

## 20.3 A Double-Precision Multiplier with Fine-Grained Clock-Gating Support for a First-Generation CELL Processor

J.B. Kuang<sup>1</sup>, T.C. Buchholtz<sup>2</sup>, S.M. Dance<sup>2</sup>, J.D. Warnock<sup>3</sup>, S.N. Storino<sup>2</sup>, D. Wendel<sup>4</sup>, D.H. Bradley<sup>1</sup>

<sup>1</sup>IBM, Austin, TX

<sup>2</sup>IBM, Rochester, MN

<sup>3</sup>IBM, Yorktown Heights, NY

<sup>4</sup>IBM, Böblingen, Germany

A double-precision multiplier for floating-point and media-streaming instructions in the first-generation CELL processor [1] on 90nm PD/SOI is reported. Multiplication by recoding and successive partial-product (PP) compression is completed in three 11FO4 cycles including merging with the aligner. Figure 20.3.3 shows the micro-architecture of the design. At 1.3V and 68°C, hardware runs at 4.76GHz (Fig. 20.3.1). The multiplier area is 0.19mm<sup>2</sup> including that of decoupling capacitors. Only regular-V<sub>t</sub> devices are used in consideration of variability, leakage, and scalability. Other noted high-speed design points in the 90nm technology are the single precision [2] and low FO4 double-precision [3] multipliers.

The first cycle starts with Radix-4 Booth logic whose inputs are two 53b operands. Booth circuits reduce the number of PP rows to 27. To minimize area and latch count, two levels of 3:2 compressions in transmission-gate (TG) style circuits are also performed in this cycle. Footless domino circuits are used for complex Booth encoding and muxing functions. Figure 20.3.4 depicts a pruned schematic diagram for the Booth encoder, Booth multiplier (MUX), and pulse-to-static converter latch.

Static cycle 2 and 3 start with low-latency pulse latches (12 unfolded and 8 folded PP rows, respectively) to maximize cycle-time utilization and minimize clock power. Cycle 2 contains third-level 4:2 compressors (CMPs) and fourth-level 3:2 CMPs. In the third cycle, the fifth-level 4:2 CMP outputs are merged with the outputs from aligners in the final 3:2 CMPs. To ensure noise immunity, no unbuffered TGs are used. Delay is reduced through customized connections between two compression levels such that the number of inversions in any given path is minimized. Interconnect penalties are minimized by splitting the wiring between the second (row folding wires) and third (buses over the aligner) cycle. Figure 20.3.5 shows exemplary 3:2 and 4:2 TG CMPs.

Input operand latches convert static inputs to clock-qualified signals for the domino stages. Booth encoders are placed in the central clock bay to minimize delay. Pulsed operand inputs to dynamic stages reduce contention current at various process and operating corners. The design tolerates 10% variation in system clock pulses, i.e., 40% evaluate or precharge duty cycle, thus enhancing the technology and frequency scalability. Besides PFET keepers for dynamic nodes, clock gated NFET keeper devices are incorporated to sustain the low state, thus allowing low-speed testing and operations under short evaluate pulse conditions. Additionally, a pulse limiter on the clock grid limits evaluation time to 20FO4 at long cycle time. This avoids keeping dynamic nodes in the evaluate state for long periods of time. Higher leakage and smaller keepers can thus be tolerated without failure. Long Booth-encoder output wires and ladder-style Booth MUX input connections are shielded from noise. Dynamic output signals are converted to static ones with a mid-cycle converter latch whose input clock is delay interlocked with the Booth MUX precharge clock to prevent early-mode race conditions. The converter latch also drives a subsequent 3:2 CMP.

A 2b sign extension on each PP row reduces latency by avoiding the slower full 5-way recoding. Correct results are generated for

all but one erroneous carry-out at MSB. Data inputs are assigned to the Booth MUX for the appended sign bit with constants so that the result becomes inverted (0 for add and 1 for subtract). Exclusive use of 4-way Booth MUX cells enables uniform cell tiling that minimizes wire length, delay, and area. A custom 3:2 TG CMP, with one inverted input, is designed with the same image as that of the regular bits. The negate function is absorbed into the first XOR in the CMP by reordering inverter outputs. Bits associated with inverted input 3:2 CMPs are never adjacent to the clock bay. Therefore, the carry outputs do not suffer the additional wire penalty. Congestion in the clock bay is alleviated with the use of smaller and faster 4-output Booth encoders.

Local clock inaccuracies need to be easily absorbed into the timing budget for dynamic pipeline stages where the many parallel timing paths are physically distant. Clock- and Signal-edge control is important in attaining performance, noise margin, and design scalability. Clocks to Booth circuits need to be tightly correlated to minimize uncertainty. This is difficult due to the >270 $\mu$ m span of 27 Booth MUX rows (total bit width >400 $\mu$ m) and many mesh clock tap points. To ensure clock alignment and pulse width control within the granular dynamic circuits, 27 Booth MUX rows are divided into 3 macros, and the 9 Booth MUX rows in each macro into 3 local clock blocks. Within each 3-row block, clocks are derived from a single mesh tap point. Hence, the precharge clock for the Booth encoder has the same source as the clocks of the Booth are divided into and dynamic-to-static conversion latch for any PP row. No dynamic signal crosses different local clock domains, resulting in high confidence in clock and data alignment within each domain. Output signals from the 9 different clock meshes are static, thus more tolerant to skew and pulse-width variation. This technique enhances dynamic circuit robustness and timing precision. Figure 20.3.6 shows the placement of 9-row building block consisting of three 3-row dynamic circuit domains, skew-tolerant static circuit domains, and subsequent pulse latches that receive static inputs. Figure 20.3.6 insert depicts the local clock buffer where four clocks are derived from the same mesh tap point.

