

Edward Herbert Designs

Saving Power by Using Switched Current Power Converters

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.

Saving power is important. Saving power not only reduces energy costs, it also results in a more economical and reliable design because of reduced heat loading.

Good efficiency of the power converter is fundamental to power savings, including at very low loads (standby). In a power supply for a microprocessor, more power savings result from operating the microprocessor at reduced voltage. Being able to make very fast and accurate transitions between voltage states allows the microprocessor to spend more time in a reduced power state, for lower total energy consumption and less heat generation.

The switched current power converter provides excellent efficiency, line-cord to processor, and low power usage in standby.

The switched current power converter with switched charge circuits provides very fast warm start. It can transition accurately from 0 volts, 0 amps to any VID voltage at full rated current in under two microseconds.

It also provides very fast recovery from a reduced voltage. As an example, it can transition accurately from 0.9 volts to 1.0 volts in less than one microseconds. Thus the microprocessor can be in a reduced power state for much more of the time.

The steady-state switching frequency is low, typically 100 kHz.

1.0. *Desirable Characteristics for Power Converters.*

1.1. Overview:

A power converter for a microprocessor must respond to very fast changes in the load current (di/dt) while maintaining tight voltage regulation, usually with specified impedance characteristic. There should be no appreciable overshoot or undershoot.

- The switched current power converter can go from full load to zero load and back again, or to any intermediate load, as fast as solid state switches can change state. There is no *theoretical* limit to its di/dt , though in a practical circuit the di/dt will be limited by parasitic inductances.

For “Demand Based Switching”, the power supply must change its voltage level very quickly, while maintaining a fast response to changes in current at each level and during the transitions.

- The switched current power converter with switched charge circuits can change its output voltage very quickly in precise steps. There is no theoretical limit to the dv/dt , though in a practical circuit the dv/dt will be controlled.

Turning off part of the server or computer reduces the power in those components to zero. Restoration to full operation requires a power supply that can return to any VID voltage very quickly and accurately, with full output current capability.

- The switched current power converter with switched charge circuits can transition from zero voltage, zero current to *any VID and full current capability* nearly instantly, with no theoretical limit on the dv/dt . In a practical circuit, the dv/dt will be controlled so as not to damage the load.

The losses should decrease as the output current decreases, to maintain good efficiency at low loads.

- In the transformer coupled switched current converter, the core losses are proportional to the output load current, going to zero as the output current goes to zero. At very low output currents, the entire transformer can operate at a low duty cycle, going to zero as the output load goes to zero.

The transformer coupled switched current converter provides isolation, so no upstream isolation is necessary. The entire server or computer can operate on one or two non-isolated fixed current loops, a very simple, efficient and economical front end, even if PFC is provided.

Following is a description of the components of the switched current power converter, with references for more information.

2.0. Power Savings using a Stepped Output Voltage:

2.1. Microprocessor Power Change with Input Voltage:

The input voltage to a processor greatly affects the power dissipation of a microprocessor. As a generalization, the power dissipation will change as the square of the input voltage, V_i^2 . As an example, if the voltage is reduced to 90 percent, the power would be to 81 percent, nearly a 20 percent savings.

The voltage on a microprocessor must be held at a maximum level to permit the fastest clock rates, but much of the time, the demand on the processor is lower, and the clock rate can be reduced. A reduced clock rate permits a lower input voltage. The reduced clock rate results in a power savings itself, as a component of the power used is directly proportional to the clock rate. There is yet a further reduction in power if fewer gates are used.

2.2. Demand Based Switching:

There has been recent work on a concept called "Demand Based Switching". Data was presented for a four gigahertz server showing four increments of voltage, clock speed and power, from a high of 1.39 volts, 4 G, 103 W to a low of 1.2 volts, 2.8 G and 65 W.

More information can be found in a paper: "[Enterprise Platform Power Efficiency for Future Server Systems](#)", by Tom Aldridge, Manager, Enterprise Power & Thermal Technology Labs, Intel Corporation, presented at IBM Platform Technology Symposium, September 14-15, 2004. (The paper is on the internet-- you can click the title while "on line" to view it).

2.3. Turning Off the Input Voltage:

Turning off the processor input voltage, either to the whole processor or one or more of its cores, saves the most power, as the power consumption goes to zero. While it is not possible to turn off all processing functions, the registers and pointers are retained if the cache is kept powered. With a multi-core processor, all but one of the cores can be shut off when not needed. To do this, the power supply must be able to restore the input voltage very quickly and accurately. The switched current power converter with switched charge circuits can do this.

