

EE222 Lecture 15 Feb. 26, 2019 Low Power Design of Logic Circuits:

- Combinational circuits
- Mixed signal (digital and analog) circuits for processing in memory (PIM)
- Sequential circuits
- clock distribution

Combinational Vs. Sequential Circuits

Combinational (or Combinatorial) Networks/Circuits

- Circuit **without storage**
- Outputs depend only on its current inputs
- Examples: NAND gate, look-up table (LUT)

Vs.

Sequential Networks/Circuits

- Circuit **with storage** elements
- Outputs depend on present inputs and also on history of inputs
- Examples: RAM, finite state machine (FSM)
- Normally, **combinational cells AND storage elements**

ISSCC 2018 / SESSION 31 / COMPUTATION IN MEMORY FOR MACHINE LEARNING / 31.1

31.1 Core-RAM: An Energy-Efficient SRAM with Embedded Convolution Computation for Low-Power CNN-Based Machine Learning Applications

Arvind K. Sivasubramanian, Anupam P. Choudhary, Massachusetts Institute of Technology, Cambridge, MA

Convolutional neural networks (CNNs) provide state-of-the-art results in a wide variety of machine learning (ML) applications, ranging from image classification to speech recognition. However, they are very computationally intensive and require large amounts of storage. Recent work shows trends reducing the size of the CNNs [1], improving energy efficiency [2], and reducing the need for large amounts of storage [3]. This work is a significant reduction in the amount of storage required for the CNNs, making it possible to store them directly in memory. However, in a conventional digital implementation [2, 3], storing the CNNs and the partial sums from the embedded SRAMs require a large amount of storage, which is not practical for low-power applications.

Figure 31.1.1 shows the overall architecture of the Core-RAM (CRAM) array. It is divided into 16 local arrays, each with 16 rows to reduce the area of the SRAMs and the local array overhead. Each local array stores the binary weights (0/1) in the 16 SRAMs (Figure 31.1.1). The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs. The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs. The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs.

charge sharing is used to integrate the base of the 2nd voltage rails with a reference local column that replicates the local bit-line capacitance. This process continues until the voltage of the 2nd voltage rail reaches the reference level, at which point the 2nd output figure. This figure represents the convolution result and no further SRAM access is required for the SRAM. Figure 31.1.4 shows the waveforms for a typical operation cycle. To reduce the effect of SRAM access, the SRAM is multiplexed to be used at the input of the SRAM to flip the inputs on alternate cycles.

The Core-RAM array is implemented in a 65nm 1P4M1S1 process. Figure 31.1.5 shows the measured SRAM access results, which is used in the 16 rows by setting the SRAM access voltage (V_{SA}) to estimate the SRAM access voltage (V_{SA}). The SRAM access voltage is compared to an external V_{SA} to reduce the SRAM access voltage. The SRAM access voltage is compared to an external V_{SA} to reduce the SRAM access voltage. The SRAM access voltage is compared to an external V_{SA} to reduce the SRAM access voltage.

Figure 31.1.6 shows the overall architecture of the Core-RAM array. It is divided into 16 local arrays, each with 16 rows to reduce the area of the SRAMs and the local array overhead. Each local array stores the binary weights (0/1) in the 16 SRAMs (Figure 31.1.1). The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs. The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs. The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs.

Figure 31.1.7 shows the overall architecture of the Core-RAM array. It is divided into 16 local arrays, each with 16 rows to reduce the area of the SRAMs and the local array overhead. Each local array stores the binary weights (0/1) in the 16 SRAMs (Figure 31.1.1). The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs. The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs. The readout circuitry for each local array is implemented in the SRAMs, which are connected to the SRAMs.

Computation in Memory (CIM)

(also called Processing in Memory (PIM))

Ref. Convolution-RAM: An Energy Efficient SRAM with Embedded Convolution Computation for Low Power CNN (convolutional neural network) - based Machine Learning Applications

ISSCC 2018 A. Biswas, et al. MIT

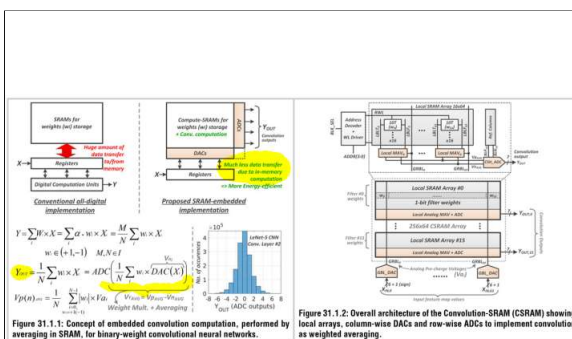



Figure 31.1.1: Concept of embedded convolution computation, performed by averaging in SRAM, for binary-weight convolutional neural networks.




