


Powering large FPGAs



J.-P. Cachemiche
F.Rethoré

FPGA power consumption

Tends to increase due to several factors

- Increasing number of logic cells (up to 10 millions in last generation devices)
- Increasing internal frequency
- Current leakage with thinner technologies

Manufacturers implement many tricks to reduce power consumption

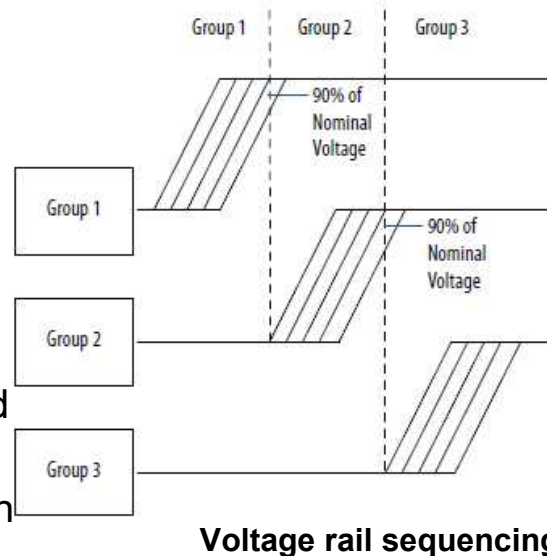
- Increase LUT size
- Decrease voltages
- Dynamic power down of scarcely used logic

Challenge

- Even if the power consumption of each logic element is lower than previous generations, the size of matrixes utilized, high-frequency designs could require core currents in excess of 100 A.

FPGA powering specificities

- Many voltages with many individual current requirements
 - o More than 25 different supplies in last generation components
 - o No real max power : depends on the firmware
- Voltage rail sequencing :
 - o Different supply voltage rails to come up in a specific sequence.
 - o Most of the time the core voltage needs to be supplied before the I/O voltages come up. Otherwise the FPGA can be damaged.
- Monotonic rise of voltage rails
 - o Ramping to help prevent excessive inrush current during startup
 - o A slower rising time could cause failure in the device configuration and incorrect operation
- Fast power transients
 - o Not supported by the power supply
 - o A complex capacitor network is required
- Voltage accuracy
 - o 2 or 3 % voltage accuracy very common
 - o Probably the most challenging issue due to voltage drop and parasitic inductances



Power group	Voltage
Group 1	VCC
	VCCP
	VCCH
	VCCL_SDM
	VCCH_SDM
	VCCPLLDIG_SDM
	VCCL_HPS
	VCCPLLDIG_HPS
	E-tile
	VCCRT_GXE
	VCCRTPLL_CR3_GX
	P-tile
	VCC_HSSI_GXP
VCCRT_GXP	
VCCFUSE_GXP	
Group 2	VCCPT
	VCCPLL_SDM
	VCCADC
	VCCPLL_HPS
	VCCM_WORD
	E-tile
	VCCH_GXE
	VCCCLK_GXE
Group 3	P-tile
	VCCH_GXP
	VCCCLK_GXP
	VCCA_PLL
	VCCIO_PIO
VCCFUSEWR_SDM	
VCCIO_SDM	
VCCIO_HPS	

