

SoC Energy Savings = Reduce+Reuse+Recycle: A Case Study Using a 660MHz DC-DC Converter with Integrated Output Filter

Guy Lemieux, Mehdi Alimadadi, Samad Sheikhaei, Patrick Palmer, and Shahriar Mirabbasi

Dept. of ECE, University of British Columbia, 2332 Main Mall, Vancouver, BC, Canada

{ lemieux, mehdi, samad, prp, shahriar } @ ece.ubc.ca

Abstract – This paper advocates ‘reduce, reuse, recycle’ as a complete energy savings strategy. While reduction has been common to date, there is growing need to emphasize reuse and recycling as well. We design a DC-DC buck converter to demonstrate the 3 techniques: reduce with low-swing and zero voltage switching (ZVS), reuse with supply stacking, and recycle with regulated delivery of excess energy to the output load. The efficiency gained from these 3 techniques helps offset the loss of operating drivers at very high switching frequencies which are needed to move the output filter completely on-chip. A prototype was fabricated in 0.18 μm CMOS, operates at 660MHz, and converts 2.2V to 0.75-1.0V at $\sim 50\text{mA}$.¹

I. INTRODUCTION

Energy savings during circuit design should follow the advice given by environmentalists: *reduce*, *reuse*, and *recycle*. Energy *reduction* is already at the forefront of most designer’s thoughts: lower the energy needed for each basic operation or component to reduce overall total energy. However, energy *reuse* and *recycling* techniques should also be considered wherever possible to help lower the energy required to operate the overall system. In this paper, we will use all 3 of these techniques in the DC-DC buck converter shown in Fig. 1.

Energy *reduction* is the common practice of finding ways to use less energy to perform an operation. In this paper, we reduce energy in the front-end drive chain of the buck converter by using half-rail instead of full-rail voltage swings. Typically, reduction follows a greedy approach. However, it may be necessary to *increase* energy use in one area to reduce energy by a greater amount in another area.

Energy *reuse* is the practice of redeploying energy in small local regions in an unregulated fashion. An LC tank resonator is a good example of local energy reuse – by exchanging energy between the capacitor and inductor, the resonator efficiency is kept high and the amount of new energy needed to operate is kept low. Another type of reuse is charge sharing, sometimes called charge recycling, charge recovery, or even energy recycling. For example, temporarily shorting all wires of a full-swing bus between accesses reuses energy by precharging all nodes to $V_{dd}/2$ (on average), saving the bus drivers some energy on the next cycle. Reuse is a localized effect that requires careful planning of the interface between the provider and receptor of the energy. In this paper, we reuse charge (and therefore energy) by supply-stacking separate front-end drive chains for the output transistors.

We define *energy recycling* as taking under-utilized energy in one part and redeploying it to other parts of the system as a *regulated current or voltage supply*. This redeployment takes a low-quality source with some residual potential energy and uses it to construct a high-quality, reliable energy source. This simplifies interface design between energy provider and receptor. Since fully integrated on-chip regulators are usually inefficient, there is little prior work in this area. In this paper, we recycle energy by taking surplus charge available from PMOS drivers and sending it to the load as a regulated supply.

Care should be taken when employing the terms energy reuse and energy recycling. Energy cannot be reused in perpetuity. Energy is consumed as it eventually turns into heat. However, CMOS logic only requires a small portion of the energy it uses to perform an operation, and it throws away the excess. For example, clock distribution charges and discharges a large clock network capacitance every cycle. However, every discharge phase wastes the energy that was stored in this capacitor. We have shown this energy can be recycled [1][2].

We’ll employ the *reduce, reuse, recycle* design technique to make practical the high-frequency operation of a DC-DC buck converter. The only way a DC-DC buck converter can fit on-chip is to reduce the inductor and capacitor values. We do this by operating at 660MHz instead of the more typical 10s of MHz. Normally, such a high frequency is avoided due to high switching losses in the drivers. With recycling, it is tolerable.

Bringing DC-DC converters on-chip is an essential part of *energy recycling*. They can take under-utilized energy stored in one part of a system and redeploy it to other parts. In [1], we used recycling to recover 50% of the energy lost in the clock tree of a high-performance microprocessor. Here, we employ reduce+reuse+recycle to lower energy consumption of the DC-DC converter front-end drivers.




