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

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

[1]Advanced Micro Devices, Inc.
[2]Cyclos Semiconductor, Inc.
[3]The University of Michigan, Ann Arbor

Outline

- Piledriver (PD) overview
- Resonant clocking
- Clock driver design
- Inductor design
- Other resonant clock components
- Putting it all together
- Measurement results
- Conclusion

Piledriver

- Extension of Bulldozer Architecture - 2-core module:
  - Shared I-fetch, Decode, Br. Predict, FP, L2, Cache unit
  - Per-core Integer schedule/execute, Load/Store

- 32nm CMOS
- HKMG, SOI
- 11 metal layers
- 33.3 mm^2 w/L2
- 216M transistors
- 0.8-1.3 V

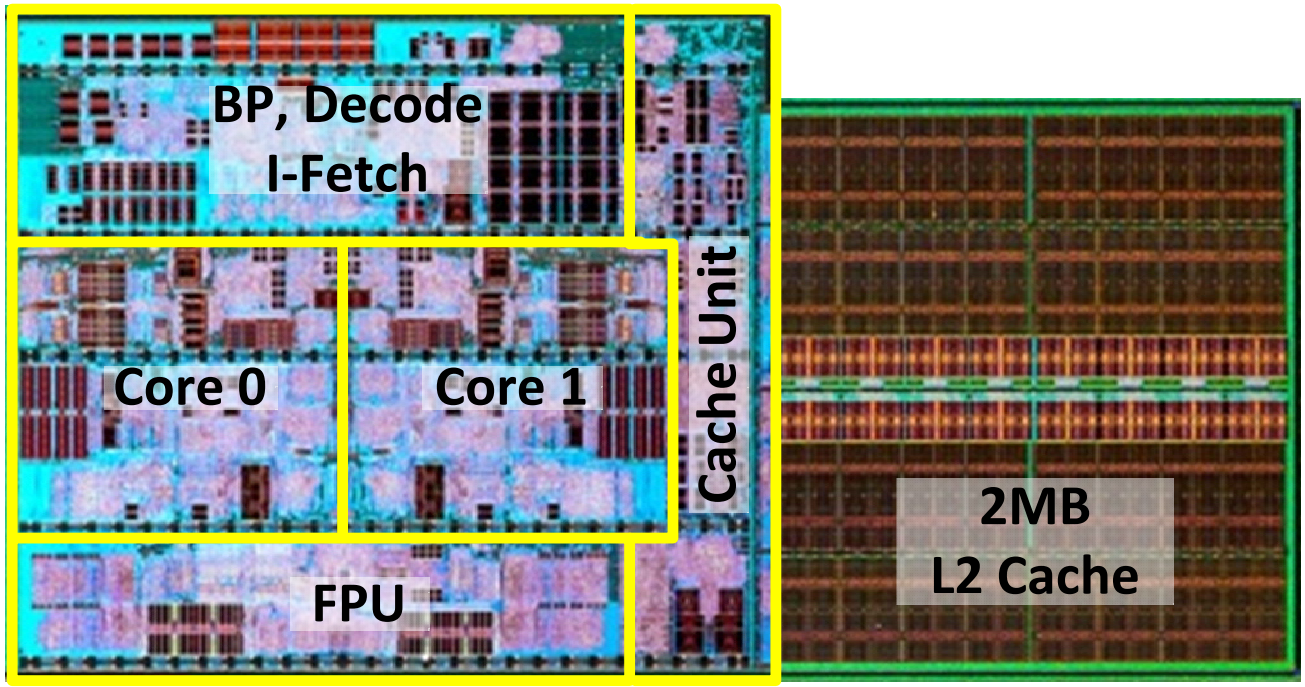

The Piledriver Global Clock

- Significant global clock loading
  - 7ps clock grid skew target across 21mm^2 core area
  - Constrained clock latency from grid to timing elements
- 24% of average application power in global clock

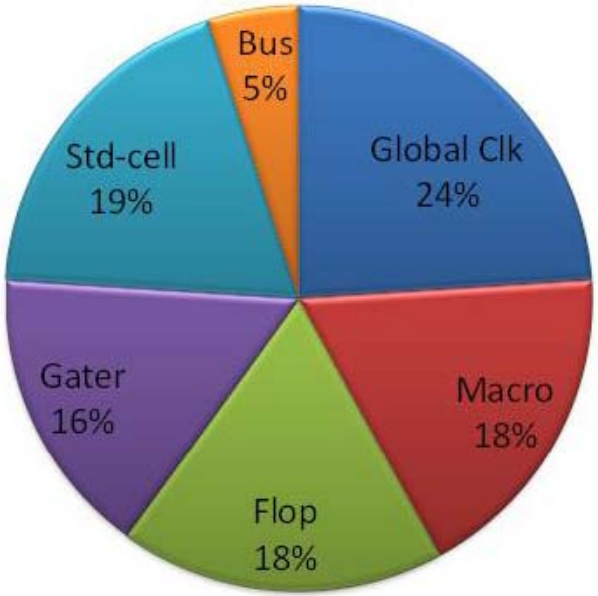

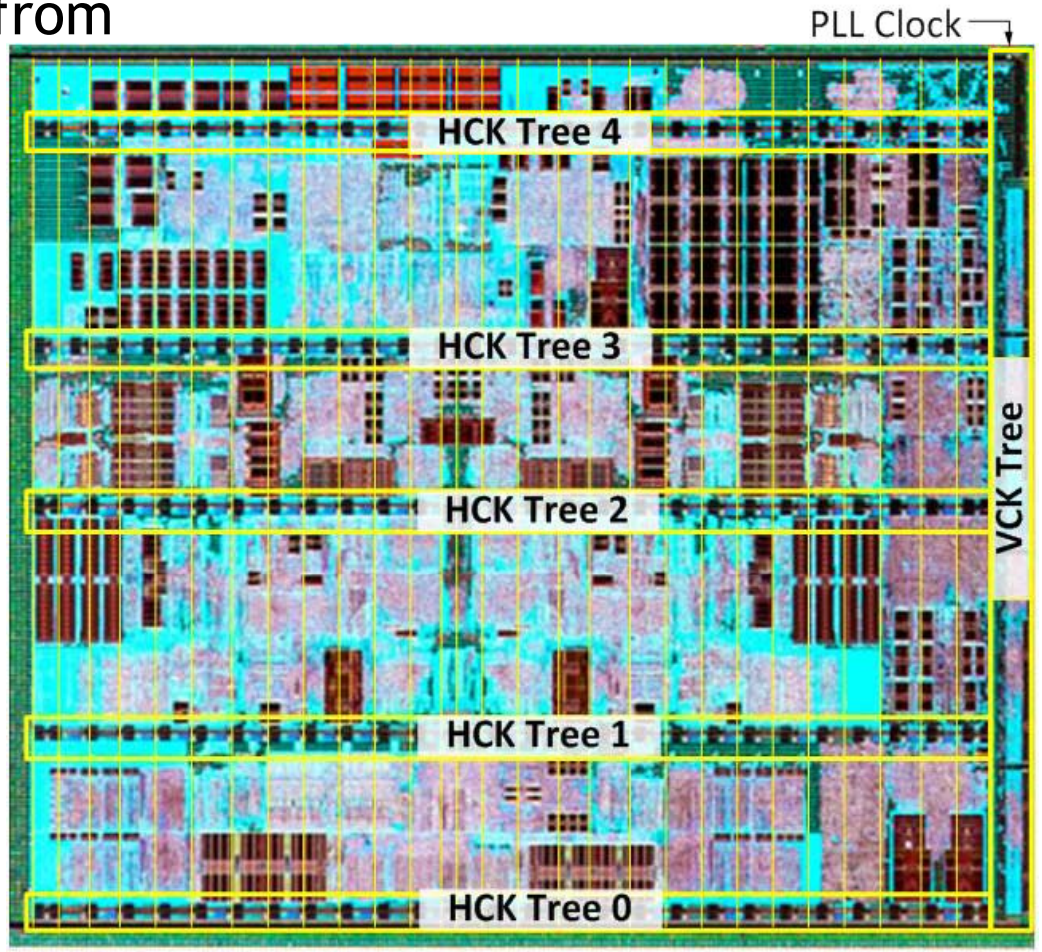

Basic Resonant Clocking Operation

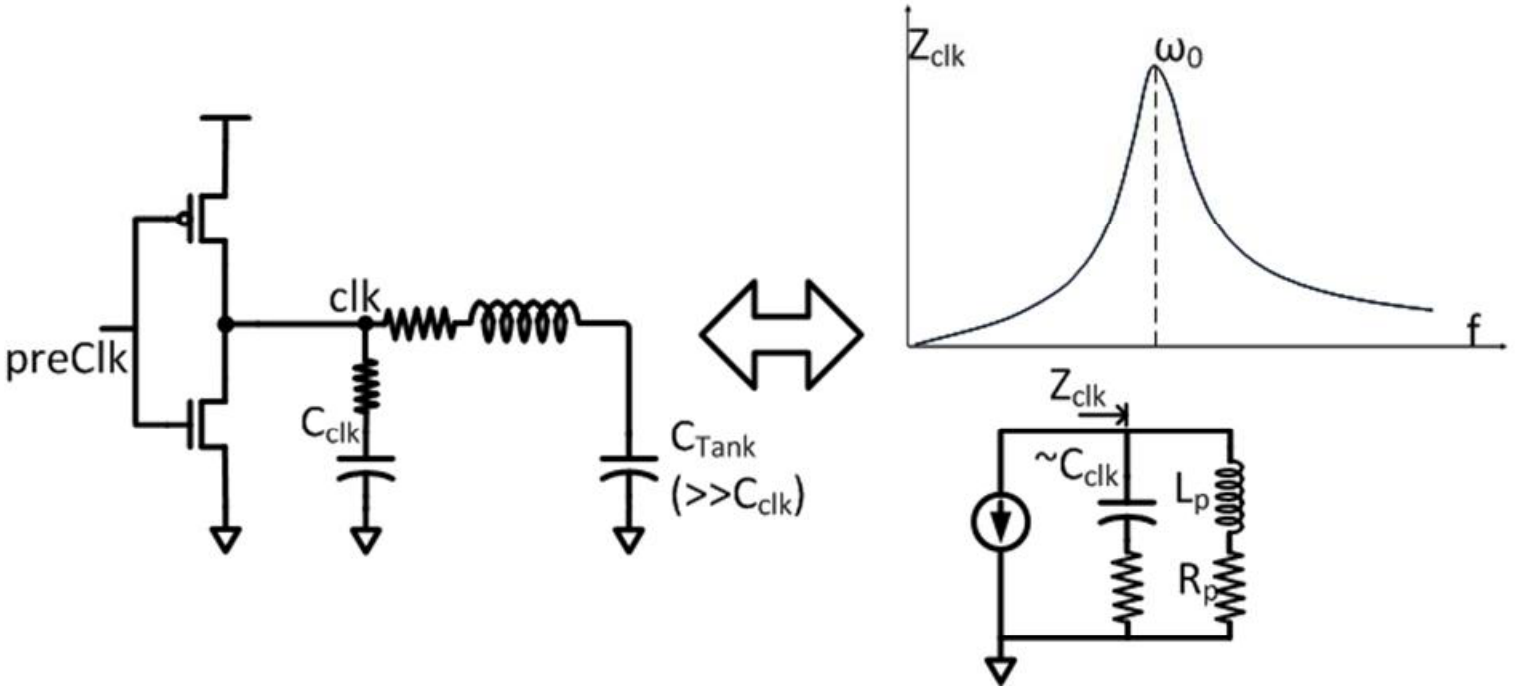

- Rely on efficient *LC* resonance between spiral inductors and grid capacitance near resonant frequency
- Efficient operation around natural frequency
- Driving clock at much lower frequencies
  - → Reduced efficiency
  - → Warped clock waveform

Resonant Clocking: Mutual Inductance

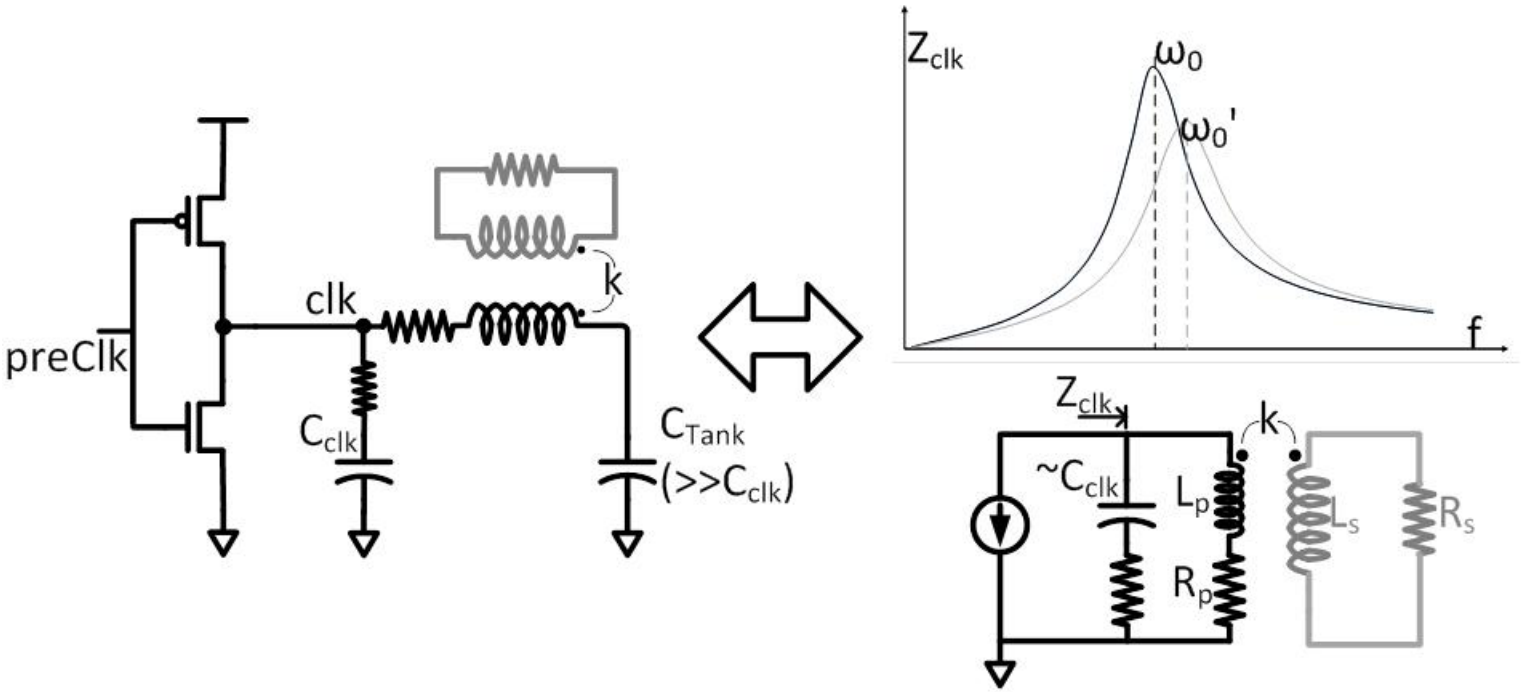

- Significant mutual inductance interaction from
  - Signal nets under and around inductor windings.
  - Power and ground nets serving circuits under inductor
- Mutual inductance causes
  - ↓ L, ↓ Q, ↑ Clock power
- Maintaining keep-out regions is prohibitive (~5% area penalty)

