Edward Herbert Designs

Switched Current Power Converter

Revised March 13, 2006

Solutions for

Advanced Performance

Power Supplies

http://anagenesis-inc.com/Herbert.html

http://eherbert.com

On behalf of Edward Herbert

For more information, contact:

Arnold Alderman
Anagenesis, Inc.
- Los Angeles Office -
222 N. Sepulveda Blvd
Suite 2000
El Segundo, CA, 90245

Tel: 310 662 4764
Fax: 310 662 4767
Mobile: 310 704 7079

Email: arnold.alderman@anagenesis-inc.com

Anagenesis, Inc., and Arnold N. Alderman, President are not patent attorneys or lawyers, and it is not expressed or implied that any legal services or legal opinions are offered on behalf of the client. Any reference to legal issues are only to help inventors and interested parties to better understand how these issues and patents effect the marketing aspects of invention commercialization. We highly recommend that you consult a qualified attorney regarding the law and patenting.

Abstract:

The switched current power converter (SCPC) can change its output current as fast as solid-state switches can change state, which is very fast. It has no theoretical di/dt limit, and is a true digital power converter.

- *Di/dt of 1000 A/us or greater.*

The switched current power converter can be augmented with switched charge circuits, which transfer a precise charge to or from the output capacitor as a pulse, to raise or lower the output voltage as a precise step function. This makes it practical to design a power converter that can change its output voltage very quickly in precise steps, even if there are simultaneous load changes. For example:

- *1.0 V, 20 A to 1.2 V, 100 A in 1 us or less*
- *1.2 V, 100 A to 1.0 V, 20 A in 1 us or less*
- *0 V, 0 A to 1.0 V, 100 A in 2 us or less*
- *1.0 V, 100 A to 0 V, 0 A in 2 us or less*

Because the switched current power converter has such a fast response time, the output capacitor can be significantly smaller, as much as an order of magnitude smaller than the output capacitors of present VRMs of comparable current output. This is a significant cost savings, and improves the MTBF as well.

- *Output capacitance: 500 uF or less.*
- *Bulk capacitors are not used.*

The design is modular, which simplifies the circuit board design and allows lower impedance interconnections.

- *For steady state conditions, the switching frequency is low typically 100 kHz. Switching losses are significantly reduced.*

High efficiency and low standby power are very important in many applications. An optimized switched current power converter provides very good efficiency, line cord or battery to processor. As important, the very fast transition times allow the processor to stay in a low power state more of the time, or turn off entirely, greatly reducing power consumption and thermal loading.

A very fast response requires very fast determination of the output voltage. A voltage measurement technique using total charge ΣQ is described.

Table of Contents

1.0. Introduction:

The switched current power converter has no theoretical limit for its di/dt. The necessity to ramp up a current in an inductor, and the time required to do so, are not a factor. The response time is determined by the propagation delay through the output voltage measurement, the digital logic and the MOSFET switching time, all of which are very fast.

With the addition of switched charge circuits, the output voltage may be stepped very rapidly and precisely, even if the output current is stepped as well.

The switched current power converter is covered by US. Patent No. 6,121,761, "Fast Transition Power Supply", issued September 19, 2000. Other patents are pending.

1.1. Switched Current Power Converters, Overview:

The principle of the switched current power converter is shown in figure 1.1.1. A number of current sources have switches on their outputs, so that the current from each current source may be directed to the output capacitor and the load, or removed, as quickly as the switches can change state.

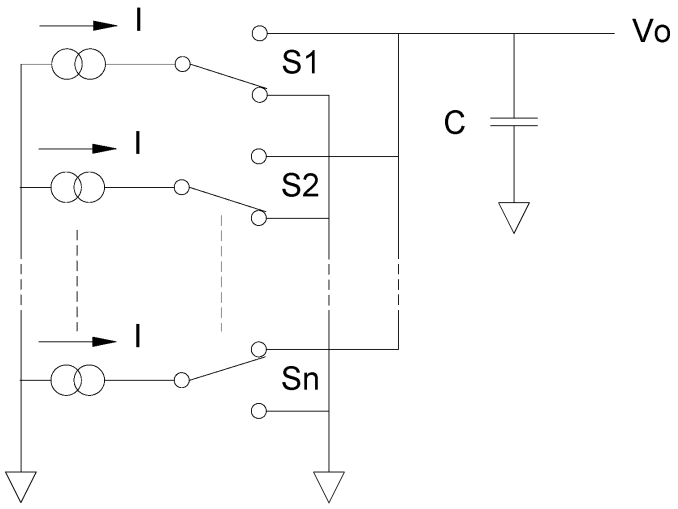

Figure 1.1.1. In the switched current power converter, the current may be directed to the output capacitor and the load, or removed, as fast as solid-state switches can change state.

The output impedance of the SCPC is less than 0.5 mΩ to 5 MHz. The SCPC was modeled in SPICE with the load and parasitic impedances of Intel ® VR10.2 (except that the bulk capacitors were not used). For each frequency on the graph in figure 1.1.2, the load was varied as a 50 A p-p ac current on a 60 A dc current, so the load varied from 35 A to 85A. The impedance at each frequency

is calculated as dV/dI. See: [Switched Current Power Converter; Transient and Frequency Response (pdf)]

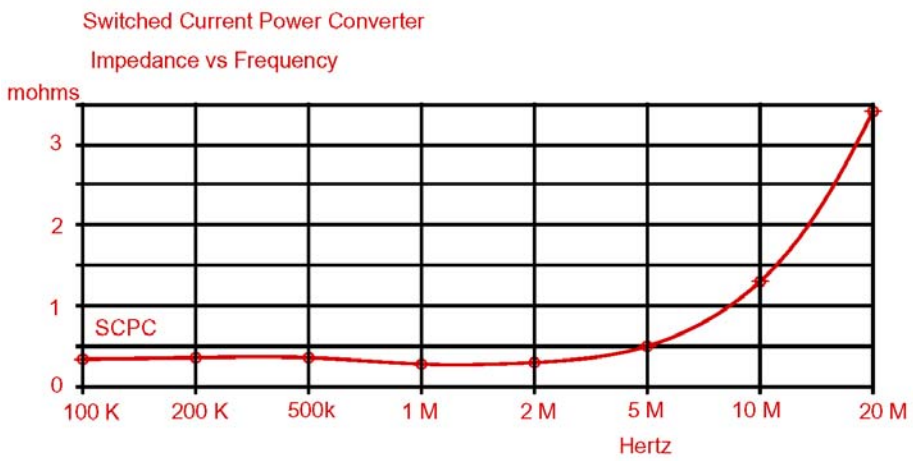

Figure 1.1.2: The response of the SCPC is flat to over 5 MHz. The large signal output impedance is less than 0.5 mΩ to 5 MHz.

With its very fast response, the bulk capacitors usually associated with a processor power converter are not used. Intel ® VR 10.2 suggests that twelve to sixteen 560 uF capacitors be used. This saves significant cost and motherboard real estate.

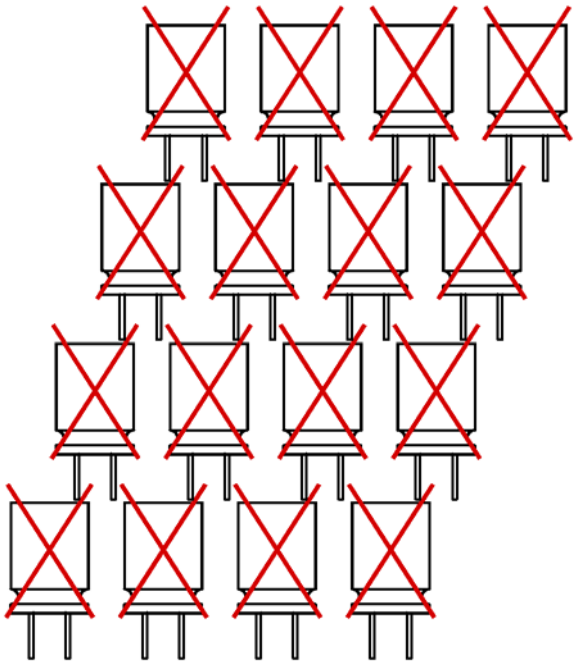

Figure 1.1.3. No bulk capacitors are used with the SCPC. With its very fast response, only the MLCCs are needed, mounted near the processor. The SCPC module itself uses only 20 uF of decoupling capacitors on its output.

1.2 Current Sources, Overview:

The choice of the current sources is very flexible and the type selected will have very little impact on the output characteristics and control. The current sources provide constant dc currents, and as long as the currents are maintained, the control and dynamics of the output stages are not affected. Accordingly, the design of the current sources can be optimized for efficiency without compromising the frequency response of the outputs. By choosing among the available current source designs, the switched current power converter can be adapted to a variety of circuit environments.

For breadboard studies, several laboratory power supplies configured in constant current mode can be used. There are current source designs based upon present multiphase buck converter designs, but for new designs, a transformer coupled SCPC is suggested.

Several representative current source circuits and modules are shown later in this presentation.

1.3. Switched Charge Circuits, Overview:

Switched charge circuits are easily added to a switched current power converter with no modification to the underlying SCPC circuits, to provide a rapid, precise step in the output voltage. A fixed charge Q is added to, or removed from, the output capacitor as a pulse when a switch changes state, which raises or lowers the output voltage very rapidly by a precise amount.

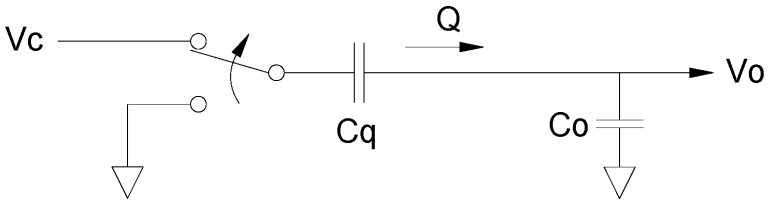

Figure 1.3.1. When the switch is closed to the charging voltage Vc, a fixed charge Q is transferred to the output capacitor Co as a pulse, raising the output voltage Vo a precise amount as a step function. When the switch is closed to ground, the same charge Q is removed, for a step decrease in the output voltage.

One or more switched charge circuits may be used. For instance, one switched charge circuit can be used to step rapidly and precisely from zero V to the minimum VID output voltage. Another can be used to step rapidly and precisely from the minimum voltage to a higher VID voltage, for “turbo” mode, and back.

Very significant power is saved in a computer system by operating the processor at reduced voltage. A power converter that steps rapidly and precisely between voltage levels allows the processor to be in a lower power state for more of the time. More power yet is saved if the processor is turned off entirely.

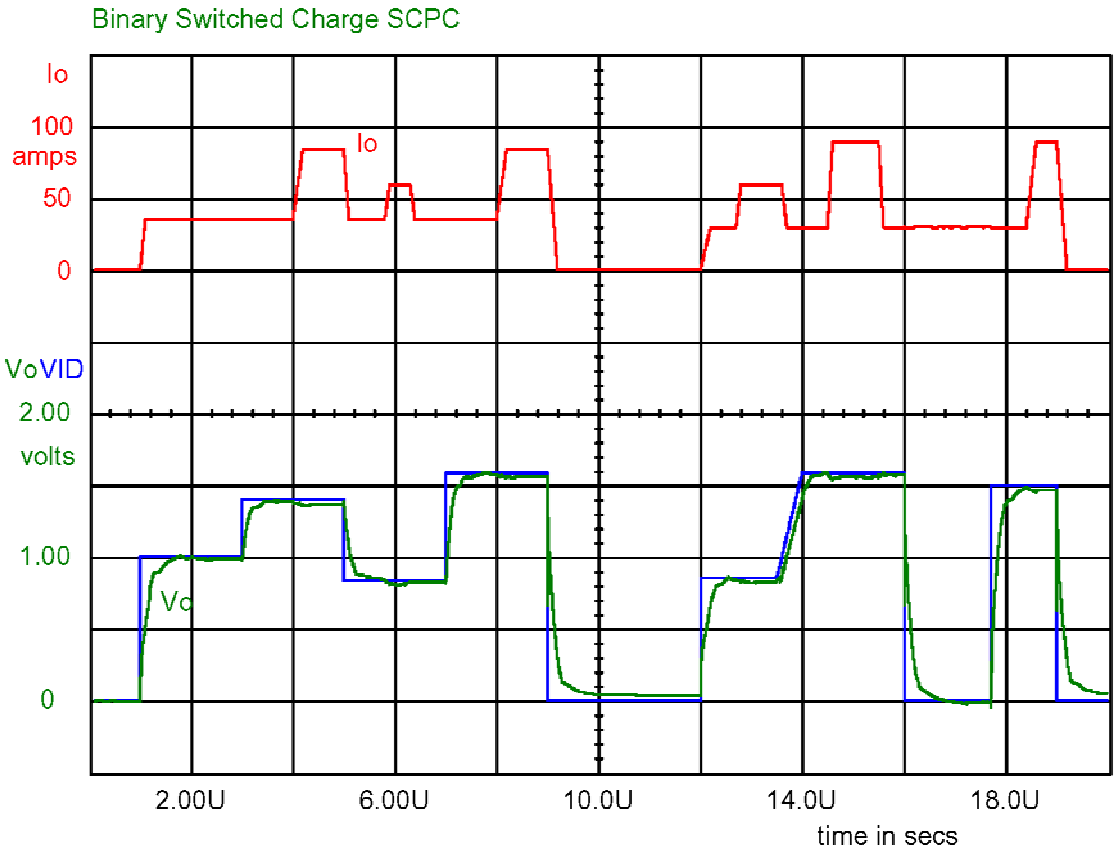

Figure 1.3.2. Switched charge can be used to step the output voltage very rapidly and precisely, irrespective of the load current. The above graph shows a SPICE simulation of the voltage response of the SCPC. The blue graph VID represents a digital VID input, and the green graph Vo shows that the output voltage of the SCPC tracks well with no overshoot. The red graph Io shows simulated load changes. Note, at 9 us, that the load can be dumped and the voltage can go to 0 V simultaneously without overshoot.

Figure 1.3.2 shows a SPICE simulation of a SCPC with binary-switched charge circuits. The SPICE model uses the output load and parasitics of Intel ® VR10.2, except that the bulk capacitors are not used. In the simulation, the blue graph VID simulates a digital VID input. The VID is stepped in various increments, including from and to 0 Vdc, and the green graph Vo shows the simulated response. At about the 13 us time, the VID is slewed at 1,500 mV/us by rapidly sequencing the simulated VID input. The red graph Io shows simulated load changes, and the load regulation due to the output impedance can be seen if the Vo graph is inspected closely.

Note in particular, at 9 us, that the VID can go to 0 Vdc simultaneously with a load dump. The output voltage follows within 2 us, with no overshoot.

The switched charge does not regulate the output voltage, it can only produce the very fast and accurate voltage step. The VID to the SCPC must be changed simultaneously, and after the step, the SCPC controls to the new voltage level.

1.4. Total Charge Measurement and Current Control:

A new voltage and current control scheme for the switched current power converter is based upon a measurement of total charge using a flash analog to digital converter, with the outputs of the flash d-a converter directly controlling the switches of the current sources. The theory of operation is shown later in this presentation.

A power converter that can go from 0 V and 0 A to any VID and full output current in 2 us, and which can respond to load changes or step changes in voltage in tens of nanoseconds requires a very fast and stable measurement and control circuit. Usual feedback circuits with their lags and compensation cannot be used. The current must be controlled precisely, yet current measurement cannot be used for the control, as it is much too slow. Voltage measurement is problematical as well, if the circuit has any parasitic inductance, as it can take a long time for the output voltage to settle down and be stable, far too long for control use.

Measurement of the total charge on the distributed output capacitors provides an elegant solution. It is very fast and unconditionally stable. The theory of operation is discussed later in this presentation.

It is also important that the control methods be adaptable to diverse specification requirements. It is usual in a processor power supply to have the output voltage drop slightly as the current increases. Fortunately, the control scheme is well adapted to this characteristic. Circuit modification can be used for a fixed voltage output or with remote sense.

1.5. Power Delivery System Parasitic Impedance:

As in any power converter for a processor, it important to minimize the parasitic impedances in the power delivery system.

The impedance of present connectors is a serious problem as well, particularly the parasitic inductance. The total charge measurement significantly reduces the effects of the series impedance.

1.6. Modular Design:

The switched current power converter uses identical parallel circuits, dividing the current among them. The power switches are closely integrated with the magnetic components as modules, simplifying the circuit board layout and providing lower impedance interconnections.

Many of the most difficult circuit problems in the design and layout of power converters having very fast di/dt relate to the square of the current (I^2), so the advantages of a modular design with a lower current per module are significant.

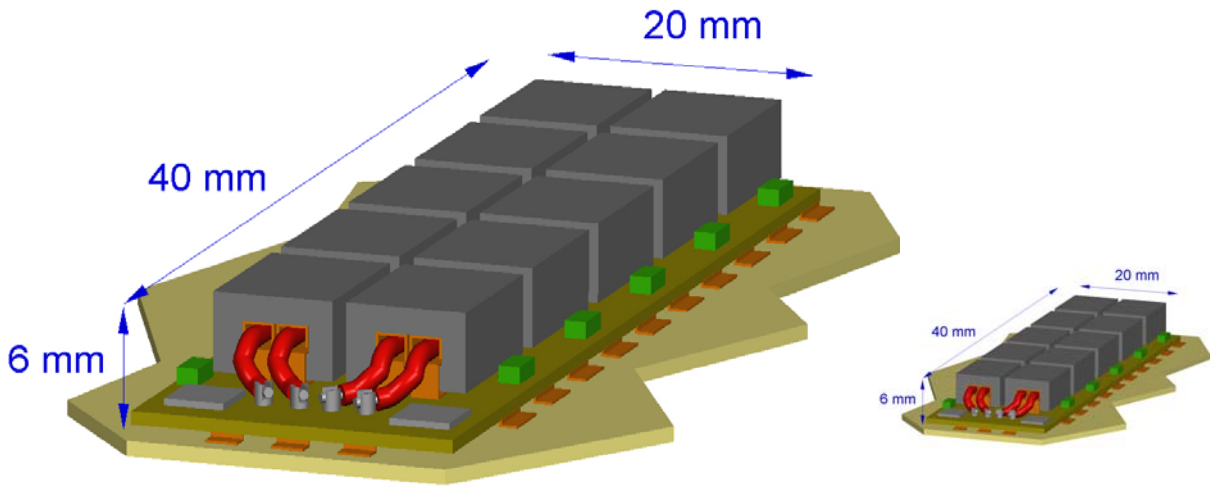

Figure 1.6.1. A module for a switched power converter contains the magnetics and its associated circuitry on a surface mount header.

