Application Bulletin
AB-24

Understanding Droop
and
Programmable Active Droop™

---Ron Lenk, Staff Applications Engineer
09/20/00

Programmable Active Droop is a trademark of Fairchild Semiconductor.

Summary

Droop adjusts a converter's output voltage based on its load current to optimize transient response and minimize the required output capacitance. Passive droop is done with a resistor. Active droop is available in the RC5057. Programmable Active Droopä makes the amount of droop settable with a resistor, so that it can be tailored individually for each application; it is available in the RC5058 family and the FAN5091 family. This bulletin discusses the operation of these methods in detail, and shows the details of component selection, including worst-case analysis.

The Fundamentals of Droop

The core voltage power supply for a microprocessor may have very tight limits on its output voltage, both for its DC regulation as well as for load transient response. This presents a difficult challenge for a conventional power supply, often resulting in an expensive Bill of Materials. To understand the problem, refer to Figure 1.

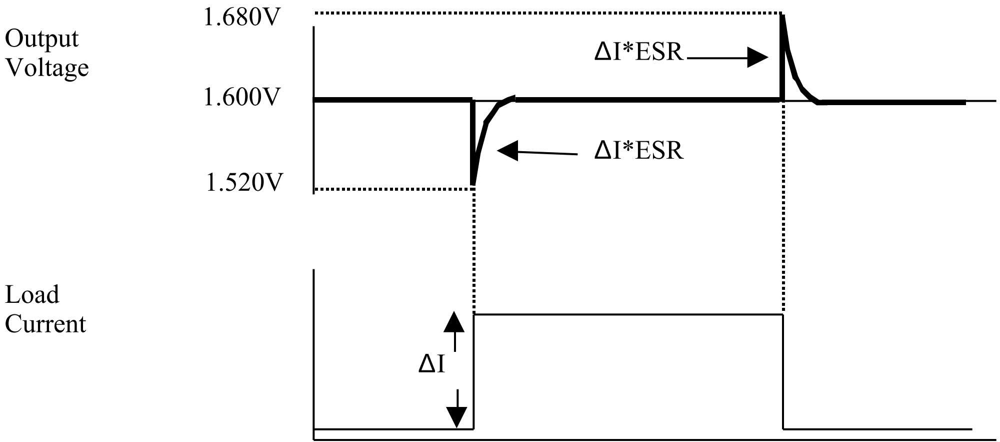

Figure 1 Transient Response of a Converter without Droop

Figure 1 shows the output voltage of a conventional converter designed to regulate at 1.600V. Since it is a "good" power supply, its DC output voltage is 1.600V regardless of the load current. When the load current suddenly changes, though, the converter's output voltage temporarily deviates from nominal. The reason that it deviates is that the converter has a limited bandwidth, that is, it cannot respond instantaneously. The output voltage then depends on the output capacitors of the converter: assuming the capacitance is large enough, the output voltage will be equal to (the nominal voltage) minus (the change in current * the ESR). To reduce the deviation of the output from nominal, the only thing that can be done is to reduce the ESR by increasing the output capacitance. There may be further effects if the capacitance is small, since then the total stored charge in the capacitors may become depleted, resulting in a further drop in output voltage.

There is another alternative, however. Suppose that instead of having a "perfect" converter that precisely regulated its output voltage to the desired value, we instead had a converter whose output voltage depended on load, as in Figure 2.