Resonant Clocking on PileDriver

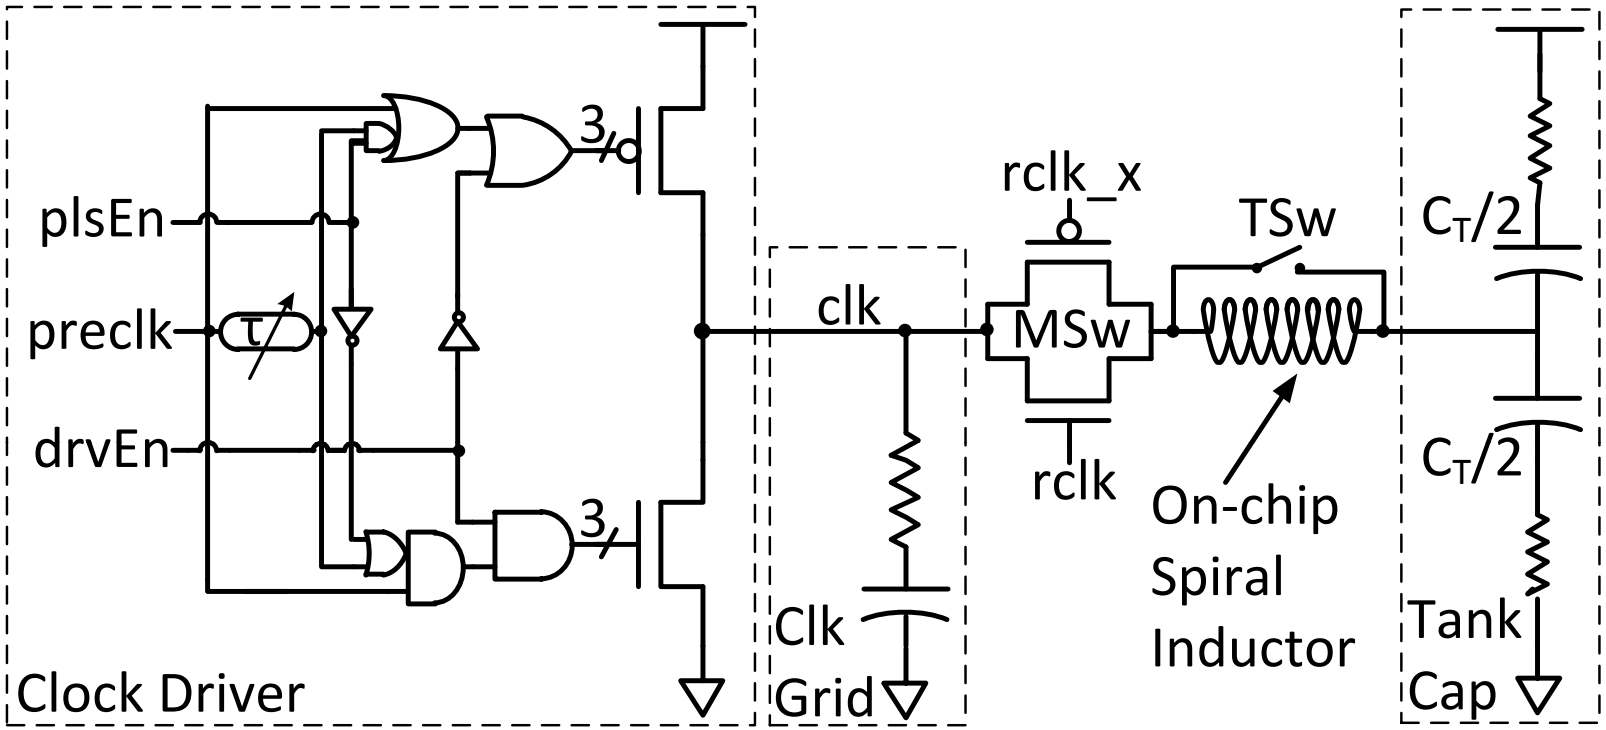

- Dual mode clock system with Mode switch (MSw)
  - resonant clocking (rclk) when MSw is closed
  - conventional clocking (cclk) when MSw is open
- $C_T \approx 6xC_{clk}$ to serve as effective AC ground
- Throttle Switch (TSw) to address transient voltage spikes

Piledriver Resonant Clocking

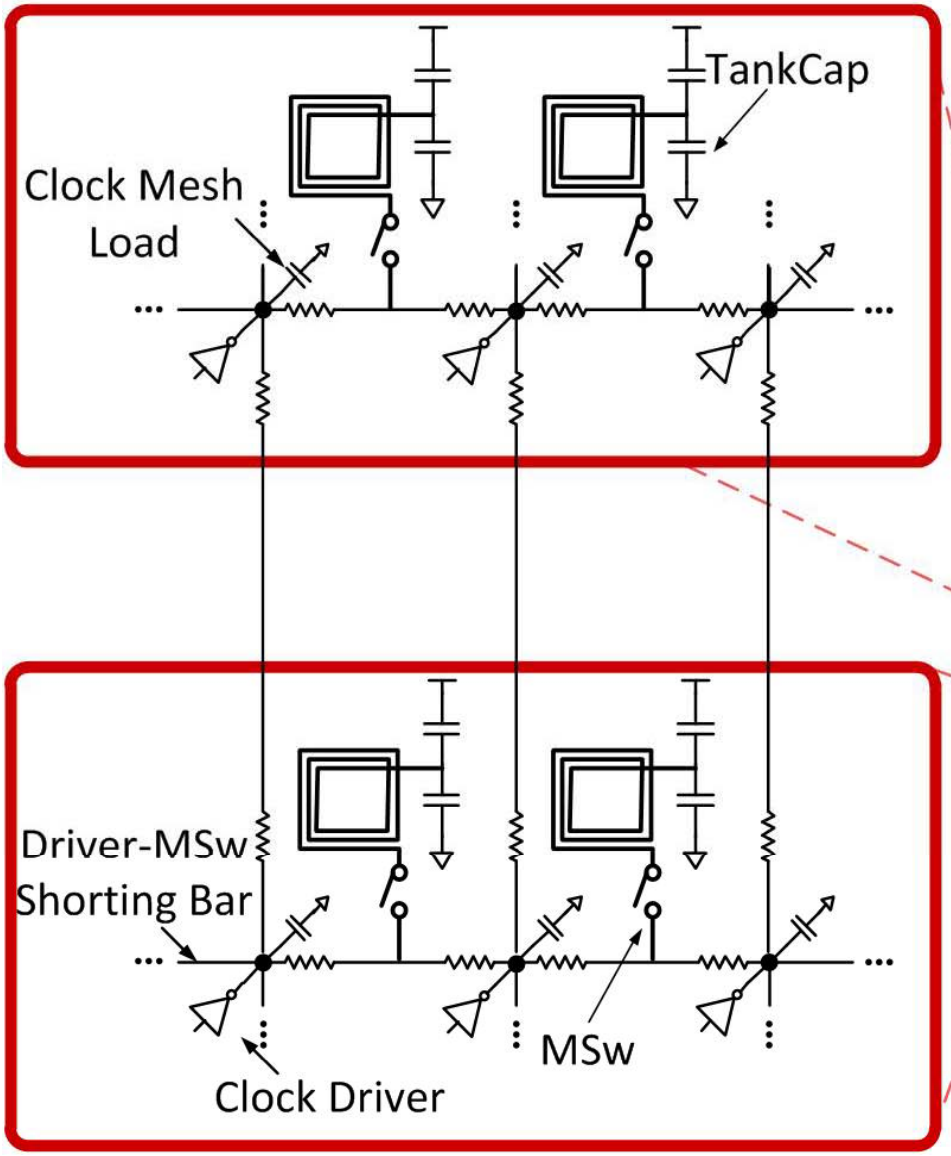

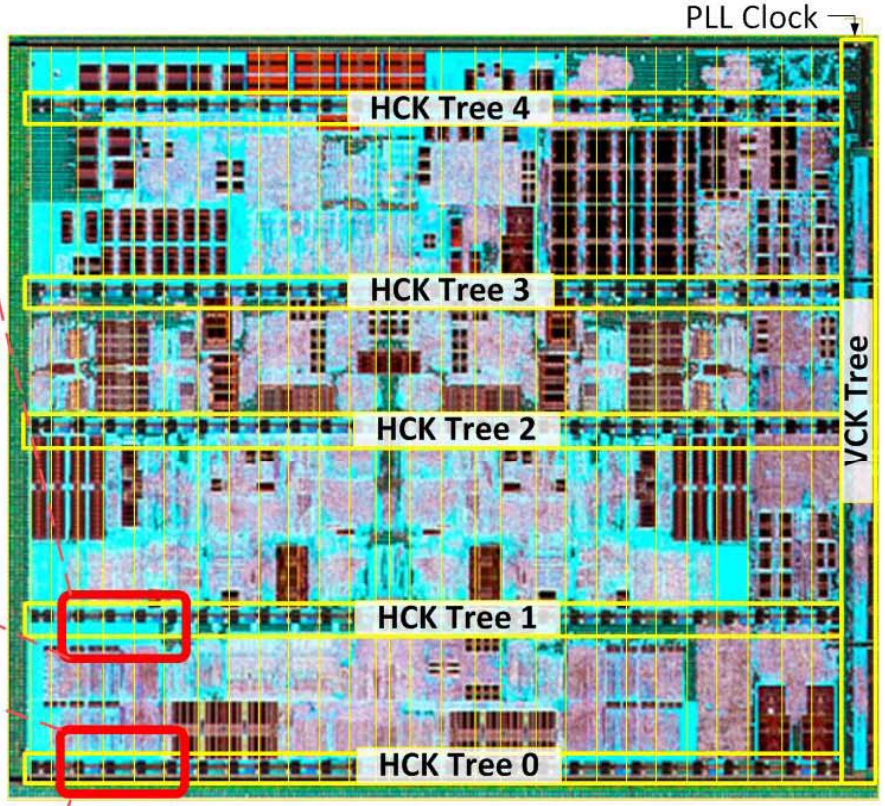

- 92 distributed inductors
- MSw, TankCap, TSw inductors contained in HCK Tree Macros
- Conventional L2CLK, NBCLK

Basic Rclk Operation (1/6)

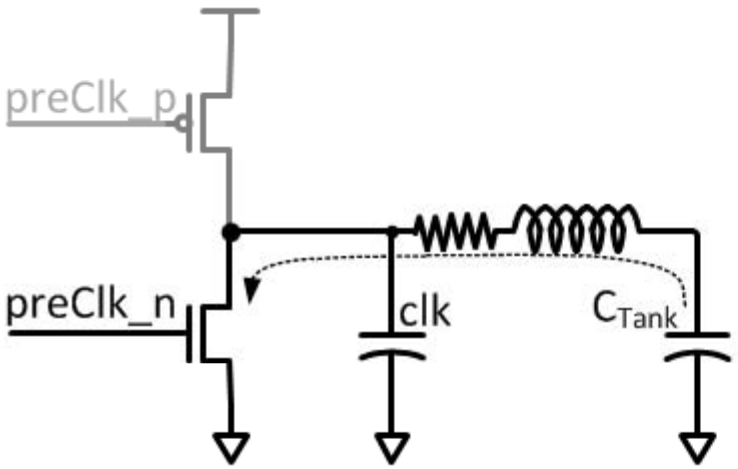

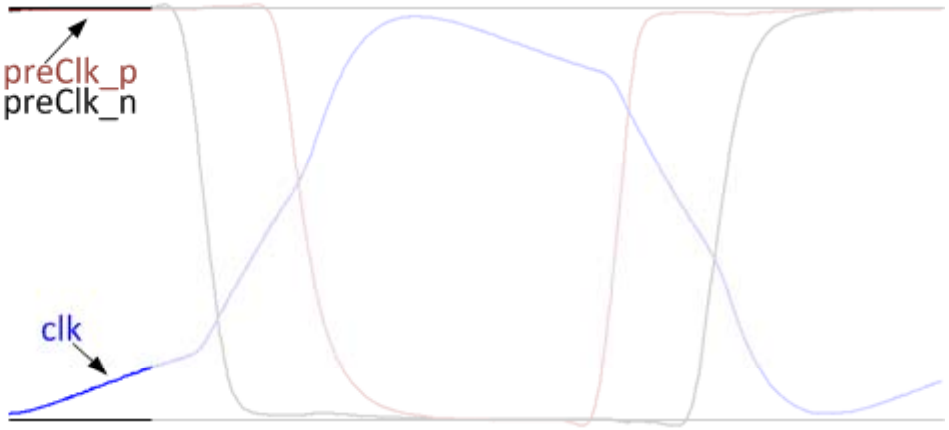

- Rclk operation can be partitioned into 6 phases
- Voltage across C_{tank} ~$Vdd/2$
- Nmos conducting, Pmos off
- *R-L* Current buildup through the Nmos
- IR drop across Nmos and grid resistance → clock voltage ↑

Basic Rclk Operation (2/6)

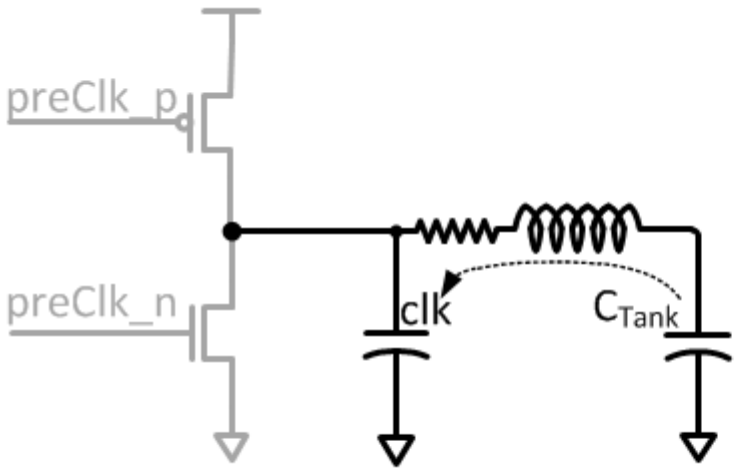

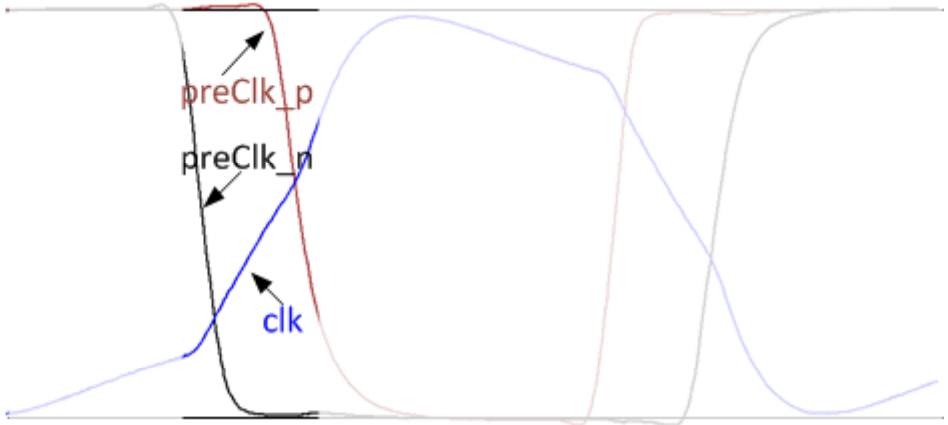

- Both, Nmos and Pmos are off
- *LC* oscillation with initial inductor current to charge clk
- Clk voltage transition a function of L,C and initial current.

