

TA 10.5 Active GHz Clock Network using Distributed PLLs

Vadim Gutnik, Anantha Chandrakasan

MIT Microsystems Technology Lab, Cambridge, MA

Most modern microprocessors use a balanced tree to distribute the clock [1]. However, at gigahertz clock speeds an increasing fraction of skew and jitter comes from random variations in gate and interconnect delay. The majority of jitter in a clock tree is introduced by buffers and inter-line coupling to the clock wires. A relatively small amount comes from noise in the source oscillator [2].

This distributed clock network generates the clock signal with phase locked loops (PLLs) at multiple points (nodes) across a chip, and distributes each only to a small section of the chip (tile) (Figure 10.5.1). Phase detectors (PD) at the boundaries between tiles produce error signals that are summed by an amplifier in each tile and used to adjust the frequency of the node oscillator.

With locally-generated clocks, there are no chip-length clock lines to couple in jitter; skew is introduced only by asymmetries in phase detectors instead of mismatches in physically separated buffers; and the clock is regenerated at each node, so high frequency jitter does not accumulate with distance from the clock source. Unlike earlier work on multiple clock domains which suggests use of multiple independent clocks, this approach produces a single fully-synchronized clock. This arbitrary network of tiles, each with its own PLL, is more general than previous active skew management approaches [3].

However, because there are many nodes, the power and size constraints on each element of a distributed clock generation architecture are even more stringent than the constraints on a single, global PLL. Furthermore, there must be a way to ensure that the multiple nodes get and stay synchronized. The oscillator, phase detector, and loop filter of a working demonstration chip, fabricated in a standard 0.35 μm , single-poly triple-metal process, are considered in turn below.

This chip uses an nMOS-loaded differential ring oscillator as a voltage-controlled oscillator (VCO) to minimize power supply noise (Figure 10.5.2). Transistors $M_4 - M_8$ comprise the differential inverter. The differential pair is $M_{5,8}$, the tail current is driven by M_6 and $M_{4,7}$ act as the nMOS load. The nMOS loads allow fast oscillation and shield the output signal from V_{DD} noise. V_{bias} is a low-pass version of V_{DD} generated by subthreshold leakage through pFET M_1 ; supply noise coupling in through C_{gs} of $M_{4,7}$ is bypassed by M_2 . The oscillation frequency is dependent on the supply voltage only through capacitor nonlinearity and the output conductance of $M_{4,7}$, and feedback of the PLL compensates drift of V_{DD} and V_{bias} .

Because phase is periodic, a set of oscillators need not all be in phase to each have zero net phase error. This phenomenon, modelock, is described in Reference [4], which notes that modelock can be avoided by using PDs whose response decreases monotonically beyond a phase difference of $\pi/2$, as from an XOR phase detector. Note that this solution precludes the use of a phase-frequency detector (PFD). Lack of a PFD is problematic because the capture bandwidth of a memoryless PLL is limited to a few percent of the center frequency, while the center frequencies of widely-spaced oscillators on a chip can easily vary by 10-20%.

The PD proposed here, shown in Figure 10.5.3, has sufficient nonlinearity, higher gain at small input phase difference and less high-frequency content than an XOR PD. The core ($M_1 - M_6$) is an nMOS-loaded arbiter which acts as a nonlinear phase detector. For no input phase difference, the output is balanced. As the phase

difference increases from zero, one output is asserted for the full duration of an input pulse, while the other output is asserted for only the remainder of the input pulse duration after the first input pulse ends, which is equal to the input phase difference. Thus the detector has high gain near zero phase error that drops off to zero as the input phase difference approaches the input pulse width (Figure 10.5.4).

The pulse generators P_1 and P_2 enable this arbiter to give frequency error feedback. If one input is at a higher frequency than the other, its output will be asserted for more input pulses than the other. Because the width of the pulses is independent of input frequency, the average output voltage corresponds to frequency. Unlike a typical phase-frequency detector, however, the strength of the error signal falls to zero as frequency difference goes to 0, so there can be no modelock problems, yet large signal frequency- (and hence, phase-) locking is enhanced. Figure 10.5.5 shows the large-signal correction and small-signal behavior of the entire array of PLLs as the already internally-locked array approaches and locks to the reference clock. The PD fits in 30x30 μm^2 .

