

A 4 GHz Non-Resonant Clock Driver With Inductor-Assisted Energy Return to Power Grid

Mehdi Alimadadi, Samad Sheikhaei, *Student Member, IEEE*, Guy Lemieux, *Senior Member, IEEE*, Shahriar Mirabbasi, *Member, IEEE*, William Dunford, *Senior Member, IEEE*, and Patrick Palmer, *Member, IEEE*

Abstract—Power consumption of a multi-GHz local clock driver is reduced by returning energy stored in the clock-tree load capacitance back to the on-chip power-distribution grid. We call this type of return *energy recycling*. To achieve a nearly square clock waveform, the energy is transferred in a non-resonant way using an on-chip inductor in a configuration resembling a full-bridge DC-DC converter. A zero-voltage switching technique is implemented in the clock driver to reduce dynamic power loss associated with the high switching frequencies. A prototype implemented in 90 nm CMOS shows a power savings of 35% at 4 GHz. The area needed for the inductor in this new clock driver is about 6% of a local clock region.

Index Terms—Charge recycling, energy recycling, energy recovery, full-bridge converter, low-power clock driver, multi-GHz clock, switching DC-DC converter.

I. INTRODUCTION

POWER consumption of digital circuits, particularly in high-performance processors, has recently increased rapidly. Coping with this increase is an important issue, not only in battery-powered applications, but also in other designs because of packaging, cooling and operating costs [1]–[3]. In these chips, the clock itself consumes a significant amount of power. For example, the 5+ GHz clock network in the IBM POWER6 processor consumes 22% (~ 22 W) of the total power and is second only to the total leakage power [4]. As another example, the Intel Itanium 2 processor uses a system clock around 2 GHz which consumes 25% (~ 25 W) of the total power [5]. Clearly, it is important to reduce clock power consumption as much as possible.

Various methods of clock energy reduction by charge recycling have been reported, including: using an additional charge reservoir with a resonant-clock scheme [6]; exchanging charge between two differential clock networks [7]; and sending charge

Manuscript received June 26, 2009; revised September 16, 2009; accepted October 31, 2009. Date of publication January 29, 2010; date of current version August 11, 2010. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC). CMC Microsystems provided CAD tools and chip fabrication. This paper was recommended by Associate Editor Y. Massoud.

M. Alimadadi, S. Sheikhaei, G. Lemieux, S. Mirabbasi, W. Dunford are with the Electrical and Computer Engineering Department, University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: mehdi@ece.ubc.ca; samad@ece.ubc.ca; lemieux@ece.ubc.ca; shahriar@ece.ubc.ca; wgd@ece.ubc.ca).

P. Palmer is with Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, U.K. (email: prp@eng.cam.ac.uk).

Digital Object Identifier 10.1109/TCSI.2009.2037850

to a switching power converter [8], [9]. This latter approach is described as energy recycling as the recycled charge is transferred to a properly regulated voltage supply. More recently, [10] applies resonant clocking to digital filters using both single and two-phase clocking schemes.

To perform recycling, these schemes require a secondary capacitor or a second clock network, or produce an auxiliary voltage supply. The resonant schemes produce a nearly sinusoidal clock waveform, limiting their use. The power converter schemes vary the clock duty cycle to regulate the output voltage. This work avoids these drawbacks by recycling this clock-stored energy back to the primary on-chip power grid. This is done in a non-resonant fashion using an on-chip inductor. In this way, clock quality is preserved: the waveform is nearly square, has sharp edges, and retains its duty ratio.

II. BACKGROUND

Modern high-performance processor designs such as the Pentium 4 and NEC SX-9 distribute a global multi-GHz clock across the chip to multiple local clock domains [11], [12]. Each local clock must drive a load and requires a buffering scheme such as the one presented in this paper. Clock gating, achieved by a *gater* circuit positioned just before the local clock driver, eliminates unnecessary activity on segments of the clock network to save power [12]. Three generations of the DEC Alpha exemplify this shift from a single centralized, global clock driver towards distributed, local, gated clocks [13].

Together, each gater and local clock driver covers up to 2 mm of clock wire length and hundreds of latches [14], [15]. The final drive stage consists of several parallel inverters, often connected to a mini-mesh at several points, covering up to 1 to 2 mm² of area (but often smaller). This approach reduces *RC* parasitics while driving the clock to the final load, which is essential for keeping low skew. In this paper, the term “mesh” refers to this local mini-mesh located after the final driver rather than a global mesh across the entire chip.

The design introduced here is a new local clock driver, consisting of an inverter chain and an inductor placed at the final load after the gater. This circuit would be copied in several locations across the chip with a separate clock driver per local, gated-clock region. Since we are dealing with providing increased drive strength, the gater circuitry itself is not included in this design. However, it cannot be ignored; as will be shown, stopping the clock introduces static power loss in the inductor, for which a solution is provided.

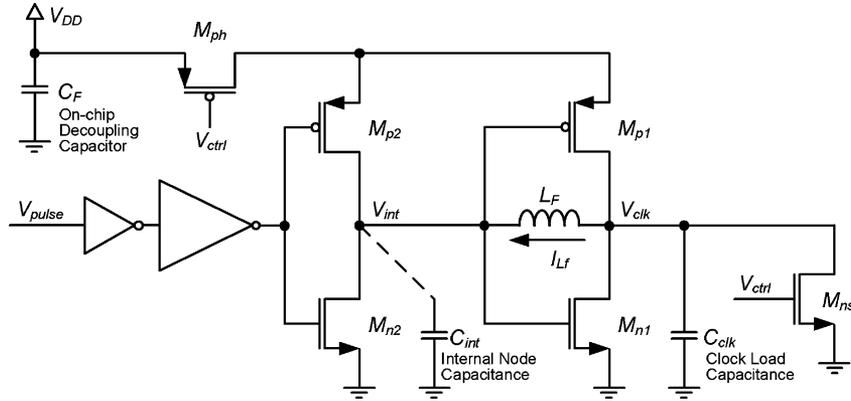


Fig. 1. Simplified low-power clock-driver.

III. CIRCUIT DESIGN

A. Simplified Circuit

A simplified version of the proposed low-power clock driver circuit is shown in Fig. 1. This circuit incorporates an inductor L_F at the clock node, but unlike resonant clocking schemes, the inductor appears in the driver side not the load side. Capacitors C_{clk} and C_{int} are the sum of wiring and transistor capacitances. Assuming a fan-out of four as the inverter taper factor, C_{int} is roughly one-fourth of C_{clk} . Capacitor C_F represents the intrinsic power-grid capacitance and the on-chip decoupling capacitances commonly added to high-performance digital designs.

The inductor L_F is used to discharge C_{clk} and transfer the stored clock energy instead of discharging it to ground. Some of this transferred energy returns to the power grid through M_{p2} . This effectively reduces the power consumption of the clock driver itself. However, we term the return of energy to the power supply as *energy recycling* because it is made available to any other circuit (not only the clock driver itself).

