

Design and Analysis of Power Distribution Networks in PowerPC™ Microprocessors

Abhijit Dharchoudhury, Rajendran Panda, David Blaauw, Ravi Vaidyanathan
Advanced Tools Group, Advanced System Technologies Lab, Motorola, Austin, TX 78730.
email: abhijit_dharchoudhury@email.sps.mot.com

Bogdan Tutuianu and David Bearden
Somerset Design Center, Austin, TX 78730

Abstract

We present a methodology for the design and analysis of power grids in the PowerPC™ microprocessors. The methodology covers the need for power grid analysis across all stages of the design process. A case study showing the application of this methodology to the PowerPC™ 750 microprocessor is presented.

Keywords

power distribution network, PowerPC™, IR-drop, reliability

1. Overview

A robust power distribution network is vital in meeting performance guarantees and ensuring reliable operation of high performance microprocessors. Higher device densities and faster switching frequencies cause large switching currents to flow in the power and ground networks which degrade performance and reliability. Excessive voltage drops in the power grid reduce switching speeds and noise margins of circuits, and inject noise which might lead to functional failures. High average current densities lead to undesirable wearing out of metal wires due to electromigration[1]. Therefore, the challenge in the design of a power distribution network is in achieving excellent voltage regulation at the consumption points notwithstanding the wide fluctuations in power demand across the chip, and to build such a network using minimum area of the metal layers. These issues are prominent in high performance PowerPC™ processors as large amounts of power have to be distributed through a hierarchy of five or six metal layers. For example, in the PowerPC™ 750 processor, the average power dissipation for a nominal V_{dd} of 2.5V is 5W.

The crux of the problem in designing a power grid is that there are many unknowns until the very end of the design cycle. Nevertheless, decisions about the structure, size and layout of the power grid have to be made at very early stages when a large part of the chip design has not even begun. Unfortunately, most commercial tools focus on post-layout verification of the power grid when the entire chip design is complete and detailed information about the parasitics of the power and ground lines and the currents drawn by the transistors are known. Power grid problems revealed at this stage are usually very difficult or expensive to fix. The methodology described in this paper is centered around an initial power grid which is refined progressively as the chip design progresses. Paramount to such a methodology is an analysis tool that has a very

wide choice of models for specifying the grid and the power consumption behavior of the chip. The models must have a range of accuracies that make them suitable for use at any stage of the overall chip design. Such an approach is necessary to identify and remedy any problems with the power grid early on so that final verification results in only minor fixes.

A critical issue in the analysis of power grids is the large size of the network (typically millions of nodes in a state-of-the-art microprocessor). Simulating all the non-linear devices in the chip together with the non-ideal power grid is computationally infeasible. To make the size manageable, the simulation is done in two steps. First, the non-linear devices are simulated assuming perfect supply voltages and the currents drawn by the devices are measured. Next, these devices are modeled as independent time-varying current sources for simulating the power grid and the voltage drops at the transistors are measured. Since voltage drops are typically less than 10% of the power supply voltage, the error incurred by ignoring the interaction between the device currents and the supply voltage is small. By doing these two steps, the power grid analysis problem reduces to solving a linear network which is still quite large. To further reduce the network size, we exploit the hierarchy in the chip design.

Note that the circuit currents are not independent due to signal correlations between blocks. This is addressed by deriving the inputs for individual blocks of the chip from the results of logic simulation using a common set of chip-wide input patterns. An important issue in power grid analysis is to determine what these input patterns should be. For IR-drop analysis, patterns that produce maximum instantaneous currents are required, whereas for electromigration purposes, patterns producing large sustained (average) currents are of interest. Determining either set of patterns is a difficult problem and has been addressed by many, including [2][3][4]. This issue will not be addressed in this paper, and we will assume that suitable input patterns are available for simulation. Moreover, in this paper, we will not discuss electromigration issues, but concentrate on the issue of computing the worst voltage drops due to IR-drop in the power network.

This paper is organized as follows. The next section describes the various analysis modes implemented in our methodology. Section 3 discusses linear solution techniques in the context of power distribution networks. A case study of IR-drop analysis on the PowerPC™ 750 processor is shown in Section 4, and the paper concludes with a discussion of open issues.

