World's Fastest Electric Vehicle's BMS Built With Home Brew Surface Mount Assembly Equipment

(picture of a "burnout" before a run)

Remember when electronic hobbyists thought that transistors would be the end of home tinkering? I don't, but I've heard stories…. Or how Integrated Circuits would do the same thing? "How can you debug the circuit if you can't even see each transistor!" I have not personally heard that one either, but I have heard comments like "What am I suppose to do when I can't buy a discontinued logic gate? Program a PAL?" and "EVERYTHING new is surface mount these days. No wonder the hobby is dying; how are you supposed to solder a part you can't even see without a microscope?" I use to believe the last comments, until I was put into a position where I had nothing to loose, and gave it a try. What I found was the reports of the death of the engineering hobby were greatly exaggerated. This contest entry describes how I currently assemble surface mount boards.

With rising fuel costs and fears of global warming, there has been renewed interest in Electric and hybrid vehicles. But Electric Vehicles still have a lot of misconceptions with the general public. When the average person thinks of an "EV", quite often images of slow golf carts and gutless "Traveling Science Fair" projects driven by hemp-wearing aged hippies come to mind. To help dispel these misconceptions, along with some sort of primordial need to go fast and use cool technology (hey, us engi-nerds gotta do _something_ different to impress each other), The National Electric Drag Racing Association (www.NEDRA.com) was formed. I also would like to show Team [www.KillaCycle.com's](http://www.KillaCycle.com) latest attempt to best their world record in the quarter mile.

First, here's a brief introduction to electric drag racing. At NEDRA sanctioned events, the goal is to reach the end of a 0.25 mile in the shortest amount of time, or to be the “quickest”. The vehicle traveling at the highest rate of speed at the end of the drag strip can claim to be the “fastest”.

Recently we had the opportunity to use www.A123Systems.com awesome lithium-ion batteries. These are the exact same cells used in some of the power tool manufacture’s line of new 36V Lithium-Ion power tools. These cells are unique in that they have a much higher power density than conventional Li-Ion batteries, have a much higher thermal conductivity which allows them to be recharged in as little as 5 minutes, and are much safer than conventional Lithium-Ion batteries. On their web site, they have some very impressive videos of what happens when you drive a nail through a charged conventional cell (can you say “road flare”?) and a charged A123 cell (mostly just steam, no flames).

Lithium-Ion batteries require a battery management system, especially when you hook 110 cells in series. This was the reason behind this project, to finish construction of a Battery Management System in time for the races.

In a nutshell, our BMS had to protect each of the 880 cells (8 in parallel, 110 in series) from over-charge, over-discharge, and over-temperature. This is not as simple as it might first appear. While all the cells may start out with the same capacity, due to aging, manufacturing variations, and the difficulties of keeping all the cells at the same temperature, their amp hour capacities and state of charge will drift apart. There always has to be one cell that is the most discharged and/or has the smallest capacity. After this cell is completely discharged, you can’t take any more power out of the battery pack without damaging this cell. Due to the high voltages involved, damaged cells can draw an arc, which will produce what is commonly referred to as a “plasma event”. If you try to completely charge the most discharged cell while charging the pack in a series string, you will end up over-charging the cells that were not completely discharged, and the ones that have lesser capacity. So you will either over-charge some cells, or under-charge others. And since the same conditions probably still exist that caused the original imbalance, this process will repeat itself next cycle, but to a greater extent.

We chose to continuously monitor each cell’s voltage and several temperatures throughout the battery pack. We log these temperatures and voltages while gong down the track, and at any time a Laptop can be used to interrogate any voltage, any temperature, or modify any parameter. If at any time any of the temperature or voltage parameters are exceeded (both too high or too low), either the charger is shut off or the motor controller is put into a low power “Valet” mode. During charge, when one cell’s voltage reaches the lower threshold, we start to PWM a power resistor across that cell to bypass some of the charging current. There is a second threshold where the PWM reaches 100%.

