Towards Estimation of Worst-case Voltage Droop in Nonlinear Power Delivery Networks

Towards Estimation of Worst-case Voltage Droop in Nonlinear Power Delivery Networks

Antonio Carlucci, Stefano Grivet-Talocia
Dept. Electronics and Telecommunications, Politecnico di Torino, Italy
antonio.carlucci@polito.it, stefano.grivet@polito.it

***Abstract*—Worst-case droop in nonlinear Power Delivery Networks is estimated by harnessing existing optimal control techniques with data-driven bilinear macromodels to find the worst-case stimulus. The approach is numerically demonstrated on a realistic Low-DropOut regulator.**

***Index Terms*—Power Integrity, Modeling, Computer-Aided Design**

I. Introduction

Power Integrity (PI) verification is a fundamental step in the design of modern electronic systems. In fact, power efficiency is becoming increasingly important in electronic design, thus leading to tighter supply voltage margins and a greater need for careful PI analyses. An important requirement for PI is that the voltage supplied to all components connected to the output of a Power Delivery Network (PDN) remains close to its nominal value according to prescribed bounds. Due to dynamic activity of the loads that draw time-varying currents from the PDN, the supply voltage is affected by droops or power noise. In design phase and, most importantly, in the final stages where the performance of the overall system is verified, the ability to establish which are the corner cases where specifications may be violated is of paramount importance. In the case of PI, this consists in finding particular excitation current patterns that result in the largest (i.e., worst-case) voltage droop.

Although PDNs are typically regarded as linear and passive, many systems of practical relevance include nonlinear components. For example, PDNs with Integrated Voltage Regulators are commonly used in multi-core microprocessors [1] and include nonlinear feedback loops to control banks of switching converters. In addition, regulators (e.g. Low-DropOut) used as the main source of supply voltage are implemented through circuits with nonlinear characteristics. These systems are typically designed to behave linearly, so that linearized models often provide good approximations of the true behavior. However, more precise analyses are enabled if (second-order) nonlinear effects are taken into account.

Estimation of the worst-case voltage droop for linear (or linearized) PDNs is discussed in [2], and an efficient method that uses rational macromodels was presented in [3]. Therein, a method is described to find the value of the worst-case droop, together with the particular current stimulus that produces it. In this work, we consider the more general case where the PDN contains nonlinear components. This case requires extending the methodology of [2], [3] to address nonlinearities.

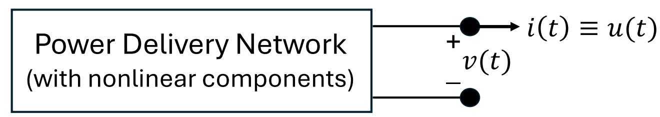

Fig. 1. Schematic depiction of a Power Delivery Network, together with the load voltage $v(t)$ and the load current $i(t)$

Let us consider a single-port nonlinear electronic system representing a source of supply voltage $v(t)$ for an arbitrary load that draws a time-varying current $i(t) \equiv u(t)$, as depicted in Fig. 1. The worst-case droop is the lowest instantaneous value of supply voltage $v(t)$ that can be reached as the current signal $u(t)$ varies in a prescribed set. In this work, we study the worst-case droop over a finite time interval, from $t = 0$ to a prescribed final time $t = T$. Any given excitation $u(t)$ corresponds to a response $v(t)$ determined by the system behavior, and the worst-case $u^*(t)$ is the one causing the response $v^*(t)$ to attain the lowest possible value in $[0, T]$.

The search for the worst-case $u^*(t)$ is carried out in a set of *admissible* signals satisfying given constraints. In practice, a constraint on the maximum amplitude $\bar{I}$ of $u(t)$ is often implied by design specifications. To fix ideas, we consider admissible signals $u(t)$ that satisfy

$$-\bar{I} \leq u(t) \leq \bar{I}, \quad \forall\, t \in [0, T].$$

Due to time invariance, it can be assumed that the lowest value of $v^*(t)$ is attained at the final time $t = T$. Therefore, $u^*(t)$ is defined as

$$u^*(t) = \underset{-\bar{I} \leq u(t) \leq \bar{I}}{\arg\min}\; v(T), \qquad v^*(T) = \min_{-\bar{I} \leq u(t) \leq \bar{I}} v(T).$$

