

Power Distribution Techniques for Dual VDD Circuits

Sarvesh H. Kulkarni

EECS Department, University of Michigan
Ann Arbor, MI 48109, USA
shkulkar@eecs.umich.edu

Dennis Sylvester

EECS Department, University of Michigan
Ann Arbor, MI 48109, USA
dennis@eecs.umich.edu

ABSTRACT

Extensive research has proposed the use of multiple on-die power supplies (VDD) for reducing power consumption in CMOS circuits. We present a detailed study and design techniques for power delivery systems in dual VDD CMOS circuits. We first show that the total current to be delivered by the voltage supplies is significantly reduced (by 27%–46%) in dual VDD circuits. This current reduction prompts various design strategies that can be employed to design the power delivery system. We describe issues that arise at the system, board and package levels and propose a high-level model for the same. We then provide a new placement driven approach for designing on-die dual VDD power grids. Compared to already existing methods, the dual VDD grids generated by our approach reduce the worst case and average voltage drop by up to 12.3% and 6.8% respectively with no area overhead and sometimes improving wire congestion. We also show that dual VDD circuits can afford lower on-die decoupling capacitance budgets.

1. INTRODUCTION

Reducing power consumption at high speed is a critical goal for VLSI designers today. The dynamic and static power consumption in CMOS circuits have a quadratic and roughly cubic dependence on the power supply voltage (VDD) respectively [1]. There is extensive work in literature that exploits this concept for reducing power by using dual (or in general, multiple) power supplies in the design. Multiple VDD design applies higher voltages to gates on critical paths and lower voltages to gates on paths with slack. In this way the power consumption reduces while timing is met. Most earlier work in this area focuses on the power supply assignment problem. In particular, [2]–[8] provide algorithms that select gates to be assigned to the available supplies. Ref. [8] and recently [9] and [10], detailed designs based on multiple power supplies.

Ref. [12] shows that using two different supplies provides near optimal power savings and adding a third supply yields little additional power reduction while further worsening the power delivery challenges. We therefore focus on dual VDD designs in this work. Also, voltage assignment can be performed at a fine-grained level (gate-level) or at the module-level. The module-level assignment somewhat simplifies the physical design and power delivery problems; however, its power savings are curtailed due to less freedom in low VDD assignment. Hence, we focus on fine-grained dual VDD assignment and also assume a standard cell row-based layout style that is commonly used in ASICs. In the remainder of this paper we refer to the lower supply in a dual VDD design as VDDL and to the higher supply as VDDH. The power supply of a reference single VDD design will simply be referred to as VDD.

Issues such as level conversion and physical design for dual VDD also arise when using this technique. Ref. [11] first proposed techniques to enable the physical design of standard cell based dual VDD circuits. The focus was on developing a new placement tool & proposing a new cell layout style that facilitated dual supply usage.

Although some dual VDD issues have received greater attention, the important issue of power delivery discussed in our paper has received little attention. Related work in this area includes approaches presented in [8], [11], [13]–[15]. Since two power supplies need to be now supplied across the die, [8] and [13] proposed physical design approaches to partition cells into regions

(rows/blocks) that contain either only VDDL cells or only VDDH cells. This allows existing standard cell libraries to be used but needs a modified standard cell placer that leads to high wirelength and core area overheads. Ref. [11] and [14] thus proposed the use of a modified standard cell layout where an extra rail is added for the added supply voltage. The existing placer can then be freely used with its full optimization power. In all approaches mentioned so far, VDDL and VDDH cells share a common ground (GND). Recently, [15] proposed the use of a separate ground GNDL (for VDDL cells) and GNDH (for VDDH cells) respectively.

Before introducing our approach for designing dual VDD power grids, we first make the observation that the power supply current required for the operation of dual VDD circuits is greatly reduced compared to a single VDD circuit with the same timing. This is so for two reasons: (i) gates assigned to VDDL need to deliver less current to charge up their load capacitances to logic state 1, and (ii) current demand on the VDDH power supply reduces since only a subset of the initial gates now draw current from it. In [6] the authors suggest that as high as 60% of total gates are assigned to the lower supply for stringent timing constraints, strengthening the claim that current drawn from the VDDH power supply should fall substantially. We use this observation and show that this quality can be harnessed to design robust power distribution systems for dual VDD circuits. Power grid design approaches should take into account the actual placement of the VDDH and VDDL cells while sizing the grid wires, i.e., if a region of the die contains more VDDH cells, more wiring resources should be dedicated to the VDDH grid, while recovering resources from the VDDL grid. The approaches in literature failed to take the placement into account. Using such ideas we show that dual VDD grids can be design to be as robust as their single VDD counterparts for no area or wire congestion penalties. Interestingly, we also show that dual VDD circuits can afford reduced decoupling capacitance due to their reduced current demand. This reduction in decoupling capacitance will improve leakage, die area and yield.

To summarize, the main contributions of this paper are:

1. We present the first detailed study of power distribution for dual VDD circuits. We explore solutions for package/board level issues as well as issues for on-die power grids.
2. We present a new placement driven power grid design methodology, *D-Place*, which improves power grid integrity.
3. We demonstrate that dual VDD power grids can be designed to be as robust as their single VDD counterparts. In fact, we show that dual VDD designs can afford lower decoupling capacitance budgets.

