METHODOLOGY TO LOWER SUPPLY VOLTAGE OF STANDARD CELL LIBRARIES

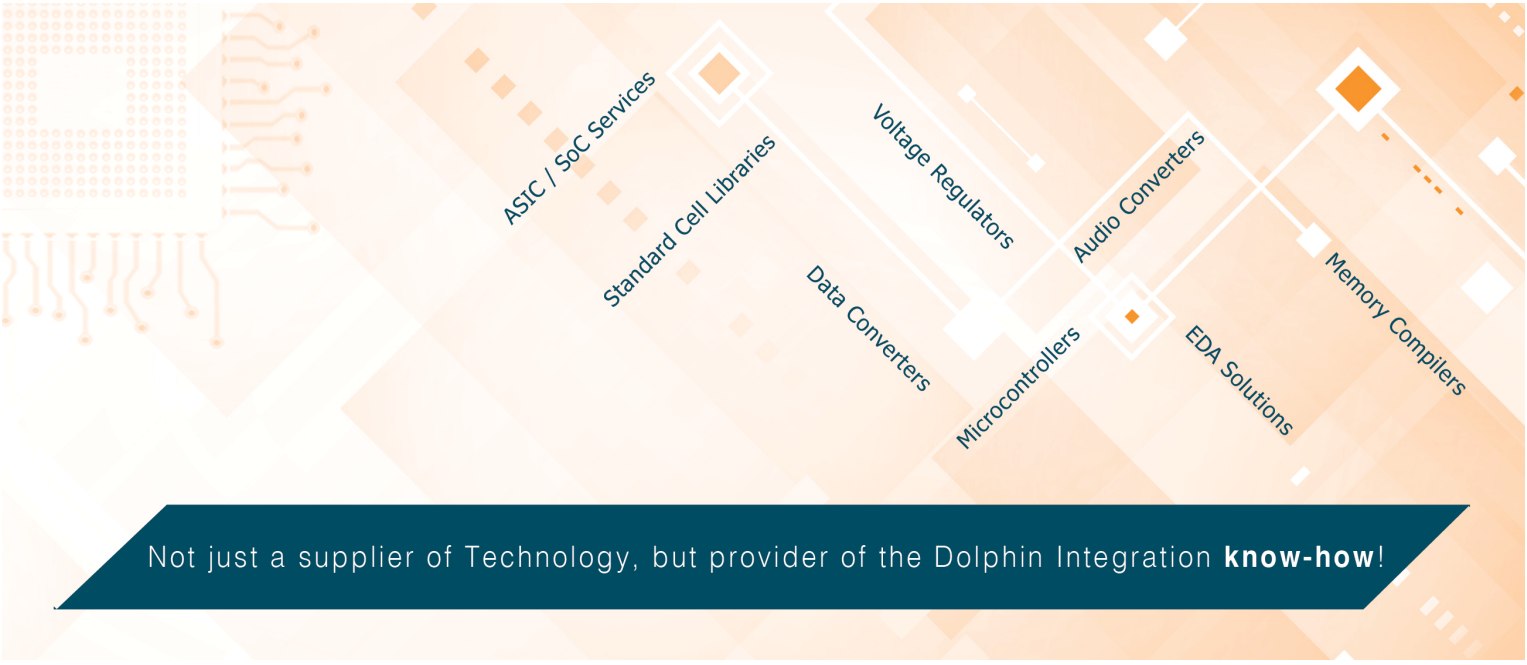

Introduction

Standard cells libraries are usually designed to operate at a specific value of supply voltage referred to as "nominal voltage".

This article details the performance trade-offs in terms of power consumption and speed when decreasing power supply voltage, as well as a methodology to determine the lowest value to use. Decreasing the supply voltage will reduce both the power consumption to some extent and the speed significantly. As the design of the cells is left unchanged, the density will remain unchanged.

In the first section, we will focus on the benefits of lowering the supply voltage on the power consumption, and then we will observe the impact of low supply voltage on the speed of the cells. After noting the importance of OCV margins, we will finally describe a methodology to obtain the best trade-off for supply voltage in terms of speed and power.

Power consumption

The main benefit of decreasing the supply voltage is to lower the power consumption of the circuit. Indeed, the power consumption is proportional to the square of the voltage. The following table shows dynamic power consumption results of a ring oscillator for different supply voltages.

