

Burst Physics: Outlaw Electric Motor Rules for Drag Racing, Speed Runs and Truck Pulling

Oct 24, 2026 · @Mike B

Purpose and scope

This book gives the math for electric motors driven to their limits for 1.4–3 seconds, not for years. Standard motor ratings assume steady state, where heat leaves as fast as it is made. A motor rated 4,200 W continuous says almost nothing about what it survives in a single pass.

The applications are RC drag racing (132 ft), RC speed runs and RC truck pulling. In each one, the motor sees extreme current for seconds, then cools between runs.

The book answers four questions for a given motor:

1. How much current and power can the winding take before the insulation fails?
2. At what current does the stator steel saturate, so extra amps stop making torque?
3. At what current and temperature do the rotor magnets start to demagnetize?
4. How fast can the rotor spin before it bursts, the bearings fail, or the voltage runs out?

Each limit is worked on a real motor and paired with a bench test readers can run.

Baseline assumptions

The reference motor is a surface-magnet permanent magnet synchronous machine (PMSM) with 0.20 mm silicon steel laminations. The hobby calls it "brushless" or BLDC. Strictly, BLDC means trapezoidal back-EMF and six-step drive; most RC motors have near-sinusoidal back-EMF but are driven six-step. Every ESC this book covers (Castle Creations, Ghost RC ACE, and XC-ESC's E6, E6 Pro and E6MAX) is a sensored six-step controller. FOC controllers exist, in crawler ESCs and in a separate XC-ESC FOC line, but they are not what drag, speed-run, truck-pull and boat racers run today. The book uses PMSM and dq-frame math to explain the physics, applied to six-step drive.

Torque model. With surface magnets, $L_d \approx L_q$, so reluctance torque drops out and torque is proportional to i_q :

$$T = \frac{3}{2} p \left(\lambda_{pm} i_q + (L_d - L_q) i_d i_q \right) \approx \frac{3}{2} p \lambda_{pm} i_q$$

In hobby units, $K_t \approx 9.55 \div kV$ (N·m per amp), with kV measured as no-load rpm per volt of DC bus. That K_t is per amp of DC-equivalent (battery-side, full-duty) current. The dq torque equation above uses peak phase current, and the ESC log reports battery current; the three are the same only at 100% duty. Everywhere this book says "amps" it names which one, because at launch they differ by 2–4×.

Lamination steel (stator and rotor). Values are typical for non-oriented grades sold as NO20 or 20JNEH; check the maker's datasheet for a specific motor.

Property	Value
Thickness	0.20 mm
Saturation, full	about 2.0 T
Design knee (torque per amp starts falling)	1.6–1.7 T
Density	about 7,650 kg/m ³
Specific heat	about 460 J/kg·°C
Core loss at 1.0 T, 400 Hz	about 8–15 W/kg by grade (NO20 near 13–15; 20JNEH near 8–12); use the maker's curve

Rotor. The rotor turns with the field, so its steel sees nearly constant flux and little core loss. The rotor is treated mainly as a mechanical problem (sleeve, burst, bearings) and a magnet-temperature problem.

Magnets. Sintered NdFeB. Grade and temperature set the demagnetization limit; data per grade is still to be gathered.

Upgrade sidebar. Cobalt-iron laminations (for example VACOFLUX 48 at 0.055 or 0.1 mm) saturate up to about 2.4 T but need annealing after cutting. The saturation chapter shows the same motor with a 2.0 T knee to show what the upgrade buys. Source: [Vacuumschmelze lamination stacks](#).

For the industrial reader

An RC PMSM is a synchronous machine with its field locked at one fixed excitation. The only way to change the effective field is from the stator, through armature reaction. This table maps big-machine terms to their RC counterparts.

Big synchronous machine	RC PMSM / ESC	Where the analogy breaks
Field winding and exciter	Rotor magnets	Excitation is fixed; it falls as magnets heat
Field current adjustment	Stator i_d , set through ESC timing	Costs stator amps and copper heat
Under-exciting / field weakening	Timing advance (negative i_d)	Also pushes the magnets toward demagnetization
V-curves	Torque and current vs timing angle	Shape changes strongly with rpm
Load (torque) angle δ	Not what ESC timing sets	Timing sets the current angle γ
Pull-out torque / slipping a pole	Can't slip; commutation follows the rotor	Limits are demagnetization and saturation
Six-step VSI drive	Most ESCs	Switches at kHz; true angle swings $\pm 30^\circ$ electrical
Closed-loop vector control	Sensored commutation: Hall sensors set the switching points	Six steps per electrical cycle, not continuous vector control; sensorless ESCs see little back-EMF at launch
Synchronous reactance	Very small L_d	Magnet thickness acts as airgap, so field weakening takes a lot of current
Thermal rating over years	Heat budget in joules over a 1.4–3 s pass	Copper is near-adiabatic during a pass
Short-time overload rating	The whole operating range	Current density runs 10–20 $\times$ continuous

Why the pull-out and reactance rows matter most. A big machine's large synchronous reactance gives the stator real authority over the field. In a surface-magnet rotor, the effective airgap is the real gap plus the magnet thickness, so L_d is tiny. Weakening the field by timing is therefore an expensive tool, paid for in current and magnet margin.

Book outline

One chapter per limit, then timing, then a chapter that combines them into a single chart of which limit hits first at each pass length.

1. **The copper heat budget.** Adiabatic heating in joules, current density, resistance rise with temperature, insulation class limits.
2. **Magnetic saturation.** Tooth and yoke flux density, the 1.6–1.7 T knee, K_t roll-off, measuring it with locked-rotor torque vs current. Sidebar: cobalt-iron.
3. **Magnet demagnetization.** NdFeB grades, temperature, the combined effect of launch i_q and timing i_d .
4. **Speed limits.** Rotor burst and sleeve stress, bearing limits, iron loss vs electrical frequency, back-EMF and $kV \times$ voltage.
5. **Timing.** Electrical vs mechanical degrees, the i_d / i_q split as $\cos \gamma$ and $\sin \gamma$, launch vs top-end, six-step angle swing, rpm-dependent boost. Includes a phasor diagram.
6. **Putting it together.** Which limit a given motor hits first for 1.4, 2 and 3 second passes, and whether an upgrade raises that limit.

Sidebar: why the hobby says BLDC; cobalt-iron upgrade path; for the industrial reader (above).

Chapter 1 draft: the copper heat budget

For a 1.4–3 second pass, the winding is limited by energy, not power: about 109 A/mm² for 2 seconds at a 150 °C rise. Heat takes tens of seconds to cross from copper into the stator, so during a pass the copper is treated as sealed (adiabatic).

The budget in joules

$$E_{\text{allow}} = m_{\text{cu}} \cdot c_{\text{cu}} \cdot \Delta T_{\text{max}}, \quad c_{\text{cu}} = 385 \, \text{J/kg} \cdot ^\circ\text{C}$$

Example: 50 g of copper allowed to rise 150 °C absorbs about 2,890 J. Over a 2 second pass that is about 1,440 W of copper loss, whatever the continuous rating says.

