STATIONARY BATTERY CHARGER SPECIFICATIONS DEMYSTIFIED

William K. Bennett
Vice President, Technology
HindlePower Incorporated, Easton, PA

Introduction

This presentation is about battery chargers. First, however, we can't avoid talking about some battery issues, as they relate to the interaction of the charger and battery. We'll try not to change the laws of battery physics.

Second, most of my experience is with lead-acid batteries. Since we're going to be discussing specifications from the charger's point of view, most of the material applies to both lead-acid and Nickel-Cadmium batteries. Discussions of specific gravity, of course, don't apply to Ni-Cd batteries, and you may know of other differences.

Finally, most of my experience is also with phase-controlled (SCR) chargers. Again, most of the material applies to other technologies, such as ferroresonant, magnetic amplifier and switchmode.

In a stationary dc system, there is normally a battery, a static ac-dc converter (charger) connected in parallel with the battery, and some parallel connected load (Fig. 1). In standby systems, the permanent load is small, but there are transient loads that exceed the capacity of the charger. We normally call the charger a "Float" charger, and take for granted that the charger controls the voltage of the dc bus.

To a certain extent this is true: when the battery is fully charged, or nearly so, the dc bus voltage is the same as the charger output voltage. But when the battery is lower than 80% state of charge (SOC), it's the battery that controls the bus voltage. At this point, the charger is a current source, and the battery is in control. The accuracy of the charger's float voltage becomes an issue only when the battery is fully charged.

Batteries, like everything else, have manufacturing tolerances. If the ratio of positive to negative active material deviates from the design, the full-charge specific gravity may be higher or lower than the factory specified value. This affects the open circuit voltage, and then, of course, the overvoltage necessary to maintain 100% SOC.

DC output performance specifications

So, just how tight does the charger's output voltage regulation have to be? A battery manufacturer gives a float voltage recommendation for a given cell type as 2.25VPC, with a range of 2.17 to 2.26 VPC. In a 60-cell string, that's 135V nominal, and a range of 130.2 to 135.6, giving a tolerance of ±2%. It would seem that it would be all right for a charger to hold 1% regulation. But let's look a little closer.

If the fully charged cell has a specific gravity of 1.215, its open circuit voltage is (empirically) 2.055V. The float voltage needs to be about 10% higher to compensate for self-discharge, because of the decreasing efficiency of the charging reaction near full charge. If it is too low, self-discharge occurs, and can lead to loss of available active material. If it's too high, the excess Ampere-hours go into electrolysis, leading to electrolyte loss and possible increased corrosion.

This is where a knowledgeable user can help the situation, by monitoring water usage and float current. The float current has to be adjusted for battery temperature (or use temperature compensation – more about that later). It's possible to customize the float voltage for a battery to optimize the trade-offs between state of charge and water usage. Once this is done, the charger needs to maintain that voltage; expect a regulation tolerance of ±0.5% or better. Fortunately, almost all modern stationary chargers meet or exceed that requirement.

Temperature compensation

I can't emphasize this enough. Every stationary battery, whether flooded or sealed, of any type, benefits from a temperature compensated float voltage. Fortunately, again, most modern chargers make this available as an option. I've seen it listed as standard on some charger specs, but I think the user needs the flexibility to decide whether it should be active in any given installation.

Temperature compensation reduces the charger float voltage as temperature increases. So how does this square with keeping ½% regulation? The trick is to set the charger float voltage when the battery (or wherever the temperature probe is located) is exactly at 25°C (77°F). Then you use the manufacturer's recommended value for float voltage.

Now, the battery is at 25°C once a year, on May 12 in the northern hemisphere (possibly more frequently in California). You could put the probe into a water bath, but it's easier to use a curve like the one in Fig. 2. This figure shows compensation curves for lead-acid and Ni-Cd batteries. There are slope limits at very low and very high temperatures. At the low end, we limit the rise in float voltage so that we don't cause false HVDC alarms; at the high end, the limit prevents the float voltage from going below the open circuit voltage. Of course, if your battery temperature approaches either of these limits, you have a lot of other problems.

Even when it's an extra-cost option, temperature compensation pays dividends in maximizing battery life and usable capacity. It's a must for VRLA and other sealed designs.

Now that you've decided to use temperature compensation, where do you mount the probe? The obvious best location is right on the battery. But at least one manufacturer (us) cautions against that, because of the possibility of incompatibilities between the cell jar material and the adhesive. If you're sure of compatibility, go ahead and mount the probe on the battery. Remember, in the VRLA case, the probe needs to be pretty close to the actual battery temperature to prevent thermal runaway.

