Modeling and Simulation of Nonlinearly Loaded Electromagnetic Systems via Reduced Order Models - A Case Study: Energy Selective Surfaces

Politecnico di Torino

ScuDo
Scuola di Dottorato ~ Doctoral School
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation
Doctoral Program in Electrical, Electronics and Communications Engineering
(XXXIV cycle)

Modeling and Simulation of Nonlinearly Loaded Electromagnetic Systems via Reduced Order Models

A Case Study: Energy Selective Surfaces

By

Marco De Stefano

******

Supervisor(s):
Prof. Stefano Grivet-Talocia

Doctoral Examination Committee:
Prof. Antonio Maffucci, Referee, Università di Cassino e del Lazio Meridionale
Prof. Dries Vande Ginste, Referee, Ghent University
Prof. Christian Schuster, Hamburg University of Technology
Prof. Giuseppe Vecchi, Politecnico di Torino
Prof. Paolo Manfredi, Politecnico di Torino

Politecnico di Torino
2022

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Marco De Stefano
2022

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

Acknowledgements

I would like to express my sincere gratitude to my supervisor Prof. Stefano Grivet-Talocia, for his invaluable patience and guidance through this journey. I would also like to thank Prof. Christian Schuster, for his counseling and for generously hosting me in his group at the Hamburg University of Technology (TUHH). This endeavor would not have been possible without the support of the German Academic Exchange Service (DAAD), who sponsored my visit to TUHH with a scholarship. Additionally, I am extremely grateful to the Deutsche Forschungsgemeinschaft (DFG), who financed the research project that started this fruitful collaboration.

I am also thankful to the members of the EMC group at Politecnico di Torino, whose practical suggestions impacted me several times. Thanks should also go to the people of the TeT group at TUHH for their kind hospitality. I would like to extend my sincere thanks to my office mates for the constructive pieces of advice and experience. Many thanks to Morten Schierholz for providing the via array dataset of Chapter 3.5.1 and for a number of helpful discussions. Special thanks to Torben Wendt for his invaluable support and contribution during these years.

Lastly, I would like to recognize the profound support of my family, especially my parents, girlfriend and friends. To them goes my grateful appreciation.

Abstract

Modern applications require sensitive electronic devices to process high-speed signals under exposure to several disturbing sources. The Electromagnetic Interference (EMI) is one of the main threats for reliable and robust electronic systems, where different external Electromagnetic (EM) sources can compromise (or even destroy) unprotected devices. An example is the High-Intensity Radiated Field (HIRF) induced either by a lightning strike or by a standard radiation transmitter (as well as an intentional EM weapon) that can induce system failure or even damage.

Recently, the Electromagnetic Compatibility (EMC) community has shown an increasing interest in energy-selective surfaces. These novel structures distinguish between high-power interference and weaker signals, enabling protection and (wireless) communication at two different energy levels. A simple energy-selective enclosure is designed by covering an aperture of a conductive shield with a grid of nonlinear devices, usually diodes. The diodes array remains transparent to weak signals, while the energy of strong fields is attenuated thanks to the nonlinear response of the diode grid.

Repeated numerical simulations are required to assess the Shielding Effectiveness (SE) of these structures, i.e. their performances, under different working conditions (in terms of shield parameters, incident field or termination type). On one hand, full-wave solvers allow computing the transient solution of the scattering problem with a relevant computational cost at every change in the system configuration. On the other hand, the unloaded enclosure is a Linear Time Invariant (LTI) system that obeys Maxwell's equations. This observation opens the investigation to hybrid simulation approaches that convert the fully coupled linear/nonlinear EM problem into an equivalent circuit formulation.

The system is represented by a linear multiport loaded with lumped nonlinear terminations, and excited with the contribution of the incident field.

A data-driven macromodeling framework fits the above procedure and enables a reusable Reduced Order Model (ROM) of the (linear) shielding enclosure. Such macromodel can be extracted during a characterization phase, and then exploited to perform multiple efficient transient simulations for performance assessment. Several challenges affect both generation of the ROM and their fast transient simulation. The objective of this work is to address such challenges, in order to establish a complete modeling and simulation framework that is robust and efficient, so that it can be used as a numerical tool during computer-aided design of energy-selective enclosures.

The first contribution of this work is a sequence of data preprocessing strategies, that combine a regularization and extrapolation procedure in a suitable asymptotic modal domain, with structured data compression approaches built on a modified Singular Value Decomposition (SVD). We show that these complementary approaches drastically improve model accuracy and robustness, while reducing model sensitivity and identification complexity. Then, we address the large-scale modeling problem by providing a structured and compressed rational fitting framework. We equip this process with an efficient passivity verification based on an adaptive-sampling strategy, since we have experienced that standard approaches are either unreliable or impractical due to high computational cost. A passivity enforcement scheme takes advantage of this algorithm and provides a final model that is suitable for guaranteed stable numerical simulations. The last contribution of this work is a robust hybrid transient solver that combines the above macromodeling framework with an efficient Waveform Relaxation (WR) based decoupling scheme. The presented result combines an inexact Newton-Krylov iteration and a time partitioning strategy to improve the well-known convergence issues of a standard WR, resulting in a fast and reliable transient solver for energy-selective shields and, in general, for nonlinearly-loaded large-scale electromagnetic structures.

Several numerical results demonstrate how the various formulations and algorithms introduced in this work effectively advance the state-of-the-art.

Contents

Notation

In this dissertation, we refer to $\mathbb{R}$ and $\mathbb{C}$ as the real and complex numerical fields, respectively. We indicate scalars, vectors and matrices with normal, lower case bold and upper case bold fonts as x, $\boldsymbol{x}$ and $\mathbf{X}$, respectively. The identity matrix of size n is denoted as $\mathbb{I}_n$, while the matrix transpose and conjugate transpose (Hermitian) are defined as $\mathbf{X}^{\mathsf{T}}$ and $\mathbf{X}^{\mathsf{H}}$. The imaginary unit is indicated with $\mathrm{j} = \sqrt{-1}$ and the Laplace variable is always referred with s. The conjugate of a complex number z is indicated with z^*. The two operators $\mathrm{Re}\{\cdot\}$ and $\mathrm{Im}\{\cdot\}$ extract the real and imaginary part of their argument. The eigenvalues and singular values of a matrix $\mathbf{X}$ are indicated as $\lambda\{\mathbf{X}\}$ and $\sigma\{\mathbf{X}\}$, respectively. The p-norm of a matrix is indicated with $\|\mathbf{X}\|_p$, where the Euclidean norm ($p = 2$) is assumed when p is omitted. Sets are denoted with calligraphic font $\mathcal{X}$ and their elements are enclosed within curly brackets $\mathcal{X} = \{\cdots\}$. The symbol $\otimes$ indicates the Kronecker product, and its property are explicitly stated when needed. With $\circledast$ we indicate the time-domain convolution operator.

Chapter 1

Introduction

This dissertation addresses the problem of numerical modeling and simulation of large-scale Linear Time Invariant (LTI) systems terminated with nonlinear loads, with a dedicated focus on energy selective shielding enclosures. This chapter introduces the context and motivations of this thesis by presenting the key challenges and providing the general structure of the manuscript.

1.1 Energy selective surfaces: an overview

Since the 1970s, when a common canonical structure was a single wire antenna loaded with a diode [1, 2], the simulation of electrical/electromagnetic structures with nonlinear (NL) terminations has been a major topic of research interest. In this context, the recent interest in energy selective surfaces provides a new challenge in the field of electromagnetic/circuit-simulations [3]. These engineered structures aim at replacing standard absorbers to protect an electronic device from signal interference, such as the Electromagnetic Interference (EMI) given by High-Intensity Radiated Field (HIRF) caused by lightning strikes or by intentional electromagnetic weapons [4]. These applications include metasurface absorbers [5–7], nonlinear impedance surfaces [8], and energy selective shielding [9, 10].

In this scenario, the increasing interest in energy selective surfaces is motivated by their ability to guarantee low-power communication while protecting sensitive devices from dangerous – or even destructive – high-power Electro-

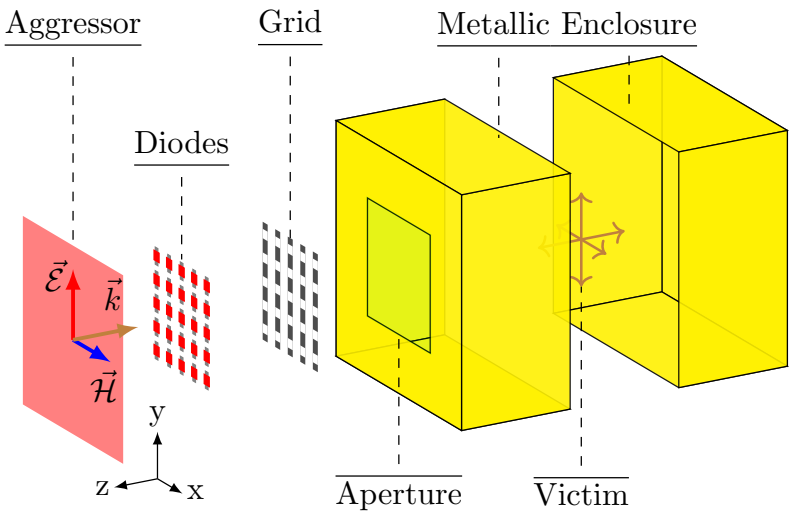

Fig. 1.1 Illustration of an energy-selective shield obtained by the combination of a metallic enclosure with a diode grid. The effect of an incident field (aggressor) on an internal device (victim) is the result of the interplay between the electromagnetic behavior of the shield and the nonlinear characteristics of diode elements. From [10], © 2021 IEEE.

magnetic (EM) sources. The simplest power-modulated protecting effect is obtained by connecting standard shielding enclosures with nonlinear devices at a set of lumped ports, which can be located across an aperture [9–13]. Figure 1.1 illustrates an example of a cubic-shaped shield, that surrounds a sensitive device (victim), with an aperture covered by a diode grid. The penetration of an incident field (aggressor) is modulated by diodes, which provide a nonlinear shielding effect by switching to a conductive mode when triggered by a high-intensity field. Conversely, the internal victim can transmit through the shield aperture at a low-power field level since the currents induced in the metallic enclosure are sufficiently small to leave the terminations in a cut-off non-conducting state. In other words, the sheet of diodes may behave as transparent or reflecting depending on the energy level of the impinging electromagnetic field.

In contrast with standard protection approaches, based on passive linear elements (e.g are wires, screens, apertures), the design phase of nonlinearly-loaded shields cannot be addressed in any analytical closed-form. For this reason, numerical simulations are the main tool to optimize the structure according to the designer requirements. Important aspects to characterize the

electromagnetic protection of an energy selective surface, as in Fig. 1.1, are attributes of both:

- the enclosure, in terms of its geometrical and physical parameters; examples are the shield and aperture size, shape, and the corresponding material properties.
- the nonlinear terminations, as the diodes type and the grid configuration/setting; examples are the number of loads per branch or the number of branches.

A thorough investigation on the effect of various design choices on the field intensity inside the shielded area has been proposed in [10]. In this case, the protection provided by the loaded enclosure is assessed by means of the nonlinear Shielding Effectiveness (SE) defined as

$$SE_{\infty}(\boldsymbol{r}) = \frac{\|e_{\mathrm{inc}}(t)\|_{\infty}}{\|e_{\mathrm{obs}}(t,\boldsymbol{r})\|_{\infty}} \tag{1.1}$$

where $e_{\mathrm{inc}}(t)$ is the incident electric field, $e_{\mathrm{obs}}(t,\boldsymbol{r})$ is the field observed by the victim located at the coordinates $\boldsymbol{r}$, and operator $\|\cdot\|_{\infty}$ extracts the largest amplitude of its argument. The SE is selected as the main design feature since it provides an estimate of the maximum penetrating E-field and enables tuning the energy selectivity of the protective covering, accordingly. To clarify, a small value of $SE_{\infty}(\boldsymbol{r})$ corresponds to a high penetration of the incident field $e_{\mathrm{inc}}(t)$ as a result of a transparent grid of diodes: in this condition, most of the energy is transmitted through the aperture. On the contrary, a large value of SE indicates that the incident field signal is reflected and/or attenuated by the presence of nonlinear terminations.

This is exemplified in Fig. 1.2, which illustrates the effect of several different types of nonlinear terminations on a 8×8 grid spread on a 25×25 cm^2 aperture for an enclosure of size $50 \times 50 \times 50$ cm^3. In the reported example, the selectivity of the SE is modulated using four types of diodes in different configurations: a pair of anti-parallel PN-junction (black), a couple of back-to-back Zener (gray) and two kind of Varactors (red and blue) are combined to obtain six different terminations for the same grid. Details on the nonlinear elements parameters are here omitted (see [10]). According to the SE levels reported in Fig. 1.2, the proposed protective structure behaves as an intensity low-pass (panels a and

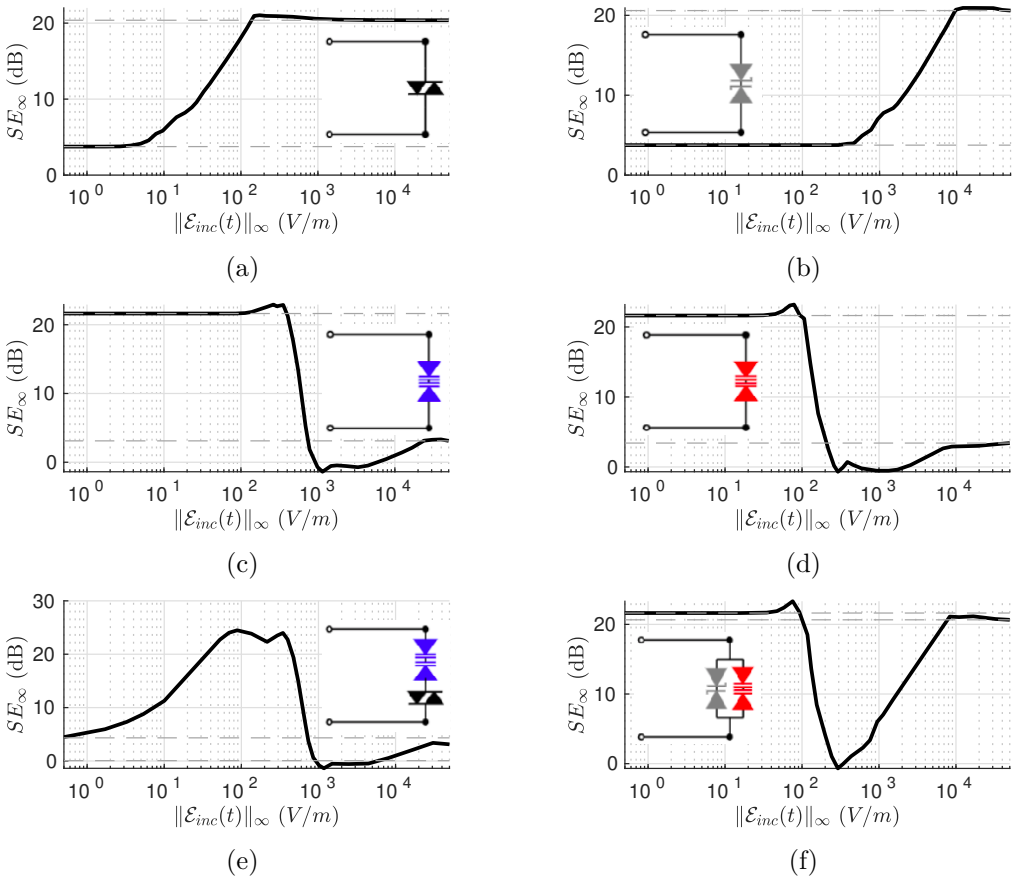

Fig. 1.2 Shielding Effectiveness (SE) plotted as a function of the incident field amplitude, for a shielding enclosure as in Fig. 1.1 with an aperture loaded by a 8×8 grid of diodes. Four types of diodes are used to modify the nonlinear characteristic: PN junction (black), Zener (grey) and Varactors (red and blue). Originally reported in [10], © 2021 IEEE.

b), high-pass (c and d), band-stop (e), and band-pass (f) filter. This example shows that the characteristics of the overall EM protection can be optimized by tuning the nonlinear elements parameters: examples are the threshold voltage, the semiconductor material, or diodes resistance and reactance. The SE behavior is always bounded by two limit cases with diodes completely OFF or ON, resulting in either an open-circuit (OC) or a short-circuit (SC) grid in the extreme situation of low or high intensity fields, respectively. Only the transition between these two limits is the SE component that is (strongly) affected by the termination choice (i.e. diodes parameters), and that requires a (nonlinear) transient simulation to be defined. This is verified for each panel of

Fig. 1.2, where the two bounds are reported as black-dashed lines. These two extreme cases (OC and SC) are fixed once the parameters of the shield are set, in particular shape and size of the enclosure, aperture size and grid density: a linear analysis of the unloaded structure is sufficient to define bounds of the SE. Nevertheless, both investigations (linear and nonlinear) are necessary to fully characterize the performances of a target nonlinearly-loaded shield.

To verify the shield performances and to satisfy the required EM protection for different working conditions, the above procedure must be iterated for several incident field waveforms. To this end, an efficient simulation framework is crucial. The main objective of this work is to present a general modeling and simulation framework, that is able to perform fast nonlinear transient simulation of energy-selective surfaces.

1.1.1 The macromodeling framework

The design of an energy selective surface requires many numerical simulations to assess the EM protection in several working conditions. The main simulation challenge arises from the coexistence of a potentially electrically large enclosure with complex and strongly nonlinear characteristics of possibly many discrete loads. For the analysis of these large-scale nonlinearly loaded systems, we can identify two classes of approaches:

1. *full-wave electromagnetic solvers*
2. *hybrid techniques.*

Several standard full-wave electromagnetic solvers are suitable for our task. Examples are the Time Domain Integral Equation (TDIE) approach [14], the Finite-Difference Time-Domain Method (FDTD) [15], the Time Domain Finite-Element Method (TDFEM) [16] and the Partial Element Equivalent Circuit (PEEC) approach [17]. All these solvers are able to retrieve the dynamic behavior of the shielding structure by coupling the solution of Maxwell's equations with the nonlinear characteristic of the attached loading devices. Nevertheless, all of them suffer from a similar drawback. Anytime a new configuration or loading condition is changed and a new transient analysis is required, all these approaches must start their evaluation from scratch. In other words, even if only part of the problem is slightly modified (as one of

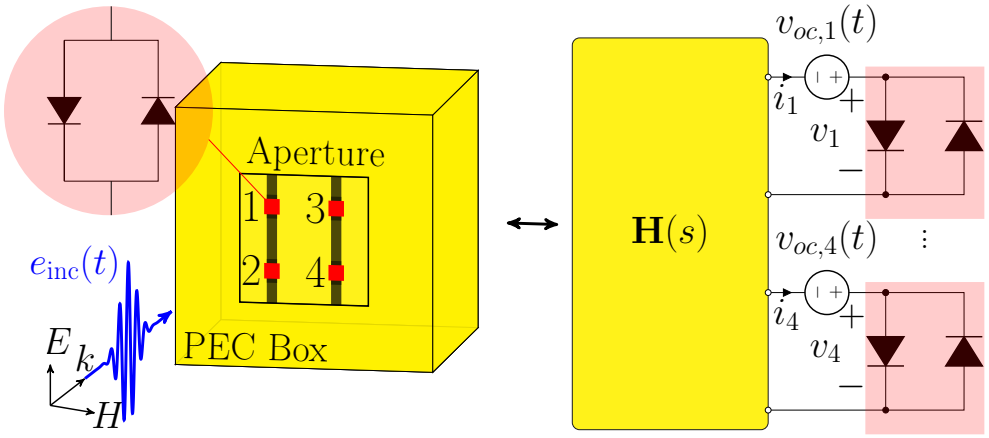

Fig. 1.3 Representation of the shielding enclosure (hybrid) simulation setup: the fully-coupled EM problem (left) is cast into an equivalent circuit representation (right). See Section 4.1 or [13] for a detailed description of the procedure. From [19] © 2022 IEEE.

the many nonlinear devices), the information among successive simulations is lost, and the solver must handle the entire system of (nonlinear) equations again. Thus, using a full-wave electromagnetic solver to perform the (iterative) design required by an energy-selective enclosure may become quite expensive and inefficient when considering several loading and excitation conditions [10].

On the other hand, removing the lumped (nonlinear) loads from the shielding enclosure opens the design investigation to hybrid simulation techniques [13, 18]. Indeed, the unloaded structure (the yellow box with dark metallic connections reported in Fig. 1.3) without the nonlinear interaction of the termination devices (the red elements) is characterized by linear Maxwell's equations. Thus, the shield can be described as a passive multiport LTI system that can be represented via reduced order (compressed) modeling approaches, discussed later in Chapter 2 and 3. The information associated with the transfer matrix $\mathbf{H}(s)$ is thus cast in a macromodel form and preserved for subsequent transient simulations. Then, the original transient scattering field problem can be solved by the scheme of Fig. 1.3 by attaching the nonlinear loads to an equivalent circuit synthesis of the surrogate model and including the incident field contribution with time-varying voltage sources. A detailed description of this hybrid framework will be provided in Section 4.1.

The literature offers several data-driven strategies for a macromodel generation from a frequency characterization of the target system [20–23]. Nowadays, the most relevant algorithm to approximate an LTI component starting from a set of frequency responses is the Vector Fitting (VF) [21]. VF is well established in the academic and industrial communities. With this approach, extracting a rational approximation of a multiport system into a passive macromodel is straightforward. A modified version of the VF [24] provides a remarkable speed-up with a relevant number of ports, while the algorithm parallelization is already well established [25, 26]. The common practice for transient simulation is to convert the resulting surrogate model to an equivalent circuit [20, 27], suitable for any circuit solver of the SPICE family [28, 29].

Nevertheless, our type of application, where the large-scale nature of the linear structure both in terms of dynamical order and number of external ports becomes a limiting factor and a modeling challenge, opens several research directions, discussed next.

1.2 Challenges

Modeling and simulation of energy-elective shielding structures poses a number of challenges from both the macromodeling and the transient simulation standpoint. As an example, let us consider Figure 1.4, which shows a $50 \times 50 \times 50$ cm^3 box-shaped enclosure with a 25×25 cm^2 frontal (squared) aperture, which is covered by a regular grid of $P = p \times p = 25$ ports (in red) connected in series by (black) metal strips of 2 mm width. This structure is used in the following as a running example to describe the above challenges.

1.2.1 Characterization and modeling

Our task is to build a model of LTI system associated to the unloaded box. Figure 1.5 reports the magnitude of impedance and admittance parameters extracted by a Method of Moments (MoM) solver [31], sampled at K frequency points. System responses are commonly available as scattering parameters $\{\mathbf{H}(\mathrm{j}\omega_k) = \mathbf{S}(\mathrm{j}\omega_k),\ k = 1, \ldots, K\}$, but a suitable conversion to admittance $\mathbf{Y}$ or impedance $\mathbf{Z}$ representation is straightforward. The discrete data-set spans

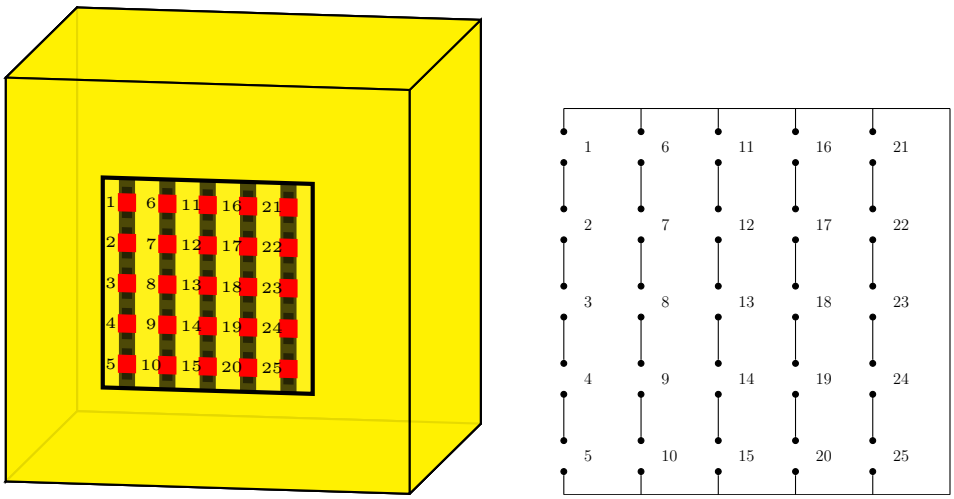

Fig. 1.4 A shielding enclosure with $P = 25$ lumped ports (in red), spread on $p \times p$ regular grid, and its DC conductive path. From [30] © 2022 IEEE.

a frequency band $\omega_k \in \Omega_D = \Omega_L \cup \Omega_H$. Note the presence of an (unknown) low-frequency region Ω_G due to inherent limitations in frequency-domain field solvers based on MoM or Finite Elements Method (FEM).

The realization of passive macromodels from band-limited tabulated frequency responses as in Fig. 1.5 poses a number of challenges:

1. *the zero-frequency (DC) sample is not available.* However, it is fundamental to have full control of the macromodel behaviour under static conditions to obtain a physically consistent time-domain simulation of the structure under nonlinear terminations. Since frequency-domain field solvers usually do not provide accurate responses below a critical $f_{\min}$, to obtain a consistent DC matrix either a static solver or a physics-based approach [32] are usually involved. Given that the shielding enclosure is designed to have very low losses or –in the limit case– lossless when made of PEC, it turns out that the DC point control is particularly challenging. In the PEC case, an equivalent static circuit of the box can be easily obtained as in the right panel of Fig. 1.4. For this structure, both the impedance and admittance matrices do not exist when the frequency vanishes, since the DC conductive path represents an ill-defined circuit when exciting all ports with independent current or voltage sources (see Chapter 2 for details). This is consistent with the low-frequency asymptotic behaviour of both $\mathbf{Y}(\mathrm{j}\omega)$ and $\mathbf{Z}(\mathrm{j}\omega)$ when $\omega \in \Omega_L$ (see Fig. 1.5),

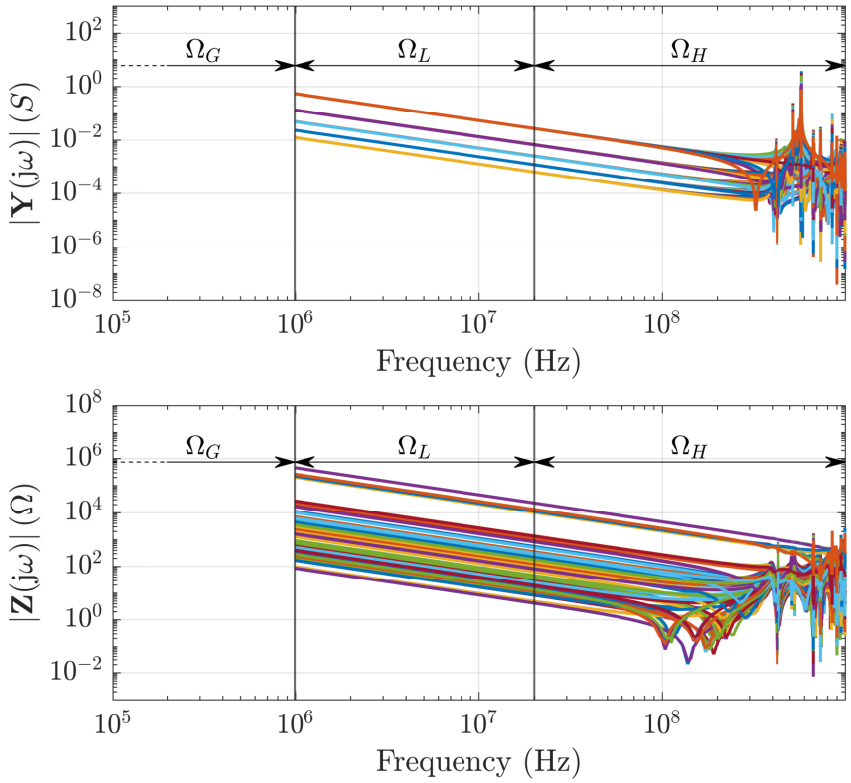

Fig. 1.5 Admittance (bottom) and impedance (top) responses of a 25-port shielding enclosure sampled with a MoM field solver. Adapted from [30] © 2022 IEEE.

which is motivated by the presence of a dominant pole at DC in both cases. Including this pole in the macromodel structure will eventually lead to asymptotic instabilities, affecting the performances of transient simulations. For the above reasons, a *regularization of the DC matrix* is necessary to guarantee the effectiveness of any time-domain simulation, especially when nonlinear terminations are involved.

2. *the absence of reliable data in a broad low-frequency range*, the *gap* Ω_G. This condition precludes any chance of controlling a macromodel behaviour in this frequency region, even if a well-defined DC matrix is included in the training data-set. Nevertheless, during a transient simulation the nonlinear terminations may emphasize the low-frequency content of ports signals, which are then affected by a possibly inaccurate response of the model, also in the case of a high-frequency excitation field. This situation is particular evident observing Fig. 1.6, which shows the accuracy of a passive surrogate model obtained with a state-of-the-art

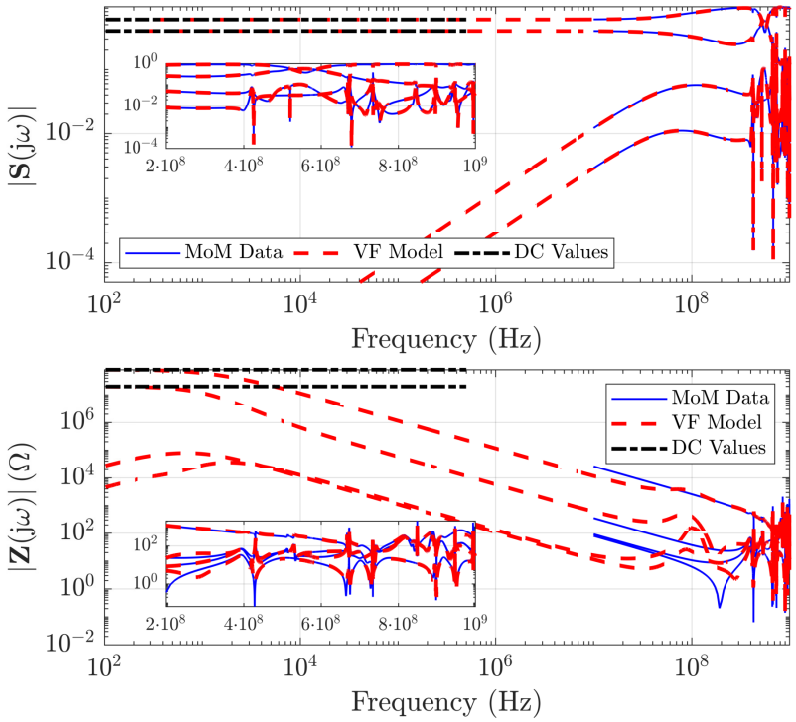

Fig. 1.6 Standard macromodeling tool [33] accuracy on a 25-port shielding enclosure. Top: scattering responses used to fit the model (with exact DC matrix). Bottom: impedance parameters show an accuracy degradation after conversion of model responses. From [30] © 2022 IEEE.

tool [33] by fitting the S-parameters. The original data-set has been enriched with a physically consistent DC matrix, that has been used to enforce an exact macromodel fit at the zero-frequency. The impedance parameters obtained after model responses conversion clearly show an inexact behaviour in the low-frequency range. Therefore, some guided smooth *data extrapolation* is needed to properly merge the DC point with (high-frequencies) field solver samples to obtain a full-band training data-set for MOR engines.

3. *the complexity of the macromodeling problem*, mostly in terms of input/output ports P. Indeed, to properly cover the shielding aperture with a suitable grid the number of nonlinear devices can easily reach hundreds or thousands of elements, each one (or each pair) attached to an electrical port. Furthermore, in shielding applications also the macromodel dynamic order (i.e. the number of poles) is likely to be signif-

icant, depending on the structure electrical size at the largest frequency of interest. Under such conditions, the required computational cost of standard techniques as VF is inappropriate or even unfeasible, especially in the critical case when a passive result is required. Indeed, the *passivity verification* is the main bottleneck in the extraction of a reliable model with a considerable size.

4. *the error magnification with a change system representation.* A broadband accurate macromodel approximation is strongly required to attenuate this challenge, especially approaching DC. On the one hand, the final model representation must retain all low-frequency features of the system behavior, like the presence of zeros at DC. On the other hand, we would like to reduce the large number of responses with a data-compression strategy to mitigate the model identification effort. Unfortunately, a standard compressed framework [34] cannot achieve these two objectives concurrently. Thus, a suitable *structured* approach is required both in the *data pre-processing* and *modeling* stages.

For these reasons, a *large-scale macromodeling* problem requires a suitable *framework* that has the capability to scale up with the number of ports, while both preserving a suitable full-band accuracy and enabling the passivity characterization of the system at hand.

1.2.2 Transient simulation

The simulation problem of Fig. 1.3 can be computed with any compatible circuit solver. Among the others, solvers of the SPICE family are well-known due to their general-purpose capability and reliability. These circuit solvers usually start from a Modified Nodal Analysis (MNA) characterization of the system under analysis, retrieve a set of Ordinary Differential Equations (ODE) equations and then find a transient solution exploiting a Newton-based approach. SPICE engines embed adaptive time sampling techniques and control parameters to suitably achieved a good tread-off between time performances and prescribed accuracy. Nevertheless, it is well-known that a MNA circuit formulation associated with a Newton approach –that requires the evaluation of a Jacobian matrix– is not likely to scale favorably with the number of circuit

components and independent variables. Thus, a general-purpose SPICE solver may not be efficient when applied to the shielding enclosure verification case of Fig. 1.3.

In this manuscript we will provide an alternative simulation framework to general-purpose circuit solvers of the SPICE family. Starting from the linear/nonlinear decoupling approach [13], we exploit a well-know technique to solve the circuit problem of Fig. 1.3, which is the Waveform Relaxation (WR) method [35, 36]. Unfortunately, the basic WR approach is likely to fail for the present application, as discussed next.

Every WR-based algorithm starts from an initial guess of port signals over a given simulation window and refines their estimate at every iteration, until convergence is detected. The simplest implementation of a WR approach involves a fixed-point iteration that evaluates sequentially the electromagnetic subsystem response and the nonlinear terminations one. Under proper conditions (in every formulation the iteration operator must be a contraction [36–40]), signals estimates stabilize. An extensive literature is available for adjusting the parameters of the decoupling scheme such that the number of WR iterations is minimized [37, 41, 39, 42]. A simple choice for this optimization is to define the decoupling scheme using scattering waves, such that a fixed reference impedance level can be used to transfer signals information among iterations.

For the addressed shielding enclosure application, impedance matching is not a reliable option:

- a constant (real) decoupling impedance is inadequate for broadband matching since the low-loss nature of the shield leads to an almost purely reactive behavior in the low-frequency range, counterbalanced by sharp resonances in the high-frequency region;
- a time-varying adaptive approach is not sufficiently robust to track the fast variations of the instantaneous load impedance provided by the switching behavior of diodes terminations, which alternates extreme configurations with very low/high impedance levels associated with conduction/cut-off modes.

For these reasons, a standard WR scheme with a constant reference impendance (e.g. $R_0 = 50\,\Omega$) requires a tremendous number of iterations to converge, as documented in Fig. 1.7. This example shows successive WR iterations performed

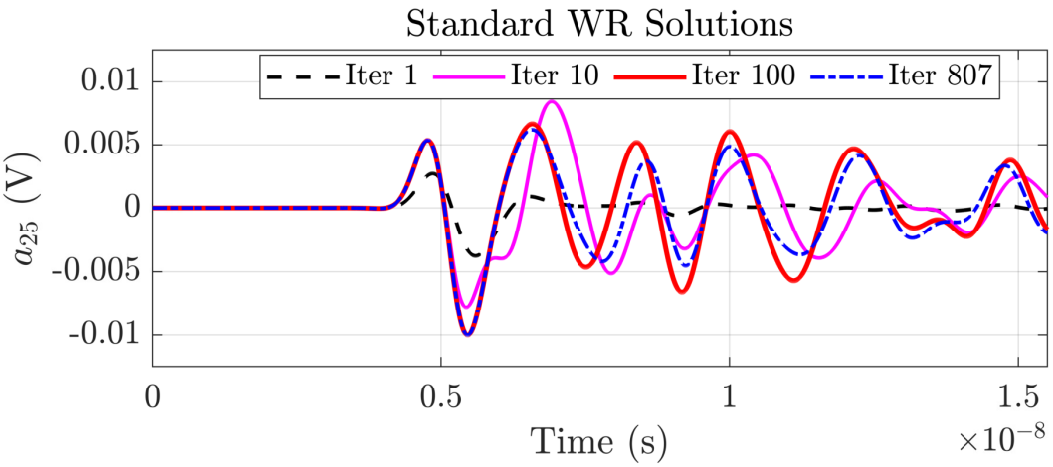

Fig. 1.7 Transient analysis of a 100-ports nonlinearly loaded shielding enclosure. Scattering signals at the 25-th port are reported for different WR iterations. Convergence is detected after 807 iterations.

on a shielding enclosure with 100-ports, loaded with anti-parallel pairs of diodes and excited with a plane wave incident field as detailed in Section 4.6.1. This sample simulation demonstrates that the number of WR iterations is exceedingly large and impractical. In order to reduce these iterations, [18] suggested an iteration-dependent decoupling scheme based on consecutive conversions of the linear macromodel into differently normalized scattering parameters, which showed, however, also the main drawback of increasing the offline cost.

1.3 Technical contributions

All presented challenges will be addressed and solved in this thesis as follows. First, Chapter 2 describes our proposed solutions to all the data-related issues (points 1, 2 and 4 of Section 1.2.1), by providing

- a regularization and extrapolation procedure based on low-frequency asymptotic modes [43], extending the preliminary results of [44, 30]. It will be demonstrated that the proposed pre-processing phase smoothly merges the high-frequency solver data reaching the DC point with a set of synthetic samples, by avoiding the low-frequency degeneracy reported in Fig. 1.5, and by providing as output a well-defined full-band data-set for modeling purposes.

- a data-compression (structured) strategy [44] to preserve low-frequency features while exploiting spatial redundancies in the system responses. Starting from the method proposed in [34] based on a modified Singular Value Decomposition (SVD), which is able to reproduce the set of $P \times P$ system responses with a linear combination of a few basis functions, Chapter 2.2 introduces two representations denoted as Block-Diagonal SVD (BD-SVD) and Hierarchical SVD (Hi-SVD). These two formulations are able to preserve low-frequency data features, such as asymptotic behaviour and expected transfer-function zeros at DC, by improving the reconstruction error and reducing the resulting sensitivity to a change of representation.

Chapter 3 introduces the structured Compressed Macromodel (CM) framework, a model generation strategy that embeds the compression strategies of Chapter 2.2 and drastically reduces the identification problem complexity. The presented procedure improves the low-frequency accuracy by including DC zeros in the model structure within the compressed formulation. The scheme includes a suitable passivity verification scheme based on a multi-stage adaptive sampling approach [45], designed to handle large-scale systems with many electrical ports. A passivity enforcement formulation concludes the presented framework. The final result is a complete procedure to extract compressed and passive macromodels, structurally imposing DC zeros and improving the low-frequency accuracy.

Chapter 4 presents a transient solver based on the above compressed macromodel representation, that allows the evaluation of transient voltages and currents at nonlinearly loaded ports in a drastically reduced runtime with respect to standard circuit solvers. The presented transient simulation algorithm is based on a decoupling Waveform Relaxation (WR) approach, strengthened by a Newton-Krylov iteration as in [46]. The proposed scheme is further enriched by a time-partitioning approach that is able to drastically improve convergence, by reducing the number of iterations per time window and consequently the overall runtime.

Numerical results on shielding enclosure of increasing complexity (up to 1024 ports) are provided along the thesis in each chapter, demonstrating that the proposed framework is as reliable as state-of-the-art tools from both the

macromodeling (as IdEM [33]) and the circuit simulation (like HSPICE [28]) standpoints. In particular, both the compressed model generation – including passivity characterization and data compression – and the proposed transient solver are shown to provide from one to three orders of magnitude speed-up with respect to standard approaches.

To summarize, the two key contributions of this work are:

- Improvement and streamlined generation of compact, passive and reliable macromodels for large-scale components, with hundreds of ports, obtained from band-limited tabulated frequency responses
- A dedicated simulation environment for nonlinearly-loaded large-scale LTI systems, with proven effectiveness for the specific case of energy-selective enclosures.

Conclusions are finally drawn in Chapter 5, which discusses the limitations and future research directions of this work.

Chapter 2

Large-scale data preprocessing

This Chapter aims at solving two major issues associated with low-loss (or lossless) large-scale tabulated frequency data.
First, in view of a macromodeling phase the system must be well characterized over the full bandwidth of interest. This requires to define a suitable training dataset, minimal in some terms, but able to cover the most relevant features of interest. Building a model on corrupted data is likely to result in an unreliable behavior, such that the original system response can not be correctly reproduced. The most (in)famous source of corruption is noise, but a very similar effect can be obtained by removing a portion of the data in a sensitive region such that important features will never be provided to the identification algorithm. It is the case of the presented shielding enclosures, which are defined by means of a Method of Moments (MoM) solver. It is well-known that this class of solvers does not allow obtaining reliable results in a low-frequency range, especially at DC. With the intent of using the resulting macromodel in a transient simulation framework with possibly many nonlinear terminations, having a well-defined DC characterization is fundamental for our purposes. On the other hand, a shielding enclosure lies in the family of low-loss (more likely loss-less) systems. We will show in the following that even if we are able to obtain a low-frequency characterization, the resulting data are very likely to be ill-defined in this band, resulting in modeling sensitivity issues when changing loading condition.
The second major problem in feeding a standard fitting algorithm with a shielding enclosure dataset is the size of the problem at hand. This kind of structures are specifically intended to take advantage of a dense grid of

lumped ports to attenuate the effect of dangerous EM incident fields. The density of this grid is one of the main design parameters, and it is reasonable to assume that a higher number of terminations will improve the overall flexibility of a nonlinear shielding characteristic. This means that to collect frequency responses of the P-port (linear) enclosure we will end up with a large-size dataset, whose memory requirements may compromise the efficiency of standard macromodeling procedures. Let us show this situation with an example. In Chapter 4 we will show results built on a shielding enclosure model in pole-residue form with $P = 1024$ ports and $\bar{n} = 71$ poles. What is the minimum requirement in terms of memory for both the training dataset and the resulting model? Let us assume that we would like to use $K = 501$ frequency samples, available from the field (MoM) solver, to run a model fitting step that guarantees a proper approximation of the system. The overall data memory consumption is $\approx$ 4GB, already considering the simplification of a symmetric transfer matrix. A full-rank model in pole-residue form with $\bar{n} = 71$ poles, either real or complex, requires $\bar{n} + 1$ (real valued) residues matrices. This means that the (symmetric) model itself would require at least 300MB of memory to store the coefficients matrices with double precision. To manipulate such amount of data/coefficients in an efficient (and parallel) way is still a considerable task, which became unfeasible when passivity come into play, as later discussed in Chapter 3. A better approach is to address the complexity in terms ports P by providing suitable data compression techniques [34] that are able to exploit redundancies in the large-scale dataset.

The above alternative strategy splits the effort of identifying model coefficients in two phases:

$\mathcal{P}_1$ *preprocessing phase*, that reduces the complexity of the training samples by means of data compression approaches

$\mathcal{P}_2$ *identification phase*, that take advantage of well-established algorithm to build a suitable model in the reduced dataset space.

These two steps are strongly related one another, since:

- suitable transformations are expected from $\mathcal{P}_1$ to consistently map the outcome of $\mathcal{P}_2$ in the original space dimension;
- the model resulting from $\mathcal{P}_2$ cannot retain any information of the original system that are lost during $\mathcal{P}_1$.

These two conditions seems to be straightforward, but we will see in this Chapter that it may be tricky to guarantee both of them overall the entire model domain of definition, i.e. for the all frequency band of interest. Indeed, particular data features such as low-frequency trends, like DC zeros, maybe not be retained while approximating system responses with a reduced number of basis functions [34].

The focus of this Chapter is the *preprocessing phase* $\mathcal{P}_1$ described above, providing two major contributions originally documented in [30]. First, a regularization and extrapolation procedure will be presented in Section 2.1 to obtain a broadband reliable dataset for macromodeling purposes, with specific focus on the DC characterization of the system. Second, the complexity in terms of number ports will be addressed in Section 2.2 by providing data compression techniques that are able to preserve low-frequency features, such as DC zeros, while approximating system behavior with a linear combination of basis functions.

2.1 Dealing with lossless solver data

Let us consider an unloaded shielding enclosure as depicted in Fig. 1.4, characterized by $P \times P$ band-limited frequency responses from a MoM solver as reported in Fig.1.5. The discrete K samples span a frequency band $\Omega_D \in 2\pi[f_{\min}, f_{\max}]$, with $2\pi f_k = \omega_k \in \Omega_D = \Omega_L \cup \Omega_H$, where responses start from a non-zero frequency $f_{\min} > 0$ due to intrinsic limitations in the solver. We can thus identify three frequency bands, from right to left in Fig. 1.5:

- Ω_H is the range at higher frequencies where resonances are mostly located
- Ω_L is the region where it is reasonable to assume that the structure behaves as electrically small, since both the admittance and impedance responses show an asymptotic trend
- $\Omega_G = (0, 2\pi f_{\min})$ is where the system characterization is not available in the dataset, and can be regarded as a missing *gap* in the data.

The presence of a low-frequency gap Ω_G prevents control on the model behavior, thus of its accuracy and sensitivity, in a broad frequency portion that has an influence on a transient simulation due to the presence of nonlinear loads.

The objective of this Section is to integrate and possibly modify the original frequency samples from the solver so that the two conditions below are satisfied:

$\mathcal{C}_1$: the system responses must be well-defined and non-singular at any frequency (including DC), irrespective to the multiport representation;

$\mathcal{C}_2$: the structure must be fully characterized over the entire frequency band of interest; any gap in the data must be avoided.

To this end, two procedures will be illustrated

- a *regularization* approach to fulfill condition $\mathcal{C}_1$ and rule out singularities on the imaginary axis at DC;
- an *extrapolation* step to recover the unknown responses in the low-frequency gap Ω_G and satisfy condition $\mathcal{C}_2$.

It will be shown that a direct extrapolation from system responses (either impedance, admittance or scattering) to extend the asymptotic behavior and fill the low-frequency gap may be ineffective. For these reasons, a more advanced and robust approach based on a regularized modal extrapolation [30, 44] is presented in the following sections.

2.1.1 DC Regularization

Above considerations are now made more precise, with reference to the DC circuit of a box structure with no metal losses, reported in the right panel of Fig. 1.4. As previously stated, both admittance $\mathbf{Y}_0 = \mathbf{Y}(s = 0)$ and impedance $\mathbf{Z}_0 = \mathbf{Z}(s = 0)$ matrices are ill-defined at DC. Indeed, by exciting all ports with independent voltage/current sources, attempting to compute admittance/impedance parameters, leads to an ill-defined circuit due to the presence of

- p independent voltage sources loops lead to a singular $\mathbf{Y}_0$, see left panel of Fig.2.1;
- $p(p-1)$ independent current sources cutsets cause a singular $\mathbf{Z}_0$, see right panel of Fig.2.1.

This observation is consistent with the presence of a pole at $s = 0$ in both representations. Therefore, the goal of this section is to apply a minimal

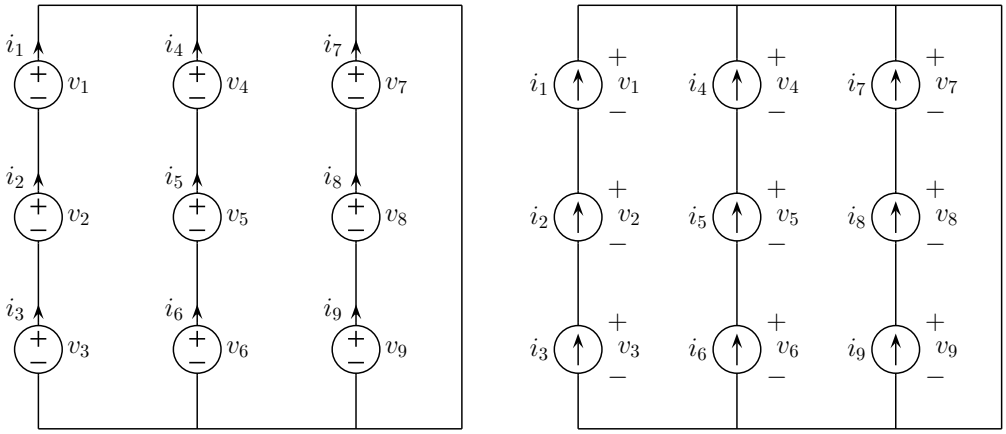

Fig. 2.1 Setup to compute admittance (left) and impedance (right) DC matrices of a $p \times p = 9$ ports shielding enclosure, as in Fig. 1.4. Both circuits are ill-defined at DC, and do not admit any solution.

modification to the ideal DC circuit of Fig. 1.4 to remove the pole on the imaginary axis. To this end, it is possible to add at each port

1. a *series* resistance r to regularize $\mathbf{Y}(s)$;
2. a *shunt* resistance R to regularize $\mathbf{Z}(s)$.

The final result is the topology presented in Fig. 2.2, which is non-singular at all frequencies for any system representation. This leads to the following DC matrices[1]

$$\mathbf{Z}_0 = \mathbb{I}_p \otimes \mathbf{Z}'_0, \qquad \mathbf{Z}'_0 = r\,\mathbb{I}_p + R\,(\mathbb{I}_p - \vartheta \boldsymbol{u}\boldsymbol{u}^\mathsf{T}) \tag{2.1}$$

$$\mathbf{Y}_0 = \mathbb{I}_p \otimes \mathbf{Y}'_0, \qquad \mathbf{Y}'_0 = \frac{(p\,r\mathbb{I}_p + R\boldsymbol{u}\boldsymbol{u}^\mathsf{T})}{p(Rr + r^2)} \tag{2.2}$$

$$\mathbf{S}_0 = \mathbb{I}_p \otimes \mathbf{S}'_0, \qquad \mathbf{S}'_0 = \frac{\Phi' - 1}{\Phi' + 1}\mathbb{I}_p - \frac{2\vartheta\Phi\boldsymbol{u}\boldsymbol{u}^\mathsf{T}}{(\Phi' + 1)(\varphi + 1)} \tag{2.3}$$

with p denoting the number of vertically aligned ports, $\vartheta = 1/p$, $\boldsymbol{u}^\mathsf{T} = [1, \cdots, 1]$, $\boldsymbol{u} \in \mathbb{R}^p$, $\Phi = R/R_0$, $\varphi = r/R_0$, $\Phi' = \Phi + \varphi$, and R_0 indicating the reference scattering impedance. These expressions assume a vertical numbering of the ports, as indicated in Fig. 1.4. In case a horizontal numbering is adopted, the

[1] A step-by-step direct computation of the DC impedance matrix $\mathbf{Z}_0$ is reported in Appendix A.

same expressions hold provided that the ordering of the terms in the Kronecker products is reverted.

