

Using Dynamic Voltage Positioning to Reduce the Number of Output Capacitors in Microprocessor Power Supplies

National Semiconductor
Application Note 1145
Dongbing Zhang
Steven Hunt
July 2000



Abstract

The relatively large steady-state window and the relatively small transient window for the core voltage of a modern advanced microprocessor make attractive implementation of the dynamic voltage positioning technique. By employing that technique, significant cost savings can be realized due to the reduced number of output bulk capacitors. Careful consideration of all sources of error is necessary for generating a sound design.

Example I

In this example, a net savings of \$0.44 is realized through the implementation of the dynamic voltage positioning (DVP below) technique.

Assume a microprocessor requires a maximal 18A. The allowed core voltage steady-state window is $\pm 100\text{mV}$. Suppose a synchronous buck converter is used and the controller's DAC tolerance is $\pm 30\text{mV}$. Also assume the output voltage ripple is set to 17mV peak-to-peak.

To accommodate the worst case load transient (i.e. the processor current changes between 0A to 18A within a few clock cycles), 14 Sanyo 6MV1500GX aluminium capacitors are needed at the output.

Or, if a $3\text{m}\Omega \pm 5\%$ IRC power resistor is put in series with the output inductor, DVP can be used and the number of output caps can be reduced to 10. Two additional signal level resistors ($1.00\text{K}\Omega$ and $75.0\text{K}\Omega$, 1%) are necessary to raise the initial output voltage by 26mV.

The cost of the power resistor is approximately \$0.20, and the cost of the capacitors is about \$0.16 each. So the total savings realized by implementing DVP is \$0.44.

The drawback is an additional maximum power loss of about 1W in the power resistor. A side benefit is the same power resistor can be used to provide an accurate current limit when a more serious over-current protection mechanism than high-side MOSFET $r_{\text{DS(on)}}$ sensing is desired.

Example II

In this example, a net savings of \$0.48 is realized through the implementation of the DVP technique.

Assume the same processor as in Example I is considered. If a PCB etch resistor is used instead of the discrete power resistor, a number of things are going to change. First, the total tolerance of the resistance will be increased to about 20% including the effect of temperature. Second, the resistor itself is free. So the amount of savings in output caps is the net savings.

By calculation, the optimal resistance of the etch resistor is $2.2\text{m}\Omega$. The initial output voltage should be raised by 16mV. This can be done by using two signal-level resistors, 100Ω and $12.4\text{K}\Omega$, 1%.

The number of output caps is now 11. So the savings is \$0.48.

The maximum power loss due to the etch resistor is 0.71W.

Introduction

For modern high-speed microprocessors such as those in the Intel Pentium® pro and Pentium® II families, there are strict load transient response requirements on the processor core voltage. Two operating windows are defined for the MPU core voltage, i.e., the transient window (or so-called AC window) and the steady-state window (or so-called DC window). The AC window is greater than or equal to the DC window. For example, the Klamath processor (Pentium® II family) requires, at the VRM connector, a DC window of 100mV, -60mV and an AC window of $\pm 140\text{mV}$ for a nominal core voltage of 2.8V. During steady-state, the core voltage is allowed to stay outside of the DC window for a short while but should never be outside of the AC window. Both windows are for instantaneous voltages, i.e. set point tolerance, ripple and noise etc. are included.

For a typical core power supply controller, the initial output voltage tolerance plus ripple is much smaller than the DC window. It therefore may be beneficial for the output voltage to be positioned at different levels within the DC window in response to different load current levels. The idea is by dynamically positioning the core voltage level according to the load current, extra window margin for the load transient response can be created.

As an illustration, Figure 1 shows two load transient response waveforms, one with DVP, the other without. Factors such as initial output voltage tolerance (typically the DAC tolerance in the case of a digitally programmable controller), ripple voltage, etc. are excluded. In the figure, the lines labeled "AC" are the transient window limits, the lines labeled "DC" are the steady state window limits, and the line marked " V_N " is the nominal core voltage.