By the way: in case you were wondering how we derive the slope we use for compensation, it comes largely from empirical data. For lead-acid, for example, the slope is an average, based on -3.3 mV/Cell/°F, at 77°F, at 2.4VPC. Not surprisingly, the slope isn't linear. Various manufacturers express this differently. Circuit designers translate it into –2.5 mV/V/°C.

Output ripple voltage and current

OK. Now you have a good accurate float voltage that also protects the battery from temperature excursions. What about filtering? Are ripple voltage and current really important to a battery? After all, the battery is a huge capacitor – the dc loads see hardly any ripple voltage, certainly not anywhere near the 5% that ANSI allows.

Many authors before me have addressed the effects of ripple current in batteries. There is cell heating, the risk of creeping discharge, excess water loss, etc. I won't belabor this subject. [1]

But I will try to look at ripple from the user's perspective. When is ripple really bad? How do you get the least bang for the buck?

First, here's the biggest no-no. Never operate an unfiltered charger in a dc system without the battery connected. This is especially true for single phase chargers. Take a look at Fig. 3. This is a picture of the output voltage of an unfiltered single phase SCR type charger, operating at light load into a resistive load (that is, no battery). You can see that the peak voltage, which occurs each half cycle, far exceeds the 150V that is typically given as the maximum input for dc switchgear.

Since batteries sometimes go offline voluntarily, I think it's a good idea always to specify some level of filtering for any system with electronic loads. Most standard filtered chargers have no more than about 2% ripple voltage when operated on a resistive load. Most manufacturers also offer what they call "battery eliminator" filtering, which guarantees some lower level of ripple when the battery is off line. You need to consider this only if some of your connected loads are particularly sensitive to ripple. [2]

One final note: in systems with active loads, such as an inverter or UPS, there may be significant ripple put onto the dc bus. The charger filtering can't do anything about this ripple, but it usually isn't of the dangerous variety. If the ripple is a problem to any other dc load, you may have to do additional filtering at the inverter or other source.

Protective devices

There are two types of protection built into a charger. First, there is active protection, such as current limit. Then there are passive protective devices, usually circuit breakers or fuses. Circuit breakers offer a lot more convenience, providing local control of input power, and they are sometimes no more expensive than fuses, at least up to a 100A trip rating.

We sometimes see a request for a high interrupting capacity (high AIC) breaker. Higher AIC breakers come at a considerable cost premium over a standard breaker, and can also add to the lead time for delivery of a charger. A high AIC breaker (on the ac input) makes sense if the ac supply has a high short-circuit rating, and cannot conveniently be protected upstream. Before you specify a dc circuit breaker with higher than 10,000 AIC, do enough legwork to be sure that the battery really can deliver such a high short circuit current. A 200AH battery typically can deliver about 2000A into a short circuit. [3]

For your high AIC input requirements, consider specifying fuses alone (usually 200,000 AIC for ac input), or fuses in series with a standard circuit breaker. You get the convenience of a circuit breaker and the interrupting capacity of the fuse for less than the cost of a high AIC breaker.

Remember that an active current limit circuit is a capable defense to protect the charger from a fault on the dc bus. But now that you're specifying filtered chargers, you need to keep in mind that the filter capacitors, discharging into a dc bus fault, will undoubtedly clear any dc fuses in the output circuit. Modern chargers have very fast-acting current limiting, so that if you specify a dc circuit breaker, the charger will tolerate the fault without clearing the output protection. If your battery short circuit capability is more that 5000 A (typical for a dc breaker AIC), you'll need to specify a high AIC breaker.

Some users like fuses on the dc side. There may be a perception that they're less expensive – that's not always the case. The charger supplier is usually willing to work with you to find the most cost-effective solution for your input and output protection requirements.

Transient response

Every charger spec sheet I've ever seen (including ours) lists transient response like this: with the battery connected, for a 20 to 80% load change, the output voltage will dip no more than X%, and recover to within 1% in, for example, 500 msec. Now go back to the first section, and remember that when the battery is not fully charged, the battery, not the charger, determines the dc bus voltage.

So if that load change took a shot of capacity out of the battery, it's now less than 100% SOC. The time it takes the system to recover to within 1% voltage regulation is a function of the charge acceptance of the battery near 100%, which is less efficient than, say, 80%. [1] The effect is exacerbated in systems with large connected loads, because the share of charger current that the battery enjoys is a lot lower.

Actually, it isn't something I would worry about. The important criteria have already been addressed in regulation, temperature compensation, ripple, and protection. The battery is sized to handle the entire load when the ac power fails. It's going to do it during a load transient.

