Efficient Parametric Assessment of Worst-Case Voltage Droop in Power Delivery Networks

(Article begins on next page)

Efficient parametric assessment of worst-case voltage droop in power delivery networks

Tommaso Bradde, Antonio Carlucci, Riccardo Trinchero, Paolo Mandredi, Stefano Grivet-Talocia
Dept. Electronics and Telecommunications, Politecnico di Torino, Italy
tommaso.bradde@polito.it

***Abstract*—Power Delivery Network (PDN) optimization is crucial for guaranteeing adequate power integrity performance in modern microprocessor systems. In this work, we introduce a novel surrogate modeling workflow for efficiently predicting the worst-case voltage droop occurring at the loading points of a PDN including a set of free design parameters. We apply the proposed approach for modeling the impact of a set of decoupling capacitors on the performance of a template PDN structure.**

***Index Terms*—Power Delivery Networks, Decoupling Capacitors, Machine Learning**

I. Introduction

Power Integrity (PI) optimization is crucial for achieving target performance in terms of efficiency and reliability of modern microprocessor systems. With current technologies, careful assessments of the Power Delivery Network (PDN) electrical performance must be necessarily performed at the system level. In this view, one challenging task in PI is the minimization of the PDN power noise, and in particular, of the voltage droop resulting from the dynamic activity of its loads. Most commonly, this performance specification is met by designing the PDN with the objective of bounding its impedance magnitude below a (possibly frequency-dependent) threshold, known as target impedance [1] [2] [3]. Optimizing a PDN to achieve minimal voltage droops can be challenging and time consuming, as any design configuration must be ultimately verified in time domain, possibly interlacing extremely costly transient analyses with the impedance shaping routine, in order to actually verify the entity of the droop under the operating conditions of interest.

This contribution introduces a surrogate modeling approach thought to alleviate the above-mentioned computational issues. In particular, we propose an efficient approach for directly predicting the worst-case voltage droop (WCVD) of a PDN as a function of the relevant design parameters, assuming that the load currents are bounded in amplitude and slew-rate. The approach expands on the recent numerical tools introduced in [4], that allow to compute directly the voltage droop without relying on any costly transient simulations. Exploiting these tools, we compute the WCVD of the PDN for a discrete number of design parameters configurations, and we exploit this information as input for generating a surrogate model which predicts the desired performance index efficiently, throughout the whole design space. The surrogate is obtained by applying Gaussian Process Regression (GPR), which is a machine learning approach that allows reproducing a complex functional relationship using a relatively low amount of training data [5]. Numerical experiments based on a template PDN description provide a proof of concept for the efficacy and efficiency of the proposed approach.

II. Problem Statement

We consider a generic Linear Time Invariant (LTI) PDN structure, subject to independent current stimuli acting in correspondence of P well-defined electrical ports, representing its loading points. We allow the PDN description to include a number ρ of free design parameters, collected in the vector $\boldsymbol{x} = [x_1, \dots, x_\rho]^T \subset \mathcal{X}$, where $\mathcal{X}$ is a hyperectangle defining the allowed domain of variation. We denote as $\mathbf{Z}(s, \boldsymbol{x}) \in \mathbb{C}^{P\times P}$ the parameterized output impedance matrix of the PDN defined at the loading points, and we assume that $\mathbf{Z}(s, \boldsymbol{x})$ can be sampled at discrete frequency-parameter configurations, via real or virtual measurements.

