The Essential Primer on the Diamond-Star Engine Control Units

Contents

Introduction

There is a great deal of confusion in the DSM community about how the engine control unit (ECU) actually operates, and why it operates the way it does. Although this information is readily available, the terminology used is unfamiliar to some, which leads to significant misunderstandings. This document is intended to provide all the requisite information necessary to properly understand the ECU and it's operation on the DSM vehicles.

For those who are familiar with this information and want more details, you will need to get the $2 DSM Technical Manual. All of the information you could ask for is contained in this reference, and I'm not attempting to duplicate it here. Get it at the Backup Manual CD home page.

Chapter 1: Open-loop Systems

Open-loop systems apply control signals without regard to the output signals being produced. To model one, take a piece of paper and draw a box on it. On the left side, draw a horizontal arrow leading into the box - this is the input. On the right side, draw a horizontal arrow leading out of the box - this is the output. Voila, an open-loop control system!

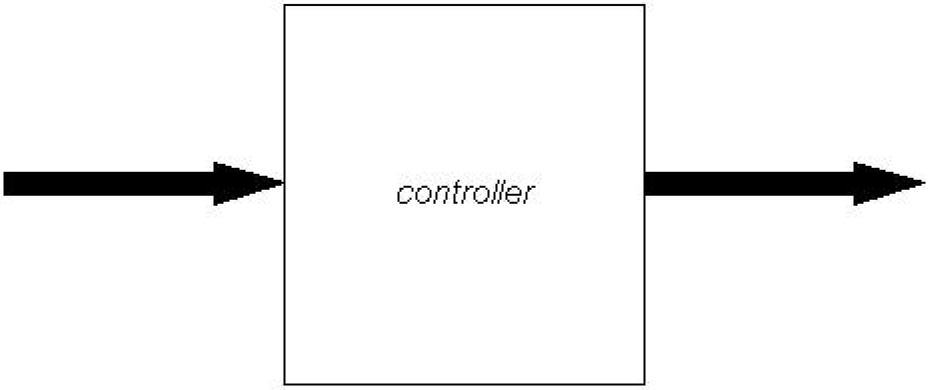

Fig.1 - Basic open-loop control system

The box is the *controller* or *control system*; the part that actually does the 'thinking', if you will. The controller can be anything - an electronic circuit, a mechanical system, a computer, whatever. It takes *input signals* (the leftmost arrow), processes them and generates *output signals* directing the action.

To get a little fancier, you can draw another box on the other end of both arrows to make 3 boxes. The leftmost box represents the *input system*, which provides the original input signals to the controller. The rightmost box is the *controlled system* - in other words, the part that the control system is running.

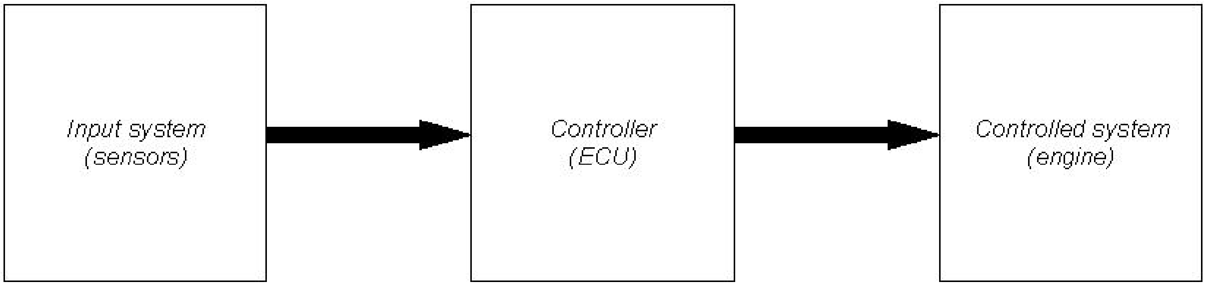

Fig. 2 - Basic DSM open-loop system

If you like, you can imagine the middle box on the page representing the ECU inside a DSM, or any other car. The input system can be any or all of the input sensors used by the DSM ECU, including airflow, throttle position, air temperature, barometric pressure, etc. The output signals consists of control outputs to fire the injectors, provide spark to the cylinders, and so on, which (naturally) run to the engine.