Figure 31.1.2: Overall architecture of the Convolution-SRAM (CSRAM) showing local arrays, column-wise DACs and row-wise ADCs to implement convolution as weighted averaging.

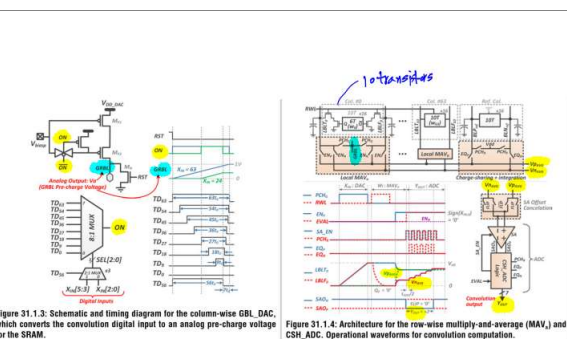


Figure 31.1.3: Schematic and timing diagram for the column-wise GBL DAC, which converts the convolution digital input to an analog pre-charge voltage for the SRAM.

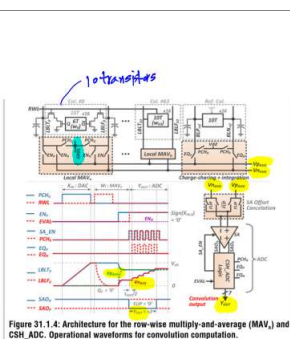
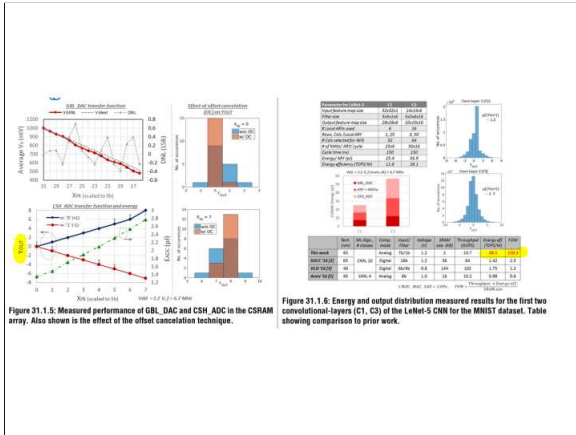


Figure 31.1.4: Architecture for the row-wise multiply-average (MAV) and CSRAM. Operational waveforms for convolution computation.



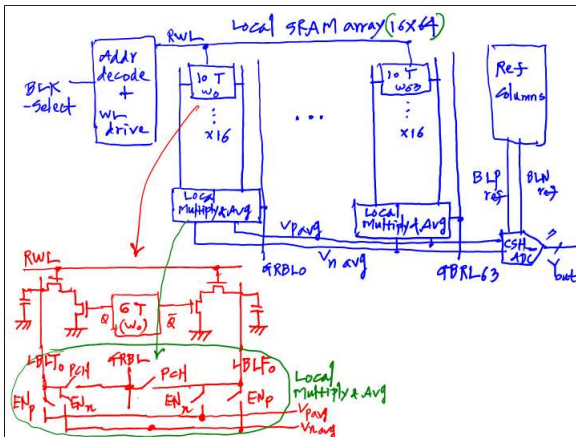
$$Y = \sum_i w_i \times x_i = \sum_i \alpha_i \times v_i \times x_i = \frac{M}{N} \sum_i \alpha_i \times x_i$$

$$Y_A = \frac{1}{N} \sum_i w_i \times x_i = \text{ADC} \left[\frac{1}{N} \sum_i \alpha_i \times \text{DAC}(x_i) \right]$$

weight multiply & average

$$V_{\text{average}} = \frac{1}{N} \sum_{i=1}^{N-1} |w_i| \times V_{A_i}$$

\leftarrow reduces the storage requirement for weights



| Technology | Machine learning | Comparison mode | Input | Voltage | SRAM size | Thru put |
|--------------------|------------------|-----------------|-------|---------|-----------|-----------|
| 65 nm | CNN | Analog | 7b | 1.2V | 2KB | 10.7 Gops |
| 65 nm (asymmetric) | CNN | Digital | 16b | 1.2V | 36 KB | 6.4 Gops |