Let us denote the impulse response as $\mathbf{z}(t, \boldsymbol{x}) = \mathcal{L}^{-1}\{\mathbf{Z}(s, \boldsymbol{x})\}$, the vector of load current sources entering the PDN as $\boldsymbol{i}(t)$, and the corresponding port voltages as $\boldsymbol{v}(t)$. Then, the convolution integral

$$\boldsymbol{v}(t, \boldsymbol{x}) = \int_0^t \mathbf{z}(\tau, \boldsymbol{x})\boldsymbol{i}(t-\tau)d\tau = (\mathbf{z}(\boldsymbol{x}) \star \boldsymbol{i})(t), \quad (1)$$

provides the instantaneous voltage droop of the PDN for any admissible design and load profile. Our objective is to generate a surrogate model for efficiently predicting the WCVD of the PDN as a function of the parameters, assuming load currents that are bounded in amplitude and slew rate. Formally, we desire a surrogate representation for the following function

$$y(\boldsymbol{x}) = \sup_{t\geq 0, \boldsymbol{i}(t)\in\mathcal{I}} ||\boldsymbol{v}(t, \boldsymbol{x})||_\infty, \quad (2)$$

where $\mathcal{I}$ denotes the set of admissible load current stimuli

$$\mathcal{I} = \{\boldsymbol{i}(t) : 0 \leq |i_j(t)| \leq I_{j,\max}, \left|\frac{di_j(t)}{dt}\right| \leq \Delta_{\max}, \forall t \geq 0, j = 1, \dots, P\}. \quad (3)$$

Notice that for step transitions $0 \leftrightarrow I_{\max}$ the slew-rate constraint implies a minimum rise time $\tau_r = I_{\max}/\Delta_{\max}$.

III. Surrogate Model Generation

The proposed modeling workflow is based on two main steps. The first consists in the generation of a dataset of pairs $\mathcal{D} = \{(\boldsymbol{x}_k, y_k)\}_{k=1}^K$, with $y_k = y(\boldsymbol{x}_k)$, obtained by evaluating (2) at discrete parameter values $\boldsymbol{x}_k \in \mathcal{X}$. The second exploits this dataset to generate the GPR model.

A. Dataset Extraction

1) Voltage Droop Computation: Given a design parameter configuration $\boldsymbol{x}_k$, the exact computation of the corresponding bound $y(\boldsymbol{x}_k)$ would require constructing the worst-case load current signal, as outlined in [6], and simulating the PDN response in the time domain applying such load profile. To avoid this expensive procedure, in this contribution we apply the simplified approach described in [4], which is equivalent for practical purposes. Specifically, we evaluate (2) as

$$y(\boldsymbol{x}_k) = \max_{i=1,\ldots,P} v_{i,\max}(\boldsymbol{x}_k), \tag{4}$$

$$v_{i,\max}(\boldsymbol{x}_k) = \sum_{j=1}^{P} I_{j,\max} \int_0^\infty |(z_{ij}(\boldsymbol{x}_k) \star g_{\tau_r})(t)|_+ dt \tag{5}$$

where $g_{\tau_r}(t)$ is a unit-area square pulse having width τ_r and $|a|_+$ is equal to a if $a \geq 0$ and is 0 otherwise. See [4, Sec.II-B] for technical details. Assuming that a closed form expression for $\mathbf{z}(t, \boldsymbol{x}_k)$ is available, the target value $y(\boldsymbol{x}_k)$ is obtained via numerical integration of (5). Most commonly, the numerical or experimental characterization of the PDN behavior is available only terms of measurements of the corresponding impedance matrix $\mathbf{Z}(s, \boldsymbol{x})$. In this scenario, a closed form approximation for $\mathbf{z}(t, \boldsymbol{x}_k)$ can be retrieved via rational fitting, using a set of measurements of the kind $\mathcal{V}_k = \{(j\omega_m, \mathbf{Z}(j\omega_m, \boldsymbol{x}_k))\}_{m=1}^M$, with $\omega_m = 2\pi f_m,\ f_m \in [f_{\min}, f_{\max}]$. The measurements are used as input for the Vector Fitting iteration [7] to generate a stable (yet not necessarily passive) approximation $\tilde{\mathbf{Z}}_k(s)$ for the PDN impedance, with structure

$$\tilde{\mathbf{Z}}_k(s) = \sum_{\ell=1}^{\bar{\ell}} \frac{\mathbf{R}_\ell}{s - p_\ell} \approx \mathbf{Z}(s, \boldsymbol{x}_{\boldsymbol{k}}). \tag{6}$$

