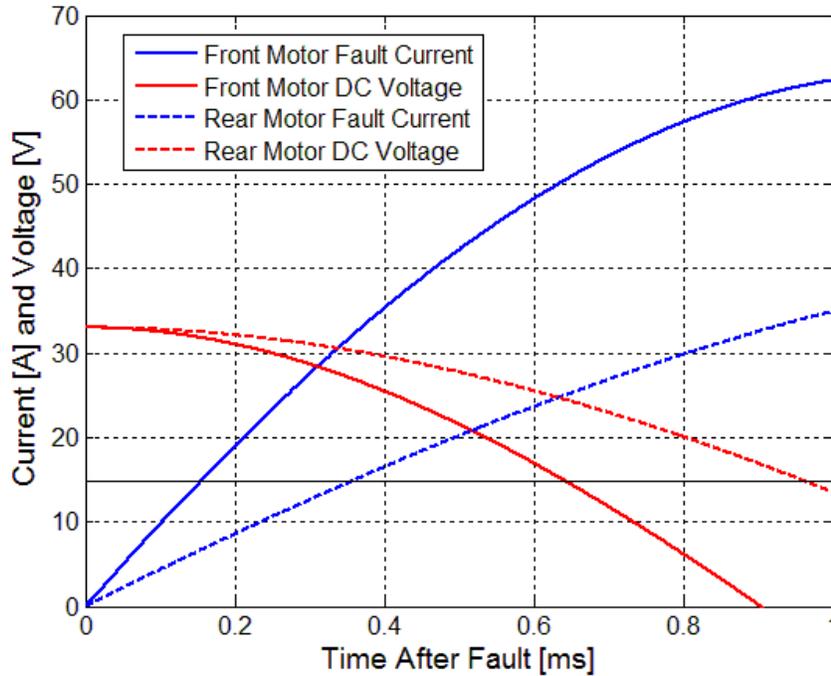


Figure 55: A simulation of the DC bus voltage and fault current for a worst-case sensor fault on the two scooter motors.

From the simulated fault, the time after which the DC bus capacitor is below the auxiliary power supply threshold of 15V is less than 1ms for faults on either scooter motor. This is below the time between Hall effect sensor transitions, even at full speed. Admittedly, the assumption that only the bus capacitor supplies current for the duration of the fault is far-fetched. The inductance and resistance of the battery cables will likely be much lower than those of the motor, so battery current is able to support the bus capacitance during the fault and keep the bus voltage above the threshold. However, the example shows that a fault of this nature could potentially put an extraordinarily large draw on the bus capacitor and should be avoided.

The likelihood of the Hall effect sensors themselves producing this particular worst-case fault is small. This is because of the nature of the Hall effect signals. Table 14 shows the six Hall effect signal states, and their relative phase angles. In order for the phase angle to be changed by 180° , all three bits would need to be inverted. It is far more likely that a single bit would be inverted during a fault, creating either a shift of 60° or an undefined state.

Table 14: The six defined signal states of the Hall effect sensors and the states required for the phase angle to be shifted by 180°. For this to occur, all three bits would need to be inverted.

Relative Phase Angle	Signal State	180° Shift
0°	001	110
60°	011	100
120°	010	101
180°	110	001
240°	100	011
300°	101	010