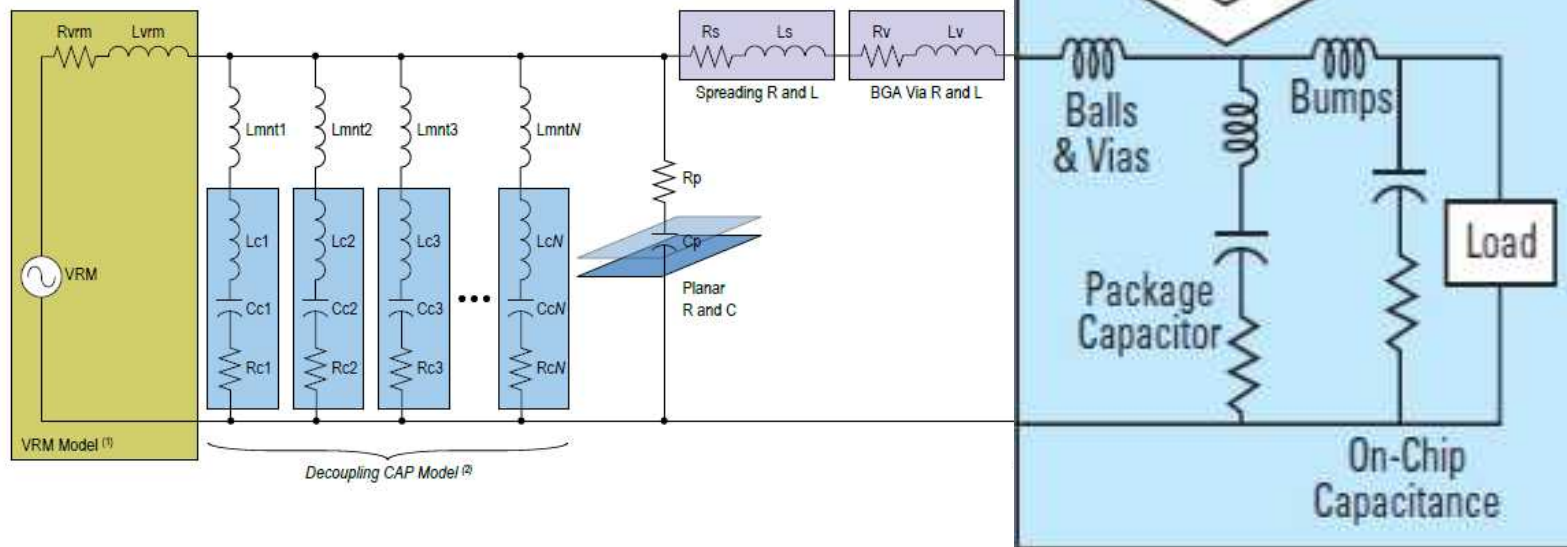
Power supplies for the Intel Agilex

Points to focus on

Power Delivery Network (PDN)

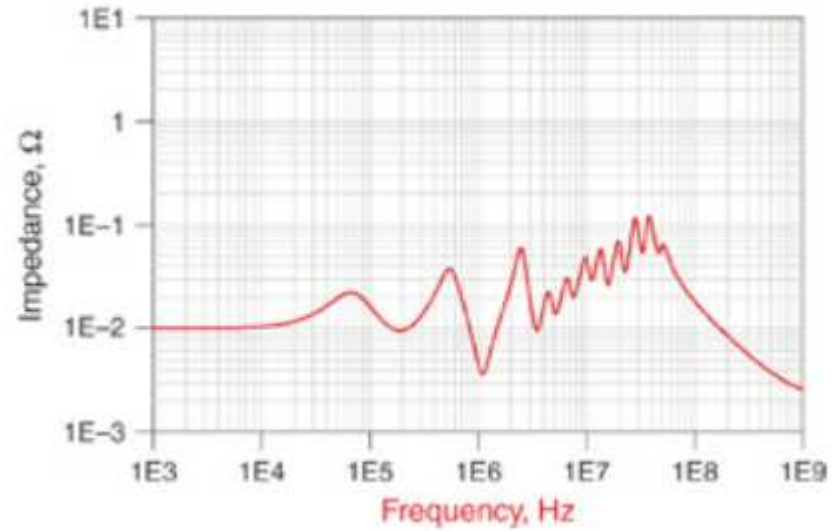
- Complex path
- It consists of all the interconnects from the Voltage Regulator Modules (VRMs) to the circuits on the die.
- Unsuccessful noise control on the PDN will contribute to contraction of the eye of any signal.
 - ➔ In particular for high speed signals : the vertical ripple noise will translate into horizontal jitter

Structure of the PDN



Typical impedance profile

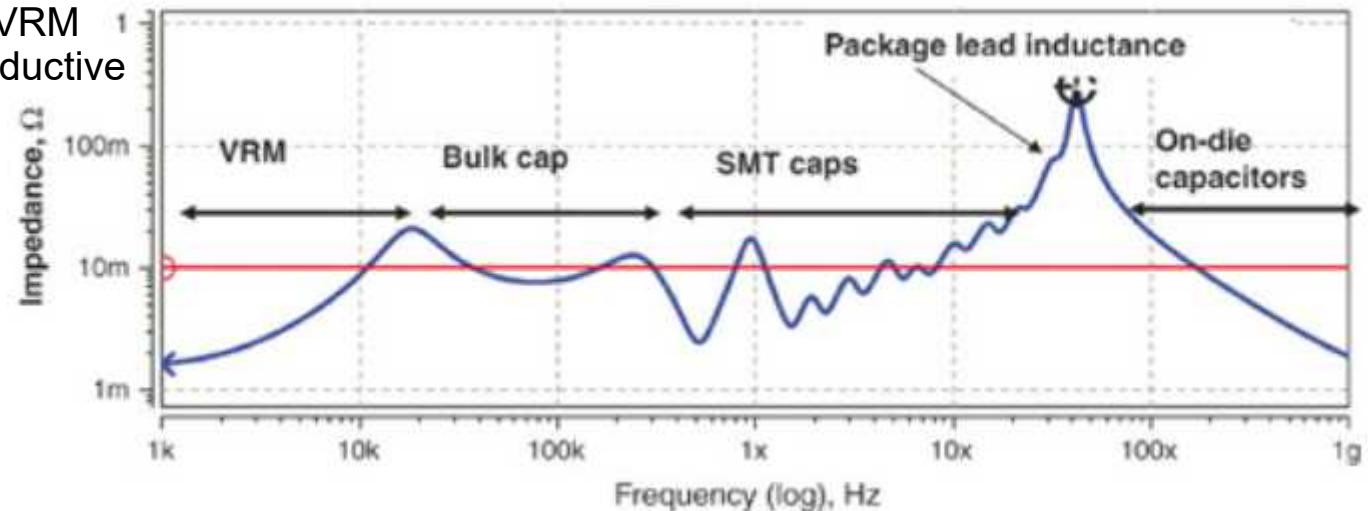
Generally looks like this :