FIGURE 1. DVP and non-DVP Output Voltage Transients Caused by Load Transients

For a non-DVP converter, steady-state core voltage doesn't vary with load current. Therefore, after a transient in either direction, the core voltage returns to V_N . See trace "a". For a DVP converter, the core voltage is a function of the load current. See trace "b". At no load, the core voltage is close to the upper limit of the DC window whereas at full load it is close to the lower limit of the DC window. This allows extra headroom for load transients in both directions. In Figure 1, β is the amount of extra transient headroom DVP creates. It is also the amount by which the nominal output voltage should be raised.

Implementation

As mentioned above, to implement DVP, the steady-state output voltage should be raised slightly at no load and it should droop as the load current increases. This characteristic can be realized by using a voltage divider in the feedback loop and adding an external droop resistor after the inductor. See Figure 2. In the figure, a synchronous buck PWM controller such as National Semiconductor's LM2635/6/7/8 is used. Now, instead of regulating point "a" as is done in non-DVP converters, point "c" is regulated. The voltage divider between points "b", "c" and ground raises the voltage at point "b" so that at no load the output voltage is slightly higher than nominal. (Notice at no load, points "b" and "a" are at the same potential). The voltage across the droop resistor is proportional to the load current during steady-state. Therefore, the heavier the load current is, the lower the output voltage will be. The corresponding voltage vs. output current characteristic is shown in Figure 3.




FIGURE 2. Implementing DVP Using a Droop Resistor and National's LM2635/6/7/8 Controllers

The resistor divider is not necessary if the internal reference voltage (typically the DAC output in the case of a digitally programmable controller) of the switching controller IC has been prebiased for DVP.

The droop resistor must be a power resistor since it is in the power path. It can be a discrete current sense resistor or it can be a PCB etch resistor. The typical resistance value is a few milli-ohms.




FIGURE 3. DVP Converter Output Characteristics

The benefits of a discrete resistor are better tolerance and an ultra-low temperature coefficient (typically $\pm 20\text{ppm}/^\circ\text{C}$). It also creates less thermal stress on the PCB. The disadvantages are component cost and availability, and very limited choice of resistance values.

The benefits of a PCB etch resistor are no cost and a flexible resistance value. The disadvantages are worse tolerance, high temperature coefficient (about $4000\text{ppm}/^\circ\text{C}$) and more thermal stress on the PCB. Which kind of droop resistor creates more savings depends on load current, DC and AC window sizes, initial output voltage tolerance, etc.

The Equations

To make a realistic comparison between a non-DVP converter and a DVP converter and to provide a design tool for DVP implementation, factors such as the DAC tolerance (assuming digitally programmable controller), output voltage ripple and droop resistor tolerance and temperature coefficient must be considered.

Figure 4 shows the distribution of the steady-state voltage of a non-DVP converter, assuming the load regulation is perfect. Note the load current is irrelevant to the distribution.




FIGURE 4. Non-DVP Steady-State Output Voltage Distribution and Transient Margins

The gray band between lines “a” and “b” corresponds to the total DAC tolerance which is 2ϵ . When voltage ripple is considered, the steady-state tolerance band must be widened to the one confined by lines “c” and “d”, i.e. $2\epsilon + \text{RIP}$. RIP is the peak-to-peak ripple voltage. The worst case full-load-to-no-load transient would occur when the steady-state voltage at full load is at line “c” because this is the case having the least transient margin (ρ in the figure). The corresponding transient margin is:

$$\rho = \delta_{AC} - \left(\frac{1}{2} \text{RIP} + \epsilon \right), \quad (1)$$

where δ_{AC} is half the size of the AC window. Assuming the AC window is symmetrical around the nominal voltage V_N , the same formula applies to load transients in the opposite direction (no-load-to-full-load).

In the case of DVP, the position of the tolerance band becomes a function of load current. It is higher at light load and lower at heavy load. See *Figure 5*.

