Title
A Stabilizing Centralized Controller for On-Chip Power Delivery Networks

Journal
IEEE Transactions on Circuits & Systems II Express Briefs, 67(4)

ISSN
1549-7747

Authors
Riad, Joseph
Li, Peng
Sánchez-Sinencio, Edgar

Publication Date
2020-04-01

DOI
10.1109/tcsii.2019.2926670

Peer reviewed

A Stabilizing Centralized Controller for On-Chip Power Delivery Networks

Joseph Riad, *Member, IEEE*, Peng Li, *Fellow, IEEE*, and Edgar Sánchez-Sinencio, *Fellow, IEEE*

***Abstract*—A major challenge currently facing integrated systems, such as microprocessors and systems-on-chip is power delivery. With increasing system size and complexity, multiple voltage regulators are often needed for a single power domain. This connection raises non-trivial stability issues that often render the system useless. It is therefore of great practical interest to tackle the stability of such systems. This brief proposes the design of a centralized controller with the aim of rendering such power delivery networks (PDNs) stable without compromising their performance. Theoretical results are confirmed with extensive transistor-level simulations and demonstrate that the controller is capable of stabilizing PDNs with much greater area efficiency than conventional techniques.**

***Index Terms*—MIMO controller, power delivery network, LDO.**

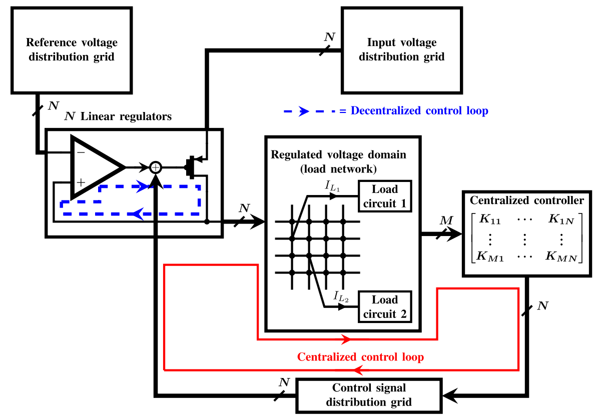

Fig. 1. Proposed centralized control structure for PDNs.

I. Introduction

BECAUSE of the ever-increasing integration density for electronic components, more and more functions are being integrated onto single chips like systems-on-chip (SoC) and microprocessors. Such systems have complicated power requirements due to the vast extent of components operating at different times and with different frequencies. It has been shown that connecting several low drop-out (LDO) linear regulators to the same power domain (see Fig. 1) can yield performance improvements in terms of load current handling capability and rejection of unwanted variations in the supply voltage (supply noise) [1]–[4]. Due to the advantages of integrated power delivery networks (PDN), they have garnered a lot of interest in the literature lately [1]–[9]. On the other hand, connecting multiple LDOs to the same PDN runs the risk of instability. This instability is due to excessive capacitance in the load network and the loading effect of other LDOs. This shifts the non-dominant poles of the system to lower frequencies, leading to phase margin loss.

Manuscript received June 4, 2019; accepted June 21, 2019. Date of publication July 9, 2019; date of current version April 1, 2020. This work was supported in part by the National Science Foundation under Grant ECCS-1810125, in part by Texas Instruments, and in part by Silicon Labs. This brief was recommended by Associate Editor Y. Zhang. *(Corresponding author: Joseph Riad.)*

J. Riad and E. Sánchez-Sinencio are with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843 USA (e-mail: joseph.riad@tamu.edu).

P. Li is with the Department of Electrical and Computer Engineering, University of California at Santa Barbara, Santa Barbara, CA 93106 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCSII.2019.2926670

Existing works in the literature tackle the stability problem by modifying the design parameters of individual LDOs such as device sizes, bias currents and/or compensation capacitors [3], [5]. Such solutions therefore tackle the global issue of PDN stability at the local level of LDO regulation loops. In some cases, as will be shown in this brief, such solutions may only be able to stabilize the network at the cost of a large area expenditure. This brief presents an attempt to explore a different approach to the problem by retro-fitting an existing PDN with a stabilizing centralized controller that manages to stabilize the network in cases where the approach in [3] requires the use of inordinately large capacitors.