An approximation for the required impulse response is obtained via analytical inverse Laplace transform

$$\tilde{\mathbf{z}}_k(t) = \mathcal{L}^{-1}\{\tilde{\mathbf{Z}}_k(s)\} = \sum_{\ell=1}^{\bar{\ell}} \mathbf{R}_\ell\, e^{p_\ell t} \approx \mathbf{z}(t, \boldsymbol{x}_k), \tag{7}$$

and can be used in place of $\mathbf{z}(t, \boldsymbol{x}_k)$ in (5) to compute the desired sample $y(\boldsymbol{x}_{\boldsymbol{k}})$.

B. Modeling via Gaussian Process regression

Once the dataset $\mathcal{D}$ is available, we use it to train a GPR model. We choose GPR because it offers a good trade-off between model accuracy and complexity, providing a flexible model with limited training cost. We consider a standard implementation with a constant prior trend β_0 and an anisotropic Matérn 5/2 covariance function, i.e.,

$$k(\boldsymbol{x}, \boldsymbol{x}') = \sigma^2 \left(1 + \sqrt{5}u + \frac{5}{3}u^2\right) \exp\left(-\sqrt{5}u\right), \tag{8}$$

where

$$u = \sqrt{\sum_{j=1}^{\rho} \frac{(x_j - x'_j)^2}{\theta_j^2}} \tag{9}$$

TABLE I
VARIABILITY RANGES OF DESIGN PARAMETERS

C_2 /nF	C_3 /μF	C_4 /μF	C_5 /μF	R_{2-5} /mΩ	L_{2-5} /pH
$[50, 150]$	$[1, 100]$	$[0.1, 10]$	$[0.01, 1]$	$25 \pm 30\%$	$1 \pm 90\%$

and $\theta_1, \ldots, \theta_\rho$ are unknown scale parameters. The covariance function describes the correlation between two points in the input space, $\boldsymbol{x}$ and $\boldsymbol{x}'$, which translates into the model smoothness. Owing to the complexity of the target function, an anisotropic kernel (in which the scale parameter differs for each input parameter) turns out to be substantially more accurate, at the expense of a slight increase of the training cost. The WCVD is predicted as

$$\hat{y}(\boldsymbol{x}) = \beta_0 + \sum_{k,m=1}^{K} \left[\tilde{\mathbf{K}}^{-1}\right]_{km} (y_m - \beta_0) k(\boldsymbol{x}, \boldsymbol{x}_k). \tag{10}$$

In (10), $\tilde{\mathbf{K}}$ is a matrix with entries $\tilde{K}_{km} = k(\boldsymbol{x}_k, \boldsymbol{x}_m) + \sigma_n^2 \delta_{km}$, for $k, m = 1, \ldots, K$, where δ_{km} is the Kronecker's delta and σ_n^2 is a noise parameter that acts as a regularizer. It should be noted that the size of $\tilde{\mathbf{K}}$ is determined by the available data samples and usually does not need to be substantially increased, even for larger input spaces [5]. The coefficient β_0 is obtained via a generalized least-square estimate as

$$\beta_0 = \frac{\boldsymbol{e}^\mathsf{T} \tilde{\mathbf{K}}^{-1} \boldsymbol{y}}{\boldsymbol{e}^\mathsf{T} \tilde{\mathbf{K}}^{-1} \boldsymbol{e}}, \tag{11}$$

where $\boldsymbol{e} \in \mathbb{R}^K = (1, \ldots, 1)^\mathsf{T}$ is a column vector of ones and $\boldsymbol{y} = (y_1, \ldots, y_K)^\mathsf{T}$ is the vector of data samples. The scale parameters $\{\theta_j\}_{j=1}^{\rho}$, the noise variance σ_n^2, and the kernel variance σ^2 are estimated via likelihood maximization.

