Versatile Framework for Power Delivery Exploration

Rassul Bairamkulov, Kan Xu, and Eby G. Friedman
Department of Electrical and Computer Engineering
University of Rochester
Rochester, New York, USA 14627
Emails: {rbairamk,kxu8}@ur.rochester.edu,
friedman@ece.rochester.edu

Mikhail Popovich, Juan Ochoa, and Vaishnav Srinivas
Qualcomm Inc.
San Diego, California, USA, 92121
Email: {mikhailp,jochoa,vaishnav}@qti.qualcomm.com

***Abstract*—Over the past decades, aggressive voltage scaling combined with increased power demands has placed stringent requirements on on-chip power quality. Unwanted voltage fluctuations and droops may cause a variety of issues, ranging from glitch power to device malfunction. If revealed at the later stages of the design process, mitigation techniques may become unbearably costly in both time and money. A framework for exploratory power delivery optimization is described to enhance the power delivery network during early stages of the design process in accordance with design specifications. The power delivery design process is converted into a constrained minimization problem, consisting of design metrics combined into objective and constraint functions. The framework supports the optimization of the power network characteristics while considering external, non-electrical design specifications, such as cost and area, providing a comprehensive network analysis capability. In one case study, a 15% reduction in decoupling capacitor placement along with a 38.6% reduction in power consumption is achieved while satisfying performance and power quality constraints.**

I. Introduction

The power delivery network plays a crucial role in system performance. With a conventional design flow, the power delivery characteristics are monitored throughout the system development process, ensuring that the parameters satisfy the required specifications. The discovery and mitigation of power delivery issues at later stages of the design process may however significantly hamper the product development flow, as little flexibility exists to make effective changes in the power network. The risk of compromising the IC design flow highlights the importance of exploration during early stages of the design process, before important parameters, such as the technology node or power network architecture, are fixed, while reducing the number of design iterations.

The large number of design parameters is an important issue in modern IC design flows. Due to an arbitrarily large number of parameters, an overconstrained system may develop, resulting in suboptimal system performance. Furthermore, a conventional design flow does not incorporate cost – an important design metric. This approach may deviate from the optimal parameters, leading to either a more expensive system with unnecessarily good performance or a system that does not satisfy the target design objectives.

A large amount of work had been described in the literature to improve the design of on-chip power networks. Optimization of IC power regulation is an extensively researched topic. In [1], for example, the authors propose two methods for optimizing DC-to-DC conversion in a smartphone, achieving a 19% reduction in power. In [2], an efficient nonlinear optimization strategy is described for minimizing the number of on-chip low-dropout regulators and the supply voltage to maximize power efficiency. A more high-level study of power regulation is presented in [3], where an efficient placement methodology of switching and low dropout regulators is discussed, achieving a heterogeneous power delivery system with up to 93% efficiency. Decoupling capacitor placement is another topic related to power network optimization. In [4], a distributed power delivery system with decoupling capacitors is optimized, achieving a 45.2% reduction in voltage drop. In [5], a link breaking methodology is proposed that sacrifices power quality of the least sensitive circuits while increasing the maximum operating frequency and reducing noise at the most sensitive nodes by placing decoupling capacitors in close proximity.

Although on-chip power delivery is a well researched problem, a significant omission in the literature is the almost exclusive focus on only optimizing the electrical parameters, only indirectly addressing external metrics. For example, manufacturing cost is an important external metric during the circuit design process. The work described in this paper closes this gap by proposing a methodology to optimize the power network during early stages of the development process, considering both external non-electrical and electrical parameters. The rest of the paper is organized as follows. The primary steps of the proposed optimization framework are described in Section II. An example of power network cost optimization which considers power consumption, power quality, and clock frequency constraints is described in Section III. The advantages and challenges of the proposed approach are discussed in Section IV. In Section V, a summary of this paper is provided.