Basic Rclk Operation (3/6)

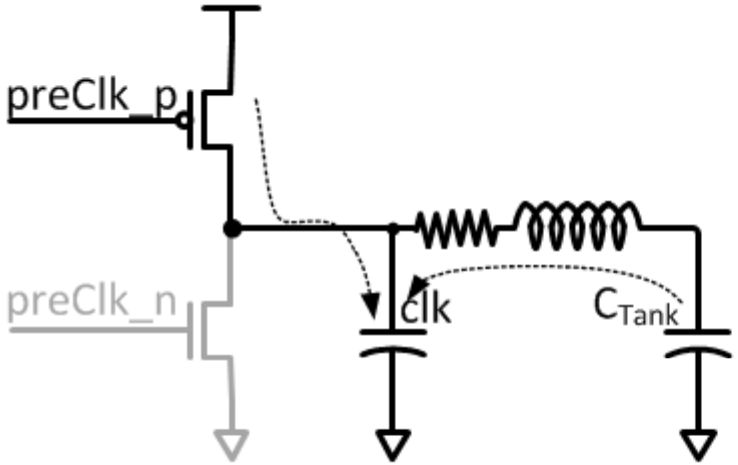

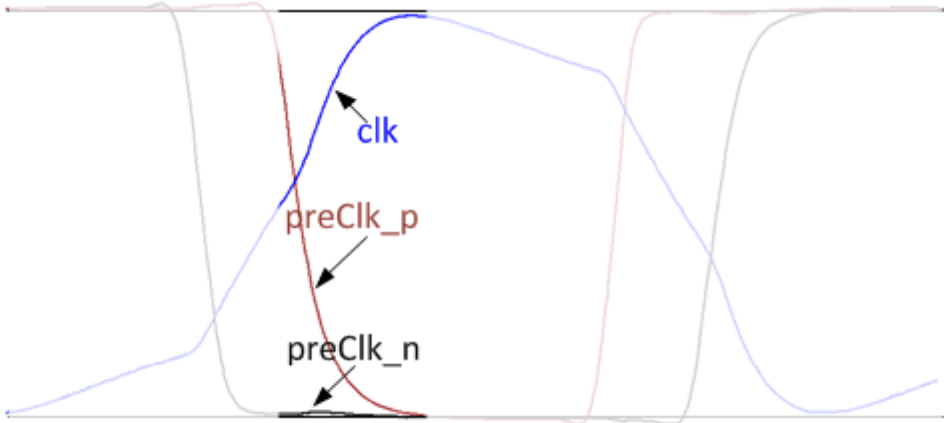

- Pmos on, Nmos off.
- Both *LC* and pull-up mechanisms in effect

Basic Rclk Operation (4/6)

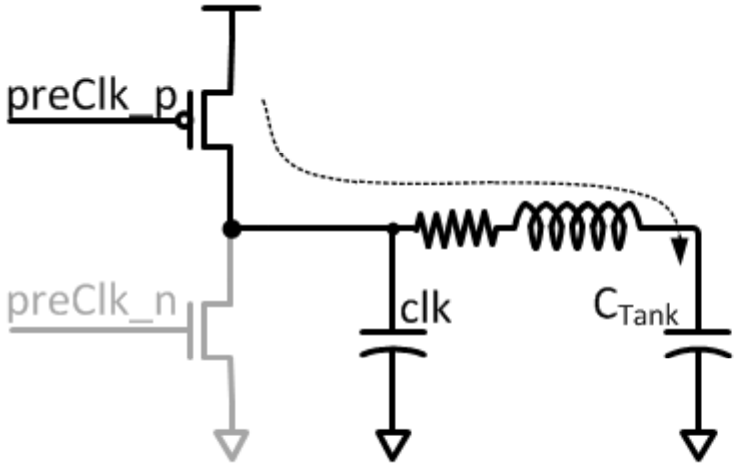

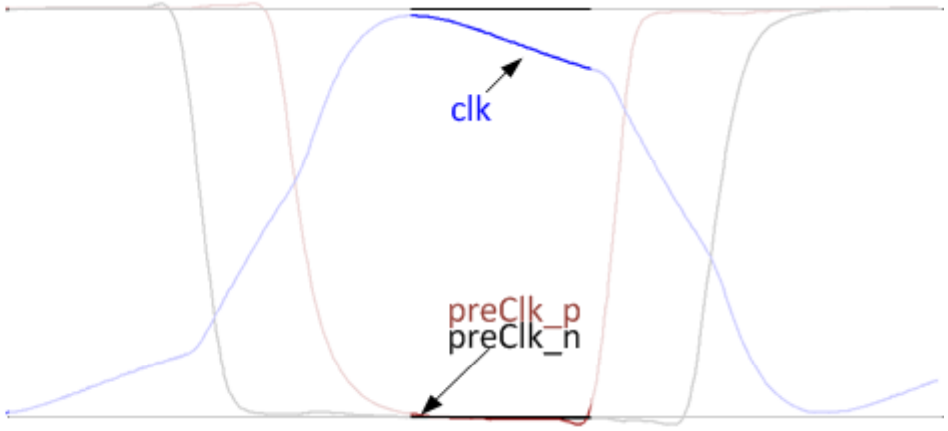

- Pmos on, Nmos off
- *RL* current buildup through the Pmos
- *IR* drop across Pmos and grid → clock voltage increase

Basic Rclk Operation (5/6)

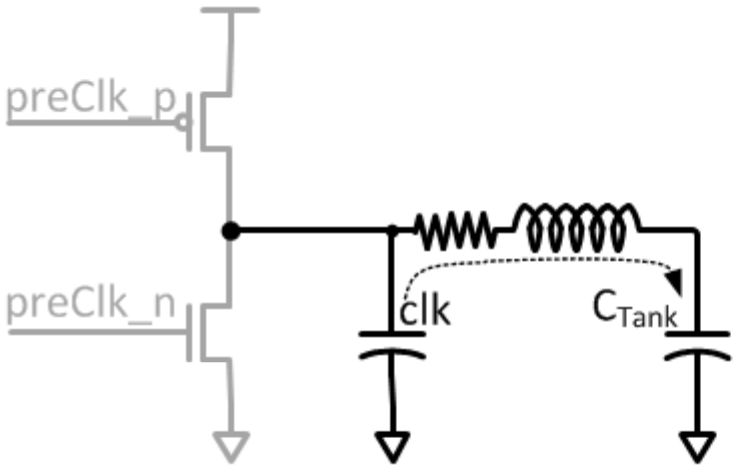

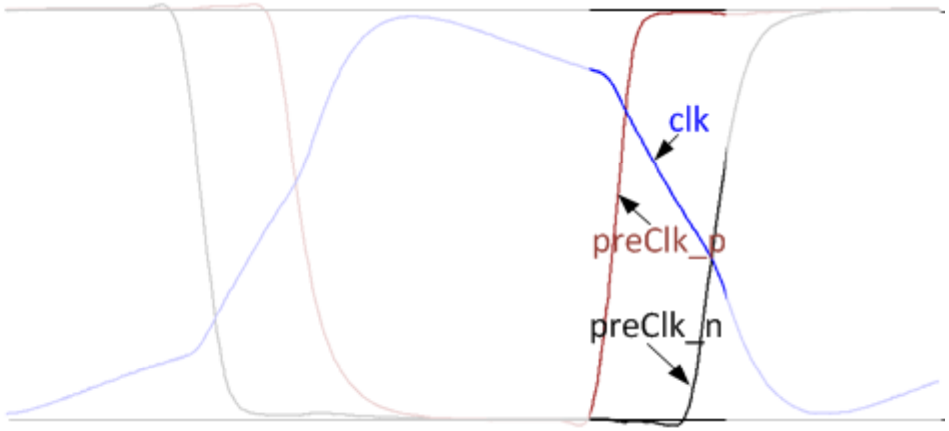

- Both driver devices are off
- *LC* oscillation with initial inductor current to discharge clk
- Clk voltage transition a function of L,C and initial current

Basic Rclk Operation (6/6)

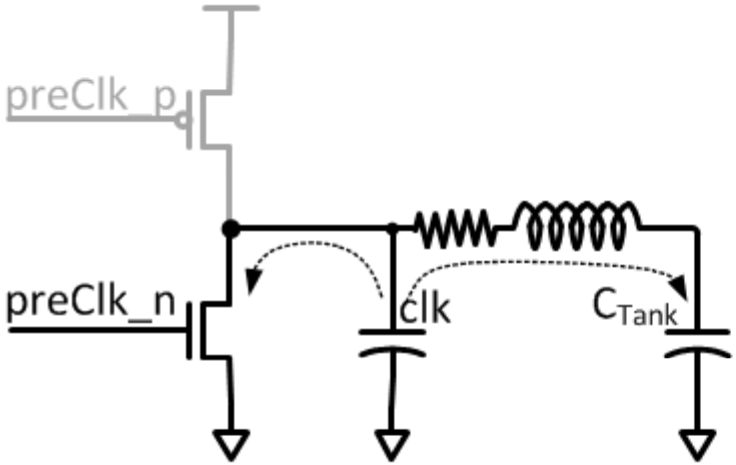

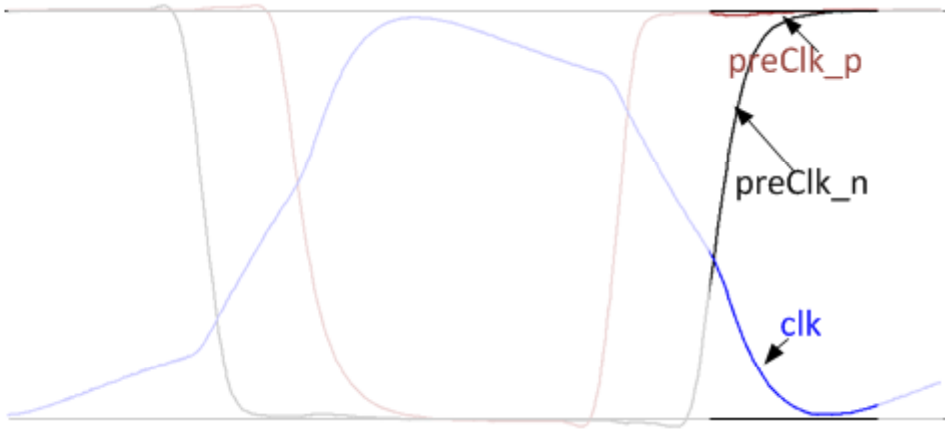

- Nmos on, Pmos off
- Both *LC* and pull-down mechanisms in effect

Cclk and Rclk Waveforms

- Reduced driver strength required for rclk
- Lower rclk slew → insertion delay increase (phase offset)
- Delayed onset of driver devices → rclk_pulse phase offset

