Link Breaking Methodology: Mitigating Noise within Power Networks

Renatas Jakushokas
jakushok@ece.rocheste.edu

Eby G. Friedman
friedman@ece.rochester.edu

Department of Electrical and Computer Engineering
University of Rochester
Rochester, New York 14627

ABSTRACT

A link breaking methodology is introduced to reduce voltage degradation within mesh structured power distribution networks. The resulting power distribution network combines a single power distribution network to lower the network impedance, and multiple networks to reduce noise coupling among the circuits. Since the sensitivity to supply voltage variations within a power distribution network can vary among different circuits, the proposed methodology reduces the voltage drop at the more sensitive circuits, while penalizes the less sensitive circuits. The proposed methodology is evaluated for two case studies, demonstrating a reduction in the voltage drop in sensitive circuits. Based on these case studies, the voltage is improved by, on average, 4% at those nodes with the highest sensitivity. The voltage after application of the link breaking methodology is, on average, 96% of the ideal power supply voltage. Lowering the noise on the power network enhances, on average, the maximum operating frequency by 11% by utilizing the proposed link breaking methodology.

Categories and Subject Descriptors

B.7.1 [**Hardware**]: Integrated Circuits—*Types and Design Styles*

Keywords

Power Distribution Networks, Power Integrity, Power Noise

1. INTRODUCTION

The increasing density and performance of integrated circuits (IC) requires advancements in design methodologies for the global interconnects, particularly the on-chip power networks, clock networks, and long distance on-chip signals. The on-chip power distribution network typically provides hundreds of amperes to the load circuits while utilizing up to 40% of the overall metal resources [1, 2]. With advancements in technology, higher current is required; therefore, more efficient on-chip power distribution networks have become an essential element of modern IC design flows.

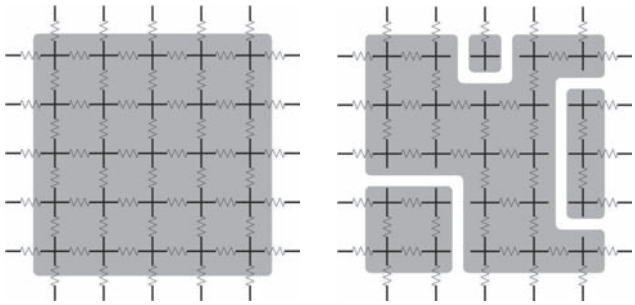

(a) (b)

Figure 1: Mesh structured power distribution network. (a) single power distribution network focused on reducing the network impedance. (b) multiple power distribution networks lower the noise at the expense of increasing the network impedance.

A change in voltage at the power node of a gate can significantly increase the delay of a logic gate [3, 4], degrading the overall performance of a system [5]. Since different circuits are affected differently by a drop in the power supply voltage, the power distribution network should be designed to satisfy multiple constraints. The voltage level for those gates along the critical path can tolerate the least voltage degradation, while the gates along a non-critical path may satisfy speed constraints despite a higher voltage drop [6]. Circuits such as PLLs (phase lock loops) and VCOs (voltage controlled oscillators) are highly sensitive to changes in the power supply [7], while digital logic circuits can tolerate much higher variations in the power supply voltage.

Separate power networks can be designed to independently supply current to different parts of a circuit; thereby shielding different parts of an IC from each other. Separate power networks are widely used in mixed-signal circuits, where the current is supplied by different power networks to the analog and digital circuits [8]. For systems requiring the same voltage level, this approach however may inefficiently utilize metal resources due to the additional area and routing constraints [6]. The number of I/O pads is also a limiting resource, preventing the use of an excessive number of separate power networks [9]. In Fig. 1, a single and multiple separate power networks are illustrated. With a single network, as shown in Fig. 1(a), the sensitive circuit (for example, a PLL) and aggressor circuit (exemplified by a large digital logic circuit) share the same power network, thereby lowering the power network impedance. A sensitive circuit can however be highly affected by the noise gen-

This research is supported in part by the National Science Foundation under Grant Nos. CCF-0811317 and CCF-0829915, grants from the New York State Office of Science, Technology and Academic Research to the Center for Advanced Technology in Electronic Imaging Systems, and by grants from Intel Corporation and Qualcomm Corporation.