In many situations, it is desirable to “stop the clock” to save power by gating the incoming clock signal. With a stopped clock, the inductor would be continuously conducting and dissipate significant static power. To solve this problem, power gating with the header transistor M_{ph} can disconnect the power supply from the driver. This also reduces standby leakage [16]. However, it introduces a new concern: the LC components can oscillate and introduce additional unwanted clock transitions until the stored energy in the system is dissipated. To address this issue, an extra nMOS transistor, M_{ns} , is added in parallel to C_{clk} to provide a discharge path for the clock, keeping V_{clk} at zero to immediately dampen any unwanted ringing. Transistors M_{ph} and M_{ns} can share the same gating signal as denoted by V_{ctrl} in Fig. 1. The size of the header transistor should be large enough to minimize the effect of its on-state series resistance on circuit operation. Compared to a regular clock driver, additional energy is required to activate the clock gate via V_{ctrl} due to the size of transistors M_{ph} and M_{ns} . However, stopping the clock is usually an infrequent operation, so a net savings can be readily achieved.

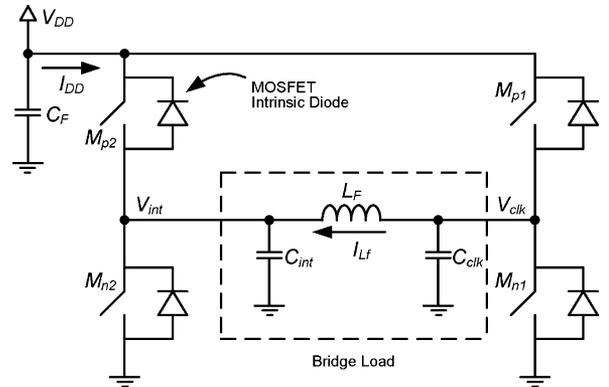


Fig. 2. Typical full-bridge DC-DC converter.

B. Full-Bridge Converter

The circuit in Fig. 1 resembles a full-bridge DC-DC converter working in boost mode, where M_{p1} , M_{n1} , M_{p2} and M_{n2} are the bridge switches, and the combination of C_{int} , C_{clk} and L_F are the bridge load.

A generic full-bridge converter is shown in Fig. 2. Here, for the purpose of clarity, the gating transistors M_{ph} and M_{ns} (which are included in Fig. 1) are omitted. The input is a fixed DC voltage, but the magnitude and polarity of the bridge load voltage ($V_{clk} - V_{int}$) can be adjusted by pulse-width modulating the gating signals. Switches (M_{p2} , M_{n1}) and (M_{n2} , M_{p1}) are usually treated as two pairs. Because of the inductive load, depending on the direction of the load voltage and current, the load may consume or return power. Since the inductor current cannot change abruptly, the load current does not become discontinuous. Instead, the input current to the bridge can change direction via the switches, so it is important that the DC source has low internal impedance. A bigger C_F would better facilitate this requirement.

If the bridge stays in a particular state long enough, the energy stored in the inductor would be large enough for charging/discharging the load capacitors. In practice, non-ideality of M_{n2} and M_{n1} results in their slow turn-on, providing the time needed for the inductor current to discharge C_{int} and C_{clk} . Similarly,

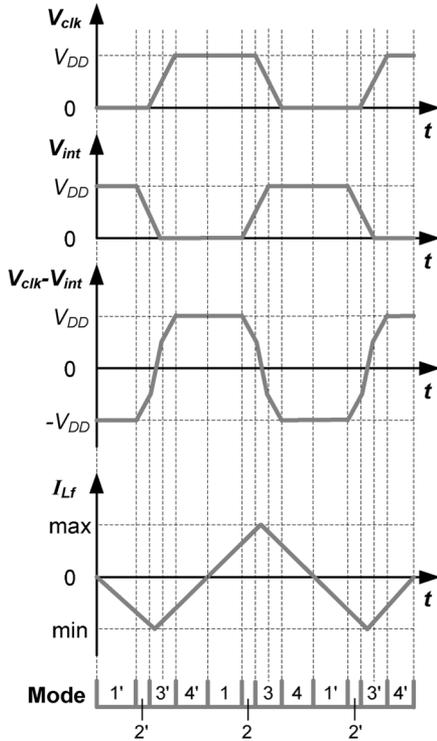


Fig. 3. Idealized timing diagram of Figs. 1 and 2.

non-ideality of M_{p2} and M_{p1} gives the inductor time to charge those capacitors.

In the simplified design of Fig. 1, the CMOS inverter propagation delay (from V_{int} to V_{clk}) helps provide time for the inductor to charge/discharge capacitor C_{clk} . This is observed, for example, after M_{p2} turns on and raises V_{int} with the assistance of the inductor before V_{clk} falls due to the turn-on of M_{n1} . The ZVS-enhanced circuit, which will be discussed later in Section III.D, utilizes zero-voltage switching to provide an even longer delay that is dynamically adjusted.

C. Modes of Operation

Operation of the circuit in Figs. 1 and 2 can be explained using the idealized timing diagram shown in Fig. 3. There are eight modes.

- Mode 1: M_{p1} and M_{n2} are on. C_{clk} is already charged up and V_{clk} is high. Inductor current is positive and is increasing linearly.
- Mode 2: M_{n2} is turned off and M_{p2} is turned on. V_{int} increases.
- Mode 3: M_{n1} is turned on and M_{p1} is turned off. V_{clk} decreases. For a short time the inductor current continues to rise. When M_{p1} is off, the inductor takes energy from C_{clk} rather than V_{DD} and helps V_{clk} to fall rapidly. The inductor will first transfer energy to C_{int} , helping M_{p2} to increase V_{int} quickly, and then transfer energy to the on-chip power grid through M_{p2} . Inductor current peaks when $V_{int} = V_{clk}$, i.e., when the voltage across L_F is zero. The inductor current starts to decrease. V_{clk} and V_{int} reach low and high values, respectively.

- Mode 4: M_{p2} and M_{n1} are on. C_{clk} is already discharged and V_{clk} is low. Inductor current is positive and is decreasing linearly.
- Modes 1'–4': With the direction of the inductor current reversed, Modes 1–4 repeat in the opposite sense to help charge capacitor C_{clk} from the stored energy in C_{int} and L_F . When C_{int} is discharged, M_{n2} keeps V_{int} at zero, providing the current path for L_F to charge up C_{clk} .

In the above discussion, whenever the absolute value of the inductor current is decreasing, the energy stored in the inductor is being delivered to another element of the circuit. Here, the destination of the energy can be C_F , C_{clk} , or C_{int} . Since the final clock load capacitance has about $4 \times$ more energy than the stage before it, the primary objective of the circuit is to recycle the energy stored in the clock node. For example, this occurs during Mode 3 when the C_{clk} charge is returned to the power grid via the inductor.

During Mode 3', the inductor also reduces the amount of energy consumed by helping to precharge C_{clk} from the energy stored in itself and C_{int} . However, as C_{int} is smaller than C_{clk} , there is no opportunity to recycle energy to the power grid in this mode. In other words, because of the asymmetry of the bridge legs, energy return occurs only in one direction; in the other direction, the energy transfer can only partially pre-charge the clock load capacitance.

