16.1 Clock Distribution on a Dual-Core, Multi-Threaded Itanium®-Family Processor

Patrick Mahoney, Eric Fetzer, Bruce Doyle, Sam Naffziger

Intel, Fort Collins, CO

Clock distribution on the microprocessor, code-named Montecito, features four distinct segments (Fig. 16.1.1) and topologies, each tuned to a specific purpose. A region-based active de-skew (RAD) system reduces the process, voltage, and temperature sources of skew across the 21.5×27.7mm^2[1] die during normal operation. Clock vernier devices (CVDs) inserted at each local clock buffer allow 70ps of adjustment via scan. The system supports a constantly varying frequency and consumes less than 25W on its 30mm route from PLL to latch.

The first segment of the clock distribution is the level-0 (L0) route which connects the PLL to the digital frequency divider (DFD). The L0 is the only segment of the clock distribution that does not adjust supplies and frequencies during normal operation. The L0 route is 20mm long consisting of four 5mm segments that are 400mV low-voltage swing differential routes. Each segment is resistively terminated at the receiver and is tapered to optimize RLC flight time and reduce power consumption. All route segments are matched in composition in both layer and length. Modeling is performed using compact loop-based RLC model extracted from custom RC tools and Raphael 3D simulation [2]. The differential nature of the route makes return-path modeling simple since all of the currents are equal and opposite. A two-stage self-biased differential amplifier (Fig. 16.1.2) is used to restore the 400mV output differential to full swing at each of the three repeaters along the route. At the DFD the route is received by a self-biased amplifier and buffered into the DLL. Self-biased amps are used in order to prevent the accumulation of common mode offsets from creating duty-cycle errors.

The second segment of the route, the level-1 (L1) route, connects the DFD to the second level cock buffer (SLCB). The DFD output varies in frequency and it operates on a varying core supply voltage. A half-frequency distribution using differential 0^0 and 90^0 clocks is used. Even at half frequency, the 90nm process technology did not support wires wide enough to achieve the full level swings required to maintain a proper duty cycle in the changing-frequency environment. In parallel with the receiver, Montecito utilizes a feedback circuit (Fig. 16.1.3) that drives the end of the wire to the supply rail based on the state of the opposite phase clock. The 0^0 and 90^0 clocks are multiplied in the SLCB receiver, with a set of self-biased amplifiers that are configured as an XOR gate. The resulting signal is a full-rate clock with no history effect due to the frequency transitions of Foxton[1] or on-die clock shrinking.

The third segment of the route, the level-2 (L2) connects the SLCB to the local clock buffers. Many contemporary designs use a grid-based clock network for this stage. While grid designs create impressively low skew numbers, the design has two major drawbacks: metal resources and power. A skew-matched RLC tree network technique is employed using an in-house tool [3]. Unit designers supply the locations of CVDs to the distribution team and locations and loads are entered into the tool. The distribution engineer takes end loads and routes the clock avoiding obstacles as necessary. The optimization engine of the in-house tool then balances the distribution using width and space to match route RLC delays. A typical SLCB drives 400 local clock buffers called clock vernier devices (CVDs) at 200 different locations across 3mm with a skew of less than 6ps between locations. The result is a clock route that is adaptable to changes in the design, and uses far less metal resources than a grid-based design and uses a minimum of power while achieving skews that are nearly as low as grid-based designs.

With 28mm of wire length and 3 repeaters, a frequency divider, and a delay line between the PLL and the gaters, even a perfectly matched system will have significant skews caused by process, temperature non-uniformity, and voltage variation across the die. Montecito implements a RAD system to null out offsets caused by these variations. RAD relies on a hierarchical collection of phase comparators between the ends of different L2 routes (Fig. 16.1.4). Each SLCB has a 128b delay line with 1ps resolution. With one zone locked as the reference, phase-comparator results are forwarded to the SLCB farthest from the reference zone. The SLCB takes the result and adjusts its delay line accordingly. This system runs continuously to null out skews in the distribution caused by temporal events like temperature and voltage variation. This is a significant improvement over previous designs which performed de-skewing statically or prior to the completion of reset. Worst-case skews for distributions to this point are now limited even under the most extreme conditions to the 10ps accuracy of the phase-comparator/delay line loop.

The fourth and final segment of the clock route, the post-gater route, is in the hands of the individual circuit designers. Clock gaters are designed by the clock team into the library in a variety of sizes. Using a dedicated in-house tool based on SPICE, circuit designers match gater loads to the load of the latches and routes that they drive. With hundreds of latches per gater, routes up to 2mm long must be engineered for delay, shielding, and load matching. Unlike the previous three stages of distribution, routes here are not corrected by the RAD system.