Typical impedance profile

Landscape sculpted by :

- VRM
 - o At very low frequencies, up to approximately 50 kHz, the VRM has a very low impedance and can respond to the instantaneous current requirements of the FPGA
 - o At higher frequency, the VRM impedance is primarily inductive
- Bulk capacitors
 - o For medium frequencies
- SMT capacitors
 - o For higher frequencies
- Parasitic lead inductance
 - o Resonance due to lead inductance, wire bonds, solder balls and on-die capacitors
 - o Very little thing you can do above the base of the rising slope of this resonance
- On-die capacitors
 - o For very high frequencies



Constraining the design



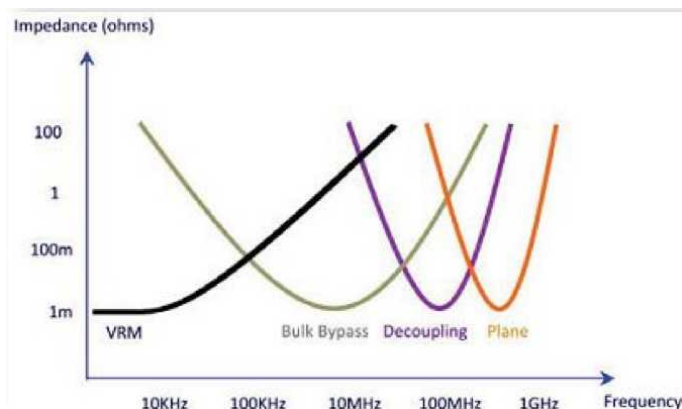
The goal will be to limit transient currents under the maximum noise accepted by the chip

- For this, the impedance of the PDN must be below a certain level: the **target impedance**.
- This impedance is usually defined as :

$$Z_{target} = \frac{\Delta V_{noise}}{I_{Max-transient}}$$

where : ΔV_{noise} is the maximum allowed ripple
 $I_{Max-transient}$ is the maximum dynamic current

Constraining the design will consist in dampening the peaks above the target impedance with capacitors and planes




Constraining the design

Decoupling using PCB capacitors becomes ineffective at high frequency

- Using PCB capacitors for PDN decoupling beyond their effective frequency range brings little improvement to PDN performance and raises the bill of materials (BOM) cost.
 - Limited by capacitor ESL and ESR
 - Range limited by path inductance between capacitors and die
- This limit is called the **effective frequency**
 - the impedance boundary is set by the vertical impedance (PCB vias and package resistance and inductance)
- At high frequencies decoupling is assured by on-die capacitors


Constraining the design will consist in pushing as much as possible this limit towards high frequencies in reducing parasitic inductances



PDN optimization

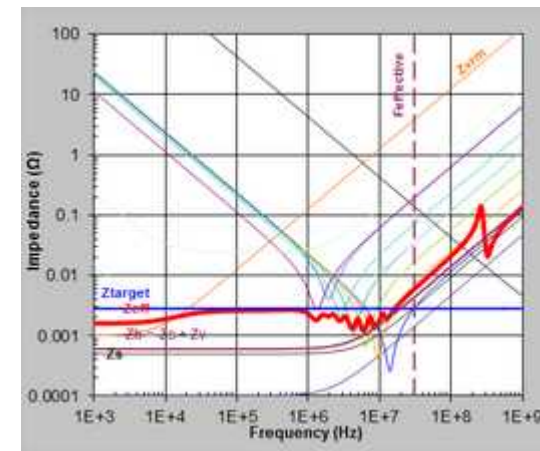


Avoid power via grouping

- Via grouping restricts current flow



Example for an Arria10

- Increasing the power vias pairs from 50 to 197 increases the effective frequency from 10.62 MHz to 30.68 MHz



