Adv. Radio Sci., 4, 197–205, 2006
www.adv-radio-sci.net/4/197/2006/

Advances in Radio Science

Timing violations due to V_{DD}/V_{SS} bounce

M. Eireiner[1], S. Henzler[1,2], J. Berthold[2], C. Pacha[2], G. Georgakos[2], and D. Schmitt-Landsiedel[1]

[1]Institute for Technical Electronics, Technical University Munich, Germany
[2]Infineon Technologies, Munich, Germany

Abstract. The effect of power supply noise in on-chip power grids and its implications on the path delay in digital circuits is examined. The simulation results show that IR-Drop and the resulting path delay are strongly affected by the layout of the circuit. Power grid design measures to reduce IR-Drop, as well as their area and performance implications are discussed.

1 Introduction

With the ongoing scaling of CMOS technologies, process and environmental variations gain significance. Since supply voltages have been scaled with the last technology nodes, power supply integrity has become a first class design issue (Benoit et al., 1998; Lin and Chang, 2001). Variations of supply voltage are either caused by the inductance of the power line times the time-deviation of the current, so called dI/dt noise, or by the voltage drop over the finite resistance of the power grid, caused by the current flowing through it.

In this the paper the implications of power noise on the path delay will be examined. In Sect. 2 an overview of the different types of variations in modern technologies is given, and the challenges imposed by environmental variations are described. Following Sect. 3 on IR-Drop, the modeling of the power grid, as well as the simulation setups are explained in Sect. 4. In Sect. 5 the propagation of the voltage drop over the power grid is explained at two different simulation setups. Simulation results for implications of the power supply noise on the path delay are shown in Sect. 6. In Sect. 7 design measures for reducing power supply noise, as well as their implications on the performance and the area overhead are shown. Finally, conclusions are drawn in Sect. 8.

Correspondence to: M. Eireiner (eireiner@tum.de)

2 Variations

In deep submicron technologies, variations, both process and environmental, gain significance. In the following, both will be discussed briefly.

2.1 Process vriations

Process variations in active and passive devices, like changes in doping concentration, oxide thickness, wire resistance and capacitance, as well as length and width of transistors, influence their performance. Process variations which are correlated between transistors on a die, like lot-to-lot, wafer-to-wafer and chip-to-chip variations, can be characterized by full-speed tests and speed monitors. After the characterization of the circuits, appropriate counter measures, such as supply voltage binning or body bias adjustment can be carried out to increase the parametric yield (Tschanz et al., 2002).

Random variations, which are not correlated, are very hard to detect and characterize. Up to now, only few approaches have been published to address this problem (Ernst et al., 2003).

Since in current technologies global or correlated variations are dominant, speed monitors and binning of frequency and supply voltage is being used successfully. Since random variations become more and more significant with ongoing technology scaling, more research is to be expected on that topic.

2.2 Environmental variations

With the scaling of technology nodes, not only process variations, but also environmental variations, such as cross talk (X-Talk), changes in temperature and power supply distortions, become more and more significant. This is due to the continuous decrease of supply voltage, accompanied with

Published by Copernicus GmbH on behalf of the URSI Landesausschuss in der Bundesrepublik Deutschland e.V.

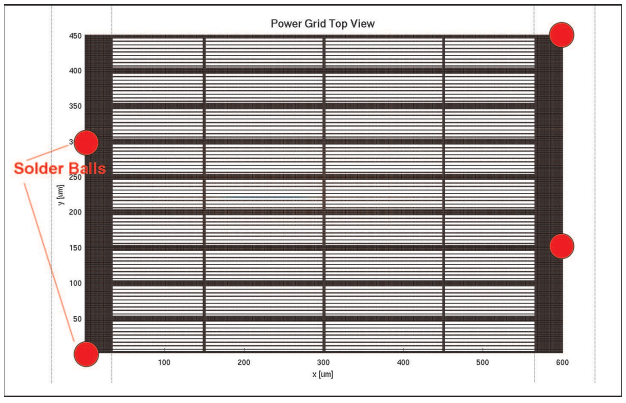

Fig. 1. Top view of an axially symmetric cut out of a four stage power grid

just a moderate scaling of threshold voltages and the resulting reduced overdrive voltage.

Actually, environmental variations are deterministic, but since the complexity is too high and most of the variations are highly influenced by the layout of the chip, these variations are modeled, if at all, statistically. Another big challenge of environmental variations is the testing or characterizing of them. This is because switching pattern and the activity of the circuit determine environmental changes. Scan tests, which are used for characterizing process variations, are not suitable for characterizing these variations, since in scan tests only few paths are triggered during one clock cycle and the switching patterns are therefore not representative for real operation. Worst case analysis and tests are one possibility to cope with these variations. Because these analyses and tests are very pessimistic the overall yield decreases.

2.3 Power supply integrity

Since the paper addresses the problem of power supply noise, its origins will be explained here. Power supply noise is caused by two different mechanisms.