To provide a motivation for this approach, note that each LDO inherently contains its own control loop. This means that connecting multiple LDOs to the same power domain essentially amounts to decentralized control of the multi-input (gates of multiple LDO power transistors) multi-output (multiple LDO outputs) (MIMO) PDN [10]. This is a good control strategy for MIMO systems with weak coupling among their constituent subsystems. For systems exhibiting strong interactions as is the case of a PDN, centralized control may be a better option [10], which is explored in this brief. The centralized controller, sensing the PDN outputs and controlling the LDOs is illustrated in Fig. 1. Note that the gains $K_{i,j}$ may, in general, be functions of frequency.

To demonstrate the viability of the centralized control technique, this brief starts by reviewing different control techniques and comparing them in the context of PDN stabilization. The modeling approach is explained briefly then

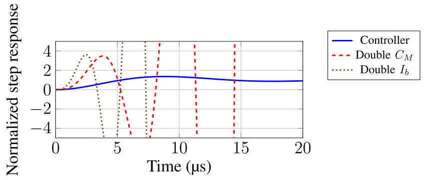

Fig. 2. Normalized step responses of a 900-node, 4-LDO network using different stabilization approaches.

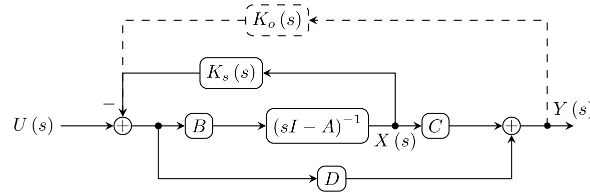

Fig. 3. Feedback control categories: state feedback (the solid path through K_s) and output feedback (the dashed path).

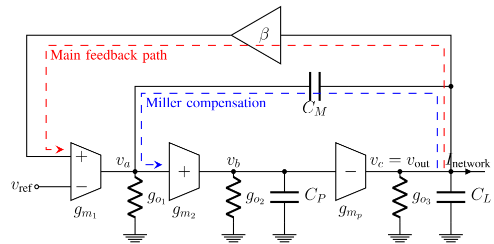

Fig. 4. Small-signal model of the LDO.

the proposed controller is discussed in detail with intuitive explanations for its stabilizing action. The effectiveness of the controller is demonstrated through transistor-level simulations. To the authors' knowledge, this is the first attempt to design a MIMO controller for stabilizing a PDN.

II. Comparison of Control Approaches

A. Decentralized vs. Centralized

Decentralized control is equivalent to using a diagonal gain matrix in the controller of Fig. 1, which is the conventional method adopted: a local loop in each LDO. This control structure may not be able to stabilize a PDN whose instability mainly arises from the strong coupling among LDOs. Centralized control accounts for these interactions by design and is therefore better able to deal with such cases. In a PDN, the coupling is strong by design since the load network needs to have a very low resistance, which causes the output of one LDO to strongly affect that of the others. As a proof of concept, a numerical model of a 900-node 4-LDO network (known to be unstable) was used for simulation. Three approaches are attempted to stabilize it: using a feedback controller like that shown in Fig. 1 (see later), doubling each LDO's compensation capacitance to increase its phase margin by reducing its bandwidth (10× more capacitance than the controller), and doubling each LDO's bias current to increase its phase margin by pushing non-dominant poles to high frequencies (30% more power than the controller). It is found that only the first approach manages to stabilize the network as shown in the normalized step responses of Fig. 2 which correspond to numerical simulation of linearized (small-signal) system models. This demonstrates the efficiency of centralized control in this case and the fact that it manages to stabilize systems where decentralized control techniques fail to do so. Next, two possible implementations of centralized control are compared.

B. State Feedback vs. Output Feedback

Feedback control may be divided into two categories: output feedback control and state feedback control. These two concepts are illustrated for a linear system in Fig. 3 where A, B, C and D are the state-space matrices of the system and $K_s(s)$ and $K_o(s)$ represent the controller transfer function for state feedback and output feedback respectively.