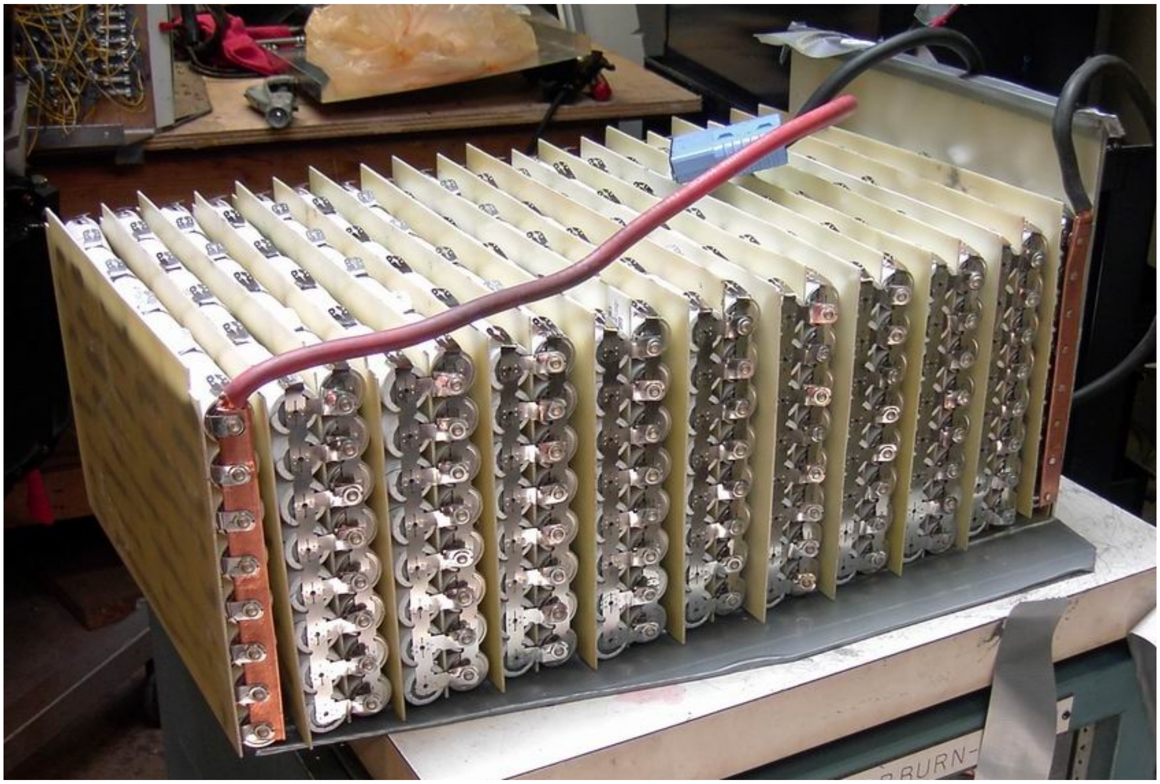

We ended up with 22 “modules” consisting of 5 cells in series, and 8 cells in parallel, for 40 cells per module. Each module was assembled on a 1/8” piece of G10 fiberglass. By the time all the dust settled, I was allowed a maximum of 1” x 13” of printed board space on each module to manage 5 cells in series. This required the almost exclusive use of surface mount components attached on both sides of the printed circuit boards.