II. Optimization Procedure

The power delivery network is a complex system for delivering a constant voltage to a vast number of time dependent loads. A simplified power network design process is shown in Fig. 1 [6]. During the preliminary analysis stage, rough

This research is supported in part by the National Science Foundation under Grant Nos. CCF-1329374, CCF-1526466, and CCF-1716091, IARPA under Grant No. W911NF-14-C-0089, AIM Photonics under Award No. 059447-007, the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI), Singapore Ministry of Education Tier 2 under Grant No. MOE2014-T2-2-105, and by grants from Cisco Systems, Qualcomm, and OeC.

978-1-5386-4881-0/18/$31.00 ©2018 IEEE

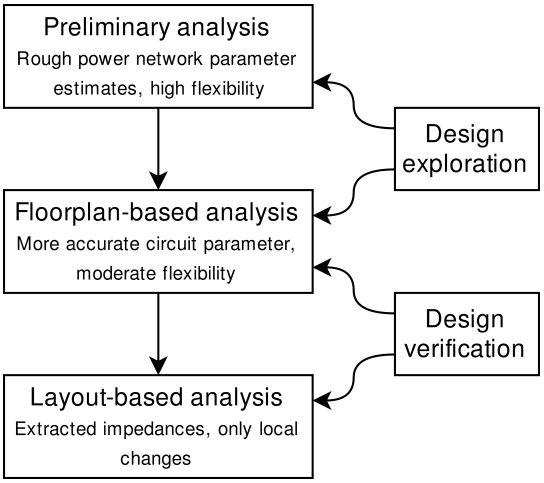

Fig. 1. Flow of power network design process [6]

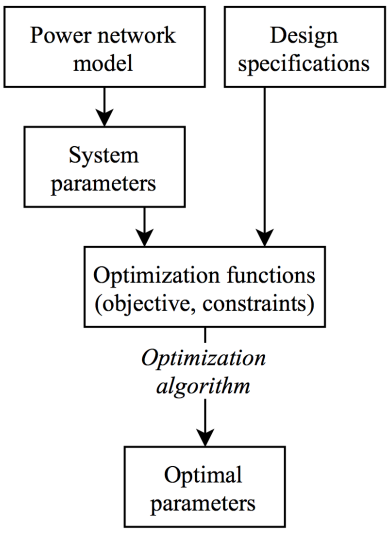

Fig. 2. Flow of system parameter optimization process

estimates of the primary power network parameters such as the current consumption of the circuit blocks and wire pitch are determined. Design exploration at this stage has the highest degree of flexibility, as most of the design characteristics such as the location of the circuit blocks are undecided. The decisions however are complicated by the lack of accurate power network parameters. During the next design step, a floorplan of the system is determined, allowing certain parameters to be extracted with higher accuracy. For example, knowing the relative location of the circuit blocks allows the length of the global interconnects to be determined. Although fixing the parameters at this stage makes the power network design process less flexible, the increased accuracy improves the quality of the analysis and exploration process. At later stages of the power network design process, the layout of the system is determined. Accurate circuit parameters are available and only minor local changes are possible.

The framework proposed in this paper is best suited for the preliminary and floor-based analysis stages, where major changes to the circuit topology and system architecture can be made to enhance the performance of the power network. The optimization procedure is summarized in Fig. 2. The power delivery network is initially analyzed to estimate the appropriate design parameters and identify degrees of freedom. Along with the design specifications, the optimization functions are generated using estimated data and passed to the optimization algorithm.

A general optimization problem can be described as

$$x_{opt} = \min f(x), \quad \text{subject to} \quad c(x) \leq 0, \tag{1}$$

where x is a vector of parameters, x_{opt} is a vector with corresponding optimal parameters, $f(x)$ is the objective function, and $c(x)$ is a vector of constraint functions. Formulating the power delivery design problem as in (1) enables the use of general optimization methods to determine the optimal parameters. Formulation of the power network optimization problem begins with identifying those variables that provide design freedom and form the vector x in (1). Examples of variable parameters in the power network design process are the number and size of the decoupling capacitors, while the ball grid array and controlled collapse chip connection pitches at printed circuit board (PCB) and package levels are typically fixed.