This problem is solved through the well-known theory of optimal control [4]. In this paper, we present a practical framework to use existing theoretical results for estimation of worst-case droop of nonlinear electronic systems. In particular, we consider a common setting in electronic design where the equations underlying the nonlinear system are not directly accessible, and only its input/output behavior is computable (e.g. using field or circuit solvers). For this scenario, we propose first building a bilinear macromodel following the methodology of [5]. Based on the particular structure of the bilinear model, we apply optimal control to derive *necessary* conditions that $u^*(t)$ should satisfy and discuss computational techniques to find it. Finally, the entire workflow is demonstrated on a Low-DropOut Regulator circuit.

II. Background: Bilinear Macromodels

In this work, it is assumed that the system under analysis is represented by a bilinear model, corresponding to the following set of state-space equations,

$$\dot{\boldsymbol{x}}(t) = \mathbf{A}\boldsymbol{x}(t) + \mathbf{N}\boldsymbol{x}(t)u(t) + \boldsymbol{b}u(t), \quad v(t) = \boldsymbol{c}\boldsymbol{x}(t) \tag{1}$$

This is motivated by the fact that, using the algorithm presented in [5], it is possible to build compact bilinear models of weakly nonlinear systems starting from frequency-domain data. In practice, the method in [5] only requires the ability to characterize the voltage response $v(t)$ of the original system by means of Harmonic Balance (HB) simulations. Starting from the results of HB analysis, the values of suitably-defined Generalized Transfer Functions (related to Volterra theory [6]) are extracted and used to construct a bilinear model like (1) that accurately represents the original system behavior.

III. Necessary conditions for worst-case stimulus

In this section, the theory of [4] is applied to the bilinear model (1) in order to find optimality conditions characterizing the worst-case $u^*(t)$. Following [4, Ch. 5], we introduce the co-state $\boldsymbol{p}(t) \in \mathbb{R}^n$ and the Hamiltonian

$$\mathcal{H}(\boldsymbol{x}(t), u(t), \boldsymbol{p}(t)) = \boldsymbol{p}(t)^{\mathsf{T}}(\mathbf{A}\boldsymbol{x}(t) + \mathbf{N}\boldsymbol{x}(t)u(t) + \boldsymbol{b}u(t)).$$

that corresponds to minimizing the cost $v(T) = \boldsymbol{c}\boldsymbol{x}(T)$.

Applying Pontryagin's minimum principle [4], the worst-case trajectories $\boldsymbol{x}^*(t)$, $u^*(t)$, $\boldsymbol{p}^*(t)$ must satisfy the following system of equations (see [4, Sec. 5.3]),

$$\dot{\boldsymbol{x}}^*(t) = \frac{\partial \mathcal{H}}{\partial \boldsymbol{p}}(\boldsymbol{x}^*(t), u^*(t), \boldsymbol{p}^*(t)), \qquad \boldsymbol{x}(0) = 0 \tag{2a}$$

$$\dot{\boldsymbol{p}}^*(t) = -\frac{\partial \mathcal{H}}{\partial \boldsymbol{x}}(\boldsymbol{x}^*(t), u^*(t), \boldsymbol{p}^*(t)), \qquad \boldsymbol{p}(T) = \boldsymbol{c}^{\mathsf{T}} \tag{2b}$$

$$\mathcal{H}(\boldsymbol{x}^*(t), u^*(t), \boldsymbol{p}^*(t)) = \min_{-\bar{I} \le u \le \bar{I}} \mathcal{H}(\boldsymbol{x}^*(t), u, \boldsymbol{p}^*(t)). \tag{2c}$$

Condition (2c) gives the worst-case input. In fact, since $\mathcal{H}$ is an affine function of u, it is minimized by choosing

$$u^*(t) = -\bar{I}\operatorname{sgn}\left\{\boldsymbol{p}^*(t)^{\mathsf{T}}\left[\mathbf{N}\boldsymbol{x}^*(t) + \boldsymbol{b}\right]\right\} \tag{3}$$

where $\operatorname{sgn}\{\cdot\}$ is the sign function. Although (3) depends on yet unknown quantities ($\boldsymbol{p}^*$ and $\boldsymbol{x}^*$), it shows that the worst-case input is either $\bar{I}$ or $-\bar{I}$ at all times, a condition that is also verified in the linear scenario considered in [3].

