LiFeCycle Battery Management System (BMS)
Theory of Operation

General Overview

In the laboratory environment, high quality lithium iron phosphate cells like the LiFeBatt or A123 batteries have demonstrated the capacity to deliver several thousand charge cycles. In real life batteries almost never deliver anything like the lab results generated from the testing of single cells. Most of the recent crop of batteries from Asia deliver very poor performance, often becoming useless after a few hundred cycles, though their advertising often quotes the A123 specification of 2,000+ cycles.

Some of the reasons cells might deliver shorter service life in serial strings in the real world than single cells deliver under lab conditions are not well understood. Some are quite well understood. Just about the worst thing you can do to these cells is discharge them below the specified voltage, nominally 2.1V for the LifeBatt cells, and somewhat higher for A123 cells. Often cells discharged this deeply never recover their full capacity, and some are rendered useless.

When several cells are connected together in a serial string, it becomes increasingly difficult as the number of cells increases to know the voltage of each individual cell. The total voltage of the string divided by the number of cells can be used to approximate the low voltage condition for any individual cell, but this assumes all the cells behave the same. Because some cells in the pack may be at different temperature than others, or may have slightly different capacity, over time a condition can develop where an individual cell can drop below its safe operating voltage before the pack cutoff voltage is reached. For this reason the LifeCycle BMS monitors the voltage of every cell in real time and disconnects the power from the load when any cell drops below the preset cutoff voltage. The BMS uses a pair of IRF4110 FETs for a combined on resistance of < .003 ohms to cut off power, or the low voltage detect signal could be used to trigger the ebrake or other controller cutoff signal and leave out the FETs.

The low voltage protection is not a latched error condition, to permit warning the operator well before there is no remaining power to safely get to a resting place or avoid a dangerous situation. Because the cells will first reach the low voltage cutoff under heavy load, by reducing the throttle setting it is then possible to continue for some time, as the cell voltage will rise when the current demand is reduced. When the cell voltage rises more than the built-in hysteresis level of .150V power will be restored. Experience has shown so far that in a well balanced pack the first low voltage warning at full throttle will occur while there is still about 10% battery life remaining. This is just a rough estimate, and this function is not intended to be routinely used as a battery level meter.

When these cells reach a state of full charge, the voltage rises rapidly. Eventually this overcharging may result in damage, shortening the service life and reducing capacity. Most simple chargers operate by shutting off at a preset peak voltage, assuming that the voltage represents the total of all the cells charged to their ideal maximum. This technique can be especially problematic with lithium iron phosphate chemistry. The tendency for a fully charged