For simple tasks, such as manual data entry, the main processor can be off most of the time, turning on only when there is a demanding task, such as updating a spreadsheet or repaginating a long document.

A video processor can shut down between screen refresh cycles.

3.0 Switched Current Power Converter:

3.1. Overview:

Some background on the switched current power converter is necessary for this discussion. A brief overview follows, and more information can be found at the web site <http://eherbert.com>.

As can be seen in figure 3.1.1, the outputs of the current sources can be switched to the output capacitor and the load as fast as the switches can change state. With MOSFETs, that is very fast indeed, permitting very fast rates of change of the current (di/dt). Because the current control is so fast, the output capacitor can be very small, perhaps an order of magnitude smaller than in a multiphase buck converter with a comparable output current rating.

As long as the current sources are reasonably accurate and stable, their dynamic characteristics have no effect whatsoever on the load dynamics, so the analysis of the switched current power converter can be separated into two entirely independent parts, control of the current sources and control of the output current.




Figure 3.1.1. In the switched current power converter, a number of current sources may be switched to the output capacitor and the load as quickly as solid state switches can change state, permitting a very fast di/dt .

3.2. Losses in the Switched Current Power Converter.

One variant of the switched current power converter used coaxial transformers to provide the multiple parallel current sources. Because the transformer provides

isolation, no upstream isolation is needed, so a much more efficient front end can be used. Unlike a transformer in a voltage source, where the core losses are nearly constant as long as the voltage is present regardless of the load, the transformer coupled switched current converter has core losses that are a function of the load current, going to zero as the load goes to zero. Figure 3.2.1 shows the relationship between the transformer core losses and the output current. This is explained in section 5.0 of this presentation.




Figure 3.2.1. The core losses in a transformer coupled switched power converter are a function of the output load, and go to zero at no load.

The secondary winding losses are more complex, but they also go to zero at no load, as seen in figure 3.2.2. Losses in the secondary switches are included.




Figure 3.2.2. The losses in the secondary winding, including the synchronous rectifiers and the current control switches, decrease to 50 percent of the full load losses at no load. With an added control for the primary switches, the losses at low load decrease to zero at no load.

4.0 Switched Charge Circuits:

4.1 Overview:

The voltage V on the output capacitor C_o depends upon its capacitance C , in farads, and the charge on the capacitor Q , in coulombs, according to the familiar formula $V = C Q$. If a fixed charge is added to or removed from the output capacitor, its voltage will change, up or down, by a fixed amount.




Figure 4.1.1. When the switch changes state, a fixed charge Q is added to or removed from the output capacitor, resulting in a fixed step in the output voltage, up or down.

The switched charge circuit transfers a fixed charge Q to or from the output capacitor when a switch changes state, causing a rapid, precise step change in the output voltage, up or down. The step change happens independently of the voltage and current controls, and the switched charge circuit has no ability to control voltage once the step has occurred. The voltage reference must change simultaneously, and the new voltage level is then maintained by controlling the current.




Figure 4.2.1. In a practical switched charge circuit, the switches will be MOSFETs.

4.2. Practical Switched Charge Circuit:

A practical switched charge circuit uses solid state switches to transfer the charge, such as the MOSFETs shown in figure 4.2.1. Because of the large peak current, a decoupling storage capacitor C_s should be used. The storage capacitor C_s may be refreshed from a small supply, perhaps the housekeeping supply. Alternatively, a controlled voltage may be used, so that the size of the step change in voltage can be adjusted for accuracy.

4.3. Losses Associated With Charge Transfer.

The transfer of charge has losses, and a representative analysis is attached to this presentation as Appendix A. The loss may be trivial, if the voltage is stepped infrequently, but it can be significant if the voltage is stepped frequently, and an energy tradeoff is necessary for each application.

Regardless of all other considerations, the loss with each step in voltage is proportional to the charge transferred. The charge transferred for a particular step in the output voltage is a direct function of the size of the output capacitor. *Therefore, the loss is directly proportional to the size of the output capacitor.* The smaller the output capacitor, the lower the lost energy per step.

