

# Fueling the Megaprocessors - Empowering Dynamic Energy Management

Bob Mammano

## ABSTRACT

The introduction of microprocessors such as the Intel PENTIUM® PRO have underscored an ongoing trend in high-end digital systems wherein performance improvements, accomplished by greater numbers of transistors switching at higher clock rates, would impose unmanageable power dissipation problems without new energy management algorithms. Solutions to these power management issues, however, have presented significant challenges to the power supply designer who must now provide increasingly accurate voltage levels while simultaneously responding to unprecedented dynamic load

excursions. In seeking design tools to apply to this task, it helps to remember that sometimes old techniques are still best, as this paper illustrates with a mix of design considerations and procedures applicable to achieving success in meeting these stringent new power supply performance goals.

## INTRODUCTION

An oft-quoted description of microprocessor development is "Moore's Law" (re: Gordon Moore, Chairman - Intel Corporation), which states that microprocessor performance roughly doubles every eighteen months. History has borne this out




Figure 1 - Microprocessor development history illustrates "Moore's Law".

with the past developments as shown in Figure 1 and with Intel's introduction of the P6 - now named the Pentium® Pro - the trend continues.

While this CPU design is a significant development and a major step forward in computer technology, the power management solutions which made it possible are not unique to this product, and either now or in the near future, will appear in introductions from many of the major microprocessor suppliers. And since these control algorithms place unprecedented demands upon their power sources, power supply designers must rapidly gain the skills necessary for successful solutions.

Microprocessor performance enhancements come from many paths. Certainly, added complexity is a contributor. The Pentium® Pro contains some 5.5 million individual transistors and compressing these vast numbers of devices onto a single silicon chip has required significant reductions in feature sizes. While these smaller geometries lead to lower device capacitance, the tighter spacings now require lower supply voltages. This trend in CPU supply voltage has been forecast as continuing to fall as shown in Figure 2, and since there has been - and most assuredly will be - no relaxation in voltage tolerances, the accuracy requirements in terms of millivolt deviation continues to diminish.

And since device loads are capacitive, higher clock speeds - increasing at some 25% per year - lead to higher power dissipation. A summary of this trend is given in Figure 3 which shows that




Figure 2 - Forecast reductions in computer power supply voltage

even with reductions in device capacitance and voltage, the number of devices and their operating frequency have served to push total chip power dissipation to ever increasing levels.




Figure 3 - Increased microprocessor power dissipation

## POWER MANAGEMENT

Internal control of power dissipation has been driven by at least two factors. The first is at the system level in the name of energy conservation. The Environmental Protection Agency, in a "Memo of Understanding" of June, 1992, defined a power management regime for desk top equipment to which they gave the designation "Energy Star". For a system to bear the Energy Star logo, it must include a "sleep" mode which automatically, after a programmable inactive period, powers down the unit to consume no more than 30 Watts of line power. At the present time this level is fixed but industry expectations are that in the near future, the values for sleep power will be reduced, and the system applicability extended. Proposals currently under consideration would extend the Energy Star range to be a fixed 30 Watt limit when operating power levels are between 50 and 300 Watts, and then add limits of 10% for power levels over 300 Watts, and 60% for less than 50 Watts of operating power. Follow on proposals include both reducing the sleep limit to 8W while adding an intermediate




“suspend” mode with a 15W limit.

Regardless of the limits, the job of the power supply designer has become more difficult. For example, a power supply designed for a 300W maximum load can fairly readily be designed with more than 80% efficiency, meaning that some 60W could be dissipated within the supply. However, that same supply must now operate with only 6W losses if one were to assume that the load could throttle back to 24W in standby.

A second, and potentially even more challenging factor driving a need for power management, relates to the increasing complexity and operating speed of the new microprocessors as described earlier. With these new devices, continuous operation of the processing circuitry would exceed the capability of extracting internally generated heat from the integrated circuit die and package. To alleviate this problem, Intel engineers have provided what they call “Dynamic Execution Technology”. This is a power management capability on board the CPU chip which powers down unneeded circuitry, even between instruction cycles. While these features effectively keep chip temperature under control, they can also cause significant load-change transients in just one or two bus clock cycles as shown in Figure 4.

In addition to the large short-duration transient switching current surges that occur on internal clock edges, the power saving modes generated by Stop Clock and Auto Halt features can impose average current swings of 10 Amps in a 10 nanosecond period. While this magnitude of load



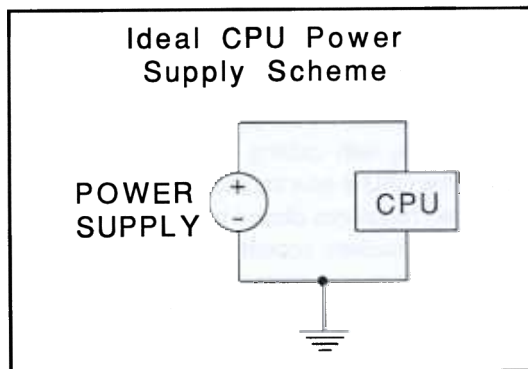
*Figure 4 - Peak and average supply current experienced in transitioning from low to high power operating modes within the Pentium® Pro*

change would always have been difficult to deal with, the reduction in supply voltage has further compounded the problem. With a 5V system, a 5% tolerance window is 250mV, but at a voltage level of 3.1V, this same percentage tolerance has shrunk to 155mV. And this limit must, of course, include the effects of both switching transient glitches as well as steady-state voltage-setting inaccuracies. It is easy to see that this issue will only get worse as further reductions in power supply voltage take place.

System engineers have been used to treating different functions as separate blocks. The power distribution system has traditionally been shown as in Figure 5 with a power supply feeding the CPU through a cable or some printed circuit traces with simple bypass capacitance at the point of use.

With average current swings of 1A/ns, and even higher clock transients, all on a low voltage bus, this new generation of “megaprocessors” will require a more exacting model such as that shown in Figure 6 where the source and load have been integrated with the distribution network as a system-level problem.

With the introduction of the Pentium® Pro, the CPU power supply designer is presented with an immediate problem which can also serve as an example of issues which - with varying degrees of intensity - are bound to be present in most future developments. Therefore, this paper will use these requirements to illustrate techniques and considerations which, hopefully, will have wider applicability.



*Figure 5 - A traditional CPU power supply distribution scheme.*

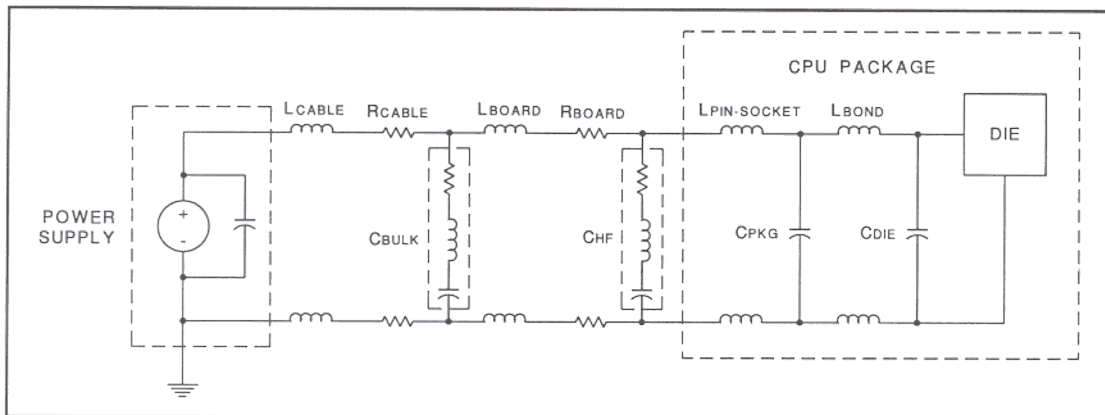


Figure 6 - A more accurate high-frequency power distribution model

### POWERING THE PENTIUM® PRO MICRO-PROCESSOR

The supply voltage for the Pentium® Pro continues the trend started when the original Pentium designers left the traditional 5V level to select 3.3V. The Pentium® Pro is currently designed for 3.1V but recognizing the critical relationship between performance and voltage, there are already expectations that future versions and/or upgrades will be specified for even lower voltage levels. Thus Intel has asked that power supply designers include the capability for digitally programming the supply voltage down to as low as 2.1V with 100mV resolution. Regardless of the voltage level, the total tolerance must be held to  $\pm 5\%$ , a value which must include the cumulative effects of initial setting accuracy, aging, temperature, ripple and noise, as well as line and load variations. And, as we have already said, load regulation must include any glitching caused by the dynamic load changes imposed by internal power management.