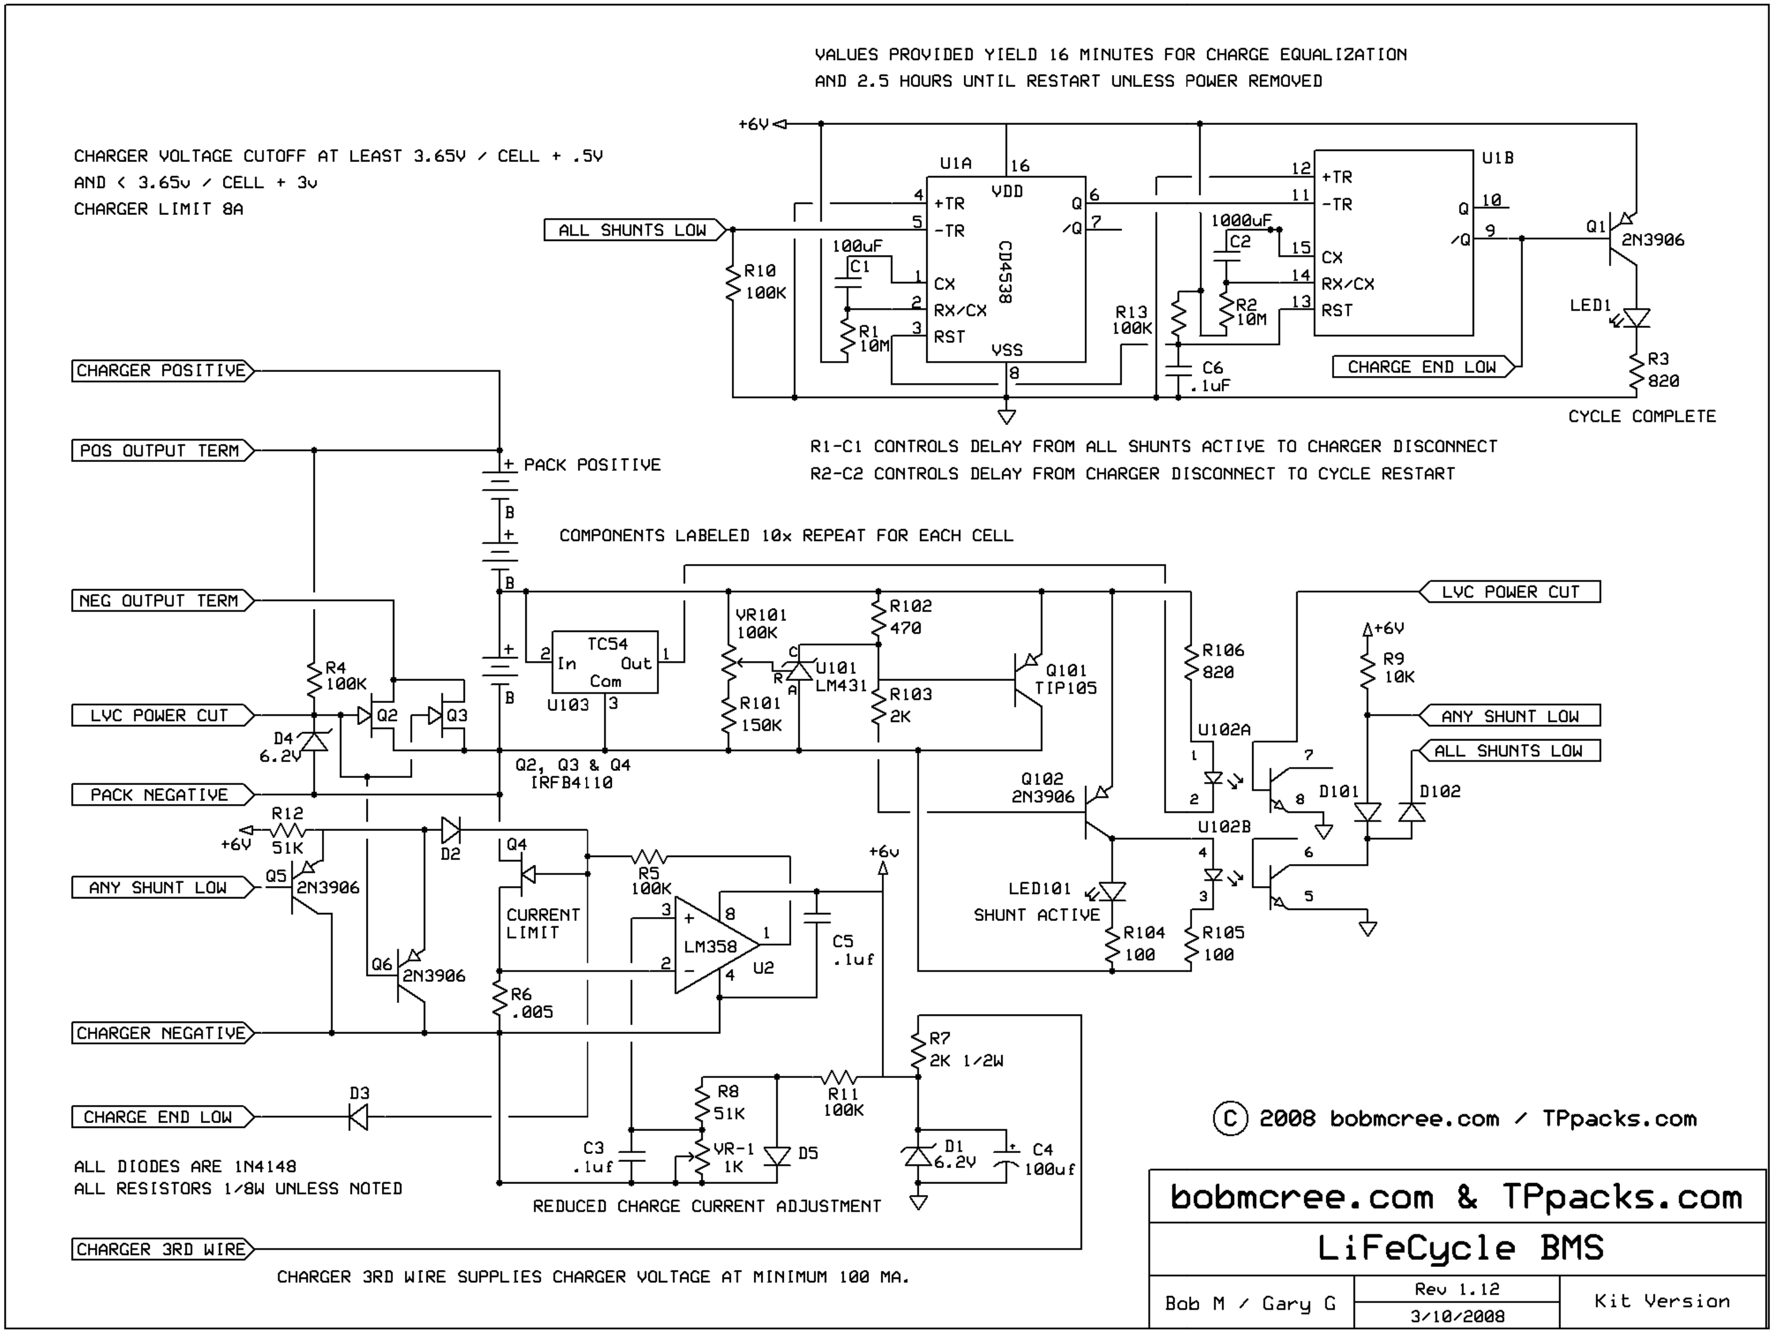

CHARGER VOLTAGE CUTOFF AT LEAST 3.65v / CELL + .5v
AND < 3.65v / CELL + 3v
CHARGER LIMIT 8A

VALUES PROVIDED YIELD 16 MINUTES FOR CHARGE EQUALIZATION
AND 2.5 HOURS UNTIL RESTART UNLESS POWER REMOVED

R1-C1 CONTROLS DELAY FROM ALL SHUNTS ACTIVE TO CHARGER DISCONNECT
R2-C2 CONTROLS DELAY FROM CHARGER DISCONNECT TO CYCLE RESTART

COMPONENTS LABELED 10x REPEAT FOR EACH CELL

ALL DIODES ARE 1N4148
ALL RESISTORS 1/8W UNLESS NOTED

REDUCED CHARGE CURRENT ADJUSTMENT

CHARGER 3RD WIRE SUPPLIES CHARGER VOLTAGE AT MINIMUM 100 MA.

Ⓒ 2008 bobmcree.com / TPpacks.com

bobmcree.com & TPpacks.com

LiFeCycle BMS

Bob M / Gary G	Rev 1.12 3/10/2008	Kit Version

cell to rapidly rise in voltage can result in the charger cutting off before all the cells achieve full charge. This can become a cumulative effect, resulting in some cells never getting fully charged, or other cells being damaged by chronic overcharging. This effect may be one of the major reasons for the discrepancy between single cell tests and performance of a multi-cell pack in the real world.

One possible solution is a charger design that charges each cell individually to the proper voltage. This approach is expensive to implement, and requires two heavy wires for each cell between the charger and the battery pack. What we have done in the LifeCycle BMS is to provide a mechanism for charging each cell individually to the proper voltage without the need for all the heavy wiring or multiple charger outputs. This is accomplished with a shunt regulator across each cell that transfers the charge current from the battery to the shunt when the battery is fully charged, continuing this process until all the cells reach proper voltage.

Detailed Operation