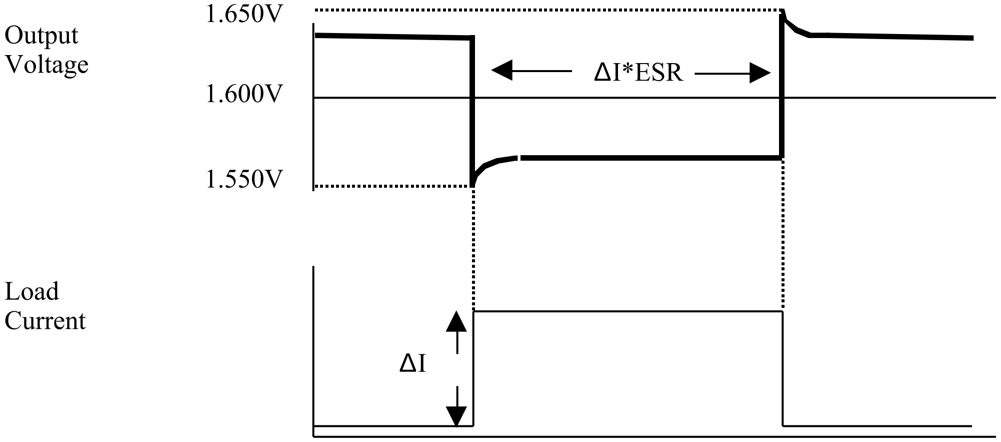

Figure 2 Transient Response of a Converter with Droop

The converter would have a DC output voltage that was higher than nominal when the load current was light, and lower than nominal when the load current was near maximum. As a result, when a load transient occurred the change in output voltage, although of the same magnitude as before, would cause the output voltage to deviate less. Essentially, the converter is "pre-positioning" itself in order to give maximum headroom for changes in load current. The change in DC output voltage with load current is called droop. The converter of Figure 1 has 0V droop; the converter of Figure 2 has ~90mV of droop. Clearly, having droop is the same as having a finite output impedance of the converter.

Types of Droop: Passive

Since the mechanism of droop is to produce a lower output voltage when there is a higher current, this is equivalent to intentionally having an output impedance of the converter. The most obvious way to do this is to introduce a physical impedance in series with the output, that is a resistor, as shown in Figure 3.

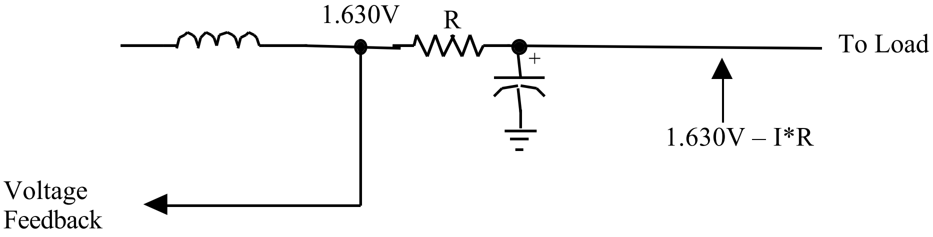

Figure 3 Passive Droop

This is called passive droop. Our "perfect" converter regulates the voltage to be exactly 1.630V; we intentionally set it higher than its nominal, say by including a resistor divider in the feedback. But the voltage that the load sees is less than 1.630V by I*R, which varies with load current as desired.

It is important to note the positioning of the feedback voltage, the resistor, and the capacitor. It seems at first thought that the resistor should be after the capacitors, as shown in Figure 4.

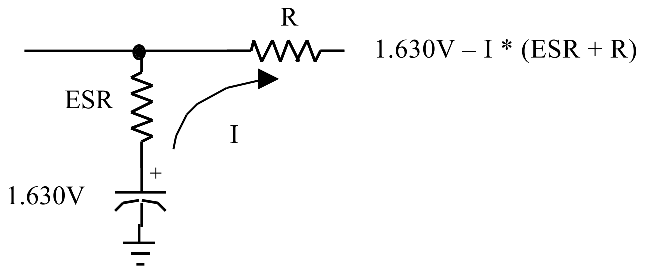

Figure 4 Incorrectly Placing Passive Droop Affects Transient Response

In this arrangement, with no load initially, the converter will charge up the output capacitor to its no-load voltage of 1.630V. When the load current suddenly increases, the output voltage immediately drops by I * (ESR + R), that is, *the droop resistor adds in to the effective ESR of the caps*. While it might seem that this is what you want (because you want the output voltage to be lower when there is more current), what it actually does is to change the DC output voltage of the converter instantly to its loaded output voltage, and then superpose the transient voltage drop on top of that. This thus actually makes the transient worse rather than better!