2. IR-Drop Analysis Modes

To apply the IR-drop analysis methodology described in this paper across all stages of the design of a complex microprocessor, we define several modes of operation of the tool. These modes are distinguished by different models of the power distribution network and of the currents being drawn by the functional blocks. In this paper, we describe three different *modes* of operation: *early* or *pre-floorplan* mode, *post-floorplan* mode and *post-layout* mode. As the design proceeds, IR-drop analysis is run in these different modes

Permission to make digital/hard copy of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 98, San Francisco, California

(c) 1998 ACM 1-58113-049-x/98/06..\$5.00

using more and more accurate models of the power grid and the block currents.

2.1 Early Mode Analysis

At the very early stages of the design of a microprocessor, there are a number of issues related to the power distribution network that have to be addressed. These include locations of the clean VDD/GND pads, nominal pitches and widths of metal layers, via styles (point or bar vias), and parameters of the chip package. Since at this early stage of the design, the power network has not yet been synthesized and the location and logic content of the blocks are not known, IR-drop analysis is performed using very simplistic models of the grid topology and the block currents.

1. A mock power grid down to the lowest metal layer is constructed using a simple uniform grid topology, where the metal lines in each layer have a user-specified pitch (separation) and width. At the areas where the metal lines of adjacent layers cross over, vias are placed according to user-specified via geometries and via styles. Other topologies such as rings can also be modeled. The clean VDD/GND pads can be placed at the periphery of the chip or on the surface of the chip using C4 pads (for flip-chip packages).
2. To model the currents drawn by the devices, a simple area-based DC estimate of the current is used. This is obtained by taking the current estimate of a previous chip and scaling it by the power supply voltage, operating frequency, complexity, size and technology variables. This estimate is inflated 3-7 times to account for differences between the average and maximum instantaneous currents and to obtain a robust grid. The current sinks are placed on the lines of the lowest metal layer at points midway between adjacent vias that connect the lowest layer to the upper layer. The value of the current is obtained by multiplying the per-unit-area current by the product of the pitches of the two metal layers.

Using simple length-based resistance formulae, a resistive electrical network is constructed from the mock grid topology. DC analysis of this network yields the IR-drops at various locations of the chip. This analysis is very fast and allows the designer to evaluate a large number of different topologies and to trade-off robustness and metal utilization in the power grid. This analysis is used to design the locations of the C4 pads and nominal pitches and widths of the metal layers. Moreover, if the processing technology allows different width and thickness combinations for some of the top metal layers, the user can determine the best values of these in terms of IR-drop. Even though the real power grid will not be as regular as the mock grid, and all the devices will not be drawing the estimated current simultaneously, important design decisions are made from the results of this simple analysis and an early picture of the robustness of the grid is obtained. An example of this analysis is given in Section 4.

2.2 Post-Floorplan Analysis

In this mode, the global power distribution network has been designed and the blocks have been placed. The locations and geometries of the power lines and the blocks are read from the design database. Even though the blocks are placed, the power grids within them have not yet been wired. The *power service terminals* (PST's) of a block are the wires in the topmost metal layer within the block that connect the global and intra-block power networks. In this mode, the PST's for a block may or may not be known - if they are not known, mock PST's are constructed. Next, the block *ports* are determined by the intersection or overlays between the global lines and the block PST's. In our

hierarchy, *blocks* are custom datapath components, synthesized random logic macros (RLM), and off-the-shelf (OTS) components (custom components that can be reused). Custom components are small but RLM's can be large. OTS components can range from small blocks (nands, nors, muxes, etc.) to large blocks (adders, comparators, etc.). A *functional block* (e.g., floating point unit, memory management unit) consists of several instances of custom, synthesized and off-the-shelf components.

Each block current can be independently described in one of the following ways, thus allowing a mixture of them (see Fig. 1).

