Energy Recovery from High-frequency Clocks using DC-DC Converters

M. Alimadadi, S. Sheikhaei, G. Lemieux, S. Mirabbasi, W. Dunford, P. Palmer.
The University of British Columbia, Vancouver, Canada
{mehdia | samad | lemieux | shahriar | wgd | prp }
@ ece.ubc.ca

Abstract

Large digital chips use a significant amount of energy to distribute a multi-GHz clock. By discharging the clock network to ground every cycle, the energy stored in this large capacitor is wasted. Instead, the energy can be recovered using an on-chip DC-DC converter. This paper investigates the integration of two DC-DC converter topologies, boost and buck-boost, with a high-speed clock driver. The high operating frequency significantly shrinks the required size of the L and C components so they can be placed on-chip; typical converters place them off-chip. The clock driver and DC-DC converter are able to share the entire tapered buffer chain, including the widest drive transistors in the final stage. To achieve voltage regulation, the clock duty cycle must be modulated; implying only single-edge-triggered flops should be used. However, this minor drawback is eclipsed by the benefits: by recovering energy from the clock, the output power can actually ***exceed*** *the additional power needed to operate the converter circuitry, resulting in an effective efficiency greater than 100%. Furthermore, the converter output can be used to operate additional power-saving features like low-voltage islands or body bias voltages.*

1. Introduction

The rapid increase in energy consumption of large digital circuits has been predominantly due to an increase in total gate capacitance combined with an increase in operating frequency. As a result, a large fraction of the total energy budget is used by the high-frequency clock network [1]. While the clock energy is obviously proportional to frequency, additional energy is also expended in modern designs to reduce clock skew and to drive increased capacitance caused by the more closely spaced wires and thinner gate oxides. Most of this energy is consumed in the final drive stage where a large load capacitance is charged and discharged every cycle.

There are several methods that are used to reduce clock distribution energy, such as gating the clock, low-swing signals, double-edge triggered flip-flops, adiabatic switching [2], and resonant clocking [3]. All of these previous techniques have attempted to reduce the power consumption of operating the clock. This paper investigates a new method, summarized in Figure 1, where the energy of the clock is not reduced directly but is instead recovered using a DC-DC converter and redeployed to another circuit in a regulated fashion. The redeployment reduces the total current draw from the primary supply. We call this concept *energy recycling* [4].

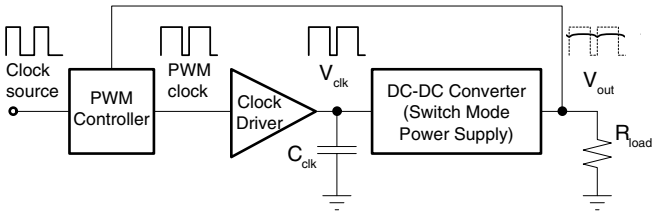

Figure 1. Recycling clock energy with a DC-DC converter.

Integrating a clock driver intended for high-performance logic with a DC-DC converter merges several compatible concepts. Using the same high switching frequency (e.g., 3GHz) for the clock and converter reduces the size of on-chip inductor and capacitor needed by the converter. Also, the final clock drivers and the DC-DC power transistors are both very wide to improve switching time of the clock and reduce output losses of the converter. These large, low-impedance transistors need to be driven by a tapered inverter chain to keep up with the high frequency. Similarly, the power used by this chain should be minimized in both cases. Lastly, many DC-DC converters use pulse width modulation for output regulation, a scheme compatible with single-edge triggered clocking.

One of the main advantages of this new approach is the efficient generation of an on-chip voltage supply which *differs* from the level offered by the primary supply. Since the DC-DC converter is small, several can be deployed across the chip to produce independent, regional power supplies. This allows several different regulated voltages to be on-chip at the same time, all powered from the same off-chip primary supply. Several power-saving techniques such as mixed-voltage islands and adaptive body bias (ABB) [5] can utilize additional supply voltages. An on-chip DC-DC converter can power these schemes without the need for external pins, external components, or board design effort.