One loop filter is associated with each VCO. To avoid the series resistor of a charge pump with passive RC compensation, a feed-forward compensation method is used. The loop filter of Figure 10.5.6 consists of two differential amplifiers. $M_1 - M_5$ make up amplifier A_1 , while $M_9 - M_{17}$ make up A_2 . The differential output currents from the PDs at the edges of each tile are summed at nodes $In+$ and $In-$, and drive both amplifiers. A_1 is a single-stage differential pair so it has relatively low gain but a bandwidth limited by g_m/C_{gs} . A_2 has a high gain cascoded stage driving a common source pFET M_{17} . M_{16} is a large gate capacitor which serves to set the dominant pole of M_2 such that the PLL network is stable. M_{15} is biased at low current to boost gain and enable time constant as low as 12kHz with a 15x15 μm^2 gate capacitor. The simple design and feed-forward compensation allow the loop filter to fit in only 15x45 μm^2 . Each clock node, consisting of an oscillator and a loop filter, takes just 45x45 μm^2 .

A chip was fabricated with a 4x4 array of nodes and PD between nearest neighbors. Counting one node and two PDs, the area overhead is approximately 0.0038mm² per tile. Another PD is between one of the nodes and the chip clock input to lock the network to an external reference. The output of the 16 oscillators is divided by 64 and driven off chip. At $V_{DD} = 3V$, the divided outputs are frequency-locked at 17 to 21MHz, corresponding to oscillator phase lock at 1.1 to 1.3GHz. An oscilloscope plot of four locked output signals is shown in Figure 10.5.7.

Long-term jitter between neighbors is less than 30ps rms. Cycle-to-cycle jitter is less than 10ps. The oscillators, amplifiers and all the biasing draws 130mA at 3V. A chip plot is shown in Figure 10.5.8. The rest of the area on the 3x3mm² chip is taken up by test circuits.

Design and measurement on this chip confirm that generating and synchronizing multiple clocks on chip is feasible. Neither the power nor the area overhead of multiple PLLs is substantial compared to the cost of distributing the clock by conventional means. Most importantly, a distributed clock network can take advantage of improved devices by shrinking the size of the cells, lowering the overall skew and jitter, so performance will scale with device speed, rather than with the much slower improvement of on-chip interconnect speed.

Acknowledgments:

The authors acknowledge support from the MARCO Focused Research Center on Interconnects funded at MIT through a subcontract from Georgia Tech. Vadim Gutnik was partly supported by a graduate fellowship from Intel Corp.

References:

- [1] Bailey, D. W. and B. J. Benschneider, "Clocking design and analysis for a 600 MHz Alpha microprocessor," *Journal of Solid State Circuits*, vol. 33, no. 11, pp. 1627-1633, November 1998.
- [2] Young, I. A., M. F. Mar, and B. Bhushan, "A 0.35μm CMOS 3-880MHz PLL N/2 clock multiplier and distribution network with low jitter for microprocessors," in *ISSCC Digest of Technical Papers*, February 1997, pp. 330-331.
- [3] Geannopoulos, G. and X. Dai, "An adaptive digital deskewing circuit for clock distribution networks," in *ISSCC Digest of Technical Papers*, February 1998, pp. 400-401.
- [4] Pratt, G. A. and J. Nguyen, "Distributed synchronous clocking," *IEEE Transactions on Parallel and Distributed Systems*, February 1995.




Figure 10.5.2: Ring oscillator schematic.




Figure 10.5.3: Phase detector.

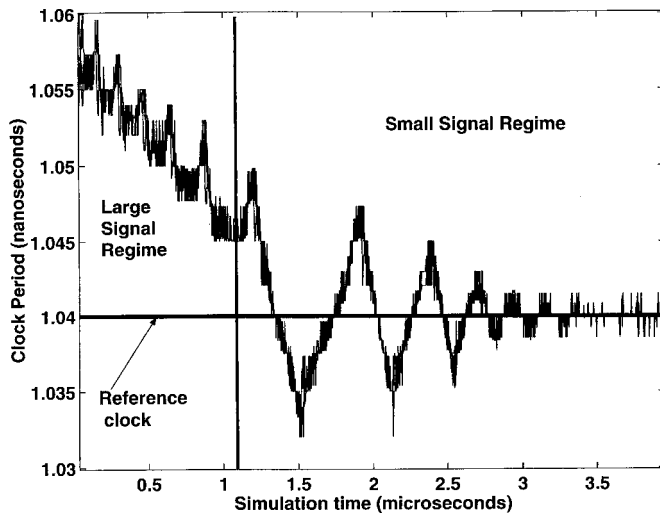


Figure 10.5.5: Locking behavior of the PLL array.