1. If the logic content of a block is not yet defined, the current model is a DC estimate based on the block area. The total block current is divided equally among all the ports. Since the area-based numbers are calculated such that they reflect peak expected currents, the analysis where every block has area-based currents is likely to be pessimistic since it assumes that each block draws this current simultaneously.
2. The next more accurate block current model is derived from a full-chip gate-level power estimation tool. Given a set of chip-wide input vectors, this tool computes the average power consumed by each block over a cycle. From the average power consumed by a block, an average block current is computed and distributed equally among all of its ports. Hence in this mode, a multi-cycle DC current signature is used for a block. Since chip-wide vectors are used for the simulation, correlation among the blocks is preserved.
3. The most accurate current model comes from a detailed transistor-level simulation of a *functional block* using PowerMill [5]. The input vectors for the functional block are derived from the chip-wide vectors through logic simulation. This ensures that correlation across functional blocks is maintained. The transistor level netlist of the block is available and capacitances are extracted for the signal nets. However, since the power grid within the block has not been designed, it is considered to be ideal. The transient current waveform drawn by each custom, RLM or OTS block within the functional block is obtained from the PowerMill simulation and is divided equally among all of their individual ports. Since the blocks are not very large and the power grid within them have not yet been wired, this block current model is quite accurate for this stage of design.

If all block current models are derived from methods 1 and 2 above, then a resistance-only electrical network is extracted from the geometrical information using length-based resistance formulae. DC analysis is then performed to yield the IR-drop values at each of the block ports (multiple DC analyses if multi-cycle DC current signatures are used). If transient current signatures are used

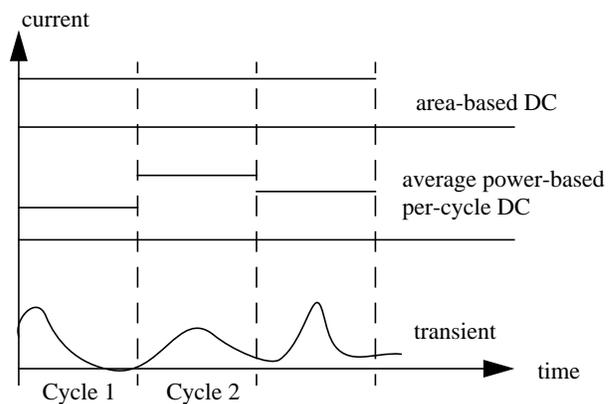


Fig. 1. Block current signature models

for some of the blocks, then an RC network is extracted from the global grid using length-based resistance extraction and statistical rule-based capacitance extraction (since global routing is not done, these statistical rules account for coupling between power and signal lines). An example of IR-drop analysis in this mode is given in Section 4.

2.3 Post-Layout Mode

This mode is used when the global and block-level grids have been completely designed. We employ hierarchical analysis, where each block is analyzed differently based on its size.

- 1. Custom blocks:** We run Railmill [6] on the custom blocks using vectors that are obtained from the common set of chip-wide vectors and clean VDD and GND locations as supplied from the design database. The RailMill analysis a) verifies that the drop in the local grid satisfies the bounds for the block and b) supplies the currents at the block ports that are then promoted to the chip level for global grid analysis.
- 2. Random Logic Macros:** RLM's are analyzed using PowerMill but unlike in Section 2.2, each RLM is broken up into its constituent standard cells. In other words, the standard cells are elevated to the chip-level in the hierarchy. The current drawn by each standard cell is measured separately and these currents are inserted into the ports of the standard cells for global grid analysis. Since the standard cells are small, this mode gives us visibility to the gate level, and the power grid within the RLM is also verified.
- 3. Off-the-shelf components:** OTS's are modeled using PowerMill in the same manner as described in the previous section. Since the OTS's are small and their PST's are at the lowest layer of metal, we get good visibility into the OTS's as well.

In this mode, RC extraction of the global grid is performed using a commercial extraction tool.

3. Linear System Solution Techniques

Several direct and iterative approaches are available [7] to solve linear systems. In this section, we analyze the relative merits and limitations of these methods as applied to solving large power networks. The PowerPC™ processors use 6 layers of metal and the power grid is a very tight mesh. This implies that “crunching” of the global grid does not yield appreciable reduction in the size of the network and DC analysis must resort to conventional matrix methods. The size and structure of the conductance matrix of the power grid is important in determining the type of linear solution technique that should be used. Typically, the power grid contains millions of nodes, but the conductance matrix is sparse (typically, less than 5 entries per row/column). This matrix is also symmetric positive definite, but for a purely resistive network, it may be ill-conditioned. Sparsity favors the use of iterative methods, but convergence is slowed down by ill-conditioning and can be mitigated to some extent by preconditioning. Iterative methods do not suffer from size limitations so long as the (sparse) matrix and some iteration vectors can fit into the memory. The single-biggest problem with direct methods is the need for large amounts of memory to store the factors of the matrix. The number of fill-ins is of the order of $O(N^2)$, where N is the number of rows/columns in the matrix. However, if fixed time steps are used for transient analysis, then the initial factorization can be reused with subsequent current vectors, thus amortizing the large decomposition time. Iterative methods do not have this feature of reusability. Among the direct methods, Cholesky factorization is best suited since the conductance matrix is symmetric and positive definite. However, for a machine with 1 GB of memory, we could