Fully on-chip power converters are not very popular because they are not very efficient. To boost efficiency, a low switching frequency and large LC components are often used. For example, [6] uses 27mm^2 at 45MHz to reach an efficiency of 65%. The DC-DC converters designed here operate at far higher frequencies to shrink area by 99%, but maintain good efficiency by taking energy stored in the clock capacitance. This extra energy is available for free, since it would normally be wasted every cycle when the clock node is discharged. As a

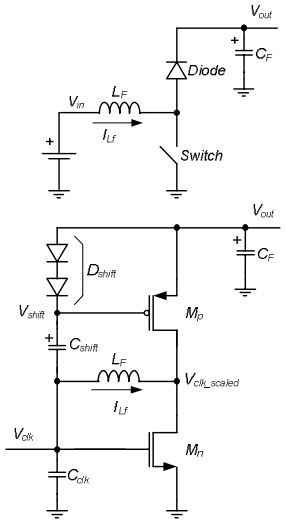

Figure 2. A boost converter (top) and the simplified diagram of the new converter (bottom).

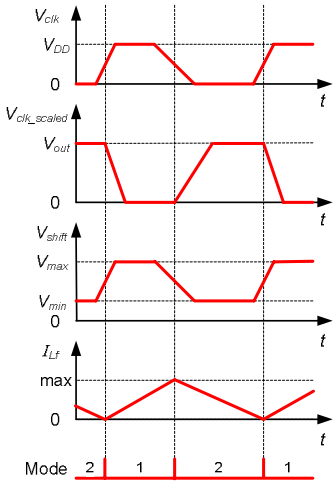

Figure 3. Idealized timing diagram of boost converter signals.

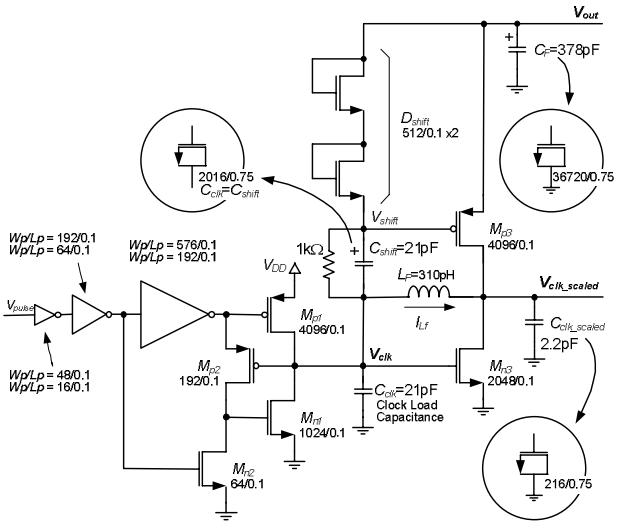

Figure 4. Fully-integrated clock driver and boost converter.

result, the power needed to operate the DC-DC converter is greatly reduced.

To compare our dual-purpose circuit with traditional on-chip power converters, we measure *effective efficiency* as the output power divided by the *incremental* input power needed to operate the converter. Effective efficiency captures how efficient a traditional converter would have to be if it was to supply the same output power using just the incremental input power used by our circuit. We consider an effective efficiency greater than 50% efficiency to be good for an on-chip converter. However, in some cases, our effective efficiency is greater than 100%. This is not evidence of the perpetual motion fallacy; it is proof that energy is being recovered from the clock tree and redirected to the converter output.

In our previous work [7], we constructed a working prototype of a circuit similar to the two circuits proposed in this paper. That work also demonstrated an effective efficiency greater than 100%. Hence, we are confident that the results shown in this paper are valid and not simulation anomalies. One drawback of the earlier work is that it employed a *buck converter* topology, where the output voltage is always less than the primary supply. This work explores two other common topologies, boost and buck-boost, because they can provide a greater output voltage range, including a magnitude greater than the primary supply level. Higher output voltages can also be useful for applications such as ABB.