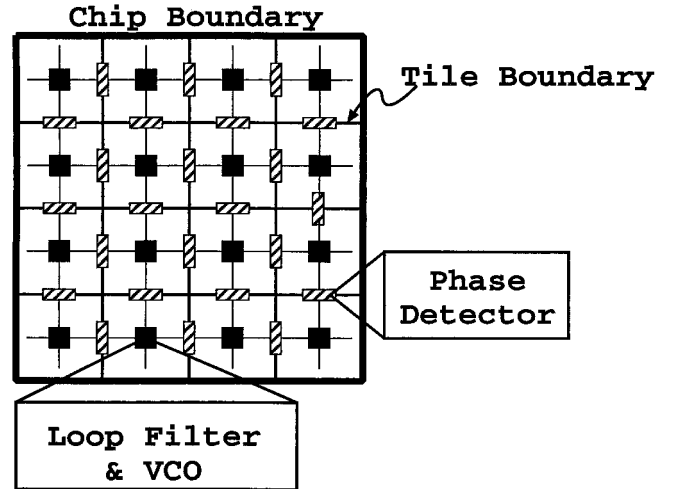


Figure 10.5.1: Distributed clocking network.




Figure 10.5.4: Simulated PD transfer curve.




Figure 10.5.6: Loop filter schematic.

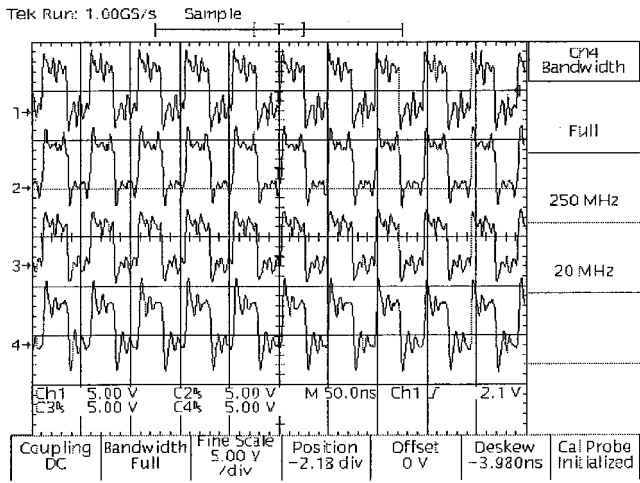


Figure 10.5.7: Frequency-locked divider outputs.




Figure 10.5.8: Distributed clock chip.