Energy Efficiency

improvement 20x



28.1 TOPS/Watt

1.42 TOPS/Watt

Figure of Merit (FOM)


= Throughput x Energy Efficiency

= 150.3 x 70 improvement

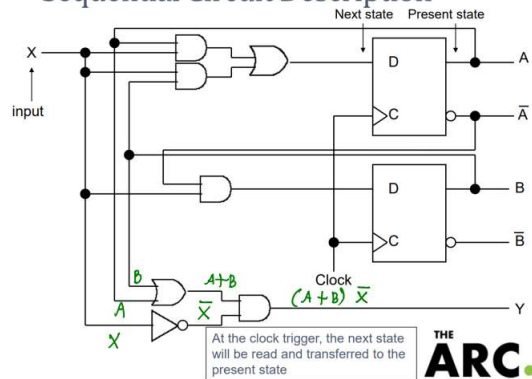


Flip-Flop Timing

- Set-up time: t_s
 - Input needs to be stable before trigger
- Hold time: t_h
 - Input needs to be stable after trigger
- Propagation delay: t_p
 - Some delay from trigger to output change
- Example: Negative edge triggered flip-flop



Sequential Circuit Description



Input Equations

$$\begin{aligned} A_{\text{next}} &= A_{\text{present}}X + B_{\text{present}}X \\ B_{\text{next}} &= A'_{\text{present}}X \\ Y &= (A_{\text{present}} + B_{\text{present}})X' \end{aligned}$$


Next state in terms of input and present state

Output in terms of input and present state

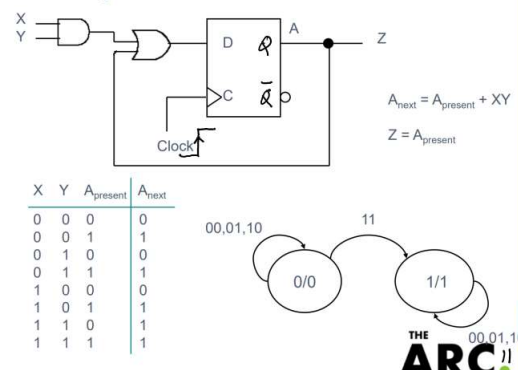
State Table

| Present State | | Input | Next State | | Output |
|---------------|---|-------|------------|---|--------|
| A | B | X | A | B | Y |
| 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 1 | 0 | 1 | 0 |
| 0 | 1 | 0 | 0 | 0 | 1 |
| 0 | 1 | 1 | 1 | 1 | 0 |
| 1 | 0 | 0 | 0 | 0 | 1 |
| 1 | 0 | 1 | 1 | 0 | 0 |
| 1 | 1 | 0 | 0 | 0 | 1 |
| 1 | 1 | 1 | 1 | 0 | 0 |

State Diagram



Example of Moore Model




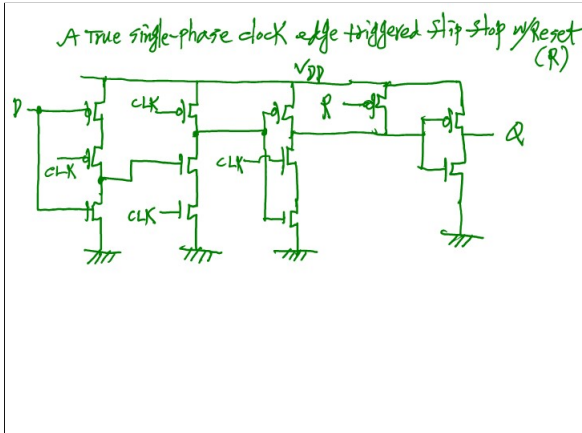


FIGURE 1: (a) CMOS logic and (b) an example of single-phase clocking.




FIGURE 2: Domino logic.




FIGURE 3: NORA logic.





FIGURE 4: (a) A single-phase frequency divider reported by Oguey and Vittoz in 1973 and (b) a latch filed for patent by Piguet in 1974.