GLSVLSI'12, May 3–4, 2012, Salt Lake City, Utah, USA.

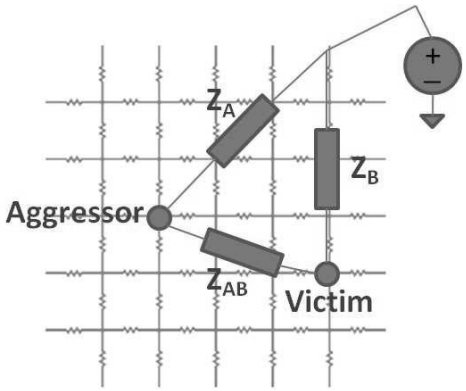

Figure 2: Aggressor and victim circuits sharing a mesh structured power distribution network. The objective is to increase Z_{AB}, while insignificantly increasing Z_A, resulting in shielding the victim from an aggressor.

erated from an aggressor circuit. With multiple power networks, as shown in Fig. 1(b), one network can be dedicated to an aggressor circuit while another network can be dedicated to the sensitive circuits, minimizing noise coupling between the aggressor and sensitive circuits. This approach however results in an increase in the power network impedance and routability constraints. The methodology proposed in this paper utilizes a single power network to provide a low network impedance and reduced routability constraints, while disconnecting (or breaking) links within the on-chip power network between an aggressor and sensitive circuit, thereby reducing noise coupling to the sensitive circuits.

The paper is organized as follows. The primary design objective for reducing voltage variations is formulated in Section 2. The sensitivity of the victim circuits to variations in the voltage within the power network is characterized in Section 3. In Section 4, the proposed link breaking methodology is described. An algorithm for breaking links for a large number of aggressor and victim circuits connected to a common on-chip power distribution network is also described in this section. In Section 5, two design cases are evaluated. The degradation in the supply voltage and propagation delay before and after applying the proposed link breaking methodology is summarized. Additional discussion related to enhancing the voltage levels within an on-chip power distribution network and the computational runtime of the algorithm is presented in Section 6. Finally, the conclusions are summarized in Section 7.

2. REDUCTION IN VOLTAGE VARIATIONS

Every circuit connected to the power distribution network can be considered as an aggressor and a victim. Two parameters are therefore assigned to each circuit, one characterizing the aggressiveness and the second the sensitivity of a circuit. The aggressor parameter is directly related to the current sunk by a circuit. Simultaneously, every circuit exhibits a different sensitivity to variations in the power network voltage. For example, a PLL is highly sensitive to voltage variations as compared to digital logic. Two circuits with a different critical path may also exhibit a different sensitivity to voltage variations: a slower critical path requires a smallest power drop, while a fast critical path can better tolerate a large voltage drop on the power network. A sensitivity factor is therefore assigned to each circuit connected to the power network. A more detailed discussion of the sensitivity factor is presented in Section 3.

In a system with multiple aggressors and victims, the objective

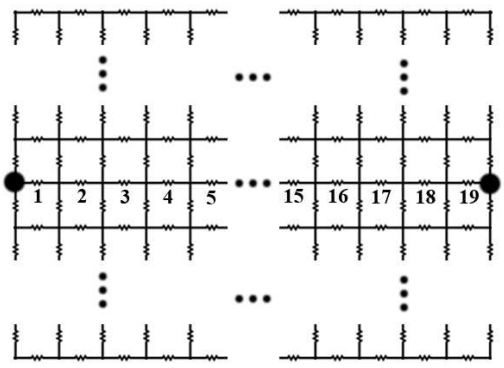

Figure 3: 20×20 node mesh structured network. The effective resistance is between the two bold nodes. The links are numbered based on the location along the x-axis.