Fig. 1. Block diagram showing energy flow in the converter.

¹ This work was partially funded by an NSERC Discovery Grant and by CMC Microsystems for CAD tools and fabrication support.




Fig. 2. Block diagram of a CMOS buck converter.

The rest of this paper describes the design of the converter, the drive chain, and modes of operation. Results from a $0.18\mu\text{m}$ CMOS 660MHz prototype show how efficiency can be improved to make on-chip conversion more practical.

II. CIRCUIT DESIGN

A. Buck Converters and ZVS

The circuit diagram of a CMOS-based buck converter is shown in Fig. 2. The PMOS and NMOS output transistors alternate between connecting the inductor to a voltage source (PMOS) to store energy in the inductor L_f , and removing the stored energy from L_f to C_f and the load (NMOS). Since large currents are required, these output transistors are very wide. The low-pass output filter, consisting of L_f and C_f , converts the square wave on node V_{inv} into a DC output signal V_{out} with a small ripple – we target a peak-to-peak output ripple of $<5\%$.

In advanced switch mode power converters, zero voltage switching (ZVS) is used to manage dynamic power loss in the

drain/source of the output transistors [3]. For ZVS operation, the output transistors are turned on when v_{ds} across the source/drain is 0V, resulting in no power loss during switching ($p_{ds} = i_{ds} \cdot v_{ds}$ when $v_{ds}=0$). This is achieved by independently driving the gates of the NMOS and PMOS to provide 4 distinct modes shown in Fig. 3. Modes 1 and 3 are ‘regular’ modes: in mode 1, the PMOS is on to provide power, resulting in increasing inductor current I_{L_f} ; in mode 3, the NMOS is on to provide a return path for the diminishing I_{L_f} . ZVS occurs in modes 2 and 4 when both NMOS and PMOS are off.

ZVS works as follows. Consider C_x in Fig. 2 which includes all parasitic capacitance at node V_{inv} including the PMOS and NMOS drain capacitances. When both PMOS and NMOS are off, a positive inductor current (mode 2) removes charge from C_x , reducing V_{inv} , while a negative inductor current (mode 4) replaces charge in C_x , increasing V_{inv} . To implement ZVS, the NMOS is turned on only when V_{inv} reaches 0V, and the PMOS is turned on only when $V_{inv} = V_{dd}$.

B. Reduce

The wide NMOS and PMOS output transistors have large input gate capacitance, requiring them to be driven by a chain of tapered inverters which are sized for minimum power, and are referred to here as the front-end drive chain. Separate drive chains are required to allow precise control of the NMOS and PMOS turn-on and turn-off times to achieve ZVS.

Despite ZVS, which reduces energy waste in the final NMOS/PMOS pair, significant losses are associated with operating the two drive chains and the gates of the output transistors at high switching frequencies.

To reduce the energy lost at every transition, each drive chain employs low-swing signaling by swinging only half-rail, between 0 and $V_{dd}/2$ or between $V_{dd}/2$ and V_{dd} for NMOS and PMOS, respectively. This saves a significant amount of energy compared to full-rail switching. However, the outputs of the low-swing drive chains must turn on their respective NMOS and PMOS output transistors, so it is essential that $V_{dd}/2 > V_{tn}$ and $V_{dd}/2 > |V_{tp}|$. To increase overdrive, it is recommended that low- V_t devices be used for the NMOS and PMOS output transistors as well as the rest of the drive chain.

C. Reuse

A half-rail swing for both drive chains offers a further advantage: the NMOS and PMOS chain can share the common reference voltage of $V_{dd}/2$. This allows energy reuse in the form of voltage supply stacking as shown in Fig. 4. Charge used by the upper PMOS drive chain still has unused potential, so it can be reused by the lower NMOS drive chain. This half-rail, supply-stacking technique was also used in [4]. A more general case of supply stacking is called charge recycling in [5]; since it employs local regulation it could also be considered a form of energy recycling by our definition.

D. Recycle

The PMOS output transistor M_3 in Fig. 4 is three times wider than NMOS output transistor M_4 . As a result, the drive chain of the PMOS (top inverter chain) is much larger and requires more charge to operate than the drive chain of the NMOS (bottom inverter chain). Charge accumulates at node V_m ,




Fig. 3. Idealized timing diagram of the internal signals.