Our paper is organized as follows. In Section 2, we describe our simulation setup and general framework. In Section 3, we demonstrate that dual VDD circuits have significantly lower supply current demands. In Section 4, we describe a study of the system board and package level issues when working with dual supplies. Section 5 presents our work for on-die power distribution grids. In Section 6 we show that dual VDD circuits can often afford lower decoupling capacitance budgets. Section 7 concludes the paper.

2. SIMULATION SETUP AND FRAMEWORK

Our work is based upon a 6 metal layer 0.13 μ m CMOS technology. The nominal voltage for this technology is 1.2V and two threshold voltages (V_{TH}) are available; 0.2V/0.1V and -0.2V/-0.1V

Table I – Dual VDD power savings and VDDL assignments
(the reported % are with respect to the original single VDD design)

	VDDL = 0.8V		VDDL = 0.6V	
	% Savings	%VDDL	% Savings	%VDDL
c880	28	65	31	55
c2670	32	65	37	56
c5315	35	58	37	49
c7552	44	91	49	71

for NMOS and PMOS respectively. Gates in the original single VDD design are sized based on an algorithm similar to TILOS [16]. We developed a static timing analysis engine and use look-up table based power/delay standard cell libraries for timing analysis and power estimation (Synopsys Library Compiler format).

For obtaining the dual VDD design from the single VDD design, we adopted the method outlined by the authors in [6] because of its simplicity in implementation; the concepts outlined in this paper are applicable to all other dual VDD assignment algorithms (such as [2]-[5], [7] or [8]). We extended the work in [6] by adding sensitivity based dual VTH assignment since dual VTH assignment is widely used in practice for further optimizing power. The details of the power optimization flow itself have been omitted due to lack of space. Literature suggests that a VDDL value that is about 50–70% of VDDH is ideal for minimizing power [17], [18]. So, for this work, we tested our algorithms for VDDL = 0.6V and 0.8V. Table I summarizes the power savings obtained for several ISCAS85 benchmark circuits [19]. The column marked “%VDDL” indicates the fraction of gates from the original single VDD design that are mapped to VDDL. The single and dual VDD designs meet the same stringent timing constraint. These results confirm that significant power savings are possible using dual VDD design and that a significant number of gates get assigned to the lower supply.

3. DUAL VDD SUPPLY CURRENT DEMAND

This section demonstrates that dual VDD circuits have significantly reduced power supply current demands. Since we use a dual VTH dual VDD process, each gate in the design can be one of four combinations, namely VDDH-low VTH, VDDH-high VTH, VDDL-low VTH and VDDL-high VTH. As a cell moves from VDDH to VDDL, it has a significantly reduced current demand. A subtle point however, is that short-circuit current that flows during switching events is also significantly reduced by such a VDD change; this also holds when a gate moves from low to high VTH. Table II summarizes these reductions in current demands for a few gates in the library. The average over all 160 cells in the library is also reported. This data was obtained through transient simulations in SPICE and thus accurately includes the changes in short-circuit current and currents charging the load.

We next extend this concept from the gate level to the circuit level. The circuit-level current demands follow the gate-level numbers presented in the previous table. Using SPICE simulations over 1000 randomly selected input vectors, Table III reports the current load on each of the power supplies (VDDL and VDDH). These results confirm that the current to be supplied by the power supplies is significantly reduced in dual VDD designs.

Table II – Normalized gate-level supply current reduction

	Single VDD		Dual VDD: VDDL=0.8V		Dual VDD: VDDL=0.6V	
	Low VTH	High VTH	Low VTH	High VTH	Low VTH	High VTH
INVX10	1.00	0.90	0.57	0.49	0.36	0.27
NAND2X2	1.00	0.85	0.54	0.45	0.34	0.23
NAND3X6	1.00	0.88	0.55	0.47	0.35	0.24
NOR2X1	1.00	0.86	0.52	0.39	0.30	0.19
NOR3X4	1.00	0.85	0.50	0.37	0.29	0.18
AVERAGE	1.00	0.88	0.54	0.44	0.33	0.23

Table III – Circuit-level current (mA) drawn from the power supplies

	Single VDD		Dual VDD: VDDL=0.8V		Dual VDD: VDDL=0.6V	
	VDD	VDDH	VDDL	VDDH	VDDL	VDDH
c880	9.7	5.6	2.2	5.9	1.3	1.3
c2670	23.6	11.9	6.5	10.1	3.0	3.0
c5315	36.7	20.9	7.2	20.9	3.6	3.6
c7552	47.9	13.9	19.4	20.4	8.5	8.5
AVERAGE %	100.0	48.5	27.7	50.7	13.5	13.5




Fig. 1. Single VDD power delivery model [20].




Fig. 2. Dual VDD power delivery model.