1.7. Simple Power IC.

The logic of the SCPC is very simple, and is easily implemented in a logic and control integrated circuit (IC). Much more important to the commercial success of the SCPC are the current handling components. All of the power MOSFETs and their low side drivers are integrated in a simple power IC that can be manufactured with just a few extra processing steps over those required to manufacture a DMOS MOSFET. For a low current SCPC, all of the power components can be integrated into a single simple power IC. For higher currents (100 A), five dual ICs are suggested. A single master control IC contains the logic and the high side drivers. See: Simple Power IC for the Switched Current Power Converter: Its Fabrication and Other Applications, Edward Herbert, March 2, 2006.

This compares very favorable to present VRMs and VRDs, such as four phase buck converters, particularly as no bulk capacitors are used.

1.8. Breadboard Test Results:

A paper was given at the IBM Platform Technology Symposium, September 14-15, 2004, Kahler Grand Hotel, Rochester, MN.:

> "Fast Transient Power Converter Using Switched Current Conversion" by Laurence McGarry, Advanced Engineering Technology Manager, Hong Kong & China, Astec Power, A Division of Emerson Network Power.

This paper is posted on the Internet, and can be found by doing an Internet search. The test results are respectable for a first breadboard, but a number of compromises were made for expediency. Much better results are anticipated with a design that is optimized for a particular application.

2.0. Switched Current Power Converters; Theory of Operation:

2.1. Current Control in Switched Current Power Converters:

In most present power converters, a change in the output current requires that the current in an inductor be changed by the same amount. Thus V/L, the ratio of the driving voltage V to the value of the inductor L, limits the di/dt.

In contrast, in the switched current power converter, the current can change, up or down, as fast as solid-state switches can change state. In a generalized design, a number "n" constant current sources can be switched either to the output capacitor and the load or to the return. See figure 2.1.1.

Control techniques are shown that provide very high bandwidth (greater than 5 MHz) and are unconditional stability. The “natural” output characteristic of the switched current power converter has a decreasing voltage with increased current (positive impedance), but modifications can reduce the impedance or provide a fixed output voltage (at the output or at a remote sense point).

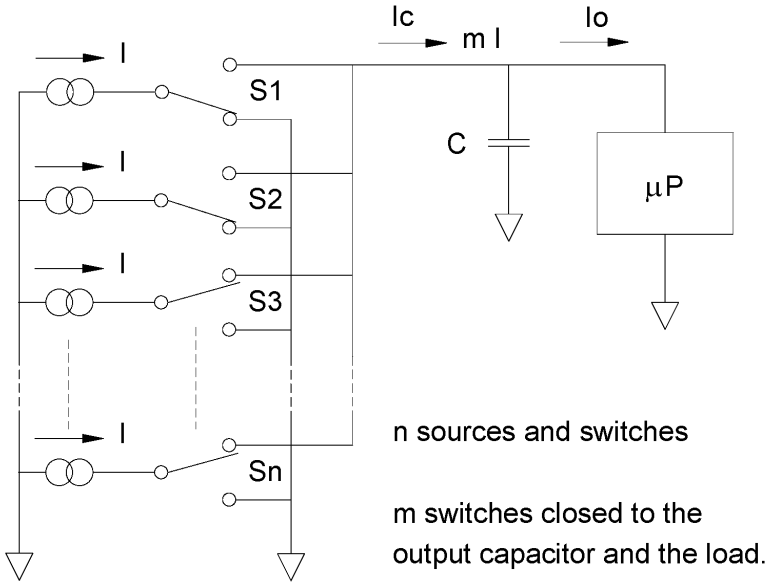

Figure 2.1.1. In a switched current power converter having n current sources and n switches of which m are switched to the output capacitor and the load, the converter current is m I, and the maximum current is n I.

2.1.1. Flash Analog to Digital Converter:

A very fast power converter current control requires very fast voltage sensing, and one of the fastest ways to measure voltage is with a flash analog to digital converter, a series of comparators with a resistor ladder network as their

references. The flash converter can directly drive the current switches of figure 2.1.1, as shown in figure 2.1.2.

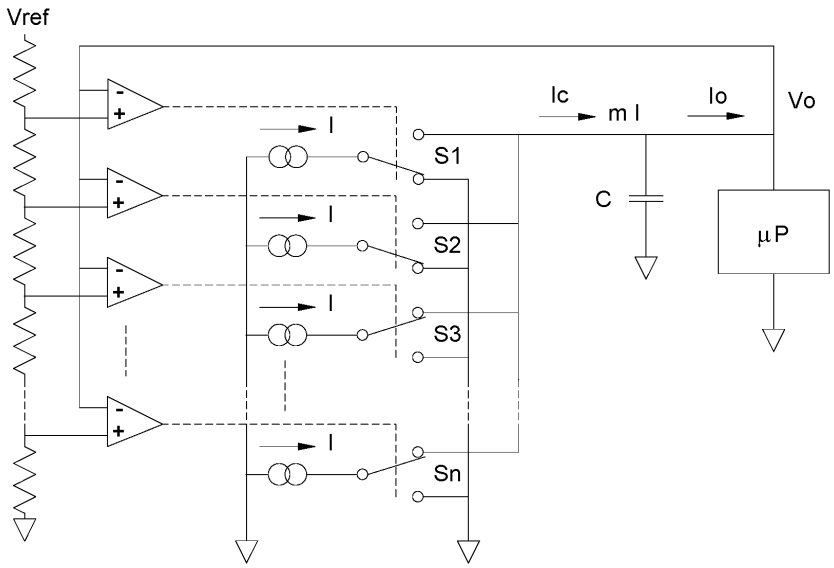

Figure 2.1.2. A flash a-d converter can directly drive the current switches to control the power converter current Ic.

(Note: although output voltage sensing is shown for illustration, the output voltage is not a suitable control measurement unless there is only one capacitor and no parasitic inductances, not a "real world" condition. In a practical SCPC, the total charge on the output capacitors is measured for control).

2.1.2. The Voltage Decreases with Increased Output Current:

The natural characteristics of a switched current power converter that is directly controlled by a flash a-d converter has a decreasing output voltage Vo with increasing load current, Io, as shown in figure 2.1.3. This characteristic is often specified for power converters for processors. By dividing the voltage droop by the maximum current, a characteristic impedance Rs is defined.

To understand this, please consider figure 2.1.2. The output voltage Vo is taken to a series of comparators, and the references of the comparators is a resistor ladder network that defines the range of the output voltage and divides it into small increments. When the output voltage Vo is below the threshold reference of all of the comparators, all of the currents will be switched to the output capacitor and the load, and the output current will be 100 percent. As the output capacitor charges up, the first threshold voltage will be reached, one switch will change state and the output current is reduced one increment.

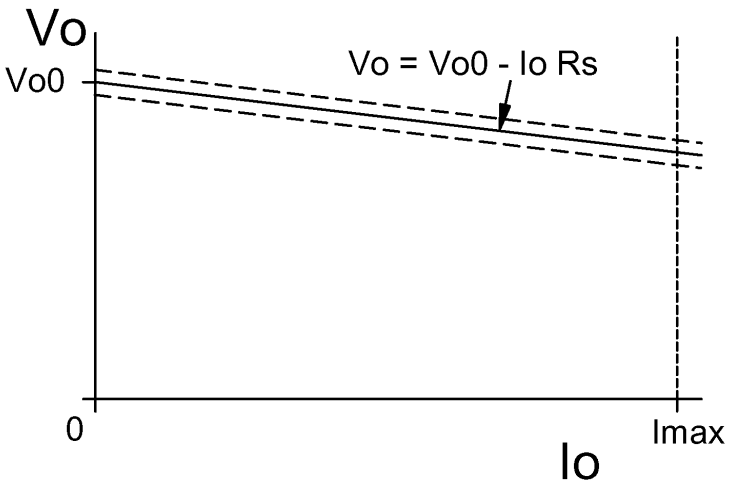

Figure 2.1.3. The natural characteristic of a switched current power converter with a flash a-d converter directly controlling the current sources has a decreasing voltage with increasing current. Often the specifications for power converters for processors require this characteristic.

As the output capacitor charges more, successive threshold voltages will be reached, and as each is reached, the output current reduces one more increment. Eventually, the converter current equals the output current, and the voltage of the output capacitor will not rise any further. If the load current increases, the voltage of the output capacitor will drop. As the voltage drops, successive current increments will be turned on until a new equilibrium is reached.

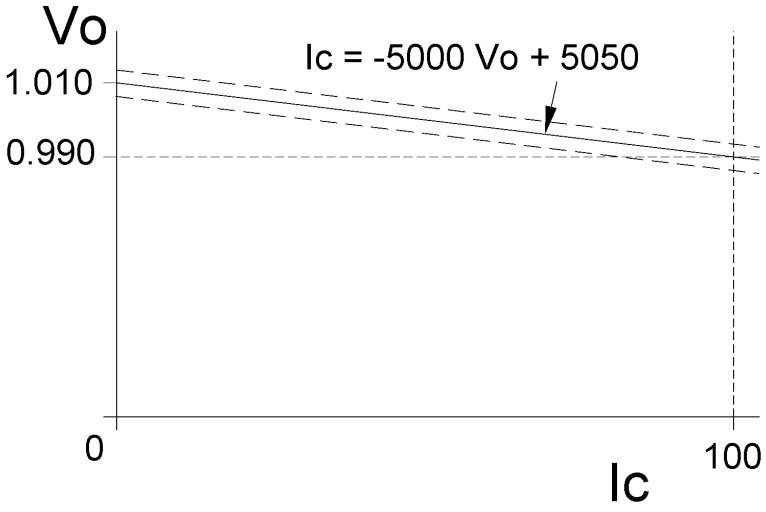

Figure 2.1.4. Characteristics of a representative switched current power converter is shown. The voltage decreases 0.020 V as the current increases to 100 A, for a conversion impedance of 0.2 mΩ.

In this manner, the output current of the switched current power converter is controlled precisely without any need to measure the current.

2.1.3 Current is a Function of the Output Voltage:

For a power converter having the characteristics of figure 2.1.3, the converter current Ic can specified as a function of the output voltage. For the example of Figure 2.1.4, the converter current Ic is given by the equation:

$$Ic = -5000\ Vo + 5050$$

This function can be approximated by the following equivalent circuit.

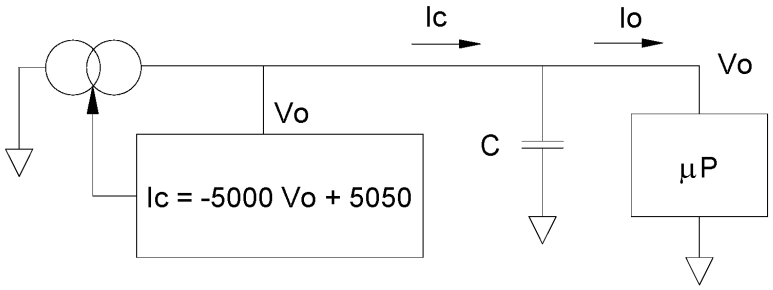

Figure 2.1.5. In a power converter having a decreasing voltage with increased current, the converter current Ic can be directly controlled as a function of the capacitor voltage, Vo.

A voltage controlled current source is an ideal circuit element, one that is implemented easily in SPICE. It is useful as a first SPICE model when getting familiar with the circuits, and as an average state model.

2.2.4. Current is in Discrete Steps:

To model power converter of figure 2.1.1, the SPICE voltage controlled current source of figure 2.1.5 can be modified to reflect the digital step-wise converter current characteristics of the SCPC. For the example of a switched current power converter having ten current sources of 10 A and ten switches, the following formula for the power converter current Ic can be used for SPICE:

$$Ic = 10\ int((-5000\ Vo + 5055)/10)$$

where “int” is the integer function, which drops the decimal portion of the expression following in parentheses. This expression constrains the power converter current Ic to whole 10 ampere increments in a SPICE model.

A complete SPICE model must include current limits at zero and the maximum current as well as hysteresis. This is beyond the scope of this presentation, but a number of SPICE models and their simulation results are presented in ”Switched Current Power Converter; Transient and Frequency Response (pdf)".

2.1.5. Load Regulation:

The characteristic of decreasing voltage with increasing load results in an apparent series impedance, and will result in the load regulation shown in figure 2.1.6.

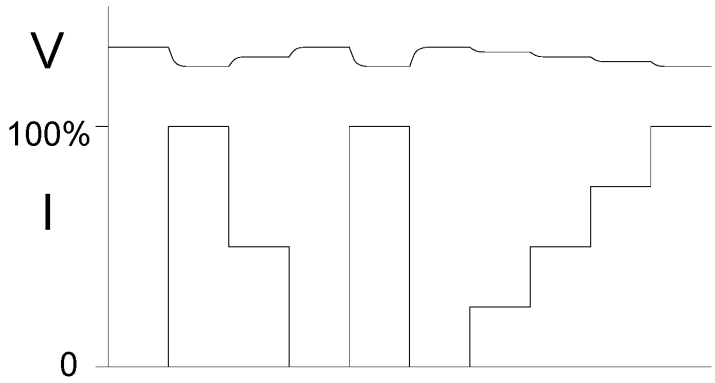

Figure 2.1.6. Following a step change in current, the output voltage transitions quickly to its new steady state value conforming to the voltage vs. current specified in figure 2.3.

2.1.7. Modulating the Output Current:

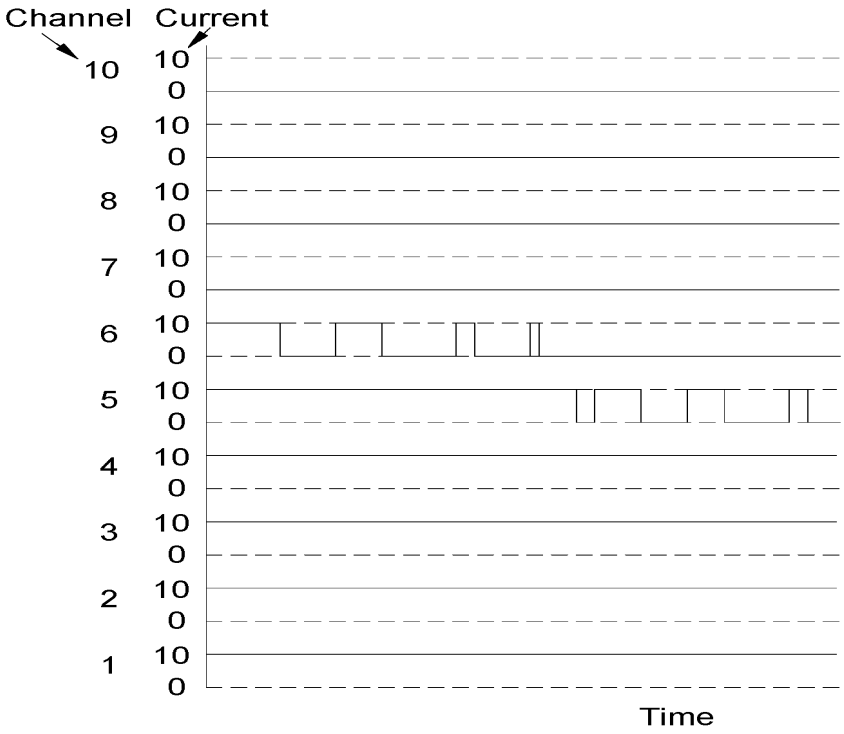

Figure 2.1.7. For an output current that is between 10 A increments, one of the ten output currents is modulated so that the time averaged output is correct. In the example shown, the current begins at 56 A and decreases to 47 A.

The output current Io is determined by the output voltage and the impedance of the load. To maintain a steady voltage on the output capacitor and the load, the switched current power converter current Ic modulates between the nearest whole 10 A increments so that the average current over time equals the output current Io. For example, with an output current of 56 A, the converter current Ic must equal 56 A on average, which is accomplished by modulating between 50 A (for 40 percent of the time), and 60 A (for 60 percent of the time).

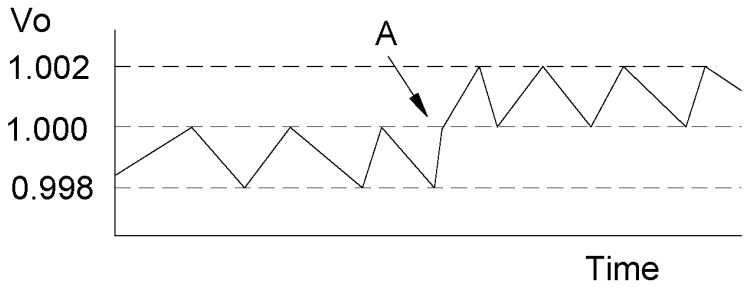

Figure 2.1.8. The ripple voltage for the currents of figure 2.1.7 is shown. Within a voltage band between comparators, the corresponding current turns on if the voltage goes below the lower voltage limit and it turns off at the higher voltage limit. Note, at point A, even though the current turned off (see figure 2.1.7), the voltage continued to rise, because the load current is decreasing. The next higher voltage threshold was reached and modulation switched to the corresponding next lower 10 A current band.

When the comparators directly control the current switches, as in figure 2.1.2, one of the switches will cycle at a fairly high rate, typically 1 MHz, while the remaining switches are not switching, they are either on or off and remain so. The one switch that is cycling conducts a fraction of the total rated current, typically 10 A. The voltage is low, the output voltage Vo, typically about 1 V.

Therefore, only one of the switches has high frequency switching losses. By comparison, in a four-phase buck converter, all four switches switch every cycle, with a varying current equal to one fourth of the output current. The voltage is higher, equal to the input voltage, typically 12 V.

2.1.8. Distributed Switch Timing.

The timing of figure 2.1.7 has the elegance of simplicity, and it is the natural timing of the circuit of figure 2.1.2. However, it has the disadvantage that, at a fixed steady state current output, one switch has the burden of the modulation, and it has all of the switching losses. Because any one of the switches could be modulating continuously at some steady state condition, all must be able to handle the stress and heat dissipation of this condition. The switches may have to be oversized, and they may require heat sinking.

It is preferred to cycle through the switches periodically, as shown in figure 2.1.9. There are a number of ways to cycle the switches, the details of which are beyond the scope of this presentation.