The switched current power converter has a very much smaller output capacitor than a multiphase buck converter of comparable current capacity, so a step voltage capacity using switched charge requires much less energy per step. The switched charge circuit is easy to add, as shown in figure 4.1.3.




Figure 4.1.3. It is straightforward to add a switched charge circuit to a switched current power converter. In the example, one switched charge circuit may step the output voltage between zero and a lower operating voltage. For “turbo” operation, the second switched charge circuit may step the output voltage precisely to a higher voltage.

4.4 Controlling dv/dt :

While a very fast voltage step is desirable, there are limits to how fast a dv/dt should be used. The microprocessor may react badly to too fast a dv/dt , and the on

resistance and saturation current of the MOSFETs will limit the current flow, which will limit the dv/dt . It is noteworthy that the step voltage and the losses are fixed, independent of the circuit resistance. Only the rate of the change (dv/dt) is affected by the resistance. Inasmuch as MOSFETs are transconductance devices, the current can be controlled by the gate to source voltage. This can provide a step voltage change with a controlled linear ramp (dv/dt), if desired.

5.0. Transformer Coupled Current Sources:

5.1. Coaxial Push Pull Transformer.

A coaxial transformer is very good source of equal currents for a switched current power converter. The coaxial transformer provides isolation, so no other isolation is needed upstream. The transformer core losses are proportional to the output current, and go to zero at no load.

The coaxial push pull transformer has a single turn push pull primary winding passing through a number of modules. Each module has its own magnetic core and a coaxial push pull secondary winding. More information on the coaxial push pull transformer may be found on the web site <http://eherbert.com/>.




Figure 5.1.1. A coaxial push pull transformer has a number of modules that are magnetic cores with secondary windings. With a single turn push pull primary winding, the turns ratio is n to 1, where n is the number of modules.

The turns ratio of a coaxial push pull transformer having a single turn push pull winding (two turns center-tapped) passing through a number n modules is n to 1. If the primary is switched at 100 per cent duty cycle (50-50, with no off time), and if the primary circuit is powered from a constant current source, the output of each module is itself a constant current source equal to the primary current, because all modules couple to the same primary winding. See figure 5.1.2.

A valuable feature of the coaxial push pull transformer is that the individual modules can be turned on or off. If a module is “on”, its output current contributes to the total output current. If the primary current is I , and a quantity m modules are “on”, then the output current is m times I . An output current between integer steps of I is derived by pulse width modulating one or more of the modules to provide an intermediate average current as needed.

A module may be turned off by turning off both of the synchronous rectifiers and turning on a shorting switch, as shown in figure 5.1.2. The enabling logic is shown in more detail in figure 5.1.3. When a module is “off”, there will be a circulating current through the shorting switch equal to one half of the primary current. The current is one half of the primary current because the ends of the push pull secondary winding are shorted, and that includes two turns. There are some losses in the circulating current, but while the module is “off”, the core is short circuited, there is no voltage

across the core and the flux change is nil. *The core loss therefore is zero for that module.*




Figure 5.1.2. If a coaxial push pull transformer is operated from a constant current source, each module is a constant current source equal to the primary current, because all of the modules couple the same primary current.

As the output current is reduced, a point will be reached where the current is less than one integer increment of the input current, I . Stated another way, this is when the output current is less than $1/n$ times the maximum rated output current. Below this current, the output current is satisfied by pulse width modulating one of the output switches, with all of the rest being “off”. (Exactly which switch is modulated “on” to satisfy the output current demand is unimportant, and they may be sequenced, to distribute the switching losses.)

When the one modulating module is “off”, all of the other modules are “off” as well, so there is no utility in generating any secondary currents. Accordingly, the whole transformer can be turned “off” during the off time, which significantly reduces the losses.

Because the primary circuit is driven by a constant current source, the primary current cannot be interrupted, but if both of the primary MOSFETs (Qap and Qbp in figure 5.1.2) are turned “on”, the primary will be short-circuited and no currents will flow in the secondary windings. This is the off state of the transformer as a whole.




Figure 5.1.3. The Clock input is synchronized to the primary switching. When the On input is high, the synchronous rectifier MOSFETs switch to produce a full wave rectified output equal to the input current I . If the On input is low, both of the synchronous rectifiers are turned off, and the shorting switch is turned on, effectively short-circuiting the module.