Current density, with resistance rising as it heats

Copper resistance climbs about 0.393% per °C, so heating speeds up as the pass goes on. Per unit volume, with J in A/mm², density d = 8,960 kg/m³ and $\rho_0 = 1.72 \times 10^{-8} \Omega \cdot m$ at 20 °C:

$$\Delta T(t) = \frac{1}{\alpha} \left(e^{\alpha J^2 t} - 1 \right), \quad \alpha = \frac{J^2}{\rho_0 d c_{cu}}$$

Solved for the allowed current density over a pass of length t:

$$J_{\max} = \sqrt{\frac{\ln(1 + \alpha \Delta T_{\max})}{\alpha t}}; \quad \alpha = \frac{J^2}{\rho_0 d c_{cu}}$$

Pass length	Allowed J, 100 °C rise	Allowed J, 150 °C rise
1.4 s	110 A/mm ²	130 A/mm ²
2.0 s	92 A/mm ²	109 A/mm ²
3.0 s	75 A/mm ²	89 A/mm ²

Continuous motors run about 5–10 A/mm², so a pass allows 10–20 times that. The resistance rise matters most on longer pulls: at a fixed 100 A/mm², the simple constant-resistance estimate gives 150 °C at 3 s, but the real rise is about 204 °C.

Phase current is not battery current

At launch the ESC runs a low duty cycle, so phase current is roughly battery current divided by duty. A 200 A battery reading at 25% duty can mean about 800 A in the winding. The heat budget must use phase current.

Limits of this model

- Proximity and skin effect in the strands at kHz electrical frequency raise AC resistance above the DC value.
- Hot spots in the end turns and solder joints run hotter than the average.
- The model starts from ambient; back-to-back passes start hot and lose budget.

Bench test (to write): measure phase resistance cold, run a timed locked-rotor pulse, measure resistance again, and back out the average copper temperature from the change.

Reading an ESC log: battery amps to winding amps

Phase current $\approx$ battery current $\div$ duty while the motor is slow enough that back-EMF is small. Use this to turn a Castle log into the number the heat budget needs.

ESC duty at the moment logged	Winding current per logged battery amp	Typical moment
20%	5.0×	First 50 ms with Start Power or a slow punch
25%	4.0×	Early launch, throttle ramp still on
33%	3.0×	Launch, ramp ending
50%	2.0×	Mid-track
75%	1.3×	Approaching trap
100%	1.0×	Full throttle near top speed

A logged 300 A peak that occurs at 25% duty means about 1,200 A in the copper for that moment. The X.BUS phase-current channel removes the guesswork; with Castle logs, read throttle % next to current and divide.

Tie to the EVUL Battery Sizing Calculator

The calculator's peak amps, average amps and "inside motor current rating" check are all battery-side current, because its pack model is battery voltage, pack IR and pack power. The copper budget here runs on phase current. Rule: take the calculator's peak amps, divide by the launch duty you run (0.25–0.4 for a typical Start Power ramp), and check that number against the J_{max} table at the pass length. A motor that passes the calculator's 350 A rating check can still see 800–1,000 A in its leads at launch.

Thermal time constants: one pass is adiabatic, a session is not

Heat leaves the copper in two steps with very different speeds. Treat the motor as two lumps:

$$C_{\text{cu}} \frac{dT_{\text{cu}}}{dt} = P_{\text{cu}} - G_1 (T_{\text{cu}} - T_{\text{fe}}), \quad C_{\text{fe}} \frac{dT_{\text{fe}}}{dt} = G_1 (T_{\text{cu}} - T_{\text{fe}}) - G_2 (T_{\text{fe}} - T_{\text{amb}})$$

For the Chapter 7 motor: 53 g of copper gives $C_{\text{cu}} \approx 20 \text{ J/}^\circ\text{C}$; about 400 g of stator steel gives $C_{\text{fe}} \approx 185 \text{ J/}^\circ\text{C}$. Through slot liner and enamel, G_1 is about $1\text{--}3 \text{ W/}^\circ\text{C}$, so the copper-to-iron time constant $\tau_1 = C_{\text{cu}} / G_1$ is $7\text{--}20 \text{ s}$. From can to moving air G_2 is about $0.3\text{--}1 \text{ W/}^\circ\text{C}$, so iron-to-air $\tau_2 = C_{\text{fe}} / G_2$ is $3\text{--}10$ minutes.

What that means at the track:

- Within a 2 s pass, less than 10% of the copper heat reaches the iron. Adiabatic is the right model for the pass.
- Between passes, the 3,000 J spreads into the whole motor and raises it about 15°C ($3,000 \text{ J} \div 205 \text{ J/}^\circ\text{C}$). The copper comes back to the iron's temperature within a minute.
- The stator holds the day's heat. Starting from 25°C with a 150°C copper ceiling, each pass that starts hot has a smaller budget: at 90°C iron the copper may rise only 60°C , so the allowed J drops to about 60% of the cold value (J scales with $\sqrt{\Delta T}$ at short times).
- A fan on the can shortens τ_2 ; nothing but time shortens τ_1 .

Bench test (to write): heat the winding with a known DC current for 2 s, log winding resistance for the next two minutes, and fit τ_1 and τ_2 from the two slopes.

Inductance sets how fast the current can rise

At the first commutation the motor is stopped, back-EMF is zero, and the winding sees the full bus through the phase inductance:

$$\frac{di}{dt} = \frac{V_{\text{bus}} - E}{L}, \quad \Delta I_{\text{ripple}} \approx \frac{V_{\text{bus}}}{D(1-D)} f_{\text{PWM}} L$$

RC inrunners have line-to-line inductance of roughly $5\text{--}20 \mu\text{H}$. On a 25 V bus with $10 \mu\text{H}$, current rises at $2.5 \text{ A/}\mu\text{s}$, so 800 A takes about 0.3 ms: far shorter than any commutation step at launch, which is why "phase current = battery current $\div$ duty" holds from the first step. The same numbers set the PWM ripple: at 25% duty and 20 kHz, about 30 A peak-to-peak, which the heat budget can ignore. Where inductance does bite is the top end: at 60,000 rpm a 4-pole motor commutates every $83 \mu\text{s}$, and with a $10 \mu\text{H}$ winding the current can no longer reach its target within a step, so torque per amp falls and timing advance starts paying for itself. That rpm-dependent boost is the Chapter 5 subject.

Chapter 2 note: saturation rolls off Kt before the copper limit

The full saturation chapter is still to write, but one correction belongs in every calculation now. Past the tooth knee (1.6–1.7 T for silicon steel, about 2.2 T for cobalt-iron), extra amps add little flux, so torque per amp falls. A usable first-order model:

$$K_{t,eff}(I) = K_t \cdot \min\left(1, \frac{I_{knee}}{I}\right), \quad I_{knee} \approx I \text{ at which } B_{tooth} = 1.65 \text{ T}$$

Below the knee, torque is $K_t \times I$. Above it, torque flattens near $K_t \times I_{knee}$ and the surplus current only heats copper and demagnetizes magnets. For the Chapter 7 example the leading tooth reaches 1.8 T at the 525 A heat limit, so I_{knee} is nearer 450 A: the last 75 A of the heat budget buys almost no launch. The EVUL calculators use $K_t \times I$ with no roll-off; treat their launch-torque figures as upper bounds once current passes the knee. Measuring I_{knee} is the locked-rotor torque-vs-current test in the saturation chapter: torque stops climbing in proportion, and the bend is the knee.

Chapter 3 draft: magnet demagnetization

Demagnetization is the one limit that is permanent. Copper cools, saturation recovers, but a magnet pushed past its knee at temperature comes back weaker and the motor keeps that loss. The limit is set by three things stacked together: the magnet's own operating point, the field the stator current throws against it, and how hot the magnet is when that happens.