Additional energy recycling occurs when L_F magnetic energy is returned to the power grid during Modes 4 and 4'.

D. ZVS-Enhanced Circuit

Ideally, all of the energy stored in C_{clk} should be recovered (by moving it to C_{int} and/or C_F) rather than being wasted by discharging C_{clk} to ground. Thus, to maximize the energy savings, the turn-on of M_{n1} should be delayed. This is achieved in the ZVS-enhanced low-power clock driver circuit of Fig. 4 with the addition of transistors M_{n3} and M_{p3} . Furthermore, M_{n3} and M_{p3} also delay the turn-on of M_{p1} , allowing C_{clk} to be precharged by the inductor. This achieves zero-voltage switching in the final drive stage and reduces switching power loss. ZVS for M_{n1} is explained in detail in [8].

The main benefit of implementing ZVS for M_{n1} is that C_{clk} will not be shorted to ground when it has a voltage across it. During the ZVS dead-time, the charge is removed (recovered) by the inductor current and consequently V_{clk} is reduced to zero. After this, M_{n1} is turned on to provide a low-loss path for current and also to keep V_{clk} around zero. If M_{n1} is not turned on, the inductor current would turn on the intrinsic body-drain diode of M_{n1} . The resulting voltage drop across this diode would contribute to the overall power consumption of the system. In the charging phase of C_{clk} , ZVS for M_{p1} causes C_{clk} to be charged initially through the inductor L_F .

IV. EQUIVALENT CIRCUIT MODELING

Assuming that the total capacitance of the reference clock driver shown in Fig. 4 is denoted by

$C_{\text{clk-chain}}$, the power consumption of the reference circuit can be estimated as $P_{\text{in2}} \cong C_{\text{clk-chain}} V_{\text{DD}}^2 f_{\text{clk}} = C_{\text{clk}}(1 + (1/4) + (1/16) + \dots) V_{\text{DD}}^2 f_{\text{clk}}$, thus

$$P_{\text{in2}} \cong \frac{4}{3} C_{\text{clk}} V_{\text{DD}}^2 f_{\text{clk}}. \quad (1)$$

To simplify the analysis of the low-power clock driver circuit in Fig. 4, it is assumed that the absolute value of the voltage across the inductor is always constant, i.e., it is assumed that modes 2, 3, 2' and 3' are very short. Therefore, I_{LF} has a triangular waveform with values limited to $-I_{\text{LF,max}}$ and $+I_{\text{LF,max}}$. If V_{drop} is the resistive voltage drop across the transistors and the inductor, then using $V_L = L_F(\Delta I_{\text{LF}})/(\Delta t)$ we have $V_{\text{DD}} - V_{\text{drop}} = L_F(2I_{\text{LF,max}})/(T_{\text{clk}}/2)$ or $I_{\text{LF,max}} = (1)/(4L_F f_{\text{clk}})(V_{\text{DD}} - V_{\text{drop}})$. Here, V_{drop} can be estimated as $V_{\text{drop}} \cong R_{\text{eq}} I_{\text{LF,rms}}$ and for the triangular L_F current $I_{\text{LF,rms}} = (I_{\text{LF,max}})/(\sqrt{3})$. Thus, $I_{\text{LF,max}} = (V_{\text{DD}})/(4L_F f_{\text{clk}} + 0.58R_{\text{eq}})$.

In the low-power clock driver circuit, there are two paths for charging and discharging C_{clk} . One path is through the inductor L_F and the other path is through the transistors M_{p1} and M_{n1} . The power consumption of the circuit is the sum of the power consumption in those two different paths

$$P_{\text{in1}} = R_{\text{eq}} I_{\text{LF,rms}}^2 + \alpha C_{\text{clk-chain}} V_{\text{DD}}^2 f_{\text{clk}}. \quad (2)$$

The first term is related to the current path through L_F and the second term is related to the current path through M_{p1} and M_{n1} transistors. Here, R_{eq} is the equivalent resistance of the inductor and the transistors in the current path, and α is a number between 0 and 1, which relates to the share of the transistors M_{p1} and M_{n1} in charging/discharging of C_{clk} . A higher inductor current would result in a faster charge/discharge of C_{clk} and as a result M_{n1} and M_{p1} would have a smaller share in charging/discharging of C_{clk} , resulting in a lower value for α . Other parameters that have an effect on α are the size and threshold voltage of the ZVS transistors which determine the turn-on time for M_{n1} and M_{p1} .

Also, P_{in1} can be rewritten as:

$$P_{\text{in1}} = \frac{R_{\text{eq}}}{3} \left(\frac{V_{\text{DD}}}{4L_F f_{\text{clk}} + 0.58R_{\text{eq}}} \right)^2 + \frac{4}{3} \alpha C_{\text{clk}} V_{\text{DD}}^2 f_{\text{clk}} \quad (3)$$

Comparing P_{in1} with P_{in2} , L_F is a new design freedom that can be used to reduce P_{in1} that is unavailable in the P_{in2} circuit. When designing a clock driver, values for L_F and α can be optimized by full circuit simulation. However, note that (2) and (3) are not valid at lower frequencies as the increased inductor current would increase the resistive voltage drop across it and the inductor current would no longer be a triangular waveform.

V. SIMULATION AND MEASUREMENT RESULTS

A. Simulation Results

By using a sufficiently large header transistor M_{ph} , the voltage drop across it would be negligible. For example,

simulation results show that a 5800/0.1 μm header transistor would have 25 mV drop across it. This drop can be reduced further by using a larger transistor at the cost of more area and increased energy to stop the clock. Simulations also show that a 1024/0.1 μm shorting transistor, M_{ns} , effectively dampens the ringing of a stopped clock. The capacitance of the shorting transistor adds capacitance to the clock node. Due to a limited number of probe pins, the V_{ctrl} input and those two transistors were not implemented on the test chip. They were also omitted in the simulations below to keep the simulated and implemented circuits the same.

Simulation results of the implemented low-power clock driver operating at 4 GHz are shown in Fig. 5. As shown in the figure, the proposed technique preserves the sharp edges of the clock in the presence of the inductor. Compared to the reference clock driver implemented in the same process, the slope of the rising clock edge in the new circuit is similar, although the falling slope is slightly slower because ZVS transistors M_{n3} and M_{p3} are in the path of charging the V_{intn} node. Thus, M_{n1} turns on slightly slower and V_{clk} has a slower falling edge. This has an effect on the observed duty cycle of the clock. Fig. 5 also shows that there is a time delay between the rising/falling edges of the three variants of the clock driver. Device variation in downstream circuits such as flip-flops may be sensitive to the falling edge rate and/or variation in observed duty cycle. In this case, the simplified circuit should be used instead of the proposed ZVS-enhanced circuit. The variation of delay in the local clock drivers can affect the hold times at receivers as well.

