Simple Analog BMS for the Tinkerer – Part II

By Victor Tikhonov
President/CTO
Metric Mind Corporation

This article is the second part of analog BMS description. The first part was published in the December 2012 issue CE. I want to reiterate that this system is not a substitute for sophisticated software-based management systems, nor is it meant for those who never held a soldering iron in their hands and don't like tinkering with electronics. It represents a proven, working idea which can be implemented for your particular circumstance. Another reason I describe in detail how it works is so that anyone can take it further or modify it to suit your particular needs.

DISCLAIMER

Even before completing this portion of the article, I had several inquiries on whether or not I plan to make this system available for purchase for private use. Presently I have no such plans, but that may change. In the past, Metric Mind Corporation allowed amateur (DIY) converters to work on their conversions using the same electronic components we supply to OEMs. However, after several misapplications of expensive hardware coupled with subsequent demands for 'gratis support' etc., I decided to restrict my attention to dealing with professional OEM engineers. If MMC were to produce MIMIC based-BMS, qualification screening of individuals would be needed. Therefore for now, readers of CE are welcome to use the idea described herein — at their own risk.

For the time being, the system described here is not available from MMC — either fully assembled and tested or as a kit. But given the schematic and a diligently checked bill of materials (BOM), success is not out of reach for a skilled person.

THE TRANSMITTER

The block diagram of the transmitter circuit complementing the MIMIC nodes forms a fully functional system. The transmitter conveys the 'ideal average' cell voltage of the traction pack to every MIMIC cell node, for comparison purposes. The block diagram of entire circuit is pretty straight forward and presented in figure 1 (page 13) below. It consists of 4 functional blocks: input isolation and scaling, beat pulse generator, voltage to time (V2T) converter and optocoupler's LEDs driver.

HOW IT WORKS

By taking each block individually, understanding the whole operation will become self-evident. The first functional block is there to provide two functions: galvanic isolation of the HV (high voltage) input from the rest of the circuit which is referenced to the ground (vehicle chassis) and to scale its output so that it represents average cell voltage.

Of course, the transfer function from input to output must be as linear as possible. However, the optocouplers used here are not high-linearity type, but the feedback technique used here removes the non-linearity of the optocouplers themselves from the equation, and the resulting measurement is more than adequate for a practical application such as this.

Detailed Schematic of the Transmitter

As shown on figure 2 (page 13), the isolator consists of the op amp U3b and the bridge formed by generic optocouplers U2 and U4 in one branch, and resistors R4 and R7 in the other. U2 is placed inside the measured voltage loop, while U4 is in the DC feedback loop.

The voltage on inverting input of the op amp is equal to ½ of the supply voltage (because R4=R7). An op amp output will swing higher or lower to maintain the same voltage (a balance) on its non-inverting input (that's basic op amp operation, called a negative feedback servo-system). That means in order to attain a balance, the drop from the bridge resistance of phototransistor in U4 has to be the same as that of U2 to also divide the two full applied supply voltages equally.

In this circuit, if the light intensity of respective LEDs is the same, that means current flowing through them is the same. The current through U2's LED is determined by (and is directly proportional to) the traction pack voltage. So, with higher voltage — more current flows — less resistance of U2's phototransistor — higher voltage on the non-inverting input of the U3b. U3b will compensate for this by raising its output so that the current flowing out and through U4's LED, R18 and R19 becomes equal to U2's current, (balance) in order for U4's phototransistor to conduct to precisely the same degree. This brings the non-inverting input voltage back to ½ of total supply voltage. This servo loop will track input current, essentially mirroring it through the feedback loop. Because output current flows through R18+R22, that voltage drop directly represents that current.