The first is called dI/dt noise. Transient changes of current cause a voltage drop over the inductance of the power grid, as described by Eq. 1, where ΔV_{dI} is the change in supply voltage due to dI/dt, L the inductance of the power grid and I the current through the power grid.

$$\Delta V_{dI} = L \cdot \frac{dI}{dt} \tag{1}$$

This voltage drop can either increase or decrease the effective supply voltage for the gates. Therefore, especially at scenarios of block activation, clock gating or frequency change, where large changes of current occur, becomes this effect present. For low power applications, on-chip inductance still can be neglected compared to the inductance of the bond wires (Piguet, 2004). Since in this paper only on-chip effects of power integrity are looked at, inductance is omitted in our analysis and is not modeled in the simulation setup.

The second effect of power supply integrity is called IR-Drop. As Ohms law states, Eq. 2, the current through the power grid causes a voltage drop over the finite resistance of it and decreases therefore the effective supply voltage seen by the gates.

$$\Delta V_{IR} = R \cdot I \tag{2}$$

For applications with low V_{DD} and high power density, high current values intensify this problem.

3 IR-Drop

3.1 Resistance of power grid

As Eq. 2 describes, the voltage drop caused by IR-Drop is proportional to the resistance of the power grid which the current must pass through. This resistance is dependent on various parameters. The first is the specific resistance of the metallization material used. For reducing the specific resistance of the metallization, aluminum has been replaced by copper in modern technologies. Another factor is the thickness of the metal layers. Together with the specific resistance, the thickness of the metal layer determines the sheet resistance R_S of a metal layer. These factors are technology dependent and can therefore not be altered by the chip designer.

The designer of a circuit has two primary possibilities to influence the resistance of the power grid, since the resistance of a wire R_{wire} is given by Eq. 3, where L_{wire} and W_{wire} are the width and length of the wire.

$$R_{wire} = R_S \cdot \frac{L_{wire}}{W_{wire}} \tag{3}$$

The length of a power line to a specific circuit is usually given by the floorplan. Therefore, the only parameter of Eq. 3, which can be directly altered by the designer is the width of the power grid wires. The topology of the power grid is the second factor through which the designer can influence the resistance of the power grid. Usually a power grid consists of three to five metal layers, each orthogonal to its neighboring metal layers. At their overlapping points, the different metal layers are connected through multiple vias in parallel, so called via stacks.

Figure 1 shows a top view of an axially symmetric cut out of a four stage power grid, as it is used in our analysis. Going from bottom to top, the lowest metal layer is contacted by the second lowest metal layer in a fixed pitch. In the same manner, the second lowest metal layer is connected to the second highest metal layer in a fixed pitch and so on. The higher the metal layer, the wider the wires are. Besides the pitches and the widths of the metal lines, the number of parallel vias per overlap point determines the resistance of the

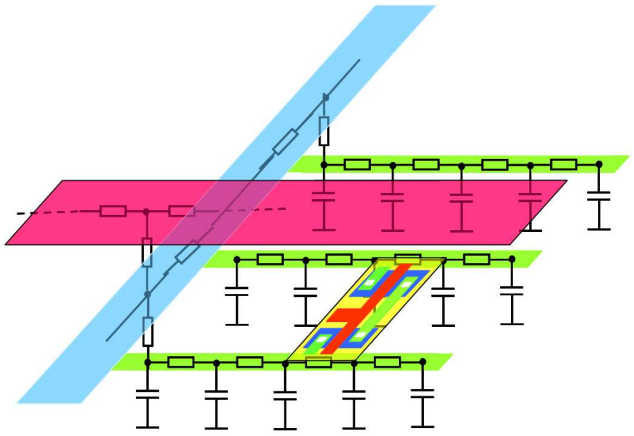

Fig. 2. Schematic of the power grid model

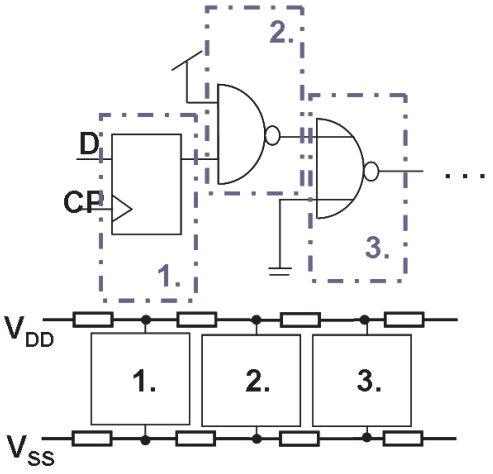

Fig. 3. Basic structure of the critical path and placing of its gates in the power grid

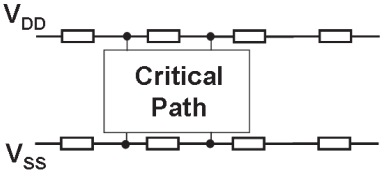

Fig. 4. Critical path placed in the power rail without additional load

power grid. Solder balls, through which the highest metal layer of the on-chip power grid is connected to the external power supply, are depicted as red circles in Fig. 1.

3.2 Outlook on future technologies

As the power density will continue to increase and the supply voltage will continue to decrease, even if slower than previously expected, on-chip currents have to rise (http://www.itrs.net, 2003). This said, it is clear that power supply integrity is a problem that will become more and more challenging with ongoing technology scaling. The most likely solution of this problem is that more metallization area and more pins will be dedicated to power distribution to reduce IR-Drop and dI/dt effects. Even today, about half of the I/O pins of a high performance microprocessor are already dedicated to power distribution (Rusu et al., 2006).

