

A Resonant Global Clock Distribution for the Cell Broadband Engine Processor

Steven C. Chan, *Member, IEEE*, Phillip J. Restle, *Member, IEEE*, Thomas J. Bucelot, John S. Liberty, Stephen Weitzel, John M. Keaty, Brian Flachs, Richard Volant, Peter Kapusta, and Jeffrey S. Zimmerman

Abstract—Resonant clock distributions have the potential to save power by recycling energy from cycle-to-cycle while at the same time improving performance by reducing the clock distribution latency and filtering out non-periodic noise. While these features have been successfully demonstrated in several small-scale experiments, there remained a number of concerns about whether these techniques would scale to a product application. By modifying the Cell Broadband Engine Processor to incorporate a large resonant global clock network, power savings with full functionality is demonstrated over a 20% range in clock frequencies, and a 6–8 Watt power savings at 4 GHz. This was achieved by changing one wiring level and adding an additional thick copper level to create inductors and capacitors.

Index Terms—Clock distribution, clock grid, clock tree, inductor, jitter, microprocessor, resonant circuit, resonant clock.

I. INTRODUCTION

CLOCK distributions on large, high-performance digital integrated circuits must meet stringent timing and power requirements. Skew and jitter from the clock distribution increases timing margins and reduces performance. The additional timing margin required for clock uncertainty prevents a chip from operating at lower supply voltages, thereby increasing power consumption. Resonant clocking techniques [1]–[6] have shown promise in reducing global clock power and timing uncertainty. By resonating the large global clock capacitance with an inductance, the energy used to charge the clock node each period can be recycled within the LC resonant tank network, resulting in lower clock power. Significant additional power savings can be realized by reducing the strength of clock buffers driving the LC load because after start-up, only losses need to be overcome at resonance. Skew and jitter

can be improved due to the band-pass characteristic of the LC network and the ability to use fewer or lower-gain clock buffering stages [1].

Numerous small test chips have demonstrated the potential benefits of resonant clocking [1]–[6], but these typically used off-chip inductors or high-quality isolated inductors which would not be practical in most applications. Other experiments used significantly different global clocking methodologies that could be difficult to incorporate into existing local clocking and timing schemes. Concerns about signal integrity and noise coupling from the inductors integrated onto the chip also remained a concern. A large-scale demonstration of an LC resonant clock on a fully-functional microprocessor was needed to address these issues and concerns, with hardware measurements to quantify the actual benefits. In this work, we address most concerns by experimentally transforming the global clock on the Cell Broadband Engine (Cell/B.E.) processor [7] into an LC resonant clock.

The Cell/B.E. processor is a high-performance multi-core microprocessor that consists of a 64-bit general purpose processor core along with eight specialized co-processors, all operating nominally at 3.2 GHz using a single high-frequency global clock that covers 85% of the chip area. The redesign of the global clock on the Cell/B.E. processor was a challenging undertaking that required modifications to an already completed commercial microprocessor in volume manufacturing. In this paper, we describe the work needed to build a resonantly clocked Cell/B.E. processor. Hardware measurements from the modified design show full functionality at 3.2 GHz, and power savings of 25% in the global clock and 5% overall in the chip power at 4 GHz. This paper is organized as follows. In Section II, the architecture of the resonant clock network is presented. Section III describes how the Cell/B.E. processor was modified. Measurement results from resonant and non-resonant control chips from the same wafer lot are presented in Section IV.

II. RESONANT-LOAD GLOBAL CLOCK ARCHITECTURE

High-performance IBM microprocessors use a tree-driven grid approach in the global clock design [7]–[10]. A buffered tree using on-chip transmission lines is used to distribute the clock across the chip, and a grid is used to short the ends of the final tree together for lower skew and better robustness with respect to process, voltage and temperature (PVT) variations [11]. This style of distribution is attractive for several reasons. First, it can be designed to have arbitrarily small nominal skew like a perfectly balanced tree design. Second, compared to a

Manuscript received April 14, 2008; revised July 28, 2008. Current version published December 24, 2008. This work was supported in part by DARPA HR0011-07-9-0002. Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.

S. Chan was with IBM T. J. Watson Research Center. He is now with TSMC, San Jose, CA 95134 USA (e-mail: SChan@tsmc.com).

P. J. Restle and T. J. Bucelot are with IBM Research Division, Thomas J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: restle@us.ibm.com; bucelot@us.ibm.com).

