Resonant Clock Design for a Power-efficient, High-volume x86-64 Microprocessor

Visvesh Sathe[1], Srikanth Arekapudi[2], Alexander Ishii[3], Charles Ouyang[2], Marios Papaefthymiou[3,4], Samuel Naffziger[1]

[1] AMD Fort Collins, CO

[2] AMD Sunnyvale, CA

[3] Cyclos Semiconductor Inc. Berkeley, CA

[4] The University of Michigan, Ann Arbor

AMD's 4+ GHz x86-64 core code-named "Piledriver" employs resonant clocking [1,2,3,4] to reduce clock distribution power up to 24% while maintaining a low clock-skew target. To support testability and robust operation at the wide range of operating frequencies required of a commercial processor, the clock system operates in two modes: direct-drive (cclk) and resonant (rclk). Leveraging favorable factors such as the availability of two thick top-level metals, high operating frequency, clock-load density, and the existing clock-design methodology [5], the rclk mode was designed to enable both reduced average power dissipation and improved peak-power-constrained performance, with minimal area impact. Fabricated in a 32nm CMOS process, this work represents the first volume production-enabled implementation of resonant clock technology.

Rclk allows power reduction by recycling charge using LC-resonance, which enables further power reduction by reducing clock driver strength. Figure 1 shows a simplified schematic of the dual-mode clock system. The mode switch MSw is closed (open) in rclk (cclk) mode. The clock driver features a pulse-drive mode for additional efficiency improvement through duty cycle control of the pull-up and pull-down switches. TSw is a throttle switch employed to reduce voltage overshoot when the MSw is turned off during frequency changes.

To operate in both modes, the clock driver needs to support frequency-dependent drive-strength and pulse modulation, both of which are efficiently implemented using a split-buffer topology. In rclk mode, drive strength is modulated through drvEn settings during P-state transitions. Pulse drive is used to enable a finer trade-off between conduction and switching losses in the driver. A local delay line delays only the asserting edge of the pull-up/down stage during pulse drive (plsEn = 1), whereas respective de-asserting edges are triggered by the non-delayed clock. Thus, the driver output duty cycle is obtained by programming the local delay to modulate the input duty cycle. This method has three advantages:(1) Enabling PLL duty cycle control of the clock to tune performance (2) Guaranteeing robust clock slew and amplitude when operating off the V-f curve and (3) Reducing susceptibility in rclk skew due to process variation in the local delay chains.

Figure 2 shows the Piledriver global clock construction in which a set of five horizontal-folded clock trees (HCK tree) drive a global clock grid [5]. Each HCK tree has up to 25 inductors

interleaved with clock drivers and can be programmed independently in test mode. The clock mode and frequency-dependent clock parameter settings (inductor connection, drive strength, pulse width) are adjusted during power-up and each P-state transition, during which time the clock mode parameters are initialized through a P-state indexed fuse based table. The power management unit accounts for the power reduction achieved in each P-State. These parameters are loaded by a sequencer in the transmit block, which distributes these parameters to the HCK trees through a source-synchronous bus inside the vertical clock tree module. Once received by the HCK trees, these parameters are broadcast to all clock drivers within each HCK tree. To avoid a circular dependence between the global clock and logic used to program the clock, all programming logic in the HCK trees is clocked by a broadly distributed intermediate stage of the clock tree. Existing clock gating mechanisms are leveraged to prevent the exposure of timing elements in the CPU to transitional clocks.

Building inductors with a good quality factor Q, is critical to rclk efficiency, and is constrained by several factors in Piledriver. The inductor windings have to be designed to share metal resources on the top two metal layers with dense power distribution. Moreover, they must accommodate a substantial number of pre-clock distribution nets and global nets that are routed through, as well as under, the inductor. Figure 3 illustrates inductor design under these constraints. At the frequencies of interest, Q is dominated by winding resistance. The inductor was therefore designed using M11 and M10 thick-metal levels, with cut-aways allowing for

maximal use of both metal layers in the presence of routes and power-supply trunks. Inductor placement was directed to ensure that power-supply trunks pass through the middle of the inductor, minimizing the impact of inductive coupling. The power grid under the inductor was designed to be “loop-less” to alleviate the Q degradation resulting from eddy currents in the power grid loops while maintaining a robust grid.