4. SYSTEM BOARD/PACKAGE DESIGN

4.1 Single VDD board

Fig. 1 shows a high level power delivery model for the system board & package of a typical integrated circuit (IC) [20]. The voltage regulator module (VRM) is situated on the motherboard and is the primary voltage source. Capacitors Cblk/Chf, Cpkg_cap and Cdie are area motherboard, package and on-die decoupling capacitors respectively. The resistances and inductances in series with the decoupling capacitors model the effective parasitic series resistance and inductance. Lmb1-Rmb1/Lmb2-Rmb2 and Lpkt-Rpkt represent the inductance and resistance of signal tracks on the motherboard and cables connecting the motherboard to the package respectively. Lpkg and Rpkg represent the inductance & resistance of the package/C4s. The values for the parameters used in this figure are from [20].

4.2 Dual VDD board

The single VDD model in Fig. 1 is extended for the dual VDD case as shown in Fig. 2. Two separate VRMs are now needed (for VDDH and VDDL). We see that this model provides two on-chip voltages (VDDH and VDDL) at nodes 3 and 2 respectively. The ground can be shared between the two voltages (node 1). In this model, the current loop shown in the upper half of the figure corresponds to that seen by the VDDH VRM and the lower one corresponds to the VDDL VRM. The current loads for each of the power supplies are also shown. As demonstrated in Section 3, each of these individual current sources is a fraction of the original current source shown in Fig. 1. On average across the benchmarks we studied, we found that the VDDH and VDDL current loads are about 49% and 28% of the original current load (VDDL = 0.8V). Since the ground path is shared, current flowing through the ground path is about 77% of that in the original design. These numbers reduce further when VDDL = 0.6V. The parameters corresponding to the motherboard will not change since similar track/planes will be still be followed by each of the currents. The socket parameter (Rskt and Lpkt) are also kept unchanged since they model the cables connecting the motherboard to the package. The package inductance & resistance (Lpkg and Rpkg), depend on the allocation of the C4s (assumed equal in number as the single VDD case). The C4s should be split equally between VDDL and VDDH, since although fewer gates remain assigned to VDDH (Table I), the current demand on VDDH still remains much more than on VDDL (Table III). Keeping the same number of C4s for the ground, these package parameters for VDDL and VDDH will be in *inverse* proportion to the fraction (=0.5) of C4s assigned to each supply.

Assuming that the dual VDD board/package allows the same real

Table IV. Power delivery model simulations.

		VDD or VDDH & VDDL		GND		
				PK	QS	PK
Single VDD	VDD	mV	92.7	65.0	92.7	65.0
		%	7.7	5.4	7.7	5.4
Dual VDD	VDDH	mV	63.0	34.0	68.9	40.7
		%	5.3	2.8	5.7	3.4
VDDL = 0.6V	VDDL	mV	18.0	9.0	68.9	40.7
		%	3.0	1.5	11.5	6.8
Dual VDD	VDDH	mV	63.0	32.0	77.8	46.0
		%	5.3	2.7	6.5	3.8
VDDL = 0.8V	VDDL	mV	37.0	18.0	77.8	46.0
		%	4.6	2.3	9.7	5.7

estate area for decoupling capacitors (as the single VDD board/package), we propose splitting the capacitances in the ratio of the current loads. The reduced current demand in dual VDD circuits proves useful here since it becomes feasible to use lesser decoupling capacitors for each of VDDL and VDDH as compared to the single VDD case. Thus, for C_{blk} (and similarly for others), we have,

$$C_{blk} = C_{blk_H} + C_{blk_L}, \text{ and} \quad (1)$$

$$C_{blk_H} / C_{blk_L} = I(VDDH) / I(VDDL)$$

where, C_{blk_H}/C_{blk_L} are the decoupling capacitances attached to VDDH/VDDL, and $I(VDDH)/I(VDDL)$ are the current demands on VDDH/VDDL respectively.

4.3 Results

Table IV reports the peak (PK) droop/bounce voltages at nodes 1, 2 and 3. The quiescent values (QS) of the voltages at these nodes are also reported and correspond to the resistive drop from the VRM up to the C4s. Absolute values in mV as well as percentages of total nominal swing (0.6V/0.8V for VDDL, and 1.2V for VDDH) are reported. The percentages reported here carry more relevance since they are essentially the fraction of the total nominal swing of the voltage supply that the on-die circuitry actually receives from the board. This percentage directly translates into delay degradation in the standard cells and is hence a good metric to follow. From Table IV, the absolute values of the droop/bounce for the dual VDD case are always better than the single VDD case. Based on a percentage metric as well, the dual VDD case does better than the single VDD case, except for the ground bounce afflicting the VDDL cells (boxes shaded gray). This subtle difference arises since, although the actual ground bounce in mV has reduced, it remains a large percentage when normalizing to VDDL (0.6V or 0.8V). The reason for this behavior lies solely in the fact that the ground path is shared by both the supplies. This points to the fact that the dual VDD board may require special care for the ground path (such as a reduced resistivity return path for ground) as compared to the single VDD board.

We also used the HSPICE circuit optimizer to optimize this circuit for minimum PK/QS values. We found that the results using the simple intuitive approach outlined above gave us results that match the optimal solution. However, from this study we also found that PK/QS are fairly insensitive to the exact ratio in which parameters such as capacitors or C4s are split among VDDL and VDDH. This is a favorable finding, since it gives the board/package designer considerable flexibility in allocating the resources to each of the two power supplies. This can also help in designing for desired values of resonance frequencies between the package inductance and on-die decoupling capacitance.

