Injector Dynamics Spreadsheet - Instructions

The Injector Dynamics spreadsheet is a powerful tool. To get the most out of it you should understand how the data is generated, and what it all means. Many of you reading this already understand the basic operation of a fuel injector, but will likely find new information in the text which will broaden your understanding, and help you to get the most from the test results.

Basic Injector Operation

At first glance, the operation of a fuel injector seems quite simple. We apply voltage to a coil, and the magnetic field generated by the coil pulls a disc, pintle, or ball off its seat allowing fuel to flow. In other words, it is an electronically controlled on/off valve.

To find the flow rate of the injector, we can supply pressurized fuel at the inlet, apply voltage to the coil, and measure the amount of fuel that flows through the injector in a given period of time. This is called static flow testing, and the result is the static flow of the injector at the stated test pressure.

This type of testing has some validity in the sense that we can clearly show that one injector flows more or less than another, but its ultimate usefulness is limited.

Once we install the injectors on the engine, we will not be holding them in the wide open or static position, we will be cycling them on and off rapidly. This is a dynamic situation, and the dynamic injector characteristics are far different than what the static flow test may predict. For this reason, static flow testing has little value to the experienced tuner.

To gain a better understanding of an injector’s dynamic flow characteristics, consider the following.

Let's say that we perform a static flow test on an injector, and find that it flows 100cc's in one minute. This information is valid, but it tells us nothing about the dynamic characteristics of the injector.

Consider what happens if we turn the injector on and off at a high rate of speed, let's say 100 times per second, and measure its flow for that same period of one minute. Obviously, the measured flow would be less than the static flow, because the injector would only be flowing fuel for a portion of a minute.

To be more exact, let's say that while the injector is cycling on and off, it spends equal amounts of time turned on, and turned off.

In this scenario, the injector would be described as running at a 50% duty cycle, which simply means that it spends 50% of its time in the on state. (More on this later)

In theory, our measured flow should equal 50% of the static flow, since the injector is turned on 50% of the time, right?

Wrong! Why you ask? As usual, when theory meets practical application Mother Nature rears her ugly head, and the immutable laws of physics get in our way.

In practice, we find that three very important things happen which result in the measured flow being less than what we estimated based on our static flow numbers.

The first is that the magnetic field which pulls the pintle off the seat takes a finite amount of time to build. While this magnetic field is building, the pintle is sitting still on its seat, and so we have an injector with voltage being applied, but absolutely no fuel flow!

Once the magnetic field becomes strong enough to pull the pintle off its seat, we run into issue number two, which is the fact that the pintle is in no big hurry to move! Instead of snapping immediately to its full open position, it lazily moves from the fully closed position (No fuel flow) to the fully open position. (Full fuel flow)

Instead of having an injector that is on 50% of the time, and off 50% of the time, we have an injector that spends a portion of its on time flowing no fuel, another portion of its on time at partial flow, and the final portion of its on time at full flow.

As a result of all this, the actual fuel flow from our injector will be less than that predicted by static flow testing.

Injector Dead Time

So now we can go to the spreadsheet for a visual representation of static vs. dynamic flow characteristics. Open the Injector Dynamics spreadsheet, and click on the Dead Time Calculator tab at the bottom of your screen. Note the difference between the red and blue lines on the chart.

Note: To the right of the chart you will find four buttons, two with an up arrow, and two with a down arrow. Click the largest of the up buttons several times to move the yellow line out of the way.

Looking at the chart, we see a group of numbers on the left side. These numbers are the scale for the Y axis of the chart, and specify flow in cc's per minute. There is also a group of numbers at the bottom of the chart. These numbers are the scale for the X axis, and specify the pulsewidth applied to the injector in milliseconds.

The blue line represents your injector if it had zero opening delay. This theoretically perfect representation is based on the static flow, and will flow 50% of static at a 50% duty cycle. (Or 30% of static at a 30% duty cycle...)

The red line represents the actual fuel flow of your injector cycled on and off 100 times per second. (Or my injector, if you have downloaded the sample Injector Dynamics spreadsheet.)