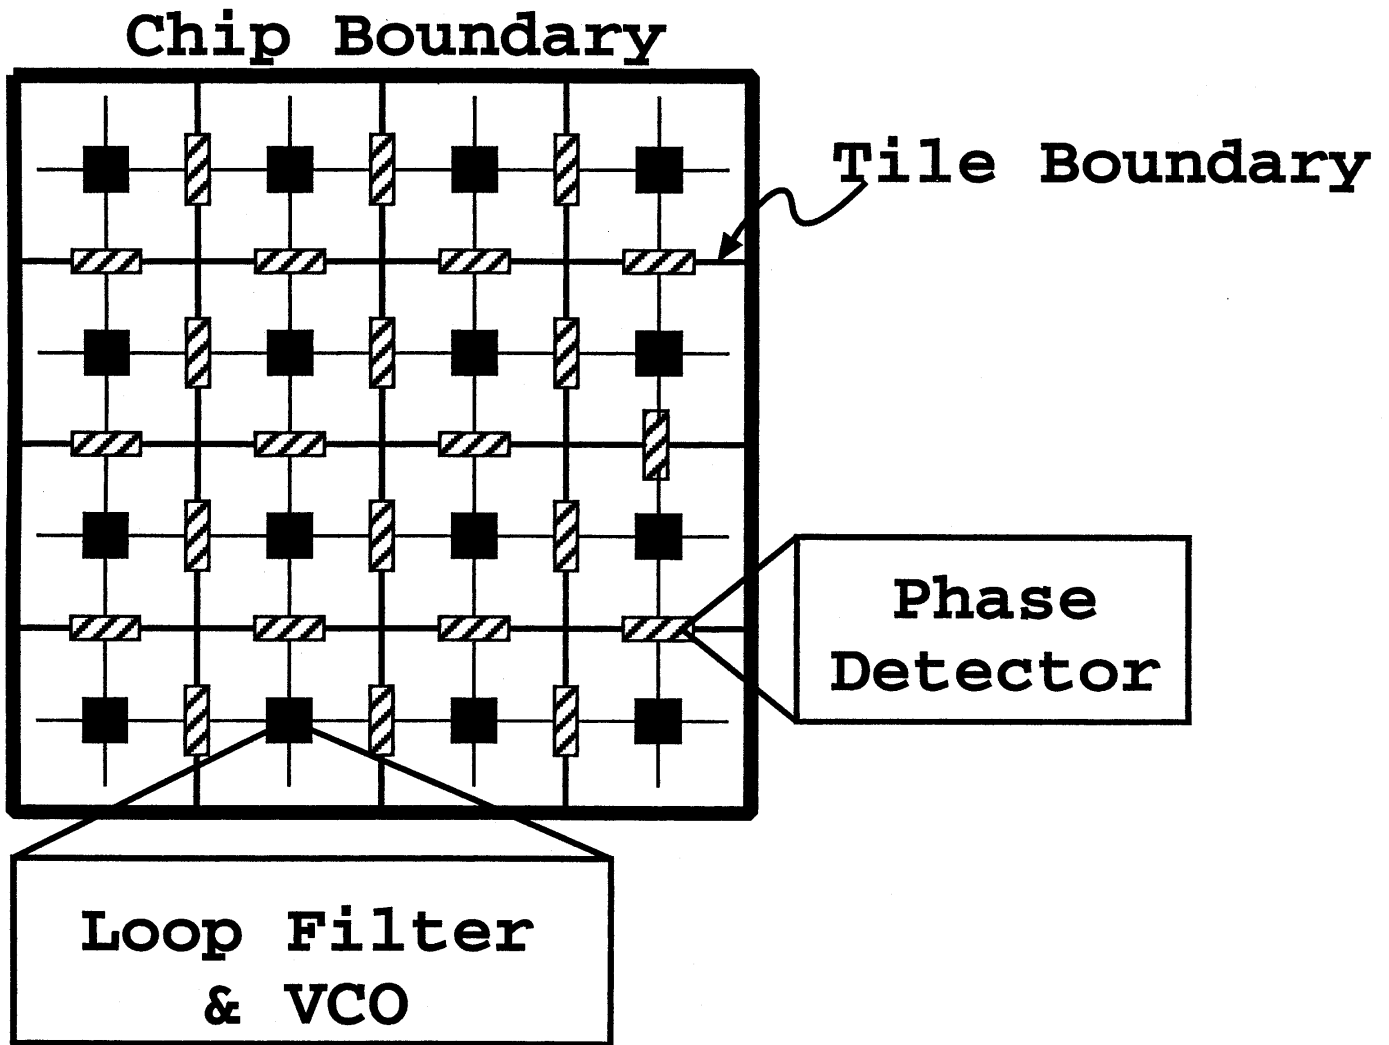


Figure 10.5.1: Distributed clocking network.




Figure 10.5.2: Ring oscillator schematic.




Figure 10.5.3: Phase detector.




Figure 10.5.4: Simulated PD transfer curve.




Figure 10.5.5: Locking behavior of the PLL array.

