

Power Distribution Paths in 3-D ICs

Vasilis F. Pavlidis

Giovanni De Micheli

LSI-EPFL

1015-Lausanne, Switzerland

{vasileios.pavlidis, giovanni.demicheli}@epfl.ch

ABSTRACT

Distributing power and ground to a vertically integrated system is a complex and difficult task. Interplane communication and power delivery are achieved by through silicon vias (TSVs) in most of the manufacturing techniques for three-dimensional (3-D) circuits. As shown in this paper, these vertical interconnects provide additional low impedance paths for distributing power and ground within a 3-D circuit. These paths, however, have not been considered in the design process of 3-D power and ground distribution networks. By exploiting these additional paths, the IR drop within each plane is reduced. Alternatively, the routing congestion caused by the TSVs can be decreased by removing stacks of metal vias that are used within a power distribution network. Additionally, the required decoupling capacitance for a circuit can be reduced, resulting in significant savings in area. Case studies of power grids demonstrate a significant reduction of 22% in the number of intraplane vias. Alternatively, a 25% decrease in the decoupling capacitance can be achieved.

Categories and Subject Descriptors

B.4.3 [Interconnections]: Physical structures, Topology.

General Terms

Performance, Design, Reliability.

Keywords

Power distribution network, 3-D integration, 3-D ICs, Through silicon vias.

1. INTRODUCTION

Three-dimensional (3-D) integration emerges as a promising solution to the many limitations of modern integrated circuits. The salient characteristics of this novel technology include the significant reduction in the interconnect length and the inherent capability to integrate heterogeneous technologies [1]. To exploit, however, these advantages specific manufacturing and design issues need to be resolved.

Power distribution is a challenging task for 3-D ICs. A significant difficulty in designing a robust power distribution network for 3-D circuits is that only one of the physical planes is directly connected to the power and ground pads of the package as illustrated in Figure 1. Consequently, power and ground are provided to the lower planes

Permission to make digital or hard copies 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 and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GLSVLSI'09, May 10–12, 2009, Boston, Massachusetts, USA.

Copyright 2009 ACM 978-1-60558-522-2/09/05...\$5.00.

through the vertical interconnects. Within each plane a power distribution network as in 2-D circuits provides sufficient current to the devices.

In planar (2-D) circuits, several power distribution networks have been developed with power grids being a commonly used topology [2]. The portion of the power grid implemented with the top metal layers is typically utilized to globally distribute current across the circuit. Multiple stacks of metal vias connecting adjacent metal layers are used to carry the current from the upper to the lower metal layers. The thin and more resistive metal layers (e.g., M1) are used to locally provide current to the transistors.




Figure 1. A four-plane 3-D integrated circuit with the I/O pads located at the uppermost physical plane.

In a 3-D circuit, each physical plane except for the uppermost plane (assuming the structure shown in Figure 1) is connected to the power and ground pads through the vertical interconnects, which are typically implemented with through silicon vias (TSVs). The through silicon vias are considerably wider than the metal vias and, therefore, less resistive. Additionally, a TSV can connect the topmost metal layer of one plane to the lowest metal layer of an adjacent physical plane providing a low impedance path for distributing the current.

Exploiting this additional path to enhance power and ground distribution within a 3-D circuit is the objective of this paper. 3-D integration technologies that support these additional paths are discussed in the following section. An analytic approach to investigate the decrease in voltage drop due to the low impedance path is described in Section 3. The effect of this path on the voltage drop within a power grid is discussed in Section 4. These additional paths can be utilized to alleviate the routing congestion due to the TSVs or to decrease the intentional decoupling capacitance. Some conclusions are offered in Section 5.

2. THROUGH SILICON VIA MANUFACTURING TECHNOLOGIES

Different fabrication approaches for TSVs are reviewed in this section. The TSVs are manufactured by etching holes into the silicon substrate or the buried oxide of a CMOS or SOI plane, respectively. A thin insulator layer is deposited (if necessary) and the openings are filled with metal, forming the interplane galvanic connections of a 3-D IC [1].

There are two main approaches in manufacturing the TSVs. The “via-first” approach where the TSVs are formed before the metallization and the “via-last” method where the TSVs are fabricated after depositing the metal layers of a plane. A “via-first” technique is depicted in Figure 2a [3]. The TSVs start from the topmost metal layer of the first plane and land on the first metal layer of the second plane. A large number of intraplane vias is used to reach the starting pad of the TSV that connects the second with the third plane.