5. DUAL VDD POWER GRID DESIGN

5.1 Framework

Our process technology has 6 metal (Cu) layers and has flip-chip package technology. The standard cells are placed in rows on the bottommost metal layer and are represented as current sources. We assume a partial electrical equivalent circuit (PEEC) model for the grid [21]. In this model, each line of each metal layer is fractured into smaller segments, and each segment is then modeled using a resistor, self inductance, mutual inductance to other segments and ground and




Fig. 3. PEEC model of power grid.

coupling capacitances. We used the following methods in calculating the parameters for each segment: (a) Resistance: is calculated simply once the length, width and sheet resistance are known. (b) Capacitance: is calculated using Wong's model [22]. Our estimation of decoupling capacitance to be added is described below. (c) Inductance: Partial self and mutual inductance depend on the geometry of the wire segments and are calculated using the method described in [23].

For simulating the grid we follow the fast and accurate R/L/C simulation method presented in [24]. The PEEC models of the grids we worked on typically have about 600,000 R/L/C elements. On a 3GHz Pentium4 2GB RAM computer, our implementation takes about 160s to generate the PEEC model and about 40s for simulation (single as well as dual VDD grids). While the PEEC model can become computationally expensive for large die sizes, we have used it for its suitability in modeling on-die power grids - the ideas presented by us can also be applied to simpler R/L/C models such as those obtained through FastHenry for the price of accuracy.

Fig. 3 shows an example of the model for the bottommost layer of the grid. The current sources shown represent the gates of the design. The vias connecting the various metal layers are modeled as resistors. The resistances and the inductances of the wire segments and some decoupling capacitors are also shown in the figure.

5.2 Single VDD grid design

Single VDD grids were held as the reference for comparison with the dual VDD grids. The single VDD grid is assumed to be regular in structure. The VDD and ground lines alternate each other and have increasing thicknesses as we move higher up in the metal layer hierarchy. The C4 diameter and pitch was assumed to be 30 μ m and 150 μ m respectively. The C4s are placed on 30 μ m wide straps on the topmost metal layer. Alternate straps for VDD and GND arrange all the C4s in a checker-board fashion such that the pitch of 150 μ m is met. We simulate grids that are $\sim 0.5\text{mm}^2$ in area and allow for 24 C4 locations (12 for VDD and 12 for GND).

Decoupling capacitors help mitigate voltage drop in the grid. We follow a simple method described in [25] to estimate the decoupling capacitance needed. We first fix a tolerance level for the permitted voltage drop on the grid ($V_{noise-limit}$). The decoupling capacitance is then estimated using Eq. 2. While this estimate appears simple, it is commonly used in practice [25]. Approaches such as [26] can be used alternatively.

$$C_{decap} = \frac{\int_0^{\tau} I(t) dt}{V_{noise-limit}} \quad (2)$$

where, $I(t)$ represents the current sunk by the switching events and τ is the switching period.

5.3 Dual VDD grid design

We first point out that we have constrained the areas of all grids studied to be the same. Since dual VDD grids need to route an additional voltage, this may imply worsened wire congestion. Thus under constant area, we finally compare the quality of the grids with respect to the voltage droop/bounce as well as wire congestion. We also assume the same number of C4s as in the single VDD case & distribute them equally between VDDH/VDDL (for the same reason




Fig. 4. Different standard cell layouts.

as described in Section 4). Since, the VDDL and VDDH cells are evenly dispersed across the die, the VDDL and VDDH C4s are also evenly dispersed. Our C4 assignment scheme is in keeping with [27] where a strong spatial locality effect is shown. The values of the current sources connected to model the gates are obtained using the analysis described in Section 3.

As mentioned briefly in Section 1, [8] and [13] proposed partitioning cells into blocks or rows, such that each block or row contained only one kind (VDDL or VDDH) of cells. This allows existing standard cell libraries to be used. However, [11] and [14] showed that this constrains the placement tool and can lead to an increase in wire length as well as core area (by up to 23% and 15% respectively) and higher post-route power dissipation. To overcome this problem, modified standard cells shown in Fig. 4 can be used. We now describe three methods (referred to as *D-Vanilla*, *DSDG* and *D-Place*) for designing dual VDD power distribution grids. We compare our approach *D-Place* against *D-Vanilla* and *DSDG*, since all these techniques allow the use of the same placement tool as used for the single VDD design (thus enabling a fair comparison).

5.3.1 D-Vanilla:

Authors in [11] proposed the 3-rail standard cell layout shown in Fig. 4(ii). (Fig. 4(i) shows a conventional 2-rail single VDD cell.) The new cell library now has two copies of each cell from the old library: one copy is powered from VDDL and the other is powered from VDDH (when using a dual VTH process there will be two more copies of each cell for low VTH and high VTH). The existing placement tool can now be used with its full optimization power since there are no constraints about where to place each kind of cell. The authors in [11] observed that since the current requirement of dual VDD circuits is reduced, the VDDL rail can be reduced in width compared to the VDDH rail, mitigating wire congestion. Thus, the rails widths for GND and VDDH are kept the same as the single VDD cell, and the rail width of the VDDL rail is scaled down by the ratio of current demand of the design when powered by VDDL to the current demand when powered by VDDH. This ratio is quite design invariant and is about 0.32 for VDDL=0.6V and 0.54 for VDDL=0.8V. Grids designed using this work look like Fig. 5(ii).