Figure 4 shows the structures required to support rclk (MSw, inductor, TankCap) that are tiled across the HCK-tree. MSw connects the inductor to the clock grid through the Driver-MSw shorting-bar. An LP formulation was used to determine inductor allocation from a palette of 5 with values in the 0.6 to 1.3 nH range. Skew was further controlled by interleaved driver/inductor placement. For each inductor, MSw size was tuned to trade-off reduced switch resistance with the increased switch parasitic that results from larger switches. For efficient rclk operation, a large low-ESR TankCap is required within a limited allocated area. To that end, a capacitor structure of approximately six times the average clock load was implemented using both metal and gate structures.

Figure 5 shows measured Cac (defined as Cac = $P_{dynamic}$ / V^2f) savings and efficiency numbers, based on power dissipation in the clock drivers and grid, in cclk and rclk modes. A test pattern with high switching activity was used for the clock power measurement. Efficiency increases up to 3.3 GHz, and declines more gradually at higher frequencies. The inherent asymmetry of energy efficiency on either side of the resonant frequency is increased due to a voltage-

dependent Q (from the series-connected MSw) and a stringent clock slew criterion that requires a stronger drive at lower frequencies. The impact of rclk on skew was minimal: Full-chip simulation analysis showed a 1 ps increase in rclk skew compared to cclk.

Figure 6 shows cclk and rclk waveforms from a full-chip clock simulation at 1.2V, 4.25GHz. Reducing clock driver strength in rclk enables greater Cac savings at the expense of reduced clock slew rates. These reduced slews result in increased cross-over current in the clock receivers. Measurements however, show a negligible change in efficiency for high-activity workloads compared to idle workloads which indicates that this effect is small. Reduced slew also causes a push-out in the 50% arrival time of the clock, potentially affecting both gater-enable paths and cross-clock domain communication. Static timing analysis with degraded slews was run on the core and resulting paths fixed.

Figure 7 shows the microphotograph of the Piledriver core. Over the frequency range 3.0 GHz to 4.4 GHz, the power savings from rclk enable either a frequency increase of about 100 MHz for the same power, or a power reduction of 5-10% for the same frequency.

Acknowledgments:

The authors thank Tom Meneghini, Kyle Viau, Manivannan Bhoopathy, Joohee Kim, Jerry Kao, Fred Brauchler, Alan Arakawa, Syed Obaidulla, Kevin Hurd, Vasant Palisetti, and Denny Renfrow for their valuable contribution to this work.

References:

[1] Drake, A.J et al. “Resonant Clocking using Distributed Parallel Capacitance,” *JSSC*, Sep 2004.
[2] Sathe, V.S et al. “Resonant-Clock Latch-Based Design,” *JSSC*, Apr 2008.
[3] Chan, S.C. et al. “A Resonant Global Clock Distribution for the Cell Broadband Engine Processor,” *JSSC*, Jan 2009.
[4] Ishii, A. et al. “A Resonant Clock 200MHz ARM926EJ-S™ Microcontroller,” *ESSCIRC* 2009.
[5] McIntyre, H. et al. “Design of the Two-core x86-64 AMD ‘Bulldozer’ Module in 32 nm SOI CMOS,” *JSSC*, Jan 2012.

Captions:

Figure 1: Simplified model of AMD’s “Piledriver” dual-mode global clock network

Figure 2: Global-clock organization and distribution. A folded clock-tree (VCK tree), and 5 horizontal folded clock trees (HCK tree) achieve a low-skew core-wide distribution to clock drivers which drive the global clock grid. The HCK-tree is also used to achieve rclk programmability.

Figure 3: Inductor design on the top two metal layers with cut-aways to accommodate power straps and global signal routes. A “loop-less” custom grid is implemented under the winding.

Figure 4: Repeated section of the HCK tree showing relative placement of final-drivers, inductor, MSw, and TankCap. The Inductor-TankCap Shorting-Bus is used to provide a distributed low-resistance TankCap connection to the inductor.

Figure 5: Measured Cac(pF) savings and clock efficiency vs. frequency. Peak efficiency is observed at 3.3GHz.

Figure 6: Simulated cclk and rclk waveforms at Vdd=1.2V, frequency=4.25GHz under different drive-strength configurations. rclk_3/8 refers to an rclk mode where the clock drivers are operating at 3/8 drive-strength. Lower-drive strengths in rclk allow for more Cac savings at the expense of lower slew rates.

Figure 7: Chip Microphotograph of the resonant-clocked 32nm AMD “Piledriver” core.

