

A Structured Krylov Subspace Projection Framework for Fast Power Integrity Verification

Original

A Structured Krylov Subspace Projection Framework for Fast Power Integrity Verification / Carlucci, Antonio; Grivet-Talocia, Stefano; Mongrain, Scott; Kulasekaran, Sid; Radhakrishnan, Kaladhar. - ELETTRONICO. - (2023), pp. 1-4. (Intervento presentato al convegno 2023 IEEE 27th Workshop on Signal and Power Integrity (SPI) tenutosi a Aveiro, Portugal nel 07-10 May 2023) [10.1109/SPI57109.2023.10145566].

Availability:

This version is available at: 11583/2979366 since: 2023-06-14T09:22:11Z

Publisher:

IEEE

Published

DOI:10.1109/SPI57109.2023.10145566

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

A Structured Krylov Subspace Projection Framework for Fast Power Integrity Verification

Antonio Carlucci*, Stefano Grivet-Talocia*, Scott Mongrain[§], Sid Kulasekaran[§], Kaladhar Radhakrishnan[§]

*Dept. Electronics and Telecommunications, Politecnico di Torino, Italy

[§]Intel Corporation, Chandler, AZ, USA

antonio.carlucci@polito.it

Abstract—This paper presents a model order reduction approach, specifically designed for the generation of compact and efficient transient simulation models of system-level power distribution networks (PDN) of multicore processor systems. The proposed approach applies a Krylov subspace projection, with a structure that is adapted to a block-coupled state-space description of individual PDN subsystems. The latter include board-package, averaged models of integrated voltage regulators switching circuitry, and individual models of all cores including regulator inductors and capacitors. Numerical results from proposed reduced-order models provide major speedup with respect to SPICE with negligible loss of accuracy.

I. INTRODUCTION

Modern multicore processor systems are equipped with sophisticated voltage regulation systems [1], aimed at stabilizing and controlling power delivery to individual cores. This is usually achieved through Fully Integrated Voltage Regulators (FIVR) [2], usually implemented as multi-phase switching power supplies with buck topology. Power transistors, switching control circuits, and the output decoupling for these FIVRs are fabricated on-die, while the inductors are placed in the package. Feedback loops through dedicated controllers determine the duty cycle of FIVR switches based on the instantaneous output voltage, which is thus adaptively filtered and stabilized. Although FIVRs provide an intended decoupling between the board/package and voltage-regulated core circuitry, a global coupling still exists and must be properly characterized and modeled for full-system power integrity verification. Unfortunately, this type of analysis requires a complete description of all system parts, which for many-core systems may become impractical due to complexity.

In this work, we attempt a complexity reduction through a suitable model order reduction approach [3], [4]. Differently from the black-box approach of [5], based on a compressed representation of the linearized output PDN impedance, in this work we propose a structured projection framework. We first cast the PDN equations in a block-structured system of coupled state-space equations. These include: models of all interconnects and linear passive components, obtained by standard rational fitting of S-parameters from electromagnetic simulations; averaged models of FIVR switches; per-core feedback loops sensing output voltage and providing duty cycle signals through dedicated controllers. The overall coupled system is nonlinear due to the FIVR feedback, however

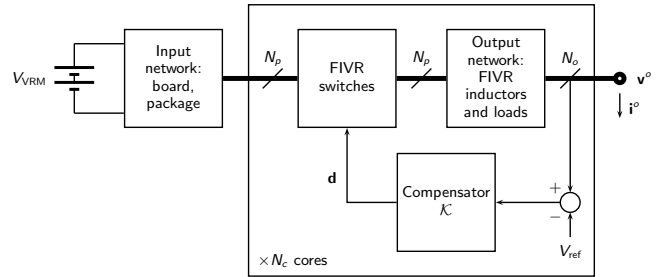


Fig. 1. Schematic illustration of the power distribution system under investigation, including N_c cores whose voltage is regulated by N_p -phase FIVRs.

such nonlinearity is concentrated at the FIVR switch models, whereas all other system parts can be seen as large-scale Linear and Time-Invariant (LTI) blocks. The size of these blocks is reduced through structured and block-partitioned Krylov subspace projection, with projection matrices based on a set of system responses (snapshots) under suitable operating points. The resulting reduced model preserves the block-interconnected structure (with nonlinear feedback) of the original system and is easily solved in time-domain by a basic discretization scheme. Numerical results obtained on a system based on a recent Intel® Core™ microprocessor show a major speedup (up to $200 \times$ with respect to reference SPICE simulations), with a negligible loss of accuracy.