5.3.2 DSDG (Dual Supply Dual Ground) [15]:

D-Vanilla discussed above shares ground between VDDL and VDDH cells. *DSDG* proposes the use of separated grounds (GNDL and GNDH) for VDDL and VDDH. Every alternate rail (from the single VDD floorplan) is now assigned to VDDL/GNDL and VDDH/GNDH. This work however did not discuss the standard cell layout that should be used. We hence propose the 4-rail standard cell layout shown in Fig. 4(iii). Each of the rails in this layout (VDDH, VDDL, GNDH and GNDL) is now half as wide as the rails in the cell shown in Fig. 4(i). Not shrinking each of the rails this way would lead to very high wire congestion for fixed area (or about 23% higher area for the same wire congestion in our studies). Grids designed using this work look like Fig. 5(iii). While the authors used simplified FastHenry based models, our analysis is more accurate since we use the PEEC model. Also, effects such as the sharing of C4s among VDDL and VDDH are considered in our implementation.

5.3.3 D-Place:

Before moving on to describing our placement driven approach, we list important points that the earlier approaches did not consider:




Fig. 5. Texture of grids designed using cell layouts in Fig. 4.

- Firstly, it is important to control voltage droop/bounce on the VDDL grid to a value that is lower than that for the VDDH grid (e.g., to ensure that <10% of rail-to-rail swing is lost due to grid losses, we need to limit droop on the VDDH (1.2V) grid to 120mV; however, for the VDDL (say 0.6V) grid, we need to limit it to only 60mV).

- Secondly, the method followed in sizing wires of each grid should take into account how much current needs to be delivered.

- Thirdly, the method followed in sizing wires should also consider the placement of the two kinds of cells.

- Fourthly, when designing the power distribution system, effects arising at the system board and package level should also be accounted for. While our work addressed these issues in Section 4, earlier work such as [15] failed to do so. This becomes especially important since [15] requires two separate grounds (GNDL and GNDH) to be supplied from outside the chip which can greatly complicate the design at the system board and package level.

The details of *D-Place* are as follows. We use 3-rail standard cells shown in Fig. 4(ii). Conventional placement tools can thus be used for the layout of the dual VDD design and the placement of the dual VDD cells is very close to the placement of the cells in the original single VDD design. The ground is common for VDDL and VDDH. Let α and β respectively be the ratios of the current demands on the VDDH and VDDL grids to the current demand of the single VDD grid, i.e., $\alpha = I(\text{VDDH})/I(\text{VDD})$ and $\beta = I(\text{VDDL})/I(\text{VDD})$. The grid droop/bounce is a result of two mechanisms: a resistive IR drop and an inductive Ldi/Dt drop, where the IR drop usually dominates the on-chip inductive effect. Let 'W' μm be the wire width that was assigned to some wire in the VDD/GND grid in the original single VDD design. We now assign the widths to corresponding VDDH, VDDL and GND wires in the dual VDD grids as follows:

$$W_{\text{VDDH}} = \alpha W \quad (3)$$

$$W_{\text{VDDL}} = \beta \frac{V_{\text{DDH}}}{V_{\text{DDL}}} W$$

$$W_{\text{GND}} = (\alpha + \beta) \frac{V_{\text{DDH}}}{V_{\text{DDL}}} W$$

The reasoning behind this sizing becomes clear when we recall our earlier comment that droop/bounce on VDDL grids needs to be very well controlled. Assuming that the original single VDD grid met a certain percentage droop/bounce budget, the VDDH grid will meet the same budget with wire widths scaled down by α . This is so since although the grid resistance goes up by $1/\alpha$, the current demand reduces by α . The IR product thus remains same. Also, the Ldi/dt effect is controlled by this wire sizing since dI/dt goes down by α and inductance exhibits a sub-linear dependence on wire width.¹ Wire sizing for the VDDL grid is done similarly with the addition of the scaling factor $V_{\text{DDH}}/V_{\text{DDL}}$. This scaling factor accounts for the fact that the VDDL grid requires tighter absolute droop/bounce

¹ Although on-chip L is typically dominated by IR, we have included it in our analysis for scenarios where this might not hold. Ldi/dt effect arising from the package is significant & was addressed by Section 4.




Fig. 6. Local and regional areas.

in order to meet the same relative voltage drop budget. The tighter design requirement of the VDDL grid must be imposed on the GND grid too (Eq. 3), since the ground path is shared by VDDL and VDDH cells, and VDDL and VDDH cells are evenly interspersed on the die.