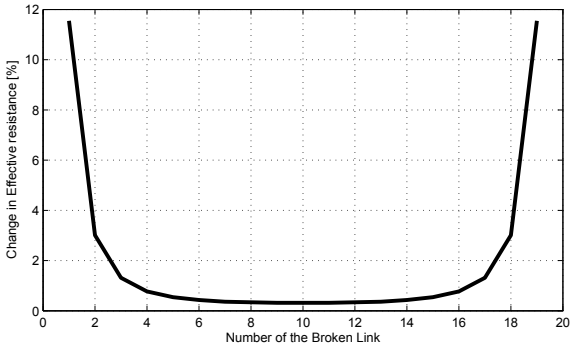

Figure 4: A change in the effective resistance between the left and right nodes within a 20×20 mesh structured power distribution network (see Fig. 3) as a function of a specific location of a disconnected link between two nodes.

LINK-BREAKING
1. Determine voltage drops at all k nodes
2. Calculate initial $delay_{ini}$ function based on (2)
3. Generate x randomly perturbed systems
4. Determine voltage drops at k nodes for x systems
5. Calculate $delay$ function based on (2) for x systems
6. For every x systems
7. Generate six different networks, where a link is broken at every direction
8. Determine new $delay$ values, maintaining network with lowest $delay$
9. Goto 7, if enchantment is achieved
10. Select system with lowest $delay$
11. If $delay_{ini} > delay$, $delay_{ini} \leftarrow delay$ and goto 3

Figure 5: Pseudocode for link breaking algorithm.

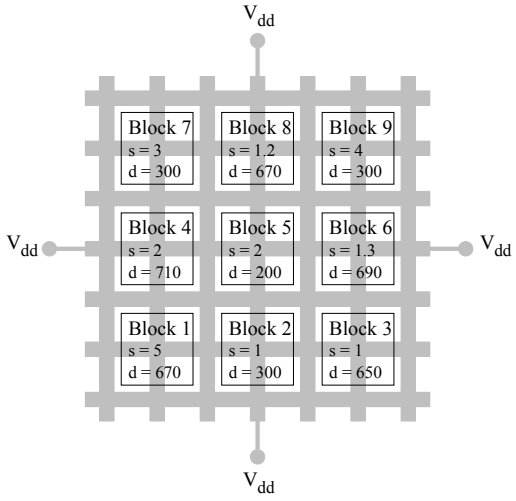

Figure 6: Nine circuit blocks are connected to a mesh structured power distribution network. Four power supplies provide the current. The numbers indicated within the blocks represent the sensitivity factor (s) and propagation delay in ps (d) when applying one volt to the block. Note that the minimum propagation delay is achieved when applying a full power supply.

is to minimize the effect of the voltage drop over the entire system. To improve the performance of an IC, the voltage drop is reduced in those circuits with high sensitivity at the expense of increasing the voltage drop in the less sensitive circuits. Three specific nodes, the victim, aggressor, and power supply, within a mesh structured power distribution network, are illustrated in Fig. 2. The objective is to increase the network impedance between the victim and the aggressor nodes (Z_{AB}), thereby reducing the influence of the aggressor on the victim node, while only minimally increasing the effective impedance between the aggressor and the power supply (Z_A).

A 20×20 node mesh structured network is illustrated in Fig 3. The normalized effective resistance between the left and right nodes as a function of a specific disconnected link at a particular location (along the x-axis) is depicted in Fig. 4. The x-axis describes the location (or link number depicted in Fig. 3) of the disconnected link between two nodes. The largest increase in the effective resistance is achieved when breaking the link closest to either node. An 11% increase in the resistance is caused by breaking a single link. This change confirms that breaking links within a mesh structured power distribution networks may result in a large change in the effective impedance; effectively shielding the victim from the aggressor.