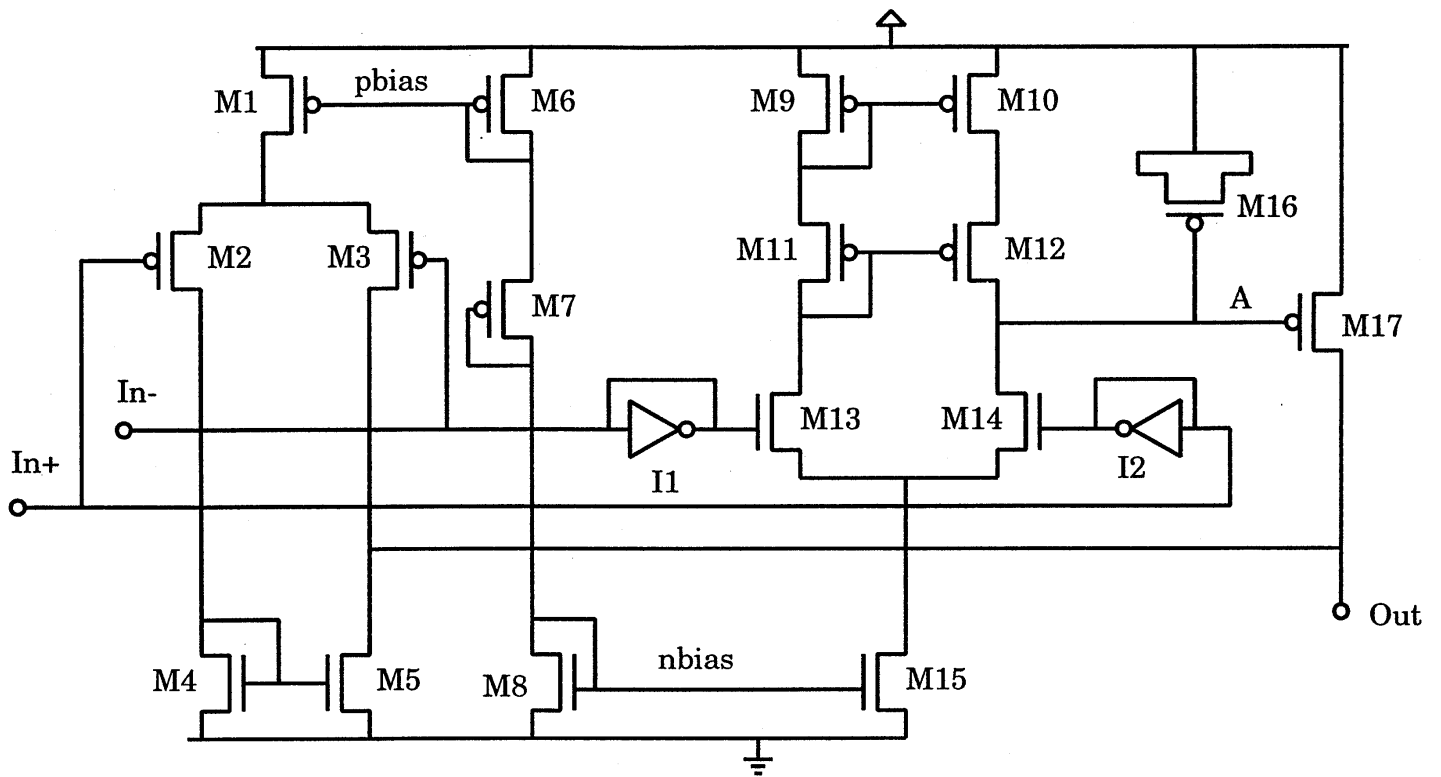


Figure 10.5.6: Loop filter schematic.