State feedback achieves arbitrary pole placement if the system satisfies the requisite controllability and observability conditions [11]. State feedback is, however, impractical in this case due to the sheer number of system states (nodes on the PDN) rendering the controller prohibitively complicated.

On the other hand, output feedback control involves only the system outputs, a much smaller number: the number of LDOs. Output feedback, however, cannot place the poles arbitrarily and it is not an easy problem to determine whether a system is stabilizable by output feedback [12]. That being said, there are cases where an output feedback controller with frequency-independent gain can stabilize an unstable system [12].

Based on this discussion, this brief focuses on designing a centralized output feedback controller for stabilizing the PDN.

III. System Modeling

The LDO used to validate the results in this brief consists of a three-stage structure, comprising a Miller-capacitance-compensated error amplifier driving a p-type metal-oxide-semiconductor field effect transistor (PMOSFET) power device as modeled by the linearized (small-signal) block diagram of Fig. 4 where g_{m_p} represents the power device with C_P as its gate capacitance and I_{network} represents the current going to the PDN. The conductance g_{o_i} represents the output conductance of stage i and the β block models the feedback factor used to set the output voltage.

Using this model and setting I_{network} to zero, circuit analysis yields an expression for the loop gain of a single LDO:

$$L(s) \simeq \beta \cdot \frac{g_{m1} g_{m2} g_{m3} \left(1 - \frac{g_{o2} C_M}{g_{m2} g_{m3}} s - \frac{C_M C_P}{g_{m2} g_{m3}} s^2\right)}{g_{o1} g_{o2} g_{o3} \left(1 + \frac{s}{\omega_{pd}}\right)\left(1 + \frac{s}{\omega_0 Q} + \frac{s^2}{\omega_0^2}\right)}$$
$$\omega_{pd} \simeq \frac{g_{o1} g_{o2} g_{o3}}{g_{m2} g_{m3} C_M} \quad \omega_0 \simeq \sqrt{\frac{g_{m2} g_{m3}}{C_L C_P}} \quad Q \simeq \frac{C_M}{g_{o1}} \sqrt{\frac{g_{m2} g_{m3}}{C_L C_P}} \tag{1}$$

To find I_{network}, the PDN must be modeled as well. The PDN represents the grid of wires used to deliver power to the load circuits [2], [3]. A simplified model of the PDN

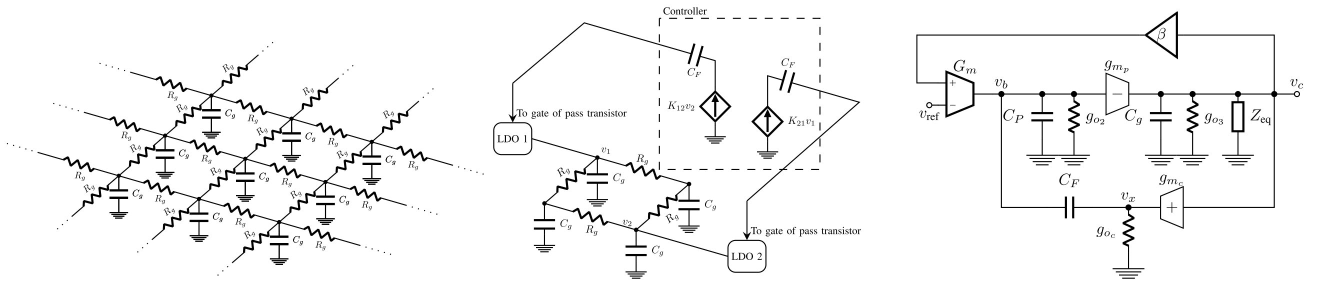

(a) (b) (c)

Fig. 5. Load network: (a) General case (b) Example network with controller (c) Small-signal model of LDO with controller.

was chosen (Fig. 5(a)) following the approach in [3] to capture its most salient features: inter-node resistance (R_g) and nodal capacitance (C_g). This ignores parasitic inductance and ground voltage variations to make the analysis more tractable and insightful without losing too much accuracy.