Meeting these requirements will require a different architecture for the power system. A central power supply with cabling or other varying feed lines to the CPU is now impractical and intervening distributed regulators placed immediately adjacent to the CPU sockets appear mandatory. The recommended power system architecture still retains the conventional off-line AC/DC power supply to provide 5V, 12V, and any other system voltages, but now the 5V bus is proposed as the power source for a dedicated DC/DC regulator designed

for and placed close to the CPU socket such that the distribution line for the CPU's Vcc voltage will be as short and well defined as possible. This, then, defines the need for a step-down DC/DC regulator module with a specification list as outlined in Figure 7.

There is, of course, one other very important but unstated requirement - that this regulator module be designed for very low cost. Thus there are many design decisions which will be discussed below in which the cost of implementation becomes a primary consideration.

#### Pentium® Pro Voltage Regulator Module (VRM) Specifications

Input Voltage: 5V  $\pm 5\%$  (12V  $\pm 5\%$  Available)

Output Voltage: 3.1V  $\pm 5\%$

Voltage ID: 16 Bit 100mV steps, 2.1 to 3.5V

Output Current: 0.5A to 11A (35 Watts)

Output ripple and noise:  $\pm 1.0\%$

Output Current Slew Rate: 30A/ $\mu$ Sec

Input current S.R.: 100mA/ $\mu$ sec

Input noise during step:  $\pm 2\%$

Efficiency: 80% full load, 40% light load

Additional Features: OVP, Pwr Good, Enable short circuit protection


Figure 7 - Major specifications for the Pentium® Pro Voltage Regulator Module (VRM)



Note that while we have said that the dynamic loading of the Pentium® Pro could impose a di/dt of 1A/ns, the slew rate specification for the regulator is listed as only 30A/us - a 33X reduction. The faster slew rate did not just disappear. The 30A/us recognizes the futility of attempting to make the regulator respond to the full needs of the CPU, and even if that were possible, the power distribution path between the two would degrade the response. Therefore this system requires a well-engineered distribution bus with adequate bypass capacitors to make the transition from 1000A/us at the CPU to 30A/us at the regulator's output. So regardless as to where the responsibility for this lies, the distribution network is a vital part of the total power system.

Recognizing the need for local regulation in close proximity to the microprocessor, Intel recommends building the regulator as a separate plug-in module as shown in Figure 8.

The envelope for this configuration defines a volume of 1.5" high, 3.1" wide, and 1.0" thick for a total volume of 4.65 cubic inches and, at 35 Watts, a power density of 7.5 Watts/cu in. While this module configuration is recommended, it is not mandatory and there will be some solutions where the regulator will be built on the motherboard. Assuming that there is room on the motherboard, the decision will probably be cost driven, trading off the value per square inch of motherboard area vs. the cost of a separate regulator PC board and its connector.



*Figure 8 - The package and mounting configuration for the VRM]*

## REGULATOR DESIGN CONSIDERATIONS

An initial assumption is made here that the most cost effective architecture for this regulator will be a non-isolated, buck switching design to step a 5V input down to 3.1V. With a load current range of 0.5A to 11A, it is highly likely that the inductor current will go into the discontinuous conduction mode at some lower current level, a characteristic which needs to be accommodated in the control design. Since 12V is available for a gate drive supply, a high-side N-channel FET switch is a logical selection, and it must be protected with some form of current limiting. Beyond this, the choices get a little more subjective and the following sections offer an overview of several independent considerations which must be made prior to the definition of a final approach.

### Circuit Control Topology:

While there are many choices possible, cost considerations will probably dictate choosing between voltage-mode, peak current-mode, and average current-mode schemes. These alternatives are shown schematically in Figure 9.

Voltage-mode control is the simplest and potentially the least costly alternative as should be obvious from the block diagram of Figure 9A. And since current sensing in a voltage-mode controller is only for protection, that portion of the circuit could possibly be even further simplified. However, there are many difficulties in using voltage-mode control for this application which are briefly outlined below:

1. It is difficult to get high gain-bandwidth due to the compensation needed to accommodate the double-pole output filter.
2. Capacitive loading affects the compensation, as does the ESR of the output capacitors.
3. The control loop compensation capacitors can take on large voltage offsets when the amplifier is driven into its stops, greatly extending the recovery from overload.
4. The closed-loop gain will change with changes in the output voltage (The input voltage is assumed constant.) which also forces a reduction in bandwidth.

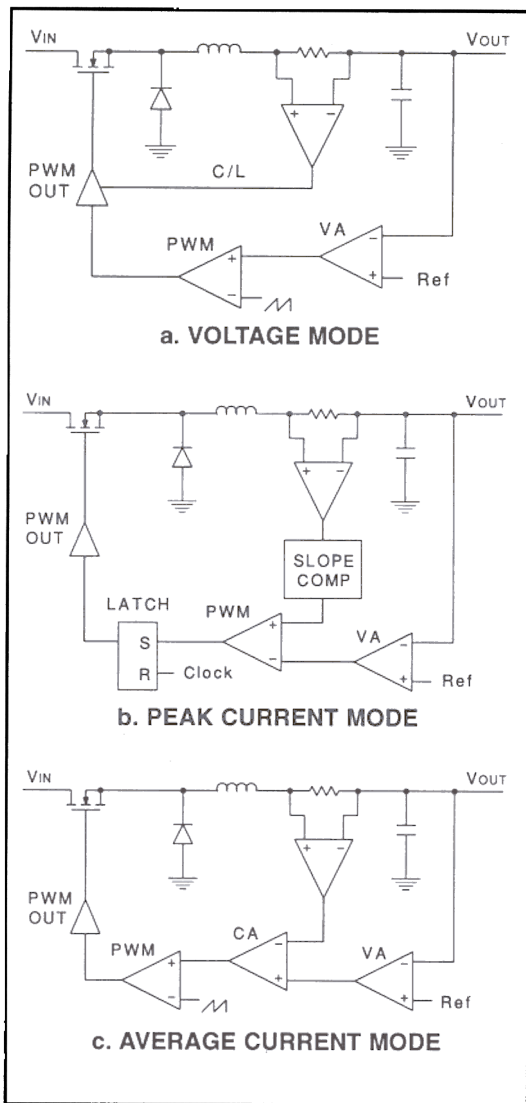


Figure 9 - Schematic configurations for three regulator control topologies

5. If an input filter is required, the input impedance of a voltage-mode regulator circuit has the potential for causing oscillations as it reacts to the output impedance of the filter.

Peak current-mode control eliminates many of the above disadvantages. As shown in Figure 9B above, there is now a current loop as well as the voltage feedback path, and current limiting can be automatically provided. Current-mode control

allows simpler compensation with higher bandwidth, freedom from the effects of variable capacitive loading, faster recovery from overload, and no interaction with an input filter. There are, however, still some disadvantages.

1. The error between the peak current measured by the control loop and the actual output current can degrade response.
2. The low gain of the current loop becomes even lower when the inductor current goes discontinuous. This can seriously degrade the performance of the voltage loop in responding to large step changes.
3. There is the added complexity caused by the need for slope compensation.
4. Current limit values typically have a tail caused by inherent circuit delays.