4 Modeling of the power grid and simulation setup

In our analysis we use a linear RC-model to model the four stage power grid. A schematic of the power grid model is depicted in Fig. 2. For simplicity reasons only the lower three stages are shown.

The values for sheet resistances, sheet capacitances, widths, pitches, and via counts are derived from a typical digital CMOS ASIC in a 90 nm low power technology. Widths of the metal lines are in the range of 250 nm – 600 nm for the lowest to 50 μm – 150 μm for the highest metal layer. Pitches range from 3 μm – 15 μm to 300 μm – 1500 μm for the lowest and the highest layer. The resistance of a via is between $0.25\Omega - 20\Omega$ and $0.02\Omega - 0.2\Omega$, the number of vias per crossing point (via count) ranges from 5 – 25 to 50 – 250 from the lowest to the highest metal layer. The sheet resistance for the different layers is between $0.02\Omega/\square$ and $0.125\Omega/\square$ from the highest to the lowest metal layer.

Since setup violations due to power noise first occur in the critical paths of a circuit, a critical path replica of an ARM9 core, with a depth of 22 stages, was taken as test vehicle in our analysis. The layout of the path has a length of 86 μm. The basic structure of the critical path and how the gates are placed in the power grid is depicted in Fig. 3. The orientation of the layout is from left to the right and all gates of the critical path are located on the same power grid. In the simulation, all gates are connected once to the power grid, so that the different gates are separated by the distance of the sum of their half lengths. Apart from the critical path, also small and medium paths are used in the simulation to model a heavier loaded power grid. The schematic of the medium and small paths, as well as the placing in the grid, is analog to the critical one shown in Fig. 3. All simulations are done in a low-power 90 nm technology, with $V_{DD} = 1.2V$ and Temp = 25°C.

As reference for the following simulations, the critical path is connected directly to a constant voltage source, and the path delay, $CP - Q_{path}$-delay, is simulated. In the following a power rail refers to a power line on the lowest metal layer. For a better understanding, we declare that a power rail ends with a low ohmic connection through a via stack to the upper metal layer.

For the first simulation setup the critical path is placed in the power grid and no additional paths are added, as shown in Fig. 4. The resulting IR-Drop at all nodes of the grid as

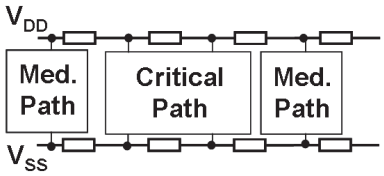

Fig. 5. Critical paths with additional medium paths in the same rail for increased load on the power rail

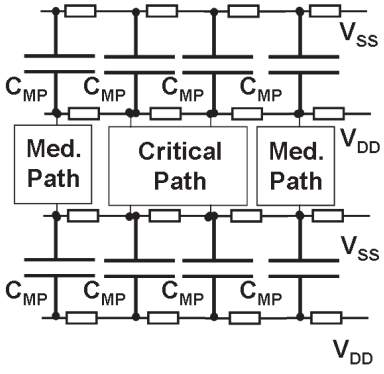

Fig. 6. Critical path with active medium path in the same and inactive paths as buffer capacitances in the shared power rails

well as the resulting $CP - Q_{path}$-delay is simulated. This results in a maximum supply voltage drop of 8.2 mV and an increase of 1% in $CP - Q_{path}$-delay.

To increase the load on the power grid, additional medium paths are placed in the same power rail, as depicted in Fig. 5. The additional currents caused by the medium paths add up to a maximum IR-Drop of 13.2 mV, which results in a $CP - Q_{path}$-delay worsening of 1.3 %.

In real circuits also inactive paths exist, also very often power rails are shared between to standard cell rows. Therefore, in the next simulation setup additional inactive medium paths are placed in the shared power rails to simulate the effect of the junction capacitances on the IR-Drop. As shown in Fig. 6, the inactive paths act as a buffering capacitance C_{MP} for the power rail of the critical path. The resulting maximum voltage drop is reduced to 10.3 mV and the resulting $CP - Q_{path}$-delay increase is only 1.1 %.

To further increase the load on the power grid, active medium paths are placed in the next simulation in the shared rail, as depicted in Fig. 7. For this simulation setup, a $CP - Q_{path}$-delay increase of 2.2 % and a maximum IR-Drop of 28.2 mV is simulated.

In a next step, the horizontal neighboring rails are loaded with medium paths, as shown in Fig. 8. In contrast to the figures shown so far, the power rails of the second metal layer and the corresponding via stacks are shown in this figure. The increased load on the power rail results in a maximum voltage drop of 27.2 mV and a $CP - Q_{path}$-delay increase of 2.4 %.