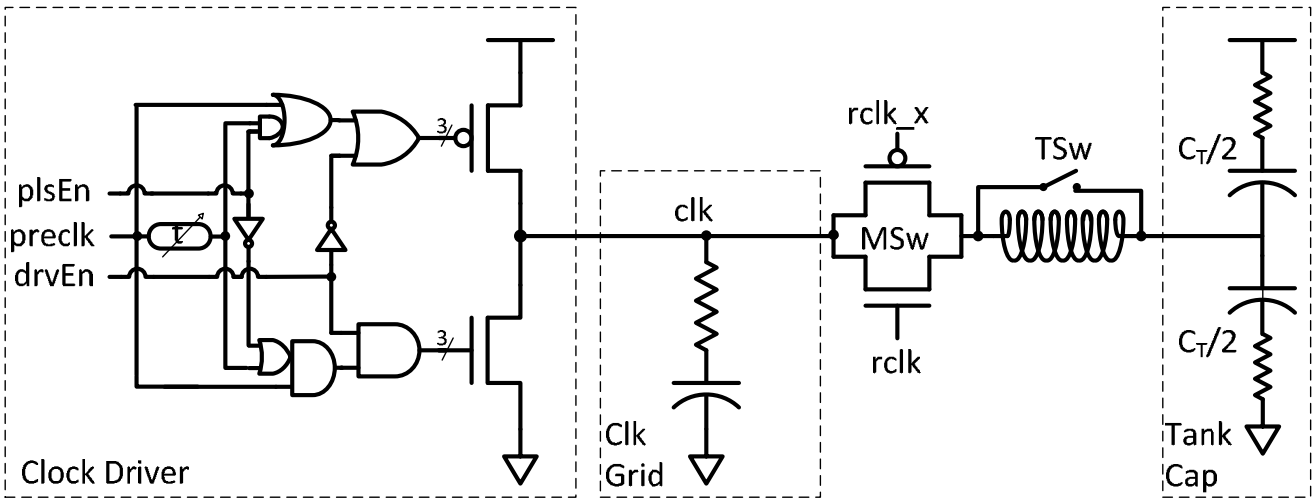

Figure 1: Simplified model of AMD's "Piledriver" dual-mode global clock network

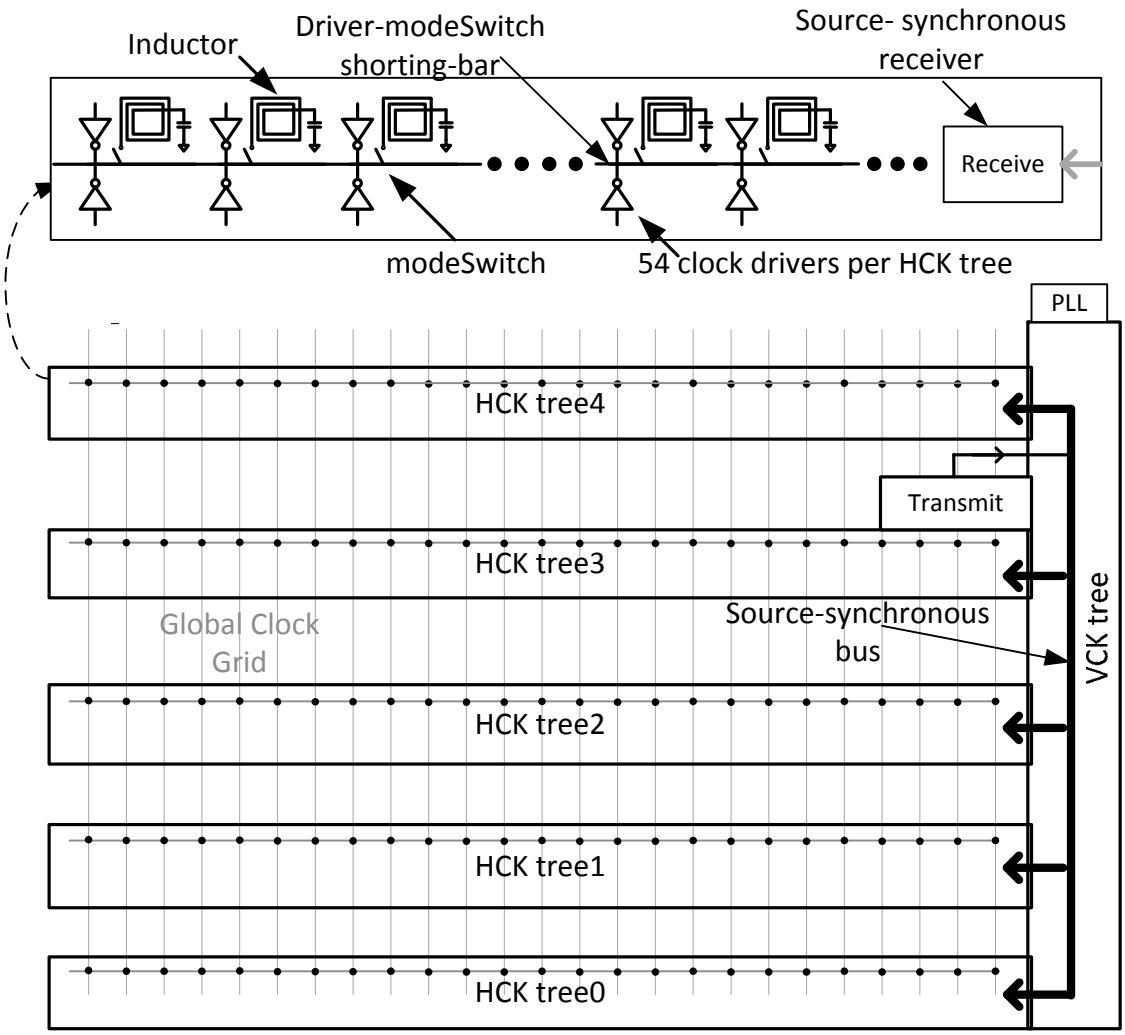

Figure 2: Global-clock organization and distribution. A folded clock-tree (VCK tree), and 5 horizontal folded clock trees (HCK tree) achieve a low-skew core-wide distribution to clock drivers which drive the global clock grid. The HCK tree is also used to achieve rclk programmability.