The difference between the theoretically perfect flow, and the actual flow is the result of the "lazy" injector opening sequence described above.

As you may have guessed, there is still more to the story, but before we get to the third and final issue which keeps our injector from having theoretically perfect flow, we need to gain a better understanding of pulsewidths, and duty cycle.

Unfortunately, this requires a bit of math. (Don't turn off your computer yet, it's actually quite simple.)

For starters, we stated in our example that we would be turning the injector on and off 100 times per second. This rate of opening and closing is called the frequency, and is normally stated in Hertz or Hz for short. In this case, 100Hz.

We can determine how long each cycle last by dividing the number of cycles into 1 second. 100 divided by 1 equals .01, which means that each cycle lasts .01 seconds. This is called the period. (The length of time required to complete one cycle.)

In our example we also stated that we would be running the injector at a 50% duty cycle. From this, we can calculate the pulsewidth.

We start by taking 50% of our period (Period of .01 times .5) and we get .005 seconds.

To convert this to milliseconds (1/1000th of a second) we multiply this by 1000, and we get 5 which is our pulsewidth in milliseconds.

To simplify: Period in seconds = 1 divided by the frequency in Hz.
Pulsewidth in seconds = Period times duty cycle.
Pulsewidth in milliseconds = Pulsewidth in seconds times 1000.

You should now have a clear understanding of the chart, and be able to determine duty cycle from the actual pulsewidth, and vice versa.

So now we get to the third and final issue that keeps our injector from having perfect dynamic flow characteristics.

When the voltage is released the pintle again moves lazily, this time from the fully open to the fully closed position. The result is that the injector continues to flow a small amount of fuel after we have removed voltage from the coil.

As opposed to what happens at the beginning of the cycle, the slow pintle movement on closing results in increased flow rather than decreased flow, but not nearly enough to offset the reduced flow on opening, and so the injector still flows less than we tell it to.

Now that we understand why an injector's dynamic flow characteristics are not perfect, let's consider how this would affect our ability to tune an engine.

Let's say that we are tuning an engine on the dyno, and we have a target air fuel ratio of 11.5:1. After the first pass, we see that at some particular point on the map, the air fuel ratio is actually 12.5:1, and we want to richen it up. One approach would be to richen it by some arbitrary amount, and make another pass to see where we end up.

A better approach would be to quantify our mixture in Lambda rather than air fuel ratio, as this allows us to make easy calculations. Converting our air fuel ratios to Lambda, we get .850 for our existing air fuel ratio, and.782 for our target air fuel ratio.

Note: Lambda = Actual air fuel ratio divided by stoichiometric air fuel ratio.

Being clever tuners, we realize that dividing the target Lambda into the existing Lambda will give us a ratio that we can use to make an exact change in pulsewidth which will give us the air fuel ratio that we are looking for. So we divide .85 by .782 and get 1.087 meaning that we need 8.7% more fuel to achieve our target Lambda value.

Looking at our map, we see that our current pulsewidth is 5 milliseconds, and so 5 milliseconds times 1.087 equals 5.44 milliseconds.

We change the map, make another pull, and find that the mixture is now much richer than we want it to be. Why? Because we were not actually getting 5 milliseconds worth of fuel to begin with, and so we applied a correction to an invalid number.

Our ECU told the injector to open for 5 milliseconds but it was effectively "dead" for a portion of that time, and so our effective pulsewidth was less than the actual pulsewidth.

When we increased the pulsewidth, the dead time of the injector remained the same, but the pulsewidth was longer, and so the dead time occupied a smaller portion of the pulsewidth.

To make this clearer, let’s put some numbers on it.

Let’s say that our injector dead time is .7 milliseconds. If we subtract .7 milliseconds from our original pulsewidth of 5 milliseconds, we are left with 4.3 milliseconds which is our effective pulsewidth.

After making the change to richen the mixture, we had a pulsewidth of 5.44 milliseconds. Subtracting the dead time of .7 milliseconds, we were left with an effective pulsewidth of 4.74 milliseconds.