To facilitate post-silicon debug and to remove skew not found in pre-silicon analysis, CVDs (Fig. 16.1.5) are inserted into each gater cluster. CVDs are controlled via scan operations and can add 70ps of delay to any clock in 8ps increments. CVDs have proven useful for finding and working around post-silicon speed-path defects. The CVDs use an attenuated drive fight to control propagation delay. The default setting is for the CVD to have no drive fight and behave like a buffer. To insert delay the scan chain asserts a 3b signal that sets the delay which causes the feedback drive to fight against the input slowing down the propagation to the output. CVDs can be programmed via scan and firmware making clock edits available in both debug and during production.

The 4-segment clock-distribution of Montecito delivers a variable-frequency clock from 100MHz to 2.5GHz over a clock network over 28mm long (Fig. 16.1.6). The route is optimized for low skew and uses active deskewing and scan-chain adjustments to reduce the total clock-network skew to less than 10ps. The clock network implements a tree-based network that reduces power and metal utilization over conventional designs.

Acknowledgements:
The authors wish to acknowledge the efforts of Steve Hall, Erin Francom, and the others whose contributions made this design possible.

References:
[1] S. Naffziger et al., "The Implementation of a 2-core, Multi-threaded Itanium®-Family Processor," *ISSCC Dig. Tech. Papers*, Paper 10.1, pp. 182-183, Feb., 2005.
[2] Raphael User Manual, Avant! Corporation
[3] Ferd Anderson et al., "The Core clock System on the Next Generation Itanium® Microprocessor," *ISSCC Dig. Tech. Papers,* vol. 1, pp. 146-148, Feb., 2002.

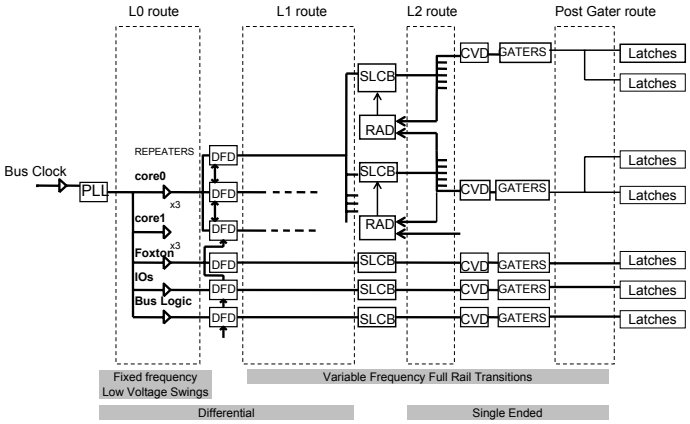

Figure 16.1.1: Clock distribution.

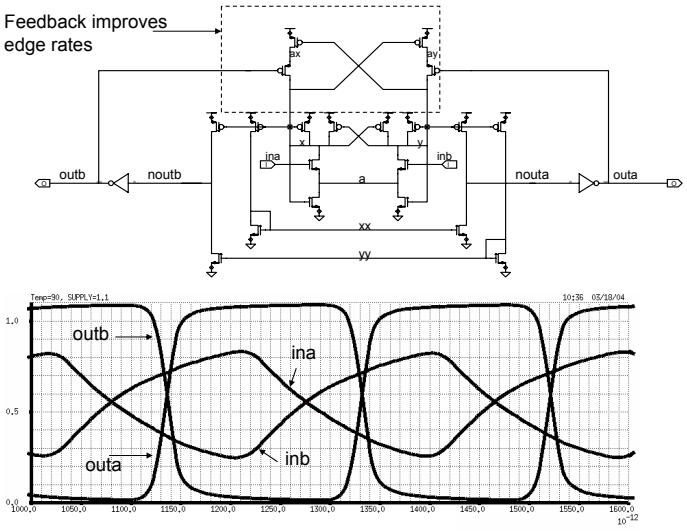

Figure 16.1.2: Level-0 route-repeater design.