3. SENSITIVITY FACTOR

The sensitivity factor describes the relative importance of a change in voltage on the performance of a circuit. A method to describe the sensitivity factor is to investigate the sensitivity of the supplied voltage on the performance (for example, the propagation delay) of a particular circuit. The sensitivity factor is [10]

$$s = \left. \frac{\frac{\Delta delay}{delay(x)}}{\frac{\Delta V}{V(x)}} \right|_{x=V_{dd}} = \frac{\Delta delay}{\Delta V} \cdot \frac{V_{dd}}{delay_{min}}, \qquad (1)$$

where $\Delta delay$ and $delay_{min}$ are, respectively, the change in the delay and the minimum delay of a circuit. The minimum delay is achieved assuming a full V_{dd} at the power rail of the circuit. ΔV is the change in the supply voltage at the node supplied to the circuit. The sensitivity factor is dependant on the type of circuit.

4. LINK BREAKING METHODOLOGY

An algorithm for determining which links should be removed, thereby shielding the sensitive circuits, is described in this section. Since each circuit within a network can be characterized as both an aggressor and a victim, each node of interest is associated with a matrix composed of two parameters $[i,s]$. Parameter i is an aggressor related parameter, and is equal to the current sunk from the network. Parameter s is related to the victim, expressing the sensitivity of the circuit connected to the node. The objective is to enhance overall performance, such as minimize the worst case delay.

$$delay_{worst} = max\left(delay_1, delay_2, ..., delay_k\right), \qquad (2)$$

where

$$delay_j = delay_{min-j} \cdot \left[\frac{s_j}{V_{dd}} \cdot \Delta V_j + 1 \right]. \qquad (3)$$

ΔV_j is a change in the voltage at node j due to load currents and the impedance of the mesh structured power distribution network. $delay_{min-j}$ is the minimum propagation delay of circuit j achieved by applying the maximum supply voltage V_{dd}. s_j is the sensitivity factor of circuit j.

Pseudocode of the LINK-BREAKING algorithm for the proposed methodology is provided in Fig. 5, with the objective of minimizing the worst case propagation delay.

In line 1, the voltage drop at k nodes (all aggressor/sensitive nodes) is determined. Based on the voltage and sensitivity of the circuits, the initial value of the delay function $delay_{ini}$ is determined, as listed in line 2. The revised number of power networks x is generated, where each network is perturbed by removing a random link. In lines 4 and 5, the voltage drop and delay function are determined for each of the perturbed networks. A search for a local minimum is evaluated for each perturbed system in lines 6 to 9. The network with the lowest $delay$ value is selected in line 10. The process is repeated until the value of the delay function cannot be further reduced.

Since k nodes of interest are typically lower than the overall number of nodes in a system, a random walk procedure can be used to efficiently determine the voltages [11], trading off accuracy with runtime. The number of parallel random walk procedures is based on the target accuracy.

5. CASE STUDIES

Two study cases are presented in this section. In both cases, the circuit is composed of nine blocks. For the first case, one block sinks significantly greater current, representing the case of a single dominant aggressor. In the second case, the current and delay of the nine blocks are varied, representing general circuits. The design objective is to minimize the worst case propagation delay.

A mesh structured power distribution network with 20×20 number of nodes is considered. A block diagram of the circuit is schematically illustrated in Fig. 6. Four one volt power supplies are connected at the center of the four edges (left, right, top, and bottom). The maximum permitted degradation in supply voltage is 0.3 V.

The current is different among the circuit blocks (see Table 1) for both study cases. The supply voltage map before and after application of the link breaking methodology, as well as the resulting power network, is illustrated in Fig. 7. The current sunk for each

Table 1: Sensitivity factor, sunk current, minimum delay, supply voltage, and propagation delay before and after application of the link breaking methodology for the nine circuit blocks. The improvement or degradation in the supply voltage, propagation delay, and maximum operating frequency are also listed. Case 1 represents the case where a single block sinks significantly higher current as compared to the other blocks. In Case 2, the sunk current, sensitivity factor, and delay are different for different blocks, representing a general design case.