S. Weitzel and B. Flachs are with the IBM Systems and Technology Group, Austin, TX 78758 USA.

J. S. Liberty is with IBM Systems and Technology Group, Customize Technology, Austin, TX 78758 USA.

J. M. Keaty is with IBM Systems and Technology Division, Austin, TX 78759 USA (e-mail: keaty@us.ibm.com).

R. Volant is with the IBM Systems and Technology Group, Hopewell Junction, VT 12533 USA.

P. Kapusta and J. S. Zimmerman are with IBM Systems and Technology Group, Essex Junction, VT 05452 USA.

Digital Object Identifier 10.1109/JSSC.2008.2007147




Fig. 1. A simplified global clock distribution with a resonant load. Four clock sectors, which form the basic building block of the distribution, are shown.

pure tree design, a tree-driven grid can achieve greater robustness to PVT as well as reduced power and very competitive nominal skew.

Fortunately, we found that such a tree driven grid design could be transformed into a resonant clock design with relatively minor modifications of the original non-resonant “control” product design, which made this experiment possible.

In order to achieve seamless integration of a resonant clock into existing tree-driven grid global clocking designs, certain features in the resonant clock are desirable. The two most important are uniform-phase and uniform-amplitude. This is because a single-phase full-rail clock drives almost all of the flip-flops and latches on a microprocessor. The resonant-load global clock distribution developed by the authors in [1] is uniform in phase and amplitude, so is an ideal candidate to be used on the modified Cell/B.E. processor.

The Fig. 1 shows a resonant-load global clock distribution as described in [1]. The clock is distributed from a single synchronous source and is buffered through a tuned-balanced tree. The tree then drives a set of clock sectors, a basic unit of the distribution driven by the lowest buffer level of the global clock tree. For simplicity, only four clock sectors are shown in Fig. 1, while on real microprocessors there are hundreds of sectors. The Cell/B.E. contains more than 800 sectors.

The sector clock buffer (SCB) associated with each clock sector provides the gain needed to drive a local tree, the clock grid, and the local clock buffers (LCBs), as shown in Fig. 2. The LCBs are the final stages driving the latches, and incorporate test and clock gating functions. To support single-phase, single-ended, LC clocking with uniform amplitude, spiral inductors are attached to the clock tree. By attaching the inductor near the end of the sector trees which drive the grid (as shown in Fig. 2) a small “treelet” with 2 leaf nodes distributes the current flowing into and out of the inductor, reducing resistive losses which would degrade the Q of the network. Connecting the inductor directly to the grid at a single point could result in local skew near that connection. Distributing the current from each inductor across 2 points on the clock grid (as opposed to just one) thus reduces the local skew on the grid. For even lower skew, a single inductor could be connected directly to the output of the sector buffer, but this reduces power savings since the sector buffer output waveform does not well match the more sinusoidal waveform natural to the inductors, and larger currents result from this mismatch. Choosing the




Fig. 2. The components and topology of a resonant clock sector with tree wires shown as dashed lines.




Fig. 3. A simple lumped circuit model of the resonant clock sector.

number of inductors and the best connection points for these inductors involves consideration of inductor area, skew, and power. The other end of the inductor is attached to a large capacitor described below. The purpose of this capacitor is to establish a mid-rail DC voltage (approximately) around which the clock network operates. The buck-converter-like topology of the resonant clock network, shown in Fig. 3, results in this capacitor (labeled C in Fig. 3) reaching a steady-state average voltage of $VDD/2$ (assuming a 50% duty cycle global clock) with a time constant and ripple determined by the ratio of the clock capacitance, C_{clock} , to the added capacitance, C. When the clock duty cycle diverges slightly from 50%, the capacitor simply averages the driving clock signal, so that a 45% duty cycle produces an average voltage $V_{ave} = 0.45 VDD$. This results in a small voltage offset in the clock grid waveform. As long as the duty-cycle is not far from 50%, which is usually the case for global clock signals, this small offset is not a problem. In fact, because of this offset, the effective duty cycle of the resonant global clock varies with the duty cycle of the clock source very much like the control non-resonant clock distribution on the product version of this chip. Since hardware sometimes performs better with a duty cycle slightly offset from the ideal 50%, this may be considered the desired behavior. In addition, if an ideal $VDD/2$ supply was used, and the desired duty cycle was not 50%, extra power would be consumed due to the resulting DC current through the inductor. In [1], the added capacitance was implemented using MOS gate capacitors (with C roughly ten times larger than C_{clock} depending on the sector) which were positioned adjacent to the spiral inductors.

