

# IBM Research Report

## The Effect of Power Islands on Delta-I Noise, Interconnect Noise, and Timing for Wide, On-chip Data-buses

**A. Deutsch, H. H. Smith<sup>1</sup>, H.-M. Huang<sup>2</sup>, A. Elfadel**  
IBM Research Division  
Thomas J. Watson Research Center  
P.O. Box 218  
Yorktown Heights, NY 10598

<sup>1</sup>IBM Systems and Technology Group  
2455 South Road  
Poughkeepsie, NY 12601

<sup>2</sup>IBM Systems and Technology Group  
2070 Route 52  
Hopewell Junction, NY 12533



Research Division  
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

# The Effect of Power Islands on Delta-I Noise, Interconnect Noise, and Timing for Wide, On-chip Data-buses

A. Deutsch, H. H. Smith<sup>1</sup>, H-M Huang<sup>2</sup>, A. Elfadel

IBM T. J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, N.Y. 10598,  
Phone: (914) 945-2858, Fax: (914) 945-2141, email:deutsch@us.ibm.com

<sup>1</sup>IBM Systems and Technology Group, 2455 South Road, Poughkeepsie, NY 12601

<sup>2</sup>IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, NY 12533

**Abstract** – A study is shown of the effect of having breaks in the power distribution on large microprocessor chips. The effect on delta-I noise, interconnect noise, and timing is illustrated through simulation results obtained with representative driver and receiver circuits and guidelines are given on how to minimize the impact of the power islands.

## Introduction

Present large microprocessors could contain more than one CPU core, large banks of data caches, and I/O macros as shown in Fig. 1. It has been found that due to the large leakage currents, need for noise isolation, and different voltage rails required, these different regions needed to have partitioned power distribution layout. Wide data-buses leaving or entering one processor unit will encounter at least one break in the power rail. The simultaneous switching of large currents into the 1-3 mm long interconnects will have a current return path that is not continuous and will generate potentially large noise contributions. It has been shown in [1], that the typical power distribution on chip will exhibit a finite impedance  $Z_{eff}(f) = R_{eff}(f) + 2\pi fL_{eff}(f)$  that will be resistive and inductive and frequency dependent. This is caused by the large difference between the solder ball pitch (200 – 400  $\mu\text{m}$ ) where current is fed into the chip and the actual device power contacts (5 – 10  $\mu\text{m}$ ), the resistive, sparse power conductors on many layers with signal-to-power line ratio of 3:1 to 10:1. This effective  $Z_{eff}(f)$  will generate the common-mode noise (CMN) on the interconnects that will add or subtract from crosstalk noise, and delta-I noise on the power rails. Due to the high integration level, the need to control the interconnect characteristics and contain noise, the power distribution will have a fairly regular pattern. This pattern then can be captured by a power-block defined in terms of lossy transmission lines with per-unit-length properties. Only two such blocks are needed, X or Y oriented, and in the direction perpendicular to the lines in the data-bus. This will provide the worst case current return scenario. The reduction in computational cost to generate these reduced size blocks versus full-chip modeling and ease of representation with distributed segmented circuit models allows for many non-linear simulations. Such simulations can be made with actual driver and receiver circuits and all the non-linear interactions between line and power noise and the effect on timing can easily be captured.

It is shown in [1] and [2] that the frequency-dependent impedance and admittance for both the long interconnects and the power blocks can be synthesized with distributed, lumped-circuit segments with Foster-type filters. Decoupling capacitors and device capacitance can be attached to the power rails in distributed manner and be very close to actual physical characteristics of the layout.

This paper will report on the findings of the interaction between the on-chip noise sources and their effect on timing in the presence of partitioned power distribution and using the same methodology as described above and in [1]. The contribution of decoupling capacitors will be discussed and the dependence on data-pattern and skew will be shown. Guidelines will also be given on how to minimize the effect of the power islands.