To investigate the effect of ZVS transistors on circuit operation, M_{p2} and M_{n1} drain currents are plotted in Figs. 6 and 7, respectively. In Fig. 6, a positive M_{p2} drain current means that C_{clk} charge is being returned to V_{DD} . For the reference circuit, the curve is completely negative as there is no energy recycling taking place. The area integrated under this curve for one clock cycle is -7.9 pC, indicating a net consumption of charge. For the simplified and ZVS-enhanced circuits, the M_{p2} drain current is positive during the falling edge of the clock, Mode 3, when energy recycling takes place. The corresponding areas are -0.7 and 1.3 pC, respectively, indicating that consumption is greatly reduced. In fact, the positive value for the ZVS-enhanced circuit means that net energy is actually returned to the power grid through M_{p2} .

In Fig. 7, a positive M_{n1} drain current means that C_{clk} is being discharged to ground. This should be as small as possible for energy recycling purposes. For the reference circuit, the area integrated under this curve for one clock cycle is 24 pC. For the simplified and ZVS-enhanced circuits, it is 19 and 16 pC, respectively. Because the M_{n1} channel current losses are resistive, the energy savings is larger than implied by these integrals. Also, during the rising edge of the clock, Mode 3', the inductor delivers some energy to C_{clk} , but M_{p1} must provide additional charge. Hence, M_{p1} drain current looks similar to that of M_{n1} rather than M_{p2} . This shows the ZVS transistors are able to reduce energy wasted by reducing the discharge of C_{clk} through M_{n1} .

In Fig. 8, the low-power clock driver is simulated at different switching frequencies to show that power savings improves as

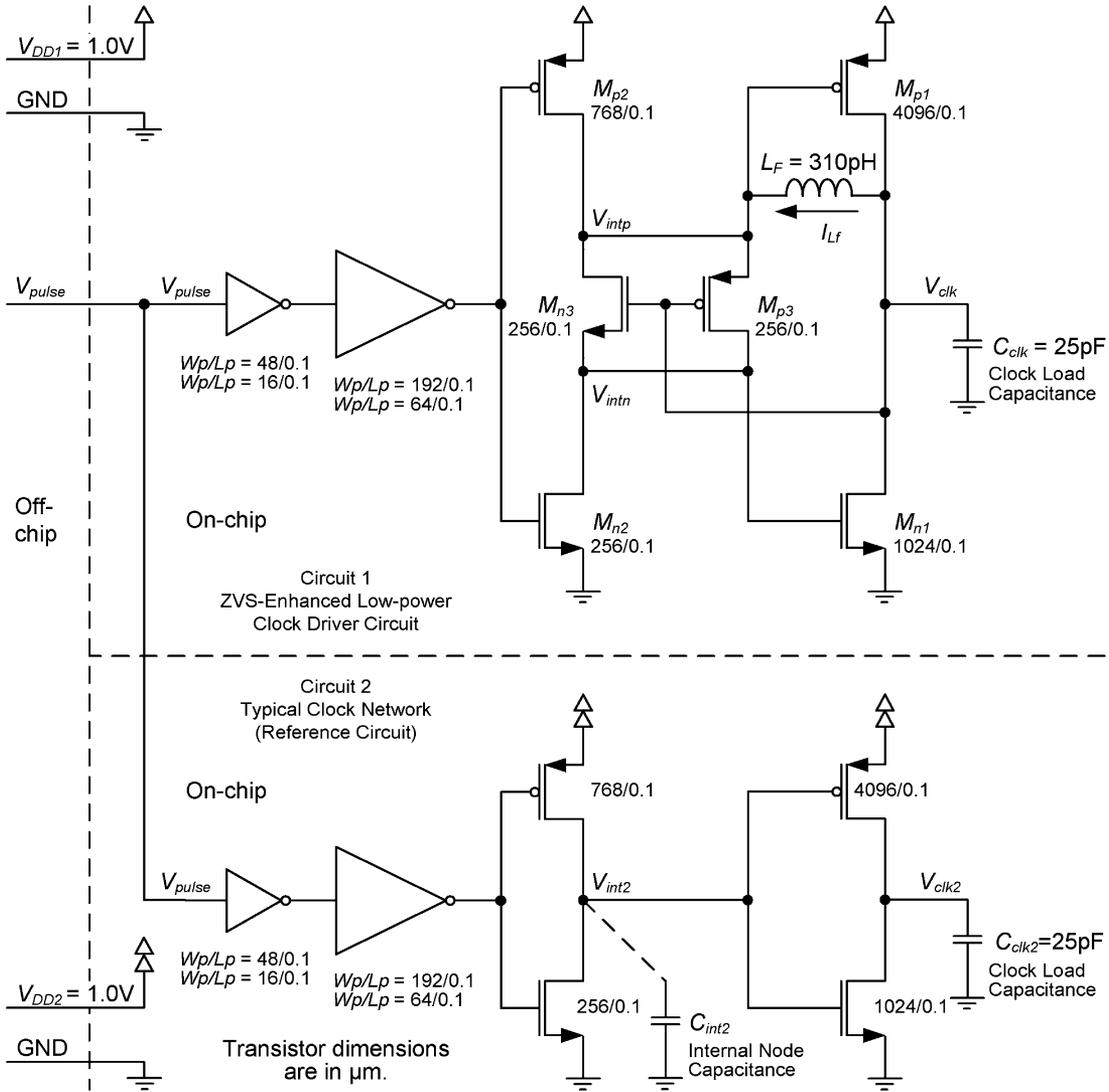


Fig. 4. Circuit diagram of the implemented prototype.

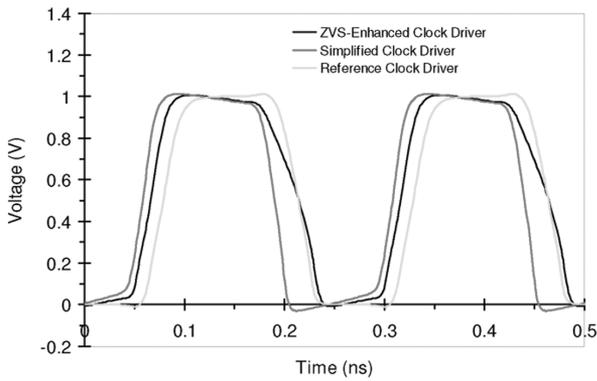


Fig. 5. Simulated clock waveforms.

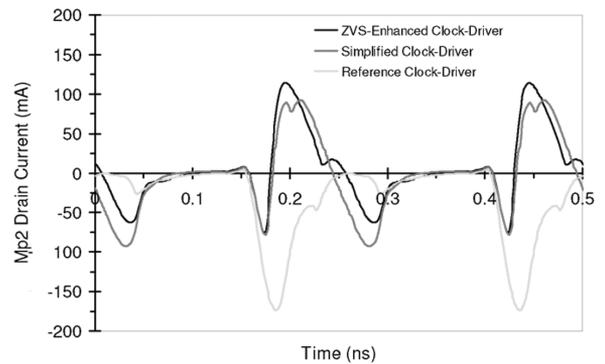


Fig. 6. Simulated M_{p2} drain currents.

the clock frequency is increased towards 4 GHz. The simplified circuit does not perform as well as the ZVS-enhanced circuit because the ZVS transistors M_{n3} and M_{p3} in Fig. 4 assist in energy return to the power grid. Also, simulation results at 4 GHz show a percentage power savings equal to $(P_{in2} - P_{in1})/P_{in2} = 37\%$.

