

A CLOSER LOOK AT THE METHODS AND LIMITS OF REGENERATIVE BRAKING

By Jeffrey Jenkins

All motors are generators, the old saying goes, with the caveat that few types are equally good at both modes of operation, and some are downright awful as generators (a classic example being the single-phase shaded pole induction type, used in myriad small appliances the world over). Fortunately, both types of AC motor commonly used in EVs - the permanent magnet synchronous and the induction asynchronous - work perfectly fine in generator mode, although each has its own quirks and practical limits of operation.

All that is required to turn any motor into a generator is to spin its rotor faster than it would spin on its own while field excitation is present. The very easiest motors to use as a generator are those with a permanent magnet for their field excitation - whether AC or DC - as they will obviously have field excitation

present at all times; just spin the shaft and connect an appropriate load and you have a generator (note that the PM AC motor produces 3-phase AC with both voltage and frequency proportional to RPM).

Slightly more difficult to use as generators are motors with an externally-excited field - such as the wound-field synchronous AC or separately-excited DC - as they require a DC power source to supply their field windings (though often there will be enough residual magnetism in the field structure to bootstrap generator operation). This does allow for a degree of control over the magnitude of the generated voltage at a given shaft RPM, however. The asynchronous induction motor is a bit more difficult still to use as a generator, mainly because its field is excited indirectly, so it needs a source of 3-phase AC present before it can generate; this issue is easily overcome as long as an inverter



Audi e-tron prototype regenerative braking testing


Photo courtesy of Audi

drives the motor, which is always the case in an EV.

Then there are the motors that rarely get used as generators, either because they are too difficult, such as the series DC motor, or too inefficient, such as the hysteresis and shaded pole AC motors, or just not terribly widespread (yet) such as the switched reluctance motor.

To understand how regen braking works, and what its practical limits are, a brief recap of how motors function may be helpful. The part of the motor called the field projects a fixed-intensity magnetic field which interacts with a rotating/variable intensity magnetic field that is projected by the armature to produce rotary motion. In AC motors the field is usually the rotor while the armature is comprised of both the stator and the inverter. If the magnetic poles of the field are locked to those of the armature the motor is said to be synchronous; when some difference between the two is allowed (or even necessary, as in the induction motor), the motor is said to be asynchronous. From the perspective of the rotor, the magnetic fields from both the armature and the field appear frozen in space in the synchronous motor, with just a few degrees of offset

Fortunately, both types of AC motor commonly used in EVs - the permanent magnet synchronous and the induction asynchronous - work perfectly fine in generator mode.



2018 Chevrolet Volt traction motor

Photo courtesy of GM



The first EVs tried to simulate the behavior of an ICE vehicle as much as possible, so as soon as the driver let off the accelerator pedal, mild regen would kick in, which felt like the familiar engine braking of an ICE.

on either side of “lock-in” separating full generating torque from full motoring torque. These fields will appear to slowly rotate with respect to each other in the asynchronous (induction) motor, however, with full motoring torque occurring at a speed several percent below synchronous speed, and full generating torque at a speed several percent higher.

Torque in both motor types used in EVs is ultimately proportional to the strength of the magnetic field from either the armature or the field, whichever is weakest. The strength of a magnetic field from a permanent magnet is fixed, of course, but it is proportional to current in an electromagnet, and since the armature is basically a collection of electromagnets arrayed in a circle about the rotor, torque can be varied by controlling the armature current. Current can't be increased arbitrarily, of course, because as one of my EE profs was all too fond of saying, “at a high enough current, everything acts like a fuse,” but long before any wires start melting another limit will kick in: magnetic saturation, or the maximum field strength a ferromagnetic material can tolerate before it abruptly loses the ability to concentrate magnetic

lines of force. Saturation reduces the torque available from the motor as well as the inductance of its windings; the latter can result in current increasing too fast for overcurrent protection to kick in. Fortunately, the magnetic circuit in most motors is designed to saturate in a progressive fashion, so the only

consequence of overpowering the armature will be a case of diminishing returns for every incremental increase in current.

The last piece of the puzzle is the inverter itself. As mentioned before, voltage and frequency in an AC motor are proportional to RPM. When the motor is operating as a motor, the current flow through the inverter hardware is such that it can only act as a buck converter - that is, step down the battery voltage to a lower RMS-value AC voltage - but when the motor is operating as a generator, current flow reverses through the inverter and the same hardware must now act as a boost converter (and rectifier). In the case of the induction motor, the inverter needs to be present to supply the field with energy before it can act as a generator, so absent any field excitation it's just a spinning lump of metal. The PM AC motor will always produce a voltage whenever its shaft is spinning, however, and if its shaft is spinning much faster than it could given the available battery pack voltage, then the inverter has to adjust the timing of its phase currents to partially suppress the PM field (aka field-weakening). Field-weakening is used in both motoring and generating modes; in the former it allows the motor to spin faster at the cost of some torque (because a weaker field means less torque, all else being equal), while in the latter it prevents destruction of the inverter from too much current flowing back into the battery.