To keep things simple, these new converters are designed and simulated in an 'open loop' mode with no voltage regulation capability. A more complete design, such as the one presented in [7], would include a feedback controller to regulate voltage by modulating the clock pulse width.

2. Integrated Boost Converter Design

In this section, we describe how a boost converter operates and show how it can be integrated with a clock driver. This circuit drives clock node V_{clk} loaded by C_{clk}. As well, it outputs a DC voltage, V_{out}, and drives another version of the clock, V_{clk_scaled}, at the same output voltage.

2.1. Background

A typical boost converter is shown in Figure 2(top). When the switch is on, voltage V_{in} will be across the inductor and it will build up current. In the next phase, when the switch is off, inductor current finds its way through the diode and charges the output capacitor.

An alternative circuit, shown in Figure 2(bottom), avoids the loss of a voltage drop across the diode. This design requires suitable gating signals for the power transistors. Since the output voltage is variable and can be higher than V_{DD}, a switched-capacitor voltage-shifter circuit is used to generate a shifted gating signal for M_p.

2.2. Basic Operation

The operation of the boost converter in Figure 2(bottom) can be described using the timing diagram in Figure 3. Ignoring the turn on/off times of the transistors, there are two basic modes of operation:

- Mode 1: At the beginning of this mode, V_{clk} goes high and M_n turns on. Consequently, voltage V_{DD} will be across the inductor L_F and the inductor current increases linearly (assuming a constant voltage across the inductor). At the same time, the voltage of capacitor C_{shift} will be added to V_{clk} so that V_{shift} reaches voltage V_{max}, a higher voltage than V_{out}. The diodes D_{shift} are reverse biased. As C_{shift} is pre-charged to $V_{out} - 2V_{diode\_drop}$ in the previous Mode 2, V_{gs} of M_p would be equal to: $V_{max} - V_{out} = (V_{DD} + (V_{out} -$

$2V_{diode\_drop}$)) – V_{out} = V_{DD} – $2V_{diode\_drop}$ which has a positive value and M_p turns off completely.

- Mode 2: As a new V_{clk} half-cycle starts, M_p turns on and M_n turns off. Capacitor C_{shift} will be charged through diodes D_{shift} to a value of $V_{out} - 2V_{diode\_drop}$. As the diodes are forward biased, V_{gs} of M_p becomes equal to $-2V_{diode\_drop}$ which has a negative value larger than the threshold voltage of M_p, turning it on completely. At this time, inductor current finds its way through M_p and will charge up the output capacitor C_F.

In the above discussion, the average voltage of C_{clk}, V_{in}, is $D \times V_{DD}$, where D is the duty cycle. This is the operating voltage available to the boost converter (and not V_{DD}). Ideally, the boost converter output voltage would be $V_{out} = \frac{1}{1-D} \times V_{in} = \frac{D}{1-D} \times V_{DD}$. Therefore, at a 50% duty cycle, the output voltage should be close to V_{DD}. To adjust and regulate the output voltage, the duty cycle of the clock signal could be changed by a control circuit.

In addition to providing output voltage levels higher than V_{DD}, the circuit produces a buffered version of the clock, V_{clk_scaled}, at the new voltage V_{out}. This clock signal can be used in the circuitry powered by the converter, but allowances for clock skew and level-conversion will need to be made in the datapath logic.

2.3. Complete Circuit

A complete implementation of the integrated clock-driver/boost-converter circuit is shown in Figure 4. Some transistors have been added to implement the capacitors. A tapered inverter chain has also been added to drive C_{clk}. C_{clk} is 21pF, a value selected to approximate the clock load in a 1mm^2 region of a high-performance processor.

One crucial change, the addition of M_{n2} and M_{p2}, has been made to delay M_{n1} from discharging the clock. By waiting for the inductor to discharge V_{clk}, the energy can be transferred to the converter instead of being wasted. This is known as zero-voltage switching (ZVS) because M_{n1} is turned on when the voltage across it is zero.