Supply voltage	0.55 V	0.75 V	0.9 V
Dynamic consumption	1.07 fW	2.06 fW	3.72 fW

Table 1 - Dynamic power consumption evolution with respect to supply voltage

The leakage power also decreases with the supply voltage, until a point where it starts increasing again because of the increase of cycle time duration. The graph in Figure 1 shows the evolution of the leakage energy per cycle for low values of supply voltage. Therefore, in terms of power consumption, there is no benefit in decreasing the supply voltage

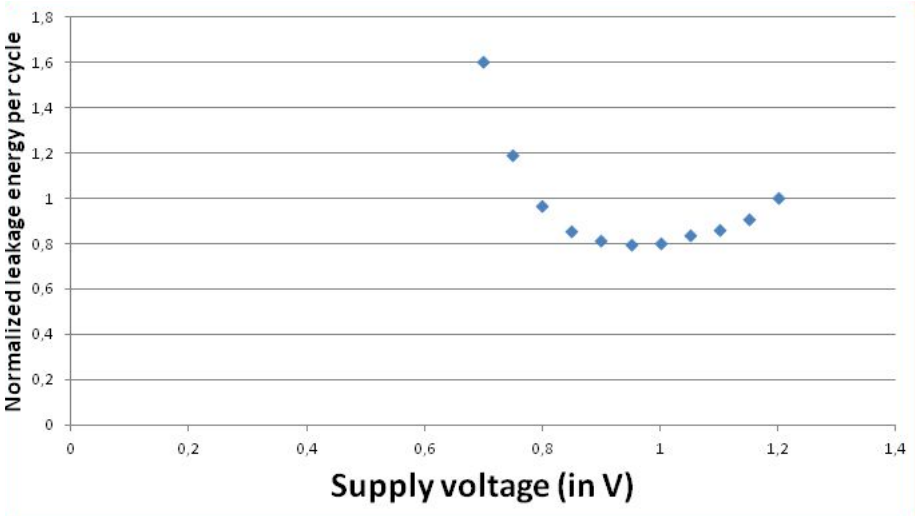

Figure 1 - Leakage energy per cycle increase at low values of supply voltage

Speed

As mentioned above, decreasing the supply voltage has a huge impact on speed. Indeed, even if the voltage swing is smaller (from 0 to the supply voltage), the signal driving the gate V_{gs} is weaker and the current in the transistor, proportional to $V_{gs} – V_{th}$, is smaller as well. V_{th} plays an important role here: high-V_{th} MOS libraries will have smaller leakage values but lower speed than low-V_{th} MOS libraries and therefore, will be less susceptible to work at very low voltage.

Evolution of speed

Decreasing the power supply voltage results in a direct increase of the propagation time of standard cells (the transition of each transistor is slower). This becomes an issue when reaching very low voltages, as some element of a system-on-chip need a minimum frequency to work properly (e.g. a 32 kHz always-on block).

The maximum frequency depends on constraint timings (setup and hold times) and propagation times. As can be observed on Figure 2 and Figure 3, the propagation time of a standard cell increases exponentially as the voltage decreases. The hard limit defined by design constraints is often reached before the transistors stop working.

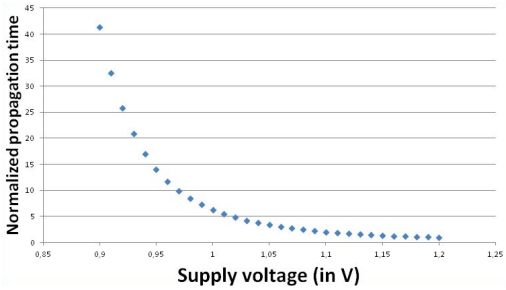

Figure 2 - Evolution of the propagation time at low supply voltage (at 55 nm)

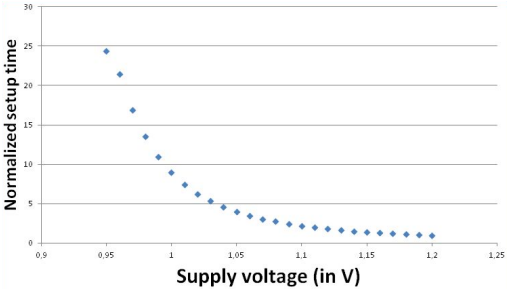

Figure 3 - Evolution of the Setup time at low supply voltage (at 55 nm)

Timing variation at low voltage

A small variation on the supply voltage will result in a large timing variation. This phenomenon is problematic as around 5 % of margin on the supply voltage is used to take IR-Drop into account. As a consequence, at such a low voltage, a variation of 5 % of VDD would induce timing variation far beyond the characterization range. Careful assessment of the dynamic IR-Drop becomes a hard requirement at low voltage.