The primary purpose of parameter estimation is evaluating the relevant design metrics. Two types of metrics are identified: electrical and non-electrical. The electrical metrics require circuit simulation using the estimated parameters and include power quality, total power, and voltage droop. To evaluate these metrics, an interface between the optimization tool (MATLAB Optimization toolbox [7]) and simulation tool (HSPICE [8]) has been developed. The second category of design metrics is non-electrical, and includes area and cost. These parameters do not typically require circuit simulation and can be determined from user specified parameters and functions.

Several design metrics are combined within a single objective function to be minimized and are denoted as $f(x)$ in (1). An example of this measure is a weighed sum of maximum clock frequency and cost. Once developed, the function is passed to the optimization algorithm, allowing the minimum value to be determined under specific constraints. For example, suppose that the power consumption of a particular system needs to be minimized given frequency and area constraints, using the wire size and supply voltage as variable parameters. By applying the optimization algorithm, the supply voltage is iteratively reduced, making the system slower until the critical path delays are comparable to the target clock period. The wire cross section is simultaneously increased to reduce the impedance and therefore the power loss in the wires until the area limit is reached. A more complete example of a power network optimization process is described in Section III.

III. Case Study

A case study is performed to demonstrate the flow of the proposed method. A typical power network represented by serially cascaded RL branches and parallel RLC branches is shown in Fig. 3. The cost is assumed to be minimized, subject

978-1-5386-4881-0/18/$31.00 ©2018 IEEE

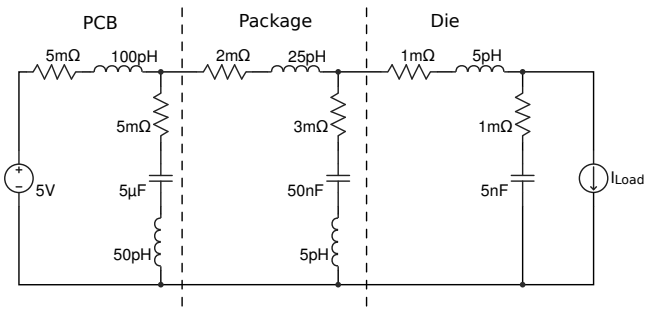

Fig. 3. Model of 1-D power delivery network with initial parameters

to power consumption, power quality, and frequency requirements. In subsection III-A, estimation and classification of the power network parameters are discussed. In subsection III-B, the optimization functions are defined. The optimization procedure and results are described in subsection III-C.

A. Parameter Estimation and Classification

A three-level power network including the PCB, package, and die levels is considered here. The series resistance and inductance of the power network are assumed fixed. The on-die parallel inductance is neglected assuming point-of-load on-die decoupling capacitors with small inductance [3]. The profile of the load current has been adapted from [9] and shown in Fig. 4(a). The load current profile models the fluctuations of the workload during system operation. The supply voltage is used as a design variable to explore the effects of supply voltage on system performance. Other controllable parameters are the number and magnitude of the decoupling capacitors within the PCB, package, and die levels.

Important tradeoffs can be noted [6]: a higher supply voltage enhances the speed but significantly increases the power consumption. Insertion of parallel decoupling capacitances is a powerful technique for reducing ripple currents since the high frequency components of the current bypass the load. Larger decoupling capacitors, however, require significant on-chip area, leading to greater system cost.

B. Definition of Optimization Functions

The cost of each system level (PCB, package, die) is assumed to be a function of the physical area which is affected by the area of the decoupling capacitors. The decoupling capacitor placement cost Q_{die} is

$$Q_{die} = w_{die} A_{die}, \tag{2}$$

where A_{die} is the area of the on-chip decoupling capacitor and w_{die} [$\$/m^2$] is the cost of the unit on-die area. The total cost of the decoupling capacitors is therefore

$$Q = \frac{1}{\varepsilon_0} \sum_{i \in S} \frac{w_i C_i d_i}{\varepsilon_i}, \tag{3}$$

where S is the set of levels in the system (e.g., PCB, package, and die), ε_0 is the permittivity of free space, C_i is the parallel plate capacitance at level i, and d_i and ε_i are, respectively, the insulator thickness and relative permittivity at level i.