All the flip-flops and latches in the design are driven by LCBs that tap into the global clock grid. LCBs provide the additional




Fig. 4. The clock distribution has 17 levels of buffering and is divided into 6 tiles—the tree to only one tile is expanded for clarity.

gain needed to drive the actual loads. While it would be desirable to save even more power by resonating the local clock signal(s) produced by the LCBs, this is not easily compatible with local clock gating, where the local clocks are often gated on a cycle-by-cycle basis. Depending on the Q of the inductor, and the size of the decoupling capacitor C in Fig. 3, the resonant clock circuit can require several cycles to converge to the desired periodic clock waveform when the clock source starts. Because the SCB is driving an LC load, much of the power savings in this resonant-load global clock distribution actually comes from the ability to reduce the drive strength of the SCBs, and the associated parasitic capacitances. In [1], total power saving in the clock distribution using a resonant-load approaches 80%, only 20% of which is the result of energy recirculation between inductor and capacitor. In this Cell/B.E. design, modeling predicts that 60% of the power savings comes from the reduction in the drive strength of the SCBs driving the LC load.

III. MODIFYING THE CELL/B.E. PROCESSOR GLOBAL CLOCK

The Cell/B.E. processor has three independent tree-driven-grid global clock distributions that support the microprocessor logic (nclk), memory interface (mick), and bus interface (bclk). The nclk operates at 3.2 GHz while the mick and bclk operate at 1.6 GHz. Since the nclk covers 85% of the chip and consumes 97% of the global clock power, only this distribution was modified to have a resonant clock.

In addition, simulations showed that expected power savings from resonant clocking increases with frequency, so that little if any power savings could be achieved for the 1.6 GHz clock distributions without higher- Q inductors.

Fig. 4 shows a cross-section of the nclk distribution. There are 17 levels of buffering in the tree from the PLL to the nclk mesh. The first 16 buffers are inverters driving long transmission lines, while the 17th buffer is a sector clock buffer (SCB). The SCB is a three stage inverter and drives the nclk mesh and local clock buffers (LCBs). There are 830 SCBs in the nclk distribution driving a total capacitive load of ~ 2 nF. The complete nclk distribution is assembled by shorting six “tiles” together. Only one tile is expanded in Fig. 4, which ends in 192 SCBs. For simplicity, the other tiles (not identical), are not expanded.

These trees consist of carefully optimized transmission lines, length and delay matched at each level to achieve low skew in the presence of PVT variations.

Since the 90-nm implementation of the Cell/B.E. processor was complete and in volume manufacturing at the time this work started, the goal was to transform the global clock into a resonant clock with as few modifications as possible to reduce design resources needed, wafer processing, and mask cost. With the processor design already completed, there was no space within the existing eight metal layers on the chip to place all the spiral inductors required, even though they would require only a small percentage of the top two existing metal layers. Instead, a new 9th level of metal was added to the chip. This new 1.2 μm -thick copper level was used to implement 830 on-chip spiral inductors. This resonant-load global clock design also requires the addition of a large capacitor to generate the $V_{DD}/2$ supply, as described above and shown in Fig. 3. Fortunately, the same new metal layer can be used to implement a large distributed metal fringe capacitor since the 830 inductors and vias needed for chip I/O occupy a small fraction of the chip area.

The area of the nclk mesh driven by a single SCB, as shown in Figs. 1 and 2 is approximately 600 μm by 400 μm . Within each of the 830 clock sectors on the chip, a 2.75 turn 1.2 nH inductor (inner diameter 100 μm) is attached to the clock tree. The clock sector load capacitance can vary across the chip by as much as 4X, which raises concerns about whether the added inductance needs to be tuned to the local sector load capacitance. However, simulations showed that a single 1.2 nH inductor design could be replicated across the chip to simplify the implementation. This simplified design was considered acceptable because it resulted in a simulated skew increase of only 3.2 ps in the resonant design compared to the original non-resonant design (8.4 ps vs. 5.2 ps) near the resonant frequency. In the original design, the SCBs were chosen from a selection of power levels to lower power and skew. The tree wires within each sector were also tuned depending on the detailed SCB placement and the distribution of clock pin loads within each sector [11]. Ideally, the inductors would also be selected to match the capacitive loads in each sector. If the Q of this distributed LC network was higher, skew and signal quality requirements would require accurate tuning of inductors and capacitors. However in this low- Q regime, with the inductors shorted together using a fairly rigid clock grid that tends to average out local variations, inductor tuning was found to be much less critical.

The inductor has a quality factor of only 1.8 at 3.2 GHz (the target resonant frequency) primarily due to eddy currents in the power grid which reduce the inductance of the spiral by a factor of two. Fig. 5 shows field solver simulations of the eddy currents in the power grid underneath the inductor (Fig. 5(a)) and the reduced magnetic field within the spiral that results (Fig. 5(b)). If the loops in the power grid could be safely removed as in Fig. 5(c), the inductor would have a quality factor of 4.8 and an inductance of 2.2 nH. Power grid cutting and power via removal techniques used to improve inductor quality factor in [1] were not used in this work because of concerns about the integrity of the resulting power grid and the costs associated with additional mask changes. Cutting nearby loops in the power grid causes two problems. First, without careful design there will be




Fig. 5. Full-wave 3D field solver simulation results.