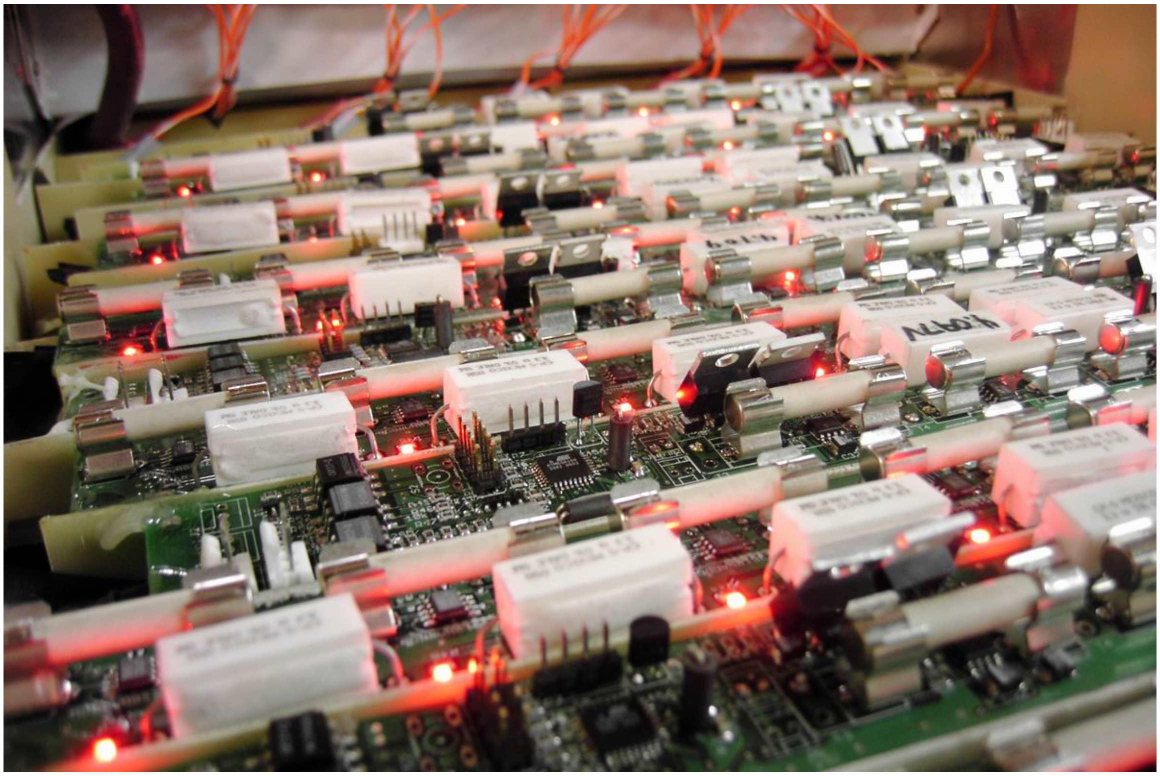

There can be almost 400V between the first cell and the last cell, and we need to be able to monitor all these voltages. To keep life simple and to minimize risk, I decided to use LM1990 differential amplifiers. These can measure with +/-250V of common mode and can tolerate +/- 350V of input voltage. . This was a good decision because there were times when a slip of the screwdriver or a misplaced voltmeter probe caused a fuse to open. If it were not for the LM1990 diff amps protecting the front end, there would have been more damage than just a blown fuse.

Over the years there have been several “Ah-Ha!” moments that have made my surface mount journey easier. I’d like to describe them in detail, but I do not have enough room. So I will quickly list some important things. First use good tweezers. As soon as I spent $10 on a pair and threw away the $0.95 pair used for plucking eyebrows, the small parts stopped doing their disappearing act. The next tip I figured out the first time I tried to solder an AD9850 DDS chip, with pins about 0.012” wide and 0.012” between pins. I worried about how to attempt this, until I just gave it a try. I had a few solder bridges between these pins, but solder wick took off the excess solder giving it an amazingly professional appearance. Another tip is to use plenty of flux from a flux pen. This really helps in both soldering and using solder wick to remove excess solder. At this point, I had the skills to do surface mount assembly, but it still took me 3 or 4 times longer than if I used all leaded parts. That is until I discovered solder paste with a home made pneumatic dispenser. This allowed me to first apply all the solder paste and then place all the components. Then all the parts are soldered at once.   I learned about solder paste and reflow techniques on the yahoo group’s E-Z_Bake group. It is a group dedicated to using conventional toaster ovens as surface mount reflow ovens (see http://groups.yahoo.com/group/E-Z_Bake/). After reading the application notes at http://www.efdsolder.com/ , I bought a 35 gram syringe of "no clean" solder paste from Kester, requested some sample syringes and tips from EFD, and built a pneumatic dispenser controller from some surplus valves, air pressure regulator, and small air compressor. This controller regulates the air pressure applied to the syringe of solder paste, and how long it is applied. After the pulse of compressed air, the syringe needs to be vented to atmospheric pressure via another valve to keep the paste from dribbling out. The syringe is then hand placed above all the surface mount pads while the valve pressure sequence is activated via a foot switch.