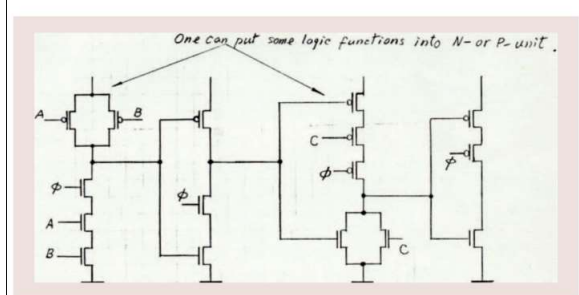


FIGURE 5: Yuan's original drawing of a TSPC circuit.




FIGURE 6: (a) A dynamic latch with a single clocked device, (b) cascaded TSPC stages, and (c) a three-stage master-slave flip-flop operating as a frequency divider.

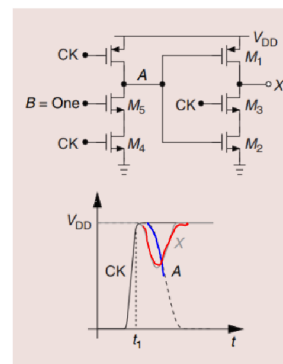


FIGURE 9: The problem of race in TSPC stages.

References

- [1] J. Yuan and C. Svensson, "High-speed CMOS circuit technique," *IEEE J. Solid-State Circuits*, vol. 24, pp. 62-70, Feb. 1989.
- [2] Y. Suzuki, K. Odagawa, and T. Abe, "Clocked CMOS calculator circuitry," *IEEE J. Solid-State Circuits*, vol. 8, pp. 462-469, Dec. 1973.
- [3] R. H. Krambeck, C. M. Lee, and H. S. Law, "High-speed compact circuits with CMOS," *IEEE J. Solid-State Circuits*, vol. 17, pp. 614-619, June 1982.
- [4] M. Shoji, *CMOS Digital Circuit Technology*. Englewood Cliffs, NJ: Prentice Hall, 1988.
- [5] N. P. Goncalves and H. J. de Man, "NORA: A racefree dynamic CMOS technique for pipelined logic structures," *IEEE J. Solid-State Circuits*, vol. 18, pp. 261-268, June 1983.

1st paper on
True single phase
clock circuit

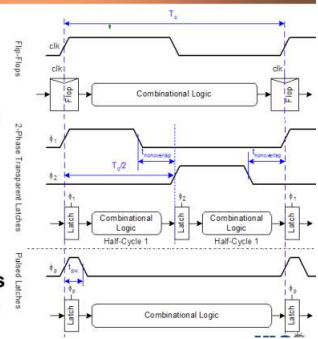
Domino Paper

NORA

Popular Clocking Disciplines with FF & Latches

- Options for timing of sequential circuits:
 - » Using Flip-Flops (1-phase single-edge-tr.)
 - » Using Transparent Latches (2-phase clocking)
 - » Using Pulsed Latches (1-phase single-edge-tr.)


Same behavior



Synchronous Clocking

Definitions:

- A (sub)circuit is said to operate synchronously iff all state transitions are restricted to occur periodically at precise moments of time that are determined by a special signal referred to as the clock.
- A clock domain is a (sub)circuit where all clock signals maintain fixed frequency and phase relationships because they are derived from a common source.



Synchronous Clocking – Design Rules

Design rule


Consistently dissociate signals into

- asynchronous reset signals (when to enter the start state)
- clock signals (when to move from one state to the next)
- information signals (what state to enter, what output to produce)


Design rule

Make the clock period long enough so that all transient effects have died out before the next active clock edge instructs registers (and other storage devices) to accept new data!

Delays - Abstraction



Timing Parameters of Combinational & Sequential Circuits



Combinational circuit with three input signals and three output signals.

t_{pd} (Propagation delay): The time required to process new input from applying a stable logic value at a data or clock input until the output has settled on its final value.

t_{cd} (Contamination delay): The time from altering the logic value at a data or clock input until the output value starts to change.