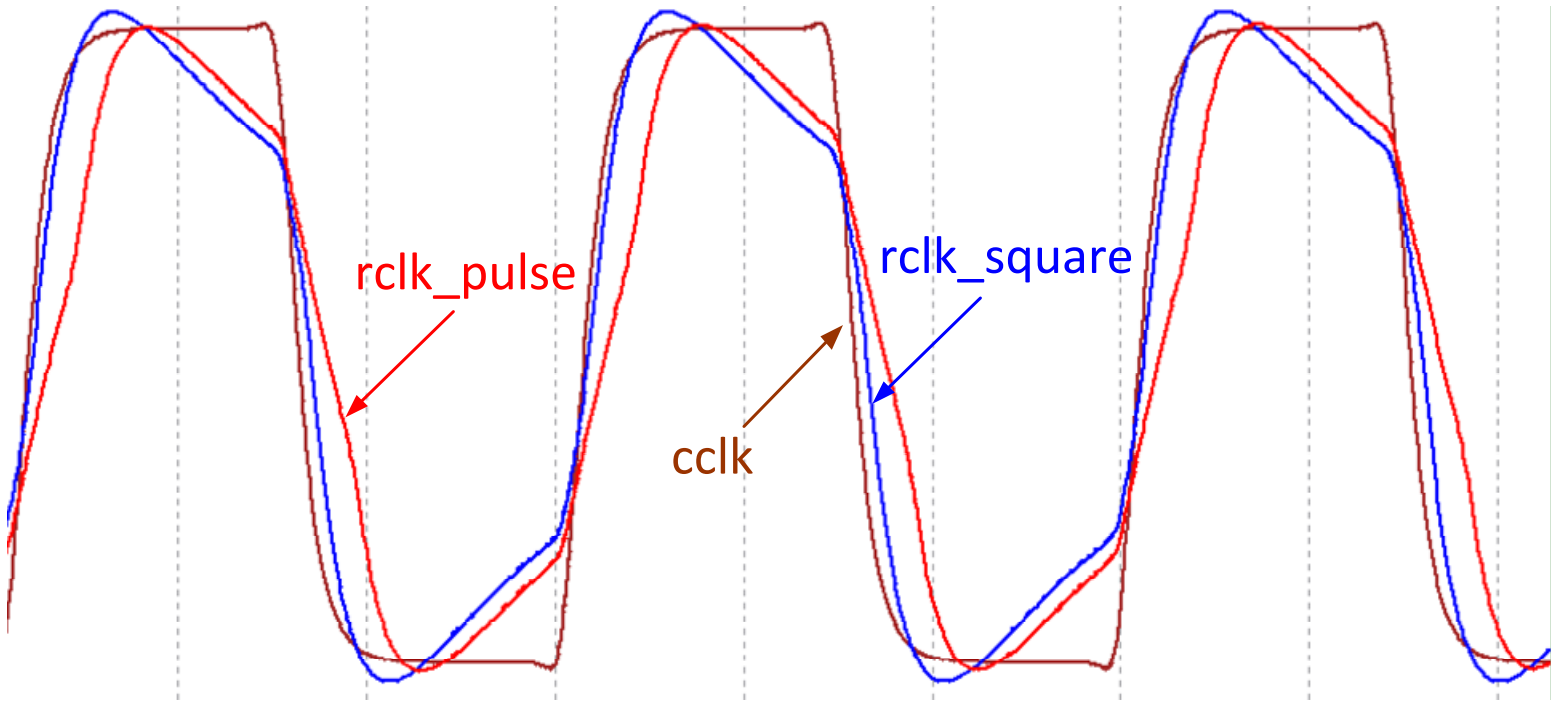

Clock Driver Design

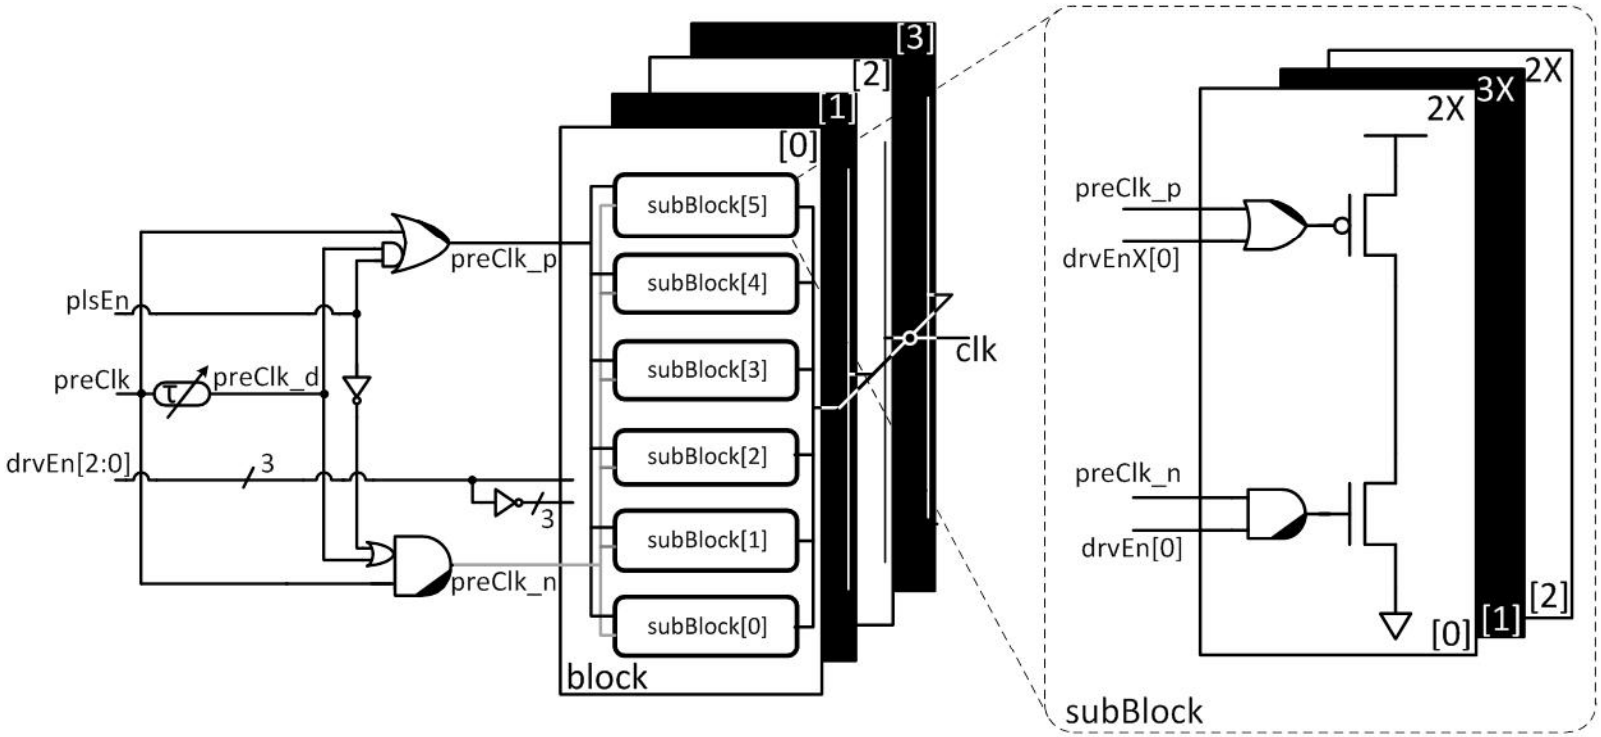

- Clock driver palette with 24 drivers
  - Up to 4 blocks. Each block can contain up to 6 subBlocks
  - Effective granularity and efficiency tradeoff
- Run-time programmable drive strength modulation support
  - Each subBlock consists of 3 banks (2:3:2 ratio)
  - drvEn[2:0] signals allow for n/7 (n=2,3,4,5,7) drive modulation

Driver Design

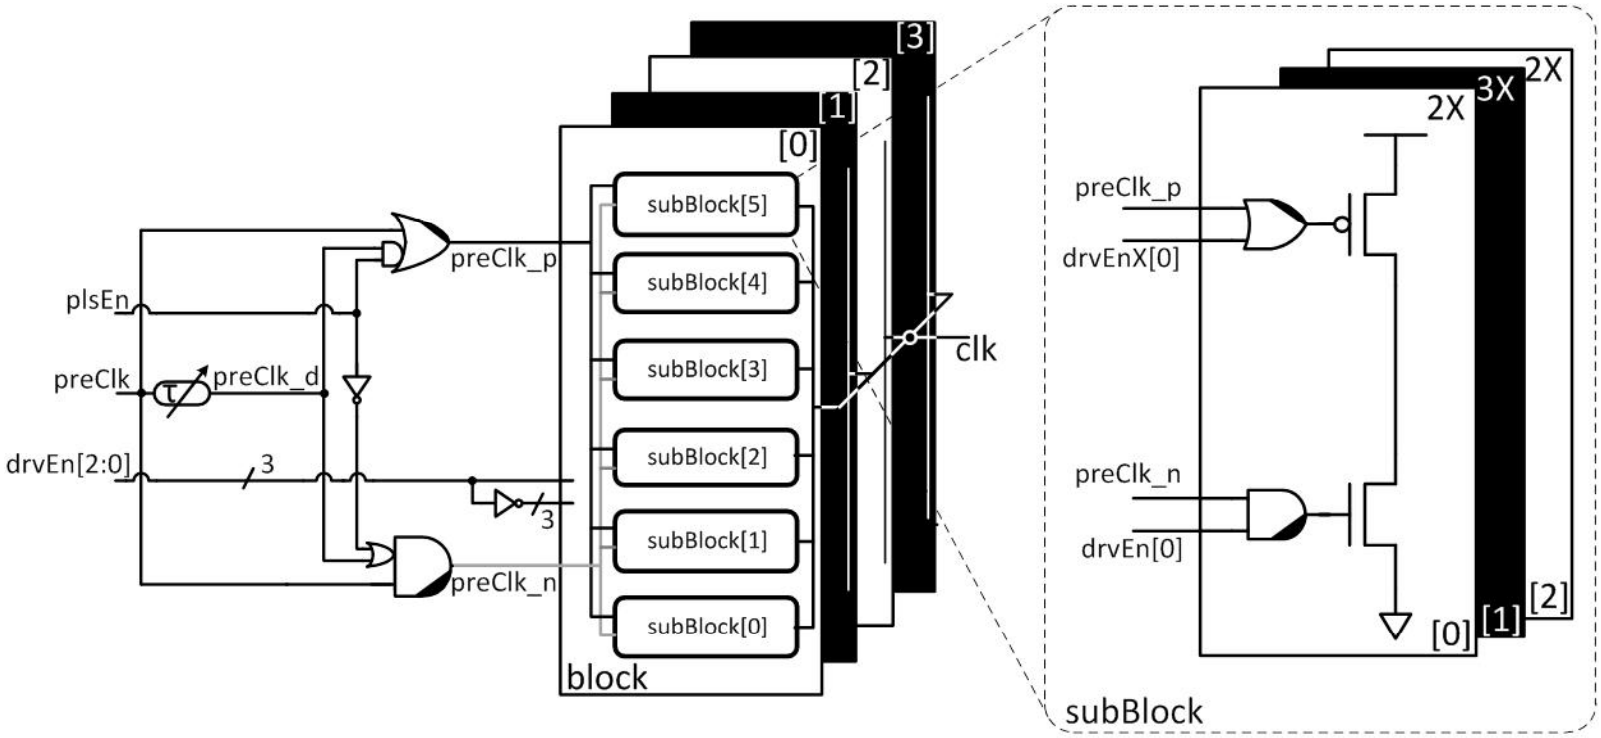

- Split buffer design
  - Skewed pre-drivers for reduced crossover current
  - Allows insertion delay management of rclk w.r.t. cclk, NB and L2 clocks

Driver Design: Pulse Mode

- Subtractive pulse-generator scheme
  - Delay chain used to delay asserting edges of nmos and pmos devices
  - De-asserting transitions not delayed
  - Ontime is a function of input duty cycle and delay amount
- Benefits over traditional pulse generation
  - Lower variation (smaller delay)
  - Support for Off P-state operation
  - Allows PLL duty cycle tuning