Passive droop does have one major problem: the droop resistor may dissipate very significant power, requiring a large, expensive component, and harming the converter's overall efficiency. For example, if a droop of 100mV is required at a load current of 50A (100mV / 50A = a 2mΩ resistor), this corresponds to a resistor power dissipation of 100mV * 50A = 5W! This requires the use of a 10W resistor; and if output voltage is 1.5V, there is an efficiency penalty of 100mV/1.5V = 7% just from the use of the droop resistor.

Types of Droop: Active

Active droop, as embodied in the RC5057, avoids the problems of passive droop. The idea of active droop is that the output current is monitored, and the converter is not "perfect": it modifies its output voltage according to the amount of current it supplies. Figure 5 shows two methods for monitoring current, for use with either the RC5057 or the RC5058 and its family (RC5059, FAN5056).

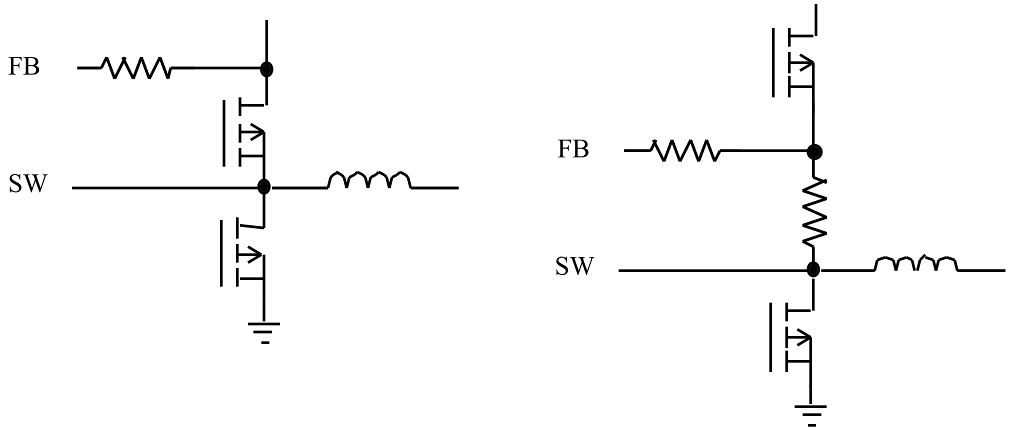

Figure 5 Two Methods for Measurement of Current

On the left side of Figure 5 is shown the normal method for monitoring current with either the RC5057 or RC5058. During the time when the high-side MOSFET is on, the output current flows through it. The IC can thus monitor the voltage across the MOSFET's drain-to-source, because it is proportional to this current: $V_{DS} = I * R_{DS,on}$.

The right side of Figure 5 shows an alternate method for monitoring current. Instead of monitoring the voltage across the MOSFET, the voltage across a resistor in series with the MOSFET is monitored instead. The advantage of this is that the resistor value is precisely known and does not change with temperature; whereas the MOSFET's $R_{DS,on}$ has considerable initial tolerance, and has a considerable temperature coefficient as well, making current measurements using the MOSFET relatively imprecise.

The current sensing in the FAN5091 is done similarly to that shown in Figure 5, but rather than sensing the voltage across the high-side MOSFET, the voltage across the low-side MOSFET is used.

With any of the methods of measuring current to implement active droop, power dissipation is reduced. In the case of measuring the voltage directly on the MOSFET, excess power dissipation is zero, since the MOSFET has to conduct anyway. When a source-side resistor is used for the current measurement, power is still less than when a resistor is used in series with the output; first, because the duty cycle is reduced, so power is I^2R*DC, and second because the sensor has gain, and so a smaller voltage needs to be developed across the sensor to develop the same droop.