The result is that we increased the actual pulsewidth by 8.7%, but the increase in effective pulsewidth was 10.2% (4.74 divided by 4.3 equals 1.102)

Since the dead time is fixed, as the pulsewidth is decreased, the dead time occupies a greater portion of the pulsewidth, and so we find that the shorter the pulsewidth, the greater the error.

(If you’re still not convinced, consider that same dead time of .7 milliseconds applied to a 2 millisecond pulsewidth!)

The end result of the all this is that we can only guess, because the injector is not delivering the amount of fuel that we tell it to. Fine you say, I'll just tune until it is right, and who cares what the actual numbers are, as long as the air fuel ratios are correct!

That would be just fine if we were tuning in a climate controlled dyno room where the barometric pressure and temperature remained constant, but what happens when we get the motor to the track, or on the street where the atmospheric conditions are different?

Our ECU will sense the change in air density, say 10%, and automatically make a 10% correction to the pulsewidth just like we did manually on the dyno. And just like our tuning session on the dyno, the actual change in fuel flow will be greater than what the ECU asked for. So now we have an engine that will only run at the proper air fuel ratio when atmospheric conditions are exactly the same as on the dyno.

We may as well have a carburetor!

So are you discouraged yet? Don't be, there's a simple answer.

The injector dead time is fixed, at least within the linear range of the injector, and can be accounted for if we know exactly what it is.

Dead Time Compensation

The manufacturer of your ECU realizes that injector dead time exists, and they have given you a way to deal with it. All quality ECU's give you the option to enter an injector lag, dead time, offset, battery compensation, etc.

Wait…did I say battery compensation???

I sure did. There are still a few more pieces to the puzzle. Both the amount of time required for the magnetic field to build, and the ultimate strength of the magnetic field are dependant on the voltage supplied to the injector.

As a result, the injector dead time increases as voltage is decreased. This is no problem, because the ECU manufacturers have taken this into account, and you are expected to enter the injector dead time for various battery voltages.

By including the injector dead time in the ECU’s internal calculations, it allows us to get the amount of fuel that we are asking for. If we want 10% more fuel, we get 10% more fuel. The ECU simply adds the injector dead time to the requested pulsewidth, resulting in accurate fuel delivery.

Once the dead time has been determined through dynamic flow testing, the ECU can properly account for changes in atmospheric conditions, and we will spend far less time chasing our tail when tuning. As an added benefit, our closed loop control will now work much more effectively because when the ECU senses a 5% error, it can accurately apply a 5% correction.

If our injector dead times are incorrect, the closed loop control will spend a fair amount of time hunting for the proper correction until it finally zeroes in on the requested air fuel ratio. Unfortunately, operating conditions may change (Due to moving to a different point on the map) before the closed loop control ever gets it right.

The end result is that we curse the ECU for having poor closed loop control, when the problem really originates with incorrect injector dead times.

Once the dead times are correct, we find that tuning becomes much easier, our air fuel ratios are more consistent, and our closed loop control is now a powerful feature.

So where do we go from here? Finally, we get to play with our graphs and get to some real world information relating to your car.

Go back to the Dead Time Calculator in the Injector Dynamics spreadsheet.

We already know what the blue and red lines are; now let's get to the yellow line labeled Calculated Flow. The yellow line represents the theoretically perfect flow of the injector, minus an applied dead time. The applied dead time is shown in the text box directly above the up arrow.

By clicking the up and down arrows, we can increase, or decrease the dead time applied to the "perfect" flow of our injector. This calculated flow is what the ECU is expecting based on the injector dead times that we enter.

The large buttons on the left change the dead time in 100 microsecond increments. The smaller buttons on the right change the dead time in 10 microsecond increments.

Note: For better resolution, the dead time is displayed in microseconds. A microsecond is one hundred thousandth of a second. Most ECU's will require the dead time compensation to be entered in microseconds, rather than milliseconds.

Microseconds times 100 equals milliseconds.
Microseconds times 100000 equals seconds.

By adjusting the dead time until the yellow line (Calculated Flow) sits on top of the red line (Actual flow) we can determine the dead time of the injectors. After playing with the up and down buttons a few things should become obvious. The first is that while the injector is fairly linear through most of its range, it deviates abruptly at either end of the scale.