As you might expect, open-loop control systems are generally not very useful, because the controller does not have any way of knowing how the controlled system is behaving. As a simple example, let's make the input a temperature sensor, and the output a set of heating elements. For any given temperature, the heater controller 'knows' to make a certain amount of heat, and it will direct the heating elements accordingly.

The system will work fine when it is put together, and probably for some time thereafter. Let's just say, however, that something goes a little bit wrong with the heating elements - one breaks, or they just wear out a bit, so the heater is now putting out less heat than it should.

In this example, the air temperature is the same, so the inputs are unchanged. The control system is still producing the correct signals to run the heater at that air temperature. The heater could have lots of extra heating capacity left, which could easily compensate for the temperature drop. The problem is that the heater is not behaving like it should, and the control system has no way of knowing this.

There are some examples where open-loop control is fine. If all that is being controlled is an indicator light, or and LCD display, maybe it doesn't matter if they work because a person will notice the fault, and correct it. Systems that are easily predictable, in other words. This isn't always the case, however. Where open-loop control fails us, the solution to this is the *closed-loop* system.

Chapter 2: Closed-loop Systems

Closed-loop systems add an extra element to allow the controller to detect changes in the controlled system, and compensate accordingly. To show this capability, you need to draw one more arrow on your diagram. This arrow begins at the top of the existing right-hand box (the 'engine'), and runs to the top of the existing middle box (the 'ECU').

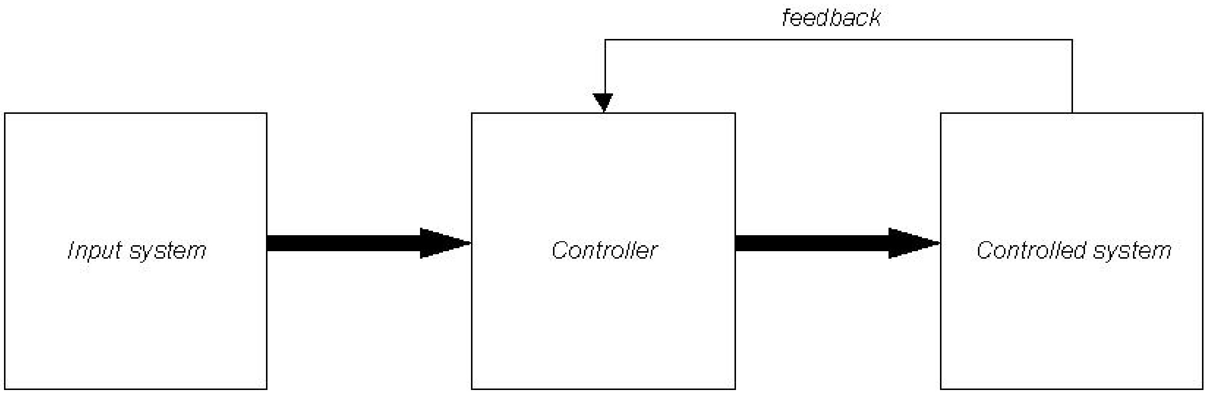

Fig. 3 - Basic closed-loop system

This extra arrow represents a quantity known as feedback. This means just what it says: some part of the operation of the controlled system is measured and 'fed back' to the controlling system. Feedback represents very useful information, because now the control system can apply a signal to the controlled system and measure the results of that signal.

Going back to the simple heater example, we already have an input system (the temperature sensor), a controller, and a controlled system (the heater). To introduce feedback into the system, we could add a temperature sensor inside the heater itself, capable of measuring the actual heat output of the unit. Prior to this, the control system would measure the air temperature and direct the heater to generate a certain output. With the addition of the extra heater sensor, the control system can actually check to see if the heater is producing the correct amount of heat, and react accordingly.

So, as the heater wears and output drops, the controller will 'notice' the drop in output and can direct the heater to produce more heat to compensate. The heating system will continue to work correctly until the heater breaks, or becomes so worn that it can't compesate for the output losses any more. Once this happens, though, a 'smart' controller can always sound an alarm or send an error message of some kind.