text continued page 14

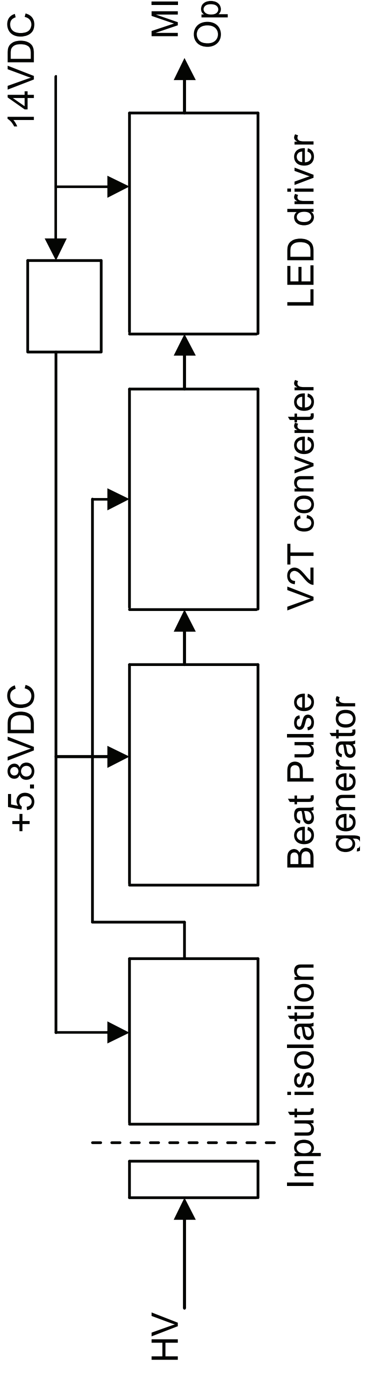

Figure 1

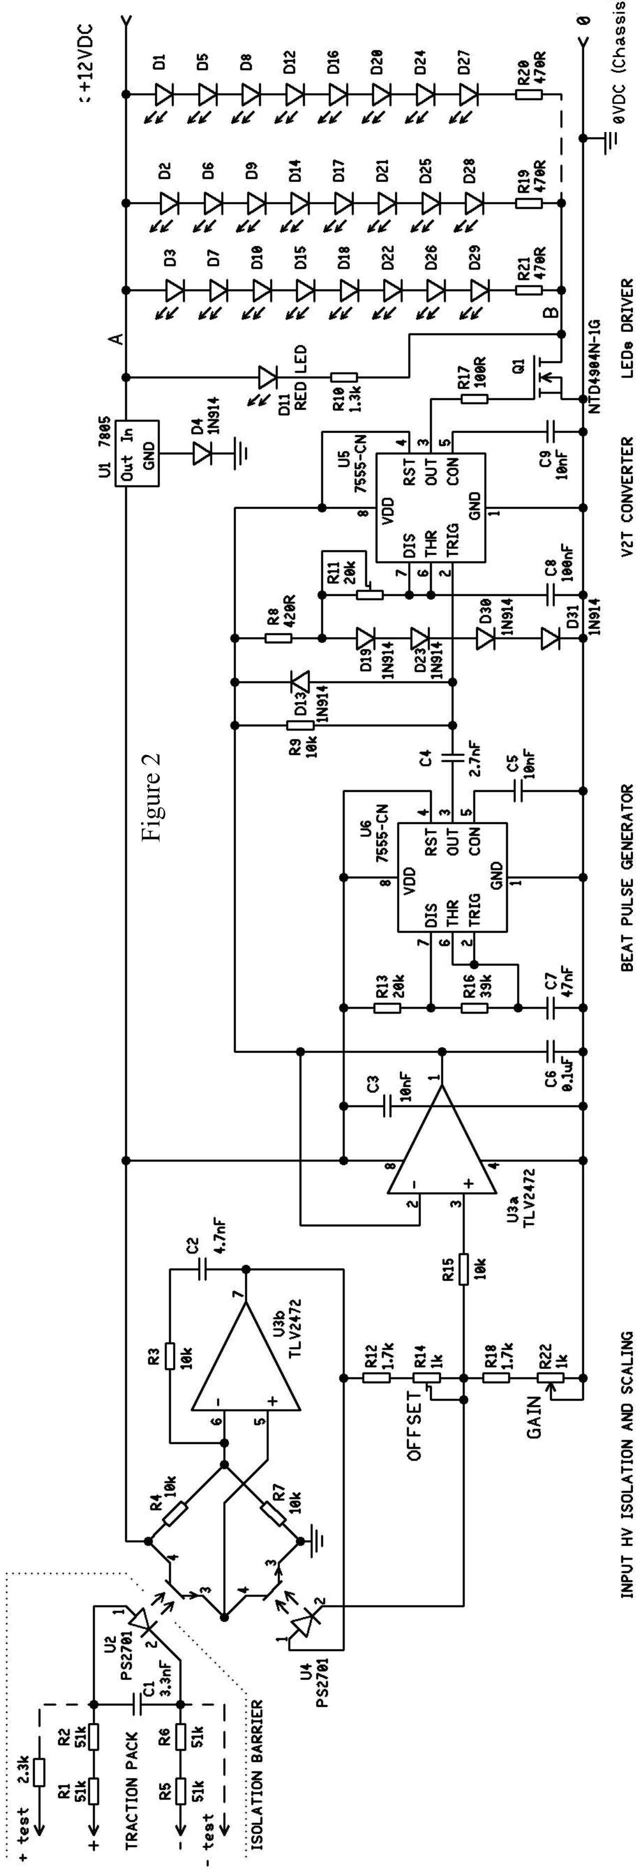

Figure 2

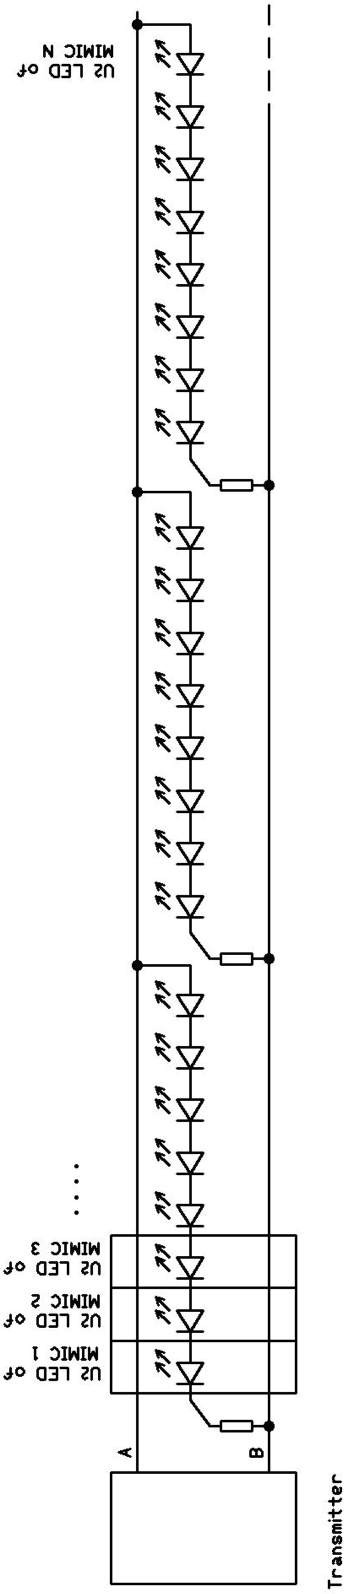

Figure 3

continued on page 14

Simple analog BMS

Continued from page 13

It is important to note that we can change the value of R22, but the current will stay the same (as it still has to match U2's current), so the output voltage will be higher or lower for any given current. In other words, we can adjust the gain of the circuit. Ideally, if input voltage is the same as output, in order to maintain the same currents through LEDs R1+R2+R5+R6 should be equal to R18, plus R22. The op amp output voltage must be about 3.2V to represent the 'average cell' voltage. However, as we have a very high voltage input, (nearly 320VDC or 100 times the output), we just have to make R1+R2+R5+R6 100 times of the value of R18+R22 in order to maintain the equal currents.