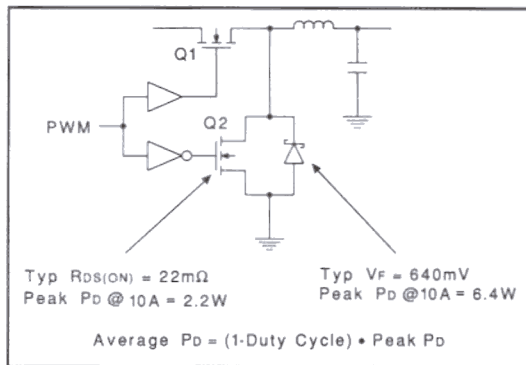
The use of average current-mode control alleviates all these problems. While the circuit of Figure 9C has become further complicated with the addition of a high-gain current amplifier, this topology has significant benefit to the specific requirements of this application. The high gain in the current loop not only provides the fastest response to changes in load current, but it also allows the circuit to traverse back and forth across the conduction mode boundary without loss in performance. Loop stability is unaffected by either changes in capacitive loading or the presence of an input filter, nor will changes in the output voltage have any impact on the voltage feedback loop.

#### Synchronous rectification:

It is easy to assume that the efficiency demands of a step-down regulator with a low output voltage could dictate the addition of a shunting switch across the rectifier as shown in Figure 10.

In this circuit, Q2 is driven with a complimentary waveform such that it is conducting during the time that Q1 is off. (Note that the rectifier is usually still required in order to allow conduction during any deadtime in the transitioning between turn-off of Q1 and turn-on of Q2.) The benefits of this circuit are apparent from the numbers given in Figure 10 where the 640mV drop of a typical Schottky





**Figure 10 - The use of a synchronous switch to reduce the voltage drop across the free-wheeling diode**

diode at 10A is exchanged for a switch which could have as little as 22mOhms of series resistance. Of course, the peak power is reduced by the duty cycle which, assuming 50%, means the addition of a synchronous rectifier reduces the 3.2W loss in the diode to 1.1W.

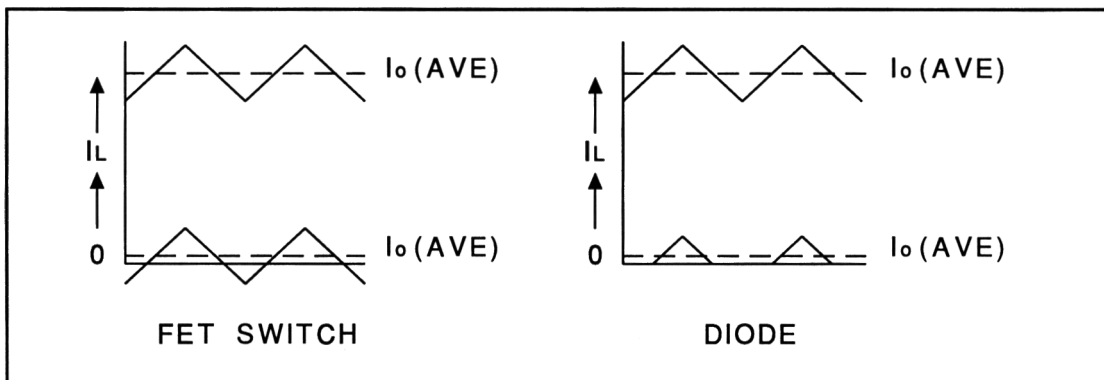
While the saving of 2.1W is not insignificant, it does amount to an overall saving in efficiency of less than 4% and comes at a cost of both the added switch and the increased complexity of the control circuit to develop its driving signals. There is also the consideration that with the use of a switch, inductor current will always be continuous. This is shown in Figure 11 where it can be seen that while there is no difference at higher load currents, at light loading a switch will allow negative current to flow if the peak-to-peak ripple current is

greater than twice the minimum load current.

While continuous inductor current could help stability if voltage-mode or peak current-mode control is used, it does have a derogatory effect on the low-current efficiency as the peak-to-peak currents are much higher than with discontinuous conduction. This is clearly an area where a cost-benefit trade-off could be made, as efficiency is not usually the most critical characteristic in the type of high-performance systems targeted by the Pentium® Pro.

### Switching frequency:

The selection of switching frequency is another area where the optimum choice may not be immediately obvious. It is easy to assume that the answer to fast dynamic load changes is high gain-bandwidth in the regulator, but it's not! The answer is in effective bypass capacitors! While a faster regulator could, perhaps, reduce the value of the output capacitors, this solution can also add substantial cost. And even if the regulator loop could respond instantly, the problem would not be solved because any inductance in the output line - and there will be some - will block an immediate response from the regulator. Additionally, a 10A change at the output is not a small-signal problem because a change of this magnitude will force the loop into saturation where all the gain-bandwidth in the world will be ineffectual - the maximum change in output current is just determined by the voltage across the regulator's inductor and its value.



**Figure 11 - A synchronous rectifier insures continuous inductor current**

Although higher switching frequencies could reduce the inductor value - and perhaps the total size of the regulator - the costs are not insignificant. Among them are the fact that the capacitors must now be high-cost tantalums or multi-layer ceramics instead of aluminum electrolytics, there will undoubtedly be an efficiency degradation due to the power FET gate drive losses, and the fact that a fast-responding regulator will present its own dynamic load change to its power source, to the extent that a more significant input filter may now be required.

### Voltage loop considerations:

While we're in the process of challenging conventional assumptions, here's another one to consider: Regardless of the type of control algorithm used, the voltage feedback loop can be shown as in Figure 12 where the box labeled Pulse Width Modulator includes the power stage and any embedded current loops if current-mode control is used.

Since the degree of voltage regulation is defined by the gain of the voltage amplifier, conventional wisdom would suggest that the gain should be as high as possible within the confines of stability. Therefore, the feedback element around the voltage amplifier,  $Z_{fb}$ , almost always includes an integrating capacitor so that no matter how the bandwidth may be limited, the circuit will have maximum gain at DC. Remember that the issue in this application is a sudden change in output loading

with a requirement that the output voltage stay within some definable tolerance window. Under these conditions, since the bandwidth of the regulator is not infinite, the output voltage is bound to take an immediate shift in response to a step change in output current and then return to close to its earlier value as the regulator drives the resultant error back toward zero. This is illustrated in Figure 13A where an integrating feedback loop is used and the resultant output voltage transient exceeds the allowable voltage window - first in one direction when the circuit is loaded, and then in the opposite polarity when the load is removed.

Note that since the voltage transient can be in either polarity, the nominal output voltage setting should be in the center of the window and, in this example, the allowable limit is exceeded in each polarity.