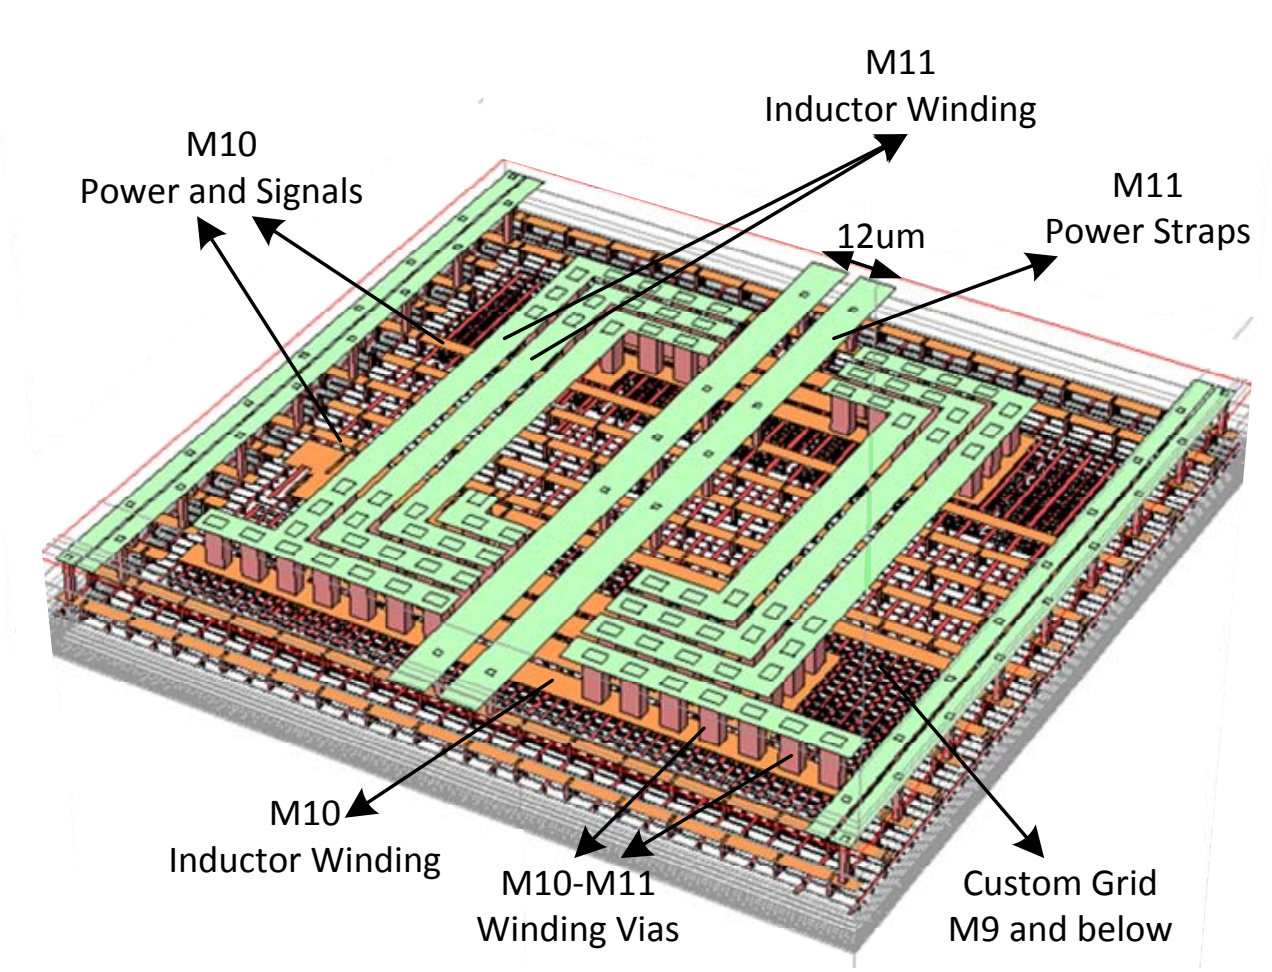

Figure 3: Inductor design on the top two metal layers with cut-aways to accommodate power straps and global signal routes. A “loop-less” custom grid is implemented under the winding.

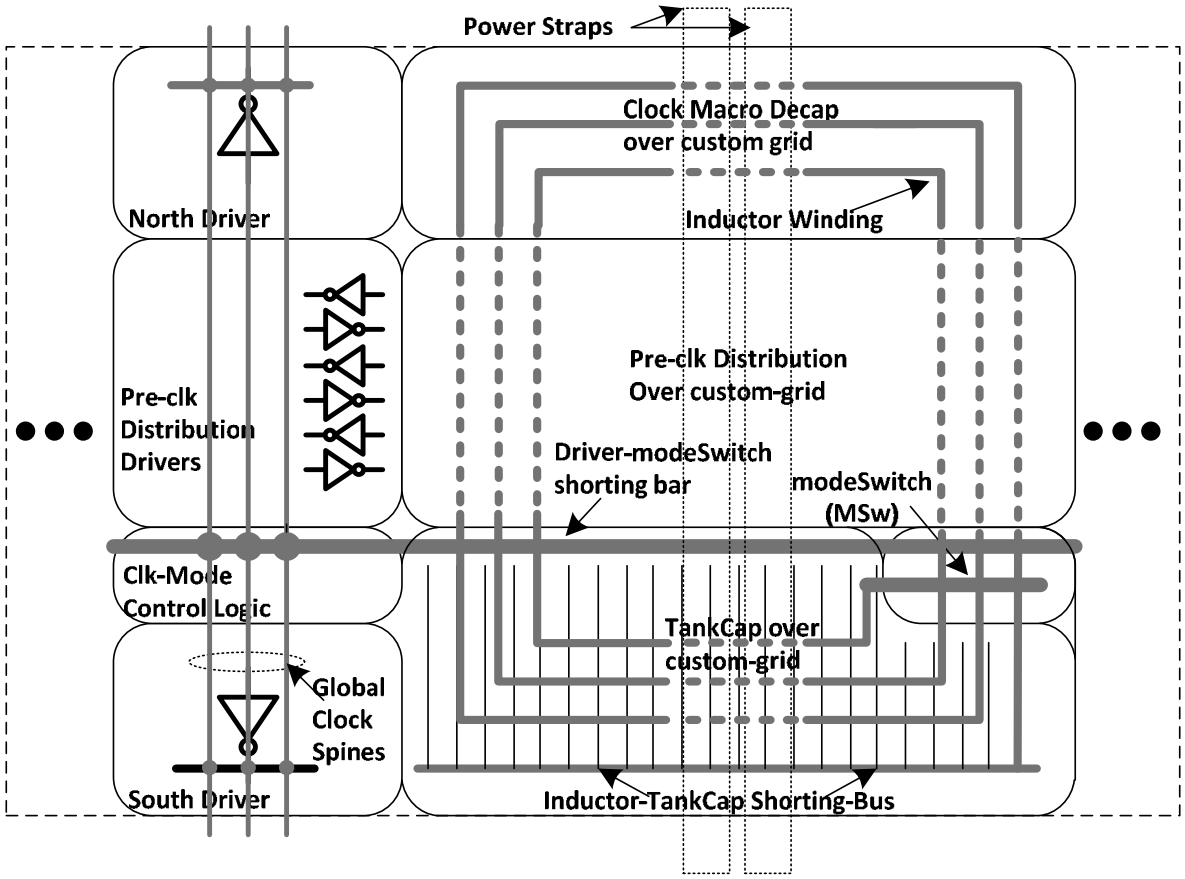

Figure 4: Repeated section of the HCK tree showing relative placement of final-drivers, inductor, MSw, and TankCap. The Inductor-TankCap Shorting-Bus is used to provide a distributed low-resistance TankCap connection to the inductor.