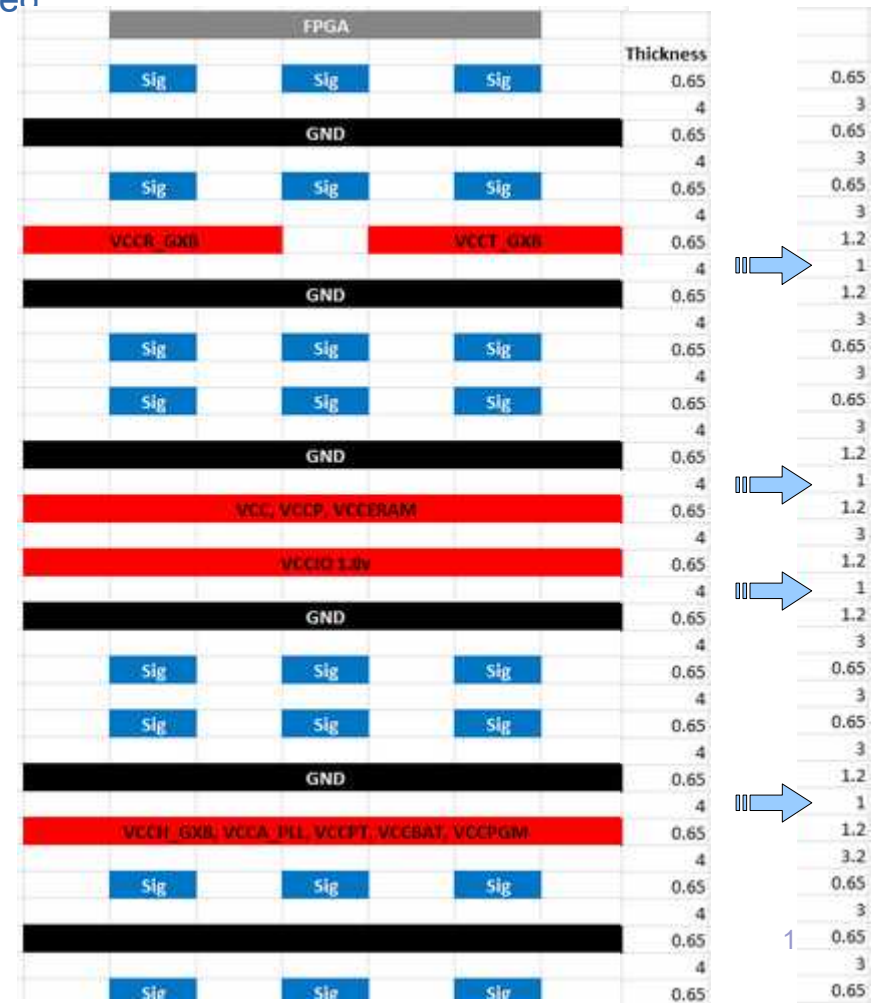
PDN optimization

Power planes

- Moving power and ground planes closer together
 - ➔ Reduces the number of discrete capacitors needed
- Increasing thickness of power planes
 - ➔ Reduce resistance

Example for an Arria10


- Reducing power plane separation from 4 to 1 mils leads to :
 - o Effective increases for VCC from 30.68 MHz to 36.71 MHz
 - o Number of capacitors reduced from more than 300 to 240 for transceiver planes



PDN optimization


Loop inductance

- The full back and forth path has to be taken into account
- Mainly two contributors:
 - o inductance due to capacitor distance
 - o inductance due to vias



Optimizing capacitor placement

- Capacitors decoupling high frequencies must be placed as close as possible from the FPGA power pins
 - ➔ Otherwise effect cancelled by track inductance
- Large capacitors can be placed quite far.
 - ➔ Path inductance negligible




Efficiency of a capacitor in function of distance

Ihsan Erdin, Ram Achar, "Placement of decoupling capacitors on power transmission lines",

PDN optimization

Other way to optimize capacitor placement

- Traditionally :
 - o The decoupling capacitors are placed at the bottom side of the PCB, beneath the FPGA and connect directly to the BGA vias.
 - o It is assumed that the best electrical position is the closest physical location to the FPGA pins.
- If the power and ground plane pair are close to the FPGA, the total vertical inductance is less if the decoupling capacitors are placed on the top surface of the PCB.
 - ➔ inductance is much lower even though the horizontal distance of the capacitors placed around the outside of the FPGA is far greater



Example for an Arria10

- Sharing capacitors between top and bottom allow to reduce their number from 255 to 180 for transceiver planes

PDN optimization

Influence of caps mounting and geometry


- Use small capacitors (0402, 0201)
- Recommended to avoid leads between capacitor and connecting vias
 - Increase parasitic inductance
- Ideal case uses 4 vias

Still better : X2Y capacitors



- Several layered capacitors
- Reduces the ESL of each capacitor by 1/n,

Example for an Arria10

- Using X2Y capacitors in the PDN Tool reduces the number of capacitors required to meet the target impedance of the transceiver supplies from 255 and 180 (TX) to 28 and 22 (RX).



Capacitor inductance
E. Bogatin



X2Y capacitors
E. Bogatin

PDN optimization

Same kind of optimization can be made for bulk capacitors

- Use of ultra-low ESR bulk capacitors
 - o Very low serial resistance : 5 m Ω instead of 50m Ω for traditional bulk resistors

Example for an Arria10

- Using ultra-low ESR decrease the number of bulk capacitors from 23 to 17


PDN optimization

Still a very large amount of caps needed for VCC

- Even if you add as many caps you can, simulation tools will not show any improvement


Possibility to derate the target frequency

- The current ramp-up period is proportional to the length of the pipelines
 - o If the pipeline are short ramp-up is abrupt
 - o If the pipeline are narrow and long the current change is proportionally smaller
- Statistical effects average all this over the core