Users should adjust the circuit by changing total input resistance accordingly for their pack voltage. The R12 and R14 are there to pre-bias the non-inverting input of U3A, thereby removing an initial ~1.1V "dead zone" output offset on U3b which is a consequence of the ~1.1V LED drop before it starts conducting current. It works because R12 and R14 start conducting current right away, unlike U4's LED which has to wait until voltage rises above its conduction threshold of about 1.1V. In this respect, R12 and R14 are optional as one would never expect output voltage drop on them to fall below 1.1V. That would represent a very dead cell. You will have permanent offset of -1.1V, but it is still a workable solution.

The rest of the circuit is just a booster stage with U3a isolating the high impedance voltage source sampled across R18+R22 from the rest of the circuit. The output of U3a following the 'average cell' voltage directly supplies the V2T converter, which is almost identical to that of every MIMIC module: it generates pulses whose length is proportional to the 'average cell's' voltage. Beat pulse generator triggers the V2T converter every 3.3ms (the generated repetition rate is 300 Hz) using its falling edge.

Because the V2T directly driven by U3a output, it produces pulses just like MIMIC module does: the duration of each is 2ms for the 3.2 Volt Vdd supply. Since there is no need for a comparator to generate a shunting pulse, there is no diode "OR" stage. Instead, output of the V2T directly drives the power FET which connects strings of LEDs to the +12V source supplied by the vehicle's low voltage accessory 12V net. On the schematic the series strings of LEDs (8 in per string) are actually composed of the LEDs inside of the optocouplers on each MIMIC module. So a 3-wire cable is needed as a bus between MIMICs to create this interconnection arrangement — See figure 3 (page 13). Because of the string of 8 series LEDs — if the number of cells to be monitored is not an even multiple of 8 — the resistor value will need to be increased to maintain the same current (~15mA) through the now shorter chain of LEDs. A simple design calculation dictates adding to the base value of 470 ohms already present works out to be approximately (1.15V/0.015mA) = 77 Ohm per every LED removed. For instance, the last chain of, say, 4 LEDs will have current limiting resistor 470+77*4=778 Ohm. Pick the nearest standard value, i.e., 750 ohm.

POWER SUPPLY COMMENT

For simplicity's sake and maintaining a low cost, this circuit uses a very common LM7805 voltage regulator "padded" by diode D4 to produce an output of about 0.7V higher (a forward diode drop) than the rating of a bare regulator. Because the maximum output voltage of the isolator can only be Vdd-Vsat-1.1V (where Vsat is saturation voltage of output stage of the op amp preventing true rail Vdd output, typically 0.25V, and where 1.1V is the forward voltage drop on U4 LED), that pencils out to be 5V- 0.25V-1.1V= 3.65V in the case of a 5V supply, which is very close to the maximum voltage of fully charged cell. To give the op amp more headroom, the diode D4 raises Vdd by 0.7V which provides plenty of margin. Alternatively, you can just use a 7806 or 7809 regulators, but these are not as commonly available as a 7805.

For the op amp you can use any part that allows single power supply operation and with sufficient output current to function as a power source for everything else. A TL2472 op amp was available to me, but there are many other suitable parts.

As a practical tip: while bundling the wiring together looks clean and neat, keep in mind that long runs of parallel and closely packed conductors make up a pretty good capacitor. Despite the operating frequency being low, you don't want to skew (shift) pulse edges and introduce delays due to a large distributed bus capacitance. So strive to minimize the total length of wiring. A star connection using a few strings directly from the transmitter is preferred, rather than having all MIMICs in one long series connection run.

The connection order of MIMICs is not important, as all LEDs carry the identical signal, and do not have to follow the path of electrical connections between cells themselves. But a predictable pattern makes for easier troubleshooting later on. Remember to document what goes where. "An ounce of prevention is well worth a pound of cure".

A final note: while it may be tempting to use dual CMOS timer package instead of two separate parts for U5 and U6, beware!

continued next page

That won't work since the two Vdd supplies for these chips are separated in this schematic, while they would not be inside dual timer chip. That would become a problem.