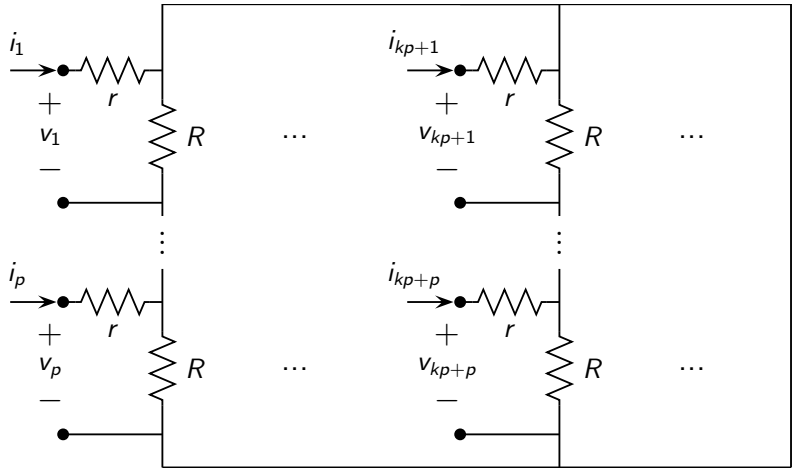

Fig. 2.2 Proposed regularized topology. Originally reported in [30] © 2022 IEEE.

2.1.2 Regularization of MoM data

The proposed regularized topology of Fig. 2.2 imposes the DC responses of the system, making the original data incompatible with the DC behavior since not included by the field solver solution in the first place. For this reason, the original data for all frequency points in Ω_D needs to be modified by ensuring compatibility of all samples with the introduced DC point. This leads to the regularization procedure defined in Algorithm 1, which is built on the assumption of a missing DC point in the available frequency data $\{\mathbf{H}(\mathrm{j}\omega_k),\, \omega_k \in \Omega_D\}$. The procedure computes a set of modified responses by connecting at each port a *shunt* resistance R and then a *series* resistance r. Notice that all matrices inversions are well-defined at any available discrete frequency sample, so that this strategy does not suffer from numerical issues.

Figure 2.3 illustrates the final result of the proposed data regularization in terms of impedance parameters. It is evident from Fig. 2.3 that the perturbation introduced to the original MoM data by this procedure is practically invisible at all frequency samples, mainly due to the particular selected values for shunt and series resistances, in this case $R = 100$ MΩ and $r = 0.1$ Ω respectively. There is not an optimal choice for these two regularization parameters, but to

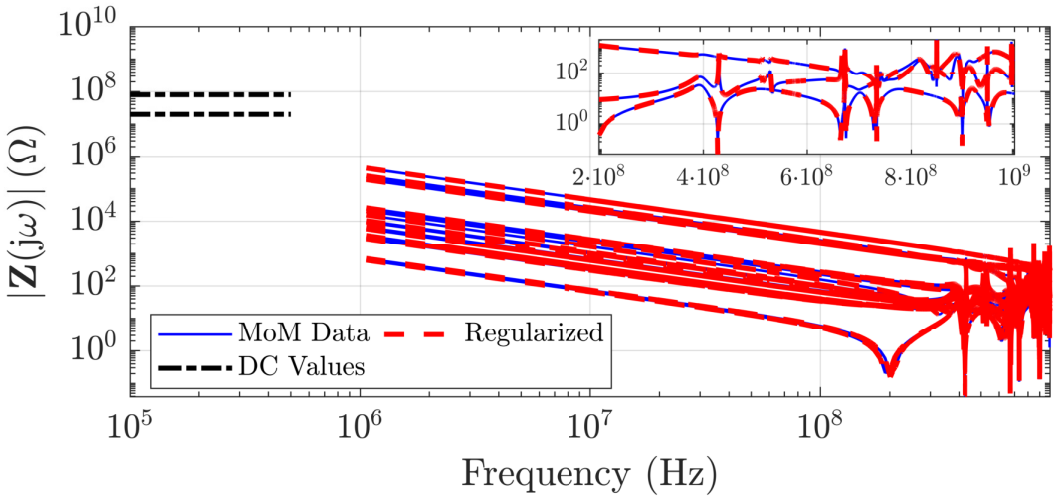

Fig. 2.3 Set of regularized impedance parameters of a 5×5 shielding enclosure ($P = 25$) and DC responses, represented as black-dashed lines. Regularization is performed by Algorithm 1, choosing $R_p = 100$ MΩ and $r = 0.1$ Ω. The high-frequency detail confirms a broad-band negligible perturbation of the proposed procedure. Originally reported in [30] © 2022 IEEE.

reduce the perturbation on the original data their value should be very large, in the case of R, and small, considering r. This is later confirmed by numerical results in Section 2.1.6.

On the other hand, Fig. 2.3 shows also that the first available frequency sample (at $f_{\min}$) does not merge smoothly with the selected DC values, reported to the low-frequency band as black-dashed lines. Indeed, the frequency gap Ω_G is still too wide to reliably characterize the system behavior from $f_{\min}$ up to DC. A procedure to provide consistent extrapolated samples in this unknown frequency region, and satisfy condition $\mathcal{C}_2$, is needed.

Algorithm 1 Data regularization as in [30]

Require: Frequency data $\{\mathbf{H}(j\omega_k),\ \omega_k \in \Omega_D\}$, R, r
1: $\mathbf{H}(j\omega_k) \leftarrow 0.5(\mathbf{H}(j\omega_k) + \mathbf{H}(j\omega_k)^{\mathsf{T}})$
2: Convert $\mathbf{H}(j\omega_k)$ to admittance $\mathbf{Y}(j\omega_k)$
3: Add *shunt* resistance R to $\mathbf{Y}(j\omega_k) \leftarrow \mathbf{Y}(j\omega_k) + \frac{1}{R}\mathbb{I}_P$
4: Add *series* resistance r to $\mathbf{Z}(j\omega_k) = [\mathbf{Y}(j\omega_k)]^{-1}$
5: Regularize impedance as $\mathbf{Z}(j\omega_k) \leftarrow \mathbf{Z}(j\omega_k) + r\mathbb{I}_P$
6: Add DC point $\mathbf{Z}_0$ from (2.1) as $\mathbf{Z}(j\omega)|_{\omega=0} = \mathbf{Z}_0$
7: **return** $\mathbf{Z}(j\omega_k)$: regularized data including DC point

2.1.3 Modal asymptotic extrapolation

Filling the low-frequency gap region with extrapolated and consistent responses may seem a simple task. Indeed, both the admittance and impedance responses show an asymptotic low frequency behavior, such that $\mathbf{Y}(\mathrm{j}\omega) \to \infty$ and $\mathbf{Z}(\mathrm{j}\omega) \to \infty$ for $\omega \to 0$ (see Fig. 1.5). A direct extrapolation of matrix elements from one of the two representations seems sufficient to obtain the required missing samples. Nevertheless, these behaviors imply the presence of both inductive and capacitive modes, responsible for the DC singularities of both system representations. This means that a direct fit of the impedance parameters would recover accurately only capacity modes, more visible in this representation, providing with a rough approximation of the inductive ones The opposite situation arises when considering the admittance parameters, as illustrated with an example in Fig. 2.5.

A better alternative to a direct extrapolation is the conversion to a modal domain, where modes are well distinguishable [44, 30]. This leads to a concurrent approximation of inductive/capacitive modes best suited for the lossless raw MoM data: the procedure presented below should be performed before applying Algorithm 1 and the regularization step of Section 2.1.2. Nevertheless, regularization and extrapolation are not independent, since the analytical DC matrix $\mathbf{Y}_0$ of (2.2) is used here to define the required modal basis. This provides two major differences with respect to other approaches based on direct computation of system modes [43]. First, these other methods are less numerically robust since the direct computation from sampled responses exposes the procedure to data corruptions due to noise or low-frequency solver inaccuracy. Second, the choice of using the real-valued (symmetric) DC matrix $\mathbf{Y}_0$ leads to a real (and orthogonal) modal transformation matrix, necessary to preserve realness and causality [47, 48] when using the data for building a model. This choice is further justified by the assumption that low-frequency modes at non-vanishing frequencies smoothly extend to the corresponding modes at DC.

The proposed procedure starts with the eigendecomposition of the DC admittance matrix $\mathbf{Y}_0$

$$\mathbf{Y}_0 = \mathbf{Q}\boldsymbol{\Lambda}_0\mathbf{Q}^\mathsf{T}, \quad \boldsymbol{\Lambda}_0 = \mathbf{Q}^\mathsf{T}\mathbf{Y}_0\mathbf{Q} = \begin{bmatrix} \frac{1}{R'}\mathbb{I}_c & \mathbf{0} \\ \mathbf{0} & \frac{1}{r}\mathbb{I}_\ell \end{bmatrix} \tag{2.4}$$

with $R' = R + r$. The eigenvalues multiplicity is $\ell = p$ for inductive modes $1/r$ and $c = p(p-1)$ for capacitive modes $1/R'$, in agreement with observations on admittance and impedance singularities of Section 2.1.1.

The orthogonality of modal matrix $\mathbf{Q} = \begin{bmatrix}\mathbf{Q}_p & \mathbf{Q}_s\end{bmatrix}$ is guaranteed by the symmetry of $\mathbf{Y}_0$, which enables to define $\mathbf{Q}_s^\mathsf{T}\mathbf{Q}_s = \mathbb{I}_\ell$, $\mathbf{Q}_p^\mathsf{T}\mathbf{Q}_p = \mathbb{I}_c$, and $\mathbf{Q}_p^\mathsf{T}\mathbf{Q}_s = \mathbf{0}$. Therefore,

$$\mathbf{Y}_0 = \frac{1}{R'}\mathbf{Q}_p\mathbf{Q}_p^\mathsf{T} + \frac{1}{r}\mathbf{Q}_s\mathbf{Q}_s^\mathsf{T}\,. \tag{2.5}$$

Correspondingly, the physical and modal DC impedance matrices read

$$\mathbf{Z}_0 = R'\mathbf{Q}_p\mathbf{Q}_p^\mathsf{T} + r\mathbf{Q}_s\mathbf{Q}_s^\mathsf{T} \quad \text{and} \quad \mathbf{\Lambda}_0^{-1} = \mathbf{Q}^\mathsf{T}\mathbf{Z}_0\mathbf{Q} = \begin{bmatrix} R'\mathbb{I}_c & \mathbf{0} \\ \mathbf{0} & r\mathbb{I}_\ell \end{bmatrix}. \tag{2.6}$$

The modal transformation into DC asymptotic domain holds for all frequencies in the original dataset $\omega_k \in \Omega_D$ as

$$\mathbf{Y}_m(\mathrm{j}\omega_k) = \mathbf{Q}^\mathsf{T}\mathbf{Y}(\mathrm{j}\omega_k)\mathbf{Q}\,, \quad \omega_k \in \Omega_D. \tag{2.7}$$

The same reasoning can be applied to both the impedance and scattering representation. As already mentioned, $\mathbf{Q}$ is a real-valued constant matrix such that the above similarity transformation (2.7) provides the modal admittance elements as linear combinations of the physical ones with real coefficients.

The result of (2.7) on both admittance and impedance parameters are reported in Fig. 2.4, which shows the typical ± 20 dB/dec slope of purely inductive/capacitive modes in the low-frequency range Ω_L. Note that those responses (modes) that tend to zero in one representation are the same that explode to ∞ for the other one.This trend confirms what mentioned before, and further illustrated later in Fig. 2.5, on the sensitivity of a direct extrapolation of the asymptotic behavior from admittance or impedance (physical) responses.

With the goal of obtaining a low-frequency extrapolation, the following analysis is restricted only to the band Ω_L where the modal responses show the ± 20 dB/dec trend. Imposing such asymptotic behavior leads to

$$\mathbf{Y}_m(\mathrm{j}\omega) \approx \frac{1}{\mathrm{j}\omega}\mathbf{\Gamma}_m + \mathrm{j}\omega\mathbf{C}_m,\ \omega \in \Omega_L \tag{2.8}$$

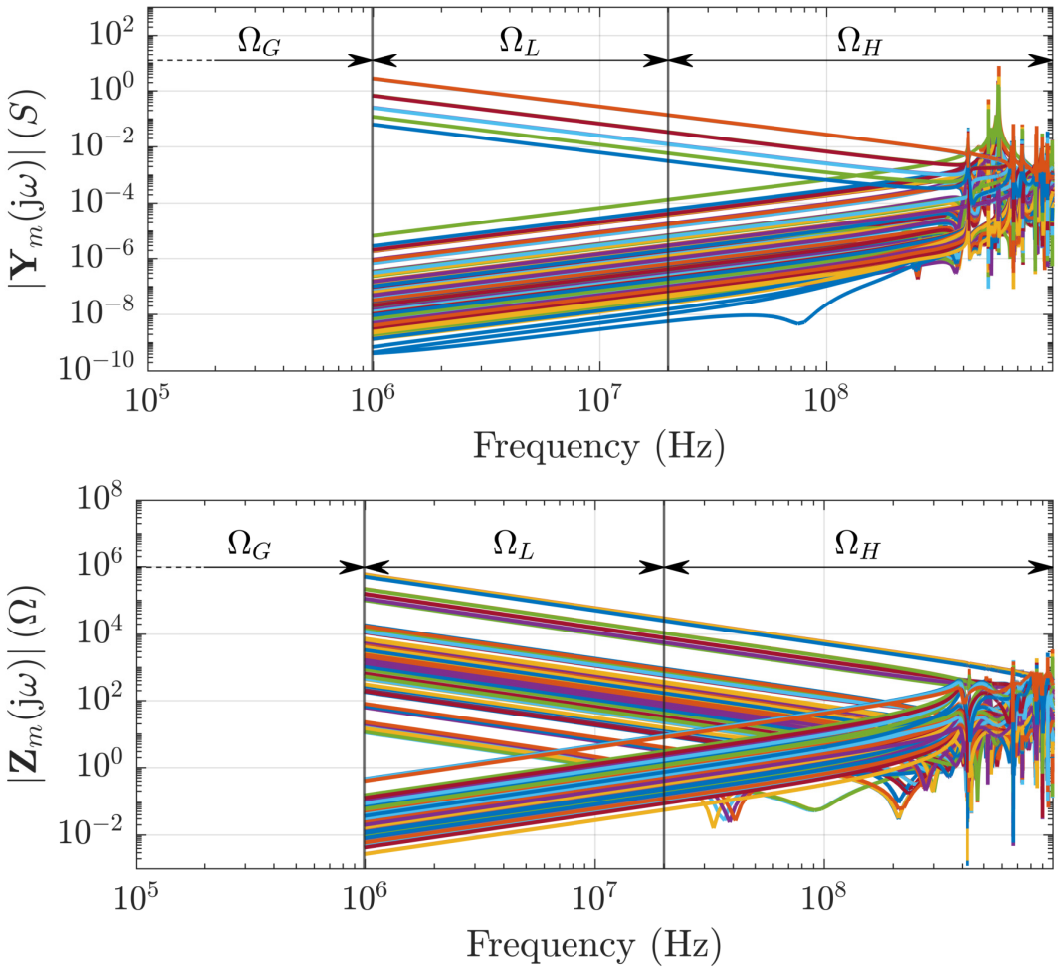

Fig. 2.4 Admittance (bottom) and impedance (top) responses of a 25-port shielding enclosure after conversion to the DC asymptotic modal domain via (2.7). Physical parameters sampled with a MoM field solver are reported in Fig. 1.5. Partially adapted from [30] © 2022 IEEE.

where $\mathbf{\Gamma}_m$ and $\mathbf{C}_m$ are constant matrices such that $(\mathbf{\Gamma}_m)_{ij} \neq 0 \Rightarrow (\mathbf{C}_m)_{ij} = 0$ for any (i,j)-th entry, and vice versa. Correspondingly, the same approximation holds in the physical domain

$$\mathbf{Y}(\mathrm{j}\omega) \approx \mathbf{Q}\left(\frac{1}{\mathrm{j}\omega_k}\mathbf{\Gamma}_m + \mathrm{j}\omega\mathbf{C}_m\right)\mathbf{Q}^{\mathsf{T}} = \frac{1}{\mathrm{j}\omega}\mathbf{\Gamma} + \mathrm{j}\omega\mathbf{C},\ \omega \in \Omega_L \tag{2.9}$$

where $\mathbf{\Gamma}$ and $\mathbf{C}$ are no more sparse but full. Note that when $\omega \to 0$ the last term in (2.9) becomes negligible and all elements of $\mathbf{Y}(\mathrm{j}\omega) \to \infty$ as expected. Taking a closer look to (2.8), the leading capacitive and inductive terms can

be separated in a matrix form as

$$\mathbf{Y}_m(\mathrm{j}\omega) \approx \begin{bmatrix} \mathrm{j}\omega\widetilde{\mathbf{C}} & \mathrm{j}\omega\widetilde{\widetilde{\mathbf{X}}} \\ \mathrm{j}\omega\widetilde{\widetilde{\mathbf{X}}}^{\mathsf{T}} & \frac{1}{\mathrm{j}\omega}\widetilde{\mathbf{\Gamma}} \end{bmatrix}, \ \omega \in \Omega_L \tag{2.10}$$

where $\widetilde{\mathbf{C}}$, $\widetilde{\mathbf{\Gamma}}$, $\widetilde{\widetilde{\mathbf{X}}}$ are constant matrix blocks.

Figure 2.5 illustrates the result of a direct fit of $\mathbf{Y}(\mathrm{j}\omega)$, limited to the range Ω_L, imposing the low-frequency approximation (2.9). The extrapolated responses in both the physical and modal domain, obtained via post-processing, show how a perfect accuracy in the admittance representation (top panel) is lost converting to the impedance parameters (bottom) due to the poor approximation of some system modes (middle). Indeed, the capacitive modes (the ones that tend to zero in $\mathbf{Y}_m(\mathrm{j}\omega)$) are visible only in the impedance representation, hence they can be merely roughly delineated observing the admittance responses. This confirms that a direct extrapolation of (2.9) is not the best possible approach.

A two-step element-wise regression can now be applied to recover each term of (2.10) from the original data samples.

As initial step, each (i, j)-th entry of (2.10) is classified either as inductive or capacitive by using the asymptotic slope of the (modal) admittance magnitude. For each frequency in the low-frequency range, a least squares fit is performed to compute coefficients μ_{ij} and ν_{ij} as

$$y_{ij;k} \approx \mu_{ij}\, \xi_k + \nu_{ij}, \quad \forall \omega_k \in \Omega_L \tag{2.11}$$

where $\xi_k = \log_{10} \omega_k$ and $y_{ij;k} = \log_{10} |Y_{m;ij}(\mathrm{j}\omega_k)|$. While ν_{ij} is disregarded, the other coefficient is rounded to $\mu_{ij} = -1$, to indicate that the (i, j)-th entry belongs to the inductive term $\mathbf{\Gamma}_m$, or to the capacitive group $\mathbf{C}_m$ for which $\mu_{ij} = 1$. The result of this step for two shielding enclosures (with 25 and 64 ports) is reported in Fig. 2.6, where filled dots denote matrix elements assigned to the capacitive term $\mathbf{C}_m$. Figure 2.6 confirms the low-frequency structure assumed in (2.10).

The second step assigns numerical values to the matrices $\mathbf{C}_m$ and $\mathbf{\Gamma}_m$. This is again achieved via a regression procedure, arranged differently for

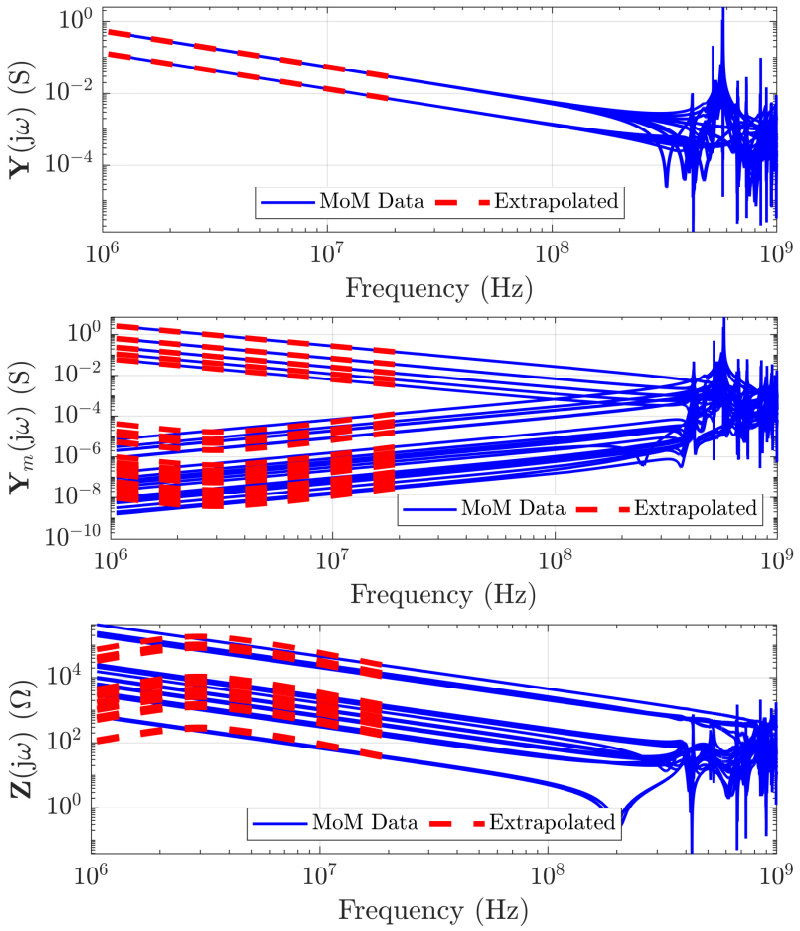

Fig. 2.5 Direct fit of admittance responses imposing (2.9) for a 25-port shielding enclosure. Only dominant inductive modes are accurately approximated, resulting in a poor representation of the impendace parameters. Top: original admittance responses; middle: after conversion to the DC asymptotic modal domain via (2.7); bottom: impedance responses, obtained as post-processing of the top panel.

capacitive/inductive modes as

$$\begin{aligned} \mu_{ij} = +1: &\qquad \omega_k C_{m;ij} \approx \eta_{ij;k},\ \forall \omega_k \in \Omega_L \\ \mu_{ij} = -1: &\qquad -\omega_k^{-1}\Gamma_{m;ij} \approx \eta_{ij;k},\ \forall \omega_k \in \Omega_L \end{aligned} \tag{2.12}$$

where now only the imaginary part of the response is used to define $\eta_{ij;k} = \mathrm{Im}\,\{Y_{m;ij}(\mathrm{j}\omega_k)\}$. The coefficients from (2.12) are assembled to form a low-frequency extrapolation model

$$\boldsymbol{\Upsilon}_m(\mathrm{j}\omega_k) = \frac{1}{\mathrm{j}\omega_k}\boldsymbol{\Gamma}_m + \mathrm{j}\omega_k \mathbf{C}_m, \quad \omega_k \in \Omega_G \cup \Omega_L\,. \tag{2.13}$$

A set of synthetic samples K_{GL} can be generated by evaluating (2.13) for $\omega_k \in \Omega_G \cup \Omega_L$ thus filling the low-frequency gap with physical consistent new samples.

The proposed two-step procedure is completely data-driven, and the only necessary prior assumption is the asymptotic modal expansion as in (2.8) and (2.13). The low-frequency matrix structure defined in (2.10) is not required a priori, but it is obtained numerically. Further, notice that the resulting low-frequency model (2.13) is evaluated both in the gap bad Ω_G and in its training band Ω_L. This enables to smoothly merge the synthetically-generated responses and the solver data as

$$\mathbf{Y}_m(\mathrm{j}\omega_k) \leftarrow \begin{cases} \mathbf{\Lambda}_0 & \omega_k = 0 \\ \mathbf{\Upsilon}_m(\mathrm{j}\omega_k) & \forall \omega_k \in \Omega_G \\ \alpha_k \mathbf{Y}_m(\mathrm{j}\omega_k) + \beta_k \mathbf{\Upsilon}_m(\mathrm{j}\omega_k) & \forall \omega_k \in \Omega_L \\ \mathbf{Y}_m(\mathrm{j}\omega_k) & \forall \omega_k \in \Omega_H \end{cases} \tag{2.14}$$

where the first line allows including the DC point. Via (2.14) all samples in the overlap region Ω_L are averaged via frequency-dependent coefficients α_k and β_k, which are linearly (or logarithmically)-spaced from 0 to 1 in Ω_L preserving the relation $\alpha_k + \beta_k = 1$. This refinement step enables to avoid a possibly inconsistent (or even sharp) transition between computed $\mathbf{\Upsilon}_m(\mathrm{j}\omega_k)$ and original data $\mathbf{Y}_m(\mathrm{j}\omega_k)$.

2.1.4 Regularization in the modal domain

The proposed data preprocessing combines the regularization and extrapolation discussed in Sec 2.1.1 and 2.1.3. The actual proposed algorithm aims at perturbing only modes that are responsible for singularities in a given representation, leaving the rest of the (modal) matrix unaltered. This approach reduces the overall perturbation with respect to a full regularization in the physical domain, which indirectly affects all modes concurrently.

The proposed modal regularization strategy is applied at all non-zero frequencies of the data samples obtained in (2.14), and its main steps are

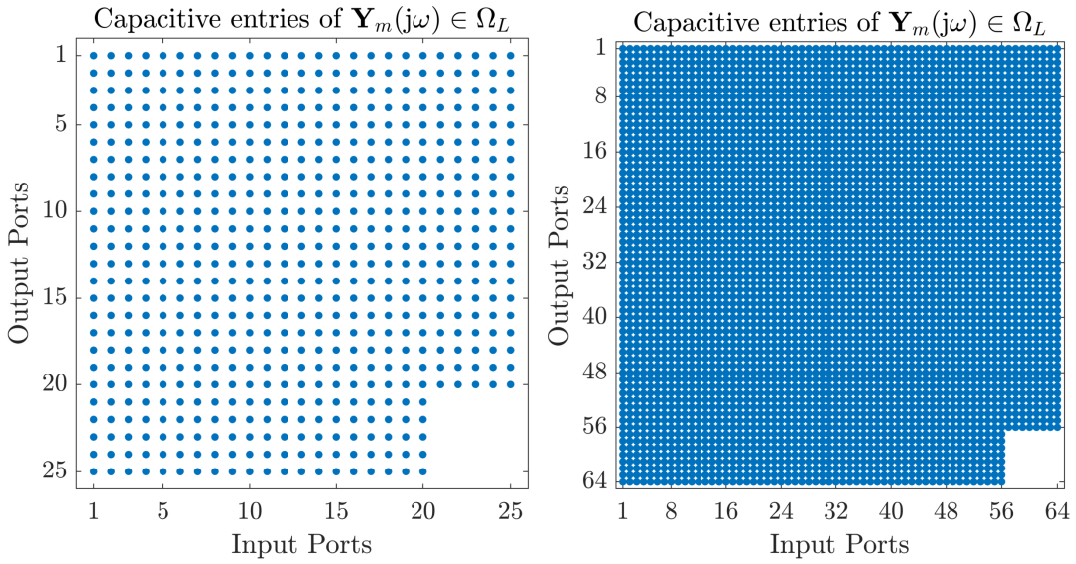

Fig. 2.6 Low-frequency modal behavior of $\mathbf{Y}_m(\mathrm{j}\omega) \in \Omega_L$: each dot corresponds to a coefficient $\mu_{ij} = +1$, which indicates membership of a transfer function entry to the capacitive group $\mathbf{C}_m(\mathrm{j}\omega)$. Results are reported for boxes with $P = 5 \times 5 = 25$ (left) and $P = 8 \times 8 = 64$ (right) ports. The modal structure (2.10) is verified for all investigated cases.

1. regularization of capacitive blocks, that cause the degeneracy of admittance matrix $\mathbf{Y}_m$ at $\omega \to 0$, by adding the contribution of DC eigenvalues R' as

$$\widehat{\mathbf{Y}}_m(\mathrm{j}\omega_k) = \mathbf{Y}_m(\mathrm{j}\omega_k) + \begin{bmatrix} \frac{1}{R'}\mathbb{I}_c & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \ \omega_k \in \Omega_G \cup \Omega_D. \tag{2.15}$$

2. regularization of inductive blocks of $\widehat{\mathbf{Z}}_m$, that cause its degeneracy at DC, by adding the contribution of DC eigenvalues r as

$$\breve{\mathbf{Z}}_m(\mathrm{j}\omega_k) = \widehat{\mathbf{Y}}_m(\mathrm{j}\omega_k)^{-1} + \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & r\mathbb{I}_\ell \end{bmatrix}, \ \omega_k \in \Omega_G \cup \Omega_D. \tag{2.16}$$

3. conversion of the regularized data samples to the physical domain

$$\breve{\mathbf{Z}}(\mathrm{j}\omega_k) = \mathbf{Q}\breve{\mathbf{Z}}_m(\mathrm{j}\omega_k)\mathbf{Q}^\mathsf{T}, \ \omega_k \in \Omega_G \cup \Omega_D. \tag{2.17}$$

Notice that step 1 and 2 leads to admittance and impedance matrices that are not singular and can be inverted at any frequencies, as implied by the regularization process. Indeed, the modal impedance dataset resulting from

the first step has the low-frequency structure

$$\widehat{\mathbf{Z}}_m(\mathrm{j}\omega) = \widehat{\mathbf{Y}}_m(\mathrm{j}\omega)^{-1} = \begin{bmatrix} (\mathrm{j}\omega\widetilde{\mathbf{C}} + \frac{1}{R'}\mathbb{I}_c)^{-1} + \star_3 & \star_2 \\ \star_2 & \mathrm{j}\omega\widetilde{\mathbf{\Gamma}}^{-1} + \star_4 \end{bmatrix}, \ \omega \in \Omega_G, \quad (2.18)$$

where matrix blocks with leading terms $(\mathrm{j}\omega)^\nu$ for $\omega \to 0$ are indicated with the generic symbol $\star_\nu$. It is easy to see that the singularity of $\widehat{\mathbf{Z}}_m(\mathrm{j}\omega)$ is provided only by the leading inductive terms in the bottom-right block of (2.18). Thus, step 2 aims at regularized only this term producing as result the modal structure

$$\breve{\mathbf{Z}}_m(\mathrm{j}\omega) = \begin{bmatrix} R'\mathbb{I}_c + \star_1 & \star_2 \\ \star_2 & r\mathbb{I}_\ell + \star_1 \end{bmatrix}, \ \omega \in \Omega_G, \quad (2.19)$$

which is consistent with the DC eigenvalues $\mathbf{\Lambda}_0^{-1}$ derived in (2.6). This further demonstrates how the above procedure leads to a compatible result with both the regularization DC circuit of Fig. 2.2 and to the direct approach presented in Section 2.1.1 (Algorithm 1), providing as additional benefit a reliable extrapolation procedure.

The above derivation assumed an initial modal admittance matrix structure (2.10), which in fact is only an approximation consisting of the sole asymptotic leading terms. However, even assuming each matrix block to be expressed as a full Taylor polynomial expansion at $s = \mathrm{j}\omega = 0$ (including only odd powers of frequency due to the assumed lossless nature of the enclosure) the same result of (2.19) holds true, although the leading powers in (2.18) may be different. Details are straightforward and here omitted.

Algorithm 2 summarizes the full preprocessing data procedure presented in this Chapter, which includes regularization and extrapolation steps in the modal domain. Figure 2.7 reports the impedance parameters after the presented data conditioning for a 25-port shielding enclosure. The final dataset includes $K = 8904$ frequency samples in the interval $[0, 1]$ GHz, considering both the DC and $K_{GL} = 903$ extrapolated points in the low-frequency range. This result shows that the proposed strategy does not affect in-band responses, which are perfectly overlapped by the original solver data also in the transition band Ω_L , and smoothly fill the gap bandwidth Ω_G up to DC.

Algorithm 2 Regularization and extrapolation in the modal domain

Require: Frequency data $\{\mathbf{H}(\mathrm{j}\omega_k),\, \omega_k \in \Omega_D\}$, R, r, K_{GL}

1: $\mathbf{H}(\mathrm{j}\omega_k) \leftarrow 0.5(\mathbf{H}(\mathrm{j}\omega_k) + \mathbf{H}(\mathrm{j}\omega_k)^{\mathsf{T}})$
2: Convert $\mathbf{H}(\mathrm{j}\omega_k)$ to admittance $\mathbf{Y}(\mathrm{j}\omega_k)$
3: Define a regularization DC circuit (e.g. Fig. 2.3)
4: Compute $\mathbf{Y}_0 = \mathbf{Y}(\mathrm{j}\omega = 0)$
5: Find modal matrix $\mathbf{Q}$ from the eigenvalue decomposition of $\mathbf{Y}_0 = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{\mathsf{T}}$
6: Convert to the modal domain $\mathbf{Y}_m(\mathrm{j}\omega_k) = \mathbf{Q}^{\mathsf{T}}\mathbf{Y}(\mathrm{j}\omega_k)\mathbf{Q}\,,\ \omega_k \in \Omega_D$
7: Assign coefficients $\mu_{ij} = \pm 1$ via regression (2.11)
8: Compute capacitive/inductive coefficients of the low-frequency model $\mathbf{\Upsilon}_m(\mathrm{j}\omega)$ via least-squares (2.12)
9: Evaluate K_{GL} samples of $\mathbf{\Upsilon}_m(\mathrm{j}\omega_k)$ with $\omega_k \in \Omega_G \cup \Omega_L$
10: Define α_k and β_k, such that $\alpha_k + \beta_k = 1\,,\ \forall \omega_k \in \Omega_G \cup \Omega_L$
11: Merge data as in (2.14)
12: Regularize admittance $\widehat{\mathbf{Y}}_m(\mathrm{j}\omega_k) \leftarrow \mathbf{Y}_m(\mathrm{j}\omega_k) + \begin{bmatrix} \frac{1}{R'}\mathbb{I}_c & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix},\ \omega_k \in \Omega_G \cup \Omega_D$
13: Regularize impedance as $\breve{\mathbf{Z}}_m(\mathrm{j}\omega_k) \leftarrow [\widehat{\mathbf{Y}}_m(\mathrm{j}\omega_k)]^{-1} + \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & r\mathbb{I}_\ell \end{bmatrix},\ \omega_k \in \Omega_G \cup \Omega_D$
14: Convert to physical domain $\breve{\mathbf{Z}}(\mathrm{j}\omega_k) = \mathbf{Q}\breve{\mathbf{Z}}_m(\mathrm{j}\omega_k)\mathbf{Q}^{\mathsf{T}},\ \omega_k \in \Omega_G \cup \Omega_D$
15: **return** $\breve{\mathbf{Z}}(\mathrm{j}\omega)$: full-band regularized and extrapolated data

Summary of key results

The main outcomes of the proposed regularization and extrapolation procedure are

1. a well-define DC point, which enables to avoid matrix singularities at all frequencies through the *regularization step*;
2. a full-band dataset, well-defined in $[0, f_{\max}]$, obtained by filling the *gap* region Ω_G with synthetic samples through the*modal extrapolation step*;
3. a (structured) modal dataset that retains DC contribution as diagonal entries of the transfer matrix.

It is sufficient to use the modal responses as training dataset to build a model that preserve a high level of accuracy in the entire (physical) low-frequency range. Notice that the modal impedance matrix in the low-frequency band Ω_G shows the rich structure of (2.19), with diagonal entries that converge to a constant valued, and off diagonal elements with first and second order zeros. The same structure is preserved in all system representations, including admittance

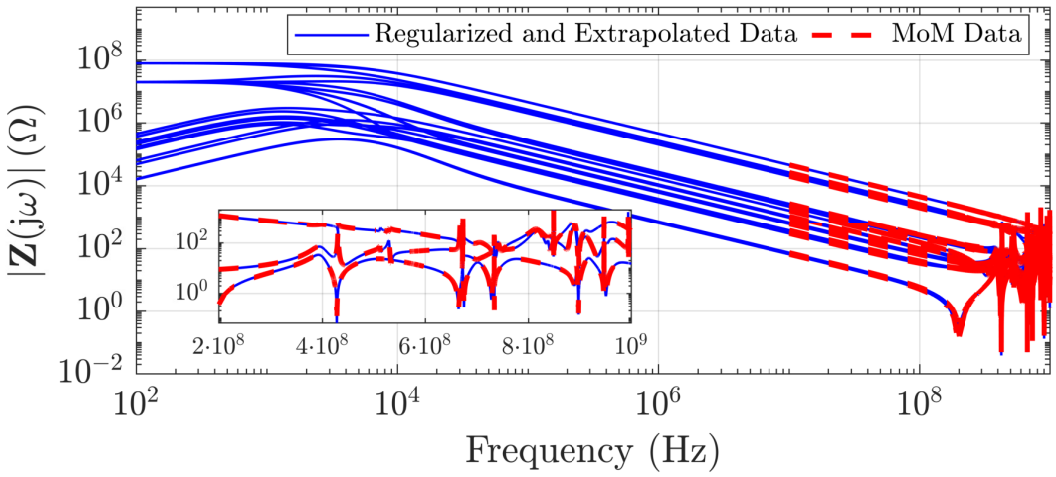

Fig. 2.7 Proposed regularization and extrapolation procedure with respect to field solver (MoM) data for a 25-port shielding enclosure ($R = 100\text{M}\Omega$, $r = 0.1\Omega$). Only a representative set of impedance responses are reported. From [30] © 2022 IEEE.

and scattering parameters. Thus, these elements must be properly considered when approximating the system behavior: they will play an important role in the proposed data and model compressed representations, to be presented next.

2.1.5 Extensions

In order to apply the proposed strategy to a general multiport system with a possibly different asymptotic behavior at DC, the only step required is to define a proper DC regularization circuit as the one provided in Fig. 2.3. In fact, the presented procedure aims at solving the most generic critical situation with both admittance *and* impedance matrices singular at DC. The case when only inductive *or* capacitive modes are responsible for an ill-defined $\mathbf{Y}_0$ *or* $\mathbf{Z}_0$ can be seen as a particular case of the above strategy. In particular, only series *or* shunt resistances are necessary to regularize the corresponding DC matrix (line 3 of Algorithm 2).

The proposed procedure is valid if the losses in the low-frequency range are sufficiently small, such that the asymptotic approximation (2.8) holds. Nevertheless, the extrapolation procedure can be extended to lossy systems either by including resistive components into (2.8) or by generating a synthetic lossless dataset by subtracting the DC matrix contribution at all frequencies

and neglecting the real part of the resulting port admittance or impedance matrix.

The proposed regularization step has been designed to deal with a lack of system characterization at DC. As already mentioned, the artificial DC circuit is aimed at providing a well-defined response when the frequency vanishes, to be used for subsequence macromodeling and simulation purposes. If a DC point is already available from an independent system characterization, either from a field solver or from direct measurements, this can be used in the proposed procedure by substituting the synthetic DC matrices (2.1)-(2.3). Only lines 3 and 4 of Algorithm 2 will be affected. The rest of the procedure applies with minor changes, necessary to adjust the modal regularization with new system eigenvalues.

2.1.6 Numerical results

This section provides numerical results that demonstrate the effectiveness of the proposed regularization and extrapolation procedure.

Changing regularization parameters

The 25-port shielding enclosure used as running examples through this section is here exploited to show the effect of the regularization parameters. Figure 2.7 shows the impedance responses resulting from a full regularization and extrapolation procedure performed in the modal domain in the range $[0, 1]$ GHz. This result is obtained by setting shunt and series resistances to $R = 100\text{M}\Omega$ and $r = 0.1\Omega$, respectively. The perturbation on the original MoM data of a direct regularization step (Algorithm 1) was tested with different values of $R \in [10^6, 10^9]\,\Omega$ and $r \in [10^{-3}, 1]\,\Omega$. A set of comprehensive admittance responses with three combinations of regularization parameters is reported in Fig. 2.8. For the selected results, the shunt resistance is fixed to $R = 1\text{M}\Omega$ while the series resistance is systematically reduced (from top to bottom panel) choosing $r = \{1, 0.1, 0.01\}\Omega$. This picture demonstrates that

- a significant perturbation in the original data is needed to obtain a DC point that smoothly merges with the low-frequency responses in Ω_L; this

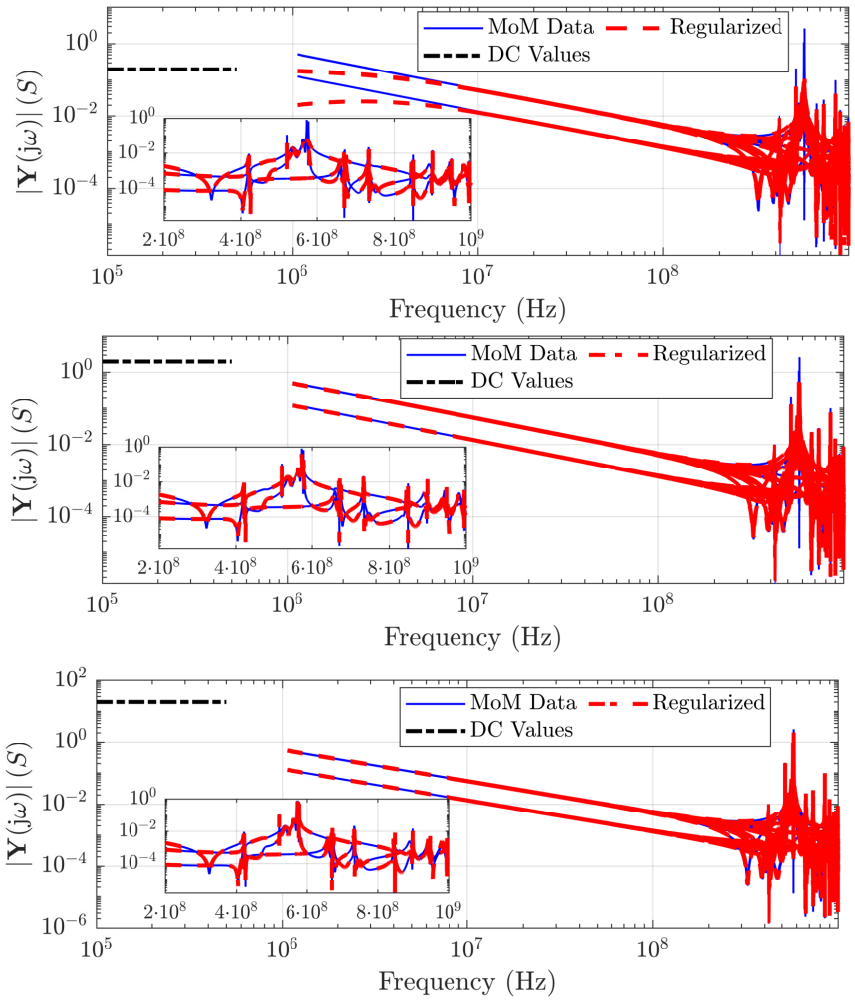

Fig. 2.8 Representative admittance responses of a 25-port shielding enclosures changing regularization parameters of Algorithm 1. Top: $R = 1\text{M}\Omega$ and $r = 1\Omega$; Middle: $R = 1\text{M}\Omega$ and $r = 0.1\Omega$; Bottom: $R = 1\text{M}\Omega$ and $r = 0.01\Omega$.

is the effect of an extreme choice of shunt and series resistances, which in this case are $R = 1\text{M}\Omega$ and $r = 1\Omega$, respectively;

- reducing r the discrepancies between original and regularized samples are attenuated, especially in the low-frequency range Ω_L, at the price of a larger DC level.

Notice that a similar reasoning can be verified observing the effect of different shunt resistances values R on the impedance responses. In this case (not shown), a larger value of R reduces the perturbation on the regularized data. The tread off between perturbation on the field solver (MoM) data and resulting

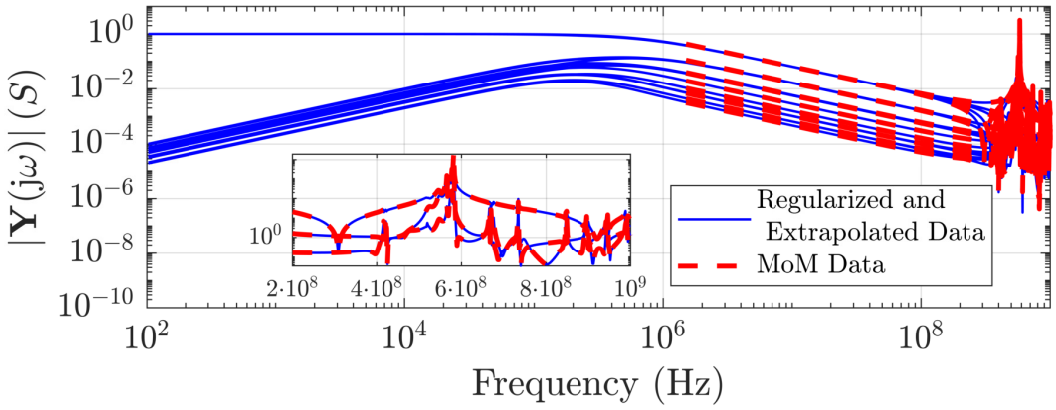

Fig. 2.9 As in Fig. 2.7 for a shielding enclosure with 100-ports

DC value is a design choice left to the user, whenever a DC characterization is not already available from a different source. A general guideline is to keep the DC value as close as possible to the system response at the first available frequency, at the price of a modest variation of the original data. To this end, a recommendation is to preserve all data features by selecting very small and large values of r and R, respectively.

From 100 to 400 ports

In this section, the shielding enclosure grid is modified by increasing the density of the $p \times p = P$ lumped ports and the proposed data conditioning approach is applied to the resulting LTI systems. This section provides results of Algorithm 2 ($R = 100\text{M}\Omega$,$r = 0.1\Omega$) on enclosures with $p = \{10, 20\}$ and $P = \{100, 400\}$ in the range $[0, 1]$ GHz. Considering that the runtime is of course case-dependent, as well as influenced on the actual setup for running the algorithms, Table 2.1 provides a summary of the selected parameters and of the overall elapsed time. In particular, Figure 2.9 compares a set of representative admittance responses of a 100-port box after regularization with the original MoM data, showing only negligible variations at the higher frequencies. Figure 2.10 demonstrates the effectiveness of the proposed data conditioning on a 400-port shielding enclosure: the full-band impedance parameters show a perfect agreement with the original data. The resulting large-scale dataset is now suitable for successive macromodeling steps, as will be discussed in Chapter 3.

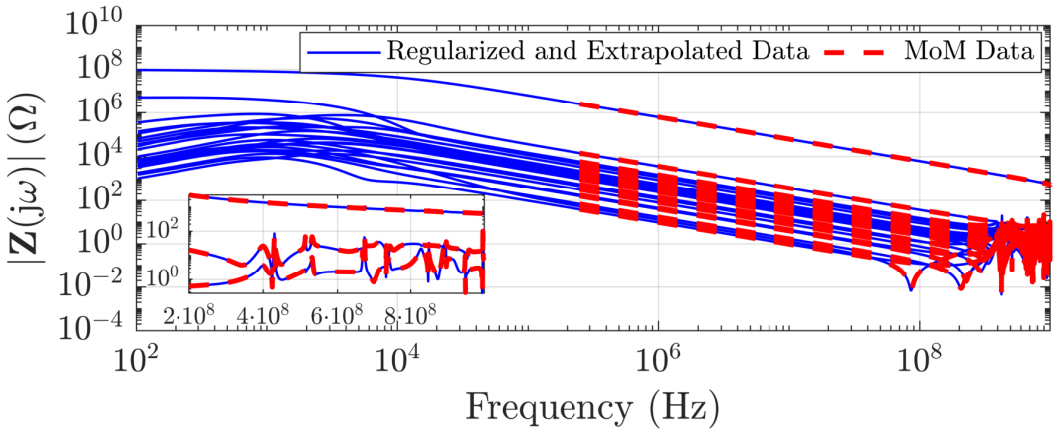

Fig. 2.10 As in Fig. 2.7 for a shielding enclosure with 400-ports.

Table 2.1 Data regularization and extrapolation results. Table summarizes frequency data information in terms of ports (P) and samples (K), the selected parameters of Algorithm 2, and the total elapsed time.

P	K	R (MΩ)	r (Ω)	K_{GL}	$Time(s)$
25	8001	100	0.1	366	2.78
100	501	100	0.1	349	3.78
400	1001	100	0.1	296	101.1

2.2 Structure-preserving data compression

In this section, we address a second major problem associated with large-scale data for macromodeling applications, namely the scalability with number of interface ports. Indeed, for the considered shielding enclosure application, when the density of the aperture grid is augmented, the corresponding number of system ports P drastically increases, leading to a complexity that may become impractical to be handled, both from macromodeling and transient simulations standpoint. The procedure originally presented in [30] is reported in the following.

Considering the model generation, even if the literature offers several ways to address this problem, there is still room for improvement. For instance, the well-know VF algorithm has been already optimized in its Fast VF implementation [24] to reduce the fitting cost by compressing the least-squares solution of the coefficients via QR factorization. Nevertheless, this approach still does

not scale favorably with the number of ports, as it will be detailed in the next section.

An approach to improve the efficiency of any identification algorithm is to reduce the size of frequency tabulated data by exploiting the natural redundancy in the system responses, and consequentially reducing the computational cost for macromodel generation. This strategy enables to optimize the model structure by avoiding the need to identify $P \times P$ responses with a rational approximation that requires P^2 sets of coefficients, one for each transfer matrix element.

A data compression technique coupled with a macromodeling identification step has been already presented in [34]. This strategy takes advantage of a standard truncated Singular Value Decomposition (SVD) to reduce the spatial redundancy in the system port responses, and builds a Compressed Macromodel (CM) via VF applied to a reduced set of basis functions. In this section, the data compression proposed in [34] will be extended to deal with the modal data structure of shielding enclosures and to preserve low-frequency features, such as zeros at DC.

The data reduction strategy of [34] is first summarized in Section 2.2.1. Two proposed improved algorithms, namely the Block-Diagonal SVD (BD-SVD) and the Hierarchical SVD (Hi-SVD), are then introduced in Section 2.2.2. Numerical results on enclosures with up to 400 ports are reported in Section 2.2.3.

2.2.1 SVD data compression: background

The data compression strategy presented in [34] requires a set of tabulated frequency responses of a P-port LTI system, defined as $\{\mathbf{H}(\mathrm{j}\omega_k) = \mathbf{H}_k,\, k = 1, \cdots, K\}$. The procedure starts stacking columns of $\mathbf{H}_k \in \mathbb{C}^{P\times P}$ in a row vector $\boldsymbol{x}_k \in \mathbb{C}^{P^2}$ defined as

$$\boldsymbol{x}_k = \mathrm{vec}(\mathbf{H}_k)^T\,. \tag{2.20}$$

The reverse mapping $(\boldsymbol{x}_k)_\ell = (\mathbf{H}_k)_{i,j}$ can be obtained defining suitable indexes as

$$\ell = i + (j-1)P, \qquad \begin{cases} i = 1 + \mathrm{mod}(\ell - 1, P) \\ j = \lceil \ell / P \rceil \end{cases} \tag{2.21}$$

where $\lceil\cdot\rceil$ rounds towards infinity and mod is the integer division remainder. With a formal terminology, defining the set of frequency responses as the tensor $\mathcal{H} \in \mathbb{C}^{K\times P\times P}$, the *unfolding* or (mode-1) *matricization* of tensor $\mathcal{H}$ [49] can be defined as

$$\mathbf{X} = \begin{bmatrix} \longleftarrow & \boldsymbol{x}_1 & \longrightarrow \\ \vdots & \vdots & \vdots \\ \longleftarrow & \boldsymbol{x}_K & \longrightarrow \end{bmatrix} = \begin{bmatrix} \uparrow & \cdots & \uparrow \\ \boldsymbol{m}_1 & \cdots & \boldsymbol{m}_{P^2} \\ \downarrow & \cdots & \downarrow \end{bmatrix} \tag{2.22}$$

where $\mathbf{X} \in \mathbb{C}^{P^2}$ collects rows vectors (2.20) in a matrix form.