This is because all fuel injectors have non linear flow at the extreme ends of their range. This is very important information to consider, but before we get to that let's continue looking at the dead time within the linear flow range of the injector.

The next thing you may notice is that even through the linear range, the injector has some slight deviances from a straight line, and you're probably wondering exactly where to place the yellow line.

To simplify this, the text box above the dead time shows the percent average dead time error. This is averaged across the linear flow range of the injector. The linear flow range is defined as the range that deviates no more than 2.5% from a straight line.

To find the most accurate dead time compensation for the linear range, simply vary the dead time until you get the smallest error.

For those of you who want to see what type of errors you may see outside the linear flow range, we have included another graph which shows the dead time error vs pulsewidth.

If you click on the Dead Time Error tab, you will see the actual percentage dead time error vs. pulsewidth throughout the range. Note that because the slope of the line changes in the non linear flow range, the dead time can only be correct at one single point in the upper and lower non linear flow range of the injector.

It is also worth noting that throughout most of the non linear flow range the slope of the line is steeper than within the linear range, and so the ECU will over compensate. This means that if you ask the injector for a 10% change in fuel flow, you may get 15%

At any point on the curve where the slope of the line is shallower than within the linear range, the ECU will under compensate. This means that if you ask the injector for a 10% change in fuel flow, you may only get 5%.

Note that this does not mean rich or lean, it simply means that the actual change in fuel flow will either be more or less than you, or your ECU asks for.

For instance, if your system is over compensating, and you ask it to go a little bit leaner, it may go a lot leaner, possibly damaging your engine!

The moral of the story? Stay out of the non linear flow range of the injector, or you will lose control of your air fuel ratios.

An engine running very large injectors will often idle in the lower non linear flow range. The result is that you have very poor control of your air fuel ratios at idle. This may be irritating, but if your engine only has to idle well enough to get up to temperature before you hit the track, it can be tolerated.

On the other hand, if you are operating in the upper non linear flow range of the injector, you could quickly get into trouble, and possibly damage your engine.

Note: There are other charts included in the Injector Dynamics spreadsheet designed to keep you out of this trouble area. We will get to those shortly.

That pretty much does it for the Dead Time Error chart, so now let’s go back to the Dead Time Calculator chart. There is one more feature located above the averaged dead time error box. This drop down box lets you choose the battery voltage.

Your injector was flow tested throughout a range of voltages to insure that your dead time compensation is correct regardless of the voltage seen at the injector.

If you click on the Flow vs. Voltage tab, you can see the actual raw flow data vs. battery voltage. This represents the average of all your injectors flowed at various voltages.

We will not be using this chart for anything in particular. The raw data is used for the calculations, and is only displayed to give you a visual indication of your injectors flow vs. voltage characteristics.

Back to the Dead Time Calculator, change the voltage from 14 to 12, and notice that the red and yellow lines no longer lay on top of one another. This is because the injector dead time at 12 volts is greater than the injector dead time at 14 volts. Now adjust the dead time again with the up and down arrows until you get the smallest error, and you are done!

The final control located on the Dead Time Calculator chart is the Set button. When you click this button, the applied dead time will be recorded in the appropriate box on the Data sheet. The values in these boxes are used to generate the Dead Time Compensation Curve chart, and to give you a simple display of the dead times vs. voltage to make it easier to enter the data into your ECU.

Recovery Time

Remember earlier when we discussed the pintle slowly moving from full flow to no flow at the end of the cycle? Would you be surprised if I told you that there was more to the story?

Probably not.

In addition to the pintle moving slowly, there is a certain amount of time required for the magnetic field in the injector coil to fully collapse. As a result of these two inter related issues, the injector needs some time to “rest” before we tell it to open again.

If we do not give the injector sufficient time to recover, it responds in a very non linear manner as you can see on the graph.

If you run your injector in this upper non linear flow range, you are begging for trouble. Many of you have probably already found trouble, and that is why you are reading this.

As far as we know, none of the current crop of ECU's allow you to account for this, and so we avoid trouble by simply staying out of this range.