The delay circuit works as follows. The gating signals of M_{p1} and M_{n1} will have nearly the same falling edge (since V_{clk} is low and M_{p2} is turned on). On the rising edge, M_{p1} is turned off first. M_{n1} is turned on after the inductor has drained V_{clk} sufficiently low to turn on M_{p2}. This allows the gating signal for M_{n1} to turn on and discharge any small amount of remaining energy in V_{clk}.

The boost converter produces V_{out} and V_{clk_scaled}. A 2.2pF capacitor C_{clk_scaled} was added to represent the clock load for a circuit powered by V_{out}; its size is determined by the strength of M_{n3} and M_{p3} as well as the desired efficiency. Some of the recovered energy is subsequently lost when this capacitor is discharged, so it should be kept small. To keep the output ripple on V_{out} < 5%, a large capacitance C_F is needed for bulk energy storage.

Except for D_{shift} and C_{shift}, all transistor body terminals are connected to their source pins. The body terminals of D_{shift} and C_{shift} are connected to ground instead. This prevents forward biasing of the body-drain intrinsic diode, in case the drain voltage goes lower than the source voltage. Also, this makes the layout easier, since no deep n-well structure is required.

Finally, a 1kΩ resistor is added in parallel to C_{shift} to bias the D_{shift} diodes and provide a DC current path to avoid floating nodes when the D_{shift} is off.

3. Integrated Buck-boost Converter Design

In this section, we describe how a typical buck-boost converter operates and show how it can be integrated with a clock tree. The converter produces a negative output voltage. One use of a negative voltage is to control the adaptive body bias of NMOS devices. However, unlike the boost converter in the previous section, this circuit does not produce another copy of the clock, V_{clk_scaled}.

3.1. Background

A typical buck-boost converter is shown in Figure 5(top). When the switch is on, voltage V_{in} will be across the inductor and its current will build up. In the next phase, when the switch is off, inductor current finds its way through the diode and, as a consequence, a negative voltage appears at node V_{out}.

In low-voltage applications such as this work, the voltage drop across the diode would result in a considerable power loss. Instead, it is more suitable to use a transistor with an appropriate gating signal, such as the circuit shown in Figure 5(bottom). However, since the source of the NMOS transistor is connected to a negative V_{out}, the gating signal for this transistor should be appropriately shifted to make sure that the transistor is off when $\overline{V_{clk}}$ is zero. Here, the required shift in the gate voltage is achieved by a switched capacitor circuit consisting of a C_{shift} capacitor and three diodes in series. The detailed mechanism for the operation of this circuit is explained in the next section.

An extra switch S_{clk} is also added between nodes V_{clk} and V_{inv}. This switch prevents V_{clk} from becoming negative as V_{inv} goes below zero when M_n is on.

3.2. Basic Operation

The operation of the buck-boost converter in Figure 5(bottom) can be described using the timing diagram in Figure 6. Ignoring the turn on/off times of the transistors, there are two basic modes of operation:

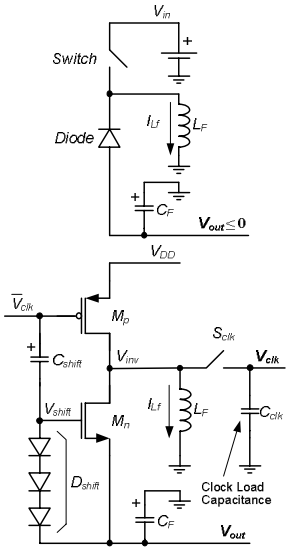

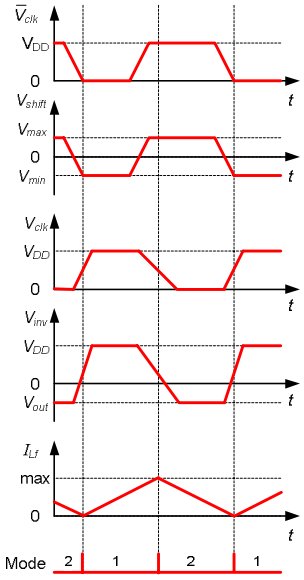

Figure 5. A buck-boost converter (top) and the simplified diagram of the new converter (bottom).