The proposition shown with Figure 13B is that by establishing a finite DC gain with non-integral feedback around the voltage amplifier, a DC error is allowed. This means that when the output is loaded, an output DC voltage droop will occur or, in other words, the circuit now has a defined non-zero output impedance. Now, by either adjusting the reference value or inserting a small DC offset in the voltage setting dividers, the output voltage is offset toward the high limit when the circuit is at minimum load, and the loop will allow it to stay near the low limit when loaded. It can now be seen that the effective window for transient effects has been

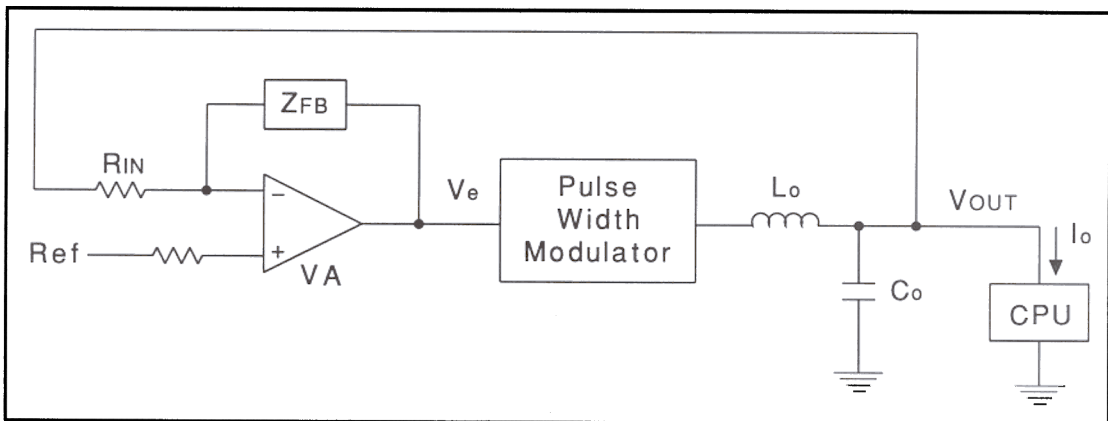


Figure 12 - A simplified representation of the voltage feedback loop in a regulator



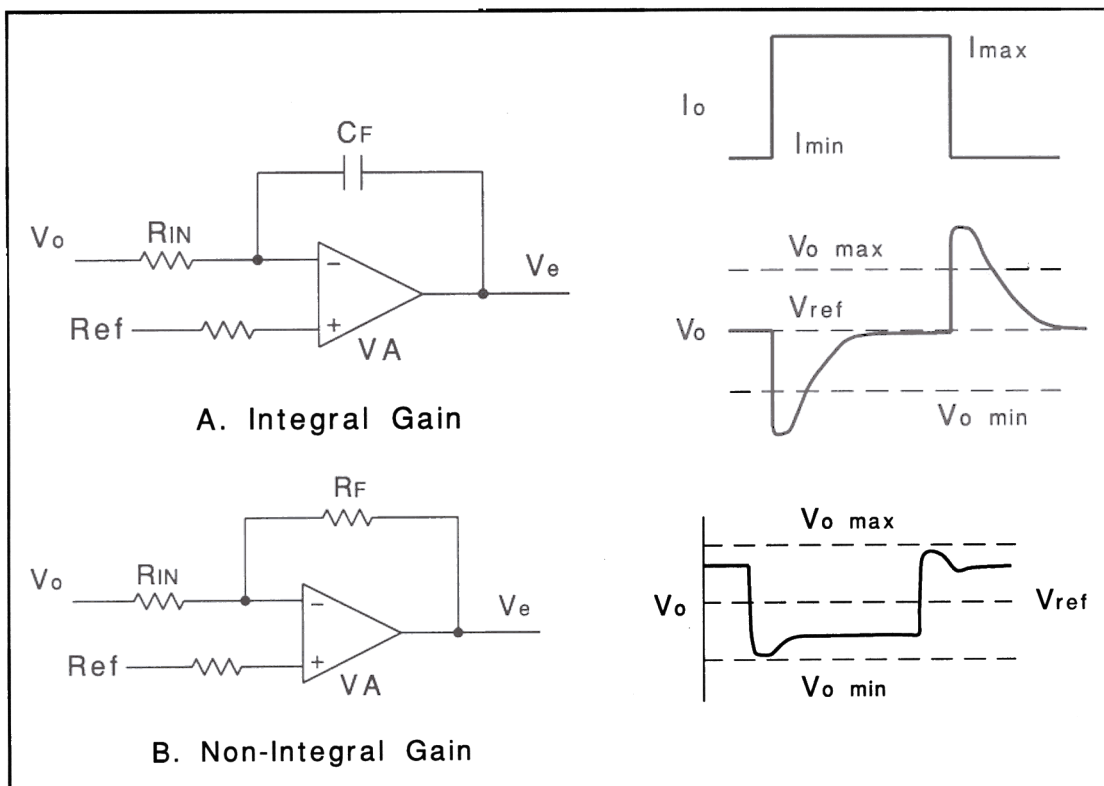


Figure 13 - Effects of non-integral voltage loop gain on voltage transient response

close to doubled and even though there is now a DC shift, the limits for this example are no longer exceeded.

#### Component parasitic elements:

One final design consideration is the recognition that between the output of the voltage regulator and the power pins of the CPU will be a power distribution network made up of conductors and capacitors, none of which can now be assumed to be ideal. It was stated earlier that the only practical approach to meeting these stringent dynamic loading requirements was through a judicious use of bypass capacitors; however, this can only be done with adequate knowledge of the characteristics of these components at the frequencies of interest. With a 10A step change in current, inductance becomes all-important - in both conductors and capacitors. It will be shown in the design example which follows that capacitors may well be selected more for their parasitic equivalent series

inductance and resistance, than for their capacitance value. And in order to define the capacitors, the series inductance of the power bus must be both known and controlled. Since inductance is defined by geometry, good mechanical control precludes the variability of power cables so printed circuit trace dimensions become a defining factor.

Therefore, before any component selection can be made, component parametric information must be obtained, either from the manufacturer or by direct measurement, and by careful analysis of the physical geometrical relationships. In all cases, it is also important to have a good understanding of the manufacturing processes in order to define applicable safety margins to allow for tolerances.

More could be said on this topic but since the purpose here is to describe a design process, we will move on to discuss how the information is used rather than how it is obtained.

## A DESIGN EXAMPLE

### Overview:

The basic architecture for this DC/DC regulator example is the simple buck configuration shown in Figure 14.

This figure shows the switching regulator which will step 5V down to 3.1V, along with a model for

the power distribution to the microprocessor's socket. This regulator topology incorporates a conventional free-wheel rectifier and uses average current-mode control switching at 200kHz fixed frequency - a value selected on the basis of cost and size. The block shown for the controller is implemented with two Unitrode products: the UC3886 which contains the pulse-width modulator,

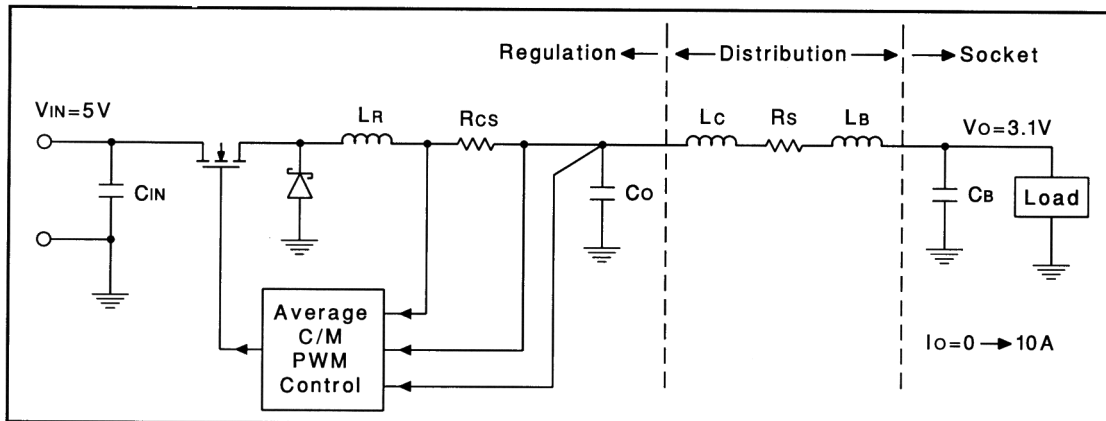


Figure 14 - A simple step-down regulator and distribution model for a Pentium® Pro application