IV. Proposed Stabilizing Controller

A. Controller Structure

Since the controller requires the addition of signals, it is simpler to implement it using currents (in transconductance-mode). The controlled input should be chosen such that it has a strong influence on the output voltage, thereby needing small controller gain. One choice is the power device gate as this directly affects the output voltage of the LDO. This way, the currents produced by many transconductors can be summed and converted to a voltage at the gate of the power device.

The proposed controller is illustrated for a 2-LDO PDN in Fig. 5(b) where K_{12} and K_{21} are implemented by transconductors and C_F is used to prevent the controller from disturbing the steady state of the LDOs. For the general case with N LDOs, the gate of each power device will receive feedback based on the output voltages of all the other $(N-1)$ LDOs.

The intuition behind this controller structure is that it manipulates the power device of each LDO based on the load current needs of all the LDOs in the network. More intuition may be gained about the benefit of using the controller by considering the case of 2 LDOs connected to a four-node network as shown in Fig. 5(b). In this case, by symmetry, the output voltages of both LDOs are the same[1] and we may use the simplified small-signal model in Fig. 5(c) to model the behavior of either LDO, where G_m models the Miller-compensated error amplifier, g_{m_c} and C_F model the controller, and Z_{eq} is the equivalent impedance of the load network.[2] From Fig. 5(c), the following model equations can be written

$$\begin{aligned} G_m(\beta v_c - v_{\text{ref}}) &= (v_b - v_x)sC_F + v_b\big(g_{o_2} + sC_p\big) \\ -g_{m_p} v_b &= v_c\big(g_{o_3} + sC_g + Y_{\text{eq}}\big) \\ g_{m_c} v_c &= v_x\big(g_{o_c} + sC_F\big) - v_b sC_F \end{aligned} \tag{2}$$

[1]This assumes that both LDOs see equal loading currents. While not always true, this assumption simplifies the analysis in order to provide insight.

[2]This model neglects the zeros in (1) since they are at much higher frequencies $\left(\frac{\sqrt{g_{m_2} g_{m_3}}}{\beta g_{m_1}} \cdot \sqrt{\frac{C_M}{C_P}} \simeq 30\times\right)$ compared to the unity-grain frequency. The numerical model used for design takes the zeros into account.

TABLE I
Pole and Zero Locations for a 2-LDO Network

Quantity	Without the controller	With the controller
$\omega_{z_1}, \omega_{z_2}$	$\infty,\ 2/\tau_g$	$g_{o_c}/C_F,\ 2/\tau_g$
ω_0	$\sqrt{\frac{\beta G_m g_{m_p}}{2C_g C_P}}$	$\sqrt{\frac{\beta G_m g_{m_p}}{2C_g C_P}}$
Q	$\sqrt{\frac{\beta G_m g_{m_p}}{2C_g C_P}} \cdot \frac{C_P}{g_{o_2}} = Q_0$	$Q_0 \cdot \frac{2C_g}{C_F} \cdot \frac{g_{o_c} g_{o_2}}{g_{m_p} g_{m_c}}$
$\omega_{p_3}, \omega_{p_4}$	$\infty,\ 4/\tau_g$	$g_{o_c}/C_F,\ 4/\tau_g$

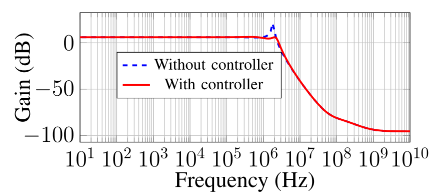

Fig. 6. Peaking reduction by the controller in a 2-LDO network.

It can be shown that $Y_{\text{eq}} = 2sC_g/(2+s\tau_g)$ where $\tau_g = R_g C_g$. Thus the closed-loop transfer function from the reference input of an LDO to its output is given by

$$H_{\text{CL}}(s) = \frac{1/\beta\big(1 + s/\omega_{z_1}\big)\big(1 + s/\omega_{z_2}\big)}{\big[1 + s/(\omega_0 Q) + s^2/\omega_0^2\big]\big(1 + s/\omega_{p_3}\big)\big(1 + s/\omega_{p_4}\big)} \tag{3}$$

where Table I shows the approximate locations of the poles and zeroes with and without the controller. The table demonstrates that the controller leaves the transfer function unchanged except for reducing Q (for large enough g_{m_c} and C_F) as the extra pole and zero cancel.[3] This Q reduction leads to reduced gain peaking and therefore improved stability as demonstrated in Fig. 6.