A transient in the other direction, say from 80% to 20%, will cause a transient voltage rise, as the output current of the charger is suddenly diverted to the battery. The charger usually senses this and reduces its output current nearly instantaneously, but the voltage bulge at the battery will persist until the new load current drains off the excess charge. If you see an unacceptably large voltage rise, investigate the battery condition.

Current limit

Current limit is one of my favorite subjects. The concept of a current limit started a long time ago, when rotating machines were used to charge batteries. A discharged battery is able to accept a very high current. A machine would try to deliver that current, overtaxing the machine, system wiring, and so forth. The idea was to provide some means to limit the current so that machine and wiring size would be economically feasible. The original current limiting means were electromechanical, and therefore weren't very accurate. About the best that the machine could do was promise that the current wouldn't "exceed 120% of the rating."

It's strange how things get inherited. That 120% stayed in specifications, many years after electronic controls were developed that could limit current sharply within 1 or 2 percent. But now the spec has assumed a different life: the charger shall have a "current limit adjustable from 80% to 120%." Users have come to expect 120% output capability from any rating as an entitlement.

Come on, are you really getting something for nothing? If a utility customer insisted that you use 500MCM conductor where 250MCM would be acceptable, are you going to give him 500MCM for the price of 250MCM?

Other manufacturers may handle this requirement differently. What we do is say that the charger is guaranteed to provide 100% output current at the highest rated dc voltage at the lowest rated ac input voltage, and incidentally has a current limit adjustable to 120%.

References

1. Harrison, A. I., "Batteries and AC Phenomena in UPS Systems." Chloride Industrial Batteries Group
2. Blohm, R., "Summary of AC ripple considerations on DC battery systems," Canadian Battery Symposium, 2001. The author is with Enersys Inc.
3. Enersys specification for CC9 cell, catalog section 51.10, rev. 01/01.

Figures

Figure 1: Block diagram – Basic SCR type charger system.

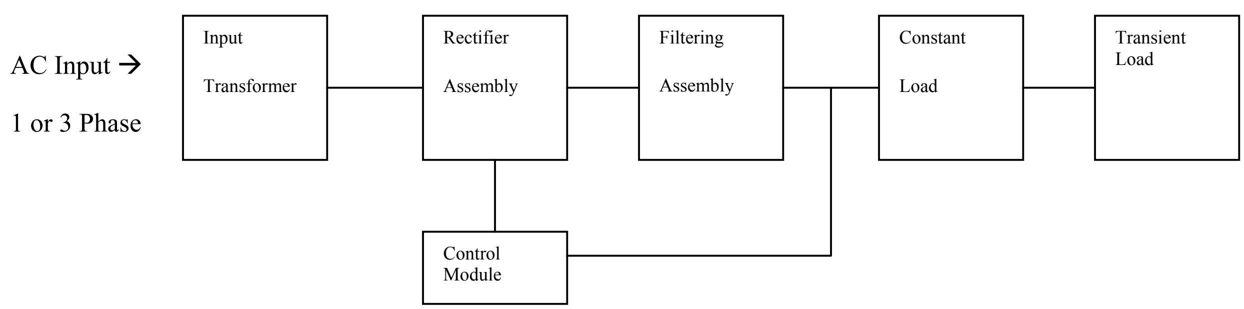

Note: AC and DC protection is included, however not specifically shown in Fig. 1.

Figure 2: Temperature compensation – You can use this to show how to estimate the proper charger float voltage for any temperature within its range. The curve is based on percentage so that it's independent of float voltage. You can also use it to point out the difference in slope between lead-acid and Ni-Cd batteries. Note that, for Ni-Cd, this means that the highest voltage allowed at the cool end is lower than for lead-acid. In our products, we base that cutoff on temperature, not voltage; that is, the voltage stops increasing at 32°F, no matter where the float is set.