Fig. 6. Metal fringe capacitor connects to inductor and a lumped model of the resonant network.

inevitable increases in local IR drops in the power grid. Second, cutting power wires would affect the on-chip transmission line properties of critical global clock wires, since these power wires can be transmission line return paths.

While one end of the inductor is connected to the clock trees driving the grid, the other end of the inductor is attached to a large distributed fringe capacitor (Fig. 6) constructed from minimum width and space wiring on the same metal layer as the inductor. With a 50% duty cycle clock, the fringe capacitor reaches a steady-state average voltage centered about $VDD/2$ in less than 3 cycles, with a simulated peak to peak ripple of 260 mV on a 1 V supply. The 12 nF fringe capacitance is 6X larger than the grid load capacitance consisting of the tree wires, the clock grid wires, the local short wires driving the gates, and finally the device gate capacitances. This 6X capacitance multiplier results in a good balance between low voltage ripple (affects power savings and skew/jitter improvement) and time to reach steady state (affects startup and low-frequency operation).

Fig. 7 is a die photo of a small section of the chip showing the columns of inductors and C4 pads on the new metal level. The inductors are small enough to avoid the C4 pad array.

In addition to the new metal layer used to implement the inductors and capacitors, one additional change was needed to build the resonant clock on the Cell/B.E. processor. A single mask level change on the fourth metal layer (M4) was used to reduce the drive strength of the 830 SCBs to half their original value. The SCBs are implemented as two parallel three stage inverters. The M4 change disconnected one of the two parallel driver's outputs and grounded its corresponding input.

In future designs, where it would be advantageous to implement a resonant clock design without the need for an additional metal layer, the inductors would need to be integrated carefully with the power grid on the top two metal layers. The large capacitors added could be implemented using MOS capacitors carefully designed to have low parasitic resistance. For many applications, a switch would be needed to allow operation in a

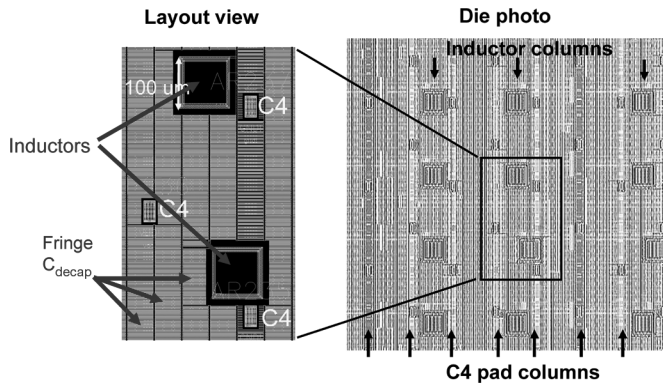


Fig. 7. Layout and die photo of a small section of the chip showing the inductor and C4 pad columns.

non-resonant mode far from resonance (such as manufacturing test). Simulations show that the addition of such a switch and MOS capacitors reduces the effective power savings by approximately 20% (for example a 5 Watt savings would be reduced to 4 Watts) and the switches would require less than 0.1% of the chip silicon area. However, if significantly higher Q inductors were desired for greater power savings and jitter reduction, this could require significant chip area ($\sim 10\%$ or more) and floor plan disruption using currently available inductor designs and technology. Alternatively, inductor Q could be increased significantly with lower cost by removing loops in the power grid near the inductors to reduce eddy currents [1]. This would impact the robustness and complexity of the power distribution, requiring more careful design and analysis of the on-chip power distributions, but would significantly reduce the cost of implementing a resonant global clock design.

IV. MEASUREMENT RESULTS

A special six wafer lot was used for this work. The six wafers were selected from a larger parent lot based on parametric data which indicated that the devices on these wafers were similar in terms of drive strength and leakage and that the modules from these wafers would have the highest yield. Four of the wafers were “resonant” and received the M4 change and the new metal level while two of the wafers received the normal design and served as “control” for comparison purposes. Additional test wafers were used to ensure that the additional metal level was added correctly, that the inductors were properly connected to the existing global clock network on the chip, and that there were no shorts in the fringe capacitor. Fig. 7 shows a section of design layout and an associated photo from an optical microscope used to inspect the quality of the new metal level. Fig. 8 shows a higher magnification photo. Metal fill is visible inside and outside the inductor in Fig. 8. Field solver simulations show that the effect on the inductor quality factor due to this metal fill is negligible since the eddy currents in the power grid dominate. The apparent “waviness” in the metal fringe capacitor in Fig. 8 is due to moiré effect.

A. Waveforms, Low Frequency Operation and Wafer Final Test

After wafer processing and C4 plating, all six wafers received low-frequency manufacturing wafer final test screening.




