International Journal of Applied Engineering & Technology

DECODE-R: DIFFERENTIABLE MULTI-OBJECTIVE CO-DESIGN OF HARMONIC FILTER TOPOLOGY AND CONTROL UNDER THD–LOSS–COST–STRESS TRADE-OFFS

Nelluri Vanajakshi [1*], Dr. M Uma Vani [2] and Dr. R Srinivas Rao [3]
[1]Research Scholar, EEE Department, Jawaharlal Nehru Technological University, Kakinada, Ramanayyapeta, Andhra Pradesh 533003, India.
[2]Professor of EEE Department, Lakireddy Bali Reddy College of Engineering, Mylavaram, Vijayawada, Krishna, Andhra Pradesh 521230, India.
[3]Professor of EEE, Jawaharlal Nehru Technological University, Kakinada, Ramanayyapeta, Andhra Pradesh 533003, India.
[1]vanajakshi.nelluri2506@gmail.com, [2]marreddy.umavani@gmail.com and [3]Srinivas.jntuece@gmail.com

ABSTRACT
Harmonic attenuation in grid-tied voltage-source converters remains difficult because filter topology, passive sizing, controller family, controller tuning, and switching frequency jointly affect total harmonic distortion, loss, cost, and stress under changing grid conditions. DECODE-R is proposed as a differentiable multi-objective co-design framework that searches relaxed filter and control structures under shared robust objectives and feasibility constraints, then extracts realizable discrete candidates for local refinement and electromagnetic transient verification. The framework combines robust scenario-based evaluation of harmonic, dynamic, and loss metrics with physics-informed StressRisk proxies for semiconductors and capacitors so that adverse deployment conditions influence both design ranking and constraint satisfaction. On a 25 kVA three-phase converter benchmark evaluated over 216 uncertainty scenarios and compared against single-tuned and damped double-tuned benchmark families, the balanced representative design DE-B1 achieved 2.38% mean point-of-common-coupling current total harmonic distortion and 738 W mean loss after full-scenario rescoring, and electromagnetic transient confirmation yielded 2.53% total harmonic distortion while preserving the defined feasibility limits. The results support DECODE-R as a device-level design workflow for navigating total-harmonic-distortion-loss-cost-stress trade-offs under scenario uncertainty.

Keywords: *Harmonic Filters, Multi-Objective Optimization, Co-Design, Differentiable Modeling, Power Quality, Switching Losses, Reliability Proxies, Scenario-Based Robustness.*

1 INTRODUCTION
Harmonic distortion remains a persistent power-quality problem in grid-tied voltage-source converters because harmonic attenuation, converter efficiency, passive-component stress, and closed-loop stability depend on the combined behaviour of the filter network and the control layer. Converter-oriented studies have already shown that modulation choice, switching frequency, and mitigation strategy shape the final harmonic response, while photovoltaic grid-connected studies further show that filter selection and control configuration cannot be treated as independent design steps (Iqbal et al., [1]; Gony et al., [2]). A device-level design workflow therefore requires a coupled view rather than a sequential view.

Existing optimisation literature has broadened harmonic mitigation design, but the strongest emphasis still falls on partial formulations. Network- and planning-oriented studies commonly balance harmonic limits against investment or allocation objectives for tuned passive filters, filter placement, or distributed-generation planning (Abbas et al., [3]; Hsiao et al., [4]; AbuZahew et al., [5]; Abubakar et al., [6]). Robust passive-filter studies and search-space reduction work have improved feasibility handling and computational tractability, yet the design logic often remains centred on filter sizing or topology within a narrower control assumption set (Daramukkala et al., [7]; Karooyee et al., [8]). Joint filter and current-control optimisation has also been explored, which confirms that harmonic attenuation, damping, and operating loss cannot be separated cleanly at the converter interface (Faria et al., [9]).

The unresolved problem is therefore a mixed discrete-continuous device-level co-design task under deployment uncertainty. Filter topology and controller-family choice are discrete, whereas component sizing, controller gains, and switching frequency are continuous. Practical feasibility further imposes power-factor, resonance-placement, thermal, ripple-current, and stability-margin limits that move with grid strength, ambient temperature, component tolerances, operating point, and background harmonics. Control and switching studies reinforce this coupling, because predictive or advanced-modulation choices change both harmonic attenuation and switching-related loss [9]. Distributed and coordinated mitigation studies show a broader system context, but they do not replace the need for an implementable device-level design method that remains robust before network-level coordination is considered (Rahman et al., [10]; Wang et al., [11]). The present scope is therefore limited deliberately to reduced-model, scenario-robust device-level co-design with EMT verification handoff rather than full network planning, full lifetime prediction, or hardware demonstration.

DECODE-R addresses this gap through a unified multi-objective co-design pipeline that jointly searches relaxed filter topology, continuous filter parameters, controller family, controller tuning, and switching frequency under shared robust objectives and shared robust constraints. The method couples a differentiable filter supergraph with a controller super-structure, embeds scenario-based aggregation directly into objective and feasibility definitions, introduces physics-informed StressRisk proxies for semiconductor and capacitor stress, and produces realizable discrete candidates through discretisation, local refinement, and EMT verification handoff. Novelty does not arise from isolated use of relaxed topology search, controller tuning, or robust multi-objective optimisation alone; novelty arises from integration of these elements in one device-level differentiable pipeline that maps deployment-oriented trade-offs to implementable filter-control designs. Comparative evaluation is then carried out against benchmark-compatible tuned-filter baselines under a common scenario set and a common rescoring protocol.

2 RELATED WORK

Multi-objective optimisation has become a standard way to manage harmonic mitigation trade-offs, but a structural difference remains between planning-oriented formulations and device-level converter-interface design. Planning studies on tuned passive filters and wider distribution-level harmonic mitigation usually balance harmonic indices against investment, siting, or allocation objectives, often under feeder-level constraints and multiple candidate locations (Hsiao et al., [4]; AbuZahew et al., [5]; Yin et al., [12]; Hamza et al., [13]). These studies are valuable for network decisions, yet the resulting design variables and feasibility logic are not identical to converter-interface co-design where resonance placement, switching frequency, current ripple, and controller interaction must be handled directly at the device boundary.

Within the passive-filter literature, recent work has extended beyond single-tuned structures toward damped and double-tuned designs, with explicit treatment of harmonic rejection, resonance behaviour, and design trade-offs under practical constraints (Lin et al., [14]; Alhaider et al., [15]; Daramukkala et al., [7]). The closest device-side stream is inverter output-filter optimisation, where component sizing is linked to feasibility, volume or cost proxies, and computational tractability. Karooyee et al., [8] are especially relevant in this respect because search-space reduction is treated as a central requirement for fast multi-objective optimisation of inverter LCL filters. That contribution clarifies an important limit of direct brute-force search at the converter interface: structural flexibility quickly enlarges the feasible-set management problem.

The main limitation of this stream is that filter topology and parameter optimisation are still often treated with a fixed or only lightly varied controller assumption. Device-level harmonic mitigation is therefore improved, but the optimisation loop does not usually carry controller-family choice, controller tuning, switching frequency, and reliability-oriented stress proxies as fully coupled decision variables under one shared robust objective set.

Joint filter and controller design is closer to the actual device-level problem because harmonic attenuation, damping action, and operating loss are shaped jointly by the electrical interface and the control law. Faria et al., [9] show this explicitly through proportional-resonant current control and output-filter optimisation for grid-tied

inverters, while [9] treat switching-harmonic attenuation through a coupled LLCL and predictive-control viewpoint. Multi-functional inverter compensation studies also connect harmonic objectives with converter-side control allocation, which reinforces the point that filter behaviour cannot be separated cleanly from controller structure and operating policy (Li et al., [16]).

Broader control literature supports the same conclusion. Advanced inverter-control surveys, hybrid optimisation studies, and photovoltaic harmonic-mitigation studies all show that waveform quality depends on the interaction among modulation, controller bandwidth, damping strategy, and switching frequency (Adnan et al., [17]; Iqbal et al., [1]; Gony et al., [2]). Stability-oriented work further shows that resonance interaction and closed-loop margin preservation are central rather than peripheral concerns in harmonic design (Pavón et al., [18]).

Even in this stronger stream, two gaps remain. First, controller tuning is often coupled to a selected filter family, but controller-family choice itself is not usually searched jointly with topology choice inside one differentiable or otherwise unified design pipeline. Second, uncertainty and reliability-related stress are rarely treated as first-class design objectives within the same co-design loop. The result is a partial co-design view rather than a complete deployment-oriented one.

Uncertainty is already recognised across harmonic-mitigation literature, but treatment is often secondary to nominal optimisation. Stochastic and scenario-based studies show that renewable penetration, operating variability, and component drift can alter harmonic feasibility, yet uncertainty is frequently appended after nominal design rather than embedded directly into both objective definitions and constraint aggregation (Abubakar et al., [6]; Lin et al., [14]; Hamza et al., [13]). A similar pattern appears in staged APF and distributed-mitigation frameworks, where deployment variability is acknowledged but device-level implementation detail and shared robust feasibility are not the main design centre (Yin et al., [12]; Rahman et al., [10]).

Reliability treatment is also usually indirect. Online tuning work and stability-oriented control studies account for operating stress and adaptation, but explicit integration of thermal and ripple-driven stress with harmonic, cost, and loss trade-offs is uncommon (Metry et al., [19]; Pavón et al., [18]). Learning-based compensation and coordination studies confirm that AI-enabled or distributed logic can improve operating behaviour, yet the dominant emphasis remains either network-level coordination or compensator behaviour rather than a realizable converter-interface co-design pipeline with interpretable stress proxies (Gaddam et al., [20]; Wang et al., [11]).

Figure 1 summarizes the three device-level streams and the central unresolved gap.

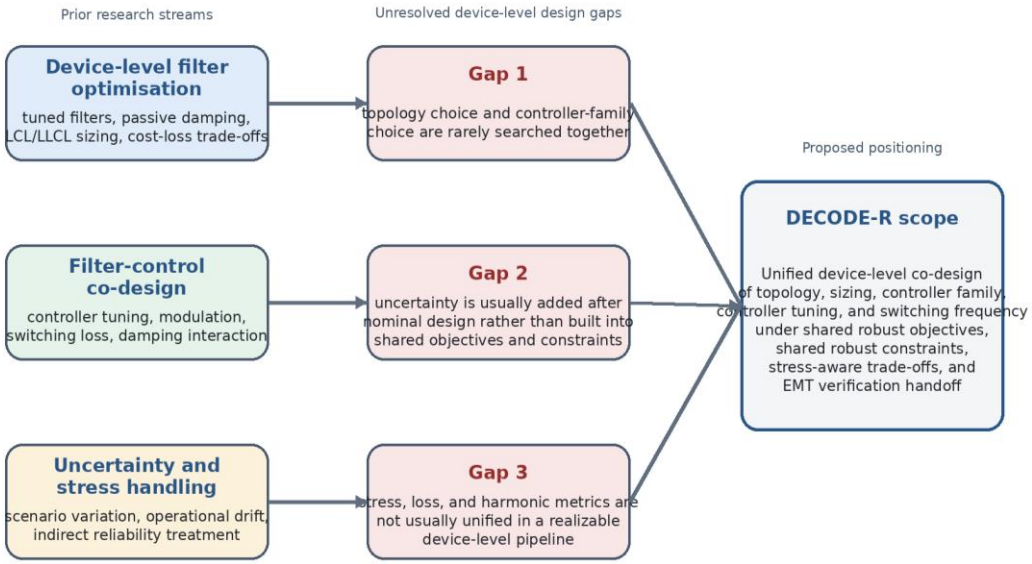

Figure 1: Positioning of DECODE-R against related research streams and unresolved device-level design gaps

The synthesis separates device-level filter optimisation, filter-control co-design, and uncertainty or stress-aware design, and it shows that no reviewed stream closes the full device-level path from joint search to realizable robust designs under one shared objective-and-constraint structure.