Block number	1	2	3	4	5	6	7	8	9	$f_{max} = \frac{1}{delay_{worst}}$
Sensitivity factor (s)	5	1	1	2	2	1.3	3	1.2	4	——
Delay [ps] @ $V_{dd} = 1V$	670	300	650	710	200	690	300	300	300	——
Case 1 (see Figs. 7(a) and 7(b))										
Sunk current	1	1	1	1	1	1	1	10	1	——
Voltage before methodology [mV]	945	944	944	937	922	934	925	850	924	——
Voltage after methodology [mV]	966	934	924	947	936	906	939	700	909	——
Voltage improvement [%]	2.1	-1.1	-2.0	-1.0	1.6	-3.0	1.5	-17.6	-1.6	——
Delay before methodology [ps]	904	336	727	847	245	794	390	375	413	1.11 GHz
Delay after methodology [ps]	855	349	762	857	243	844	387	445	446	1.17 GHz
Delay improvement [%]	5.4	-3.9	-4.8	-1.2	0.8	-6.3	0.8	-18.7	-8.0	5.4
Case 2 (see Figs. 7(d) and 7(e))										
Sunk current	1	5	5	2	2	3	1.3	4	1.2	——
Voltage before methodology [mV]	907	861	850	901	874	875	907	876	901	——
Voltage after methodology [mV]	958	825	781	928	838	864	852	716	890	——
Voltage improvement [%]	5.5	-4.3	-8.1	2.9	-4.1	-1.2	-6.1	-18.2	-1.3	——
Delay before methodology [ps]	1050	366	800	910	268	860	411	369	448	0.95 GHz
Delay after methodology [ps]	870	378	847	870	283	869	463	430	463	1.15 GHz
Delay improvement [%]	17.1	-3.2	-5.8	4.5	-6.0	-1.1	-12.7	-16.5	-3.3	17.1

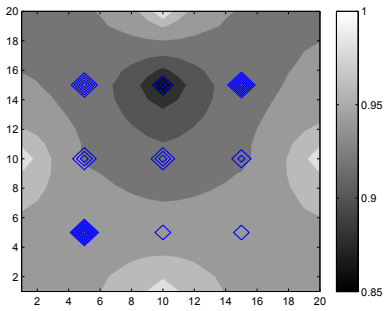

(a) Case 1: voltage before optimization

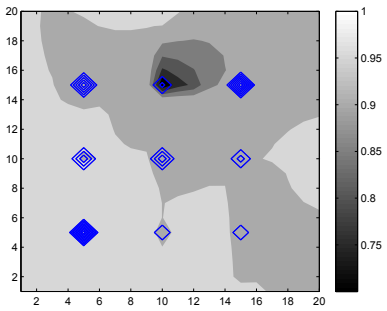

(b) Case 1: voltage after optimization

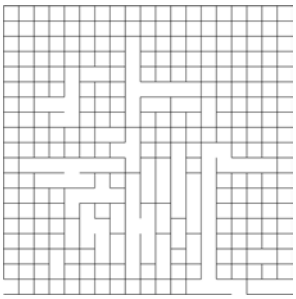

(c) Case 1: power network after optimization

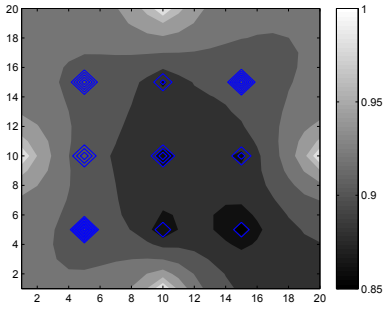

(d) Case 2: voltage before optimization

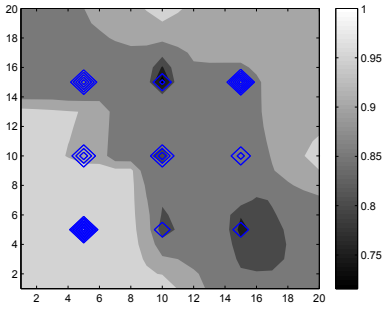

(e) Case 2: voltage after optimization

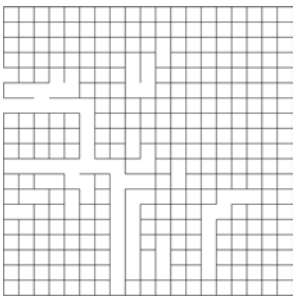

(f) Case 2: power network after optimization

Figure 7: Map of voltage variations before and after application of the link breaking methodology for two case studies. The diamond shapes represent the location of the aggressor/victim circuit blocks. The size of the diamond represents the relative sensitivity factor of a particular block. The resulting power network after the link breaking methodology is also illustrated. The first case represents the case where a single block sinks significantly higher current as compared to the other blocks. In the second case, the sunk current, sensitivity factor, and delay are different for different blocks, representing a general design case.