With a fixed bit-width dataflow image for area efficiency, the CMPs are folded into 2 physical rows at the fourth and subsequent levels. The MSB and LSB rows have a large bit index offset. As a result, timing of the MSB CMP carries into the LSB CMP is challenging, since it is in a different row and occupies an extreme position. In the long-wire point-to-point drive situation where the third level CMP outputs feed the fourth level inputs, wire loads make up a significant portion of the total load. Delay and slew are improved by replicating a destination CMP. The slight load increase is mitigated by using wide wires for long connections perpendicular to the dataflow. Six bits are replicated for each folded row.

Besides device sizing, cell placement, floorplan, stage shortening, and latch-count reduction for power management, this design supports fine-grained clock gating where the clocks for dynamic and static circuits can be controlled in 3 and 5 gating domains, respectively. Each gating domain can be independently configured from the control macro, allowing pipelined turn-on and minimizing disruption to the clock and power grid. Figure. 20.3.2 shows an example of active power management versus gating configuration.

### Acknowledgment:

The authors thank the management and technical support from the CELL design team, IBM E&TS and Research Division.

### References:

- [1] D. Pham et al., "The Design and Implementation of a First-Generation CELL Processor," *ISSCC Dig. Tech. Papers*, Paper 10.2, pp. 184-185, Feb., 2005.
- [2] S. Vangal et al., "A 5GHz Floating Point Multiply-Accumulator in 90nm Dual Vt CMOS," *ISSCC Dig. Tech. Papers*, pp. 336-337, Feb., 2005.
- [3] W. Belluomini et al., "An 8GHz Floating-Point Multiply," *ISSCC Dig. Tech. Papers*, Paper 20.1, pp. 374-375, Feb., 2005.

Heat sink temp set to 12C		
Vdd (V)	Freq (GHz)	Chip Temp (C)
1.3	4.76	68
1.2	4.60	52
1.1	4.32	41
1.0	3.74	32
0.9	3.14	27
Chip temp fixed at 85C (by varying heat sink control)		
Vdd (V)	Freq (GHz)	
1.3	4.60	
1.2	4.36	
1.1	4.06	
1.0	3.55	
0.9	3.00	

Figure 20.3.1: Frequency vs. supply voltage during daxpy benchmark executions.

Power (mW)	25% CG 10% SF	25% CG 30% SF	50% CG 10% SF	50% CG 30% SF	No CG 10% SF	No CG 30% SF
Dynamic Domain 1 and Static Domain 1 (Top portion of cycle 1 circuits)	24	29	46	56	91	109
Dynamic Domain 2 and Static Domain 2 (Mid portion of cycle 1 circuits)	24	29	46	56	91	109
Dynamic Domain 3 and Static Domain 3 (Bot portion of cycle 1 circuits)	24	29	46	56	91	109
Static Domains 4 and 5 (Cycle 2 and 3 circuits)	28	71	34	80	48	99
Total Active Power	100	158	172	248	321	426

Figure 20.3.2: One exemplary configuration showing active power vs. clock gating for dynamic and static clock domains at 4GHz based on hardware calibrated models.