The reviewed literature therefore supports four direct design implications. First, topology choice and controller-family choice need to be treated jointly at device level rather than through sequential tuning. Second, uncertainty must influence both objective aggregation and feasibility definitions, not only post hoc sensitivity reporting. Third, reliability-related stress needs a design-stage representation that is simple enough for optimisation yet still tied to physics-relevant thermal and ripple drivers. Fourth, the optimisation output must remain convertible to realizable discrete designs that can be checked in higher-fidelity EMT analysis before deployment claims are made.

Karooyee et al., [8] and Faria et al., [9] remain the closest conceptual device-level reference points because they expose the computational burden of inverter-side filter search and the importance of joint filter-control tuning. Even so, those streams do not provide a unified pipeline that searches topology family, controller family, continuous sizing, continuous tuning, and switching frequency together under one shared robust objective-and-constraint structure with explicit stress-aware trade-offs.

These studies remain important for evaluation logic, but not in the same way. Karooyee et al., [8] and Faria et al., [9] serve mainly as conceptual device-level anchors for structural coupling and joint filter-control tuning, whereas tuned passive-filter studies such as Hsiao et al., [4] and Alhaider et al., [15] define the main implemented benchmark families that can be aligned under a shared scenario set and a shared feasibility envelope. That distinction matters because later experimental comparison must preserve reproducibility and benchmark symmetry without implying one-to-one baseline equivalence across studies that use different topology libraries, control assumptions, or reporting detail. Accordingly, later benchmarking uses directly executable tuned-filter families as implemented baselines, whereas the Karooyee- and Faria-type studies mainly anchor the interpretation of device-level structural coupling and joint filter-control search burden.

DECODE-R is therefore positioned as a device-level, deployment-oriented co-design workflow rather than as a network-coordination or compensator-allocation framework. Distributed mitigation and incentive-driven coordination remain relevant as wider system context, but they sit outside the direct method scope addressed here (Rahman et al., [10]; Wang et al., [11]). The resulting design implication is a unified differentiable pipeline that links relaxed search, robust scenario aggregation, physics-informed StressRisk, discrete extraction, local refinement, and EMT verification handoff within one consistent path from formulation to implementable design.

3 METHODS

DECODE-R formulates harmonic-filter and controller co-design as a single robust multi-objective problem over structure, sizing, tuning, and switching frequency. The retained harmonic-balance evaluator, averaged dynamic model, physics-informed stress layer, and preference-conditioned generator are fixed before the optimization campaign begins. Figure 2 summarizes the full information flow, Table 1 records the fixed implementation settings used in the reported study, Table 2 states the robust objective and constraint aggregation rules. Figure 3 makes the admissible supergraph explicit, Table 3 defines the exact filter and controller library, and Table 4 collects the loss, thermal, and discretization constants used for the reported study.

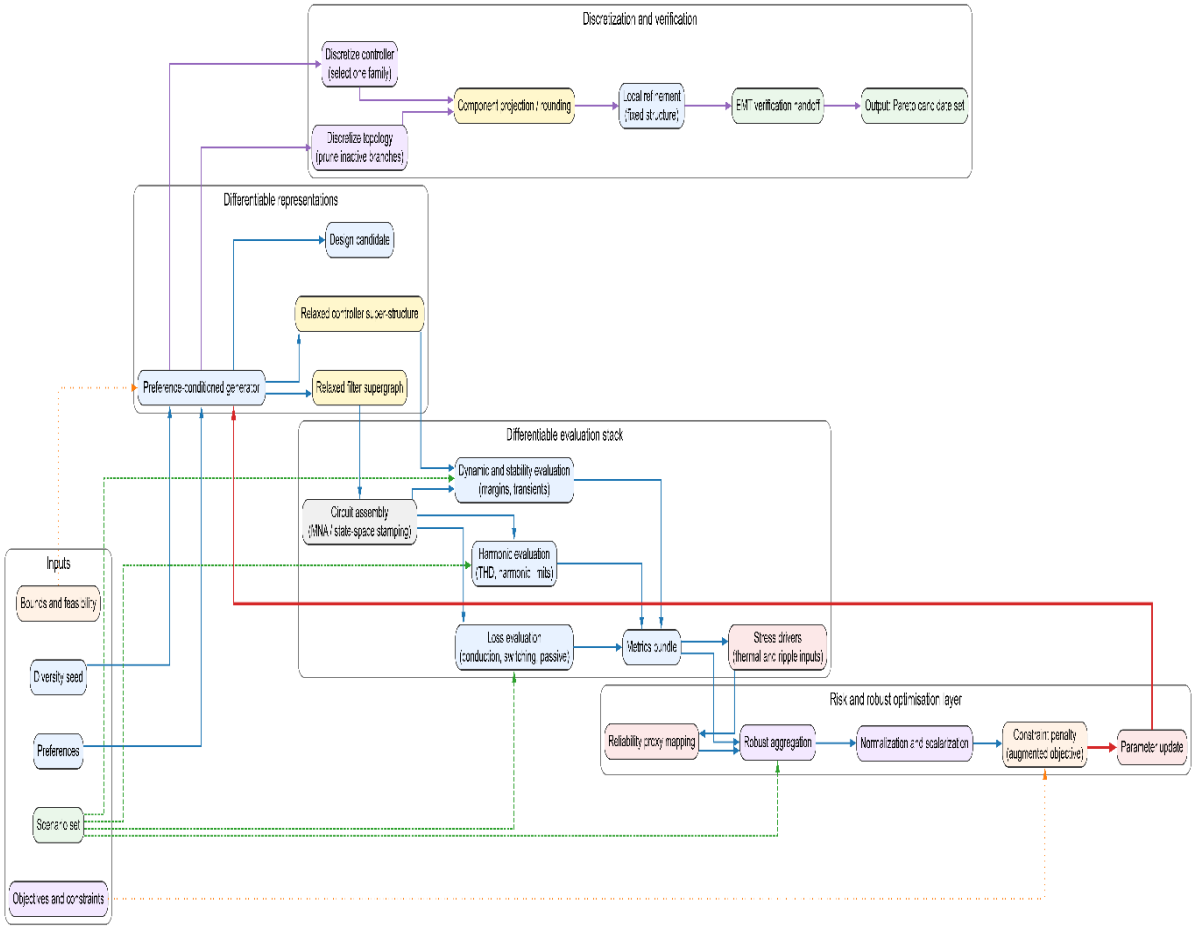

Figure 2: DECODE-R architecture for robust multi-objective co-design of filter topology and control

Figure 2 DECODE-R architecture and information flow. A shared generator maps a preference vector and an optional latent code to a relaxed co-design state consisting of filter-topology weights, continuous passive-component parameters, controller-family gates, controller parameters, and switching frequency. Scenario batches feed a retained-band linear harmonic-balance evaluator, an averaged dynamic and stability model, and a physics-informed stress layer. The resulting robust objectives and feasibility checks drive gradient-based updates. Dominant relaxed selections are then discretized, projected to realizable component values, locally refined under fixed structure, and handed off to electromagnetic transient verification.

Table 1: Fixed implementation choices used in the Section 3 optimization loop.

Block	Adopted choice	Fixed value(s) used in this study
Harmonic evaluator	Linear harmonic-balance solver	Retained orders $h = 1, \ldots, 25$; dense LU solve; relative residual tolerance 10^{-8}
Dynamic evaluator	Continuous-time averaged converter-filter-grid model	Aggregate digital delay $T_d = 1.5 / f_{sw}$; fixed horizon $T_h = 60\,\text{ms}$; RK4 step $25\,\mu\text{s}$
Objective aggregation	Mixed mean/CVaR	$\text{CVaR}_{0.95}$ for THD and StressRisk; scenario

Block	Adopted choice	Fixed value(s) used in this study
	formulation	mean for total loss; deterministic cost proxy
Constraint aggregation	Worst-case feasibility rules	Full-scenario max over Ω or dynamic-subset max over Ω_{dyn}, as stated in Table 2
Generator architecture	Shared-trunk neural generator	Two hidden layers of 128 GELU units; four task-specific heads; latent dimension $d_z = 4$
Optimizer	Adam with augmented-Lagrangian penalties	Learning rate 10^{-3}; $\beta_1 = 0.9$; $\beta_2 = 0.999$; batch size $B = 24$; iteration budget $T = 12,000$
Temperature schedule	Joint topology/controller annealing	$\tau, \tau_c : 1.50 \rightarrow 0.15$ over the first $9,600$ iterations
Local refinement	Fixed-structure continuous re-optimization	L-BFGS-B; 80 steps maximum; relative objective tolerance 10^{-6}
Stress-map constants	Semiconductor and capacitor proxies with conservative smooth-maximum aggregation	Peak-temperature weight 0.65; cycling weight 0.35; cycling reference 25 °C swing; cycling exponent 1.7; capacitor ripple exponent 2.1; capacitor reference temperature 70 °C; capacitor temperature coefficient 0.035 per K; smooth-maximum sharpness 6.

3.1 System definition and robust co-design formulation

The plant comprises a three-phase grid-connected voltage-source converter, an AC-side interface filter, the PCC measurement point, and the grid equivalent viewed from the PCC. The optimization boundary ends at the PCC: current distortion, power factor, passive-component ratings, thermal stress, and small-signal margins are all assessed at or downstream of that boundary. Two complementary models are used inside the loop: a retained-band harmonic model for steady-state distortion and ripple-driven losses, and a continuous-time averaged model for dynamic metrics and stability margins.

The harmonic view resolves the retained frequency band only; the method does not attempt to enforce electromagnetic-interference compliance beyond the retained harmonics, resolve device-level d_v / d_t stress, or model stray high-frequency parasitics inside the optimization loop. Those effects are checked only after discretization during EMT verification. Magnetic saturation and long-horizon aging are not modeled explicitly; instead, parameter drift and ESR growth enter through the scenario distribution.

The co-design decision vector is Eq 1

$$x = \{\pi, \theta_f, \theta_c, f_{sw}\}, x \in X \quad ...(Eq\ 1)$$

where π collects relaxed filter-structure weights, θ_f collects continuous passive-component values, θ_c collects controller-family gates and continuous controller parameters, and f_{sw} is the switching frequency. The feasible set X enforces engineering bounds on passive values, controller gains, admissible controller families, and switching frequency.