With the primary short-circuited by turning on both of the primary MOSFETs Q_{ap} and Q_{bp} , no secondary currents will flow and there will be no flux switching in any core. The transformer losses will be zero, except for the I^2R loss in the primary current, and that will be divided between the two sides of the push pull winding. Nonetheless, the transformer can revert to 100 percent output current in a few nanoseconds, simply by restoring normal switching to the primary MOSFETs and turning on all of the modules.

6.0 Isolated SCPC Power Distribution:

6.1.1. Switched Current Power Converter

The variable turns ratio coaxial push pull transformer of figure 5.1.2 is powered from a constant current source, and the output current may be varied from zero to n times the input current (where n is the number of modules) by controlling whether the various modules are conducting to the output, as a switched current power converter. Several switched current power converters can be connected in series to a single current source, as seen in figure 6.1.1.




Figure 6.1.1. Several switched current power converters can be powered in series from a single constant current source. Because the coaxial push pull transformer provides isolation, upstream isolation is not necessary.

The constant current source may be a buck converter without an output capacitor, as shown in figure 6.1.1. The buck converter is inherently a current device, and it is easily controlled for constant current with a simple control that senses the output current. A hysteretic control works well, and is preferred for its simplicity.

Figure 6.1.1 also shows that several switched current power converters can be powered from a single constant current source as long as the input voltage is sufficiently high. The input voltage must be higher than the sum of the primary voltages of the transformers that are in series. The primary voltage of each switched current power converter is proportional to its load current, so it is maximum at full load and will equal n times the output voltage, where n is the number of modules.

Because the coaxial push pull transformer provides isolation, the constant current source may be an unisolated source. Having the isolation at the switched current power converters increases efficiency, as the transformers are operating at 100 percent duty cycle and the modules that are switched out during light load operation have no core losses.

7.0. Rapid Recovery from Standby:

Representative power converter states defined:

There are a number of possible states for a switched current power converter with switched charge circuits for a microprocessor. Some representative states, and their transition times to a stable output voltage are discussed.

1. **Power is disconnected**, so the power consumption is zero.
2. When the **power is first connected**, the PFC front end will charge the storage capacitors. When that is finished, almost no power is consumed, but there is energy stored on capacitors. This probably takes tens of milliseconds to be stable.
3. **Enable “stand by mode”**. Current is established in the current loop, but the processor power supply is off. This takes tens of microseconds to be stable. Power is low, but a goal of ½ watt may not be realistic.
4. **Turn on processor power supply**, low power mode, 0.900 volts, full current capable. This takes a few microseconds. Power is as required by the load, plus losses.
5. **Step processor to turbo mode**, 1.000 volts, full current capable. This takes tens of nanoseconds. Power is as required by the load, plus losses.
6. **Step processor to low power mode**, 0.900 volts, full current capable. This takes tens of nanoseconds. Power is as required by the load, plus losses.
7. **Turn off processor power supply**, 0 volts, no current, no power. This takes a few microseconds. Reverts to standby power.

Steps 4, 5, 6 and 7 can be repeated as desired, to enable the processor to be in the lowest power state consistent with the computing demand. The transition times given are the minimum practical but can be increased if it is desired to limit dv/dt.

Appendix 1:

A.1. Energy Lost when Switching Charge:

Some energy is lost when transferring charge between capacitors, and understanding this loss is important for thermal and efficiency calculations.

A.1.1. A Step Voltage of 0.1 Volts:

Consider that a processor operates at 1.000 volts for a “turbo” mode and 0.900 volts the rest of the time. The circuit of figures A.1.1 and A.1.2 is used to step the output voltage by removing or injecting a charge Q from or into the output capacitor. With a switch impedance of $10\text{ m}\Omega$, the transition takes a couple hundred nanoseconds. Care must be taken to reduce the parasitic inductance to an absolute minimum, so that it does not ring.

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




Figure A.1.1. At time = t_0 , the output voltage V_o is 1.000 volts. At time = t_1 , the switch changes state and removes charge from the output capacitor C_o to step the output voltage V_o to 0.900 volts.




Figure A.1.2. At time = t_2 , the output voltage V_o is 0.900 volts. At time = t_3 , the switch changes state and adds charge to the output capacitor C_o to step the output voltage V_o to 1.000.

Let $V_c = 12\text{ v}$, $C_o = 250\text{ }\mu\text{F}$ and $C_q = 2.1\text{ }\mu\text{F}$.