Consequently, a TSV can directly interconnect the topmost with the lowest metal layer of two adjacent planes. In other words, both the starting and landing pads of a TSV can be connected to the power distribution network. These interconnections form a path with low impedance characteristics as compared to a traditional stack of intraplane vias connecting these two metal layers. This additional path, however, has not been considered in distributing power and ground within a 3-D circuit [4], [5].

In “via-last” approaches, the TSVs usually connect only the uppermost metal layers of successive planes. An example of a “via-last” technology is illustrated in Figure 2b [6]. In this technology, a TSV is formed between the back side metal layers of two adjacent planes through the M3 of the upper plane. By using a BVIA in Figure 2b), the TSV can also be connected to M1. Consequently, for specific “via last” approaches additional paths can also be formed to facilitate power and ground distribution. In the following section, a simple analysis is used to demonstrate the benefits resulting from these secondary paths.




Figure 2. Different manufacturing approaches for 3-D ICs (not to scale), (a) “via-first” approach [3] where the TSVs are formed immediately after the fabrication of the transistors and (b) “via-last” approach [6] where the TSVs are formed after the deposition of the metal layers on each plane.

3. EFFECT OF TSVs ON POWER DISTRIBUTION NETWORKS

In this section, the effect of a TSV on the produced voltage drop, where the TSV is connected both to the topmost and lowest metal layers of a power distribution network, is investigated. Consider, for example, the one-dimensional segment of a power distribution network of a physical plane within a 3-D circuit shown in Figure 3. The equivalent circuits of these two segments are illustrated in Figure 4. Note that each circuit contains two paths but with different impedance characteristics.

The primary difference between the two power delivery systems depicted in Figures 3a and 3b is that the TSV in Figure 3b connects only to the topmost metal layer (MT), while in Figure 3a the TSV also connects to the first metal layer (M1). With the latter approach additional paths are formed, which are shown by the thick solid curves in Figure 3a. These paths are called “TSV paths” in the remainder of the paper for brevity. To investigate the physical behavior of the TSV path, a current source is assumed to be connected to MT through a stack of intraplane vias. Note that devices can only be placed at a certain distance from a TSV due to manufacturing limitations.

The current is considered to flow only over two metal layers (e.g., MT and M1). The inclusion of any other metal layer would increase the impedance of the path since MT is the least resistive layer. The on-chip inductance is omitted in the analysis due to the local (within a few micrometers) and fast decay of the effect that inductance has on the power distribution networks [7]. The inductive component of the TSV impedance is, however, included in the analysis. Furthermore, the capacitance of the TSV is not considered, since for power and ground TSVs this capacitance behaves as decoupling capacitance improving the impedance characteristics of the power delivery system. Consequently, including the capacitance of the TSV in the analysis further demonstrates the usefulness of the TSV path.




Figure 3. Current paths within a 3-D circuit, (a) where the TSV is connected to the power lines on both the uppermost (MT) and the first (M1) metal layers and (b) where the TSV is connected only to the topmost (MT) metal layer.




Figure 4. Equivalent circuits of the current flow paths illustrated in Figure 3.

Based on the circuits shown in Figure 4, a first order analysis is used to determine the greatest current that each of the circuits can carry while satisfying the voltage drop constraints. Assuming an allowed ripple of $V_{ripple} = 5\%$ of the nominal power supply, the maximum current that the circuits in Figures 4a and 4b can carry is, respectively, described by

$$I_L \leq \frac{(1 - V_{ripple})V_{dd}}{(R_{M1} + Z_{TSV}) \parallel (R_{MT} + R_{via})}, \quad (1)$$

and

$$I_L \leq \frac{(1 - V_{ripple})V_{dd}}{(R_{M1} + R_{via}) \parallel (R_{MT} + R_{via})}. \quad (2)$$

Z_{TSV} and R_{via} are the impedance of a TSV and the resistance of a stack of vias connecting the topmost and lowest metal layers, respectively (see Figure 3). The resistance of the topmost and lower metal layers are $R_{MT} = r_{MT}l$ and $R_{M1} = r_{M1}l$, respectively, where r_{MT} and r_1 are the resistance per length and l is the distance of the current source from the TSV. The width of the power and ground lines on the topmost layer is twenty times wider than the minimum supported width. The width of the lowest metal layer is twice the minimum supported width. Consequently, for an industrial $0.18 \mu\text{m}$ technology with six metal layers, the resistance per length is $r_{M1} = 29.78r_{MT}$ and $r_{MT} = 4.62 \Omega/\text{mm}$. In addition, the resistance of a stack of vias between M6 and M1 is $R_{via} = 32.5 \Omega$ [8]. Note that at advanced technology nodes the resistance of a via will scale. The number of interconnect layers; however, will increase requiring a larger number of vias to connect to the topmost layer of a power distribution network. Consequently, the total resistance of a stack of vias R_{via} remains significantly greater than the impedance of a TSV.