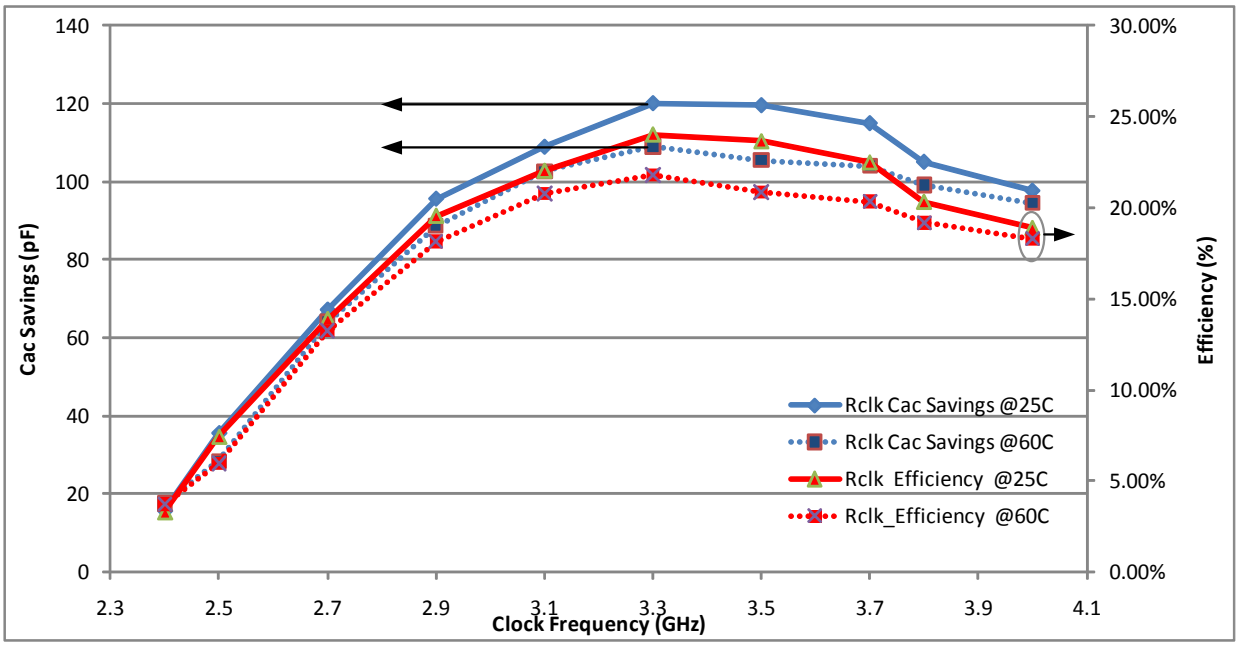

Figure 5: Measured Cac(pF) savings and clock efficiency vs. frequency. Peak efficiency is observed at 3.3GHz.