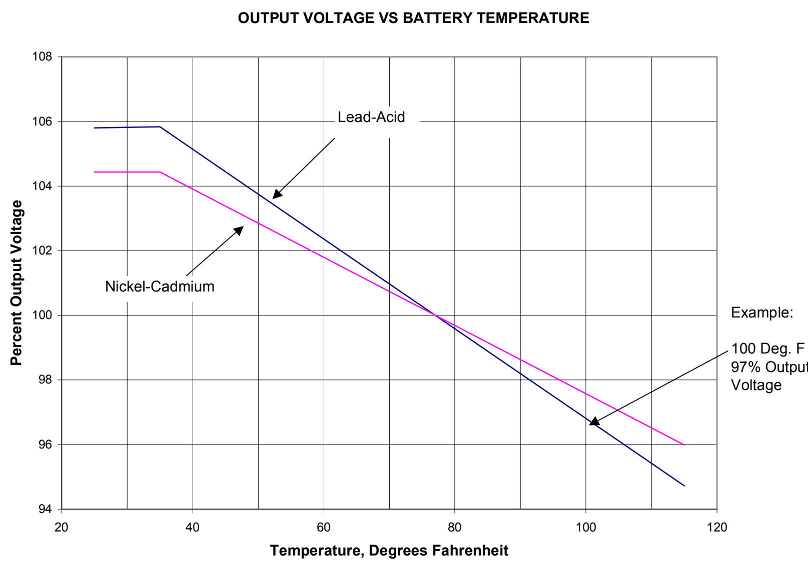

Figure 3: Waveforms for an unfiltered charger without battery – Shown for a single phase charger for the first cycle after startup. This curve shows the charger raw rectified voltage delivered by the SCR bridge (blue), and the resulting voltage across a resistive load (red). In both cases, the SCR firing angle is adjusted for and average dc output of about 130V. Figure 3A shows the condition for a full load, about 25A. Notice that the ripple content is large, but in the steady state won't exceed about 150 Vdc. Some connected loads can tolerate this.

Figure 3B shows the condition for very light load, in this case about 5A. This is the condition you might find in a substation, where most of the connected load operates only during an emergency. In this case, note that the peak voltage across the load almost reaches the peak of the transformer secondary voltage.

Why is the peak voltage so high? The transformer secondary voltage has to be at least 180Vac rms in order to guarantee that the charger can reach a usable equalize voltage (140Vdc) with a low ac line condition (88%). At 180V rms, the peak is 255V.

There are few dc loads that can tolerate a peak voltage this high. Of course, the waveforms shown are for a resistive load; equipment with different input characteristics, such as switchmode power supplies, may produce different waveforms.

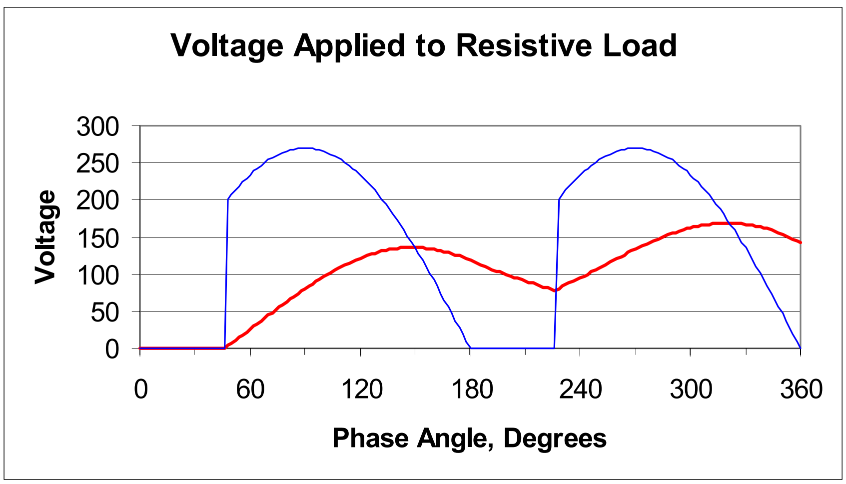

Figure 3A

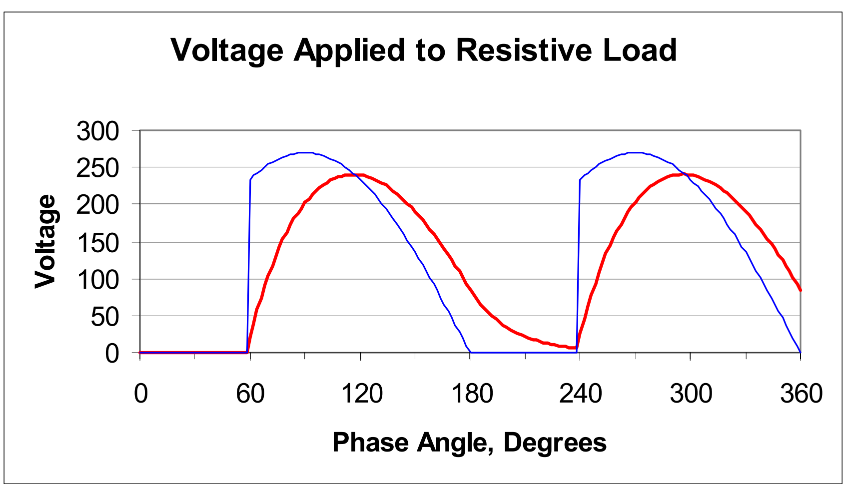

Figure 3B