Figure 6. Idealized timing diagram of buck-boost converter signals.

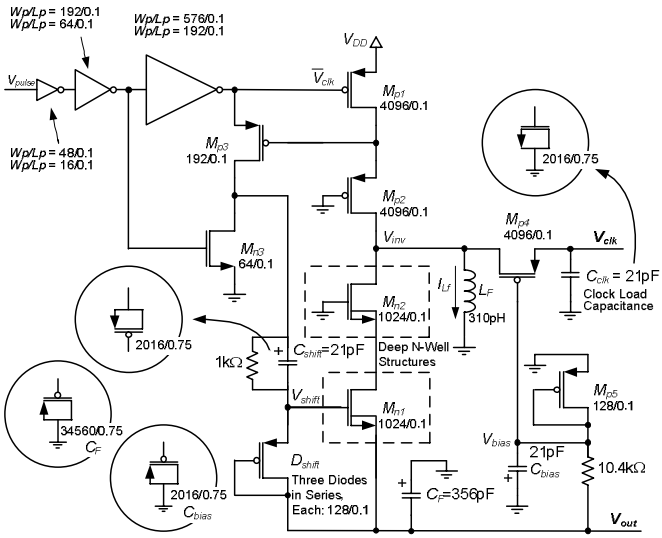

Figure 7. Fully-integrated clock driver and buck-boost converter.

- Mode 1: At the beginning of this mode, $\overline{V_{clk}}$ goes to zero and M_p turns on. The switch S_{clk} is closed and C_{clk} is charged. Consequently, voltage V_{DD} is across the inductor L_F and the inductor current increases linearly assuming a constant voltage across the inductor. At the same time, voltage of the capacitor C_{shift} from the previous Mode 2 is added to $\overline{V_{clk}}$ and V_{shift} reaches a lower value than V_{out} as diodes D_{shift} are reverse biased. Since C_{shift} is pre-charged to $V_{out} + 3V_{diode\_drop}$ in the previous Mode 2, the V_{gs} of M_n would be equal to $V_{shift} - V_{out} = -V_{DD} + 3V_{diode\_drop}$, which has a negative value and M_n turns off completely.
- Mode 2: As $\overline{V_{clk}}$ is high, M_p is off and M_n is on. At the same time, capacitor C_{shift} will be charged through diodes D_{shift} to a value of $V_{DD} - (V_{out} + 3V_{diode\_drop})$. Since the diodes are forward biased, the V_{gs} of M_n is equal to $3V_{diode\_drop}$, which is a positive value larger than the threshold voltage of M_n, thus ensuring it turns on completely. At this time, inductor current finds its way through M_n and will charge up the output capacitor C_F to a negative voltage value. The switch S_{clk} is closed at the beginning of Mode 2 to allow the inductor to discharge C_{clk}. However, when V_{inv} starts to go negative, the switch is opened to keep V_{clk} at zero.

The operating voltage available to the buck-boost converter is $V_{in} = D \times V_{DD}$. Ideally, the output voltage of the converter would be $V_{out} = \frac{-D}{1-D} \times V_{in} = \frac{-D^2}{1-D} \times V_{DD}$. Like the earlier boost converter, the output voltage can be regulated by modulating the duty cycle D.

3.3. Complete Circuit

A complete implementation of the integrated clock driver/ buck-boost converter is shown in Figure 7. Many of the changes are similar to those in the boost circuit, *e.g.*, M_{p3} and M_{n3} to delay the energy-wasting discharge of C_{clk}.