For the general case, the small-signal model is used for controller design. For this purpose, a numerical model is developed wherein the controller parameters g_{m_c} and C_F are swept and the system pole with the largest real part is used to assess stability. This can be used to determine the stabilizing range of values for each parameter. Within this range, circuit simulations are used to select parameter values that yield

[3]In reality, ω_0 is changed as well but, with $C_F \ll C_g, C_P$, this change is negligible.

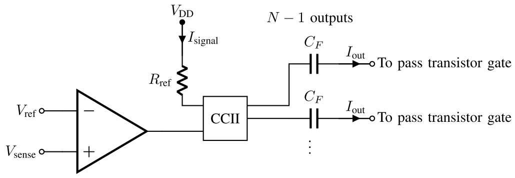

Fig. 7. Controller unit cell repeated for each of the N LDOs.

TABLE II
TESTING CASES. IN EACH CASE, $C_g = C_{g_t}/900$

		4 LDOs	8 LDOs	16 LDOs
C_{g_t}(nF)	5	Case 1	Case 3	Case 5
	45	Case 2	Case 4	Case 6

TABLE III
CONTROLLER PARAMETERS FOR THE 6 DIFFERENT CASES

Case number	1	2	3	4	5	6
R_{ref} (kΩ)	150	100	300	100	500	70
C_F (pF)	0.15	0.5	0.2	0.3	0.1	0.2

a settling time of less than 10 µs. In this process, the complexity of the eigenvalue analysis can be reduced by using a reduced-order model for the PDN.

B. Circuit Implementation

To make the controller implementation scalable, the sensed voltages are first amplified using a folded cascode amplifier and then a second-generation current conveyor (CCII) [13] is used as a transconductor. This way, the transconductance of the CCII can be small compared to the required g_{m_c}. The output of the CCII can then be mirrored as needed to send feedback signals to all the LDOs. This is shown in Fig. 7. Note that R_{ref} defines g_{m_c} and can be connected off-chip to allow the controller to be tuned post-silicon.

V. SIMULATION RESULTS

To validate the proposed approach, several PDNs are tested using transistor simulations in a 0.5 µm CMOS technology. Two things are varied: total grid capacitance and LDO number. The PDN is a 900-node square grid. Table II lists the 6 cases.

For each case, the numerical model was used to pick values for controller parameters; these are summarized in Table III. It should be noted that $g_{m_c} = 1000/R_{ref}$ where the factor of 1000 is the gain provided by the folded cascode amplifier.

A. LDO Arrangements

Fig. 8 shows the different arrangements of LDOs around the grid for the different cases. The loads used in the load sharing test will be addressed in a later subsection.

B. Controller Efficiency

To demonstrate the efficiency of the controller, each of the networks described in Table II is tested with a load current connected near the center of the grid and pulsed from 0 to

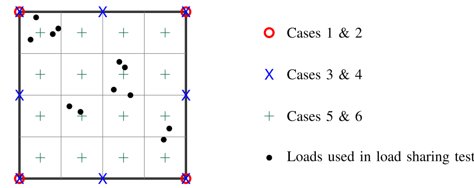

Fig. 8. Different arrangements of LDOs around the grid.

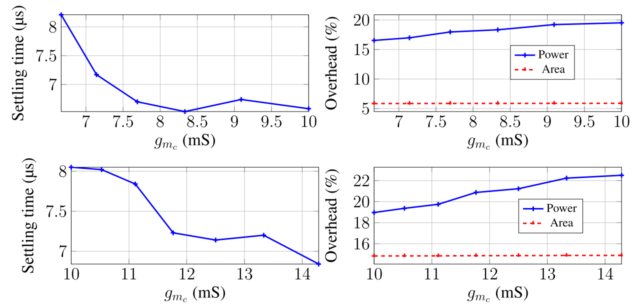

Fig. 9. Settling time and controller overheads as functions of controller gain for cases 1 (top row) and 2 (bottom row).