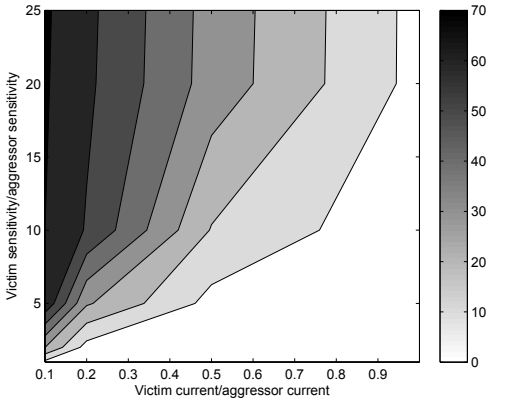

(a)

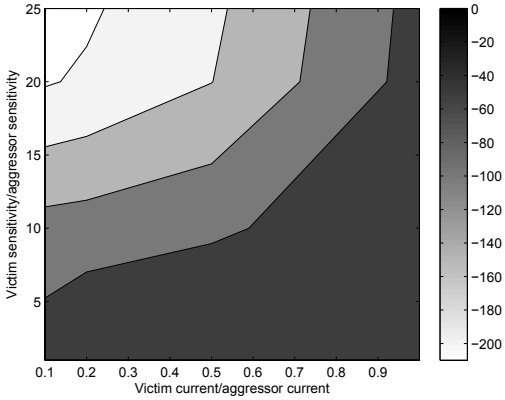

(b)

Figure 8: Change in voltage drop at the (a) victim and (b) aggressor circuit. The darker shade represents a greater reduction in the voltage drop at the victim and a lower increase in the voltage drop at the aggressor.

of the cases, voltage before and after application of the methodology, sensitivity, propagation delay, and improvement in the supply voltage and propagation delay are listed in Table 1.

Case 1 (Figs. 7(a) and 7(b)) illustrates the case where the current sunk by block 8 (the aggressor) is significantly higher as compared to the other circuit blocks. The highest degradation in the supply voltage is within the aggressor circuit; however, the voltage drop is reduced in those circuit blocks with a higher sensitivity and minimum delay, resulting in a reduction in the worst case delay and a higher maximum operating frequency. The increase in the supply voltage at block 1 is 2%, achieving 97% of the ideal power supply voltage and resulting in an improvement in the propagation delay of 5%. Note that the improvement in the propagation delay is greater than the supply voltage due to the high sensitivity factor. In Case 2 (Figs. 7(d) and 7(e)), the voltage at block 1 is increased by 5%, achieving 96% of the ideal power supply voltage and resulting in an improvement in the propagation delay of 17%.

6. DISCUSSION

The voltage drop within a power distribution network is investigated for circuit blocks with different current levels and sensitivities. The minimum propagation delay ($delay_{min}$) is maintained

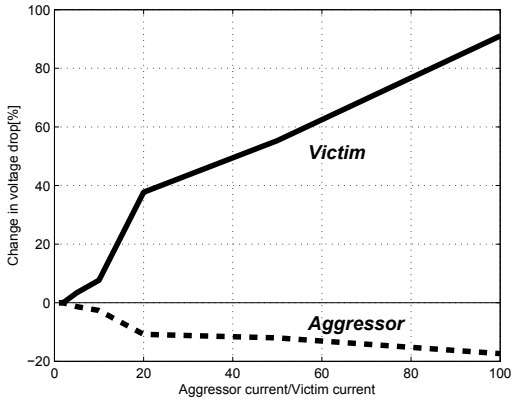

Figure 9: Change in voltage within the power distribution network for the victim and aggressor circuits as a function of the ratio of the current sunk by the aggressor and victim circuits. The sensitivity factor is assumed equal for both circuits.