Here, the power consumption of the ZVS-enhanced and reference circuits are $P_{in1} = 86$ mW and $P_{in2} = 136$ mW, respectively.

To evaluate the effect of inductor value on power consumption, the ZVS-enhanced circuit is simulated with different inductor values by varying a factor K such that $L_F = K \times 310$ pH.

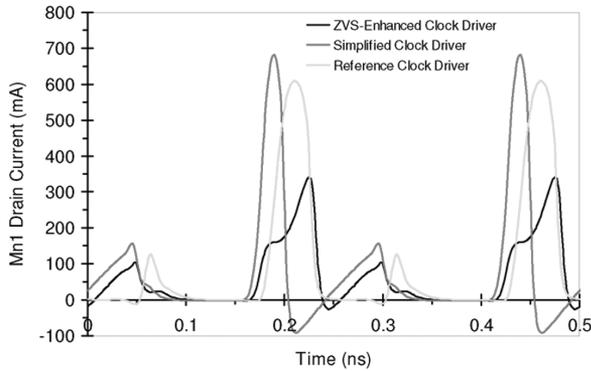
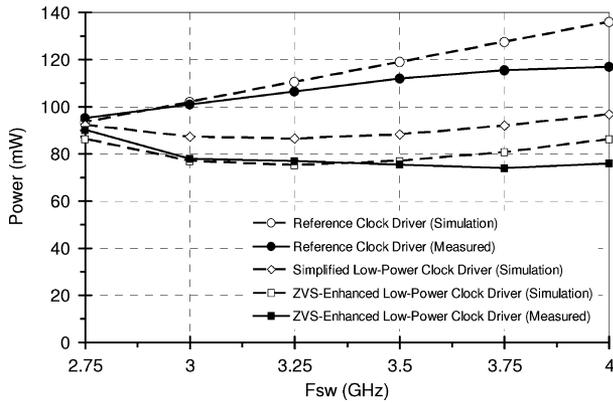
Fig. 7. Simulated M_{n1} drain currents.

Fig. 8. Test and simulation results of the circuits in Fig. 4.

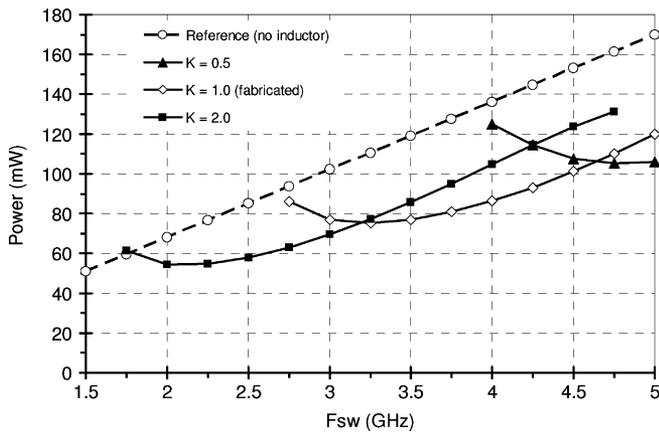
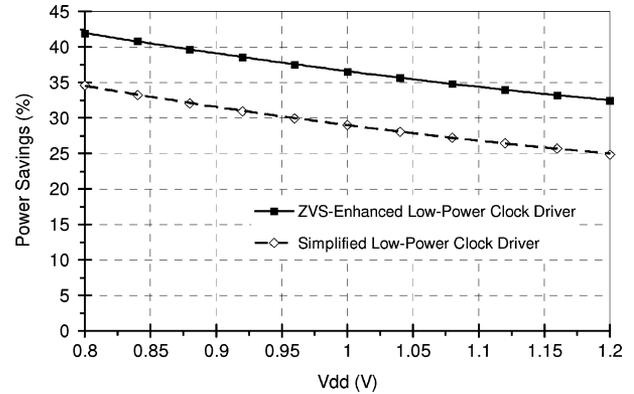
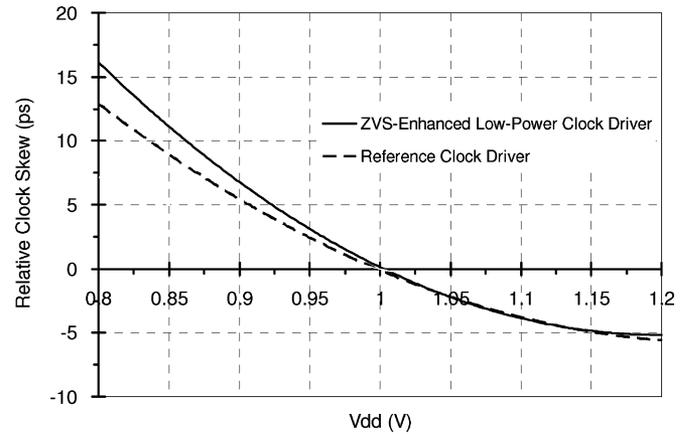
Fig. 9. Effect of changing inductor value on power savings ($L_F = K \times 310$ pH).

Fig. 9 shows the results and suggests an optimum inductor value is needed for different frequency ranges. For example, at $K = 1$, minimum power consumption is achieved over the clock frequency range of 3 to 4 GHz. This value of inductance corresponds to the fabricated prototype.

Also, the effect of V_{DD} variation on power savings is simulated in Fig. 10. It shows the percentage power savings slightly decreases as V_{DD} increases. At higher V_{DD} values, more energy is available to be recycled from the clock load capacitor, however because of higher current levels, the energy dissipated

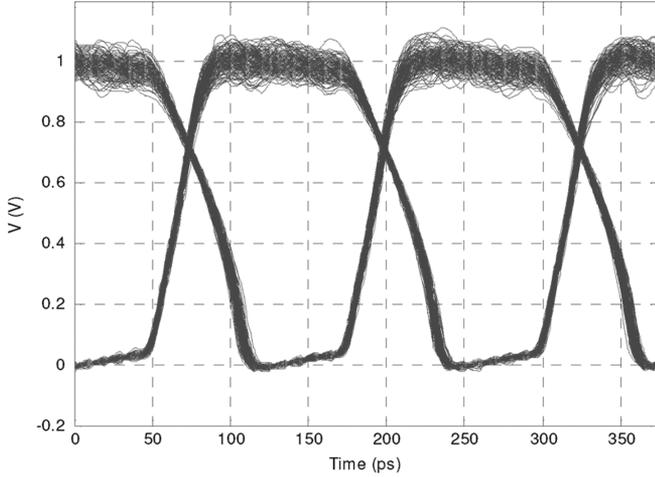
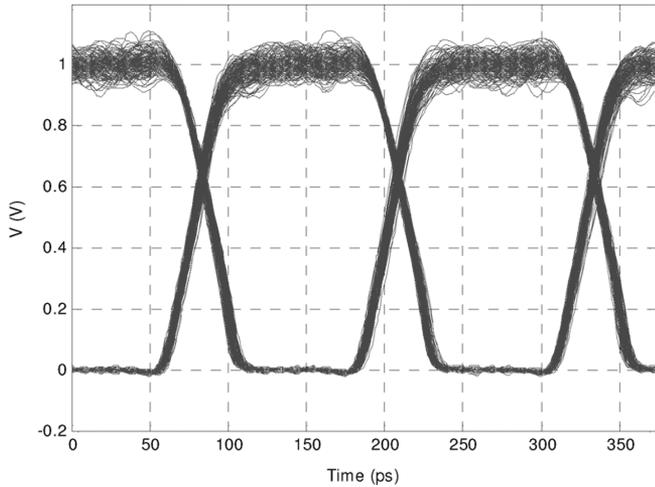
Fig. 10. Effect of V_{DD} variation on power savings at 4 GHz.Fig. 11. Effect of V_{DD} variation on clock skew at 4 GHz (relative to $V_{DD} = 1$ V).

in the series resistance of the circuit reduces the percentage of power savings.