A truncated SVD decomposition is applied to the real matrix obtained stacking real and imaginary part of $\mathbf{X}$, as

$$\begin{bmatrix} \mathrm{Re}\,\{\mathbf{X}\} \\ \mathrm{Im}\,\{\mathbf{X}\} \end{bmatrix} \approx \bar{\mathbf{U}}\bar{\mathbf{\Sigma}}\bar{\mathbf{V}}^{\mathsf{T}} \tag{2.23}$$

where $\bar{\mathbf{\Sigma}} \in \mathbb{R}^{\rho\times\rho}$ is a diagonal matrix that collects the ρ leading singular values, $\bar{\mathbf{U}} \in \mathbb{R}^{2K\times\rho}$ and $\bar{\mathbf{V}} \in \mathbb{R}^{P^2\times\rho}$ are the left and right singular vector matrices. By definition, $\bar{\mathbf{U}}^{\mathsf{T}}\bar{\mathbf{U}} = \mathbb{I}_\rho$ and $\bar{\mathbf{V}}^{\mathsf{T}}\bar{\mathbf{V}} = \mathbb{I}_\rho$. Notice that the number of retained singular values ρ is chosen such that $\rho \ll \min\{2K, P^2\}$, ensuring that (2.23) is a low-rank approximation and consequentially $\bar{\mathbf{V}}\bar{\mathbf{V}}^{\mathsf{T}} \neq \mathbb{I}_{P^2}$. Next, the real and imaginary components of the scaled (left) singular vectors are recombined as

$$\bar{\mathbf{W}} = \begin{bmatrix} \mathbb{I}_K & \mathrm{j}\mathbb{I}_K \end{bmatrix} \bar{\mathbf{U}}\bar{\mathbf{\Sigma}} \tag{2.24}$$

in order to approximate the matricization of system responses as

$$\mathbf{X} \approx \bar{\mathbf{X}} = \bar{\mathbf{W}}\bar{\mathbf{V}}^{\mathsf{T}}\,. \tag{2.25}$$

On the one hand, splitting real and imaginary elements in (2.23) enables to preserve (by construction) the causality [47] of $\bar{\mathbf{X}}$, since relationships between real/imaginary components are preserved by (2.24) and (2.23). On the other hand, the ℓ-th entry of the $P \times P$ original transfer matrix can be estimated as

$$\boldsymbol{m}_\ell \approx \sum_{q=1}^{\rho} v_{\ell,q}\bar{\boldsymbol{w}}_q\,. \tag{2.26}$$

which corresponds to the ℓ-th column of $\mathbf{X}$, while $\bar{\boldsymbol{w}}_q$ indicates the q-th column of $\bar{\mathbf{W}}$. The set $\{\bar{\boldsymbol{w}}_q, q = 1, \cdots, \rho\}$ collecting these vectors is regarded as a set of *basis functions*. In short, the original P^2 responses can be effectively approximated as a linear combination of a reduced set of ρ frequency-dependent basis functions, as detailed in both (2.25) and (2.26).

A rational approximation of all $\bar{\boldsymbol{w}}_q$ enables to build a surrogate model to reconstruct the $P \times P$ system behavior with a major saving in terms of computational costs, whenever $\rho \ll P^2$. Furthermore, a bound on the approximation error (2.25) is available [34] as

$$\mathcal{E} = \left\| \mathbf{X} - \bar{\mathbf{X}} \right\|_2 \leq \sqrt{2}\sigma_{\rho+1} \,. \tag{2.27}$$

with $\sigma_{\rho+1}$ first neglected singular value. Note that the truncation error (2.27) cannot be recovered by next steps of the modeling process: the behavioral model build on top of $\bar{\mathbf{X}}$ will show an error with respect to the original $\mathbf{X}$ at least of $\sqrt{2}\sigma_{\rho+1}$. Therefore, the error introduced by the data compression approximation will set a bound for the accuracy of all subsequent macromodeling steps.

The above compression strategy does not preserve the responses structure of a shielding enclosure data, especially in the modal domain. Indeed, the truncated SVD approximation (2.26) does not maintain the presence of zeros at DC even in the case of a small threshold $\sigma_{\rho+1}$, as illustrated in Fig. 2.11 for a $P = 25$ port enclosure. Before compression, the reported data have been extrapolated and regularized by means of the procedure proposed in Section 2.1.3, following Algorithm 2 ($r = 0.1\Omega$, $R = 100\text{M}\Omega$ and $K_{GL} = 903$). The final dataset provides system responses in the frequency range $[0, 1]$ GHz with $K = 8904$ samples. A total of $\rho = 73$ basis functions have been obtained imposing a threshold $\sigma < 10^{-6}$ to the SVD truncation (2.23). The bottom panel of Fig. 2.11 shows how all basis functions saturates to a different non-vanishing level at low-frequencies: the presence of a DC zero is not embedded in the basis structure and typical low-frequency decays of -20dB/dec (and -40dB/dec) are retained only up to a given point by the linear combination of nonzero basis (2.26). Hence, after an SVD the DC zero behavior might only be recovered by numerical cancellation, which after truncation is prevented by the selected threshold $\sigma_{\rho+1}$. This trend must be interpreted as an additional loss introduced by the data compression, which will lead to a non-physical modeling, and to

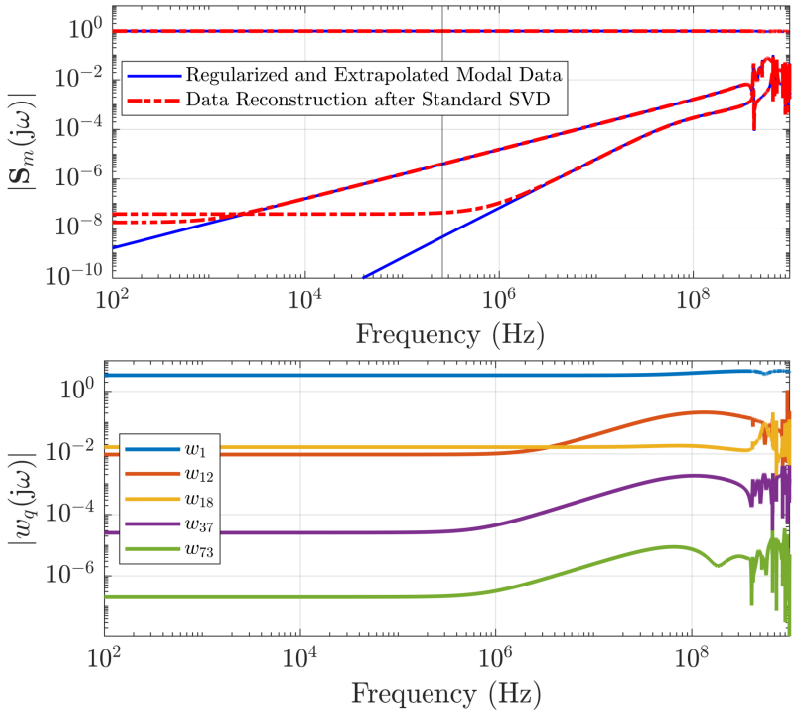

Fig. 2.11 Representative responses of a 25-port shielding enclosure after data compression via basic SVD [34]. Top: DC zeros of different orders are not preserved by the low-rank approximation (2.32). Bottom: resulting basis functions ($\rho = 73$ with threshold $\sigma_{\rho+1} = 10^{-6}$) do not retain the low-frequency behavior of the enclosure and are non-vanishing at DC. From [30] © 2022 IEEE.

a consequent unreliable transient simulation. Thus, the DC zeros behavior must be retained through the all modeling procedure. This includes the data compression step, and motivates the extension of the standard SVD reduction of [34] to more sophisticated compression approaches that are able to preserve a suitable data structure.

2.2.2 Structured SVD compression

The proposed data compression approach is based on the idea of using a specific set of basis functions that inherits the same structural properties of the original system, in order to preserve these features while reducing the data complexity. The aim of following procedures is to preserve the presence of DC zeros in the system responses.

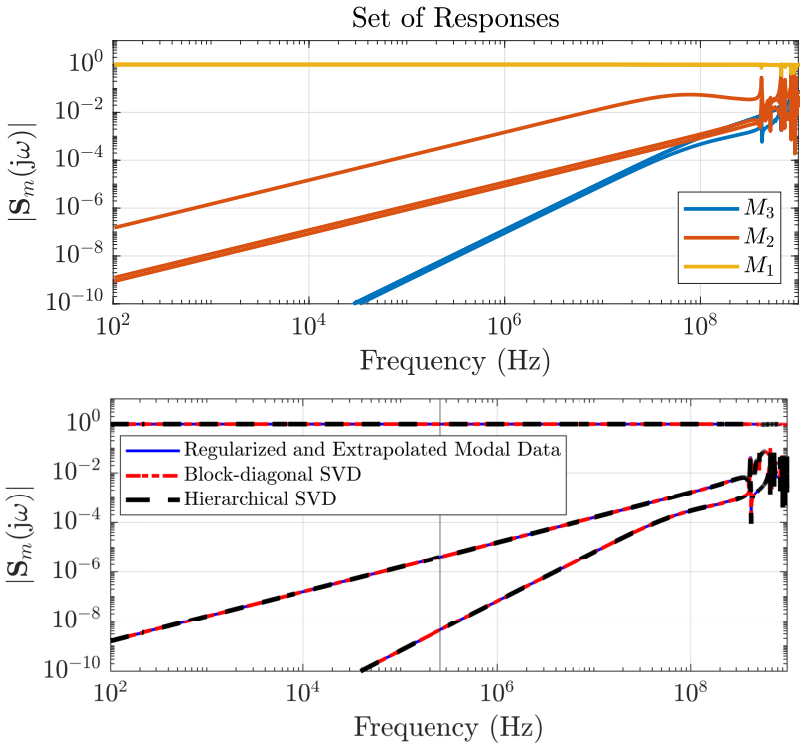

Fig. 2.12 Preserving low-frequency features while compressing the data of a 25-port box. Top: modal responses are clustered according to the order of DC zeros. Bottom: all set of responses are suitably reconstructed by presented compression strategies. From [30] © 2022 IEEE.

Let us start with the system responses matricization as in (2.22), which stacks the $P \times P$ transfer matrix entries following the column-based ordering defined by (2.20). First, the columns of $\mathbf{X}$ are rearranged by defining a permutation matrix P such that

$$\mathbf{X} = \mathbf{MP} = \begin{bmatrix} \mathbf{M}_1 & \mathbf{M}_2 & \mathbf{M}_3 \end{bmatrix} \mathbf{P}\,, \quad \mathbf{PP}^{\mathsf{T}} = \mathbb{I} \tag{2.28}$$

to identify group of responses $\mathbf{M}_\nu$, with P_ν number of columns for each set such that $\sum_\nu P_\nu = P^2$. For the specific case of (modal) shielding enclosures responses, $\nu = \{1, 2, 3\}$ and $\mathbf{M}_1 \in \mathbb{C}^{K \times P_1}$ stores the entries of $\mathbf{H}(s)$ that saturate to a nonzero DC value, while the other two blocks $\mathbf{M}_{2,3}$ gather all entries with DC zero of order 1 or 2, respectively. Figure 2.12 (and later Fig. 2.14d) shows an example of the three groups of responses on a 25-port box, automatically identified by means of a procedure similar to the one of Section 2.1.3.

The basic SVD reduction of Section 2.2.1 on the reordered matrix $\mathbf{M}$ leads to a reconstruction of each group $\mathbf{M}_\nu$ of responses that relies on *all* basis functions. To see this, the (full) decomposition (2.26) can be combined with (2.28) to obtain

$$\begin{aligned}\mathbf{M} = \begin{bmatrix}\mathbf{M}_1 & \mathbf{M}_2 & \mathbf{M}_3\end{bmatrix} = \begin{bmatrix}\mathbf{W}_1 & \mathbf{W}_2 & \mathbf{W}_3\end{bmatrix}\mathbf{V}^\mathsf{T} = \\ = \begin{bmatrix}\mathbf{W}_1 & \mathbf{W}_2 & \mathbf{W}_3\end{bmatrix}\begin{bmatrix}\mathbf{V}_{11} & \mathbf{V}_{12} & \mathbf{V}_{13}\\ \mathbf{V}_{21} & \mathbf{V}_{22} & \mathbf{V}_{23}\\ \mathbf{V}_{31} & \mathbf{V}_{32} & \mathbf{V}_{33}\end{bmatrix}.\end{aligned} \tag{2.29}$$

where the 3×3 block expansion enables to see how each group of basis functions $\mathbf{W}_\nu$ is associated to sets $\mathbf{M}_\nu$, with $\nu = \{1,2,3\}$. It is clear from (2.29) that

$$\begin{aligned}\mathbf{M}_1 &= \mathbf{W}_1\mathbf{V}_{11} + \mathbf{W}_2\mathbf{V}_{21} + \mathbf{W}_3\mathbf{V}_{31}\\ \mathbf{M}_2 &= \mathbf{W}_1\mathbf{V}_{12} + \mathbf{W}_2\mathbf{V}_{22} + \mathbf{W}_3\mathbf{V}_{32}\\ \mathbf{M}_3 &= \mathbf{W}_1\mathbf{V}_{13} + \mathbf{W}_2\mathbf{V}_{23} + \mathbf{W}_3\mathbf{V}_{33}\end{aligned} \tag{2.30}$$

which confirms that each $\mathbf{M}_\nu$ is affected by all sets $\mathbf{W}_{1,2,3}$ through independent blocks of the transformation matrix $\mathbf{V}^\mathsf{T}$.

The objective now is to create a direct relationship between the structure of one data responses group and one basis functions set. To this end, it is convenient to start assuming that the last subset $\mathbf{W}_3$ is built using only responses from group $\mathbf{M}_3$, which collect double order zeros at DC: thus $\mathbf{W}_3$ will inherit the same low-frequency behavior. To preserve this property and reconstruct $\mathbf{M}_3$ in (2.30), both blocks $\mathbf{V}_{13}$ and $\mathbf{V}_{23}$ must vanish. Similarly, the requirement of having single DC zeros behavior on basis set $\mathbf{W}_2$ is achieved by setting $\mathbf{V}_{12} = 0$. This means that imposing a lower-triangular structure to $\mathbf{V}^\mathsf{T}$ preserve DC zeros while reducing the data complexity in terms of number of responses. The above condition requires building each basis subset $\mathbf{W}_\nu$ from a restricted set of responses. Two different strategies can be considered

- *Block-Diagonal SVD (BD-SVD):* each subset $\mathbf{M}_\nu$ is treated independently, giving a block-diagonal $\mathbf{V}^\mathsf{T}$;
- *Hierarchical SVD (Hi-SVD):* each subset $\mathbf{M}_\nu$ is processed iteratively with ordering $\nu = 3,2,1$; at each step, the projection of the current responses

onto the basis functions generated at previous steps is removed, and the resulting deflated responses are processed by SVD; this process results in a block-triangular $\mathbf{V}^\mathsf{T}$.

Both strategies were implemented and tested. Although under the pure data processing standpoint these two approaches are equivalent, with the hierarchical approach potentially leading to a smaller total number of basis functions, it turns out that the block-diagonal approach leads to better results in terms of overall accuracy after macromodel generation (see Sec. 3.2) and is preferred. Both approaches are summarized below for completeness.

Block-Diagonal SVD

The first strategy aims at obtaining a block-diagonal $\mathbf{V}^\mathsf{T}$, by processing each subset of responses $\mathbf{M}_\nu$ for $\nu = \{1, 2, 3\}$ individually and by means of disjoint SVDs. This implies computing the low-rank approximation

$$\begin{bmatrix} \mathrm{Re}\,\{\mathbf{M}_\nu\} \\ \mathrm{Im}\,\{\mathbf{M}_\nu\} \end{bmatrix} \approx \bar{\mathbf{U}}_\nu \bar{\mathbf{\Sigma}}_\nu \bar{\mathbf{V}}_\nu^\mathsf{T}, \qquad \nu = 1, 2, 3 \tag{2.31}$$

such that $\bar{\mathbf{\Sigma}}_\nu$ gathers the leading ρ_ν singular values, and $\bar{\mathbf{U}}_\nu \in \mathbb{R}^{2K \times \rho_\nu}$, $\bar{\mathbf{V}}_\nu \in \mathbb{R}^{P_\nu \times \rho_\nu}$, with $\bar{\mathbf{U}}_\nu^\mathsf{T} \bar{\mathbf{U}}_\nu = \mathbb{I}$ and $\bar{\mathbf{V}}_\nu^\mathsf{T} \bar{\mathbf{V}}_\nu = \mathbb{I}$. The individual set of basis functions can now be defined as

$$\bar{\mathbf{W}}_\nu = \begin{bmatrix} \mathbb{I}_K & \mathrm{j}\mathbb{I}_K \end{bmatrix} \bar{\mathbf{U}}_\nu \bar{\mathbf{\Sigma}}_\nu, \qquad \bar{\mathbf{V}}'_\nu = \bar{\mathbf{V}}_\nu^\mathsf{T}, \qquad \nu = 1, 2, 3 \tag{2.32}$$

that allows to approximate (2.31) as

$$\mathbf{M}_\nu \approx \bar{\mathbf{M}}_\nu = \bar{\mathbf{W}}_\nu \bar{\mathbf{V}}'_\nu, \qquad \nu = 1, 2, 3. \tag{2.33}$$

The overall Block-Diagonal SVD (BD-SVD) approximation reads as

$$\mathbf{M} \approx \bar{\mathbf{M}} = \begin{bmatrix} \bar{\mathbf{M}}_1 & \bar{\mathbf{M}}_2 & \bar{\mathbf{M}}_3 \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{W}}_1 & \bar{\mathbf{W}}_2 & \bar{\mathbf{W}}_3 \end{bmatrix} \begin{bmatrix} \bar{\mathbf{V}}'_1 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \bar{\mathbf{V}}'_2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \bar{\mathbf{V}}'_3 \end{bmatrix}. \tag{2.34}$$

The major advantages of the BD-SVD are

- a major flexibility in the accuracy control for individual responses; this enables to specifically tune the accuracy at DC when the data is expressed in the modal domain, since the diagonal matrix entries corresponds to the system eigenvalues when $f = 0$.
- a sparse transformation $\bar{\mathbf{V}}$, which is reflected in the resulting macromodel coefficients.

On the other hand, a compression strategy applied to the entire dataset may result in a smaller number of overall basis functions ρ with respect to the proposed BD-SVD, for which $\rho = \sum_\nu \rho_\nu$ and no correlations hence redundancies between different subsets is exploited.

Hierarchical SVD

The main difference of the Hi-SVD strategy with respect to the BD-SVD is that subsets of responses are not approximated independently. Indeed, processing each set in a specific order $\nu = \{3, 2, 1\}$ and removing the contribution of already processed responses by projection, guarantees that data features are equivalently retained in the resulting basis functions. The procedure starts with the group $\mathbf{M}_3$ that collects responses with double DC zeros. Following the same procedure reported in Sec. 2.2.2, the SVD approximation is computed (2.31) for $\nu = 3$ as

$$\widetilde{\mathbf{M}}_3 \approx \bar{\mathbf{U}}_3 \bar{\mathbf{V}}'_3 \tag{2.35}$$

where $\bar{\mathbf{V}}'_3 = \bar{\mathbf{\Sigma}}_3 \bar{\mathbf{V}}_3^\mathsf{T}$. The real-valued matrix $\bar{\mathbf{V}}'_3$ embeds also the scaling factors induced by the leading ρ_3 singular values so that the columns of $\bar{\mathbf{U}}_3$ are orthonormal.

The next step is to remove from $\widetilde{\mathbf{M}}_2$ its projection onto the column space of the orthogonal (and real) matrix $\bar{\mathbf{U}}_3$. The result is approximated through a second truncated SVD via (2.31)-(2.33), obtaining

$$\widetilde{\mathbf{M}}_2 - \bar{\mathbf{U}}_3(\bar{\mathbf{U}}_3^\mathsf{T}\widetilde{\mathbf{M}}_2) \approx \bar{\mathbf{U}}_2 \bar{\mathbf{\Sigma}}_2 \bar{\mathbf{V}}_2^\mathsf{T} = \bar{\mathbf{U}}_2 \bar{\mathbf{V}}'_2. \tag{2.36}$$

With a simple manipulation this expression can be rewritten as

$$\widetilde{\mathbf{M}}_2 \approx \bar{\mathbf{U}}_2 \bar{\mathbf{V}}'_2 + \bar{\mathbf{U}}_3(\bar{\mathbf{U}}_3^\mathsf{T}\widetilde{\mathbf{M}}_2) = \bar{\mathbf{U}}_2 \bar{\mathbf{V}}'_2 + \bar{\mathbf{U}}_3 \bar{\mathbf{V}}'_{3,2} \tag{2.37}$$

where $\bar{\mathbf{V}}'_{3,2} = \bar{\mathbf{U}}_3^{\mathsf{T}}\widetilde{\mathbf{M}}_2 \in \mathbb{R}^{\rho_3 \times P_2}$. Finally, the same procedure is replicated for the last group, obtaining

$$\widetilde{\mathbf{M}}_1 - \bar{\mathbf{U}}_2\bar{\mathbf{V}}'_{2,1} - \bar{\mathbf{U}}_3\bar{\mathbf{V}}'_{3,1} \approx \bar{\mathbf{U}}_1\bar{\mathbf{V}}'_1 \tag{2.38}$$

with $\bar{\mathbf{V}}'_{2,1} = \bar{\mathbf{U}}_2^{\mathsf{T}}\widetilde{\mathbf{M}}_1 \in \mathbb{R}^{\rho_2 \times P_1}$ and $\bar{\mathbf{V}}'_{3,1} = \bar{\mathbf{U}}_3^{\mathsf{T}}\widetilde{\mathbf{M}}_1 \in \mathbb{R}^{\rho_3 \times P_1}$. All steps from (2.35)-(2.38) can be defined in an iterative fashion [30] for $\ell = \{3, 2, 1\}$ as

$$\widetilde{\mathbf{M}}_\ell - \sum_{j=\ell+1}^{3} \bar{\mathbf{U}}_j\bar{\mathbf{V}}'_{j,\ell} \approx \bar{\mathbf{U}}_\ell\bar{\mathbf{\Sigma}}_\ell\bar{\mathbf{V}}_\ell^{\mathsf{T}} \tag{2.39}$$

with $\bar{\mathbf{V}}'_\ell = \bar{\mathbf{\Sigma}}_\ell\bar{\mathbf{V}}_\ell^{\mathsf{T}}$ and $\bar{\mathbf{V}}'_{j,\ell} = \bar{\mathbf{U}}_j^{\mathsf{T}}\widetilde{\mathbf{M}}_\ell$. Collecting all approximations in a block-triangular structure reads

$$\widetilde{\mathbf{M}} = \begin{bmatrix}\widetilde{\mathbf{M}}_1 & \widetilde{\mathbf{M}}_2 & \widetilde{\mathbf{M}}_3\end{bmatrix} \approx \tag{2.40}$$

$$\approx \begin{bmatrix}\bar{\mathbf{U}}_1 & \bar{\mathbf{U}}_2 & \bar{\mathbf{U}}_3\end{bmatrix} \begin{bmatrix}\bar{\mathbf{V}}'_1 & \mathbf{0} & \mathbf{0} \\ \bar{\mathbf{V}}'_{2,1} & \bar{\mathbf{V}}'_2 & \mathbf{0} \\ \bar{\mathbf{V}}'_{3,1} & \bar{\mathbf{V}}'_{3,2} & \bar{\mathbf{V}}'_3\end{bmatrix}. \tag{2.41}$$

Finally, reconstructing complex-valued responses as in (2.32) leads to the Hierarchical SVD approximation

$$\mathbf{M} \approx \bar{\mathbf{M}} = \begin{bmatrix}\bar{\mathbf{M}}_1 & \bar{\mathbf{M}}_2 & \bar{\mathbf{M}}_3\end{bmatrix} = \tag{2.42}$$

$$= \begin{bmatrix}\bar{\mathbf{W}}_1 & \bar{\mathbf{W}}_2 & \bar{\mathbf{W}}_3\end{bmatrix} \begin{bmatrix}\bar{\mathbf{V}}'_1 & \mathbf{0} & \mathbf{0} \\ \bar{\mathbf{V}}'_{2,1} & \bar{\mathbf{V}}'_2 & \mathbf{0} \\ \bar{\mathbf{V}}'_{3,1} & \bar{\mathbf{V}}'_{3,2} & \bar{\mathbf{V}}'_3\end{bmatrix}. \tag{2.43}$$

With respect to the unstructured SVD [34] and to the BD-SVD, the Hi-SVD constructs set of unitary basis functions such that the euclidean norm of all q-th basis $\bar{\boldsymbol{w}}_{\nu,q}$ is always $\|\bar{\boldsymbol{w}}_{\nu,q}\|_2 = 1$ (and not $\|\bar{\boldsymbol{w}}_{\nu,q}\|_2 = \sigma_{\nu,q}$). Coherently, the rows in the superposition matrices $\bar{\mathbf{V}}'_\nu$ are no longer orthonormal due to the scaling in (2.39) by the singular value matrices $\bar{\mathbf{\Sigma}}_\nu$. This is not a problem since such condition is not required in the following.

The Hierarchical SVD strategy retains all benefits of the BD-SVD approach, while generally resulting in a smaller number of basis functions $\rho = \sum_\nu \rho_\nu$ at the same approximation level.

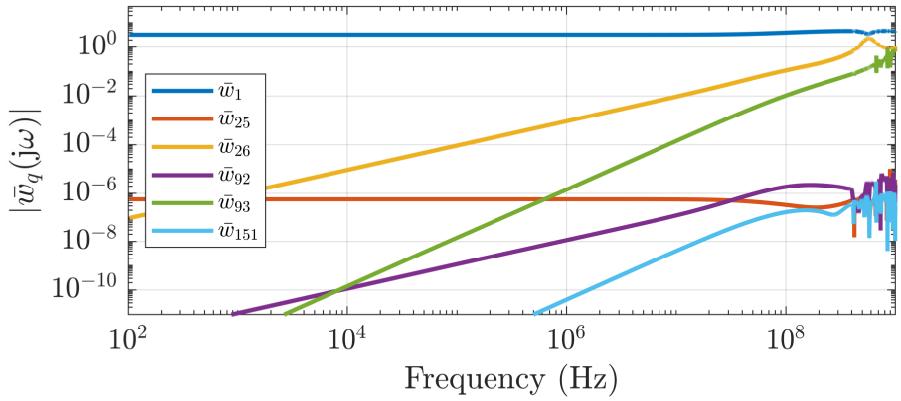

Fig. 2.13 Block-Diagonal SVD (BD-SVD) basis functions ($\rho = 151$) retain low-frequency data features, such as DC zeros. Only first and last basis of each $\bar{\mathbf{W}}_\nu$ with $\nu = \{1, 2, 3\}$ are reported. From [30] © 2022 IEEE.

2.2.3 Numerical results

This section reports numerical results that confirm the effectiveness of proposed compression strategies. All results have been obtained using a prototype MATLAB code on a Workstation based on Core i9-7900X CPU (3.3 GHz) with 64 GB of RAM.

A 25-port enclosure example

The same 25-port enclosure in Section 2.2.1 is here used to compare proposed compression strategies with the standard unstructured SVD approach, demonstrating their advantages. The BD-SVD and the Hi-SVD are applied to the full-band (modal) scattering samples by grouping system responses as in the top panel of 2.12, according to the DC zeros order. In both cases, the same thresholds $\{10^{-8}, 10^{-6}, 10^{-6}\}$ are used to approximate the corresponding group of responses $\mathbf{M}_1$,$\mathbf{M}_2$ and $\mathbf{M}_3$. Since the compression is applied in the modal domain, the truncation error bound for $\mathbf{M}_1$ is set to guarantee an accurate DC response while reconstructing the diagonal elements of the transfer matrix $\mathbf{S}_m(\mathrm{j}\omega)$, which are nonzero at low-frequencies. Figure 2.12 demonstrates that all data features are always retained by the proposed strategies, for all groups of responses. This is confirmed by the subset of basis functions obtained with the BD-SVD reported in Fig. 2.13, which inherits the order of DC zeros for each block. As expected, the overall number of basis functions ($\rho_{\text{BD-SVD}} = 151$

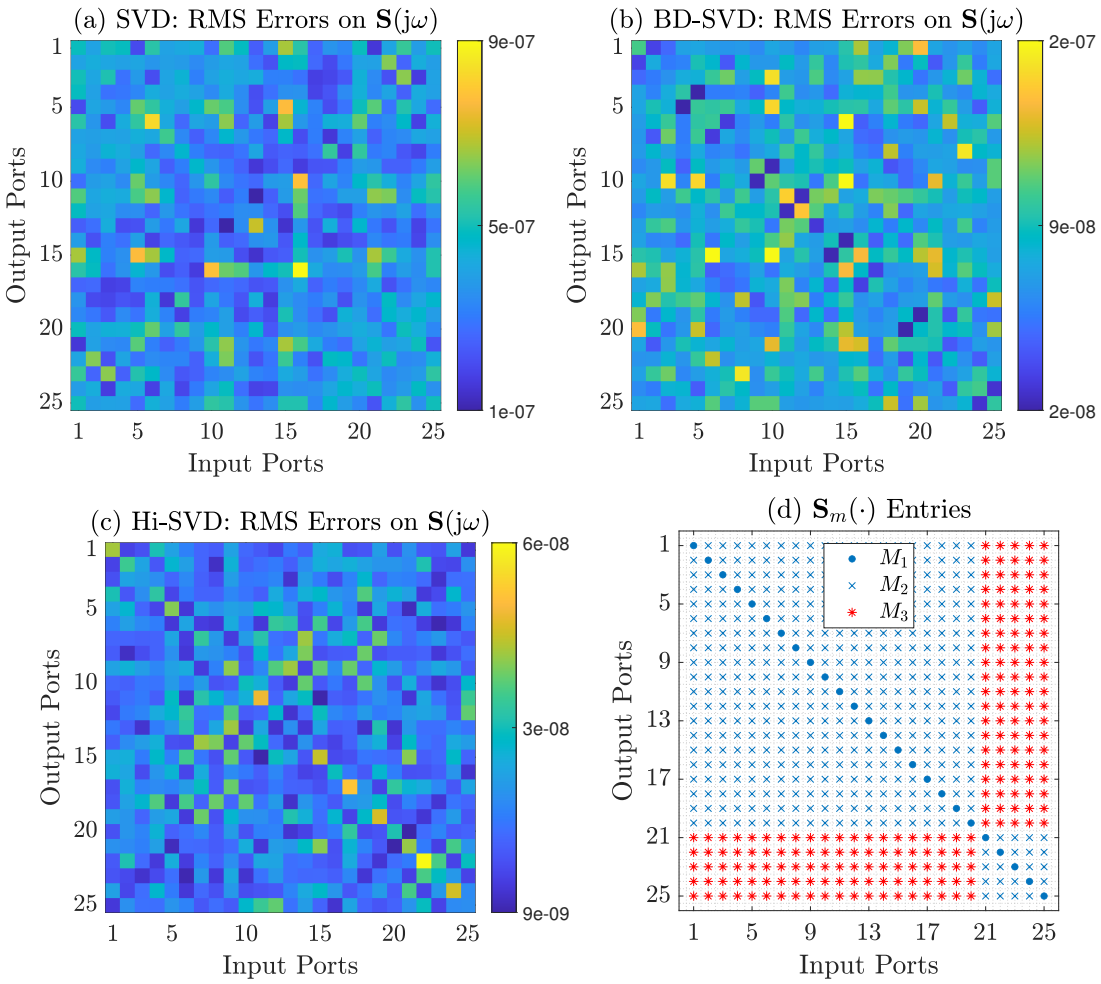

Fig. 2.14 Comparison of data reduction techniques on a 25-port box. Compression accuracy on all transfer function entries via: Standard SVD(a), proposed Block-Diagonal SVD (b) and Hierarchical SVD (c). Panel (d) indicates the grouping of transfer matrix elements according to the DC zeros responses.

and $\rho_{\text{Hi-SVD}} = 108$) is almost double with respect to a standard SVD ($\rho = 73$). Yet, the compression factor is close to 50%, already considering as reference only the upper triangular $P(P+1)/2$ entries of the transfer function. All data compression techniques provide a similar reconstruction error for all scattering elements, as reported in Fig. 2.14. Nevertheless, the proposed structured strategies provide a drastic improvement in terms of DC accuracy, as reported in Table 2.3.

Scaling up to 400 ports

This section summarizes the data compression result on shielding enclosures with increasing size in terms of ports. With reference to the $p \times p$ regular grid covering the enclosure aperture, results are presented for boxes with $P = 64$, 100, 225 and 400 ports, correspondingly $p = 8, 10, 15, 20$ grid branches.

The data conditioning procedure provided in Section 2.1.3 was applied to all cases as in the previous example to obtain a set of regularized parameters defined in the band $[0, 1]$ GHz. A compression of all (modal) scattering responses was obtained by means of the proposed BD-SVD, applied with truncation thresholds $\{10^{-8}, 10^{-6}, 10^{-6}\}$ to each set $\mathbf{M}_\nu$, with $\nu = \{1, 2, 3\}$. Table 2.2 shows a summary of the compression results. What is striking about this table is the dramatic reduction in terms of basis functions ρ when the number of response (i.e. ports) increases. Indeed, considering as reference only the upper-triangular responses of the transfer function, the compression factor defined as $\frac{2\rho}{P(P+1)}$ reaches an impressive value of 0.3% for the largest ports count ($P = 400$). On the other hand, even if the basic SVD [34] result in almost half the quantity of basis functions to obtain an analogous accuracy level, the total elapsed time required by the two algorithms is similar, with a minor memory requirement for the BD-SVD while computing the approximation (2.31). To conclude, Table 2.3 compares the absolute errors on the DC matrix reconstruction for the main system representations (scattering, admittance and impedance). As expected, all scattering errors of the proposed compression strategy are coherent with the accuracy threshold selected for the set $\mathbf{M}_1$ (10^{-8}), which collects the DC eigenvalues contribution. The presented BD-SVD outperforms the standard SVD [34] of at least an order of magnitude for all cases, even if the same threshold error is used in both SVD approximations, with a major advancement considering the mean value of the error due to the proposed reconstruction potentiality to guarantee that all vanishing matrix entries remains null at DC.

2.3 Conclusions

This Chapter provided a complete strategy to preprocess large-scale low-loss frequency data of electromagnetic systems with a large ports count, with

Table 2.2 Data compression results. The second column indicates the number of upper-triangular responses of the transfer function, used to denote compression level. The RMSE is computed on $\mathbf{S}(\mathrm{j}\omega)$. Adapted from [30] © 2022 IEEE.

P	$P(P+1)/2$	SVD as in [34]				Block-Diagonal SVD			
	Total	ρ	% of Total	Time (s)	RMSE	ρ	% of Total	Time (s)	RMSE
25	325	73	22.46	0.81	$9.16 \cdot 10^{-7}$	151	46.46	0.43	$1.79 \cdot 10^{-7}$
64	2080	80	3.85	13.7	$9.55 \cdot 10^{-7}$	198	9.52	4.89	$3.20 \cdot 10^{-7}$
100	5050	82	1.62	2.13	$8.22 \cdot 10^{-7}$	205	4.05	2.60	$1.93 \cdot 10^{-7}$
225	25425	79	0.31	14.5	$1.07 \cdot 10^{-6}$	212	0.83	18.51	$2.38 \cdot 10^{-7}$
400	80200	85	0.11	40.0	$1.04 \cdot 10^{-6}$	238	0.30	42.89	$2.76 \cdot 10^{-7}$

Table 2.3 DC accuracy of data compression. $\delta^{\mathbf{H}}$ is the absolute error between original and compressed DC matrices, with $\mathbf{H} = \{\mathbf{S}, \mathbf{Y}, \mathbf{Z}\}$ indicating the scattering, admittance or impedance representation, respectively.

	$\delta^{\mathbf{S}}_{\max}$	$\delta^{\mathbf{S}}_{\text{mean}}$	$\delta^{\mathbf{Y}}_{\max}$	$\delta^{\mathbf{Y}}_{\text{mean}}$	$\delta^{\mathbf{Z}}_{\max}$	$\delta^{\mathbf{Z}}_{\text{mean}}$
P	SVD as in [34]					
25	$6.02 \cdot 10^{-8}$	$1.50 \cdot 10^{-8}$	$1.97 \cdot 10^{-5}$	$8.34 \cdot 10^{-6}$	$4.00 \cdot 10^{6}$	$1.07 \cdot 10^{6}$
64	$6.83 \cdot 10^{-8}$	$1.03 \cdot 10^{-8}$	$3.73 \cdot 10^{-5}$	$1.35 \cdot 10^{-5}$	$3.74 \cdot 10^{6}$	$6.90 \cdot 10^{5}$
100	$1.58 \cdot 10^{-8}$	$2.56 \cdot 10^{-9}$	$1.13 \cdot 10^{-5}$	$3.34 \cdot 10^{-6}$	$1.42 \cdot 10^{6}$	$1.63 \cdot 10^{5}$
225	$3.87 \cdot 10^{-8}$	$2.71 \cdot 10^{-9}$	$7.42 \cdot 10^{-6}$	$2.06 \cdot 10^{-6}$	$3.89 \cdot 10^{6}$	$2.30 \cdot 10^{5}$
400	$5.20 \cdot 10^{-8}$	$3.02 \cdot 10^{-9}$	$2.27 \cdot 10^{-5}$	$5.00 \cdot 10^{-6}$	$4.56 \cdot 10^{6}$	$1.86 \cdot 10^{5}$
P	Block-Diagonal SVD					
25	$4.90 \cdot 10^{-14}$	$2.06 \cdot 10^{-15}$	$1.83 \cdot 10^{-11}$	$3.33 \cdot 10^{-12}$	4.33	0.256
64	$1.25 \cdot 10^{-9}$	$2.89 \cdot 10^{-11}$	$5.39 \cdot 10^{-7}$	$4.54 \cdot 10^{-8}$	$1.32 \cdot 10^{5}$	$2.38 \cdot 10^{3}$
100	$8.79 \cdot 10^{-10}$	$8.74 \cdot 10^{-12}$	$9.96 \cdot 10^{-8}$	$5.57 \cdot 10^{-9}$	$8.66 \cdot 10^{4}$	774
225	$2.48 \cdot 10^{-9}$	$9.03 \cdot 10^{-12}$	$3.70 \cdot 10^{-7}$	$1.04 \cdot 10^{-8}$	$2.44 \cdot 10^{5}$	786
400	$2.46 \cdot 10^{-9}$	$9.11 \cdot 10^{-12}$	$7.11 \cdot 10^{-7}$	$1.31 \cdot 10^{-8}$	$2.22 \cdot 10^{5}$	702

the main objective of improving numerical performance of macromodeling algorithms in terms of processing time and memory requirements, accuracy, and robustness. The main outcomes of the presented preprocessing phase are:

1. a procedure to define a suitable physic-based DC circuit, necessary to remove possible system degeneracy at all frequencies due to the low-loss (or in the limit case loss-less) nature of the electromagnetic structure; this result is particularly appealing when a static characterization of the device is not available from other sources, as field measurements or specialized solvers;

2. a full-band dataset for macromodeling purposes, filling unknown frequency regions with consistent synthetic responses that smoothly merged with the original ones; this solution is achieved by means of a robust regularization and extrapolation procedure, designed on a suitable data projection into an asymptotic modal domain;
3. a modal representation of the system to structurally isolate the dominant low-frequency content in different transfer matrix elements; this step enables to take proper care of the fitting accuracy in this frequency region, which is particularly relevant (and critical) when the final model is employed for transient simulations with nonlinear components;
4. two structured data compression strategies based on a customized Singular Value Decomposition (SVD) approximation to reduce the overall number of responses necessary to characterize the (full) original system, while preserving low-frequency features of the data.

The overall result is a robust preprocessing strategy for large-scale data that guarantees a suitable starting point for any modeling algorithm, by reducing model sensitivity to terminations as will be documented in Chapter 3. Indeed, it will be demonstrated that the above steps enable to strongly attenuate the effect of loading conditions to the full-band (compressed) representation of the original system.

Several numerical results based on shielding enclosures of increasing size (up to 400 ports) were presented in this Chapter. All examples support the above observations, demonstrating the effectiveness of proposed algorithms. Beyond the considered shielding enclosures modeling and simulation problem, the proposed data conditioning approaches (regularization, extrapolation and structured compression) are completely general and can be applied to other applications with similar characteristics.

Chapter 3

Macromodeling of large-scale LTI systems

This Chapter addresses the problem of constructing a low-complexity model of a passive large scale LTI system starting from frequency tabulated data. This goal is addressed by splitting the overall macrodeling challenges in two sub-problems:

1. compact model extraction of a low-loss system with massive ports count, and strong sensitivity to the loading condition, especially at low-frequency;
2. passivity characterization and enforcement of a large-scale model.

As already discussed in the introduction, shielding enclosures applications drive the modeling complexity to a higher level with respect to state-of-the-art solutions. On the one hand, the number of lumped ports can easily reach hundreds or even thousands of elements, which are used to close the shield aperture with a diode grid, see Fig. 1.4. On the other hand, the large number of poles necessary to reach the required full-band accuracy leads to a representation of the system that is far beyond the standard definition of "reduced-order". Hence, trough this work the term *large-scale* refers to both dynamic order and number of input/output ports.

Even if formulations of standard rational identification as Vector Fitting (VF) [21] are already available in a form that is able to scale favorably with the number of ports, as the Fast VF [24, 34], or its parallel implementation for

multicore [25] and GPU hardware [50], various challenges remain, as discussed below.

Error magnification due to the change of representation is a well-known problem of macromodeling techniques. A reduced-order model usually provides a high level of accuracy in the representation used for the fitting, usually scattering with a reference resistance level R_0, while loading the model with a termination different from R_0 may degrade the approximation quality [51]. This situation is exacerbated in case of low-loss (and in the limit case lossless) electromagnetic components due to the poor conditioning in the low-frequency range. A significant number of ports aggravates this situation and further stresses the macromodeling procedure due to the increased level of (computational) complexity. It is thus fundamental to have an efficient fitting framework to handle the modeling complexity while at least mitigating the sensitivity to loading condition.

Once a proper model of the multiport structure is somehow obtained, the second major macromodeling challenge comes into play: the passivity characterization and enforcement of a large-scale LTI system.
The addressed shielding enclosures are passive structures that by definition are unable to generate energy on their own [47, 52–54]. Thus, a representative macromodel must not only reproduce the electromagnetic behavior of the enclosure, but also inherit the passivity property. To this end, state-of-the-art macromodeling techniques are equipped with suitable perturbation schemes that process the (non-passive) outcome of an identification approach (as VF) and enforce the model passivity by solving a constrained problem [55–67]. Indeed, the model passivity qualification (and enforcement) is a fundamental step to guarantee the stability of any transient numerical simulation [68].

Regardless of the recent developments offered by the literature [69], standard passivity enforcement schemes are still not suitable to scale with the complexity offered by shielding enclosures. On the one hand, modifying the macromodel coefficients and guaranteeing an overall passive behavior is a straightforward task with either standard [64] or advanced perturbation approaches [70, 71], if a suitable passivity characterization is available. On the other hand, finding the location of passivity violations represents the main criticality of the enforcement procedure. A suitable perturbation scheme must provide both the model

qualification (passive/non-passive) and the exact location of passivity violations, the latter necessary for the successive perturbation step. For this reason, the passivity characterization is the most demanding phase from both the computational and storage perspectives. Before [45], to the best of the author's knowledge, robust approaches compatible with the large-scale application of shielding enclosures and requiring limited computing resources were not available.

These challenges motivated a novel interest in a passivity verification technique based on an adaptive-sampling approach [45]. The latter strategy, reported in Section 3.3, completely avoids both Hamiltonian and LMI conditions in favor of an efficient and reliable scheme, here validated with an extensive regression test. This kind of framework does not allow for a theoretical assurance of finding all passivity violations but can be equipped with a number of strategies to make the occurrence of missed violations very unlikely.

This chapter is organized as follows. Section 3.1 introduces the background material for rational approximation and passivity verification of LTI systems. Section 3.2 focuses on the first macromodeling challenge by introducing a complete framework that addresses the complexity of the system in terms of electrical ports while preserving the low-frequency behavior [30]. This result is obtained by integrating standard fitting techniques as Vector Fitting (VF) with the compressed data representation presented in Chapter 2. Section 3.3 addresses the passivity characterization of a large-scale model by providing an efficient and reliable sampling-based passivity verification scheme [45]. Section 3.4 provides a passivity enforcement algorithm compatible with the presented structured and compressed macromodeling framework. Section 3.5 validates the overall procedure, both in terms of passivity verification and fitting. We will demonstrate that the outcome of this Chapter is an efficient strategy to build a compressed passive macromodel, with a DC-preserving reduction approach and a reduced sensitivity to loading conditions.

3.1 Background on model fitting and passivity

This chapter provides a complete macromodeling framework for passive Linear Time Invariant (LTI) systems. To this end, we include the following fundamental background material for later use. Among the several rational-approximation fitting approaches [72, 23, 20, 22], we select and summarize the VF algorithm in Section 3.1.1. Notice that the provided framework is nevertheless not restricted to this identification scheme, and it is generally suitable for any rational modeling approach. Thus, we introduce in Sec. 3.1.2 the Compressed Macromodel (CM) framework presented in [34] as starting fitting strategy to address the large-scale nature of the shielding enclosure application. Section 3.1.3 addresses the surrogate model passivity characterization to suitably retain this system property through the modeling step and ensure stable transient simulation results. Section 3.1.4 concludes this background introduction with a comparison on the three main classes of passivity verification approaches, motivating the addressed problem.

3.1.1 Rational fitting: the VF algorithm

This section presents the well-known Vector Fitting (VF) algorithm [21] in its most efficient sequential implementation to deal with large-scale multiport structures, namely the Fast-VF [24]. VF is one of the most popular data-driven strategies for LTI systems identification [73–75], and over the last decades it has been already exploited for multicore parallelization [25] and optimization for GPU architectures [50].

Starting from a set of tabulated input/output frequency responses $\{\breve{\mathbf{H}}(s_k) \in \mathbb{C}^{P\times P}, k = 1, \cdots, K\}$, with P electrical ports and $s_k = \mathrm{j}\omega_k$, this algorithm builds a rational approximation of the target system transfer matrix and solves the fitting problem

$$\breve{\mathbf{H}}(s_k) \approx \mathbf{H}(s_k) \quad \text{where} \quad \mathbf{H}(s) = \sum_{n=1}^{\bar{n}} \frac{\mathbf{R}_n}{s - p_n} + \mathbf{R}_0 \tag{3.1}$$

where the set of $\bar{n}$ poles p_n and residues $\mathbf{R}_n$ are iteratively estimated. Note that (3.1) assumes a common set of poles for all transfer matrix elements.

The VF algorithm performs a two-step procedure. First, a pole relocation scheme is applied to iteratively estimate model poles p_n. Then, the corresponding residues $\mathbf{R}_n$ are evaluated. The fitting problem (3.1) is solved by improving an initial guess of the model poles p_n every iteration.
We start defining the rational model representation (3.1) similarly to [20] as

$$\mathbf{H}(s;\boldsymbol{x}) = \frac{\mathbf{N}(s;\boldsymbol{x})}{D(s;\boldsymbol{x})} = \frac{\sum_{n=0}^{\bar{n}} \mathbf{C}_n \phi_n(s)}{\sum_{n=1}^{\bar{n}} d_n \phi_n(s) + 1} \tag{3.2}$$

where partial fractions are collected as

$$\phi_0(s) = 1 \quad \text{and} \quad \phi_n(s) = \frac{1}{s - p_n}\,, \text{ with} \quad n = 1, \cdots, \bar{n}\,. \tag{3.3}$$

The model decision variables of (3.2) are gathered in

$$\boldsymbol{x} = (\boldsymbol{c}_1, \cdots, \boldsymbol{c}_{P^2}, \boldsymbol{d})^{\mathsf{T}} \tag{3.4}$$

with

$$\boldsymbol{c}_\ell = (\boldsymbol{c}_{\ell,0}, \cdots, \boldsymbol{c}_{\ell,\bar{n}})^{\mathsf{T}}\,, \quad \boldsymbol{d} = (d_1, \cdots, d_{\bar{n}})^{\mathsf{T}} \tag{3.5}$$

where $\boldsymbol{c}_{\ell,n}$ denotes the (i,j)-th component of matrix $\mathbf{C}_n \in \mathbb{R}^{P\times P}$ mapped through the linear index $\ell = i+(j-1)P$. With this formulation, the denominator $D(s;\boldsymbol{x})$ is the classic VF weight function, used to multiply both sides of (3.1) and to obtain the linearized optimization problem

$$\sum_{n=0}^{\bar{n}} \mathbf{C}_n \phi_n(s_k) - \left(\sum_{n=1}^{\bar{n}} d_n \phi_n(s_k) + 1\right) \breve{\mathbf{H}}(s_k) \approx \mathbf{0} \tag{3.6}$$

that must be solved for all frequencies $s_k = \mathrm{j}\omega_k$ in a least square sense. We now collect the partial fraction basis evaluated in the set of available frequency samples $\{s_k = \mathrm{j}\omega_k, k = 1, \cdots, K\}$ to obtain matrices

$$\boldsymbol{\Phi}_1 = \begin{pmatrix} \phi_1(s_1) & \cdots & \phi_{\bar{n}}(s_1) \\ \vdots & \ddots & \vdots \\ \phi_1(s_K) & \cdots & \phi_{\bar{n}}(s_K) \end{pmatrix} \quad \text{and} \quad \boldsymbol{\Phi}_0 = \begin{pmatrix} \mathbf{1}_K & \boldsymbol{\Phi}_1 \end{pmatrix} \tag{3.7}$$

where $\mathbf{1}_1 = (1, \cdots, 1)^\mathsf{T} \in \mathbb{R}^K$ is a column vector of ones. Then, we cast the ℓ-th entry of each transfer matrix $\breve{H}_\ell(s_k) = \breve{H}_{i,j}(s_k)$ into

$$\breve{\mathbf{H}}_\ell = \operatorname{diag}\left\{\breve{H}_\ell(s_1), \cdots, \breve{H}_\ell(s_K)\right\}, \quad \boldsymbol{b}_\ell = \left(\breve{H}_\ell(s_1), \cdots, \breve{H}_\ell(s_K)\right)^\mathsf{T}. \tag{3.8}$$

The fitting problem (3.6) formulated for each individual element ℓ of the transfer matrix and for all K frequency samples reads

$$\begin{pmatrix}\boldsymbol{\Phi}_0 & -\breve{\mathbf{H}}_\ell\boldsymbol{\Phi}_1\end{pmatrix}\begin{pmatrix}\boldsymbol{c}_\ell \\ \boldsymbol{d}\end{pmatrix} \approx \boldsymbol{b}_\ell \tag{3.9}$$

and can be stacked for all P^2 responses to obtain the complete linear system

$$\begin{pmatrix}\boldsymbol{\Phi}_0 & \mathbf{0} & \cdots & \mathbf{0} & -\breve{\mathbf{H}}_1\boldsymbol{\Phi}_1 \\ \mathbf{0} & \boldsymbol{\Phi}_0 & \cdots & \mathbf{0} & -\breve{\mathbf{H}}_2\boldsymbol{\Phi}_1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \boldsymbol{\Phi}_0 & -\breve{\mathbf{H}}_{P^2}\boldsymbol{\Phi}_1\end{pmatrix}\boldsymbol{x} = \begin{pmatrix}\boldsymbol{b}_1 \\ \boldsymbol{b}_2 \\ \vdots \\ \boldsymbol{b}_{P^2}\end{pmatrix}. \tag{3.10}$$

Solving (3.10) for $\boldsymbol{x}$ (3.4) gives coefficients of the denominator (the weight function) $\boldsymbol{d}$ and all residues $\boldsymbol{c}_\ell$, which are disregarded. The initial set of poles p_n is replaced with zeros of the denominator in (3.1), denoted with z_n. The main idea is that the zeros of $D(s;\boldsymbol{x})$ must simplify the (unknown) exact system poles when (3.10) is solved exactly. Thus, the denominator of (3.1) is unitary if the set p_n matches the poles of the original structure. The zeros z_n are computed solving an eigenvalue problem associated to the state-space realization $\{A_D, B_D, C_D, D_D\}$ of the denominator such that

$$z_n = \operatorname{eig}\left\{A_D - B_D D_D^{-1} C_D\right\} \tag{3.11}$$

and the (relocation) process $p_n \leftarrow z_n$ is iterated until convergence. When an unstable zero appears, the real part of the corresponding pole is changed in sign to ensure the model stability (*pole flipping* step). At last, when the poles stabilize the original fitting problem (3.1) is solved and model coefficients $\mathbf{R}_n$ are computed. Standard modifications of (3.1) and (3.9) ensure real-valued residues also in the case of complex-conjugate pairs of poles. We refer to [20] for implementation details.