TABLE I
PARAMETERS OF DECOUPLING CAPACITOR COST

	Die	Package	PCB
Cost per m^2, normalized	20.25	4.5	1
Insulator thickness	0.9 nm [10]	12 μm [11]	250 μm [12]
Insulator permittivity	3.9 [10]	4.6 [11]	4.5 [12]

The oxide thickness and dielectric constant are described in [10]–[12]; however, the cost per area is not as clear. Based on the review of publicly available cost information [13]–[17], the cost per unit area of a package is approximately 3 to 6 times greater than the cost of unit PCB area, and approximately 3 to 10 times lower than the cost of unit die area. To simplify the cost estimate, the cost per unit area [m^2] of a PCB is normalized to 1, the package area cost is assumed to be 4.5, and the cost per unit on-die area is assumed to be 20.25, 4.5 times greater than the cost per unit area of the package. The normalized cost estimates used in this case study are listed in Table I.

The target constraint metrics are power consumption, power quality, and speed. The power consumption is directly measured through simulation, and the corresponding constraint function is

$$c_1(x) = P - P_{max}, \tag{4}$$

where $c_1(x)$ is the initial constraint function, P is the measured power, and P_{max} is the upper bound on the power consumption. Since the constraint function is negative, (4) ensures that the power dissipation does not exceed the maximum allowable power level.

For frequency, the constraint is

$$t_{p,CP} \leq T_{min}, \tag{5}$$

where $t_{p,CP}$ is the propagation delay of the critical path and T_{min} is the lower bound on the clock period. Evaluation of this metric, however, is computationally expensive and requires identification of the critical paths and extensive parameter extraction. In this case, accuracy is sacrificed for higher computational efficiency. The voltage at the load is, therefore, used as the speed metric,

$$c_2(x) = V_{min} - min(V_L(t)), \tag{6}$$

where $V_L(t)$ is the instantaneous voltage at the load, and V_{min} is the minimum voltage to maintain reliable high speed operation.

The third design constraint is power quality, described as voltage fluctuations, and is formulated as

$$c_3(x) = \frac{max(V_L(t)) - min(V_L(t))}{V_{rail}} - \Delta V_{max}, \tag{7}$$

where V_{rail} is the supply voltage, and ΔV_{max} is the maximum allowed fluctuation. The optimization constraints are listed in columns two and three of Table II.

C. Optimization Results

The Interior Point Algorithm, part of MATLAB Optimization Toolbox [7] and HSPICE [8], is used in this case study. The optimization functions, circuit parameters, and external