When hand placing the components, you don't need to worry too much about getting them exactly lined up, because of the amazing part that comes next. As you heat up your printed circuit board enough to melt the solder, the surface tension of the solder will pull most of the parts into alignment!

It is best to follow the recommended temperature profile in the solder paste data sheets (see http://www.kester.com/Data%20Sheets%5CSolder%20Pastes%5CNo-Clean%5CR276%20Global%20(10Nov04).pdf for the temperature profile). The goal is to pre-heat the board at a rate fast enough so the solder paste flux does not dry out, but slow enough not to thermally shock the board and components. The "Soaking Zone" is to allow all the components to equalize in temperature at just below their maximum safe storage temperature. Finally the board is heated above the solder melting point as quickly as possible, and then quickly cooled down to the Soaking Zone. The board is gradually brought back to room temperature at a slow rate to prevent thermal shock. My attempts at temperature controlling a toaster oven and a hot plate taught me that they did not have enough power to follow this profile. Now I just put the board in the toaster oven, crank it up to about 250 deg. C, and watch until all the solder paste has turned shiny. I then turn off the oven, open the door about half way, and let it cool for a few minutes.

You may have noticed that this BMS has components soldered on both sides. How is it possible to prevent the bottom components from falling off if in a toaster oven? It's not ----- unless the bottom components are glued on. For this project I finally splurged on a $200 hot air SMD rework station (http://www.sparkfun.com/commerce/categories.php?cPath=46_48). When I had to reflow solder the second side of the PCB, I would pre-heat the PCB on the hotplate set for about 200 deg. C and use the hot air reflow wand, set to about 350 deg C, to reflow a few parts at a time. If the board is not preheated, the solder paste and component leads come up to melting temperature before the PCB does, and by the time the PCB comes up to temperature, most of the flux has evaporated leaving what looks like a cold solder joint. This also hinders the surface tension of the molten solder from doing it's trick.