α and β discussed up to this point are chip level (global) current ratios and do not include placement information. Indeed, the current demands across various regions of the chip can differ substantially. In order to include the placement information while sizing each wire segment, we thus introduce local and regional variations of these α and β . We first divide the die into several small “local” areas. The exact size of this local area can be freely chosen; we took it to be the area bound by adjacent ground (or equivalently power) lines on consecutive metal layers. The “regional” area around each “local” area is then defined as the area bounded by its neighboring “local” areas (Fig. 6). Now we compute α and β for all local and regional areas. Finally, each wire segment inside each local region is sized using Eq. 3, where the α and β are replaced by “effective” α and β . The effective α (similarly β) is defined as follows:

$$\alpha_{\text{effective}} = \frac{\alpha_{\text{local}} + \alpha_{\text{regional}} \frac{\text{Area}_{\text{local}}}{\text{Area}_{\text{regional}}} + \alpha_{\text{global}} \frac{\text{Area}_{\text{local}}}{\text{Area}_{\text{global}}}}{1 + \frac{\text{Area}_{\text{local}}}{\text{Area}_{\text{regional}}} + \frac{\text{Area}_{\text{local}}}{\text{Area}_{\text{global}}}} \quad (4)$$

The ratio of the areas used in Eq. 4 act as scaling factors in order to ensure that the local α are weighted more while allowing for neighboring regions to be taken into account when sizing wires. This heuristic approach effectively guides the sizing of the wires by thickening wires in areas of higher current demand and shrinking them down in areas of lower current demand for each grid (VDDH/VDDL/GND). Fig. 7 shows the flowchart for D-Place.

Although the newly proposed 3-rail/4-rail standard cells require library modifications, this can be accomplished using existing design automation tools such as Cadence Abstract Generator. We emphasize that as wires are sized in proportion to the current (maintaining current density), they will not violate electromigration constraints.

5.4 Results

5.4.1 Voltage drop across grids

Table V presents results for the grids when VDDL = 0.6V and 0.8V. The MAX/AVG rows correspond to maximum/average voltage drops across all nodes in the design. The percentages reported are the sum of the power droop and the ground bounce, representing the potential difference available to the cell. Also, the percentages for each cell are taken with respect to its nominal rail to rail swing (1.2V for VDDH gates and 0.6V or 0.8V for VDDL gates).




Fig. 7. Flowchart for D-Place.

Table V. Power grid % voltage drop comparisons.

(A) VDDL = 0.6V					
		Single VDD	DSDG	D-Vanilla	D-Place
c880	MAX	16.9%	30.9%	16.4%	18.6%
	AVG	9.5%	14.7%	9.6%	9.5%
c2670	MAX	25.6%	35.5%	32.2%	25.5%
	AVG	15.9%	19.8%	15.2%	14.5%
c5315	MAX	29.6%	38.2%	37.4%	32.0%
	AVG	21.6%	23.4%	20.2%	19.8%
c7552	MAX	26.8%	34.2%	34.5%	29.4%
	AVG	22.2%	21.0%	21.1%	18.7%

(B) VDDL = 0.8V					
		Single VDD	DSDG	D-Vanilla	D-Place
c880	MAX	16.9%	30.3%	16.3%	19.5%
	AVG	9.5%	15.9%	9.7%	9.8%
c2670	MAX	25.6%	36.1%	27.6%	27.0%
	AVG	15.9%	22.1%	15.8%	15.3%
c5315	MAX	29.6%	38.1%	33.0%	31.8%
	AVG	21.6%	25.4%	20.1%	20.3%
c7552	MAX	26.8%	31.4%	31.6%	28.7%
	AVG	22.2%	24.9%	22.3%	20.1%

From these tables, it can be seen that dual VDD grids can be designed to be as robust as their single VDD counterparts in terms of average voltage drop with some cases showing better results in the dual VDD design. D-Place has slightly inferior (<2.6%) results compared to single VDD in terms of the MAX values (in terms of absolute values in mV this corresponds to <15mV and can be easily compensated by techniques such as locally widening wires if desired by the designer). Also, since the MAX values are singularities, the AVG values as discussed above better depict the general trend. We found that, although D-Place has poorer MAX values in rare cases, the voltage at a majority of other locations on the grid was in fact better for D-Place than for single VDD. This fact is borne out by the fact that although the MAX values for D-Place can be poorer, the AVG values are in fact better (e.g., c7552 in Table V.A). Voltage drop contours shown in Fig. 8 (please note the different scales) show that the dual VDD grid is better off across most of the die. This is also evident in the gate count histogram in Fig. 9.

With respect to the AVG values, D-Place outperforms DSDG and D-Vanilla by up to 6.8% and 2.4% respectively. Looking at the MAX values, the results obtained using D-Place are generally comparable to the single VDD case. On the other hand, D-Vanilla and in particular DSDG frequently have poor performance as compared to single VDD. D-Place outperforms DSDG and D-Vanilla by up to 12.3% and 6.7% with respect to MAX values.

Finally, comparing the VDDL = 0.8V and VDDL = 0.6V cases, we see that the 0.6V case behaves better than 0.8V. This is a favorable finding as the 0.6V VDDL also has lower power (Table I).

We attribute the poor performance of D-Vanilla and DSDG to their failure in considering the important points listed in Section 5.3.