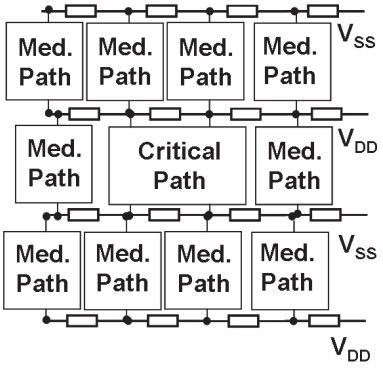

Fig. 7. Critical path with active medium paths in same and the shared power rails

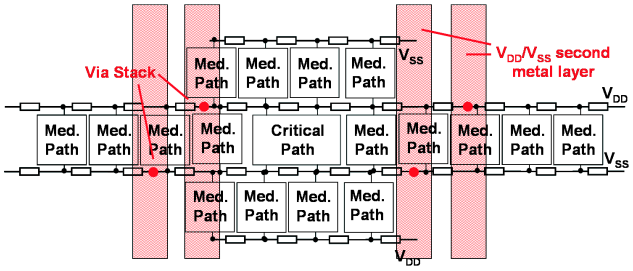

Fig. 8. Critical path with active medium path in shared and neighboring rails

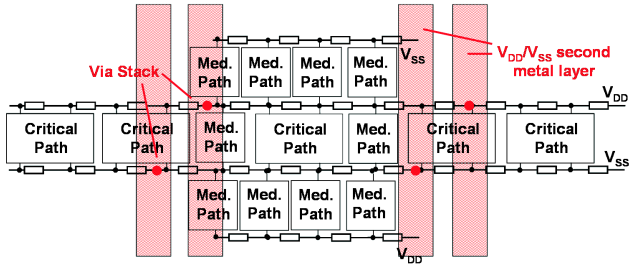

Fig. 9. Critical path with medium paths in the shared rail and critical paths in the horizontal neighboring rails

In the next simulation the medium paths in the horizontal neighboring rails are replaced by critical paths, see Fig. 9. The maximum IR-Drop is simulated to be 30.2 mV, the path delay worsening does not increase any further and stays at 2.4 %, as in the case for medium paths in the neighboring rails.

Since in real circuits it happens that a lot of critical paths are placed next to each other, e.g. in data paths, this is modeled in the next simulation setup. Therefore, critical paths are not only placed in the next 10 neighboring rails, but also in their shared power rails, as illustrated in Fig. 10. A total number of 60 critical paths is used in this simulation. The resulting $CP - Q_{path}$-delay increase is 4.1 %, the maximum voltage drop is 49.8 mV.

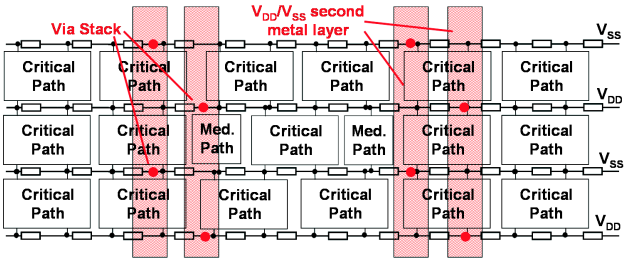

Fig. 10. Critical path surrounded with medium and critical paths in the neighboring and shared rails

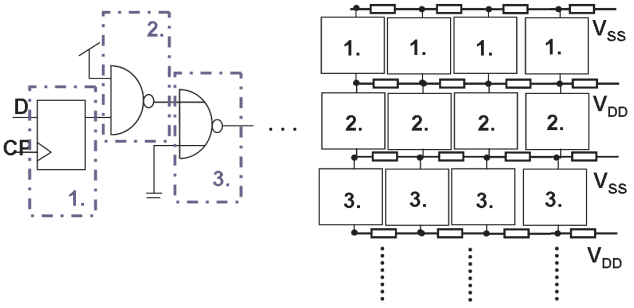

Fig. 11. Change of orientation of the critical path, with other critical paths in parallel

In the next simulation setup, the orientation of the critical path in the power grid is changed. Now the gates of the path are not all placed in one power rail, as shown in Fig. 3, but distributed over more rails. Assuming parallel critical paths, now all flipflops of the critical paths are in the first power rail, in the next power rail all first logic gates are placed and so on. This is illustrated in Fig. 11. In our simulation we place 16 critical paths next to each other, about one quarter of the critical paths we used in the last simulation. The change of orientation causes that now all gates in a rail switch at the same time, this in turn increases the voltage drop and path delay. Since the layouts of the gates have different lengths, the longest cell, flipflop, is taken as reference and the logic gates are grouped to best fit the length of the flipflop, to best fill the available area between power rails. This setup results in a maximum IR-Drop of 60.8 mV, and a delay increase of 5.8% is observed.

In the last simulation the number of parallel critical paths, with changed orientation, is increase to 48, to match better the number of critical paths used in the last simulation with regular orientation. This simulations results in a increase in path delay of 11.3% and a maximum supply voltage drop of 93.6 mV.

For easier understanding, all abbreviations which will be used later are shortly explained and the respective figures are given in Table 1.