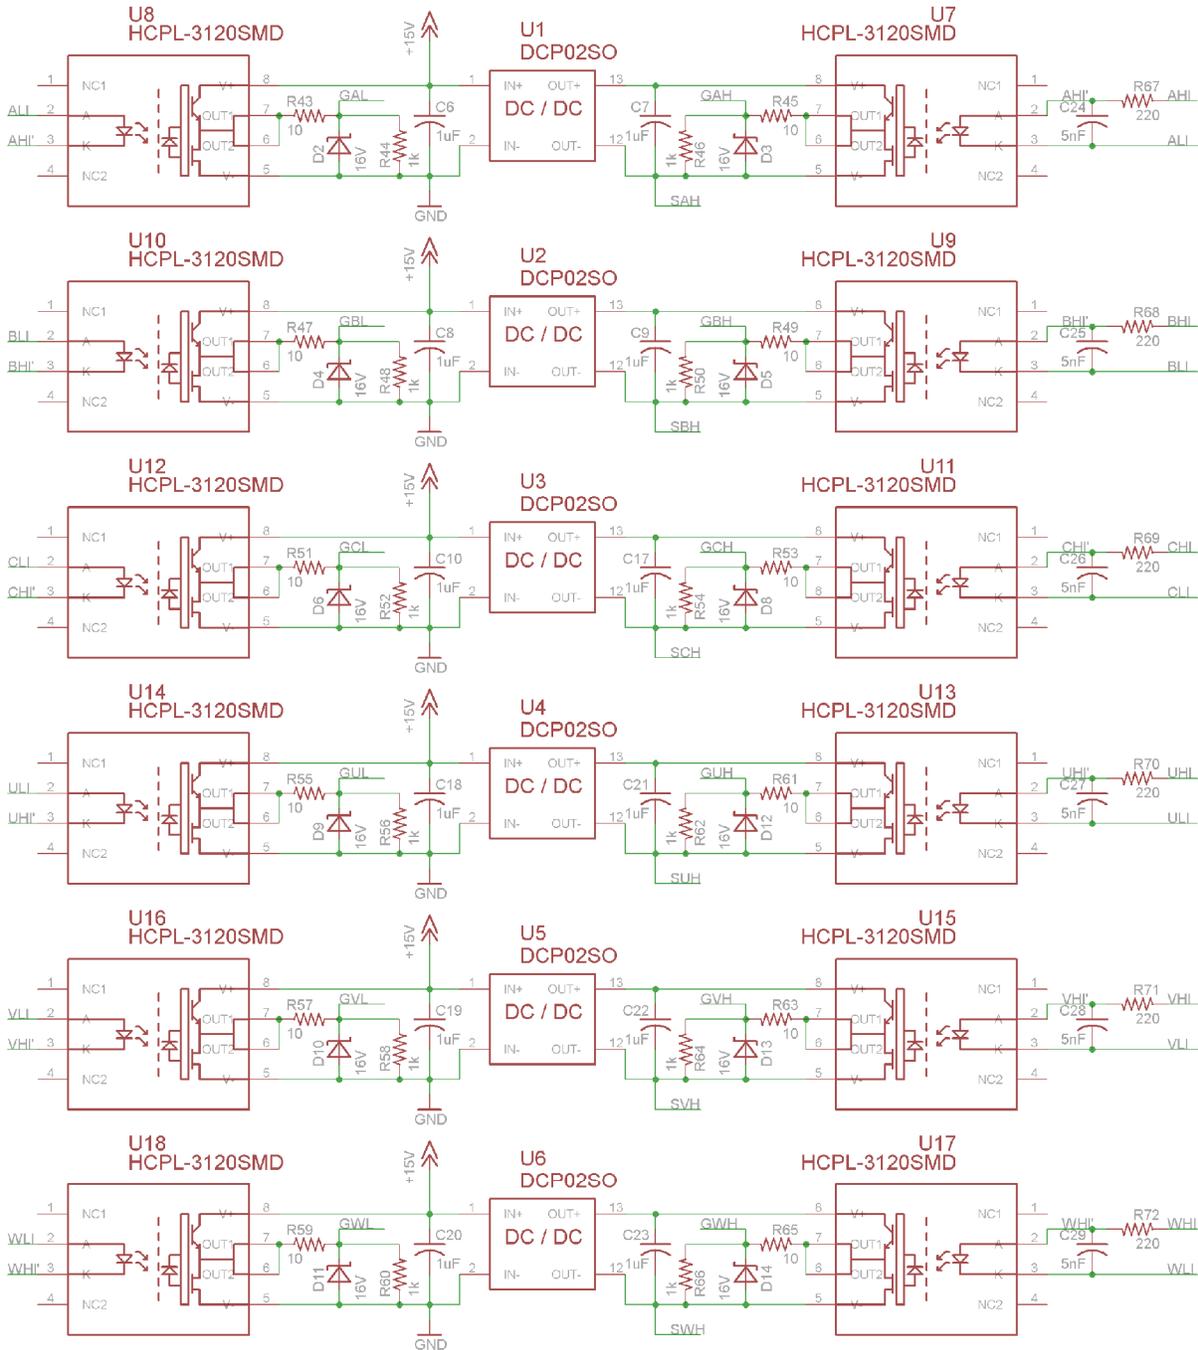
In the sine wave drive mode, the Hall effect sensors and open-loop timing are used to interpolate the exact position of the rotor (see Section TK). The transition of Hall effect sensor states sets an initial index in the sine table. Based on the estimated speed of the motor, the sine table index increments by dead-reckoning until the next transition. If this open-loop incrementing is left unchecked and a bad Hall effect signal creates a bad speed estimate, a fault such as the one described above can occur. Obviously, any speed estimates that are above the maximum speed of the motor can be thrown out. But in many applications, including vehicles, the speed range is so large that this alone will not prevent a bad speed estimate from causing unchecked progression through the sine table.