TABLE IV
EFFICIENCY OF THE PROPOSED CONTROLLER

Case	Proposed approach			Increasing C_M	
	C_{extra} (pF)	Overhead Area	Overhead Power	C_{extra} (pF)	Overhead Area
1	0.45	5.85%	16.54%	25	213%
2	1.5	14.84%	18.97%	37[a]	316%
3	1.4	14.07%	21.09%	80	683%
4	2.1	20.15%	27.66%	26[a]	222%
5	1.5	15.13%	26.45%	40	341%
6	3	28.07%	36.81%	16[a]	137%

[a] C_{extra} only stabilizes the network; it is not possible to achieve the same performance as the proposed approach in this case.

500 µA. This small value is chosen since the LDOs used here are output-capacitorless and so are less stable at smaller load currents [14]. The controller is then removed and C_M inside each LDO is increased until the same settling time is obtained. The results are reported in Table IV. The table reports the extra capacitance per LDO and reports area overheads as percentages of LDO area. Evidently, the proposed approach provides vastly superior area efficiency to the conventional approach while incurring a relatively low power overhead. Note also that in some cases, the conventional approach cannot achieve the same performance as the proposed one.

C. Design Trade-Offs

This section illustrates the trade-offs in the design of the controller. First, at a fixed LDO number (case 1), different g_{m_c} values were used. The same load current in the previous section was used. The results, along with the controller overheads are shown in Fig. 9. As expected, increasing g_{m_c} improves stability at the cost of increased overheads, while increasing C_g requires increased g_{m_c}. Similar trade-offs can be observed for other cases.

Another trade-off to investigate is varying the number of LDOs with a fixed controller gain. For this test, different

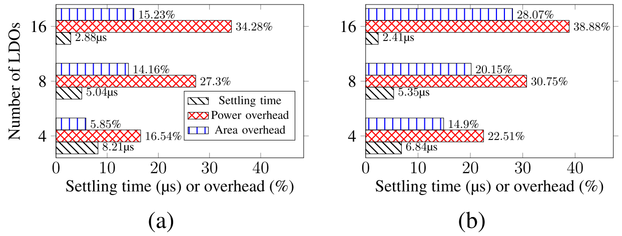

Fig. 10. T_s and controller overhead for cases (a) 1, 3 and 5 ($R_{\text{ref}} = 150\,\text{k}\Omega$) and (b) 2, 4 and 6 ($R_{\text{ref}} = 70\,\text{k}\Omega$).

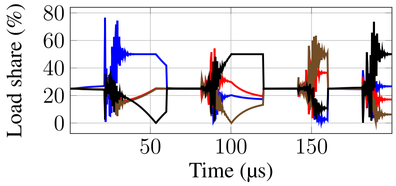

Fig. 11. LDO load shares with the controller enabled.

networks were simulated at fixed controller gain. The results are shown in Fig. 10. As expected, increasing the number of LDOs yields a better settling time as it improves load handling capability and provides a stronger feedback signal in the controller. Note that the controller overhead scales sub-linearly with the number of LDOs in the network.

D. Load Sharing

For this test, the loads shown in Fig. 8 were used. They are divided into groups (separated by the vertical lines in Fig. 8) and the groups were pulsed starting from the left. This test is performed for the network of case 1 in 3 cases:

1) All load pulses are small (0 to 100 μA)
2) Various loads; group 1: 0 to 100 μA, group 2: 10-15 mA, group 3: 35-40 mA and group 4: 45-50 mA
3) All load pulses are large (45-50 mA)

The load groups were pulsed at times 0, 60, 120 and 160 μs with a pulse width of 20 μs each. In the first case, the uncontrolled PDN was unstable and the results for the controlled network are shown in Fig. 11. The figure shows that the LDOs have different shares of the load at each pulse event depending on the proximity of the load to each LDO. Without the controller, all LDOs have equal load shares as their output currents oscillate in unison. For the two other tests, the uncontrolled PDN was stable and the controller did not have much effect and may thus be disabled to save power.