The system is designed to handle chargers from 2-8A, and to reduce heating in the shunts and reduce voltage drop in the wiring that might make sensing the cell voltage difficult, the charge current is reduced to a nominal 1A when the first cell in the pack reaches full voltage and its shunt activates. This provides a good compromise between a very fast charge cycle and a slower more complete charge, by charging at the higher rate until the first cell hits the peak, then reducing the current. As each cell reaches full charge its shunt will activate and continue to pass current on to the rest of the cells, until all are fully charged. The provision for a delay between the time the last shunt activates and the termination of the charge cycle assures a complete charge for each cell. The full charged cell voltage is adjustable, nominally from 3-4v. An individual control for each cell permits adjusting the voltage with the cells attached thus eliminating errors due to voltage drop in the wires. By measuring the voltage directly at each cell while adjusting the controls optimum accuracy can be obtained.

The low voltage detection is performed with a Microchip TC54 IC, which is really designed to reset a laptop CPU on low battery voltage. The part draws only 1 microamp from the cell, and is capable of sourcing or sinking up to 10 mA when the input drops below the set point. The voltage divider used to set the shunt voltage draws about another 15 microamps. from each cell. The charge circuitry is powered only when the charger is connected. In the case of the 2.1v parts this means the output of the TC54 will go low when the input drops below 2.1v. It will not return to the high state until the voltage rises to about 2.25v. This is called hysteresis, and keeps the signal from bouncing around when the level is close to the trigger point.