In figure A.1.1, the voltage V_o is 1.000 V at t_0 . The switch is thrown at t_1 to remove some charge and reduce the output voltage to 0.900 V.

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

At t_0 , the energy is as follows:

$$E_{c_q} = \frac{1}{2} (2.1 \times 10^{-6}) (11^2) = 127\text{ }\mu\text{J}$$

$$E_{co} = \frac{1}{2} (250 \times 10^{-6}) (1.000^2) = 125 \mu\text{J}$$

$$E_{total} = 252 \mu\text{J}$$

At t_1 ,

$$E_{cq} = \frac{1}{2} (2.1 \times 10^{-6}) (-1^2) = 1.1 \mu\text{J}$$

$$E_{co} = \frac{1}{2} (250 \times 10^{-6}) (0.900^2) = 101 \mu\text{J}$$

$$E_{total} = 102 \mu\text{J}$$

The energy E_{co} in the output capacitor has been reduced by $24 \mu\text{J}$, and the total energy is reduced by $150 \mu\text{J}$.

The difference is the energy lost, $126 \mu\text{J}$.

In figure A.1.2, the voltage V_o is 0.90 V at t_2 . The switch is thrown at t_3 to add some charge and increase the output voltage to 1.000 V . At t_2 , the energy is as follows:

$$E_{cq} = \frac{1}{2} (2.1 \times 10^{-6}) (-1^2) = 1.1 \mu\text{J}$$

$$E_{co} = \frac{1}{2} (250 \times 10^{-6}) (0.900^2) = 101 \mu\text{J}$$

$$E_{total} = 102 \mu\text{J}$$

At t_3 ,

$$E_{cq} = \frac{1}{2} (2.1 \times 10^{-6}) (11^2) = 127 \mu\text{J}$$

$$E_{co} = \frac{1}{2} (250 \times 10^{-6}) (1.000^2) = 125 \mu\text{J}$$

$$E_{total} = 252 \mu\text{J}$$

The energy E_{co} in the output capacitor has been increased by $24 \mu\text{J}$, and the total energy is increased by $150 \mu\text{J}$.

The difference is the energy lost, $126 \mu\text{J}$.

Therefore, the energy lost per cycle is $252 \mu\text{J}$. This is about 5 watts for twenty thousand cycles per second, or once per $50 \mu\text{sec}$ on average. This is a very worthwhile expenditure of power if it allows the processor to be in a reduced power state for a significant part of the time.

The energy lost may be reduced 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 volts, the losses are reduced by 75 percent. As a tradeoff, the R_{ds} of the MOSFET and the stray inductance has to be much lower to achieve the same step time (dv/dt), or else the

rise and fall time will be longer. If a lower voltage source is already available, using it may be worthwhile. Whether it would be worthwhile to add a voltage source would be a tradeoff of the particular application.

A.2.1. 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 A.2.1 and A.2.2 shows the charge transfer. The voltage can be brought up to 0.900 volts in a few microseconds. If it is desired to bring the voltage up to 1.000 volts, this circuit can be used in combination with the circuit of figures A.1.1 and A.1.2. See also figure 4.1.3. Care must be taken to reduce the parasitic inductance to an absolute minimum, so that it does not ring.




Figure A.2.1. At time = t_0 , the output voltage V_o is 0.0 volts. At time = t_1 , the switch changes state and adds charge the output capacitor C_o to step the output voltage V_o to 0.900 volts.




Figure A.2.2. At time = t_2 , the output voltage V_o is 0.900 volts. At time = t_3 , the switch changes state and removes charge from the output capacitor C_o to step the output voltage V_o to 0.0.

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

Let $V_c = 12$ v, $C_o = 250$ μ F and $C_q = 20.3$ μ F.

In figure A.2.1, the voltage V_o is 0.0 V at t_0 . The switch is thrown at t_1 to add charge and increase the output voltage to 0.900 V.

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

At t_0 and t_3 , the stored energy is zero.

At t_1 and t_2 ,

$$E_{cq} = \frac{1}{2} (20.3 \times 10^{-6}) (11.1^2) = 1251 \mu\text{J}$$

$$E_{co} = \frac{1}{2} (250 \times 10^{-6}) (0.900^2) = 101 \mu\text{J}$$

$$E_{total} = 1,352 \mu\text{J}$$

This energy is required to bring the output voltage up to 0.900 volts, 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 watts 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.