The first two conditions (2a) and (2b) translate into

$$\dot{\boldsymbol{x}}^*(t) = \mathbf{A}\boldsymbol{x}^*(t) + \mathbf{N}\boldsymbol{x}^*(t)u^*(t) + \boldsymbol{b}u^*(t) \tag{4}$$

$$\dot{\boldsymbol{p}}^*(t) = -\left[\mathbf{A} + \mathbf{N}u^*(t)\right]^{\mathsf{T}}\boldsymbol{p}^*(t), \qquad \boldsymbol{p}^*(T) = \boldsymbol{c}^{\mathsf{T}} \tag{5}$$

To summarize, conditions (3), (4) and (5) must be necessarily satisfied by the worst-case trajectories $\boldsymbol{x}^*(t)$, $\boldsymbol{p}^*(t)$, $u^*(t)$.

IV. Computation of worst-case stimulus by gradient descent

The search for the stimulus $u(t)$ that minimizes the final value $v(T)$ can be carried out using a *gradient descent* method that iteratively perturbs the input $u(t)$ to decrease the cost function $v(T)$. Consider the responses $\bar{\boldsymbol{x}}(t)$ and $\bar{v}(t)$ of (1) corresponding to a given input $\bar{u}(t)$. Upon applying a small perturbation $\delta u(t)$, so that $u(t) = \bar{u}(t) + \delta u(t)$, it will reflect in the output with a perturbation $\delta v(t)$, so that $v(t) = \bar{v}(t) + \delta v(t)$. A first-order approximation to $\delta v(t)$ is obtained by considering the linearization of the system (1) around $\bar{\boldsymbol{x}}(t)$, that is the linear time-varying system

$$\delta\dot{\boldsymbol{x}} = [\mathbf{A} + \mathbf{N}\bar{u}(t)]\delta\boldsymbol{x} + [\mathbf{N}\bar{\boldsymbol{x}}(t) + \boldsymbol{b}]\delta u(t), \;\; \delta v(t) = \boldsymbol{c}\delta\boldsymbol{x}(t) \tag{6}$$

In the following, let us discretize $[0, T]$ into K uniformly-spaced timesteps $t_k = k\Delta t$ with $\Delta t \triangleq {}^{T}\!/_{K}$. We approximate $\bar{u}(t)$ and $\delta u(t)$ with piecewise-constant signals with values $\bar{u}_k \triangleq u(t_k)$ and $\delta u_k \triangleq \delta u(t_k)$. Using the impulse response $h(t, \tau)$ associated with (6), the solution at final time reads

$$\delta v(T) = \int_0^T h(T, \tau)\delta u(\tau)\,\mathrm{d}\tau \approx \Delta t \sum_{k=1}^{K} h\left(T, t_k\right)\delta u_k.$$

This shows that the derivative of $v(T)$ with respect to the value $\bar{u}_k$ that the input takes at time t_k is $h(T, t_k)\Delta t$. This implies that perturbing the input according to the rule $\delta u_k = -\lambda h(T, t_k)$ with a sufficiently small step $\lambda > 0$ will decrease the cost function $v(T)$. Since this update might lead to an inadmissible $u(t)$ that exceeds the maximum amplitude $\bar{I}$, the result needs to be *projected* back onto the admissible set by the operator $\Pi[\cdot]$ that clips the signal at the amplitude $\bar{I}$, that is $\Pi[u(t)] = \max(\min(u(t), \bar{I}), -\bar{I})$. This results in the gradient-based algorithm reported in Algorithm 1.