There are three implementation decisions in Figure 7 that warrant further discussion. First, transistors M_{p2} and M_{n2} are added to protect M_{p1} and M_{n1} from large voltage drops across them. Connecting the gates of transistors M_{p2} and M_{n2} to the ground will provide for automatic on-off timing and proper operation of the circuit.

Second, transistor M_{p4} acts as the switch to prevent V_{clk} from going negative. The gate of M_{p4} is connected to V_{bias}, which is set at the threshold voltage of PMOS transistor M_{p5}. When V_{inv} is positive, M_{p4} is on and provides the path for the inductor current to discharge C_{clk}. When V_{inv} falls below zero, M_{p4} turns off and nodes V_{inv} and V_{clk} are disengaged. Meanwhile, M_{n2} turns on and provides a path for the inductor current. In this design, V_{bias} is generated by a small DC current passing through the diode-connected PMOS transistor M_{p5}. To stabilize the voltage, capacitor C_{bias} is added to the node V_{bias}.

Third, the body terminals of all NMOS transistors need to be connected to their source or the most negative voltage in the system to prevent forward biasing of body-source intrinsic diodes. For M_{n1} and M_{n2}, their body is connected to the (non-ground) source node, so these transistors need to be isolated inside a deep n-well layout structure. All other transistor body terminals are connected to their sources. C_F and C_{bias} are implemented as PMOS transistors; NMOS transistors would require a deep n-well since V_{bias} and V_{out} are both negative.

Table 1. Specification summary of the integrated clock driver / converter circuits.

Converter type	Boost	Buck-Boost
Technology	90nm CMOS	90nm CMOS
Layout Area (mm^2)	0.26	0.20
Switching frequency, F_{sw} (GHz)	3	3
Inductor, L_F (pH)	310	310
Capacitor, C_F (pF)	378	356
Supply voltage, V_{in} (V)	1	1
Output voltage, V_{out} (V)	0.75 ~ 1.75	-0.5 ~ -1.4
Output voltage ripple	< 5%	< 5%
Nominal output current, I_{out} (mA)	50	50
Effective efficiency, η_{eff} (%) at nominal output	98 ~ 24	66 ~ 20
Duty cycle, D (%)	40 ~ 80	20 ~ 70

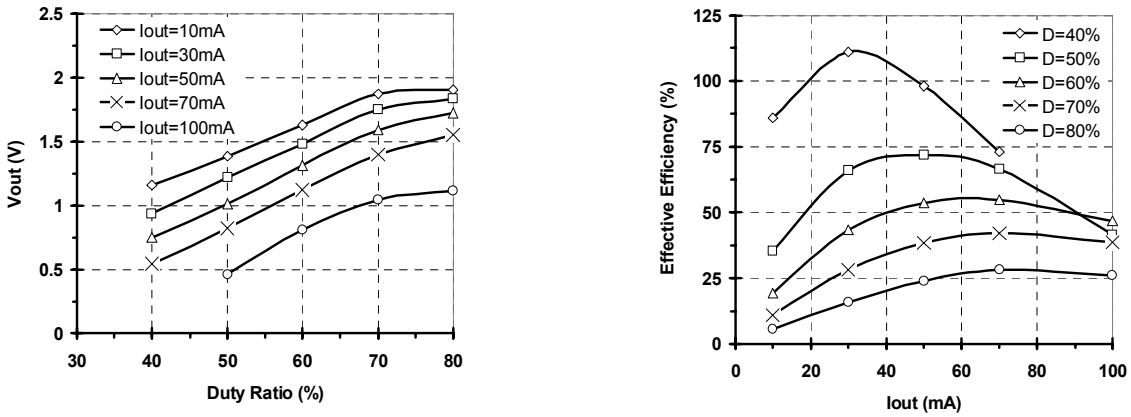

Figure 8. Simulation results of the boost converter in Figure 4.