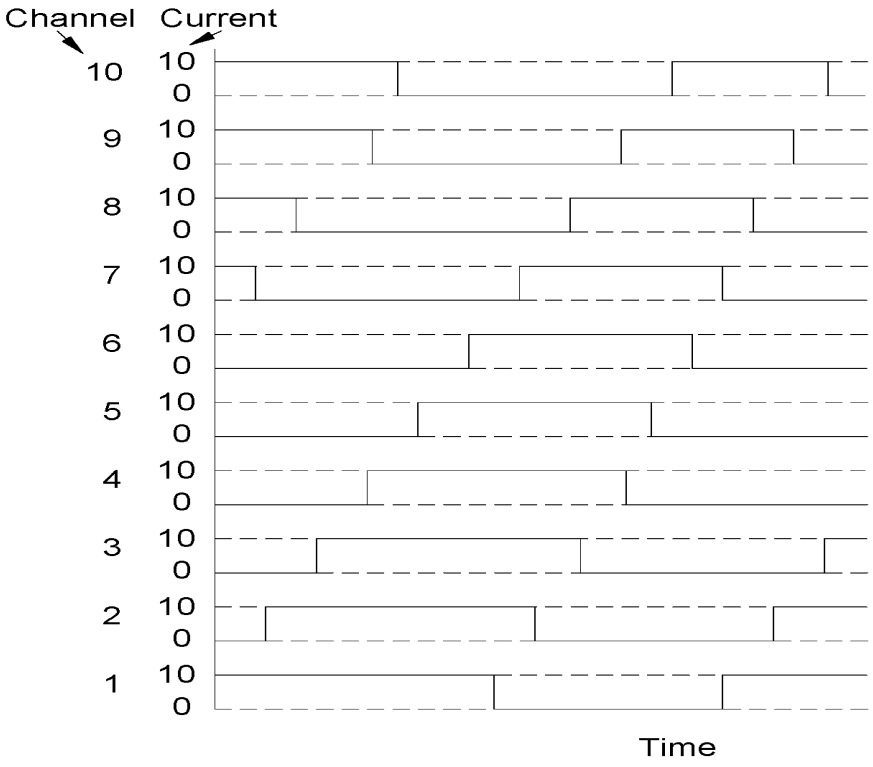

Figure 2.1.9. Using the steering logic, the switching is distributed among the switches so that no one switch has the burden and the entire switching losses of the current modulation as in figure 2.1.7.

If the natural switching is 1 MHz, then by using cyclic ring-counter timing, each switch has one tenth of the switching frequency, typically 100 kHz. The duty-cycle of each is proportional to the output current (percent of full load), and the voltage is low, typically about 1 V.

The very low switching frequency reduces switching losses. Further, since the switched current is always the same, the gate drive can be optimized for that current.

2.2. Whole Ladder Network Hysteresis:

It is common to use hysteresis in a comparator switching circuit, to prevent cycling at a high rate around the set-point. "Whole ladder hysteresis" is a technique used to provide reasonable hysteresis around any incremental set-point on the resistor ladder network without spreading the outer limits of the control band too much.

2.2.1. Hysteresis Around a Comparator:

The circuit of figure 2.2.1 is a familiar analog circuit. When the voltage Vx is below the threshold voltage and rising, the comparator will trip at a higher voltage than when the voltage Vx is above the threshold and falling. When the output of the comparator is low, it pulls down the voltage on the non-inverting input of the comparator, changing the trip point to provide the desired hysteresis.

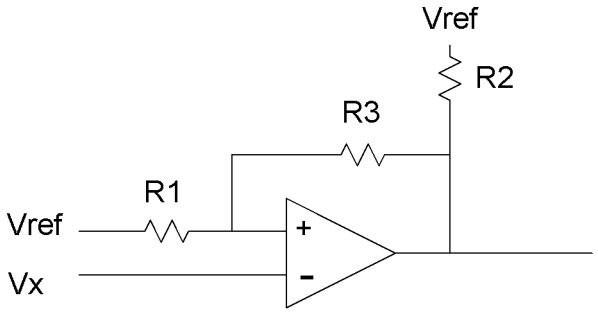

Figure 2.2.1. A comparator having hysteresis is shown. It is assumed that the comparator has an open drain output and the load impedance is much greater than R2.

2.2.2. Whole Ladder Hysteresis:

When the voltage reference for the comparator is a resistor ladder network, the hysteresis affects the voltage of the entire ladder network, a characteristic that can be used to advantage. See figure 2.2.2.

A number n comparators have as their references the nodes of a resistor ladder network. The outputs of the comparators are assumed to be open drain for this discussion. The impedance of the ladder network at the nodes of the ladder network is largely determined as the parallel combination of Ran and Rbn, as the resistors Rd1, Rd2, Rd3 etc., are much lower in value. The voltage relationships in the circuit are straightforward, but because there are so many nodes, the computations are involved and best done by a computer routine. The equations can be written and solved by a mathematics program, but trial values in a SPICE model works, too.

2.2.3. Hysteresis and the Voltage Increments of the Ladder Network:

The switched current power converter, when controlled directly with a flash a-d converter, is a hysteretic control. Hysteresis is added to keep the hysteretic control from cycling rapidly about the set points. The hysteresis also defines the ripple voltage in the output, so an early design trade-off is the acceptable ripple voltage versus the output capacitor value and the switching frequency..

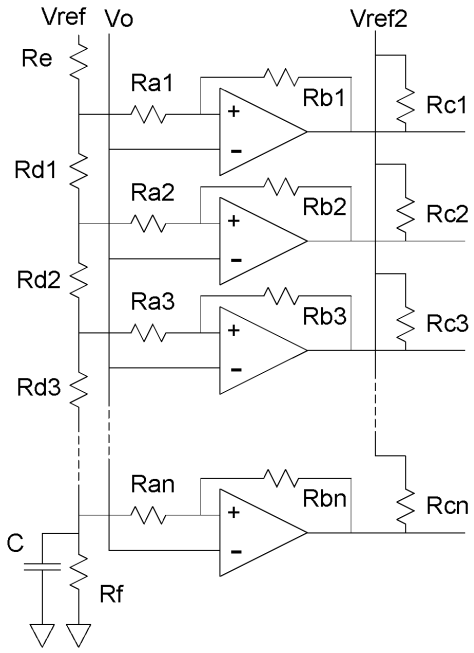

Figure 2.2.2. When the voltage reference for the comparators is a resistor ladder network, the hysteresis resistors affect the voltages of the entire resistor ladder network, a characteristic that can be used to advantage.

The switch cycle frequency changes with load, but the highest switching frequency occurs when the output current is midway between the step current values. As an example, for a SCPC having a maximum current of 100 A divided into ten current sources of 10 A each, the maximum switching frequency occurs at the “5’s”, or 5 A, 15 A, 25 A, etc. A larger output capacitor reduces the switching frequency, but it is preferred to keep the output capacitor small.

As an example, consider a hysteresis voltage of 5 mv was chosen as the best trade-off. The resistor network will be designed such that, as any comparator switches, its non-inverting input will change 5 mv. For a steady state output current the converter current switch between integer steps with a 5 mv hysteresis band, and there will be 5 mv output ripple.

With a changing current, for example, a decreasing output current Io, there will be a time when switching the comparator does not cause a voltage reversal, as the converter current will continue be too large and the output voltage will

continue to drift higher until the next higher threshold is met. See figures 2.1.7 and 2.1.8 as an example. The voltage steps in the ladder network should equal the hysteresis voltage, so, for our example, the next threshold will be 5 mv higher.

Accordingly, the chosen hysteresis determines the voltage increments between steps of the resistor ladder network. This, in turn, determines the amount the voltage will droop as the current goes to maximum, which determines the "impedance" of the power converter. In the hypothetical example, with ten increments of 10 A each, and a 5 mv hysteresis, the natural voltage droop will be 50 mv, and the impedance is 0.050 v/100 a, or 0.5 mΩ.

2.2.4. Effective Voltage Increments with Whole Ladder Hysteresis:

In figure 2.2.2, each of the n comparators has a resistor Ra1, Ra2, Ra3 --- Ran on its non-inverting input. If this resistor is very large compared to the resistance of the ladder network, the hysteresis is applied to the individual comparator with little effect on the ladder network as a whole. If it is smaller, however, the hysteresis will affect the voltage of the whole resistor ladder network, resulting in an effective compression of the voltage steps. As each comparator goes low in turn, the voltages of the nodes of the whole ladder network are reduced by about the same amount. This has the effect of "flattening" the voltage-current characteristic curve, reducing the converter impedance. However, if excessive, it will make the power converter unstable.

To understand this, consider an example in which the third comparator of figure 2.2.2 switches low. Its hysteresis resistor Rb3 will pull down the non-inverting input of the third comparator by 5 mv. This will also pull down the voltages of the entire resistor ladder network through the resistor Ra3. If the incremental voltage between rnodess of the ladder network is 5 mv, if when a comparator goes low, it reduces the voltage of the ladder network by 4 mv, the effective step voltage of the ladder network is reduced to 1 mv. The droop in going from no load to full load will be 10 mv, and the converter impedance is effectively to 0.1 mΩ.

Note: The circuit of figure 2.2.2 is a simplified circuit, to explain the principle of operation. It is not a practical circuit for implementation as an integrated circuit, as the required resistor values are much too high. The detail design of practical circuits is beyond the scope of this presentation.

2.2.4. Stability Concerns:

Taken to an extreme, if the whole ladder hysteresis voltage equals the step voltage of the ladder network, the voltage current characteristic curve becomes flat. The next comparators reference voltage is pulled down to equal the output voltage Vo, each in turn, and all of the comparators will trip. SPICE simulation shows that the switches will bang between all on and all off for this condition.

Therefore, the whole ladder hysteresis must always be less than the step voltage increment of the resistor ladder network. As the whole ladder hysteresis approaches the step voltage increment of the resistor ladder network, the system becomes overly sensitive to noise and stability is compromised.

The capacitor C in figure 2.3.2 provides increased noise immunity, as it prevents the ladder network node voltages from changing instantaneously. Dynamically, the comparators see the steps of the ladder network, but as the R-C time constant settles to its limit, a flatter voltage vs. output current characteristic (lower impedance) can be achieved for steady state.

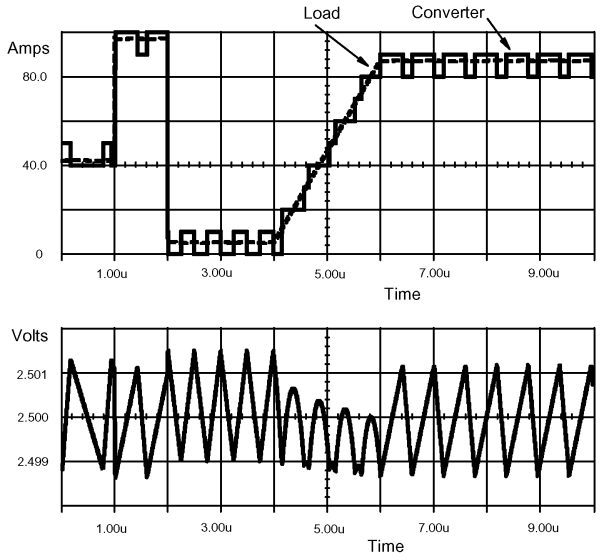

Figure 2.2.3. SPICE model oscillograph showing the output load current, the power converter current and the output voltage. The control is a flash analog to digital converter with hysteresis acting on the whole ladder network. The output voltage droop is nearly flat.

Figure 2.2.3 shows the simulated oscillographs of a representative SPICE model of a switched mode power converter with whole ladder network hysteresis. The natural voltage increments of the resister ladder network is about 2.5 mV, but as each comparator trips, its hysteresis pulls the whole ladder down by about 2.4 mV, delayed by the time constant of the voltage on the capacitor C.

In the SPICE simulation, the load current starts at 42 A, then steps to 97A after 1 µS, then steps to 5A at 2 µS, then ramps to 87 A from 4 to 6 µS, a where it remains for the remainder of the simulation, 10 µS. The power converter current equals the output load current, on average, by stepping between whole 10 a increments. The output voltage is nominally 2.5 v, has a hysteresis of 2.5 mV, and is about 1 mV higher at no load than at full load.

Note in particular that with a step decrease in the output current from 95 A to 5 A there is no voltage overshoot. This shows that there are *no theoretical limits to the di/dt in the switched current power converter topology*.

2.3. "Total Charge" Measurement and Control:

The most important function of a control circuit for a power converter is to adjust the input current so that it equals the output current. This is a necessary condition at steady state. Once current equilibrium is achieved, the converter current may then be adjusted somewhat, up or down, to correct errors in the voltage. Conceptually, the output current could be measured and the input current could be adjusted to match. Unfortunately, measuring the output current is not useful as a control input, because it is much too slow.

2.3.1. Capacitor Charge as a Control Function:

Current, in amperes, is the flow of charge, in coulombs per second. The voltage on a capacitor is the charge stored in the capacitor times the capacitance. If the current into a capacitor does not equal the current out of the capacitor, the stored charge changes and so does the capacitor voltage. Thus change in the capacitor voltage is an easy and very fast way to determine that there is an error in the converter current.

2.3.2 Parasitic Impedances:

In a practical power converter for a processor, the power delivery system will have parasitic impedances. Of particular concern are the series parasitic inductances, because they cause the model of figure 2.1.2 above to become unstable. This is because the voltage on the output capacitor does not immediately change in response to an increase in power converter current. The upstream capacitors, on the input side of the parasitic inductance, can store significant excess charge before any of the charge (current) flows to the output capacitor. Thus there is a lag in the output voltage Vo. While the circuit can be compensated, the compensation slows down the circuit response to transients and a much larger output capacitor is needed.

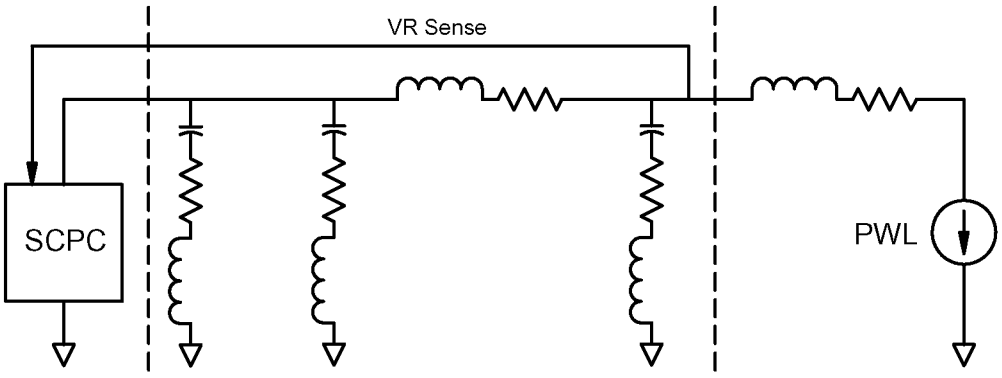

Figure 2.3.1. With series inductance, even a small parasitic inductance, the direct control of the converter current as a function of the output voltage becomes unstable, and large oscillations build very rapidly. While the circuit can be compensated, the compensation slows down the circuit response to transients and a much larger output capacitor is needed.

2.3.3. Total Charge Sensing Improves Stability and Response:

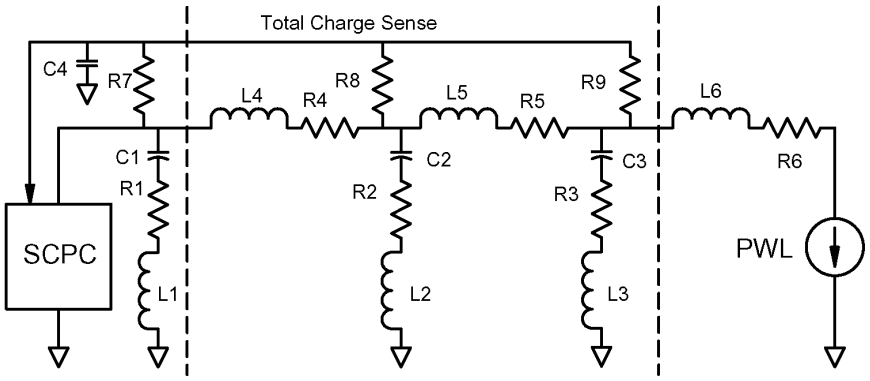

Figure 2.3.2. Total charge sensing can be substituted for output voltage sensing, and it is unconditionally stable. The response need not be compromised by a compensation network, but the series parasitic resistance does affect the output voltage slightly.

This problem of instability due to the parasitic inductances is overcome if the charges on all of the capacitors in the power converter are measured and summed as the control input. Once again the circuit is exquisitely sensitive to differences between the converter current and the output current. The converter current is precisely controlled, the bandwidth is very high (greater than 5 MHz) and it is unconditionally stable.

If the power distribution system is under-damped, there may be oscillations and ringing in the voltages on the various capacitors, but this is invisible to the control system. In SPICE simulations, using representative parasitic impedances from Intel® VR10.2, it seems that a typical layout is fairly well damped, but this is an area of concern for designers.

2.3.4. Total Charge Measurement, ΣQ (Theory):

If the total charge in the various capacitors is measured and summed, a much faster and more accurate control is possible. To introduce the concept of total charge, ΣQ, please see figure 2.3.3, which implements the voltage measurement of figure 2.3.2, except that differential measurement is used to account for ground differences.

The voltages of the various capacitors C1 through C3 are measured, preferably differentially to account for ground differences. The charge on a capacitor is determined as the voltage times its capacitance. If the summing resistors R1 through R3 have their *conductivity* (1/R) proportional to the capacitance, the output of the summing amplifier is proportional to the total charge, ΣQ. Because the total charge ΣQ is linearly proportional to the voltage, it can be used for control purposes as a substitute for the output voltage, with appropriate scaling. $\Sigma Q = \Sigma(V_n C_n)$ (the total capacitance ΣC_n is assumed to be constant).