Algorithm 1: Solution by gradient-based algorithm

Input: Initial guess $u(t)$, tolerance ϵ, step λ

repeat

  Set $\bar{u}(t) \leftarrow u(t)$;

  Compute $\bar{\boldsymbol{x}}(t)$ by solving (1) with input $\bar{u}(t)$;

  Compute impulse response $h(T, t)$ of (6);

  Update $u(t) = \Pi[\bar{u}(t) - \lambda h(T, t)]$;

until $||u - \bar{u}|| > \epsilon$;

return $u^*(t) = u(t)$

Note that the impulse response $h(T, t)$ can be computed numerically by considering that

$$h(T, t) = \boldsymbol{p}^{\mathsf{T}}(t)[\mathbf{N}\bar{\boldsymbol{x}}(t) + \boldsymbol{b}] \tag{7}$$

where $\boldsymbol{p}(t)$ solves $\dot{\boldsymbol{p}}(t) = [\mathbf{A} + \mathbf{N}\bar{u}(t)]^{\mathsf{T}}\boldsymbol{p}(t)$, $\boldsymbol{p}(T) = \boldsymbol{c}^{\mathsf{T}}$. Therefore, the values $h(T, t_k)$ that are needed to update $\bar{u}_k$ in a practical implementation are found by integrating (7) using standard methods (e.g. Euler's implicit method). Comparing (7) with (3), (5) shows that, if the candidate $\bar{u}(t)$ satisfies (3), that is $\bar{u}(t) = -\bar{I}\operatorname{sgn} h(T, t)$, the projected gradient algorithm stops because the updated $\Pi[\bar{u}(t) - \lambda h(T, t)] = \bar{u}(t)$.

V. Numerical Results

In this section, worst-case droop estimation is demonstrated on the Low-DropOut regulator initially introduced in [7]. We considered a PDN system made up of an LDO circuit followed by a passive network representing a PCB. A model was built of the voltage response $v(t)$ to the load current $i(t) \equiv u(t)$. The nonlinear model represents the behavior around a nominal operating point defined by 20 mA load bias current, for which the load voltage is $V_0 = 2.710$ V. Following [5], samples of the Generalized Transfer Functions were extracted up to degree three using HSPICE HB analysis, and a bilinear model like (1) was built by fitting the data. As discussed in [5], a bilinear model built in this way is only valid as long as the system operates in the *weakly nonlinear* regime. To ensure model validity, the analysis is carried out assuming a maximum amplitude $\bar{I} = 8\,\text{mA}$, intended as variation from the bias.

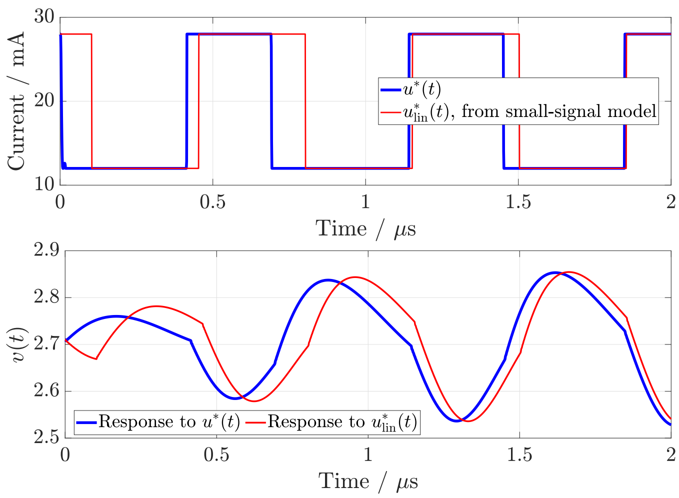

Fig. 2. Top: Worst-case stimulus found using gradient descent (blue), or small-signal linearization (black). Bottom: voltage response to the stimuli above.

A MATLAB implementation of Algorithm 1 was used to estimate the worst-case stimulus $u^*(t)$. The algorithm was run with $\lambda = 10^3$. Considering a time interval up to $T = 2\,\mu$s and a timestep $\Delta t = 1\,$ns, the optimization method runs for 85 iterations before $||u - \bar{u}||_2$ falls below the tolerance $\epsilon = 10^{-3}$. The algorithm is initialized with $\bar{u}(t) = 0$ and it takes 55 s on a 3.9 GHz computer. The worst-case input found by the algorithm is depicted in Fig. 2 (top,blue line).