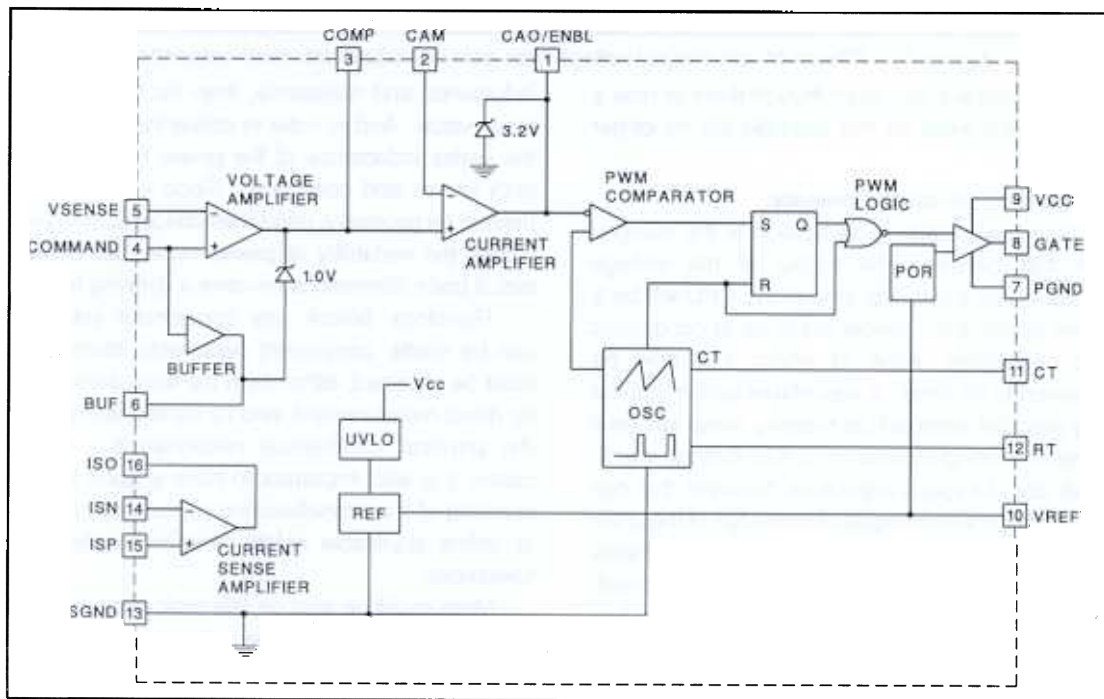


Figure 15 - The UC3886 average C/M PWM controller

FET driver, and the analog control functions; while the UC3910 provides a four-bit digitally defined reference voltage, along with the monitoring and protection circuitry. The block diagrams for these devices, and for their combination to form the complete regulator, are shown in Figures 15, 16, and 17.

This regulator design and the IC's allowing its implementation were defined to meet all the Pentium® Pro specifications as listed in Figure 7, with low cost being second only to performance as a defining objective. Unitrode Application Note U-157 contains the details of this design and will not be repeated here. Instead, we will concentrate on the power distribution network which, as it turns out, is actually the prime defining element for the specific requirements for the regulator. It should also be noted that some of the numbers given below have been simplified as an aid in describing the design process.

The regulator itself is packaged as a stand-

alone, plug-in module with the feedback loop closed within the module. While remote sense could have been used to sense the voltage at the CPU socket, it was not deemed necessary due to the well-defined distribution path between them. All this means is that the DC resistance of the connector pins and PC trace will add to the closed loop output impedance of the regulator - there is negligible impact on the dynamic response.

### Establish an error budget:

The overall tolerance for the Vcc voltage - at the CPU socket - is specified as  $\pm 5\%$ . With the UC3910 reference capable of 0.5% accuracy, and no error contribution from resistive dividers,  $\pm 2\%$  was allocated for the combined effects of setting accuracy, temperature, line regulation, and ripple, leaving  $\pm 3\%$  for load regulation. To maximize the utilization of this window, non-integral feedback was applied to the voltage error amplifier in the control circuit as follows: (Refer to Figure 18 for definition of the circuit parameters.)

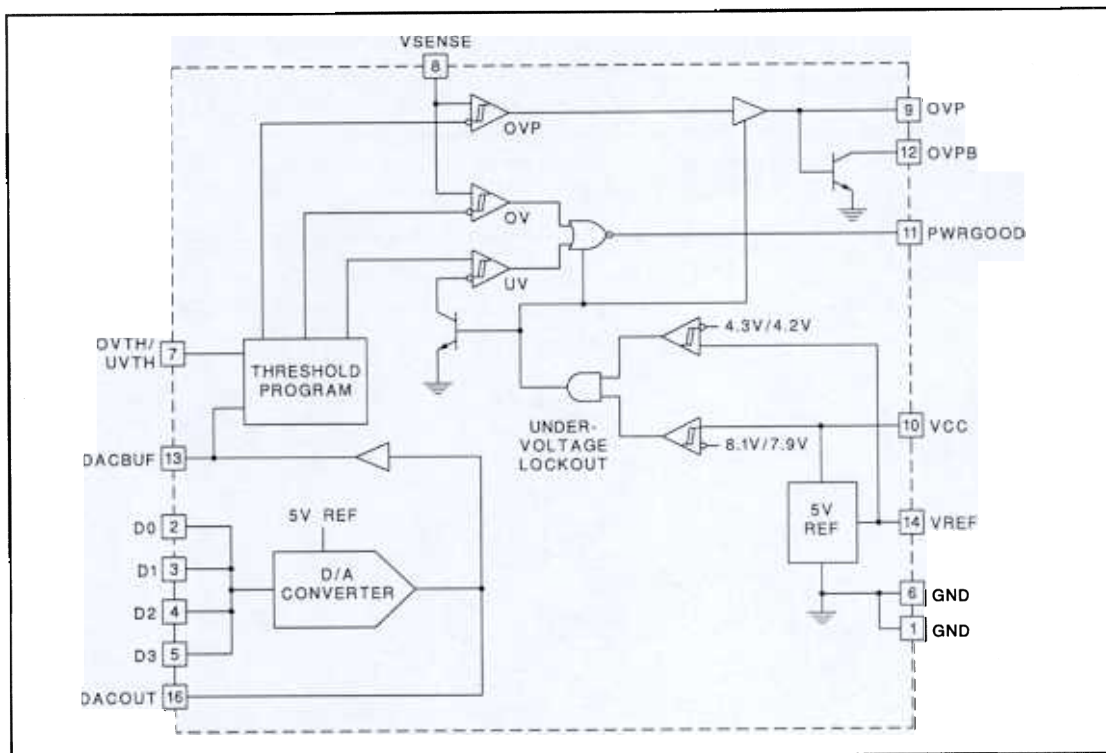



Figure 16 - The UC3910 contains a digitally controlled reference and the protection functions

Subtracting the voltage loss outside the loop caused by the DC resistance of the power bus and

The current amplifier will force  $v_e$  to equal  $v_c$  and thus, with  $K_c$  set arbitrarily at 8 and  $R_{cs}$  at



0.01Ω, it follows that

$$K_{va} = \frac{10A \cdot 0.1 \cdot 8}{0.150V} = 5.3$$

The allowable output voltage shift with load current also establishes the closed-loop output impedance of the regulator as

$$R_o = \frac{\Delta V_o}{\Delta I_o} = 15.0m\Omega$$

Note that if remote sense were used, the distribution series resistance,  $R_s$ , would be inside the loop but the voltage amplifier gain would be set slightly different so that the impedance at the CPU socket would still be 18.6 mΩ.

One other important point is that while we do not yet know the bandwidth requirements for these elements, the low closed-loop gain settings for both  $K_c$  and  $K_{va}$  would indicate that it should not be too difficult to obtain flat response up to at least one-fourth the switching frequency which, as it will turn out in this example, is well more than necessary.