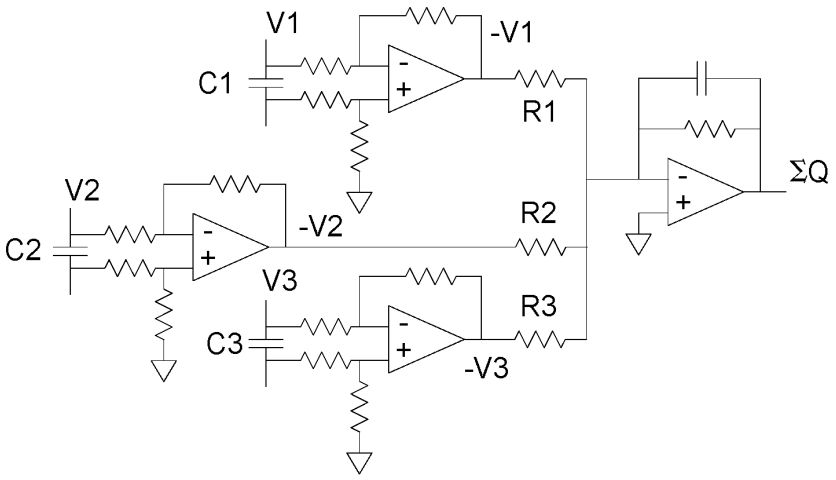

Figure 2.3.3. The voltages of the various capacitors in the system are measured differentially, to remove ground differences. The voltages can then be summed, factored by the value of each capacitor. The sum is the total charge, ΣQ.

2.3.5. Simpler, Faster Charge Measurement:

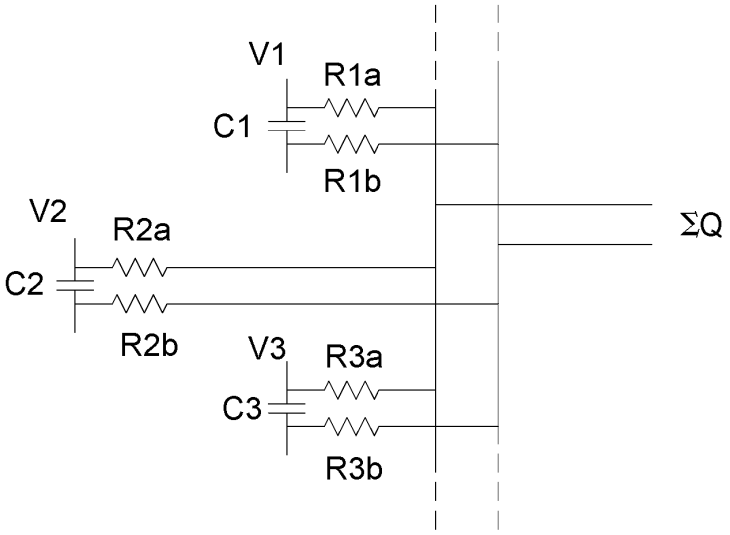

Figure 2.3.4. If the *conductivity* (1/R) of the resistors is proportional to the capacitance of the associated capacitors, the output is proportional to the total charge, ΣQ.

The circuit of figure 2.3.3 is useful for explaining the concept of total charge measurement, but it has the disadvantage that a large number of precision resistors and very fast amplifiers are needed, especially if a large number of points in the circuit are to be measured. Further, the settling time through the amplifiers in series, each with its response and slew rate limitations, results in a significant lag. A much simpler and faster circuit is shown in figure 2.3.4.

It is suggested that two layers of the current distribution bus be devoted to the total charge measurement. Because these two layers do not conduct significant current, their foils may be thin. Measurement points consisting of two resistors can connect the various capacitors throughout the power distribution to the total charge measuring system, as shown, one on the power side and one on the return side at each capacitor (or group of capacitors). Three resistor pairs are shown in figure 2.3.4, for simplicity, but it is contemplated that a large number of measurement points would be used, including points on the processor die. The resistors need not be precise. Likely, screened on resistive ink would suffice.

The relationship of the resistors can be expressed by the following equations:

C1 R1a = C2 R2a = C3 R3a = ----- = Cn Rna
C1 R1b = C2 R2b = C3 R3b = ----- = Cn Rnb

There may be transmission line effects, and it may be advantageous to consider impedance matching in choosing the value of the resistors, particularly the ones that are furthest from the control circuits. Obviously, they cannot all be the ideal value for impedance matching, but the more critical resistors can be optimized, and the others can be selected as ratios of them using the above formulae.

Note that the equations do not require that the high side resistors and the low side resistors be equal, and it may be desirable to use smaller resistors in the return side, for a lower impedance ground return. Such considerations are a tradeoff of a particular application.

2.3.6. Calculating the Output Voltage, Vo:

Control of the total charge ΣQ does not control the output voltage Vo directly, unless the output voltage Vo is taken directly from the capacitor and there is only one capacitor (or there is zero impedance between capacitors). With distributed capacitors having parasitic impedance between them there will be voltage drops in the system, which will affect the output voltage Vo. At steady state conditions, only the resistances are applicable, and the output voltage calculations are straightforward. For transient conditions, it is suggested that a SPICE simulation be used with all identifiable parasitic impedances included.

If the stray inductances are removed, the schematic of figure 2.3.2 reduces to the schematic of figure 2.3.5. This is valid for steady state conditions, and is used to calculate the "error" introduced by using total charge instead of Vo as the control input. (SPICE modeling suggests that the transient voltages due to the parasitic inductances can be rolled off without compromising the control.)

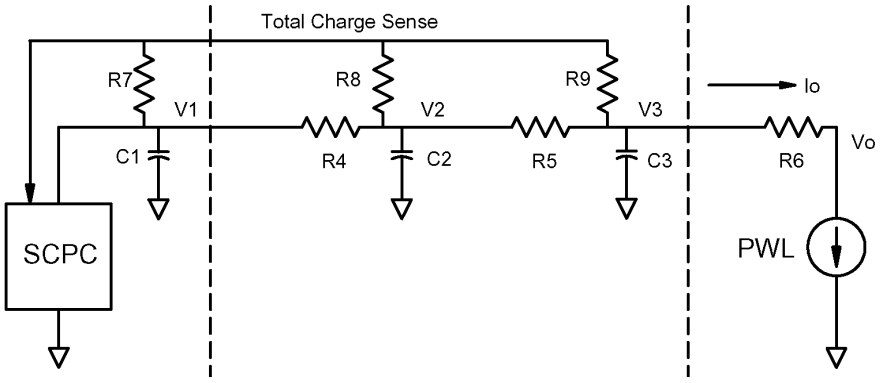

Figure 2.3.5. For steady state conditions, the circuit of figure 2.3.2 reduces to this equivalent circuit. Because R7, R8 and R9 are large, they conduct negligible current, so the average current from the SCPC equals the output current.

To calculate the output voltage Vo at steady state conditions (with reference to figure 2.3.5), let V1, V2 and V3 equal, respectively, the voltages on the capacitors C1, C2 and C3.

By inspection, the charge on the three capacitors is

$$Q1 = C1\ V1$$
$$Q2 = C2\ V2$$
$$Q3 = C3\ V3.$$

Therefore, the total charge ΣQ is given by

$$\Sigma Q = C1\ V1 + C2\ V2 + C3\ V3$$

Also, by inspection, given an output current Io,

$$V1 = Vo + Io\ (R4 + R5 + R6)$$
$$V2 = Vo + Io\ (R5 + R6)$$
$$V3 + Vo + Io\ R6$$

To simplify the final expression, let us define a total capacitance ΣC, a total charge reference ΣQR, and an equivalent resistance RE as follows:

$$\Sigma C = C1 + C2 + C3, \text{ in farads}$$
$$\Sigma QR = \Sigma Q / \Sigma C, \text{ in V}$$
$$RE = (C1/\Sigma C)\ R4 + ((C1 + C2)/\Sigma C)\ R5 + R6, \text{ in ohms}$$

Substituting, collecting terms and rearranging yields the following expression for the output voltage Vo:

$$Vo = \Sigma QR - Io\ RE$$

It can therefore be seen that ΣQR equals the no load output Voltage, Vo0, and RE is a resistive impedance term. It can be seen that if C1 and C2 are small compared to the total capacitance ΣC, then RE is dominated by R6.

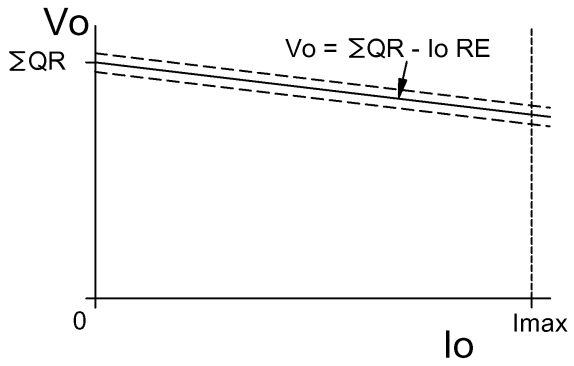

Figure 2.3.6. The output voltage Vo as a function of the output current Io using total charge ΣQ control.

2.3.6. Output Voltage Vo for a SCPC Using Total Charge Control:

When total charge control ΣQ is used with a switched current power converter that uses a flash a-d to directly control the current switches, as in figure 2.1.2, the total circuit impedance is the sum of the impedance Rs due to the resistor ladder network plus the impedance RE due to the total charge ΣQ control.

$$Vo = Vo0 – Io\ (Rs + RE)$$

Whole ladder hysteresis, as explained above in section 2.2, can be used to reduce the converter impedance Rs, and improved layout and components can improve the parasitic impedance RE.

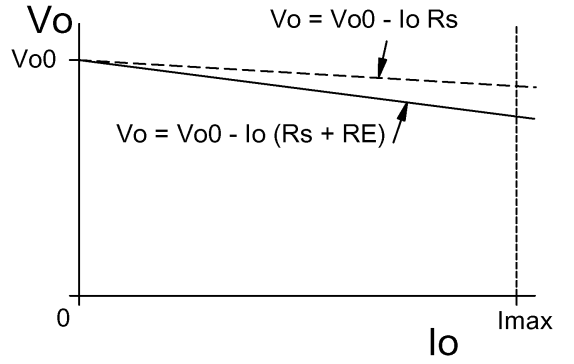

Figure 2.3.7. When total charge ΣQ control is used with a switched current power converter in which a flash a-d directly controls the current switches, Rs and RE combine to define the total circuit impedance.

2.3.7. Overcoming Parasitic Impedance.

A serious impediment to using a plug-in VRM is the problem of the connector's parasitic impedance, particularly the parasitic inductance. The total charge control largely overcomes the effects of this parasitic impedance, making it unimportant as a design tradeoff.

For a low voltage, high current power supply of conventional design, parasitic impedances in the path of current flow is a serious problem. It takes a significant driving voltage to increase the rate of current flow in an inductor, and a higher driving voltage just is not available with present VRMs. This problem is largely overcome with a SCPC using total charge measurement and control.

The equations from 2.3.6, repeated below, show that the impedance R4 is significantly attenuated if C1 is small compared to the total capacitance, ΣC. In the circuit of figure 2.3.5, the impedance R4 simulates an impedance at the connector of a remote VRM. The decoupling capacitors are assumed to be on the motherboard near the processor.

$$RE = (C1/\Sigma C) R4 + ((C1 + C2)/\Sigma C) R5 + R6, \text{ in } \Omega$$

$$Vo = \Sigma QR – Io\ RE$$

An expression in an equation makes much sense with a feel for what is happening in the circuit, empirically. The capacitor C1 may be, as an example, the output capacitor of a VRM that is remote from the MLCC capacitors and the processor. However, it is part of the total charge sensing circuit, and, being directly connected to the VRM, it will be the first to see a response to a drop in the output voltage and the total charge. As an example, let us consider a two capacitor system where the VRM decoupling capacitor C1 is 20 μF, and the output capacitance C3 is 380 μF, for a total capacitance ΣC of 400 μF. If the output voltage decreases 5 mV, the total charge Q is decreased by 2 μC. If this charge is replaced initially entirely on C1, it will result in a voltage rise of 100 mV. This is sufficient voltage to drive the current through the parasitic impedance. Yet the total charge sensing and control prevents the voltage from rising excessively. As the charge is transferred to the output capacitor C3, the voltage on the input capacitor C1 decreases so that the total charge ΣQ remains correct.

The total charge system measures total charge ΣQ, and it is immune to voltage oscillations within the power distribution network. It will maintain the correct total charge ΣQ even if the voltages at the various capacitors are ringing. It is very desirable that the power distribution be a damped system. Because the imaginary part of the complex inductance is lossy, and the losses increase rapidly with frequency, it is likely that most real power distribution systems will be damped at the critical frequency.

2.3.8. Modified Total Charge Sensing

Performance may be improved with the addition of a resistor from the total charge sense to the output voltage Vo, as shown by R3 in figure 2.3.8. (In the SPICE models, a lead-lag network is used here, resulting in a lower output impedance to a higher frequency).

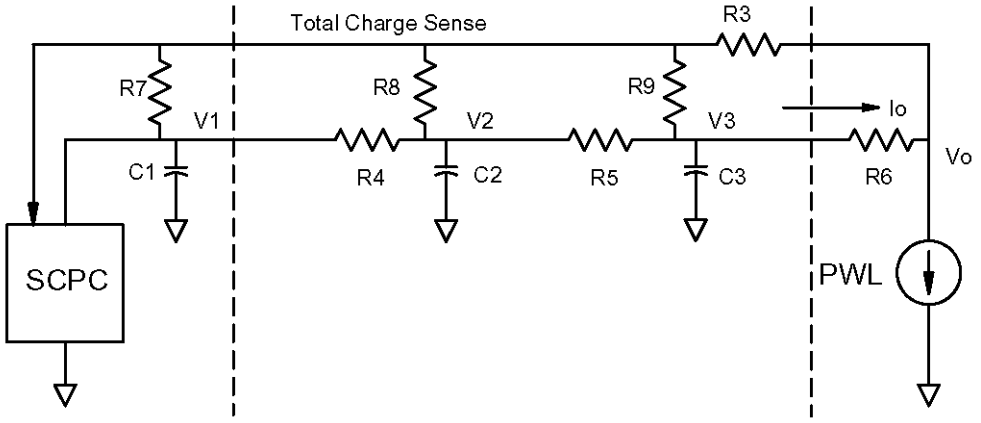

Figure 2.3.8. The total charge sense can be modified with a resistor to the output voltage. In the SPICE models, a lead-lag network was used for R3.

Let us define a resistance Rp equal to the parallel combination of R7, R8, and R9. Using R3, the equation of 2.3.7 (Vo = ΣQR – Io RE) becomes

$$Vo = \Sigma QR - Io * RE\left(\frac{R3}{R3 + Rp}\right)$$

In the SPICE models, R3 is approximately equal to Rp and has a lead-lag around it. The impedance attributable to the total charge sense is reduced by one half and the response at 5 MHz is significantly improved. In the limit, with a very small R3, the circuit reduces to a feedback from Vo, and stability issues will be seen as the benefits of the total charge sensing are attenuated.

ΣQR was shown in 2.3.7 to be the no load voltage, which in many processor power supplies is a digital input VID.

2.4. Using Feedback, for Constant Voltage or Remote Sense

Some applications require a fixed output voltage, either at the power converter output or at a remote sense. For these applications, feedback from the output or the remote sense can be used. Because the lags in the switched current power converter are very small, the bandwidth can be very high.

2.4.1. Feedback Control from Vo or a Remote Sense:

The theory of operation in this section is based upon section 2.1, and the switched current power converter with a flash a-d converter shown in figure 2.1.2. Understanding these circuits is fundamental to this discussion.

If it is desired to use the switched current power converter of figure 2.1.2, but to have a voltage output which is constant, either at the output or at a remote sense point, feedback around the switched current power converter can be used.

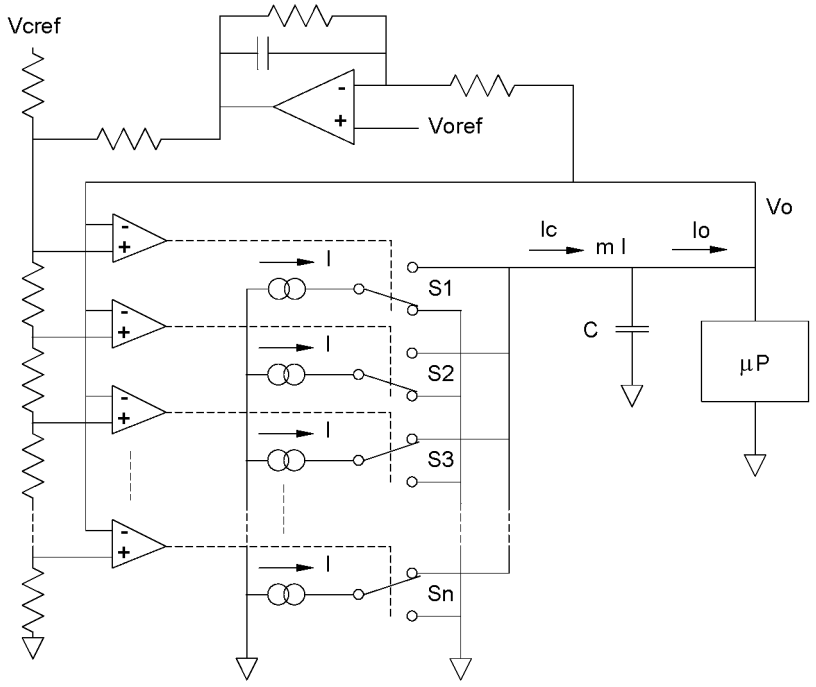

Figure 2.4.1. By adding feedback around the power converter of figure 2.1.2 to change slightly the reference for the ladder network, the switched current power converter can be adapted for a fixed voltage at the output or at a remote sense.

For stable operation, it is necessary to have a negative slope (impedance) in the voltage-current characteristics, at least dynamically. For stable operation with distributed parasitic inductance, the total charge measurement system for control is preferred, and this introduces some additional negative slope (impedance).

To provide a fixed output voltage, a switched current power converter having sufficient negative slope in its voltage-current characteristics for stable operation is designed. Then a feedback loop from the output (or a remote sense) is added to change slightly the reference voltage for the ladder network so that the output voltage settles to the correct state after a transient. See figure 2.4.1.