So let me say this loud and clear. **Stay away from the upper non linear flow range of your fuel injector!!!**

The 80% Duty Cycle Myth

So now we get to the fun part of the spreadsheet which is making horsepower predictions for your injectors. Note that we can make fairly accurate predictions of the horsepower potential of your injectors, but that doesn't mean you will ever get there. It only means that your injectors will supply enough fuel to allow for that amount of horsepower.

In the past, the horsepower potential of a set of injectors was calculated by taking the static mass flow at 80% duty cycle, divided by the BSFC of your engine. As you have already seen, the dynamic flow at 80% duty cycle is no where near 80% of static flow.

Additionally, the recovery time of your injector requires that you give it sufficient time to rest, and at high rpm the recovery time can take up a considerable percentage of the total cycle time. That means that you either enter the non linear upper flow range of your injector, or you limit your duty cycle to a number that gives the injector a sufficient recovery time.

In either case, the horsepower potential of your injectors can be a long ways off from the 80% duty cycle prediction.

Still, the 80% rule has become gospel, and most of us assume that if we "follow the rule" we will stay out of trouble.

It's time to change all of that.

Let's start by putting some real numbers on things. Let's say for instance that you have a V8, and your ECU only has 4 injector outputs. To run all eight injectors, you must wire, and fire them in pairs. To do this, you must fire the injectors once per engine revolution. This is typically referred to as a "batch fire" system.

Now let's say that you will be running this V8 all the way up to 9,000 RPM. This may seem like a stretch, but I'm sure that someone reading this is revving their batch fire V8 this high, so we will use this for our example.

If we divide 9,000 revolutions per minute by 60, we get 150 revolutions per second. This means that we will have 150 complete cycles per second which as you remember is a frequency 150 Hz.

Hopefully you also remember that if we divide 1 second, by 150 Hz, we get the period, which is .0067 seconds, or 6.7 milliseconds. That means that we have only 6.7 milliseconds to inject fuel into the engine and let the injector recover before we fire it again. So let’s say that our injectors have a recovery time of 2 milliseconds. If we subtract the 2 millisecond recovery time from our period of 6.7 milliseconds, we are left with only 4.7 milliseconds to inject fuel into the motor. This is our maximum allowable pulsewidth if we are to stay within the linear flow range of our injectors.

If we divide our period of 6.7 milliseconds into our maximum allowable pulsewidth of 4.7 milliseconds, we get a duty cycle of only 70.2%

Once our duty cycle exceeds 70.2%, we have stepped over the line, and the injector will respond in a very non linear manner. Add to that the fact that a portion of our injector on time will be consumed by the dead time, and we end up with a very short effective pulsewidth.

To put a few more numbers on this, let’s say that our injector dead time is .7 milliseconds. This leaves us with an effective pulsewidth is only 4.0 milliseconds, and our injector is only flowing 59.7% of its static flow!

In case you missed that, here’s the math. The period is 6.7 milliseconds. From this we subtract the dead time of .7 milliseconds, and the recovery time of 2 milliseconds and we are left with a pulsewidth of 4.0 milliseconds. Then we divide our period of 6.7 milliseconds into our pulsewidth of 4.0 milliseconds and we get .597. We then move the decimal point two places to the right to get 59.7%.

If you’re stubborn like me, you look closely at your Dead Time Calculator chart, see that the injector flow increases dramatically just outside of the upper linear flow range, and realize that the mixture will go rich which is not dangerous at all.

You are entirely correct, but what happens if you dyno tuned the engine in that non linear flow range, and then you get to the track, atmospheric conditions change, and the ECU tries to compensate by leaning the mixture slightly?

You guessed it, the mixture leans substantially.

On a good day, you may just make less than optimum horsepower. On a bad day, you might run over your own crankshaft.

This is worth saying again. **Stay away from the upper non linear flow range of your injector!**

Staying Out of Trouble

So now let's go to the section of the Injector Dynamics spreadsheet designed to keep you out of trouble.

Click on the Max Duty Cycle tab in the injector dynamics spreadsheet.