II. PROBLEM STATEMENT AND NOTATION

The structure under analysis is depicted in Fig. 1. The PDN provides the supply voltage to N_c microprocessor cores. In this work, all cores are considered identical, so that the subsystem in the box in Fig. 1 is repeated N_c times. The leftmost voltage source represents the mainboard voltage supply. This is connected to the *input network* block, which represents the electrical behavior of the board and package as a passive LTI system that models the effect of parasitics, on-board decoupling capacitors, and it includes a linear model of the VRM. The *output network* box is another passive linear system describing integrated passive components and the die. The microprocessor units to which the power is supplied are represented as N_o current sources loading each output network. The load currents for all cores are indicated by i^o and the corresponding load voltages are v^o .

Core-level voltage regulation is implemented with FIVRs, depicted in the middle box in Fig.1. These are switching converters with time-varying duty cycle and N_p phases. The i -th core has a dedicated controller that senses the output voltage and steers the duty cycle signal $d_i(t)$ of the respective switches so that the load voltage tracks a constant reference value V_{ref} . In this work, the FIVR switches are represented with ideal transformers, which is a well-known averaged approximation of buck converters. In particular, each core is associated with N_p ports on the output side corresponding to an equal number of ports on the input network side; a per-core bank of N_p ideal transformers connects these two groups of ports.

The overall PDN system can be represented by the following system of differential-algebraic equations

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}_w\mathbf{w} + \mathbf{B}_u\mathbf{u} \quad (1a)$$

$$\mathbf{z} = \mathbf{C}_z\mathbf{x} + \mathbf{D}_{zw}\mathbf{w} + \mathbf{D}_{zu}\mathbf{u} \quad (1b)$$

$$\mathbf{y} = \mathbf{C}_y\mathbf{x} + \mathbf{D}_{yw}\mathbf{w} + \mathbf{D}_{yu}\mathbf{u} \quad (1c)$$

$$\mathbf{w} = \mathbf{\Delta}(\mathbf{d})\mathbf{z} \quad (1d)$$

$$\dot{\mathbf{x}}_k = \mathbf{A}_k\mathbf{x}_k + \mathbf{B}_k(\mathbf{N}\mathbf{y} - \mathbf{V}_{\text{ref}}), \quad \mathbf{d} = \mathbf{C}_k\mathbf{x}_k \quad (1e)$$

where all quantities are defined as follows. Let $\mathbf{v}_1, \mathbf{i}_1$ be vectors collecting all voltages and currents at the $N_c N_p$ ports of the input network (to the left of the FIVRs in Fig.1) and $\mathbf{v}_2, \mathbf{i}_2$ be the corresponding quantities on the output network side. The characteristic equations of the FIVR switches are

$$\mathbf{v}_2 = \mathbf{\Delta}_1(\mathbf{d})\mathbf{v}_1 \quad \mathbf{i}_1 = -\mathbf{\Delta}_1(\mathbf{d})\mathbf{i}_2$$

$$\mathbf{\Delta}_1(\mathbf{d}) \triangleq \text{blkdiag}(\mathbf{\Delta}_{1,1}, \dots, \mathbf{\Delta}_{1,N_c}), \quad \mathbf{\Delta}_{1,i}(\mathbf{d}) \triangleq d_i \mathbb{I}_{N_p}$$