Table 1. Descriptions of used abbreviations for simulation setups

Abbreviation	Description
GridOnly	Critical path, without additional loadand see Fig. 4
RailActive	Critical path and active medium paths in the railand see Fig. 5
RA-VDDBuff	like RailActive, with buffered shared railsand see Fig. 6
RA-VDDActive	like RailActive, with loaded shared railsand see Fig. 7
RA-VDDA-LRnoiseMP	like RA-VDDActive, horizontal rails loaded with med. pathsand see Fig. 8
RA-VDDA-LRnoiseCP	like RA-VDDActive, horizontal rails loaded with critical pathsand see Fig. 9
CPallover	All rails are loaded with critical paths see Fig. 10
OriChange16bit	Change of orientation, 16 critical paths in paralleland see Fig. 11
OriChange48bit	Change of orientation, 48 critical paths in paralleland see Fig. 11

5 Propagation of V_{DD} bounce

5.1 Reduction of complexity

Looking at our simulation setup, five independent variables can be identified. These are the two power supply rails voltages V_{DD} and V_{SS}, the coordinates, x and y, of the power grid nodes and the time.

In our analysis, the power grids for V_{DD} and V_{SS} are modeled identically. Therefore, IR-Drops will be corresponding, even if the two power rails are stressed at different transitions of a gate, the V_{DD} rail at a low-high, the V_{SS} rail at a high-low transition. Hence, only the V_{DD} rail is observed in our analysis. It should be noted that equal modeling of the two rails is only correct for a triple well process, in which the bulk is not connected to the V_{SS} rail. In a twin well process, there exists an additional high ohmic path parallel to the V_{SS} rail through the bulk. However, this effect is assumed to be higher order and therefore is neglected in this work.

In general, no implications can be made on the path delay just by knowing the maximum occurring IR-Drop (Henzler et al., 2005). However, within one rail the shape of the IR-Drop is quite similar with respect to different locations in the rail, as shown in Fig. 12. The transients of two nodes within one rail are shown. The first is placed directly under a via stack, $x = 150\,\mu$m, the other transient is taken at the node at which the highest IR-Drop, worst case (WC), occurs. One sees, that both shapes are similar and only differ in amplitude. Therefore, the maximum supply voltage drop can be used as metric for comparison between different nodes of the power grid, within one simulation setup. Together with the change in path delay, these are the figures of merit used in our work.

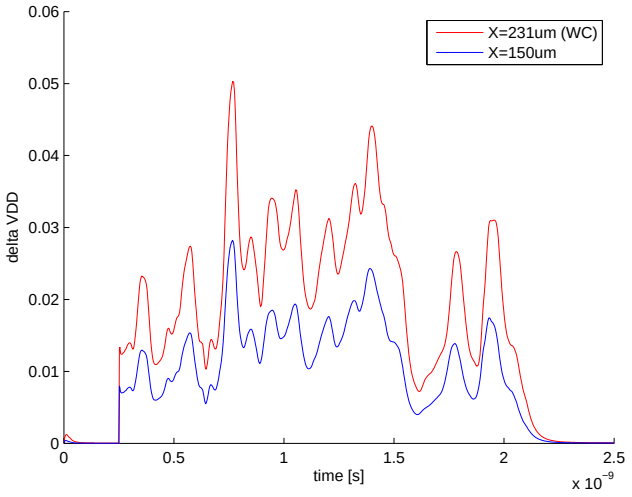

Fig. 12. Transient of IR-Drop, ΔV_{DD}, for two nodes within one power rail

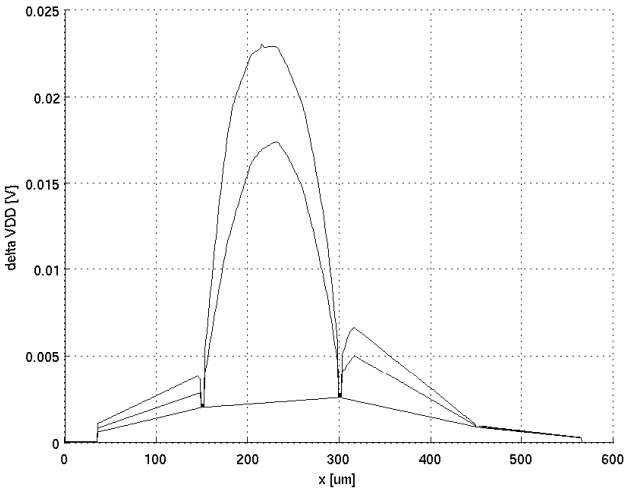

Fig. 13. Propagation of the peak voltage drop in dependence on x-coordinate for the rail loaded by the critical path and its neighboring rails, for the RA-VDDActive simulation

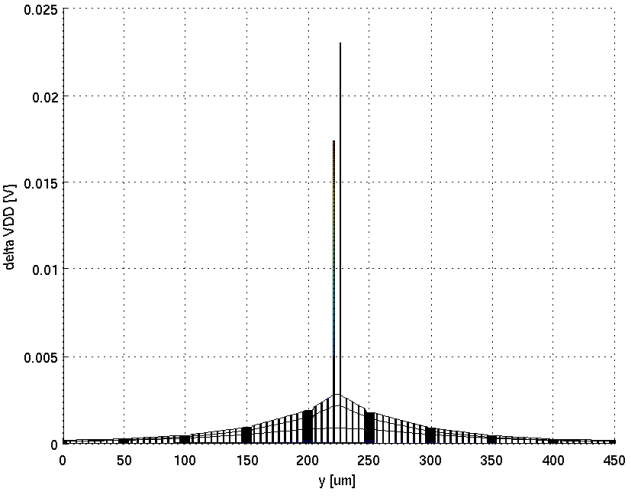

Fig. 14. Propagation of the peak voltage drop in dependence on y-coordinate for the RA-VDDActive simulation

5.2 Propagation of peak

In the following, a rail is taken which is terminated by via stacks at $x = 150\,\mu$m and $x = 300\,\mu$m. The y coordinate is $y = 225\,\mu$m. The coordinates correspond to those in Fig. 1.