Assuming that $V_{dd} = 1$ Volt and $l = 30 \mu\text{m}$ [1], the voltage drop at the current source, V_L for both circuits is plotted in Figure 5. The TSV is assumed to have a total resistance of 1Ω , including the resistance of the large group of metal vias (see Figure 2a) and the inductance is $L_{TSV} = 20 \text{ pH}$ [5]. The maximum switching frequency is set to $f = 10 \text{ GHz}$ to consider the high frequency components constituting a digital signal. Additionally, the selected TSV resistance is larger than typically reported resistances depicting a worst case resistance for the TSV [1], [3].




Figure 5. Voltage drop at the current source as a function of the current drawn by the power supply.

From Figure 5, when the TSV path is exploited a significantly larger current can be supplied to the transistors without exceeding the allowed voltage ripple. Thus, for the circuit shown in Figure 4b the limit of the voltage drop is reached where $I_{Lmax} = 3.3 \text{ mA}$, while for the circuit that includes the TSV path the maximum current that can be sustained is $I_{Lmax} = 9.4 \text{ mA}$; a considerable $2.8\times$ increase.

Alternatively, the maximum distance of the current source from the TSV, where the circuit sinks a fixed current is plotted in Figure 6. When the TSV path is considered, the current source can be placed farther away from the power and ground pads or, equivalently, a smaller number of stacks of vias can be used to sufficiently distribute the current within the circuit.




Figure 6. Voltage drop as a function of the distance of the current source from the TSV.

The difference in the produced voltage drop for a fixed current load of $I_L = 3$ mA between the circuits of Figure 4 is significant for small distances and gradually decreases with distance. For larger distances, in the order of a few millimeters, both of the circuits result in a similar voltage drop. This behavior can be explained by considering the current flow through the different paths that exist within the circuit shown in Figure 4a. In this power distribution network, the current is propagated through the low resistance TSV and M1, which exhibits a higher resistance as compared to M6. Alternatively, the current can be distributed through the less resistive M6 as compared to M1 and the stack of vias, which in turn, is at least an order of magnitude more resistive than the TSV.

For short distances (*i.e.*, hundreds of micrometers), the path over the TSV and M1 (*e.g.*, the TSV path) exhibits a lower impedance as compared to the path consisting of M6 and the stack of vias. Consequently, most of the current flows through the TSV path resulting in a substantially smaller voltage drop as compared to the circuit shown in Figure 4b. For the circuit in Figure 4b, the greatest portion of the current flows through M6 and the stack of vias, since the alternative path comprising M1 and a stack of vias exhibits a considerably greater impedance.

As the distance between the current source and the TSV in Figure 4a increases, the resistance along M1 also increases faster than the resistance along M6. Beyond a specific distance, which depends on the impedance characteristics of the interconnects, the path that consists of M6 and a via stack begins to exhibit a lower impedance than the path that includes M1 and the TSV. Beyond this distance most of the current flows through M6 and, eventually, in both circuits the voltage drop is approximately the same. Note that the inclusion of the on-chip inductance would further emphasize the use of the TSV path since the impedance of layer M6 would increase.

This behavior, therefore, suggests that the TSV path when used for power and ground distribution has a local effect and is efficient for those transistors included within a specific region around the TSV. This distance is essentially determined by the current demand in the vicinity of the TSV and the interconnect impedance characteristics of the TSV and intraplane metal layers.

In other words, these ancillary paths should not be perceived as another means to globally distribute power and ground within a 3-D circuit. These paths, however, can be used to locally enhance the power distribution within the circuit. There are different ways to

exploit this advantage. For example, the TSVs are a crucial element of a 3-D circuit since these interconnects provide the interplane communication and power. The size of the TSVs, however, considerably increases the routing congestion. To mitigate this increase, several stacks of vias within the power grid in a nearby region from each TSV can be removed, since most of the current flows through the TSV and M1. Alternatively, the required intentional decoupling capacitance can be reduced, since the voltage drop from the power supply to the transistors is considerably improved. Case studies of power grids are discussed in the following section highlighting these points.

4. POWER GRIDS WITH TSV PATHS

Although the circuits investigated in Section 3 are suitable for a first-order analysis, the effect of the TSV path is more accurately captured by exploring the behavior of a power grid. Portions of a power delivery system for one plane of a 3-D IC are analyzed in this section.