Fig. 3. Block diagram of the fully integrated DC-DC converter prototype with charge-recycling diodes.

which is stored in the middle capacitor C_m , and operates near $V_{dd}/2$. The accumulated charge powers the NMOS chain. In [4], excess charge is dissipated to V_{ss} through an additional regulator, forcing node V_m to $V_{dd}/2$. Instead, we *recycle* this excess charge by delivering it to the converter output load.

The recycling task is performed by two series diode-connected NMOS transistors, D_1 and D_2 . These diodes automatically deliver charge to the load when $V_m > V_{inv} + 2V_{tn}$ without the need for additional gating signals. Two diodes in series are needed to limit the drop of V_m when M_4 is ON and V_{inv} is low. The goal is to keep V_m near $V_{dd}/2$. Hence, accumulated charge at C_m is removed through the diodes by inductor L_f instead of an external regulator. The voltage divider R_1 and R_2 puts V_m near $V_{dd}/2$ at startup and does not significantly contribute to operational power.

In this design, weak negative feedback helps keep V_m near a stable operating point of $V_{dd}/2$. If V_m increases, the bottom chain receives a higher supply voltage, which increases its power intake and causes V_m to drop. At the same time, M_4 turns on with a higher V_{gs} and V_{inv} is pulled closer to V_{ss} , giving D_1 and D_2 higher V_{gs} , facilitating charge removal from C_m . Similarly, if V_m decreases, the top chain receives a higher supply voltage, which results in increasing its power intake and causing V_m to increase. Also, a lower V_m causes D_1 and D_2 to receive lower V_{gs} , facilitating accumulation of charge in C_m .

E. Prototype Implementation

We have constructed a prototype chip in $0.18\mu\text{m}$ CMOS. Node V_m is made available off-chip to be externally adjusted or probed. Resistors R_3 and R_4 are 50Ω terminators so V_{pp} and V_{pn} can be driven by external test equipment. Capacitance C_m was chosen to be 20 times larger than the NMOS C_{gate} to limit ripple at V_m . L_f and C_f values were chosen to be 4.38nH and 1.1nF , respectively, to operate at a switching frequency of

660MHz with a voltage ripple of less than 5% at 50mA load.

The NMOS transistors in the top inverter chain for M_3 need to have zero body voltage with respect to their source, so they are isolated from the p-substrate using n-well and deep n-well implantation as described in [6]. The same procedure is used for D_1 and D_2 , where the body should be connected to the drain to reverse bias the intrinsic body diode.

It is difficult to employ zero voltage switching (ZVS) at a high frequency. One method has been successfully reported by the authors in [1]. For this paper, test equipment generates inputs V_{pn} and V_{pp} with appropriate delays to achieve ZVS. Fig. 3 shows the ideal timing diagram. In the figure, the two time periods when both transistors are OFF are characterized as T_{delay1} and T_{delay2} , corresponding to the dead-time needed to implement ZVS operation for M_4 and M_3 , respectively.

A chip micrograph is shown in Fig. 5. The integrated inductor L_f is implemented with the top four metal layers put in parallel to reduce series resistance. A patterned ground shield is used to reduce substrate loss. The integrated capacitor C_f is implemented using gate capacitance of an array of NMOS transistors. The 4mm^2 total die area uses 2.5mm^2 for the converter. Even at 660MHz , the inductor dominates the area. Designed for a current of 50mA at 1V , the power converter achieves a power-to-area ratio of $20\text{mW}/\text{mm}^2$.

F. Prototype Limitations

There are a few limitations with the implemented prototype. First, the low- V_t transistors required for M_3 , M_4 and the drive chains were not available in the CMOS process that was used. Using them would help the drivers fully turn ON with the low-swing voltage supply, thereby reducing power consumption in the drive chains and improving power delivery to the output load. Instead, regular transistors were used, resulting in degraded efficiency of the prototype converter.




Fig. 5. Chip micrograph.

Second, power is lost due to the voltage drop across diodes D_1 and D_2 . The diodes were used to keep it simple for proof-of-concept, but a more complex circuit could be devised.

Third, we used a typical $4\times$ sizing ratio in the tapered drive chain instead of a high value (e.g., $10\times$) typical for low power.

Fourth, a more practical circuit would be more complete: ZVS timing, voltage regulation, and convert a battery voltage (e.g., 2-3V) to a useful voltage (e.g., 1.8V). We kept it simple.