IV. NUMERICAL RESULTS

We provide a proof of concept for the proposed approach considering a 2-D distributed structure, representative of a template PDN on a printed circuit board. The board consists of two parallel square planes with side length $l = 7.5$ cm, separated by a dielectric material of width $d = 0.5$ mm. The dielectric has relative permittivity $\epsilon_r = 5.5$ and loss tangent $\tan\delta = 0.01$. Five ideal (lumped) ports are defined between top and bottom planes at coordinates $\{(0,\ 3.8), (3,\ 3), (2.7,\ 3), (3,\ 2.7), (3.3,\ 3)\}$, defined in cm taking the bottom left corner of the board as the origin.

We close port #1 on a Voltage Regulator Module (VRM), modeled as a RL series circuit with $\mathrm{R_{VRM}} = 2\ \mathrm{m}\Omega$ and $\mathrm{L_{VRM}} = 1$ nH. Port #2 is considered as the loading point of the PDN, and closed on a RC series circuit representing a simplified silicon die model (with $\mathrm{R_{die}} = 50\ \mathrm{m}\Omega$ and $\mathrm{C_{die}} = 5$ nF) in parallel with a decoupling capacitor. With the remaining ports left open, the impedance of the structure is shown in Fig. 1 (blue line). We assume ports #2 to #5 be shunted with four decoupling capacitors. The latter are modeled as series RLC circuits, whose lumped element values are the design parameters of interest. Table I reports the admissible parameter ranges, using subscripts to identify the

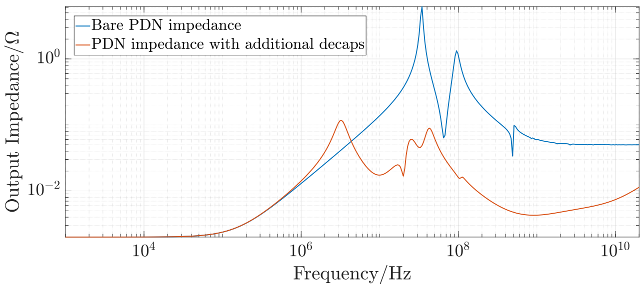

Fig. 1. The output impedance of the considered template PDN. The blue line shows the bare impedance without decoupling capacitors. The red line shows the impedance loaded by with a random admissible choice of such capacitors.

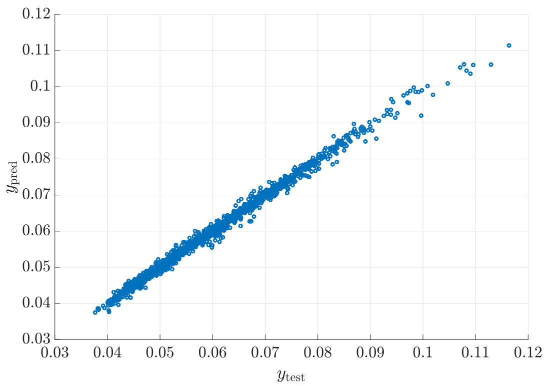

Fig. 2. Correlation plot of the reference data against the predictions obtained via the GPR model built with K_{train}=1000 training samples.

components reference ports. The PDN impedance obtained with one random design choice is shown in Fig. 1 (red line).

Our objective is to predict the WCVD of the PDN as a function of these parameters when the loading point is subject to a load current with maximum amplitude $I_{\max} = 1$ A and minimum rise time $\tau_r = 3$ ns. To this aim, we generate the dataset $\mathcal{D}$ with $K = 2500$, defining the sampling points $\{\boldsymbol{x}_k\}_{k=1}^{K}$ using a Sobol sequence. For each sampling point, we retrieve the rational model (6) approximating the output impedance of the PDN and we compute the corresponding worst-case droop according to (4), (5). The data are then divided into disjoint training and test sets, denoted as $\mathcal{D}_{\text{train}} = \{(\boldsymbol{x}_j, y_j)\}_{j=1}^{K_{\text{train}}}$ and $\mathcal{D}_{\text{test}} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^{K_{\text{test}}}$, where we set $K_{\text{train}} = 1000$ and $K_{\text{test}} = 1500$. Using $\mathcal{D}_{\text{train}}$, the GPR model is trained in 40 s on a common laptop.