A direct solution of (3.10) is computationally expensive due to the large dimension of the regressor matrix. Indeed, one should notice that the dimension of the left-side matrix is $KP^2 \times (P^2(\bar{n}+1)+\bar{n})$, which does not scale favorably with the model complexity, both in terms of ports P and poles $\bar{n}$. Nevertheless, an efficient implementation of the VF takes advantage of a "thin" QR factorization to speed up and optimize the relocation process. Starting from (3.9), we can write

$$\begin{pmatrix} \boldsymbol{\Phi}_0 & -\breve{\mathbf{H}}_\ell \boldsymbol{\Phi}_1 \end{pmatrix} = \mathbf{Q}_\ell \mathbf{R}_\ell = \mathbf{Q}_\ell \begin{pmatrix} \mathbf{R}_\ell^{11} & \mathbf{R}_\ell^{12} \\ \mathbf{0} & \mathbf{R}_\ell^{22} \end{pmatrix} \tag{3.12}$$

where the only matrix associated to the denominator residues is the lower block $\mathbf{R}_\ell^{22} \in \mathbb{R}^{\bar{n}\times\bar{n}}$. Stacking all terms $\mathbf{R}_\ell^{22}$ for all ports contribution leads to a *compressed* set of equations

$$\begin{pmatrix} \mathbf{R}_1^{22} \\ \vdots \\ \mathbf{R}_{P^2}^{22} \end{pmatrix} \boldsymbol{d} = \begin{pmatrix} \mathbf{Q}_1^{\mathsf{T}} \boldsymbol{b}_1 \\ \vdots \\ \mathbf{Q}_{P^2}^{\mathsf{T}} \boldsymbol{b}_{P^2} \end{pmatrix} . \tag{3.13}$$

The regressor dimension is now $\bar{n}P^2 \times \bar{n}$ and the denominator coefficients are computed with a major saving of memory. Nevertheless, the presented Fast-VF scheme [24] requires P^2 computations of QR factorizations in (3.12). Thus, this operation is significantly demanding when P reaches thousands of elements.

3.1.2 The Compressed Macromodeling framework

Section 2.2 introduced a SVD data compression technique based on the established method of [34], and presented structured SVD approaches [30] that are able to preserve particular features in groups of responses, such as single or multiple zeros at DC. All reduction methods lead to an approximation of the matricized system responses, which can be written as

$$\mathbf{X} \approx \bar{\mathbf{X}} = \bar{\mathbf{W}} \bar{\mathbf{V}}^{\mathsf{T}} , \tag{3.14}$$

where $\bar{\mathbf{V}} \in \mathbb{R}^{P^2\times\rho}$ is either full or block-sparse depending on whether the standard method [34] or one of the structured approaches (BD-SVD or Hi-SVD)

is adopted. In the latter, $\bar{\mathbf{V}}^{\mathsf{T}}$ accounts also for the permutation matrix $\mathbf{P}$ used in (2.28) to reorder responses for individual group processing. The matrix $\bar{\mathbf{W}} \in \mathbb{C}^{K\times\rho}$ collects in its columns the various *basis functions* $\bar{\boldsymbol{w}}_q$ obtained by SVD approximation, which may be further separated in groups characterized by different low-frequency behavior.

The *compressed macromodeling* framework of [34] applies to each basis function $\bar{\boldsymbol{w}}_q$, building a set of rational macromodels

$$w_q(s) = r_{q\infty} + \sum_{n=1}^{N_w} \frac{r_{qn}}{s - p_n}\,, \quad q = 1, \cdots, \rho \tag{3.15}$$

where p_n are a set of N_w common poles, r_{qn} are the associated model coefficients (residues), with $r_{q\infty}$ direct coupling term. This *Compressed Macromodel (CM)* can be identified by feeding to the Fast VF algorithm [24] the set of all basis functions, concurrently, and by solving the fitting problem

$$w_q(\mathrm{j}\omega_k) \approx (\bar{\boldsymbol{w}}_q)_k, \quad q = 1, \ldots, \rho,\ k = 1, \ldots, K \tag{3.16}$$

where each k-th element $(\bar{\boldsymbol{w}}_q)_k$ of every basis is associated to the frequency sample ω_k. The resulting CM (3.15) can be cast in a row vector form as

$$\boldsymbol{w}(s) = \begin{bmatrix} w_1(s) & w_2(s) & \cdots & w_\rho(s) \end{bmatrix}, \tag{3.17}$$

which is equivalent to the state-space realization

$$\boldsymbol{w}(s)^{\mathsf{T}} = \mathbf{C}_w(s\mathbf{I} - \mathbf{A}_w)^{-1}\boldsymbol{b}_w + \boldsymbol{d}_w\,. \tag{3.18}$$

where $\mathbf{A}_w \in \mathbb{R}^{N_w\times N_w}$ is a diagonal matrix that retains poles p_n, $\boldsymbol{b}_w = \mathbf{1}_{N_w}$ is a vector of ones, $\mathbf{C}_w \in \mathbb{R}^{\rho\times N_w}$ stores model residues r_{qn}, and $\boldsymbol{d}_w \in \mathbb{R}^{\rho}$ collects the direct coupling terms $r_{q\infty}$. Classic adjustments [20] enable to modify the above representation to deal with complex conjugate pair of model poles while preserving the realness of all coefficients. The original $P \times P$ system transfer function can be retrieved as

$$\mathbf{H}(s) = \sum_{n=1}^{\bar{n}} \frac{\mathbf{R}_n}{s - p_n} + \mathbf{R}_0 = \mathbf{C}(s\mathbb{I} - \mathbf{A})^{-1}\mathbf{B} + \mathbf{D} = \mathrm{mat}(\bar{\mathbf{V}}\boldsymbol{w}(s)^{\mathsf{T}}) \tag{3.19}$$

where mat reconstructs a matrix of proper size from a vector collecting its entries, and $\{\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}\}$ is a set of compatible state-space matrices. For later use, notice that the number of system states N, i.e. the dynamic order and size of matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, is in general $N \gg P$, and in the particular case of a minimal representation (with full-rank $\mathbf{R}_n$) can be defined as $N = \bar{n}P$ [20, 76, 77].
The above compressed framework enables a dramatic reduction of model coefficients, to be identified by enforcing (3.16) and necessary for the subsequent passivity enforcement (discussed next). Indeed, the number of required decision variables to retrieve the system responses is related to the overall count of basis functions, and $\rho \ll P^2$ by the low-rank approximation (3.14). This leads to a drastic reduction of the computational cost required to build a full passive behavioral model. On the other hand, if the macromodel structure does not inherit all features of basis functions groups, the effort spent in the DC-preserving data compression is lost at this stage.

On the computational cost of the CM framework

The main bottleneck of the SVD-based compression method is the memory cost since the overall number of frequency samples K and responses P could result in a matrix $\begin{bmatrix} \mathrm{Re}\{\mathbf{X}\} \\ \mathrm{Im}\{\mathbf{X}\} \end{bmatrix} \in \mathbb{R}^{2K \times P^2}$ of an excessive size for the available resources. In terms of operations, this SVD requires $O(4P^2K^2)$ flops, assuming that $P^2 > K$, and leads to ρ basis functions that can be fed to an identification algorithm such as Fast-VF. On the other hand, the Fast-VF algorithm requires a set of QR factorizations to speed up the pole relocation process at each iteration, see (3.12). This operation is usually less expensive in terms of memory since applied P^2 times on smaller matrices of size $\mathbb{C}^{K \times (2\bar{n}+1)}$, where $\bar{n}$ indicates the number of poles. An effective implementation would perform these QR factorization using parallel computing threads. Notice that also the cost of the least squares (LS) to compute the denominator coefficients at each iteration, see (3.13), is dominated by the number of responses. This cost is $O(2P^2\bar{n}^3)$ flops if all P^2 responses are processed concurrently.

Table 3.1 summarizes the overall cost in terms of operations for the proposed CM framework and the standard Fast-VF. Notice that we reasonably assumed that $2K \geq (2\bar{n}+1)$ for the QR, that $P^2 \geq K$ for the SVD, and that $\rho \ll P^2$ for

Table 3.1 Summary of the number of operations required by the standard Fast-VF and by the proposed CM framework, the latter applied on ρ basis functions. Residues computation cost is omitted.

	QR (flops)	LS (flops)	SVD (flops)	Total for I iterations
Fast-VF	$O(P^2 4K(2\bar{n}+1)^2)$	$O(2P^2\bar{n}^3)$	-	$O(4IP^2(K(2\bar{n}+1)^2)+\bar{n}^3/2)$
CM	$O(\rho\, 4K(2\bar{n}+1)^2)$	$O(2\rho\bar{n}^3)$	$O(P^2 4K^2)$	$O(4I\rho(K(2\bar{n}+1)^2+\bar{n}^3/2)+4P^2K^2)$

the SVD truncation. Considering a number of VF iterations I and comparing the overall cost of Fast-VF and CM, the latter is preferable when

$$O(4IP^2(K(2\bar{n}+1)^2+\bar{n}^3/2)) > O(4I\rho(K(2\bar{n}+1)^2+\bar{n}^3/2)+4P^2K^2) \quad (3.20)$$

is satisfied. Eliminating common factors and removing the terms that under our running assumptions are negligible, this condition simplifies to $4I\bar{n}^2 > K$. The larger the number of VF iterations, the more costly is Fast-VF with respect to the single instance of the SVD. Correspondingly, the larger is the number of required system poles for a fixed number of frequency samples, the more convenient is the CM approach (with quadratic scaling). For the systems that are the subject of this work, it is always verified that $\rho \ll P^2$, which is in any case the dominant and key condition that supports applicability of the CM approach.

3.1.3 Passivity of LTI systems

This section provides a background on multiport systems passivity, summarizing well-known results [20] and setting notation for later sections. This work considers only models in the scattering representation, but the extended framework for the passivity characterization of other system representation (immitance or hybrid) applies with minor modifications, as already detailed in the literature [20].

The passivity characterization of a P-port system starts with the definition of the absorbed instantaneous power as

$$p(t) = \boldsymbol{v}(t)^\mathsf{T}\boldsymbol{i}(t) = \boldsymbol{i}(t)^\mathsf{T}\boldsymbol{v}(t) = \boldsymbol{a}(t)^\mathsf{T}\boldsymbol{a}(t) - \boldsymbol{b}(t)^\mathsf{T}\boldsymbol{b}(t) \quad (3.21)$$

where $\boldsymbol{v}(t) = [v_1(t), \ldots, v_P(t)]^\mathsf{T}$ and $\boldsymbol{i}(t) = [i_1(t), \ldots, i_P(t)]^\mathsf{T}$ are time-varying column vectors storing ports voltages and currents, respectively, and $\boldsymbol{a}(t)$ and $\boldsymbol{b}(t)$ collect input/output scattering waves with reference impedance R_0.

The cumulative net energy absorbed by the multiport system at a given instant of time t is defined as

$$\mathcal{E}(t) = \int_{-\infty}^{t} p(\tau) d\tau \,. \tag{3.22}$$

For any passive multiport system, the cumulative net energy (3.22) is non-negative for any time t so that

$$\mathcal{E}(t) \geq 0 \,, \, \forall t \,. \tag{3.23}$$

The above definition of passivity is now made more precise for the particular case of lumped P-ports LTI systems. For this family of Multi Input Multi Output (MIMO) devices, the input/output behavior can be represented in terms of a transfer function $\mathbf{H}(s) \in \mathbb{C}^{P \times P}$, where s is the Laplace variable. We recall that Section 3.2 provides a compressed identification framework to obtain a model of $\mathbf{H}(s) \in \mathbb{C}^{P \times P}$, ultimately represented in a pole-residue form (3.19) or as an equivalent state-space realization. Thus, we will assume in the following two working conditions:

- $\mathcal{A}_1$ the asymptotic stability of (3.1) is guaranteed by construction so that all model poles p_n have strictly negative real part, as obtained with the VF, see Section 3.1.1;
- $\mathcal{A}_2$ all matrices in (3.1) are real-valued, as guaranteed via standard manipulations [20].

We have the following result.

Theorem 3.1 *[20, 53, 54, 47] An LTI system in scattering representations with transfer matrix* $\mathbf{H}(s)$ *is passive if and only if* $\mathbf{H}(s)$ *is Bounded Real (BR).*

In the case of a rational transfer matrix $\mathbf{H}(s)$ the above theorem simplifies to a set of simple conditions.

Theorem 3.2 *[20, 53, 54] A rational matrix* $\mathbf{H}(s)$ *is BR if and only if*

$\mathcal{C}_1$. $\mathbf{H}(s)$ *has no poles in* $\mathbb{C}^+$ *such that the system is asymptotically stable, with* $\mathrm{Re}\{p_n\} < 0\ \forall n$*;*
$\mathcal{C}_2$. $\mathbf{H}^*(\mathrm{j}\omega) = \mathbf{H}(-\mathrm{j}\omega)$*;*
$\mathcal{C}_3$. $\mathbb{I} - \mathbf{H}(\mathrm{j}\omega)^{\mathsf{H}}\mathbf{H}(\mathrm{j}\omega) \geq 0,\ \forall\omega \in \mathbb{R}$.

Under the working conditions $\mathcal{A}_{1,2}$ of (4.8), $\mathcal{C}_{1,2}$ are guaranteed by construction and only $\mathcal{C}_3$ must be verified for all frequencies.
A reformulation of $\mathcal{C}_3$ based on the Singular Value Decomposition (SVD) of $\mathbf{H}(\mathrm{j}\omega)$ at each frequency denoted as

$$\mathbf{H}(\mathrm{j}\omega) = \mathbf{U}(\mathrm{j}\omega)\mathbf{\Sigma}(\mathrm{j}\omega)\mathbf{V}(\mathrm{j}\omega)^{\mathsf{H}} \tag{3.24}$$

leads to the condition

$$\sigma_{\max}\{\mathbf{H}(\mathrm{j}\omega)\} = \|\mathbf{H}(\mathrm{j}\omega)\|_2 \leq \gamma = 1 \quad \forall\omega \in \mathbb{R}. \tag{3.25}$$

The above sufficient condition for passivity implies that the maximum singular value of the scattering transfer matrix must be below a critical threshold $\gamma = 1$ for all (real) frequency ω. Notice that (3.25) enables to define the *passivity metric* function

$$\varphi(\omega) = \sigma_{\max}\{\mathbf{H}(\mathrm{j}\omega)\} \tag{3.26}$$

which has the interesting property of being continuous, since $\mathbf{H}(\mathrm{j}\omega)$ is regular in an open subset of the complex plane containing the imaginary axis (i.e. there are no poles in $\mathbb{R}^{\mathrm{j}\omega}$), and is generally smooth but not differentiable in a finite number of points, due to crossing of singular values trajectories. The first attribute of (3.26) is particularly appealing since enables the use of sampling approaches for checking the system passivity, as it will be further discussed next.

An alternative passivity characterization of a state-space system involves the well-known Bounded Real Lemma (BRL) or Kalman-Yakubovich-Popov (KYP) Lemma [20, 54, 78].

Lemma 3.1 *A scattering LTI system is passive (dissipative) if and only if*

$$\exists \mathbf{P} = \mathbf{P}^{\mathsf{T}} > 0 : \quad \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{u} \end{pmatrix}^{\mathsf{T}} \begin{pmatrix} \mathbf{A}^{\mathsf{T}}\mathbf{P} + \mathbf{P}\mathbf{A} + \mathbf{C}^{\mathsf{T}}\mathbf{C} & \mathbf{P}\mathbf{B} + \mathbf{C}^{\mathsf{T}}\mathbf{D} \\ \mathbf{B}^{\mathsf{T}}\mathbf{P} + \mathbf{D}^{\mathsf{T}}\mathbf{C} & -(\mathbb{I} - \mathbf{D}^{\mathsf{T}}\mathbf{D}) \end{pmatrix} \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{u} \end{pmatrix} \leq 0 \tag{3.27}$$

The passivity verification problem can be cast in a closed algebraic form. Using the state-space realization of the scattering model (3.19), we define the so-called *Hamiltonian* matrix

$$\boldsymbol{\mathcal{M}} = \begin{pmatrix} \mathbf{A} + \mathbf{B}\mathbf{R}^{-1}\mathbf{D}^{\mathsf{T}}\mathbf{C} & \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathsf{T}} \\ -\mathbf{C}^{\mathsf{T}}\mathbf{S}^{-1}\mathbf{C} & -\mathbf{A}^{\mathsf{T}} - \mathbf{C}^{\mathsf{T}}\mathbf{D}\mathbf{R}^{-1}\mathbf{B}^{\mathsf{T}} \end{pmatrix} \tag{3.28}$$

with $\mathbf{R} = \mathbb{I} - \mathbf{D}^{\mathsf{T}}\mathbf{D}$ and $\mathbf{S} = \mathbb{I} - \mathbf{D}\mathbf{D}^{\mathsf{T}}$. It is well known [79, 64, 20] that a sufficient passivity condition for a scattering model is that $\boldsymbol{\mathcal{M}}$ has no purely imaginary eigenvalues $\mu_k = \mathrm{j}\hat{\omega}_k$. These eigenvalues provide all frequency locations where trajectories of singular values of $\mathbf{H}(\mathrm{j}\omega)$ cross the passivity threshold $\gamma = 1$, as detailed in [64, 79]. Thus, the imaginary eigenvalues of (3.28) can be used to localize model passivity violations on the frequency axis.

Note that the above theorem requires the asymptotic passivity of $\mathbf{H}(\mathrm{j}\omega)$, since the non-singularity of $\mathbb{I} - \mathbf{D}^{\mathsf{T}}\mathbf{D}$ implies that $\|\mathbf{D}\|_2 < \gamma = 1$. The Hamiltonian matrix $\boldsymbol{\mathcal{M}}$ is only defined when $\mathbb{I} - \mathbf{D}^{\mathsf{T}}\mathbf{D}$ is non-singular, equivalently when $\|\mathbf{D}\|_2 < \gamma = 1$. These conditions imply the asymptotic passivity of $\mathbf{H}(\mathrm{j}\omega)$ for $\omega \to \infty$. If the singular values of the direct coupling term $\mathbf{D}$ are too close to the passivity threshold γ, the above non-singularity condition can be relaxed by means of an extended eigenvalue problem as

$$\boldsymbol{\mathcal{M}}_e \boldsymbol{v} = \mathrm{j}\omega_0 \boldsymbol{\mathcal{K}} \boldsymbol{v}, \tag{3.29}$$

where we can define the *extended Hamiltonian pencil* $(\boldsymbol{\mathcal{M}}_{\mathrm{e}}, \boldsymbol{\mathcal{K}})$ as in [80]

$$\boldsymbol{\mathcal{M}}_{\mathrm{e}} = \begin{bmatrix} \mathbf{A} & \mathbf{0} & \mathbf{B} & \mathbf{0} \\ \mathbf{0} & -\mathbf{A}^{\mathsf{T}} & \mathbf{0} & -\mathbf{C}^{\mathsf{T}} \\ \mathbf{0} & \mathbf{B}^{\mathsf{T}} & -\mathbb{I} & \mathbf{D}^{\mathsf{T}} \\ \mathbf{C} & \mathbf{0} & \mathbf{D} & -\mathbb{I} \end{bmatrix} \quad \boldsymbol{\mathcal{K}} = \begin{bmatrix} \mathbb{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbb{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \tag{3.30}$$

Its purely imaginary generalized eigenvalues μ_k can be computed to localize passivity violation of the system. Since obtaining the spectrum of the extended pencil (3.30) increases the computational effort with respect to (3.28), this approach is recommended only if the standard Hamiltonian matrix (3.28) is ill-defined. In this work, the extended pencil (3.30) is used only when $|\sigma_{\max}\{\mathbf{D}\} - 1| < 10^{-4}$, as will be documented in some of the examples reported in Section 3.5.1.

3.1.4 Discussion on passivity verification

To summarize, the three approaches for passivity verification of a (scattering) system are:

1. *Sampling of local passivity conditions*, by verifying whether the passivity metric exceeds the allowed limit $\varphi(\omega) > \gamma = 1$ over a finite number of frequency samples $\{\omega_k,\ k = 1, \ldots, K\}$. This is the computationally less expensive verification in terms of both memory and number of elementary operations, since each frequency ω_k can be processed individually. Indeed, the complexity in computing the maximum singular value $\sigma_{\max}$ at a given frequency point is $O(P^3)$, while only considering the overall number of operations for K samples the complexity grows to $O(KP^3)$. On the other hand, the main disadvantage of this approach is that by checking passivity over a finite number of points some *narrow* violation may be missed by the checking algorithm, since very close to the passivity threshold $\gamma = 1$ or spread over a small frequency region. Nevertheless, another advantage of this approach is that, in view of passivity enforcement, all frequency bands where passivity violations are located and the corresponding local singular value maxima $\sigma_{\max}$ are available for postprocessing.

2. *Linear Matrix Inequalities*, provide a fully algebraic approach, thanks to the BRL presented in Lemma 3.1, thus is very robust and does not require any sort of sampling. Unfortunately, checking the feasibility of (3.27) is a very expensive procedure that requires $O(N^6)$ operations. Even if a particular formulation [81] reduces its cost to $O(N^4)$, this approach is not suitable for a massive number of ports P. On the other, a passivity check based on (3.27) provides only a boolean result, and it is of no

Characterization method	Operations	Storage
Linear Matrix Inequalities [81]	$O(N^4)$	$(N+P)^2$
Hamiltonian (full)	$O((2N)^3)$	$4N^2$
Sampling	$O(K(P^3+P^2\bar{n}))$	P^2
N: model order (states)	K: frequency samples	
P: input-output ports	$\bar{n}$: (common) poles	

Table 3.2 CPU cost and required storage for different passivity characterization strategies. Adapted from [45] © 2021 IEEE.

use for localizing passivity violations in view of a passivity enforcement loop. Those enforcement schemes that embed (3.27) as constraint [55, 82] provide the same scalability of the corresponding check [67], hence are not suitable for a large-scale system.

3. *Spectral properties of Hamiltonian Matrices*, are the state-of-the-art method for low/medium size models since provide an algebraic approach to pinpoint passivity violations. As already mentioned, the solution of eigenproblem (3.28) enables to retrieve passivity violations location along the frequency axis, at the price of finding purely imaginary eigenvalues of $\boldsymbol{\mathcal{M}}$ with $O((2N)^3)$ operations. Even if some improvement has been demonstrated [65, 66] for medium size models, taking advantage of spare-representations and an Arnoldi process attempting to compute only purely imaginary eigenvalues (instead of the all spectrum of $\boldsymbol{\mathcal{M}}$), the extension to a large-scale system is still limited. Indeed, these approaches require the inversion of a $P \times P$ matrix while computing the Krylov subspace iteration, leading to an approach that alleviates only partially the computational effort of a large-scale model. When the number of ports P reaches thousands of elements, this strategy is not appropriate.

A comparison of these strategies in terms of both number of required floating points operations, up to a constant factor associated to the selected implementation/algorithm, and of required storage memory is reported in Table 3.2.

To motivate the proposed approach, we recall here the same example presented while introducing Chapter 2. Supposing that we want to build a rational passive macromodel of a shielding enclosure with $P = 1024$ lumped ports and $\bar{n} = 71$ poles, for a total of $N = 72704$ states, as the one illustrated

in Section 3.5.2. In the following, we will try to provide some estimate in terms of performances for the three classes of passivity characterization approaches. While for both the LMI and the Hamiltonian strategy is possible to know in advance the overall required memory, approximately 43GB and 160GB respectively, for the adaptive sampling scheme we can provide only the single-point evaluation that requires storing 16MB. We recall that for the latter it is not possible to anticipate the total number of required samples K, since the (scalar) investigated function, the passivity metric (3.26), changes accordingly to singular values trajectories. Nevertheless, even considering a very large number $K = 10^5$ of frequency samples, the relative CPU cost for both LMI and Hamiltonian approaches is $2 \cdot 10^5\times$ and $26\times$ higher with respect to the sampling approach.

We conclude that the only approach that is able to manage the large-scale nature of the investigated models is the passivity characterization framework based on sampling strategies. This motivates the focus of the Section 3.3, which will provide a passivity verification scheme [45] aimed at reducing the chances of missing violations via sampling, hence at providing a wrong classification of a passive/non-passive model.

3.2 Structured Compressed macromodeling

This section provides fundamental steps for the construction of a large-scale macromodel, where the term large-scale is mostly related with the electrical ports count. The final goal is to provide a flexible compressed modeling procedure, equipped with a full-band error control strategy, to reduce the induced approximation error while changing system representation, and ultimately attenuate the sensitivity of the model to loading conditions. Thus, the compressed structured macromodeling framework [30] is introduced by integrating the DC-preserving data reduction techniques of Chapter 2.2 with an ad hoc procedure to ensure an accurate model extraction. The procedure is reported enriching the same steps of [30]. Numerical results on shielding enclosures with increasing ports count are presented later in Section 3.5.2, after the introduction of a reliable passivity verification and enforcement scheme in Sections 3.3 and 3.4, respectively.

3.2.1 Fitting structured basis functions

The goal now is to integrate the proposed data-compression techniques of Chapter 2.2 into a suitable identification framework, which generalizes previous approaches by providing full control on the system DC behavior at all macromodeling phases.

The novel contribution of [30] is the realization of a *compressed structured macromodel*

$$\widehat{\boldsymbol{w}}(s) = \begin{bmatrix} \widehat{w}_1(s) & \widehat{w}_2(s) & \cdots & \widehat{w}_\rho(s) \end{bmatrix} \tag{3.31}$$

where each model $\widehat{w}_q(s)$ preserves the feature of each basis $\bar{\boldsymbol{w}}_q$ of $\bar{\mathbf{W}}$. In essence, DC zeros of any order are imposed by construction in the macromodel generation. This condition is obtained by imposing

$$\widehat{\boldsymbol{w}}(s)^{\mathsf{T}} = \boldsymbol{\Gamma}(s) \cdot \boldsymbol{w}(s)^{\mathsf{T}} \tag{3.32}$$

where $\boldsymbol{\Gamma}(s)$ is the block-diagonal matrix

$$\boldsymbol{\Gamma}(s) = \mathrm{diag}\{m_1(s)\mathbb{I}_{\rho_1}, \cdots, m_\nu(s)\mathbb{I}_{\rho_\nu}, \cdots, m_{\bar{\nu}}(s)\mathbb{I}_{\rho_{\bar{\nu}}}\} \tag{3.33}$$

with frequency-dependent weighting factors $m_\nu(s)$ defined for each ν-th group of ρ_ν bases as the high-pass filter

$$m_\nu(s) = \prod_{i=1}^{\nu} \frac{s}{s - p_i^\infty}, \quad \nu = 1, \ldots, \bar{\nu}. \tag{3.34}$$

The particular case of shielding enclosures (to be identified in the modal domain) requires of imposing DC zeros with multiplicity from 0 to 2 for $\nu = \{1, 2, 3\}$, respectively. This can be achieved defining

$$\boldsymbol{\Gamma}(s) = \begin{bmatrix} m_1(s) \cdot \mathbb{I}_{\rho_1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & m_2(s) \cdot \mathbb{I}_{\rho_2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & m_3(s) \cdot \mathbb{I}_{\rho_3} \end{bmatrix}. \tag{3.35}$$

with $\bar{\nu} = 3$ and

$$m_1(s) = 1, \quad m_2(s) = \frac{s}{s - p_2^\infty}, \quad m_3(s) = \frac{s}{s - p_2^\infty} \cdot \frac{s}{s - p_3^\infty}. \tag{3.36}$$

These weighting functions enable to reproduce the required DC behavior, preserving both the rational model form and the unitary bounding condition of scattering systems, since the high-pass filter structure (3.34) guarantees that $|m_\nu(\mathrm{j}\omega)| \leq 1$ for $\nu = 1, \ldots, \bar{\nu}$ and all frequencies. On the other hand, the effect of $\boldsymbol{\Gamma}(s)$ in the modeling band is minimized by selecting all poles p_ν^∞ after the maximum fitting frequency, such that

$$\omega_{\max} < |p_\nu^\infty| \ , \ \forall \nu \tag{3.37}$$

where $\omega_{\max}$ is the maximum frequency used for the model identification. Furthermore, the behavior for $s = \infty$ of the two compressed representations, structured (3.32) and standard (3.17), can be identically tuned by the direct coupling coefficients $r_{q,\infty}$. Indeed, being

$$\lim_{s\to\infty} \boldsymbol{\Gamma}(s) = \mathbb{I}_\rho \tag{3.38}$$

the two asymptotic responses of $\boldsymbol{w}(s)$ and $\widehat{\boldsymbol{w}}(s)$ are the same.

The identification of a macromodel with the proposed structure (3.32) is a task that can be achieved with standard algorithms, such as the Fast VF. Two possible approaches are available. On the one hand, we can adapt the core of VF itself defining modified frequency-dependent basis functions that include the low-frequency structure, i.e. the zero at DC and an additional high frequency pole p_ν^∞. A simpler approach is presented here, following [30]. Indeed, the identification problem can be modified rescaling each basis at all frequencies by the inverse of the corresponding weighting matrix $\boldsymbol{\Gamma}_k^{-1} = \boldsymbol{\Gamma}(\mathrm{j}\omega)^{-1}$. With this second approach, a standard rational representation (3.15) is suitable to the modified fitting problem rewritten as

$$\boldsymbol{w}_q(\mathrm{j}\omega_k) \ \approx \ (\widetilde{\boldsymbol{w}}_q)_k \ = \ (\boldsymbol{\Gamma}_k)_{qq}^{-1}(\bar{\boldsymbol{w}}_q)_k, \quad q \ = \ 1, \ldots, \rho, \ \ k \ = \ 1, \ldots, K \tag{3.39}$$

which can be solved with a standard modeling engine, as the Fast VF. Indeed, all rescaled basis vectors $\widetilde{\boldsymbol{w}}_q$ show a non-vanishing value at DC, since all DC zeros are canceled out when rescaling for the inverse high-pass filter of (3.34). The final structured compressed model is obtained as in (3.32) by multiplication of the rational model with the filtering term $\boldsymbol{\Gamma}(s)$, and by enforcing the prescribed low-frequency behavior by construction.

The result of rescaling procedure (3.39) on a 25-port shielding enclosure is reported in Fig. 3.1. The proposed strategy enables to further improve the fitting process accuracy by including standard weighting schemes [20] directly in the filtering term. For example, all bases can be normalized to a common magnitude level by including an additional weight, in terms of the 2-norm or RMS value of each q-th basis, to each diagonal entry of the filtering term $\mathbf{\Gamma}_{qq}(s)$ in (3.32) and (3.39).

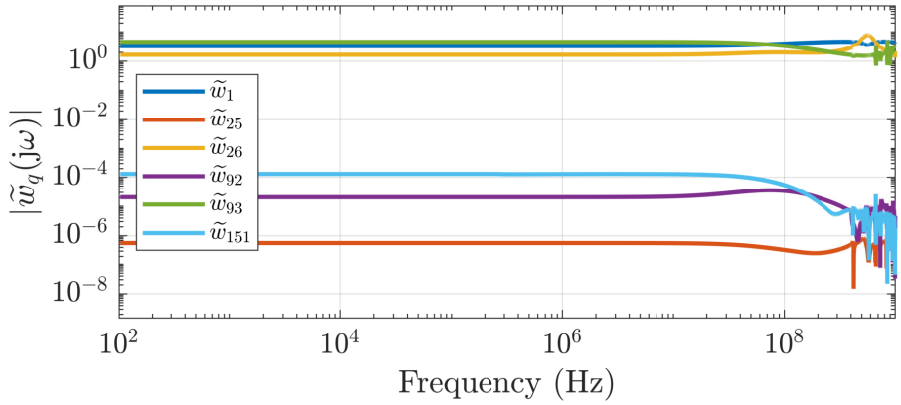

Fig. 3.1 Subset of BD-SVD basis functions rescaled via (3.39). Original bases are reported in Fig. 2.13. From [30] © 2022 IEEE.

To guarantee an exact model behavior at DC also for non-vanishing responses, the standard least square problem (3.39) is enriched with a constraint at $s_1 = \mathrm{j}\omega_1 = 0$ (the first available frequency point with index $k = 1$), defined as

$$(\widetilde{\boldsymbol{w}}_q)_1 = r_{q\infty} + \sum_{n=1}^{N_w} \frac{r_{qn}}{-p_n} \qquad q = 1, \ldots, \rho, \tag{3.40}$$

and suitably embedded in the identification step, i.e. in the VF iteration.

3.2.2 Compressed Model representations

In view of passivity enforcement, it is shown here how the compressed structured macromodel (3.31) can be mapped to the original $P \times P$ multiport scattering representation $\mathbf{H}(s)$ through a constant algebraic transformation. First, matrix

$\boldsymbol{\Psi} \in \mathbb{R}^{P\times \rho P}$ is defined as in [34]

$$\bar{\mathbf{V}} = \begin{bmatrix} \bar{\mathbf{V}}_1 \\ \vdots \\ \bar{\mathbf{V}}_P \end{bmatrix}, \quad \boldsymbol{\Psi} = \begin{bmatrix} \bar{\mathbf{V}}_1 & \bar{\mathbf{V}}_2 & \cdots & \bar{\mathbf{V}}_P \end{bmatrix} \tag{3.41}$$

where each $\bar{\mathbf{V}}_j \in \mathbb{R}^{P\times\rho}$ is obtained rearranging the blocks of rows in $\bar{\mathbf{V}}$ indexed by $\{j(P-1)+1,\dots,jP\}$. It can be demonstrated that the equivalent representations hold

$$\mathbf{H}(s) = \sum_{n=1}^{\bar{n}} \frac{\mathbf{R}_n}{s-p_n} + \mathbf{R}_0 = \mathrm{mat}(\bar{\mathbf{V}}\widehat{\boldsymbol{w}}(s)^{\mathsf{T}}) = \boldsymbol{\Psi}(\mathbf{I}_P \otimes \widehat{\boldsymbol{w}}(s)^{\mathsf{T}}) = \sum_{q=1}^{\rho} \mathbf{K}_q \widehat{w}_q(s) \tag{3.42}$$

where the mat operator reshapes a vector of P^2 elements into a $P\times P$ matrix, and where $\mathbf{K}_q \in \mathbb{R}^{P\times P}$ is a set of matrices defined as

$$\mathbf{K}_q = \begin{bmatrix} \bar{\boldsymbol{v}}_q^1 & \bar{\boldsymbol{v}}_q^2 & \cdots & \bar{\boldsymbol{v}}_q^P \end{bmatrix} \tag{3.43}$$

with vector $\bar{\boldsymbol{v}}_q^j$ selecting the q-th column of $\bar{\mathbf{V}}_j$. It is important to notice that each $\mathbf{K}_q$ inherits the same sparsity pattern of $\bar{\mathbf{V}}$ in (2.34). Finally, the model response at a given frequency can be denoted as $\mathbf{H}_k = \mathbf{H}(\mathrm{j}\omega_k)$, and let $\boldsymbol{h}_k$ be a column vector of length P^2 stacking its columns. This *vectorized macromodel response* can be written as

$$\boldsymbol{h}_k = \mathrm{vec}(\mathbf{H}_k) = \bar{\mathbf{V}}\widehat{\boldsymbol{w}}(\mathrm{j}\omega_k)^{\mathsf{T}} = \bar{\mathbf{V}}\boldsymbol{\Gamma}_k \boldsymbol{w}(\mathrm{j}\omega_k)^{\mathsf{T}} = \bar{\mathbf{V}}\boldsymbol{\Gamma}_k(\mathbf{I}_\rho \otimes \boldsymbol{\varphi}_k^{\mathsf{T}})\,\mathrm{vec}(\mathbf{R}_w) \tag{3.44}$$

where $\mathbf{R}_w \in \mathbb{R}^{(N_w+1)\times\rho}$ stores all model coefficients r_{qn}, $r_{q\infty}$ in (3.15), and $\boldsymbol{\varphi}_k = \boldsymbol{\varphi}(\mathrm{j}\omega_k) \in \mathbb{C}^{N_w+1}$ stacks the partial fraction basis functions $\varphi_0(s_k) = 1$, $\varphi_n(\mathrm{j}\omega_k) = (\mathrm{j}\omega_k - p_n)^{-1}$. The last step exploited the known properties of the Kronecker product.[1]

[1]The Kronecker product can be used to rewrite a matrix equation like $\mathbf{AXB} = \mathbf{C}$ in a vectorized form as $(\mathbf{B}^{\mathsf{T}} \otimes \mathbf{A})\mathrm{vec}(\mathbf{X}) = \mathrm{vec}(\mathbf{AXB}) = \mathrm{vec}(\mathbf{C})$.

3.3 Passivity check via adaptive sampling

This Chapter aims at providing a complete framework for the generation of passive large-scale macromodels. In this context, the model passivity characterization plays a crucial role as a prerequisite for each iteration of a passivity enforcement loop. As the passive structure that is intended to replicate, the compact black-box model must not be able to generate energy. This simple idea is verified in this section through a multi-stage adaptive sampling scheme [45].

This section illustrates the passivity verification scheme originally reported in [45], providing an efficient tool for checking passivity of large-scale macromodels with a large ports count.
The first step of the proposed methodology requires to further re-formulate the passivity condition of a scattering system (3.25) in terms of the passivity metric (3.26) *observations*, by stating the passivity verification problem as

Find all local maxima larger than a given threshold γ of a continuous univariate function of frequency $\varphi(\omega)$ over the entire frequency axis.

The solution of this (classic) optimization problem seems straightforward, however:

1. It is necessary to detect *all* local maxima of $\varphi(\omega)$, to verify whether ones of them exceed the passivity threshold γ. Unfortunately, $\varphi(\omega)$ commonly shows multiple maxima.

2. The passivity metric variation is characterized by *sharp* peaks, narrowed in a very small frequency band, usually close to the critical threshold. Hence, the exact location of all local maxima is hard to be detected via sampling, and it is equally difficult to verify with a finite number of samples whether a given point or its neighbors correspond to a passivity violation.

An example of these behaviors is reported in both Fig. 3.2 and Fig. 3.3, the latter depicting a detailed view in a zoomed frequency region.

These challenges motivate the main passivity-driven features of the presented multi-stage adaptive sampling scheme [45], which tunes its resolution

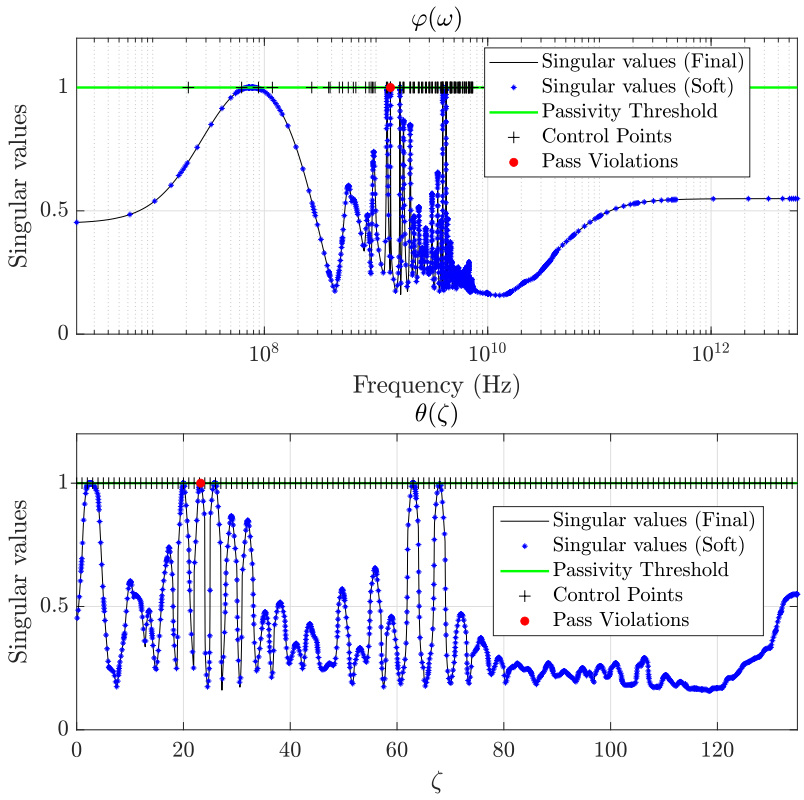

Fig. 3.2 Top panel: graphical illustration of the adaptive sampling-based passivity characterization. Bottom panel: rescaled passivity metric $\theta(\zeta)$ after adaptive frequency warping. From [45] © 2021 IEEE.

by monitoring fast variations of $\varphi(\omega)$ and its distance to the critical threshold γ, while limiting the overall number of required observations K.

The same structure provided by [45] is used in the following to illustrate the two-stage passivity verification scheme. Section 3.3.1 provides the initial partitioning approach that leads to the definition of (potentially many) frequency subbands, aimed at rescaling the passivity metric into a smoother equivalent function through a warping procedure. Section 3.3.2 illustrates a hierarchical refinement scheme based on a tree-search approach to identify all possible peaks in each (independent) subband. The result of these two steps is a reliable and efficient passivity characterization scheme, presented in Sec. 3.3.3, that is able to overcome most of the limitations of standard methods presented in Section 3.1.4. Numerical results supporting this statement will be provided both in Section 3.5.1 and through the following algorithm description.

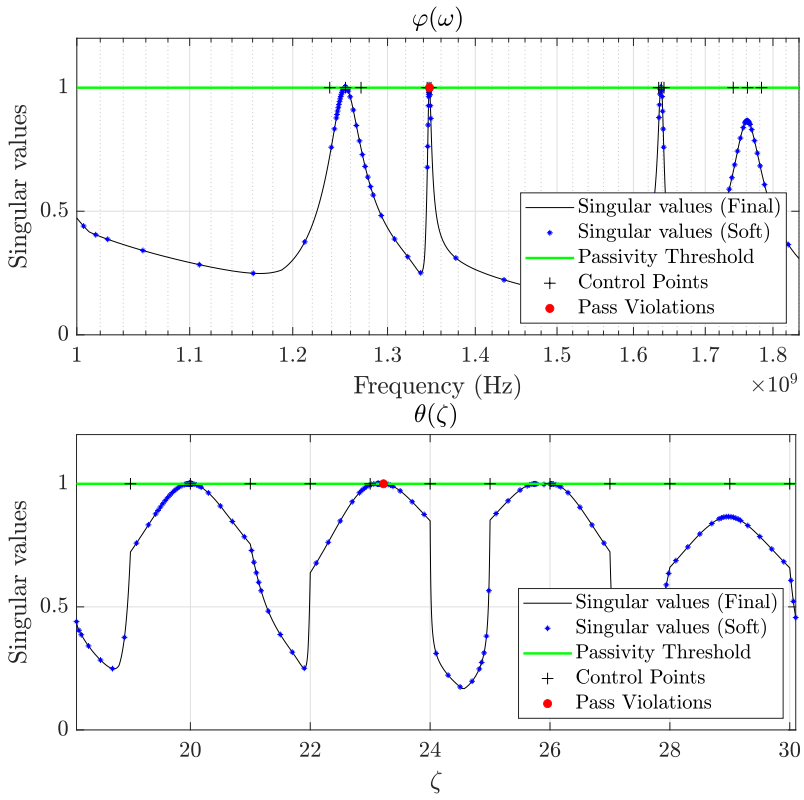

Fig. 3.3 Enlarged view of Figure 3.2. From [45] © 2021 IEEE.

3.3.1 Stage 1: adaptive frequency warping via poles location

The first step of the procedure defines an invertible nonlinear transformation

$$\theta(\zeta) = \varphi(\mathcal{W}^{-1}(\zeta)) \tag{3.45}$$

inducing the change of variable $\zeta = \mathcal{W}(\omega)$. The goal of (3.45) is to obtain a *smoothed* version of the passivity metric $\varphi(\omega)$, by "flattening" sharp peaks so that the resulting $\theta(\zeta)$ exhibits a uniform variation in the band of interest. This condition is reached by means of a suitable *warping function* $\mathcal{W}$, obtained as follows. First, a set of control points $\Omega = \{\hat{\omega}_\ell,\ \ell = 0, \dots, L\}$ are defined as

$$0 = \hat{\omega}_0 < \hat{\omega}_1 < \cdots < \hat{\omega}_\ell < \hat{\omega}_{\ell+1} < \cdots < \hat{\omega}_L = \infty \tag{3.46}$$

The nonlinear change of variable $\mathcal{W}$ is constructed assembling piecewise linear mapping of each subband bounded by two control points $\omega \in [\hat{\omega}_\ell, \hat{\omega}_{\ell+1}]$ to the corresponding normalized interval $\zeta \in [\ell, \ell+1]$. This condition formally reads

$$\zeta = \ell + \frac{\omega - \hat{\omega}_\ell}{\Delta_\ell} \quad \forall \omega \in [\hat{\omega}_\ell, \hat{\omega}_{\ell+1}], \quad \ell = 0, \ldots, L-2 \tag{3.47}$$

with $\Delta_\ell = \hat{\omega}_{\ell+1} - \hat{\omega}_\ell$ length of the ℓ-th subband. To deal with the asymptotic (infinite) frequency, the last subband is defined as

$$\zeta = \ell + \frac{\omega - \hat{\omega}_\ell}{\omega} \quad \forall \omega \in [\hat{\omega}_\ell, \hat{\omega}_{\ell+1}], \quad \ell = L-1 \tag{3.48}$$

The effect of the *warping function* $\mathcal{W}$ is to map the (infinite) frequency axis $\omega \in [0, +\infty)$ into the normalized interval $[0, L]$, producing the result of Fig. 3.2 and Fig. 3.3. Both figures show how nonuniform controls points (indicated with crosses) in the original frequency axis (top panels) are mapped to an equivalent uniform distribution into the normalized interval (bottom), while retaining the same singular value information. The *warping effect* of the set Ω is provided by the concentration and location of control points, which must be selected to enclose sharp peaks of the passivity metric $\varphi(\omega)$, such that in the (normalized) domain ζ all tight spikes are enlarged while broad maxima are tightened. The result is that the warped passivity metric $\theta(\zeta)$ show local minima in each (normalized) subband with a similar extension, while fast variations and maxima around the passivity threshold are practically eliminated. The optimal situation is characterized by single (or a very limited number of) peaks in each interval, mostly located at one of the subband edge. This condition is reached in Fig. 3.2 with a proper choice of control points, detailed next.

Control points selection

The control points selection procedure [45] is based on the assumption that variations in the passivity metric $\varphi(\omega)$, i.e. in the singular values trajectories, of rational model (4.8) are ultimately induced by poles/residues location, which turns out to be responsible for variations of the transfer matrix responses.

Notice that the final goal is to tune the size of each interval Δ_ℓ, choosing control points $\hat{\omega}_\ell$ such that

- a tight subband (small Δ_ℓ) encloses sharp changes in $\varphi(\omega)$, so that the warping effect is able to flatten (enlarge) these maxima in the normalized axis ζ
- a wide interval (large Δ_ℓ) captures gradual (in the limit case nearly constant) variations of $\varphi(\omega)$, inducing a relaxation of $\theta(\zeta)$.

Thus, a set of points is obtained by taking advantage of poles location, as originally presented in [66], defining

$$\omega_{n,r} = \beta_n + \alpha_n \tan\frac{r\pi}{2(R+1)} \quad r = -R, \cdots, R \tag{3.49}$$

for each pair of complex conjugate poles $p_n = \alpha_n \pm \mathrm{j}\beta_n$, where R is a control parameter necessary to tune the resulting number of samples for each pole. Points defined as in (3.49) include negative values, which are disregarded. Thus, the set of control points Ω includes all $\omega_{n,r} \geq 0$ sorted in ascending order. At last, set Ω is processed to avoid an excessive number of samples by setting a maximum resolution $\Delta\omega$ defined as in [45]

$$\Delta\omega = \frac{p_{\max}}{N\rho} \tag{3.50}$$

with $p_{\max} = \max(\omega_{\max}, \max_n |p_n|)$ and $\rho \gg 1$ additional control parameter.

Using the set of control points Ω defined blindly as in [45] leads to a poor preprocessing step for passivity verification purposes, considering that

$\mathcal{C}_{p1}$ an excessive number of subbands are created if the control parameters values are set as in [45]. This is a conservative choice that may compromise the efficiency of a sampling scheme, which ideally requires a small value of points per pole R. On the other hand, reducing this parameter may compromise the spacing of control points at the highest frequencies, in particular when $\omega > \omega_{\max}$.

$\mathcal{C}_{p2}$ sampling via (3.49) in case of real poles leads to a poor distribution of control points.

$\mathcal{C}_3$ the pruning procedure by means of (3.50) may remove poles with a small imaginary part, that show a tight interval Δ_ℓ.

A set of control parameters R_ν with $\nu = \{\text{cp}, \text{rp}, \text{hf}\}$ are defined to compensate for $\mathcal{C}_{p1,p2}$. These are selected such that:

- R_{cp}, for complex (conjugate) poles, and can be safely small, even $R_{\text{cp}} = 1$;
- R_{rp}, for real poles, so should be at least $R_{\text{rp}} \geq 2$;
- R_{hf}, for poles with magnitude of real/imaginary component strictly larger than $\omega_{\max}$, and should be even bigger, with at least $R_{\text{hf}} \geq 3$.

The third condition $\mathcal{C}_3$ is mitigated by artificially changing the real component of highly resonant poles in (3.49), so that

$$\alpha_n \leftarrow c \cdot \alpha_n \quad \text{if} \quad Q_n \approx \frac{|\beta_n|}{2\,|\alpha_n|} > Q_{\max} \tag{3.51}$$

with common settings $c = 50$ and $Q_{\max} = 500$.