Example



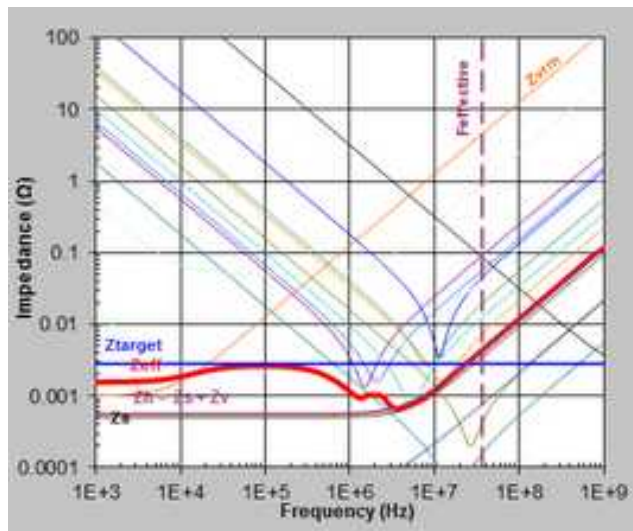
- Running frequency = 350 MHz, average pipeline length = 10
- Maximum transient current frequency = 35 MHz
- Above 35 MHz, the worst-case transient current spectrum would drop off at -20 dB/decade and the resulting target impedance would increase with frequency.



PDN optimization

Example for Arria 10

- For a running frequency of 300 MHz and a current ramp up period of 25 clock cycle the number of decoupling capacitors required to meet the VCC supply is reduced from 301 to 37
 - Large improvement



How to estimate the power consumption ?

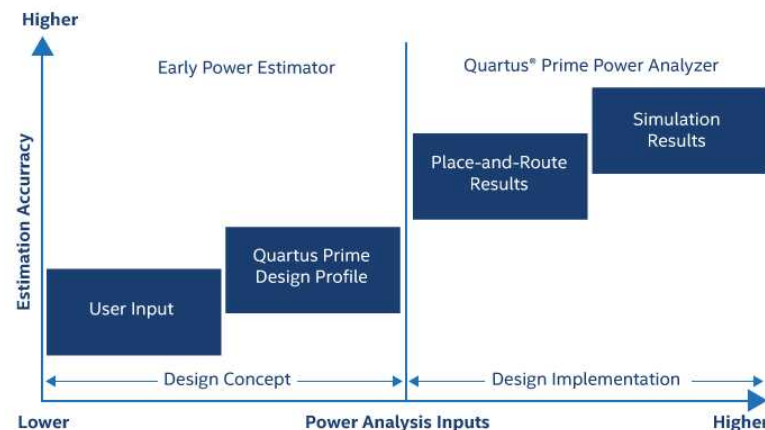
Manufacturers provide tools

Ideal case (... that never happens !):

- The HDL code is completely known when you start designing your board
 - ➔ Use of a post routing power estimator

Normal case:

- You start designing your board ...
- But have only an rough estimation of what the size of the final firmware will be
- You must make engineering decisions to guess the needs
 - ➔ Use of an early power estimator (Excel based) where all the parameters have to be entered manually



Early power estimator

Allows to :

- Evaluate the power consumption for each power rail
- Calculate a target impedance
- Calculate an effective frequency
- Can automatically propose the number and type of capacitors required
- Calculate an effective impedance vs frequency in function of number and type of decoupling capacitors
- Calculate the die temperature in function of the cooling air flow

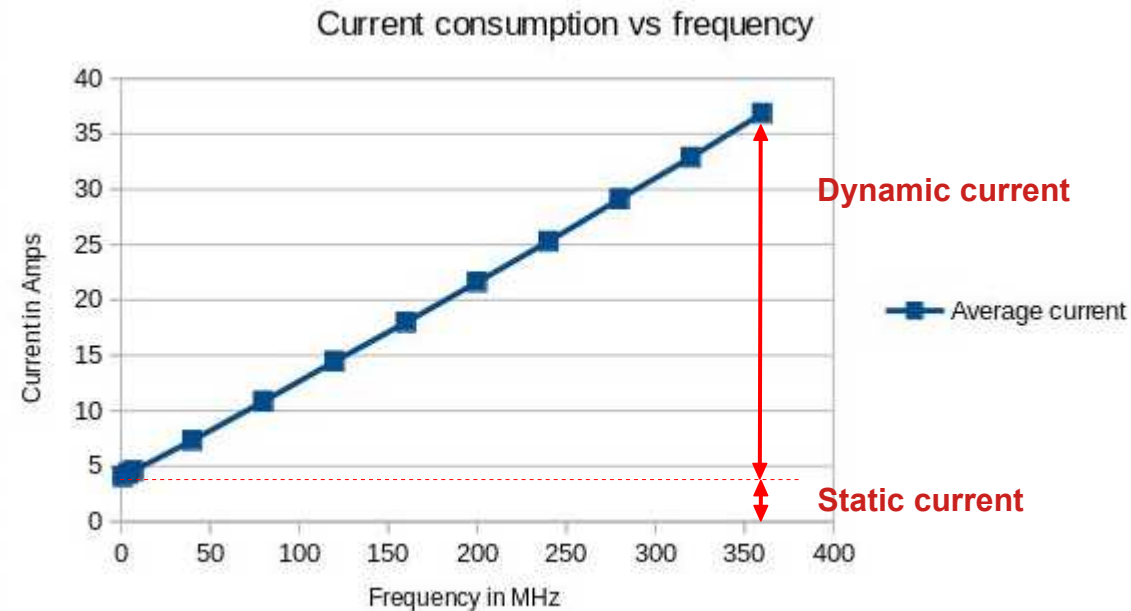
Input data for each rail

- Operating frequency
- Toggle rate
- Occupancy
- Power plane structure and dimensions
- Dielectric materials
- Capacitor type and geometry
- Via sizes and geometry
- VRM parasitics
- Estimated current Ramp Up Period
- Air flow