Let us further introduce

$$\mathbf{w} \triangleq \begin{pmatrix} \mathbf{i}_1 \\ \mathbf{v}_2 \end{pmatrix}, \quad \mathbf{z} \triangleq \begin{pmatrix} \mathbf{v}_1 \\ \mathbf{i}_2 \end{pmatrix}, \quad \mathbf{\Delta}(\mathbf{d}) \triangleq \begin{pmatrix} \mathbf{0} & -\mathbf{\Delta}_1(\mathbf{d}) \\ \mathbf{\Delta}_1(\mathbf{d}) & \mathbf{0} \end{pmatrix}$$

so that we can compactly represent the FIVR switches with the equation $\mathbf{w} = \mathbf{\Delta}(\mathbf{d})\mathbf{z}$ in (1d). Furthermore, the behavior of the linear and passive subsystems in Fig. 1, i.e. the input network and all the output networks, can be viewed as a single coupled dynamical system mapping the inputs $V_{\text{VRM}}, \mathbf{i}^o, \mathbf{i}_1, \mathbf{v}_2$ to the outputs $\mathbf{v}_1, \mathbf{i}_2, \mathbf{v}^o$. Hence, by stacking V_{VRM} and \mathbf{i}^o in the vector \mathbf{u} and letting $\mathbf{y} \triangleq \mathbf{v}^o$, we can use the state-space representation for this subsystem in (1a)-(1b)-(1c). Finally, the N_c compensators can be grouped into another linear subsystem whose input is the vector of error signals for each core $\mathbf{e} \triangleq \mathbf{N}\mathbf{y} - \mathbf{V}_{\text{ref}}$, where $\mathbf{N} \in \{0, 1\}^{N_c \times N_o N_c}$ is a matrix that selects N_c entries out of all load voltages to be compared with the vector of constant reference voltages \mathbf{V}_{ref} . The compensators output is the vector of duty cycle signals $\mathbf{d} \in \mathbb{R}^{N_c}$, whose i -th entry is $d_i(t)$. The state-space equations of the compensator subsystem are (1e). The main objective of this work is the definition of a fast transient simulation algorithm that can solve (1) under realistic current excitation of all cores.

III. FORMULATION

Let us analyze the structure of (1). Due to (1d), the overall system is nonlinear, as it involves products of state variables.

Conversely, the first three equations (1a)–(1c) are linear due to the adopted models of input and output networks. The complexity in (1) mainly comes from the first state equation, which has as many components as the dynamic elements used to model the electrical behavior of board, package and die. In the application at hand, \mathbf{x} is high-dimensional and the remaining states \mathbf{x}_k of the duty cycle controllers are few ($\approx 5 - 10N_c$). This motivates a structured approach in which we derive a reduced-order model of (1) by reducing the dimensionality of \mathbf{x} only and leaving the controller equations and the other algebraic constraints untouched. Specifically, we derive the following reduced subsystem by projection of the first three equations in (1)

$$\begin{cases} \dot{\mathbf{x}}_r = \hat{\mathbf{A}}\mathbf{x}_r + \hat{\mathbf{B}}_w\mathbf{w} + \hat{\mathbf{B}}_u\mathbf{u} \\ \mathbf{z} = \hat{\mathbf{C}}_z\mathbf{x}_r + \mathbf{D}_{zw}\mathbf{w} + \mathbf{D}_{zu}\mathbf{u} \\ \mathbf{y} = \hat{\mathbf{C}}_y\mathbf{x}_r + \mathbf{D}_{yw}\mathbf{w} + \mathbf{D}_{yu}\mathbf{u} \\ \mathbf{w} = \mathbf{\Delta}(\mathbf{d})\mathbf{z} \\ \dot{\mathbf{x}}_k = \mathbf{A}_k\mathbf{x}_k + \mathbf{B}_k(\mathbf{N}\mathbf{y} - \mathbf{V}_{\text{ref}}), \quad \mathbf{d} = \mathbf{C}_k\mathbf{x}_k \end{cases} \quad (2)$$