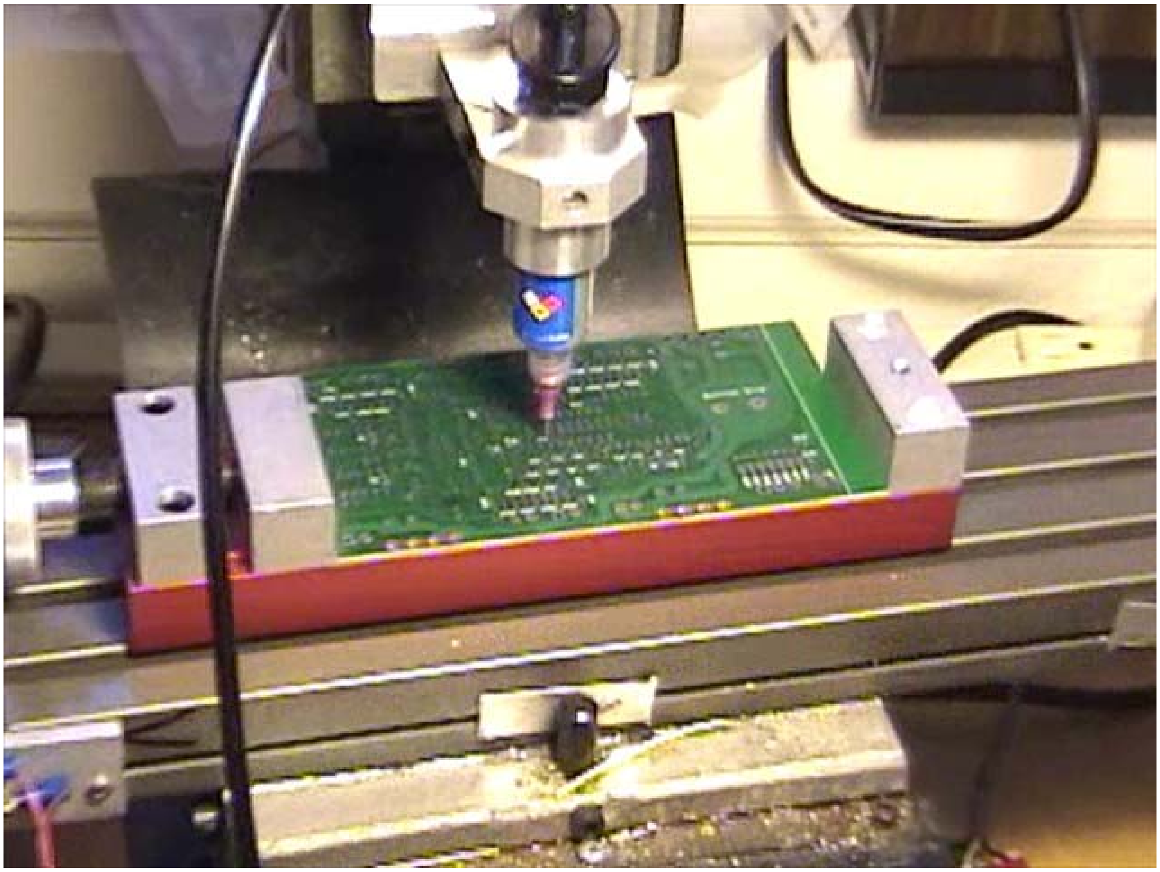

My final big advancement was the development of a CNC based solder paste dispensing system, and a CNC based surface mount pick and place machine. I think it was mostly the “way cool”-ness factor involved that motivated me to work on these projects, but I’m glad I did, because it was a huge time saver when I had to make 10s of the same PCB.

My home-brew CNC controller drives stepper motors using the windows based Mach3 controller software (www.ArtOfCNC.com). These motors are attached to my desktop sized Taig mill (www.TaigTools.com). I wish I could talk more about how useful it is to have a small CNC mill in the electronic hobbyist’s toolkit, but that would have to be another paper entirely by itself. I use it for milling out custom enclosures for my electronic projects, to engrave front panels, and to engrave and drill prototype printed circuit boards. It’s hard to justify making PCBs at home unless you either a) enjoy the process, b) have much more time than money (in which case I’m surprised you can afford a small mill), or c) need the board in less than the 48 hours it takes to have an outfit like www.BareBonesPCB.com manufacture and deliver it to your door.

For the solder paste dispenser, I added a way of holding the syringe of solder paste in the spindle motor location, and a way of controlling the dispenser valves under CNC control. This was accomplished using a spare output bit on the PC’s printer port and writing a really simple M-code macro in the Mach3 control program. To generate G-Code to tell the mill where to dispense solder paste dots, I used www.CadsoftUSA.com’s Eagle PCB cad software with a modified ULP (smd-coordinate.ulp) to output an Excel Comma

Separated Variable (csv) format data file. This file can be sorted in Excel to minimize the path traveled by the solder paste tip. Then this file is read by a LabVIEW program which calculated the length traveled, determined how many dots of paste to dispense based on pad size, added in offsets and mirrored axis as necessary, and then generated the G-Code file to be read my Mach3. Please see www.ciciora.com/surface_mount.html for a video clip of the dispenser in action on a different project.