On Chip Variation (OCV) impact

The values of propagation and setup times given in the liberty files provide the theoretical speed at which a library can be used. However, the maximum frequency must be reached in worst-case scenarios, with worst-case variations. OCV margins are used to model timing variations induced by temperature, IR-Drop, and process local variations across the chip. Even with advanced OCV techniques, these margins increase significantly when the supply voltage is lowered. This is an important factor when targeting a specific frequency as the OCV margins can account for half of this frequency.

Low voltage determination methodology

As described previously, the notion of lowest voltage possible is not absolute. It depends on criteria that have to be chosen wisely in order to ensure a correct behavior of the library.

Two different values of low voltage are significant when using a library for low power purposes:

- The lowest voltage at which the library can work properly with acceptable speed performances and reduced power consumption,
- The lowest voltage at which the output of the cells is maintained when there is no activity at the input (this scenario happens when a power island is switched to retention mode where its main voltage supply is lowered to reduce the overall leakage consumption).

Criteria for speed and power

As the main interest in decreasing the supply voltage is to lower power consumption, the criteria depend on what kind of power the focus is put. If the main interest is the leakage power reduction, the criterion depends on the duration of the cycle time. For low values of supply voltages, the increase of the cycle time is more important than the reduction of leakage power and the overall energy increases. If the main goal is the dynamic power, the supply voltage can still be decreased even if the leakage consumption increases. At a certain point, the leakage power becomes more important than the dynamic power. The supply voltage should not be lowered below that point.

As discussed before, the speed of the circuit is highly impacted by the decreasing of the supply voltage. Most of the time, speed is the decisive factor in the determination of the low voltage. Two criteria are important:

- In the case where a specific frequency must be reached, the speed of the cells must ensure that the library can be used at this frequency,
- The variability of the timing characteristics of the cells (propagation and constraint timings) must not exceed a certain threshold.

In the case where a specific frequency must be reached with the library, it is important to take the OCV margins into account. A maximum operating frequency can be obtained

through a place and route trial at nominal voltage. With basic ring oscillator spice simulations, ratios can be used to assess the value of the frequency at the lowest supply voltage.

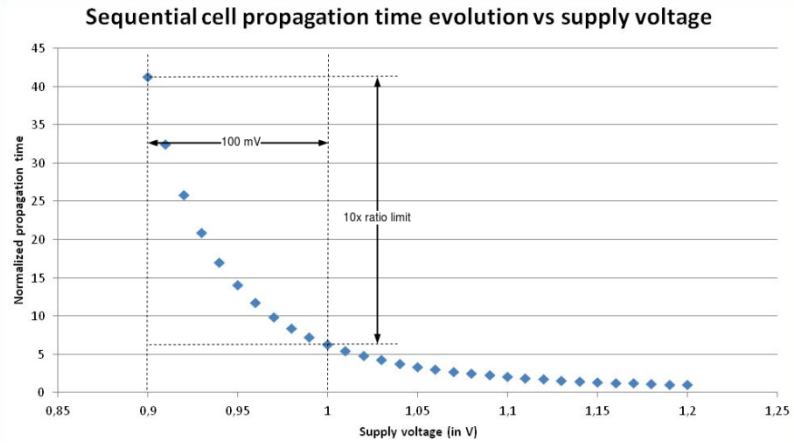

Figure 4 - Criteria on propagation time variation

The other limit to consider is the variability of the propagation timings. This variation should be limited in order to reduce the sensitivity of the cells to a small change of supply voltage so that the characterization models (liberty files) are still valid. As shown in Figure 4, a simple criterion is used: a variation of 100 mV on the supply voltage must not result in an increase of the propagation time greater than a factor of 10.

Estimating the low voltage from a ring oscillator circuit

In order to simulate the propagation time, we propose to use a simple circuit containing a sequential element and a ring oscillator of inverters to model the transition of the signals.