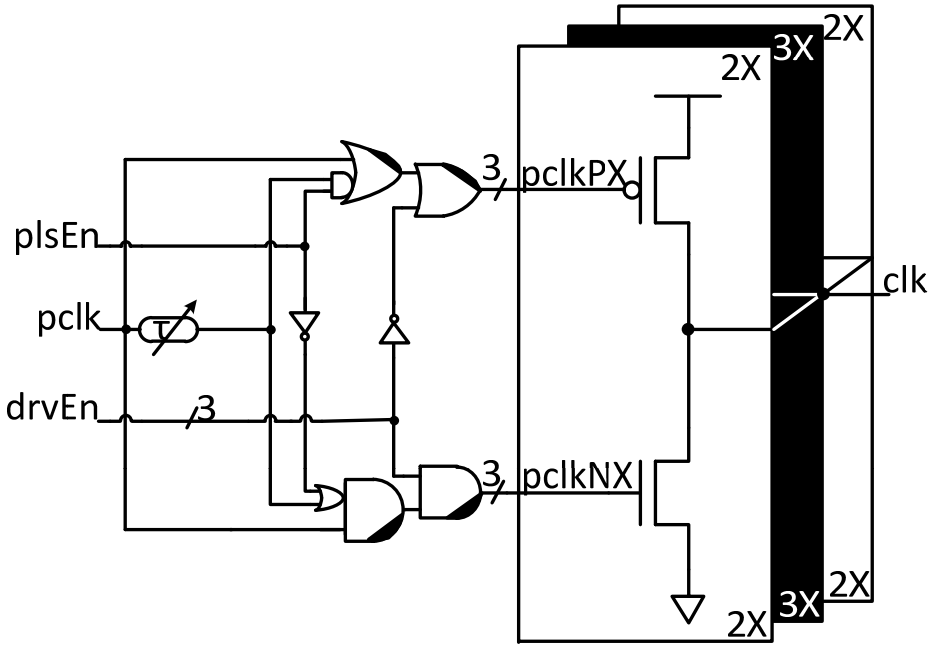

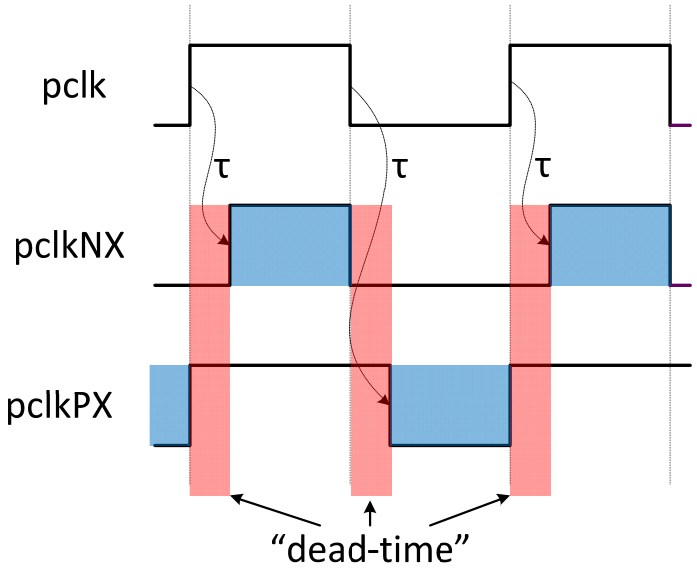

Inductor Design

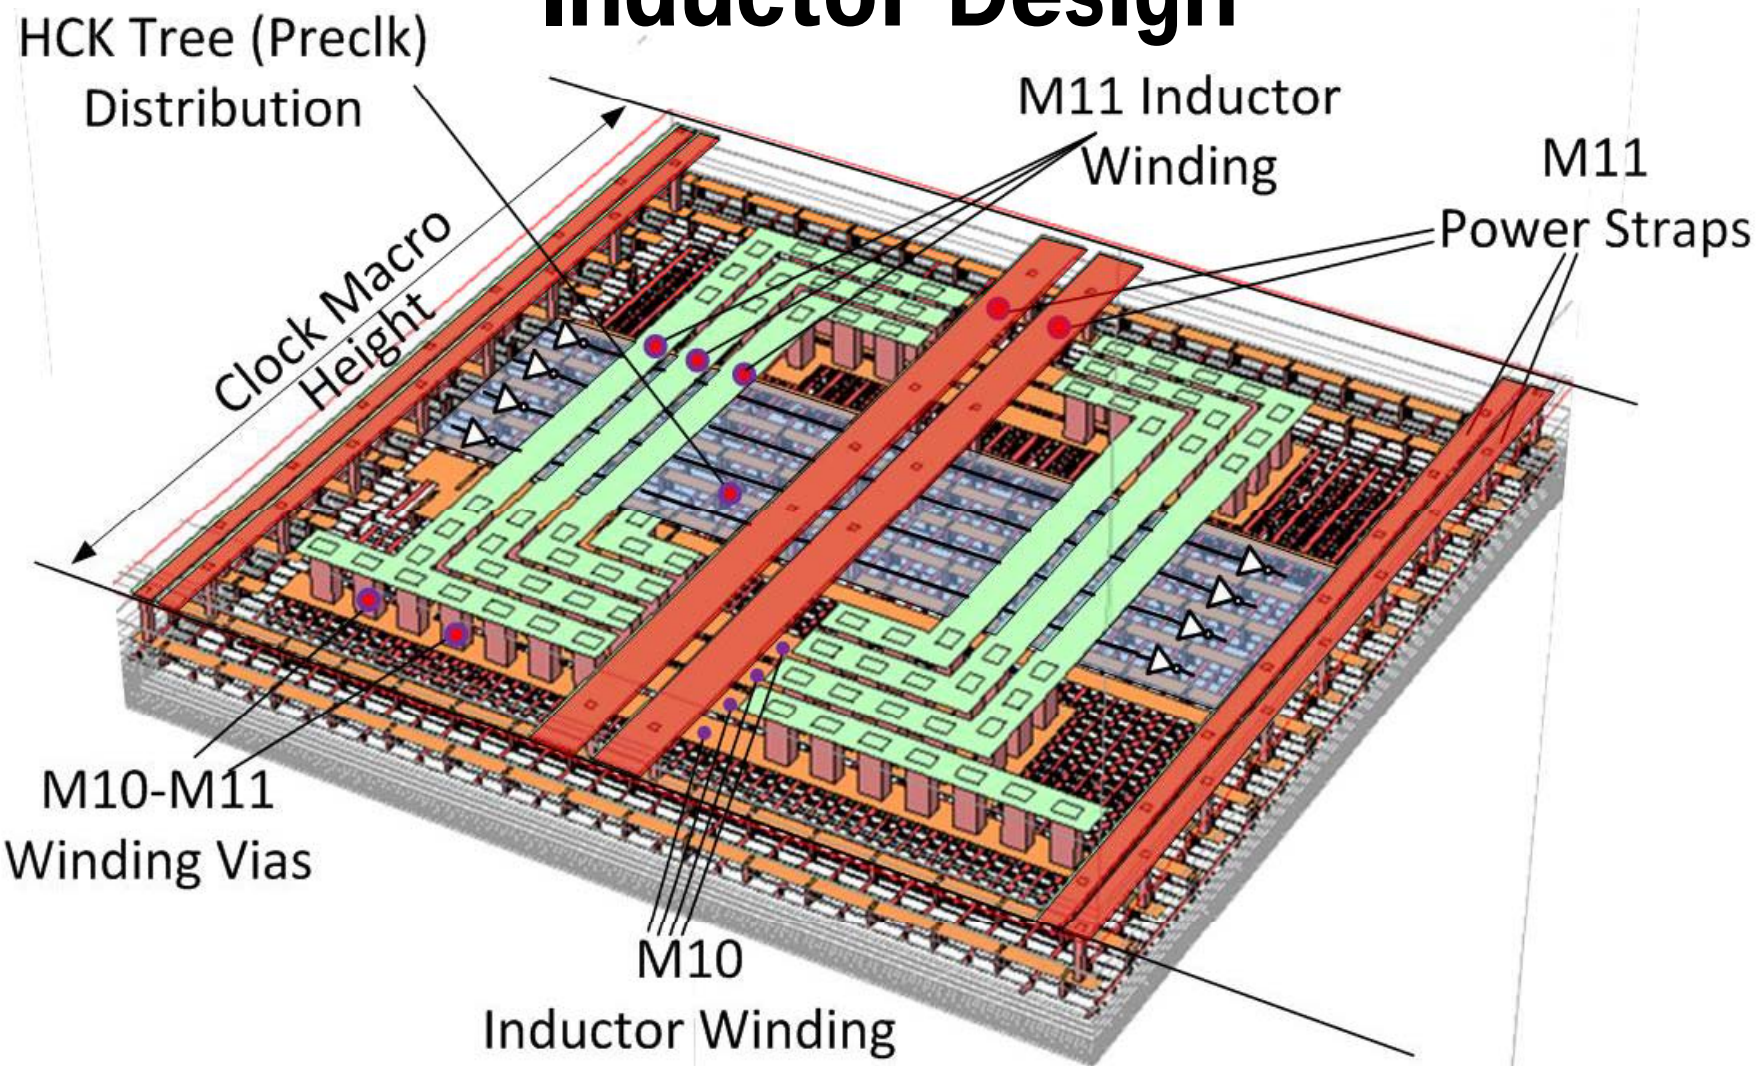

- Clk macro, bump pitch constrain inductor size

Inductor Design (cont'd)

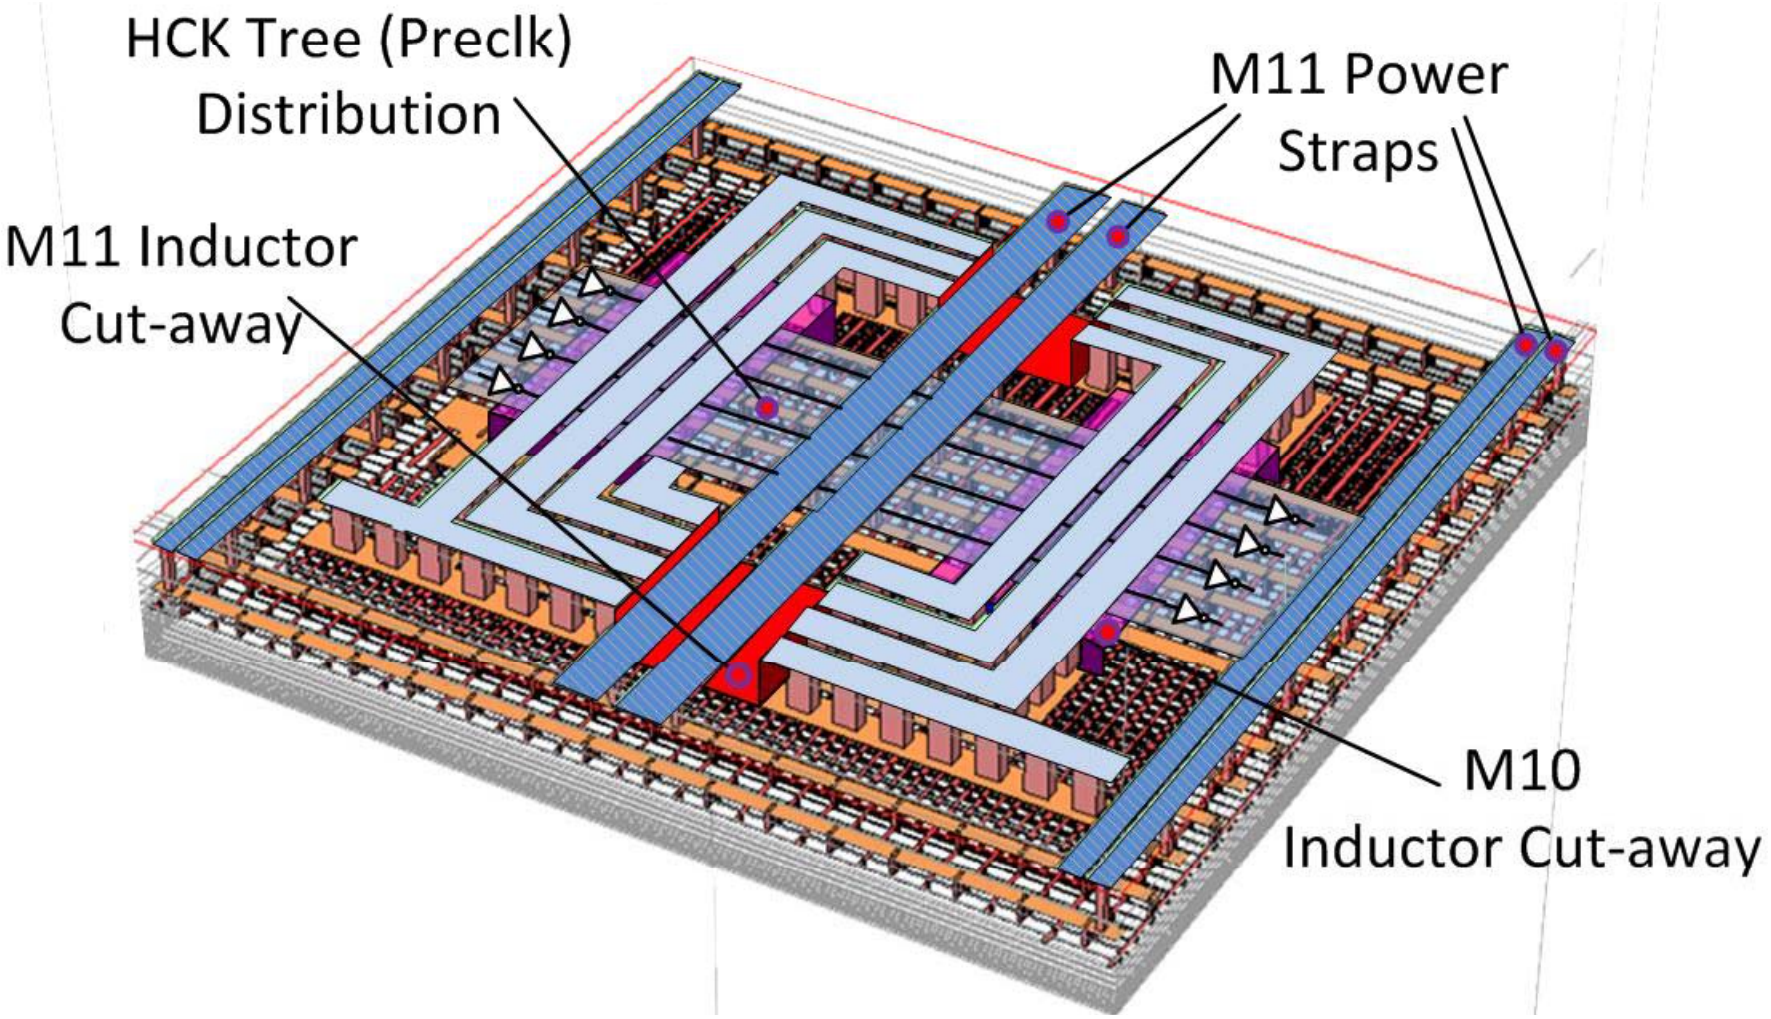

- Metal sharing with existing power → cut-aways
- Center power straps, HCK tree through inductor for mutual inductance cancellation

Inductor Design (cont'd)

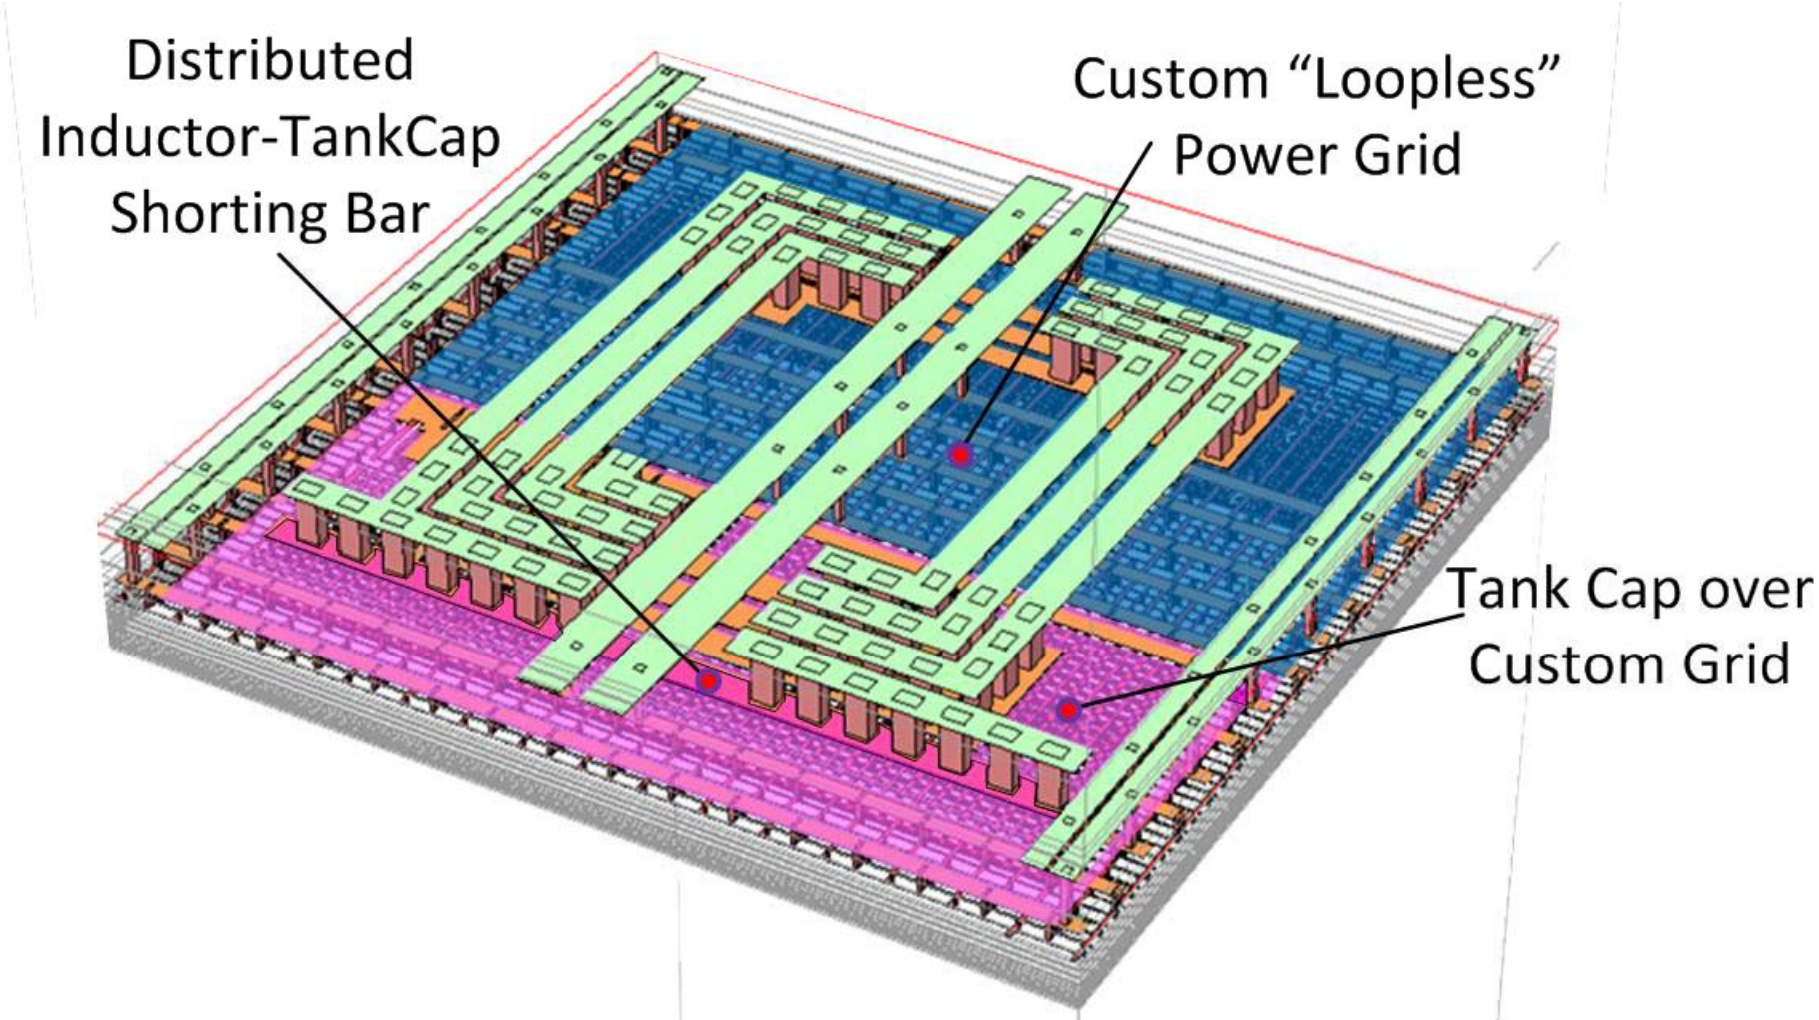

- Custom "loopless" power under inductor to avoid Q degradation due to power grid eddys
- TankCap built in Si, metal to meet capacitance, ESR target