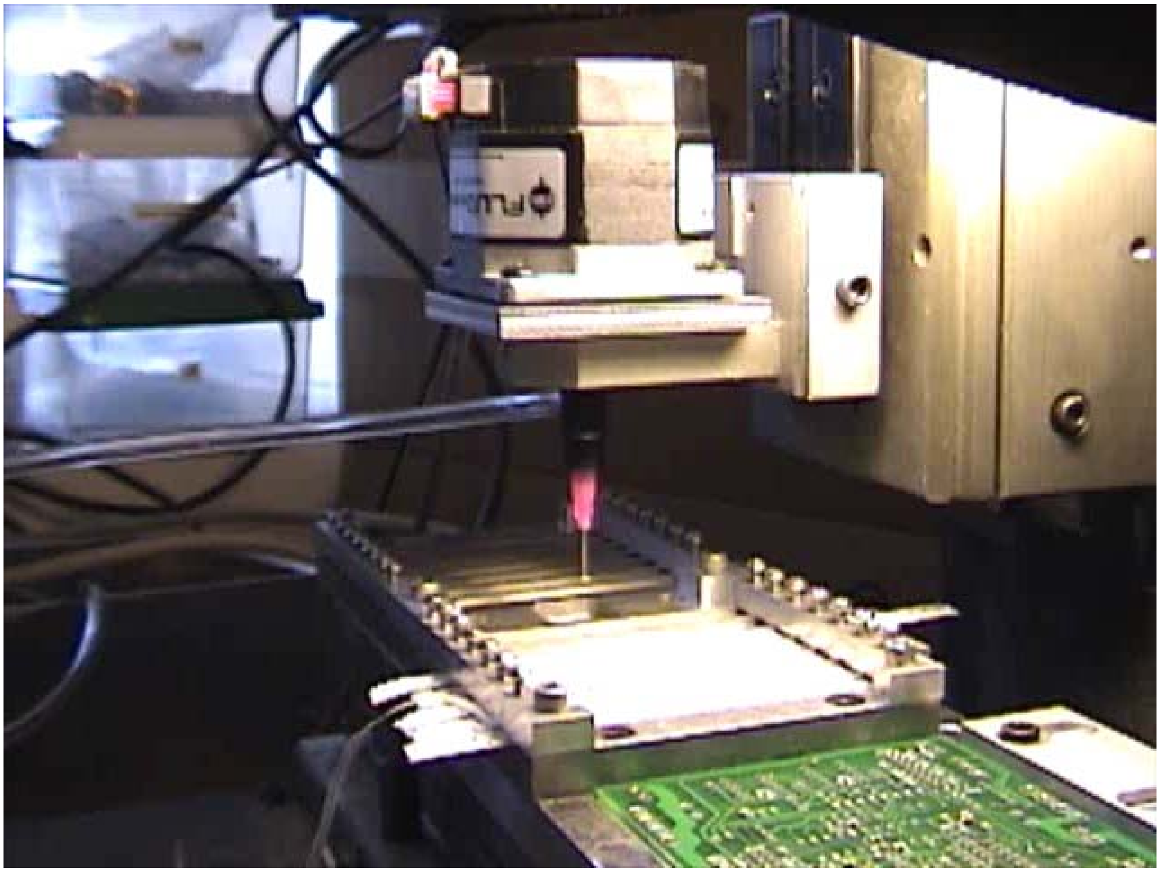

The pick and place fixture was a bit more fun. In place of the spindle motor, I made a bracket to hold a stepper motor for the 4th axis. This is needed to rotate the components after they have been picked up. On the stepper motor shaft, I made an adaptor to hold a 10ml syringe, which held the same tips used for solder paste dispensing. Using the same valve box used in solder paste dispensing, I added a third valve that applied vacuum to the pick and place tip.

Two more fixtures had to be made, one to hold the surface mount components at a known location, and another one to hold the printed circuit board in place. To make this as flexible as possible, I attached the surface mount tape fixture to www.HighTechSystemsLLC.com 's Modular Holding System. Even though the surface tension of the molten solder helps center the components, the tape that surface mount components come on does not hold the parts accurately enough. So in the part fixture I included a "centering pit". First the pick and place head is moved above the part to be picked up. Then vacuum is applied, attaching the part to the tip. The tip is then raised and rotated to the correct orientation. The part is then moved to the center of the centering pit, and the side of the component is bumped up against each of the 4 sides. This aligns the part square with the axis of motion, and insures the part is perfectly centered on axis of rotation. Notice I didn't say perfectly centered on the pick and place tip; the tip can be out of alignment with it's axis of rotation. By rotating first, then centering, this error can be taken out. The part is then lowered to its destination, and sticks to the solder paste already applied to the printed circuit board. At this point a puff of compresses air insures that the part does not stick to the tip, and then the tip is vented to atmospheric pressure. This is repeated until all components are placed.