only simulate a network with 300K nodes with this technique. Because of this severe size limitation, we are currently using the conjugate gradient iterative scheme with incomplete Cholesky preconditioning to solve our matrices.

As mentioned earlier, a flat analysis of the entire power distribution network where each transistor that connects to the power lines is modeled as a current source would be computationally infeasible due to the large size. This size limitation can be mitigated by hierarchical analysis, in which each block has a macromodel which is used for the analysis of the global grid. A block macromodel consists of current sources at the block ports and an admittance matrix relating the currents and the voltages at the ports. For an exact equivalence with flat analysis, the admittance between every pair of ports must be modeled resulting in a dense admittance matrix for each block. However, this adversely affects the speed of an iterative method, reducing the efficiency gained by a hierarchical approach. To preserve the efficiency of the global analysis, we ignore the admittance between the ports of a block and model the blocks purely as current sources. For chip-level analysis, the error due to this assumption can be kept within bounds if the blocks are small. In our methodology, the blocks used for global power grid analysis are OTS's, standard cells and custom elements, all of which are small. An important mechanism implicit in this hierarchical analysis is that each block has an IR-drop budget, i.e., a maximum allowed voltage drop at its ports. If the voltage at the port of a block violates its budget, then the problem is fixed by making the power grid that supplies this block more robust.

4. Case Study

In this section, we provide examples of the application of the various stages of the IR-drop analysis methodology described above. As mentioned in Sec 2.1, the mock global grid in the early mode analysis consists of uniformly spaced power lines and clean VDD sources (C4 pads). Since the analysis will result in a symmetric voltage distribution, we confine the analysis to a representative area of the chip. The mock global grid is shown in Fig. 2(a): the uniformly spaced lines and C4 pads are clearly seen. The current sources are inserted at the lowest metal layer midway between the vias connecting it to the upper layer. We assume that the total current drawn by the chip is 6A, which makes the power of the chip 15W (three times the average power of 5W). By dividing the total current by the estimated area of the chip, we get the current per unit area. The voltage map using this current value is shown in Fig. 2(b). As expected, the worst voltage drop occurs at the points with the largest distances from the C4 pads; for this analysis, the worst voltage drop was measured to be 70mV. Now, if the pitch and width of the second layer of metal is halved (keeping the same amount of power grid metal utilization), the voltage drop reduces to 45mV. However, the trade-off with the reduced voltage drop is that the congestion for routing signal lines will be significantly more. This example illustrates the type of trade-offs that can be explored using the early analysis mode.

Fig 3(a) shows the placement of the blocks and the power grid in the PPC750 processor - there are approximately 15000 blocks. As mentioned earlier, these blocks consist of customs, RLM's and OTS's. The floating point unit (FPU) functional block is located at the bottom right corner of the chip, and is shown highlighted. The first run of the post-floorplan analysis mode was done with area-based DC current estimates for each of the blocks in the chip. The per-unit area current estimate was modified to yield a less pessimistic value of 4.5A for the total chip current. The voltage map for this run is shown in Fig. 3(b) - the worst drop in this case was mea-

sured to be 170 mV and was occurring in the IO pads near the left edge of the chip. This value is worse than the worst drop predicted by the early analysis even though the average pitches and widths of the metal layers are similar in both analyses. This is because the real grid is not as uniform as the mock grid and the “distance” from the nearest C4 pad is more for the IO pads. However, within the FPU, the worst drop was measured to be 110mV near the bottom right corner of the chip. A close-up of the voltage map within the FPU is shown in Fig. 3(c). The sum of the area-based DC currents for all of the blocks within the FPU functional unit was computed to be 0.7A.