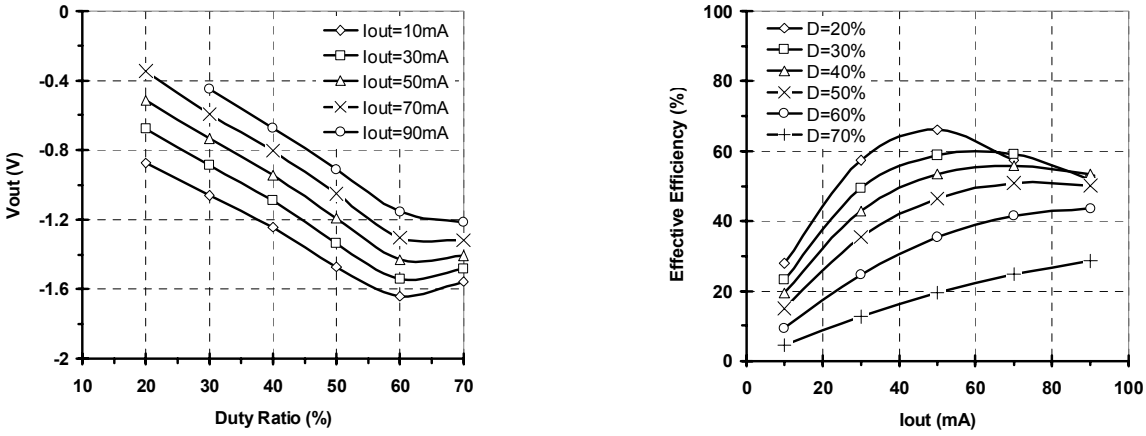

Figure 9. Simulation results of the buck-boost converter in Figure 7.

4. Simulation Results

To evaluate the performance of the combined clock-driver/converter circuits, a reference circuit containing only the tapered inverters to form a clock driver was designed using similar transistor sizes. All circuits were simulated in 90nm CMOS technology at F_{sw} = 3GHz. In all circuits, the transistors were all of low-threshold voltage type to facilitate operation at lower V_{DD} levels.

Figure 8 shows the output voltage and the effective efficiency of the boost converter at different duty cycles and output currents. We define the effective efficiency as $\eta_{eff} = \frac{P_{out1}}{P_{in1} - P_{in2}} \times 100$, where P_{in1} and P_{in2} are the power consumption of the new combined driver/converter circuit and the reference clock circuit, respectively, and P_{out1} is the output power of the converter. For the reference circuit, simulations determined P_{in2} was 100mW. High power consumption is primarily due to a large C_{clk}.

Figure 8 shows the output voltage and effective efficiency of the integrated boost-converter circuit. The clock duty cycle, D, and output current are both varied. The output voltage increases as D is increased. At the same time, the effective efficiency decreases because the

amount of energy that is being recycled (available in the clock capacitor) is constant while the amount of energy that is delivered to the output is increased (because of the increase in the output voltage). For example, at 50mA output current, the output voltage changes from 0.75 to 1.75V by varying the duty cycle from 40 to 80%. The corresponding effective efficiency ranges from 98 down to 24%. By varying the duty cycle, the output current corresponding to the highest value of effective efficiency is changed and maximum effective efficiency of 111% is achieved at D = 40% with I_{out} = 30mA and V_{out} = 0.95V. As mentioned earlier, achieving an effective efficiency above 100% is definitive proof that "free energy" is being recovered from the clock tree.

Similarly, Figure 9 shows the output voltage and effective efficiency of the integrated buck-boost converter circuit. In this case, at 50mA output current, the output voltage changes from -0.5 to -1.4V when varying the duty cycle from 20 to 70%. The corresponding effective efficiency ranges from 66 down to 20%. The lower efficiency is a result of more transistors in the main current path. While the effective efficiency result does not prove that clock energy is being recovered like the boost circuit, looking at Figure 7 reveals that during recycling time (when M_{n2} is off because $V_{inv} \geq -V_{th\_nmos}$) there is no path from C_{clk} to ground except through L_F. This means charge in C_{clk} is being recycled to current in L_F.