where $\hat{\mathbf{A}} = \mathbf{W}^T \mathbf{A} \mathbf{V}$, $\hat{\mathbf{B}}_w = \mathbf{W}^T \mathbf{B}_w$, $\hat{\mathbf{B}}_u = \mathbf{W}^T \mathbf{B}_u$, $\hat{\mathbf{C}}_z = \mathbf{C}_z \mathbf{V}$, $\hat{\mathbf{C}}_y = \mathbf{C}_y \mathbf{V}$ are the reduced state-space matrices obtained by *Petrov-Galerkin* projection via the biorthogonal matrices \mathbf{V} and \mathbf{W} (i.e., $\mathbf{W}^T \mathbf{V} = \mathbb{I}$). These matrices are here constructed so that the input-output mapping from \mathbf{u} to \mathbf{y} is reproduced accurately, using a specific linearized moment matching approach to deal with the nonlinearity of (1).

A. Local linearization

The main tool we use to overcome the fact that (1) is a nonlinear system is local linearization around an operating point. Given a constant input $\bar{\mathbf{u}}$, the corresponding operating point is $(\bar{\mathbf{x}}, \bar{\mathbf{x}}_k, \bar{\mathbf{w}}, \bar{\mathbf{z}}, \bar{\mathbf{d}}, \bar{\mathbf{u}}, \bar{\mathbf{y}})$, that is a steady state solution of (1) obtained by letting $\mathbf{u}(t) = \bar{\mathbf{u}}$ and $\dot{\mathbf{x}} = \dot{\mathbf{x}}_k = \mathbf{0}$. Let us identify the operating point with $\mathbf{p}(\bar{\mathbf{u}}) \triangleq (\bar{\mathbf{x}}, \bar{\mathbf{w}}, \bar{\mathbf{z}}, \bar{\mathbf{d}}, \bar{\mathbf{u}}, \bar{\mathbf{y}})$, and let us split each variable as $\mathbf{x} = \bar{\mathbf{x}} + \tilde{\mathbf{x}}$, where $\bar{\mathbf{x}}$ is the bias and $\tilde{\mathbf{x}}$ is the small-signal component. The linearized system around $\mathbf{p}(\bar{\mathbf{u}})$ can be easily obtained by replacing (1d) with its first-order approximation $\mathbf{w} = \mathbf{\Delta}(\mathbf{d})\mathbf{z} \approx \bar{\mathbf{w}} + \tilde{\mathbf{w}}$, with $\bar{\mathbf{w}} = \mathbf{\Delta}(\bar{\mathbf{d}})\bar{\mathbf{z}}$, $\tilde{\mathbf{w}} = \mathbf{\Delta}(\bar{\mathbf{d}})\tilde{\mathbf{z}} + \mathbf{\Delta}(\bar{\mathbf{d}})\tilde{\mathbf{z}}$.

The resulting set of equations can be assembled in the compact linearized descriptor form (we omit the detailed expression of the descriptor matrices due to lack of space)

$$\begin{cases} \mathcal{E}\dot{\boldsymbol{\xi}} = \mathcal{A}_p\boldsymbol{\xi} + \mathcal{B}_p\tilde{\mathbf{u}} \\ \tilde{\mathbf{y}} = \mathcal{C}_p\boldsymbol{\xi} + \mathcal{D}_p\tilde{\mathbf{u}} \end{cases} \quad \text{with} \quad \boldsymbol{\xi} = \begin{pmatrix} \tilde{\mathbf{x}} \\ \tilde{\mathbf{x}}_k \\ \tilde{\mathbf{z}} \end{pmatrix}.$$

The transfer function of the linearization of (1) around the operating point associated to $\bar{\mathbf{u}}$ can be written as

$$\mathbf{H}(s; \mathbf{p}(\bar{\mathbf{u}})) = \mathcal{C}_p (s\mathcal{E} - \mathcal{A}_p)^{-1} \mathcal{B}_p + \mathcal{D}_p. \quad (3)$$