The axes are labeled clearly, and should require no explanation. There are 2 lines on the graph. The red line labeled Maximum Linear Duty shows the absolute maximum duty cycle vs. rpm. If you exceed the duty cycle shown at any corresponding rpm, you have stepped over the line, and the injector flow will become non linear.

The blue line labeled Recommended Max Duty gives you a 10% safety factor. This means that if your duty cycle is high enough to meet the blue line, you can still increase fuel flow by 10% before the injector output becomes non linear.

It is recommended that you do not exceed the duty cycle shown by the blue line while tuning on the dyno. This gives you a 10% safety factor in case the atmospheric conditions at the race track require richening the mixture.

Note that as explained earlier, the injector recovery time is fixed and not duty cycle or rpm dependant and takes up a greater amount of the available period as rpm increases. As a result, the maximum allowable duty cycle decreases as rpm increases.

Lets play around with this a bit so that you can become more familiar with the spreadsheet, and maybe learn a few interesting things while you are at it.

If you roll your mouse pointer on to either of the lines where they cross one of the x axes, a window pops up which shows the actual value. This is handy, and easier on your eyes than trying to decide if the line is at 82%, 83% etc.

This popup window is active on all the charts in the Injector Dynamics spreadsheet.

Now click on the data tab, and note the box labeled firing arrangement. Try changing the firing arrangement, and see how it affects the maximum allowable duty cycle. By changing the firing arrangement from batch to sequential, the available period is doubled for a given rpm, and so the recovery time of the injector takes up less of the period. The result is that there is more time available to inject fuel, and a higher allowable duty cycle.

Now change the recovery time shown on the Data page and note how it affects the maximum allowable duty cycle.

Cool huh?

Horsepower

Now click on the Max Horsepower tab to see what your injectors are capable of.

Again, the chart axes are clearly labeled and should require no explanation. The upper box shows the specific gravity of the fuel, and can be changed with the up and down arrows.

This is a very important parameter, because a fuel injector is a volume flow device, not a mass flow device. It couldn't care less what mass of fuel it is flowing. To know the actual mass fuel flow into the engine, we have to know how dense the fuel is. In this case we are stating the density in grams per cubic centimeter which is referred to as the specific gravity.

Give this a few clicks, and see how it affects the horsepower output.

Note that running fuel with a higher specific gravity does not mean that you will make more horsepower; it simply means that a lesser volume of fuel will be required to make a given amount of horsepower.

The second number you see is the Brake Specific Fuel Consumption. This can vary from as low as .35 at the torque peak on a high compression race motor, to as high as .75 on a turbocharged rotary running pump gas.

The BSFC is the mass of fuel in pounds required to make one horsepower for a period of one hour. That complicated sentence boils down to this. Fuel flow in lbs per hour divided by BSFC equals horsepower.

Unfortunately, we cannot give you any guidelines here because all engines are different. (Hopefully you have previous dyno numbers, and have a fairly accurate estimate.)

Like the Max Duty Cycle chart, the red line represents the absolute maximum while staying within the linear flow range of your injectors, and the blue line represents what is allowable with a 10% safety factor.

If you enter accurate Specific Gravity, and BSFC numbers, the horsepower predictions should be quite accurate.

Injector Peak and Hold Currents

If you go back to the Data page, you will see a peak, and hold current listed separately. This describes the peak current reached during injector opening, and the hold current used to keep the injector open for the remainder of the cycle. These are the current settings used during the flow testing process.

If you have high impedance injectors, these boxes will display NA, because this doesn't apply to your injectors.

Both the peak and hold current settings have been optimized for the broadest possible linear flow range. Some ECU's, like the Motec MX00 series give you complete control of both peak and hold current. Others, like the Motec MX, and MXX series only allow you to specify the peak current, with the hold current set at 25% of the peak. Still others like the Haltech E8 for give you a few choices of peak current settings covering a broad range, with the hold current set at 25% of the peak.

As you can see, there is great variation from one ECU to another, and this is why we need to know what ECU you are using before we test your injectors.