Table 1 provides a summary of the specifications and results for the two converter circuits. We have performed a layout of both circuits. To save time, we re-used many circuit components between the two circuits. The reference clock driver circuit (drivers and capacitive load) is 0.03mm^2. The boost and buck-boost circuits in this paper require 0.26mm^2 and 0.2mm^2 area, respectively, including the on-chip inductor and filter capacitor. Although the buck-boost circuit has more components, it was implemented in less area than the boost circuit because of additional layout effort.

In contrast, [6] implemented a fully on-chip buck converter in 0.18µm SiGe RFBiCMOS technology was 65% efficient. It also used an area of 27mm^2 to fit the large passive components. The circuits in this paper achieve an effective efficiency in the same range using only 1/100th of the area.

5. Conclusions

In this paper, the idea of energy recovery from a high-speed clock load in digital circuits is investigated by exploring the integration of two DC-DC converter topologies, namely boost and buck-boost, with a high-speed clock driver. These converters recycle the energy (in the form of charge stored) from the clock capacitance. Several characteristics make integration of these two different circuits very promising: the multi-GHz switching frequency of the clock allows the use of very small passive components in the power converters, enabling full on-chip integration; the clock driver and power converter can share the tapered buffer chain, which can be optimized for low power; and energy wasted during clock discharge can be recovered using the converter.

Simulation results of the circuits in a standard 90nm technology show an effective efficiency up to 111%. The only way for this to exceed 100% is to recover the "free energy" from the clock that would otherwise be wasted. Comparing the two converters, the boost converter has a simpler implementation and results in higher efficiencies than the buck-boost converter.

These results are very promising, but they are only scratching the surface of a brand new area. A significant amount of work remains to be done to improve these circuits. No doubt the circuit implementations are sub-optimal, and higher efficiencies are possible. Only 'open loop' power measurements were done; a full circuit including a controller for power regulation is needed. Although simulations show good-quality, quasi-square clock waveforms, there is concern that clock jitter may be increased as a result of the power converter. Also, pulse width modulation adds its own problems: it increases jitter, and high-speed designers often use more complex clocking schemes than single-edge-triggered clocking.

References

[1] J. Friedrich, B. McCredie, N. James, B. Huott, B. Curran, E. Fluhr, G. Mittal, E. Chan, Y. Chan, D. Plass, S. Chu, H. Le, L. Clark, J. Ripley, S. Taylor, J. Dilullo and M. Lanzerotti, "Design of the Power6 Microprocessor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2007, pp. 96-97.

[2] M. Stan and W. Burleson, "Low-power CMOS clock drivers," in *Proc. ACM/IEEE Int. Workshop on Timing Issues in the Specification and Synthesis of Digital Systems (TAU Workshop),* 1995, pp. 149-156.

[3] S. C. Chan, K. L. Shepard, and P. J. Restle, "Uniform-Phase Uniform-Amplitude Resonant-Load Global Clock Distributions," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, Jan. 2005, pp. 102-109.

[4] G. Lemieux, M. Alimadadi, S. Sheikhaei, P. Palmer, S. Mirabbasi, "SoC Energy Savings = Reduce+Reuse+Recycle", to appear at *IEEE Canadian Conference on ECE*, Niagara Falls, ON, May 2008.

[5] J. W. Tschanz, J. T. Kao, S. G. Narendra, R. Nair, D. A. Antoniadis, A. P. Chandrakasan and V. De, "Adaptive Body Bias for Reducing Impacts of Die-to-Die and Within-Die Parameter Variations on Microprocessor Frequency and Leakage," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, Nov. 2002, pp. 1396-1402.

[6] S. Abedinpour, B. Bakkaloglu, and S. Kiaei, "A Multi-Stage Interleaved Synchronous Buck Converter with Integrated Output Filter in a 0.18µm SiGe Process," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2006, pp. 356–357.

[7] 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," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2007, pp. 532-533.