The maximum voltage across the droop resistor (V_{droop}) determines by how much the output voltage should be raised at no load. If the error introduced by the droop resistor is ignored (as in the case of *Figure 5*), the center of the output voltage tolerance band should be half of V_{droop} above V_N at

no load, and half of V_{droop} below V_N at full load, as shown in *Figure 5*. The extra transient margin created is thus half of V_{droop} .




FIGURE 5. DVP Steady-State Output Voltage Distribution with Zero-Tolerance Droop Resistor

However, there is inevitably a tolerance and temperature coefficient associated with the droop resistor. Assume the tolerance and variation due to the temperature change of the droop resistor is $\pm\lambda$. See *Figure 6* below.




FIGURE 6. DVP Steady-State Output Voltage Distribution with Finite-Tolerance Droop Resistor

This error will make the upper boundary of the tolerance band at full load (line "c") 1λ closer to V_N than the lower boundary of the tolerance band at no load (line "d"). To keep the two boundaries (lines "c" and "d") equidistant from V_N so that transients in both directions have the same margin, the no-load offset β should be adjusted to a value smaller than half of V_{droop} .

The appropriate offset of the no-load tolerance band is:

$$\beta = \frac{1}{2} V_{droop} (1 - \Sigma), \quad (2)$$

where V_{droop} is the maximum droop voltage, i.e. the droop voltage at full load, and Σ is droop resistor tolerance ($\lambda = V_{droop} \times \Sigma$).

The new transient margin ρ' is now the distance between the line "c" and the upper boundary of the AC window or that between the line "d" and the lower boundary of the AC window. The formula to calculate the margin is:

$$\rho' = \rho + \beta \quad (3),$$

so the difference between the new and old transient margins is the offset β .

It is necessary to check whether the line "d" is still above the lower boundary of the DC window. The reason line "d" is more important is that it is farther away from V_N than line "c" and in some cases the lower half of the DC window is smaller than the upper half. The criterion is:

$$V_{droop} \leq \frac{2\delta_{DCL} - RIP - 2\epsilon}{1 + 3\Sigma}, \quad (4)$$

where δ_{DCL} is the size of the lower half of the DC window.

The percentage savings in the amount of output capacitors is:

$$SVGS\% = \frac{\beta}{\rho'}. \quad (5)$$

The cost savings in output capacitors can be determined by the following equation:

$$\Delta\$ = \left(\frac{1}{\rho} - \frac{1}{\rho'} \right) \cdot ESR \cdot PRICE \cdot I_{CORE}, \quad (6)$$

where ESR is the equivalent series resistance of each output capacitor, PRICE is the unit price of a capacitor and I_{CORE} is the maximum load current.

The actual savings may be slightly more or less than the result of Equation (6) because the number of capacitors is always an integer.

Once the type of output capacitors is known, the number of them can be determined by the following criterion:

$$N \geq \frac{ESR \cdot I_{CORE}}{m}, \quad (7)$$

where m is transient margin, i.e. ρ or ρ' .

Example I Revisited

Assumptions:

$$RIP = 17\text{mV}, \epsilon = 30\text{mV},$$

$$\delta_{AC} = 100\text{mV}, \delta_{DCH} = \delta_{DCL} = 70\text{mV},$$

$$R_{droop} = 3\text{m}\Omega \pm 5\%, I_{CORE} = 18\text{A}.$$

Thus,

$$V_{droop} = 3\text{m}\Omega \times 18\text{A} = 54\text{mV}.$$

Then by Equation (1), the non-DVP load transient margin is:

$$\rho = 100\text{mV} - (\frac{1}{2} \times 17\text{mV} + 30\text{mV}) = 61.5\text{mV}.$$

The DAC offset in this case will be, by Equation (2):

$$\beta = 0.5 \times 54\text{mV} \times (1 - 5\%) = 26\text{mV}.$$

By Equation (3), the DVP transient margin is:

$$\rho' = 61.5\text{mV} + 26\text{mV} = 87\text{mV}.$$

Check if the tolerance band is out of the DC window by using Equation (4):

$$(2 \times 70\text{mV} - 17\text{mV} - 2 \times 30\text{mV}) \div (1 + 3 \times 5\%) = 54.8\text{mV}.$$

This value is greater than the 54mV actual maximum droop voltage so the condition is satisfied.