To evaluate the sensitivity of clock latency (skew) to V_{DD} , the supply voltage to the circuits is slowly increased from 0.8 V to 1.2 V, in 100 ns total time. The trend line of the clock rising edge time, relative to the reference point of $V_{DD} = 1$ V, is shown in Fig. 11. That figure shows that the ZVS-enhanced circuit and the reference clock behave somewhat similar in terms of clock skew, when the supply voltage is varied.

In an actual processor die, because of resistance and inductance of the power distribution network, noise may be coupled to V_{DD} . Although on-chip decoupling capacitances are used to circumvent this, some residual noise still exists. The effect of V_{DD} noise on the clock waveform is studied here by simulation. In these simulations, additive white Gaussian noise (AWGN) with standard deviation of 0.05 V is added to a V_{DD} of 1 V. The eye diagram for the ZVS-enhanced and reference clock circuits are shown in Figs. 12 and 13, respectively, for a 4 GHz clock (a clock period of 250 ps). The height of the eye shows the effect of supply noise on the noise margins of the clock voltage levels, and the width of the eye shows the effect of supply noise on clock jitter [17].

Table I summarizes the simulated peak-to-peak and RMS clock jitter values for the three circuits. The results suggest that adding the inductor has reduced the jitter due to filtering of the

Fig. 12. Effect of V_{DD} noise with $\sigma = 0.05$ V on ZVS-enhanced clock driver.Fig. 13. Effect of V_{DD} noise with $\sigma = 0.05$ V on reference clock driver.TABLE I
SIMULATED CLOCK JITTER VALUES

Circuit Configuration	Peak-to-Peak Jitter (ps)	RMS Jitter (ps)
Simplified Low-Power Clock Driver	6.6	1.1
ZVS-Enhanced Low-Power Clock Driver	8.1	1.4
Reference Clock Driver	17.6	2.2

noise current. The slight increase in jitter of the ZVS-enhanced driver compared to the simplified driver can be attributed to the extra ZVS transistors.

B. Implementation

As a proof of concept, the two circuits in Fig. 4 have been fabricated in a 1P7M2T 90 nm CMOS process using low- V_t transistors to facilitate operation at lower V_{DD} levels. Although low- V_t devices dissipate more leakage power than high- V_t devices, clock drivers usually need to be fast. While the clock is

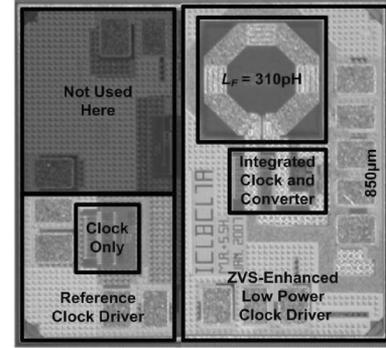


Fig. 14. Chip micrograph.

stopped, leakage can be reduced using the power gating transistor M_{ph} in Fig. 1. Across multiple instances of the circuit, there is likely very little variation in the threshold voltage of the ZVS transistors M_{n3} and M_{p3} , as they are very large transistors and the effect of a small number of dopant atoms, dopant atom locations, and local densities within a channel have negligible effect compared to their effect in a minimum-size transistor.

The coreless inductor L_F is made with a single loop using the four top metal and one extra aluminum (ALUCAP) layers in parallel. Placing a metal patterned ground shield (PGS) in metal layer 1 next to the substrate reduces the loss due to eddy currents as well as substrate noise [18]. By using strings of ground substrate contacts, induced current in the substrate is shorted to the system ground [19]. The value of inductance is extracted using ASITIC [20]. Its value is 310 pH, at 4 GHz, with lumped π model capacitances of 210 fF and a Q -factor (quality factor) of 22 at a resonant frequency around 21 GHz. A series resistance of 0.2Ω is also extracted. Although very wide metal with multiple layers in parallel is used to reduce the series resistance of the inductor as much as possible, the circuit performance also depends on the series resistance of the transistors in the bridge. In general, the performance of the circuit improves as the series resistance of the bridge transistors is lowered.

In the chip, the total capacitance connected to node V_{clk} (shown as C_{clk} in Fig. 4) is 25 pF. The amount of capacitance due to flip-flops, representing a fanout-of-four load to the driver, is 21 pF. This is implemented as the gate capacitance of 2016/0.75 μm nMOS transistor array. All transistor bodies are connected to their sources, except for M_{n3} whose body is connected to ground. This prevents forward biasing of body-drain intrinsic diode and avoids the need for isolation with a deep n-well structure.

The chip micrograph is shown in Fig. 14. The inductor area is 0.1 mm^2 . The ZVS-enhanced low-power clock driver (including the inductor) and the reference circuit occupy 0.15 mm^2 and 0.03 mm^2 , respectively.

The power gating transistor M_{ph} and the nMOS shorting transistor M_{ns} were not implemented in the test chip. Otherwise, the area occupied by them would have been approximately 0.0035 and 0.0006 mm^2 , respectively, which are very small compared to the total circuit area.

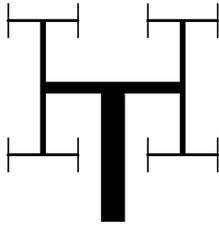


Fig. 15. Tapered H-tree clock distribution network.

C. Test Results

Chip measurement results in Fig. 8 show energy savings for a clock frequency range of 2.75 to 4 GHz. The measurements show increasing power savings as clock frequency increases to 4 GHz. At lower frequencies, the inductor current will have more time to build-up, which results in an increased resistive voltage drop across the inductor. Thus, the energy savings are reduced. To improve this, a larger inductance is needed as shown in Fig. 9.

The simulation results show very good agreement with the measured results below 3.5 GHz, but begin to deviate at higher frequencies. Measurements above 4 GHz were not possible due to limits of our test equipment. At 4 GHz, measurements confirm the power consumption is reduced from 117 mW (in the reference circuit) to 76 mW (in the ZVS-enhanced circuit), a net power savings of 35%.