## Simulation Conditions

Fig. 2 shows the circuit diagram for the simulations. 24-signal data-bus is shown with additional three signals that represent the three power rails under consideration, Vdd, Vcs, and Vgnd. A two-stage driver buffer bank is shown at the drive end with drive impedance of 25  $\Omega$ . Similarly, there is a two stage receiver bank to simulate a multi-stage buffered path. All the signal lines were 3-mm long and having 0.8  $\mu\text{m}$  widths and spacings on the topmost layer, with  $R = 229 \Omega/\text{cm}$ . Two power-blocks are shown, one at the driver and one at the receiver end. Vdd, Vcs, and Vgnd represent the three lossy transmission lines in the power-block, referenced to an ideal ground plane. These blocks connect to the package side where the power is fed into the chip. In this analysis, the chip carrier is assumed to have an ideal impedance and supply a clean 1.0 V voltage level. The actual package model can easily be added in the simulation as shown in [1]. Fig. 2 is a schematic representation of the line and power-block models. The power-block in this case contained all the power and ground conductors on the top four layers with a unit cell of 64 x 200  $\mu\text{m}$ . The solder ball pitch was 200  $\mu\text{m}$ . Decoupling capacitor density was assumed to be 5  $\text{fF}/\mu\text{m}^2$  and the device intrinsic capacitance density was 0.8  $\text{fF}/\mu\text{m}^2$ . It was assumed that the capacitors had a 20 ps time constant and only half of the available devices were active. Thus for a 3 mm length and 200  $\mu\text{m}$  span, the device decoupling was 240 pF and for the unit power-block we had 64 pF for Vdd and 16 pF for Vcs decoupling with respect to the local Vgnd at the driver end. The target delta-I noise budget was 100 mV with the interconnect noise

budget of 400 mV. The effect of noise on delay was targeted to be around 10 ps and anticipated skews in buffer timings were  $\pm 15$  ps. Targeted propagated risetimes were 50 – 75 ps.

### Simulation Results without Power Islands

The analysis started with all the power rails continuous. Six, twelve, and twenty-four buffers were switched with a data pattern that assured the worst case summation of crosstalk and common-mode noise at the input to the receiver, as was explained in [1] and [2], namely --- --- +V+ ---. This means that the victim line V is monitored and the + + buffers switch in opposite direction to the --- buffers.

Fig. 3 shows the line noise at the receiver input, the signal at the output of the driver, the signal at the end of the 3-mm lines, and the delta-I noise monitored on the Vdd rail. In this case, the line noise increases from 360.5 to 459.0, to 499.3 mV. Delta-I noise values were 117.4, 231.3, and 375.7 mV at the driver power-block and 55.7, 143.8, and 204.0 mV at the receiver block. Delta-I noise increases nearly linearly with the number of drivers, half of it propagates to the receiver buffers, while the line noise saturates around 12 buffers. If device decoupling is added along the lines or decoupling capacitors at the driver, or both, delta-I reduces to 55.8 mV for 24 lines. The interconnect noise, however, is only minimally reduced. It was explained in [1] that CMN noise increases with the number of drivers. As the number of in-phase buffers increases, the effective line impedance increases and the driver risetime becomes faster. On the other hand, delta-I noise increases with the number of drivers. This will then slow-down the driver risetime and thus generates less crosstalk and CMN. There is a non-linear compensation occurring between line noise and power supply noise that is hard to predict and requires this type of non-linear simulations to assess. Introducing 15 ps skew between the center buffers and the rest of the 21 buffers did not affect the line noise or delta-I noise.

The effect of noise on timing was investigated for three data patterns, namely 000 000 0+0 000, +++ +++ +- +++ , and - - - -+- ---. This would correspond to no interconnect noise, worst case line noise, and smallest line noise but highest delta-I noise. Fig. 4 shows the driver end and line end signals with all the decoupling present for a 24-line case. Delays (buffer and interconnect) were 32.0, 73.2, and 44.1 ps, while risetimes were 38.8, 91.5, and 121.6 ps. A comparison of the simulations with and without decoupling, revealed that decoupling had very little effect on timing. It is noticed in Fig. 4 that the risetime distortion is the largest for the case when the interconnect noise would be the largest due to the very high coupling between the lines.