Figure 13 shows the propagation of the peak in x direction for the RA-VDDActive simulation. The start position for the critical path is $x = 170\,\mu$m. The signal propagation in the critical path is from left to the right, i.e. from lower to higher x-values. For better visualization only three rails are shown, the rail used by the critical path and the two neighboring ones. The highest voltage bounce is at the rail which is loaded by the critical path and the medium paths in the same and in the shared rail. The second highest voltage peak is from the power rail which shares the V_{SS} and is loaded with medium paths. The third rail shown is not loaded at all and the resulting voltage bounce is imposed from the loaded rails over the higher metal layers and via stacks.

In Fig. 13, clearly the fast fading away of the bounce towards the low ohmic via stacks is shown. The asymmetry in the loading of the horizontal neighboring rails and of the peak bounce can be explained by the direction of signal propagation within the critical path, which is from left to the right. Therefore, the point of highest stress appears in the right part of the power rail and the grid is therefore loaded asymmetrically.

For the same simulation, RA-VDDActive, the propagation of the maximum voltage drop in y direction is shown in Fig. 14. Again, the two loaded rails are shown clearly. In contrast to the previous figure, the fading away of the peak is much faster and abrupt. This is due to the fact that in y-direction the rails are connected through the second metal layer, which is low ohmic compared to the lowest metal layer, and the via stacks. The fading on the second metal layer is corresponding to the fading on the lowest metal layer. The peak decreases about linearly from one via stack to the next one.

A 3-dimensional plot of the peak propagation is shown in Fig. 15. For better illustration, a zoom-in at the via stacks at $x = 150\,\mu$m and $x = 300\,\mu$m is shown. The low ohmic contacts through the via stacks, as well as the fading away over the higher metal layers can be seen.

After the IR-Drop caused by two loaded rails has been examined, now the IR-Drop for the CPallover simulation will be looked at. In Fig. 16 the voltage drop in dependence on the x-direction is plotted for the three power rails at $y = 200\,\mu$m, $y = 205\,\mu$m, and $y = 210\,\mu$m, where the rail at $y = 200\,\mu$m is the first unloaded, the rails at $y = 205\,\mu$m and $y = 210\,\mu$m

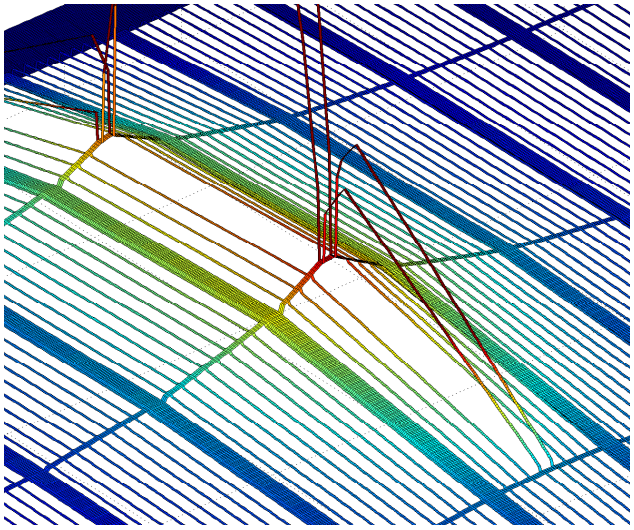

Fig. 15. 3-dimensional plot of the peak propagation over the power grid for the RA-VDDActive simulation, zoomed in to the via stacks at $x = 150\,\mu$m and $x = 300\,\mu$m

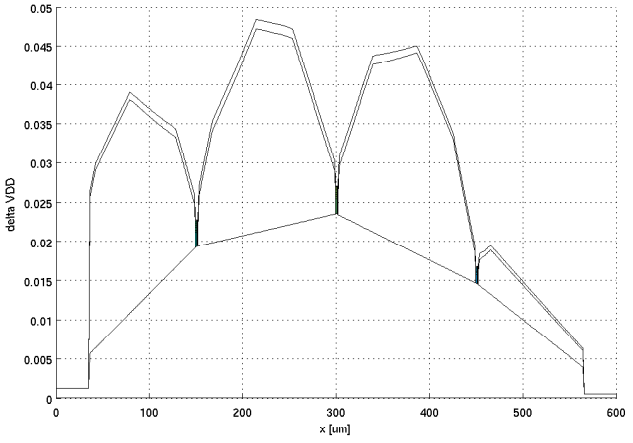

Fig. 16. Propagation of the peak voltage drop in dependence on x-coordinate for two loaded and one neighboring unloaded rail, for the CPallover simulation

are both loaded with critical paths.