E. Effect of Line Delay

Due to the large size of PDNs, the delay of the feedback signal may be a concern; so a test that included equal delays on all the feedback lines was carried out. Fig. 12 shows the settling time of the network as a function of delay. It can be seen that delay reduces the stabilizing effect of the controller until it is no longer able to stabilize the PDN. Such curves can be used to estimate the maximum size of the network that the controller can stabilize. Note also that changing LDO number makes the PDN more robust to delay (except in case 6) while increasing C_g makes it more vulnerable, as expected.

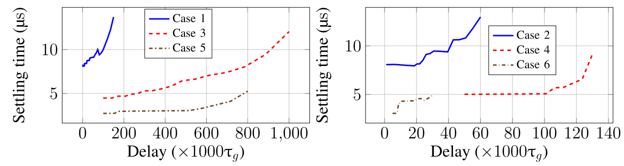

Fig. 12. Load step settling time vs. delay on the feedback line.

VI. Conclusion

This brief explored the feasibility of retro-fitting a PDN with a stabilizing centralized controller. A controller implementation was suggested and its stabilizing effect was demonstrated by circuit simulations where it was shown to have superior area efficiency to conventional techniques. The results suggest the performance benefits of centralized control and warrant further exploration of different controller structures.

References

[1] J. F. Bulzacchelli *et al.*, "Dual-loop system of distributed microregulators with high DC accuracy, load response time below 500 ps, and 85-mV dropout voltage," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 863–874, Apr. 2012.
[2] S. Lai, B. Yan, and P. Li, "Localized stability checking and design of IC power delivery with distributed voltage regulators," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 32, no. 9, pp. 1321–1334, Sep. 2013.
[3] I. Vaisband, B. Price, S. Köse, Y. Kolla, E. G. Friedman, and J. Fischer, "Distributed LDO regulators in a 28 nm power delivery system," *Analog Integr. Circuits Signal Process.*, vol. 83, no. 3, pp. 295–309, 2015.
[4] Z. Toprak-Deniz *et al.*, "Distributed system of digitally controlled microregulators enabling per-core DVFS for the POWER8^{TM} microprocessor," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers (ISSCC)*, Feb. 2014, pp. 98–99,
[5] S. Lai, B. Yan, and P. Li, "Stability assurance and design optimization of large power delivery networks with multiple on-chip voltage regulators," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, 2012, pp. 247–254.
[6] I. Vaisband and E. G. Friedman, "Stability of distributed power delivery systems with multiple parallel on-chip LDO regulators," *IEEE Trans. Power Electron.*, vol. 31, no. 8, pp. 5625–5634, Aug. 2016.
[7] X. Zhan, P. Li, and E. Sánchez-Sinencio, "Distributed on-chip regulation: Theoretical stability foundation, over-design reduction and performance optimization," in *Proc. 53rd ACM/EDAC/IEEE Design Autom. Conf. (DAC)*, Jun. 2016, pp. 1–6.
[8] X. Zhan, J. Riad, P. Li, and E. Sánchez, "Design space exploration of distributed on-chip voltage regulation under stability constraint," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 26, no. 8, pp. 1580–1584, Aug. 2018.
[9] X. Zhan, P. Li, and E. Sánchez-Sinencio, "Taming the stability-constrained performance optimization challenge of distributed on-chip voltage regulation," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, to be published.
[10] M. S. Mahmoud, *Decentralized Control and Filtering in Interconnected Dynamical Systems*. Boca Raton, FL, USA: CRC Press, 2011.
[11] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, vol. 2. New York, NY, USA: Wiley, 2007,
[12] Q.-G. Wang, Z. Ye, W.-J. Cai, and C.-C. Hang, *PID Control for Multivariable Processes*, vol. 373. Heidelberg, Germany: Springer-Verlag, 2008,
[13] W. Surakampontorn, V. Riewruja, K. Kumwachara, and K. Dejhan, "Accurate CMOS-based current conveyors," *IEEE Trans. Instrum. Meas.*, vol. 40, no. 4, pp. 699–702, Aug. 1991.
[14] J. Torres *et al.*, "Low drop-out voltage regulators: Capacitor-less architecture comparison," *IEEE Circuits Syst. Mag.*, vol. 14, no. 2, pp. 6–26, May 2014.