#### Passive component definition:

The first step in even beginning to understand the dynamic problem for the regulator is to define the passive components in the power distribution path. Referring back to Figure 14, we have shown capacitors  $C_o$  and  $C_b$ , along with inductances  $L_c$  and  $L_b$ , between the regulator and the CPU socket. These parameters are defined as:

$C_o$  = Regulator output capacitance

$C_b$  = Board mounted capacitance at the CPU socket

$L_c$  = Connector pin inductance  
(All pins in parallel)

$L_b$  = PC board power trace inductance  
between regulator and CPU.

The impedance of each of these components will play a decisive part so their parasitic elements need to be defined. To illustrate this, consider the

operation of the bypass capacitor,  $C_b$ , positioned to provide the first level of accommodation for the sudden change in load current. Figure 19 shows




Figure 19 -The voltage transients on a bypass capacitor caused by a sudden change in load current

the effects of the parasitic series inductance and resistance, along with the true capacitance of this component, and it can be seen how the ESL becomes the dominant parameter in determining the initial response to a sudden change in load.

In choosing capacitor types it should be remembered that the parasitic elements are more strongly defined by the package style, shape, and size rather than the value of capacitance. Selected capacitors considered for this application are listed in the table shown in Figure 20.

Equally important as the bypass capacitors, the inductance of the power conductor between the regulator and the load must be defined. A recognition of the importance of this parameter was the main consideration in locating the regulator adjacent to the CPU with only a PC trace as a distribution bus, but this arrangement still does not make the series inductance go to zero. The geometry of a conductive strip over a ground plane is illustrated in Figure 21 which also defines the important dimensions. With the assumption that the width of the trace, *w*, is more than 10 times the separation, *d*, the series inductance of this conductive path can then be calculated from the physical dimensions as:




Figure 21 - Inductance of a conductive trace over a ground plane

$$L_b = 12 \frac{d \cdot l}{w} \text{ nH}$$

When *w* > 10*d* and all dimensions are in centimeters.

It is now instructive to examine Intel's recommendations for powering the Pentium® Pro. The physical layout of the socket, the regulator connector, and the interconnecting PC trace are shown in Figure 22.

Note that the socket power pins and the decoupling capacitors are highly distributed with many parallel paths making an exact analysis extremely difficult, however, the values are such that simplifying assumptions can be made with the results still helpful. In this case, we will assume that all the

| Capacitor Options                  |             |       |       |
|------------------------------------|-------------|-------|-------|
|                                    | Capacitance | ESR   | ESL   |
| <b>Ceramic</b>                     |             |       |       |
| X7R 0603 Pkg                       | 0.1µF       | 100mΩ | 3.0nH |
| X7R 1206 Pkg                       | 1.0µF       | 120mΩ | 1.9nH |
| <b>Aluminum Electrolytic</b>       |             |       |       |
| 10mm Dia 12.5mmH                   | 270µF       | 120mΩ | 4.0nH |
| 10mm Dia 20mmH                     | 1500µF      | 40mΩ  | 5.0nH |
| <b>Sold Tantalum</b>               |             |       |       |
| 16V D case                         | 47µF        | 100mΩ | 2.0nH |
| 6.3V E case                        | 220µF       | 100mΩ | 2.7nH |
| Note: ESL includes termination via |             |       |       |

Figure 20 - Selected capacitor options and their parasitic values








Figure 22 - Recommended socket layout for the Pentium® Pro

bypass capacitors are in parallel at the center of the CPU socket, which is 5cm from the center of the regulator connector, and that the interconnecting trace is 3.3cm wide and .025cm above the ground plane. This yields a value for  $L_b = 0.5$  nH which we will use in the equivalent schematic for the power distribution path shown in Figure 23.

In addition to the 0.5 nH inductance of the PC trace, another 1.0 nH is added for the inductance of the regulator connector pins. (This is measured

data with 11 pins in parallel for power and return.)

The equivalent circuit of Figure 23 assumes the regulator as a voltage source of 3.1V with an output impedance of 15 m $\Omega$ . The regulator output capacitor,  $C_o$ , is shown with its parasitics included, along with the socket bypass capacitance,  $C_b$ . And, of course, the load is shown as a current sink with a value switching between zero and 10A in 10 nsec.

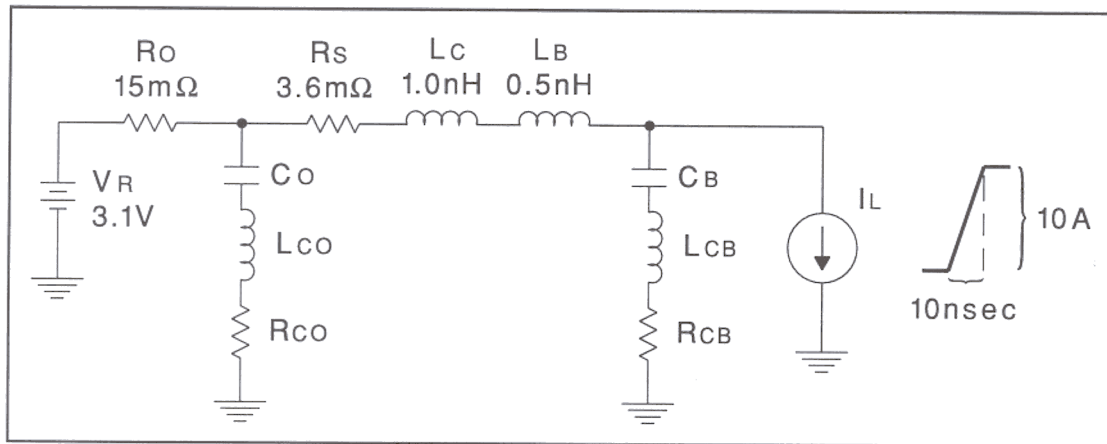


Figure 23 - A discrete equivalent circuit for the power distribution network

### The design process:

Remember that the power system as we have defined it herein includes all the elements between the power source and the load. We have simplified this to a two-part analysis - The DC/DC converter operating from a 5V bus, and the power distribution bus between the regulator and the CPU socket. Because each of these elements impacts the other, the logical method of analysis is to start with the load and define the distribution network, which will then characterize the more specific requirements for the regulator.

With the definition of the distribution network, and the component elements of the equivalent circuit known, this part of the problem becomes simplified to the point of determining the type, quantity, and location for the capacitors which will be able to reduce the 1 Amp/nanosecond risetime of load current to something that the bandwidth of the regulator can handle while keeping the voltage deviation within the limits defined. While modern-day computer simulation programs could readily provide very accurate results for a multitude of possible solutions, such an approach involves a lot of "cut and try" without the understanding to insure that an "acceptable" solution is not an overkill and therefore too costly.

The problem of a power distribution bus with parasitic elements is a problem in network analysis and tools to do this manually before the days of computers are still effective today. Moreover,

these techniques are both quick and easy and, more importantly, provide an intuitive appreciation of the problem that assists in the solution.

The design of networks and the selection of their components are most easily done by translating the problem from the time domain into the frequency domain and from a problem of voltage and current to a problem of impedance. This transformation removes one variable from the analysis, simplifying the process while making potential solutions more obvious. An important principle which will be used in this analysis is that if the impedance of the network that supplies power to the CPU stays below some definable value at all frequencies, then the voltage deviation due to the current step will remain within specification. Utilizing this principle is all the more eased by the fact that in most cases, only the magnitude of the impedance vector is needed as the poles and zeros are well damped and the phase relationships can be ignored.