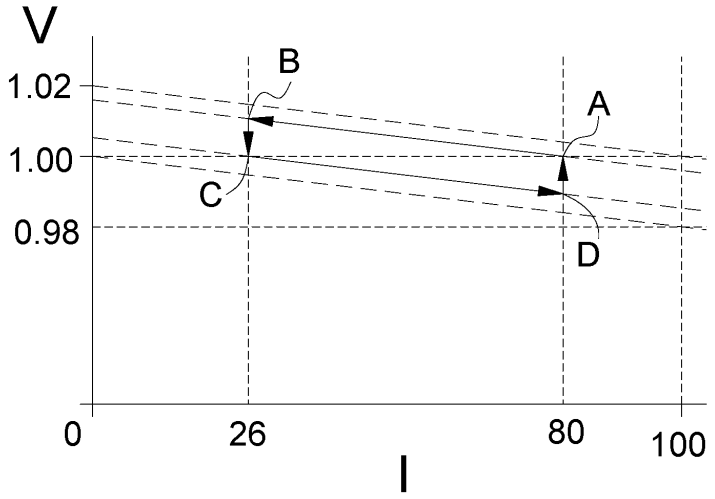

Figure 2.4.2. For a step change in current from 80 A to 26 A, the voltage first rises from 1.00 V at point A to 1.011 V at point B, then settles to 1.00 V at point C. For a step change from 26 A to 80 A, the voltage first falls from 1.00 V at point C to 0.989 V at point D, then settles to 1.00 V at point A**.**

To understand this circuit, consider a step change in current from 80 A to 26 A, as shown in figure 2.4.2. Starting at a point A, the voltage first rises from 1.00 V to 1.011 V at point B (the "natural" voltage of the SCPC without feedback), then settles to 1.00 V at point C as the feedback circuit adjusts the reference voltage. The step from point A to point B has the dynamic response of the fundamental switched current power converter, and is very fast. The settling time from point B to point C is the response of the feedback amplifier and its compensation, which is slower, but still very fast as compared to a conventional power converter. For a step change from 26 A to 80 A, the voltage first falls from 1.00 V at point C to 0.989 V at point D, then settles to 1.00 V at point A.

The feedback must be rolled off sufficiently for stable operation, but the lags to be compensated are in the voltage sense, the logic and the switching speed of the current switches, all very fast compared to the di/dt of an inductor, so the transients are very short.

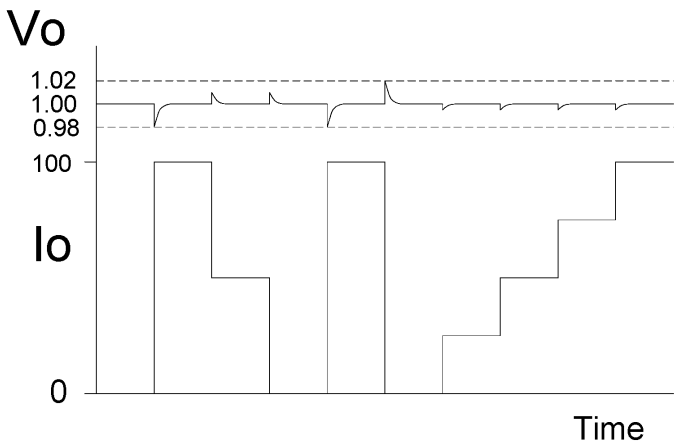

Figure 2.4.3. If a switched current power converter with feedback to the comparator reference follows the characteristics of figure 2.2.2, then the maximum overshoot is 1.02 V, and the maximum undershoot is 0.98 V, for 100 A step changes in the output current.

2.4.2. Using Feedback with Total Charge ΣQ Control:

As was true with the circuit of figure 2.3.1, the inner loop of the circuit of 2.4.1 is unstable if there are multiple capacitors with parasitic inductance between them. The solution is the same, use total charge ΣQ control for the inner loop.

3.0. Varying the Output Voltage:

3.1. Voltage Control by Varying the VID.

With reference to figure 2.1.2, the output voltage Vo of a switched charge power converter can be changed by changing the reference voltage, Vref. The reference may be an analog input signal or it may be a digital VID with a d-a converter. If the VID is changed, the output voltage will follow if the change is not too fast. In the SPICE model used for this presentation, with a 35 A load, a reasonable dv/dt limit is about 90 mV/µs increasing, and about half that decreasing.

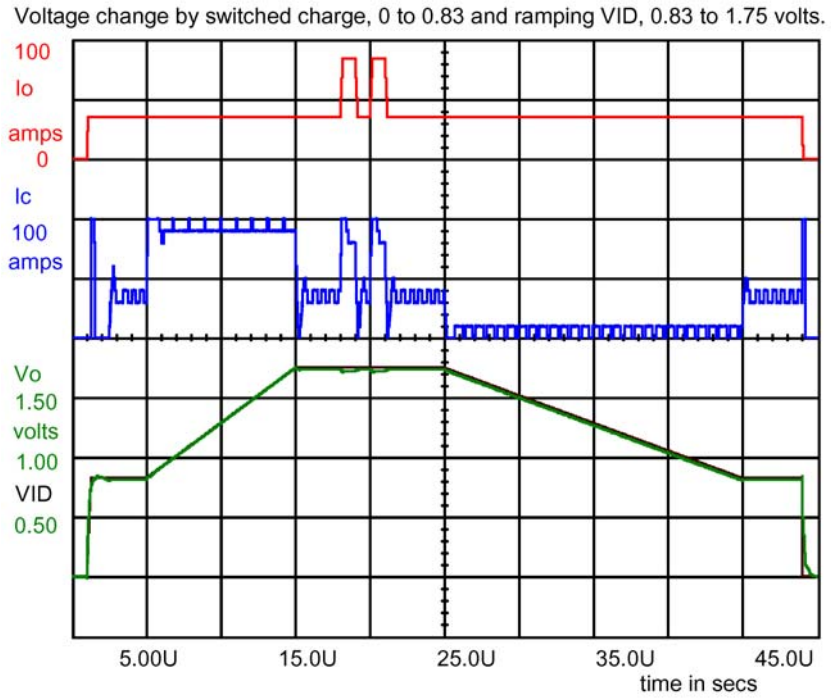

Figure 3.1.1. The output voltage of a SCPC can be varied by varying the voltage reference, or VID. In the graph of a SPICE simulation, the black curve is a commanded VID input, and the green curve shows the resulting output voltage Vo. The red curve is a simulated load current Io, and the blue curve is the SCPC converter output current Ic.

The dv/dt limitation is due to the current available to charge the output capacitors, and the size of the output capacitance. For the example of a SCPC with 100 A rated current, if the load is 35 A, there are 65 A available to charge the capacitor to a higher voltage. Reducing the voltage is slower, as the SCPC does not sink current. Decreasing the voltage relies upon the load absorbing charge from the output capacitor. A smaller output capacitance will allow a faster dv/dt by VID ramping. SPICE simulations and the models used are shown in Switched Current Power Converter; Transient and Frequency Response (pdf).

3.2. Voltage Control with Switched Charge circuits.

The switched charge circuit operates entirely separately from the switched current circuits, and can be added as an accessory with no other hardware change. Only the VID input and the connection to the output capacitors are common. A separate connector or solder pads for a switched charge module can be next to the SCPC module, and the power system can be enhanced by adding the switched charge module with no other hardware changes. There are several choices for implementing a switched charge voltage control. The same interchangeability is true for any of them, so, any "flavor" of switched charge circuit can be used. The user can "plug and play" as desired.

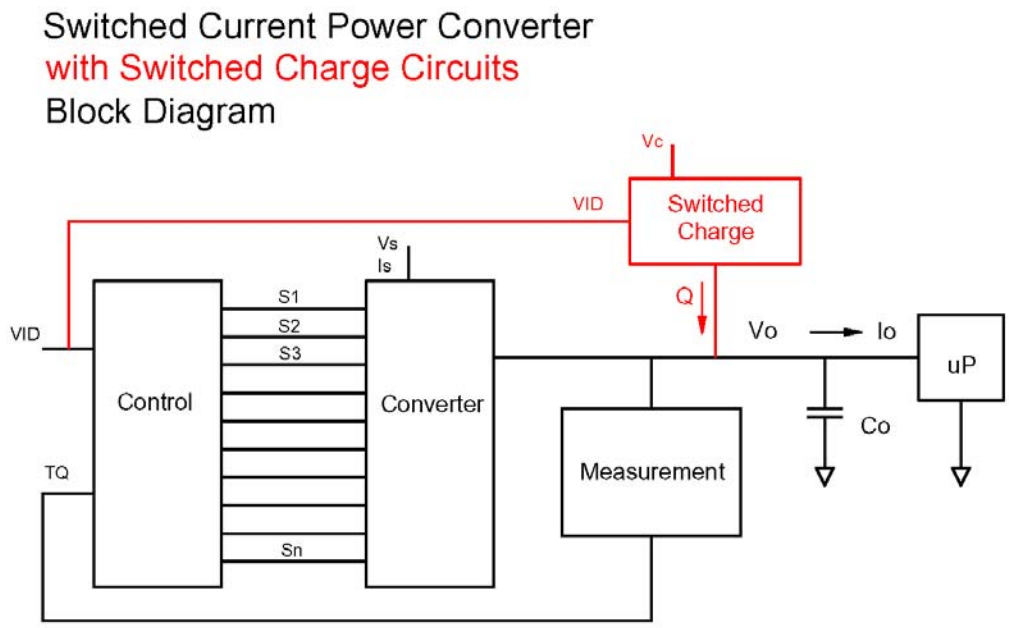

Figure 3.2.1. The switched charge circuits can be added to an SCPC with no change to the other circuits.

3.2.1. Charge Transfer:

The voltage V on a capacitor equals the charge Q times the capacitance C. If the charge is changed, the voltage will change proportionately. If a precise charge is added to, or removed from, a capacitor, its voltage will step by a precise amount, and this can be achieved very quickly with solid state switching.

Figure 3.2.2 shows a basic switched charge circuit, consisting of a totem pole driver and a capacitor Cq. When the driver goes high, as precise charge is transferred to the output capacitor Co, and when it goes low, a precise charge is removed, causing a precise and very fast output voltage step.

The switched charge circuit has no ability to regulate the output voltage, only to step it rapidly. Voltage control is done by the basic SCPC control, by controlling the current. Accordingly, the voltage reference VID must step simultaneously by the same amount so that the ongoing voltage regulation is at the new level.

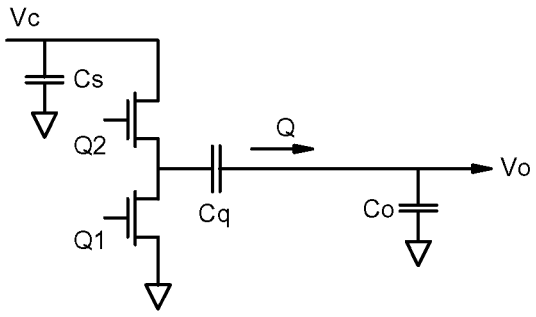

Figure 3.2.2. When the switch changes state, a fixed charge Q is added to or removed from the output capacitor Co, resulting in a precise fixed step in the output voltage, up or down.

A stepped charge circuit may provide one step, for power up, or a few steps, for specific discrete voltage changes, or it may be a binary sequence. With five binary stages, any of 32 discrete voltage steps may be commanded, in any combination, from any VID to any other VID, extremely rapidly and precisely. SPICE models are shown in Switched Current Power Converter; Transient and Frequency Response (pdf).

3.2.2. Accuracy Considerations:

Transferring a precise charge requires precision capacitors, not an attractive option in production manufacturing. Given the large variations in capacitance within the allowable tolerance, if precise steps are needed, it may be necessary to trim the total capacitance. Capacitance also changes with the operating conditions, and this must be accommodated.

Several characteristics of the design are helpful. One is that the absolute accuracy of the capacitors is unimportant if the relative accuracy is good, as it is the capacitor ratio that determines the voltage step. By using the same style capacitor for all of the capacitors and by having them mounted together so that their temperatures are the same can compensate for most environmental variables.

An important consideration is that the accuracy of the output voltage is controlled ultimately by the VID input to the SCPC circuits. If a step change from the switched charge circuit is not precise, the voltage will slew rapidly to the voltage commanded by the VID input to the SCPC. In a power converter, a voltage overshoot may be more serious than undershoot, so the step change circuits may be designed somewhat low, maybe 10 percent. In this scenario, the voltage would step 90 percent of the transition using the switched charge circuit, then

slew the remaining 10 percent by normal SCPC control. This is a little slower, but still an order of magnitude faster than ramping for the entire voltage transition.

If needed, an auto-compensation circuit may be used, and it is fairly simple. Errors in the capacitor ratio may be corrected by changing the driving voltage Vc for the charge transfer capacitor Cq, with reference to figure 3.2.2. A simple calibration scheme uses trial pulses with increasing driving voltage Vc until the step voltage is correct. A calibration circuit is modeled in SPICE, and the simulation results are in Switched Current Power Converter; Transient and Frequency Response (pdf).

3.2.3. Circuit Examples:

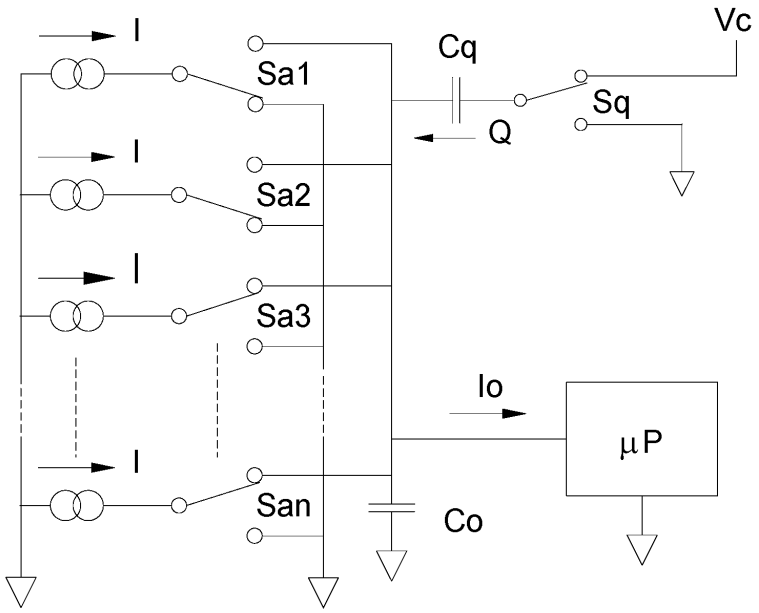

Figure 3.2.3. The addition of a charge switching capacitor Cq and switch Sq enables this power converter to step its output voltage nearly instantaneously and very precisely.

If the supply voltage for a processor is reduced, the losses in the processor are much lower. This power savings can be realized for more of the time when the power converter can step its output voltage very rapidly and precisely.

As an example, a processor might have an idle mode with a supply voltage of 0.80 V, and a turbo mode with a supply voltage of 1.00 V. The switch from idle mode to turbo mode may also coincide with a step increase in the current, and the return to idle mode may coincide with a step reduction in current. The transitions must be very fast and precise, and the voltage must not overshoot or undershoot. See figure 3.2.4.

The addition of a switched charge circuit accomplishes this, as shown in figure 3.2.3. The size of the charge switching capacitor Cq is a function of the driving

voltage Vc, the step in voltage dV desired and the size of the output capacitor Co. As an example, if the driving voltage Vc equals 12 V, the desired step voltage dV is 0.20 V and the output capacitor Co is 250 uF, the value of the charge switching capacitor Cq is given by the following equation:

$$Cq = (dV/Vc)\ Co = (0.20/12)\ 250\ uF = 4.2\ uF.$$

The charge Q that must be transferred can be calculated by the size of the step voltage dV and the size of the output capacitor Co.

$$Qc = dV\ Co = 0.20 \times 250 \times 10^{-6} = 50 \times 10^{-6}\ \text{coulombs}$$

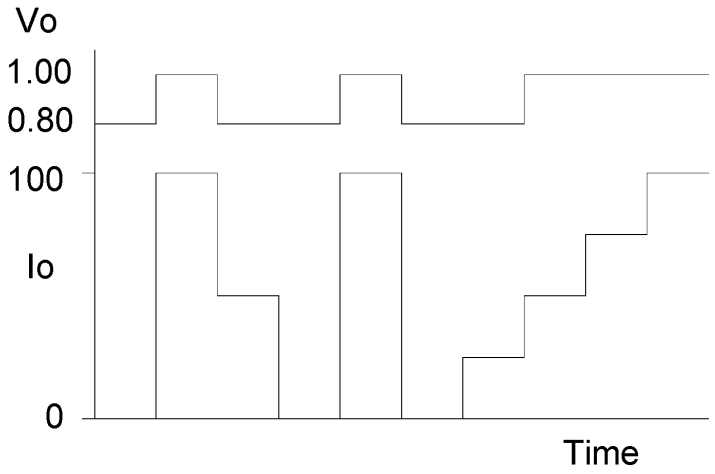

Figure 3.2.4. The output voltage Vo can step very quickly and precisely if a charge is added to or removed from the output capacitor as a pulse, even if the load changes at the same time.

Given that current in amperes equals coulombs per second, it is apparent that the transition time is a function of the charging current, so a low impedance circuit is important for a fast transition. If the charging current averages 100 A, the transition time is 500 ns. This requires an impedance of under 0.1 Ω, which is not difficult to achieve with an inexpensive MOSFET. The MOSFETs must also have a sufficiently high pinch-off current.

If a precise charge is transferred, and the capacitors have reasonable precision, then the step voltage is precise. A good layout is necessary to prevent ringing. Added circuit impedances do not change the final voltage, only the speed of the transition.

3.2.4. Binary Switched Charge Circuits:

The fastest and most flexible voltage control option is the switched current power converter with a binary switched charge circuit. In the example, five switched charge circuits have a binary relationship, that is, each subsequent switched charge circuit has half the capacitance of the previous one, so that charge is transferred to, or removed from, the output capacitor in binary increments. With

five binary switched charge circuits, 32 discrete increments of voltage may be commanded, as with a VID input. In addition, a sixth switch charge circuit may be used for turn on. This stage adds sufficient charge to the output capacitor to establish the minimum operating VID, if used alone.

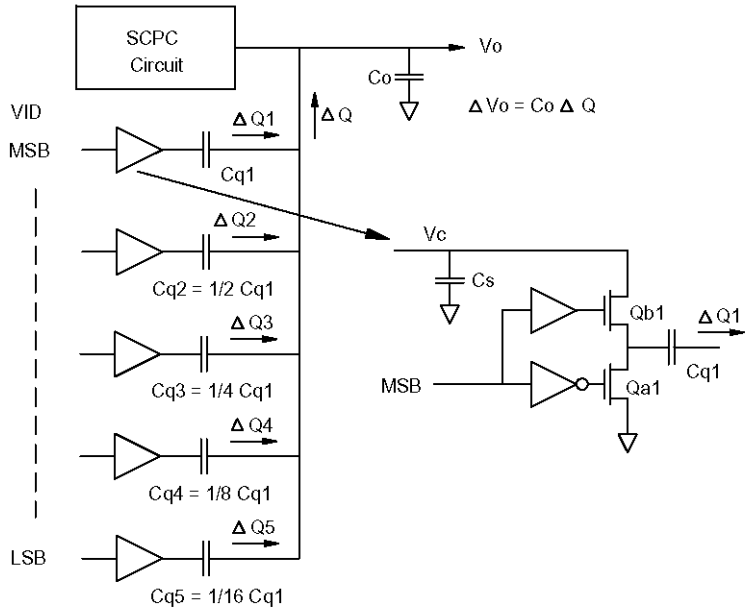

Figure 3.2.3. The binary switched charge circuit has five charge transfer capacitors in a binary sequence. Each successive capacitor is one half the capacitance of the one before it. For the the time constants to be the same, the series resistors (Rdson of the MOSFETs) double in a binary sequence. This suggests that the MOSFETs in an IC have their active areas successively smaller by half.