Two pairs of ten by ten resistive grids are used to model a portion of a power distribution network. Each pair corresponds to the topmost (*i.e.*, M6) and lowest (*i.e.*, M1) metal layers, which are depicted in Figures 7 and 8, respectively. At each grid node except for the node at the middle, a stack of vias is used to connect the two grids. At the center of the grids a TSV connects to a power pad. The only difference between the two pairs of grids is that in one of the grids the TSV is connected to both M6 and M1 in addition to the power pad (*i.e.*, the TSV path). Alternatively, in the other grid, the TSV is connected only to the power pad and M6, while the node at the center of the M6 and M1 grids are connected with a stack of interplane vias. The ground distribution network is similarly modeled.

Each stack of vias is modeled as a resistor, while the impedance of the TSV includes a resistive and an inductive component. The resistances and inductance provided in Section 3 are also used in this section. Specific nodes of the grid that models layer M1 are connected to a current source are shown in Figure 8. Furthermore a decoupling capacitor notated as C_{dec} is connected to every grid node, as illustrated in Figure 8. Each current source is modeled with a triangular waveform [9]. The rise and fall times are 30 ps and 70 ps, respectively, and the switching period is 100 ps. In other words, no intermediate time between successive switching is assumed. A voltage ripple of 5% of the power supply is assumed and $V_{dd} = 1$ Volt.

To demonstrate the effect of the low impedance path within the power grid formed by connecting the TSV to both M6 and M1, two different switching scenarios are considered. Initially, all of the sources shown in Figure 8 draw current, while in the second scenario only three sources switch. The length of the grid segments is notated by l , and is varied to explore the resulting voltage drop.

For the first scenario, each current source draws a peak current of $I_L = 0.8$ mA. In addition, the decoupling capacitors are set to zero to investigate the voltage drop across the entire grid caused by the switching of the current sources. Both of the grids are simulated with SPICE.




Figure 7. Resistive grid used to model a segment of a power delivery system on the uppermost (M6) metal layer.




Figure 8. Resistive grid used to model a segment of a power delivery system on the lowest (M1) metal layer.

The voltage drop at specific nodes of the M1 grid (including the node where the maximum voltage drop occurs) are plotted in Figure 9 for increasing length l of the grid segment. These nodes are located at the upper left corner (S1), at a $4l$ distance to the right of the TSV (S2), and at the TSV (S3). The voltage drop on these nodes is illustrated in Figure 9 by the curves denoted by circles, squares, and triangles, respectively. For the grid where the TSV path is present (depicted in Figure 9 by the family of solid curves), the voltage drop is considerably lower as compared to the grid where the TSV path is not considered. Note that the voltage drop at the current source located at the upper left corner of the grid (S1) (*i.e.*, pair of curves denoted by circles) is affected less by the TSV path as compared to the other two nodes. This situation demonstrates the locality of the effect caused by the TSV path.

In addition, there is a negligible increase (~ 5 mV) of the voltage at the TSV location shown by the solid curve with increasing l . This counterintuitive behavior can be explained by considering the

currents that flow through the TSV path and the neighboring paths through the stacks of vias. As l increases, the impedance of each M1 segment becomes comparable to that of a stack of vias. Consequently, the current that flows through the TSV and M1 starts decreasing. Alternatively, the current that flows over M6 and the neighboring to the TSV stacks of vias increases. This change in the flow of the current causes the small voltage increase at the TSV node with increasing l .




Figure 9. SPICE simulations of the voltage drop on M1 grid for different nodes and with no stack of vias removed ($d = 0$) with (solid curves) and without (dashed curves) the TSV path.

The decrease in the voltage drop due to the TSV path can be used to improve routability within a plane. This improvement is important since TSVs increase routing congestion. To demonstrate that a smaller number of stack of vias within the power grid is required when the TSV path is considered, stacks of vias are removed within an increasing radius from the TSV, which is notated as d . The resulting voltage drop is depicted in Figure 10 where $l = 30$ μm . Note that although there are fewer paths providing current to the transistors, since stacks of vias are removed from the grid, the TSV path supports a greater amount of current and, therefore, the voltage drop specification is maintained up to $d = 3$ or, equivalently, with 22% fewer intraplane vias (see dashed curves in Figure 10).

Stacks of vias are also removed from the grid where the TSV path is not present. The voltage drop on this grid is also shown in Figure 10 by the solid curves. In this grid, the TSV is not connected to M1 and, consequently, the voltage drop increases rapidly as stacks of vias are removed in an effort to reduce routing congestion.