So, closed-loop operation allows the controller to compensate for external variations that may affect the operation of the controlled system. Since very few parts in the real world are absolutely precise, closed-loop operation is a necessity in virtually all control systems.

Chapter 3: Closed-loop operation on DSMs

Relating the closed-loop system to the DSM ECUs, the newly-added feedback arrow represents the output of various sensors which are placed in the engine to monitor its operation. The primary sensor of concern is the *oxygen sensor*, a device which measured oxygen in the exhaust gases produced from the engine. Others include the RPM sensor and crankshaft angle sensor, which are also crucial elements in engine control.

[I suppose to be more accurate, you should divide the 'engine' box into two halves. The top half would represent the sensors used to monitor the engine, while the bottom half could represent the spark plugs, fuel injectors and other controlled components.

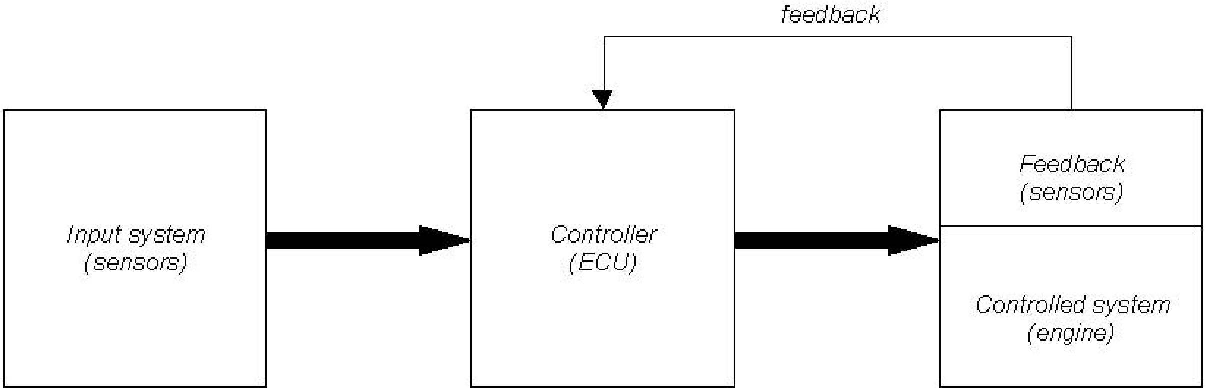

Fig. 4 - DSM closed-loop system

Still, you get the idea.]

When people talk about open or closed-loop operation on DSM ECUs, they are invariably talking about the specific function of fuel delivery. The ECU has other functions as well, but the most prominent task is the requirement to inject the correct amount of fuel into the engine, based on the amount of air entering the cylinders. Too little fuel and the car won't run - too much and the car will not burn it all, leading to waste and poor fuel economy.

The ECU has several input sensors that help it determine the correct amount of fuel. The most widely recognized is the *mass air flow sensor (MAS or MAF)*. The MAS consists of the air volume sensor, the intake air temperature sensor, and the barometric pressure sensor. These three sensors form a single component located directly behind the air filter, and it provides the ECU with the information required to calculate air mass.

[Why mass? Well, the ideal ratio of air to fuel depends on the number of molecules of each type of substance. Knowing the volume doesn't get you anywhere, since the number of air molecules in a given volume changes with temperature and pressure. If you know air volume, temperature and pressure, though, you can get air mass, which in turn can be used to determine the correct amount of fuel.]

So the input signals consist of the air volume, temperature and pressure from the MAS cluster. Feedback is provided by the oxygen sensor, located in the cars downpipe.

[Why oxygen? Measuring the unburned oxygen in the exhaust is a practical way of figuring out how much fuel was really burned in the engine. Very little oxygen in the exhaust indicates that too much fuel was added to the air/fuel mixture. This is known as a *rich* condition. On the other hand, lots of oxygen in the exhaust indicates that there is too much air in the mixture, a condition known as a *lean* mixture.