For reasons of brevity, we will not show the results of IR-drop analysis when the area-based current estimate of the blocks is replaced by current estimates based on full-chip multi-cycle power simulation. Instead, we will show the results with transient current signatures obtained from PowerMill and also concentrate on the FPU area of the chip. As discussed earlier, PowerMill simulation is done on the transistor level netlist of all the blocks within the FPU with extracted parasitics. The input vectors for the FPU are obtained from the results of a full-chip logic simulation, using “high stress” chip-wide vectors. A subset of the total current waveform of the FPU is shown in Fig. 3(d). The first current spike peaks at approximately 2.1A, whereas the second spike peaks at approximately 0.7A. This shows that the area-based estimates agree well with the typical peak values (0.7A); however, there are input patterns which cause the peak values to be significantly higher as in the case of the first spike (2.1A). Since the sum of the FPU currents using the area-based estimate was also 0.7A, we expect that the worst voltage drop within the FPU for this time point will be comparable to the area-based analysis. In fact, the worst drop within the FPU in this case was measured to be 130mV. However, as shown in Fig. 3(e), the voltage map is different from the area-based analysis. In fact, the worst drop region is much more localized in this analysis. This is due to the non-uniform spatial distribution of the block currents, i.e., at this time point, the blocks near the worst drop region are drawing much larger currents compared to the other blocks. Fig. 3(f) shows the voltage map for the analysis at the time point when the total FPU current is 2.1A. In this case, the worst drop within the FPU is 300mV and the worst drop region has shifted upward. This is due to the fact that at this time instant, the locations of the high current blocks are different from the previous case.

This analysis shows that because of the safety margins built into the early and post-floorplan modes, the typical voltage drops in the final verification stage are well within the block budgets. However, as seen in this example, transient analysis may reveal cases when the worst drop is worse than expected. In these cases, the voltage drops can be improved by making local alterations to the power grid.

5. Conclusions

We presented a coherent design and analysis flow for designing the power grids of large, high performance processors. The usefulness of the multi-mode analysis capability in progressively refining the design of a power grid design was demonstrated through case studies. Several issues related to power grid analysis were discussed and a case study of the PPC750 processor was presented.

Areas of future work involve investigating fast solution techniques for power and ground networks, calibrating and reducing the error due to hierarchical analysis, incorporating inductance effects into the IR-drop analysis, and determining optimal locations of decoupling capacitors.

References

- [1] Black, J.R. Electromigration failure modes in aluminum metalization for semiconductor devices. *Proc. IEEE*, pp. 1587-1594, Sept 1969.
- [2] S. Chowdhry et al. Estimation of maximum currents in MOS IC logic circuits. *IEEE Trans. CAD*, pp. 642-654, June 1990.
- [3] H. Kriplani et al. Pattern independent maximum current estimation in power and ground buses of CMOS VLSI circuits: algorithms, signal correlations, and their resolution. *IEEE Trans. CAD*, pp. 998-1012, Aug 1995.
- [4] A. Krstic et al. Vector generation for maximum instantaneous current through supply lines for CMOS circuits. *Proc. 34th DAC*, pp. 383-388, June 1997.
- [5] *PowerMill User Guide*, Synopsys Inc., 1997.
- [6] *RailMill User Guide*, Synopsys Inc., 1997.
- [7] Golub, G and Van Loan, C. *Matrix Computations*. The Johns Hopkins Univ. Press, 1989.