the same. A 20×20 mesh structured power distribution network with two power supplies and two current sources (one aggressor and one victim) is considered. The voltage improvement at the victim and degradation at the aggressor are illustrated, respectively, in Figs. 8(a) and 8(b). Note that by assigning a higher sensitivity to the victim circuit, the voltage drop on the power network at the victim is reduced. Simultaneously, the voltage drop at the aggressor is increased, while the aggressor is less sensitive to voltage variations. The tradeoff between reducing the voltage drop at the victim while increasing the voltage drop at the aggressor is an important aspect of the proposed link breaking methodology.

The improvement and degradation of the voltage drop at, respectively, the victim and aggressor are depicted in Fig. 9 for different ratios of the current sunk by the victim and aggressor, assuming the two circuits have equal sensitivity. Note that a higher change in voltage is achieved at the victim when the current sunk by the aggressor is greater. This effect is due to the dominance of the aggressor on the victim circuit before applying the link breaking methodology.

The computational runtime of the algorithm, depicted in Fig. 5, is evaluated for differently sized power distribution networks. The algorithm has been executed on a Linux eight-core with 8 GB RAM system. The runtime as a function of the number of nodes in the power network is depicted in Fig. 10. The runtime of the link breaking methodology can also be accelerated by utilizing multigrid-like techniques [12] and ignoring those current sources located far from the target nodes. The number of aggressor and/or victim circuits is not a dominant factor affecting the runtime of the algorithm, as illustrated in Fig. 11. Initially, the runtime increases exponentially with the number of aggressor and victim circuits. With a further increase in the number of circuits, the computational runtime decreases due to the smaller number of links that can be disconnected. For those cases where only a small number of circuit are evaluated within a large power distribution network, the random walk method [11] can be used to estimate the voltage variations, significantly accelerating the link breaking methodology.

The worst case voltage drop (located at the aggressor) cannot be reduced by utilizing the link breaking methodology, since the methodology always increases the worst case power network impedance. However, the effect of the aggressor on other circuits with a higher sensitivity and propagation delay can be reduced, resulting in enhanced overall system performance.

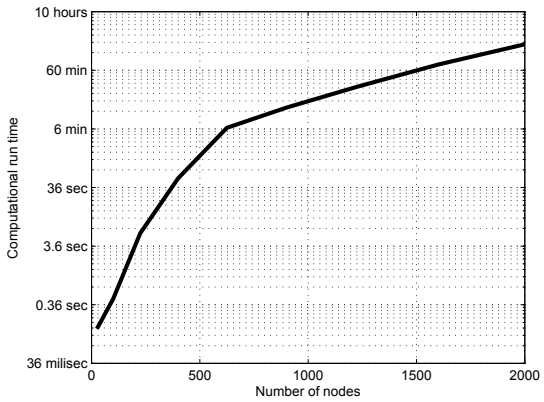

Figure 10: Computational runtime of the link breaking methodology as a function of the number of nodes within the power distribution network.

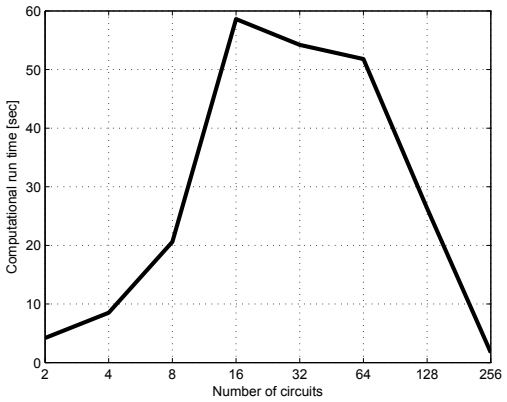

Figure 11: Computational runtime of the link breaking methodology as a function of the number of victim and aggressor circuits. The runtime initially increases with a higher number of circuits. After reaching a peak, the runtime decreases due to the smaller number of links that can be removed.