Finally, a dedicated sample distribution is used for the last frequency band beyond the last training point $\omega_{\max}$. Indeed, the adaptive sampling must cover the all frequency axis, including $\omega = \infty$, so that this sample is included into Ω as well as other $\kappa + 1$ logarithmically spaced points defined as

$$\omega_\nu = \omega_{\max} \cdot 10^{d\frac{\nu}{\kappa}} \quad \nu = 0, \cdots, \kappa \tag{3.52}$$

with the purpose of extending to d decades after $\omega_{\max}$. The last two parameters (out of eight) of the warping procedure are normally set to $d \geq 0.5$ and $\kappa \geq 2$.

To summarize, the first stage of the presented passivity verification scheme includes a warping strategy based on a fixed set of control points, chosen on the basis of poles location. This procedure relies on eight control parameters, whose recommended values are reported in Table 3.3 after a reliability assessment through an extensive heuristic study on a large number of testcases, see Section 3.5.1. On the other hand, these parameters enable a strong algorithm flexibility when embedded in a passivity enforcement scheme. Indeed, the number of control points (i.e. subbands) can be suitably reduced to speed up early stages of the passivity enforcement loop, when violations are easily detectable with few subbands, and increased to guarantee that later iterations will not miss smaller peaks, which are closer to the passivity threshold and

hard to detect. This condition is reached changing R_ν, as demonstrated with an example in Sec. 3.3.3.

3.3.2 Stage 2: A Modified Naive Multi-scale Search Optimization

The second stage of the presented scheme provides an adaptive sampling algorithm to possibly identify all local maxima of $\theta(\zeta)$ above the passivity threshold $\gamma = 1$, equipped with a hierarchical strategy to be applied on each normalized interval $\zeta \in \mathcal{X}_\ell = [\ell, \ell + 1]$.

Among the several sparse sampling techniques, the presented passivity verification builds on a tree-based algorithm that follows a divide-and-conquer approach equipped with restarting strategies, namely the Naive Multi-scale Search Optimization (NMSO) [83]. Notice that the presented problem – *find a (global) maxima of a monovariate objective function as* $\theta(\zeta)$ *with a finite number of samples*– lies in the class of well-known optimization problems extensively addressed by the literature [84–88]. Nevertheless, as most of the alternatives already available, also the effectiveness of the NMSO may be compromised by the presence of multiple local maxima in the search domain. This motivates the first stage of the presented procedure, the warping and frequency splitting approach of Section 3.3.1, whose objective is to minimize the number of peaks included in each subband $\mathcal{X}_\ell = [\zeta_\ell, \zeta_{\ell+1}]$. Thus, the second stage of the passivity verification applies on each individual $\mathcal{X}_\ell$ the (modified) NMSO algorithm, which is usually able to perform at its best thanks to the provided working conditions (mostly only one maximum per subband).

The following framework simplifies the notation by referring to the search space as the normalized interval $\mathcal{X} = \mathcal{X}_0 = [0, 1]$.

The presented algorithm expands nodes of a tree $\mathcal{T}$ at a given level $h \geq 0$ every iteration, with the goal of obtaining a sparse segmentation of the search space $\mathcal{X}$ into a number of M^h cells defined as

$$\mathcal{X}_{h,i} = [iM^{-h}, (i+1)M^{-h}], \quad i = 0, \ldots, M^h - 1. \tag{3.53}$$

where h indicates the refinement *level* (or *scale*). The center of every cell, here indicated as *leaf* or *node*, is analytically defined as $\zeta_{h,i} = (i + \frac{1}{2})M^{-h}$. Each expansion sequence increases the leaves number of a factor $M \geq 2$, by adaptively creating new children at a deeper scale. Our implementation is based on an odd expansion factor M.

As already stated, the final goal is to find all maxima of $\theta(\zeta)$ in $\mathcal{X}$ with the minimum number of evaluations $\theta_{h,i} = \theta(\zeta_{h,i})$ of the objective function. This condition is achieved by means of the adaptive hierarchical refinement strategy illustrated in Fig. 3.4. Notice that, once a prescribed scale h is defined, two set of leaves can be identified:

$\mathcal{L}^h$ the set of all possible nodes of the complete tree, meaning the maximum number of evaluations at the scale h, indicated by (all) dots of Fig. 3.4

$\mathcal{E}^h \subseteq \mathcal{L}^h$ the (sub)set of evaluated leaves, corresponding to actual samples of the objective function, the colored dots of Fig. 3.4.

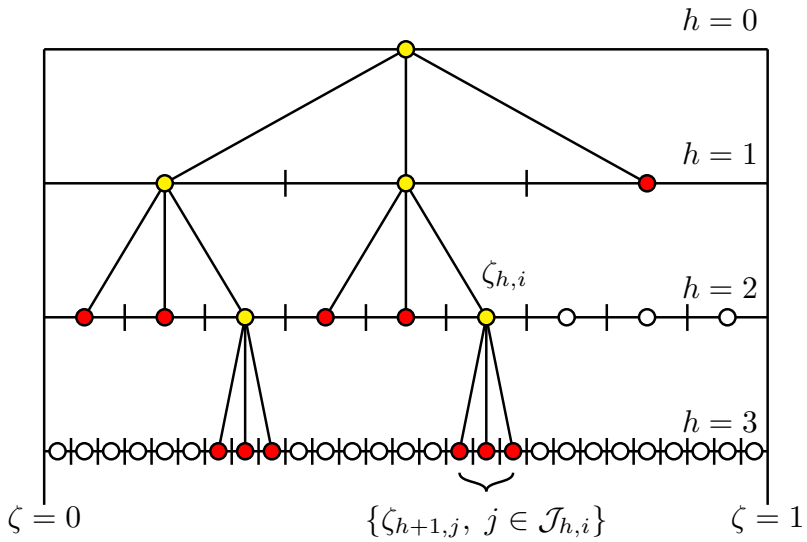

Fig. 3.4 Representation of a M-tree, with odd $M = 3$, up to the level $h = 3$. Leaves are indicated with dots, colored ones showing the set of evaluated nodes $\mathcal{E}^h$. Yellow dots denote expanded leaves, that cannot be part of a new refinement, while red dots indicate candidates for future refinements (elements of $\mathcal{C}$). From [45] © 2021 IEEE.

3.3.3 The adaptive sampling scheme

Algorithm 3 reports a summary of the procedure, which is detailed through this section. The tree expansion process follows a global iteration index μ. Observing the set of computed leaves $\mathcal{E}_\mu$ at a given iteration μ, three scales of the tree $\mathcal{T}$ can be identified

- $h_{\min}$: the highest level among all leaves in $\mathcal{E}_\mu$
- $h_{\max}$: the deepest level among all leaves in $\mathcal{E}_\mu$
- h (*current level*): the level of the leaf under investigation, the latter indicated as *current node*.

Furthermore, the set of evaluated leaves $\mathcal{E}_\mu$ can be partitioned as

$$\mathcal{E}_\mu = \mathcal{C}_\mu \cup \mathcal{B}_\mu \tag{3.54}$$

where

$\mathcal{C}_\mu$ gathers all (candidates) nodes that can be exploited for an expansion, e.g. the yellow dots of Fig. 3.4;

$\mathcal{B}_\mu$ (the *basket*) collects leaves that cannot be a candidate for refinement since fulfilling a stopping criterion (see next).

Implementation assumptions

With the goal of keeping under control the overall count of samples, an (initial) maximum number of function evaluations (samples) is set through the *budget* n^e. Depending on tree evolution, this budget may be updated dynamically through iterations, see later (3.72). Therefore, the algorithm is parameterized by a sequence of tentative budgets collected in vector $\boldsymbol{n^e}$. If the initial budget needs to be increased, the next element in vector $\boldsymbol{n^e}$ will be used as current budget. This procedure will be formally detailed later in a suitable paragraph. We recall that our implementation uses an odd expansion factor M.

Algorithm initialization

The algorithm starts by initializing the iteration index $\mu = 0$ and computes all tree leaves up to a given initial refinement level $h_0 \geq 0$. This requires to

Algorithm 3 Find $\theta_{\max} = \arg\max \theta(\zeta)$ for $\zeta \in [0,1]$

Require: h_0, M, $\delta\zeta$, $\delta\theta$, $\delta\eta$, ε, ϱ, n^e, δn^e
1: Initialize $\mu = 0$, $h \leftarrow h_0$, $\mathcal{B}_0 = \emptyset$
2: Initialize $\mathcal{C}_0 = \{\zeta_{h,i},\ i = 0, .., M^h - 1\}$, $K = M^h$
3: Initialize $\theta_{\max} = -\infty$
4: **while** $\mathcal{C}_\mu \neq \emptyset$ **do**
5: Choose current point $\zeta_{h,\imath_*}$ via (3.58)
6: Expand current point via (3.57)
7: $K \leftarrow K + M - 1$
8: Update $\theta_{\max} \leftarrow \max\{\theta_{\max}, \{\zeta_{h+1,j},\ j \in \mathcal{J}_{h,\imath_*}\}\}$
9: **if** $K > n^e$ **then**
10: **if** $\mathcal{U}_1$ `AND` $(\mathcal{U}_2$ `OR` $\mathcal{U}_3)$ **then**
11: update budget via (3.72)
12: update $\varepsilon \leftarrow \varrho\,\varepsilon$
13: **else**
14: go to 27
15: **end if**
16: **end if**
17: **if** $\mathcal{S}_1$ `OR` $\mathcal{S}_2$ `OR` $\mathcal{S}_3$ **then**
18: insert new points in basket via (3.67)
19: $h \leftarrow h_{\min}$
20: reset ε
21: **else**
22: flag new points for refinement via (3.65)
23: $h \leftarrow h + 1$
24: **end if**
25: $\mu \leftarrow \mu + 1$
26: **end while**
27: **return** samples $\mathcal{E}_\mu = \mathcal{C}_\mu \cup \mathcal{B}_\mu$ and maximum $\theta_{\max}$

compute samples of the target function θ (the passivity metric after rescaling) at all points in the set $\{\zeta_{h_0,i},\ i = 0, \ldots, M^h - 1\}$. These new nodes are then gathered in the set $\mathcal{E}_0 = \mathcal{C}_0$, while the basket is initialized as $\mathcal{B}_0 = \emptyset$.

Leaves refinement (expansion process)

The refinement of a cell $\mathcal{X}_{h,i}$ into M leaves (children) at scale $h+1$ is defined as

$$\mathcal{X}_{h+1,j}, \quad j \in \mathcal{J}_{h,i} = \{Mi, \ldots, M(i+1) - 1\} \tag{3.55}$$

where indexes of leaves inside the *cone of influence* of leaf $\zeta_{h,i}$ are collected in $\mathcal{J}_{h,i}$. Notice that the implementation assumption of an odd partition factor M enables to re-use the information of the node under refinement, since it coincides with the middle point of its children

$$\zeta_{h+1,Mi+\lfloor M/2 \rfloor} = \zeta_{h,i}, \tag{3.56}$$

thus avoiding one evaluation. A new labeling of the node can be simply applied as in the illustrative example of Fig. 3.4. Consider also that relabeling $\zeta_{h,i}$ (after its expansion) can lead to an update of the minimum level $h_{\min}$ among candidates for the next refinement. Thus, the refinement of a node can be formally expressed as

$$\mathcal{E}_{\mu+1} = (\mathcal{E}_{\mu} - \{\zeta_{h,i}\}) \cup \{\zeta_{h+1,j} : \ j \in \mathcal{J}_{h,i}\} \tag{3.57}$$

Best candidate for expansion

The main step of a tree-based adaptive sampling procedure is the refinement of a leaf (3.57), which requires being able of selecting the best candidate from the set $\mathcal{C}_{\mu}$. With the goal of identifying all local maxima, this corresponds to the node defined as

$$\zeta_{h,i_*} = \arg\max\{\theta(\zeta_{h,i}) : \ \zeta_{h,i} \in \mathcal{C}_{\mu}\} \tag{3.58}$$

readily the leaf at the h-th level with the biggest target function value. In other words, our framework is based on the simple strategy of expanding nodes with the largest passivity metric among all available points, following the same principle that rules any optimistic optimization algorithm [85, 86] as the NMSO.

Restarting criteria and nodes grouping

The algorithm presented in [45] proposes several stopping conditions to restart a sequence of expansion, i.e. to decide if leaves created during the last refinement (3.57) do not require being further exploited. This choice is based on a set of control parameters such that three main restart criteria are checked:

$\mathcal{S}_1$: a maximum resolution is achieved at the refinement scale $h+1$ when

$$M^{-h-1} < \delta\zeta \tag{3.59}$$

with $\delta\zeta$ control parameter.

$\mathcal{S}_2$: the maximum change of the target function within neighbor children (leaves expanded at level $h+1$) is below a prescribed threshold

$$\Delta_{h,i} < \delta\theta \tag{3.60}$$

defining

$$\Delta_{h,i} = \max\{|\theta_{h+1,j+1} - \theta_{h+1,j}|,\ j = Mi, \ldots, M(i+1)-2\} \tag{3.61}$$

with $\delta\theta$ control parameter.

$\mathcal{S}_3$: the maximum change of the target function within neighbor children is smaller than their distance from the passivity threshold

$$\Delta_{h,i} < |\widehat{\theta}_{h,i} - \gamma|, \tag{3.62}$$

with

$$\widehat{\theta}_{h,i} = \max\{\theta(\zeta_{h+1,j}),\ j \in \mathcal{J}_{h,i}\} \tag{3.63}$$

This last criterion (3.62) is checked only after a minimum resolution has been reached, so that

$$M^{-h-1} < \delta\eta \tag{3.64}$$

with $\delta\eta$ control parameter.

These three conditions enable a classification procedure on the resulting nodes from an expansion process, assigning leaves to either $\mathcal{C}_{\mu+1}$ or $\mathcal{B}_{\mu+1}$.

In particular, if all criteria $\mathcal{S}_1$, $\mathcal{S}_2$, $\mathcal{S}_3$ are not verified (false) all leaves $\zeta_{h+1,j}$ obtained from the latest expansion are classified as potentially critical and are assigned to the set of candidates for refinements as

$$\mathcal{C}_{\mu+1} = (\mathcal{C}_\mu - \{\zeta_{h,i}\}) \cup \{\zeta_{h+1,j},\ j \in \mathcal{J}_{h,i}\}\,. \tag{3.65}$$

Notice that one of these $\zeta_{h+1,j}$ is most likely to be refined at the next iteration (if the budget is not exceeded), while the rest are still kept in consideration for possible future expansions. This step increases the tree depth with an increment of the current level as

$$h \leftarrow h + 1\,. \tag{3.66}$$

In the other case, even if only one of the conditions $\mathcal{S}_1$, $\mathcal{S}_2$, $\mathcal{S}_3$ is verified all leaves $\zeta_{h+1,j}$ are classified as non-critical and are assigned to the basket

$$\mathcal{B}_{\mu+1} = (\mathcal{B}_\mu - \{\zeta_{h,i}\}) \cup \{\zeta_{h+1,j},\ j \in \mathcal{J}_{h,i}\}. \tag{3.67}$$

This situation corresponds to a *restart* of the exploration path so that the current level is reset to

$$h \leftarrow h_{\min} \tag{3.68}$$

in order to improve the exploration feature of the overall sampling strategy by neglecting further refinements of nodes $\zeta_{h+1,j}$.

Budget improvement

When the maximum number of samples, the *budget*, is reached the adaptive sampling process is terminated. Nevertheless, it may be the case that the algorithm stops either while exploiting a critical node or before exploring a portion of the space never expanded. To avoid these situations, a budget improvement (equivalently denoted as *increase* or *update*) is triggered by three criteria

$\mathcal{U}_1$: all children from the latest expansion are passive

$$\widehat{\theta}_{h,i} < \gamma \tag{3.69}$$

$\mathcal{U}_2$: one of the children is within a relative distance ε to the passivity threshold

$$\frac{\gamma - \widehat{\theta}_{h,i}}{\widehat{\theta}_{h,i}} < \varepsilon \tag{3.70}$$

with ε control parameter.

$\mathcal{U}_3$: the maximum change of the target function (3.63) within neighbor children is larger than their distance to the passivity threshold

$$|\gamma - \widehat{\theta}_{h,i}| < \Delta_{h,i} \tag{3.71}$$

Equivalently to $\mathcal{S}_3$, this criterion is checked only when a target resolution is exceeded (3.64).

A budget update is triggered by the verification of $\mathcal{U}_1$ and at least one between $\mathcal{U}_2$ and $\mathcal{U}_3$. Notice that this situation enables to avoid further refinements when a passivity violation is already detected, boosting performances in view of a passivity enforcement scheme as in Section 3.4. The budget improvement is in general defined as

$$n^e \leftarrow n^e + \delta n^e \tag{3.72}$$

with δn^e a control parameter that can varies with the iteration count.

Notice that criterion $\mathcal{U}_2$ may be become ineffective (since always true) while increasing the refinement level h when the peak under investigation is very close to the passivity threshold. To avoid this situation, any budget improvement (3.72) is associated to a tightening of ε in (3.70) by defining the gain $\varrho < 1$ as $\varepsilon \leftarrow \varrho\varepsilon$, while every *restart* of the expansion sequence restores the original value of ε.

Recycling basket elements

The algorithm proposed in [45] differs from the original NMSO also for the use of the basket $\mathcal{B}_\mu$. Indeed, the passivity-driven sampling always precludes the use of nodes in the basket to start new expansion sequence when all possible candidates have been exploited and $\mathcal{C}_\mu = \emptyset$. This implementation choice is performed in favor of runtime at the price of a reduced exploration capability. On the other hand, simple modifications apply to restore the original NMSO feature, the so-called *basket reuse*, that relocates elements of $\mathcal{B}_\mu$ to the set of possible candidates for refinement $\mathcal{C}_{\mu+1}$ when needed. The recommendation of [45] is to apply this strategy only if a very rigorous verification is needed, as a final step of a very meticulous passivity enforcement procedure.

Table 3.3 Stage 1: frequency warping. Algorithm parameters as in [45] © 2021 IEEE.

Step 1: Frequency Warping									
Mode	ρ	R_{cp}	R_{rp}	R_{hf}	c	$Q_{\max}$	κ	d	Emphasis
soft	10^3	1	2	5	50	500	3	0.5	*speed*
hard	∞	3	3	6	50	500	3	0.5	*accuracy*
final	∞	3	3	6	50	500	3	0.5	*qualification*

Table 3.4 Stage 2: adaptive sampling. Algorithm parameters as in [45] © 2021 IEEE.

Step 2: Modified NMSO							
Mode	M	$\delta\zeta$, $\delta\theta$	$\delta\eta$	ε	ϱ	n^e	Emphasis
soft	5	10^{-8}	10^{-3}	10^{-3}	0.1	$7, 10, 20, \ldots, 100$	*speed*
hard	5	10^{-8}	10^{-2}	10^{-3}	0.1	$10, 20, \ldots, 100$	*accuracy*
final	3	10^{-8}	10^{-3}	10^{-4}	0.1	$50, 100, \ldots, 250$	*qualification*

Three parameters combinations are recommended to improve speed (*soft*), accuracy (*hard*) and to stress model qualification (*final*).

Parameters settings

A summary of the two set of control parameters necessary for the stage 1 and 2 of the presented scheme are reported in Table 3.3 and 3.4, respectively. Both of them illustrate three sets of recommended settings from [45], depending on the required level of accuracy while performing the passivity verification. Three operating modes are suggested:

1. *soft*, suitable for initial passivity enforcement iterations when passivity violations are easily detectable and the sampling accuracy is not an issue
2. *hard*, to be used when the passivity metric is expected to be very close to the critical threshold γ, since the model is either very accurate or it is the result of a possible incomplete passivity enforcement scheme
3. *final*, for the model characterization after completing passivity enforcement, as a final passivity certification stage.

Figure 3.5 shows the result of both the *soft* and *final* mode on a significant example, the benchmark #444, selected among the full database provided in Sec. 3.5.1. Starting with rational model of a $P = 92$ ports system with $\bar{n} = 16$ poles (and a resulting number of states $N = 1472$), the presented adaptive

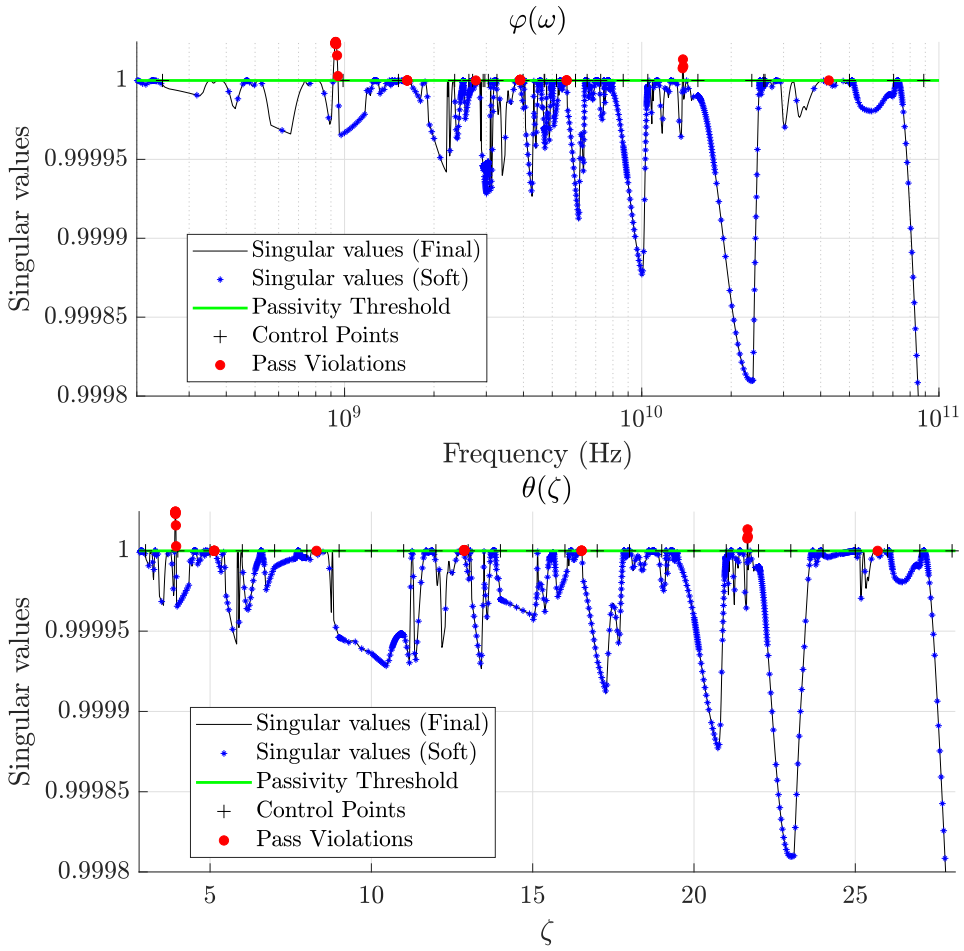

Fig. 3.5 Passivity characterization via adaptive sampling: benchmark #444. While the top panel shows the passivity metric $\varphi(\omega)$, the bottom panel reports the rescaling effect introduced by the Stage 1 of the procedure. From [45] © 2021 IEEE.

sampling scheme with the *soft* parameters configurations required $K = 2255$ function evaluations (samples) to identify 28 points with $\varphi(\omega) > \gamma$, by splitting the frequency axis with $L = 30$ control points. The other two modes, *hard* and *final*, provided respectively $K = 4245$ samples, of which 47 over the threshold, and $K = 9827$ evaluations, with 271 violations. Both cases take advantage of the same number of control points $L = 55$.

Considering that the worst-case cost in terms of operations of the proposed algorithm is $O(K_{\max}(P^3 + P^2\bar{n}))$, with $K_{\max} = (L-1) \cdot \max\{n^e\}$, all results confirm an efficient adaptive sampling capability also for this critical example. Indeed, both the *soft* and *hard* modes required $\approx 78\%$ of the corresponding

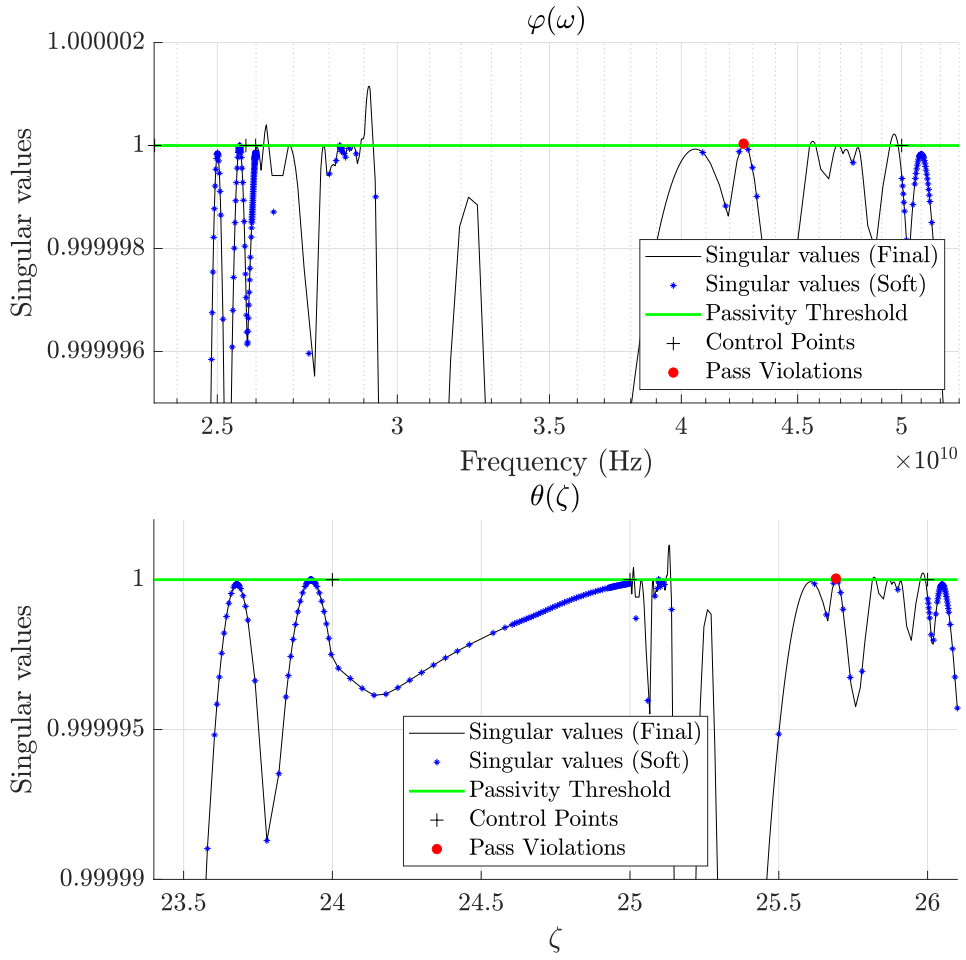

Fig. 3.6 Detailed view of Fig. 3.5 on a band with critical passivity violations, close to the passivity threshold, and hard to be identified with relaxed parameters settings (*soft* mode). From [45] © 2021 IEEE.

$K_{\max}$, while the verification with the *final* setting stopped with $\approx 73\%$ of the upper limit cost.

Notice that local maxima identified by the presented algorithm in each subband $\mathcal{X}_\ell$ are not necessarily maxima of the passivity metric $\varphi(\omega)$ over the frequency axis. Thus, a suitable postprocessing phase is needed to remove among all identified points those that do not correspond to an actual peak of the passivity metric. The final adaptive sampling result is the identification of 7, 14, and 21 local maxima after *soft*, *hard*, and *final* mode, respectively.

Notice that these outcomes are consistent with running the adaptive sampling in *soft* configuration, allowing an infinite number of function evaluations by setting $n^e = \infty$ and enabling the basket reuse.

3.4 Enforcing structured model passivity

The macromodel of a passive structure must not generate energy of its own. Therefore, the model passivity must be verified and enforced through a suitable iteration process. The focus of this section is the passivity enforcement of the structured compressed macromodel (3.42), assuming that an appropriate passivity characterization is available, as the one presented in Section 3.3. The asymptotic model stability is a requirement that is not questioned here, since provided by construction from standard VF. Only the scattering multiport representation will be considered in the following, but the generalization of this procedure for other representations is straightforward [20].
For the moment, it is sufficient to see that a passive scattering model requires a bounded real transfer function. Since model asymptotic stability and realness of coefficients matrices is guaranteed by construction, this condition is implied by

$$\sigma_{\max}\{\mathbf{H}(\mathrm{j}\omega)\} \leq 1 \qquad \forall \omega \in \mathbb{R}, \tag{3.73}$$

where $\sigma_{\max}$ indicates the largest singular value of the system response $\mathbf{H}(\mathrm{j}\omega)$.

The literature offers several approaches [55, 64, 57, 80] to enforce (3.73). The iterative scheme adopted here is based on the perturbation of local singular values where the passivity condition (3.73) is violated. This procedure, that iteratively remove these passivity violations by perturbing model coefficients, is well-established in most of state-of-the-art modeling tools [33]. Considering a single passive violation that occurs at frequency ω_k, the singular value decomposition of the model response $\mathbf{H}(\mathrm{j}\omega_k) = \mathbf{H}_k \in \mathbb{C}^{P\times P}$ reads

$$\mathbf{H}_k = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^{\mathsf{H}} \tag{3.74}$$

where $\mathbf{U}_k$,$\mathbf{V}_k$ are (unitary orthogonal) singular vector matrices, and $\mathbf{\Sigma}_k$ collects the P singular values. To find the passivity violations, i.e. all frequencies ω_k, we use the adaptive sampling scheme [45] provided in Section 3.3.

A standard procedure is to write a first-order perturbation of each singular value that violates (3.73) as

$$\widehat{\underline{\sigma}}_k \approx \underline{\sigma}_k + \Delta\underline{\sigma}_k = \underline{\sigma}_k + \mathrm{Re}\left\{\boldsymbol{u}_k^{\mathsf{H}} \Delta\mathbf{H}_k \boldsymbol{v}_k\right\} \tag{3.75}$$

where $\underline{\sigma}_k > 1$ is a critical singular value (hence a passivity violation), $\boldsymbol{u}_k$ and $\boldsymbol{v}_k$ are the associated left and right singular vectors, while $\Delta\mathbf{H}_k$ is the resulting perturbation of the model response at the frequency ω_k. To guarantee local passivity at this frequency, condition (3.73) requires that $\widehat{\underline{\sigma}}_k < 1$ and leads to

$$\mathrm{Re}\left\{\boldsymbol{u}_k^{\mathsf{H}} \Delta\mathbf{H}_k \boldsymbol{v}_k\right\} \leq 1 - \underline{\sigma}_k \,. \tag{3.76}$$

The condition (3.76) can be used to enforce passivity at all frequencies, including the asymptotic point at $s = \infty$, by perturbing all model coefficients concurrently. Note that in the compressed framework the perturbation required to obtain a passive direct coupling is not straightforward: all strategies suggested by [34] are either sub-optimal or based on the definition of a Linear Matrix Inequality (LMI), hence computationally very expensive (see Section 3.3). On the contrary, the approach proposed in [30] perturbs all coefficients concurrently while enforcing (3.76) at all frequencies, including $s = \infty$. The direct coupling terms $r_{q\infty}$ are then locked only when the asymptotic passivity is reached, in order to decrease the number of free variables with a (minor) improvement in terms of runtime and a smaller complexity.

The first step of the passivity enforcement relates model coefficients with the (vectorized) induced perturbation as

$$\boldsymbol{x}_w = \mathrm{vec}(\Delta\mathbf{R}_w) \tag{3.77}$$

where $\Delta\mathbf{R}_w$ is the model coefficients perturbation, gathered as in (3.44) by collecting also direct coupling terms. Using (3.44), the local singular value perturbation can be defined as

$$\begin{aligned} \mathrm{Re}\left\{\boldsymbol{u}_k^{\mathsf{H}} \Delta\mathbf{H}_k \boldsymbol{v}_k\right\} &= \mathrm{Re}\left\{(\boldsymbol{v}_k^{\mathsf{T}} \otimes \boldsymbol{u}_k^{\mathsf{H}})\mathrm{vec}(\Delta\mathbf{H}_k)\right\} \\ &= \underbrace{\mathrm{Re}\left\{(\boldsymbol{v}_k^{\mathsf{T}} \otimes \boldsymbol{u}_k^{\mathsf{H}})\bar{\mathbf{V}}\,\boldsymbol{\Gamma}_k(\mathbf{I}_\rho \otimes \boldsymbol{\varphi}_k^{\mathsf{T}})\right\}}_{\boldsymbol{p}^{\mathsf{T}}} \boldsymbol{x}_w \end{aligned} \tag{3.78}$$

showing that the local passivity constraints (3.76) can be written in terms of model coefficients as

$$\boldsymbol{p}_k^{\mathsf{T}} \boldsymbol{x}_w \leq 1 - \underline{\sigma}_k \,. \tag{3.79}$$

Notice that in practical situation (3.79) is cast in a matrix form since aggregate constraints are needed to compensate for multiple passivity violations.

The constrained least-squares problem iteratively solved by the passivity enforcement aims at minimizing the cost function

$$\mathcal{E}^2 = \sum_{q=1}^{\rho} \mathcal{E}_q^2 = \sum_{q=1}^{\rho} \sum_{k=1}^{K} |\Delta \widehat{w}_q(\mathrm{j}\omega_k)|^2 \tag{3.80}$$

subject to a set of inequalities constraints (3.79) and, in our case, to equality constraints (3.40) to enforce the prescribed level of accuracy at DC.

Notice that (3.80) is quadratic in the free variables $\boldsymbol{x}_w$, which are only the $\rho \cdot (N_w + 1)$ instead of the $P^2 \cdot (N_w + 1)$ associated to the full model (3.42). Indeed, working with the compressed macromodel coefficients allows for a drastic reduction in decision variables while equivalently enforcing the passivity of the full-size $P \times P$ system. Thus, the complexity of the passivity enforcement scheme is independent from the number of system electrical ports P and it is only associated with the dimension of the basis functions set ρ used for the data compression. Furthermore, the low-frequency structure introduced by means of the filtering term $\boldsymbol{\Gamma}(s)$ does not introduce extra unknowns into the process, since poles p_ν^∞ are constant, and this matrix plays the role of a frequency weight in the constraint definition (3.78).

3.5 Results on passivity and modeling

We now report a set of results to demonstrate the efficiency and scalability of the proposed macromodeling framework. First, Section 3.5.1 presents numerical results to estimate the performance of the passivity verification scheme of Section 3.3. Then, Section 3.5.2 demonstrates the scalability of the passive macromodeling procedure on shielding enclosure structures with increasing complexity, up to 1024 ports. All numerical examples are obtained with the

same Workstation of Section 2.2.3, based on an i9-7900X CPU running at 3.3 GHz with 64 GB RAM.

3.5.1 The passivity verification scheme

The sampling-based passivity verification presented in Section 3.3 is here validated on numerous benchmarks, in order to assess both reliability and efficiency of the algorithm. Results originally documented in [45] are reported here with the intent of providing a complete overview of the scheme. At first, a comprehensive result on a large number of benchmarks obtained from a vast database of macromodels is reported. At last, two significant large-scale examples are presented, namely a 400-port shielding enclosure and a 640-port via array (the latter to show that proposed algorithm is general purpose and applicable to any rational macromodel).

Quality check: assessing consistency and efficiency

This section provides results of the passivity check scheme on a total of 447 models with different complexity, both in terms of ports and dynamic order. These benchmarks are obtained from a database of 243 different structures, processed by means of a state-of-the-art commercial macromodeling tool (IdEM [33]) to obtain a standard VF rational model as in (3.1). After the fitting step, 206 of these models were non-passive, depicting a variety of passivity metric behaviors exceeding the passivity threshold γ. Since the violation extent of some models was sufficiently large to be identified with a trivial search, we postprocessed 204 of these non-passive models with few iterations of a passivity enforcement loop. Since the goal of the presented algorithm is to adaptively sample the passivity metric, which is strongly modified during each enforcement iteration, these (almost passive) models can be regarded as an independent set of test cases. Furthermore, since at each enforcement iteration the leading violations are removed or contracted toward the passivity threshold, the latest set of models show a most challenging framework due to the presence of tight violations that are hardly detectable with a poor verification approach. The final outcome is a set of 447 models, used to stress the presented passivity verification scheme.

Table 3.5 Summary of sampling-based passivity check. From [45] © 2021 IEEE.

# Tests	Passive	TP	FP	FN	Passive but FN
447	179	446	0	1	1

Each test is classified as: True Positive (**TP**), False Positive (**FP**), False Negative (**FN**), see text.

Table 3.5 illustrates a first summary of this investigation, by comparing the sampling-based passivity results with the standard Hamiltonian approach based on (3.28) or (3.30). To this end, each testcase can be classified as

- **TP**(True Positive), if both passivity verification approaches provided the same result, either passive or not passive;
- **FP** (False Positive), if the sampling scheme flagged the model as passive while the Hamiltonian verification did not;
- **FN** (False Negative), if the Hamiltonian verification flagged a model as passive while the sampling scheme did not;
- **Passive but FN** reports the number of non-passive models that the Hamiltionian check failed to classify, missing some passivity violation.

The two algorithms provided the same model classification for 99.9% of the cases, while only once the presented sampling algorithm identified a non-passive model while the Hamiltonian-based characterization provided a positive (passive) outcome. A further investigation shows that the only False Negative was produced by a non-passive model with a very narrow and small violation ($|\sigma_{\max} - 1| \approx 1 \cdot 10^{-10}$), reported in Fig. 3.7. In this case, imaginary eigenvalues of the Hamitolnian matrix (3.28) were not detected so that the corresponding check provided a passive flag. From the outcome of this intensive testing campaign, two interesting conclusions can be stated in terms of reliability. First, the provided sampling approach seems to be as robust as standard Hamiltonian approaches, with the advantage of identifying also those critical small violations that may be lost in the uncommon (and extraordinary) event of a Hamiltonian test fail the due to an ill-conditioned matrix.

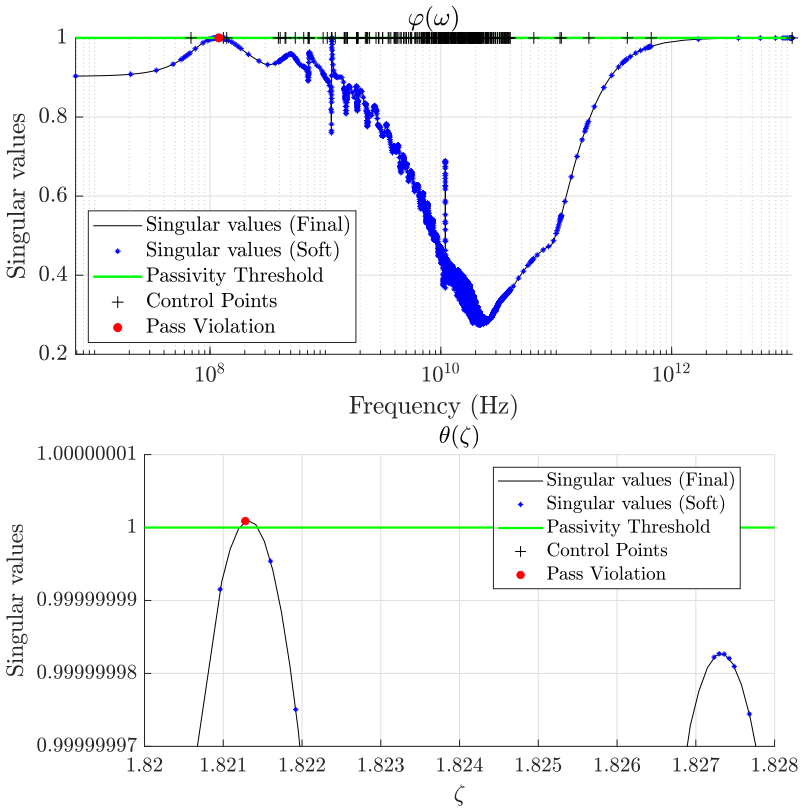

Fig. 3.7 False Negative case (#260): the tight passivity violation is not identified with a standard Hamiltonian-based check. Bottom: detail of the rescaled metric $\theta(\zeta)$ after warping of $\varphi(\omega)$. From [45] © 2021 IEEE.

Next, the two passivity verification approaches are compared in terms of runtime. Figure 3.8 shows the result produced with a prototype MATLAB [89] code, requiring a full Hamiltonian solver, and compares the required elapsed time (top panel) of the two strategies for an increasing model complexity, reported in terms of ports P, poles $\bar{n}$ and states N in the bottom panel. From this figure it is evident that the Hamiltonian check is still unbeatable in the small-size models case, even if both strategies provided the same outcome in less than a second. Turning to the large-scale investigated examples, the opposite situation arises. In this configuration, the proposed scheme is able to provide results in less than 10 seconds for all examples, providing a speedup of 50× with respect to the standard Hamiltonian characterization.

Notice also that, as expected, the runtime of the proposed algorithm is case-dependent and can not be inferred a priori looking at the model complexity. This is clearly confirmed by Fig. 3.8, which in fact demonstrates how two models

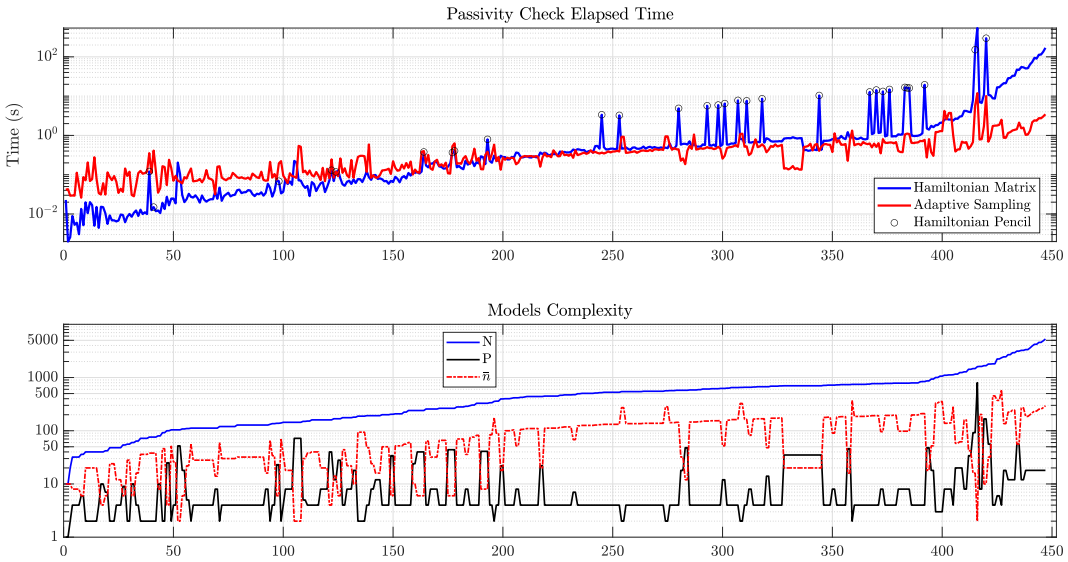

Fig. 3.8 Regression test results: elapsed time (top) and model complexity (bottom) are reported for all 447 benchmarks. From [45] © 2021 IEEE.

with similar complexity may lead to quite different results, as reported for most of the medium-size tests. Indeed, some observations are in order:

1. the extendend Hamiltonian pencil (3.30), more numerically stable and robust, was necessary to handle models with almost singular direct coupling term $\|\mathbf{D}\|_2 \approx 1$, at the price of an increased runtime with respect to the classic Hamiltonian formulation (3.28);

2. the Hamiltonian approach is able to detect only the crossing of γ by singular values trajectories, while the presented sampling technique increases the resolution at local maxima which are close to the threshold up to the point to distinguish whether they are associated to a passivity violation or not.

The latter further supports previous consideration on the False Negative testcase reported in Fig. 3.7. Indeed, this tight violation is the result of a narrowband singular value crossing of threshold γ in two nearby frequency points, that find their couterpart in two (almost defective) imaginary eigenvalues of an ill-conditioned Hamiltonian matrix. In the same condition, the adaptive sampling scheme improves the refinement level until the local maxima characterization is satisfactory. On the other hand, if the peaks of the passivity metric stay very

close to the threshold without exceeding it, the presented strategy will perform in the same way with the different outcome of not detecting any violation. This may provide a slight computational overhead when the passivity metric is consistently near γ. Conversely, the computational effort of a Hamiltonian approach remains constant, in the latter case without identifying any imaginary eigenvalues, with the possible risk of encounter the situation of Fig. 3.7, thus providing an unreliable result.

This further supports the need of a localized search strategy for each subband, even if this choice may increase the sampling effort in region of the space where passivity violations are not an issue. To this end, an illustrative example is presented in Fig. 3.9. This toy case is built on a system with two complex conjugate pairs of poles, defined to be very close one another, such that $p_n^1 = \alpha_n^1 \pm \mathrm{j}\omega_n^1$ and $p_n^2 = \alpha_n^2 \pm \mathrm{j}\omega_n^2$ where $\omega_n^i = \omega_0^i + n \cdot \delta\omega^i$ and $\delta\omega^i \ll \omega_0^i$. The passivity metric of this system changes unpredictably according to the residues, and it may show strong variations. On the other hand, it is likely that the two trajectories centered at ω_0^1 and ω_0^2 will not be distinguishable using a coarse sampling approach or with the only use of control points, as documented in Fig. 3.9. This motivates the presence of samples in frequency regions that may seem unnecessarily over-sampled, as at those frequencies higher than 10 GHz in Fig. 3.2. Notice that without knowing in advance the passivity metric variations, assuming the worst case scenario enables avoiding to miss narrow maxima as it already happened in the very simple example of Fig. 3.9.

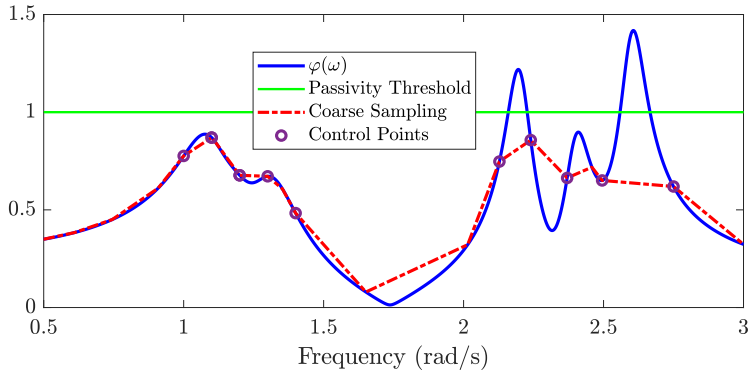

Fig. 3.9 A synthetic example to show how control points alone cannot provide full information of local peaks (and their variation): a local refinement is needed. From [45] © 2021 IEEE.

A 400-ports shielding enclosure

In this section a shielding enclosure with $P = 400$ is used as benchmark to analyze the sampling-based passivity verification with several Hamiltonian strategies available in the commercial macromodeling tool IdEM [33].

For this example, the unprocessed MoM dataset, defined as a set of 500 frequency responses in the range $[0.25, 990]$MHz, is used to build a standard VF model via IdEM [33]. The result is a rational representation with $\bar{n} = 26$ poles, that can be suitably converted into an equivalent state-space realization with order $N = 10400$. Model validation for a selected set of responses is reported in Fig. 3.10, proving a sufficient level of accuracy. Table 3.6 summarizes the comparison result of four different passivity check strategies:

1. *Full Hamiltonian*: computes imaginary eigenvalues of the standard Hamiltonian matrix (3.28);
2. *Sparse Hamiltonian*: takes advantage of an iterative multishift Krylov subspace iteration within the Hamiltonian solver, as reported in [65];
3. *Adaptive Hamiltonian*: mixes a sparse Hamiltonian eigensolver with a sampling strategy, as described in [66];
4. The two-stage adaptive sampling scheme of [45].

All Hamiltonian-based check are applied as available in [33], using two operating modes with $T = 1$ and $T = 8$ parallel threads, respectively. The adaptive

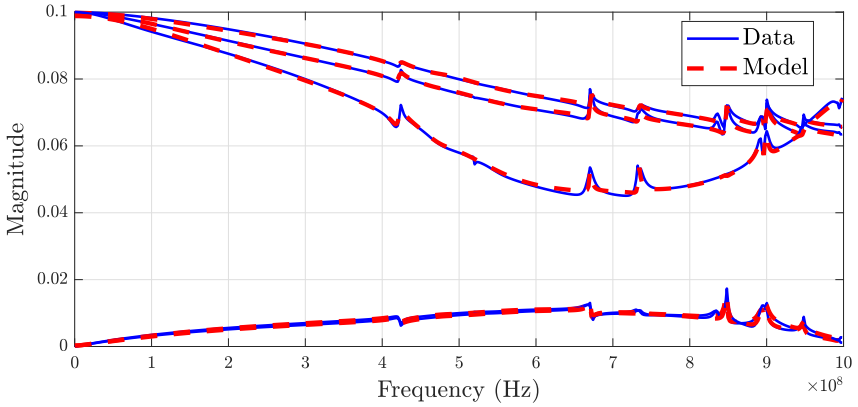

Fig. 3.10 Shielding enclosure model validation. From [45] © 2021 IEEE.

	Hamiltonian [33]			Proposed in [45]		
# Threads	Full	Sparse	Adaptive	*soft*	*hard*	*final*
8	180	76	75	0.4	1.41	4.49
1	226	189	183	1.03	3.19	9.82

Table 3.6 Passivity check of a shielding enclosure model. Runtime (minutes) of four strategies, under both single- and multi-threaded mode. From [45] © 2021 IEEE.

sampling approach was implemented as a prototype serial MATLAB [89] code, run by limiting the MATLAB environment to the required multithread level, affecting the low-level performances of the SVD, therefore slowing down the single-point passivity metric evaluation (3.26).

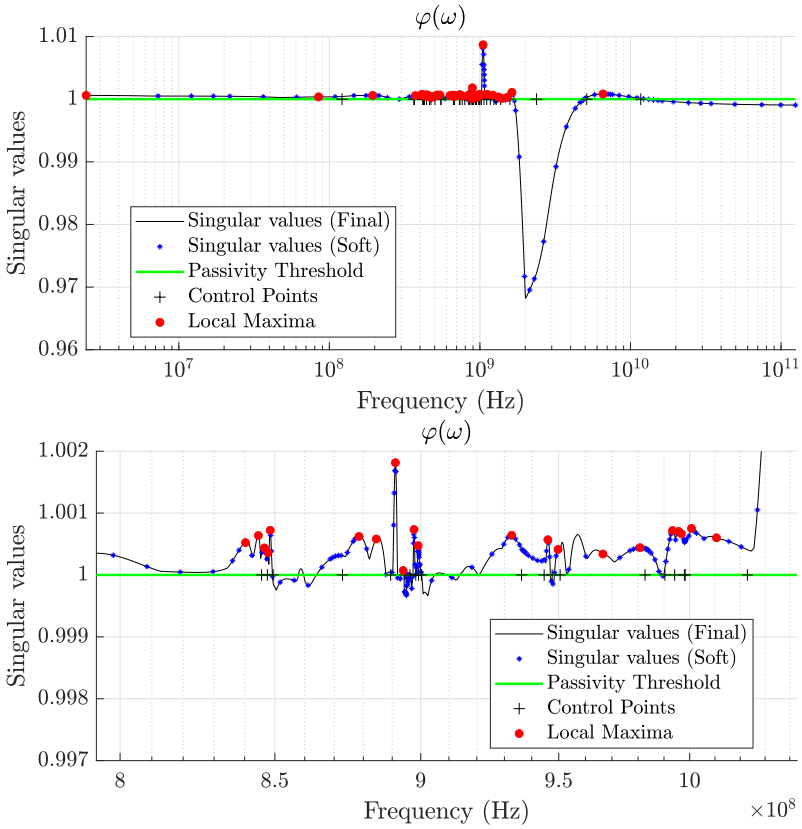

Fig. 3.11 Shielding enclosure: passivity metric (top) and a relevant band detail (bottom). From [45] © 2021 IEEE.