A deployment scenario is denoted by $\xi \in \Omega$ and bundles grid strength, ambient temperature, component tolerance and ESR drift, operating point, and background harmonic content. For any pair (x, ξ), the evaluator returns the scenario-conditioned packet Eq 2

$$y = g(x, \xi) = \{THD, PF, PM, GM, J_{dyn}, P_{loss}, StressRisk, \ldots\} \quad ...(Eq\ 2)$$

The robust objective vector is Eq 3

$$F(x) = \begin{bmatrix} f_{THD}(x) \\ f_{loss}(x) \\ f_{cost}(x) \\ f_{risk}(x) \end{bmatrix} = \begin{bmatrix} CVaR_{0.95}[THD(x, \xi)] \\ E_{\xi \sim \Omega}[P_{loss}(x, \xi)] \\ Cost(x) \\ CVaR_{0.95}[StressRisk(x, \xi)] \end{bmatrix} \quad ...(Eq\ 3)$$

To keep the hardware-cost term explicit and reproducible, the manuscript uses the fixed dimensionless proxy Eq 4

$$Cost(x) = c_b N_b(x) + c_L \sum_{\ell \in L} \frac{L_\ell I_{\ell,rated}^2}{L_{base} I_{base}^2} + c_C \sum_{c \in C} \frac{C_c V_{c,rated}^2}{C_{base} V_{base}^2} + c_R \sum_{r \in R} \frac{P_{R,r}^{rated}}{S_n} \quad ...(Eq\ 4)$$

with fixed weights $c_b = 0.20$, $c_L = 0.35$, $c_C = 0.30$, and $c_R = 0.15$. Here $N_b(x)$ is the number of retained passive branches after discretization, $L_{base} = 1 mH$, $C_{base} = 10 \mu F$, $I_{base} = S_n / (\sqrt{3} V_{LL})$, and $V_{base} = V_{LL} / \sqrt{3}$. The proxy is scenario-independent by construction and is used only for relative comparison within the admissible design envelope.

Feasibility is enforced through robustly aggregated constraints. In the reported study, every feasibility constraint uses a deterministic max-type aggregation rather than a CVaR relaxation: full-scenario worst-case aggregation is used for power factor, passive-component ratings, capacitor ripple, resonance placement, and thermal limits, while worst-case aggregation over the dynamic subset $\Omega_{dyn} \subset \Omega$ is used for phase margin, gain margin, and the transient score. Table 2 states the exact aggregation rule for each constraint block.

The resulting co-design problem is Eq 5

$$\min_{x \in X} F(x) \text{ subject to } \phi_j(x) \le 0,\ j = 1, \ldots, J \quad (3.5) \quad ...(Eq\ 5)$$

This formulation keeps structure search, continuous sizing, controller selection, controller tuning, and switching-frequency choice inside one decision space while making the uncertainty treatment explicit rather than post hoc.

Table 2: Objective and robust-constraint aggregation rules used in the reported study

Quantity or constraint block	Definition used in the study	Robust operator and scope
f_{THD}	Grid-current distortion objective	$CVaR_{0.95}[THD(x, \xi)]$ over full Ω
f_{loss}	Total-loss objective	$E_{\xi \sim \Omega}[P_{loss}(x, \xi)]$ over full Ω
f_{cost}	Hardware-cost proxy	$Cost(x)$; scenario independent

Quantity or constraint block	Definition used in the study	Robust operator and scope
f_{risk}	Stress objective	$\mathrm{CVaR}_{0.95}\left[\mathrm{StressRisk}\left(x,\xi\right)\right]$ over full Ω
ϕ_{PF} , ϕ_{rate} , ϕ_{therm} , ϕ_{ripple} , ϕ_{res}	Power-factor, rating, thermal, ripple, and resonance feasibility	$\max_{\xi\in\Omega}(\cdot)$
ϕ_{PM} , ϕ_{GM} , ϕ_{dyn}	Margin and transient-feasibility checks	$\max_{\xi\in\Omega_{dyn}}(\cdot)$

3.2 Differentiable topology-control Representation

The filter library is represented by a supergraph and the controller library by a parallel super-structure. Figure 3 makes the retained electrical branch set explicit, and Table 3 lists the exact filter families and controller families searched in the reported study; no other structural options are available to the optimizer.

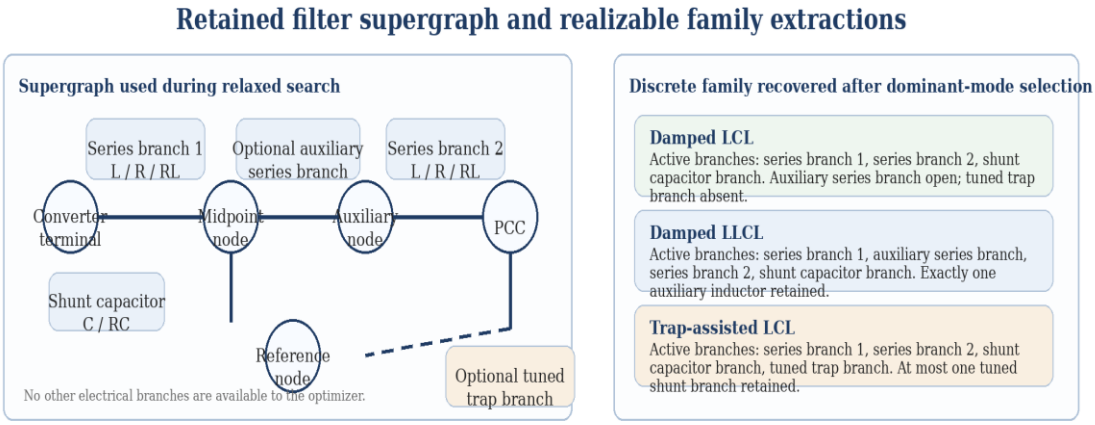

The supergraph contains one converter-side series branch, one optional auxiliary series branch, one PCC-side series branch, one shunt capacitor branch, and one optional tuned shunt branch.

Figure 3: Retained filter supergraph and realizable family extractions

Figure 3 retained filter supergraph and realizable family extractions. The supergraph uses one converter-side series branch, one optional auxiliary series branch, one PCC-side series branch, one shunt capacitor branch, and one optional tuned shunt branch. Damped LCL, damped LLCL, and trap-assisted LCL are obtained by activating the corresponding branch subsets under the realizability rules stated in Table 3.

The actual supergraph used in the reported study contains five electrical branches: a converter-to-midpoint series branch, an optional midpoint-to-auxiliary series branch, an auxiliary-to-PCC series branch, a midpoint shunt capacitor branch, and an optional PCC shunt trap branch. No bypass branches, hybrid active-filter elements, or additional interior nodes are available to the optimizer.

Table 3: Admissible filter and controller library used by the relaxed search.

Search block	Admissible option	Tunable quantities retained after selection	Realizability rule

Search block	Admissible option	Tunable quantities retained after selection	Realizability rule
Filter family	Damped LCL	L_c, L_g, C_f, R_d	One connected series path and one shunt capacitor group
Filter family	Damped LLCL	L_c, L_m, L_g, C_f, R_d	Midpoint inductor retained only when LLCL is selected
Filter family	Trap-assisted LCL	L_c, L_g, C_f, L_t, C_t, R_d	At most one tuned trap branch, tuned in the 5th/7th-harmonic vicinity
Controller family	dq-PI with capacitor-current feedback	K_p, K_i, K_{ad}, ω_{pll}	Active damping retained only if capacitor-current sensing is present
Controller family	PR with notch compensation	K_p, K_r, ω_n, ζ_n, ω_{pll}	Notch block retained only for the PR family
Controller family	One-step predictive current law	Predictive weight coefficients, current-penalty gain, ω_{pll}	Only one controller family is retained after discretization

For each supergraph edge e with admissible mode set M_e, discrete mode selection is relaxed by a temperature-controlled softmax over edge logits $a_{e,m}$ Eq 6:

$$\pi_{e,m} = \frac{\exp\left(a_{e,m}/\tau\right)}{\sum_{m' \in M_e} \exp\left(a_{e,m'}/\tau\right)}, \sum_{m \in M_e} \pi_{e,m} = 1 \quad ...(\text{Eq } 6)$$

Each active mode carries a local parameter block $\theta_{e,m}$ containing the associated inductance, capacitance, resistance, or tuned-branch value. Positive physical quantities are emitted in log space and mapped back through exponentiation so that passivity and positivity are preserved throughout optimization.

The relaxed edge admittance is assembled as a convex mixture of mode-level admittances, Eq 7

$$Y_e(s) = \sum_{m \in M_e} \pi_{e,m} Y_{e,m}\left(s;\theta_{e,m}\right) \quad ...(\text{Eq } 7)$$

which allows the full circuit to be stamped by modified nodal analysis without enumerating each candidate topology explicitly. Temperature τ is linearly annealed during training so that diffuse edge selections sharpen gradually toward a discrete structure.

The controller super-structure uses the admissible family set listed in Table 3. Each family has its own parameter block and controller-selection logit, and the relaxed controller gate and blended control law are

The implemented controller variants are fixed as follows. The d_q frame family is a synchronous-frame current controller with capacitor-current feedback active damping. In this block, i_{dq}^* denotes the d_q frame current reference, i_{dq} denotes the measured converter current, $e_{dq} = i_{dq}^* - i_{dq}$ denotes the current-tracking error, $u_{ff,dq}$

denotes the feedforward voltage term formed from PCC-voltage and cross-coupling compensation, and $i_{C,dq}$ denotes the measured capacitor current used by the active-damping path Eq 8.

$$u_{dq}(s) = K_p e_{dq}(s) + \frac{K_i}{s} e_{dq}(s) - K_{ad} i_{C,dq}(s) + u_{ff,dq}(s) \quad ...(Eq\ 8)$$

The stationary-frame family is a proportional-resonant current controller with one retained notch compensation block in the current-feedback path. Here, e_{ab} denotes the stationary-frame current error, and $C_n(s)$ denotes the retained second-order band-stop block centered at the selected notch frequency with the bandwidth and gain listed in Table 3. Eq 9

$$u_{\alpha\beta}(s) = \left[K_p + K_r \frac{2\omega_c s}{s^2 + 2\omega_c s + \omega_0^2} \right] e_{\alpha\beta}(s) - C_n(s) i_{\alpha\beta}(s) ...(Eq\ 9)$$

The retained notch block is therefore fully determined by its center frequency, bandwidth, and gain; no additional hidden controller elements are introduced outside the library stated in Table 3.

The predictive family uses one-step dq-current prediction with a quadratic tracking term and an explicit switching-penalty weight. Its internal prediction model is written in discrete time, where A_d, B_d, *and* E_d are the retained plant matrices, $v_{g,dq}[k]$ is the sampled grid-voltage disturbance, and $i_{dq}[k+1|k](u[k])$ is the predicted dq current associated with the candidate control move $u[k]$ Eq 10.

$$i_{dq}[k+1] = A_d i_{dq}[k] + B_d u_{dq}[k] + E_d v_{g,dq}[k] \quad ...(Eq\ 10)$$

The implemented control move is selected from the admissible averaged-voltage set by minimizing $J_{pred}(u[k]) = \left\| i_{dq}^*[k+1] - i_{dq}[k+1|k](u[k]) \right\|_2^2 + \lambda_{sw} \left\| u[k] - u[k-1] \right\|_2^2$. These three laws define the exact controller variants referred to in Table 3; only their gain, notch, predictive-weight, and delay parameters vary during search.

In the reported study, this admissible averaged-voltage set is the continuous dq-voltage region imposed by the regulated dc link and space-vector modulation, so the predictive branch searches only averaged voltage commands that remain inside the modulation-feasible disk Eq 11.

$$\rho_k = \frac{\exp(b_k / \tau_c)}{\sum_{k' \in K} \exp(b_{k'} / \tau_c)}, u(t) = \sum_{k \in K} \rho_k u_k(t; \theta_{c,k}) \quad ...(Eq\ 11)$$

where τ_c is annealed jointly with τ. During the relaxed phase, the blended controller is used only as a differentiable surrogate for structure selection; after discretization, the design retains one controller family only.