### Simulation Results with Power Islands

Simulations were performed by assuming that the interconnects were crossing over a complete break in all the three power rails, Vdd, Vcs, and Vgnd. This break was placed 200  $\mu\text{m}$  away from the driver so the lines had two separate sections, 200  $\mu\text{m}$  and 2.8 mm long. In this case, the delta-I noise increased significantly and linearly with the number of drivers, 6, 12, 24, to values of 193.4, 351.0, and 491.0 mV, respectively. At the same time, the interconnect noise no longer showed saturation. It increased to 402.3, 676.3, and 953.8 mV for 6, 12, and 24 lines. When decoupling and device capacitance was included, delta-I noise was again reduced to 21.8, 42.2, and 100.5 mV as seen in Fig. 5. Even though it is increasing linearly, it is very small and does not exceed the budgeted level. On the other hand, the interconnect noise is very little reduced by decoupling to only 371.5, 592.0, and 800.4 mV, respectively. The drivers see the open section of line very close, the signal shows large overshoots because of this, and the interconnect noise is very high. The break in the power grid was also moved from being close to the driver, at 200  $\mu\text{m}$ , to 1.0 mm, 2.0 mm, and 2.8 mm distance. The interconnect noise levels were 592.0, 621.4, 548.8, and 501.7 mV, respectively, for a 12-line switching case. This indicated that the worst location of the power island was close to the driving end of the lines.

Fig. 6 shows the simulation for 12 lines with two data patterns, namely +++ +++ +V+ +++ and --- --- +V+ ---. In the first case, the crosstalk and common-mode noise would subtract resulting in only 230.6 mV of noise. In the second case, the two noise contributions add and result in 592.0 mV. This same trend was found before, without power breaks, in [1], and the effects remained the same. All the available decoupling was included and the power breaks were again at 200  $\mu\text{m}$  distance. This behavior is very important to understand because it affects the timing results.

### Compensating for the Negative Effect of Power Islands

Various attempts were made to reduce the excessively large interconnect noise due to the presence of power rail breaks. It was found that the line noise dropped from 800.4 mV to only 744.4 mV if the Vcs rail was continuous. Removing the decoupling for the Vcs rail also had only a minimal effect. Adding some connecting straps on the Vgnd rail of 1  $\Omega$  to 10  $\Omega$  reduced the line noise only to 534.1 or 666.5 mV, respectively. Having Vgnd fully continuous, however, reduced the line noise to 488.9 mV for 24 line switching, which is very close to the level without any breaks, namely 476.3 mV. Similar effects were found for the Vdd rail. Inserting some resistive straps of 10  $\Omega$  to 100  $\Omega$  had no effect at all. Having Vdd continuous with Vgnd open, Vcs open, and even no decoupling for Vcs, dropped the line noise to 454.6 mV. Having Vgnd or Vdd fully continuous alleviates the effect of any other breaks. Moreover, when Vgnd or Vdd are continuous, the line

noise will once again start to saturate with the number of drivers switching. With Vgnd connected, the line noise went from 448.3 to 490.3 mV for 12 and 24 active drivers. Delta-I noise doubled from 50.3 to 101.2 mV but was still within the desired budget.

### Effect on Timing

The total delay and propagated risetime degradation were monitored for 24 active lines, with decoupling included, and for 000 000 0+0 000, +++ +++ +-+ +++, and --- --- +- - data patterns. The center + line was monitored. Fig. 7a contrasts the propagated signals with and without power breaks for the data pattern that has the highest line coupling while Fig. 7b shows the response for the lowest line noise case, or --- --- +- - . With power breaks, total delay (line and buffer) goes from 32.3, to 100.1, to 29.9 ps, respectively, for the three data patterns which is a very large range and will result in large timing jitter. Propagated risetimes were 41.2, 97.0, and 24.9 ps, respectively. The break in the power rail is very close to the drivers and thus the devices are driving an open-ended interconnect and overshoots due to doubling are noticed in Fig. 7b. For this case, the effective line impedance is higher and driver response is faster and thus delay is smaller. The risetime in Fig. 7b exhibits distortions due to interconnect noise above the receiver switching threshold. The large interconnect noise is responsible for the large risetime distortion in Fig. 7a. Even 15 ps of skew between drivers in the bank can generate additional risetime distortion of 2-7 ps for the pattern of Fig. 7a.

Finally, having either Vgnd or Vdd continuous, also restored the timing to the values obtained without any power islands. For Vgnd connected, delay was 82.3 ps, for Vdd connected, delay was 74.5 ps, while for no-break case, the delay was 73.2 ps. Risetimes were 88.2, 91.5, and 91.5 ps, respectively. Comparing the timing with the presence of power rail breaks showed that adding decoupling did not affect the timing response more than 3-6 ps in delay. This is happening because, as it was shown earlier, decoupling is not helping reduce the interconnect noise. All the noise on the bus is circulating between the lines with no clear return path due to the breaks and is affecting the propagated waveform shapes.