With the basics of motor/generator operation in hand, next up is how and when to apply regen in the EV. The first EVs tried to simulate the behavior of an ICE vehicle as much as possible, so as soon as the driver let off the accelerator pedal, mild regen would kick in, which felt like the familiar engine braking of an ICE. This didn't return much energy to the battery, however, and in many such instances it would be more efficient overall to simply coast, saving regen for when stronger braking was required or when going

downhill. Another tweak employed by some OEMs is a separate knob or paddle for controlling regen intensity; this vastly improves the flexibility of deploying regen. A more natural way to employ regen is to link its operation to the brake pedal. This can be as simple as activating regen alongside the brake lights, to a more sophisticated scheme in which regen intensity is controlled by the pressure in the mechanical braking system. Besides recapturing energy that would otherwise be wasted as heat, regen can also assist with traction control, especially in the ideal case of a separate motor at each wheel, as electric motors are capable of much finer torque control than any ICE or anti-lock braking system.

There are some limits to regen braking. For starters, it can only recapture the energy used to accelerate the vehicle or climb an incline, minus some inevitable losses. Speaking of which, the energy recaptured by regen has to go through the full conversion process - from chemical to electrical to mechanical to road - twice. Typical efficiencies for each major step in this process are 99% for lithium chemistry batteries, 96-98% for inverters, 80-95% for motors (though this can drop much lower, especially at either extreme of the power range), 95% for hypoid gear differentials (which tends to get overlooked) and, finally, 85-95% for tires. Even taking the best-case values for each figure, that comes out to an overall efficiency from battery to road of 83%, and a round-trip efficiency of 69%; in other words, you can't always regen brake your way to longer range.

Another, more obvious, limit to regen braking is tire adhesion. This is less of an issue in front-wheel drive EVs, but applying too much braking torque to the rear wheels in rear- or all-wheel drive EVs (especially motorcycles) can, shall we say, make for an exciting driving experience. This is due to the phenomenon called "load transfer," in which the center of mass on any wheeled vehicle shifts due to acceleration or deceleration forces, and this shift is proportional to the height of the center of mass above ground and

The energy recaptured by regen has to go through the full conversion process - from chemical to electrical to mechanical to road - twice.

Once speed drops below a certain point, regen will fail to return enough energy to be worth the bother even if the rate of deceleration is high.

inversely proportional to the wheelbase (note that in a 4-wheel vehicle, load transfer can occur front to back, as during acceleration and braking, or side to side, as during turning). Load transfer can be exacerbated by weight transfer, which is the same effect except caused by suspension travel or actual shifting of liquids, cargo, etc, in the vehicle. Together these phenomena lower the weight over the rear wheels during braking, making them less effective at transferring force without skidding.


The maximum torque from regen is generally the same as during motoring, and while this torque can be impressive, it still isn't anywhere close to what the mechanical brakes can achieve, thus regen cannot substitute for the mechanical brakes. The equation to calculate the effective power of the brakes to decelerate a given weight at a given g-force is no different from its more familiar use in acceleration:

$$P = m * a * v$$

where P is in Watts, m is weight in kg, a is acceleration in m/s², and v is velocity in m/s. For example, the 2019 Audi e-tron Quattro has an estimated weight of 2,400 kg, and should be easily capable of braking at 0.8 g of deceleration (1 g = 9.81 m/s²); if traveling at 100 kph (62 mph, 27.78 m/s), that works out to a brake power of 523 kW. Compare that to the estimated maximum power output of its electric drivetrain of 300 kW. There is a flip side to this as well: using regen to decelerate at too low of a rate might not produce enough power to overcome the losses in the electrical components of the drivetrain. As crazy as it might seem, this would lead to regen actually decreasing range compared to mechanical braking. This is because the electrical components incur little to no loss when the motor is not active (either motoring or generating) - there are no resistive losses in the battery or wiring, no switching and conduction losses in the inverter, no iron and copper losses in the motor - but once the motor is put into operation all of those losses come

into play. Nailing down exact numbers for these losses is notoriously difficult, but a reasonable range would be 5-20 kW for just the electrical components (higher for heavier/higher-power/less-efficient vehicles). Dropping the rate of deceleration in the above example to 0.01 g would result in ~6.5 kW of power from regen, but if the combined electrical losses are, say, 8 kW, then ~1.5 kW of extra drain on the battery will occur. Similarly, once speed drops below a certain point, regen will fail to return enough energy to be worth the bother even if the rate of deceleration is high. The kinetic energy equation:


$$K_e = 0.5 * m * v^2$$

where K_e is in joules (or watt-seconds), m is weight in kg, and v is velocity in m/s, shows that energy drops off dramatically with decreasing velocity. A 2,400 kg vehicle traveling at 100 kph has 926 kJ (0.26 kW-hr) of recoverable energy, but at 10 kph that drops to a paltry 0.0026 kW-hr, or 100 times less. Even with no losses to contend with that's a mere 2.6 watt-hours of energy. Even if decelerating at a sufficient rate to produce more power than the electrical losses, that's not exactly worth writing home about. Still, regen is a net positive most of the time, and something ICE vehicles can't do at all. 




TDK Technology Driving new mobility.


TMR angle and motion sensors for high-precision demands



Hall sensors and integrated controllers for small brushless DC and stepper motors



Pressure sensor dies with long-term stability



PiezoHapt™ and PowerHap™ piezo actuators with haptic feedback for multiple waveforms