Discrete extraction enforces four realizability rules. First, the retained passive branches must form one connected series path from the converter terminals to the PCC. Second, exactly one shunt capacitor group is permitted, and at most one trap branch may remain active. Third, branches assigned to open or none are pruned before parameter projection. Fourth, active-damping terms are retained only when the selected controller family includes the corresponding measured signal. These rules prevent the relaxed search from generating algebraically valid but physically unusable filter-control combinations.

3.3 Differentiable Performance Evaluation Stack

The optimization loop uses one implemented evaluator rather than a menu of interchangeable models. Harmonic quantities are computed with a linear harmonic-balance formulation over a fixed retained harmonic set, dynamic quantities are computed from a continuous-time averaged model with an explicit aggregate digital delay, and fixed electrical-loss and thermal coefficients are applied consistently across all candidate designs. Table 1 summarizes the solver, optimizer, and horizon settings used in the reported study, and Table 4 records the component-level coefficients used by the loss and stress submodels.

Let $\mathbf{X}_H$ collect the retained harmonic phasors of grid current, PCC voltage, and the internal branch quantities required by the loss and stress models. For each scenario, the harmonic state is obtained from Eq 12

$$A_H(x,\xi)\mathbf{X}_H = b_H(x,\xi) \quad \text{...(Eq 12)}$$

where A_H is assembled from the relaxed filter admittances, the grid equivalent, background voltage harmonics, and the closed-loop linearized response in the retained band. The solve uses dense LU factorization with a relative residual tolerance of 10^{-8}; harmonic coupling beyond the retained band is excluded from the optimization loop.

Once $\mathbf{X}_H$ is known, grid-current total harmonic distortion at the PCC is Eq 13

$$\mathrm{THD}(x,\xi) = \frac{\sqrt{\sum_{h=2}^{25} I_{g,h}^2(x,\xi)}}{I_{g,1}(x,\xi)} \quad \text{...(Eq 13)}$$

with $I_{g,h}$ denoting the RMS magnitude of the h th grid-current harmonic. Fundamental active and apparent power are obtained from the fundamental PCC voltage and current phasors, and the corresponding power factor is used directly in the feasibility checks.

Dynamic behavior is evaluated on an averaged plant of the form Eq 14

$$\dot{x}_p(t) = A_p(x,\xi)x_p(t) + B_p(x,\xi)u(t),\ y_p(t) = C_p(x,\xi)x_p(t) + D_p(x,\xi)u(t) \quad \text{...(Eq 14)}$$

together with an aggregate control delay represented by $T_d = 1.5/f_{sw}$. Phase margin and gain margin are extracted from the linearized loop around the scenario operating point. Short-horizon transient quality is measured over a fixed $60\,\mathrm{ms}$ horizon by Eq 15

$$J_{dyn}(x,\xi) = w_s t_s(x,\xi) + w_o M_p(x,\xi) + w_e \int_0^{T_h} \square e(t) \square_2^2 \, dt, T_h = 60\,\mathrm{ms} \quad \text{...(Eq 15)}$$

using weights $w_s = 0.35$, $w_o = 0.25$, and $w_e = 0.40$.

Loss calculations are coupled directly to the harmonic and averaged-model outputs. Total loss is decomposed as Eq 16

$$P_{loss}(x,\xi) = P_{cond}(x,\xi) + P_{sw}(x,\xi) + P_L(x,\xi) + P_C(x,\xi) \quad \text{...(Eq 16)}$$

The four loss components are instantiated explicitly, and the fixed coefficients used in the reported study are collected in Table 4 so that the reduced-model loss calculations are fully specified Eq 17 to Eq 20.

$$P_{cond}(x,\xi) = \Sigma_d \left(V_{0,d} I_{d,avg}(x,\xi) + r_{on,d} I_{d,rms}{}^2(x,\xi) \right) \quad \text{...(Eq 17)}$$

$$P_{sw}(x,\xi) = f_{sw}\Sigma_d\left(E_{on}(i_d, v_d) + E_{off}(i_d, v_d)\right) \quad ...(Eq\ 18)$$

$$E_{on}(i,v) = a_{on} + b_{on}i + c_{on}v + d_{on}iv, E_{off}(i,v) = a_{off} + b_{off}i + c_{off}v + d_{off}iv \quad ...(Eq\ 19)$$

$$P_L(x,\xi) = \Sigma_\ell\left(R_{cu,\ell}I_{\ell,rms}{}^2(x,\xi) + k_{core,\ell}\Sigma_h(hf_0)^{\alpha_\ell} I_{\ell,h}{}^2(x,\xi)\right), P_C(x,\xi) = \Sigma_c\left(ESR_c I_{c,rms}{}^2(x,\xi)\right) \quad ...(Eq\ 20)$$

The switching-energy coefficients, core-loss coefficient and exponent, ESR reference, and thermal parameters are fixed once for the entire study; only the electrical design variables and deployment scenario change during optimization.

3.4 Physics-informed stress-to-risk Mapping

The reliability-oriented objective is intentionally framed as a stress proxy rather than as an absolute lifetime predictor. Direct mission-profile life estimation would require component-specific aging data, long-horizon duty-cycle processing, and electro-thermal calibration beyond the scope of the present design loop. The reported study therefore fixes one semiconductor thermal network, one capacitor hotspot model, and one conservative smooth-maximum aggregation so that the stress objective is fully specified rather than left as a family of alternatives.

The semiconductor thermal path used in the reported study is fixed to a two-pole Foster network. Its branch resistances and time constants are listed in Table 4, and the same network is used for every candidate design and every scenario Eq 21.

$$Z_{th}(t) = R_{th,1}\left(1 - e^{\frac{-t}{\tau_1}}\right) + R_{th,2}\left(1 - e^{\frac{-t}{\tau_2}}\right) \quad ...(Eq\ 21)$$

Semiconductor stress is derived from junction temperature and cycling amplitude. With $P_{semi}(t; x, \xi)$ denoting the time-varying semiconductor loss power and Z_{th} the lumped thermal impedance, the junction-temperature trajectory is Eq 22

$$T_j(t; x, \xi) = T_a(\xi) + (Z_{th} * P_{semi})(t) \quad ...(Eq\ 22)$$

The semiconductor-risk term is Eq 23

$$Risk_{semi}(x,\xi) = w_T\,\mathrm{softplus}\left(\frac{T_{j,max}(x,\xi) - T_{j,lim}}{s_T}\right) + w_\Delta\left(\frac{\Delta T_j(x,\xi)}{\Delta T_{ref}}\right)^p \quad ...(Eq\ 23)$$

so that one term penalizes limit proximity and the other penalizes cycling intensity.

Capacitor stress is driven by ripple-current heating. The hotspot estimate is Eq 24

$$T_C(x,\xi) = T_a(\xi) + R_{\theta,C}P_C(x,\xi) \quad ...(Eq\ 24)$$

and the capacitor-risk term is Eq 25

$$Risk_{cap}(x,\xi) = \left(\frac{I_{C,rms}(x,\xi)}{I_{C,rated}}\right)^q \exp\left[\gamma\left(T_C(x,\xi) - T_{ref}\right)\right] \quad ...(Eq\ 25)$$

This form increases smoothly with both ripple-loading ratio and hotspot temperature and is intentionally interpretable: hotter or higher-ripple operation always raises the proxy.

The component-level terms are aggregated conservatively through a smooth maximum Eq 26,

$$\mathrm{StressRisk}(x,\xi)=\frac{1}{\kappa_r}\ln\left(\exp\left[\kappa_r \mathrm{Risk}_{\mathrm{semi}}(x,\xi)\right]+\exp\left[\kappa_r \mathrm{Risk}_{\mathrm{cap}}(x,\xi)\right]\right) \quad \text{...(Eq 26)}$$

and the robust risk objective is the scenario-tail statistic already defined in the robust objective definition above. This fixed two-term aggregation makes the reported risk values unambiguous: the objective ranks semiconductor and capacitor stress exposure only, while inductor feasibility is handled separately by worst-case constraints.

Table 4: Loss-model coefficients, thermal constants, and catalog-rounding rules used for the reported discrete designs.

Block	Adopted form	Fixed values or rule used in this study
Semiconductor conduction model	Linear on-state loss	On-state voltage intercept 1.85 V; on-state resistance 21 mΩ per switch; resistance temperature coefficient 0.0039 per °C referenced to 25 °C.
Switching-energy fit	Bilinear turn-on and turn-off energy surfaces	Turn-on coefficients [0.12, 0.018, 5.5e-4, 1.2e-5] and turn-off coefficients [0.09, 0.015, 4.8e-4, 1.0e-5] in millijoule units for [constant, current, voltage, current×voltage].
Inductor and capacitor passive-loss model	Copper loss plus frequency-weighted core-loss proxy; ripple-current ESR loss	Core-loss coefficient 2.8e-5; frequency exponent 1.35; capacitor ESR reference 62 mΩ at 25 °C with 0.25% per °C increase in thermal scenarios.
Semiconductor thermal network	Two-pole Foster network	Branch 1: 0.12 K/W and 0.025 s; branch 2: 0.18 K/W and 0.42 s; junction limit 120 °C.
Capacitor hotspot and risk mapping	Single-resistance hotspot model with exponential temperature acceleration	Capacitor thermal resistance 2.6 K/W; ripple-current rating 12 A RMS; reference temperature 70 °C; ripple exponent 2.1; temperature coefficient 0.035 per K.
Stress aggregation constants	Semiconductor peak-temperature and cycling terms combined with capacitor term through a smooth maximum	Peak-temperature weight 0.65; cycling weight 0.35; cycling reference 25 °C swing; cycling exponent 1.7; smooth-maximum sharpness 6.
Catalog rounding after discretization	Rule-based engineering quantization after dominant-mode selection	Inductors nearest 0.01 mH; shunt and tuned capacitors nearest 0.2 μF; damping resistors nearest 0.1 Ω; tuned branch re-tuned by nearest feasible pair with less than 2% tuning-frequency deviation; no vendor-specific catalog lock-in was applied in the reported study.

3.5 Preference-conditioned optimization, discretization, and EMT handoff

DECODE-R uses a preference-conditioned generator so that one trained model can trace multiple trade-off regions of the Pareto set. Let $w=[w_1,w_2,w_3,w_4]^{\cdot}$ be a nonnegative preference vector with $\sum_i w_i=1$, and let $z\sim \mathrm{N}\left(0,I_4\right)$ be an optional latent code used to generate diversity under the same preference. The generator is a shared-trunk multilayer perceptron with two hidden layers of 128 GELU units and four task-specific output heads for topology logits, passive parameters, controller logits and controller parameters, and switching frequency. The mapping is Eq 27

$$x=G_\phi\left(w,z\right) \quad ...(\text{Eq } 27)$$

where the topology head returns the logits that define π, the passive-parameter head returns bounded continuous values, the controller head returns controller-family logits and controller parameters, and the switching-frequency head returns f_{sw}. The latent dimension is fixed at four. Training uses Adam with learning rate 10^{-3}, $\beta_1=0.9$, and $\beta_2=0.999$.