Because any combination of the switches can be turned on, the output voltage can step to any VID within its operational range upon turn on or at any time thereafter. By changing the VID binary command as a step, the voltage will step to any other VID within its operational range or it may be stepped to zero.

The VID binary steps may also be commanded very rapidly in binary sequence, to achieve a very fast and accurate voltage slew between the first and final binary values.

In the graph following, a number of scenarios are played out. Between 1 and 5 µs, the output voltage steps, following the VID, as the load is pulsed. At 9 µs, the voltage goes to zero simultaneously with a load dump. Between 13 and 14 µs, the voltage slews very rapidly (1,600 mV/ µs) from the minimum operational VID to the maximum operational VID. At 15 µs, it turns off.

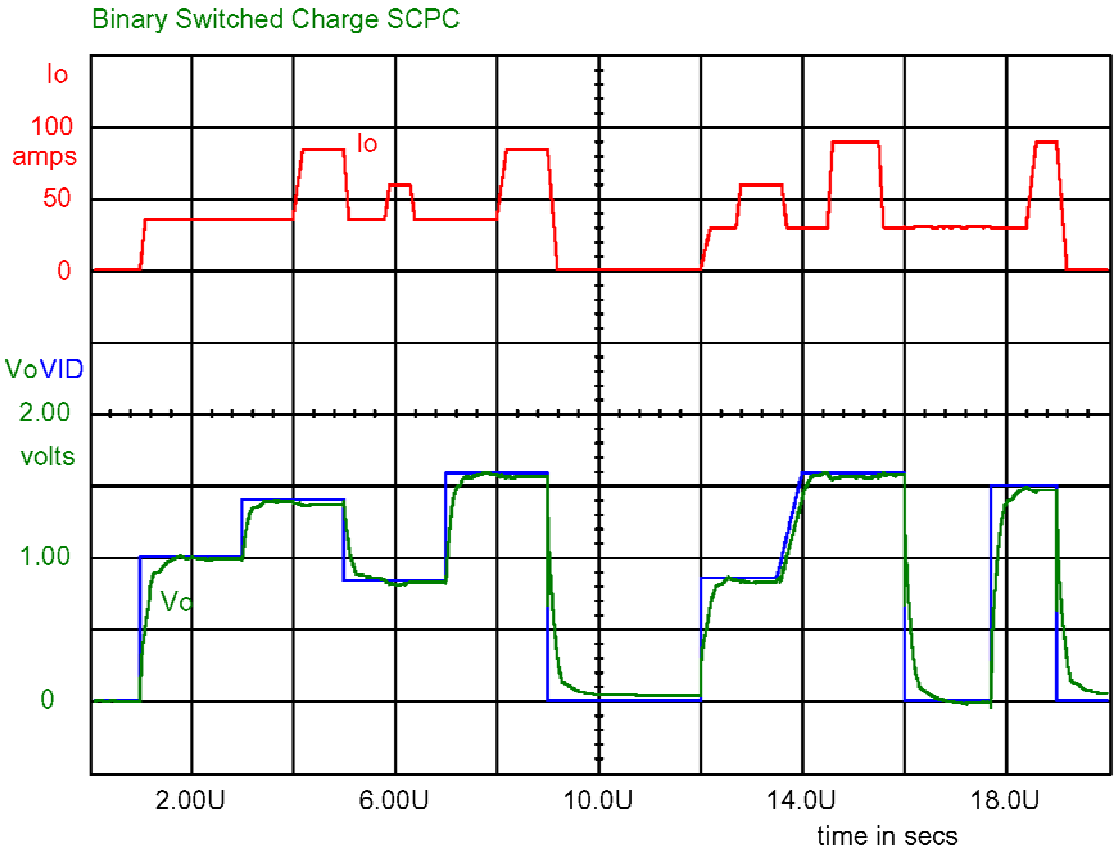

Figure 3.2.6. The graph shows a SPICE simulation of the voltage response of a SCPC with binary switched charge circuits. The blue graph VID represents a digital VID input, and the green graph Vo shows that the output voltage of the SCPC tracks well with no overshoot. The red graph Io shows simulated load changes. Note, at 9 us, that the load can be dumped and the voltage can go to 0 V simultaneously without overshoot.

Between 17 and 19 µs, the circuit turns on, a pulse of maximum current is applied and removed, then it turns off, all within about 1 ½ µs. This last scenario may be important for power reduction for a processor that has to respond rapidly and intensively when an interrupt occurs but is idle most of the time. Its net power consumption is very low, for cool operation or long battery life.

3.3. Energy Loss in Transferring Charge.

When charge is transferred, some of the energy is lost. Over a number of cycles, this could represent a significant power loss. A processor uses much less power at a lower voltage, and none at all if the voltage is zero. If a processor can remain in a lower power state, or off, for more of the time by using switched charge voltage transitions, the trade-off is favorable. Following, in 3.2.1, there is an analysis of the power loss, with examples. If the power is cycled on and off fewer than 1,000 times per second, and the voltage level is stepped fewer than 10,000 times per second, the net power savings may be significant.

This suggests the possibility that the processor (or part of it) can be off most of the time. Much of the use of processors involves manual data entry from a keyboard, with most of the time being wasted, waiting for the next keystroke. The processor can turn off between keystrokes, extending battery life significantly.

Many systems are maintained in a high state of readiness because a fast response is needed. If a μs or two of delay is tolerable, the switched charge circuit allows the processor to be turned off when idle. An interrupt controller can signal the SCPC with switched charge to wake up, and full voltage at full operating current is available in about a microsecond.

3.3.1. A Step Voltage from 0.9 V to 1.0 V:

Energy is lost for both transitions, that is, when reducing the charge to lower output voltage Vo and when increasing charge if the increase the output voltage.

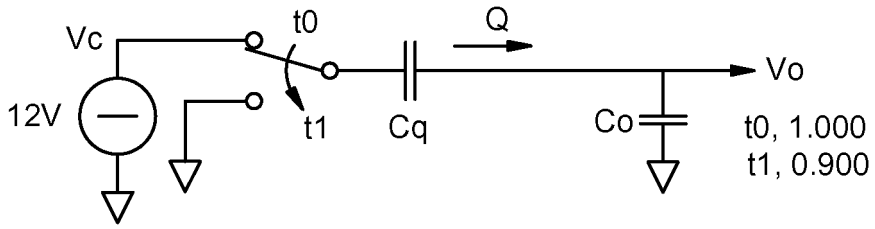

Figure 3.3.1. At time = t0, the output voltage Vo is 1.000 V. At time = t1, the switch changes state and removes charge from the output capacitor Co to step the output voltage Vo to 0.900 V.

Using the schematic in figures 3.3.1, let Vc = 12 v, Co = 250 uF and Cq = 2.1 uF.

In figure 3.3.1, the voltage Vo is 1.000 V at t0. The switch is thrown at t1 to remove some charge and reduce the output voltage to 0.900 V.

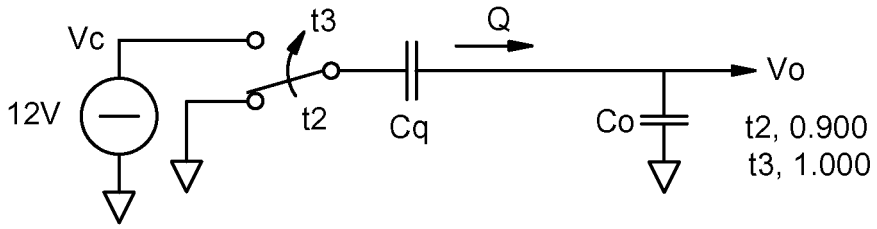

Figure 3.3.2. At time = t2, the output voltage Vo is 0.900 V. At time = t3, the switch changes state and adds charge to the output capacitor Co to step the output voltage Vo to 1.000.

The energy E on a capacitor is given by $E = \frac{1}{2} C V^2$.

At t0, the energy is as follows:

$Ecq = \frac{1}{2} (2.1 \times 10^{-6}) (11^2) = 127$ uJ

$Eco = \frac{1}{2} (250 \times 10^{-6}) (1.000^2) = 125$ uJ

Etotal = 252 uJ

At t1,

$Ecq = \frac{1}{2} (2.1 \times 10^{-6}) (-1^2) = 1.1$ uJ

$Eco = \frac{1}{2} (250 \times 10^{-6}) (0.900^2) = 101$ uJ

Etotal = 102 uJ

The energy Eco in the output capacitor has been reduced by 24 uJ, and the total energy is reduced by 150 uJ.

The difference is the energy lost, 126 uJ.

In figure 3.3.2, the voltage Vo is 0.90 V at t2. The switch is thrown at t3 to add some charge and increase the output voltage to 1.000 V. At t2, the energy is as follows:

$Ecq = \frac{1}{2} (2.1 \times 10^{-6}) (-1^2) = 1.1$ uJ

$Eco = \frac{1}{2} (250 \times 10^{-6}) (0.900^2) = 101$ uJ

Etotal = 102 uJ

At t3,

$Ecq = \frac{1}{2} (2.1 \times 10^{-6}) (11^2) = 127$ uJ

$Eco = \frac{1}{2} (250 \times 10^{-6}) (1.000^2) = 125$ uJ

Etotal = 252 uJ

The energy Eco in the output capacitor has been increased by 24 uJ, and the total energy is increased by 150 uJ.

The difference is the energy lost, 126 uJ.

Therefore, the energy lost per cycle is 252 uJ. This is about 2.5 W for ten thousand cycles per second. Whether this is a problem or not depends upon the frequency of the step voltage changes. With an active processor, this may be a problem, so caution is advised.

The energy lost is reduced significantly by using a lower voltage as the charge supply, as the largest contributor to the loss equations is the square of the voltage difference between charge supply and the output voltage. A larger capacitor is required, however, so the net gain is approximately as the inverse of the voltage difference. As an example, if the charging voltage is reduced to 3.3 V, the losses are reduced by 75 percent. As a tradeoff, the Rds of the MOSFET and the stray inductance has to be much lower to achieve the same step time (dv/dt). If a lower voltage source is already available, using it is probably worthwhile. Whether it would be worthwhile to add a voltage source would be a tradeoff of the particular application.

3.3.2. A Step Voltage from 0 to 0.900 Volts:

Consider an example in which a large part of the processor is turned off when it is not needed. To be able to do this, the power converter must be able to provide the normal operating voltage with full load current capability very rapidly and accurately. The switched current power converter with switched charge circuits is able to do this in a few microseconds.

The circuit of figures 3.3.3 and 3.3.4 shows the charge transfer. The voltage can be brought up to 0.900 V in a few microseconds. If it is desired to bring the voltage up to 1.000 V, this circuit can be used in combination with the circuit of figures 3.3.1 and 3.3.2. Care must be taken to reduce the parasitic inductance to an absolute minimum, so that the circuit does not ring.

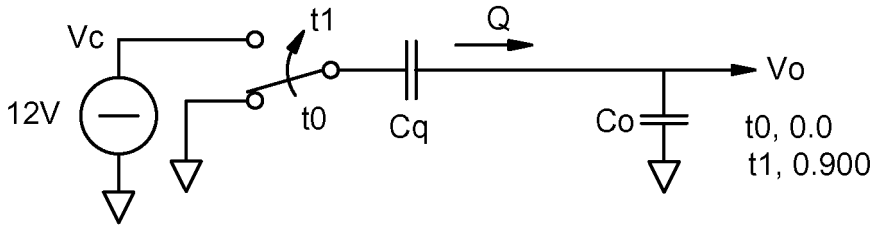

Figure 3.3.3 At time = t0, the output voltage Vo is 0.0 V. At time = t1, the switch changes state and adds charge the output capacitor Co to step the output voltage Vo to 0.900 V.

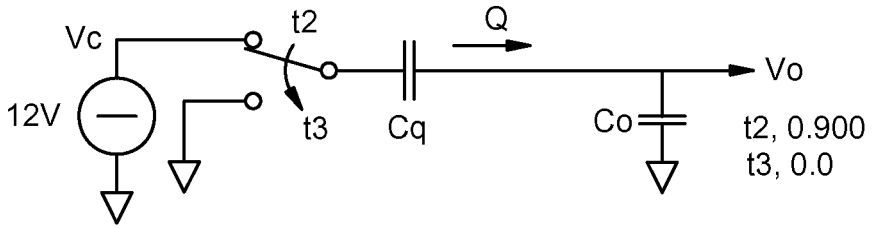

Figure 3.3.4 At time = t2, the output voltage Vo is 0.900 V. At time = t3, the switch changes state and removes charge from the output capacitor Co to step the output voltage Vo to 0.0.

Energy is lost for both transitions, that is, when applying the charge to turn on the output voltage Vo and when removing the charge.

Let Vc = 12 v, Co = 250 µF and Cq = 20.3 µF.

In figure 3.5.3, the voltage Vo is 0.0 V at t0. The switch is thrown at t1 to add charge and increase the output voltage to 0.900 V.

The energy E on a capacitor is given by $E = ½ C V^2$.

At t0 and t3, the stored energy is zero.

At t1 and t2,

$Ecq = ½ (20.3 \times 10^{-6}) (11.1^2) = 1251$ µJ

$Eco = ½ (250 \times 10^{-6}) (0.900^2) = 101$ µJ

Etotal = 1,352 µJ

This energy is required to bring the output voltage up to 0.900 V, and the energy wasted is the difference, 1,251 µJ. On discharge, all of the energy is lost, 1,352 µJ. The total per cycle is 2,603 µJ, or approximately 2.6 W for one thousand cycles per second.

This suggests that the processor could be turned on and off one thousand times per second with very little lost energy, far less than would be saved. For manual data entry, it seems that it would be practical to turn off the bulk of the processing functions between key strokes, as an example.

4.0. Constant Current Sources; Modular Design:

The n constant current sources used in the switched current power converters may be inductors in which the inductor current I is controlled at a constant value (though there is some ripple). A multi-phase buck converter is a suitable circuit. The preferred circuit uses a matrix transformer with a constant current source as its primary input power

In theory, there is zero loss in a short circuit, but in practice, there are losses associated with the circulating current. In modern MOSFETs, the resistance of the package is greater than the on resistance of the MOSFETs. The success of the design relies on optimizing the packaging design. It is for this reason that a modular design is recommended, with the switching MOSFETs integrated into the module and located very close to the inductor terminals.

4.1. Simple Buck Converter as a Constant Current Source:

The current source shown in figure 4.1.1 is based upon a multiphase buck converter design. However, instead of the output current In going to an output capacitor and the load, it is switched either by S3 to the output or by S4 to the return. The Inductor current is refreshed periodically by turning on S1. The rest of the time, S2 is on.

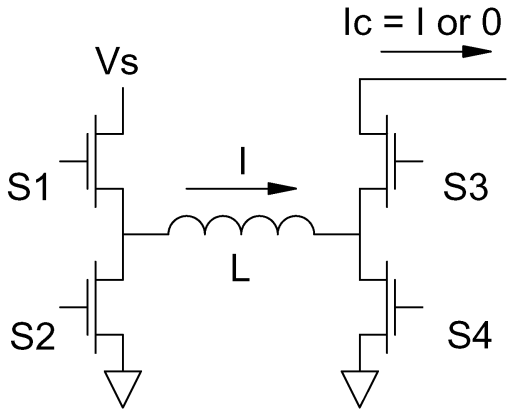

Figure 4.1.1. The current I in the inductor L is controlled with switches S1 and S2 to be a constant value I (ignoring ripple). If S3 is on, the inductor current I is directed to the load, and Ic = I. If S4 is on, the current is switched to the return, and Ic = 0.

4.2. Modified Buck Converter as a Constant Current Source:

A modified buck converter, shown in figure 4.1.2, has only one switch, S1 or S2, in series with the output during the on time. During the off time, the current circulates in a tight loop around the inductor. The output current cannot be zero

when the current in the inductor is refreshed, but the minimum output current needed is very low, so this is not a problem for most applications.

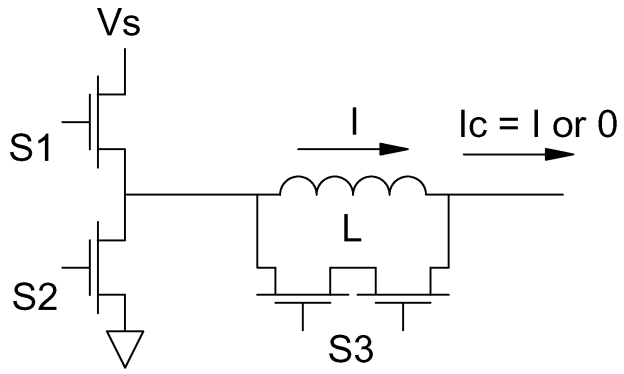

Figure 4.2.1. In this constant current source, the current Ic goes to the load if either switch S1 or S2 is on, so only one switch is in the circuit for the on state. During the off time, the switch S3 conducts the current I in a tight loop around the inductor, and Ic = 0.

4.3. Transformer Coupled Constant Current Sources:

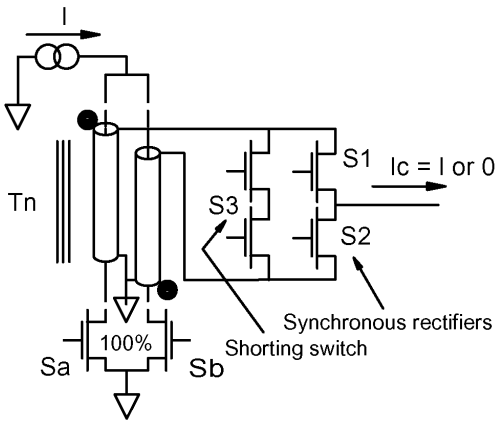

Figure 4.3.1. A matrix transformer provides the parallel constant current sources. The primary is driven by a constant current source, and push-pull switches Sa and Sb conduct alternately at 100% duty cycle. If the synchronous rectifiers S1 and S2 are switching, the secondary current goes to the output In. If S3 is conducting, a current ½I circulates in a tight loop within the transformer module.

