

16.2 A 90nm Variable-Frequency Clock System for a Power-Managed Itanium®-Family Processor

Tim Fischer, Ferd Anderson, Ben Patella, Sam Naffziger

Intel, Fort Collins, CO

The Montecito[1] processor contains two Itanium-family cores with Foxtan[2] technology on a 1.7B-transistor die. The variable-frequency clock system (Fig 16.2.1) consists of a single PLL [3] that generates a multiple $6 \leq M \leq 20$ of the system clock frequency distributed to 14 digital frequency dividers (DFDs) for division to the proper zone frequency. Each DFD (Fig. 16.2.2) consists of a DLL and a state machine that dynamically selects among 64 DLL phases generated from the PLL clock. This allows the DFD output frequency to vary according to $F_{DFD} = F_{PLL} / (1 + D/64)$ where $0 \leq D \leq 63$, yielding a range of 1.0 to $0.504 F_{PLL}$ in $1/64^{\text{th}}$ increments.

Clock zones consist of 2 cores each with 3 DFDs; one 1GHz DFD for Foxtan technology control; one DFD for each of 6 front-side bus (FSB) stripes; one DFD for bus logic. Each DFD output clock is distributed to second level clock buffers (SLCBs) for delay tuning to 1ps resolution via active deskew. 35 regional active deskew (RAD) phase comparators are distributed in each core to actively deskew neighboring SLCBs, yielding <10ps of skew across the 21.5mm by 27.7mm die [4]. SLCB clocks are distributed to 7536 clock vernier devices (CVDs) per core for local delay fine-tuning via scan. Gaters provide a final gain stage, power-saving enables, and pulse shaping for low latching overhead and skew compensation through transparency [5].

To configure the clock-system components, a "translation table" determines PLL, DFD, and aligner divisors from pin-selected system clock frequencies (200, 266, 333, 400MHz) and fuse-selected bus-logic and core clock frequencies. Fuses determine the core startup ("safe") and limit frequencies. The clock system has two frequency modes: *fixed* and *variable* (FFM and VFM, respectively). The clock system starts in FFM and is then placed into VFM by firmware.

In FFM, 13 of 14 DFDs are frequency and phase aligned; the 14th is always a fixed 1GHz for Foxtan-technology power-management algorithms. The 13 aligned DFDs have identical fixed divisors: 0 for maximum FFM frequency, and > 0 to achieve a "safe" startup frequency before entering VFM. On power-up or reset, DFD DLLs start and lock on the PLL clock autonomously. Once all 13 DLLs lock, DFD dividers start synchronously and remain phase/frequency locked to the PLL clock.

In FFM, the core, FSB, and bus logic clocks align to the external system clock by a phase aligner system. This aligner adjusts DFD phase selection using up/down controls, sliding the phase around without changing frequency. At startup the aligner eliminates built-in core/FSB route mismatch [4] and aligns both to the system clock to within 20ps across PVT. DFD clock synthesis allows phase adjustment in uniform $1/64$ cycle steps with virtually infinite range. In fact, an inversion error on first silicon in the bus logic clock tree (due to logic equivalence escape) is transparently corrected by the aligner at startup with no added skew or functional impact due to the adaptive design.

In VFM, core DFD frequency (F_{CORE}) dynamically tracks core voltage (V_{CORE}) via a programmed regional voltage detector (RVD) voltage-frequency (V-F) response in the voltage-to-frequency converter (VFC) loop. The RVD consists of a one-cycle delay-line with a programmable mix of RC and FET delay. This delay-line output is applied to a phase comparator to produce T_{CORE} adjust signals UP, DN (down) and DZ (deadzone). The DZ capability controls VFC loop stability. The RVD delay, its RC composition and the DZ width are all scan programmed at startup by hardware and system software.

High VFC bandwidth can track power-managed V_{CORE} modulation as well as high-frequency switching transients. A new frequency is selected with 1.5 cycle average response to a local voltage change event. This frequency change is distributed to latches in ~ 700 ps [4] (Fig. 16.2.3). In each VFC cycle, a DFD utility clock edge: 1) propagates 2400 μ m to an RVD; 2) a comparator produces an UP, DN, or DZ request; 3) routes 2400 μ m back to DFD; 4) PCSM arbitrates and resolves comparator meta-stability; and 5) produces a divisor adjust set up to next clock edge at the DFD. The high bandwidth greatly reduces CPU exposure to voltage-transient-induced timing issues enabling F_{CORE} to track a voltage transient of up to 30mV/ns with 700ps of lead time on average.

In VFM, DFDs synthesize an F_{CORE} range of F_{PLL} to $F_{PLL}/2$ in 1.6% steps over a V_{CORE} range of 0.8V to 1.2V. DFDs receive inputs from 4 local RVDs and from other same-core DFDs. The DFD phase compensator state machine (PCSM) arbitrates RVD requests and same-core DFD inputs to derive local DFD divisor adjusts which (a) preserve intra-core DFD phase lock and (b) track programmed V-F response. All DFDs start synchronously in safe mode using phase 0, and same-core DFD phase lock is maintained in VFM by PCSM arbitration.