Optimization is performed on normalized objectives, Eq 28

$$\hat{f}_i\left(x\right)=\frac{f_i\left(x\right)}{s_i},S\left(x;w\right)=\sum_{i=1}^{4}w_i\hat{f}_i\left(x\right) \quad ...(\text{Eq } 28)$$

with fixed engineering scales $\mathrm{s_{THD}}=1\%$, $\mathrm{s_{loss}}=1\mathrm{kW}$, $\mathrm{s_{cost}}=1$, and $\mathrm{s_{risk}}=0.25$. Robust constraints are converted to smooth nonnegative violations $\mathrm{v_j}\left(\mathrm{x}\right)=\mathrm{softplus}\left(\phi_\mathrm{j}\left(\mathrm{x}\right)\right)$ and combined with the scalarized objective through the augmented objective Eq 29

$$\mathrm{L}\left(\mathrm{x};\mathrm{w},\lambda,\rho\right)=\mathrm{S}\left(\mathrm{x};\mathrm{w}\right)+\sum_{\mathrm{j}=1}^{\mathrm{J}}\lambda_\mathrm{j}\mathrm{v_j}\left(\mathrm{x}\right)+\frac{\rho}{2}\sum_{\mathrm{j}=1}^{\mathrm{J}}\mathrm{v_j^2}\left(\mathrm{x}\right) \quad ...(\text{Eq } 29)$$

Scenario robustness is estimated on mini-batches $\{\xi_\mathrm{b}\}_{\mathrm{b}=1}^{\mathrm{B}}\subset\Omega$ with batch size $\mathrm{B}=24$. For any scenario-dependent quantity $\mathrm{Z}\left(\mathrm{x},\xi\right)$, the batch expectation is Eq 30

$$\hat{\mathrm{E}}\left[\mathrm{Z}\right]\left(\mathrm{x}\right)=\frac{1}{\mathrm{B}}\sum_{\mathrm{b}=1}^{\mathrm{B}}\mathrm{Z}\left(\mathrm{x},\xi_\mathrm{b}\right) \quad ...(\text{Eq } 30)$$

and the differentiable batch estimator used for conditional value at risk is Eq 31

$$\mathrm{CVaR}_\alpha\left[\mathrm{Z}\right]\left(\mathrm{x}\right)=\min_{\eta\in\mathrm{R}}\left[\eta+\frac{1}{\left(1-\alpha\right)\mathrm{B}}\sum_{\mathrm{b}=1}^{\mathrm{B}}\mathrm{softplus}\left(\mathrm{Z}\left(\mathrm{x},\xi_\mathrm{b}\right)-\eta\right)\right] \quad ...(\text{Eq } 31)$$

At each iteration, the method samples $\left(\mathrm{w},\mathrm{z}\right)$ and a scenario batch, generates a relaxed design, evaluates the robust objectives and constraints, updates the generator parameters ϕ with Adam, and updates the penalty multipliers λ by the standard monotone augmented-Lagrangian rule. The iteration budget is fixed at $\mathrm{T}=12{,}000$ per run. The topology and controller temperatures τ and τ_c are linearly annealed from 1.50 to 0.15 over the first $9{,}600$ iterations, while the penalty parameter ρ grows geometrically from 1 to 20 across the run.

After training, candidate extraction proceeds in four stages. First, the trained generator is sampled over the preference simplex and each relaxed design is discretized by dominant-mode selection as defined in Section 3.2. Second, passive values are projected to admissible bounds and then mapped to a rule-based engineering quantization grid used in the reported study rather than to a vendor-specific catalog. Inductors are rounded to the nearest 0.01 mH value on the retained design grid, shunt and tuned capacitors are rounded to the nearest 0.2 μF value, damping resistors are rounded to the nearest 0.1 Ω value above the dissipation requirement, and the tuned trap branch is re-tuned by the nearest feasible inductance-capacitance pair so that the realized tuning frequency remains within 2% of the pre-rounded target. Third, a fixed-structure local refinement stage re-optimizes the continuous passive values, controller parameters, and switching frequency within the selected family to recover performance lost during discretization. Fourth, only the quantized discrete candidates are transferred to EMT simulation under the representative nominal, stressed, and dynamic-event cases defined later in the manuscript. The representative designs reported later in the manuscript therefore use post-quantization discrete values rather than pre-rounded relaxed quantities.

The output of DECODE-R is therefore not a single design but a filtered archive of feasible non-dominated candidates, each with a discrete filter family, realizable passive values, one controller family, tuned controller parameters, switching frequency, and both reduced-model and EMT-verified metrics. This separation between relaxed optimization and discrete verification is central to the method: it preserves the differentiability needed for efficient search while ensuring that only physically realizable designs are carried forward to final comparison.

4 EXPERIMENTAL STUDY

This section validates DECODE-R on a three-phase grid-connected converter benchmark under structured deployment uncertainty. Evidence is reported in four layers: representative-design results, archive-level aggregate performance across repeated runs, stressed-corner robustness, and electromagnetic transient confirmation. Every repeated claim is tied to one fixed design ID, and the balanced DECODE-R design is identified consistently as DE-B1 throughout this section.

4.1 Test system definition and benchmark configurations

Experimental assessment used a 25 kVA three-phase grid-connected voltage-source converter connected to a 400 V, 50 Hz utility interface through an AC-side harmonic filter placed at the point of common coupling. The benchmark fixed the converter stage, synchronization structure, sensing channels, modulation framework, and point of harmonic evaluation for all compared methods so that only filter and control design freedom varied across methods.

The DECODE-R search space retained three admissible filter families within one common design envelope: damped LCL, damped LLCL, and trap-assisted LCL. The baseline methods were adapted to the same converter and scenario space but retained their original structural restrictions. PSHF-IFA [4] was evaluated as a single-tuned passive-filter baseline with fixed dq current control and a limited switching-frequency sweep. DR-DDTF [15] was evaluated as a double-resistor damped double-tuned passive-filter baseline with the same fixed dq control backbone and a limited switching-frequency sweep.

A common feasibility envelope was enforced across all methods. Grid-current total harmonic distortion at the point of common coupling was limited to 5.0%, power factor was constrained to at least 0.98, minimum phase margin and gain margin were fixed at 40° and 6 dB, semiconductor junction temperature was limited to 120 °C, capacitor ripple current was limited to 0.90 per unit of rated capability, and the dominant filter resonance was kept above the control-band interaction region and below 0.45 of the switching frequency. The same full-scenario rescoring protocol and the same objective definitions were used for all retained candidates.

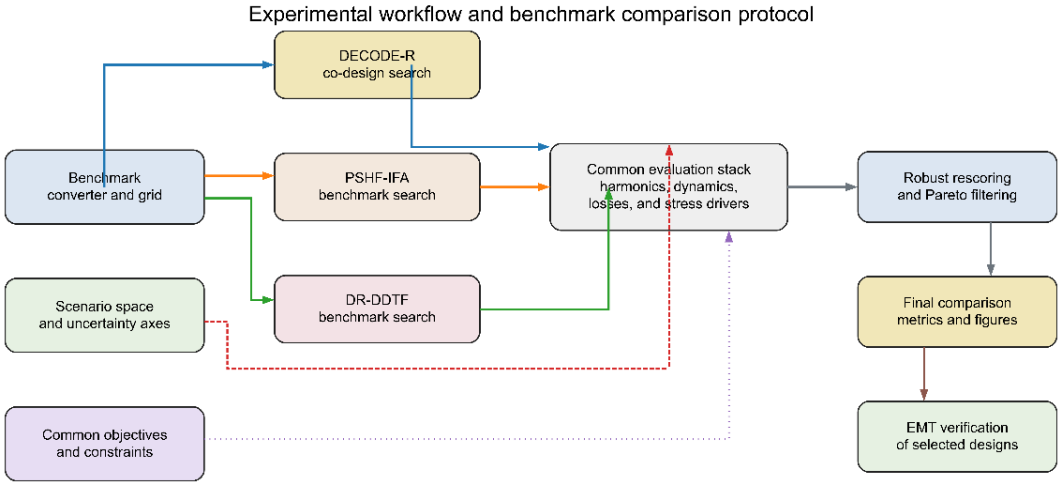

Figure 4: Experimental workflow and benchmark comparison protocol

Figure 4 Experimental workflow and benchmark comparison protocol. Common converter assumptions, method-specific candidate generation, full-scenario rescoring, and final electromagnetic transient validation used for all compared methods. Table 5 common electrical ratings, control assumptions, filter parameter bounds, switching-frequency limits, and feasibility constraints adopted for all compared methods.

Table 5: Benchmark converter, grid, and design bounds used in the comparison.

Category	Parameter	Value or range	Notes
Converter	Rated apparent power	25 kVA	Three-phase grid-connected voltage-source converter
Converter	Grid voltage	400 V RMS	Point-of-common-coupling line-line voltage
Converter	Grid frequency	50 Hz	Utility frequency
Converter	Direct-current link voltage	760 V	Regulated direct-current bus
Converter	Modulation framework	Space-vector modulation	Common switching framework
Filter bounds	Converter-side inductance	0.6 to 2.5 mH per phase	Common lower and upper search bounds
Filter bounds	Grid-side inductance	0.2 to 1.5 mH per phase	Used where the selected topology contains a grid-side branch
Filter bounds	Filter capacitance	4 to 35 μF per phase	Applied to shunt filter branches
Filter bounds	Damping resistance	0.5 to 15 Ω	Applied to passive damping branches
Filter bounds	Tuned-branch range	230 to 430 Hz	Fifth- and seventh-harmonic vicinity when enabled

Switching range	Switching frequency	6 to 16 kHz	Continuous in DECODE-R and limited discrete sweep in baselines
Control set	DECODE-R controller family	dq-PI with capacitor-current feedback; PR with notch compensation; one-step predictive current law	Searchable only in DECODE-R
Control set	Baseline controller	Fixed dq current control	Common sensing and synchronization structure
Constraint	Point-of-common-coupling current THD limit	5.0%	Feasibility threshold
Constraint	Minimum power factor	0.98	Grid-side requirement
Constraint	Minimum phase margin	40°	Small-signal stability requirement
Constraint	Minimum gain margin	6 dB	Small-signal stability requirement
Constraint	Junction-temperature limit	120 °C	Semiconductor stress limit
Constraint	Capacitor ripple-current limit	0.90 per unit	Applied relative to rated ripple capability
Constraint	Resonance placement band	Above 450 Hz and below 0.45 of switching frequency	Used to avoid control-band interaction

4.2 Scenario set design and uncertainty axes

Robust assessment used a structured uncertainty space built from grid strength, ambient temperature, component tolerance drift, operating-point variation, and background voltage distortion. The full scenario set contained 216 cases formed by the Cartesian combination of four short-circuit-ratio levels, three ambient-temperature levels, three tolerance bundles, three operating states, and two background-harmonic levels.

The dynamic subset used for transient constraints and step-response reporting consisted of representative load-step and grid-impedance-jump cases drawn from the same uncertainty space. Full-scenario means and 0.95 conditional value at risk were reported for total harmonic distortion and StressRisk, while phase-margin, gain-margin, thermal, ripple, and resonance-placement limits were enforced through worst-case or robust constraint aggregation during rescoring.