If the mixture is too rich, the engine is eating more fuel than it can use, which is a waste of gasoline. Also, unburned gas and soot from the incomplete burn will be expelled from the engine, coating exhaust parts and leading to increased air pollution. When the mixture is too lean, though, the mix will burn much hotter than normal, leading to the possibility of literally melting engine components. Bad bad bad.]

With the MAS providing the input and the oxygen sensor the feedback, our fuel control system can run in closed-loop mode. Note I said 'can', not 'will'. Just because the feedback sensor is present doesn't mean the ECU has to use it. More on this later.

Before we can continue, a little background on the oxygen sensor. The sensor output is inversely proportional to oxygen - the less oxygen, the higher the sensor output. Also, the sensor output does not increase steadily with decreasing oxygen. Instead, the sensor can be thought of as an 'oxygen thermostat', which tends to change a great deal at a certain oxygen level.

Not by coincidence, the DSM oxygen sensor is designed for a large change in output right around the oxygen level that results from the ideal 14.7:1 air/fuel ratio. Add a little bit of air, and the signal goes way down, while a little bit more fuel makes the signal go way up. Outside this 'magic' range, the oxygen sensor is roughly linear, although it is difficult to predict the actual output for any given air/fuel ratio because of variations in the sensors themselves.

Under normal operating conditions, the engine is running at a near-constant speed with the throttle in a near-constant position. This can be at idle, or while cruising down the street or highway.

Under these conditions, the engine is doing a steady amount of work, well within its power capabilities. The ECU, therefore, is programmed to maximize fuel economy and minimize emissions. To this end, it uses closed-loop operation to try to optimize the amount of fuel being delivered to each cylinder.

So what does this mean??

Well, the ideal air/fuel mixture is 14.7:1. From the MAS information, the ECU knows the mass of air entering each cylinder, and can determine the correct amount of fuel to deliver. It translates these into injector open and close times to control the fuel delivery.

Good enough? Well, not really. Variations in injectors, fuel pressure, and the sensors themselves (all of 'em) can lead to significant inaccuracies in this system. So the injector open time that the ECU has figured out is really not much more than an educated guess.

So what happens? Well, the ECU will fire the injectors, and then check the oxygen sensor output to see if the fuel delivered was too much, or too little. As described before, the oxygen sensor output changes sharply at the special 14.7:1 point. If the ECU is a little off in one direction or the other, the oxygen sensor feedback signal changes a lot, letting the ECU know that it missed the correct fuel amount.

Knowing this information, the ECU can change the next set of injector times to be a little closer to ideal. If too much fuel was delivered, the times are shortened; too little, and the injectors are fired a bit longer. The ECU will check the oxygen sensor output again after this round, readjust the injector times, and repeat over and over again. In this manner, the ECU is continually adjusting the fuel quantity to try to get it just right, and hold that oxygen sensor signal at the magic point of 500 mV.

Owners of air/fuel gauges will know, however, that the ECU never gets it 'just right'. In fact, the oxygen sensor signal changes so much around that perfect 14.7:1 point that the ECU cannot ever hit the 500 mV point and hold it steady. Instead, the ECU keeps hitting around that point, never quite getting it right.

In more detail, what happens is the ECU is a little off one way (too rich by a bit), and the oxygen sensor signal shoots up. The next round, the ECU trims the fuel down a little to fix it. This time, the ECU is off a little the other way (a tad too lean), and the oxygen sensor signal shoots down. The ECU trims the fuel up a bit in response, and keeps missing in this fashion all the time. This is what leads to the 'bouncy' oxygen sensor signal output that all A/F gauge users are familiar with. Far from being a sign of malfunction, this up-and-down signal behavior is perfectly normal and, in fact, is an excellent sign of a healthy oxygen sensor.

To sum up, closed-loop operation means that the ECU uses the MAS information to determine an initial amount of fuel to deliver. The oxygen sensor signal provides feedback to help correct this fuel amount for variations in the individual sensors and engine. This leads to better fuel economy and emissions, as there is very little wasted fuel. Closed-loop operation occurs when the engine is running at a near-constant speed and a near-constant throttle position.

So closed-loop operation works great! Why change it?

Chapter 4: Open-loop operation on DSMs

As mentioned before, closed-loop operation is designed to produce maximum fuel economy under constant operating conditions, When conditions are not constant, fuel economy may not be the primary goal.