Suppose the output capacitors are from the Sanyo MV-GX series, the unit price is \$0.16, and the ESR of each capacitor is 47m Ω , then the cost savings in output capacitors is:

$$\Delta\$ = (1/61.5\text{mV} - 1/87\text{mV}) \times 47\text{m}\Omega \times \$0.16 \times 18\text{A} = \$0.64.$$

After subtracting the cost of the droop resistor, the net savings is \$0.44.

Example II Revisited

Since the droop resistance value is now flexible, it is necessary to find out the optimal value first. The larger the droop voltage, the larger the offset β and the larger the extra margin will be. The largest allowable droop voltage can be determined by Equation (4):

$$V_{droop} = (2 \times 70\text{mV} - 17\text{mV} - 2 \times 30\text{mV}) \div (1 + 3 \times 20\%) = 39.4\text{mV}.$$

Therefore, the optimal droop resistance is:

$$R_{droop} = V_{droop} \div I_{CORE} = 39.4\text{mV} \div 18\text{A} = 2.2\text{m}\Omega.$$

The initial offset β is, by Equation (2):

$$\beta = 0.5 \times 39.4\text{mV} \times (1 - 20\%) = 16\text{mV}.$$

The new transient margin is, by Equation (3):

$$\rho' = \rho + \beta = 61.5\text{mV} + 16\text{mV} = 77\text{mV}.$$

Therefore, by Equation (5), the savings in output caps is:

$$\Delta\$ = (1/61.5\text{mV} - 1/77\text{mV}) \times 47\text{m}\Omega \times \$0.16 \times 18\text{A} = \$0.45.$$

To calculate the actual savings, calculate the number of caps needed in the non-DVP case:

$$N1 = 47\text{m}\Omega \times 18\text{A} \div 61.5\text{mV} \approx 14.$$

And calculate the number of caps needed in the DVP case:

$$N2 = 47\text{m}\Omega \times 18\text{A} \div 77\text{mV} \approx 11.$$

So the difference is 3 caps, and thus \$0.48.

Comments

1. As discrete resistors go lower in resistance value, it is harder and harder for resistor vendors to make them very accurate. IRC has resistance values down to 3mΩ. Sometimes paralleling two discrete resistors is the only solution.
2. When designing an on-board supply, the DC window can be relaxed to slightly larger than that written in the VRM specifications. Refer to the processor specifications instead. Typically a relaxation of ±10mV is possible.
3. When switching from a non-DVP design to a DVP design, a slightly larger inductor might be necessary to keep the output ripple voltage to same.
4. Normally the top and bottom layers of a PCB are plated and copper thickness on those two layers is inaccurate. Try to use an inner layer to place an etch resistor.
5. As a good rule, use 20mil/A current density for 1 oz. copper when designing an etch resistor.
6. A DVP design tool in the form of a Microsoft Excel spreadsheet is available from the authors to automate the design process. The tool is good for designing DVP for National Semiconductor's LM2635/6/7/8 family of products.
7. To realize the deadbeat type of response similar to trace 'b' in *Figure 1* during a fast load transient, the converter's loop characteristic needs to be fine tuned.

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT AND GENERAL COUNSEL OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.



National Semiconductor Corporation
Americas
Email: support@nsc.com

www.national.com

National Semiconductor Europe

Fax: +49 (0) 180-530 85 86
Email: europe.support@nsc.com
Deutsch Tel: +49 (0) 69 9508 6208
English Tel: +44 (0) 870 24 0 2171
Français Tel: +33 (0) 1 41 91 8790

National Semiconductor Asia Pacific Customer Response Group

Tel: 65-2544466
Fax: 65-2504466
Email: ap.support@nsc.com

National Semiconductor Japan Ltd.

Tel: 81-3-5639-7560
Fax: 81-3-5639-7507