The magnet's operating point with no current

A surface magnet of thickness l_m facing an effective airgap g_{eff} (mechanical gap plus sleeve plus a Carter allowance for the slot openings) sits on a load line set by its permeance coefficient:

$$P_c \approx \frac{l_m}{g_{eff}}, \quad B_m = B_r \frac{P_c}{P_c + \mu_r}, \quad H_m = -\frac{B_m}{\mu_0} \frac{P_c}{P_c + \mu_r}$$

with $\mu_r \approx 1.05$ for NdFeB. For the Chapter 7 rotor (2.0 mm magnets, 0.5 mm gap, 0.3 mm sleeve): $P_c \approx 2.0 \div 0.8 = 2.5$, $B_m \approx 1.25 \times 2.5 \div 3.55 = 0.88$ T (the airgap flux already quoted in Chapter 7), and $H_m \approx -280$ kA/m. That 280 kA/m is the magnet's own reverse field before the stator does anything. Thin magnets with big gaps start deep in the second quadrant; thick magnets start near B_r .

What stator current adds

Stator current produces a magnetomotive force per pole that drives extra reverse field through the magnet. For a three-phase winding with N_{ph} series turns per phase, winding factor k_w , p pole pairs and peak phase current I :

$$F_a \approx 0.9 \frac{N_{ph} k_w p I}{l_m + \mu_r g_{eff}} \quad \Delta H \approx \frac{F_a}{g_{eff}}$$

Chapter 7 motor, 1 turn per coil, 2 paths ($N_{ph} = 2$, $k_w \approx 1$, $p = 2$) at the 525 A heat limit: $F_a \approx 0.9 \times 2 \times 525 \div 2 \approx 470$ A-turns per pole, and $\Delta H \approx 470 \div (2.0 + 1.05 \times 0.8) \text{ mm} \approx 165$ kA/m. Note that rewinding changes nothing here: 2 turns at 263 A is the same ampere-turns. The worst-case field at the magnet is then

$$H_{\text{worst}} = |H_m| + \Delta H \approx 280 + 165 = 445 \text{ kA/m}$$

Where timing puts it. With zero timing (pure i_q) the armature field crosses the magnet sideways: it strengthens the leading edge and weakens the trailing edge, so the trailing edge of each magnet sees roughly the full ΔH while the middle sees little. That is why partially demagnetized motors lose flux at one edge of every pole first. Timing advance rotates the current toward negative i_d , and the component $\sin \gamma \times \Delta H$ now pushes straight through the whole magnet. Six-step drive swings the real angle $\pm 30^\circ$ around the set timing, so the design case is $\gamma_{\text{set}} + 30^\circ$. At 30° of Turbo on top of 15° base timing the entire magnet sees about $0.7 \times \Delta H$ plus the edge effect. The practical rule: the trailing edge is always at H_{worst} ; advance moves the rest of the magnet toward it.

Temperature takes the margin away

Both B_r and the intrinsic coercivity H_{cj} fall with temperature, and H_{cj} falls about five times faster:

$$B_r(T) = B_r(20) [1 - 0.0012 (T - 20)], \quad H_{cj}(T) = H_{cj}(20) [1 + \beta (T - 20)]$$

with $\beta \approx -0.0055 / ^\circ\text{C}$ for N and M grades and about $-0.0045 / ^\circ\text{C}$ for SH, UH and EH (the high-coercivity grades are less temperature-sensitive as well as starting higher). The usable limit is not H_{cj} itself but the knee H_k where the demagnetization curve bends, typically $0.85\text{--}0.9 \times H_{cj}$ for a well-made grade; use the datasheet H_k when it is given.

Typical catalog minimums at 20°C (confirm the grade actually in the motor):

Grade suffix	Max working temp	H _{cj} at 20 °C	H _k at 80 °C (est.)	H _k at 120 °C (est.)	H _k at 150 °C (est.)
N (none)	80 °C	955 kA/m (12 kOe)	545	360	230
M	100 °C	1,115 kA/m (14 kOe)	635	420	270
H	120 °C	1,355 kA/m (17 kOe)	805	575	405
SH	150 °C	1,590 kA/m (20 kOe)	985	745	560
UH	180 °C	1,990 kA/m (25 kOe)	1,235	930	700
EH	200 °C	2,390 kA/m (30 kOe)	1,485	1,115	845

H_k estimates use $0.85 \times H_{cj}(T)$ with $\beta = -0.0055 / ^\circ\text{C}$ for N and M, -0.005 for H, and -0.0045 for SH, UH and EH. Values in kA/m.

The rule

$$|H_m| + \Delta H_{\text{launch}} \leq \frac{H_k(T_{\text{magnet,max}})}{1.2}$$

The 1.2 is margin for the six-step angle swing, local field concentration at magnet edges and slot openings, and datasheet-minimum scatter. Reading the table against the Chapter 7 motor's 445 kA/m:

- **N grade** fails at 120 °C (360 kA/m) and is marginal even at 80 °C ($545 \div 1.2 = 454$). Any N-grade magnet that has seen a hot session at launch current has probably lost flux.
- **H grade** holds to 120 °C with $575 \div 1.2 = 479$ against 445: a thin margin that disappears with Turbo advance.
- **SH grade** holds to 150 °C ($560 \div 1.2 = 467$) and is the minimum for a drag motor that will see timing advance.
- **UH grade** holds to 150 °C with $700 \div 1.2 = 583$ against 445, a $1.3\times$ margin on top of the 1.2. This is the choice in the Chapter 7 design table, and this is why.

Two design levers besides grade: thicker magnets raise P_c (less negative H_m) and lengthen the path (smaller ΔH per ampere-turn), so going from 2.0 to 3.0 mm cuts H_{worst} by roughly a quarter at the cost of rotor inertia; and a Halbach array lowers the

stator's reverse field reaching the magnet body because the flux is steered outward rather than driven through it.

How hot the magnet actually gets

In a single 2 s pass the magnets barely warm: their eddy loss (Chapter 3 sidebar) is a few watts against a rotor mass of 30–50 g, so the rise is under 5 °C. The magnet temperature that matters is the session temperature, set by heat conducted and radiated from the stator across the airgap plus the rotor's own iron and eddy loss at high electrical frequency. Expect the magnets to run 10–30 °C below the stator iron in a drag motor and close to the iron in a long speed run. The two-node model in Chapter 1 gives the iron temperature; take the magnet at iron minus 20 °C for drag and iron for speed runs until measured.

A motor temperature sensor on the ESC reads the can or end bell, which lags the iron by another 10–20 °C during a session. Treat a 100 °C can reading as 120 °C iron and about 100 °C magnets.

What partial demagnetization looks like

- Back-EMF at a fixed rpm drops; kV reads higher than it used to.
- Torque per amp drops in the same proportion, so the car needs more current for the same launch and the copper budget shrinks.
- Sensorless ESCs cog or lose sync at launch because the flux they listen for is weaker.
- The loss is uneven: trailing edges first, so back-EMF becomes less sinusoidal and the delta loop current (Chapter 8) rises.

It is never recovered by cooling. Remagnetization needs a pulse magnetizer with the magnets in place.

Does reversing rotation help?

A reasonable thought, once you know the trailing edge takes the damage: run the motor backwards for a while, reverse it at the transmitter, and put the fresh edge to work. The direction dependence behind the idea is real. Making torque requires the stator field to lead the magnet, which strengthens the leading edge and weakens the trailing one, so reversing the rotation does swap which edge is attacked.

It still does not help, for three reasons.