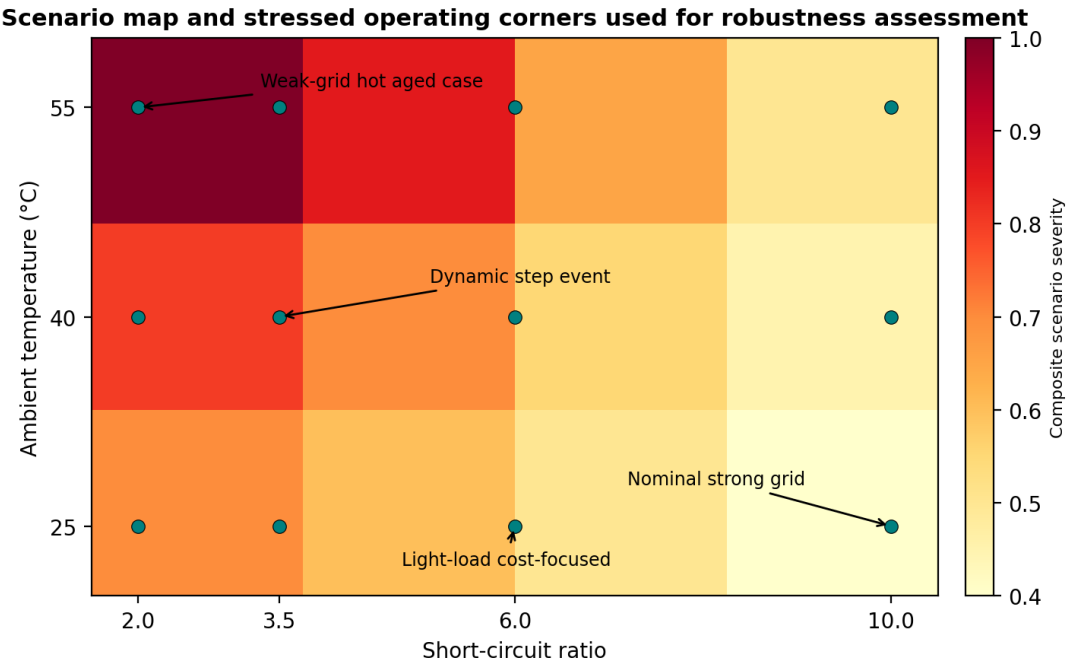

Figure 5: Scenario map and stressed operating corners used for robustness assessment.

Nominal conditions were defined by short-circuit ratio 10, ambient temperature 25 °C, zero parameter drift, unity-power-factor export near 0.5 per unit active power, and 1.0% background voltage distortion. Moderate-stress and high-stress conditions were introduced to cover weak-grid behavior, hot-environment operation, capacitor aging, and operating-point variation under lagging reactive demand. The same 216-scenario set was used for final rescoring of all retained candidates from all three methods.

Figure 5 scenario map and stressed operating corners used for robustness assessment. Structured uncertainty space used for full-scenario rescoring and stressed-corner analysis across nominal, weak-grid, hot-environment, aging-related, and dynamic-event conditions. Table 6 definition of nominal and stressed operating conditions across grid strength, temperature, parameter tolerance, operating point, and background harmonic content. The table establishes the uncertainty space used for robust scoring and scenario sweeps.

Table 6: Scenario space and uncertainty levels used in the experimental campaign.

Scenario factor	Nominal level	Lower or weaker corner	Upper or stronger corner	Sampling rule	Expected impact
Grid strength	Short-circuit ratio 10	Short-circuit ratio 2	Short-circuit ratio 10	{2, 3.5, 6, 10}	Lower short-circuit ratio increases resonance sensitivity and weak-grid interaction
Ambient temperature	25 °C	25 °C	55 °C	{25, 40, 55} °C	Higher temperature increases junction and capacitor hotspot stress

Tolerance bundle	L and C = 0%; ESR = 0%	L and C = ±10%; ESR = +20%	L and C = ±10%; ESR = +20%	3 bundles: nominal, moderate, stressed	Changes resonance location, ripple current, and damping effectiveness
Operating point	0.5 per unit active power, unity power factor	0.25 per unit active power	1.0 per unit active power with lagging reactive power	3 fixed operating states	Changes current magnitude, loss level, and controller loading
Background voltage THD	1.0%	1.0%	3.5%	{1.0, 3.5}%	Raises point-of-common-coupling harmonic burden and filter current stress
Dynamic event	None	Load-step event	Grid-impedance jump	Included in dynamic subset	Reveals margin reserve and transient robustness

4.3 Compared methods and evaluation protocol

Three methods were compared: DECODE-R, PSHF-IFA, and DR-DDTF. DECODE-R searched filter family, filter sizing, controller family, controller tuning, and switching frequency within the admissible benchmark space. PSHF-IFA was adapted as a single-tuned planning and sizing baseline with fixed current-control structure and a limited switching-frequency sweep. DR-DDTF was adapted as a damped double-tuned passive-filter baseline, again with fixed current-control structure and a limited switching-frequency sweep. Two directly related device-side studies discussed in Sections 1 and 2, namely the fast multi-objective inverter-filter optimization line of Karooyee et al. and the joint filter-plus-controller optimization line of Faria et al., were not reproduced as full baselines because the published papers do not provide a reusable common implementation, a shared uncertainty stack, or a directly transferable constraint envelope for the present benchmark. Accordingly, no direct numerical claim is made here against those studies within this benchmark.

All methods used the same objective set, the same scenario space, the same feasibility rules, and the same final rescoring protocol. Twenty-four independent runs were executed for each method, and each run used a budget of 12,000 evaluator calls. All archived candidates were rescored on the full 216-scenario set before infeasible points were removed and non-dominated filtering was applied. Hypervolume, spacing, and additive epsilon were computed on the fully rescored feasible union archive under a common normalization range.

Reduced-model optimization and archive post-processing were executed in Python 3.11 with automatic differentiation, sparse linear algebra, and fixed random seeds on a 16-core workstation with 64 GB memory. Electromagnetic transient verification used MATLAB/Simulink R2024b with Simscape Electrical, fixed-step discrete integration, a 1 microsecond electrical time step, and a final-window fast Fourier transform taken over the last 10 fundamental cycles after the transient interval.

Representative designs were selected after full rescoring. Cost-lean and harmonic-lean points were taken from the feasible non-dominated set by minimum cost proxy and minimum nominal total harmonic distortion, respectively. The balanced point was defined as the feasible non-dominated design with minimum normalized distance to the ideal point Eq 32.

$$d_{bal}(x) = \sqrt{\Sigma\left(g_i(x) - g_i^*\right)^2} \quad \text{...(Eq 32)}$$

In Eq. (4.1), $g_i(x)$ denotes the i^{th} normalized objective coordinate of design x after full rescoring, and g_i^* denotes the corresponding ideal-point coordinate from the feasible union archive. The four objective coordinates correspond to fully rescored total harmonic distortion, total loss, cost proxy, and StressRisk. Table 7 summary of optimization scope, structural flexibility, optimized variables, repeated-run settings, and evaluation budget for DECODE-R, PSHF-IFA, and DR-DDTF under a common comparison protocol.

Table 7: Compared methods, design freedom, and evaluation budget.

Method	Structure freedom	Design variables optimized	Runs	Evaluation budget	Output type	Expected strength	Expected limitation
DECODE-R	Search over admissible filter families and controller families	Topology, inductance, capacitance, resistance, controller choice, controller tuning, switching frequency	24	12,000 evaluations per run	Dense Pareto archive	Strong robust trade-off balance across scenarios	Higher optimization overhead
PSHF-IFA	Single-tuned passive family only	Single-tuned branch values, limited switching-frequency sweep, fixed dq current-control gains	24	12,000 evaluations per run	Sparse archive in low-cost region	Low hardware cost and simple structure	Narrow structural flexibility and weaker stressed-case robustness
DR-DDTF	Double-resistor damped double-tuned family only	Tuned branches, damping resistors, limited switching-frequency sweep, fixed dq current-control gains	24	12,000 evaluations per run	Archive within tuned family	Strong nominal harmonic rejection near tuned conditions	Higher passive loss and higher cost proxy

4.4 Main comparative results, sensitivity, and stressed-corner analysis

Results are interpreted in four linked layers. Figure 6 shows the Pareto-front structure under the common robust objective set. Table 8 reports representative non-dominated designs. Table 9 reports archive-level mean and standard-deviation values across repeated runs after full rescoring. Figure 7 and Table 10 report distribution-level and stressed-corner robustness, while Figure 8 and Figure 9 summarize representative-design trade-offs and sensitivity behavior.

The Pareto-front structure showed three distinct regions. DECODE-R occupied most of the low-total-harmonic-distortion and low-StressRisk region while retaining moderate cost and moderate loss. PSHF-IFA occupied the lowest-cost region and part of the low-loss region, but that region shifted toward higher total harmonic distortion and lower stressed-case feasibility. DR-DDTF remained competitive in a narrow nominal harmonic corner, but higher passive loss and higher cost proxy reduced overall multi-objective balance after full-scenario rescoring.

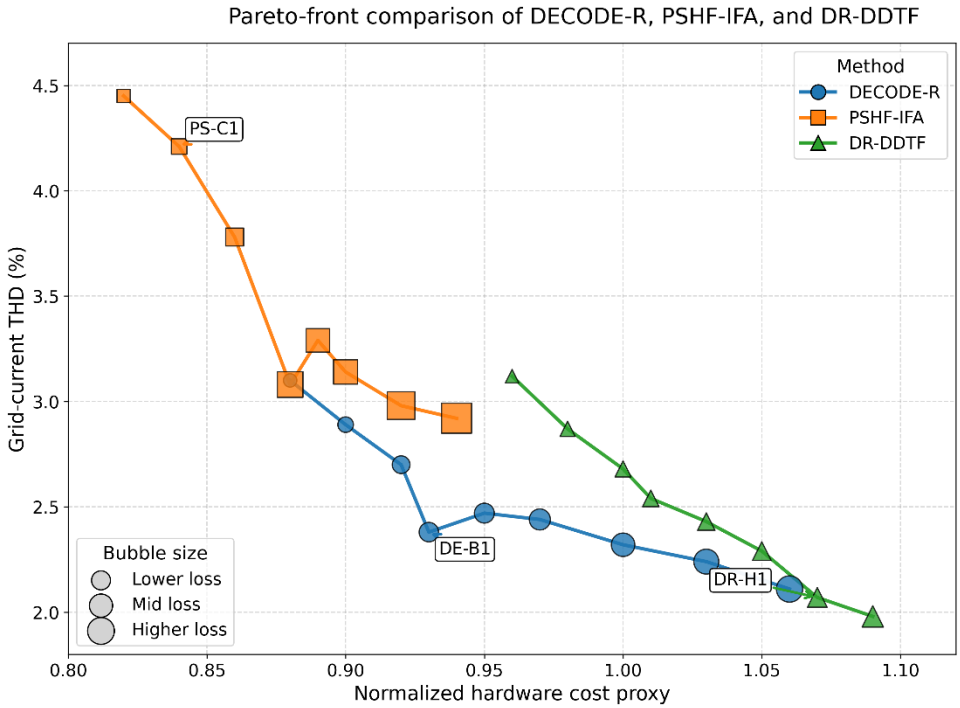

Figure 6: Pareto-front comparison of DECODE-R, PSHF-IFA, and DR-DDTF. Non-dominated solutions under the shared robust objective set, showing cost-lean, balanced, and harmonic-lean regions across the three compared methods.

Representative non-dominated designs are reported in Table 8. The balanced DECODE-R design, DE-B1, used a damped LLCL structure with dq current control and capacitor-current feedback at 10.8 kHz. After full-scenario rescoring, DE-B1 produced 2.38% mean point-of-common-coupling current total harmonic distortion, 738 W mean loss, 97.13% efficiency, cost proxy 0.93, and StressRisk 0.24. The harmonic-lean DECODE-R point, DE-T1, lowered mean total harmonic distortion to 2.11%, but that gain came with higher cost proxy, higher loss, and slightly higher StressRisk. PS-C1 remained the lowest-cost representative point and the lowest-loss point in the comparison set, whereas DR-H1 remained the nominal harmonic-lean baseline point with 2.07% mean total harmonic distortion at the expense of higher passive loss and higher cost proxy.