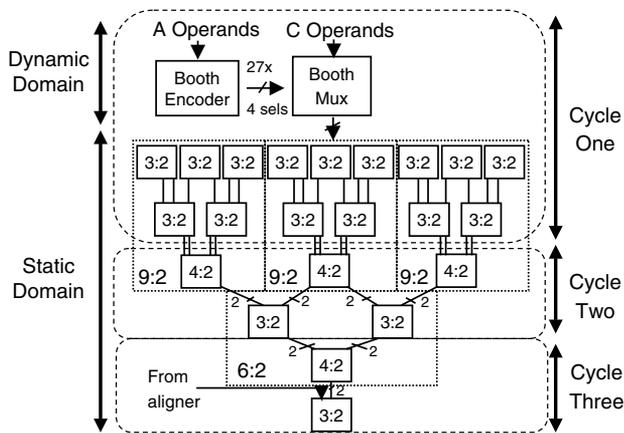


Fig. 20.3.1 Architectural diagram of the multiplier Wallace tree structure

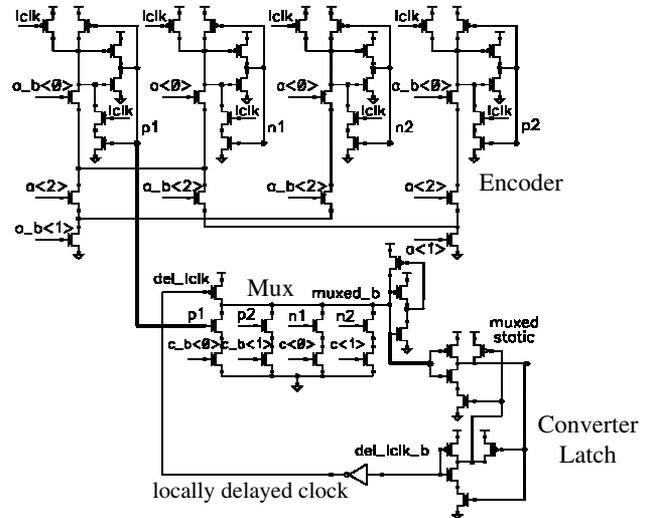


Figure 20.3.4: Pruned schematic diagram for Booth encoder, Booth mux, and cross-coupled NAND2 converter latch.

Figure 20.3.3: Architectural diagram of the multiplier Wallace tree structure.

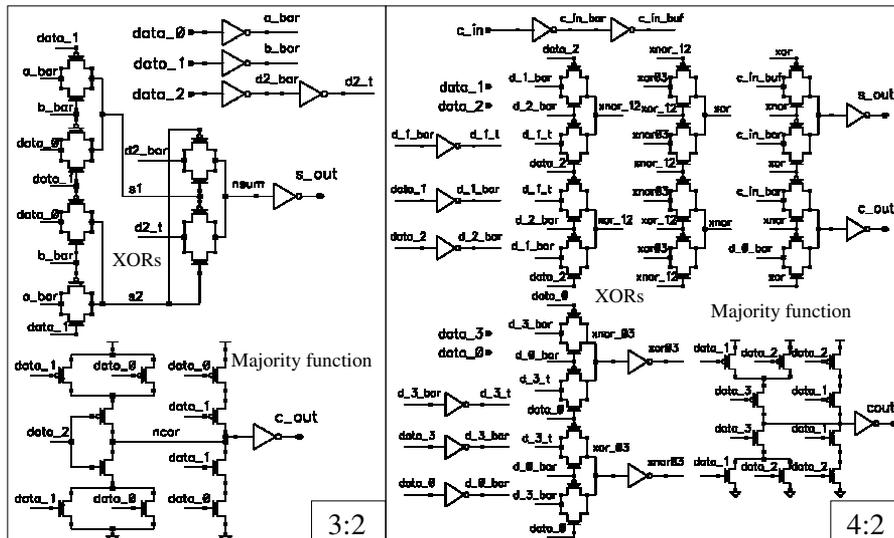


Figure 20.3.5: Exemplary compressor circuits.

Continued on Page 605

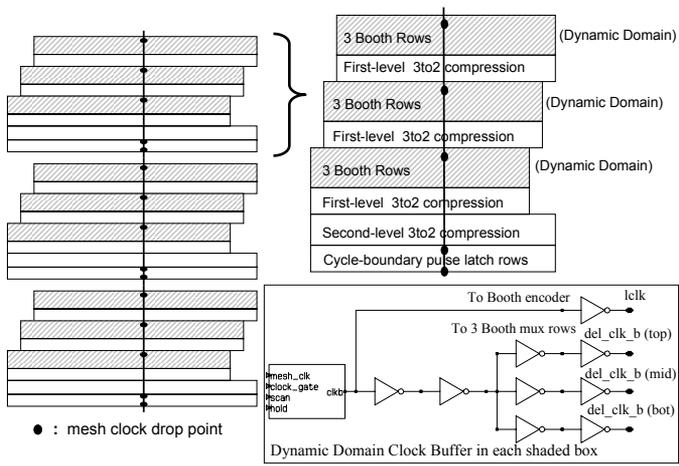


Figure 20.3.6: Local clock arrangement.

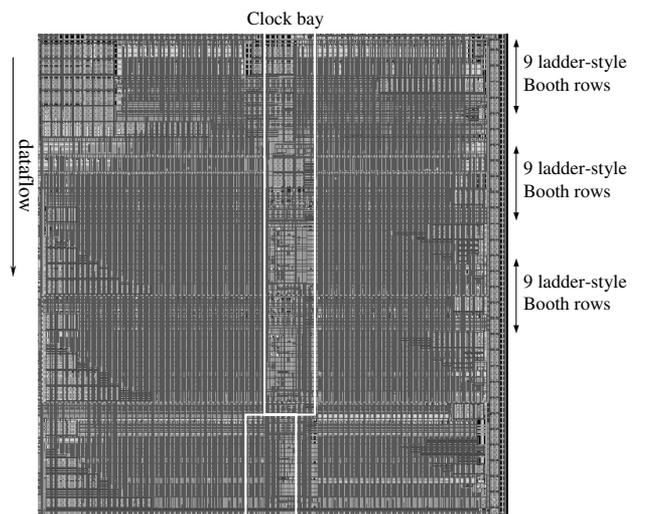


Figure 20.3.7: Physical view of the multiplier. Decoupling capacitors are populated outside trapezoidal areas and inserted at the vicinity of dynamic stages and input latches for power grid robustness.