Power consumption

Two types

- **Static power** : the power that the configured device consumes when powered up but no user clocks are operating (leakage current) ;
- **Dynamic power** : the additional power consumption of the device due to signal activity.
Dynamic power is dependent on :
 - o the **logic cells occupancy**,
 - o the **operating frequency** of the design,
 - o and its **toggling rate** (number of state change per clock period).
 - ➔ Very dependent on the firmware
 - ➔ The most important contribution
- Example :
 - o Arria10 power consumption vs frequency
 - o 90 % logic cells used
 - o Toggle rate ~ 50 %
 - o Frequency sweep between 0 and 360 MHz
 - ➔ Only limited by cooling



Methodology used for the PCIe40 design

Estimation of the 3 principal features

- Toggle rate ~50 % for 90 % of the design (by default the PDN tool set it to only 12%)
 - o Because we mainly handle compressed data, similar to random data for which the toggle is 50 %
- Operating frequency = 240 MHz
 - o Actually 200 MHz used in the LHCb firmware
 - o ... But 280 MHz for the ALICE firmware
 - o We do not know for Mu3e nor Belle II
- Occupancy = 90 %
 - o Because physicists will not stop adding features as long as some room is left ! ☺
 - o Above 90 % we cannot route the FPGA

➔ **Core power estimation for PCIe40 : 52 A**



All the optimizations explained before for routing the card (except X2Y)

Derating with average pipeline (25 clock cycles)

Final number of decoupling capacitors for the FPGA

- VCC : 139, VCCT : 71, VCCR : 54, Others : 62, Total : 326

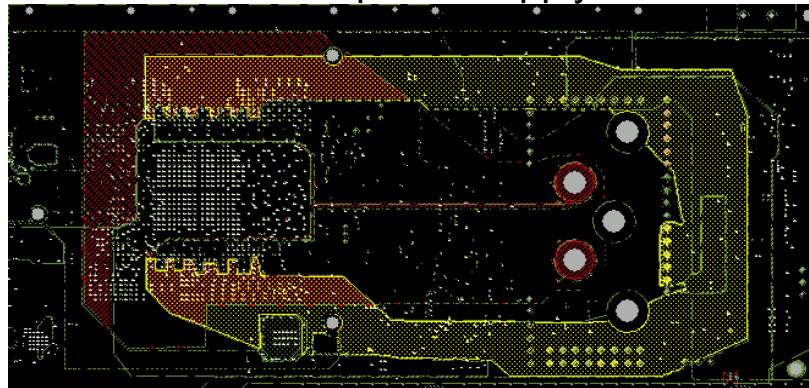
VRM sense for the PCIe40

Mainly on voltages with high currents

- To compensate voltage drop and respect accuracy

Geometry issue

- Due to a large number of power planes, sometimes the plane geometry is a bit weird
- Where to put the sense for the receivers power supply in this case ?



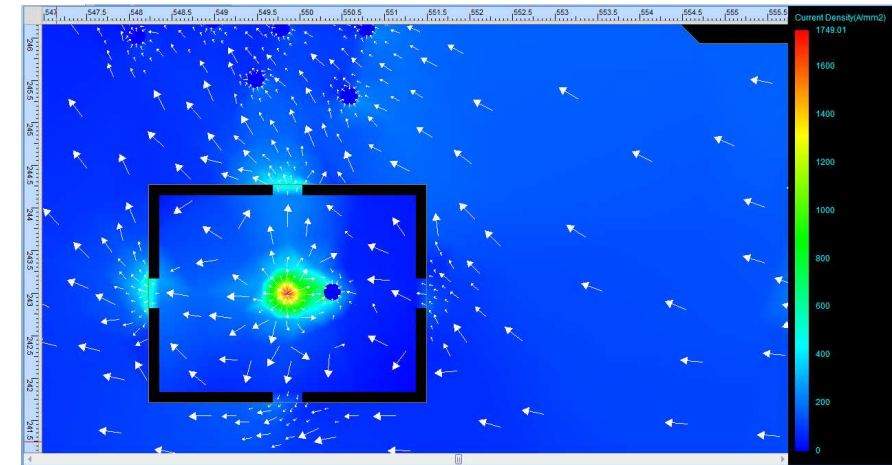
- If sense placed on one side the voltage drop is only compensated on this side
 - ➔ Use two VRMs ?
 - ➔ Fortunately transceivers are always on in the Arria10, even if not used, to avoid aging