Demagnetization is not wear. It is a one-time, self-limiting loss. A magnet driven past its knee drops onto a new, lower recoil line and stays there; run the identical pass again and nothing further happens, because the operating point no longer crosses the knee. More is lost only when the field or the temperature exceeds what that magnet has already seen. So the damaged edge is not getting worse pass after pass, and it is not waiting to be rested. Reversing hands the same hit to the one edge that was still intact, and the first edge does not come back. The motor ends with two weak edges instead of one.

Both edges always carry flux. Torque and kV come from the pole's whole flux linkage, not from whichever edge happens to be leading. A weakened trailing edge lowers torque per amp in both directions; there is no orientation in which the strong edge works alone.

Timing advance is not edge-selective anyway. The cross-magnetizing part of armature reaction is what loads one edge. The $\sin \gamma$ component that timing advance introduces pushes straight through the magnet body, uniformly, and does not care which way the rotor turns. On a setup running real turbo advance that uniform component dominates, so reversing changes nothing at all.

One related effect worth knowing: under braking the stator field lags rather than leads, so regen stresses the opposite edge from launch. In drag racing the braking currents are far smaller than the launch currents, so this is a rounding error, but a car that launches hard and brakes hard does work both edges a little.

The test this idea does produce. Edge demagnetization makes the flux around the rotor asymmetric, and that asymmetry has a direction. Spin the motor with a drill in one direction and capture the line-to-line back-EMF waveform on a scope, then spin it the other way and capture it again. A healthy motor looks the same both ways. An edge-demagnetized motor shows a skewed waveform whose skew mirrors when the direction reverses. This separates the two failure modes: a skew that mirrors is armature-reaction damage at the pole edges, while an even drop in amplitude with the waveform shape unchanged points at heat or timing advance having taken the whole magnet down. Amplitude alone cannot tell those apart.

Bench test

1. **Baseline:** spin the motor with a drill at a fixed, measured rpm and record peak line-to-line voltage. Record winding temperature.
2. **Expose:** run the launch current for the pass length at the intended motor temperature (preheat in an oven or by repeated passes), with the timing you race.

3. **Repeat the baseline** after the motor cools. A drop in volts per rpm is flux loss, one for one. More than 2% is demagnetization; under 1% is measurement noise.
4. Repeat at the next higher temperature or timing step until the drop appears. That temperature and current pair is the real limit for that motor and grade; put it on the motor's spec sheet.

Limits of this model. F_a and ΔH here are per-pole averages; finite-element work shows edge fields 1.3–1.5× higher at slot openings, which the 1.2 margin only partly covers. μ_r and β vary by maker. Hybrid-bonded and plastic-bonded magnets have far lower H_{cj} and are outside this chapter.

Chapter 3 sidebar: segmenting the magnets

Cutting a 100 mm × 5–8 mm magnet into four 25 mm pieces cuts its eddy loss by only 10–16%. Large cuts need segments about as long as the magnet is wide: 16 pieces of 6.2 mm cut loss by 48–67%.

Why magnets carry eddy currents

Sintered NdFeB conducts, at roughly $1.5 \mu\Omega \cdot \text{m}$. The magnets see flux ripple that does not turn with the rotor: stator slot harmonics, six-step current harmonics and PWM ripple. That ripple drives current loops inside each magnet, running along its length and across its width.

Why width sets the limit

For a thin magnet of width a and length b in uniform ripple, the loss per unit volume is the conductivity times $(dB/dt)^2$ times a shape factor:

$$f(a,b) = \sum_{m,n \in \mathbb{N}} \frac{64}{\pi^4 m^2 n^2} \left[\left(\frac{m\pi}{a} \right)^2 + \left(\frac{n\pi}{b} \right)^2 \right]$$

For a long strip this approaches $a^2/12$, set by width alone. A 5–8 mm × 100 mm strip is already close to that limit, which is why short cuts along the length barely help. Each segment needs to be roughly square before the factor falls steeply.

Assumptions: the magnet is thin (about 1–1.2 mm), the ripple is uniform across it, and the currents are resistance-limited. The figures include the length lost to 0.003 in (0.076 mm) spacers.

Where the model stops holding

At high frequency, currents crowd toward the magnet surface. NdFeB skin depth is about 7.8 mm at 6 kHz and 4.3 mm at 20 kHz. Slot harmonics in a typical motor sit in the resistance-limited range, so the numbers hold for them. PWM ripple at 20 kHz and above is near the skin depth, so expect smaller gains there.

Spacer material

Carbon fiber conducts, so fibers can bridge the gap and partly reconnect the segments. Use an insulator:

- G10/FR4 fiberglass shim
- Kapton film
- An epoxy bond line, if the magnets have an epoxy coating

Nickel-plated magnets conduct at their surface and need a true insulating layer between segments.


Where it pays off

In a 2 second drag pass, magnet eddy heating is modest because the magnets have a lot of mass. The gain is larger for speed runs, where electrical frequency is highest, and for back-to-back passes, where heat builds. Cooler magnets also keep more demagnetization margin.

Loss cut vs segment count

Four segments cut eddy loss only 10-16%; about 16 cut it in half

Eddy loss cut vs one solid 100 mm magnet, about 1.2 mm thick




Resistance-limited thin-plate model run for this book · 0.076 mm spacers included

The 8 mm magnet gains faster because its segments turn square sooner; a 5 mm magnet needs about 20 segments to match 12 on an 8 mm one.

The segmented 4-pole rotor

Each 25 mm ring carries all four poles; the cuts run across the length



segmented 4-pole rotor · end view and side view, spacers not to scale

The poles run around the rotor, not along it. Cutting the magnets lengthwise makes four identical rings, each with the full N, S, N, S pattern.

Chapter 4 note: speed limits, with numbers

For RC-size rotors the order in which speed limits arrive is: voltage runs out, then iron loss and bearings, and only far later the sleeve. The full chapter will add bearing and iron-loss data; the burst calculation can be settled now.

Sleeve hoop stress

A ring of magnets of density ρ_m and thickness t_m at radius r , spinning at ω , pushes outward on its sleeve with pressure p ; the sleeve of thickness t_s carries that as hoop stress plus its own inertia:

$$p = \rho_m t_m \omega^2 r, \quad \sigma_{\text{hoop}} = \frac{p r}{t_s} + \rho_s \omega^2 r^2$$

Chapter 7 rotor: $r = 10\text{ mm}$, 2.0 mm NdFeB at $7,500\text{ kg/m}^3$, $0.3\text{ mm carbon sleeve}$ at $1,600\text{ kg/m}^3$.

Speed	Tip speed	Magnet pressure on sleeve	Sleeve hoop stress	Electrical frequency (4-pole)
60,000 rpm	63 m/s	5.9 MPa	about 205 MPa	2.0 kHz
100,000 rpm	105 m/s	16 MPa	about 565 MPa	3.3 kHz

Hoop-wound carbon fiber holds 1,500–2,500 MPa, so the sleeve has a 7× margin at 60,000 rpm and still about 3× at 100,000. The practical sleeve rule is not strength but preload: the sleeve must keep the magnets in compression at top speed so the adhesive never carries tension, which means an interference fit sized for p at the top rpm plus thermal growth. A sleeve that goes slack lets the magnets shift, and a shifted magnet cracks on the next launch.

What actually limits speed on these motors

- **Voltage.** Top rpm = kV × sagged volts × loaded factor. For most setups this binds first; see the EVUL Gearing Calculator.