A simple way to prevent this scenario is to allow the open-loop incrementing to span only 60°. Under normal operation, the limit should be reached at almost the same time as a new signal transition comes in. If a new signal transition does not occur, the sine wave drive waits at that angle until a new signal comes in. This implicitly favors the Hall sensor transition points over the speed estimator. On the other hand, if a transition point is missed, it might be better to trust the speed estimator for one extra 60° span.

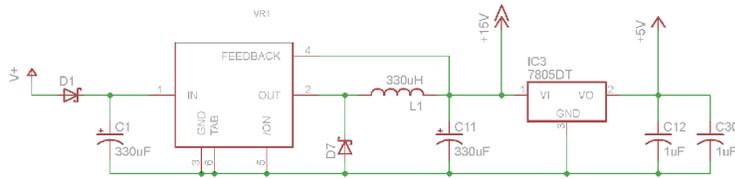
The most likely fault that lasts longer than one transition period is a complete loss of Hall effect signals. A disconnected signal cable could cause this, and this is a scenario that should be safely handled by the controller. The on-board pull-up resistors will send the sensor signal to {111}, an undefined state, in the event of a cable disconnect. When this undefined state is detected, the safest course of action is to open-circuit all three phases and allow the motor to coast. Any other action, such as shorting the three phases, could result in large currents flowing if the motor is driving an inertial load. Any system in which braking is critical should be equipped with a mechanical back-up, even if dynamic motor braking is the primary method.

5.3 Full Controller Schematic

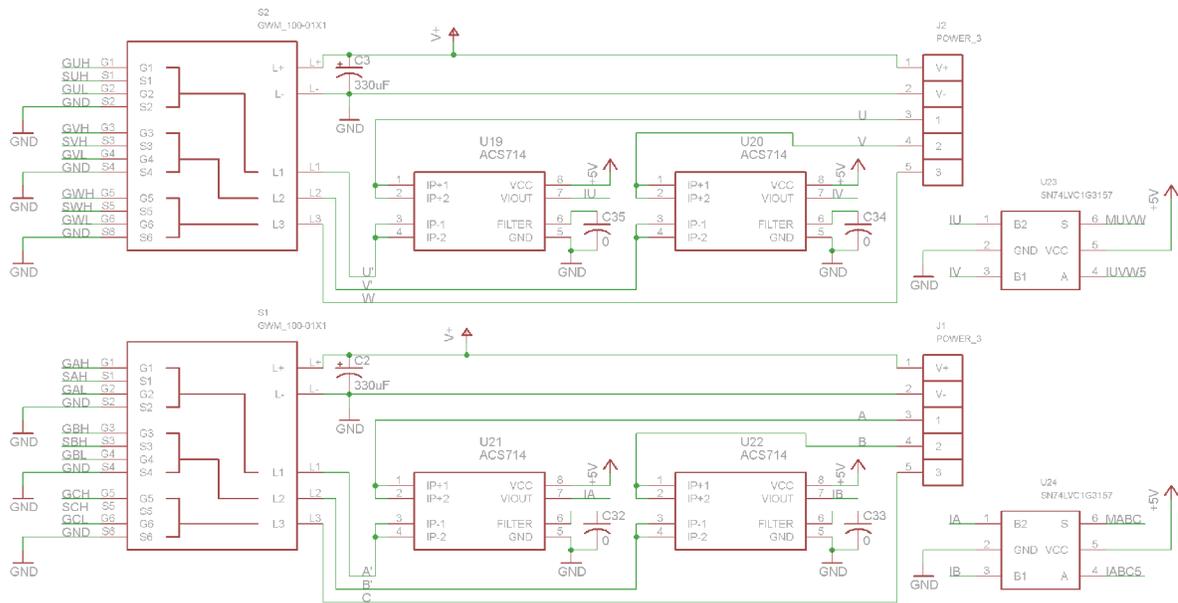
ISOLATED GATE DRIVE:



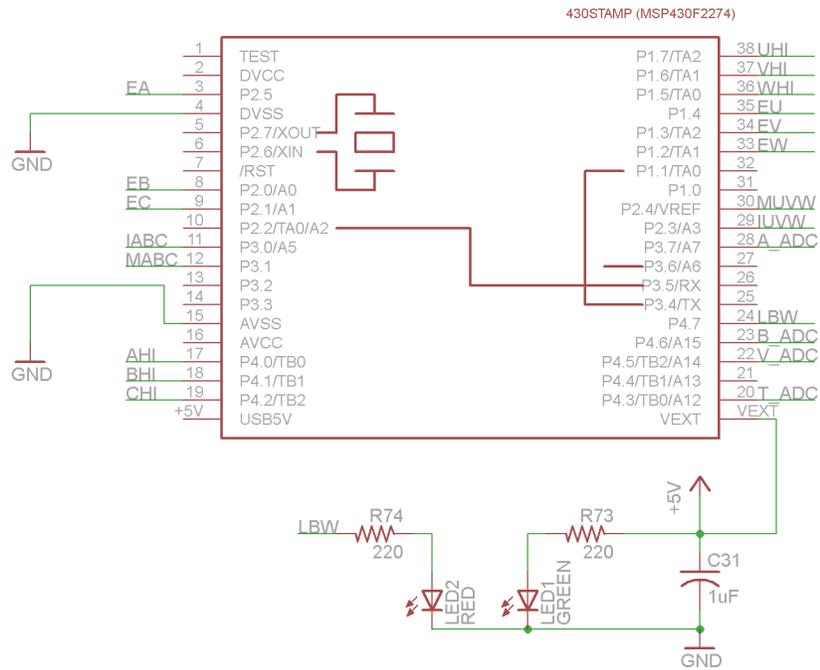
POWER SUPPLIES:



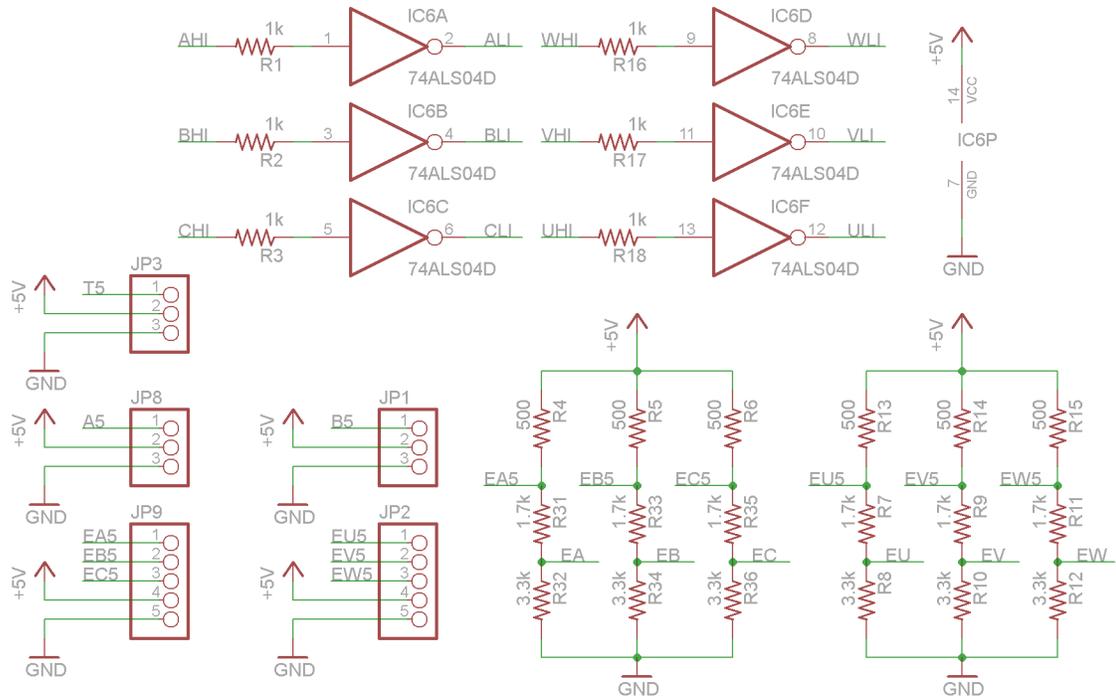
POWER INVERTERS:



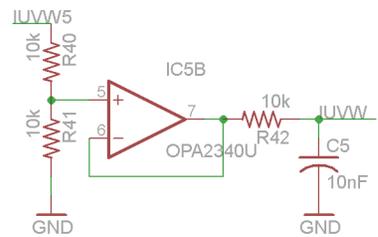
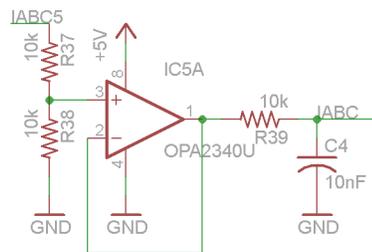
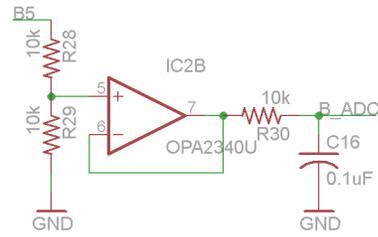
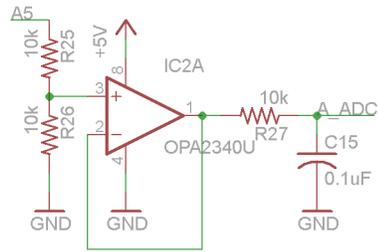
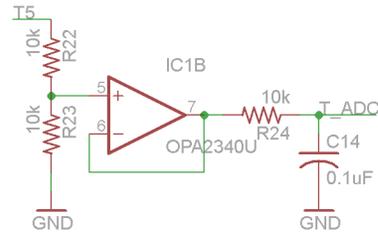
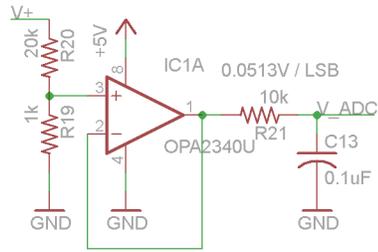
MICROPROCESSOR:



DIGITAL:



ANALOG:



5.4 Table of Parameters for Test Motors

Table 15: Properties of the two test motors.

Parameter (Symbol) [Units] Description	Value for Scooter Front Motor	Value for Scooter Rear Motor
Phase Resistance (R_a) [Ω] The resistance of a single phase (out of 3) of the motor windings. Since the motor is connected in a “wye” configuration, this is the line-to-neutral resistance of the winding.	0.111	0.167
Phase Inductance (L_a) [mH] The contribution to winding inductance of a single phase (out of 3) of the motor windings. The rotor is slightly salient, i.e. the inductance is weakly a function of rotor position, but for this analysis only the average phase inductance is considered.	0.22	0.50
BLDC Motor Constant (K_t) [N-m/A] This constant applies only to BLDC drive, as described in section 4.1. It is identically the “torque constant,” in [N-m/A], and the back EMF constant in [V/(rad/s)]. (Those units are equivalent.)	0.20	0.30
BLDC Motor Constant (K_v) [RPM/V] This is just the inverse of the constant given above, but in units compatible with the typical rating for BLDC motors in the hobby market.	47	32
Rotor Outer Radius (r) [m]	0.0635	0.0635
Rotor Inertia (J_R) [kg-m²] The moment of inertia of a single rotor.	0.0040	0.0040
Vehicle Inertia (J) [kg-m²] The combined inertia (rotational and linear) of two rotors, vehicle, and rider, reflected into the rotational domain of either rotor. A mass of 75kg is assumed for the rider and vehicle.	0.31	0.31

References

- [1] Mevey, J. (2009). *Sensorless Field Oriented Control of Brushless Permanent Magnet Motors*. M.S. Thesis. Kansas State University: USA.
- [2] *HCPL-3120/J312, HCNW3120: 2.5 Amp Output Current IGBT Gate Drive Optocoupler*. Datasheet. Avago Technologies. URL: <http://www.avagotech.com/docs/AV02-0161EN>.
- [3] *DCP02 Series: Miniature, 2W, Isolated Unregulated DC/DC Converters*. Datasheet. Texas Instruments. URL: <http://focus.ti.com/lit/ds/symlink/dcp020503.pdf>.
- [4] *PWM Speed Control*. 4QD-TEC Electronics Circuit Reference Archive. URL: <http://www.4qdtype.com/pwm-01.html>.
- [5] *LM1575/LM2575/LM2575HV: Simple Switcher 1A Step-Down Voltage Regulator*. Datasheet. National Semiconductor. URL: <http://www.national.com/ds/LM/LM1575.pdf>.
- [6] *GWM 100-01X1: Three phase full Bridge with Trench MOSFETs in DCB isolated high current package*. Datasheet. IXYS. URL: http://download.siliconexpert.com/pdfs/2009/10/13/10/35/2/360/ixy_/manual/gwm10001x1.pdf.