Because there is a finite risetime for the current step, the spectrum of this waveform does not extend to infinite frequency but decreases at a constant rate (20 dB per decade) up to a cutoff frequency which is defined by the rise time. Above this frequency, the rate increases to 40 dB per decade. This means that if the impedance of the network is constant at low frequencies up to the cutoff frequency, and beyond that rises at a rate of 20 dB per decade, the voltage deviation will remain

within a constant limit for all frequencies. This cut-off frequency is related to the current risetime as:

$$f_c = \frac{0.35}{tr} = 35 \text{ MHz}$$

So the problem for this example is to maintain the impedance below 18 milli $\Omega$  from DC to 35 MHz while allowing a rise of 20 dB / decade above that value. The design of the optimum network within

the restraints given will start at the high end of the frequency spectrum and work downward, adding the appropriate capacitance as needed to meet the impedance objectives. This means that high frequency capacitors without much energy storage capability will be placed closest to the CPU to minimize series inductance while larger valued units with higher inductance can be placed further away where they will be effective at lower frequencies.


### REACTANCE PAPER by Phil Todd - Power Innovations

Reactance paper is a versatile tool to an analog circuit designer. It is a relative of the slide rule and uses logarithms to perform calculations. Reactance paper is simply a log-log plot of impedance versus frequency which has been overlaid with the impedance curves of both inductors and capacitors. A sample of this paper is shown below.

With this tool, the impedance of an inductor or capacitor at any frequency can be read from the Y-axis at the level where the curve for that value of inductance or capacitance intersects the desired frequency on the X-axis. One of the

great advantages of reactance paper is that it allows the visualization of network impedance as a function of frequency. It is also self-documenting when the calculation is finished.

Reactance paper can be used for many other functions such as op amp gains and control loop compensation but these are beyond the scope of this paper. One disadvantage of reactance paper is that it has apparently disappeared from drafting supply houses but a sample for copying for personal use is included with these instruction materials.



### Step-by-step design example:

1. Plot the low frequency impedance requirement from DC to 35 MHz as a horizontal line at 18 m $\Omega$ . The 10 nsec risetime of the load current causes the impedance curve to break upward at 35 MHz, with a slope equivalent to that of an inductance of 80 pH. This curve (or line segments) now represent the upper limit of output impedance for the power system for both the regulator and distribution network.

2. To be lower than the high frequency rising impedance, the inductive component of the board capacitance must be lower than 80 pH. The selection of 30 paralleled 1.0 uF ceramic X7R capacitors in 1206 packages for Cb will yield an equivalent inductance, Lcb, of 63 pH. This impedance is plotted the reactance chart.
3. Thirty 1.0 uF capacitors also have an effective ESR of 4 mΩ for Rcb. With decreasing frequency, the decreasing inductive impedance curve becomes resistive at 10 MHz and then capacitive due to Cb at 1.3 MHz, rising such that it would cross the impedance limit at 300 kHz. With this solution for Cb, the reactance plot shows that we have accommodated the dynamic loading requirements at all frequencies above 300 kHz.
4. Additional capacitance, Co, is now needed but its ESR (Rco) and ESL (Lco) adds to the series resistance and inductance. The total



resistance is now  $R_{co} + R_s$  and the inductance is  $L_{co} + L_c + L_b$ . A variety of capacitor types could be used for  $C_o$  but, in the interests of low cost, we will first explore the use of aluminum electrolytics. Four 1500  $\mu F$ , 10 x 20 mm units in parallel will provide  $C_o = 6000 \mu F$ ,  $L_{co} = 1.25$  nH, and  $R_{co} = 10$  mOhms. The total inductance is now  $1.25 + 1.0 + 0.5 = 2.75$  nH while the resistance is  $10 + 3.6 = 14$  m $\Omega$ . Placing these values on the reactance plot shows that the rising impedance from  $C_b$  intersects the new series resistance at 400 kHz where ( $R_{co} + R_s$ ) now dominates until  $C_o$ , at 6000 $\mu F$ , begins to raise the impedance at 2500 Hz. (The inductance, which would raise the impedance above 800kHz, is not a factor as its effect is negated by the shunting of  $C_b$ .) It is now below 2500Hz that the regulator must act to hold the total impedance below 18 m $\Omega$  but, before leaving the higher frequency range, one further step must be performed:

5. Any potential resonant peaks must be examined to insure that damping is adequate to maintain the desired impedance value. In this case, there is a resonance at 550 kHz between ( $L_{co}+L_c+L_b$ ) and  $C_b$ . At this point, the resonant impedance is,

$$Z_{res} = \sqrt{\frac{L}{C}} = \sqrt{\frac{2.75nH}{30\mu F}} = 9.5m\Omega$$

The allowable Q at this point is the ratio between the maximum allowable impedance and the resonant impedance,

$$Q_{max} = \frac{Z_{limit}}{Z_{res}} = \frac{18m\Omega}{9.5m\Omega} = 1.9$$

To insure that Q is less than this maximum, the damping must be greater than,

$$R_d > \frac{Z_{res}}{Q_{max}} = \frac{9.5}{1.9} = 5.0 \text{ m}\Omega,$$

and since the resistive damping at this frequency is ( $R_{co} + R_s$ ), at 14 m $\Omega$  this criteria is satisfied.

6. Below 2500 Hz, the impedance is rising due to the capacitance of  $C_o$  and it is here that we need the regulator to take over in supplying the load current. To evaluate the regulator response characteristics, we need to determine the bandwidth of the current control loop.

Referring back to Figure 18 where we defined the elements of an average current-mode control system, the parameters which are significant to the current loop include the following:

The switching frequency,  $f_s = 200$  kHz

Inductor value,  $L = 15 \mu H$

Current sense resistor,  $R_{cs} = 10$  m $\Omega$

Differential current buffer,  $K_c = 8$

Peak-to-peak ramp waveform,  $V_s = 1.8$  V

The small-signal inductor current,  $i_L = V_{in} d / X_L$ , where  $d$  is the small-signal duty-cycle

$$d = \frac{v_{ca}}{V_s} \text{ and } X_L = 2\pi f L.$$

Therefore,

$$i_L = (V_{in} \cdot v_{ca}/V_s) \frac{1}{(2\pi f L)}$$

The transconductance gain of the modulator and power stage is then

$$\begin{aligned} G_{pwr} &= \frac{i_L}{v_{ca}} = \frac{V_{in}}{2\pi f L V_s} \\ &= \frac{5V}{6.28} \cdot 15\mu H \cdot 1.8V \cdot f = \frac{29,500}{f} \end{aligned}$$

The total current loop gain is then

$$K_{cl} = G_{pwr} \cdot R_{cs} \cdot K_c \cdot K_{ca}$$

And if we set this expression equal to one at the crossover frequency,  $f_c$ , we can solve for the current amplifier gain needed to provide that bandwidth, under the assumption that the bandwidth of



Kc and Kva are both significantly higher than fc. Since the reactance plot calculations have shown that fc need be no higher than 2500 Hz, then

$$K_{ca(min)} = \frac{2500}{29500} \cdot .01 \cdot 8 = 1.04 \text{ at } f_c$$

Since obtaining an open loop amplifier gain of only slightly more than one at 2.5 kHz is a relatively trivial task, the results of this analysis have shown that a properly engineered power path between the regulator and the load has turned what might have initially appeared a very daunting design task, into a much more manageable problem.

## ALTERNATE SOLUTIONS

It was stated earlier that many solutions to this problem are possible. The approach described above emphasizes low cost but other criteria may be more important. For example, if small size was

the primary issue, higher switching frequencies and solid tantalum capacitors might be considered. To show how easy it is to consider these alternatives, we can draw another reactance plot as shown in Figure 25. Here we have assumed the same paralleled ceramic capacitors for Cb at the CPU socket to accommodate the highest frequencies, but defined a grouping of 47uF, 16V, type 593D, surface mount chip capacitors for Co.