When the TC54 output goes low, it sinks current, turning on the LED in the optic isolator. Current to the LED is limited by a resistor to under a milliamp. The isolators used have a high current transfer ratio, so this is enough to pull down a couple of milliamps. The isolators permit measurement of each individual cell voltage and producing a signal referenced to a common ground. The resulting signal is used to pull the gate of the FETs low when any optic isolator is turned on, cutting power from the battery to the load or signaling the controller to shut down if that function is desired instead, thus eliminating any voltage drop between the battery and the load.

The second channel of the dual optic isolator is used to combine the signals from the shunts. Diodes D101 and D102 are used to generate two different logic signals from the shunt isolator signals. One signal, ANY SHUNT LOW, goes low as soon as the first shunt is activated when its cell reaches the target voltage. This signal is used to reduce the charge current to 1A (adjustable nominally from .5-2.0A) to avoid overheating the shunts. The other signal ALL SHUNTS LOW, goes low when all the shunts have activated, and is used to trigger the timer which generates the delay until the end of the charge cycle. A second timer provides a “time-out” before the charge cycle will automatically restart. It is intended that the charger be removed when the cycle is complete. During this last period the cycle complete LED will be on and the charge current will be cut off.

Each shunt regulator consists of a voltage divider resistor network, an LM431 shunt regulator, and a TIP105 Darlington PNP transistor. Q102 turns on the led for each shunt when activated, and turns on the optic isolator used to combine the shunt signals. The LM431 acts like a comparator with a 2.5v reference on one input and a PNP output stage, so that when the REF input is driven to 2.5v the device acts like a Zener diode and sinks current. When U101 is off, the base of Q101 is held high by R102, then when U101 is on it pulls down the base of Q101 which causes it to sink current.

The LM431 is specified to operate properly between 1-10 ma and the minimum current gain of the TIP105 is 1000, so if an adequate heat sink was provided and the values of R8 and VR1 were modified the shunt current could be raised up to at least 6A. The heating from the shunts at this current would be unmanageable without a large heat sink and fan.

At the input of each shunt circuit is a voltage divider consisting of VR101 and R101 which is used to adjust the desired battery voltage to the 2.5v trigger level for the LM431. D101 and D102 are used to produce control signals for the charge current regulation system. When the signal from any shunt channel turns its optic isolator on, it pulls down the signal ANY SHUNT LOW. The control circuit uses this signal when high to force Q4 on through R12 and D2, so that when this signal is high the full charge current is permitted to go to the pack. When the signal is low it turns on Q5 which pulls down the anode of D2 to CHARGER NEGATIVE so that it cannot force the FET on and control of the current regulation loop returns to U2. This signal is also driven low by the signal LVC CURRENT LIM which is created by Q6 which buffers the LVC POWER CUT signal and uses it to pull down R12. When all the shunts are low there is no longer a pullup path through R9 to the 6V through D101 and D102 as there is when any of the optic isolators are not conducting, so the signal ALL SHUNTS LOW is generated.

The signal ALL SHUNTS LOW triggers a positive pulse out of U1A, length determined by R1*C1, and the falling edge of that pulse triggers U1B. When the /Q output of U1B goes low, it pulls down the gate of Q4 through D3, cutting off the charge current for the delay set by R2*C2. When Q4 is not turned full on by the signal ANY SHUNTS LOW being high (the initial state), or when it is not cut off completely by the signal CHARGE END LOW pulling down the gate through D3, cutting off the charge current completely, the charge current is controlled by U2. This op-amp circuit limits the charge current to a value set by comparing the voltage on the negative input, which is the voltage across the sense resistor R6, with a DC voltage of 5 mV that is applied to the positive input.

The DC voltage of about 0-10 mV (depending on the setting of VR1) reference for the op-amp negative input is generated from a voltage divider comprised of R8 and VR1, which functions as a 50:1 divider of the 0.5v generated across D5, a 1N914. This voltage of 0-10 mV is referenced to the charger negative, which is also the ground for the op-amp. The op-amp will drive the output positive when the + input is greater than the – input, so it will work to achieve the voltage set byVR1 across the shunt, R6.