Timing Quantities – Comb. and Seq. Circuits

Combinational and sequential circuits


Propagation delay t_{pd}

- New stable input (data or clock) – output settled on final value
- Example NAND: A-to-Z, and B-to-Z
- Example latch: D-to-Q, and/or CLK-to-Q

Contamination delay (or retain delay) t_{cd}

- Altering input (data or clock) – first change of value at output
- By definition: $0 \leq t_{cd} \leq t_{pd}$

Synchronous Timing




Synchronous timing:
All registers synchronized with same CLK

Timing Constraints (Setup & Hold)


- There are two main problems that can possibly arise in synchronous logic:
 - » **Max Delay:** The data doesn't have enough time to pass from one register to the next before the next clock edge.
 - » **Min Delay:** The data path is so short that it causes a hold violation in capturing register.
- Max delay violations are a result of a slow (long) data paths, including the register's t_{su} , therefore they are often called "Setup violations".
- Min delay violations are a result of a fast (short) data path, causing the data to change before the t_{hold} of the reg has passed, therefore they are often called the "Hold violations".

Setup (Max) Constraint

- Let's see what makes up our clock cycle:
 - » After the clock rises, it takes t_{cq} for the data to propagate to point A.
 - » Then the data goes through the delay of the logic to get to point B.
 - » The data has to arrive at point B t_{su} before the next clock edge.
- In general, our timing path is a race:
 - » Between the Data Arrival, starting with the launching clock edge.
 - » And the Data Capture, one clock period later.




Setup (Max) Constraint



$$T > t_{CQ} + t_{logic} + t_{SU}$$


Hold (Min) Constraint

- Hold problems occur due to the logic changing before t_{hold} has passed.
- This is not a function of cycle time – it is relative to a single clock edge!
- Example of meeting the hold constraint:
 - » The clock rises and the data at A changes after t_{cq} .
 - » The data at B changes $t_{pd}(\text{logic})$ later.
 - » Since the data at B had to stay stable for t_{hold} after the clock (for the second register), the change at B has to be at least t_{hold} after the clock edge.



VLSI

Hold (Min) Constraint



$$t_{CQ} + t_{logic} > t_{hold}$$

Summary

- For **Setup** constraints, the clock period has to be longer than the data path delay:

» This sets our maximum frequency.

$$T > t_{CQ} + t_{logic} + t_{SU}$$

- » If we have setup failures, we can always just **slow down the clock**.


- For **Hold** constraints, the data path delay has to be longer than the hold time:

» This is **independent of clock period**.

$$t_{CQ} + t_{logic} > t_{hold}$$

- » If there is a hold failure, your chip will never work!

Fundamental Timing Conditions – Reminder



The formulas to the right apply to the data input of register U2.

Setup Condition

$$t_{clk} + t_{setup} > t_{pdff} + t_{pdc} + t_{suff} (-t_{sk})$$


Hold Condition

$$t_{hoff} \leq t_{cdff} + t_{cdc} (-t_{sk})$$

Clock Skew

Positive clock skew t_{sk} relaxes the setup condition, while having detrimental effects on the hold condition.

Clock Skew and Jitter



Skew: different physical places, different clock arrival time

Jitter: cycle to cycle variation in one physical location

VLSI

□ Clock skew

- » Spatial variation in temporally equivalent clock edges; deterministic (can control it) + random, t_{SK}

□ Clock jitter

- » Temporal variations in consecutive edges of the clock signal; purely random
- » Cycle-to-cycle (short-term) t_{JS}
- » Long term t_{JL}


Both skew and jitter affect the effective cycle time

□ Variation of the pulse width

- » Important for level-sensitive clocking (latches)

VLSI

Sources of Clock Uncertainties



Sources of clock uncertainty

Sources of skew include:


- Different buffers & wires between source and leafs

Sources of jitter include:

- Temperature gradient
- VDD variations/noise (data-dependent, varies from cycle to cycle)
- Capacitive coupling from nearby signals

VLSI

Sources of Clock Uncertainties




- Disparities of switching thresholds across clocked cells translate ramp times into skew.
- Bistables get characterized for fast clock ramps. Setup and hold times tend to grow when a cell is being clocked with slow ramps.
- Correct behavior and accurate timing are put at risk.

Avoid slow transitions on clock nets




Clock Non-Idealities – Waveform Examples




Digital System Clocking: Oklobdzija, Stojanovic, Markovic, Nedovic VLSI

Positive and Negative Skew



Positive skew
Data & clock propagate in the same direction. Relaxed setup constraints (longer effective clock cycle).



Negative skew
Opposite directions. Shorter effective clock cycle.