Fig. 8. Die photo of a small section of the chip showing a corner of one inductor and the metal fringe capacitor.




Fig. 9. Simulated waveforms at start-up, near resonant frequency.

Although high yield was achieved on the two control wafers as expected, all four resonant wafers had zero yield from this manufacturing test. More in-depth analysis and simulation revealed that low frequency operation of the resonant parts is hindered by mismatch between the LC resonant frequency and the low frequency input test clock. Fig. 9 shows simulated waveforms near resonance at the buffer output, the clock mesh, and the large capacitor. Only 2 or 3 cycles after start-up are required to achieve low jitter operation, although several cycles are needed for the large decoupling capacitor to reach steady state. Simulation waveforms in Fig. 10 show that at a 1.2 V supply voltage, there is double switching at the output of an LCB at 200 MHz and slight double switching at 1 GHz. At low supply voltages, there is less double switching at 200 MHz but more at 1 GHz. Because wafer final test measurements are done at these low frequencies, the double switching resulted in all the resonant wafers receiving a manufacturing screen yield of zero. However, because simulations indicated that the clock would function better at frequencies above 1 GHz, modules were still built from a selection of dies from the resonant wafers based on the assumption that they would work at normal operational frequencies.

B. Test Setup and Chip Booting

After resonant and non-resonant control modules were built from diced wafers, tests were performed on a Cell/B.E. evaluation board with a socket for the processor module. Module temperature control was through a chilled water cooling solution and an on-module temperature sensor. Normally, the Cell/B.E. processor boots using an external 400 MHz reference clock that




Fig. 10. Simulations show double switching below 1.6 GHz. At low supply voltages, there is less double switching at 200 MHz.

is applied to the chip. Until the on-chip PLL is activated, the clock grid runs at this reference clock frequency. When activated, the on-chip PLL multiplies this frequency by 8X to produce the core clock grid standard 3.2 GHz operating frequency. Since the resonant clock does not function well at very low frequencies, the reference clock was raised to 1.6 GHz to successfully boot the part thereby resolving the low frequency operation problems. The on-chip PLL was then set to a 2X multiplier to get to the standard 3.2 GHz operating frequency. A high performance function generator was used to provide this reference clock, since the evaluation board clock generator was not capable of operating at frequencies greater than 1 GHz. Multiple frequencies required during testing were achieved through a combination of module PLL divider settings and function generator settings. Correct operating frequency was validated using an oscilloscope connected to a Cell/B.E. test port set up to measure the chip global clock signal. A module test interface (JTAG) was used in initial module communications and setup. This was used for scan based program load and evaluation. Full Cell/B.E. Linux program boot and operation was achieved on the resonant parts using evaluation board flash firmware and an external hard drive. The Linux testing on the resonant Cell/B.E. parts used the I/O communication ports and XDR memory ports running at full speed. An external calibrated power supply with on-die voltage sense feedback control was used to make power measurements in place of the evaluation boards VRM (Voltage Regulator Module).

C. Test Operation and Clock Jitter

Multiple test programs were used to evaluate resonant chip functionality and performance. To confirm full chip functionality, a Cell/B.E. based architectural test program was utilized. This program was designed to test the software architectural compliance of the resonant modules. The program was designed to have very high functional test coverage of the complete Cell/B.E. module. Scan (JTAG interface) based testing was also used to evaluate resonance clock modules. The scan based testing included special purpose programs designed to run on the Cell/B.E. processor. Control programs running on an external Linux based computer were used to automatically load each test programs through JTAG, run it, and check whether

the test ran successfully. This control program would automatically change the voltage and frequency to determine functional operation region of the module under test (voltage/frequency shmoos). Power versus frequency results were also determined using a Cell/B.E. power reference program run using this same methodology. Also, array and logic built in self tests (BIST) that are part of the Cell/B.E. logic were exercised by this external control computer. These BIST functions were also exercised at multiple frequencies and core voltages to determine functional operation region of the module under test.

Voltage/frequency shmoos were made to determine if there is any difference in the minimum passing voltages (V_{min}) between the resonant and non-resonant parts. Statistically significant lower V_{min} on the resonant chips would indicate larger margins on critical timing paths due to improved clock jitter. However, the sample size of approximately 30 parts tested did not show a significant increase or decrease in V_{min} . Subsequent simulations of various sources of jitter showed that with this low-Q resonant clock design, only a small reduction in cycle-compression (short cycles) is expected, and the jitter reductions depends on the characteristics of the jitter source. As one example, there are often many circuits receiving a half-frequency clock, which can tend to collapse V_{dd} slightly every other cycle, which can in turn affect clock buffer delays and cause the global clock signal to consist of alternating long and short cycles. The resonant clock will tend to average these long and short cycles, reducing jitter and improving performance. However in this low-Q design, this jitter reduction is small, and in any case this kind of easily filtered jitter is not expected to be significant on the Cell/B.E. processor.