- **Iron loss.** Stator core loss grows between $f^{1.5}$ and f^2 . At 60,000 rpm a 4-pole stator runs at 2 kHz, five times the 400 Hz the lamination data is quoted at, so expect 8–20× the 1.0 T loss figure. On a speed run this heat comes out of the same stator the copper is dumping into, which is why thinner laminations are on the speed-run build list.
- **Bearings.** Hobby 8 mm bore bearings run dN (bore mm × rpm) of about 480,000 at 60,000 rpm, at the top of the grease-lubricated range; above that, hybrid ceramic balls and light oil, and expect bearing heat to show on the can.
- **Six-step commutation time.** At 2 kHz electrical, each 60° step lasts 83 μ s against an L/R current rise of similar length, so the phase current cannot reach its target within a step. That is the regime where Boost/Turbo advance earns its keep, and where it costs magnet margin (Chapter 3).
- **Rotor burst.** Last on the list for anything under about 120,000 rpm on a 20 mm rotor with a proper sleeve.

Chapter 5 note: timing and pole count

Pole count decides *when* in a pass timing advance starts doing anything, because everything the ESC times is measured in electrical degrees and a 6-pole rotor runs three electrical cycles per turn where a 2-pole runs one.

What XC-ESC actually exposes

The E6, E6 Pro and E6MAX are tuned through the XC-Link app over Bluetooth; the same values appear as X.BUS registers, readable over the bus but not writable, so an outside controller can log the setting but not change it mid-pass.

Setting	Range on the E6 family	X.BUS register	Notes
Turbo timing	0–26° electrical (E6 and E6MAX product pages; the E6MAX page quotes the same 0–26°)	app only	XC claims about 20% more rpm at full turbo on a road speed run
Turbo trigger delay	0–1000 ms	app only	How long after the trigger before the advance is applied

Setting	Range on the E6 family	X.BUS register	Notes
Start power	app scale	app only	Holds launch duty down, which is what makes phase current exceed battery current
Punch	0–7	app only	Throttle-to-duty rate at launch
Drive (PWM) frequency	XC's race guide: 8 kHz modified, 12–16 kHz stock	—	Not a timing setting, but it sets ripple and switching loss
Drag brake	5–10% outdoor, 10–15% indoor (XC race guide)	—	—

XC's own race-programming guide gives starting points of 25–30° turbo for modified and 30–38° for stock, which exceeds the 0–26° the E6 product pages quote; the guide is written across their ESC range, so confirm the ceiling on the unit in hand. **Every one of those degrees is electrical.**

Sensored and sensorless: the same degrees, different confidence

Both kinds of controller advance timing, and both measure it in electrical degrees, so everything in this chapter applies to either. What differs is how the ESC knows where the rotor is, and how well that holds up at launch.

A sensored ESC reads position directly from hall sensors, which give an edge every 60 electrical degrees. Advance is a shift of the switching point relative to that edge, and the position is known even at a standstill.

A sensorless ESC has no sensors. It watches the back-EMF on whichever of the three phases is not being driven and looks for its zero crossing, which falls 30 electrical degrees before the ideal commutation point. The controller then waits a delay computed from the previous commutation interval and switches. Advance means shortening that delay.

Three things follow, and all of them bite hardest in the first half second of a pass:

- **There is no back-EMF at zero rpm.** Below a few thousand rpm a sensorless ESC runs open loop: it aligns the rotor to a known position, then ramps blind commutation until the signal appears. That is exactly the launch.

- **The delay is a prediction from the last interval.** Under hard acceleration the rpm rises substantially within a single commutation, so the predicted delay is stale and the switch lands late. The timing effectively retards itself in proportion to acceleration, and the error is worst at low rpm and high g — a 132 ft launch is the worst case in the hobby.
- **High current blinds the detector.** PWM switching noise and freewheeling spikes swamp the zero crossing, so the controller blanks a window after each switch. At the currents in Chapter 1's heat budget that window has to be generous, which costs timing resolution.

Demagnetization makes all three worse (Chapter 3): weaker magnets mean a weaker signal to find, so a tired motor loses sync before it loses much torque.

One asymmetry worth noting: a sensorless ESC does not need the pole count at all, because commutation is purely electrical — it needs that number only to display mechanical rpm. A sensed motor's hall sensors are physically placed for one pole count, so sensors and magnets have to match.

For 132 ft drag and truck pulling the launch is the whole event, which is why these classes run sensed. Every controller this book covers — the XC-ESC E6, E6 Pro and E6MAX, the Castle Cobra 8 and Cobra 5 HV, and the Ghost RC ACE versions built on them — is a sensed six-step unit. A sensorless setup can still be timed, and the degree math above is unchanged, but the advance it delivers at launch is the least trustworthy part of the pass.

The same degree is a different angle on every rotor

$$\theta_{\text{mech}} = \frac{\theta_{\text{elec}}}{p}, \quad f_e = \frac{\text{rpm}}{60}, \quad t_{\text{per elec degree}} = \frac{1}{360 f_e}$$

with p = pole pairs: 1 for a 2-pole, 2 for a 4-pole, 3 for a 6-pole.

Quantity	2-pole	4-pole	6-pole
Electrical cycles per turn	1	2	3
26° of turbo, in rotor degrees	26.0°	13.0°	8.7°
One hall step (60° elec), in rotor degrees	60.0°	30.0°	20.0°
Electrical frequency at 30,000 rpm	500 Hz	1,000 Hz	1,500 Hz
Time per electrical degree at 30,000 rpm	5.6 μs	2.8 μs	1.9 μs

Quantity	2-pole	4-pole	6-pole
Time the full 26° is worth at 30,000 rpm	144 μ s	72 μ s	48 μ s


Two consequences. First, the same turbo number is a much smaller physical rotation on a 6-pole, so sensor placement error and hall jitter cost proportionally more: a 2° mechanical sensor misalignment is 2° of electrical error on a 2-pole and 6° on a 6-pole. Second, the 6-pole reaches any given electrical frequency at a third of the rpm, so everything that depends on frequency — inductive lag, iron loss, the need for advance — arrives three times earlier in the pass.

Why the advance has to come in earlier on a 6-pole

A 6-pole runs out of bus volts at a third of the rpm a 2-pole does

Volts the winding inductance alone demands at 400 A · 10 microhenry line-to-line

Volts



model run for this book · 10 μ H line-to-line, 400 A, 25.2 V bus · $V = \omega_e L I$

Inductance resists current change, and the voltage it demands grows with electrical frequency: $V_L = \omega_e L I$. Once that demand approaches the pack voltage, the controller can no longer drive the commanded current within a step, torque per amp falls, and advance is what buys it back by starting the current earlier. With a 10 μ H winding at 400 A on a 6S pack, the 6-pole crosses that line at about 20,100 rpm, the 4-pole at 30,100 and the 2-pole not until 60,200. That is the whole reason a 6-pole wants its timing in early and a 2-pole can stay near zero advance for most of a 132 ft pass.

The same ratio shows up in the Surpass F-series 2-pole motors covered elsewhere in this book: their progressive feel is partly pole count and partly that they simply are not frequency-limited until far higher rpm.

How fast you can come up on timing

Turbo trigger delay spreads the advance over a window, so the ramp rate is the range divided by the delay.