The same linearization procedure can be carried out on the projected nonlinear system (2) to obtain the reduced linearized transfer function

$$\hat{\mathbf{H}}(s; \hat{\mathbf{p}}(\bar{\mathbf{u}})) = \hat{\mathcal{C}}_{\hat{\mathbf{p}}} (s\hat{\mathcal{E}} - \hat{\mathcal{A}}_{\hat{\mathbf{p}}})^{-1} \hat{\mathcal{B}}_{\hat{\mathbf{p}}} + \hat{\mathcal{D}}_{\hat{\mathbf{p}}} \quad (4)$$

where $\hat{\mathbf{p}}(\bar{\mathbf{u}})$ is the operating point corresponding to the input $\bar{\mathbf{u}}$ in the reduced system (2). Note that, by enforcing that the operating points match as $\mathbf{p}(\bar{\mathbf{u}}) = \hat{\mathbf{p}}(\bar{\mathbf{u}})$ for a fixed $\bar{\mathbf{u}}$, then

$$\hat{\mathbf{A}}_{\hat{\mathbf{p}}} = \mathcal{W}^T \mathcal{A}_p \mathcal{V}, \quad \hat{\mathbf{B}}_{\hat{\mathbf{p}}} = \mathcal{W}^T \mathcal{B}_p, \quad \hat{\mathbf{C}}_{\hat{\mathbf{p}}} = \mathcal{C}_p \mathcal{V}, \quad \hat{\mathbf{D}}_{\hat{\mathbf{p}}} = \mathcal{D}_p$$

where $\mathcal{V} \triangleq \text{blkdiag}(\mathbf{V}, \mathbb{I}, \mathbb{I})$ and $\mathcal{W} \triangleq \text{blkdiag}(\mathbf{W}, \mathbb{I}, \mathbb{I})$.

B. Projection matrices

This section discusses how to construct the projection matrix \mathbf{V} . The general idea proposed here is that the linearized transfer function of the reduced nonlinear system (4) should (approximately) interpolate the linearization of the original system (3). We consider separately matching at DC and at other frequencies.

DC accuracy: Steady-state conditions are of particular importance in the verification of power delivery networks. Hence, we would like our reduced model to reproduce the exact steady-state response of the full system. In particular, for any constant $\bar{\mathbf{u}}$, the equilibrium solution $\mathbf{p}(\bar{\mathbf{u}})$ of (1) should be the same as $\hat{\mathbf{p}}(\bar{\mathbf{u}})$. This is ensured by taking

$$\mathcal{R}(\mathbf{V}) \supset \mathcal{R}(\mathbf{A}^{-1}(\mathbf{B}_w \ \mathbf{B}_z)) \quad (5)$$

where \mathcal{R} denotes the column space (range). This condition implies that, for any steady-state solution of the original system $(\bar{\mathbf{x}}, \bar{\mathbf{x}}_k, \bar{\mathbf{w}}, \bar{\mathbf{z}}, \bar{\mathbf{d}}, \bar{\mathbf{u}}, \bar{\mathbf{y}})$, the reduced system has a steady-state solution at $(\bar{\mathbf{x}}_r, \bar{\mathbf{x}}_k, \bar{\mathbf{w}}, \bar{\mathbf{z}}, \bar{\mathbf{d}}, \bar{\mathbf{u}}, \bar{\mathbf{y}})$ with $\bar{\mathbf{x}} = \mathbf{V}\bar{\mathbf{x}}_r$. Consequently, $\mathbf{p}(\bar{\mathbf{u}}) = \hat{\mathbf{p}}(\bar{\mathbf{u}})$.

Transfer function moments: In order to enforce $\mathbf{H}(s; \mathbf{p}(\bar{\mathbf{u}})) \approx \hat{\mathbf{H}}(s; \hat{\mathbf{p}}(\bar{\mathbf{u}}))$, we consider a finite set of operating points induced by the inputs $\bar{\mathbf{u}}_1, \dots, \bar{\mathbf{u}}_M$ and a finite set of points $s_1, \dots, s_K \subset \mathbb{C}_+$. For any given $\bar{\mathbf{u}}_j$ with $\mathbf{p}_j = \mathbf{p}(\bar{\mathbf{u}}_j)$, consider the block matrix \mathbf{M}_j defined as