When the output voltage of U2 goes positive above about 2-3V, the FET will conduct, causing a voltage to develop across R6 relative to the charger – which is also the op-amp ground. When the voltage across R6 > the voltage on the + input of U2, the op-amp will reduce its output voltage, thus reducing the current, and the loop will regulate at the current set by VR1 as long as the charger supplies enough voltage. By using an op-amp loop rather than a current regulator we can work with the lowest charger voltage possible. The output of U2 goes through R5 to the gate of the FET. The gate of the FET draws very little current, so by using a high value for R5 we can force Q4 on or off through D2 (on) or D3 (off) while the op-amp tries in vain to maintain its programmed output.

The signal ANY SHUNT LOW, when in the high state, forces Q4 on by turning off Q5 thus permitting R12 to pull up the voltage to the gate of Q4 through D2. If any cell reaches the programmed peak voltage its optic isolator will pull this signal low, causing Q5 to turn on pulling down R12 which can no longer force the gate of Q4 high through D2, so control of the charge current returns to U2 unless the signal CHARGE END LOW is low, in which case the gate of Q4 is pulled down through D3 by the /Q output of U1B and Q4 is shut off. When the charge current is cut off most chargers will cut off. A nominal 2.5 hour delay is provided for users with supplies that do not cut off at peak voltage, after which time the charge cycle will restart.

The op-amp circuit and the control logic for charge related signals are only powered when the charger is connected, by using a 3 pin plug with two positive pins. Current drawn from the second positive pin is just a few milliamps, but enough to shorten the pack shelf life if we did not do this. This makes the circuit a bit difficult to understand at first glance. One must consider that the op amp circuit only functions when the charger is connected, and that it functions referenced to the charger negative.

The 6V supply is generated by Zener D1 with current limited by R7. Power is only supplied to the + end of R7 when the charger is connected. The op-amp derives its power from this source, but uses the charger negative as ground.

To derive the reference voltage for the + input of the op-amp a stable source is provided by dividing the voltage drop of D5 by 50 with VR1 at maximum setting, thus providing about 10 mV. This voltage is relatively stable over the range of battery voltage the charger is designed to handle. The 1N914 diode will have about 0.5V drop across it at <100 uA where it is being used, and the voltage is not critical. If the charge current varies by as much as 10% from the nominal 1A value it will not be a problem, as this only occurs during the final few minutes of the cycle with cells that are balanced to within a few percent. With the range of VR1 providing 0-10 mV, it should be possible to adjust the limited charge current between about .5A and 2A, there is a 2-3mV input offset that makes the pot necessary to adjust the current to a repeatable value, nominally 1A with the standard heat sink. With a charge current of 1A, with all the shunts active, the heat sink should not exceed 135 F in a normal room temp environment.

The TIP105 is capable of handling up to 8A so the current reduction would not really be necessary, but a fairly good sized heat sink would be required. The dissipation of each shunt could be as high as 30W, which the device can easily dissipate but if half the shunts are active in a 16 cell pack that would be 240W which would probably require a heat sink of a kilogram or fan cooling. By reducing the current during this final phase the charge cycle is not extended excessively and much more accurate control over the actual cell voltage is possible, due to the lower voltage drops in the wiring. The system has been designed with all wires to the cells equal length and the adjustment can be made while observing the cell voltage directly.

The TIP105 is a Darlington pair with a minimum current gain of 1000. The LM431 can sink 10 mA so that makes it possible to sink up to 10A in the shunts IF there was a huge heat sink, but their maximum rating is about 8A. . There seems to be no need to complete the charge cycle at the full current, and lowering the current makes it much easier to detect the accurate cell voltage for cycle termination.

The 6v supply is generated from the charger second + wire referenced to the pack negative. It is possible during normal operation for several volts to develop between pack negative and charger negative, and this might appear at first glance to be problematic, but it is not. First, one must consider that the circuit cannot function without the charger producing, so we don't care what it does when there is no charge current. Second, if the supply voltage to the op-amp varies because of this difference in ground potentials it will not affect operation. If the current is 1A that will develop 5 mV across R6, compared to 5 mV generated by R8 & R11 & D5, both voltages are referenced to the op-amp ground, which is the charger negative.

A power-on reset circuit R13 and C6 hold the reset pin of both one-shots low until power has time to settle, thus preventing early shutdown or starting in the wrong part of the charge cycle.

The LiFeCycle BMS is not perfect. We hope that with some feedback from users we can get it a bit closer. The circuit is “dumb” but it does what we need it to. The parts are all fairly inexpensive.

Bob Mcree – March 2008