Turbo delay	Ramp rate for 26°	What it means on a 2 s pass
0 ms	step change	All the advance at the trigger: hardest on magnets, most likely to break traction
200 ms	130 °/s	Advance complete in the first tenth of the run
500 ms	52 °/s	Advance complete about a quarter of the way down
1000 ms	26 °/s	Advance still arriving at half-track

The useful rule is to match the ramp to the rpm where the motor actually needs it, from the chart above:

- **6-pole:** needs advance from roughly 20,000 rpm. On a 2 s pass a 6-pole drag motor is usually past that within the first half second, so a short delay (0–200 ms) and most of the range is right.
- **4-pole:** needs it from roughly 30,000 rpm, typically reached a third to half way down. A 300–500 ms delay puts the advance where the frequency problem starts.
- **2-pole:** often never reaches its limit in 132 ft. Advance here buys top-end rpm, not launch, so a long delay (500–1000 ms) or a low turbo setting keeps the launch clean and the magnets cool.

What the advance costs

Advance is not free, and the bill is in Chapter 3. Rotating the current toward negative i_d puts the component $\sin \gamma$ of the stator field straight through the magnet instead of across it. At 26° of turbo with the six-step angle swing of $\pm 30^\circ$, the design case is 56°, where $\sin \gamma = 0.83$: the magnet sees about 83% of the full armature reaction as a direct demagnetizing field rather than the edge-loaded case at zero timing. For the Chapter 7 example that

moves the magnet from roughly 445 kA/m to near 500 kA/m, which is what pushes an H-grade magnet past its margin at 120 °C and makes SH or UH the sensible grade for any motor that runs turbo.

A 6-pole makes this worse in one specific way: because it needs the advance earlier in the pass, the magnet sees the demagnetizing field while the copper is still dumping its launch heat, so the peak field and the peak temperature arrive closer together.

Bench test

- 1. Log a pass with X.BUS at 20–100 Hz: phase current, rpm, throttle and the timing registers together.
- 2. Plot rpm against time and mark where it crosses the frequency limit from the chart for your pole count.
- 3. Compare that moment with where the logged turbo advance actually arrives. If advance arrives well after the crossing, the motor spent that stretch making less torque per amp than it could; if it arrives well before, it was weakening field and heating magnets for nothing.
- 4. Measure back-EMF before and after a session at the chosen setting (Chapter 3 bench test) to confirm the magnets held.

Limits of this model. The 10 μ H and 400 A are placeholders until the reference motor is measured; the crossing rpm scales inversely with both. The model uses line-to-line inductance and ignores saliency and saturation of the inductance itself, which falls at high current and pushes the crossings higher.

Chapter 7 draft: the clean-sheet burst motor

A burst motor is designed around its copper heat budget and its saturation knee, not a continuous rating. In a first-pass 40 × 120 mm 4-pole design, the copper allows about 3.3 N·m for a 2 second pass, but the leading stator teeth saturate before that current is reached.

Design choices

Part	Choice	Why
Winding	Highest slot fill; compressed or rectangular wire	Copper mass is the heat budget

Part	Choice	Why
End connections	Bus bars instead of end-turn loops	Shorter, lower-resistance connections and room for a longer stack
Magnet wire	220 °C enamel	About 50% more temperature rise to spend than 155 °C wire
Slot insulation	Epoxy powder coat	Thinner than paper liners, so more room for copper
Teeth and yoke	Sized so the 1.6–1.7 T knee sits just above launch current	Wider steals copper; narrower saturates early
Laminations	0.20 mm silicon steel, bonded rather than interlocked	Interlocks join sheets electrically and add stator eddy loss
Magnets	SH or UH grade, about 2 mm thick	Demagnetization margin over peak strength; less sensitive to airgap
Magnet layout	Halbach array (stretch option)	Stronger, smoother gap flux with less back iron and inertia
Retention	Carbon-fiber sleeve	Strong, and conducts too little to matter as a thin wrap
Rotor shape	Long and thin	Torque grows with D^2L ; inertia grows with D^4L
Drive	Sensored six-step, near-zero timing at launch, Boost/Turbo advance ramping with rpm	Launch current goes to torque, not field weakening
Pre-pass	Start cold	Every degree below ambient adds to the heat budget

Worked example: 40×120 mm, 4-pole, 12-slot (first pass)

Assumed geometry: stator OD 37 mm, bore 21 mm, stack 92 mm, airgap 0.5 mm, sleeve 0.3 mm, magnets 2.0 mm at Br 1.25 T, teeth 3.4 mm wide, yoke 4.0 mm, slot fill 45%, bus-bar end connections adding 10% copper. Every value is a placeholder until a reference motor is measured.

Result	Value
Airgap flux density (peak)	0.88 T
Tooth flux from magnets alone	1.42 T
Yoke flux	1.47 T
Copper mass	about 53 g
Heat budget at 150 °C rise	about 3,000 J, or 1.5 kW for 2 s
Ampere-conductors per slot at the 2 s limit	about 525 A
Heat-limited torque, 2 s pass	about 3.3 N·m
Leading-tooth flux at that current	about 1.8 T, past the knee
Rotor inertia	about 9.9 kg·mm ²

The turns count only picks the voltage. Heat-limited torque comes out the same for any number of turns; the winding just trades kV against current.

Conductors per slot	kV (approx.)	Phase current at the 2 s limit
1	1,545	525 A
2	772	263 A
3	515	175 A
4	386	131 A

What the example says.

- Saturation comes before heat at 2 seconds. Torque per amp will roll off at launch unless the teeth get wider or the magnets weaker.
- A 4-pole motor on a 37 mm stator is iron-heavy: the yoke needs about 4 mm to carry half a pole's flux, leaving only about 53 g of copper. A 6-pole layout halves the yoke flux and frees room for copper, at the cost of 50% higher electrical frequency.
- Bus bars matter mainly through length: every millimeter saved at the ends can go into stack, and torque and heat budget both grow with stack length.


Three builds on one platform

- **Drag:** 4-pole, long and thin for low inertia, teeth sized for launch current.
- **Speed runs:** 4-pole, thinnest laminations, strongest sleeve, voltage headroom for top rpm; pre-laminated magnets worth pricing here.
- **Truck pulling:** 6–8 poles for torque at low rpm and the most copper, since pulls run longest.

Benchmark motors to measure: Tenshock EVO (bus-bar stator), and Neu Motors as an American-made reference.


The clean-sheet rotor

The clean-sheet rotor: Halbach magnets, carbon sleeve, light hollow hub



End view

Arrows show each segment's magnetization



Side view

Long and thin: torque grows with D^2L , inertia with D^4L

clean-sheet rotor · 4-pole Halbach end view and side cutaway, not to scale

The Halbach segments turn their magnetization around the rotor, so flux concentrates outward toward the stator and the hub needs little iron.

Chapter 8 draft: windings — wye, delta, turns and paths

The winding sets the motor's voltage-to-current ratio, not its torque capability. Wye vs delta, turns per coil and parallel paths all move kV up or down; the heat-limited torque from Chapter 1 stays the same. Choose the winding last, to match the battery, the ESC and the gearing.

Wye vs delta with the same coils

Quantity	Wye (star)	Delta
Voltage across each coil	Line voltage $\div$ 1.73	Full line voltage
Current in each coil	Equal to line current	Line current $\div$ 1.73
kV	1×	1.73×
Torque per amp of line current	1×	$1 \div 1.73 = 0.58\times$
Line-to-line resistance	$2 \times$ coil-phase R	$\frac{2}{3} \times$ coil-phase R
Heat-limited torque	Same	Same
Circulating current at no load	None possible	Possible (third harmonic)

Switching wye to delta is electrically the same as rewinding with 1.73× fewer turns. Both ends of the torque-speed line move out by 1.73×: no-load speed and stall torque at a given voltage.