Fig. 8. Voltage drop (%) contours across die for c7552 [VDDL = 0.6V].




Fig. 9. Statistics of gate voltage drop for c7552 [VDDL = 0.6V].

Table VI. Additional power grid metrics.

(A) Voltage variation metric.							
	Single VDD	DSDG		D-Vanilla		D-Place	
		0.6V	0.8V	0.6V	0.8V	0.6V	0.8V
c880	10.4%	24.5%	21.1%	11.2%	11.0%	13.8%	13.5%
c2670	14.9%	26.6%	25.2%	26.3%	22.4%	18.7%	19.7%
c5315	13.7%	28.2%	23.8%	28.4%	22.6%	21.9%	20.2%
c7552	10.8%	19.9%	16.3%	24.5%	23.9%	19.1%	18.3%

(B) Wire congestion metric.							
	Single VDD	DSDG		D-Vanilla		D-Place	
		0.6V	0.8V	0.6V	0.8V	0.6V	0.8V
c880	0.17	0.17	0.17	0.19	0.20	0.17	0.16
c2670	0.17	0.17	0.17	0.19	0.20	0.16	0.16
c5315	0.17	0.17	0.17	0.19	0.20	0.18	0.16
c7552	0.17	0.17	0.17	0.19	0.20	0.15	0.15

In addition, referring to Fig. 5, we can observe that DSDG (in contrast to D-Place and D-Vanilla) results in longer current return paths thus leading to more severe I_{dl}/dt effects.

5.4.2 Additional comparison metrics

Other metrics usually followed when studying power grid performance consider wire congestion and the variation of the voltage across the die. The variation metric is defined as the difference between the maximum and minimum voltage droop/bounce at a given time and is important when performing static timing analysis. This metric should ideally be small. The wire congestion metric amounts to comparing the fraction of routing tracks used by power grid. Again, this metric should ideally be small since the remaining signal wires will have more space for routing. We have ensured that technology imposed rules on minimum wire width and spacing are obeyed by all grids studied. Table VI.A and VI.B compare the various grids with respect to the voltage variation metric and the wire congestion metric for $V_{DDL} = 0.6V$ and $0.8V$. From Table VI.A, among the dual VDD grids, D-Place grids have minimum variation. The voltage variation due to D-Place is somewhat inferior compared to single VDD. From Table VI.B, the wire congestion in the D-Place grids is seen to be superior. This is due to the fact that D-Place adaptively shrinks and widens the grid wires depending on the current demand. DSDG on the other hand is invariant to the value of the V_{DDL} (and the design itself) and hence has uniformly worse congestion. Since D-Vanilla only shrinks down the V_{DDL} rail, it has more congestion.

6. DECOUPLING CAPACITANCE BUDGET

Up to this point we have calculated the decoupling capacitance for the single VDD grid once and held it fixed across all the dual VDD grids. We now relax this constraint to examine how the reduced current demand in dual VDD circuits can be used to reduce decoupling capacitance. Recalling Eq. 2, due to the lower switching current for V_{DDL} gates the required decoupling capacitance corresponding to V_{DDL} gates is also lower. Care must be taken when dealing with the denominator of Eq. 2 however. $V_{noise-lim}$ is an absolute voltage value, and if a constraint of 10% of nominal is considered, the value of $V_{noise-lim}$ differs between V_{DDH} and V_{DDL} gates. We employed this technique and scaled the decoupling capacitances of c2670 accordingly. Table VII summarizes the results for this case including the MAX/AVG voltage droop/bounce of the resultant power grids. Numbers reported in brackets are the values from Table V.B (i.e., with the original decoupling capacitance).

From this table and Table V.B, we see that decoupling capacitance in the dual VDD grid can be reduced from 2.36nF to 1.93nF (18%) while resulting in only 0.6% and 1.6% increase in MAX and AVG

Table VII. Decoupling capacitance (Decap) reduction [$V_{DDL} = 0.8V$].

Scaled Decap Dual VDD		
Decoupling Capacitance	Decap (V_{DDH})	1.02nF (1.06nF)
	Decap (V_{DDL})	0.91nF (1.30nF)
	Total Decap	1.93nF (2.36nF)
Grid integrity metrics		
	MAX	27.6% (27.0%)
	AVG	16.9% (15.3%)

voltage droop/bounce respectively. This reduction will improve leakage, area and yield (arising from oxide defects).

7. CONCLUSIONS

In conclusion, we presented the first detailed study of power delivery issues in dual VDD design. We first showed that dual VDD circuits lead to highly reduced current demands on the power supplies mitigating the power delivery problem. We began with board/package level issues and moved on to describe a placement driven method for designing on-die power grids. We demonstrated that the dual VDD power delivery scenario is no worse than for single VDD circuits. We also showed that dual VDD circuits can afford reduced decoupling capacitance budgets.

We have presented a practical approach for dual VDD grid design that provides superior results as compared to prior approaches. Future work could include the application of more rigorous single VDD power grid optimization approaches such as [28] to dual VDD designs, further enabling this powerful power optimization technique.

8. ACKNOWLEDGMENTS

The authors would like to thank Y. Kim and S. Pant for their valuable inputs and assistance with the PEEC model extractor.