$$\begin{pmatrix} \mathbf{M}_{j,1} \\ \mathbf{M}_{j,2} \\ \mathbf{M}_{j,3} \end{pmatrix} \triangleq \left((s_1 \mathcal{E} - \mathcal{A}_{p_j})^{-1} \mathcal{B}_{p_j}, \dots, (s_K \mathcal{E} - \mathcal{A}_{p_j})^{-1} \mathcal{B}_{p_j} \right)$$

where $\mathbf{M}_{j,1}, \mathbf{M}_{j,2}, \mathbf{M}_{j,3}$ provide a partitioning conformal with ξ . Now if \mathbf{V} is such that $\mathcal{R}(\mathbf{V}) \supset \mathcal{R}(\mathbf{M}_{j,1})$, then $\mathcal{R}(\mathcal{V}) \supset \mathcal{R}(\mathbf{M}_j)$, implying

$$\hat{\mathbf{H}}(s_k; \hat{\mathbf{p}}(\bar{\mathbf{u}}_j)) = \mathbf{H}(s_k; \mathbf{p}(\bar{\mathbf{u}}_j)) \quad k = 1, \dots, K. \quad (6)$$

In words, if the range space of \mathbf{V} contains the image of the first block-row of \mathbf{M}_j , then the linearization of the reduced system (2) around the operating point induced by $\bar{\mathbf{u}}_j$ matches the linearized transfer function of the original system (1) around the same operating point for the prescribed set of frequencies $\{s_k\}$. Finally, multiple operating points are considered at once by collecting all $\mathbf{M}_{j,1}$ in a single matrix Ψ_1

$$\Psi_1 \triangleq (\mathbf{M}_{1,1} \ \cdots \ \mathbf{M}_{j,1} \ \cdots \ \mathbf{M}_{M,1})$$

C. Approximate interpolation via SVD

If the matrix \mathbf{V} is constructed to satisfy

$$\mathcal{R}(\mathbf{V}) \supset \mathcal{R}(\Psi_1), \quad (7)$$

then the interpolation condition (6) will hold for all $j = 1, \dots, M$. However, the dimension of $\mathcal{R}(\Psi_1)$ can quickly grow

large if many operating points and frequencies are considered. Therefore, we give up the exact interpolation condition and construct \mathbf{V} by taking only the first ρ principal components of Ψ_1 as given by its singular value decomposition

$$\Psi_1 = (\mathbf{P}_1 \ \mathbf{P}_2) \text{blkdiag}\{\Sigma_1, \Sigma_2\} \mathbf{Q}^T$$

with the condition (7) relaxed to $\mathcal{R}(\mathbf{V}) \supset \mathcal{R}(\mathbf{P}_1)$. Combining this with the condition for DC accuracy in (5), we can finally state our choice of \mathbf{V}

$$\mathbf{V} = \text{orth}\{\mathbf{P}_1, \mathbf{A}^{-1}(\mathbf{B}_w \ \mathbf{B}_z)\}.$$

D. Stability

In order to establish stability results, let us focus on the first four equations in (1). By viewing it as a linear system with time-varying parameters $\mathbf{d} = \mathbf{d}(t)$, we can use a result analogue to [6, p.63], stating that if the LTI system in (1a)–(1c) is passive and $\Delta(\mathbf{d}) + \Delta(\mathbf{d})^T \preceq 0$, then the time-varying system is stable for any trajectory $d_i(t) \in [0, 1]$. Since the original subsystem represented by (1a)–(1c) is passive and the duty cycle signals are positive and smaller than 1 by construction, the same property can be preserved in the reduced model (2) by preserving passivity in the projection. As explained in [7], this can be done by first finding a positive definite matrix \mathbf{X} that certifies the passivity of the original subsystem based on the Positive Real Lemma. This matrix is then used to construct \mathbf{W} as $\mathbf{W}^T = (\mathbf{V}^T \mathbf{X} \mathbf{V})^{-1} \mathbf{V}^T \mathbf{X}$.