Table 8: Representative non-dominated designs and key performance metrics.

Method	Representative design ID	Filter family or topology	Key sizing values	Controller family	Switching frequency	Mean THD	Mean loss	Efficiency	Cost proxy	StressRisk
DECODE-R	DE-B1	Damped LLCL	Lc = 1.35 mH; Lg = 0.58 mH; Cf = 18 µF; Lt = 0.16 mH; Rd = 6.2 Ω	dq-PI with capacitor-current feedback	10.8 kHz	2.38%	738 W	97.13%	0.93	0.24
DECOD	DE-T1	Trap-	Lc = 1.12	PR with	13.6 kHz	2.11%	808	96.87%	1.06	0.29

E-R		assisted LCL	mH; Lg = 0.51 mH; Cf = 16 µF; fifth-harmonic trap branch enabled	notch compensation			W			
PSHF-IFA	PS-C1	Single-tuned	L = 1.02 mH; Ct = 9.5 µF; tuned-Q = 32	Fixed dq current control	8.0 kHz	4.21%	684 W	97.34%	0.84	0.39
PSHF-IFA	PS-B1	Single-tuned	L = 1.28 mH; Ct = 12.2 µF; tuned-Q = 38	Fixed dq current control	10.0 kHz	3.08%	777 W	96.89%	0.88	0.34
DR-DDTF	DR-H1	Double-resistor damped double-tuned	Ls = 1.18 mH; C5 = 7.2 µF; C7 = 10.6 µF; Rd1 = 2.9 Ω; Rd2 = 4.8 Ω	Fixed dq current control	9.6 kHz	2.07%	826 W	96.75%	1.07	0.33
DR-DDTF	DR-B1	Double-resistor damped double-tuned	Ls = 1.34 mH; C5 = 6.6 µF; C7 = 9.8 µF; Rd1 = 2.6 Ω; Rd2 = 4.4 Ω	Fixed dq current control	9.0 kHz	2.54%	801 W	96.84%	1.01	0.30

Archive-level aggregate performance is reported in Table 9. Those values summarize the retained Pareto archives after 24 repeated runs and full-scenario rescoring. DECODE-R achieved the lowest mean total harmonic distortion among the archives, the lowest mean StressRisk, the highest feasibility rate, the strongest normalized hypervolume, the smallest spacing value, and the best additive epsilon indicator. PSHF-IFA retained the lowest mean cost proxy, the lowest mean loss, and a slight mean-efficiency advantage, but archive feasibility and stressed-case behavior remained weaker. DR-DDTF remained intermediate in most archive-level metrics and retained stronger nominal tuned-harmonic behavior than PSHF-IFA, but passive-loss burden and cost proxy shifted the archive away from the most balanced region.

Table 9: Aggregate multi-objective performance comparison across methods after 24 repeated runs and full-scenario rescoring.

Method	THD mean ± std	Loss mean ± std	Efficiency mean ± std	Cost proxy mean ± std	StressRisk mean ± std	Power factor mean ± std	Feasibility rate	Hypervolume	Spacing	Additive epsilon
DECODE-R	2.63 ± 0.41%	746 ± 39 W	97.05 ± 0.19%	0.92 ± 0.05	0.24 ± 0.05	0.994 ± 0.003	93.6 ± 3.7%	0.812 ± 0.024	0.118 ± 0.021	0.087 ± 0.016
PSHF-IFA	3.21 ± 0.58%	731 ± 33 W	97.11 ± 0.16%	0.88 ± 0.04	0.35 ± 0.07	0.989 ± 0.004	75.8 ± 5.1%	0.634 ± 0.041	0.157 ± 0.026	0.146 ± 0.024
DR-DDTF	2.74 ± 0.49%	812 ± 46 W	96.79 ± 0.22%	1.03 ± 0.06	0.31 ± 0.06	0.991 ± 0.004	82.4 ± 4.6%	0.701 ± 0.033	0.141 ± 0.023	0.121 ± 0.019

Distribution-level robustness is summarized in Figure 7. Median scenario total harmonic distortion values were 2.54% for DECODE-R, 2.63% for DR-DDTF, and 3.14% for PSHF-IFA. The corresponding 95th-percentile values were 3.55%, 4.10%, and 4.78%, and the corresponding 0.95 conditional value at risk values were 3.43%, 4.06%, and 4.72%. The same ordering was observed for tail StressRisk values, with DECODE-R at 0.33, DR-DDTF at 0.41, and PSHF-IFA at 0.48. Figure 8 shows the representative-design loss, efficiency, and risk trade-off space. PSHF-IFA preserved the lowest-cost and slightly highest-efficiency point in the light-load region, while DE-B1 retained the lowest combined loss-risk envelope among the balanced designs and DR-H1 retained the narrow nominal harmonic advantage with a higher passive-loss burden.

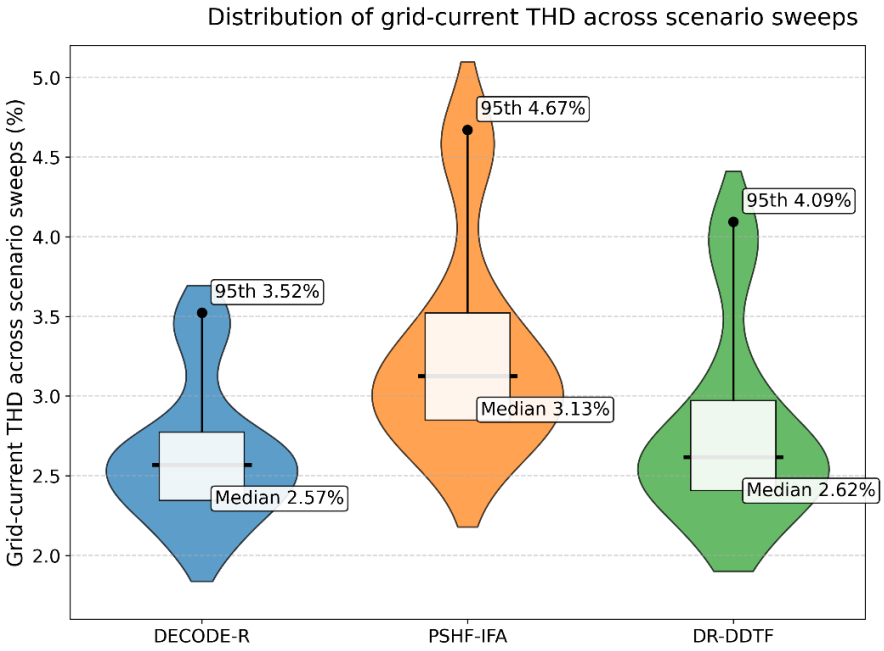

Figure 7: Distribution of grid-current THD across scenario sweeps. Scenario-level total harmonic distortion distributions after full rescoring, showing median behavior, upper-tail separation, and spread under nominal and stressed conditions.

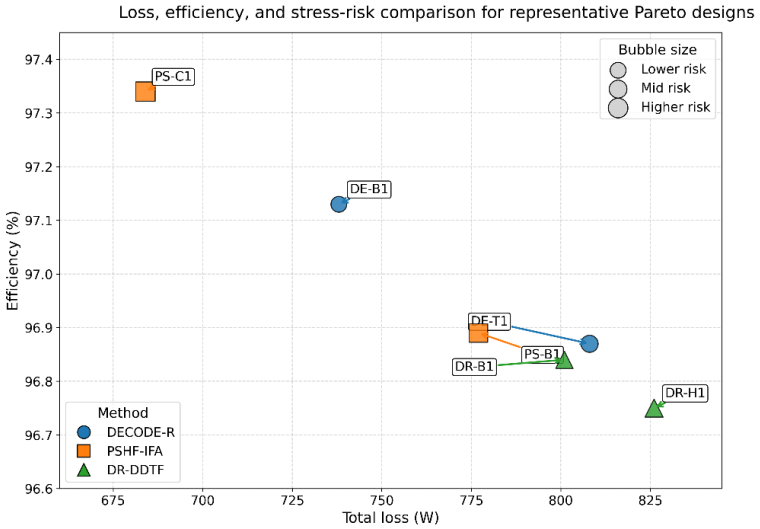

Figure 8: Loss, efficiency, and StressRisk comparison for representative Pareto designs. Representative-design trade-offs among total loss, efficiency, and stress-related risk across cost-lean, balanced, and harmonic-lean operating points.

Sensitivity behavior is summarized in Figure 9. An increase in switching frequency from 8 to 14 kHz reduced median total harmonic distortion by about 0.56 percentage points for DECODE-R, 0.41 percentage points for PSHF-IFA, and 0.49 percentage points for DR-DDTF, while increasing mean loss by about 96 W, 83 W, and 112 W, respectively. A negative 10% capacitance deviation increased median total harmonic distortion by 0.23 percentage points for DECODE-R, 0.61 percentage points for PSHF-IFA, and 0.44 percentage points for DR-DDTF. A short-circuit-ratio drop from 10 to 2 reduced phase margin by 11.7°, 18.4°, and 14.8%, respectively. A positive 20% equivalent-series-resistance drift increased median StressRisk by 0.05, 0.11, and 0.08, respectively. Those trends explain why DECODE-R retained stronger stressed-case behavior without removing the low-cost region preserved by PSHF-IFA or the narrow nominal tuned advantage preserved by DR-DDTF.

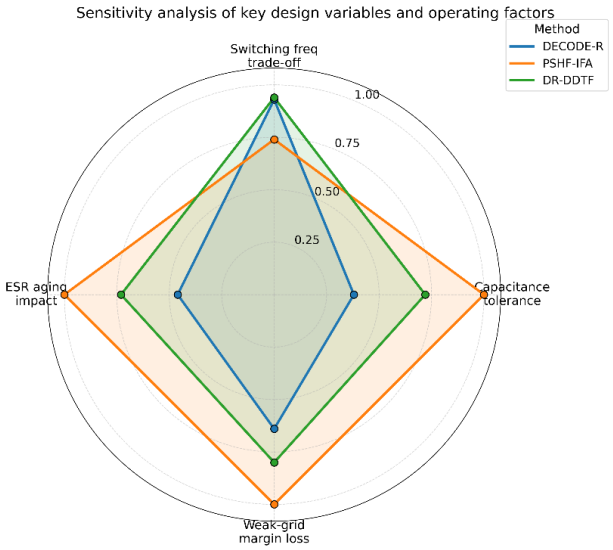

Figure 9: Sensitivity analysis of key design variables and operating factors. Relative method response to switching-frequency variation, capacitance deviation, weak-grid operation, and equivalent-series-resistance aging.

Stressed-corner robustness is reported in Table 10. DR-DDTF preserved the lowest nominal-corner total harmonic distortion by a narrow margin. Under the weak-grid, hot, aged-equivalent-series-resistance, high-background-harmonic corner, DECODE-R retained the lowest total harmonic distortion and the largest thermal reserve. Under the short-circuit-ratio 3.5 load-step event, DECODE-R retained the fastest settling time and the strongest dynamic margin. In the equivalent-series-resistance-aged capacitor corner, DECODE-R also retained the lowest ripple loading and the lowest risk score among the three compared methods.