9. REFERENCES

- [1] R. Krishnamurthy, *et al.*, "High-performance and low-power challenges for sub-70nm microprocessor circuits," *Proc. CICC*, pp. 125-128, 2002.
- [2] K. Usami and M. Horowitz, "Clustered voltage scaling technique for low-power design," *Proc. ISLPED*, pp. 3-8, 1995.
- [3] C. Chen, A. Srivastava, and M. Sarrafzadeh, "On gate level power optimization using dual-supply voltages," *IEEE TVLSI*, vol. 9, pp. 616-629, 2001.
- [4] C. Yeh, *et al.*, "Gate-level design exploiting dual supply voltages for power-driven applications," *Proc. DAC*, pp. 68-71, 1999.
- [5] D. Nguyen, *et al.*, "Minimization of dynamic & static power through joint assignment of threshold voltages and sizing optimization," *Proc. ISLPED*, pp. 158-163, 2003.
- [6] S. H. Kulkarni, *et al.*, "A new algorithm for improved VDD assignment in low power dual VDD systems," *Proc. ISLPED*, pp. 200-205, 2004.
- [7] W. Hung, *et al.*, "Total power optimization through simultaneously multiple-VDD multiple-VTH assignment and device sizing," *Proc. ISLPED*, pp. 144-149, 2004.
- [8] K. Usami, *et al.*, "Automated low-power technique exploiting multiple supply voltages applied to a media processor," *IEEE JSSC*, pp. 463-472, 1998.
- [9] S. Mathew, *et al.*, "A 4GHz 300mW 64b integer execution ALU with dual supply voltages in 90nm CMOS," *Proc. ISSCC*, pp. 162-519, 2004.
- [10] K. Zhang, *et al.*, "A 3GHz 79Mb SRAM in 65nm CMOS technology with integrated column based dynamic power supply," *Proc. ISSCC*, 2005.
- [11] C. Yeh, *et al.*, "Layout techniques supporting the use of dual supply voltages for cell-based designs," *Proc. DAC*, pp. 62-67, 1999.
- [12] M. Hamada, Y. Ootaguro, and T. Kuroda, "Utilizing surplus timing for power reduction," *Proc. CICC*, pp. 89-92, 2001.
- [13] M. Igarashi, *et al.*, "A low-power design method using multiple supply voltages," *Proc. ISLPED*, pp. 36-41, 1997.
- [14] J.-S. Wang, *et al.*, "Design of standard cells used in low-power ASIC's exploiting the multiple-supply-voltage scheme," *Proc. IEEE ASIC Conf.*, pp. 119-123, 1998.
- [15] M. Popovich, *et al.*, "On-chip power distribution grids with multiple supply voltages for high performance integrated circuits," *Proc. GLVLSI*, pp. 2-7, 2005.
- [16] J. Fishburn and A. Dunlop, "TILOS: a posynomial programming approach to transistor sizing," *Proc. ICCAD*, pp. 326-328, 1985.
- [17] M. Takahashi, *et al.*, "A 60-mW MPEG4 video codec using clustered voltage scaling with variable supply-voltage scheme," *IEEE JSSC*, pp. 1772-1780, 1998.
- [18] T. Kuroda and M. Hamada, "Low-power CMOS digital design with dual embedded adaptive power supplies," *IEEE JSSC*, pp. 652-655, 2000.
- [19] F. Brglez and H. Fujiwara, "A neural netlist of 10 combinational benchmark circuits and a target translator in Fortran," *Proc. ISCAS*, pp. 695-698, 1985.
- [20] Intel, "Intel Pentium 4 processor in the 423 pin/Intel 850 Chipset Platform," 2002.
- [21] A. E. Ruehli, "Inductance calculations in a complex integrated circuit environment," *IBM J. R&D*, pp. 470-481, 1972.
- [22] S. C. Wong, G. Y. Lee, and D. J. Ma, "Modeling of interconnect capacitance, delay and crosstalk in VLSI," *IEEE Trans. Semiconductor Manufacturing*, pp. 108-111, 2000.
- [23] C. Hoer and C. Love, "Exact inductance equations for rectangular conductors with applications to more complicated geometries," *J. Res. Nat. Bureau Standards*, 1965.
- [24] T.-H. Chen, *et al.*, "INDUCTWISE: inductance-wise interconnect simulator and extractor," *Proc. ICCAD*, pp. 215-220, 2002.
- [25] S. Zhao, K. Roy, and C-K Koh, "Decoupling capacitance allocation and its application to power-supply noise-aware floorplanning," *IEEE TCAD*, pp. 81-92, 2002.
- [26] H. Su, S. S. Sapatnekar, and S. R. Nassif, "Optimal decoupling capacitor sizing and placement for standard-cell layout designs," *IEEE TCAD*, pp. 428-436, 2003.
- [27] E. Chiprout, "Fast flip-chip power grid analysis via locality and grid shells," *Proc. ICCAD*, pp. 485-488, 2004.
- [28] J. Singh and S. Sapatnekar, "Congestion-aware topology optimization of structured power/ground networks," *IEEE TCAD*, pp. 683-695, 2005.