In the schematic below (Figure 5), the sequential cell used is a flip-flop. Its output is connected to its D input through a chain of inverters in order to create a ring oscillator. The clock input goes through a couple of clock buffers before reaching the flip-flop. Finally, an AND gate is used to initialize the input of the flip-flop at the beginning of the simulation. Capacitance loads are added at the output of each cell to emulate a real circuit. A SPICE simulation is run on this testbench to extract the propagation times (clock to output) and power consumption values of the flip-flop and of the inverter chain for a given power supply voltage target (V_{targ}). These measurements are performed at V_{targ} + 100 mV and the timing degradation between V_{targ} + 100 mV and V_{targ} is computed. From these values, the supply voltage is either validated or rejected. By dichotomy, the lowest supply voltage complying with the selected criteria is determined.

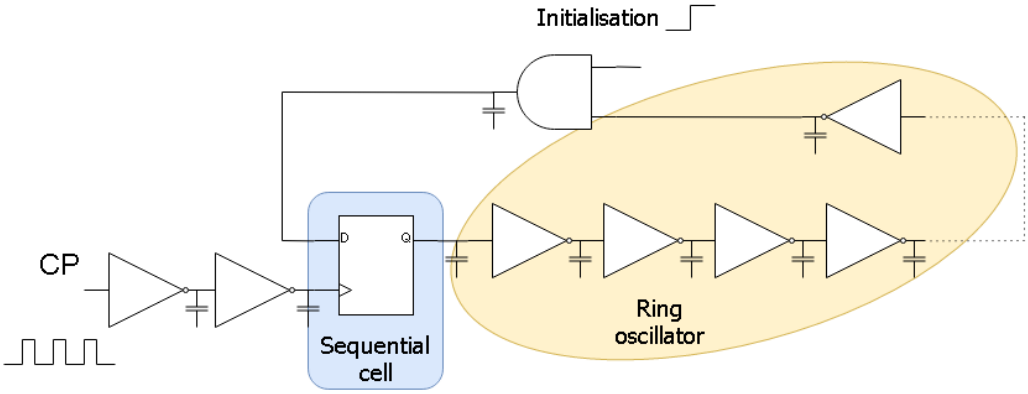

Figure 5 - Ring oscillator circuit for flip-flop low voltage determination

Estimating the retention supply voltage

Another interesting voltage characteristic that can be measured is the retention supply voltage of the memory element of the sequential cell, which can be the only part still supplied in retention mode.

In order to get this value, two options are considered. First, a static retention supply voltage can be obtained with simple DC simulations on all the cells. If a valid operating point is found, the voltage is validated for retention purposes. Thus, the limiting cells are identified. Using this method is optimistic as it does not take into account failures occurring during voltage changes while entering or exiting the retention state. To increase the reliability of this value, the limiting cells can be tested with a dynamic stimulus using the schematic of Figure 5 to check the state of the cell output when switching the supply voltage from its nominal value to the tested retention voltage.

Results at 28 nm

In a high density library at 28 nm using pulsed latches, we obtain a value of low supply voltage equal to 0.58 V where the nominal voltage is 0.9 V, using the speed criteria described before. The dynamic retention supply voltage is measured at 0.22 V.

Conclusion

This article describes benefits and issues occurring when decreasing the supply voltage of standard cell libraries. Using an existing standard cell library, we demonstrated that, with a lower supply voltage, a smaller value of power consumption could be reached. It is important to be aware of the related difficulties: lower speed performance, variability of timings and OCV margins increase. We proposed methods based on simulations of the standard cell library to determine the best performance trade-offs as well as the limit values of dynamic and static retention supply voltage.

Eventually, when the constraint on power consumption and speed is not reachable with an existing library, the design must be modified (limitation of the number of stacked MOS, variation of MOS length). Libraries can be optimized to reach lower power consumption with reasonable speed performances, but with an impact on area.

About the autor

Bastien Arricca joined Dolphin Integration in 2012 as standard cells designer. He is now involved in standard cells design for low power architectures. Bastien holds Master of Science in micro and nanotechnologies for integrated systems from EPFL (Switzerland), Grenoble-INP PHELMA (France) and Politecnico di Torino (Italy).

Dolphin Integration IPs:

- SESAME BIV : https://www.chipestimate.com/log.php?from=%2Fip.php%3F9%2Btracks%2BExtreme%2BLow%2BVoltage%2Bstandard%2Bcell%2Blibrary%2B%26id%3D39029%26partner%3DDolphin%2520Integration&logerr=1

- SESAME eLV : https://www.chipestimate.com/log.php?from=%2Fip.php%3FStandard%2Bcell%2Blibrary%2Bdesigned%2Bwith%2B3.3%2BV%2BIO%2B%26id%3D35803%26partner%3DDolphin%2520Integration&logerr=1