Active droop has some problems also. Since there must of necessity be gain stages in the IC associated with the conversion of the measured current into droop, the amount of droop will have a tolerance to it. Taken along with the tolerance and temperature coefficient of the MOSFET, this results in rather substantial margins being required to ensure that tight specifications on voltage are being met.

Types of Droop: Programmable Active Droopä

The RC5058 family and the FAN5091 go one step further than the RC5057 by offering Programmable Active Droopä . The RC5057 has a maximum droop that is directly tied to the current limit setting. Specifically, the current limit of the RC5057 is set by an external resistor to a certain value of current I_{CL}; maximum droop then occurs at $2/3*I_{CL}$, and is equal to 40mV. In the RC5058 and FAN5091, the functions of current limit and droop are separated, each having its own external resistor, so that they may be programmed independently. It is thus possible to obtain more than 40mV of droop (maximum is 60mV for the RC5058 family and 10% of V_{out} for the FAN5091).

The droop of the RC5058 and its family is set according to the equation:

$$R_{Droop} = \frac{14.4K\Omega * I_{max} * R_{sense}}{18 * V_{Droop}}$$

Here, I_{max} is the current at which the full droop voltage is desired, R_{sense} is the resistance of the sensing element as discussed below, and V_{droop} is the desired amount of droop. For many applications using the RC5058 family, V_{droop} will be the full available 60mV, as the more droop available the fewer the required output capacitors.

The current sense element in designs using the RC5058 family is typically the high-side MOSFET, or alternately a sense resistor, both in the same configuration as with the RC5057 as shown in Figure 5. In the case of MOSFET current sensing, again, care must be used in selection of the value of R_{droop} due to the change in value of R_{sense} with temperature and component tolerance, and the 10% tolerance of the 14.4KΩ internal resistor.

The droop of the FAN5091 is set by an external resistor from pin 21 to ground. The FAN5091 has a wider range of droop than the RC5057 or RC5058 family. It can be programmed to have any droop up to a maximum of 10% of the programmed output voltage, according to the equation:

$$R_{Droop} = \frac{V_{Droop} * RT}{I_{max} * R_{sense}}$$

All of the terms have their same meaning, and RT is the value of the resistor that sets the switching frequency of the FAN5091. Observe that since the value of the droop resistor depends on RT, the switching frequency must be selected before the droop resistor value. The reason that the value of R_{Droop} was made dependent on RT was to remove the inaccuracy of depending on an internal resistor, such as the 14.4KΩ in the equation for the droop of the RC5058; depending on RT instead means that the tolerance will typically be only 1% instead of 10% for this component of the equation.

Unlike the RC5057 and the RC5058 family, the FAN5091 senses current on the low-side MOSFET rather than the high-side. Thus, precision sensing can be accomplished with a sense resistor in series with the source, and requires an extra connection pin, available as the FAN5093, as shown in Figure 6.

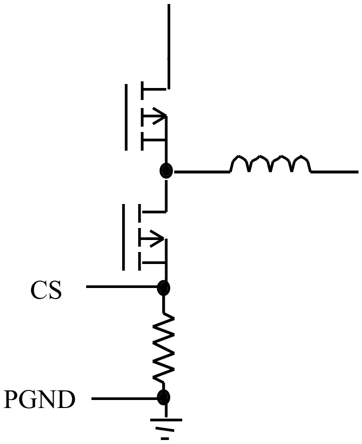

Figure 6 Connection of FAN5093 for Precision Current Sensing

How Much Droop?: RC5058 Family

Now that the methods of generating droop have been reviewed, it remains to consider the question of how much droop to use in a particular application, and how many output capacitors will be required as a result. Two examples will show the application of the RC5058 and the FAN5091 in typical situations.