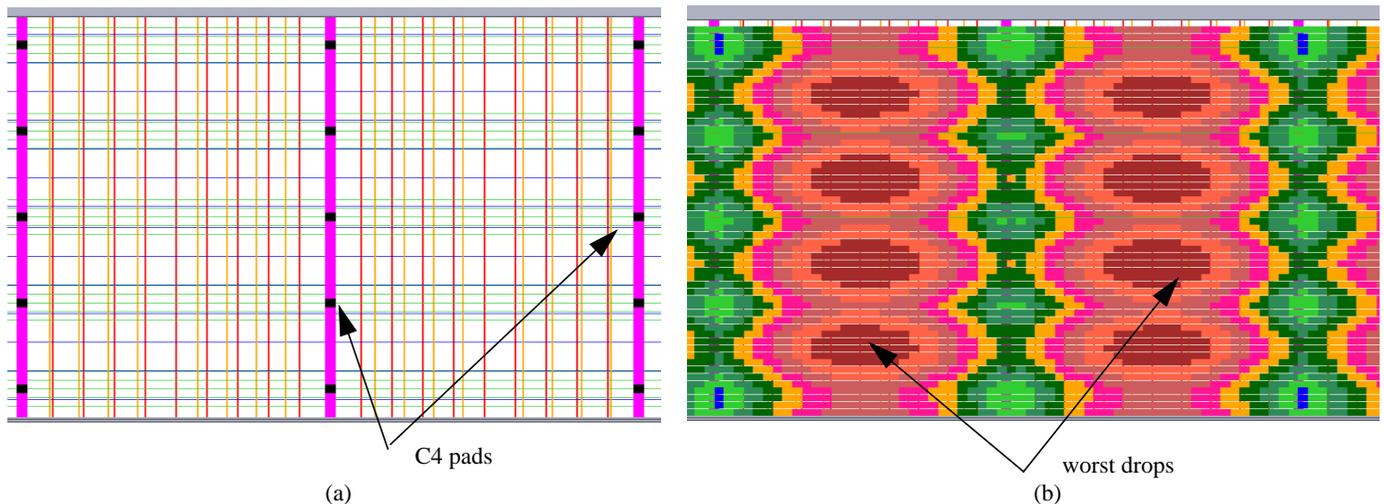
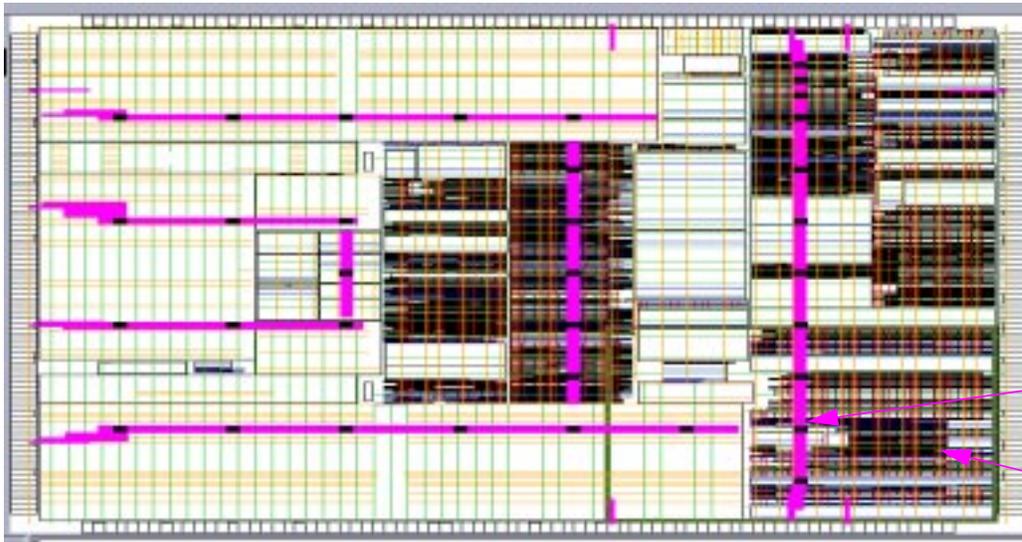


Fig. 2: Early mode analysis of PPC750: (a) mock grid, (b) corresponding voltage map.

(a)