Setup (Max) Constraint

- The **Launch path** (still) consists of:

$$t_{cq} + t_{logic} + t_{setup}$$

- But if jitter makes the launch clock later, we need to add it to the data path delay.

$$t_{launch} = t_{CQ} + t_{logic} + t_{SU} + t_{jitter}$$

- The **Capture path** consists of:

- The clock period (T)

- Positive skew means the capture clock path is longer.

- If jitter makes the capture clock earlier, we need to subtract it.

$$t_{capture} = T + \delta - t_{jitter}$$

- Our max constraint is:

$$t_{capture} \geq t_{launch}$$

- So we get:


$$T \geq t_{CQ} + t_{logic} + t_{SU} + 2t_{jitter} - \delta$$

Jitter: always worse T



Setup (Max) Constraint

- Data has to arrive before *next* active clock edge.



$$T + \delta > t_{CQ} + t_{logic} + t_{SU} + 2t_{jitter}$$



Hold (Min) Constraint

- The **Launch path** (still) consists of:

$$t_{cq} + t_{logic}$$

- But if jitter makes the launch clock *earlier*, we need to subtract it from the data path delay.

$$t_{launch} = t_{CQ} + t_{logic} - t_{jitter}$$

- The **Capture path** consists of:

- Skew that makes the clock edge arrive at the capture register later than at the launch register.

- Actually, since it is a single clock edge, jitter should effect the capture clock the same as the launch clock.

- But as a worst case, we will add it as spacial jitter.

Clock non-idealities:
More difficult to meet hold constraint, need longer contamination delay

- Our min constraint is:

$$t_{launch} \geq t_{capture}$$


- So we get:

$$t_{CQ} + t_{logic} \geq \delta + t_{hold} + 2t_{jitter}$$




Hold (Min) Constraint

- Data has to arrive *after* same clock edge has arrived at capturing register




$$t_{CQ} + t_{logic} > t_{hold} + \delta + 2t_{jitter}$$

Clock Skew in Alpha Processor




- Terrible clock skew
- Zero skew by clock drivers
- Data path in middle has intermediate skew
- Peripherals for I/O has high skew

EV6 (Alpha 21264) Clocking



Global clock waveform
(target spec)




Next version of Alpha processor had less skew: **target >50ps**

Insert clock from PLL into center, then H-tree, and more local final clock buffers


Also, implemented more clock gating for power reduction

VLSI

EV6 Clock Skew



Clock Distribution – H-Tree



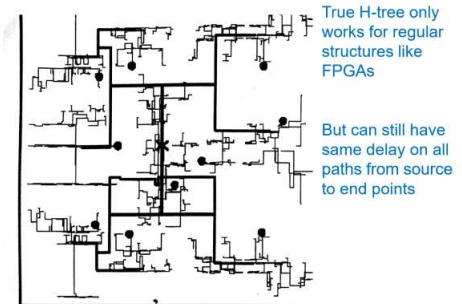
Clock tree minimizing skew

Clock insertion delay to every FF identical (ideally)

Sometimes: work with “**useful (intentional) skew**”, e.g. if many FFs have hold violation, to avoid insertion of a large amount of buffers

Equal wire length/number of buffers to get to every location

More realistic H-tree



True H-tree only works for regular structures like FPGAs


But can still have same delay on all paths from source to end points

[Restle98]

Dealing with Skew and Jitter


- Balance clock paths using regular distribution network, such as H-tree
- Use local clock GRIDS (increased cap and power)
- Route data and clock in opposite directions to improve hold at the expense of setup.
- Shield clock wires (minimize capacitive coupling)
- Use dummy metal density fillers for regular wires
- Use decoupling capacitors (for stable VDD)
- Time borrowing (or cycle stealing): long path borrows time from subsequent short path, accomplished using latches
- "Useful skew" to avoid expensive buffering for hold fix

Additional Timing Parameters of Sequential Circuits



Rising-edge-triggered D-type flip-flop

The value at the input D becomes available at the output Q at the rising edge of the clock



t_{su} (**Set-up time**): The lapse of time before the active clock edge during which an input is required to assume a fixed logic value of either 0 or 1 at the input of a clocked circuit.

t_{ho} (**Hold time**): The lapse of time after the active clock edge during which data are required to remain logically unchanged at the input of a clocked circuit.