978-1-5386-4881-0/18/$31.00 ©2018 IEEE

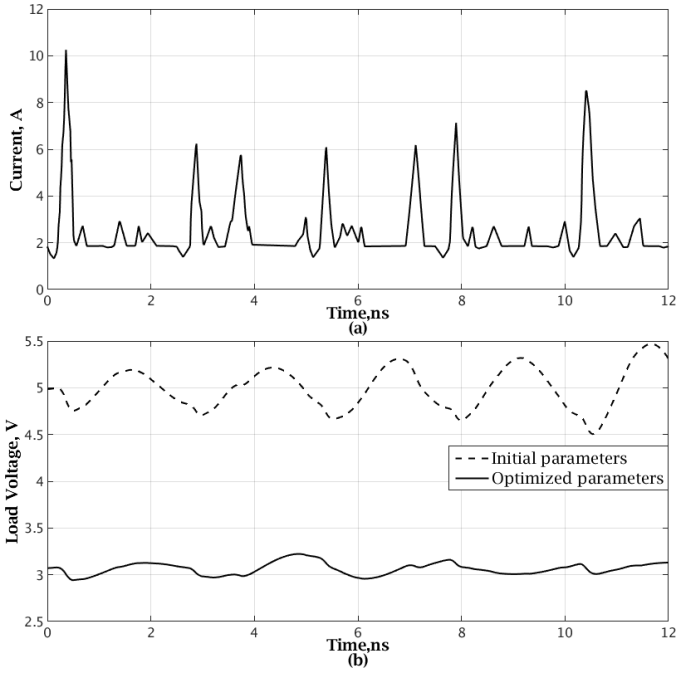

Fig. 4. Waveform of power network, a) load current adapted from [9], and b) load voltage with initial and optimized parameters

parameters are inputs to the optimization algorithm. The optimization procedure has been run on an Intel Core i7-6700 3.40 GHz 8-core computer using different initial conditions to avoid any local minima. The initial parameters that produce the lowest cost under specified constraints are listed in column four of Table II.

The optimization process is completed in 28 seconds, requiring 66 function evaluations to converge. The load voltage waveforms are shown in Fig. 4(b). The power network initially exhibits an underdamped response, resulting in relatively large droops and overshoots. After optimization, the voltage fluctuations are reduced in the optimized power network by choosing an appropriate decoupling capacitor. The reduction in the load voltage fluctuations allows the supply voltage to be scaled since fluctuations are less likely to drop below the minimum allowed level. Reducing the supply voltage, in turn, leads to lower power dissipation.

The optimization results are listed in column five of Table II. As compared to the initial suboptimal parameters, the cost has decreased by almost 15% from 0.317 to 0.270. The initial parameters do not satisfy the power dissipation and load voltage constraints. A 38.6% reduction in power consumption is achieved, from 10.6 watts to 6.51 watts. Most of the reduction in power originates from the reduced supply voltage, from 5 volts to 3.09 volts. In addition, a 53% decrease in fluctuations is achieved, from 19.3% to 9.07%. As a result, the optimized parameters satisfy the target requirements, including the power and voltage constraints.

IV. Discussion

This case study demonstrates the potential of the proposed tool. Given design specifications and estimated parameters, the power network design process is formulated as a nonlinear optimization problem solved using fast and robust algorithms. The primary advantage of the proposed method is the flexibility in the choice of parameters and functions. This approach constructs optimization functions using a combination of electrical and external parameters to achieve a more comprehensive power network analysis process. Another advantage of this procedure is the ability to handle any circuit topology. Arbitrarily complex circuits can be analyzed given ample computational time. Finally, the speed of the algorithm is controlled by adjustable tolerances, offering a tradeoff between accuracy and computational efficiency.

TABLE II
Optimization constraints, with initial and optimal parameters

Parameter/Metric	Lower Bound	Upper Bound	Initial Value	Optimized Value
Supply voltage	1.4 volts	10.0 volts	5.0 volts	3.09 volts
PCB decap	25.0 nF	10.0 μF	5.00 μF	2.71 μF
Package decap	50.0 pF	100 nF	50.0 nF	9.77 nF
Die decap	2.00 pF	10.0 nF	5.00 nF	9.32 nF
Minimum load voltage	1.40 volts	—	2.96 volts	2.94 volts
Power dissipation	—	10.0 watts	10.6 watts	6.51 watts
Load voltage	—	10.0%	19.3%	9.07%
Normalized cost	—	—	0.317	0.270

Several limitations of the proposed framework exist. To evaluate the external metric, a function for that metric needs to be provided, requiring appropriate assumptions. Another limitation of the proposed methodology is the dependence on simulation time. With a significant increase in the number of nodes and tightening of the constraints, the optimization time may dramatically increase. More efficient circuit solvers, such as [18], may improve the simulation time by several orders of magnitude. Finally, the quality of the results depends upon the initial values and the optimization algorithm. If the objective function exhibits many local minima, the optimization algorithm may produce a suboptimal solution. The optimization method and initial values, therefore, need to be carefully chosen.