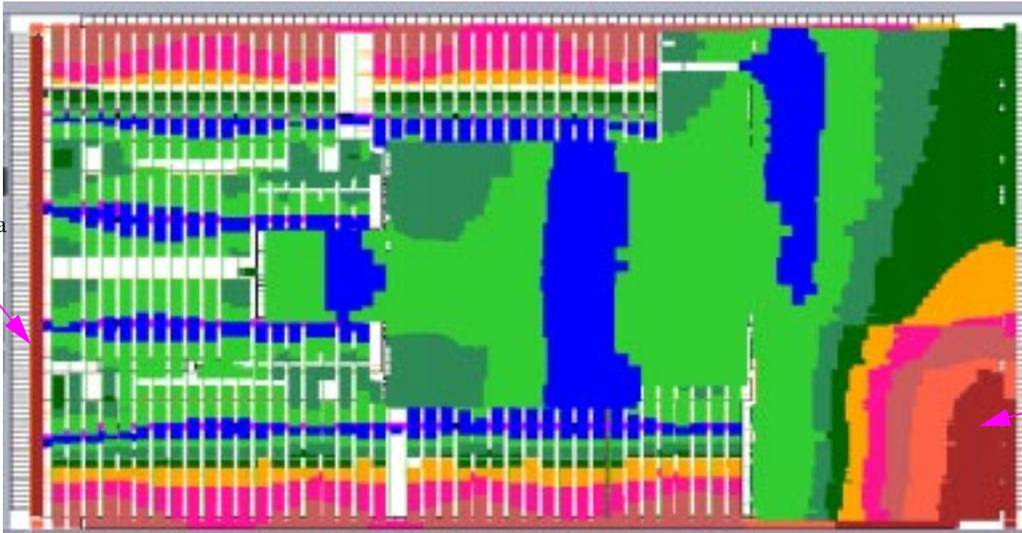


C4 pad
inside FPU

FPU

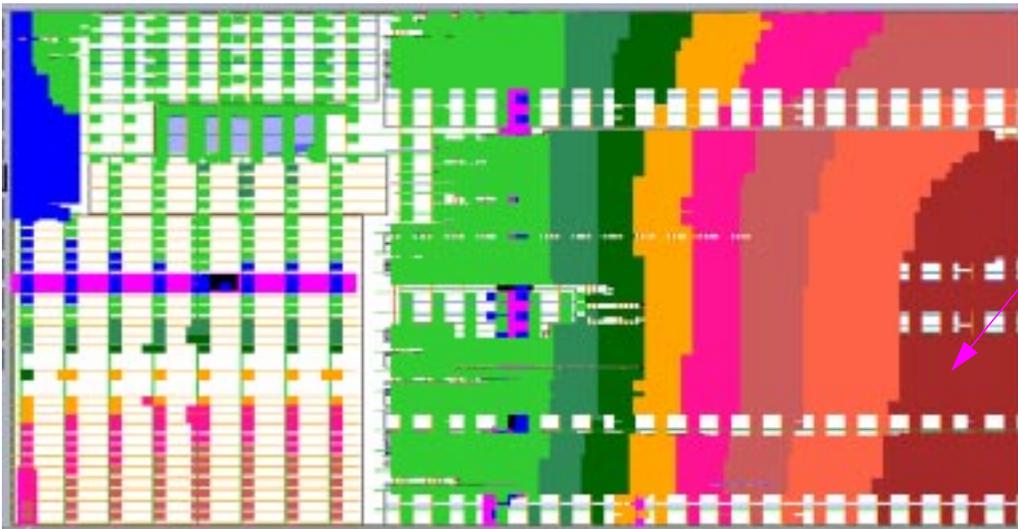
worst
drop area

(b)



worst
drop area

(c)