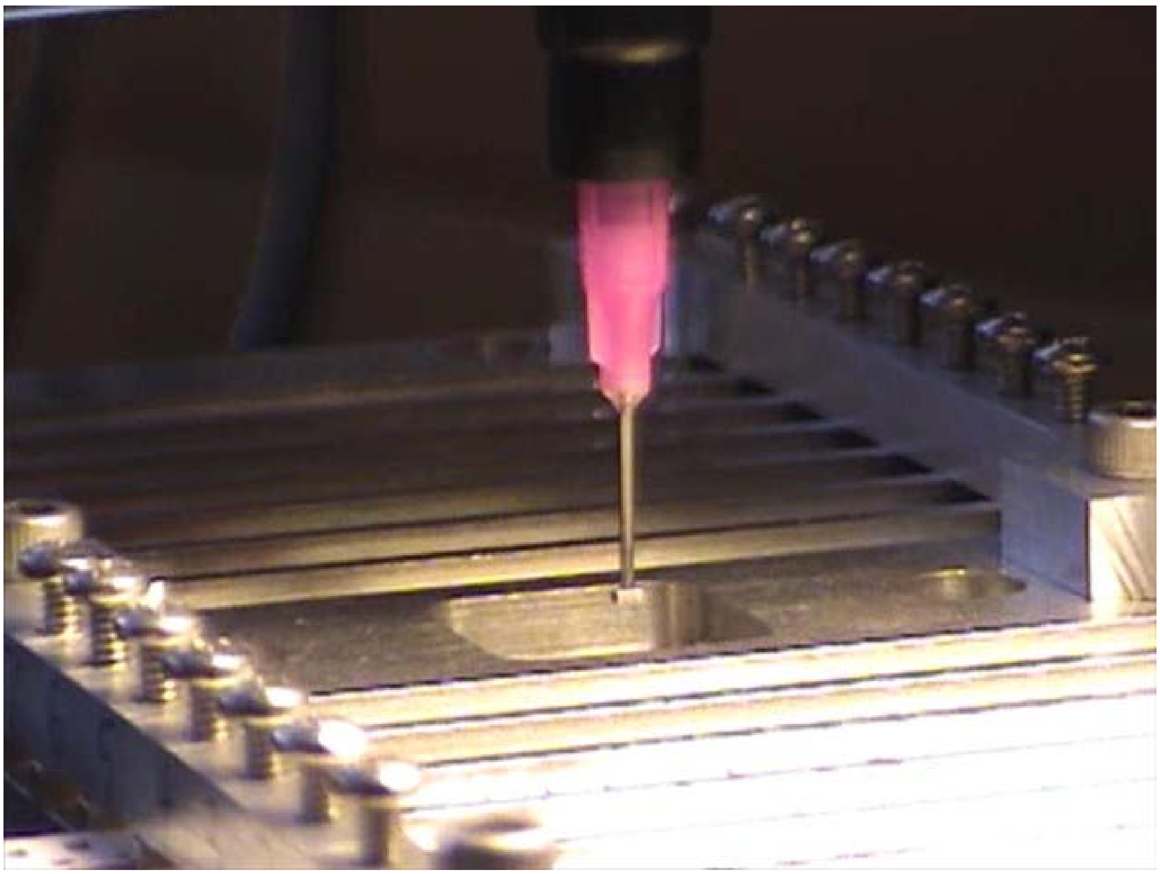

Without my Homebrew Surfacemount assembly techniques I couldn't have finished this project in time for the races.