Testing of the resonant modules indicated that the resonant chips fail below 1.6 GHz. This is due to the double switching of the LCB inputs due to frequency mismatch between the LC resonant frequency and the input clock frequency. As shown in the simulations waveforms in Fig. 10, this double switching is the same reason that wafer final test (using a 200 MHz clock) showed zero yield on the resonant parts. Fig. 10 also shows that at low supply voltages (<0.8 V), some low frequency functionality is restored (confirmed in the hardware) because the center point of the glitching is no longer around the threshold voltage of the receiving LCBs. Future implementations of resonant clocking could include a switch to disable the resonance for low frequency operation and wafer test.

D. Power Savings

Fig. 11 shows the measured power dissipation (normalized) of the resonant and non-resonant chips with only the global clock running and leakage current subtracted. There is a local minimum in the resonant chip power at 3.2 GHz (the resonant frequency) and a power savings of 5–25% compared to the non-resonant chip power between 4 and 5 GHz. The power savings is larger at higher frequencies because even at half strength, the SCBs are still too strong. As previously published [1], maximum power saving is achieved when the buffers driving the resonant clock are sized just strong enough to achieve a full-rail, approximately sinusoidal, clock signal. When the sector clock buffers are stronger, power is wasted driving the LC circuit harder than

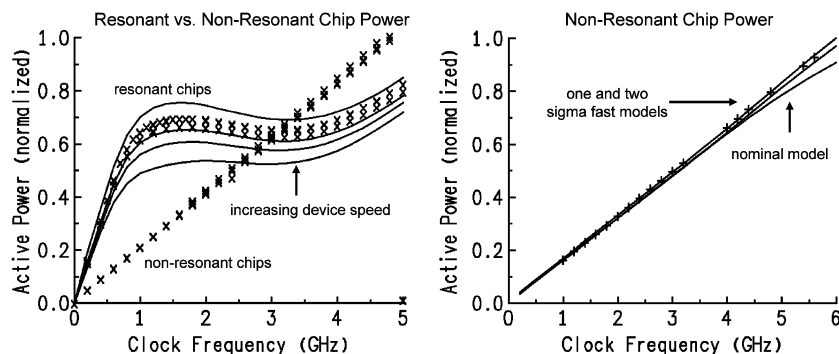


Fig. 11. Measured power and simulated power for different device speeds (solid curves). Non-resonant power used to determine actual device speed.

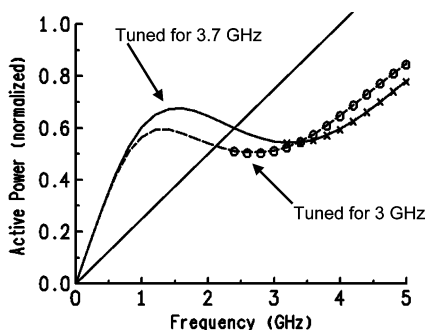


Fig. 12. Simulation results showing simulated effect of tuning for 3 GHz vs. 3.7 GHz. Symbols highlight frequencies with best signal quality.

necessary, fighting the natural sinusoidal waveform. When running actual workloads, the measured power savings is approximately 5% of the total chip power at 4 to 5 GHz. At frequencies below 3 GHz, the resonant chips use more power because of the extra loading of the fringe capacitor. Fig. 11 also shows that the resonant chip power is sensitive to device speed, increasing with increasing device speed (solid curves from simulation). The nearly linear slope of the measured non-resonant chip power above 4 GHz shown in Fig. 11 confirms that the lot is at least one sigma fast, and most likely two sigma fast, which results in reduced power savings in the resonant chips. With nominal hardware, simulation shows that the expected power savings would have been closer to 10% of total chip power between 4 to 5 GHz.

Since chip operation and hardware power savings appeared to improve with increased frequency, while low frequency operation was limited, a simulation study was done to study the effects of tuning a clock distribution for different frequencies. Fig. 12 shows the predicted global clock power vs. frequency for a hypothetical design where the inductors and buffers were tuned for either 3 GHz or 3.7 GHz. For both design frequencies, the operating frequencies that resulted in good signal quality are marked with symbols on the curves. Good signal quality for this simulation study was defined as a single 20% to 80% transitions in less than 40 ps with no glitches entering this range of 20% to 80% Vdd. Note that tuning for lower frequency extends the range of good signal quality to lower frequencies, but as expected slightly reduces the power savings achieved at higher frequencies. Note also that good signal quality was achieved at all higher frequencies (up to 5 GHz) regardless of the design frequency.