Table 3.6 demonstrates the efficiency of the presented adaptive sampling technique with respect to standard Hamiltonian approaches, providing a better

runtime in all cases and an unchanged result in terms of characterization. Figure 3.11 illustrates the passivity metric (envelop of maxima singular value trajectories), and indicates with red dots the identified local maxima associated to passivity violations

A 640-ports via array

The next example is a 20×20 vias array (4:1 signal to ground ratio) constructed on a 8-layer Printed Circuit Board (PCB). The overall structure is a 640-ports LTI system, where the first 320 port are defined in the top layer and the others on the bottom one. The MoM field solver [90] was used to build a set of responses $\{\breve{\mathbf{H}}(\mathrm{j}\omega_k) \in \mathbb{C}^{P\times P}, k = 1, \cdots, K\}$ with $K = 150$ samples up to 30 GHz. A rational model with $\bar{n} = 40$ poles and an overall number of states $N = 25600$ was obtained by fitting this dataset with the VF-based identification tool IdEM [33].

The same passivity verification comparison of the previous example is provided, in order to assess performances of the adaptive sampling passivity verification against standard Hamiltonian approaches available in IdEM [33]. Model validation on a set of representative responses is reported in Fig. 3.12. Figure 3.13 provides the passivity metric (3.26) characteristic obtained with the sampling-based check of Sec. 3.3, indicating with red dots local maxima associated to passivity violations. As it can be noticed from the bottom panel of Fig. 3.13, the different sampling operating modes, *soft*, *hard* (not shown) and *final*, identified a slightly different number of maxima (10 for the first, and 11 for the other two) motivating the need of different parameters setting. Indeed, the peak close to 4 GHz is not detected by the fast (*soft*) mode, that performs a check in the band up to 4.3 GHz, and it stops after detection of the (largest) violation close to 2 GHz. This result is consistent with the observations on the parameters settings provided in Sec. 3.3.3.

Table 3.7 summarizes execution time of the four passivity verification algorithms, showing a drastic improvement with respect to standard Hamiltonian approaches, at least for the two operating mode *soft* and *hard*. Notice that an excessive memory demand prevents the full Hamiltonian characterization.

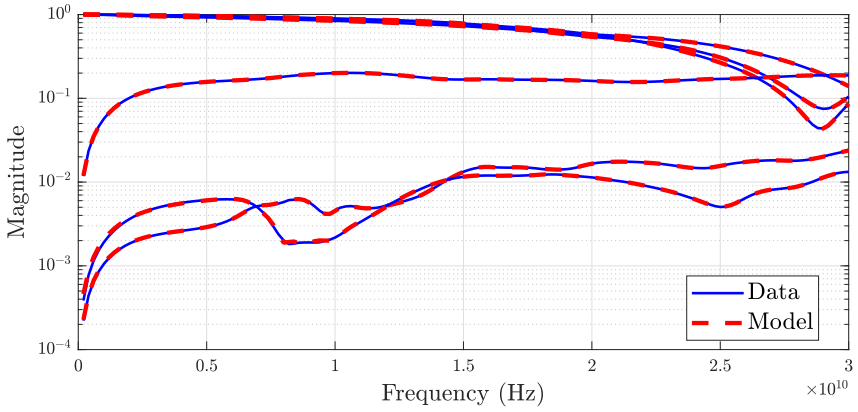

Fig. 3.12 Via array model validation on a set of representative responses. From [45] © 2021 IEEE.

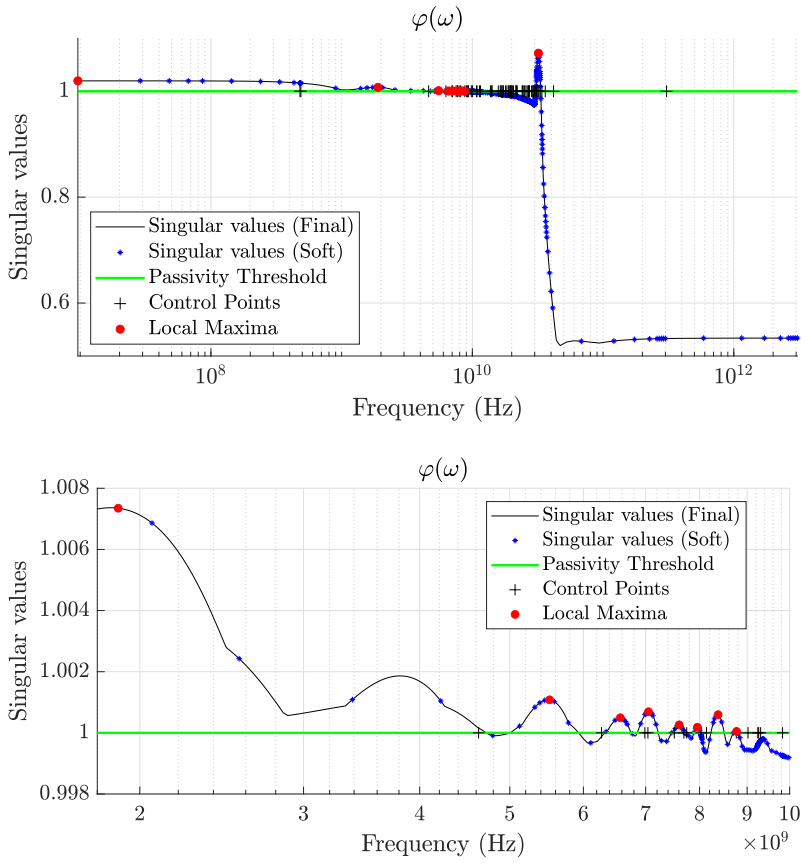

Fig. 3.13 Sampling-based passivity check results on the via array model. Both *hard* (not displayed) and *final* parameters setting enable to locate all local maxima (see bottom panel). From [45] © 2021 IEEE.

	Hamiltonian [33]			Proposed in [45]		
# Threads	Full	Sparse	Adaptive	*soft*	*hard*	*final*
8	Fail	19	29	1.53	8.15	17.35
1	Fail	67	68	4.41	23.1	47.05

Table 3.7 Passivity check of a via array model with $P = 640$. Elapsed time (minutes) of four passivity verification strategies, under both single- and multi-threaded mode. From [45] © 2021 IEEE.

Discussion of sampling-based passivity check results

The results originally presented in [45] prove that the sampling-based passivity check is the only approach that is able to manage a large-scale macromodel characterization. Some observations on the algorithm behavior are in order.

- The frequency warping procedure of stage 1 may lead to a large number of frequency intervals when the number of model poles $\bar{n}$ increases. Nevertheless, this condition does not limit the algorithm capabilities since all subbands are treated independently by the actual sampling (stage 2). Thus, a data-based parallelization of the presented scheme is trivial, and it is expected to further enhance the already competitive serial implementation that produced the presented results. Code parallelization has not been exploited yet, and it is left for later investigations.
- A too soft parameters setting may not allow detecting more than one maximum per subband. Specifically, this condition may be verified with a conservative (small) budget selection or with excessively strong stopping criteria.
- The passivity metric smoothness is the main responsible for the time performances of the passivity-based check, whose efficiency is strictly problem dependent and can only be inferred on a statistically base. If local maxima are sharp and close to the passivity threshold, the adaptive refinement strategy will possibly require many function evaluations to estimate if any peak is associated with a passivity violation, as in Fig. 3.13. This condition is exacerbated by the presence of many local maxima, reducing the overall speed of the algorithm. On the other hand, if the passivity metric shows large passivity violations, these are easily detected

and the algorithm will prevent further refinements, providing a very fast result.

- Since the algorithm of stage 2 is based on the NMSO, whose asymptotic consistency is guaranteed [83], by allowing both an infinite budget and basket reuse the same asymptotic converge is inherited by the passivity-based sampling. This means that relaxing these two conditions, all local maxima can be properly identified. This approach was never exploited since not necessary in all available test cases, that were correctly characterized by the presented implementation.

3.5.2 Passive modeling of shielding enclosures

This section provides numerical results on the proposed macromodeling strategy applied to shielding enclosures of increasing size in terms of ports. The original MoM data is pre-processed with the full-band regularization and extrapolation routine described in Chapter 2, which provides a suitable starting point for the modeling phase. This section enriches the results of [30] with additional information of the large-scale examples. To conclude, a shielding enclosure with 1024 ports and 72704 states is presented to further stress the scalability of the entire macromodeling procedure, including fitting and passivity enforcement [19].

A 25-port box

The final result on the 25 port shielding enclosure used as running example through Chapter 2 is presented here. Two models are constructed on top of the full-band (modal) scattering responses now defined in the range $[0, 1]$GHz, and are used to illustrate the difference of a standard compressed model [34] against the proposed structured-compressed model of Section 3.2.1. Both models are fit with the same number of poles $N_w = 89$, but using two different set of basis functions. The first model is trained by solving the fitting problem (3.16) on the unstructured bases obtained from [34] and reported in the bottom panel of Fig. 2.11. The other model is built on the set of basis resulting from the BD-SVD of Section 2.2.2, and reported in Fig. 3.1 after rescaling via (3.39) to enforce the proper low-frequency structure. In addition to the DC-zeros

multiplicity, model responses level at DC is enforced through the additional constraint (3.40) to guarantee an exact match with the fitting data.

The enforcement scheme provided in Section 3.4 is used to finally obtain a passive structured-compressed macromodel of the 25-port box. Details on performances are provided in the next section in terms of iterations and runtime. Panel (a) of Fig. 3.14 shows the envelope of maximum singular values before and after the proposed enforcement loop, confirming a final passive result. The two models responses are compared in Fig. 3.14, using as reference the three main system representations: scattering, admittance and impedance are depicted from panel (b) to (d), respectively. Notice that these results are obtained after post-processing of scattering (modal) responses of both macromodels.

Figure 3.14 clearly shows a superior accuracy of the proposed structured compressed model under different loading conditions, as confirmed by the model behavior in different representations, in the entire modeling bandwidth up to DC. A similar result is obtained for all examples reported next.

Scaling up to 400 ports

In this section, the compressed dataset presented in Chapter 2.2.3 are used as a starting point for the proposed macromodeling scheme, which retains the low-frequency behavior of the data. In particular, models of shielding enclosures with 64, 100, 225 and 400 ports will be presented here. Notice that all models are constructed on scattering (modal) parameters, in the range $[0, 1]$GHz, fitting the set of ρ structured basis functions obtained via BD-SVD. Table 3.8 summarizes the result of the structured macromodeling procedure. Regarding execution time, the VF step shows that, thanks to the data compression phase, the model extraction is not affected by the number of ports P but only by the dimension of the basis set ρ, as expected. On the other hand, most of the elapsed time is spent for the passivity enforcement. Unfortunately, the adopted iterative scheme does not allow estimating the required number of iterations (i.e. the elapsed time) in advance, since the importance of passivity violations are related to the particular trajectories of all singular values. Nevertheless, the maximum RMS error clearly shows that all models are extremely accurate throughout the whole bandwidth.

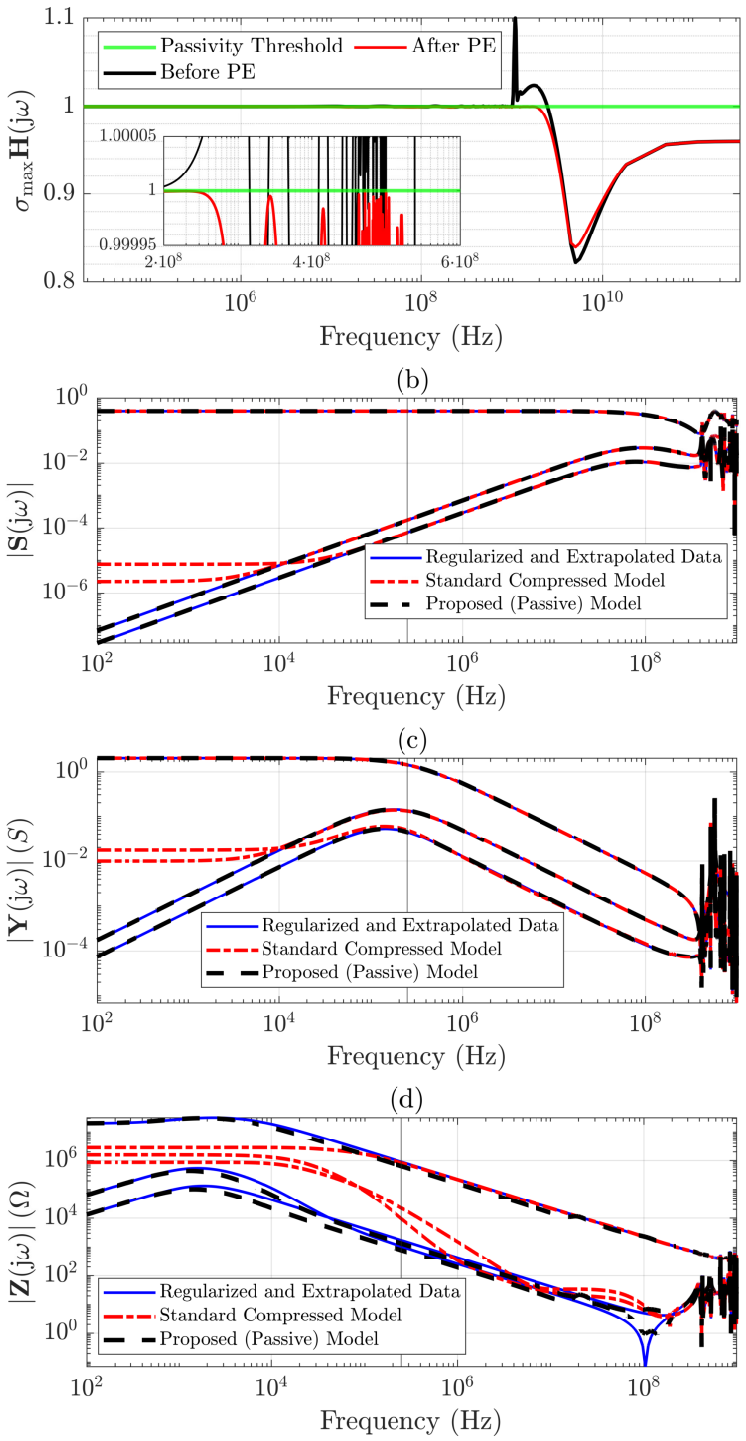

Fig. 3.14 A 25-port shielding enclosure: (a) passivity characterization before and after Passivity Enforcement (PE); model validation on a set of representative responses in scattering, admittance and impedance representation(b)-(d). From [30] © 2022 IEEE.

Ports	ρ	Order	VF Time (s)	Passivity Enforcement Time (min)	Iterations	RMS error on $\mathbf{S}(\mathrm{j}\omega)$
25	73	89	86.11	3.5	53	$6.77 \cdot 10^{-4}$
64	80	85	76.0	13.0	101	$3.62 \cdot 10^{-3}$
100	82	79	41.5	9.65	43	$7.79 \cdot 10^{-4}$
225	79	79	58.1	30	41	$4.78 \cdot 10^{-4}$
400	85	79	65.2	152	52	$9.57 \cdot 10^{-4}$

Table 3.8 Structured compressed macromodeling on shielding enclosures, built on the data-compression results of Table 2.2. From [30] © 2022 IEEE.

Figures 3.15 and 3.16 confirm that a similar level of accuracy is preserved for all representations, after conversion of (modal) models responses to the physical domain. To conclude, a comparison of standard compressed models as in [34] with the proposed structured ones confirms how the latter are able to preserve the same accuracy level down to DC, with a major improvement of the overall macromodeling phase.

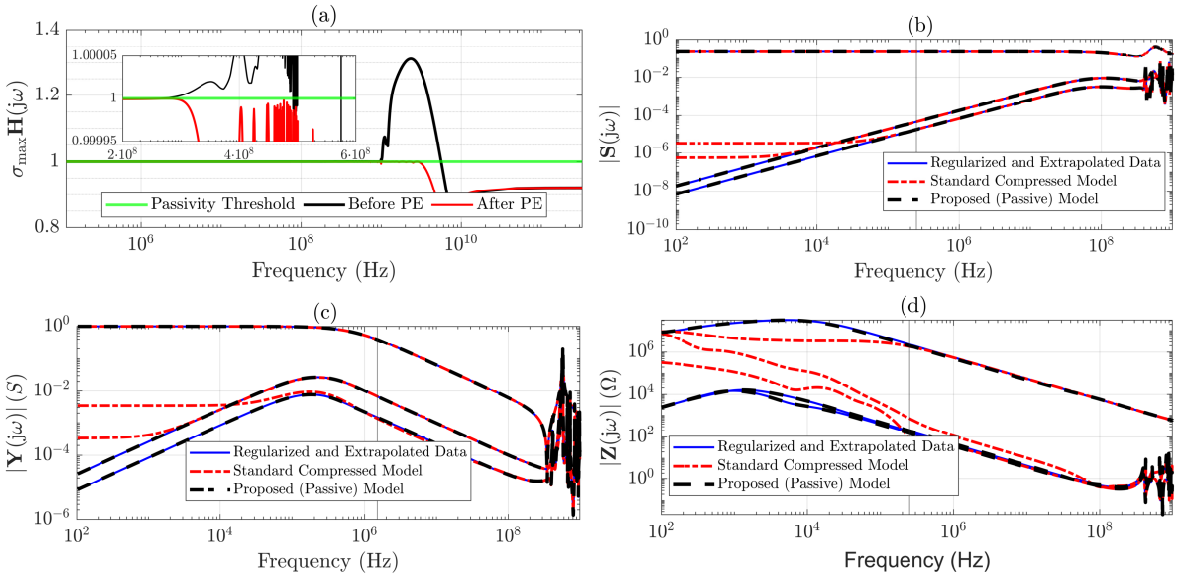

Fig. 3.15 As in Fig. 3.14 for shielding enclosures with $P = 64$ (panels a and b), $P = 100$ (c), and $P = 225$ (d). Panel (c) from [30] © 2022 IEEE.

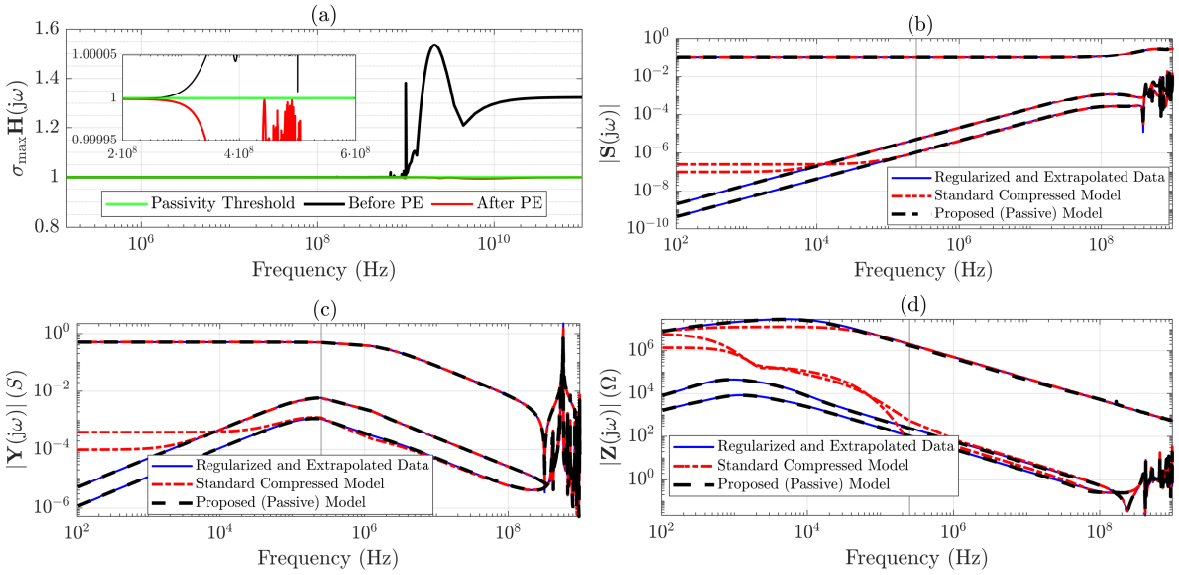

Fig. 3.16 As in Fig. 3.14 for a 400-port box. From [30] © 2022 IEEE.

A passive 1024-ports shielding enclosure

The last result presented in this chapter is aimed at demonstrating the effectiveness of the overall macromodeling procedure, including the passivity enforcement. For this reason, a passive model of a shielding enclosure with 1024 ports is presented here, which is the largest test case considered in this work.

The starting point is the usual frequency characterization consisting of $K = 750$ samples in the range $[0.02, 1]$GHz obtained with the MoM field solver. The original dataset has been first regularized and extrapolated up to DC following the procedure of Chapter 2, and a compressed macromodel has been realized following [34], since requirements on low-frequency accuracy were relaxed in favor of a simpler result obtained with a reduced number of basis, hence of model coefficients. This should not be intended as a limitation, but instead considered as a further confirmation of the presented strategy flexibility that still guarantees a certain level of accuracy in the low-frequency range thanks to the additional synthetic samples and the modal representation. The final rational representation has $\bar{n} = 71$ common poles, with a total number of $N = 72704$ states, and it is built on a reduced set of 41 basis functions. Building a full-size model as in (3.19) would have required $\approx 5 \cdot 10^6$ coefficients,

against the 2952 identified with the compressed framework, with an overall saving of $\approx 1000\times$ in terms of free variables to be identified during the fitting step.

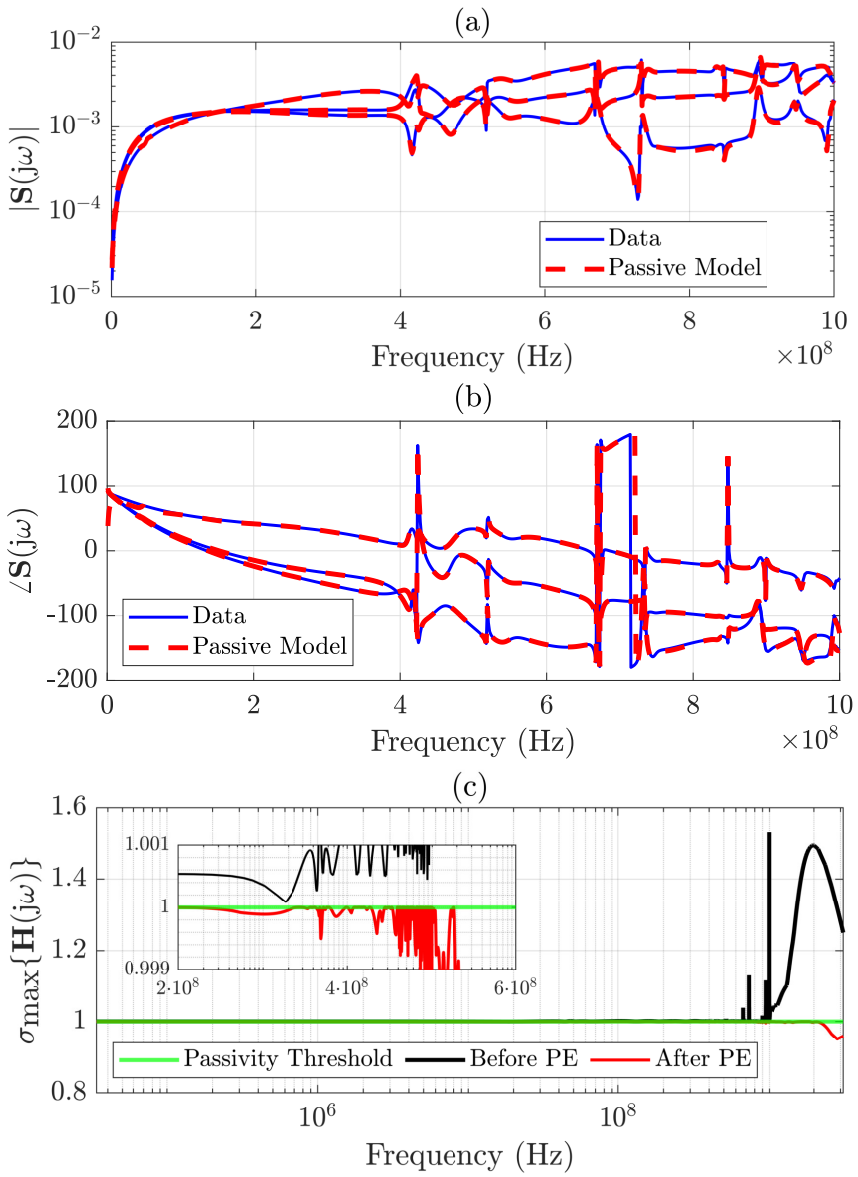

Fig. 3.17 Passive model of a 1024-port enclosure: model validation on a set of representative response, magnitude (a) and phase (b); passivity metric before and after passivity enforcement (c). From [19] © 2022 IEEE.

The model validation for the scattering parameters, after conversion from the modal responses, is reported in panels (a) and (b) of Fig. 3.17. Panel (c) shows the effect of a passivity enforcement scheme on the passivity metric (3.26),

demonstrating that all singular values are confined just below the passivity threshold $\gamma = 1$, hence the final outcome is passive. To reach this result, the model extraction required a total of 119 seconds, including data compression, and around 44.3 hours to enforce the model passivity by taking advantage of the presented sampling-based check. The overall enforcement loop required 82 iterations, with the first 70 performed using the Algorithm 3 in *soft* mode while only 12 rounds required the *hard* parameters setting, see Tables 3.3 and 3.4. Consider that the passivity enforcement was setup to automatically switch between the two modes whether the first reach a final passive outcome or when the maximum iteration count (of 70) is hit, in favor of a faster check during the first phase. The *final* check mode at last confirmed the passivity qualification outcome. The documented result shows how the passivity characterization (and enforcement) is still the most demanding phase of the all macromodeling procedure.

Notice that this level of complexity (1024 ports and 72704 states) with a standard macromodeling tool as IdEM [33] is practically not affordable without decreasing the overall complexity, i.e. the number model of poles, and ultimately affecting the final accuracy.

3.6 Conclusions

This Chapter provided a complete framework aimed at improving the state-of-the-art of macromodeling approaches for the construction of large-scale low-loss multiport passive systems. The main results of the presented strategy involved two fields of applications:

1. the identification process in case of high interface ports count and strong sensitivity to the change of system representation, due to the low-loss nature of the structure under analysis.
2. the passivity characterization of large-scale macromodels.

The corresponding main outcomes can be summarized as:

1. a complete strategy for reducing the model complexity in terms of coefficients, by controlling low-frequency accuracy and imposing particular features as DC zeros of several orders. The provided approach builds

on structured data-compression strategies, that exploit port responses redundancies, and preserve specific features of the data (as zeros at DC), which are here included in both the fitting and passivity enforcement stage. The final result is a procedure that enables to control full-band modeling accuracy up to DC and reduces the overall model sensitivity to loading condition, as demonstrated through numerical results.

2. a flexible and sparse compressed model representation, to be further exploited in transient simulation solvers that are aware of the sparse configuration to speed-up time-domain numerical solutions. Indeed, as it will be demonstrated in Chapter 4, the synthesis as equivalent circuit of large-scale rational model is not the most efficient choice for the integration in system-level simulations. As a matter of fact, when both the dynamic order and the count of electrical ports are relevant, the number of circuit elements associated to model coefficients may drastically affect the efficiency of reliable transient solvers, as those of the SPICE family.

3. a hierarchical adaptive-sampling passivity characterization algorithm for large-scale macromodels. The two-stage scheme, based on a pole-based warping frequency strategy followed by a passivity-driven tree-search approach, is proven to be as reliable as state-of-the-art approaches based on Hamiltonian matrices, with the advantage of a much better scalability with the model complexity. The presented implementation, combined with a suitable passivity enforcement scheme, enables building accurate passive macromodels with many ports and high dynamic order, pushing the limit of large-scale modeling.

Several numerical results have been presented to support the above statements. Models up to 1024 ports and 72704 states have been shown through this Chapter, always providing a satisfactory approximation of the related electromagnetic structure behavior.

Future investigation in the direction of a further increased system complexity are still needed. First, the code parallelization of the presented schemes have never been addressed so far, especially regarding the passivity characterization algorithm. A naive implementation of the sampling-based check that takes

advantage of multicore or GPU architectures would further push the potentiality of this approach, providing a sensible speed-up to the main bottleneck of the overall procedure, hence the passivity enforcement. To conclude, the investigation of power-integrity modeling applications, that are similarly characterized by a relevant number of ports and strong sensitivity to the change of termination, is left for future research.

All presented results will be used as starting point for Chapter 4, whose aim is the application of these compact large-scale models in a suitable hybrid transient simulation framework for the time-domain characterization of shielding enclosures.

Chapter 4

Transient analysis

This Chapter provides a transient hybrid solver formulation that combines a macromodeling framework and an iterative decoupling approach based on an advanced Waveform Relaxation (WR) method. Thus, we continue our investigation on the challenging problem associated with energy-selective shielding enclosures [3, 7, 10] moving the focus to their time-domain analysis. The reference structure is a multiport Linear Time Invariant (LTI) system (the metallic box) terminated with nonlinear components at a set of lumped ports, spread across an aperture, as illustrated in Fig. 1.3. The objective of a shielding enclosure design is to verify the shielding performances in terms of Shielding Effectiveness (SE) with different incident fields, i.e. we want to estimate the electric field that can be observed inside the protective frame under the effect of a (possibly dangerous) incident signal. The analysis of an energy-selective nonlinear characteristic requires a time-domain simulation since this behavior is achieved due to the loading condition provided by the nonlinear devices. Full-wave electromagnetic solvers are suitable for solving the coupled linear/nonlinear Electromagnetic problem the in time domain. Nevertheless, these approaches turn out to be very inefficient when repeated analyses are required. Hybrid techniques can overcame this limitation by providing an equivalent formulation suitable for circuit solvers, see Fig. 1.3.

The main goal of this chapter is to provide a better alternative to a SPICE engine for performing transient simulations of large-scale nonlinearly loaded systems. To this end, we take advantage of the hybrid decoupling approach [13]

to split the problem into linear/nonlinear sub-systems, and we use a time-domain method for computing transient signals at the lumped (decoupled) ports. The presented procedure is a Waveform Relaxation (WR) scheme [35, 36], which is merely a method that iteratively improves an initial approximation of the target signals. The standard WR is implemented as a fixed-point iteration that corrects an estimate of the port signals by evaluating the response of the linear system and the nonlinear loads alternatively. This approximation stabilizes if the iteration operator is a contraction [36–40]. To improve the algorithm convergence, the decoupling network parameters can be suitably optimized [37, 41, 39, 42]. These approaches are however ineffective for the present application, and better solutions are needed.

This chapter illustrates the results originally reported in [19] and presents an efficient WR-based hybrid solver to overcome the well-known convergence difficulties induced by a poor match of the decoupling impedance level. The scheme of [19] builds on the result of [46] and introduces a novel algorithm by combining an inexact Newton-Krylov iteration and a time partitioning (windowing) approach [36]. The combination of these two ingredients drastically improves the local convergence in each interval of time and reduces the required computational effort, both in terms of time and memory. The presented scheme provides an integrated framework specifically designed to handle large-scale systems, equipped with suitable sub-systems evaluations and a restarting process. A large set of numerical results is used to validate the improved WR solver against the state-of-the-art HSPICE engine, demonstrating its reliability and efficiency for shielding enclosures applications with up to 1024 electrical ports.

This chapter documents the results of [19] as follows. First, background elements concerning the hybrid strategy and the standard WR scheme are reported in Section 4.1 and Sec. 4.2, respectively. Then, advanced WR formulations to handle the windowing strategy and the quasi-Newton iteration are summarized in Sec. 4.3 and Sec. 4.4. The fully integrated algorithm of [19] is presented in Sec. 4.5, while a complete validation of the scheme through numerical results is provided in Sec. 4.6. Conclusions are finally drawn in Section 4.7.

4.1 Background: model-based hybrid solvers

This section summarizes the three-step hybrid method proposed in [13] to compute the transient analysis of scattered electromagnetic fields by a metallic structure loaded with lumped nonlinear elements. This procedure is essential to evaluate the performance of any shielding enclosure, and it corresponds to finding a solution to the electromagnetic problem illustrated in Fig. 1.3. The main goal is to estimate the total electric field $e_{\text{obs}}(t, \boldsymbol{r})$ observed at the coordinates $\boldsymbol{r}$, located inside the shielding enclosure. To this end, we start assuming a Perfectly Electrically Conducting (PEC) system (the shield), electrically large, with multiple nonlinear devices (electrically small) connected to a set of lumped ports. The structure is subject to an incident plane wave electrical pulse $e_{\text{inc}}(t)$. Under such conditions, part of the incident field is scattered by the structure due to the induced currents that flow through the lumped terminations, which provide a nonlinear loading effect. To compute the observed field $e_{\text{obs}}(t, \boldsymbol{r})$, the scheme proposed in [13] suggested the following three steps:

1. *Replacing of the metallic enclosure under field excitation with a generalized Thevenin equivalent circuit.* To convert the metallic structure excited by an electric field to a classic Thevenin equivalent form as in the right panel of Fig. 1.3, two ingredients are necessary, $\mathbf{H}(s)$ and $\boldsymbol{v}_{oc}(t)$. These elements can be obtained by solving two independent field problems addressed by a frequency domain solver as MoM [31]. First, the impedance transfer matrix $\mathbf{Z}(s) \in \mathbb{C}^{P\times P}$ of the unloaded structure must be characterized in a prescribed frequency bandwidth. To compute the impedance parameters all loads are detached from the enclosure, ports are excited with current sources and the corresponding ports voltages are computed for a set of frequency points K. This sub-system corresponds to the multiport yellow box of Fig. 1.3, where $\mathbf{H}(s)$ indicates a generic system representation suitably obtained after postprocessing of $\mathbf{Z}(s)$. The second ingredient of the Thevenin equivalent is the set of voltage sources $\boldsymbol{v}_{oc}(t) \in \mathbb{R}^P$, which translates the effect of the incident electric field $e_{\text{inc}}(t)$. To this end, the corresponding frequency domain transfer function under the excitation field $\mathbf{H}_{oc}(s) \in \mathbb{C}^P$, such that $\mathbf{V}_{oc}(s) = \mathbf{H}_{oc}(s)E_{\text{inc}}(s)$, is obtained by running a proper field solver with all electrical ports left open. Thus, the

simulation parameters in terms of the electric field are accounted for only by voltage sources of a generalized Thevenin equivalent circuit. Note that the frequency characterization of the enclosure $\mathbf{H}(s)$ is independent of the field excitation $e_{\text{inc}}(t)$, hence these parameters are the ideal candidate for a macromodeling approach as the one proposed in Chapter 3. On the other hand, changing the polarization or orientation of $e_{\text{inc}}(t)$ requires re-evaluating $\mathbf{H}_{oc}(s)$.

2. *Transient analysis of port voltages with nonlinear terminations.* This step is the actual (transient) numerical simulation of the circuit reported in Fig. 1.3, and it requires assembling all elements of the Thevenin equivalent circuit, including nonlinear terminations, to properly retain an equivalent of the nonlinearly loaded electromagnetic structure excited with the incident electric field. To this end, the two frequency domain datasets from the previous step must be suitably converted into a compatible time-domain form. First, a macromodel of $\mathbf{H}(s)$ is built and converted to an equivalent representation suitable for the selected transient solver. Then, a model of $\mathbf{H}_{oc}(s)$ is synthesized and exploited in a time-domain convolution with the electric incident field wave $e_{\text{inc}}(t)$ to compute the time-domain voltage sources $\boldsymbol{v}_{oc}(t)$ as final outcome. Thus, the solution of the equivalent Thevenin circuit loaded with nonlinear devices, as illustrated in Fig. 1.3, is obtained through a proper transient-solver. It can be demonstrated that the transient port voltages computed with this approach are equivalent (up to modeling errors) to the one observed solving for the full EM field coupling problem.

3. *Calculation of observed electric field.* The original radiation problem is cast into a set of equivalent linear sub-problems solved by taking advantage of the superposition theorem. Indeed, considering the contribution of an individual incident plane wave $\boldsymbol{E}_{\text{inc}}(s, \boldsymbol{r})$, the total observed electric field can be computed as the sum of three contributions

$$\boldsymbol{E}_{\text{obs}}(s, \boldsymbol{r}) = \sum_{i=1}^{P} \mathcal{C}_i V_i + \mathcal{D}\, \boldsymbol{E}_{\text{inc}} + \boldsymbol{E}_{\text{inc}}(s, \boldsymbol{r}) \tag{4.1}$$

where operator $\mathcal{C}_i$ represents the radiated field from a unit voltage source V_i located at the i-th port, with all other ports short-circuited, $\mathcal{D}$ is the scattering operator of the short-circuited enclosure.

The combination of these three steps is a robust and efficient procedure to compute the scattered EM problem with an equivalent circuit-based formulation, as confirmed by several results presented in [13].

We remark that the target of this dissertation lies within the first two steps of the above procedure, while Step 3 will not be considered in this work since only related to a (straightforward) post-processing phase. Indeed, Chapter 2 aimed at solving issues that characterized the dataset $\mathbf{H}(s)$ from Step 1, and Chapter 3 provided a complete framework for the generation of the models that are employed during the numerical simulation of Step 2. The latter is finally improved by this Chapter, which focuses on providing an efficient and reliable transient solver to compute the solution of the nonlinearly loaded Thevenin equivalent circuit of Fig. 1.3. For these reasons, the decoupling approach (Steps 1 and 2) is further detailed next, setting notation for later use.

4.1.1 Modeling of the decoupled LTI enclosure

This section provides a brief overview of the decoupling scheme [13] used to compute scattered EM fields of the shielding enclosures loaded with nonlinear terminations, formalizing the results of Steps 1 and 2 described above. Regarding the simulation setting illustrated in Fig. 1.3, the goal of this section is to compute voltages $\boldsymbol{v}(t)$ and currents $\boldsymbol{i}(t)$ at each lumped port.

To this end, we start defining the exchanged scattering waves at the ports

$$\boldsymbol{a}(t) = (\boldsymbol{v}(t) + \mathbf{R}_0 \boldsymbol{i}(t))/2 \tag{4.2}$$

$$\boldsymbol{b}(t) = (\boldsymbol{v}(t) - \mathbf{R}_0 \boldsymbol{i}(t))/2 \tag{4.3}$$

with $\mathbf{R}_0$ diagonal matrix storing the reference resistance level, and $\boldsymbol{a}(t), \boldsymbol{b}(t) \in \mathbb{R}^P$ incident/reflected scattering (voltage) waves.

Using the Laplace variable $s = \mathrm{j}\omega$, the reflected waves at the lumped ports can be defined in the frequency domain as

$$\mathbf{B}(s) = \mathbf{H}(s)\mathbf{A}(s) + \mathbf{\Theta}(s) \tag{4.4}$$

where $\mathbf{B}(s) = \mathcal{L}\{b(t)\}$ and $\mathbf{A}(s) = \mathcal{L}\{a(t)\}$, with $\mathcal{L}$ Laplace transformation. Notice that $\mathbf{H}(s) \in \mathbb{C}^{P\times P}$ is the transfer function (scattering matrix) of the enclosure with detached loads, while $\mathbf{\Theta}(s) \in \mathbb{C}^{P}$ stores contributions of the incident field under matching conditions. The latter can be written as

$$\mathbf{\Theta}(s) = \frac{1}{2}(\mathbb{I}_P - \mathbf{H}(s))\mathbf{V}_{oc}(s) \tag{4.5}$$

with $\mathbf{V}_{oc}(s) = \mathcal{L}\{\boldsymbol{v}_{oc}(t)\}$ denoting the voltages across the ports when all loads are detached and the structure is excited with the incident plane field $e_{\mathrm{inc}}(t)$. These set of voltages can be retrieved as

$$\mathbf{V}_{oc}(s) = \mathbf{H}_{oc}(s)E_{\mathrm{inc}}(s) \tag{4.6}$$

where $E_{\mathrm{inc}}(s) = \mathcal{L}\{e_{\mathrm{inc}}(t)\}$ and $\mathbf{H}_{oc}(s) \in \mathbb{C}^{P}$ is the open-circuit transfer function obtained following Step 1.

As already anticipated, a full-wave frequency domain solver as MoM [31] can be suitably called to characterize the two transfer functions $\mathbf{H}_{oc}(s)$ and $\mathbf{H}(s)$ in the band of interest, creating two frequency datasets $\breve{\mathbf{H}}_{oc}(\mathrm{j}\omega_\ell)$ and $\breve{\mathbf{H}}(\mathrm{j}\omega_\ell)$ in a set of points $\ell = 1, \cdots, L$. We further stress that only responses of $\breve{\mathbf{H}}_{oc}(\mathrm{j}\omega_\ell)$ should be re-evaluated when a different incident field $e_{\mathrm{inc}}(t)$ is applied to the structure, while $\breve{\mathbf{H}}(\mathrm{j}\omega_\ell)$ is fixed by the enclosure geometry.

These datasets are then exploited in a macromodeling step, either with a standard approach [21–23], or using the one proposed in Chapter 3, by enforcing the fitting condition

$$\mathbf{H}(\mathrm{j}\omega_\ell) \approx \breve{\mathbf{H}}(\mathrm{j}\omega_\ell), \quad \mathbf{H}_{oc}(\mathrm{j}\omega_\ell) \approx \breve{\mathbf{H}}_{oc}(\mathrm{j}\omega_\ell)\,, \tag{4.7}$$

where the surrogate rational models are defined as

$$\mathbf{H}(s) = \sum_{n=1}^{\bar{n}} \frac{\mathbf{R}_n}{s - p_n} + \mathbf{R}_0\,, \quad \mathbf{H}_{oc}(s) = \sum_{n=1}^{\bar{n}} \frac{\mathbf{\Upsilon}_n}{s - p_{oc,n}} + \mathbf{\Upsilon}_0 \tag{4.8}$$

with $\mathbf{R}_n \in \mathbb{R}^{P\times P}$ and $\boldsymbol{\Upsilon}_n \in \mathbb{R}^P$ model residues, and p_n and $p_{oc,n}$ corresponding poles. The passivity of $\mathbf{H}(s)$ is enforced in this work by means of the scheme proposed in Chapter 3.

When performing a transient simulation, the macromodel in (4.8) should be suitably embedded in a circuit solver. A common practice converts the rational form (4.8) in a state-space set of Ordinary Differential Equations (ODE), synthesizes an equivalent circuit representation of the model, and then runs the time-domain analysis with the selected SPICE engine. On the other hand, advanced SPICE solvers such as HSPICE [28] enable a direct definition of the pole-residue form (4.8), the *Foster form*, via suitable behavioral elements. The latter representation avoids the state-space conversion and reveals its efficiency in a numerical integration phase, taking advantage of a recursive convolution formulation, as detailed later.

4.2 The Waveform Relaxation method for time-domain analysis

Waveform Relaxation (WR) algorithms [35, 36] are a class of transient simulation approaches to analyze the behavior of a dynamic nonlinear system in time-domain. The main feature of any WR scheme is the ability to decompose a set of differential equations into a manageable number of (smaller) sub-problems, to be solved almost independently by iteratively updating relaxation sources. WR methods may differ by the definition of the relaxation process or by the strategy used to split the original system. Several examples are available in the literature [35, 36, 38, 39, 42].

This section provides a compact introduction to a scattering-based method belonging to the family of WR algorithms, with a specific focus on the Longitudinal Partitioning (LP) approach [39], which takes advantage of the scattering waves exchanged between a large-scale structure (usually linear) and its terminations (usually nonlinear) to set up a decoupling network. The general iterative setting is depicted in Fig. 4.1, where the relaxation sources at any iteration ν are the incident/reflected scattering waves of (4.2).

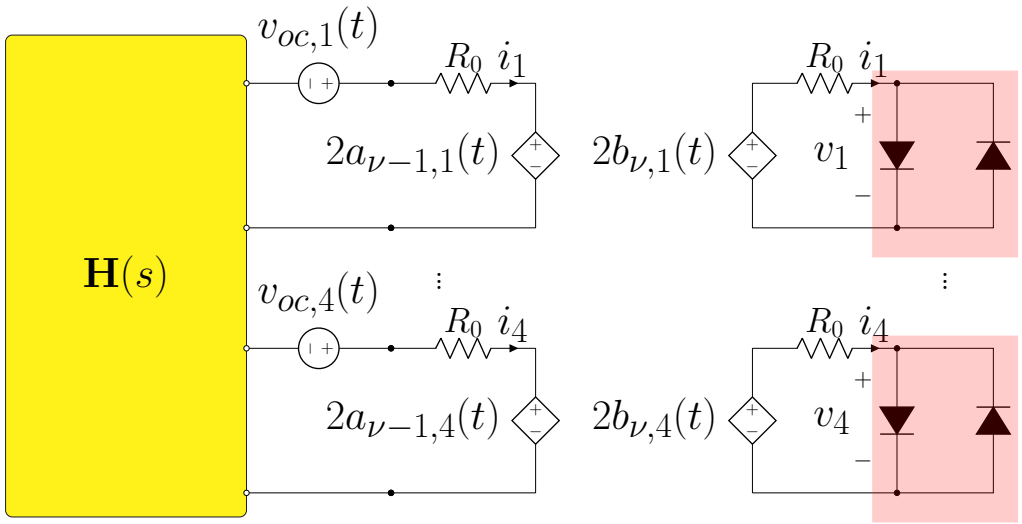

Fig. 4.1 WR-LP decoupling approach applied to the circuit problem of Fig. 1.3. Originally reported in [19] © 2022 IEEE.

This figure represents the WR decoupling applied to the circuit problem of Figure 1.3 the latter can be characterized by the following system of nonlinear equations

$$\begin{cases} \boldsymbol{b}(t) = (\mathcal{H}\boldsymbol{a})\,(t) + \boldsymbol{\vartheta}(t) \\ \boldsymbol{a}(t) = \mathcal{G}(\boldsymbol{b}(t)) \end{cases} \tag{4.9}$$

with time-varying signals $\boldsymbol{a}(t)$, $\boldsymbol{b}(t)$ and $\boldsymbol{\vartheta}(t)$ defined as port scattering waves and forcing term, respectively. The two operators $\mathcal{H}$ and $\mathcal{G}$ indicate the convolution with the shielding enclosure impulse response and the nonlinear map of the scattering fields provided by the terminations, respectively. Note that both operators are purely algebraic under the assumption of a suitable characterization of both enclosure (the discretized convolution can be represented as a matrix operator acting on a vector of discrete time samples) and loads.

The WR-LP scheme computes a solution of (4.9) solving at each ν-th iteration the following problem

$$\begin{cases} \boldsymbol{b}_\nu(t) = (\mathcal{H}\boldsymbol{a}_{\nu-1})\,(t) + \boldsymbol{\vartheta}(t) \\ \boldsymbol{a}_\nu(t) = \mathcal{G}(\boldsymbol{b}_\nu(t))\,. \end{cases} \tag{4.10}$$

Starting from an initial guess, which is commonly $\boldsymbol{a}_0 = \mathbf{0}$, the two equations of (4.10) are solved sequentially until suitable stopping criteria are reached. A common choice is to monitor the iteration error and to define a convergence threshold as

$$\|\boldsymbol{a}_\nu - \boldsymbol{a}_{\nu-1}\|_\infty < \epsilon \tag{4.11}$$

where the ∞-norm $\|\cdot\|_\infty$ selects the element with the largest amplitude among all signal components.
The WR-LP algorithm of (4.10) shows its full potentiality if

- C_1 the scheme converges;
- C_2 the computation of both $\mathcal{H}$ and $\mathcal{G}$ is fast;
- C_3 the number of iterations to obtain convergence is small.

All these conditions will be discussed next. In particular, C_1 is the topic of Section 4.2.1, C_2 is detailed in Sec. 4.2.2 and 4.2.3, while C_3 is the main objective of the full scheme provided by [19] and reported in Section 4.5.

4.2.1 Convergence of WR

Any fixed-point iteration algorithm as WR converges when the mapping operator among iterations is a contraction in a prescribed norm. Verifying this condition in the case of the WR-LP scheme is a straightforward task when considering linear terminations and converting (4.10) to the frequency domain. Following the same procedure of [39], we can rewrite all vectors and operators of (4.10) as functions of the Laplace-domain

$$\begin{cases} \mathbf{B}_\nu = \mathbf{H}\mathbf{A}_{\nu-1} + \boldsymbol{\Theta} \\ \mathbf{A}_\nu = \boldsymbol{\Gamma}\mathbf{B}_\nu(t)\,. \end{cases} \tag{4.12}$$

with $\mathbf{B} = \mathbf{B}(s) \in \mathbb{C}^P$ and $\mathbf{A} = \mathbf{A}(s) \in \mathbb{C}^P$ scattering waves, $\mathbf{H} = \mathbf{H}(s) \in \mathbb{C}^{P\times P}$ and $\boldsymbol{\Gamma} = \boldsymbol{\Gamma}(s) \in \mathbb{C}^{P\times P}$ scattering matrices of the (linear) electromagnetic structure and terminations, respectively. The last term $\boldsymbol{\Theta} = \boldsymbol{\Theta}(s) \in \mathbb{C}^P$ denotes the incident field contribution (4.5).

Defining $\boldsymbol{\Upsilon} = \boldsymbol{\Gamma}\boldsymbol{\Theta}$, the solution after $\mathcal{I}$ iterations is

$$\mathbf{A}_{\mathcal{I}} = \boldsymbol{\Gamma}\mathbf{H}\mathbf{A}_{\mathcal{I}-1} + \boldsymbol{\Upsilon} = \sum_{\nu=0}^{\mathcal{I}-1} (\boldsymbol{\Gamma}\mathbf{H})^{\nu}\boldsymbol{\Upsilon} \tag{4.13}$$

and matches the result of [39]. The solution when $\mathcal{I} \to \infty$

$$\mathbf{A}_{\infty} = \sum_{\nu=0}^{\infty} (\boldsymbol{\Gamma}\mathbf{H})^{\nu}\boldsymbol{\Upsilon} \tag{4.14}$$

corresponds to the solution of (4.12) satisfies the condition

$$\rho_{\max}\{\boldsymbol{\Gamma}\mathbf{H}\} = \max_{\omega,i} |\lambda_i\{\boldsymbol{\Gamma}(\mathrm{j}\omega)\mathbf{H}(\mathrm{j}\omega)\}| < 1 \tag{4.15}$$

where $\lambda_i\{\cdot\}$ indicates the i-th eigenvalue of the enclosed matrix. Note that a faster convergence rate is associated with a smaller $\rho_{\max}$. Recalling that

$$\rho_{\max}\{\boldsymbol{\Gamma}\mathbf{H}\} \leq \|\boldsymbol{\Gamma}\mathbf{H}\| \leq \|\boldsymbol{\Gamma}\| \, \|\mathbf{H}\| \tag{4.16}$$

a sufficient condition for the convergence of the WR-LP scheme [39] is

$$\|\mathbf{H}(\mathrm{j}\omega)\|_2 < 1\,, \quad \|\boldsymbol{\Gamma}(\mathrm{j}\omega)\|_2 < 1 \quad \forall\omega \tag{4.17}$$

which corresponds to the requirement of the strict dissipativity of both the electromagnetic structure and terminations. From a simplistic point of view, one can notice that the frequency-dependent convergence rate decreases the more passive are the two components of the decoupling scheme, meaning that the algorithm requires a smaller number of iterations to reach convergence.