### Summary

In conclusion, it was shown that the effect of having power islands on large microprocessor chips can be easily analyzed with the power-block methodology. Even without power rail breaks, delta-I noise is increasing with the number of simultaneously switching drivers while interconnect noise saturates around 12 active devices. Decoupling capacitors help reduce delta-I noise but timing is not improved. This was the case even for the power island condition. When power breaks are introduced, both delta-I noise and interconnect noise will increase with the number of drivers, even with decoupling. Decoupling is not reducing the line noise, only the power rail noise. The worst impact the breaks have is when they are close to the driver end of the lines.

Introducing resistive straps for Vgnd or Vdd rail did not improve the response significantly. Having one of the rails fully continuous assured, however, that both the line noise and the timing were restored to the level without any power islands, and the line noise showed again a saturating trend.

Power islands introduce significant risetime distortions especially for data patterns that generate the highest interconnect coupling. There is a large difference in timing between extreme data patterns of 28.2 to 102.6 ps, in delay, for example, which will result in large timing tolerances or jitter. Decoupling capacitors or device intrinsic capacitances are quite as efficient in reducing delta-I noise. The availability of intrinsic capacitance is, however, harder to predict. While they do not help with the power island problem, they do isolate the chip power distribution from the negative effects of the finite impedance of the package power rails [1]. Power rail breaks have negative impact both on timing and interconnect noise but they will probably increase in number with the multi-core designs. Careful control of at least one power rail in the data-bus path being continuous was shown to be necessary in order to deliver multi-GHz type of operation on such large microprocessor chips.

### References

- [1] A. Deutsch, H. H. Smith, B. J. Rubin, B. L. Krauter, G. V. Kopcsay, "Methodology to Simulate Delta-I Noise Interaction with Interconnect Noise for Wide, On-chip Data-buses Using Lossy transmission-line Power-blocks", Digest of 13<sup>th</sup> IEEE Topical Meeting on Electrical Performance of Electronic Packaging, Oct. 25-27, 2004, Portland, OR, pp. 295-298.
- [2] A. Deutsch, H. H. Smith, G. V. Kopcsay, B. L. Krauter, C. W. Surovic, A. Elfadel, D. J. Widiger, "Understanding Common-Mode Noise on Wide Data-Buses", Digest of 12<sup>th</sup> IEEE Topical Meeting on Electrical Performance of Electronic Packaging, Oct. 27-29, 2003, Princeton, NJ, pp. 309-312.




Fig. 1 Multi-core chip layout with various power islands.

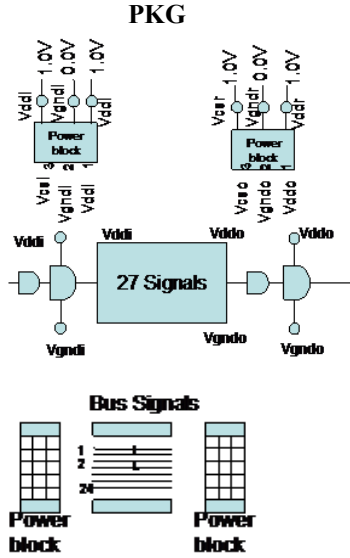


Fig. 2 Circuit diagram used in simulations with two power blocks connected to driver and receiver circuits.




Fig. 3 Simulation results with 6-24 lines with  $R = 229 \Omega/\text{cm}$ ,  $l = 3\text{mm}$ , without power islands, and without decoupling.




Fig. 4 Simulation for 300-ps pulse without power islands and 24 lines with all decoupling included.

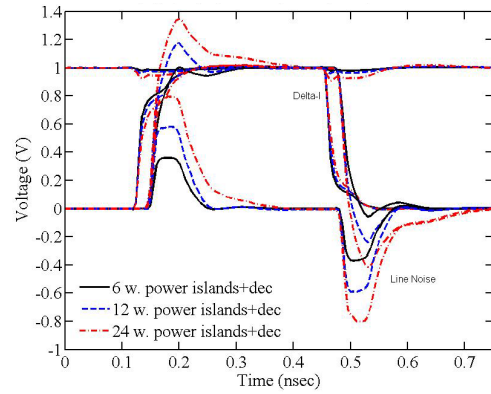


Fig. 5 Simulations for 6-24 lines with power islands and all decoupling included.




Fig. 6 Simulations for 23 lines with power islands and all decoupling included for two extreme switching patterns.




Fig. 7 Simulations for 24 lines, with and without power islands, with decoupling, and two extreme data patterns.