TUNING THE CIRCUIT

This is simple — first remove the U3b output offset by adjusting the R14 trimming pot. In my prototype this happened when total resistance of R12+ the active portion of R14 was 2.26k ohms. The simplest way is to set the wiper at about the half-way mark initially (at 500 Ohm). Then adjust the gain so that input voltage increment causes the same *percentage* output voltage increment. The easiest way to accomplish this is before installation in the EV, on the bench — apply a precise 3.2V input voltage — the output should follow the input one for one. For this to work, connect a variable power supply directly to "+ test" and "– test" temporary inputs (connection shown with dashed lines, this will bypass R1 R2 R5 and R6. Also use a 2.3k resistor instead.)

Now set input voltage to 3.0VDC. Adjust gain pot such that the output of U3a reads a corresponding 3.0V. Next, increase test supply voltage to 3.5V and watch for the incrementing of the U3a output — it should rise to 3.5V as well. If it is less, increase the gain by turning up R22 (do the opposite if output rises to more than 3.5VDC) until 0.5V input increment causes 0.5V output increment — this will happen when actual input and output voltages do not match.

Finally, adjust offset trim pot R14 to make actual voltage match. Since the combination of R14+R12 is connected in parallel to U4's LED, adjusting R14 will shunt this LED to a different degree, which causes the U3b output current to split in two paths (through LED and R12+R14 in different proportions. (This, in turn, means you have essentially changed U4's Current Transfer Ratio (CTR). That will slightly affect the gain adjustment you just tuned. So with a couple of iterations of adjusting R22 and R14 you will achieve the condition where U3a output follows test input precisely. At that time the output of U5 will generate 2ms long pulses.

Now you can disconnect your temporary power supply from '+test' and '-test' inputs and connect real high voltage traction battery to the + and — inputs marked on schematic "TRACTION PACK". Pay close attention — the "test" inputs are for 3.2V input voltage ONLY! If you accidentally connect high voltage pack to "test" inputs, the current through U2 be is restricted only by low value 2.3k resistor and U2 will instantly blow up! To avoid potential troubles, after initial tuning the circuit on the bench with 3.2V supply on "test" inputs it is better to physically remove 2.3k resistor (in +test lead) from the PCB.

Now the system is ready for commissioning. Total resistance of R1+R2+R5+R6 should be about as many times larger than the sum of R18+the active portion of R22, as many cells you have in your pack. The total input resistance affects gain the same way as R22 does.

When everything is in place and functions, you would want to make sure the "turn-on threshold" of each MIMIC is the same (set with by R4 — refer back to part 1 of the article). While turning R22 from limit to limit, you will change output of U3a by at least +/- 0.5V, making the 'average cell' value artificially 0.5V lower or higher than it really is. This in turn will cause all or none MIMICs to shunt their cells, which is a good way to verify if you have any abnormally high or any too low voltage cells. You'd want to set R22 just under threshold so that the 'average cell' is slightly larger voltage than real cells, so that none is normally shunted. As soon as any cell exceeds this 'average', it will gradually increase the time its FET is on and shunt the cell — this can be observed by watching D2 LED on each MIMIC module.

IMPROVING THIS SYSTEM

You may notice that the maximum shunting duty cycle can never be much more than 50%, and this value decreases when it's needed most — near the end of a recharge. This is because the timer must "expend" first 2ms of the total 3.3ms cycle to sample input voltage; i.e., the "information" about the 'average cell' value is contained in this 2ms time slot.

The circuit cannot start shunting until it understands if it is really 2ms or not, and if you start shunting right away you have only 3.3-2.0 or 1.3ms (out of the slot of 3.3ms) to do it. What can be done is to add a second timer to the MIMIC module (using a dual 555 CMOS timer chip) in a monostable fashion (t=~3.3ms), triggering it by the output of the first one. This way, for as long as the cell voltage exceeds 'average cell' voltage, the shunt turns on hard while the shunting duty cycle is 100%, providing more aggressive control. What is lost is the fine proportional degree of shunting but the action will happen faster.