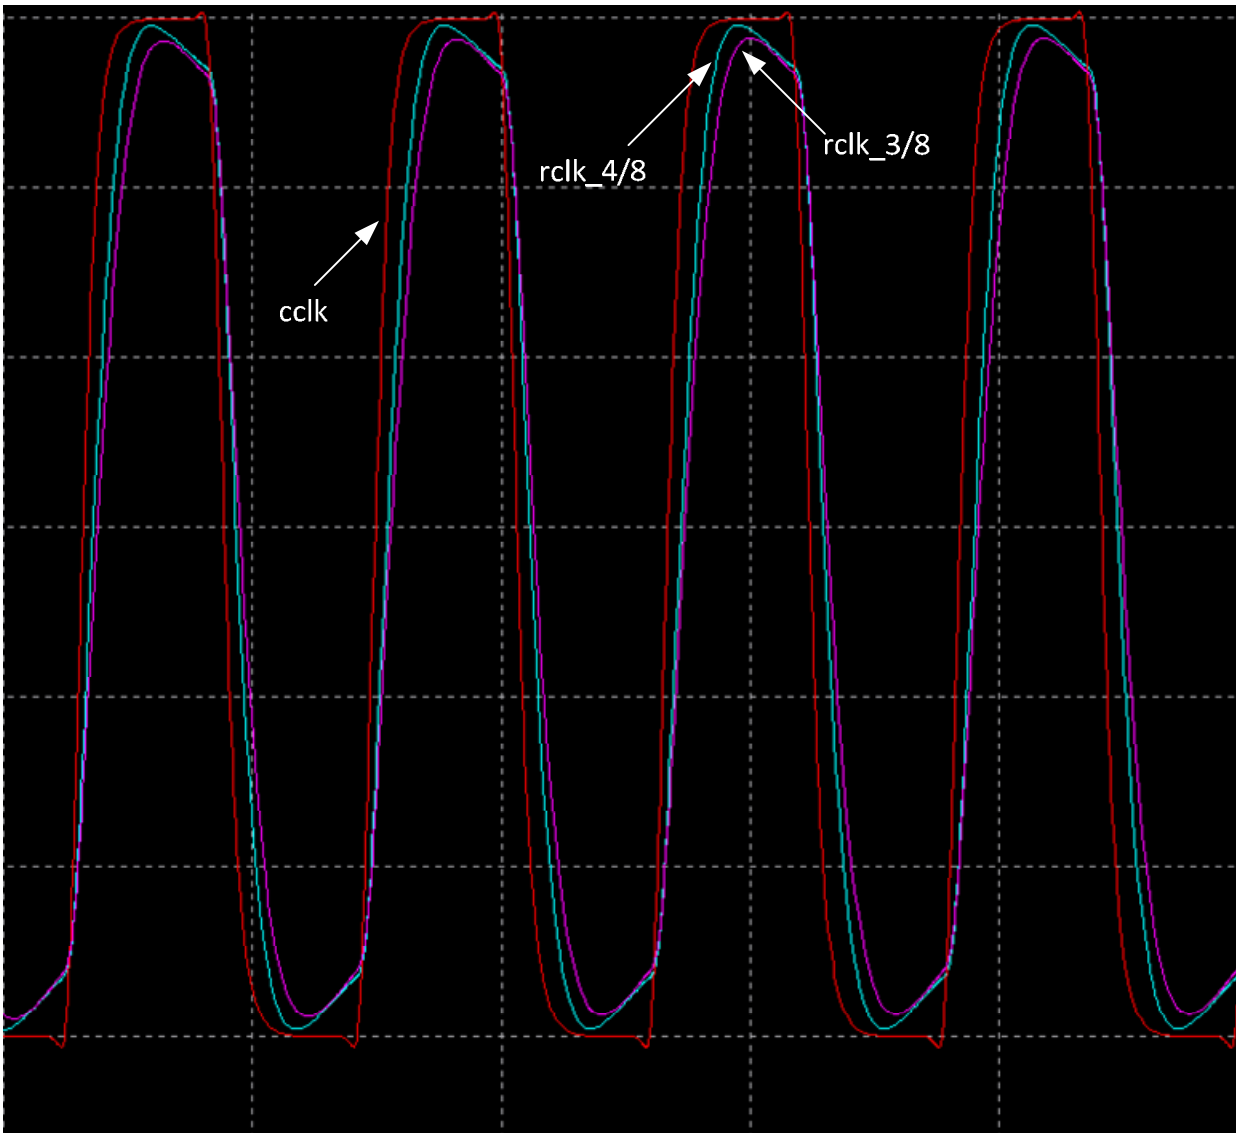

Figure 6: Simulated cclk and rclk waveforms at Vdd=1.2V, frequency=4.25GHz under different drive-strength configurations. rclk_3/8 refers to an rclk mode where the clock drivers are operating at 3/8 drive-strength. Lower-drive strengths in rclk allow for more Cac savings at the expense of lower slew rates.