worst
drop area

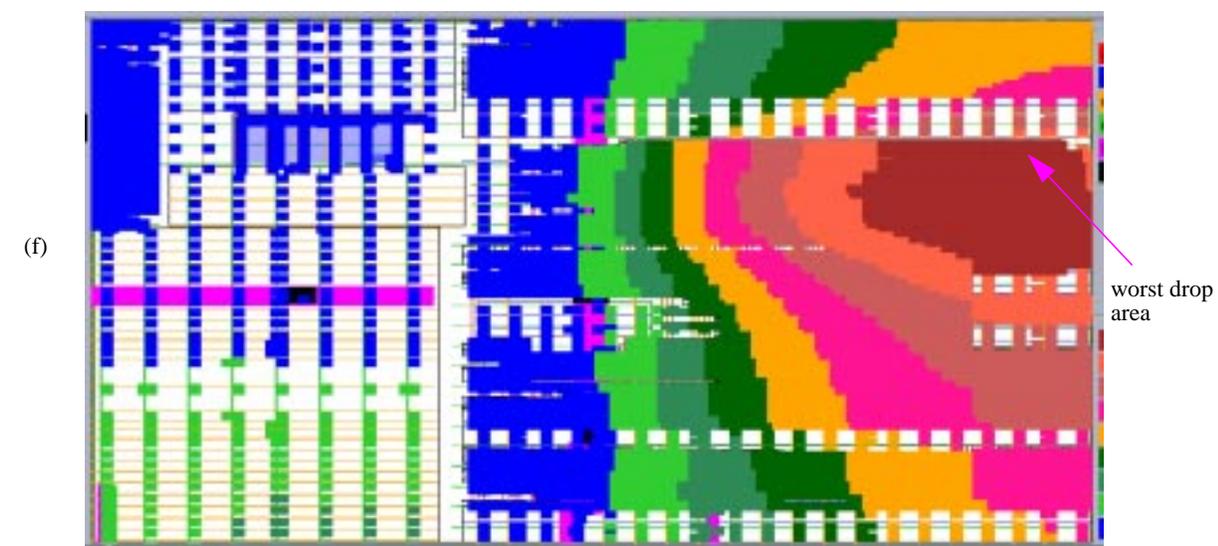
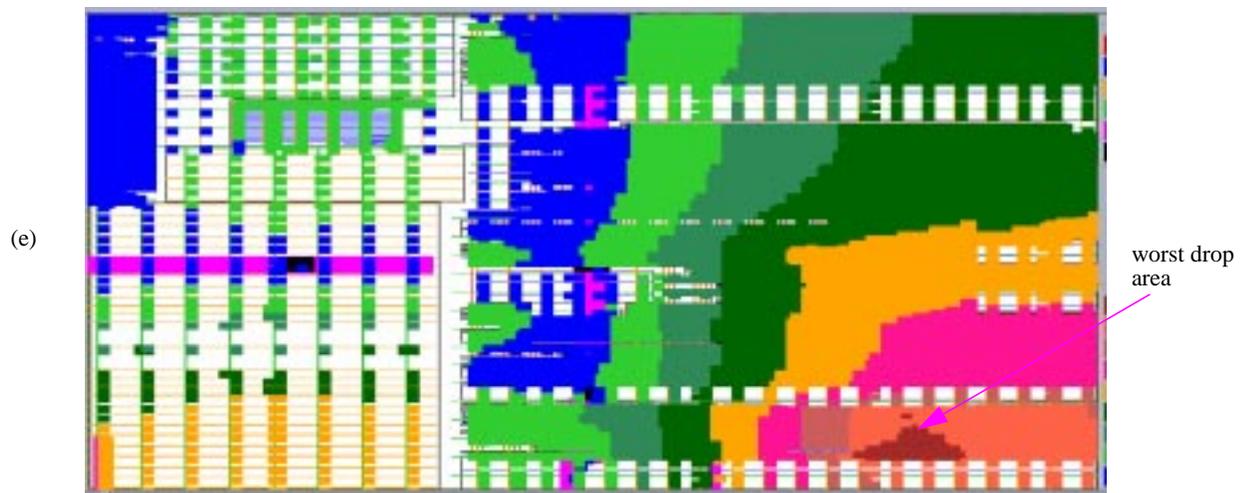
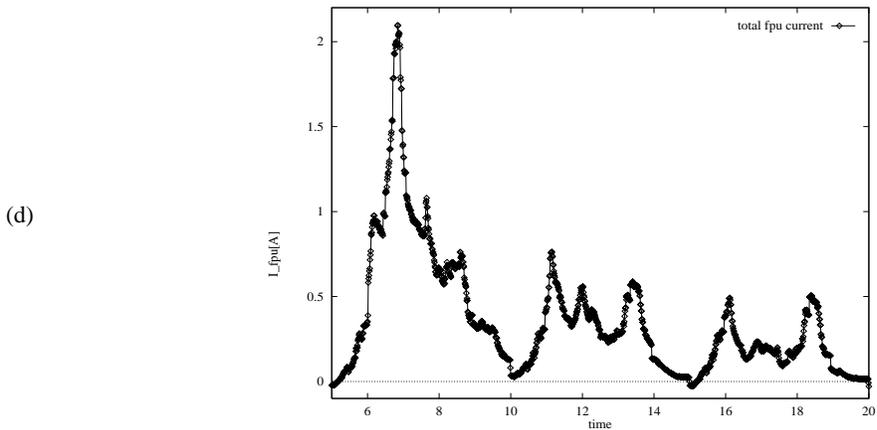


Fig. 3. PPC™ 750 case study: (a) power grid and block layout, (b) voltage map for area-based block currents, (c) voltage map inside the FPU for area-based block currents, (d) total transient currents in the FPU, (e) voltage map inside the FPU for total current of 0.7A, (f) voltage map inside the FPU for total current of 2.1A.