The proposed hybrid framework guarantees that both the shielding enclosure $\mathbf{H}(\mathrm{j}\omega)$ and the nonlinear terminations (diodes) are passive by construction, thus ensuring the overall convergence of a WR approach. Nevertheless:

- the selected shielding enclosure (made of PEC) is almost lossless in the lower bandwidth and strongly resonant in the highest frequency region (see Chapter 2), meaning that the overall losses are minimal and that $\|\mathbf{H}(\mathrm{j}\omega)\|_2 \lesssim 1$ for most of the band of interest;

- the selected nonlinear loads (diodes pairs) are likely to have a small dissipation during the transition between the two extreme behaviors associated with their conduction and cutoff modes. Thus, also the terminations are in general $\|\boldsymbol{\Gamma}(\mathrm{j}\omega)\|_2 \lesssim 1$ with the (most likely) worst-case situation given by the extreme modes, when they do not dissipate energy since they can be characterized as short and open circuits with $\|\boldsymbol{\Gamma}(\mathrm{j}\omega)\|_2 = 1$.

The above working conditions of the selected application may discourage the use of any WR approach due to the poor spectral radius $\rho_{\max} \lesssim 1$ and the guarantee of a slow convergence rate. Nevertheless, following the procedure of [19], we will demonstrate that a proper configuration of WR provides a promising approach that gains both convergence and scalability under the worst-case scenario given by the large-scale problem at hand.

4.2.2 Evaluation of the linear system

An essential prerequisite of a competitive WR algorithm is a rapid evaluation of both terms obtained via the selected decoupling scheme (4.10), illustrated in Fig. 4.1. In this section, the first equation of (4.10) is addressed, namely the computation of linear operator $\mathcal{H}$, following the same strategy of [13, 18] and in particular of [19]. Starting from the frequency domain definitions of scattering waves and forcing term (4.4)-(4.6), we can define all operations cast in $\mathcal{H}$ as

$$\boldsymbol{b}(t) = \boldsymbol{h}(t) \circledast \boldsymbol{a}(t) + \boldsymbol{\vartheta}(t) \tag{4.18}$$

$$\boldsymbol{\vartheta}(t) = (\mathbb{I}_P - \boldsymbol{h}(t)) \circledast \boldsymbol{v}_{oc}(t)/2 \tag{4.19}$$

$$\boldsymbol{v}_{oc}(t) = \boldsymbol{h}_{oc}(t) \circledast e_{inc}(t) \tag{4.20}$$

with $\circledast$ time-domain convolution operator. Notice that the contribution of the forcing term $\boldsymbol{\vartheta}(t)$ in (4.10) is constant through WR iterations and can be computed during the algorithm initialization.

All steps of (4.18) can be efficiently evaluated by taking advantage of a *recursive convolution* formulation valid for the two impulse responses $\boldsymbol{h}(t)$ and $\boldsymbol{h}_{oc}(t)$, which are suitably described by a pole-residue model (4.8). Under this condition, these signal convolutions reduce to the solutions of an IIR filter as

defined in the following. We start assuming a uniform integration time-step Δt such that each instant of time can be defined as $t_k = k\Delta t$ for $k = 0, \cdots, \bar{k}$. First, we analyze a single port/pole toy system identified by a unique pole p with corresponding residue R. The resulting output $y(t)$ under excitation $u(t)$ at every instant of time t_k can be defined as

$$\widehat{y}_k = R\widehat{x}_k, \quad \widehat{x}_k = \alpha\widehat{x}_{k-1} + \beta_0\widehat{u}_k + \beta_1\widehat{u}_{k-1} \tag{4.21}$$

where $\widehat{x}_k = x(t_k)$ indicates the discrete state associate to the pole p with corresponding residue R. As detailed in [20], the invariant coefficients $\alpha = e^{p\Delta t}$ and $\beta_{0,1}$ can be defined according to the selected time discretization method.

To handle the general multiport case, we need to consider that the number of both input/output (ports) P and poles $\bar{n}$ affects the overall count of system states N. Thus, we start associating to each model pole (4.8) an auxiliary vector collecting the time-domain state evolution $\boldsymbol{x}_n(t) \in \mathbb{R}^P$, and we take advantage of the recursive convolution formulation (4.21) to approximate time samples as

$$\boldsymbol{x}_n(t_k) \approx \alpha_n\boldsymbol{x}_n(t_{k-1}) + \beta_{0,n}\boldsymbol{a}(t_k) + \beta_{1,n}\boldsymbol{a}(t_{k-1}) \quad \forall k = 1, \ldots, \bar{k} \tag{4.22}$$

with α_n and $\beta_{j,n}$ as in (4.21). Discrete reflected waves of (4.18) are then retrieved adding contributions of all states through model coefficients as

$$\boldsymbol{b}(t_k) \approx \sum_{n=1}^{\bar{n}} \mathbf{R}_n\boldsymbol{x}_n(t_k) + \mathbf{R}_0\boldsymbol{a}(t_k) + \boldsymbol{\vartheta}(t_k) \quad \forall k = 1, \ldots, \bar{k}\,. \tag{4.23}$$

Notice that both (4.22) and (4.23) require a proper initialization at the first instant of time t_0 of the auxiliary state vector $\boldsymbol{x}_0$, which stores all contributions $\boldsymbol{x}_n(t_0)$, and of the input scattering wave $\boldsymbol{a}_0 = \boldsymbol{a}(t_0)$. Note also that in case of complex conjugate poles coefficients $\mathbf{R}_n$ of (4.23) are complex-valued. In this case, only the (doubled) real-valued contribution of the product $\mathbf{R}_n\boldsymbol{x}_n(t_k)$ is added to the output vector $\boldsymbol{b}(t_k)$.

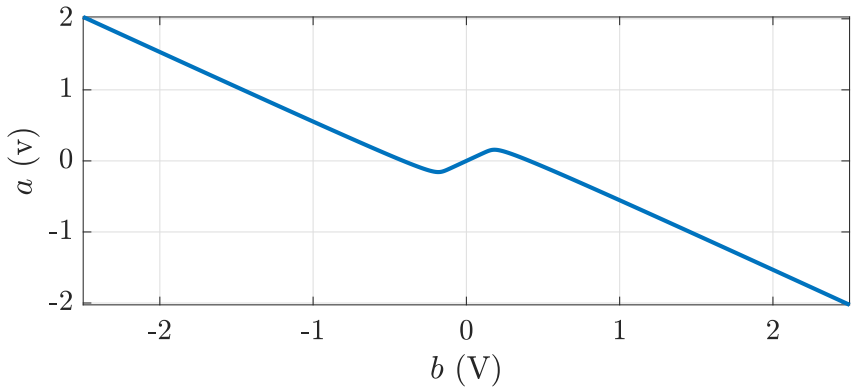

Fig. 4.2 Illustration of the characteristic equation (the operator $\mathcal{G}$) for a pair of ideal diodes in anti-parallel configuration, with reference resistance level $R_0 = 50\,\Omega$.

4.2.3 Evaluation of nonlinear loads

This dissertation focuses only on static nonlinear terminations, which are identical and not coupled through electrical ports. Under these conditions, that fit the case of diodes pairs, the algebraic operator $\mathcal{G}$ is a map between the incident/reflected scattering waves defined as

$$\boldsymbol{a}(t) = \mathcal{G}(\boldsymbol{b}(t)) \tag{4.24}$$

which can be applied independently to each component. Figure 4.2 illustrates an example of (4.24) when the selected termination is a pair of ideal diodes in anti-parallel configuration, like the one used through the numerical results of Section 4.6. Note that the input of $\mathcal{G}$ is the reflected scattering wave $\boldsymbol{b}(t)$ from the shielding enclosure, computed via recursive convolution. The practical implementation of the operator $\mathcal{G}$ in this work relies on a unique lookup table, defined sampling with a fine sweep of the scattering termination equation. Extensions to general nonlinear terminations will be discussed in Chapter 5.

4.3 Accelerating WR with windowing

In the WR framework, time partitioning techniques are well-established [36, 91] and already applied for efficient parallelization implementations [38]. In this work, we will further exploit a time *windowing* approach to break the simulation

window into small chunks with the objective of reducing the number of iterations required to reach convergence, ultimately improving the resulting runtime.

This property of the WR is illustrated with an example in Fig. 4.3, where the WR-LP error evolution among iterations is reported for a shielding enclosure with 100 ports. The simulation setting includes the electromagnetic structure terminated with identical pairs of ideal diodes in antiparallel configuration and excited with a plane incident electrical field, as detailed later in Section 4.6.1. Breaking the transient simulation interval, up to $t_{\max} \approx 15.5$ ns, in short portions provides a drastic improvement of the convergence rate. Indeed, the error at each iteration ν computed reducing the observation window on different pieces of the simulation interval $[0, \sigma t_{\max}]$, with $\sigma \in \{0.25, 0.5, 0.75, 1\}$, demonstrates how the WR convergence is reached faster for the shortest chunks, while a larger number of runs is required for the full-size window. Figure 4.3 illustrates the error evolution among iterations for both the approximation accessible from the previous iteration $\nu - 1$, reported in the bottom panel, and the exact solution computed imposing an error threshold of $\varepsilon = 10^{-8}$ and depicted in the top panel.

This example further stresses the opportunity of splitting the simulation interval to reach convergence with a prescribed target threshold, as already investigated since the very beginning of the WR approaches [36]. The simple strategy in [36] tried to both minimize the number of time points processed at each iteration and to optimize the window length to ensure a uniform converge in each time interval. In this dissertation, we focus only on the first partition effort by setting an upper bound to the computational burden required in terms of memory and by enabling to process all chucks independently, adding an outer iteration loop to the standard WR-LP scheme. To this end, we proposed in [19] to apply a uniform partition by means of M control points

$$0 = T_0 < T_1 < \cdots T_m < \cdots T_M = t_{\max} = K\Delta t \tag{4.25}$$

fixing the number $\bar{k}$ of uniformly sampled time points in each window, with $T_m - T_{m-1} = \bar{k}\Delta t$, simplifying the implementation effort.
Notice that this configuration includes two extreme cases. If the window is composed of only one sample (with $\bar{k} = 1$ and $M = K$), the transient evaluation is similar to the one applied by standard SPICE solvers. On the contrary, in

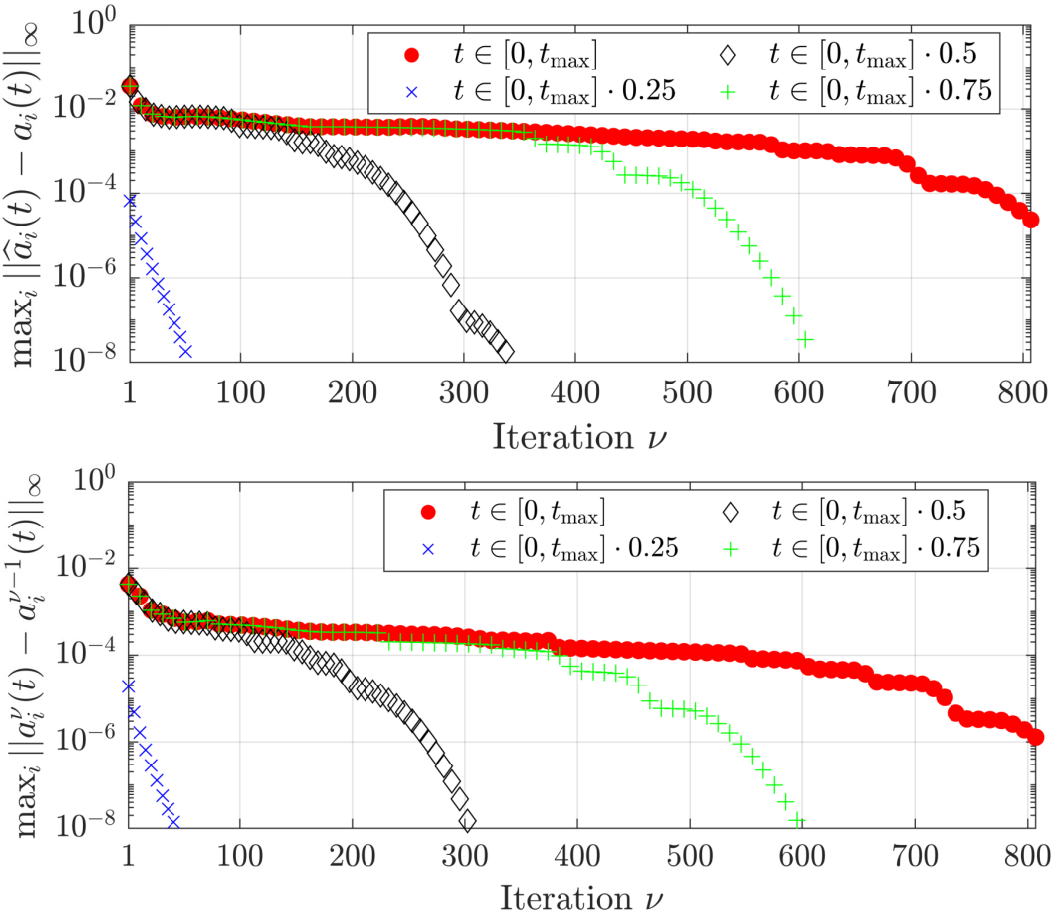

Fig. 4.3 Monitoring WR error evolution among iterations for different portions of the simulation window. Ports scattering waves $\boldsymbol{a}(t)$ are compared with the exact WR solution $\widehat{\boldsymbol{a}}(t)$ (top panel), and with the signals available from the previous iteration (bottom). Reference $\widehat{\boldsymbol{a}}(t)$ computed with error threshold $\varepsilon = 10^{-8}$. From [19] © 2022 IEEE.

the limit case when only a time window $M = 1$ (with $\bar{k} = K$) is considered the algorithm simplifies to the standard WR scheme, and all samples are processed at the same time for each WR iteration. As a result of the latter configuration, the convergence rate of the scheme is the one associated with the red dots in Fig. 4.3.

The solution proposed in [19] lies in between these two extreme situations, limiting the WR iteration to $\bar{k} > 1$ time points and imposing an upper bound to the number of samples to be handled simultaneously by improving the overall convergence of the scheme. Thus, the presented approach integrates a simple strategy to suitably initialize the (restarted) WR when dealing with time

partitioning. The initial conditions for sub-systems evaluation are combined with the windowing approach in Section 4.5, while a WR-based iterative solver is presented next to replace the basic WR loop with a significant advantage in terms of global convergence.

4.4 WR and inexact Newton-Krylov methods

We start introducing a discrete and vectorized version of the continuous WR-LP scheme (4.9), by imposing a uniformly sampled time vector $t_k = k\Delta t$ with $k = 0, \ldots, \bar{k}$. Under this condition, all time-varying signals $\boldsymbol{a}(t)$, and $\boldsymbol{b}(t)$ and $\boldsymbol{\vartheta}(t)$ can be cast as

$$\widetilde{\boldsymbol{a}} = \mathrm{vec}\,(\boldsymbol{a}(t_1), \ldots, \boldsymbol{a}(t_{\bar{k}})) \in \mathbb{R}^{\bar{k}P}\,, \quad \widetilde{\boldsymbol{b}} = \mathrm{vec}\,(\boldsymbol{b}(t_1), \ldots, \boldsymbol{b}(t_{\bar{k}})) \in \mathbb{R}^{\bar{k}P} \tag{4.26}$$

collecting all ports contribution and leading to

$$\begin{cases} \widetilde{\boldsymbol{b}} = \widetilde{\mathbf{H}}(\widetilde{\boldsymbol{x}}_0, \boldsymbol{a}_0) \cdot \widetilde{\boldsymbol{a}} + \widetilde{\boldsymbol{\vartheta}} \\ \widetilde{\boldsymbol{a}} = \widetilde{\mathbf{G}}(\widetilde{\boldsymbol{b}}) \end{cases} \tag{4.27}$$

where algebraic operators $\widetilde{\mathbf{H}}$ and $\widetilde{\mathbf{G}}$ are now discrete. We recall that $\widetilde{\mathbf{G}}$ denotes the evaluation of the nonlinear function of the loads. Whereas the discrete convolution $\widetilde{\mathbf{H}}$ as defined in (4.27) is formally associated with matrix multiplication, while the actual implementation takes advantage of a recursive convolution computation to obtain the auxiliary states vectors via (4.22) and to compute the scattering waves as in (4.23). The notation of $\widetilde{\mathbf{H}}(\widetilde{\boldsymbol{x}}_0, \boldsymbol{a}_0)$ in the first row of (4.27) expresses the dependency of the linear operator to a set of initial conditions, as available at the first time instant $t_0 = 0$, which are required to define both (4.22) and (4.23). These correspond to the input scattering wave $\boldsymbol{a}_0 = \boldsymbol{a}(t_0)$ and the discrete auxiliary state vector $\widetilde{\boldsymbol{x}}_0$, respectively.

We are now able to replace the standard WR iteration and solve (4.27) by properly defining a nonlinear multivariate problem. Indeed, we can reformulate (4.27) eliminating $\widetilde{\boldsymbol{b}}$ as

$$\widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}) = \mathbf{0}, \quad \widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}) = \widetilde{\boldsymbol{a}} - \widetilde{\mathbf{G}}(\widetilde{\mathbf{H}}(\widetilde{\boldsymbol{x}}_0, \boldsymbol{a}_0) \cdot \widetilde{\boldsymbol{a}} + \widetilde{\boldsymbol{\vartheta}}) \tag{4.28}$$

where $\widetilde{\mathbf{F}}$ is the nonlinear residual map $\mathbb{R}^{P\bar{k}} \to \mathbb{R}^{P\bar{k}}$ of the discrete $\bar{k}$ samples. A solution of (4.28) can be obtained with a standard (iterative) Newton method as

$$\widetilde{\boldsymbol{a}}_{\eta+1} = \widetilde{\boldsymbol{a}}_{\eta} - [\widetilde{\mathbf{J}}(\widetilde{\boldsymbol{a}}_{\eta})]^{-1}\widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}_{\eta}) \tag{4.29}$$

with Jacobian denoted as $\widetilde{\mathbf{J}}(\widetilde{\boldsymbol{a}}_{\eta})$ and where $\widetilde{\boldsymbol{a}}_{\eta}$ is the approximated signal available at the η-th iteration.

The Newton iteration can be summarized as

1. find an approximation of $\widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}_{\eta})$
2. solve for $\boldsymbol{s}_{\eta}$ the system of equations

$$\widetilde{\mathbf{J}}(\widetilde{\boldsymbol{a}}_{\eta})\boldsymbol{s}_{\eta} = -\widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}_{\eta}) \tag{4.30}$$

3. find $\widetilde{\boldsymbol{a}}_{\eta+1} = \widetilde{\boldsymbol{a}}_{\eta} + \lambda_{\eta}\boldsymbol{s}_{\eta}$ where λ_{η} is the iteration step length.

Following the procedure proposed in [46], we can avoid computing the Jacobian from Step 2, which represents the most demanding operation due to the large-scale nature of $\widetilde{\boldsymbol{a}}_{\eta}$. To this end, we replace the exact solution of the descent direction $\boldsymbol{s}_{\eta}$ with the approximation

$$\left\|\widetilde{\mathbf{J}}(\widetilde{\boldsymbol{a}}_{\eta})\boldsymbol{s}_{\eta} + \widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}_{\eta})\right\| \leq \gamma \left\|\widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}_{\eta})\right\| \tag{4.31}$$

known as *inexact Newton condition*. The solution of (4.31) can be efficiently computed avoiding the factorization of the Jacobian by taking advantage of an iterative scheme to approximate $\widetilde{\mathbf{J}}(\widetilde{\boldsymbol{a}}_{\eta})$ with a proper Krylov subspace definition.

To solve the inexact Newton iteration we selected the GMRES method as provided in [92]. Following the results of [46], we applied the WR-NGMRES scheme to strengthen the Newton-based iteration with an initial condition obtained with a reduced number of WR-LP iterations ν_i as in Section 4.2.

The iterative algorithm is terminated when the nonlinear residual drops below an accuracy threshold defined as

$$\left\|\widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}_{\eta})\right\| < \tau_r \left\|\widetilde{\mathbf{F}}(\widetilde{\boldsymbol{a}}_{\eta-1})\right\| + \tau_a \tag{4.32}$$

where control terms τ_r and τ_a are used to tune the relative and absolute accuracy, respectively.

This iterative solver provides the following advantages:

1. the global convergence of the scheme is ensured when GMRES is properly initialized, therefore a small number ν_i of standard WR fixed point iteration is run as a preprocessing step.
2. the number of iterations to achieve convergence is drastically reduced if the WR-NGMRES is applied instead of the standard WR-LP.

Both these conditions will be supported by numerical results in Section 4.6. The main drawbacks of this formulation are that:

- several function evaluations may be necessary to correctly approximate (4.31) and construct a basis of the Krylov subspace;
- the dimension of the nonlinear multivariate function (4.28) scales badly with the length of $\widetilde{\boldsymbol{a}}$, thus with the number of ports P or the window size $\bar{k}$.

To overcome these limitations, a suitable combination of windowing, evaluation restart, and inexact Newton-Krylov iterations is suggested in [19] and reported in the following.

4.5 Integrating WR-NGMRES and windowing

We now present the transient simulation scheme proposed in [19], whose outcome is an efficient and reliable framework that can handle the large-scale nature of the shielding enclosure verification under nonlinear load terminations. The presented method combines the windowing approach of Section 4.3 and the inexact Newton-Krylov iterations of Section 4.4 through a suitable restarting process. This approach, reported in Algorithm 4, exploits a WR-based decoupling method also in the case of poor matching conditions with remarkable convergence properties.

We start by defining a set of control points (4.25) to obtain M windows, with $\bar{k}$ time-points in each partition $\mathcal{I}_m = [T_{m-1}, T_m]$. To discriminate signals belonging to different time frames $\mathcal{I}_m$, we enrich the discretized signals (4.26) notation as

$$\widetilde{\boldsymbol{a}}_{\nu,m} = \mathrm{vec}\left(\boldsymbol{a}_\nu(t_1^m), \ldots, \boldsymbol{a}_\nu(t_{\bar{k}}^m)\right) \in \mathbb{R}^{\bar{k}P} \tag{4.33}$$

where m refers to the outer iteration on time intervals, and ν indicates the inner loop performed via either fixed-point WR or quasi-Newton iterations. Notice that the last time point of each window $\mathcal{I}_{m-1}$ should match the initial step of $\mathcal{I}_m$, ensuring a continuous partition of the simulation interval. Thus, samples of (4.33) can be formally defined as

$$t_k^m = T_{m-1} + k\Delta t, \quad \text{with} \quad t_0^m = t_{\bar{k}}^{m-1} = T_{m-1}\,. \tag{4.34}$$

We can now focus on the core of the presented scheme, summarized in pseudo-code form in Algorithm 4. Inputs of the Algorithm 4 are all quantities required for the evaluation of the two algebraic operators $\widetilde{\mathbf{H}}$ and $\widetilde{\mathbf{G}}$ in (4.27), and all parameters necessary to tune the windowing step and the convergence criteria (both in terms of accuracy and iterations). At first, zero vectors state/solution are suitably initialized (line 1 and 2). Then, the outer loop on each m-th time windows is started (line 2) and the (constant) initial solution value $\widehat{\boldsymbol{a}}_{m-1}$ is propagated to all $\bar{k}$ samples (line 4). Thus, the recursive convolution operator is suitably set to the initial conditions associated to the processed window (line 5), as summarized by the brief notation $\widetilde{\mathbf{H}}_m = \widetilde{\mathbf{H}}(\widehat{\boldsymbol{x}}_{m-1}, \widehat{\boldsymbol{a}}_{m-1})$.

The WR-NGMRES algorithm is then run (line 6). At first, the standard WR-LP method is applied solving each step of the system

$$\begin{cases} \widetilde{\boldsymbol{b}}_{\nu,m} = \widetilde{\mathbf{H}}_m \cdot \widetilde{\boldsymbol{a}}_{\nu-1,m} + \widetilde{\boldsymbol{\vartheta}} \\ \widetilde{\boldsymbol{a}}_{\nu,m} = \widetilde{\mathbf{G}}(\widetilde{\boldsymbol{b}}_{\nu,m}) \end{cases} \tag{4.35}$$

for ν_i iterations or until convergence is detected (line 10). Then, the NGMRES is applied to solve

$$\widetilde{\mathbf{F}}_m(\widetilde{\boldsymbol{a}}_{\nu,m}) = \widetilde{\boldsymbol{a}}_{\nu,m} - \widetilde{\mathbf{G}}(\widetilde{\mathbf{H}}_m \cdot \widetilde{\boldsymbol{a}}_{\nu,m} + \widetilde{\boldsymbol{\vartheta}}) = \mathbf{0} \tag{4.36}$$

for a maximum of $\nu_{\max}$ iterations or until convergence is reached (line 17). Since (4.36) provides a final approximation of the discrete incident wave $\widetilde{\boldsymbol{a}}_{\nu,m}$, at last $\widetilde{\boldsymbol{b}}_{\nu,m}$ is evaluated with final linear convolution (line 20). Notice that the latter step enables also to update all state variables, including values at the last time instant T_m, necessary to initialize the iteration on the next $m+1$ time window. Thus, initial conditions for the consecutive outer iteration are

Algorithm 4 Pseudocode of the WR-based scheme of [19].

Require: $\widetilde{\boldsymbol{\vartheta}}$, $\{\alpha_n, \beta_{0,n}, \beta_{1,n}\}$, $\{\mathbf{R}_n\}$, $\widetilde{\mathbf{G}}$
Require: M, $\bar{k}$, ν_i, ϵ, $\nu_{\max}$, τ_r and τ_a
1: Initialize states $\widehat{\boldsymbol{x}}_0 = \mathbf{0} \in \mathbb{R}^{P\bar{n}}$
2: Initialize solution $\widehat{\boldsymbol{a}}_0 = \mathbf{0} \in \mathbb{R}^{P}$
3: **for** $m = 1$ to M **do**
4: Initialize $\widetilde{\boldsymbol{a}}_{0,m} = (\mathbf{1}_{\bar{k}} \otimes \widehat{\boldsymbol{a}}_{m-1})$
5: Initialize operator $\widetilde{\mathbf{H}}_m = \widetilde{\mathbf{H}}(\widehat{\boldsymbol{x}}_{m-1}, \widehat{\boldsymbol{a}}_{m-1})$
6: **for** $\nu = 1$ to ν_i **do**
7: Apply convolution $\widetilde{\boldsymbol{b}}_{\nu,m} = \widetilde{\mathbf{H}}_m \cdot \widetilde{\boldsymbol{a}}_{\nu-1,m} + \widetilde{\boldsymbol{\vartheta}}$
8: Apply terminations $\widetilde{\boldsymbol{a}}_{\nu,m} = \widetilde{\mathbf{G}}(\widetilde{\boldsymbol{b}}_{\nu,m})$
9: **if** $\|\widetilde{\boldsymbol{a}}_{\nu,m} - \widetilde{\boldsymbol{a}}_{\nu-1,m}\|_\infty < \epsilon$ **then**
10: Break
11: **end if**
12: **end for**
13: **while** $\nu < \nu_{\max}$ **do**
14: $\nu \leftarrow \nu + 1$
15: Solve $\widetilde{\mathbf{F}}_m(\widetilde{\boldsymbol{a}}_{\nu,m}) = \mathbf{0}$ via NGMRES
16: **if** $\|\widetilde{\mathbf{F}}_m(\widetilde{\boldsymbol{a}}_{\nu,m})\| < \tau_r \|\widetilde{\mathbf{F}}_m(\widetilde{\boldsymbol{a}}_{\nu-1,m})\| + \tau_a$ **then**
17: Break
18: **end if**
19: **end while**
20: Update $\widetilde{\boldsymbol{b}}_{\nu,m} = \widetilde{\mathbf{H}}_m \cdot \widetilde{\boldsymbol{a}}_{\nu-1,m} + \widetilde{\boldsymbol{\vartheta}}$
21: Store final states $\widehat{\boldsymbol{x}}_m = \widetilde{\boldsymbol{x}}_{\nu,m}(T_m)$
22: Store final solution $\widehat{\boldsymbol{a}}_m = \boldsymbol{a}_{\nu,m}(T_m)$
23: **end for**
24: Merge converged signals for all m windows
25: **return** Converged port signals $\widetilde{\boldsymbol{a}}$ and $\widetilde{\boldsymbol{b}}$

retrieved both in terms of state variables (line 21) and input signal (line 22). To conclude the loop on time windows, all chunks of the resulting ports signals are gathered in a unique output result.

4.6 Transient simulations results

This section provides numerical results to assess the reliability and performances of the presented transient solver. These examples have been originally documented in [19], and are here reported supporting the presented WR-based

solver. The structure used for this investigation is the standard shielding enclosure of Fig. 1.3 loaded with nonlinear devices. The selected nonlinear terminations are anti-parallel ideal diodes pairs, each one defined by the relation $i_D = I_s(e^{v_D/V_T} - 1)$ with thermal voltage $V_T \approx 25$ mV and saturation current $I_s = 1$ nA. Thus, the comprehensive nonlinear characteristic of every load connected to one of the P electrical ports is $i = I_s(e^{v/V_T} - e^{-v/V_T})$.

The excitation is an incident plane wave, with electric polarization parallel to the diode grid branches and with orthogonal orientation. The selected electric field signal is a Gaussian modulated pulse defined as

$$e_{\text{inc}}(t) = \sin(2\pi f_c(t - t_0))e_g(t), \quad e_g(t) = Ee^{-\frac{(t-t_0)^2}{2\sigma^2}} \tag{4.37}$$

with center frequency f_c, time delay t_0, amplitude E and pulse width σ. Notice that we will modify both the center frequency and the signal amplitude among test cases to create a comprehensive set of simulation benchmarks, stressing resonances of the linear structure and exacerbating the nonlinear characteristic of the loads while defining performances of the shielding structure. This condition aggravates the lack of matching conditions at the WR decoupling ports, thus demonstrating the potential of an integrated WR simulation framework.

This section provides numerical results with four specific aims. Section 4.6.1 compares different reference WR schemes and confirms the improvements of the presented method in terms of convergence and number of iterations. Section 4.6.2 provides a comparison concerning a state-of-the-art SPICE solver, confirming that a suitable Waveform Relaxation (WR) can improve performances in terms of both accuracy and computational time. Section 4.6.3 illustrates the outcome of a testing campaign obtained by sweeping both center frequency and amplitude of the excitation signal $e_{\text{inc}}(t)$ for a total of 50 simulations. This systematic test enables us to confirm the previous comparison with a standard SPICE approach from a statistical standpoint. Section 4.6.4 concludes the numerical results, testing the scalability of the method in a large-scale setting and using, at last, the model with 72704 states and 1024 electrical ports presented in Chapter 3.

All benchmarks are tested using $\nu_i = 1$ iteration of the standard WR-LP and imposing a maximum of $\nu_{\max} = 100$ inexact-Newton iterations, even if this hard limit has never been reached, as illustrated next. All results are obtained

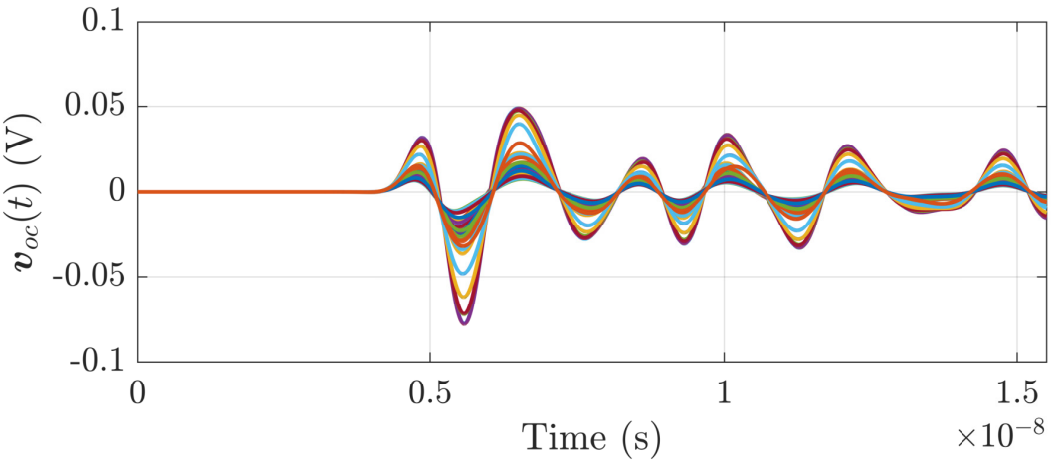

Fig. 4.4 Open-circuit voltages used in Sec. 4.6.1 to excite the 100-ports shielding enclosure and to compare WR schemes results. The selected incident field is a Gaussian Modulated Pulse centered at 450 MHz and with unitary amplitude.

with a prototype MATLAB code run on a Workstation with 64 GB of RAM and an i9-7900X CPU core (3.3 GHz), limiting the simulation environment to one computational thread.

4.6.1 Analysis of WR iterations

This section compares different WR iterations results on a shielding enclosure with $P = 100$ ports and a corresponding 10×10 grid of diodes. To estimate performances of the shielding we followed the procedure summarized in Section 4.1 (see [13] for details), and we started with a frequency characterization via MoM solver [31] of both the unloaded structure transfer function $\mathbf{H}(s)$ and of the open-circuit responses $\mathbf{H}_{oc}(s)$ under the incident field influence. Thus, we generated reduced order models from both sampled responses via the procedure described in Chapter 3. This step enabled us to suitably get all ingredients to define the discrete recursive convolution operator $\widetilde{\mathbf{H}}$ of (4.27), and to finally run the WR-based study.

The excitation electrical field $e_{\text{inc}}(t)$ is defined as in (4.37) with $f_c = 450\text{MHz}$, $\sigma \approx 0.415$ ns and unitary amplitude $E = 1$. The resulting open-circuit voltage sources $\boldsymbol{v}_{oc}(t)$ after convolution (4.18) are reported in Fig. 4.4. These are necessary to set the simulation benchmark as illustrated in 1.3, which is then used to run a series of transient simulations with different solvers. In all cases,

a time-domain analysis has been performed up to $t_{\max} \approx 15.5$ ns with a uniform integration time step $\Delta t = 2.22$ ps.

The fully integrated scheme presented in Section 4.5 is then compared with the standard WR-LP scheme [39], the WR method improved with the time partitioning approach of Sec. 4.3, and the WR-NGMRES method [39], i.e. the WR-based solver improved with a quasi Newton-Krylov iteration applied to the full simulation window as described in Sec. 4.4. All these WR algorithms are implemented in a suitable MATLAB environment and compared with a common setting, i.e. one computational thread and same accuracy thresholds to verify convergence $\varepsilon = \tau_r = \tau_a = 10^{-6}$. Thus, this test enables to split contributions of all separated ingredients that are integrated into the presented WR-based solver.

WR	Standard LP	NGMRES	Windowing		Proposed [19]	
Iter (mean)	807	25 (2+23)	90.6	25.45	7.6	6.52
Iter [min,max]	-	-	[1-127]	[1-41]	[1-9]	[1-9]
Windows (M)	1	1	10	100	10	100
Time (s)	650	779	88.9	30.3	45.4	15.6

Table 4.1 Comparison of four WR-based implementations, imposing convergence with stopping threshold $\varepsilon = \tau_r = \tau_a = 10^{-6}$. The number of iterations is documented both in terms of mean and minimum/maximum values when windowing is applied (last four cases). As in [19] © 2022 IEEE.

Table 4.1 summarizes the result of the four algorithms in terms of convergence (iterations) and computational time. Some observations are in order.

- The standard WR-LP decoupling method required the largest amount of iterations to converge (807), with a consequent significant runtime of 650 seconds due to the long sequence of evaluations of both operators $\widetilde{\mathbf{H}}$ and $\widetilde{\mathbf{G}}$.
- The inexact Netwon-Krylov implementation of the WR-NGMRES method improved the convergence rate since only 25 iterations are necessary to obtain an accurate result, but it deteriorated the overall runtime requiring 779 seconds. This behavior is not surprising due to the (large) number of evaluations needed to build a suitable basis for the Krylov subspace that solves the inexact-Newton step avoiding the system Jacobian computation.

To support this observation, we verified that most of the algorithm effort is spent for the NGMRES iterations (23), which represent approximately the 99.6% of the total runtime.

- The time windowing approach drastically improved the standard WR result in terms of runtime with more than 20× of speed-up, even if the overall number of iterations (2545, with an average of 25.45 evaluations for each of the $M = 100$ windows) is quite large. This result is reasonable if we consider that the convergence rate is improved (see Sec. 4.3 and [36]) but also that a reduced number of points are iteratively evaluated by the two discrete operators. Intuitively, each sequential evaluation is much faster if applied to a smaller time partition and not to the entire simulation interval. Thus 807 operations on the full window are more expensive than 2545 iterations on smaller chunks.

- The presented approach, which combines the WR-NGMRES and the time windowing methods, provided the best results both in terms of iterations (with average values of 7.6 and 6.52, considering $M = 10$ and $M = 100$ windows respectively) but also in terms of runtime (15.6 seconds), enabling an overall speed-up of about 40× with respect to the standard WR iteration. Notice that the smallest time partition with $M = 10$ degraded the algorithm efficiency, as expected. Thus, all results in the following sections are obtained by fixing $M = 100$ windows as reference choice.

Figure 4.5 shows an example of the converged ports voltages, demonstrating the accuracy of all methods with respect to a state-of-the-art SPICE solver.

4.6.2 A comparison with SPICE

We now compare the presented WR solver with a state-of-the-art SPICE engine, assessing performances related to computational time and final accuracy. The shielding enclosure selected for this investigation is the $P = 400$ ports system already presented in Section 3.5.2, and modeled by means of a Compressed Macromodel representation (3.19). The excitation electrical field is the Gaussian pulse $e_{\text{inc}}(t) = e_g(t)$ of (4.37), with standard deviation $\sigma \approx 0.47$ ns and

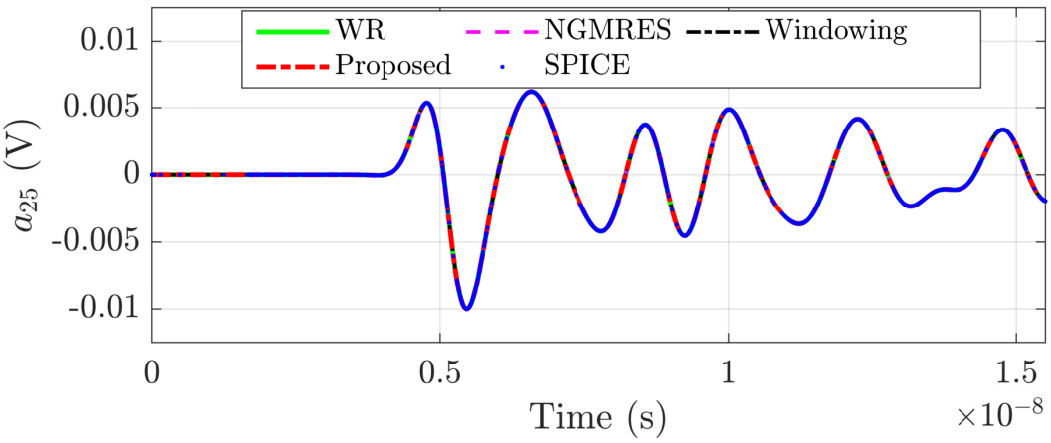

Fig. 4.5 Transient results of a 100-port shielding enclosure with different WR solvers. Table 4.1 summarizes all results for this example. Notice that only signals obtained with $M = 100$ windows are reported. From [19] © 2022 IEEE.

amplitude $E = 10^4$ selected to enhance the nonlinear behavior of all anti-parallel diodes. The time-domain verification of this system is performed up to $t_{\max} = 20$ ns, setting for the WR approach a uniform integration time step of $\Delta t = 25$ ps. This setting enables us to stress the limits of both transient solvers, as detailed in the following.

The SPICE engine selected for this investigation is a solver commonly available in the industry environment (HSPICE–L-2016.06-SP2-1 win64) that enables tuning several control parameters to perform time-domain simulations. Since we want to assess the efficiency and reliability of the presented solver, we collected solutions of the SPICE engine for the same transient analysis with different options settings, and we compared these results with the WR outcomes. In order to guarantee a fair comparison we should consider that:

- All presented WR schemes rely on a fixed integration time-step Δt, while HSPICE performs an adaptive stepping analysis to optimize the number of samples. Thus, the output signals of the two solvers are likely to be sampled at different instants of time, and they cannot be compared without a suitable interpolation procedure. Thus, the resulting induced interpolation error should be considered when comparing the two solutions in terms of accuracy.

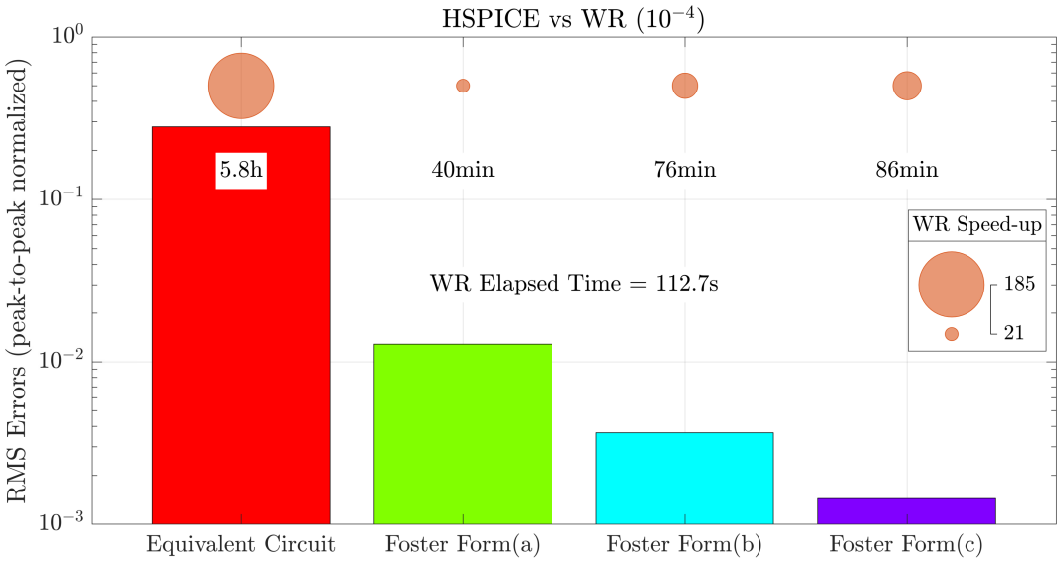

Fig. 4.6 Comparing WR and HSPICE transient results for a 400-port shielding enclosure with nonlinear terminations. Summary of the results in terms of accuracy, runtime and speed-up including four HSPICE settings (details in the text). From [19] © 2022 IEEE.

- The model representation used in the solver affects the outcome, especially in terms of runtime [27]. HSPICE enables to define surrogate models both as equivalent circuits and as pole-residue (4.8) transfer functions (via the *Foster* form). In this investigation, we compared both model representations. Nevertheless, notice that not all SPICE engines embed a Foster form and that the most common synthesis requires using a (sparse) state-space system representation [27, 20].

- To limit the number of control parameter changes, we focus on the most relevant in terms of final sensitivity. To this end, the HSPICE parameters modified in the following are:

  – `OPTION.DELMAX`, which sets an upper bound on the maximum allowed integration time-step, limiting the (automatic) adaptive stepping procedure;
  – `OPTION.ACCURATE`, which gathers several internal settings to boost the final solution accuracy.

Figure 4.6 provides a graphical summary of the result from this benchmark, comparing the outcome of the presented framework ($\tau_a = \tau_r = 10^{-4}$)

concerning HSPICE in terms of RMS error (vertical bars), speed-up (circles) and computational time (text notes). Notice that the graph shows the results of four HSPICE configurations, indicated with colored bars:

1. *Equivalent Circuit* (red bar): the SPICE engine embeds the macromodel in a standard equivalent circuit representation [20, 27], and the transient simulation runs with default HSPICE options;
2. *Foster Form (a)* (green): a suitable pole-residue (Foster) form includes the macromodel in the HSPICE environment, and the simulation runs with default options;
3. *Foster Form (b)* (cyan): as in case (a) but enhancing the HSPICE accuracy with the control parameter `OPTION.ACCURATE`;
4. *Foster Form (c)* (purple): as in case (b) but adding an upper bound to the integration time step imposing `OPTION.DELMAX`$=10^{-11}$.

For this specific example, Figure 4.6 indicates the following observations. First, regarding the solution accuracy, different HSPICE configurations strongly modify the simulation result. In particular, it is interesting to notice that:

- two model representations with the same simulation settings provide different results in HSPICE (red and green bars) with respect to the same reference signal, in this case the presented WR scheme;
- improving the simulation accuracy, the HSPICE result with the model in pole-residue form approaches the WR solution.

The latter point motivates the choice of using the Waveform Relaxation (WR) signal as reference instead of the SPICE solution, at least for this example. Regarding runtime, the presented Waveform Relaxation (WR) iteration always provides the fastest solution (112.7 seconds), showing a speedup that ranges from 21× up to 185× with respect to all HSPICE configurations. Notice that the Foster model form is always faster and more precise than the equivalent circuit representation, by at least an order of magnitude in all cases.

Figure 4.7 depicts the first two HSPICE Foster results for the transient voltage with the worst-case error concerning the WR solution, i.e. the signal at port 399. These results confirm how the HSPICE solutions get closer to the WR signal when improving the transient solver accuracy by tuning its control parameters. The benchmarks provided in the following sections take advantage

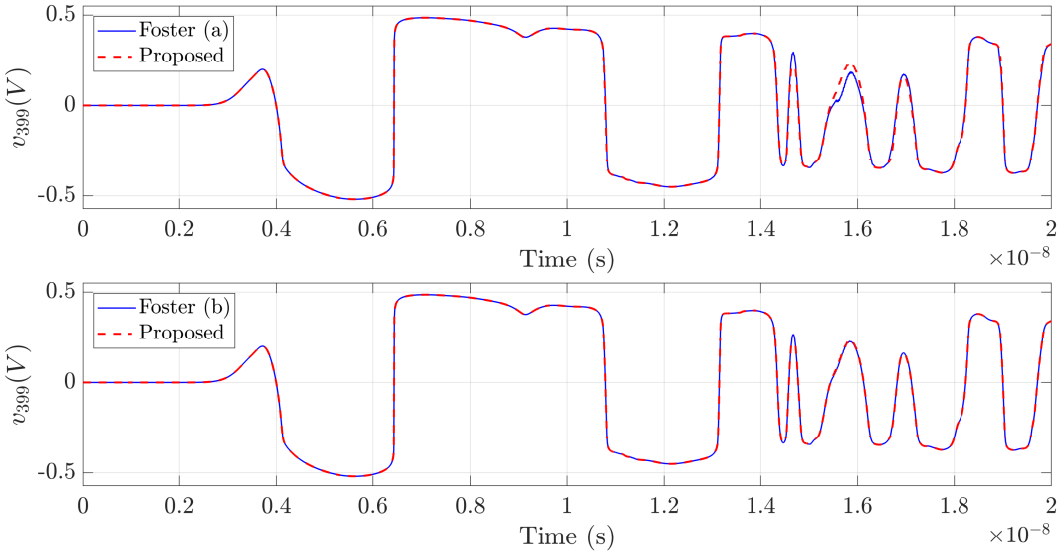

Fig. 4.7 Comparing WR and HSPICE transient results for a 400-port shielding enclosure with nonlinear terminations. The same port voltage signal is reported for two different HSPICE settings (details in the text). From [19] © 2022 IEEE.

of these observations, with all tests enabling `OPTION.ACCURATE` and adopting the Foster model form.

4.6.3 Systematic analysis against SPICE

We now focus on a systematic assessment of the presented WR scheme performances with respect to HSPICE, using as a running example the same 100-port model reported in Sec. 4.6.1 but modifying the excitation field waveform (4.37) in terms of frequency and amplitude. This test includes 25 independent transient simulations obtained from the combination of 5 (linearly spaced) center frequencies $f_c \in [100, 800]$ MHz and 5 (logarithmic spaced) values of amplitude $E \in [1, 10^4]$. Note that we adjusted the Gaussian pulse width σ of (4.37) to each center frequency. This procedure allows preserving the overall waveform shape through a time-domain stretching effect, conserving the center frequency location but modifying the pulse spectrum accordingly.

The final result reported in Fig. 4.8 gathers 50 transient simulations performed via HSPICE with the two leading model representations presented in Section 4.6.2, the equivalent circuit synthesis and the pole-residue (Foster) form. Regarding the WR algorithm options, all simulations are computed with

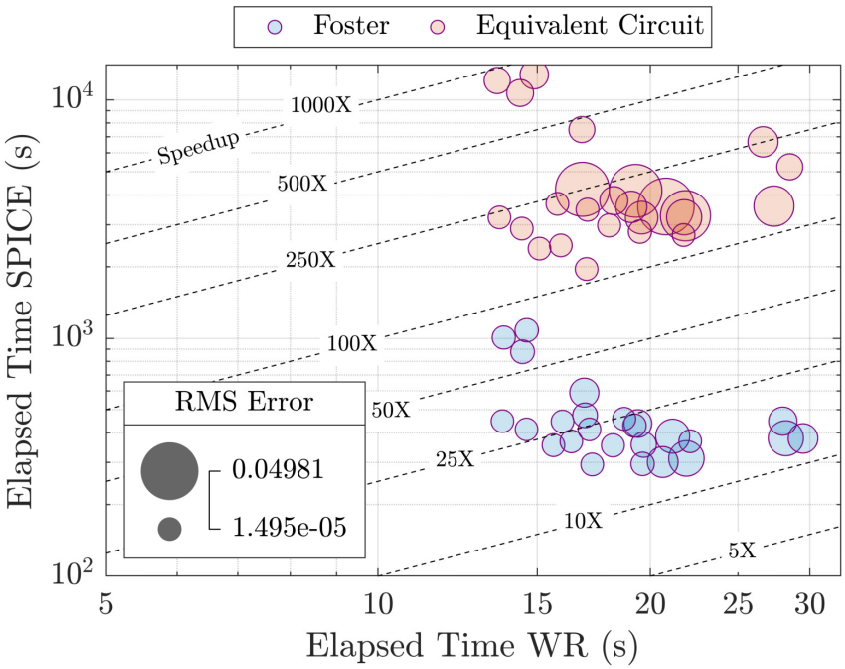

Fig. 4.8 Systematic transient analyses on a $P = 100$ nonlinearly loaded shielding enclosure changing incident field, center frequency and amplitude. Details in the text. From [19] © 2022 IEEE.

accuracy threshold $\varepsilon = \tau_a = \tau_r = 10^{-6}$ for approximately $K \approx 7000$ time samples and adapting the integration time step to the field center frequency as $\Delta t = 0.001/f_c$. The latter setting is necessary to ensure a suitable description of the open-circuit voltages $\boldsymbol{v}_{oc}$ that are used as excitation sources in the equivalent circuit problem formulation, see Fig 1.3.