The delta loop problem

A delta winding is a closed loop. If the back-EMF contains a third harmonic, and non-sinusoidal RC motors usually do, the three third-harmonic voltages add up around the loop instead of canceling. They drive a current around the triangle that the ESC never sees, heating the copper even at no load. Wye has no loop, so that current has nowhere to flow.

Rule: default to wye. Use delta only to reach a kV that turns and paths cannot, and test the loop first (shop tools below).

The three knobs

```
kV \approx kV_1 \times \frac{a}{N} \times \begin{cases} 1 & \text{wye} \\ 1.73 & \text{delta} \end{cases}
```


N is turns per coil, a is the number of parallel paths, and kV₁ is the one-turn, one-path, wye value for the stator. For the Chapter 7 example, kV₁ is about 772. Parallel strands do not change kV; they only fill the slot.

What the ESC sees

At the same ESC current limit, a higher-kV winding gives less launch torque and more top speed. Peak power stays close, because power is torque times speed. The winding decides where along the track that power arrives.

Wye and delta at a glance

Same three coils: delta gives 1.73× the kV, wye gives 1.73× the torque per amp




Wye (star)

$$\text{kV} \times 1$$

Each coil sees line volts $\div 1.73$

Coil current = line current



Delta

$$\text{kV} \times 1.73$$

Each coil sees full line volts

Coil current = line current $\div 1.73$


wye and delta connections of the same three coils

In wye, two coils in series share the line voltage; in delta, each coil takes the full line voltage and the triangle forms a closed loop.

Torque vs speed, wye vs delta

Same 400 A limit: wye launches with 1.73× the torque, delta reaches 1.73× the rpm

Torque (N·m) vs rpm, same coils, 14.8 V, ESC capped at 400 A




Chapter 7 example motor, 1 turn per coil, 2 paths · ignores sag, inductance and saturation

Peak power is close for both. Wye suits a heavy launch such as truck pulling; delta suits a long pull to top speed, if the delta loop test passes.

The kV ladder

One 12-slot stator spans about 100 to 2,700 kV through turns, paths and wye or delta

kV (log scale) vs turns per coil, Chapter 7 example stator



$kV \approx 772 \times \text{paths} \div \text{turns}, \times 1.73 \text{ for delta}$ · Chapter 7 example stator

Below about 400 kV, wye steps are 25% or smaller. Above that they jump (515, 772 and 1,545 kV), and delta fills the gaps at 669, 892 and 1,338 kV.

Shop tools

1. Is this motor wye or delta? Resistance won't tell you: both read equal across every pair of leads. Look at the end turns instead.

- **Wye:** one star point where three coil ends are joined, often sleeved and tucked into the end turns. Each lead connects to one coil end.
- **Delta:** no star point. Each lead connects to two coil ends.

2. The delta loop test. Before trusting a delta winding at full power:

1. Unsolder one corner of the delta to open the loop.
2. Spin the rotor with a drill at a steady speed.
3. Measure the voltage across the open gap with a scope or true-RMS meter.
4. A pure-sine motor reads close to zero. Whatever you read is mostly third harmonic, and when the loop is closed it drives circulating current limited only by the coils' own resistance and inductance.
5. Compare it with the line-to-line voltage at the same speed. More than a few percent means real loop loss; stay wye.

3. Rewind conversions (same kV).

From	To	Turns per coil
Wye	Delta	$\times 1.73$ (3T wye $\approx$ 5T delta)
Delta	Wye	$\div 1.73$ (5T delta $\approx$ 3T wye)
1 path	2 paths	$\times 2$, with half the strands per path
2 paths	1 path	$\div 2$, with double the strands

4. Measure kV. Run the motor unloaded on the ESC at full throttle with zero timing advance: $kV \approx \text{rpm} \div \text{battery volts}$. Or spin it with a drill: $kV \approx \text{rpm} \div \text{peak line-to-line voltage}$.

5. Choose the winding, in order.

1. Set the top-speed target from gearing, rollout and trap speed.
2. Estimate loaded pack voltage under sag.

3. Pick $kV \approx \text{top rpm} \div (\text{loaded volts} \times 0.85\text{--}0.9)$, leaving margin for load.
4. Check launch current: $I = \text{launch torque} \div K_t$, with $K_t = 9.55 \div kV$. If it exceeds the ESC or battery limit, lower kV and gear taller.
5. Read the nearest N, paths and wye or delta off the kV ladder.

Racing ESCs: Castle Creations, Ghost RC and XC-ESC

Most drag, speed-run and pulling racers run Castle Creations controllers or Ghost RC's ACE versions built on them, and XC-ESC's E6MAX adds a live control bus. Castle's logs record battery current, so an 800 A log is battery current and the winding can see more at launch. XC's X.BUS reports phase current directly.

The controllers

Controller	Built on	Input	Timing tools	Launch tools	Logging
Ghost RC Cobra 8 ACE (1/10 no-prep drag)	Castle Cobra 8	2–4S	Boost and Turbo timing, up to 120° total, sensed	Start Power with ramp and hold time, launch throttle curve, RPM limiter	Adds motor timing and acceleration channels; 8× logging time
Ghost RC Cobra 5 HVACE (drag and speed runs)	Castle Cobra 5 HV	6–12S	Boost up to 45° (by rpm or throttle); Turbo up to 45° (by power, time-based)	Perfect Pass throttle ramping up to 25 s, Advanced Start Power, RPM limiter	Adds motor timing and acceleration channels; 8× logging time
XC-ESC E6MAX (1/6–1/8 speed, drag, bashing), 550 A	XC's own design	3–8S	Turbo 0–30° electrical with 0–1000 ms trigger delay (X.BUS register table)	Start power, punch (0–7), drag brake	XC-Link app over Bluetooth; live X.BUS telemetry including phase current

Both ACE units are sold with no warranty, with some of Castle's protection features removed or restricted for racing, and the Cobra 8 ACE requires an external capacitor pack. Neither listing gives a current rating.

What the log's current means

Castle lists [battery current](#) among its logged channels, alongside rpm, voltage, ESC and motor temperature and throttle. Motor phase current is not logged.

- At full throttle near top speed, duty is near 100%, so phase current is close to battery current.
- At launch, duty is low and phase current is roughly battery current $\div$ duty. Start Power and throttle ramps hold duty down on purpose, so the winding's peak current can exceed the logged peak.
- The heat budget in Chapter 1 runs on phase current, so treat a logged 800 A as the minimum the copper saw.

What 800 A demands of the copper

For 800 A to survive a 2 second pass at a 150 °C rise (109 A/mm²), each phase lead needs about 7.3 mm² of copper. With 100 strands, that is about 0.30 mm per strand. A thinner lead would melt within the pass, so a motor that logs 800 A and survives has at least that much copper at its leads.

Timing at these settings

Boost and Turbo stack on top of base timing as rpm or power rises, which matches the Chapter 5 rule: little advance at launch, more as speed builds. At high totals, most of the stator current goes into field weakening rather than torque ($\sin \gamma$ against $\cos \gamma$), and that current also pushes the magnets toward demagnetization. Logged timing and rpm together let a racer see exactly where the advance comes in on a pass.

Sources: [Ghost RC Cobra 8 ACE listing](#), [Ghost RC Cobra 5 HV ACE listing](#), [Castle Creations technology page](#).