The clock waveforms are made available off-chip using open-drain pMOS buffers. At 4 GHz, the RMS clock jitter is measured to be 1.25 ps and 1.17 ps for the ZVS-enhanced low-power clock driver and the reference clock driver, respectively. The difference is near the limitations imposed by the lab equipment. The jitter is in the same range predicted by simulation, i.e., 1 to 2 ps. Although simulation disagreed with measurements and predicted the ZVS-enhanced low-power circuit to have lower jitter, noise from many different sources may contribute to the measured results.

VI. DISCUSSION

The design introduced here benefits from the charge stored in the clock load capacitance. It relies upon a low-RC metal network between the driver and the final flip-flop loads. This network is designed manually for low delay, shielding, and load matching [11], [14], [15]. Thus, the exact location for the proposed clock driver circuits depends upon the configuration of the local clock distribution network. Although our prototype and simulations did not model metal RC, we did use a multi-fingered transistor layout to represent distributed flip-flops.

Some clock distribution networks form a tree structure as shown in Fig. 15 [21], [22] rather than the gated mesh described earlier. If the ends of the tree branches are connected to each other, a mesh-like structure is formed which has reduced interconnect resistance and lower skew within the clock tree. In this case, a single lumped driver can be used to drive the entire clock tree. The driver needs to provide enough current to drive the clock network load capacitance while keeping the clock waveform intact.

In the example of Fig. 15, the last inverter that drives the clock tree trunk is the biggest inverter, which drives the capac-

itance of the whole H-tree. Alternatively, inverters can be distributed along the path of the H-tree, with the last inverters in the chain representing numerous distributed drivers. These distributed drivers can be shorted together to reduce metal resistance and eliminate any final local clock skew. In both of these cases, the last inverter stage is the ideal location to include the inductor and use the simplified or ZVS-enhanced low-power driver.

A practical concern is the area overhead imposed by the large inductor when integrated into a real processor. The clock network in IBM's POWER6 consumes about 22 W. With an overall area of 341 mm², this results in an estimated clock power consumption of 65 mW/mm². Using $P = F_{sw}CV_{DD}^2$, the overall clock capacitance is estimated to be 4.4 nF or 13 pF/mm². In this work, a load gate capacitance of 21 pF has been assumed, corresponding to a 1.6 mm² region of the POWER6 die. In contrast, the area needed to implement the inductor in this new clock driver is 0.1 mm² which is an increase of about 6% in chip area.

There will be some variation across the die for clock capacitance, as some regions have little clocking and others are densely clocked. Although we assume a 21 pF capacitance covers a 1.6 mm² region on average, there are likely regions with much more dense clocking. Hence, a region somewhat less than 1 mm² is probably a better estimate. Even with this adjustment, the area 'cost' of power savings is still attractive.

A resonant clocking test-chip, fabricated in 90 nm CMOS, achieved 20% energy savings in [6]. That chip uses $C_{clk} = 7.5$ pF and four parallel sets of LC passives resulting in an effective C of 80 pF and effective L of 250 pH with a target resonance of 3.7 GHz. The total area of the LC passives is ~ 0.06 mm². One key difference, however, is that resonant clocking applies only to clock *distribution* energy, which is only a small fraction (less than one-quarter) of the clock *load* energy required to drive the end latches [23]. By recycling the energy of the final clock load, this work achieves much larger net energy savings, but also requires wider metal paths for low resistance in the inductor (more layout area).

This new driver is scalable for processors with large clock loads. With each generation, the clock network is divided up into more and more local regions. However, the "capacitance per local region" depends mostly upon the granularity of clock gating used in the design to save power.

VII. CONCLUSION

The two circuits presented in this paper are intended to replace local clock drivers in modern processor designs. Each gated region would get one of these new drivers, resulting in several inductors being distributed across the chip.

Power consumption of a 4 GHz clock driver is reduced by 35% by recycling energy stored in the clock tree load capacitance and delivering it back to the on-chip power distribution grid. The energy is transferred in a non-resonant way using an on-chip inductor in a configuration resembling a full-bridge DC-DC converter. A ZVS technique is used to further increase the power savings. In particular, this design provides a nearly square clock with fast edges, which is an improvement to the sinusoidal waveform with slow edges produced by resonant clocking.

A 90 nm chip prototype is fabricated and measured to confirm the energy savings. Simulations predict a 37% power reduction while measurements show a 35% power reduction. Clock measurements also suggest that jitter performance is slightly deteriorated using this technique, despite simulations predicting that it should be enhanced. The layout area of the driver is about 6% of the area over which it operates.

Investigating the source of the power savings by simulation, transistor currents and clock waveforms clearly show that energy recycling is being performed. The effect of varying the supply voltage and the inductor value on power consumption is also presented, showing that an optimized design may be achieved for a particular clock frequency.

The ZVS-enhanced driver circuit achieves the best power savings, but it also degrades the falling clock edge rate. The simplified circuit offers a good compromise solution.