Test features include: on-die clock shrink (ODCS) [6], clock-edge manipulation in the DFD; 4 self-calibrated salmon ladders for deterministic test trigger transport between clock domains; a 2-pin clock-observation test port.

Figure 16.2.4 shows simulated VFC response to a V_{CORE} transient. The fast VFC response allows reduction of frequency guard-banding normally used to insure critical path timing during supply transients. This increases VFM performance over FFM as a function of on-die supply noise (Fig. 16.2.5) which has been observed using an on-die power-measurement circuit [9] to be about 70mV_{pp}. Figure 16.2.6 shows silicon waveforms of core and bus-logic clocks at startup in FFM and later in VFM: the bus logic voltage and frequency remain fixed at 1.2V/1.6GHz while in this case the cores are running at 1.2V/2.14GHz.

Clock-system results on first silicon included full functionality of FFM, VFM, and all units described above. The clock system has been shown to operate at up to 2.5GHz at 1.2V, and enabled first silicon boot of Linux, HPUNIX, and Windows on multiple platforms with Foxtan technology.

Acknowledgements:

The authors recognize the dedicated efforts of a talented and many-skilled team in designing, verifying, and debugging the Montecito clock system.

References:

- [1] S. Naffziger et al., "The Implementation of a 2-core, Multi-Threaded 64b Itanium®-Family Processor," *ISSCC Dig. Tech. Papers*, Paper 10.1, Feb., pp. 182-183, 2005.
- [2] C. Poirier et al., "Power and Temperature Control on a 90nm Itanium®-Family Processor," *ISSCC Dig. Tech. Papers*, Paper 16.7, pp. 304-305, Feb., 2005.
- [3] K. Wong et al., "Cascaded PLL Design for a 90nm CMOS High-Performance Microprocessor," *ISSCC Dig. Tech. Papers*, pp. 422-424, Feb., 2003.
- [4] E. Fetzer et al., "Clock Distribution on a Dual-Core Multi-Threaded Itanium®-Family Processor," *ISSCC Dig. Tech. Papers*, Paper 16.1, pp. 292-293, Feb., 2005.
- [5] S. Naffziger, et al., "The Implementation of the Itanium2 Microprocessor," *IEEE J. Solid-State Circuits*, pp. 1448-1459, Nov., 2002.
- [6] S. Tam et al., "Clock Generation and Distribution for the First IA64 Processor," *IEEE J. Solid-State Circuits*, Nov., 2000.
- [7] F. Anderson et al., "The Core Clock System for the Next Generation Itanium Processor," *ISSCC Dig. Tech. Papers*, pp. 146-148, Feb., 2002.
- [8] S. Tam et al., "Clock Generation and Distribution for the Madison Processor," DTTC2002.
- [9] E. Alon et al., "Circuits and Techniques for High-Resolution Measurement of On-Chip Power Supply Noise," *Symp. VLSI Circuits*, pp. 102-105, June, 2004.

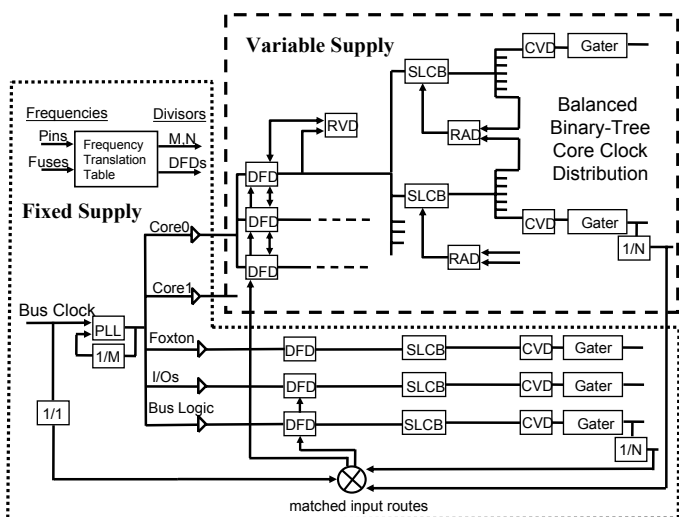


Figure 16.2.1: Montecito clock system topology.

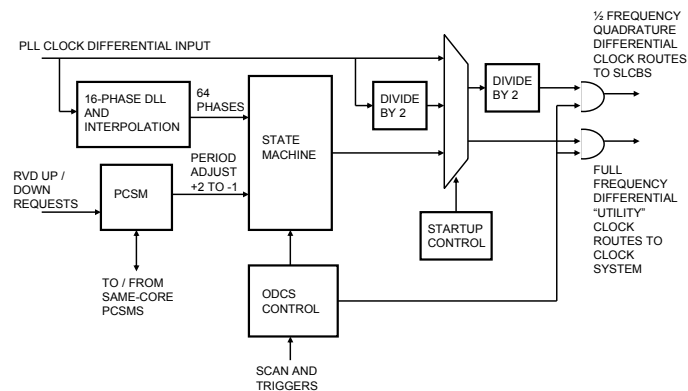


Figure 16.2.2: DFD / PCSM block diagram.