In contrast to Fig. 13, in Fig. 16 three peaks can be observed for each loaded rail. Additionally, the highest peak increased from 24 mV in Fig. 13 to 48 mV in Fig. 16, which is about a factor of two. The peak value at the unloaded rail increased from 2.5 mV in Fig. 13 to 23 mV in Fig. 16, which is a factor of about 10. This reflects the heavier loading of the second metal layer, through which the unloaded rail is connected to the loaded ones.

In Fig. 17 the propagation of the voltage drop in dependence on the y-direction is shown. The major difference to Fig. 14 is that there are not two, but 10 rails loaded. As said before, the maximum voltage drop increases due to the heavier loading of the total power grid. The fading away

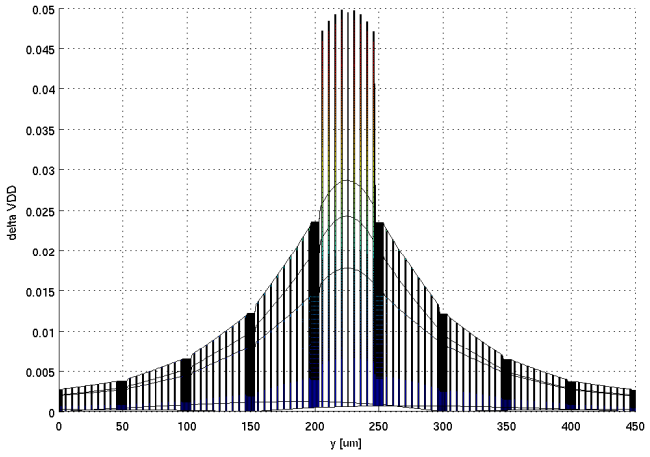

Fig. 17. Propagation of the peak voltage drop in dependence on y-coordinate for 10 loaded and the neighboring unloaded rails, for the CPallover simulation

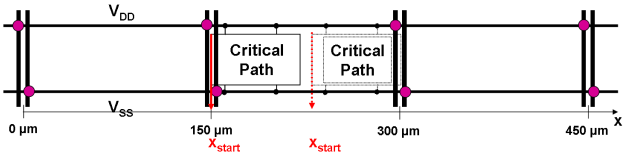

Fig. 18. Schematic of varying starting positions of the critical path within one rail

of the peak on the second layer is illustrated even better in this figure.

6 Delay implications

6.1 Placing in the rail

The layout of the critical path has a length of 86 μm, a rail in our setting has a length of 150 μm. To examine the impact of the placing of the critical path within one rail, all simulations were done for the starting positions $x = 150\,\mu$m, $x = 170\,\mu$m, $x = 190\,\mu$m, and $x = 210\,\mu$m in the rail ranging from $x = 150\,\mu$m to $x = 300\,\mu$m, as depicted in Fig. 18. In Table 2 the path delays for varying starting positions within one rail are displayed. The results are given for the falling edge, which is the slower one in the considered path.

Two effects can be seen. Firstly, the starting condition at $x = 190\,\mu$m is always the worst case starting point. The other observation is that the results do only differ by less than 1% between the best delay at $x = 150\,\mu$m and the worst case delay. Therefore, it can be said that the delay of the critical path is almost insensitive to the placing within the rail.

Table 2. Simulation results of the path delay for varying placing of the critical path within one rail

Simulation	Delay [ns] for $x_{start}[\mu m]$ 150	170	190	210	worst case x
GridOnly	1.364	1.366	1.366	1.365	190
RailActive	1.366	1.369	1.370	1.369	190
RA-VDDBuff	1.365	1.367	1.368	1.368	190
RA-VDD-Active	1.375	1.381	1.384	1.383	190
RA-VDDA-LRnoiseMP	1.378	1.384	1.386	1.386	190
RA-VDDA-LRnoiseCP	1.377	1.384	1.386	1.386	190
CPallover	1.405	1.409	1.409	1.407	190

6.2 Increased power supply loading

As already mentioned and expected, the path delay of the critical path increases with increased loading of the power grid. The simulation results for the path delay as well as the relative change of the path delay for the falling edge, which is the more critical one, are displayed in Table 3. The simulation without power grid serves as reference. As can be seen, the path delay increases up to 4.1% for the CPallover simulation. In this simulation setup, the power grid is loaded with 60 critical paths, all placed horizontally in the power grid.

However, if we compare the CPallover simulation with the OriChange16bit setup, we see that even if in the latter case the grid is only loaded with about a quarter of the gate count, the path delay degrades up to 5.8%. This is due to that fact that in the OriChange16bit setup, all gates within one power rail switch at the same time. Through the synchronous switching within one rail, the lowest metal layer is stressed more than in the CPallover case. This higher stress of the lowest metal layer causes a higher IR-Drop, which in turn increases the delay degradation. For the OriChange48bit simulation, which uses again less critical paths than the CPallover simulation, the path degradation increases even to 11.1%. Even if this setup might seem a bit unlikely to occur in real circuits, layouts like this can exist in data paths of semi-custom designed circuits, where the placing of the cells is done automatically.