The dimensions for these units are 7.3 x 4.3 x 2.8 mm, significantly smaller than the aluminum electrolytics chosen above, but it will still take six of these tantalum capacitors in parallel to keep the ESR below 18 mΩ. The chart makes it readily obvious that ESR, rather than capacitance or inductance, is the determining parameter. On the other hand, with a capacitance of only 282uF total for Co, the bandwidth of the regulator must now be moved up to 40kHz and, while this will take a little more effort in optimizing the gain stages, it is still a

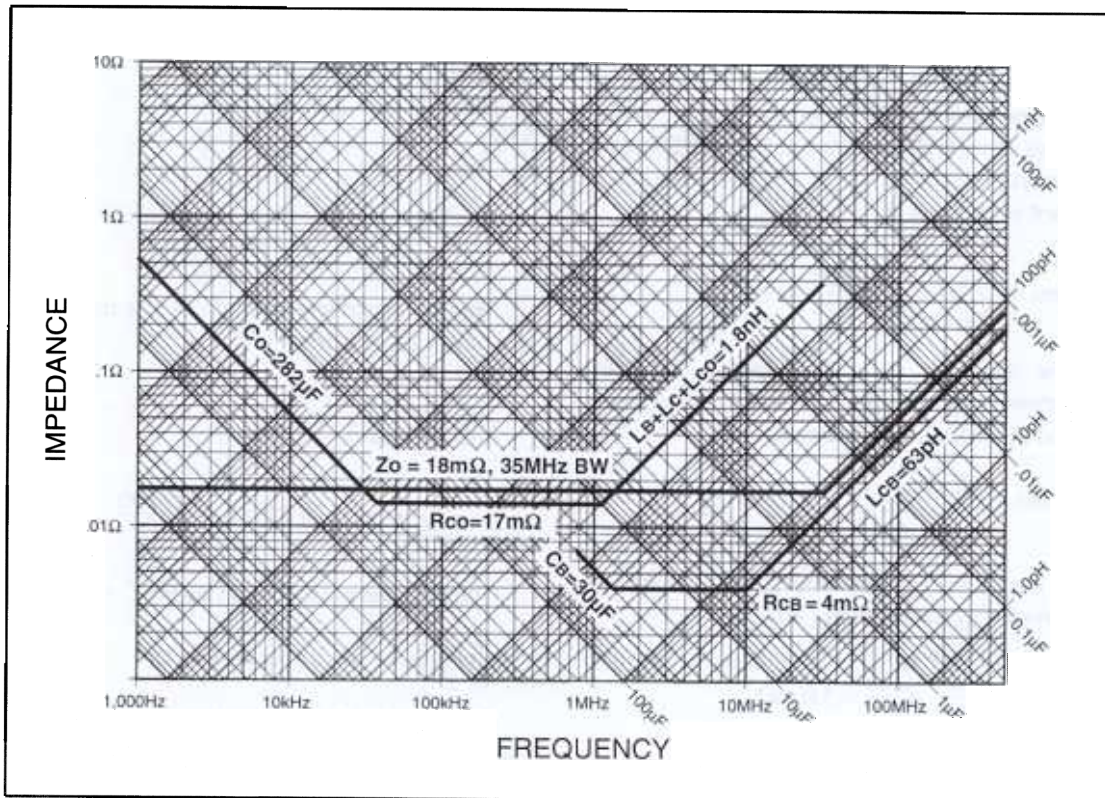


Figure 25 - A reactance plot using solid tantalum chip capacitors for the regulator's output capacitors



viable option, and the resultant module will be considerably smaller.

## CONCLUSIONS

The point of these discussions has been to present a quantitative picture of both the requirements which the new "megaprocessors" have imposed upon their power systems along with a design process which, by taking a system-level point of view, can reduce design tasks to a series of very manageable steps. This has shown that the most critical parts of this solution lie not with the power conditioning circuitry but with the interconnecting network and, by properly engineering that portion of the power system, the actual regulator design becomes much more manageable.

While the use of reactance paper offers a simple and insightful approach to defining an acceptable power distribution network, this is not to say that computer simulation has no place in this analysis. In fact, it is always worthwhile to solve the same problem by different means to confirm the results. In this case, the parameters of our first example were incorporated into the simulation model shown in Figure 26 and the results - now in the time domain - are plotted in Figure 27. Note that it takes three different time scales to view the range of events caused by the load transient but the overall goal of allowing less than  $\pm 3\%$  load voltage deviation has been substantiated.

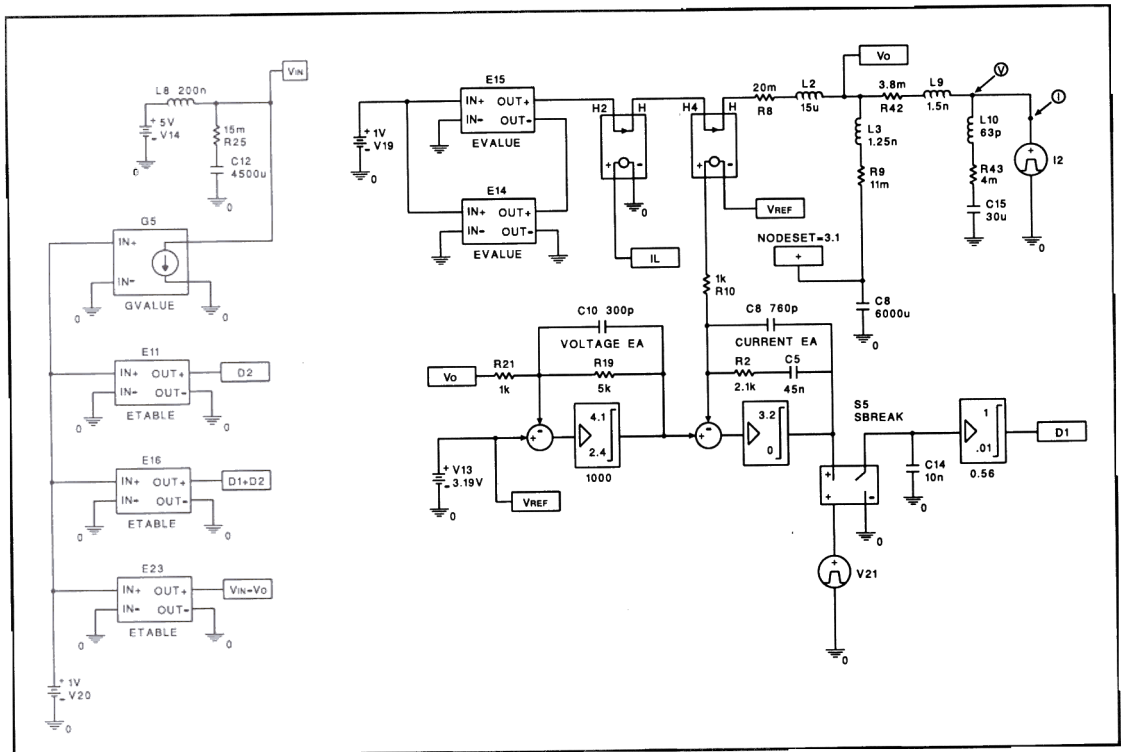


Figure 26 - A simulation model for the voltage regulator and its distribution network




Figure 27 - Load voltage deviation caused by a 10 Amp, 10 nsec change in load current

## REFERENCES

1. Pentium® Pro Processor Power Distribution Guidelines, Application Note AP-523, Intel Corporation, Nov 95.
2. Electronic Circuits, Mohammed S. Ghausi, D. Van Nostrand Company, pp 235-261.
3. Principles of Feedback Control, Volume 1, Feedback System Design, George Biernson, John Wiley & Sons, Section 1.1.
4. Care and Feeding of the CPU, Unpublished manuscript, Phillip C. Todd, Power Innovations, Long Beach, CA.
5. Converter Optimization for Powering Low-Voltage, High-Performance Microprocessors, John O'Conner, APEC '96 Conference Proceedings.
6. Fueling the Megaprocessors - A DC/DC Converter Design Review Featuring the UC3886 and UC3910, Larry Spaziani, Unitrode Application Note U-157.