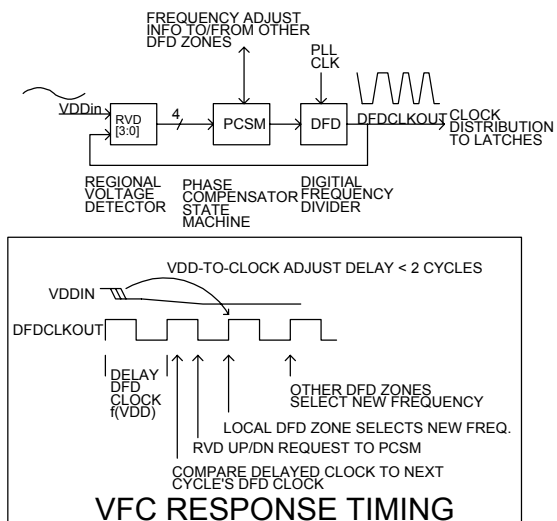


Figure 16.2.3: Voltage-to-frequency converter loop (VFC).

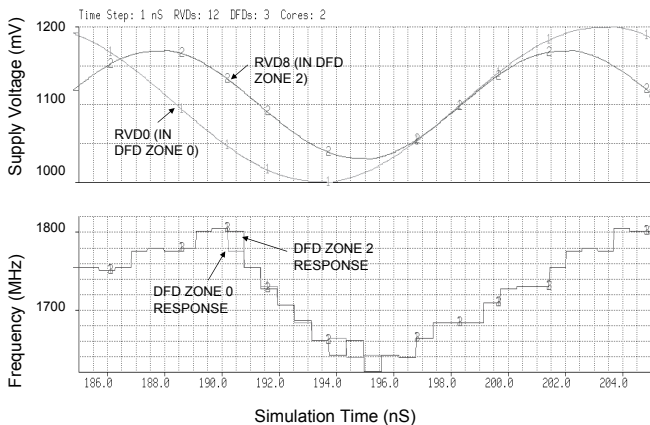


Figure 16.2.4: VFM supply transient tracking response (simulated).

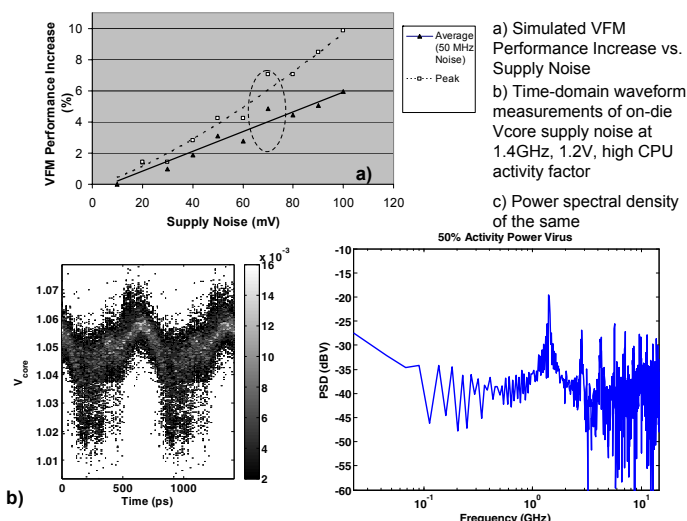


Figure 16.2.5: VFM performance versus supply noise.



Figure 16.2.6: FFM/VFM core/bus clock oscilloscope traces.

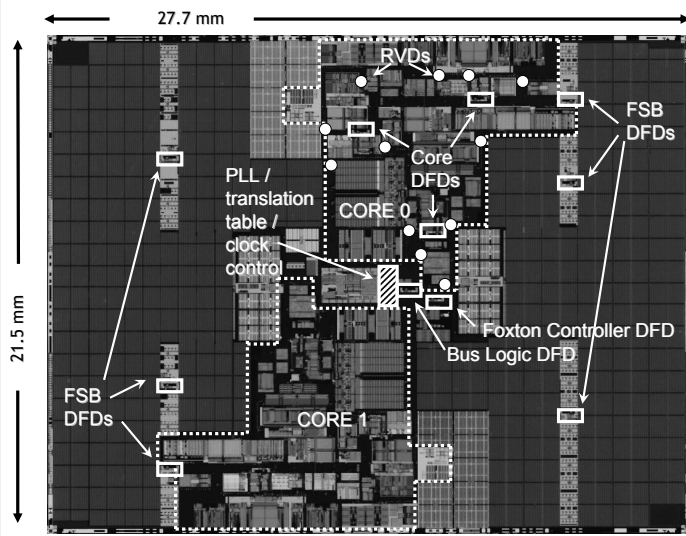


Figure 16.2.7: Die micrograph.