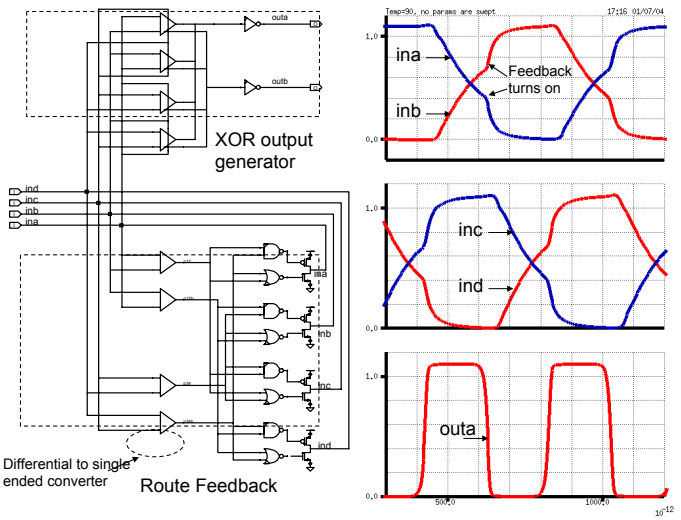

Figure 16.1.3: L1 route topology.

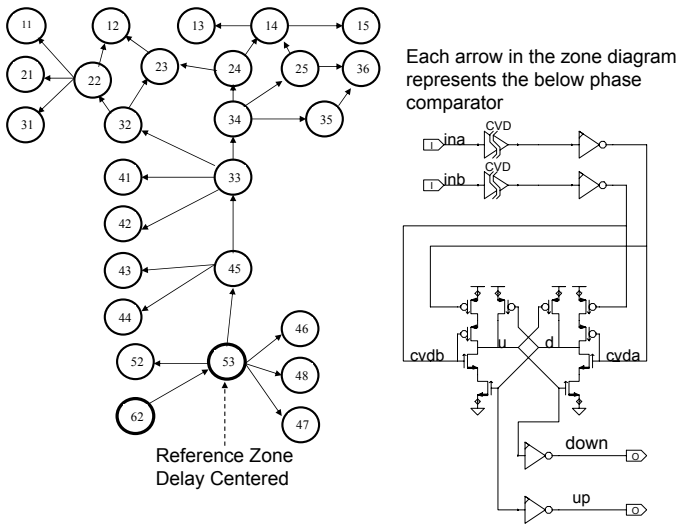

Figure 16.1.4: RAD phase comparator and connectivity.

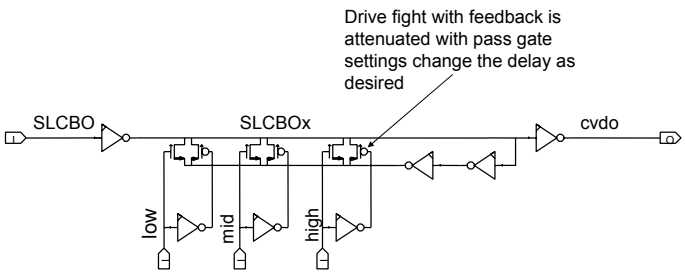

Low, mid, and high are scan chain outputs.

Their values and corresponding delays are binary weighted.

Figure 16.1.5: CVD design.

Route	Terminals	Distance	Power	Delay
L0	14	20mm	600mW	640ps
L1	71	5mm	700mW	215ps
L2	14500	2.0-3.5mm	3.3W	60ps
Post Gater	~5million	0-1.5mm	20W	12ps

Figure 16.1.6: Route statistics.

Continued on Page 599

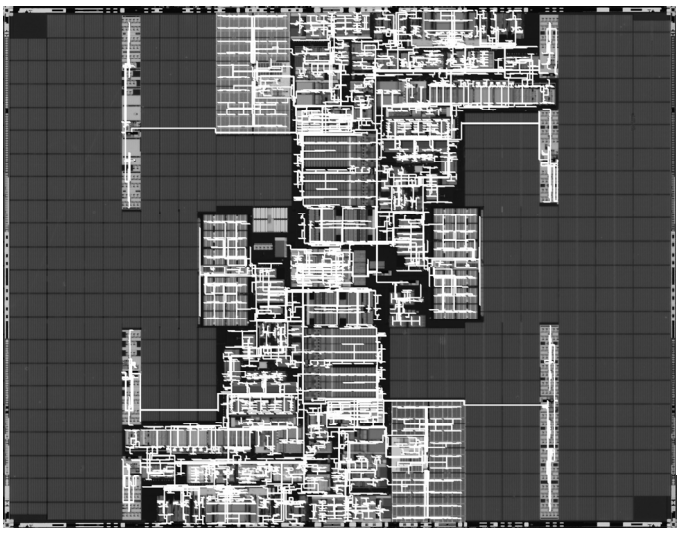

Figure 16.1.7: Die micrograph showing clock distribution.