As owners of sports cars, DSM owners expect a certain amount (read: lots!) of power from their cars. And, as any racer will tell you, power and fuel economy are uneasy bedfellows.

In fact, in order to get maximum power from the engine, you have to make sure to keep it cold. In fact, even running the ideal 14.7:1 air/fuel ratio under high load, high RPM conditions can be a very bad thing for your engine. You want to make certain the engine runs cool enough to prevent it from self-destructing.

One way of achieving cooler temperatures is to inject more fuel into the cylinders. This extra fuel acts as a coolant, reducing combustion temperatures. The cooler temperatures also reduce the possibility of the air/fuel mixture igniting in the cylinder before the spark plug is fired, a condition called *preignition* or *knock*. This lets the engine controller advance the timing on the engine, resulting in more power.

To achieve this cooling effect, the programmers of the DSM ECUs decided to program in very conservative values for the required amount of fuel, which guarantees the (stock) engine will run safely no matter how hard the car is driven. As a result, the ECU is only concerned with delivering enough fuel to each cylinder - how much more than enough is not a concern.

To this end, the ECU still determines the amount of fuel to deliver from the signals delivered by the MAS. Because the engine is under acceleration, however, the fuel quantity selected by the ECU is larger than before. Also, since the fuel amount is guaranteed to be enough (by design), the ECU no longer checks the oxygen sensor to see how close it was to the 'correct' output. (The signal is still there - the ECU simply doesn't bother doing anything with it.)

This failure to check the oxygen sensor for sufficient fuel is the difference between open-loop and closed-loop operation. While under 'normal' conditions, the ECU checks the oxygen sensor for the results of its own actions - monitors the feedback, in other words. While under acceleration, the ECU does not use the feedback information, and is 'demoted' to open-loop operation for the duration of the acceleration period. When the acceleration ceases, the ECU returns to closed-loop operation once again.

While running in open-loop operation, the ECU is designed to provide an air/fuel ratio of about 11:1; less air, more fuel. (This will change somewhat with RPM and airflow, but that's close enough at the moment.) This is well above the switch point for the oxygen sensor. Consequently, the oxygen sensor remains fairly constant during open-loop operation, on the rich side. Viewing the signal with a voltmeter or air/fuel ratio meter will show a very high, relatively steady oxygen sensor output, in contrast to the 'bouncy' signal seen during closed-loop operation. This signal is normal and, in fact, eminently desireable - the higher the better (to a point).

Just as an aside, the ECU designers went one better than this in the cooling department. While under high loads, the ECU actually fires *all four* injectors *in between* the normal firing times. This extra fuel has no chance of being burned in the engine (at least until the valves open), but helps to keep things cool.

The behavior of the ECU during open-loop operation has serious implications for those who modify their engines. More on this in a bit.

Chapter 5: Determination of open-loop vs. closed-loop operation

So how does the ECU decide when 'normal' conditions exist and when 'acceleration' conditions prevail? The primary trigger that switches the ECU from one control method to another is a simple lookup table, which compares throttle (accelerator) position against engine speed. These signals are both provided to the ECU from various sensors in the engine. A second trigger is change in throttle position (delta-TPS) - a large enough change in throttle position over a short enough period of time will also cause the ECU to enter open-loop mode.

Closed loop operation occurs when the engine is idling, or running at near-constant throttle position at near-constant engine speed. During idle, the ECU actually has two goals: one, to deliver the correct fuel, and two, to maintain the idle speed at 750 RPM. The fuel is handled by controlling the injectors and using the oxygen sensor for feedback. The idle speed is controlled by the idle speed control motor (ISC), using the crankshaft angle sensor as feedback.

[Why the crankshaft angle sensor, and not the tachometer? The tachometer signal is generated by the ECU, not from the engine. The ECU determines RPM by watching the crankshaft angle sensor, and runs the tachometer in the instrument cluster.]

Relatively large changes in throttle position at a given RPM will cause the ECU to change to open-loop operation. Once the throttle position and RPM are back "in sync" with each other, the ECU will generally switch back to closed-loop operation. This can occur if the throttle position moves again, or if the engine speed begins to level out at the new throttle position. However, there are some special cases that may change this behavior.