Data call window: $t_{su} + t_{ho}$

Timing Quantities (Cont'd) – Seq. Circuits Only

Sequential circuits only

Setup time t_{su}

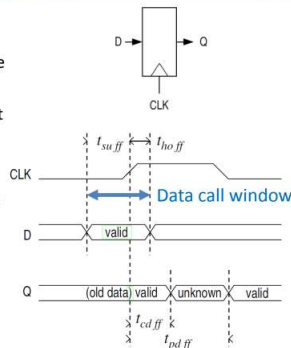
- Time lapse before active clock edge during which input of bistable element (flip-flop, latch, SRAM, ROM, ...) needs to settle for correct sampling

Hold time t_{ho}

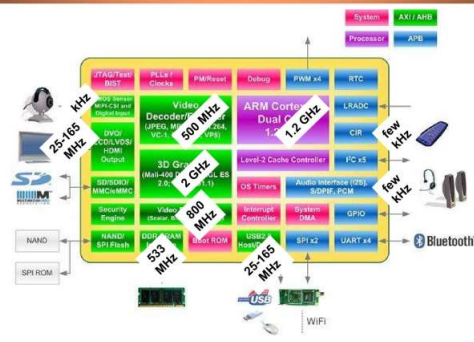
- Time lapse after active clock edge during which input data of bistable element needs to remain unchanged

Data call window

- Sum of t_{su} and t_{ho}
- Typically centered around active (positive) clock edge




A Modern SoC has Many Clock Domains




Designs with Multiple Clocks

- **Clock domain:** (Sub)circuits in which all clock signals maintain fixed frequency and phase relationship
 - » A single clock domain may use multiple divided clocks
- **Clock box:** top-level entity that generates all clock signals




Clock Dividers

- Derive a slow clock from a fast clock by dividing it by an integer number (counter)
- Remember: all clocks (including divided clocks) need to be free of hazards and glitches
 - » Must assume that any logic can produce glitches
 - » Generated (divided) clocks MUST come directly from a FF output



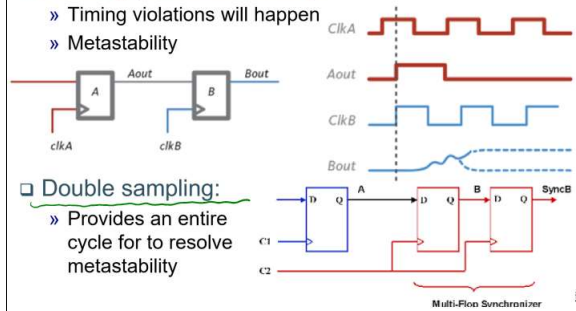
Clock Dividers in the Clock Tree

- Clock and derived clock must be phase aligned at the clock roots
- Clock divider has a delay (FF propagation delay)
 - » Need to consider clock divider during clock-tree generation
- Modern clock tree synthesis tools can trace through clock dividers




Crossing Clock Domains

- Arbitrary phase relationship
 - » Timing violations will happen
 - » Metastability



Clock Gating


- Flip-Flops with ENABLE: prevent FlipFlop(s) from capturing new data
 - » Functional: preserve content over multiple cycles
 - » Power savings: avoid activity in a block that is inactive



Clock Gating

- Advantages:
 - » Saves a multiplexer per FlipFlop with Enable (Area, Delay, and Power advantage)
 - » Avoids activity on the clock net leading to the FlipFlop
 - » Avoids activity on the FlipFlops clock pin reducing internal power consumption
- Disadvantages:
 - » Need for additional logic to suppress the clock while FlipFlop is disabled (area and power penalty)
 - » Need to ensure that the clock signal is free of glitches


- A safe strategy to realize clock gating



- » Latch on the Enable signal shields glitches during the sensitive period of the clock
- » Implemented with individual cells: some timing constraints need to be observed
- » Often realized as single dedicated library cell


Clock Gating

- Coarse grained: disable entire blocks of the design
 - » Inserted manually
- Fine grained: enable for small groups of FlipFlops
 - » Automatic insertion



Improving IO Timing

- Chip provides a clock output that is aligned with the clock at the leaves of the FlipFlops
 - » Output of the clock output pad is declared as clock leaf
 - » Delayed clock is used as a reference for rest of system



Improving IO Timing

- Delay locked loop (DLL)
 - » Generates a phase shifted clock such that the reference input is phase aligned with the input clock
 - » Clock reference is taken from the leaves of the clock tree
- Internal clock at the leaves is aligned with clock input