Designing is making a lot of compromises

Power flow simulation

Important to simulate the power flow

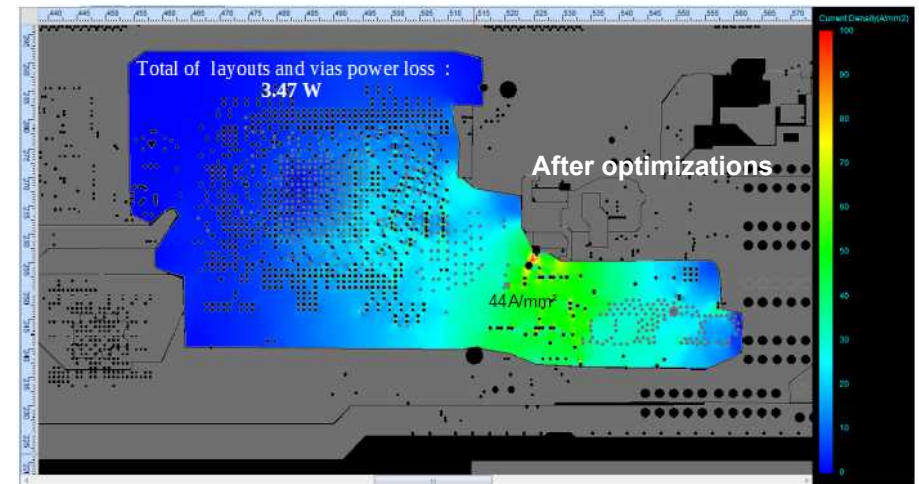

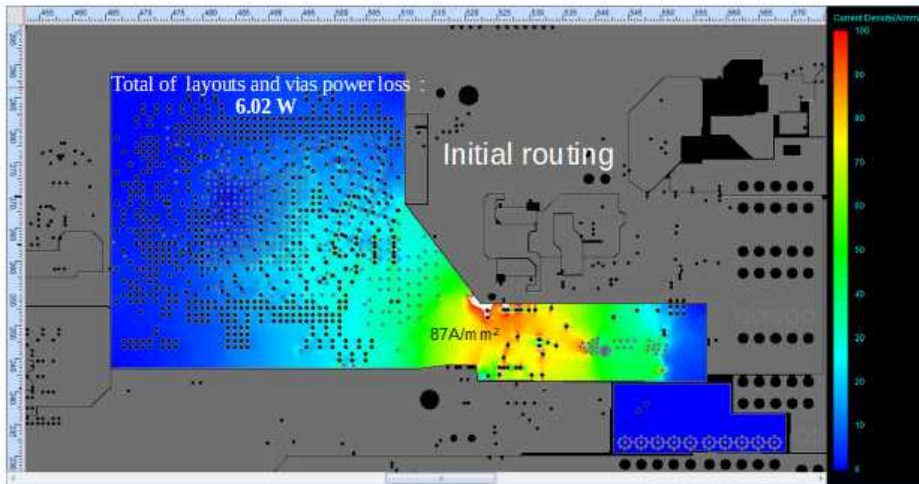
- High currents on the core and badly placed vias can generate very hot points
 - ➔ Vias can melt
 - ➔ Delamination issues
- Energy lost in the planes increases power consumption unnecessarily



Sigrity simulation :

Via in pad : 30A → 450 A/mm²

Thermal connection : 8,34 A in the via (max authorized = 7.5 A)



Conclusion

Power estimation is difficult when you do not know the final firmware

- Many most likely hypothesis have to be made
- Specially if the card can run many applications or firmwares

Many design tricks can decrease the parasitic elements : but simulations required

PDN tools are not perfect

- Manufacturer tools know nothing about the capacitor placement
- EDA tools know nothing about the internal parasitics of the FPGA package