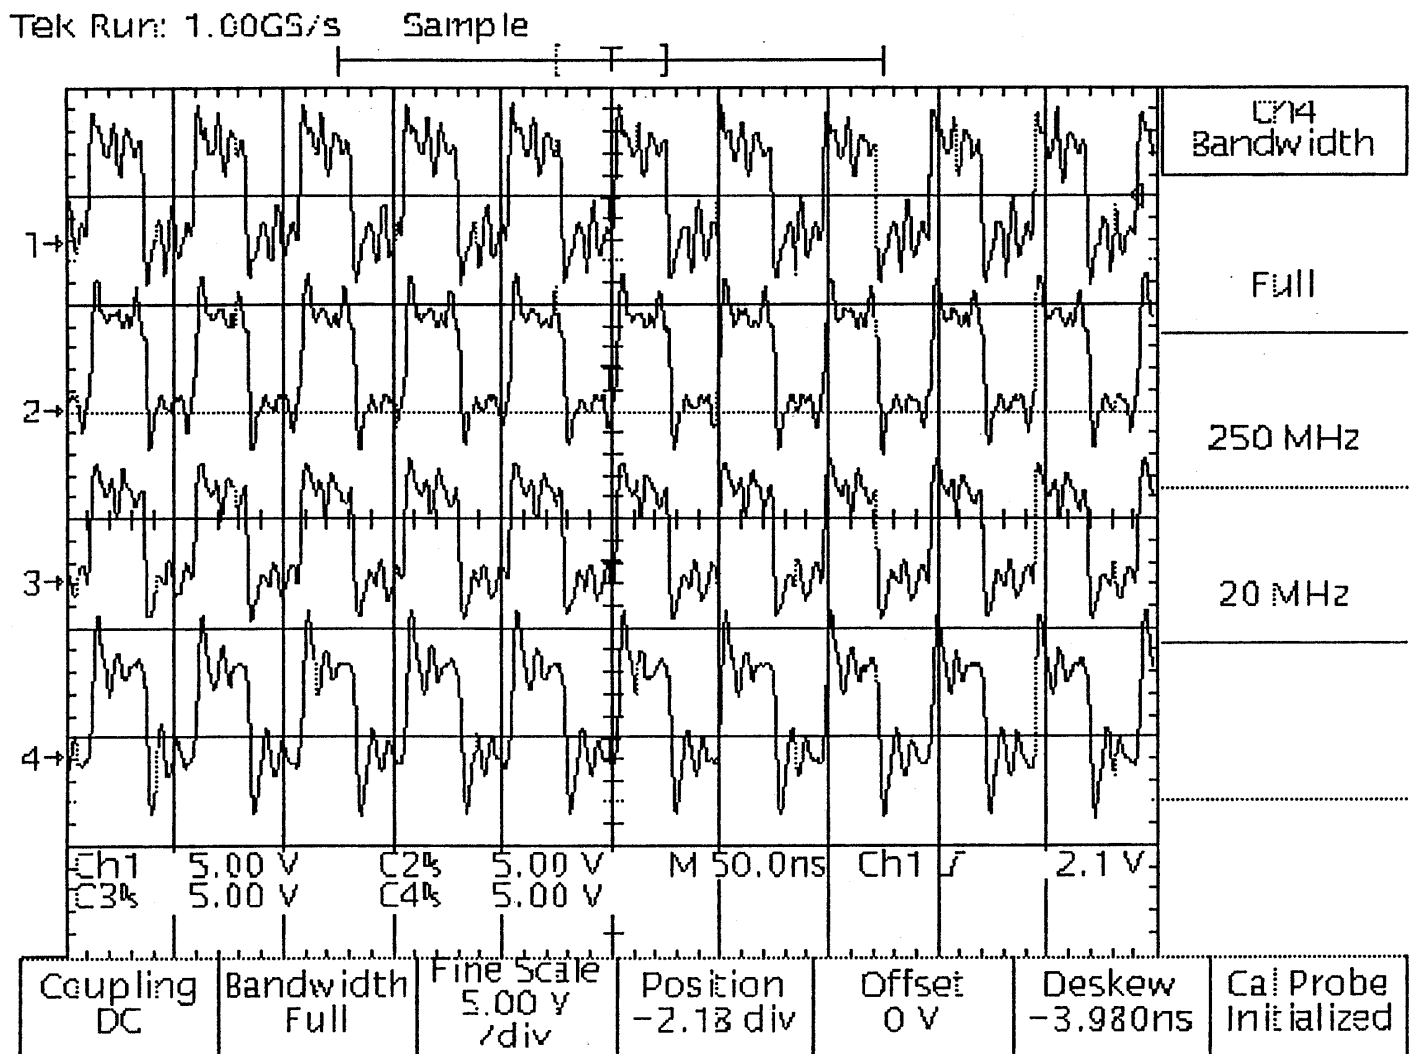


Figure 10.5.7: Frequency-locked divider outputs.




Figure 10.5.8: Distributed clock chip.