Other rclk Components: Mode Switch

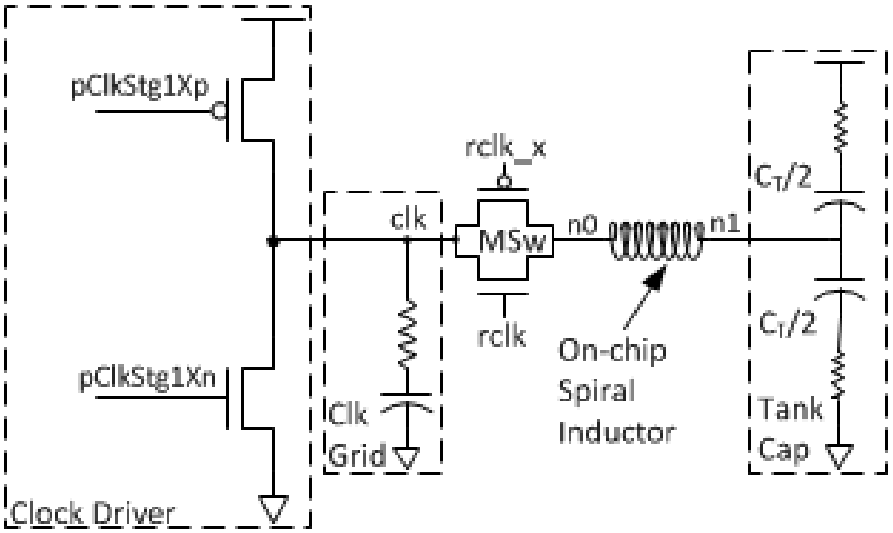

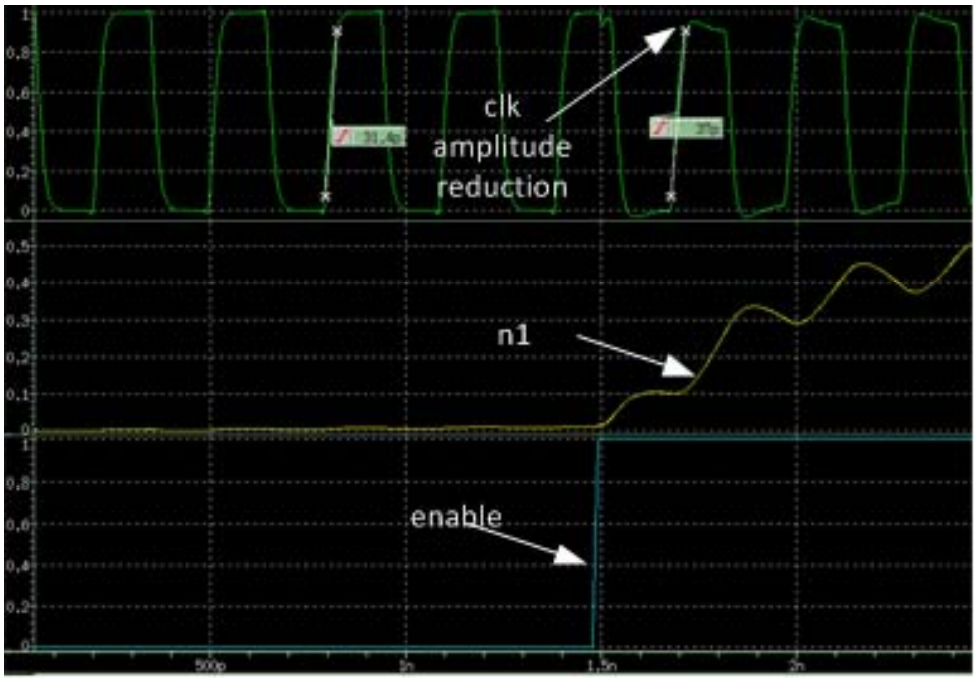

- rclk resistance ↔ cclk loading tradeoff
- Mode switch offers fet resistance
  - Voltage dependent
  - Lower overdrive → Negative Temperature Coefficient
- cclk → rclk mode causes excessive clock grid loading
  - V(n1) not always well defined in cclk mode
  - Results in reduced clock amplitude and degraded slew → timing impact

Mode Switch (Contd.)

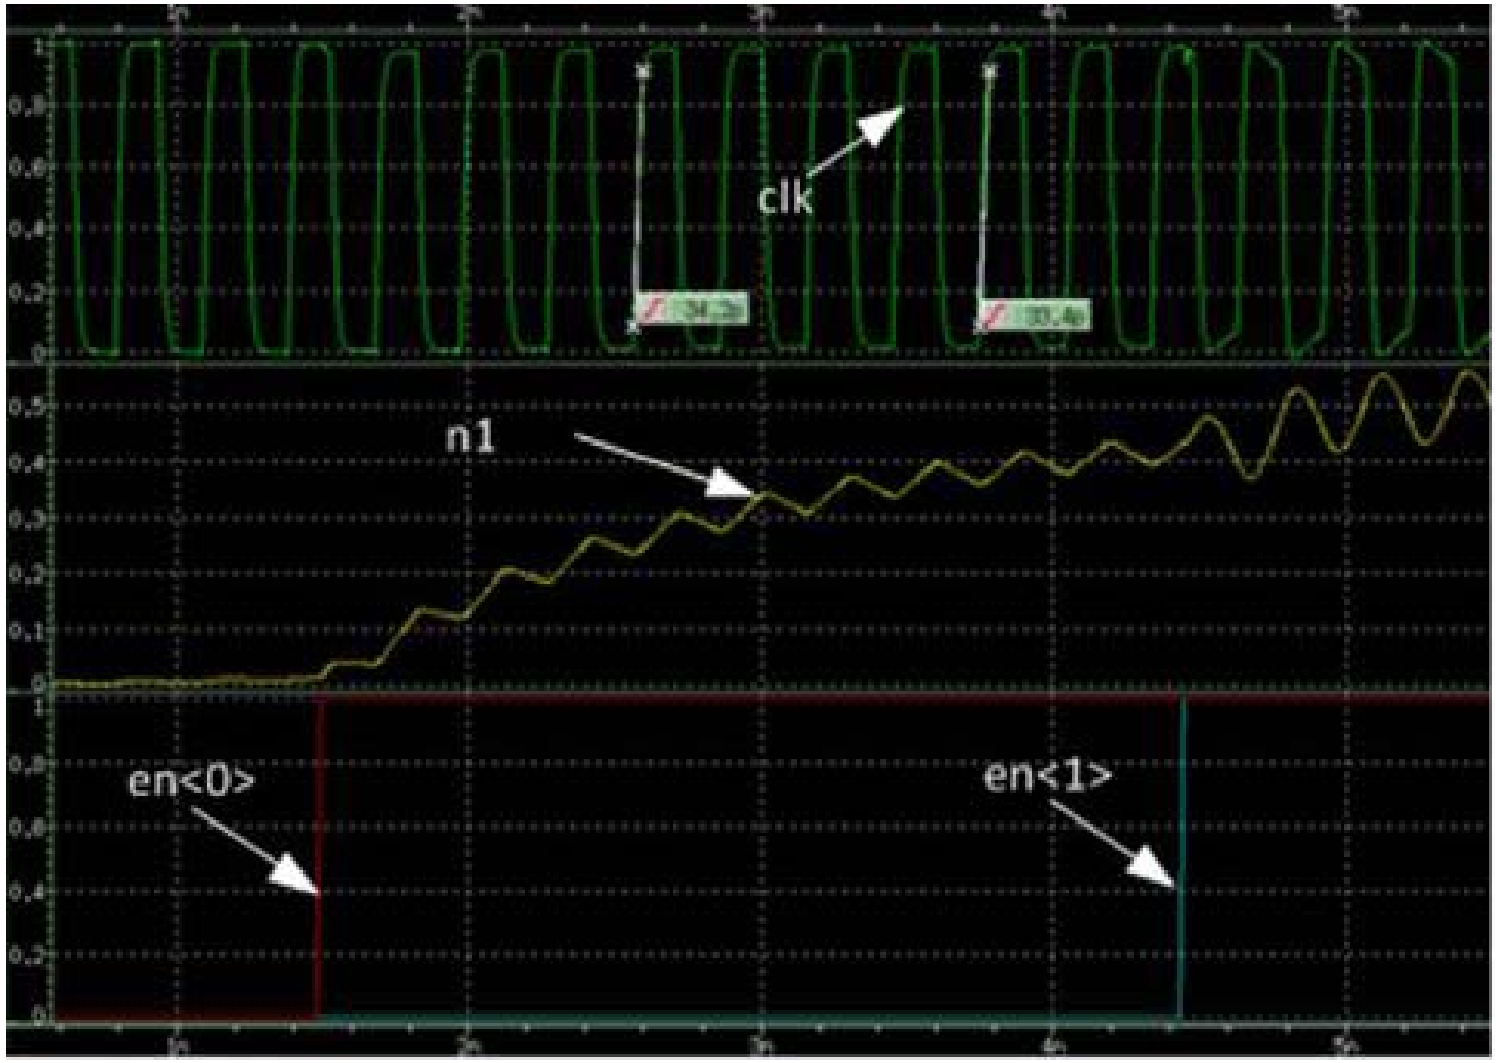

- Staging techniques used
  - Turn on Mode Switch in stages (like in power gating)
  - “Warm” up n0 before making a low resistance connection to grid.

Voltage Overshoot on Mode Switch

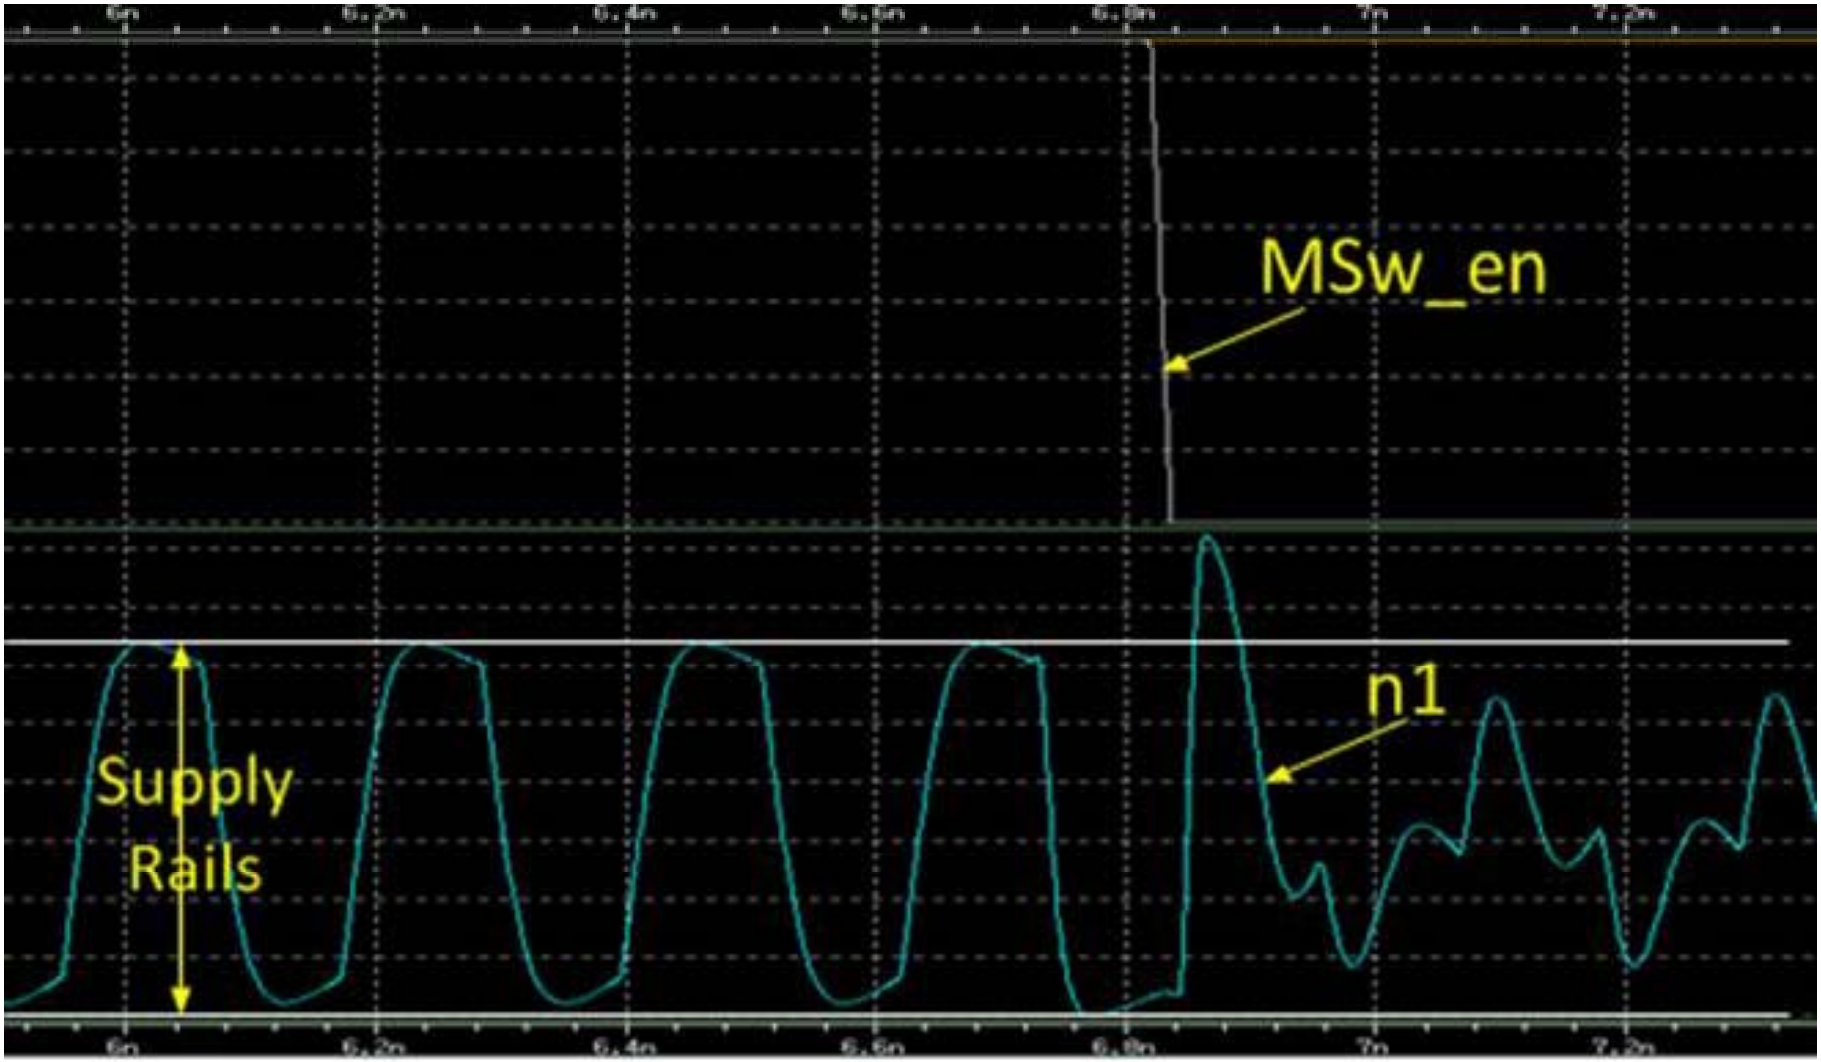

- rclk → cclk transition can result in voltage overshoot on n1.
- Oxide stress poses a reliability issue.