In the second simulation scenario, only three sources switch. These three sources are located at the upper left corner of the grid (S1) and at the adjacent nodes (in the east and west direction) of the node where the TSV is connected (S2). The peak current of these sources is $I_L = 8$ mA, while the rise and fall times are the same as in the previous scenario. The voltage drop at these current sources (*i.e.*, S1, S2) and the center of the M1 grid (S3) is depicted in Figure 11. The voltage drop within the two grids with (solid curves) and without (dashed curves) the TSV path is illustrated. The simulation results indicate that the additional path can again decrease the voltage drop for this current source configuration. The locality of the efficiency of the TSV path is also demonstrated, since the voltage drop at the remote current source does not change considerably with varying the length of the grid segment. Note the solid and dashed curves denoted by circles in Figure 11, which are practically indistinguishable.




Figure 10. SPICE simulations of the maximum voltage drop on M1 grid by successively removing stacks of vias (*i.e.*, increasing d) with (dashed curves) and without (solid curves) the TSV path.




Figure 11. SPICE simulations of the voltage drop on M1 grid for different nodes and with no stack of vias removed ($d = 0$) with (solid curves) and without (dashed curves) the TSV path. Only three current sources switch.

The proposed power distribution paths can also be utilized to decrease the extrinsic or intentional decoupling capacitance, which is used to compensate for the voltage drop on the power grid. For the first simulation scenario, the peak current of the sources is increased to $I_L = 1$ mA, while the decoupling capacitors are varied so that the voltage drop constraint is approximately met. The required decoupling capacitance is listed in Table 1 where the TSV path is present and where this path is not considered. The grid that includes the TSV path requires 25% less capacitance to satisfy the voltage drop constraints. This decrease results in important savings in the area of a plane within a 3-D system.

Table 1. Minimum node voltage within the power grids with and without considering the TSV path.

Total decoupling capacitance [nF]	Minimum node voltage [mV]	
	with TSV path	without TSV path
0.05	952	945
0.0625	956	952

5. CONCLUSIONS

3-D integration technologies support secondary paths for distributing power and ground within a 3-D circuit. A first-order analysis demonstrates the usefulness of including these paths in the design process of 3-D power distribution networks. These paths exhibit particularly low impedance characteristics supporting the distribution of large amount of current in the vicinity of a TSV without exceeding the pre-determined voltage drop. Consequently, stacks of vias can be removed from the power grid, since a small percentage of current flows through these vias, decreasing routing congestion within a 3-D circuit. Alternatively, the intentional decoupling capacitance can be decreased resulting in a considerable savings in area. For a case study of a power grid, this capacitance is decreased by 25% when the TSV is connected to M1 in addition to the topmost metal layer. Alternatively, 22% fewer intraplane vias can be utilized while satisfying the allowed voltage ripple.

6. REFERENCES

- [1] V. F. Pavlidis and E. G. Friedman, *Three-Dimensional Integrated Circuit Design*, Morgan Kaufmann Publishers, 2009.
- [2] M. Popovich, A. V. Mezhiba, and E. G. Friedman, *Power Distribution Networks with On-Chip Decoupling Capacitors*, Springer, 2008.
- [3] R. S. Patti, "Three-Dimensional Integrated Circuits and the Future of System-on-Chip Designs," *Proceedings of the IEEE*, Vol. 94, No. 6, pp. 1214-1224, June 2006.
- [4] J. Sun *et al.*, "3D Power Delivery for Microprocessors and High-Performance ASICs," *Proceedings of the IEEE Applied Power Electronics Conference*, pp. 127-133, February 2007.
- [5] G. Huang *et al.*, "Power Delivery for 3D Chip Stacks: Physical Modeling and Design Implication," *Proceedings of the IEEE Electrical Performance of Electronic Packaging Conference*, pp. 205-208, October 2007.
- [6] C. K. Chen *et al.*, "Characterization of a Three-Dimensional SOI Integrated-Circuit Technology," *Proceedings of the IEEE SOI Conference*, pp. 109-110, October 2008.
- [7] S. Pant and E. Chiprout, "Power Grid Physics and Implications for CAD," *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 199-204, July 2006.
- [8] R. Doering and Y. Nishi (eds.), *Handbook of Semiconductor Manufacturing Technology*, CRC Press, 2008.
- [9] M. Popovich, M. Sotman, A. Kolodny, and E. G. Friedman, "Effective Radii of On-Chip Decoupling Capacitors," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 16, No. 7, pp. 894-907, July 2008.