IV. NUMERICAL RESULTS

The proposed approach was tested on a mobile system based on a 4-core Intel® Core™ microprocessor. This is the same system already used in [5], for which each core has $N_o = 36$ output ports where voltage is to be sensed and regulated, as well as $N_p = 4$ ports connecting to the four phases of each FIVR. One output voltage for each core is fed to a feedback controller which determines the duty cycle of the corresponding FIVR switches. Application of proposed MOR scheme led to a reduced model with 527 states, with $\rho = 350$ and 177 additional states to enforce DC accuracy, whereas the full-order system has 2673 state variables. The global basis \mathbf{V} was constructed based on $M = 4$ distinct operating points defined by a sequential activation sequence for the cores.

The resulting system is excited by a sequence of current pulses, whose cumulative waveforms for each core are depicted in the bottom-left panel of Fig. 2. Such currents are uniformly distributed across the ports of each core, so that each core port is excited with the respective waveform divided by 36. The rise time of each pulse is 5 ns. The top panels of Fig. 2 compare the results obtained by proposed reduced-order solver to reference SPICE results, for two different port voltages located on two different cores. Bottom-right panel of Fig. 2 reports the maximum instantaneous deviation between our proposed solution and SPICE, computed for each time step as the largest error among all $N_c N_o = 144$ output voltages. This error is uniformly less than 5 mV, which is well below the engineering accuracy that is requested for this transient




Fig. 2. Transient simulation results of a 4-core microprocessor system. Top: load voltage at selected ports of two different cores (left: port 1 of core 1; right: port 18 of core 3); bottom-left: total load currents applied to each core; bottom-right: maximum instantaneous error across all load voltages.

power integrity verification. Within this good accuracy level, larger instantaneous errors are related to multiple concurrent core switching events.

In terms of runtime, the SPICE simulation required 6880 seconds whereas the proposed model could be solved under the same loading conditions and on the same workstation in 33 seconds. Both simulations were performed with a maximum step size of 50 ps, resulting in about $2 \cdot 10^5$ time steps. The speedup in runtime is about $200\times$. These results confirm the excellent potential of proposed approach for efficient and comprehensive transient power integrity verification.

V. CONCLUSIONS

This paper presented a structured Krylov subspace projection framework for the model order reduction of large-scale descriptions of system-level power distribution networks from VRM to core loads, including per-core integrated voltage regulators. The proposed approach preserves the structure of the initial system through a block-partitioned projection, constructed by assembling snapshots of the linearized system states at suitably chosen operating points. The resulting reduced model can be simulated (including feedback voltage stabilization) very efficiently, as demonstrated by the $200\times$ speedup with respect to SPICE that was achieved on a mobile system based on an Intel® Core™ microprocessor. In order to be ready for routine application, some work is still required for automating the various modeling steps, including selection

of number and placement of the snapshots. This ongoing work will be documented in a future report.

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] E. A. Burton, G. Schrom, F. Paillet, J. Douglas, W. J. Lambert, K. Radhakrishnan, and M. J. Hill, "FIVR — Fully integrated voltage regulators on 4th generation Intel® Core™ SoCs," in *2014 IEEE Applied Power Electronics Conference and Exposition - APEC 2014*, Mar. 2014, pp. 432–439.
- [3] W. H. A. Schilders, H. A. Van Der Vorst, and J. Rommes, *Model order reduction: theory, research aspects and applications*. Springer Verlag, 2008.
- [4] A. C. Antoulas, *Approximation of large-scale dynamical systems*. Society for Industrial and Applied Mathematics, 2005.
- [5] A. Carlucci, S. Grivet-Talocia, S. Mongrain, S. Kulasekaran, and K. Radhakrishnan, "Towards accelerated transient solvers for full system power integrity verification," in *2022 IEEE 31st Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*, 2022, pp. 1–3.
- [6] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear matrix inequalities in system and control theory*. Society for Industrial and Applied Mathematics, 1994, vol. 15.
- [7] A. Megretski, "6.242, Fall 2004: Model Reduction Course Notes, ch. 4, MIT, Department of Electrical Engineering and Computer Science." [Online]. Available: web.mit.edu/6.242/www/images/lec4_6242_2004.pdf