Above a certain RPM level, the ECU will run open-loop regardless of throttle position, except when the throttle is completely closed (full deceleration). Those with tendancies towards high highway cruising speeds will have noted the large decrease in fuel economy above certain speeds. Above a certain RPM, the ECU is providing all that extra gas to the engine, some of which is blown straight out the exhaust pipe. As you might imagine, this radically decreases fuel economy.

The critical RPM depends on the model of car, air temperature, atmospheric pressure, and a few other factors. It is identical for front-wheel-drive turbo DSMs than all-wheel-drive turbos, and tends to hang around 4000-4500 RPM. Because FWD cars are geared slightly differently than AWD cars, FWD cars will achieve a slightly higher cruising speed before hitting open-loop operation than AWD cars.

In the event of a sensor failure, the ECU may have no choice but to use the open-loop method in order to keep the engine running. An excellent example of this is when the MAS air volume sensor fails. Although the temperature and barometric pressure sensors may continue to function in the MAS, the ECU now does not know the volume of air is entering the engine. This causes the ECU to run in open-loop mode, which provides the guaranteed minimum required amount of fuel and the maximum safety for the engine.

In this case, the process is a little more sophisticated than that. In order to have some idea of the air volume, the ECU uses a small pre-programmed table to guess at the correct airflow based on throttle position and RPM. Otherwise, the engine could not run at all. This table, of course, is programmed conservatively in order to safeguard the engine. This value is then used to further processed to pick the open-loop fuel quantity. Still, the process is open-loop, as the ECU does not use the oxygen sensor signal as feedback to regulate fuel flow.

Chapter 6: Open-loop operation on modified vehicles

As mentioned before, the open-loop operation of the ECU under acceleration conditions is 'guaranteed by design'. This is only possible if all the assumptions that went into the ECU programming are true. Of course, since the designers were programming for a factory car, they took the design parameters from a stock car and built them into their programming.

In order to guarantee that sufficient fuel reached the cylinders, the programmers had to consider fuel pressure, injector flow rate, injector dead time, spark time, spark duration, valve timing and a host of other engine operational parameters, down to such minutae as battery voltage. Using these values, they figured out how much fuel could be delivered in a given time, added a hefty safety margin and based their ECU program on that.

This is great, until the owner of the car begins to modify the engine to produce more power. By doing this, at least one (probably more) of the engine operational parameters is changed, and the ECU program is no longer as accurate as it was. Under closed-loop control, the changes will be detected by the ECU, which will automatically compensate for the altered conditions and maintain the correct air/fuel mixtures.

Open-loop control, however, is a different story. Without feedback control, there is no way for the ECU to compensate for (or even be aware of) the modifications and the changes they have wrought to engine operation. If enough modifications are done, it is possible that the ECU will simply not provide enough fuel to the engine, resulting in more generated heat. Enough extra heat can easily cause severe damage to an engine.

We have already seen that the ECU begins open-loop operation when the car is under acceleration. This is perhaps the worst time for racers, who often run at full accleration for prolonged periods of time. Since the ECU does not do it, it is up to the operator to monitor the engine condition and make sure that enough fuel is getting into the engine. The factory provided no gauges to achieve this, and it is necessary for the driver to have auxiliary instruments installed that let him or her monitor the important aspects of engine operation.

The most popular gauge is an air/fuel gauge. This handy little device is essentially a voltmeter with a display. When hooked into the oxygen sensor signal (the same one the ECU itself uses) the operator can watch the oxygen sensor signal to be certain that the engine has sufficient fuel. Other possible instruments include an exhaust gas temperature (EGT) gauge, an injector duty cycle monitor, and a pressure (boost) gauge to directly monitor the turbocharger.

When it comes to how rich DSMs must run during full acceleration, the problem isn't a lack of information but a surplus of information. That is, there is lots of information out there, but it's all different. There are many reasons for the discrepancies, including differences in location, weather conditions, individual automobiles, fuel quality, measuring instruments and techniques, and so forth. Because every setup is different, there is no 'magic' number that can definitively be stated as the 'perfect' number for DSMs.