V. Conclusions

A flexible procedure and framework for early power delivery design exploration is presented in this paper. The framework supports power delivery optimization, including both electrical and external parameters. The advantages of the method include versatility of application and flexibility in circuit topology. The limitations include the dependence on the simulation engine and the need for external parameter estimation. With these methods, a more comprehensive power network analysis process is produced. The validity of the method is demonstrated by a case study, where the cost of the power delivery network is reduced by 15% while achieving target power consumption and voltage fluctuation requirements.

References

[1] W. Lee, Y. Wang, D. Shin, N. Chang, and M. Pedram, "Optimizing the Power Delivery Network in a Smartphone Platform," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 33, No. 1, pp. 36–49, January 2014.

978-1-5386-4881-0/18/$31.00 ©2018 IEEE

[2] Z. Zeng, X. Ye, Z. Feng, and P. Li, "Tradeoff Analysis and Optimization of Power Delivery Networks with On-Chip Voltage Regulation," *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 831–836, June 2010.
[3] I. Vaisband and E. G. Friedman, "Heterogeneous Methodology for Energy Efficient Distribution of On-Chip Power Supplies," *IEEE Transactions on Power Electronics*, Vol. 28, No. 9, pp. 4267–4280, September 2013.
[4] S. Kose and E. G. Friedman, "Distributed On-Chip Power Delivery," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, Vol. 2, No. 4, pp. 704–713, December 2012.
[5] R. Jakushokas and E. G. Friedman, "Power Network Optimization Based on Link Breaking Methodology," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 21, No. 5, pp. 983–987, May 2013.
[6] I. P. Vaisband, R. Jakushokas, M. Popovich, A. V. Mezhiba, S. Köse, and E. G. Friedman, *On-Chip Power Delivery and Management*, Springer International Publishing, 2016.
[7] MathWorks, "Optimization Toolbox User's Guide", September 2017.
[8] Synopsys, "HSPICE Quick Reference", March 2017.
[9] B. Ko, J. Kim, J. Ryoo, C. Hwang, J. Song, and S. W. Kim, "Simplified Chip Power Modeling Methodology without Netlist Information in Early Stage of SoC Design Process," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, Vol. 6, No. 10, pp. 1513–1521, October 2016.
[10] H. Iwai, "Future of Integrated Devices," *Proceedings of the International Workshop on Junction Technology*, p. 43, June 2015.
[11] J. Kim, "Power Integrity of SiP (System In Package)," *IEEE Video Distinguished Lecturer Program*, August 2010.
[12] B. Archambeault, "Effective Power/Ground Plane Decoupling for PCB," *IEEE Video Distinguished Lecturer Program*, October 2007.
[13] Or-Bach, Zvi, "Moore's Law Has Stopped at 28nm," March 2014. [Online]. Available: http://electroiq.com/blog/2014/03/moores-law-has-stopped-at-28nm/. [Accessed: 2018-01-25].
[14] Eric Esteve, "Why SOI is the Future Technology of Semiconductor," January 2014. [Online]. Available: https://www.semiwiki.com/forum/content/3077-why-soi-future-technology-semiconductor.html. [Accessed: 2018-01-25].
[15] Anysilicon.com, "IC Package Price Estimator," [Online]. Available: http://anysilicon.com/package-price-estimator/. [Accessed: 2018-01-25].
[16] Pcbshopper.com, "A Price Comparison Site for Printed Circuit Boards," [Online]. Available: https://pcbshopper.com/. [Accessed: 2018-01-25].
[17] Pcbcart.com, "Printed Circuit Board Calculator," [Online]. Available: https://www.pcbcart.com/quote. [Accessed: 2018-01-25].
[18] J. Yang, Z. Li, Y. Cai, and Q. Zhou, "PowerRush: A Linear Simulator For Power Grid," *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pp. 482–487, November 2011.

978-1-5386-4881-0/18/$31.00 ©2018 IEEE