The scatter plot of the GPR predictions against the test samples is shown in Fig. 2 and proves the remarkable accuracy of the surrogate, exhibiting a coefficient of determination of

$$R^2 = 1 - \frac{\sum_{i=1}^{K_{\text{test}}} \left(y_i - \hat{y}(\boldsymbol{x}_i)\right)^2}{\sum_{i=1}^{K_{\text{test}}} \left(y_i - \bar{y}\right)^2} = 0.98, \tag{12}$$

where $\bar{y}$ is the dataset mean. Using the model, a batch of 1000 test samples is computed in 66 ms, against the 54 s required by directly computing (5). This confirms that also in case of the considered academic example the proposed approach provides major improvements in terms of efficiency. To conclude, following the approach presented in [4], for each test sample, we build the worst-case current giving rise to the WCVD, and compute the associated voltage responses (1), over a timespan of 5 μs. Figure 3 shows the ensemble of these responses over a restricted time window. The solid red line marks the WCVD for the best design among the test configurations, and shows that the bound is actually attained by the corresponding voltage droop waveform (blue solid line); the GPR prediction deviates by 0.5% from the reference.

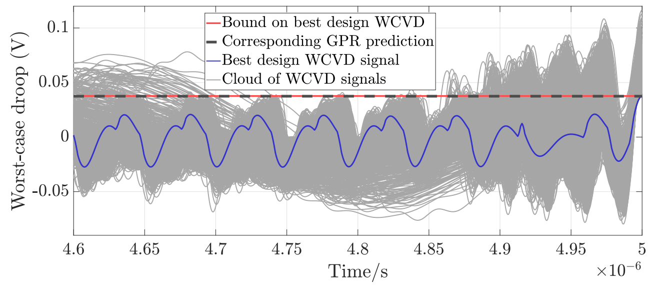

Fig. 3. Grey lines: WCVD waveforms associated to the test samples. The blue line is the voltage droop signal corresponding to the best design. Red and black dashed lines are the corresponding exact and surrogate WCVD bounds.

V. Conclusions

This work introduced a novel approach for fast parametric assessment of the WCVD occurring in a PDN under constrained load current profiles. The proposed method combines GPR surrogate modeling with established numerical tools, enabling efficient prediction of the considered performance index for the sake of design verification and optimization.

References

[1] J. Chen and M. Hashimoto, "A frequency-dependent target impedance method fulfilling voltage drop constraints in multiple frequency ranges," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 10, no. 11, pp. 1769–1781, 2020.

[2] Y. Kim, K. Kim, J. Cho, J. Kim, K. Kang, T. Yang, Y. Ra, and W. Paik, "Power distribution network design and optimization based on frequency dependent target impedance," in *2015 IEEE Electrical Design of Advanced Packaging and Systems Symposium (EDAPS)*, 2015, pp. 89–92.

[3] S. Liang, B. Zhao, S. Bai, S. Connor, M. Cocchini, S. Scearce, D. Becker, M. Cracraft, M. S. Doyle, A. Ruehli, and J. Drewniak, "Decoupling capacitor optimization to achieve target impedance in pcb pdn design," in *2021 IEEE International Joint EMC/SI/PI and EMC Europe Symposium*, 2021, pp. 967–972.

[4] 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.

[5] P. Manfredi and R. Trinchero, "A probabilistic machine learning approach for the uncertainty quantification of electronic circuits based on Gaussian process regression," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 8, pp. 2638–2651, 2021.

[6] W. Reinelt, "Maximum output amplitude of linear systems for certain input constraints," in *Proceedings of the 39th IEEE Conference on Decision and Control*, vol. 2, 2000, pp. 1075–1080.

[7] S. Grivet-Talocia and B. Gustavsen, *Passive Macromodeling: Theory and Applications*. New York: John Wiley and Sons, 2016.