Throttle Switch

- Throttle switch connected across inductor
  - Low resistance help damp overshoot
  - Turns on as Mode Switch turns off

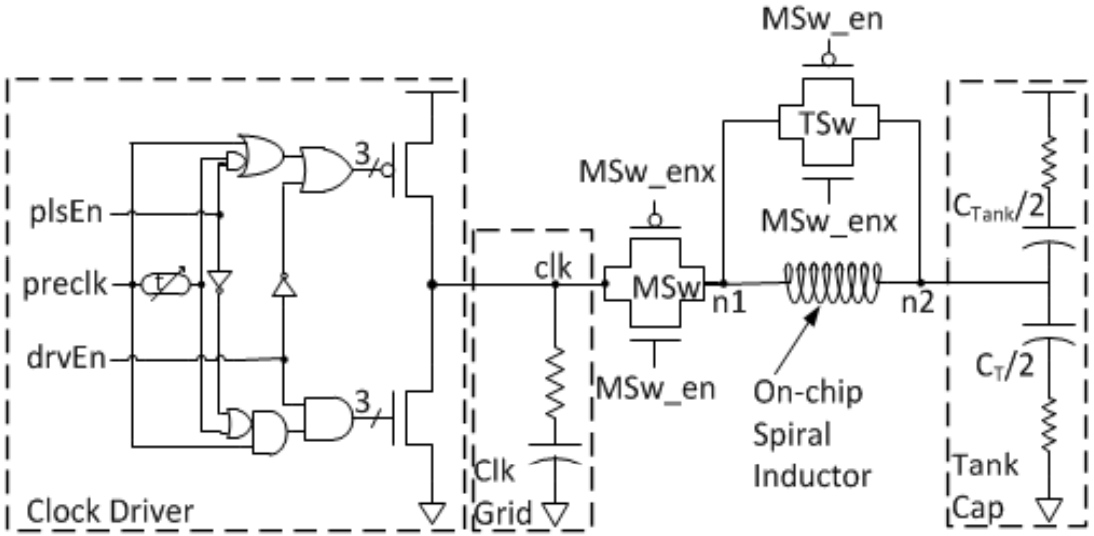

Other rclk Components: C_{Tank}

- C_{Tank} acts as an A.C ground connection.
- ↑C, ↓ESR to serve as a low loss ground connections
  - Low resistance power, ground, inductor connection
  - High bandwidth decap
- Implemented underneath inductor using Si and metal
- Distributed C_{Tank} connection using inductor winding for low resistance contact
- Metal cap built to meet C, ESR requirements

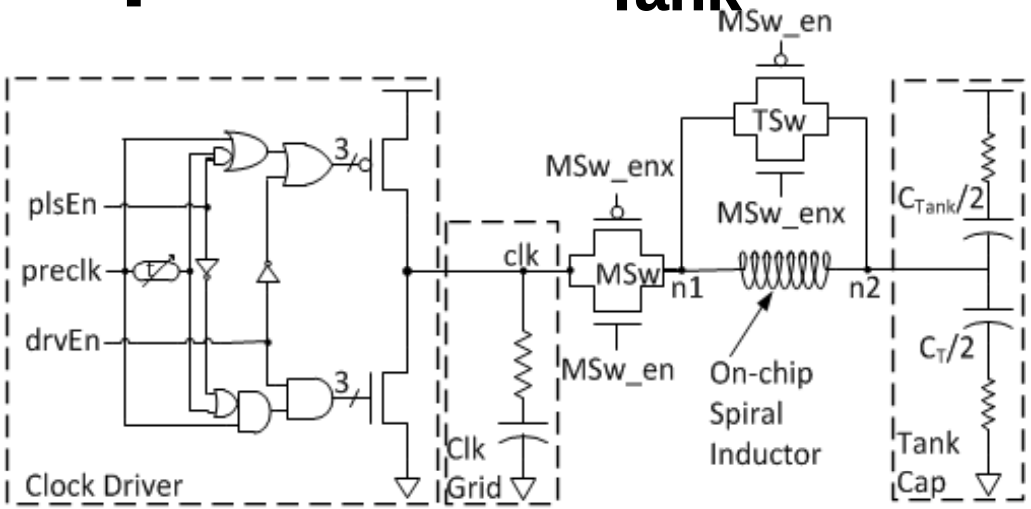

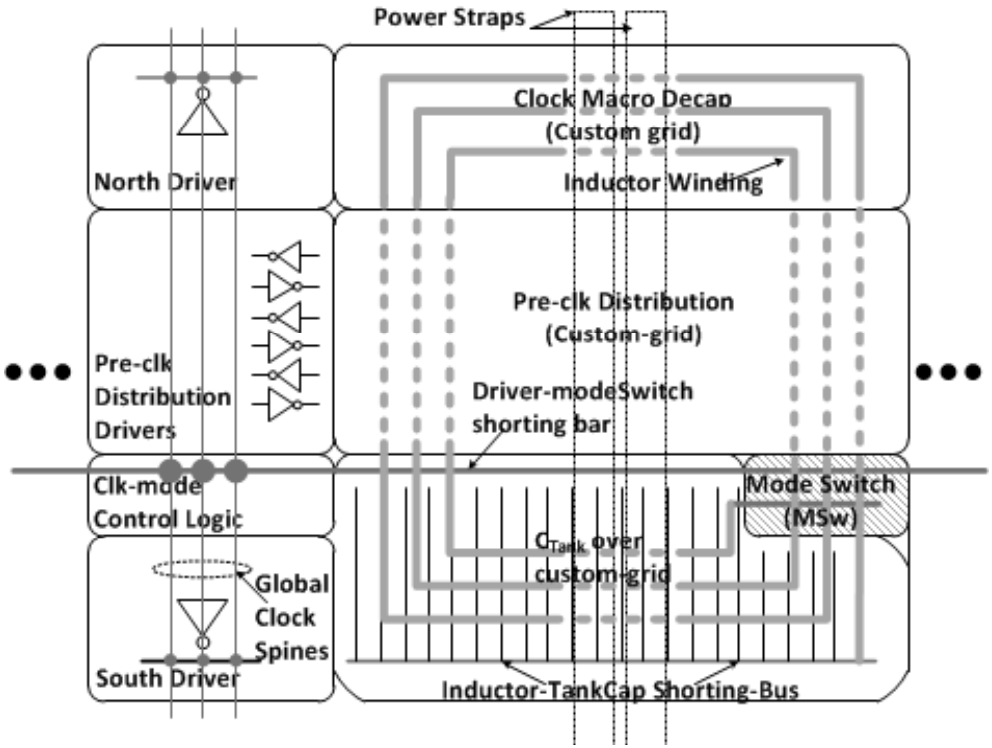

Clock Tuning

- 7ps grid skew target across 22mm^2
- Heavy clock grid loading requires effective strategies for:
  - Grid wire tuning
  - Clock driver tuning
  - Inductor tuning
- Elmore delay-aware local wire routing solution.
- Clock wire tuning algorithm meets target skew with
  - Inductance-aware clock spine geometries
  - Minimal clock spine capacitance

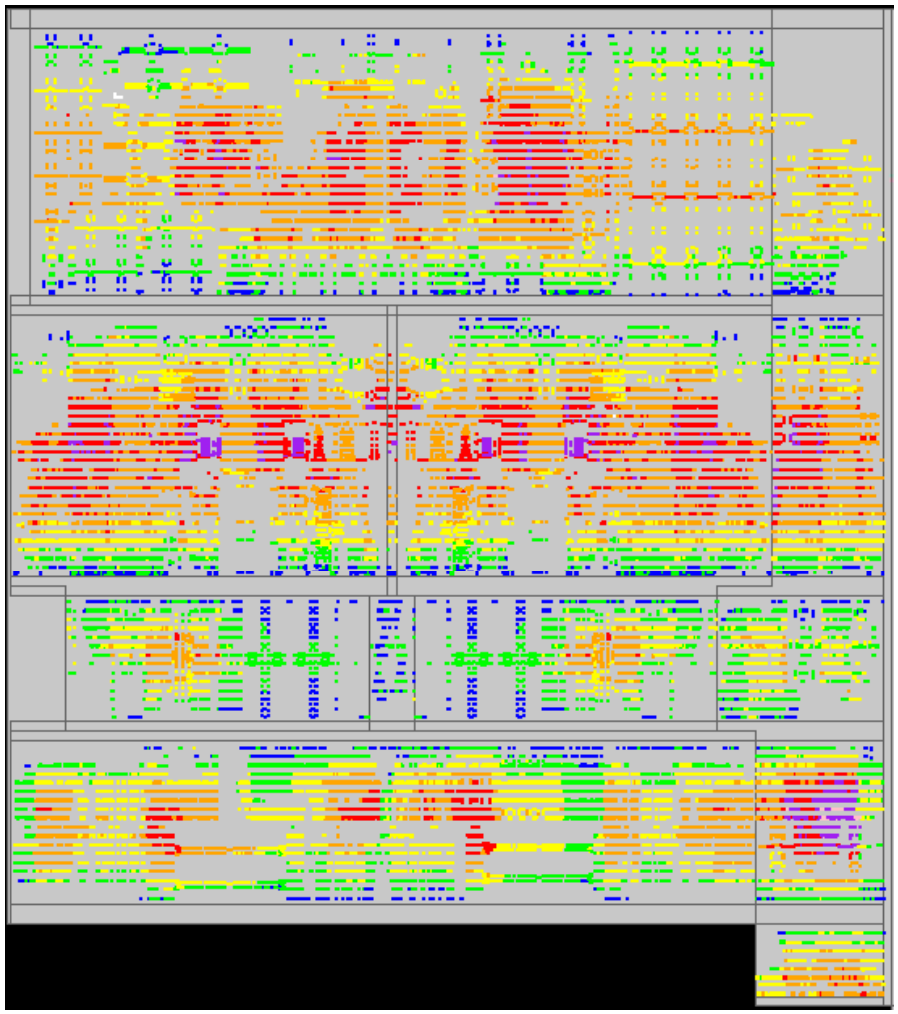

Clock Tuning (contd.)

- Global clock load varies significantly across core
- Effective driver and Inductor allocation key to maintaining clock skew
  - Driver palette size : 24
  - Inductor palette size : 5 (0.5—1.3nH range)
- Iterative Linear Programming-based algorithm for driver and inductor allocation
  1. Start with initial driver/inductor assignment
  2. Linearize problem – Obtain sensitivity matrix for each driver/inductor location
  3. Setup L.P, solve for optimal assignment
  4. Run full chip clock skew analysis
  5. If skew budget not met, goto 2

Clock Tuning (contd)

- Clock skew control for cclk and rclk
  - Wire tuning algorithm to constrain clock latency
  - Iterative LP formulation for optimal driver and inductor assignment
  - Interleaved driver and inductor placement
  - Additional rclk skew impact <1ps

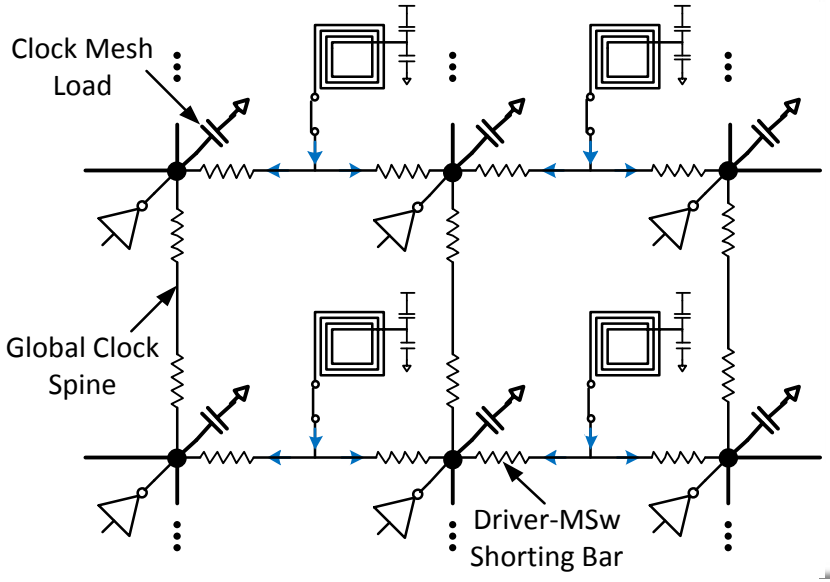

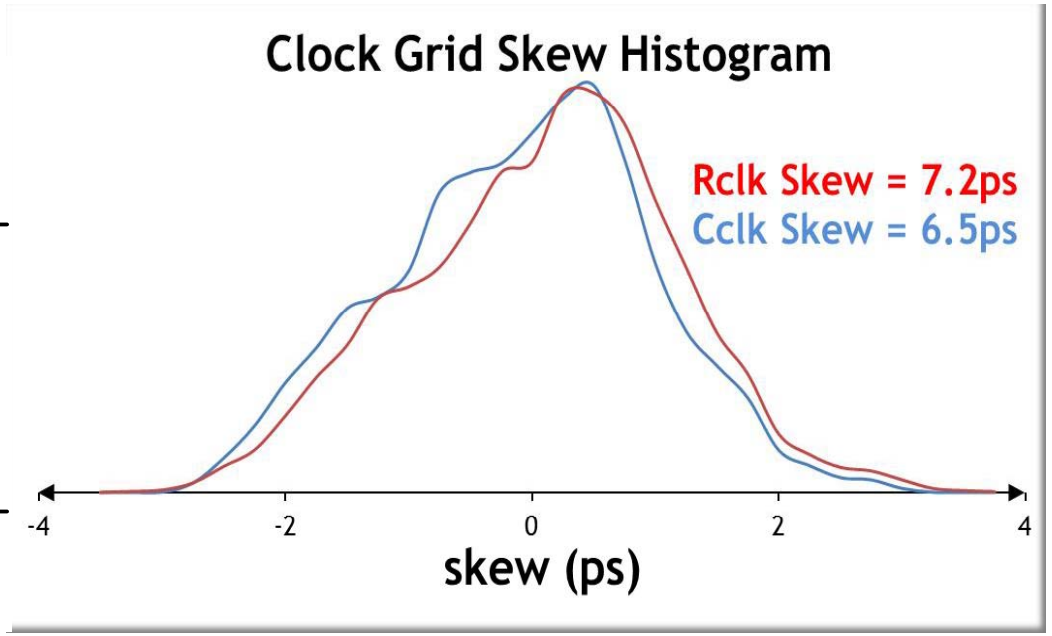

Putting It All Together

- Clock configuration programming during P-State transitions
- Frequency-indexed fuse table to access configuration bits
  - Mode selection (rclk,cclk)
  - Driver strength
  - Pulse_en, Pulse duty cycle
- Source-synchronous transfer
- Cclk mode during P-State transition