However, there have been some **general** guidelines laid down over the years. Authorities in the field have stated that DSMs must run an absolute minimum of 850 mV from the oxygen sensor. Those who use EGT sensors have recommended a maximum operating temperature of less than 1700 degrees Farenheit (925 degrees Celcius). Different owners have different preferences, but the majority tend towards oxygen sensor readings of at least 900 mV, and EGTs below 1600 degrees F. Several racers have expressed a serious preference for even more conservative values. However, it is possible that some cars under some conditions would be able to safely exceed these values without harm.

Chapter 7: DSMs and fuel cut

Turbochargers are positive feedback devices - the faster they spin, the more air they push into the engine. The more air they push into the engine, the more exhaust gases are produced. The more exhaust gases that are produced, and the faster the turbo spins. The faster they spin....

See the problem? The faster they go, the faster they go! Eventually something would have to break. To prevent this situation from occuring at all, turbos have a device on them called a wastegate that limits the turbo speeds to manageable levels. This is a pressure-activated valve that opens at a preset pressure and vents additional intake pressure either to the atmosphere or back into the exhaust stream.

Modern turbo design notwithstanding, it is still very possible for the wastegate on the turbo to malfunction, leading to exactly the ever-faster turbo behavior described above. Recognizing this, the designers of the DSM ECU built a fail-safe limiting mechanism into the DSM ECU modules. And it goes like this:

Should the ECU ever see an intake air mass greater than a certain preset level, it will stop fuel delivery and spark to the engine cylinders. This 'critical mass' depends on air volume, temperature and pressure, as measured by the MAS sensors. The act of the ECU cutting off fuel delivery to the engine is known as the infamous 'fuel cut'.

Why'd they do this?? Well, should the wastegate ever malfunction in such a way as to allow the turbo to spin out of control, the ECU will stop firing the cylinders to save the turbo from exploding. This is not the only method the ECU has to limit the turbo operation, but it is the last and most desperate. The mass air limit at which the ECU will fuel cut was based off of the maximum amount of fuel that the stock fuel pump and injectors could deliver to the engine cylinders.

Unfortunately, those owners who purposefully modify their engines to provide greater power are deliberately forcing more air through their engines in order to produce more power. Recognizing the limitations of the stock fuel delivery system, serious modders usually upgrade the fuel pump and/or injectors in order to provide enough fuel capacity for the hopped-up engine.

However, the fuel cut limit, based on the stock fuel system, is pre-programmed into the ECU and cannot be changed. There is no method by which the ECU can be made aware of the improvements made to the fuel system, so it will blithely cease fuel delivery once the mass air intake reaches a predetermined level *regardless of how much fuel capacity is really available*. This characteristic of the DSM ECUs, while understandable, has been a source of nearly endless frustration for power-hungry owners.

There are solutions to this problem. Aside from reprogramming the ECU (which is sometimes possible) to eliminate the fuel cut, all of these solutions utilize the same basic principle - fool the ECU into believing there is less air entering the engine than, in fact, there is. This is often accomplished by altering one or more of the input signals sent to the ECU by the various air measurement sensors present in the MAS.

For example, a volume of air at higher temperature contains a smaller mass of air than an equivalent volume at lower temperature. The ECU determines the intake air temperature from a sensor in the MAS. Electrically altering this signal can make the ECU read a higher air temperature than actually exists. Based on this erroneous temperature, the ECU will calculate the mass air intake as being less than the actual amount.

Even expensive fuel control systems utilize this technique to keep the ECU from activating fuel cut. These systems, however, often intercept more than one ECU signal, and can provide their own monitoring equipment to allow the operator to properly observe the condition of the operating engine. This allows the owner to 'fine-tune' the car to achieve the desired characteristics while keeping the engine safe from harm.

Obviously, doing this type of modification is a relatively dangerous process. By changing the ECU inputs, the operator runs the risk of altering both open and closed-loop ECU operation to the point where the engine does not receive enough fuel. As always, it is up to the operator to monitor the engine operation to ensure that a damaging situation does not occur.