The first step in the dynamic flow testing process is to find the optimum peak and hold currents based on the options available in your ECU. This is an important step, because it has a substantial effect on the width of the linear flow range of your injectors.

To explain why, we first need to understand the operation of a peak and hold drive circuit (Used for low impedance injectors) and a saturated drive circuit (Used for high impedance injectors)

Let's start with a saturated drive circuit. A saturated drive circuit typically drives an injector with a resistance in the 10 to 16 ohm range. The saturated drive circuit is a simple on/off switch that interrupts the ground while the injector receives battery voltage at all times.

This “switch” is normally in the form of a MOSFET or bipolar transistor, but for the sake of this conversation, we will call it a switch.

When the switch closes the injector the coil has both a ground, and battery voltage, and so the coil is energized, and the pintle lifted off the seat.

At the beginning of the cycle, the current builds slowly, finally reaching its maximum at which point the coil is “saturated” This current is held constant until the switch is turned off at which point the magnetic field collapses, and the pintle drops back on to its seat. The maximum current is determined by the resistance of the injector, and can be calculated by dividing the injector resistance into the battery voltage. For instance, a 12 ohm injector supplied with 12 volts will result in a maximum current flow of 1 amp.

Note: This is over simplified because we are not considering the resistance of the switch, the wires, or any other resistances in the circuit. Still it is fundamentally correct, and will do for this conversation.

The disadvantage of this type of drive circuit is that the strength of the magnetic field used to pull the pintle off the seat is dependant on current, and the resistance of the injector limits us to a relatively low current level.

Saturated drive circuits generally result in large dead times, and poor control of the injector when the pulsewidths are short, because the pintle moves so slowly. This broadens the lower non linear flow range of the injector.

Additionally, the recovery time of the injector is influenced greatly by the amount of current flowing through the coil when it is turned off. Generally speaking, as current at the end of the cycle is increased, the recovery time is also increased which broadens the upper non linear flow range of the injector.

The end result is that a saturated drive circuit/injector typically has a much narrower linear flow range.

Note: There are many factors contributing to everything stated above, but again the description is fundamentally correct, and sufficient for this conversation.

So now we get to the peak and hold drive circuit which is typically used to drive an injector with a resistance in the 1 to 4 ohm range. The peak and hold driver is wired in the same way as the saturated driver, but uses a “smart switch” to interrupt the ground.

How smart is it? Well, here’s what happens. At the beginning of the cycle, the “smart switch” supplies a ground to the injector just like with a saturated drive circuit. The big difference here is that the injector resistance is low, and so the current builds rapidly, and to a higher level. This builds a stronger magnetic field in the injector resulting in the injector going from no flow, to full flow in a much shorter period of time. This makes the lower non linear flow range of the injector narrower, giving better control of the injector when pulsewidths are short.

If the driver were to continue supplying a ground to the injector for the rest of the cycle, the current would build to a very high level. Once the pintle has been pulled all the way open, additional current flow will only waste energy, and generate heat. Additionally, the amount of current required to quickly lift the pintle off its seat is far greater than that required to hold it there for the rest of the cycle.

This is where it gets interesting. Once the current reaches a pre defined maximum, the driver will then reduce the current to a lower pre defined level which is just high enough to hold the injector in the fully open position.

Since the current required to hold the injector open is quite small, the magnetic field collapses quickly, the pintle drops back on to its seat quickly, and the injector's recovery time is reduced.

This narrows the upper non linear flow range.

By optimizing these two parameters, the linear flow range of the injector can be made as broad as possible allowing the injector to have good control when the pulsewidths are short, and a shorter recovery time which results in a higher maximum allowable duty cycle.

Conclusion

So there you have it. Everything you need to know to effectively utilize the information, and give yourself a better more consistent tune. Probably even more important is that fact that you know your safe limits, and should be able to keep yourself out of trouble.

Authors Note: Every time I re-read the document, I think of things to add, or better descriptions of what already exists. For this reason the Injector Dynamics spreadsheet and this text will be evolving over time. If you have any suggestions or questions, please send an email to injectordynamics@yawpower.com

Have fun, and thanks for using our service.

Paul Yaw
Yaw Power Products