An alternative approach to obtain an estimate of the worst-case excitation would be to neglect the nonlinearity and use a linearized small-signal model. In that case, the worst-case input can be directly found by considering the sign of the impulse response, as discussed in [2], [3]. This alternative is here considered and the worst-case excitation $u^*_{\text{lin}}(t)$ thus estimated is compared with the $u^*(t)$ resulting from the gradient-based approach in Fig. 2 (top). The response of the original system was computed through HSPICE simulations with the stimuli of Fig. 2 (top), and the voltage responses are reported in Fig. 2 (bottom). Using the stimulus $u^*(t)$, the voltage droop with respect to the nominal voltage V_0 is 182 mV. On the other hand, using the estimate $u^*_{\text{lin}}(t)$ based on the small-signal model, the simulated worst-case droop is 174 mV (attained at $t = 1.33\,\mu$s), and the droop at the final time $t = T$ is 170 mV . Therefore, in this case, using the small-signal model would lead to an underestimation of the voltage droop of at least 8 mV (4%).

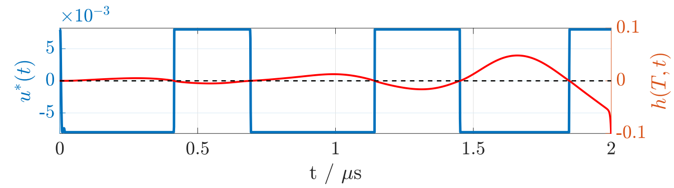

Fig. 3. Worst-case input found by Algorithm 1 and the corresponding $h(T, t)$.

To give an intuitive picture of the condition (3) satisfied by $u^*(t)$, Fig. 3 shows the solution $u^*(t)$ found in the last iteration of Algorithm 1, along with the impulse response $h(T, t)$ computed in the same iteration. The sign of $h(T, t)$ is opposite to that of the candidate solution, whose value is either $-\bar{I}$ or $\bar{I}$. This means that any perturbation $\lambda h(T, t)$ yields an inadmissible solution, hence the optimization terminates.

VI. Discussion and Conclusion

The material presented in this paper describes a full methodology to estimate the worst-case droop and stimulus for nonlinear PDNs. This is made efficient by resorting to a compact data-driven bilinear model used as surrogate for the original system. The approach is based on formal theoretical tools that allow tackling the problem in a non-heuristic way. In view of an exact and complete methodology, these results should be improved by investigating whether the worst-case found by the adopted local method is indeed a *global* optimum. In addition, the limitation of weak nonlinearity entailed in the modeling procedure instrumental to this methodology constitutes a bottleneck for systems that are driven into strongly nonlinear regime, possibly by the action of the worst-case stimulus that is being sought.

References

[1] K. Radhakrishnan, M. Swaminathan, and B. K. Bhattacharyya, "Power delivery for high-performance microprocessors—challenges, solutions, and future trends," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 11, no. 4, pp. 655–671, 2021.

[2] V. Drabkin, C. Houghton, I. Kantorovich, and M. Tsuk, "Aperiodic resonant excitation of microprocessor power distribution systems and the reverse pulse technique," in *2002 IEEE 11th Topical Meeting on Electrical Performance of Electronic Packaging*, 2002, pp. 175–178.

[3] A. Carlucci, T. Bradde, and S. Grivet-Talocia, "Fast prediction of worst-case voltage droops in power distribution networks," in *2024 IEEE 28th Workshop on Signal and Power Integrity (SPI)*, 2024, pp. 1–4.

[4] D. Kirk, *Optimal Control Theory: An Introduction*, ser. Dover Books on Electrical Engineering Series. Dover Publications, 2004.

[5] A. Carlucci, I. V. Gosea, and S. Grivet-Talocia, "Data-driven modeling of weakly nonlinear circuits via generalized transfer function approximation," *IEEE Access*, vol. 13, pp. 2746–2762, 2025.

[6] W. Rugh, *Nonlinear System Theory: The Volterra/Wiener Approach*, ser. Johns Hopkins series in information sciences and systems. Johns Hopkins University Press, 1981.

[7] R. J. Milliken, J. Silva-Martinez, and E. Sanchez-Sinencio, "Full on-chip CMOS low-dropout voltage regulator," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 54, no. 9, pp. 1879–1890, 2007.