XC-ESC E6MAX and X.BUS

The E6MAX is rated 550 A continuous for 3–8S, with 18 MOSFETs, active freewheeling, a 25 A BEC (6, 7.4 or 8.4 V), 7 AWG power and motor wires, IP67 sealing and a weight of about 510 g. It is set up and updated through the XC-Link app over Bluetooth. Its X.BUS port lets an outside controller command it and read live telemetry during a pass.

Note for this public draft: the X.BUS frame format, register table and startup handshake come from XC-ESC's partner documentation and are withheld here pending confirmation from XC. What follows describes what the bus makes possible, without reproducing the protocol. If you are building against X.BUS, ask XC-ESC for the current specification.

Phase current is the key channel. It is the quantity the Chapter 1 heat budget runs on, and Castle's logs do not record it.

What a bus controller can do with it

- **Heat-budget governor:** track winding temperature rise from live phase current with the Chapter 1 model, and taper throttle before the copper reaches its limit.
- **Saturation-aware launch:** cap launch current near the saturation knee, where extra amps stop making torque.
- **Traction control:** compare motor speed with wheel speed and trim throttle on wheelspin.
- **Phase-current logging:** record what the winding actually sees, at 20–100 samples a second.

Live timing changes need a firmware build from XC that makes the timing registers writable; that request is with XC. A Python bench implementation of the bus, with the startup handshake and the heat-budget governor, is in the project files as `xbus.py`.

Sources: [XC-ESC E-MAX series](#), [E6MAX parameters](#), [Red RC on the E6MAX](#), XC-ESC X.BUS Bus Control Protocol V1.0.0-251203 (partner document).

How the ESC buses compare

X.BUS is the only built-in multi-ESC bus on a car ESC here, and the only one that reports phase current. That makes it the bus to build a multi-motor controller and heat-budget governor around.

Maker and link	Multi-ESC bus	Control over the bus	Current reported
XC-ESC X.BUS	Built in: one wire, up to 16 ESCs	Throttle and brake	Battery and phase
Castle Link Live	None: one ESC per throttle wire	Throttle (the normal signal)	Battery
Castle Serial Link	Add-on adapter per ESC; up to 64 on TTL serial, I2C or SPI	Throttle, failsafe, emergency stop	Battery (2015 manual lists a 250 A maximum)
Hobbywing car ESCs	None found	None	Battery, one way to the radio

Maker and link	Multi-ESC bus	Control over the bus	Current reported
Hobbywing Platinum (aircraft)	None: one-way serial on the programming wire	None	Battery
Hobbywing XRotor CAN (drone)	DroneCAN, CAN at 1 Mbit	Throttle	Not stated

Castle Serial Link in brief. Each ESC gets its own adapter. Commands are 5 bytes (address, register, 16-bit value, checksum); replies are 3 bytes; the checksum makes each packet sum to zero. Readable registers cover voltage, ripple, current, throttle, output power, electrical rpm, temperature and BEC; writable ones are throttle, failsafe and emergency stop, with a watchdog that falls back to failsafe if packets stop. Values scale as $\text{register} \div 2042 \times \text{a per-channel factor (20 for volts)}$.

Hobbywing in brief. Car ESCs send telemetry one way through Hobbywing's Telemetry Adapter to Futaba, FrSky or Jeti radios. Platinum aircraft ESCs send one-way serial telemetry at 19,200 baud (V4) or 115,200 baud with CRC (V5). Only the XRotor CAN drone ESCs offer a true control bus.

Sources: [Castle Serial Link manual v1.5](#), [freeCLL \(Castle Link Live converter\)](#), [ArduPilot: Hobbywing DroneCAN ESC](#), [ExpressLRS: Hobbywing telemetry](#), [Hobbywing Telemetry Adapter](#).

Capacitors and ripple on the DC bus

At a part-throttle launch, the ESC's input capacitors carry more RMS current than the battery does. The full explanation, with the ESR curve and worked numbers, is in the companion document [Ripple, ESR and the DC Bus](#).

Where ripple comes from. Below full throttle, a six-step ESC pulses one high-side FET at its PWM frequency (typically 8–32 kHz). The battery supplies only the average current, $D \times I$. The battery leads' inductance stops the battery from following the pulses, so the capacitors supply the rest:

$$I_{\text{cap,rms}} = I_{\text{phase}} \sqrt{D(1-D)}$$

At 600 A phase current and 30% duty, the battery delivers 180 A but the caps carry 275 A RMS. At full throttle the PWM ripple disappears and only the smaller commutation ripple is left. The launch is the hard part for the caps.

ESR, not microfarads, sets most of the ripple. For four 1,000 μF low-ESR caps on 8S at that launch, the capacitance term is about 2.0 V and the ESR term about 2.6 V, together over the 10% guideline (3.4 V) that racers quote from Castle. Eight caps halve both.

ESR changes with frequency. In an aluminum electrolytic, the measured ESR includes dielectric loss, which behaves like $\tan \delta \div (2\pi fC)$. So ESR is about 12 $\times$ higher at 100 Hz than at 20 kHz, then flat through the PWM band, then rising again past self-resonance (about 40 kHz for a 1,000 μF can). The PWM frequency is a fixed setting on most ESCs, so the switching ripple stays on the flat, low-ESR floor all pass. What changes with speed is the commutation ripple: its frequency is $6 \times \text{pole pairs} \times \text{rpm} \div 60$, so it climbs from the steep, high-ESR part of the curve toward the floor as rpm rises. Temperature moves ESR more than frequency does: a cap at -10°C has about 2.5 $\times$ the ESR it has at 20°C .

Heat is a burst budget, like copper. That launch puts about 325 W into the bank. A 2 s pass adds roughly 225 J, about 11°C on 16 g of caps. The continuous ripple rating (about 2.5 A per can) means nothing for a 2 s pass; the limit is the heat the cans soak up and how many passes stack before they cool.

Smoothing, in order:

1. Short battery leads; lengthen the motor leads instead.
2. Keep + and – twisted or taped together.
3. Add about 220 μF of low-ESR capacitance per 10 cm of extra battery lead (the Retzbach rule), mounted at the ESC.
4. Several small low-ESR or polymer caps rather than one big one.
5. Leave voltage margin above a fully charged pack.
6. Check cap temperature after passes and log ESR when the caps are new.

The X.BUS advantage again. X.BUS reports battery current and phase current together, so $\text{duty} \approx \text{battery} \div \text{phase}$, and the cap ripple current can be calculated from every logged pass.

Open questions and data needed

- ☐ Which reference motors: 1/8-scale 4-pole inrunners (4274 / 4268 class), larger truck-pull motors, or both?

- ☐ Real data from one raced motor: kV, phase resistance, weight, can size, turns, strand count and wire gauge
- ☐ Magnet grade and temperature rating in the reference motor (Chapter 3 now gives the method; needs the grade and the back-EMF before/after test)
- ☐ Insulation class of the magnet wire, to set the allowed temperature rise
- ☐ Bench-check each ESC's six-step timing behavior (base timing, Boost and Turbo) on the reference motor
- ☐ Logged pass data (current, voltage, rpm over time) to check the models against
- ☐ Phase inductance (line-to-line) of the reference motor, for the current-rise and ripple figures in Chapter 1
- ☐ Copper-to-iron and iron-to-air thermal time constants from the 2 s DC pulse test in Chapter 1
- ☐ Locked-rotor torque vs current on the reference motor, to find I_{knee} for the Chapter 2 roll-off
- ☐ Duty cycle at the logged peak current on a typical Castle launch, to validate the battery-to-phase table