Nearly impossible to design a card that can sustain any kind of application

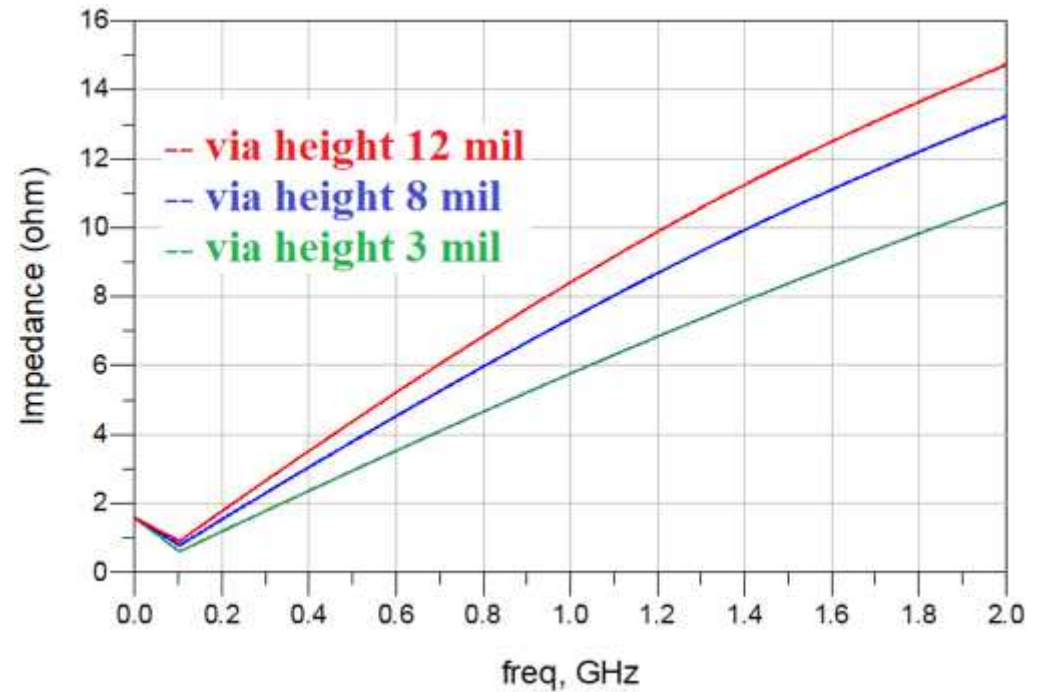

- The more you know about the final firmware the better you can optimize
 - ➔ Important to have early simulations of the final firmware, specially for the most consuming tasks
- At some time, limits exist and firmware people have to cope with it
 - ➔ Example : possibility to run several shifted clock domains, or increase pipeline speed

Put some automatic power cut security (just in case ...)

More information

Path length vs via height

Resulting impedance



Power consumption of Arria10 example shown in the slides

Estimated power per rail

Supply Name	Nominal Voltage (V)	Allowable Ripple (%)	I _{max} (A)	Current Transient (%)	Target Impedance (mR)
VCC	0.9	5	30	50	3
VCCR_GXB	1.0	3	6	30	16.67
VCCT_GXB	1.0	2	2	50	20

Main power supplies


Power Supply Pin	Voltage(V)	Current (mA)	Power Group
VCC	0.9	32000.00	1
VCCP	0.9	13500.00	1
VCCERAM	0.9	0.045	1
VCCR_GXBL1C	1.0	224.09	3
VCCR_GXBL1D	1.0	403.75	3
VCCR_GXBL1E	1.0	991.02	3
VCCR_GXBL1F	1.0	1021.69	3
VCCR_GXBR4C	1.0	944.15	3
VCCR_GXBR4D	1.0	955.29	3
VCCR_GXBR4E	1.0	667.44	3
VCCR_GXBR4F	1.0	757.41	3
VCCT_GXBL1C	1.0	56.42	2
VCCT_GXBL1D	1.0	78.86	2
VCCT_GXBL1E	1.0	297.65	2
VCCT_GXBL1F	1.0	372.85	2
VCCT_GXBR4C	1.0	356.23	2
VCCT_GXBR4D	1.0	356.23	2
VCCT_GXBR4E	1.0	281.97	2
VCCT_GXBR4F	1.0	298.58	2

Detail

Measuring the PDN

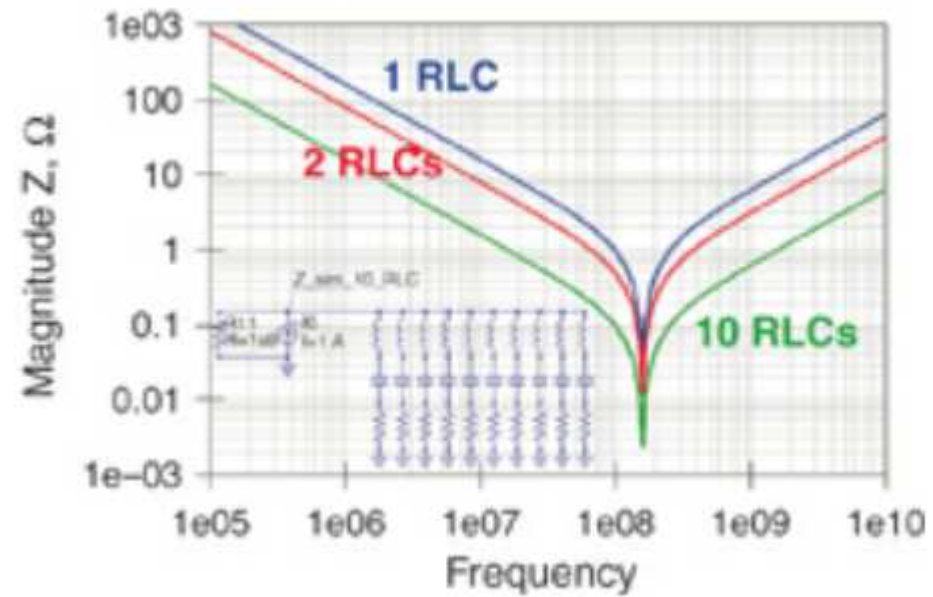
Possibility to use a VNA to extract FPGA internal parasitics

- 2-port S-parameter measurement
- Expensive :
 - ➔ A dedicated prototype has to be built



X2Y capacitors

Allow to increase capacitance and decrease the parasitic inductance




E. Bogatin

Current ramp up

5000 pseudo-random generators running simultaneously

➔ Not so random




Simulation

Measurement

Introducing 16 different seeds and more feed backs

➔ Closer to reality



Simulation