In a transformer, the net ampere-turns is zero, neglecting magnetizing currents. Therefore, if the primary of a transformer has a constant current, the secondary coupled with it has a constant current as well, determined by the turns ratio of the transformer. 100 percent duty cycle switching of the primary push-pull MOSFETs is assumed.

The matrix transformer has a number of single turn modules with a common primary winding, so each of the modules has a constant current output equal to the primary current. For simplicity, and for optimum coupling, coaxial modules are preferred, and that is what is shown in the schematic of figure 4.3.1. The

single turn coaxial module is just two formed metal inserts in a hole through a block of ferrite. The ends of the metal inserts are returned at the ends as surface mount terminals, and the primary winding is a single wire through a row of modules, a very inexpensive construction.

In figure 4.3.1, it can be seen that when S3 is closed, it shorts both turns of the secondary winding. Therefore, the secondary current during the off time is reduced to one half, and the I^2 component of the power loss is reduced to one fourth.

The transformer coupled switched power converter provides isolation, and can operate directly from a 48 V input by using a four turn primary winding.

4.4. Inductor Modules:

There are a number of good reasons to modularize the switched current power converter. Probably the most important is to keep the circuit very tight, for minimum parasitic inductance in its interconnections, both within the module and on the printed wiring board. Another is the economy of producing a large number of identical, very simple sub-assemblies.

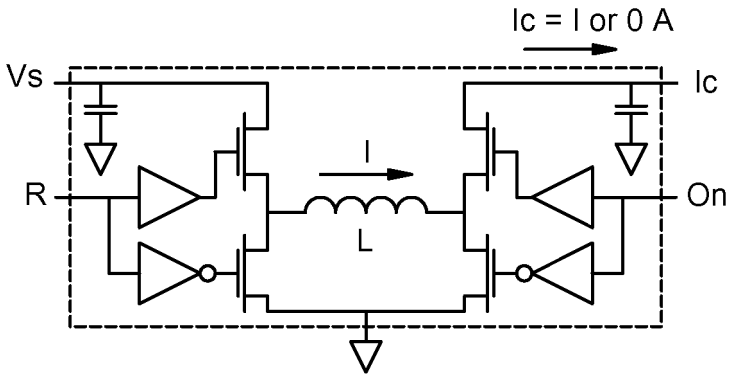

Figure 4.4.1. The buck converter switched current source module may contain the inductor, the power MOSFETs and their drivers, an input decoupling capacitor and a module output capacitor. There are three power connections and two signal connections.

For any switched current power converter, there is the question of how many constant current sources to use. There are some advantages to using more current sources, as the current modulation is between smaller step increments of current, and the current in each module is reduced. If the current is refreshed in a multiphase sequence, the number of phases should be less than Vs/Vo, the input voltage divided by the output voltage, so that there is no overlap of the phases. As an example, with a 1 V output and a 12 V source, twelve is the maximum number of phases needed to avoid overlap. Ten phases is suggested for a practical power converter. The current can be sensed during the refresh cycle with a single current sense in the high side.

The basic buck converter current source of figure 4.1.1, when made as a module, may contain the inductor, the power MOSFETs and their drivers, a voltage source decoupling capacitor and a module output capacitor. See figure 4.4.1. The alternate current source of figure 4.2.1 has the same logical connections and similar components. See figure 4.4.2.

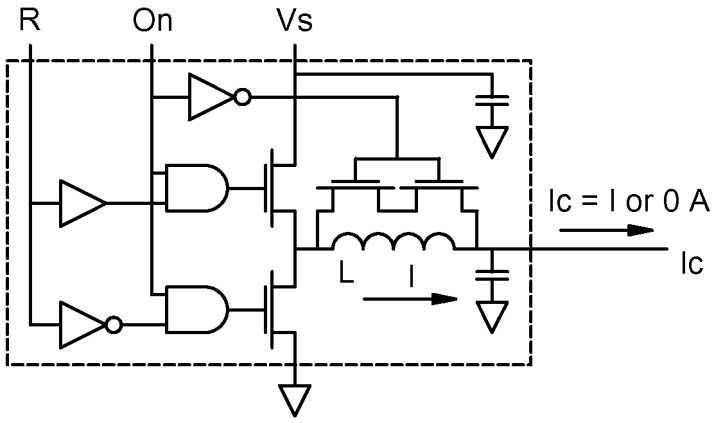

Figure 4.4.2. If the On input is high, one of the MOSFETs on the left is on, and current flows to the output Ic. If the refresh input R is on, the MOSFET that is connected to the voltage source Vs is on, to recharge the inductor current. It is necessary for the On input to be high during this time, so the minimum output current cannot go to zero. If the On input is low, the current circulates in a tight loop around the inductor, and the output current Ic is zero.

The physical implementation of a representative current source module is shown in figure 4.4.3. The inductor sits above a small circuit board that contains the logic, the switching MOSFETs and the capacitors. Surface mount terminals provide a connection to a circuit board for the power, the ground, the output and the logic.

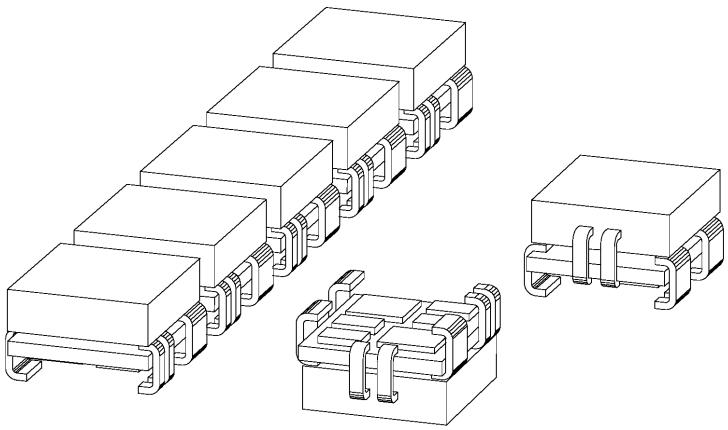

Figure 4.4.3. A representative current source module has a small circuit board with the components mounted on it. An inductor sits on top of the board, and the surface mount terminals extend below it. In a typical switched current power converter, ten modules may be used, perhaps five on each side of the printed wiring board.

4.5. Transformer Coupled Modules

The transformer coupled current source provides isolation, and can be powered from a 48 V power source by using a four turn primary winding.

A matrix transformer is a transformer using a number of cores in a row with a common primary winding passing through all of the cores. The cores usually have two single-turn secondary windings already installed in a core. The secondary windings may be connected externally as a push-pull secondary winding. In figure 4.5.2, it can be seen that the secondary windings are very simple, made of formed metal with folded self-leads.

A switched current power converter might use ten transformer modules, five in a row on each side of the printed circuit board. The primary winding may be installed after assembly to the circuit board, and are simply a pair of "U" shaped wires that are slipped through the aligned through holes and terminated on the board.

In operation, an ON input either enables a CLOCK input to toggle the synchronous rectifiers if high or holds them both off if low. If the ON input is low, a shorting switch (back to back MOSFETS) shorts the secondary windings, diverting the current from the output of the module and reflecting a short circuit to the primary. Note that the shorting switch shorts the ends of the push-pull windings, so it is shorting two turns. That means the current is one half, so the losses due to the circulating current are reduced. With the core shorted, the flux is suppressed, so the core losses are reduced as well.

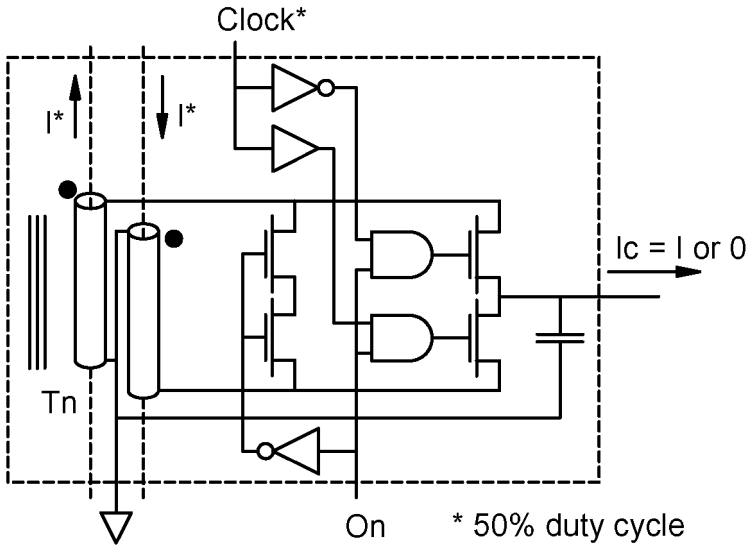

Figure 4.5.1. A module for a transformer coupled current source has the tubular secondary windings in a transformer core, MOSFETs (synchronous rectifiers and a shorting switch) and their drivers, and an output capacitor. The primary winding is added later.

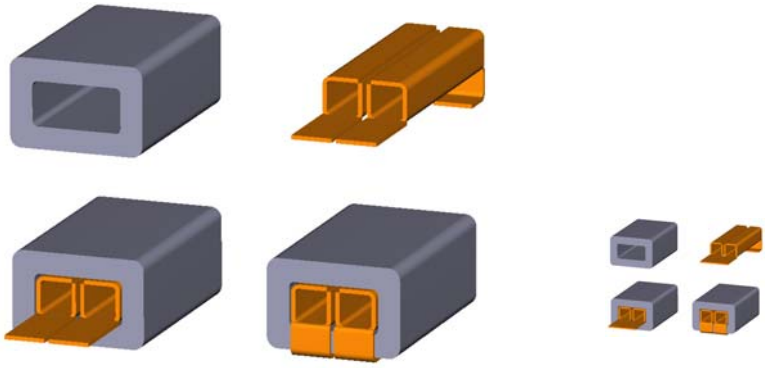

Figure 4.5.2. The transformer is a simple pressed core with two stamped metal inserts. The ends of the inserts are bent back, as self-leads. The core is 3 x 6 x 6 mm.

The power MOSFETs of the secondary circuit can be integrated into a "Simple Power IC", see http://eherbert.com/simpleic.pdf. For a low power SCPC, all of the secondary channels can be in one IC, but for a VRM or VRD (100 A or so), the preferred partition is five dual channels of 20 A each. The dual channel is ideal for mounting on a transformer module with two cores, as it optimizes the secondary connections. The module in figure 4.5.3 is shown with one of the transformer cores lifted so that the power IC can be seen on the little daughterboard.

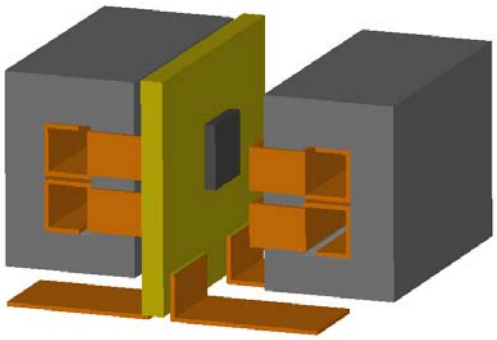

Figure 4.5.3. For a VRD, a small module with two transformer cores is built on a vertical daughterboard with surface mount feet. One of the transformer cores is shown lifted, to reveal the simple power IC, as dual channel circuit containing all of the secondary power MOSFETs for two transformer sections.

Several of the modules are placed on the motherboard with their through holes aligned, and the primary winding is added. A single master IC has the control and logic, as well has the high side drivers. Alternatively, the entire transformer can be mounted and pre-wired on a larger daughterboard, as in figure 4.5.4.

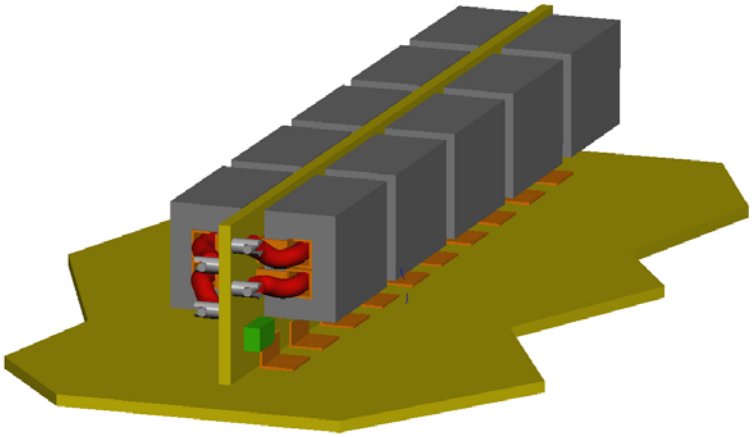

Figure 4.5.4. The switched current power converter can be mounted on a small vertical daughterboard, to minimize board real estate. Components are mounted on both sides of the daughterboard, under the transformer cores. Small decoupling capacitors are on the module, but most of the output capacitance is MLCCs mounted near the processor. The bulk capacitors are not used.

For a lower profile, the transformers can be mounted on one side of a surface mount daughterboard. Although not optimum for the secondary winding connections, the connections are still tight and reasonably low inductance.

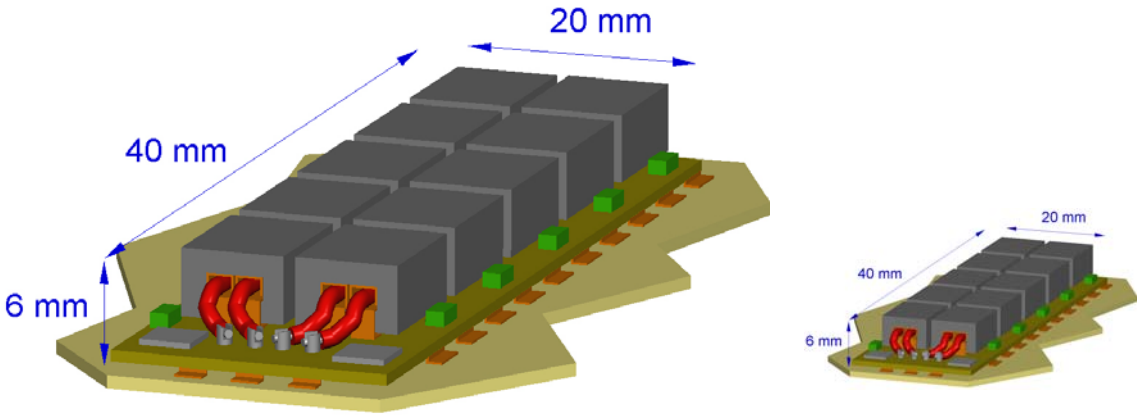

Figure 4.5.5. For a lower profile, the transformer cores can mount flat on one side of the daughterboard. The power ICs are under the transformer cores, and other circuits are on the bottom side. Small decoupling capacitors can be seen on the edge of the module, but most of the output capacitance is MLCCs mounted near the processor. The bulk capacitors are not used.

5.0. Capacitor Size, Propagation Delay, Hysteresis and Switching Frequency:

Because of its very fast response, the SCPS needs much less capacitance than a buck converter of comparable rating. The bulk capacitors are not used, and the total capacitance of the MLCCs needed is lower.

Considerations of parasitic impedance (ESR and ESL) may require using many MLCCs in parallel, but each is smaller if the total capacitance is lower. Cost savings are significant, because lower-value capacitors are much less expensive. The lower value parts are smaller, saving board real estate as well.

5.1. Charging Currents and dV/dt:

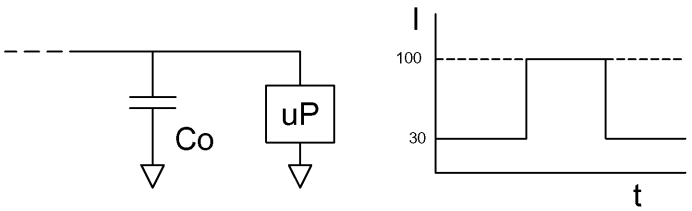

Figure 5.1.1. The first step in determining the value of the output capacitor Co is to determine the maximum current steps, up and down.

It is important to know the dV/dt when a load change occurs. The maximum dV/dt occurs when there is a maximum load change. For our discussion and examples, let us assume that a processor has a maximum current of 100 A, and a leakage current of 30 A. Therefore, the maximum current step is 70 A, positive or negative.

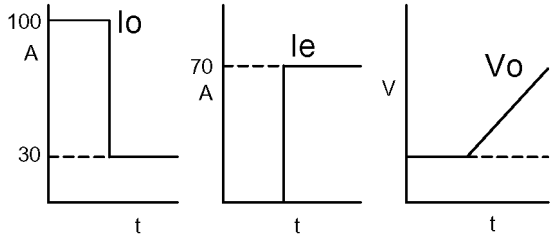

Figure 5.1.2. The instantaneous difference between the converter output current Ic and the load current Io is the "error current," Ie. The error current Ie causes the output voltage Vo to rise.

The "error current" Ie is the difference between the converter current Ic and the output load current Io. The dV/dt is determined by the error current Ie and the size of the output capacitor Co, as:

$$dV/dt = Ie/Co$$

There is a dV/dt as well with a constant load current, as the converter output current modulates between steps so that the average converter output current matches the load current. Figure 5.1.2 shows the converter current Ic modulating between 30 A and 40 A to produce an average output load current Io of 35A.

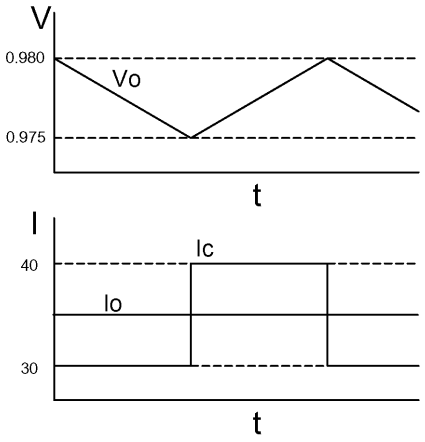

Figure 5.1.2. The output current of the SCPC is in discrete steps. In the graphs, the correct output current 35 A is the average resulting from modulating between 30 A and 40 A. The "error current" is – 5 A or + 5 A, resulting in dV/dt. The comparator hysteresis determines the threshold voltages and defines the ripple voltage.