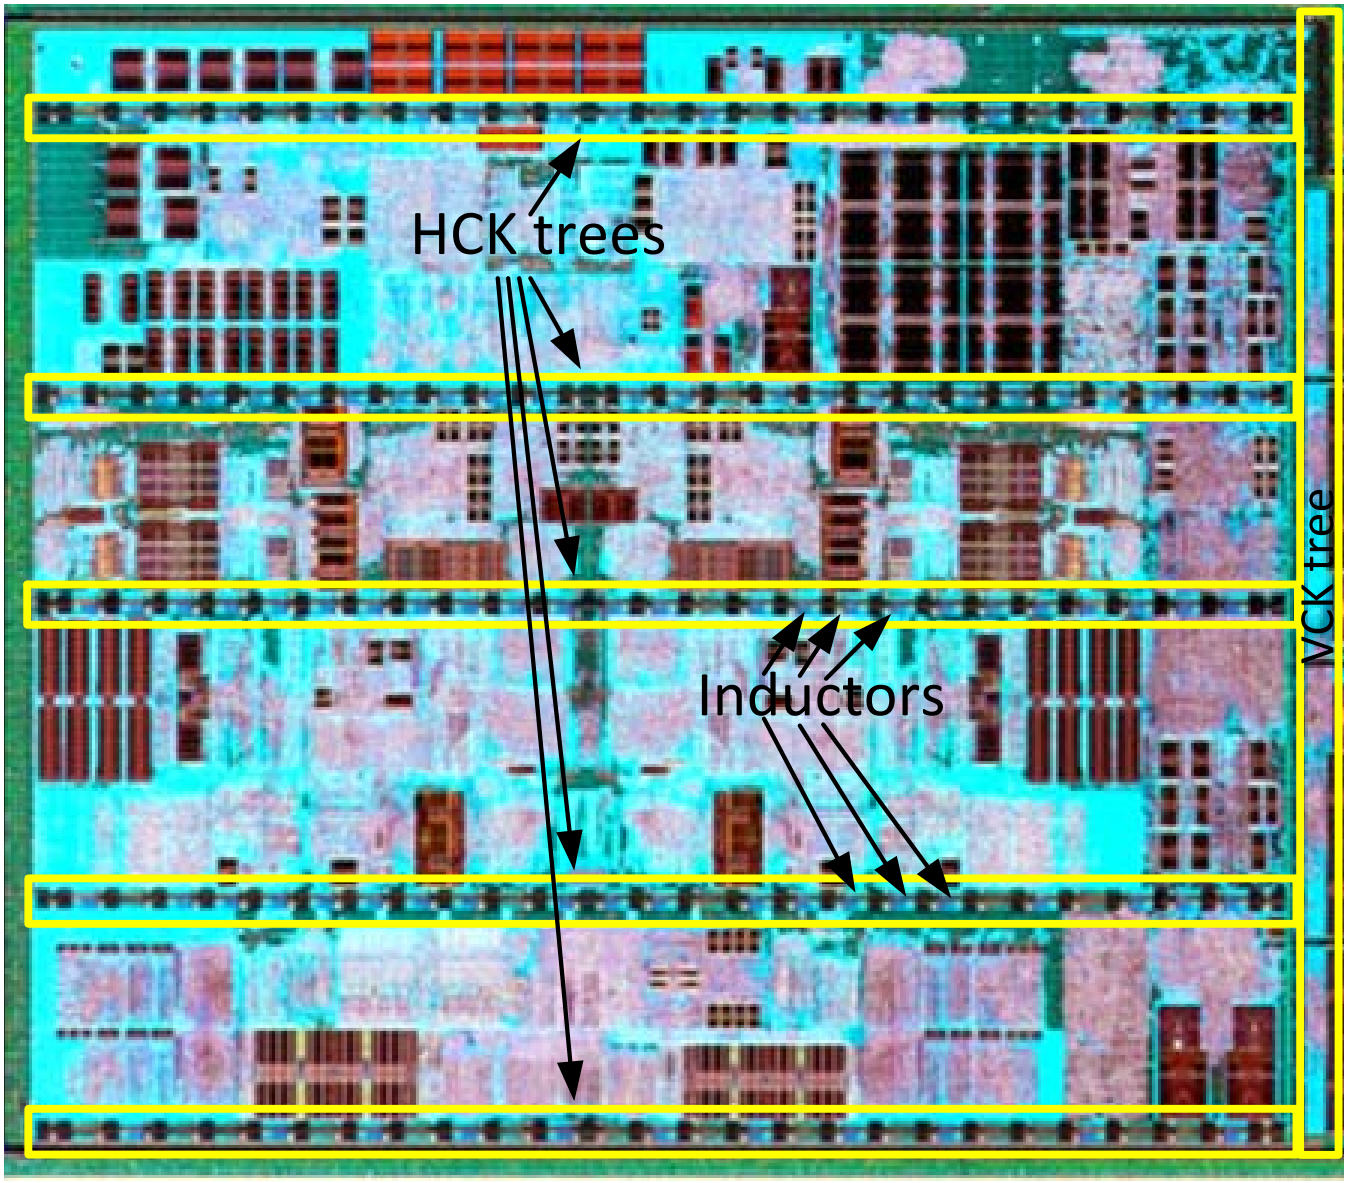

Figure 7: Chip Microphotograph of the resonant-clocked 32nm AMD “Piledriver” core.