V. CONCLUSION

In this paper, we describe how the Cell/B.E. processor was modified to demonstrate the first resonantly-clocked high-performance microprocessor. Power savings of approximately 5% of total chip power is realized while running actual workloads, and full functionality is achieved between 1.6 and 5 GHz. Unfortunately, the fast hardware fabricated (1–2 sigma fast) reduced the measured power savings by roughly a factor of two near the resonant frequency. Resonant modules were running after a few hours of test software debug. Linux boot was achieved using resonant Cell/B.E. chips and an architecture test suite program was executed successfully. Low frequency operation issues remain a concern in the resonant-load global clocking scheme because of double switching at the local clock buffers due to frequency mismatch between the LC network and the input clock frequency. Future implementations will require a switch to disable the resonance for low frequency test and debug. Modeling shows that resonant clocking becomes more attractive at clock frequencies above 3 GHz using the integrated inductors described in this paper. To achieve significantly greater power savings, and jitter reduction, and extend these benefits to below 2 GHz, would require the design and integration of higher-Q inductors. This enhancement can be realized either by dedicating chip area to the inductors, modifying the nearby power grid to remove eddy-current loops, or other inductor innovations.

ACKNOWLEDGMENT

The authors would like to thank the Cell/B.E. designers of the Sony Toshiba IBM design center, as well as many IBM personnel at three other sites who were generous with their time and expertise.

REFERENCES

- [1] 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.
- [2] S. Chan, P. Restle, T. Bucelot, S. Weitzel, J. Keaty, J. Liberty, B. Flachs, R. Volant, P. Kapusta, and J. Zimmerman, "A resonant global clock distribution for the Cell broadband engine processor," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 512–513.
- [3] L.-M. Lee and C.-K. Yang, "An adaptive low-jitter LC-based clock distribution," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 182–183.
- [4] F. O'Mahony, C. P. Yue, M. A. Horowitz, and S. S. Wong, "A 10-GHz global clock distribution using coupled standing-wave oscillators," *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 1813–1820, Nov. 2003.

- [5] J. Wood, T. C. Edwards, and S. Lipa, "Rotary traveling-wave oscillator arrays: A new clock technology," *IEEE J. Solid-State Circuits*, vol. 36, no. 11, pp. 1654–1665, Nov. 2001.
- [6] A. J. Drake, K. J. Nowka, T. Y. Nguyen, J. L. Burns, and R. B. Brown, "Resonant clocking using distributed parasitic capacitance," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1520–1528, Sep. 2004.
- [7] D. Pham *et al.*, "The design and implementation of a first-generation cell processor," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 184–185.
- [8] C. J. Anderson *et al.*, "Physical design of a fourth-generation POWER GHz microprocessor," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2001, pp. 232–233.
- [9] J. Clabes *et al.*, "Design and implementation of the POWER5 microprocessor," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2004, pp. 56–57.
- [10] J. Friedrich *et al.*, "Design of the Power6 Microprocessor," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 96–97.
- [11] P. J. Restle *et al.*, "A clock distribution network for microprocessors," *IEEE J. Solid-State Circuits*, vol. 36, no. 5, pp. 792–799, May 2001.



Steven C. Chan received the B.S. degree with high honors and the M.S. degree in electrical engineering and computer science from the University of California at Berkeley in 1996 and 1998, respectively, and the Ph.D. degree in electrical engineering from Columbia University in 2005. His doctoral work was supported by fellowships from IBM and the Semiconductor Research Corporation.

He is currently a technical manager in the Design and Technology Platform Division at TSMC Technology, Inc in San Jose, CA. His research interests include development of system-on-chip design methodologies, clock design and interconnect analysis, and electronic design automation. He has authored or co-authored more than a dozen technical papers and holds two U.S. patents related to the design of high-frequency low-power resonant global clock distributions. From 2005–2007, he was with the VLSI Design Department at the IBM T.J. Watson Research Center as a research staff member. While at IBM, Dr. Chan received the 2005 Pat Goldberg Memorial Best Paper Award in Electrical Engineering from the IBM Research Division. From 1997 to 2001, he was with CadMOS Design Technology, now part of Cadence Design Systems, where he and colleagues developed the first commercial signal integrity analysis tools, Pacific and CeltIC, for digital integrated circuits.



Phillip J. Restle received the B.A. degree in physics from Oberlin College in 1979, and the Ph.D. degree in physics from the University of Illinois at Urbana in 1986.

He then joined the IBM T. J. Watson Research Center as a Research Staff Member, where he initially worked on CMOS parametric test and modeling, CMOS oxide-trap noise, package testing, and DRAM variable retention time. Since 1993 he has concentrated on tools and designs for VLSI clock distribution networks contributing to more than a dozen server and game microprocessors, including all recent high-performance IBM servers such as the POWER6 processor, the z10 mainframe processor, the Xbox 360 processor, and the Cell Broadband Engine.