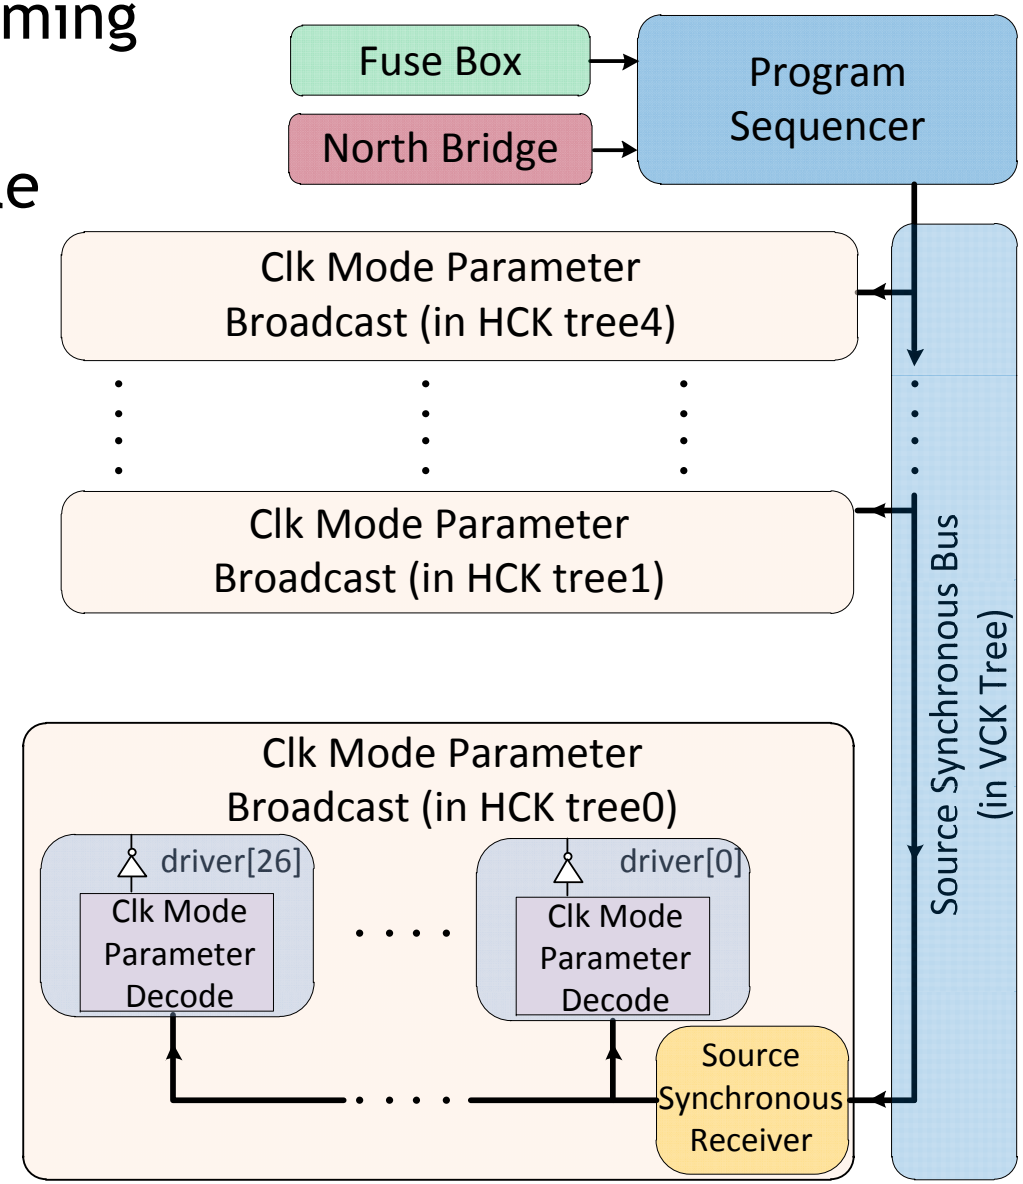

Measured Efficiency (%) vs. Frequency

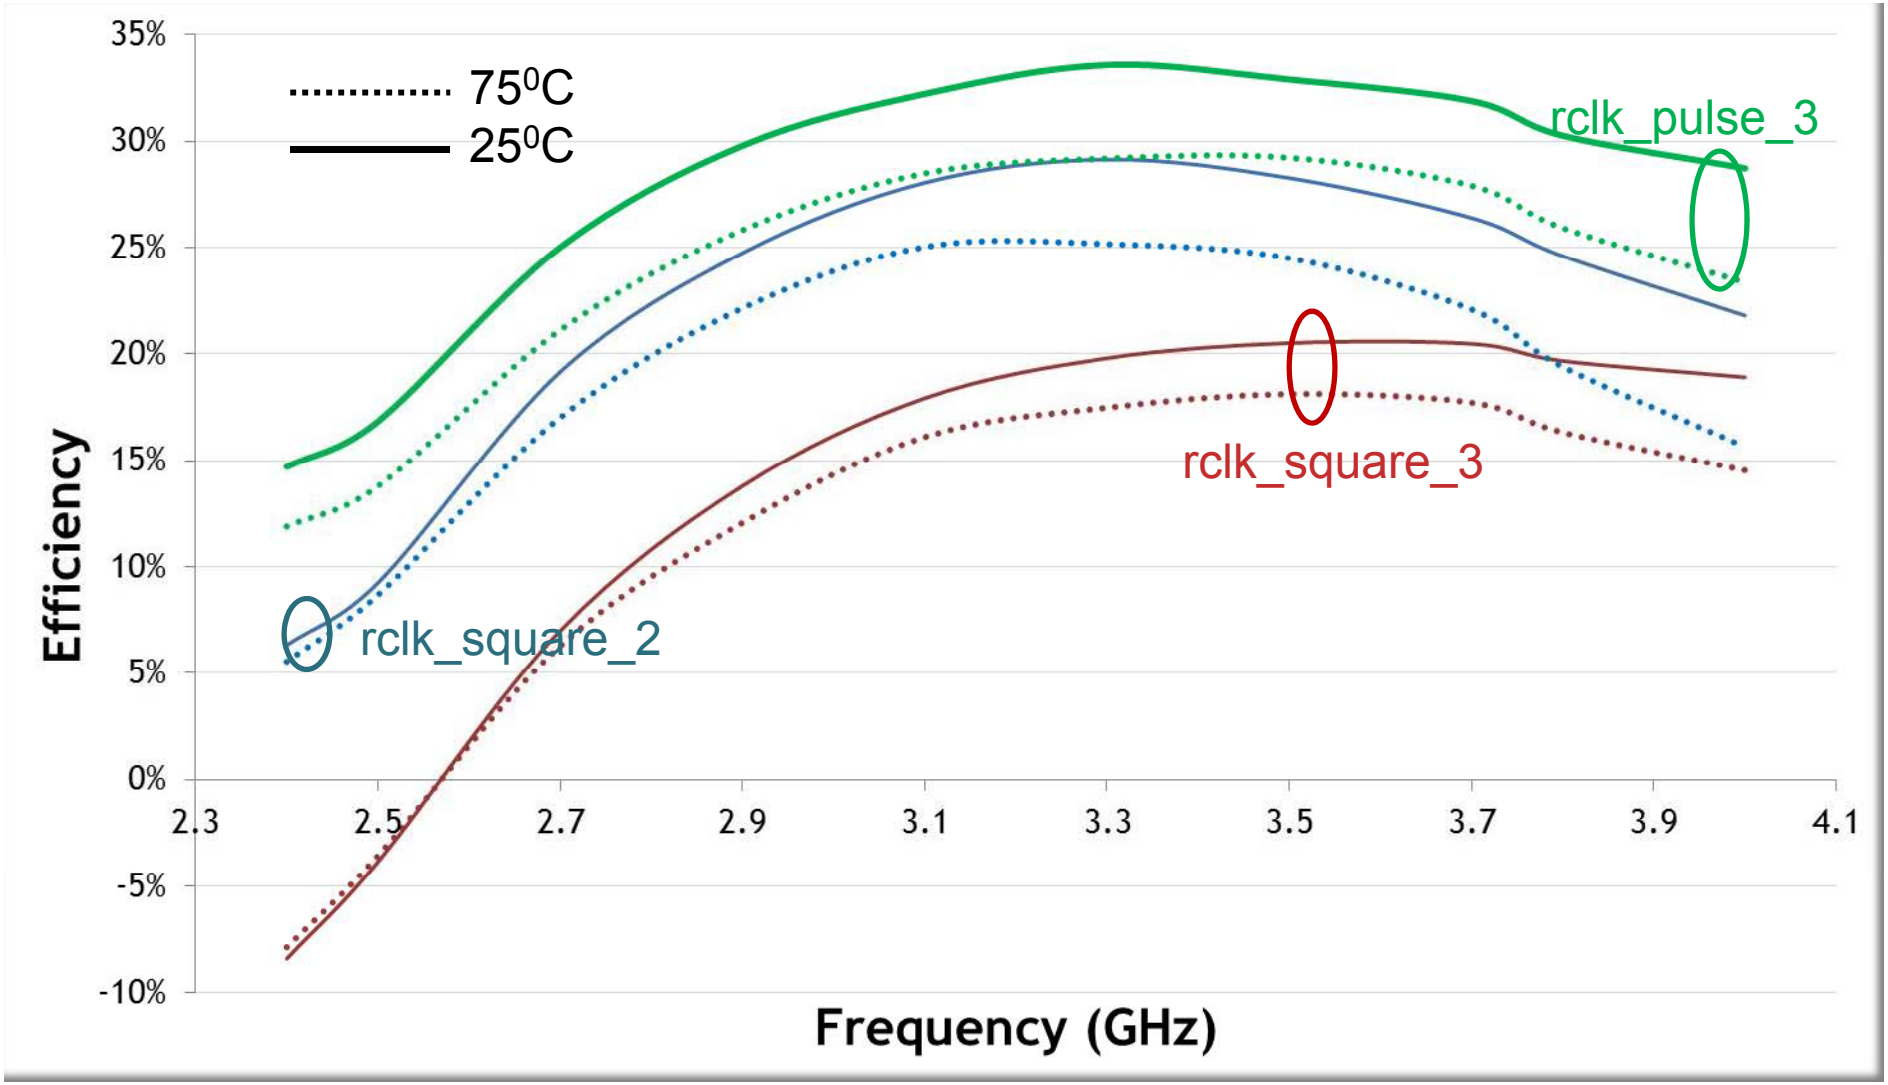

- Efficiency : Percentage clock power savings over cclk
- rclk_square_x → Clock driver strength modulation of x/7

f_{max} Measurements

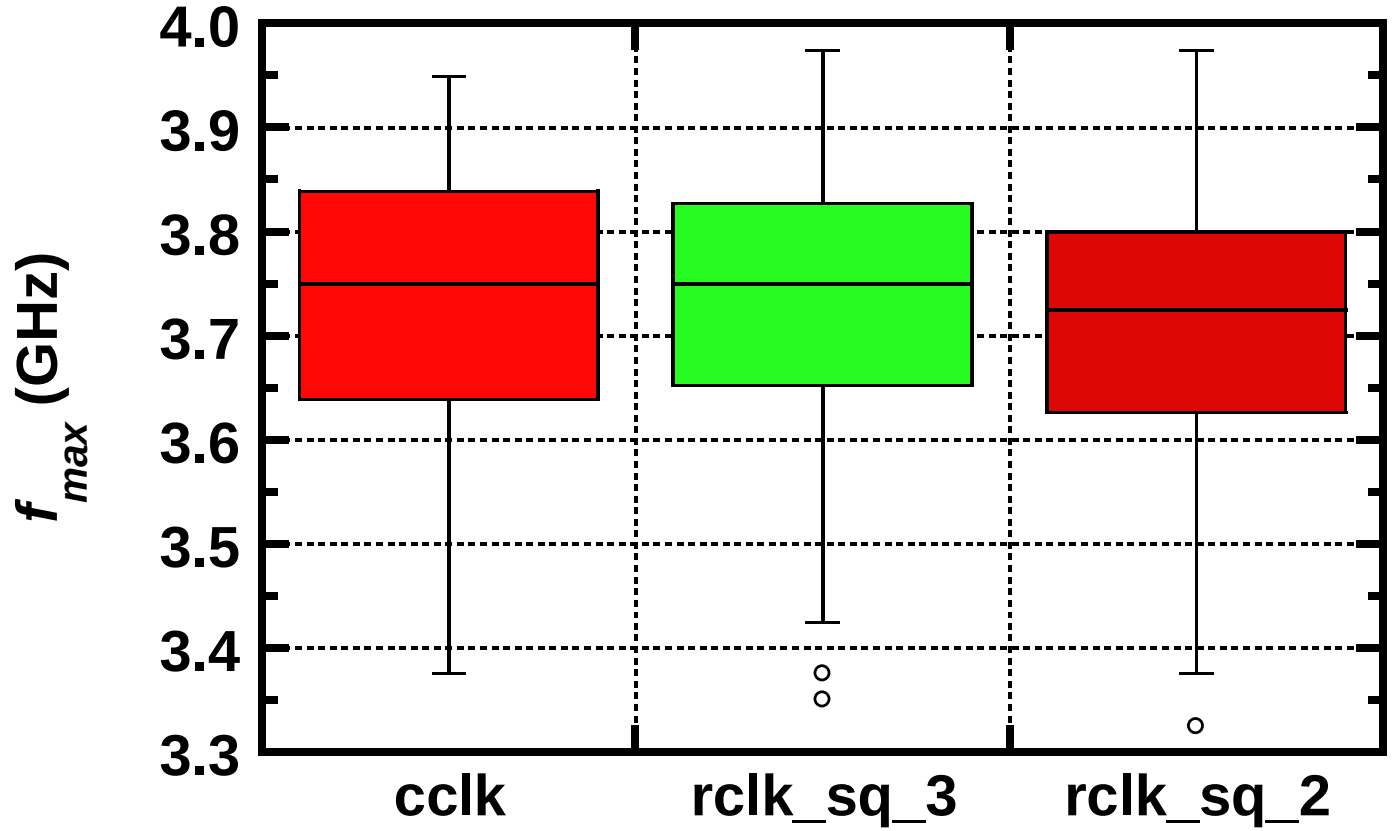

- Frequency-limiting patterns in HST setup
- 0 MHz median, 5 MHz mean (0.13%) frequency overhead in rclk
  - rclk ⇔ cclk phase offset
  - Low-slew rclk waveforms → ↑ variation in timing elements

Rclk Measurement Summary

- Successfully ran SST (System Stress Test) over 2 weeks
- Latest Fmax impact data on a larger set of parts shows an Fmax overhead of ~0.2%
- Up to 34% energy efficiency achieved in the global clock using Pulse Mode
  - Not production ready due to excessive Fmax impact driven by phase offset to NB and L2 clock interface
  - Phase offset issue resolved in current design
- Temperature effect : Overall efficiency degradation with T
  - Positive temperature coefficient due to metal resistance
  - Negative temperature coefficient due to low-overdrive MSw
- Traces with higher activity provide additional efficiency
  - Increased clock load dominates additional crossover current.

Conclusion

- Dual-mode resonant clock design in 32nm SOI
  - Conventional mode : < 2.9GHz
  - Resonant mode : > 2.9GHz to f_{max}
- Power savings
  - Clock power: ↓ 25%
  - Average application power (core): ↓ 4.5%
  - Idle power (core): ↓ 10%
- Built upon existing clock infrastructure
- No CMOS technology modification

Acknowledgments

AMD

- Manivannan Bhoopathy
- Kyle Viau
- Alan Arakawa
- Tom Meneghini
- Denny Renfrow
- Vasant Palisetti
- Kevin Hurd
- David Johnson
- Rohit Kumar[1]
- Alvin Loke
- Bill McGee

[1] Currently with Apple Computers

Cyclos Semiconductor

- Fred Brauchler
- Joohee Kim
- Jerry Kao