REFERENCES

- [1] J. Yu, C. Chung, and C. Lee, "A symbol-rate timing synchronization method for low power wireless OFDM systems," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 9, pp. 922–926, Sep. 2008.
- [2] S. Baeg, "Low-power ternary content-addressable memory design using a segmented match line," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 6, pp. 1485–1494, Jul. 2008.
- [3] Y. Kim, J. Kim, J. Kim, and B. Kong, "CMOS differential logic family with conditional operation for low-power application," *IEEE Trans. Circuits System. I, Reg. Papers*, vol. 55, no. 5, pp. 437–441, May 2008.
- [4] 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 *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Feb. 2007, pp. 96–97.
- [5] S. Naffziger, B. Stackhouse, T. Grutkowski, D. Josephson, J. Desai, E. Alon, and M. Horowitz, "The implementation of a 2-core, multi-threaded itanium family processor," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 197–209, Jan. 2006.
- [6] 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, pp. 102–109, Jan. 2005.
- [7] S. C. Chan, K. L. Shepard, and P. J. Restle, "Distributed differential oscillators for global clock networks," *IEEE J. Solid-State Circuits*, vol. 41, no. 9, pp. 2083–2094, Sep. 2006.
- [8] M. Alimadadi, S. Sheikhaei, G. Lemieux, S. Mirabbasi, and P. Palmer, "A 3 GHz switching DC-DC converter using clock-tree charge-recycling in 90 nm CMOS with integrated output filter," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Feb. 2007, pp. 532–533.
- [9] M. Alimadadi, "Recycling clock network energy in high-performance digital designs using on-chip dc-dc converters," Ph.D. dissertation, Dept. Elect. Comp. Eng., Unive. British Columbia, BC, Canada, 2008.
- [10] V. S. Sathe, J. C. Kao, and M. C. Papaefthymiou, "Resonant-clock latch-based design," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 864–873, Apr. 2008.
- [11] N. A. Kurd, J. S. Barkatullah, R. O. Dizon, T. D. Fletcher, and P. D. Madland, "A multigigahertz clocking scheme for the pentium 4 microprocessor," *IEEE J. Solid-State Circuits*, vol. 36, no. 11, pp. 1647–1653, Nov. 2001.
- [12] K. Yoshihiro, I. Yasuhiro, S. Tomoki, K. Keisuke, O. Koki, and K. Masahito, "CAD technology of the SX-9," *NEC Tech. J.*, vol. 3, no. 4, pp. 34–38, 2008.
- [13] P. E. Gronowski, W. J. Bowhill, R. P. Preston, M. K. Gowan, and R. L. Allmon, "High-performance microprocessor design," *IEEE J. Solid-State Circuits*, vol. 33, no. 5, pp. 676–686, May 1998.
- [14] S. Naffziger, G. Colon-Bonet, T. Fischer, R. Riedlinger, T. Sullivan, and T. Grutkowski, "The implementation of the itanium 2 microprocessor," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, pp. 1448–1460, Nov. 2002.
- [15] P. Mahoney, E. Fetzer, B. Doyle, and S. Naffziger, "Clock distribution on a dual-core multi-threaded itanium-family processor," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Feb. 2005, pp. 292–293.
- [16] J. W. Tschanz, S. G. Narendra, Y. Ye, B. A. Bloechel, S. Borkar, and V. De, "Dynamic sleep transistor and body bias for active leakage power control of microprocessors," *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 1838–1845, Nov. 2003.
- [17] B. Mesgarzadeh, M. Hansson, and A. Alvandpour, "Jitter characteristic in charge recovery resonant clock distribution," *IEEE J. Solid-State Circuits*, vol. 42, no. 7, pp. 1618–1625, Jul. 2007.
- [18] C. Yue and S. Wong, "On-chip spiral inductors with patterned ground shields for Si-based RF ICs," *IEEE J. Solid-State Circuits*, vol. 33, no. 5, pp. 743–752, May 1998.
- [19] J. N. Burghartz, "Progress in RF inductors on silicon—Understanding substrate losses," in *Proc. IEEE Int. Electron Devices Meeting (IEDM)*, Dec. 1998, pp. 523–526.
- [20] A. M. Niknejad and R. G. Meyer, "Analysis, design, and optimization of spiral inductors and transformers for Si RF IC's," *IEEE J. Solid-State Circuits*, vol. 33, no. 10, pp. 1470–1481, Oct. 1998.
- [21] E. G. Friedman, "Clock distribution networks in synchronous digital integrated circuits," *Proc. IEEE*, vol. 89, no. 5, pp. 665–689, May 2001.
- [22] B. H. Calhoun, Y. Cao, X. Li, K. Mai, L. T. Pileggi, R. A. Rutenbar, and K. L. Shepard, "Digital circuit design challenges and opportunities in the era of nanoscale CMOS," *Proc. IEEE*, vol. 96, no. 2, pp. 343–365, Feb. 2008.
- [23] N. Ranganathan and N. Jouppi, "Evaluating the potential of future on-chip clock distribution using optical interconnects," Hewlett-Packard Development Company, Tech. Rep. HPL-2007-163, Oct. 2007.



Mehdi Alimadadi received the B.A.Sc. degree from Iran University of Science and Technology, Tehran, Iran, in 1989 and the M.A.Sc. and Ph.D. degrees from University of British Columbia, Vancouver, BC, Canada, 2000 and 2008, respectively.

He worked as an Electrical Design Engineer for a couple of companies in Toronto, ON, Canada, during years 2000 to 2004. In 2007, he also worked as a part time consultant for a local company. His research interests include on-chip power management, switching power converters and DSP based digital

controllers.

Dr. Alimadadi has been a recipient of MITACS ACCELERATE grant and a Postdoctoral Research Fellow at UBC ECE department working on advanced high frequency power converters. He is also a registered Professional Engineer (P.Eng.) in the Province of British Columbia, Canada.



Samad Sheikhaei (S'02) received the B.Sc. and M.Sc. degrees in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1996 and 1999, respectively, and the Ph.D. degree from the University of British Columbia, Vancouver, BC, Canada, in 2008.

He was engaged in research and design engineering at Sharif University of Technology. He also worked in industry for a couple of years. In August 2009, he joined the Department of Electrical and Computer Engineering at the University of Tehran, Iran, where he is an Assistant Professor. His research interests are analog, mixed-signal, and RF integrated circuits design, and his current research projects include high-speed analog-to-digital converters, high-speed serial links, and on-chip dc-dc power converters.



Guy Lemieux (M'03–SM'08) received the B.A.Sc. degree from the division of engineering science at the University of Toronto, and the M.A.Sc. and Ph.D. degrees in electrical and computer engineering at the University of Toronto, Toronto, ON, Canada.

In 2003, he joined the Department of Electrical and Computer Engineering at The University of British Columbia, Vancouver, BC, Canada, where he is an Associate Professor. He is coauthor of the book *Design of Interconnection Networks for Programmable Logic* (Kluwer, 2004). His research interests include

computer-aided design algorithms, VLSI and SoC circuit design, FPGA architectures, and parallel computing.

Dr. Lemieux was a recipient of the Best Paper Award at the 2004 IEEE International Conference on Field-Programmable Technology.



Shahriar Mirabbasi (S'95–M'02) received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1990, and the M.A.Sc. and Ph.D. degrees in electrical and computer engineering from the University of Toronto, Toronto, ON, Canada, in 1997 and 2002, respectively.

Since August 2002, he has been with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC, Canada where he is currently an Associate Professor. In 2008,

he was a visiting Professor at the Swiss Federal Institute of Technology (ETH) Zurich, and subsequently in 2009 at the Laboratoire d'Intégration du Matériau au Système (IMS Lab), Bordeaux, France. His current research interests include analog, mixed-signal, RF, and mm-wave integrated circuit and system design with particular emphasis on communication, sensor interface, and biomedical applications.



William G. Dunford (S'78–M'81–SM'92) was a student at Imperial College, London, U.K., and the University of Toronto, Toronto, ON, Canada. He has also been a faculty member of both institutions and is now on the faculty of the University of British Columbia.

Industrial experience includes positions at the Royal Aircraft Establishment (now Qinetiq), Schlumberger and Alcatel. He has had a long term interest in photovoltaic powered systems and is also involved in projects in the automotive and energy

harvesting areas. He is a director of Legend Power Systems, where he has also been active in product development.

Dr. Dunford has served in various positions on the Advisory Committee of the IEEE Power Electronics Society and chaired PESC in 1986 and 2001.



Patrick R. Palmer (M'87) received the B.Sc. and Ph.D. degrees in electrical engineering from Imperial College of Science and Technology, University of London, U.K., in 1982 and 1985, respectively.

He joined the faculty at the Department of Engineering, University of Cambridge, UK in 1985 and St. Catharine's College, Cambridge in 1987. He became an Associate Professor in the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC, Canada in 2004 and Reader in Electrical Engineering, University of

Cambridge, U.K., in 2005. His research is mainly concerned with the characterization and application of high-power semiconductor devices, computer analysis, simulation and design of power devices and circuits and he has further interests in fuel cells. He has extensive publications in his areas of interest and is the inventor on two patents.

Dr. Palmer is a Chartered Engineer in the U.K.