Dr. Restle received IBM awards for the Mainframe G4, G5, and G6 microprocessors, for the Power4 and Power5 microprocessors, for the PowerPC 970 used in the Apple G5 machine, as well as an IBM corporate award for VLSI clock distribution design and methodology. He received the 2005 Pat Goldberg Memorial Best Paper award. He holds 11 patents, has written 21 papers, and has given keynotes, invited talks, and tutorials on clock distribution, high frequency on-chip interconnects, and technical visualizations in VLSI design.



Thomas J. Bucelot received the B.S. in physics from the State University of New York, College at Cortland, in 1969 and the Ph.D. degree in low temperature physics from the University of Virginia, Charlottesville, VA, in 1979.

He is a Research Staff Member and Manager of the Design Systems Group at the IBM Thomas J. Watson Research Center. He joined IBM Research in 1979 to work on the Josephson super conducting computer project. From 1983 to 1992, he worked with and managed a group doing research in device and process measurements for liquid nitrogen and room temperature optimized CMOS technologies. In 1992, he joined the VLSI Design Department. His current interests include VLSI Design Tool Development for high performance, low power microprocessors, power characterization, global clocks, and improved design efficiency and data analysis utilizing databases and integrated design environments.



John S. Liberty received the B.S. and M.S. degrees in electrical engineering from North Carolina State University, Raleigh, NC.

He is an advisory engineer in the Sony, Toshiba, and IBM (STI) design center. He is responsible for helping architect and develop the Cell Broadband Engine. He main focus within the Cell BE is the SPU SIMD processor. He was a lead in designing the SPU Channel interface and the SPU's inherent security architecture. Before working on Cell BE, he was a designer working on Graphic Processing Units. He has

6 patents granted with more than 10 patents pending and has coauthored 4 technical articles.



Stephen Weitzel (M'75) received the B.S. degree in electrical engineering from Pennsylvania State University, University Park.

He joined IBM as a Test Equipment Engineer in 1974 and held various positions in test engineering and circuit design in the East Fishkill development center, Noyce design center, Somerset design center, and STI design center. He is currently a Senior Technical Staff Member working on high-frequency clock distributions for IBM in the high performance microprocessor center, Austin, TX. He has 11 patents and

has co-authored nine papers.



John M. Keaty received the B.S. degree in mathematics from the State University of New York at Plattsburgh in 1974, the M.A. degree in mathematics in 1977 and the M.S. degree in computer science in 1980, both from the University of Wisconsin-Madison.

He then joined the IBM Microelectronics Division in Burlington, Vermont to work on automated diagnostic systems for semiconductor logic products. Since that time he has worked in ASIC Product Development on several CMOS logic families and also on Industry Standard (x86) Microprocessor Development in the IBM Microelectronics Division before transferring to the IBM Server Group in Austin, Texas in 1996. There he was responsible for timing of the Power4 processor in the R/S 6000 Regatta system. He then went on to lead the Chip Integration of the Cell processor for the Sony, Toshiba, IBM (STI) Design Center. He is currently a Senior Technical Staff Member in the IBM Systems and Technology Division responsible for global integration of the next generation Power7 processor.



Brian Flachs received the B.S.E.E. degree in 1988 from New Mexico State University and the M.S. and Ph.D. degrees in 1994 from Stanford University where he developed interests in computer architecture, image processing and machine learning.

He served as architect, microarchitect and unit logic lead for the SPU Team and is interested in low latency-high frequency processors. Previously serving as microarchitect for IBM Austin Research Laboratories' 1 GHz PowerPC Project.



Peter Kapusta joined IBM in 2005 after receiving the B.S.E.E. degree from the University of Vermont. He currently works as a Product Engineer in support of the Cell BE microprocessor.



Richard Volant has over 20 years experience in semiconductor fabrication through IBM T. J. Watson Research Center, IBM's 200 mm Advanced Silicon Technology Center and the 300 mm fab at Hudson Valley Research Park in East Fishkill, NY. Projects included thin film metals, electron beam lithography and SiGe. Most work has focused on advanced interconnect technology, integrated passives as well as MEMS (Micro Electro Mechanical Systems). He currently holds over 55 patents.



Jeffrey S. Zimmerman received the B.S. degree in electrical engineering from the Pennsylvania State University, University Park, PA, in 1989.

He joined IBM in East Fishkill, NY, in 1989 where he design BJT memories. In 1994 he moved to IBM in Burlington, VT where he worked on IO design, chip integration, and yield improvements for the PowerPC processors. Since 2006 he has been developing IP for the IBM Foundry. He has coauthored three papers and holds 8 patents.