7. CONCLUSIONS

The design of the power distribution network is an essential part of an IC design flow. The network is typically designed as a single network or multiple separate networks. The advantages of a single network are a reduced network impedance and fewer routability constraints, while multiple separate networks have the advantage of lower noise coupling. The proposed link breaking methodology utilizes a single network, disconnecting links between the aggressive and sensitive circuits; thereby, isolating the victim from the aggressor. This approach reduces the noise, while maintaining a low network impedance.

Sensitivity to changes in the supply voltage will vary for different circuits. Voltage variations at the more sensitive circuits need to be reduced at the expense of increased voltage variations at the less sensitive circuits. A smaller voltage drop is also important in long critical paths as compared to shorter less critical logic paths. The aggressiveness and sensitivity of circuits are considered during the link breaking process. The methodology is evaluated for two cases. The objective for these case studies is reduced worst case propagation delay by increasing the supply voltage at blocks with high propagation delay. An average enhancement of 4% in power supply voltage at nodes with high sensitivity and high propagation delay is achieved, resulting in, on average, 96% of the ideal power supply voltage at these nodes. As a result, an average improvement of 11% in the maximum operating frequency is achieved when utilizing the proposed link breaking methodology.

8. REFERENCES

[1] D. Blaauw, R. Panda, and R. Chaudhry, "Design and Analysis of Power Distribution Networks," *Design of High-Performance Microprocessor Circuits*, Chandrakasan, Bowhill, and Fox, Eds., Chapter 24, pp. 499–522, IEEE Press, 2001.

[2] R. Jakushokas, M. Popovich, A. V. Mezhiba, S. Kose, and E. G. Friedman, *Power Distribution Networks with On-Chip Decoupling Capacitors, Second Edition*, Springer, 2011.

[3] C. Tirumurti, S. Kundu, S. Sur-Kolay, and Y.-S. Chang, "A Modeling Approach for Addressing Power Supply Switching Noise Related Failures of Integrated Circuits," *Proceedings of the IEEE Design, Automation and Test in Europe Conference and Exhibition*, Vol. 2, pp. 1078 – 1083, February 2004.

[4] L. H. Chen, M. Marek-Sadowska, and F. Brewer, "Buffer Delay Change in the Presence of Power and Ground Noise," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 11, No. 3, pp. 461–473, June 2003.

[5] M. Saint-Laurent and M. Swaminathan, "Impact of Power-Supply Noise on Timing in High-Frequency Microprocessors," *IEEE Transactions on Advanced Packaging*, Vol. 27, No. 1, pp. 135–144, February 2004.

[6] V. Kursun and E. G. Friedman, *Multi-Voltage CMOS Circuit Design*, John Wiley & Sons, Inc., 2006.

[7] X. Lai and J. Roychowdhury, "Fast, Accurate Prediction of PLL Jitter Induced by Power Grid Noise," *Proceedings of the IEEE Custom Integrated Circuits Conference*, pp. 121 – 124, October 2004.

[8] P. Larsson, "Measurements and Analysis of PLL Jitter Caused by Digital Switching Noise," *IEEE Journal of Solid-State Circuits*, Vol. 36, No. 7, pp. 1113 –1119, July 2001.

[9] L. A. Arledge Jr. and W. T. Lynch, "Scaling and Performance Implications for Power Supply and Other Signal Routing Constraints Imposed by I/O Pad Limitations," *Proceedings of the IEEE Symposium on IC/Package Design Integration*, pp. 45 –50, February 1998.

[10] M. Alioto and G. Palumbo, "Impact of Supply Voltage Variations on Full Adder Delay: Analysis and Comparison," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 14, No. 12, pp. 1322 –1335, December 2006.

[11] P. G. Doyle and J. L. Snell, *Random Walks and Electric Networks*, Mathematical Association of America, 1984.

[12] J. N. Kozhaya, S. R. Nassif, and F. N. Najm, "Multigrid-Like Technique for Power Grid Analysis," *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pp. 480–487, November 2001.