

8.4 The Clock Distribution of the Power4 Microprocessor

Phillip J. Restle, Craig A. Carter², James P. Eckhardt³, Byron L. Krauter², Bradley D. McCreddie², Keith A. Jenkins, Alan J. Weger, Anthony V. Mule⁴

IBM Research, Yorktown Heights, NY, ²IBM Enterprise Systems Group, Austin, TX, ³IBM Enterprise Systems Group, Poughkeepsie, NY, ⁴Georgia Institute of Technology, Atlanta, GA

Considering the complexity of this 174M-transistor dual-processor chip, the clock distribution is relatively simple. A single chip-wide clock domain is used, with no active or programmable skew-reduction circuitry. Figure 8.4.1 shows a schematic diagram of the clock generation and distribution. Figure 8.4.2 shows a 3D visualization of every wire and buffer, where the Z-axis represents delay. A single PLL is used near the center of the chip to minimize the global clock distribution delay. An analog power supply is generated on chip for the analog PLL circuits using a capacitor mounted on the chip module. The PLL oscillator typically runs at a programmable 8 times the reference clock frequency, and twice the chip clock frequency, and is divided by 2 to generate a 50% duty-cycle global clock.

In the system, 10-hour maximum jitter measurements show the PLL produces <5ps cycle compression (jitter) from the PLL and 30ps total compression from the PLL plus clock distribution. The chip-to-chip jitter, defined as the maximum phase difference between clock edges on different processor chips, is <120ps.

The design of the PLL is SOI-specific, and requires several modifications to achieve the same chip to chip jitter as PLLs fabricated in a similar bulk CMOS process, as shown in Figure 8.4.3. The main focus of the SOI design is to add body contacts to the analog circuits, while minimizing both the body resistance and the device width uncertainty that accompany a body-contacted device. In addition, the capacitor structures used in the PLL are modified to maintain leakage levels acceptable for correct circuit operation. The SOI capacitor structures are improved by adding a thicker oxide structure, and strategic layout of the capacitor recessed-oxide openings.

The basic strategy chosen for the global clock distribution is an extension of a previously successful topology using a number of tuned trees driving a single full-chip clock grid [1]. The methodology is upgraded significantly to handle the increased complexity and speed of this chip. The priorities of schedule and performance force the global clock distribution to be designed, routed, and tuned with minimal impact on the turnaround time required for macro design, unit integration, chip integration, chip timing, late bug-fixes, and hardware bring-up. However, no compromise on performance is acceptable.

A strategy with multiple global clock domains would have allowed active or programmable de-skewing during operation. Having multiple domains would also allow coarse clock gating, and could result in lower skew within each small domain. Inevitably, however, with multiple domains there is increased skew and uncertainty between domains. In addition, multiple clock domains complicate early and late-mode timing, and degrade critical paths that cross multiple domain boundaries. Extensive simulations of the Power4 chip and test-chip hardware measurements support the simplifying decision to maintain a single domain global clock grid for the entire chip, with no programmable or active de-skewing. Figure 8.4.4 shows the smooth grid skew modeled from a worst-case across-chip buffer-

delay variation. While across-chip process variations causes skew as large as 70ps simulation in early test-chip hardware [2], even this skew causes no functional problems and insignificant performance impact because of the unique waveform smoothing characteristics of the grid driven at many points by tuned tree networks. This global clock strategy keeps the local skew within acceptable bounds even with large across-chip variations.

To achieve a methodology with minimum impact on schedule, a combination of top-down and bottom-up design strategies is used. The global clock-routing from the PLL to the single clock grid proceeds top-down, with all wiring and trees designed and optimized at the chip level. Maximum (before tuning) contract blockages are passed down to functional units. Within each of the 6 functional units in each core, the detailed clock routing (twig routing) proceeds from the bottom-up to connect to the single grid.

Macro designers route the local clock buffers to the latches, and create clock pins very near each local clock buffer. To avoid poor clock pin placement, the design guidelines allow only very short clock wires in macros and units, which results in 15,200 global clock pins at the chip level. While this complicates global distribution, it allows a chip-level wiring tool to optimize the wiring from the grid to each pin. Finally, the detailed load capacitances are passed up and used to tune the 64 trees driving the grid at 1024 points [1,3].

Figure 8.4.5 shows global clock waveforms measured using Picoprobes at 19 places showing 25ps maximum measured skew. Picosecond imaging for circuit analysis (PICA) measurements are also conducted to study packaged chips [4]. Light emission from 9 of the 64 sector buffers (buffer level 4 in figures 8.4.1,2) in different halves of the chip is measured. The results shown in figure 8.4.6 confirm <18ps skew at the leading edge of the photon pulses. The larger skew in the falling edge is expected due to wiring and load variations on these buffers and is corrected by the tree tuning.

Performance and schedule challenges of the global clock for the Power4 microprocessor are met using an SOI-specific PLL and a simple streamlined global clock distribution methodology that produces a robust high-quality global clock requiring no adjustment circuitry for skew reduction.

Acknowledgments:

The authors thank James Warnock, Thomas Bucelot, Charles Kircher, David Heidel, Michael Rosenfield, Carl Anderson, Mark Papermaster and the entire Power4 team for making it work.

References:

- [1] P. J. Restle et al, "A Clock Distribution Method for Microprocessors", IEEE JSSC, vol. 36, pp. 792-799, May 2001.
- [2] B. D. McCreddie et al, "A 1 GHz Power4 Testchip Design" Hot Chips 11, Symp. On High-Perf. Chips, Palo Alto CA, Aug. 1999.
- [3] P. Camporese, et al, US Patent 1998000222143, Issued March 20, 2001.
- [4] J. C. Tsang, J. A. Kash, and D. P. Vallett, Proc. IEEE, vol. 88, pp. 1440-1459, Sept. 2000.

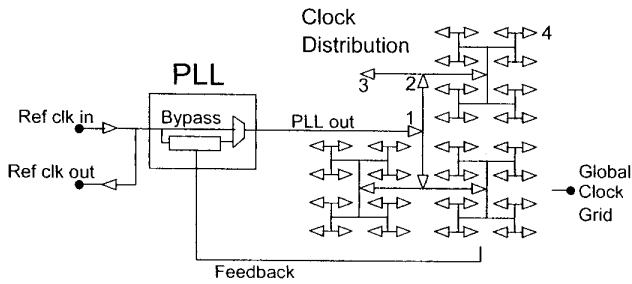


Figure 8.4.1: Schematic diagram with PLL and 4 buffer levels.




Figure 8.4.2: 3D visualization of the global clock distribution.

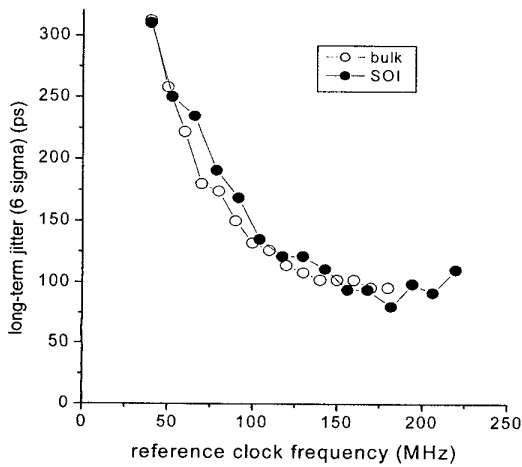


Figure 8.4.3: Measured PLL long-term jitter.




Figure 8.4.4: Clock distribution with buffer-delay variation.




Figure 8.4.5: Measured clock waveforms at 19 points on grid.




Figure 8.4.6: PICA waveforms.