In general — if using good quality, closely matched cells, the original method described above will yield better results. If the cells are mismatched, more crude but brutal on/off control may more effectively keep cells from overcharging. It is critical to remember — **ANY BMS IS NOT THERE TO FIX BAD CELLS!** It is there to compensate for inevitable "normal" variations in SOC due to manufacturing tolerances and different environmental conditions (such as front and back partial-pack enclosures' temperatures), and for each cell

continued page 16

Paying to Power Our Cars

By Gina Coplon-Newfield, Sierra Club's Director of Green Fleets & Electric Vehicles Initiative.

According to a new report "Where Your Money Goes" by Union of Concerned Scientists (UCS), the average American spends $22,000 to pay for gasoline over the lifetime of the vehicle. That's about as much as some people spend on the cars themselves!

President Obama doesn't like this equation at all. In this week's State of the Union address, the president proposed a new Energy Security Trust that would use funds generated by oil and gas revenues on public lands to invest in advanced vehicle technology that will "shift our cars and trucks off oil for good."

UCS further found that only about 81 cents per $50 fill-up goes to the local service station, and the vast majority of the money goes directly to U.S.-based Big Oil companies and state-owned foreign oil companies.

Plug-in electric vehicles, however, are powered partly or fully by domestic electricity sources. "Freedom from oil" is a reason that many EV drivers cite for switching to plug-in cars. Iraq and Afghanistan war veteran Tim Goodrich in a previous Sierra Club blog post said he bought his Nissan Leaf because the cost of filling up with gas is just too much. I'm not just writing about the price we're paying at the pump; I am also referring to the cost to our future generations, our national security, and our economy. As a veteran, I have seen the toll these costs take and I am doing what I can to stop contributing to the problem.

In terms of dollars and cents, some are wary of an EV purchase price that is often more expensive than gasoline-powered counterparts. However, our friends at Plug In America contend that EVs are actually cheaper than many gas-powered vehicles when you factor in all the costs, including purchase/lease price, maintenance, and fuel. They did the math and found that over a ten-year period, a Nissan Leaf would cost about $38,430, a Toyota Matrix about $58,613, and a Toyota Prius V about $51,359. And of course for the electric car, no money goes to Big Oil.

Out-going U.S. Energy Secretary Steven Chu recently announced a new Workplace Charging Challenge, an effort to provide more people with the opportunity to charge their EVs at work. Google, Verizon, and the Big Three U.S. automakers are among the 21 founding partners of the workplace charging program (at no tax-payer expense) that pledge to "assess workforce plug-in electric vehicle charging demands, and then develop and implement a plan to install workplace charging infrastructure for at least one major worksite location."

These workplace EV chargers will help as more people do the math and figure out how to free themselves from oil.
http://sierraclub.typepad.com/compass/2013/02/paying-to-power-our-cars.html

---

Simple analog BMS

Continued from page 15

causing that gradual SOC divergence — there is simply no mechanism to inherently equalize cells as the pack ages and is repeatedly cycled. In a series string, each cell doesn't know about existence of other cells (other than sometimes sensing heat of adjacent cells).

COOLING

Make sure there is adequate cooling for all shunt resistors. Do not buy 5W rated shunt resistors if your shunt current is 1.5A (which means power dissipation is at cell voltage of 3.2V is 3.2V*1.5A= ~5W. That resistor will not glow red but will be hot enough to unsolder itself from the PCB and the circuit will fail in a hurry. Study the data sheet for those power resistors carefully to see the temp raise rating vs. dissipating power. It might be tempting to choose low resistance shunts to tame misbehaving cells, but remember, using 100 individual shunts that are all on will produce ~500W of waste heat for this implementation example. That's nearly like having ten regular 50W incandescent light bulbs on at the same time inside the [normally] enclosed box. Aside being wasteful, give thought as to how best to get rid of all that heat!

It is common for application circuit diagrams to omit such components such as bypassing capacitors or safety fuses, commonly accepted as default. They serve a purpose, so don't forget them. By bypassing the incoming 12V supply with 1,000uF electrolytic cap (watch its low temp. ratings!) — that will be sufficient. Consider adding reverse polarity protection by using a forward biased diode in series +12V supply rail.

If you have any questions about described system, always welcome to drop me a line to tech.support@metricmind.com.

Victor Tikhonov, MMC