5.2. Dv/dt and Propagation Delay:

Propagation delay is the time between ideal switch actuation (open or close) and actual switch actuation. The propagation is measured from a triggering event, which, for the SCPC, is when a voltage crosses the threshold voltage of a comparator until the switch has actuated and the current has changed accordingly. The total propagation delay has a number of contributors:

1. Any lag through the voltage sensing circuit or its compensation.
2. The propagation delay through the comparator.
3. The delay through the logic.
4. The delay through the MOSFET gate drivers.
5. The time that it takes to charge or discharge the gate capacitance of the MOSFET and turn it on or off.

During the propagation delay, the current error persists, and the dV/dt continues until the switch actually changes state. This results in an overshoot, as can be seen in figure 5.2.1. To maintain the ripple voltage within specification, the comparator hysteresis may be reduced.

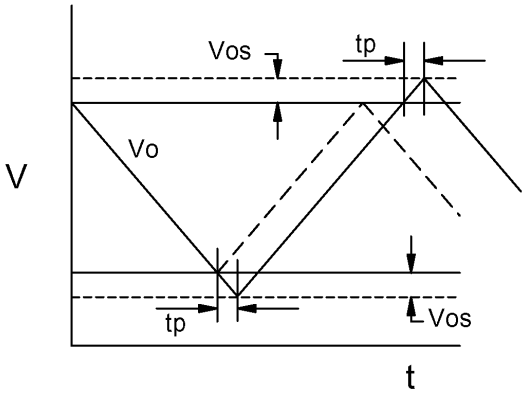

Figure 5.2.1. The propagation delay tp causes an overshoot voltage Vos. This increases the ripple voltage and decreases the hysteretic switching frequency. To compensate, the hysteresis is reduced.

The voltage overshoot Vos is symmetrical if the load current is half way between current steps, that is, X +5 A, as seen in figure 5.2.1. For other values, the offset voltage Vo is biased, as seen in figure 5.2.2. For a given propagation time tp, the voltage Vos continues further if the dV/dt is higher.

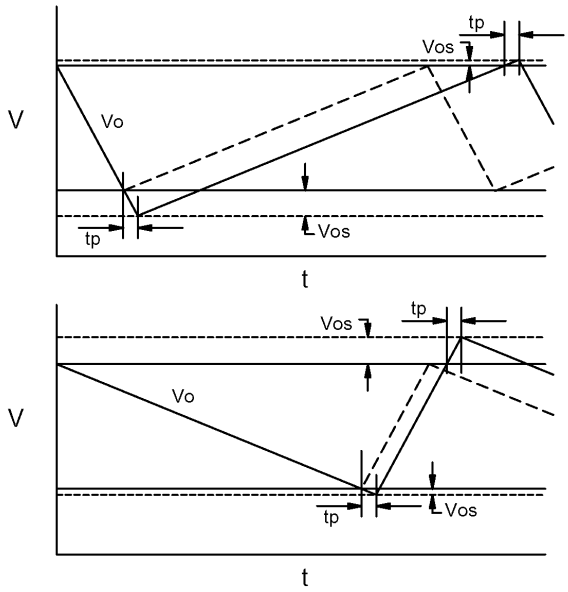

Figure 5.2.2. If he load current is higher than the mid-point of the current step (top graph), the discharging error current is larger and the negative overshoot voltage Vos is also larger. However, the charging current is lower, so the positive overshoot voltage Vos is lower. This biases the output voltage slightly lower, which is good – the voltage is supposed to be lower at higher current. The second graph shows the graph for a load current that is below the mid-point of the current step.

In the top graph of figure 5.2.2, the load current is 8 A above the current step. When the step current is off, the error current is – 8 A, and the discharge is faster, resulting in a large negative overshoot Vos during the propagation delay

tp. When the step current turns on again, the error current is + 2 A. The dV/dt is lower and the positive overshoot voltage Vos is lower. This has the effect of biasing the output voltage somewhat lower if the current is higher. This voltage bias actually is beneficial, as the voltage is supposed to be lower at higher current.

The conditions are reversed for a load current that is below the mid-point of the current step.

5.3. Maximum Load Step and Output Impedance:

It is usual for a power converter for a processor to have a specified output impedance Rs, for example, in the order of 1.2 mΩ (Intel ® VR 10.2). The impedance is often expressed graphically, as in figure 5.3.1. This high impedance has no utility for the processor, it is an accommodation to the response time of the power converter. An output impedance Rs = 0.5 mΩ is a reasonable value for a SCPC.

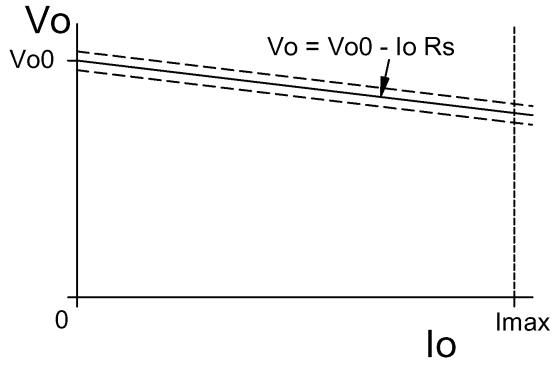

Figure 5.3.1. Usually, a power converter for a processor has a specified output impedance, Rs. 1.2 mΩ is typical for a multi-phase buck converter (Intel ® VR 10.2). A reasonable value for the SCPC is 0.5 mΩ.

It takes a finite time for the output capacitor Co to change voltage, even with a maximum current step. If there is significant propagation delay or other lags, the voltage may overshoot or undershoot. A larger output capacitor Co slows the dV/dt and prevents overshoot.

The buck converter, even a multi-phase buck converter, had a very slow response, so it requires a very high capacitance even with a high output impedance. A typical output impedance (from Intel ® VR 10.2), is 1.2 mΩ, and a typical output capacitance is 7,100 to 9,500 uF.

The SPCP, with its faster response, may have an output impedance of 0.5 mΩ, and requires only 250 uF of output capacitance. This assumes a propagation delay of 50 ns.

5.3.1 Example: SCPC with a Current Step, 100 A to 30 A:

Let's assume a SCPC with the following characteristics and calculated values:

Load, Io = 100 A. At t0, the current steps to 30 A
Channels = 10
Step current, Is = 10 A
Output capacitor, Co = 250 uF
Output impedance, 0.5 mΩ
Propagation delay: tp = 50 ns

For a maximum step current dIo of -70 A, given the output impedance of 0.5 mΩ, the output voltage Vo changes by -35 mV in response to the current change.

$$dVo = dIo\ Rs$$

$$dVo = -70\ A \times 0.5\ m\Omega$$

$$dVo = -35\ mV$$

The timing in the SCPC is complex as the charging slope changes each time that a switch changes state. A SPICE model is suggested for analysis. The operation of the circuit can be explained empirically, however, to show how the propagation delay tp affects the answer. See figure 5.3.2, which illustrates a graphical analysis.

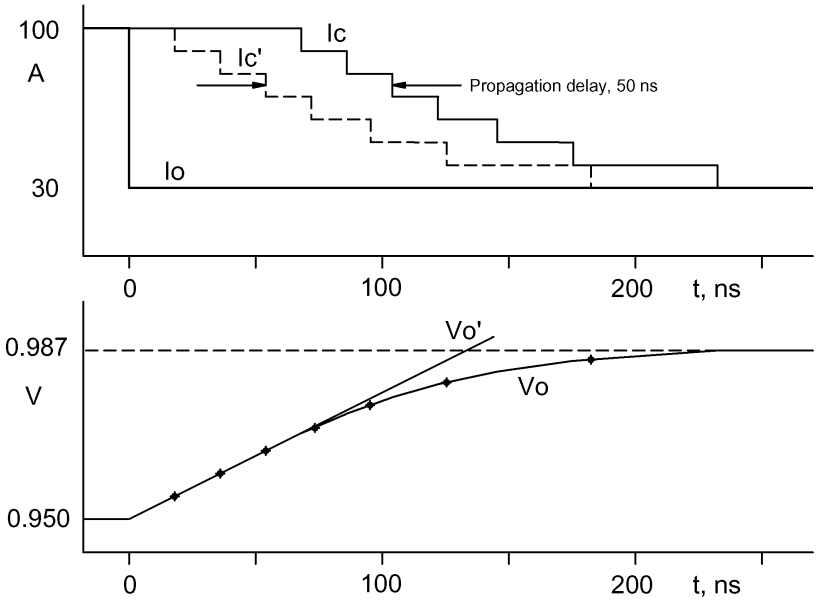

Figure 5.3.2. In response to a step change in the output current Io, the voltage Vo begins to rise in the output capacitor. The tick marks indicate when the voltage crosses the several thresholds of the resistor ladder network, and the dashed curve in the top graph indicates the "ideal" converter current. The actual converter output current Ic is delayed by 50 ns. Once the switches actual begin switching, the error current is reduced, and the charging slope is lower.

A t0, the current steps instantly from 100 A to 30 A, a step of -70 A. The output voltage Vo begins to rise, and in response, the SCPC decreases the converter current Ic to match. Because the error current Ie must discharge the output capacitor Co, the voltage does not change instantly, but ramps up at a steady rate, initially, of approximately 2.8 mV per ns.

The nodes of the resistance ladder network are 5 mV apart, to define the output impedance of 0.5 mΩ. As the voltage rises, as each node is reached, 10 A less current is provided, but only after the propagation delay of 50 ns. Therefore, the charging rate dVo/dt is a straight line Vo' at first, until the currents actually change at around 70 ns.

With the error current Ie = + 70 A, the output voltage Vo on the output capacitor Co rises 5 mv in 18 ns.

$$tx = Co\ dVo / Ie$$

$$tx = 250 \times 10^{-6} \times 5 \times 10^{-3} / 70$$

$$tx = 18\ ns$$

So, the first threshold is reached in 18 ns, the second threshold is reached in 36 ns, the fourth is reached in 54 ns.

Given the propagation time of 50 ns, the current Ic begins to decrease from 100 A in 10 A increments, 50 ns after the first threshold is reached, or at 68 ns, and so forth, as shown in figure 5.3.2.

The voltage actually overshoots slightly, to 0.987 whereas the ideal voltage is 0.985. The capacitor value of 250 uF is right at its limit to control overshoot for a propagation delay of 50 ns. In the SPICE modeling, the recommended value of the MLCCs from Intel ® VR 10.2 is used, 540 uF. This provides a good margin for component tolerance and environmental changes.

5.3. Capacitor Value and Switching Frequency:

The ripple voltage and the current steps likely are known design parameters. If so, the value of the output capacitor determines the switching frequency. As in any hysteretic control, the switching frequency is not constant, but it will be maximum when the output load current is half way between current steps, or Io = X + 5 A, for the example of a SCPC with 10 A steps. If the charging time tc is the time for the voltage to go from one threshold to the other, then

$$C = Ie\ tc / Vpp = Ie / 2\ Vpp\ fs$$

$$tc = C\ Vpp / Ie$$

$$fs = Ie / 2\ C\ Vpp$$

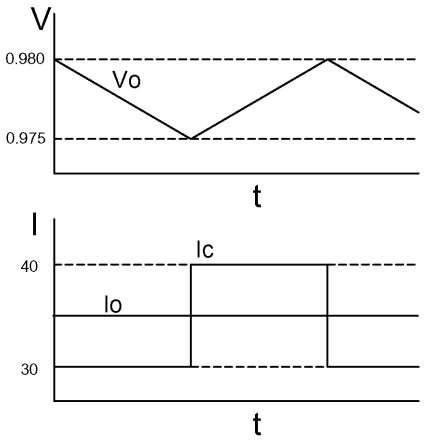

Figure 5.4.1. The switching frequency is the reciprocal of the time that it takes to cycle from one comparator threshold to the other and back. The control is hysteretic in nature, so the switching frequencies varies, and is maximum when the switching is symmetrical, that is, when the output current is half way between steps, or X + 5 A.

As an example, find the switching frequency, given a 10 A step and a 5 mV p-p ripple voltage with an output capacitance of 540 uF:

$$fs = Ie / 2\ C\ Vpp$$

$$fs = 5 / 2 \times 540 \times 10^{-6} \times 5 \times 10^{-3}$$

$$fs = 926\ KHz$$

The switching frequency is the frequency of the active switch, if the switch timing of figure 2.1.7 is used. If the cyclic "ring-counter" timing of figure 2.1.9 is used, then the frequency of any one switch is the switching frequency fs divided by the number of channels. With 10 channels, given the example, each switch has a nominal switching frequency of about 100 kHz (or less, as it is hysteretic) with a steady state load current Io.

This low frequency ensures low switching losses, especially as each channel switches a constant current of 10 A. With the constant current, the gate drive can be optimized for that current.

6.0. Reduced Losses for Standby:

When one of the SCPC outputs is not switched on, a current circulates to the return. While this feature allow the current to be switched to the output nearly instantly, there are some losses associated with the circulating current. This has raised some concern about efficiency with a low current output.

While a high efficiency at low power is meritorious in a power converter, more important is the total system energy consumption. *The greatest opportunity for power savings is to allow the processor to be in a lower power state (or off) for more of the time.* However, a processor cannot enter a lower power state (or turn off) if it cannot recover quickly enough to meet system timing demands. The very fast response of the SCPC allows the processor to go to sleep in more circumstances, and stay there for longer, more than compensating for the small losses in the circulating currents. However, there are strategies for reducing this power, and some of them are explored in this section.

Theoretically, a short circuit dissipates zero power, but in practical circuits, there is always some resistance. Every effort should be made to reduce the resistance as much as practical. This improves the efficiency for all modes of operation.

In the transformer configuration of the switched current power converter, the circulating currents are one half of the output current for a given section of the transformer, because the shorting switch is across both sides of the secondary push-pull winding. The I^2 component of the loss is one fourth.

During the off time, the shorting switch in the simple power IC has a very low on resistance, reducing losses further. See: http://eherbert.com/simpleic.pdf.

6.1. Reduced Current for Idle Modes:

If there are modes of operation where the output current can be reduced, the currents in the current sources can be reduced. Because the losses are proportional to I^2, a modest reduction in current can result in a significant reduction in losses. If the current is reduced by one half, the losses attributed to the current squared (I^2) are reduced by 75%.

After the processor has entered an idle state, the current in the current sources can be reduced. If the current source is a buck converter, the current bleeds down slowly, as the V driving the di/dt in the inductor L is low, but that is unimportant, because the output switches accommodate any state of the source currents as long as the output current demand does not exceed the total of the available currents.

A design tradeoff has to be made for the recovery mode. With a basic buck converter current source, if the current is reduced for standby, the energy is not

immediately available for an "instant" recovery to full output current. The recovery scenario is one of the following:

1. The load may have a di/dt limit initially, during recovery.

2. A signal in advance of recovery can be sent to the power converter to restore full output current capability. The power converter then refreshes the currents to their normal value.

6.2. Current can Track Demand:

A more generalized power saving mode is realized by letting the current sources track actual demand rather than just decreasing to a fixed value. Please review the timing diagram of figure 2.1.9. The switched current power converter of this example has a full power capability of 100 A, and if only 30 A is needed, the switches have a 30 percent duty cycle. The total current, when the switches are on, is three times 10 A, or 30 A, and the circulating current is seven times 10 A, or 70 A.

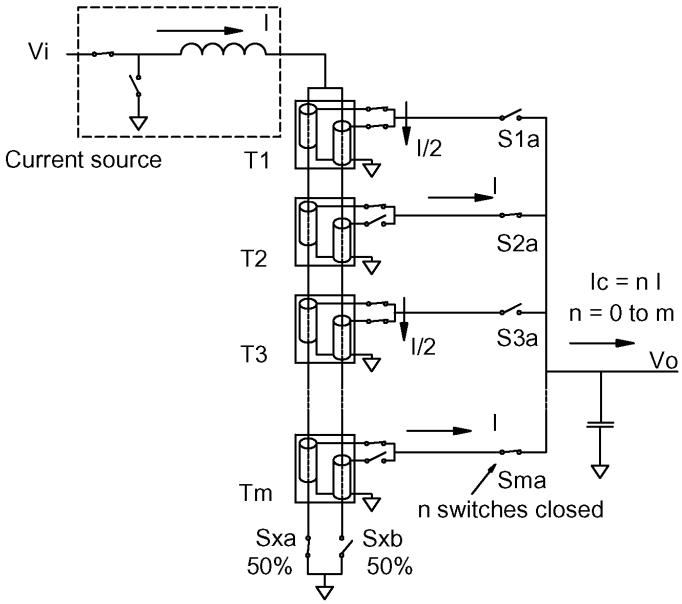

6.1.1. If the current is reduced in the current source, the power lost in circulating currents is reduced by two mechanisms. In all circuits, the power is reduced as the square of the current. And, with a lower input current, more of the switches will be closed to the load to generate the required output current Ic, reducing the number of sections that have circulating current.

Let the constant currents track the actual output such that all channels operate at 70 percent duty cycle. At steady state operation, seven switches are on to the output capacitor and the load, so the current in each is 4.3 A, or 30 A total, as before even though the individual currents are less. The circulating current is

reduced to three times 4.3 A, or 13A. The losses attributable to I^2 are reduced by 80 percent.

In such an operational mode, the instant full current capability is reduced to 43 A. If a transition to full power is necessary, either advance notice to the power converter to change mode is required, or the di/dt of the processor has to be restrained (in the positive direction only) until the currents recover. For many applications, this constraint would not be burdensome.

6.3 Reduced Power Operation with "Instant" Recovery.

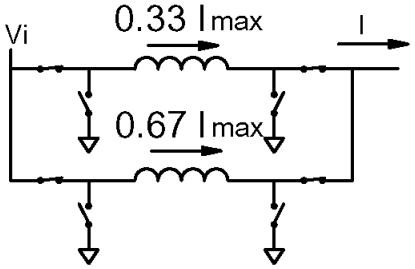

Figure 6.3.1. The current source for the transformer can be two current sources, for example 0.33 I and 0.67 I. Currents of 33 %, 67 % or 100 % can be selected.

If the power converter must be able to recover to full output load "instantly", the circuit may be modified to provide that capability.

A simple way to reduce power for standby is to have a two level current capability in the prime power source, as shown in figure 6.3.1. Upon command, the power converter changes the current reference for the current sources to a specific lower level for standby. If full current capability is required, both prime current sources are switched on, full current flows in the transformer primary winding and full current is available to the secondary circuits. This transfer can be very rapid, a couple of microsecons.