Starting with the RC5058, suppose that the static limits on the output voltage are required to be +89/-79mV, and the transient limits ±134mV. These numbers are set by the load's requirements, in the case under consideration a CPU. Nominal output voltage is 2.000V, maximum output current is 14.2A, and we will be using MOSFET current sensing because of cost constraints. Assume the high-side MOSFET is an FDB6030L, with a nominal $R_{DS,on}$ of 9.5mΩ and a maximum of 13.5mΩ at 25°C, increasing 38% at 100°C to a maximum of 18.6mΩ. (A higher or lower temperature may be assumed as appropriate.)

Since the static negative limit is larger than the maximum droop of the RC5058 (|-79mV| > 60mV), we will set the droop to be maximum when the parameters in the numerator of the equation for setting droop are minimum; that way, when the parameters increase, we will still be at maximum droop.

Thus, we want to have 60mV of droop when $R_{DS,on}$ = 9.5mΩ. Remembering that the 14.4KΩ resistor has a tolerance of ±10%, we select

$$R_{Droop} = \frac{13.0K\Omega * 14.2A * 9.5m\Omega}{18 * 60mV} = 1.62K\Omega$$

With this value of droop resistor, we are now assured of having 60mV of droop whenever the current is maximum.

Now we turn to calculation of the number of output capacitors required to meet the transient limits. Taking the case when the current rises from zero to maximum first, the change in current causes the output voltage to drop by ΔI * ESR of the caps. The worst case will occur, then, when the RC5058 voltage is initially (at zero current) sitting at its minimum. The RC5058 is factory trimmed to be 1% higher than code, with a tolerance of ±1%. Thus its minimum output voltage is the VID code setting, 2.000V in this case. The output is allowed to drop 134mV from this value. Since this is caused by a current step of 14.2A, the required ESR for this transition is just ESR = 134mV / 14.2A = 9.4mΩ. Using Rubycon 6.3V, 1500μF ZL caps with an ESR of 23mΩ maximum at 25°C, 23mΩ / 9.4mΩ = 2.4 or rounding up, 3 caps.

On the opposite transition, from maximum to zero current, the output voltage will rise by ΔI * ESR, and so we must calculate the maximum output voltage from the RC5058. Factory trim and tolerance is +2% over nominal. Additionally, the RC5058 reference has a (positive) temperature variation of +8mV at 2V, which is +0.4%. So the output may be targeted as high as 2V + 2.4% = 2.048V. Since we are at maximum current initially, the output will be drooped by 60mV, yielding 1.988V. Since we are allowed to rise 134mV, there is 2.134V – 1.988V = 146mV of headroom; this being more than the 134mV of the previous paragraph, the number of caps required for this edge is less. We can thus conclude that for this particular set of parameters, 3 outputs caps are required. Note that this result is compatible with the calculation in the appendix to the datasheet of the RC5058, where seven 44mΩ Sanyo capacitors were used rather than the three 23mΩ Rubycon capacitors used here.

One thing to note, though not relevant to the worked example, is that if the allowable negative static limit were smaller than the 60mV droop available from the RC5058, it might be necessary to bias the output voltage upwards by using a resistor divider to the VFB pin, to allow maximum droop to be used. In this case, it must be checked that the maximum static limit isn't exceeded. In the example, the output may be as high as 2.048V and the maximum limit is +89mV, so that the bias cannot be more than an additional 41mV. If this technique is used, the bottom of the divider should be a 1KΩ resistor, in order to avoid offset due to the input leakage current of the VFB pin.

How Much Droop?: FAN5091

Calculating droop and output capacitors for the FAN5091 is essentially the same as the calculation for the RC5058, except that the FAN5091 can have almost any practically encountered amount of droop desired. As a specific example, suppose that the static limits on the output voltage are required to be
+40/-70mV, and the transient limits +50/-80mV. Nominal output voltage is 1.350V, and maximum output current is 60A. Assume the low-side MOSFET is a pair of FDS6690, each with a nominal $R_{DS,on}$ of 11mΩ (5.5mΩ total) and a maximum of 13.5mΩ at 25°C, increasing 38% at 100°C to a maximum of 18.6mΩ (9.3mΩ total). Further, assume that the switching frequency of each slice is 300KHz, so that the RT = 41.2KΩ.