III. EXPERIMENTAL RESULTS

A. Fabricated Prototype Measurements

Ten chips were tested to verify consistency. Due to the lack of low- V_t transistors, the chips were tested at 2.2V instead of the typical 1.8V for this technology. Conversion efficiency and V_{out} measurements with standard error (S_E) bars are presented in Fig. 6. Node V_m was above the 1.1V predicted by simulation, so we applied external regulation to sink excess charge, reducing efficiency. The output is adjustable between 0.75V to 1V by varying duty cycle D from 45 to 64% with a fixed load. Conversion efficiency, P_{out}/P_{in} , ranges 25 to 31%.

B. Simulation Results

Four variants of the circuit were simulated with compounding changes: (i) baseline converter, full-swing drivers; (ii) add low-swing/stacked drive chain to reduce+reuse energy; (iii) add diodes and C_m to recycle energy, similar to the prototype; and (iv) add low- V_t transistors. Low- V_t transistors were not in the CMOS design kit, so it is mimicked by a 0.1V voltage shift at transistor gates. Efficiency results from simulation are shown in Fig. 7. As expected, the circuit with all the options has the highest efficiency. Indeed, the efficiencies show improvement with each additional change. For example, at a 30% duty cycle, the efficiency of the circuits are (i) baseline 19%, (ii) low-swing 26%, (iii) recycling diodes 30%, and (iv) low- V_t 34%. At 40% duty cycle, the efficiency improves from 22% to 46% with the reduce, reuse, and recycle methodology.

IV. CONCLUSIONS

Design of a fully integrated DC-DC converter, implemented in a $0.18\mu\text{m}$ CMOS process, is used to demonstrate three techniques for energy savings. The design process saves energy in the front-end drive chain by following the environmentalist's mantra of reduce, reuse, and recycle. Energy was first reduced by employing ZVS and modifying the front-end to use half-swing instead of full-swing signaling. Energy was reused by employing separate PMOS and NMOS drive chains and supply-stacking them – an additional $V_{dd}/2$

simultaneously serves as output-low voltage for the PMOS chain and as the output-high voltage for the NMOS chain. Energy was recycled by diverting excess charge not needed by the NMOS chain and delivering it in a regulated fashion to the output load of the converter. Together the switching losses are sufficiently reduced to make it practical to operate a DC-DC converter at high frequency and bring the LC filter on-chip.

Energy recycling is shown with the use of on-chip DC-DC converters. By converting unused potential energy into a useful regulated supply, we are able to power other parts of a circuit instead of wasting energy by simply dissipating unwanted charge to ground. We believe many applications can benefit from this new design technique.

REFERENCES

- [1] M. Alimadadi, S. Sheikhaei, G. Lemieux, S. Mirabbasi, and P. Palmer, "A 3GHz Switching DC-DC Converter Using Clock-Tree Charge-Recycling in 90nm CMOS with Integrated Output Filter," *ISSCC*, Feb., 2007.
- [2] M. Alimadadi, S. Sheikhaei, G. Lemieux, S. Mirabbasi, W. Dunford and P. Palmer, "Energy Recovery from High-frequency Clocks using DC-DC Converters," *IEEE Int'l Symp. on VLSI*, Montpellier, France, April, 2008.
- [3] A.J. Stratakos, S.R. Sanders, and R.W. Brodersen, "A Low-voltage CMOS DC-DC Converter for a Portable Battery-operated System," *Power Electronics Specialists Conference*, June, 1994, pp.619-626.
- [4] J. Xiao, A. Peterchev, J. Zhang, S. Sanders, "A $4\mu\text{A}$ -Quiescent-Current Dual-Mode Buck Converter IC for Cellular Phone Applications," *IEEE Solid State Circuits Conference*, Feb., 2004, pp. 280-283.
- [5] S. Rajapandian, K.L. Shepard, P. Haucha, and T. Karnik, "High-voltage Power Delivery Through Charge Recycling," *IEEE JSSC*, June, 2006.
- [6] J. Lin, "Challenges for SoC Design in Very Deep Submicron Technologies," National Semiconductor Corporation, Oct., 2003.




Fig. 6. Measured performance of 10 chips with standard error (S_E) bars.




Fig. 7. Efficiency variation with duty cycle for various designs.