Figure 4.8 compares runtime, accuracy and speed-up of the presented WR scheme with HSPICE. This illustration identifies each simulation with a circle, whose size specifies the worst-case RMS error between port voltages computed with the two solvers.

Figure 4.9 reports the transient voltages that correspond to the two largest circles, i.e. the simulation test cases with the worst agreement between HSPICE and the WR signals. It is evident from these results that the WR and Foster solvers always provide the same solution, while the equivalent circuit outcomes are the ones responsible for the dominant RMS errors of Fig. 4.8, confirming the observations of Section 4.6.2.

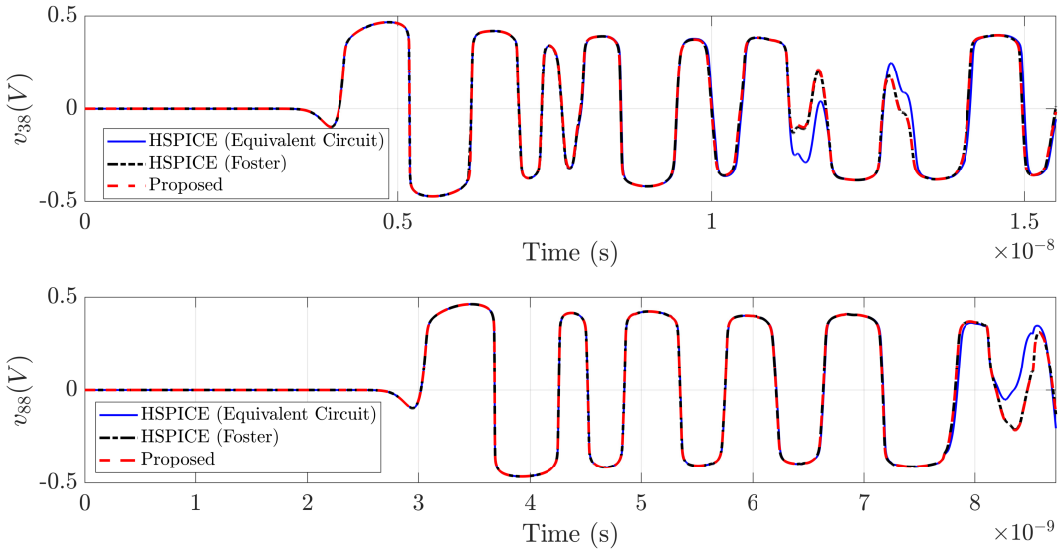

Fig. 4.9 Systematic analysis results: port voltages with worst-case RMS errors between WR and HSPICE (Foster or equivalent circuit representation) among all simulations of Fig. 4.8. From [19] © 2022 IEEE.

On the other hand, Fig. 4.8 also demonstrates that the main advantage of the presented WR iteration is provided in terms of speed-up, ranging from $10\times$ to $1000\times$ in the case of HSPICE-Foster and equivalent circuit form, respectively. From an average perspective, the speed-up factor is bounded between $10-25\times$ running Foster-based simulations and $100-250\times$ embedding in HSPICE an equivalent circuit representation of the model.

4.6.4 Scaling with system complexity

This last section concludes the numerical results by assessing the scalability of the presented transient solver when the number of electrical ports (and nonlinear terminations) reaches thousands of elements. To this end, the largest shielding enclosure model with $P = 1024$ ports already documented in Chapter 3 is here included in the full WR simulation setting, providing a comparison with the two other structures reported in this Chapter with $P = 100$ and 400 ports. Thus, these three shielding enclosures are excited with the same incident electrical field pulse defined as in (4.37), with amplitude $E = 10^4$ and centered at $f_c = 400$MHz.

Ports	HSPICE (Foster) Time	Proposed Time (s)	RMS Error	Speed-Up Factor
100	306.5 s	8.2	$2.26 \cdot 10^{-3}$	37×
400	76 min	112.7	$3.67 \cdot 10^{-3}$	41×
1024	6.4 h	257.8	$1.87 \cdot 10^{-3}$	89×

Table 4.2 Scaling with ports: transient solvers comparison. From [19] © 2022 IEEE.

Table 4.2 provides a comparison of the resulting transient simulation outcomes for both HSPICE and the WR solver. Figure 4.10 and Fig. 4.11 illustrate a set of representative port voltages for the 400 and 1024-port enclosures, respectively. Both cases demonstrate a high level of accuracy of the presented WR scheme with respect to the HSPICE solution, with a worst-case RMS deviation of $3.67 \cdot 10^{-3}$ and $1.88 \cdot 10^{-3}$, respectively. Note that Table 4.2 confirms both that the RMS error is close to 10^{-3} for all cases, and the presented approach is likely to scale favorably with the model size since the speed-up factor increases with the port count, reaching a maximum of 89×.

To conclude, two additional observations are in order. First, for this investigation only the HSPICE-Foster configuration has been exploited, which is very likely to represent the best possible setting for this engine, i.e. the fastest way to obtain a solution with a SPICE-based solver. Second, our prototype WR implementation can be efficiently used to solve the largest benchmark with 1024-ports on a standard notebook with 16GB of RAM and a Core i5-10210U CPU. In this case, the MATLAB environment would require approximately 3 minutes to solve a simulation problem that cannot be solved with the standard HSPICE engine due to excessive memory requirements.

4.7 Conclusions

This Chapter presented an iterative transient solver for the time-domain analysis of large-scale systems terminated with (static) nonlinear loads. The provided strategy belongs to the family of Waveform Relaxation (WR) approaches and constructs a suitable decoupling framework on top of a macromodeling scheme. The provided approach enables us to:

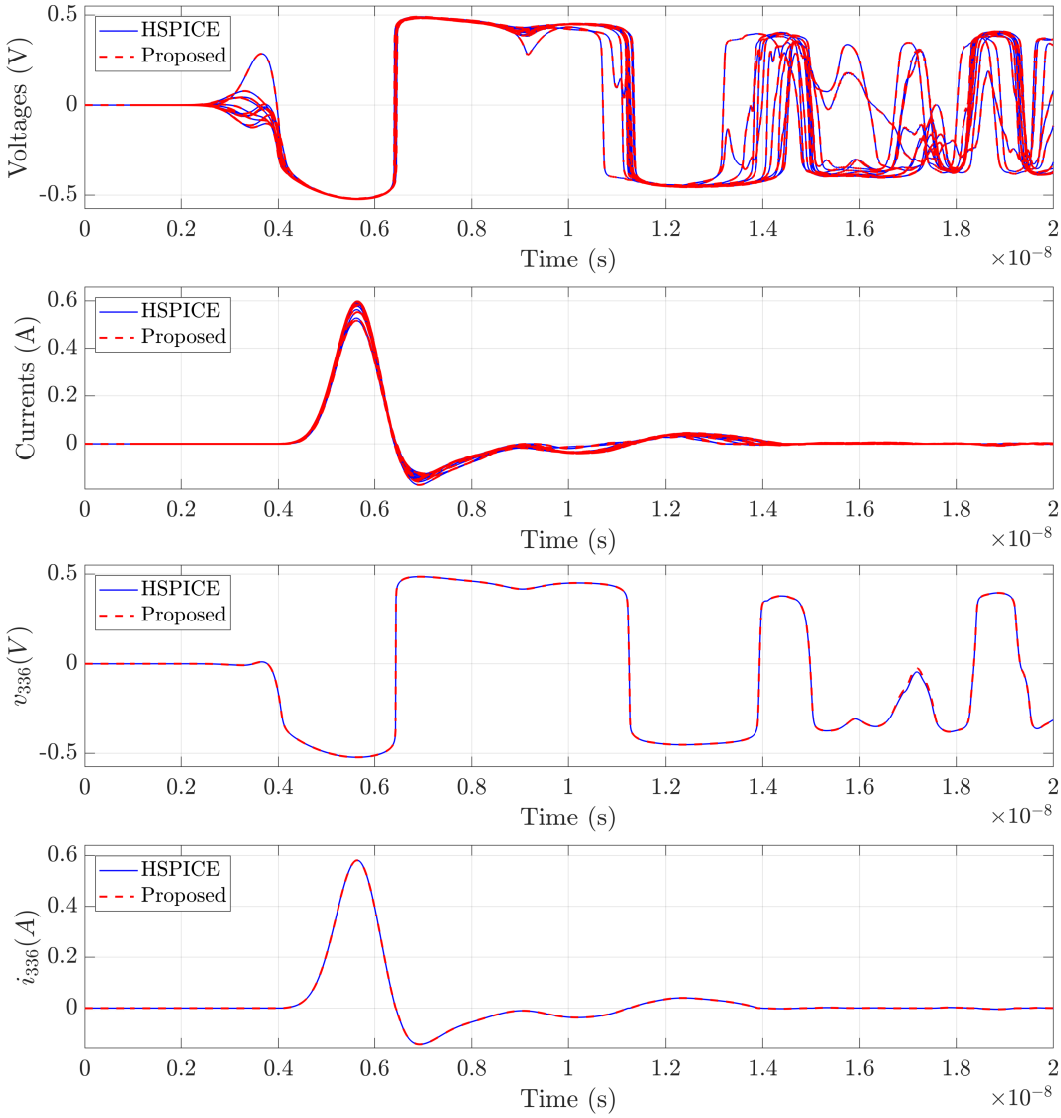

Fig. 4.10 Transient analysis of a nonlinearly loaded shielding enclosure with 400-ports. Top panels illustrate a representative set of ports voltages and currents, while the last two focus on the port 336, which shows the worst-case RMS error $3.67 \cdot 10^{-3}$ among all responses. From [19] © 2022 IEEE.

1. Decouple the simulation of the (linear) macromodel of the electromagnetic structure and the evaluation of nonlinear devices, which can be performed separately due to a standard Longitudinal Partitioning (LP) approach [39], see Section 4.2.

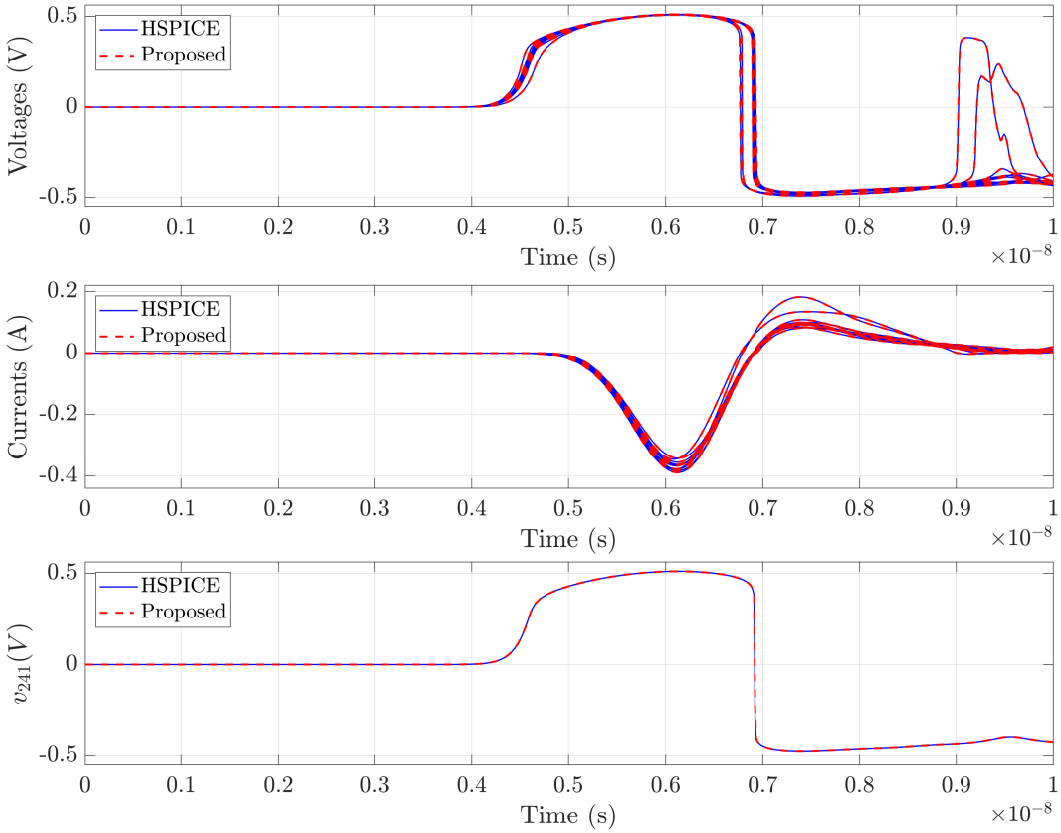

Fig. 4.11 Transient analysis of a nonlinearly loaded shielding enclosure with 1024-ports. As in Fig. 4.10, but only voltage at port 241 is reported (worst-case RMS error of $1.87 \cdot 10^{-3}$). From [19] © 2022 IEEE.

2. Control (and limit) the memory requirements of the individual WR iterations with a suitable segmentation of the simulation window [36, 38], see Sec. 4.3.

3. Overcome the well-known lack of (or poor) convergence of the WR decoupling method defining an inexact Newton-Krylov step [46], see Sec. 4.4.

4. Improve the overall convergence of the scheme, thanks to a fully integrated environment that combines the time partitioning (windowing) scheme, the formulation of the problem based on a quasi-Newton iteration, and a suitable restarting procedure of the WR evaluations operators, see Sec. 4.5.

5. Outperform prior state-of-the-art iterative solver methods belonging to the same WR class, as well as standard circuit solvers like HSPICE, both in terms of accuracy and runtime, see Section. 4.6.

6. Formulate a general framework based on surrogate models, which is likely to scale favorably with the number of electrical ports and that ultimately enables a divide-and-conquer macromodeling approach to solve highly complex and coupled linear/nonlinear circuit problems.

Numerical results support the above conclusions illustrating transient analyses of nonlinearly loaded shielding enclosures structures with up to 1024 electrical ports.

Future investigations will be devoted to addressing the remaining open issues of the presented strategy. From an implementation standpoint, the algorithm can be improved by embedding an adaptive time partitioning approach [36] and exploiting parallelization both in terms of windows [38, 93] and sub-systems evaluations. On the other hand, the framework validation is currently limited to static nonlinear terminations. Nevertheless, the presented formulation can be extended with minor modification to deal with more complex and general nonlinear components by embedding proper Model Order Reduction (MOR) techniques [94, 95] to describe dynamic nonlinearities and to obtain a fast evaluation of the nonlinear sub-system. To conclude, the extension of the presented solver to other structures beyond the proposed shielding enclosure application is expected to be applicable with minor modifications.

Chapter 5

Conclusions

This dissertation investigated the numerical challenges of modeling and simulating energy selective enclosures. The particular features and challenges induced by this application enabled the development of data pre-processing, modeling and simulation strategies for large-scale LTI systems coupled with nonlinear devices. The final result of this thesis is a *general* and robust integrated framework that starts from tabulated frequency data or measurements of an unloaded large-scale shielding enclosure, delivers a compact model and validates the performances of the nonlinearly loaded structure in the time domain with an efficient simulation strategy. Several numerical results supported the effectiveness of all steps in the presented strategies. The main contributions of this thesis are summarized in the following.

Large-scale data preprocessing Chapter 2 presented several data preprocessing methods for improving and/or enabling successive macromodeling steps. At first, we compensated for an initial incomplete frequency characterization of the structure due to full-wave field solver limitations. To fill the low-frequency (LF) gap with well-defined synthetic responses, we designed an extrapolation and regularization procedure in a suitable asymptotic modal domain. Then, we introduced customized compression strategies to reduce the data complexity while preserving LF data features. These approaches, code-named Hierarchical SVD (Hi-SVD) and Block-Diagonal SVD (BD-SVD), drastically improved the full-bad accuracy during reconstruction, down to DC.

Structured and compressed macromodeling framework In Chapter 3, we introduced a macromodeling framework suitable for the shielding application, where the resulting models are large-scale both in terms of dynamic order (i.e. the number of poles) and input/output ports. At first, we reviewed rational approximation strategies, we introduced the VF algorithm and the classic Compressed Macromodel (CM) framework based on a standard SVD, and we provided background notions on the passivity characterization. Then, we presented a compressed and structured macromodeling strategy to embed the data-reduction methods of Chapter 2 in a suitable model representation that inherited low-frequency data features, such as DC zeros of various orders. We equipped the proposed strategy with a sampling-based passivity verification algorithm and a consistent enforcement scheme to ensure the certification of a final passive model.

Transient analysis via WR-based decoupling scheme In Chapter 4, we focused on the transient analysis of nonlinearly loaded shielding enclosures. To this end, we adopted a hybrid simulation approach that converted the fully-coupled electromagnetic problem into an equivalent circuit formulation. This strategy decoupled the linear/nonlinear sub-problems and enabled both a suitable macromodeling framework and a Waveform Relaxation (WR) approach. In this context, we proposed a simulation scheme that improved the convergence rate of the standard WR solver by combining a time windowing strategy and an inexact Newton-Krylov iteration. We validated the presented scheme on shielding enclosures macromodels with up to 1024 ports and 72704 states, demonstrating the effectiveness and reliability of the proposed solution with respect to state-of-the-art SPICE engines.

5.1 Open investigations

The strategies provided through this work open the investigation of further research directions to develop a more general, robust and reliable framework for the next generation's applications.

Data extrapolation of lossy system The extrapolation procedure presented in Chapter 3 assumed a low-loss asymptotic behavior. Nevertheless, a more general approach may be needed in the case of lossy systems. Indeed, the proposed strategy applies with no relevant modifications if the real component of the responses in the low-frequency range is sufficiently small to be approximated as null. This condition is not always satisfied in the case of significant losses. In this context, there is no guarantee that the low-frequency asymptotic behavior assumed in (2.9) still holds. Thus, two possible approaches can be investigated. First, we can improve the low-frequency approximation (2.9) by including a resistive frequency-dependent block and extend the component-wise fit (2.12) also to the real part of each response. Alternatively, we could generate a synthetic low-loss dataset by subtracting the contribution of the DC matrix (when known) from all frequencies and apply the proposed extrapolation procedure as is. These alternative approaches are left as future investigations.

Parallelization of the passivity verification algorithm A parallel implementation of the passivity verification scheme presented in Chapter 3 would improve the main bottleneck of the macromodeling procedure, hence the passivity enforcement phase. Indeed, Stage 2 (Section 3.3.2) of the proposed scheme naturally fits the parallel computation paradigms since the tree-based adaptive sampling scheme runs on each subband independently. Thus, we expect to provide a drastic speed-up by tackling the search effort on a multicore or GPU architecture. Nevertheless, this implementation requires careful consideration in terms of resources. On the one hand, the memory requirement is affected by the macromodel size, whose coefficients are necessary to compute the system responses for all subbands. Thus, the $P \times P$ residue matrices should be streamed among the parallel architecture. On the other hand, the required sample distribution among subbands is not uniform, and some frequency intervals will be more expensive and time-consuming than others. Thus, the precise impact of a parallel framework cannot be estimated a priori. We might estimate an ideal speedup condition (never achievable in practice), corresponding to a cut of the adaptive sampling cost by a factor of N_T (the number of computing threads), in case the number of subbands is an integer multiple of N_T, and in case the cost for adaptive subsampling each subband is identical (balanced). If

these conditions are not met, the speedup would be dominated by the slowest subband requiring more internal samples.

Optimization of transient simulation via WR The presented WR implementation still leaves open several future research directions. First, we expect to improve the algorithm efficiency at each iteration with a suitable multicore or GPU-based parallelization [38, 93]. On the other hand, in the direction of a better convergence rate, an automatic or dynamic time windowing strategy can be adopted. From the formulation standpoint, the presented application is limited to static nonlinearities, and a suitable extension to more advanced and general nonlinear devices is needed. To conclude, the presented solver can be modified to take advantage of possible time-varying or spatial redundancies in the decoupling signals, thus using a projection-based approach to compress the dimension of exchanged signals.

Extension to standard applications The last open question is related to the application side. Through this work, we mostly presented results associated with the shielding enclosure case. Nevertheless, the presented strategies are general and applicable to a wide range of problems. We recall that the standard WR solver is already proven to be effective for transmission lines (TLs) [39] and Power Distribution Networks (PDNs) [38] analysis. Thus, we expect the proposed framework to perform even better due to the combination of a suitable compressed macromodeling approach. The validation of the presented strategies on standard applications, such as TLs, PDNs and PCBs, is thus considered a future research activity of great interest.

Acronyms

BD-SVD Block-Diagonal SVD. 15, 38, 43–50, 58, 70, 102, 103, 144

BR Bounded Real. 62

BRL Bounded Real Lemma. 63, 65

CM Compressed Macromodel. 15, 38, 55, 59–61, 133, 145

EM Electromagnetic. iv, 2, 3, 5–7, 18, 110, 113, 114

EMC Electromagnetic Compatibility. iv

EMI Electromagnetic Interference. iv, 2

FDTD Finite-Difference Time-Domain Method. 6

FEM Finite Elements Method. 9

GMRES Generalized Minimal RESidual. 126–128, 132, 133

Hi-SVD Hierarchical SVD. 15, 38, 43, 45–48, 58, 144

HIRF High-Intensity Radiated Field. iv, 2

KYP Kalman-Yakubovich-Popov. 63

LMI Linear Matrix Inequality. 54, 66, 67, 90

LP Longitudinal Partitioning. 116–119, 123–128, 132, 141

LS least squares. 60, 61

LTI Linear Time Invariant. iv, 2, 7, 8, 16, 36, 38, 52–55, 62, 63, 98, 110, 144

MIMO Multi Input Multi Output. 62

MNA Modified Nodal Analysis. 12

MoM Method of Moments. 8–10, 17–19, 22, 24, 26, 33–36, 96, 98, 102, 106, 112, 115, 131

MOR Model Order Reduction. 11, 143

NL nonlinear. 2

NMSO Naive Multi-scale Search Optimization. 78, 82, 85, 86, 101

OC open-circuit. 5, 6

ODE Ordinary Differential Equations. 12, 116

PCB Printed Circuit Board. 98, 147

PDNs Power Distribution Networks. 147

PEC Perfectly Electrically Conducting. 9, 112, 119

PEEC Partial Element Equivalent Circuit. 6

RMS Root Mean Square. 70, 103, 136, 138–142

RMSE Root Mean Square Error. 50

ROM Reduced Order Model. v

SC short-circuit. 5, 6

SE Shielding Effectiveness. iv, 4–6, 110

SVD Singular Value Decomposition. v, 15, 38–41, 43–49, 51, 58–61, 63, 97, 145

TDFEM Time Domain Finite-Element Method. 6

TDIE Time Domain Integral Equation. 6

TLs transmission lines. 147

VF Vector Fitting. 8, 12, 37, 38, 52–56, 58–62, 69, 70, 88, 92, 96, 98, 103, 145

WR Waveform Relaxation. v, 13–15, 110, 111, 116–140, 142, 143, 145, 147

Appendix A

DC characterization

This appendix provides a step-by-step computation of the impedance matrix $\mathbf{Z}_0$ as in (2.1) obtained with the proposed regularized DC topology presented in Fig. 2.2. The final result is a full characterization of the system DC response, including admittance (2.2) and scattering (2.3) representations, irrespective to the port numbering arrangement.

Since in the proposed regularized topology of Fig. 2.2 identical series resistances r are added at each lumped port of a $p \times p = P$ regular grid, it is straightforward to see that

$$\mathbf{Z}_0 = r\,\mathbb{I}_P + \widehat{\mathbf{Z}} \tag{A.1}$$

where $\widehat{\mathbf{Z}}$ is the multi-port $P \times P$ impedance matrix obtained by removing all resistors r from the circuit of Fig. 2.2. Elements of $\widehat{\mathbf{Z}}$ are computed following their definition, by taking advantage of the regular grid configuration and addressing each branch separately.

The calculation of $\widehat{\mathbf{Z}}$ starts considering an individual $p \times p$ block corresponding to a generic m-th branch, in the following indicated with $\widehat{\mathbf{Z}}_{mm}$ and obtained as in Fig. A.1 by

- exciting all ports located at the m-th grid branch with current sources
- leaving all ports in other branches open.

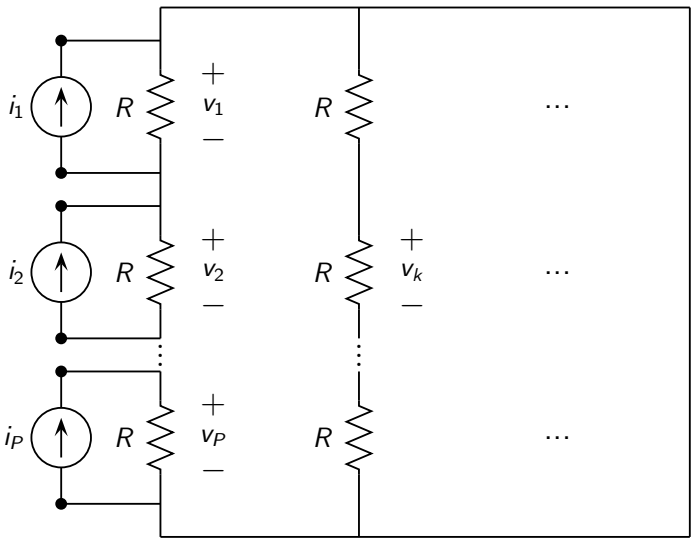

Fig. A.1 Circuit to compute the impedance matrix $\widehat{\mathbf{Z}}_{mm}$ associated to the m-th grid branch. Since all resistances on other branches are short-circuited, each $\boldsymbol{v}_k = 0 \, \forall k$ that does not belong to the m-th branch is vanishing. The trans-impedance matrix block $\widehat{\mathbf{Z}}_{mn}$ with $m \neq n$ is null.

Computing the voltage drop on each port of the m-th branch enables to define diagonal and off-diagonal elements of $\widehat{\mathbf{Z}}_{mm}$ as

$$(\widehat{\mathbf{Z}}_{mm})_{i,i} = \frac{(p-1)R^2}{pR} = \frac{p-1}{p}R \qquad \forall i = 1, \cdots, p \qquad \text{(A.2)}$$

$$(\widehat{\mathbf{Z}}_{mm})_{i,j} = -R\frac{R}{R+(p-1)R} = -\frac{R}{p} \qquad \forall i \neq j. \qquad \text{(A.3)}$$

The resulting $p \times p$ impedance matrix can be defined in a compact form as

$$\widehat{\mathbf{Z}}_{mm} = R \begin{bmatrix} 1-\vartheta & -\vartheta & \cdots & -\vartheta \\ -\vartheta & 1-\vartheta & \ddots & \vdots \\ \vdots & \ddots & \ddots & -\vartheta \\ -\vartheta & \cdots & -\vartheta & 1-\vartheta \end{bmatrix} = R(\mathbb{I}_p - \vartheta \boldsymbol{u}\boldsymbol{u}^{\mathsf{T}}) \qquad \text{(A.4)}$$

where $\vartheta = 1/p$, $\boldsymbol{u}^{\mathsf{T}} = [1, \cdots, 1]$, and $\boldsymbol{u} \in \mathbb{R}^p$.

Consider now blocks $\widehat{\mathbf{Z}}_{mn} \in \mathbb{R}^{p\times p}$ collecting trans-impedances between ports located in different branches. Due to the presence of a return path through the box enclosure (the short circuit on the right side of Fig. A.1), all voltages at the m-th branch are vanishing when the current excitation sources are located in another n-th branch, so that the trans-impedance matrix reads

$$\widehat{\mathbf{Z}}_{mn} = \mathbf{0} \quad \forall m \neq n. \tag{A.5}$$

Combining (A.4) and (A.5) with (A.1) to include (regularization) series resistances r, each block in the DC impedance matrix $\mathbf{Z}_0$ associated to a given branch can be defined as

$$\mathbf{Z}'_0 = \mathbf{Z}'_{mm} = r\,\mathbb{I}_p + \widehat{\mathbf{Z}}_{mm} = r\,\mathbb{I}_p + R\,(\mathbb{I}_p - \vartheta \boldsymbol{u}\boldsymbol{u}^{\mathsf{T}}) \qquad m = 1,\ldots,p \tag{A.6}$$

$$\mathbf{Z}'_{mn} = \widehat{\mathbf{Z}}_{mn} = 0 \qquad \forall m \neq n \tag{A.7}$$

Therefore, the complete $P \times P$ DC impedance matrix $\mathbf{Z}_0$ is obtained by collecting all blocks $\mathbf{Z}'_0$ corresponding to individual branches, with an arrangement that depends on the global numbering of the ports. Using a *vertical* numbering (superscript v) the ports are numbered consecutively in each branch and then all branches are listed consecutively, as in the left panel of Fig. A.2. The complementary *horizontal* numbering (superscript h) indicates a global listing of all ports, as depicted in the right panel of Fig. A.2. These two cases correspond to

$$\mathbf{Z}^v_0 = \mathbb{I}_p \otimes \mathbf{Z}'_0, \qquad \mathbf{Z}^h_0 = \mathbf{Z}'_0 \otimes \mathbb{I}_p \tag{A.8}$$

where $\otimes$ indicates the Kronecker product. Using the Sherman-Morrison formula, the DC admittance and scattering matrices with vertical port alignment (superscript v) can be analytically computed as in

$$\mathbf{Y}^v_0 = \mathbb{I}_p \otimes \mathbf{Y}'_0, \qquad \mathbf{Y}'_0 = \frac{(p\,r\mathbb{I}_p + R\boldsymbol{u}\boldsymbol{u}^{\mathsf{T}})}{p(Rr + r^2)} \tag{A.9}$$

$$\mathbf{S}^v_0 = \mathbb{I}_p \otimes \mathbf{S}'_0, \qquad \mathbf{S}'_0 = \frac{\Phi' - 1}{\Phi' + 1}\mathbb{I}_p - \frac{2\vartheta\Phi\boldsymbol{u}\boldsymbol{u}^{\mathsf{T}}}{(\Phi' + 1)(\varphi + 1)} \tag{A.10}$$

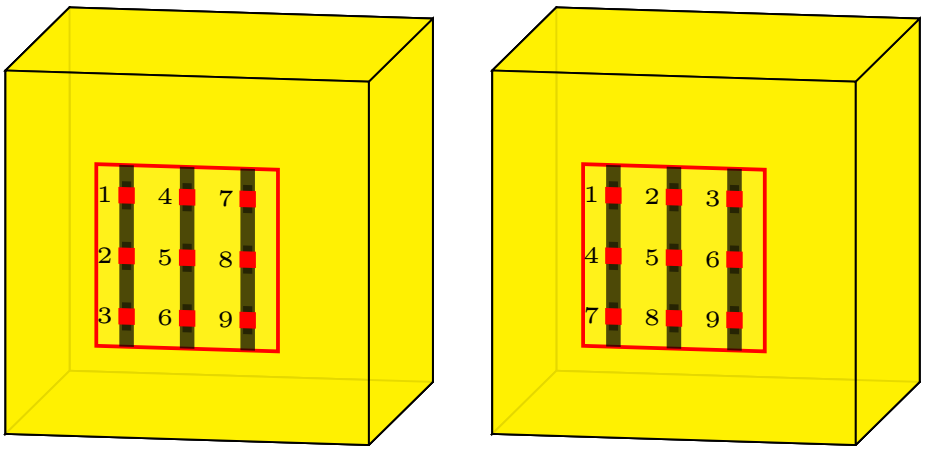

Fig. A.2 Vertical (left) and horizontal (right) ports numbering for a $P = 3 \times 3 = 9$ shielding enclosure.

where $\Phi = R/R_0$, $\varphi = r/R_0$, $\Phi' = \Phi + \varphi$, and R_0 denoting the scattering port reference impedance, usually $50\,\Omega$. For the horizontal port numbering scheme, the order of the Kronecker product factors is reversed as in (A.8).

References

[1] FM Tesche and TK Liu. Transient response of antennas with nonlinear loads. *Electronics Letters*, 11(1):18–19, 1975.

[2] T. Liu and F. Tesche. Analysis of antennas and scatterers with nonlinear loads. *IEEE Transactions on Antennas and Propagation*, 24(2):131–139, 1976.

[3] George V Eleftheriades. Protecting the weak from the strong. *Nature*, 505(7484):490–491, 2014.

[4] A.E. Spezio. Electronic warfare systems. *IEEE Transactions on Microwave Theory and Techniques*, 50(3):633–644, 2002.

[5] Hiroki Wakatsuchi, Sanghoon Kim, Jeremiah J. Rushton, and Daniel F. Sievenpiper. Circuit-based nonlinear metasurface absorbers for high power surface currents. *Applied Physics Letters*, 102(21):214103, 2013.

[6] Hiroki Wakatsuchi. Waveform-selective metasurfaces with free-space wave pulses at the same frequency. *Journal of Applied Physics*, 117(16):164904, 2015.

[7] Sanghoon Kim, Hiroki Wakatsuchi, Jeremiah J. Rushton, and Daniel F. Sievenpiper. Switchable nonlinear metasurfaces for absorbing high power surface waves. *Applied Physics Letters*, 108(4):041903, 2016.

[8] Zhangjie Luo, Xing Chen, Jiang Long, Ryan Quarfoth, and Daniel Sievenpiper. Nonlinear power-dependent impedance surface. *IEEE Transactions on Antennas and Propagation*, 63(4):1736–1745, 2015.

[9] Cheng Yang, Peiguo Liu, and Xianjun Huang. A novel method of energy selective surface for adaptive HPM/EMP protection. *IEEE Antennas Wirel. Propag. Lett.*, 12:112–115, 2013.

[10] Cheng Yang, Torben Wendt, Marco De Stefano, Marc Kopf, Christopher Marc Becker, Stefano Grivet-Talocia, and Christian Schuster. Analysis and optimization of nonlinear diode grids for shielding of enclosures with apertures. *IEEE Transactions on Electromagnetic Compatibility*, 63(6):1884–1895, 2021.

[11] Cheng Yang, Heinz-D. Brüns, Peiguo Liu, and Christian Schuster. Impulse response optimization of band-limited frequency data for hybrid field-circuit simulation of large-scale energy-selective diode grids. *IEEE Transactions on Electromagnetic Compatibility*, 58(4):1072–1080, 2016.

[12] Cheng Yang, Peiguo Liu, Heinz-D Brüns, and Christian Schuster. Design aspects for HIRF protection of a rectangular metallic cavity using energy selective diode grids. In *2016 Asia-Pacific International Symposium on Electromagnetic Compatibility (APEMC)*, volume 1, pages 35–37. IEEE, 2016.

[13] T. Wendt, C. Yang, H. D. Brüns, S. Grivet-Talocia, and C. Schuster. A macromodeling-based hybrid method for the computation of transient electromagnetic fields scattered by nonlinearly loaded metal structures. *IEEE Transactions on Electromagnetic Compatibility*, 62(4):1098–1110, Aug 2020.

[14] K. Aygun, B.C. Fischer, Jun Meng, B. Shanker, and E. Michielssen. A fast hybrid field-circuit simulator for transient analysis of microwave circuits. *IEEE Transactions on Microwave Theory and Techniques*, 52(2):573–583, 2004.

[15] Haiyan Xie, Jianguo Wang, Ruyu Fan, and Yinong Liu. A hybrid FDTD-SPICE method for transmission lines excited by a nonuniform incident wave. *IEEE Transactions on Electromagnetic Compatibility*, 51(3):811–817, 2009.

[16] Rui Wang and Jian-Ming Jin. Incorporation of multiport lumped networks into the hybrid time-domain finite-element analysis. *IEEE Transactions on Microwave Theory and Techniques*, 57(8):2030–2037, 2009.

[17] Sohrab Safavi and Jonas Ekman. A hybrid peec–spice method for time-domain simulation of mixed nonlinear circuits and electromagnetic problems. *IEEE Transactions on Electromagnetic Compatibility*, 56(4):912–922, 2014.

[18] Torben Wendt, Marco De Stefano, Cheng Yang, Stefano Grivet-Talocia, and Christian Schuster. Iteration dependent waveform relaxation for hybrid field nonlinear circuit problems. *IEEE Transactions on Electromagnetic Compatibility, accepted*, 2022.

[19] Marco De Stefano, Torben Wendt, Cheng Yang, Stefano Grivet-Talocia, and Christian Schuster. A waveform relaxation solver for transient simulation of large-scale nonlinearly loaded shielding structures. In *IEEE Transactions on Electromagnetic Compatibility, accepted*, 2022.

[20] S. Grivet-Talocia and B. Gustavsen. *Passive Macromodeling: Theory and Applications*. John Wiley and Sons, New York, 2016 (published online on Dec 7, 2015).

[21] B. Gustavsen and A. Semlyen. Rational approximation of frequency domain responses by vector fitting. *Power Delivery, IEEE Transactions on*, 14(3):1052–1061, jul 1999.

[22] Yuji Nakatsukasa, Olivier Sète, and Lloyd N. Trefethen. The aaa algorithm for rational approximation. *SIAM Journal on Scientific Computing*, 40(3):A1494–A1522, 2018.

[23] S. Lefteriu and A. C. Antoulas. A new approach to modeling multiport systems from frequency-domain data. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 29(1):14–27, jan. 2010.

[24] D. Deschrijver, M. Mrozowski, T. Dhaene, and D. De Zutter. Macromodeling of multiport systems using a fast implementation of the vector fitting method. *Microwave and Wireless Components Letters, IEEE*, 18(6):383–385, june 2008.

[25] A. Chinea and S. Grivet-Talocia. On the parallelization of vector fitting algorithms. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 1(11):1761–1773, November 2011.

[26] Srinidhi Ganeshan, Naveen Kumar Elumalai, Ramachandra Achar, and Wai Kong Lee. Gvf: Gpu-based vector fitting for modeling of multiport tabulated data networks. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 10(8):1375–1387, 2020.

[27] Chiu-Chih Chou and José E. Schutt-Ainé. Equivalent circuit synthesis of multiport s parameters in pole–residue form. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 11(11):1971–1979, 2021.

[28] HSPICE, Synopsis [online] `http://www.synopsis.com`.

[29] Spectre, Cadence [online] `http://www.cadence.com`.

[30] Marco De Stefano, Torben Wendt, Cheng Yang, Stefano Grivet-Talocia, and Christian Schuster. Regularized and compressed large-scale rational macromodeling: Theory and application to energy-selective shielding enclosures. In *IEEE Transactions on Electromagnetic Compatibility*, pages 1–15, 2022.

[31] CONCEPT-II, 2019.

[32] Stefano Grivet-Talocia and Flavio Canavero. Dc-compliant macromodels based on the method of characteristics for frequency-dependent transmission lines. In *2006 1st Electronic Systemintegration Technology Conference*, volume 1, pages 56–61, 2006.

[33] IdEM R2018, Dassault Systèmes.

[34] S. B. Olivadese and S. Grivet-Talocia. Compressed passive macromodeling. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 2(8):1378–1388, August 2012.

[35] E. Lelarasmee, A.E. Ruehli, and A.L. Sangiovanni-Vincentelli. The waveform relaxation method for time-domain analysis of large scale integrated circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 1(3):131–145, 1982.

[36] Jacob K. White and Alberto Sangiovanni-Vincentelli. *Relaxation Techniques for the Simulation of VLSI Circuits*, volume 20 of *The Kluwer International Series in Engineering and Computer Science*. Springer US, Boston, MA, 1986.

[37] M.J. Gander and A.E. Ruehli. Optimized waveform relaxation methods for rc type circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 51(4):755–768, 2004.

[38] Ramachandra Achar, Michel S. Nakhla, Harjot S. Dhindsa, Arvind R. Sridhar, Douglas Paul, and Natalie M. Nakhla. Parallel and Scalable Transient Simulator for Power Grids via Waveform Relaxation (PTS-PWR). *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 19(2):319–332, feb 2011.

[39] V. Loggia, S. Grivet-Talocia, and H. Hu. Transient simulation of complex high-speed channels via waveform relaxation. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 1(11):1823–1838, November 2011.

[40] Tarik Menkad and Anestis Dounavis. Convergence of the resistive coupling-based waveform relaxation method for chains of identical and symmetric circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 68(12):5120–5133, 2021.

[41] Martin J. Gander, Mohammad Al-Khaleel, and Albert E. Ruchli. Optimized waveform relaxation methods for longitudinal partitioning of transmission lines. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 56(8):1732–1743, 2009.

[42] Tarik Menkad and Anestis Dounavis. Using strictly dissipative impedance coupling in the waveform relaxation method for the analysis of interconnect circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 68(3):1283–1296, 2021.

[43] Jaeyong Cho, Jangyong Ahn, Jongwook Kim, Jaehyoung Park, Yujun Shin, Kibeom Kim, Junsung Choi, and Seungyoung Ahn. Low- and high-frequency extrapolation of band-limited frequency responses to extract delay causal time responses. *IEEE Transactions on Electromagnetic Compatibility*, 63(3):888–901, 2021.

[44] M. De Stefano, S. Grivet-Talocia, T. Wendt, C. Yang, and C. Schuster. Low-frequency modal extrapolation and regularization for full-bandwidth macromodeling of electromagnetic structures. In *2022 IEEE 26th Workshop on Signal and Power Integrity (SPI)*, pages 1–4, 2022.

[45] Marco De Stefano, Stefano Grivet-Talocia, Torben Wendt, Cheng Yang, and Christian Schuster. A multistage adaptive sampling scheme for passivity characterization of large-scale macromodels. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 11(3):471–484, March 2021.

[46] S. B. Olivadese, S. Grivet-Talocia, C. Siviero, and D. Kaller. Macromodel-based iterative solvers for simulation of high-speed links with nonlinear terminations. *Components, Packaging and Manufacturing Technology, IEEE Transactions on*, 4(11):1847–1861, Nov 2014.

[47] P. Triverio, S. Grivet-Talocia, M. S. Nakhla, F. Canavero, and R. Achar. Stability, causality, and passivity in electrical interconnect models. *IEEE Trans. Advanced Packaging*, 30(4):795–808, Nov 2007.

[48] P. Triverio and S. Grivet-Talocia. Causality-constrained interpolation of tabulated frequency responses. In *IEEE* 15^{th} *Topical Meeting on Electrical Performance of Electronic Packaging (EPEP 2006), Scottsdale, Arizona*, pages 181–184, October 23–25, 2006.

[49] Tamara G. Kolda and Brett W. Bader. Tensor Decompositions and Applications. *SIAM Review*, 51(3):455–500, aug 2009.

[50] S. Ganeshan, N. K. Elumalai, R. Achar, and W. K. Lee. Gvf: Gpu-based vector fitting for modeling of multiport tabulated data networks. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 10(8):1375–1387, 2020.

[51] A. Ubolli, S. Grivet-Talocia, M. Bandinu, and A. Chinea. Sensitivity-based weighting for passivity enforcement of linear macromodels in power integrity applications. In *DATE 2014 - Design, Automation and Test in Europe, Dresden, Germany*, pages 1–6, March 24–28, 2014.

[52] J. C. Willems. Dissipative dynamical systems part I: General theory. *Archive for Rational Mechanics and Analysis*, 45(5):321–351, 1972.

[53] M. R. Wohlers. *Lumped and Distributed Passive Networks*. Academic press, 1969.

[54] B. D. O. Anderson and S. Vongpanitlerd. *Network analysis and synthesis*. Prentice-Hall, 1973.

[55] C. P. Coelho, J. Phillips, and L. M. Silveira. A convex programming approach for generating guaranteed passive approximations to tabulated frequency-data. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 23(2):293 – 301, feb. 2004.

[56] D. Saraswat, R. Achar, and M. S. Nakhla. Fast passivity verification and enforcement via reciprocal systems for interconnects with large order macromodels. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 15(1):48–59, Jan 2007.

[57] D. Saraswat, R. Achar, and M. S. Nakhla. Global passivity enforcement algorithm for macromodels of interconnect subnetworks characterized by tabulated data. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 13(7):819–832, July 2005.

[58] C. S. Saunders, J. Hu, C. E. Christoffersen, and M. B. Steer. Inverse singular value method for enforcing passivity in reduced-order models of distributed structures for transient and steady-state simulation. *Microwave Theory and Techniques, IEEE Transactions on*, 59(4):837–847, April 2011.

[59] T. Brull and C. Schroder. Dissipativity enforcement via perturbation of para-Hermitian pencils. *Circuits and Systems I: Regular Papers, IEEE Transactions on*, 60(1):164–177, Jan 2013.

[60] S. Gao, Y.-S. Li, and M.-S. Zhang. An efficient algebraic method for the passivity enforcement of macromodels. *Microwave Theory and Techniques, IEEE Transactions on*, 58(7):1830–1839, July 2010.

[61] T. D'haene and R. Pintelon. Passivity enforcement of transfer functions. *Instrumentation and Measurement, IEEE Transactions on*, 57(10):2181–2187, Oct 2008.

[62] B. Porkar, M. Vakilian, R. Iravani, and S. M. Shahrtash. Passivity enforcement using an infeasible-interior-point primal-dual method. *Power Systems, IEEE Transactions on*, 23(3):966–974, Aug 2008.

[63] T. Wang and Z. Ye. Robust passive macro-model generation with local compensation. *Microwave Theory and Techniques, IEEE Transactions on*, 60(8):2313–2328, Aug 2012.

[64] S. Grivet-Talocia. Passivity enforcement via perturbation of Hamitonian matrices. *IEEE Trans. Circuits and Systems I: Fundamental Theory and Applications*, 51(9):1755–1769, September 2004.

[65] S. Grivet-Talocia and A. Ubolli. On the generation of large passive macromodels for complex interconnect structures. *IEEE Trans. Advanced Packaging*, 29(1):39–54, February 2006.

[66] S. Grivet-Talocia. An adaptive sampling technique for passivity characterization and enforcement of large interconnect macromodels. *IEEE Trans. Advanced Packaging*, 30(2):226–237, May 2007.

[67] S. Grivet-Talocia and A. Ubolli. A comparative study of passivity enforcement schemes for linear lumped macromodels. *IEEE Trans. Advanced Packaging*, 31(4):673–683, Nov 2008.

[68] S. Grivet-Talocia. On driving non-passive macromodels to instability. *International Journal of Circuit Theory and Applications*, 37(8):863–886, Oct 2009.

[69] A. Chinea, S. Grivet-Talocia, S. B. Olivadese, and L. Gobbato. High-performance passive macromodeling algorithms for parallel computing platforms. *IEEE Transactions on Components, Packaging, and Manufacturing Technology*, 3(7):1188–1203, July 2013.

[70] Bjørn Gustavsen. Passivity assessment and enforcement utilizing eigenpairs information. *Electric Power Systems Research*, 194:107041, 2021.

[71] Bjørn Gustavsen. Passivity enforcement by residue perturbation via constrained non-negative least squares. *IEEE Transactions on Power Delivery*, 36(5):2758–2767, 2021.

[72] A. C. Antoulas. *Approximation of large-scale dynamical systems*. Society for Industrial and Applied Mathematics, 2005.

[73] B. Gustavsen. Improving the pole relocating properties of vector fitting. *Power Delivery, IEEE Transactions on*, 21(3):1587–1592, july 2006.

[74] S. Lefteriu and A. C. Antoulas. On the convergence of the vector-fitting algorithm. *Microwave Theory and Techniques, IEEE Transactions on*, 61(4):1435–1443, 2013.

[75] Chiu-Chih Chou and José E. Schutt-Ainé. On the acceleration of the vector fitting for multiport large-scale macromodeling. *IEEE Microwave and Wireless Components Letters*, 31(1):1–4, 2021.

[76] B. Gustavsen. Computer code for rational approximation of frequency dependent admittance matrices. *Power Delivery, IEEE Transactions on*, 17(4):1093–1098, oct 2002.

[77] E. G. Gilbert. Controllability and observability in multivariable control systems. *Journal of the Society for Industrial & Applied Mathematics, Series A: Control*, 1(2):128–151, 1963.

[78] C. Scherer and S. Weiland. Linear matrix inequalities in control. *Lecture Notes, Dutch Institute for Systems and Control, Delft, The Netherlands*, 2000.

[79] S. Boyd, V. Balakrishnan, and P. Kabamba. A bisection method for computing the H_∞ norm of a transfer matrix and related problems. *Mathematics of Control, Signals and Systems*, 2(3):207–219, 1989.

[80] Z. Ye, L. M. Silveira, and J. R. Phillips. Extended Hamiltonian pencil for passivity assessment and enforcement for S-parameter systems. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2010*, pages 1148–1152, March 2010.

[81] Lieven Vandenberghe, Venkataramanan Balakrishnan, Ragnar Wallin, Anders Hansson, and Tae Roh. *Interior-Point Algorithms for Semidefinite Programming Problems Derived from the KYP Lemma*, volume 312, pages 579–579. Springer, Berlin, Heidelberg, 05 2005.

[82] H. Chen and J. Fang. Enforcing bounded realness of S parameter through trace parameterization. In *Electrical Performance of Electronic Packaging, 2003*, pages 291–294, Oct 2003.

[83] Abdullah Al-Dujaili and S. Suresh. A Naive multi-scale search algorithm for global optimization problems. *Information Sciences*, 372:294–312, dec 2016.

[84] Donald R Jones, Cary D Perttunen, and Bruce E Stuckman. Lipschitzian optimization without the lipschitz constant. *Journal of optimization Theory and Applications*, 79(1):157–181, 1993.

[85] Rémi Munos. Optimistic optimization of a deterministic function without the knowledge of its smoothness. *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011, NIPS 2011*, pages 783–791, 2011.

[86] Rémi Munos. The optimistic principle applied to games, optimization and planning: Towards foundations of monte-carlo tree search. *Foundations and Trends in Machine Learning*, 7(1):1–130, 2014.

[87] Ziyu Wang, Babak Shakibi, Lin Jin, and Nando Freitas. Bayesian multi-scale optimistic optimization. In *Artificial Intelligence and Statistics*, pages 1005–1014. PMLR, 2014.

[88] Kenji Kawaguchi, Leslie P Kaelbling, and Tomás Lozano-Pérez. Bayesian optimization with exponential convergence. *Advances in neural information processing systems*, 28, 2015.

[89] The Mathworks, Inc., `www.mathworks.com`. *MATLAB User's Guide*.

[90] S. Müller, F. Happ, X. Duan, R. Rimolo-Donadio, H. Brüns, and C. Schuster. Complete modeling of large via constellations in multilayer printed circuit boards. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 3(3):489–499, 2013.

[91] Jacob White and A. L. Sangiovanni-Vincentelli. Partitioning Algorithms and Parallel Implementations of Waveform Relaxation Algorithms for Circuit Simulation., 1985.

[92] Carl T Kelley. *Solving nonlinear equations with Newton's method*. SIAM, 2003.

[93] Sebastian Schöps, Innocent Niyonzima, and Markus Clemens. Parallel-in-time simulation of eddy current problems using parareal. *IEEE Transactions on Magnetics*, 54(3):1–4, 2018.

[94] Saifon Chaturantabut and Danny C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM Journal on Scientific Computing*, 32(5):2737–2764, 2010.

[95] Behzad Nouri and Michel S. Nakhla. Model order reduction of nonlinear transmission lines using interpolatory proper orthogonal decomposition. *IEEE Transactions on Microwave Theory and Techniques*, 66(12):5429–5438, 2018.