Since the amount of droop is not limited, our strategy must differ from that for the RC5058 family. We must ensure that in worst case the droop does not cause the output voltage to go lower than the static voltage limit. The FAN5091 has no initial offset from its VID code, but it does have a tolerance of ±1%, and a temperature variation of +2mV, corresponding to about 0.1%. The output voltage can be low by as much as 1.350V * 1% = 13.5mV below nominal, and since we are only allowed to be 70mV below, the maximum permissible droop is 70mV – 13.5mV = 56mV. The only contribution to error in the droop equation for the FAN5091 is from the sense element, the MOSFET (which is why it is advantageous to use a sense resistor, using the FAN5093). The droop voltage will be maximum when R_{sense} is, and so we select

$$R_{Droop} = \frac{56mV * 41.2K\Omega}{60A * 9.3m\Omega} = 4.13K\Omega$$

Now we can calculate the number of output caps. In the load current transition form zero to sixty amps, the voltage is initially sitting high; its minimum high value has already been calculated to be 1.350V – 1% = 1.336V, or 13.5mV down. Since it's allowed to go 80mV down, we have 80mV – 13.5mV = 66mV of headroom. At a current of 60A, this means ESR = 66mV / 60A = 1.1mΩ. Using the Rubycon caps, we require 23mΩ / 1.1mΩ = 21 caps.

For the load current transition from sixty to zero amps, we calculate the FAN5091's maximum low value. The part has 1% tolerance, and so it can be 13.5mV low. Additionally, there is droop. To find minimum droop (to make the output voltage maximum) we select R_{sense} to be minimum:

$$V_{Droop} = \frac{4.13K\Omega * 60A * 5.5m\Omega}{41.2K\Omega} = 28mV$$

exactly half of maximum droop. Maximum output voltage at 60A is thus 13.5mV + 28mV = 42mV below nominal; since the transient is allowed to go +50mV positive, the headroom is 42mV + 50mV = 92mV, which is less than the other transition, and therefore does not determine the number of caps required.

From the description given of the calculations used to determine the number of output caps required, it becomes clear how the number could be minimized. The headroom on the two transitions should be equal: now one is 66mV, the other 92mV. The average headroom is (66mV + 92mV) / 2 = 79mV, and this can be achieved by biasing the output voltage up by 79mV – 66mV = 13mV with a resistor divider in the feedback path. This then requires only 79mV / 60A = 1.32mΩ, or 23mΩ / 1.32mΩ = 18 caps.

Of course, this upward bias with a resistor divider is subject to the limitation that it not exceed the positive static limit. Since the upward tolerance of the FAN5091 is +1.1% of 1.350V = 14mV, and the upper static limit is +40mV, the maximum upward bias is 40mV – 14mV = 26mV, so that 13mV is OK. Again, the divider impedance should be relatively low to avoid offsets due to input current.

For such stringent requirements as these, it is clear that the 2:1 tolerance in the amount of droop is costing many capacitors. Use of the FAN5093 permits use of a sense resistor, rather than the $R_{DS,on}$ of the MOSFET. Repeating the calculations, droop is now always 56mV, so that the output is 13.5mV + 56mV = 70mV below nominal, and headroom available is 70mV + 50mV = 120mV. Average headroom is (66mV + 120mV) / 2 = 93mV, which requires biasing up the output by 93mV – 66mV = 27mV, or 26mV to comply with static specs. This is now ideal performance, since full use of both the static and transient limits is now being made. ESR required is 93mV / 60A = 1.55mΩ, and 23mΩ / 1.55mΩ = 15 caps are needed, a reduction of 25%. If board space were critical, use of 10mΩ Oscons would further reduce the number of capacitors required to 7, though at a substantial cost penalty.