7 Design strategies

In the last Section, the delay implications of IR-Drop were drawn for different simulation setups. In the following Section, two design measures to reduce IR-Drop, and their implications of path delay and area overhead are discussed. As noted earlier, there exist numerous possibilities to change the resistance of the power grid. There we examined two examples to show their different impact on path delay and area overhead.

Table 3. Simulation results for the path delay for different simulation setups

Simulation	Delay [ns] Rising	Falling	Rel. Change	ΔV_{DD} [mV]
No Grid	1.240	1.353	0.0%	0.0
GridOnly	1.252	1.366	1.0%	8.2
RailActive	1.257	1.370	1.3%	13.0
RA-VDDBuff	1.256	1.368	1.1%	10.3
RA-VDDActive	1.276	1.384	2.3%	28.2
RA-VDDA-LRnoiseMP	1.279	1.386	2.4%	27.2
RA-VDDA-LRnoiseCP	1.279	1.386	2.4%	30.2
CPallover	1.290	1.409	4.1%	49.8
OriChange16bit	1.312	1.419	4.9%	60.8
OriChange48bit	1.378	1.505	11.2%	93.6

The first design measure discussed in this Section is the increase of the line width of the lowest metal layer. Through the increased width of the power rail, its resistance decreases proportionally, which consequently results in a reduced IR-Drop.

A widening of the lowest power rails can be achieved in two ways. The first would be to increase the height of the standard cells. This has the advantage that the routing ability on the lowest level is not affected. However, the increase of line width results in a direct increase of the chip area. The other possibility would be to keep the standard cell height constant, but decrease the routing capability in the lowest metal layer. The area impact of this method can hardly be estimated in general, since the area overhead is dependent on whether a reduction of routability in the lowest metal layer is acceptable and if the routing can be done on higher metal layers. A widening of about 17% of the lowest power rail results for the OriChange16bit simulation in a reduction of 15% of the maximum voltage drop to 51.4 mV. The path delay increase due to IR-Drop is reduced to 5.1%, compared to 5.8% with the smaller power rail.

The second design measure to reduce IR-Drop is the doubling of the via count. Through this, the higher metal layers are better connected to the lower ones and the overall grid resistance is reduced.

The area impact of this design change is quite small, since on the lowest metal layer a via stack occurs only every 150 μm. Therefore, the routing capability on this layer is little affected. On the higher metal layers, things are even better, since routing usually is more relaxed on those layers and the pitches get even wider. The doubling of the via count results in a path delay increase of only 5.0% compared to 5.8% in the OriChange16bit simulation. The maximum voltage drop is reduced from 60.8 mV to 50.2 mV.

Therefore, better IR-Drop performance can be achieved by less area impact. This shows that for power grid optimization

a detailed sensitivity analysis, which is beyond the scope of this work, has to be carried out for all design parameters, to find the best trade-off between IR-Drop reduction and area impact.

8 Conclusions

In this paper, the effect of IR-Drop in on-chip power grids and its implications on the path delay of a critical path has been examined. The simulations show performance degradation of up to 4% respectively 11% for a critical path in typical topologies. It should be noted that in a chip, effects of local IR-Drop are superposed, giving rise to higher worst-case voltage bounces. It was also shown, that the layout of circuit plays an important role, which caused differences in IR-Drop by more than a factor of two in our settings. Finally, design counter measures for reducing IR-Drop and a brief discussions of the area implications were presented. Altogether, the paper underlines the increasing importance of power supply noise in digital integrated circuit design.

References

Benoit, M., Taylor, S., Overhauser, D., and Rochel, S.: Power Distribution in High-Performance Design, in: International Symposium on Low Power Electronics and Design, 1998.

Ernst, D., Austin, D. B. T., Flautner, K., Mudge, T., and et al: Razor: A low-power pipeline based on circuit-level timing speculation, in: Proc. 36th Ann. Int'l Symp. Microarchitecture (MICRO-36), IEEE CS Press, 7–18, 2003.

Henzler, S., Nirschl, T., Berthold, J., Georgakos, G., and Schmitt-Landsiedel, D.: Design and Technology of Fine-Grained Sleep Transistor Circuits in Ultra-Deep Sub-Micron CMOS Technologies, in: Proceedings of International Conference on Integrated Circuit Design and Technology, 2005.

http://www.itrs.net: International Roadmap for Semiconductors, Tech. rep., ITRS, 2003.

Lin, S. and Chang, N.: Challenges in Power-Ground Integrity., in: ICCAD, 651, 2001.

Piguet, C., ed.: Low-Power Electronics Design, CRC Press, 2004.

Rusu, S., Tam, S., Muljono, H., Ayers, D., and Chang, J.: A Dual-Core Multi-Threaded Xeon Processor with 16MB L3 Cache, in: Digest of Technical Papers. ISSCC, 102–103, 2006.

Tschanz, J., Kao, J., Narendra, S., Nair, R., Antoniadis, D., Chandrakasan, A., and De, V.: Adaptive body bias for reducing impacts of die-to-die and within-die parameter variations on microprocessor frequency and leakage, in: Digest of Technical Papers. ISSCC, 2, 2002.