Table 10: Scenario-robustness summary under stressed operating corners.

Scenario corner	DECODE-R	PSHF-IFA	DR-DDTF	Main observation
Strong-grid nominal corner	2.24% THD / 91 °C / pass	2.98% THD / 93 °C / pass	2.18% THD / 95 °C / pass	DR-DDTF showed the lowest nominal total harmonic distortion by a narrow margin
Weak-grid, 55 °C, aged equivalent-series resistance, high background harmonics	3.58% THD / 108 °C / pass	4.87% THD / 116 °C / marginal	4.12% THD / 113 °C / pass	DECODE-R retained the lowest stressed-case total harmonic distortion and the largest thermal reserve
Light-load, cost-focused corner	2.71% THD / 97.10% / cost 0.93	3.48% THD / 97.28% / cost 0.84	2.83% THD / 96.88% / cost 1.02	PSHF-IFA remained the lowest-cost and slightly highest-efficiency option in the light-load region
Short-circuit-ratio 3.5 load-step event	24 ms / 52° / pass	37 ms / 43° / marginal overshoot	29 ms / 48° / pass	DECODE-R retained the fastest settling time and the strongest dynamic margin
Equivalent-series-resistance-aged capacitor corner	ripple 0.78 per unit / risk 0.29	ripple 0.94 per unit / risk 0.43	ripple 0.85 per unit / risk 0.36	DECODE-R lowered capacitor stress under aging-related conditions

4.5 EMT verification and computational considerations

Electromagnetic transient verification transferred one representative balanced design from each Pareto archive to high-fidelity waveform simulation. The validated designs were DE-B1, PS-B1, and DR-B1. The electromagnetic transient suite contained three cases: nominal steady-state operation, a weak-grid hot-environment case, and a load-step case. Fixed-step discrete electromagnetic transient simulation used a 1 microsecond electrical time step. Steady-state total harmonic distortion was extracted from the last 10 fundamental cycles of each case after removal of the initial transient interval.

Reduced-model predictions and electromagnetic transient values are compared in Table 11. Agreement remained close but not identical, which is appropriate for a reduced harmonic-dynamic evaluator. For DE-B1, mean total harmonic distortion changed from 2.38% in the reduced model to 2.53% in electromagnetic transient simulation, and mean loss changed from 738 W to 761 W. For PS-B1, the corresponding deviations were 3.08% to 3.32% and 777 W to 801 W. For DR-B1, the corresponding deviations were 2.54% to 2.61% and 801 W to 829 W. Electromagnetic transient feasibility remained 3 out of 3 cases for DE-B1 and DR-B1, while PS-B1 retained feasibility in 2 out of 3 cases because the weak-grid hot case became marginal.

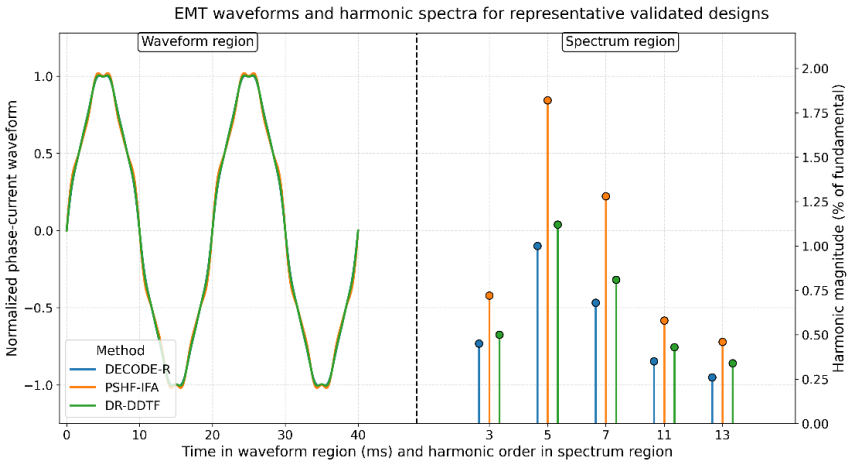

Figure 10: EMT waveforms and harmonic spectra for representative validated designs. Time-domain current quality and harmonic spectra used to confirm reduced-model ordering under nominal and stressed conditions.

Figure 10 supported the same qualitative ordering seen in the reduced-model study. DE-B1 retained the smoothest current waveform in the weak-grid case, DR-B1 retained strong nominal harmonic attenuation, and PS-B1 retained the broadest transient envelope during the step event. The reduced-model-to-electromagnetic-transient deviation remained small enough to support the use of the differentiable evaluator inside the optimization loop while preserving a separate high-fidelity confirmation stage.

Computational effort was dominated by repeated scenario evaluation of harmonic metrics and stability checks rather than by electromagnetic transient simulation, because electromagnetic transient validation was applied only to a small set of final candidates. Under the fixed evaluation budget, DECODE-R required 4.3 hours per run, DR-DDTF required 3.4 hours per run, and PSHF-IFA required 3.1 hours per run. The same budget of 12,000 evaluator calls per run produced a retained feasible Pareto archive of 46 ± 6 solutions for DECODE-R, 24 ± 5 solutions for DR-DDTF, and 19 ± 4 solutions for PSHF-IFA. Runtime therefore favored PSHF-IFA, but retained feasible archive quality favored DECODE-R.

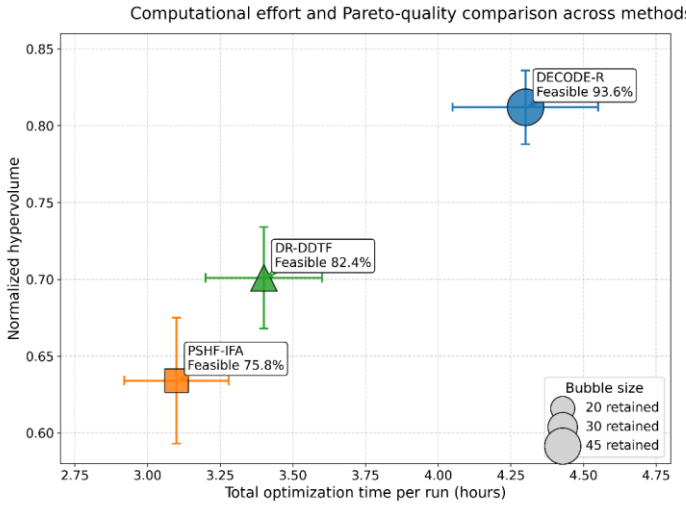

Figure 11: Computational effort and Pareto-quality comparison across methods. Runtime-quality trade-off across methods under the shared evaluation budget, using retained feasible hypervolume and archive size after full-scenario rescoring.

Computational workload was dominated by repeated scenario evaluation of harmonic metrics and stability checks rather than EMT simulation, because EMT was applied only to a small set of selected candidates. PSHF-IFA had the shortest total optimization time per run, followed by DR-DDTF, while DECODE-R required the largest initial optimization effort because of the coupled topology-control search. Figure 11, Computational effort and Pareto-quality comparison across methods, therefore showed a nontrivial trade-off: lower initial computational burden for PSHF-IFA, moderate burden for DR-DDTF, and higher burden for DECODE-R.

Table 11: EMT verification and computational summary.

Method	Representative EMT design	Predicted THD	EMT THD	Predicted loss	EMT loss	Feasibility after EMT	Evaluator calls	Total optimization time	Retained Pareto solutions	Main computational note
DECODE-R	DE-B1	2.38%	2.53 %	738 W	761 W	3 out of 3 cases pass	12,000 per run	4.3 h per run	46 ± 6	Highest initial overhead and strongest retained feasible archive
PSHF-IFA	PS-B1	3.08%	3.32 %	777 W	801 W	2 out of 3 cases pass	12,000 per run	3.1 h per run	19 ± 4	Lowest setup burden and lowest-cost region, but weaker stressed-case robustness
DR-DDTF	DR-B1	2.54%	2.61 %	801 W	829 W	3 out of 3 cases pass	12,000 per run	3.4 h per run	24 ± 5	Moderate runtime, strong nominal tuned response, and narrower thermal reserve in the hot weak-grid case

Across the shared evaluation budget, Section 4 supports one consistent interpretation. PSHF-IFA preserves the lowest-cost region and the simplest passive structure. DR-DDTF preserves a narrow nominal tuned-harmonic corner. DECODE-R preserves the strongest overall balance when feasibility rate, upper-tail total harmonic distortion, stress-related risk, and retained Pareto quality are considered together under structured deployment uncertainty.

5 Conclusion

DECODE-R provides a device-level co-design workflow for grid-tied converters under shared THD-loss-cost-stress objectives and shared feasibility constraints. The framework couples differentiable topology-control search with stress-aware scenario evaluation, realizable discrete extraction, fixed-structure refinement, and electromagnetic transient verification. The contribution therefore lies in one deployment-oriented pipeline for coupled filter and control design rather than in separate sequential tuning steps. Across the shared benchmark and common feasibility envelope, the study supports one consistent empirical interpretation. PSHF-IFA retains the lowest-cost region and the simplest passive structure. DR-DDTF preserves a narrow nominal harmonic advantage near tuned operating conditions. DECODE-R is most useful when the design target requires simultaneous control of total harmonic distortion, loss, cost proxy, and StressRisk under deployment uncertainty. At the representative-design layer, the balanced design DE-B1 delivered 2.38% mean point-of-common-coupling current total

harmonic distortion and 738 W mean loss after full-scenario rescoring. Electromagnetic transient validation yielded 2.53% total harmonic distortion, and all three validation cases remained feasible. At the archive level, DECODE-R retained the highest feasible hypervolume and the strongest feasibility rate across repeated runs. This result indicates stronger trade-off balance under uncertainty rather than isolated nominal superiority on a single metric. Several limitations remain. StressRisk is a physics-informed stress proxy derived from reduced electro-thermal mappings and should not be interpreted as a direct lifetime estimate. The uncertainty space, although structured, remains finite and does not exhaust rare grid events, long-horizon degradation, or feeder reconfiguration effects. The harmonic-balance and averaged dynamic models improve optimisation tractability but do not replace hardware validation, detailed high-frequency parasitic modelling, or full mission-profile aging assessment, and discrete extraction can still introduce mild mismatch before local refinement and electromagnetic transient confirmation. Future work should extend the admissible topology library, strengthen electro-thermal fidelity, incorporate richer scenario weighting or distributional uncertainty, connect device-level co-design with feeder-level constraints, and move toward hardware-in-the-loop or experimental validation of the retained designs.

REFERENCES

[1] Iqbal, Hasan, and Arif Sarwat. “Design and Implementation of Hybrid GA-PSO-Based Harmonic Mitigation Technique for Modified Packed U-Cell Inverters.” Energies 18, no. 1 (December 31, 2024): 124. https://doi.org/10.3390/en18010124.

[2] Ali I. Gony, Hashim, Ghamgeen Izat Rashed, Ansuman Badjan, Ahmed O. M. Bahageel, and Husam I. Shaheen. “Advanced Harmonic Mitigation in Photovoltaic Grid-Connected Systems: Jellyfish Search Optimization Approach.” ELEKTRIKA- Journal of Electrical Engineering 23, no. 3 (December 29, 2024): 6–18. https://doi.org/10.11113/elektrika.v23n3.544.

[3] Abbas, Ahmed S., Ragab A. El-Sehiemy, Adel Abou El-Ela, Eman Salah Ali, Karar Mahmoud, Matti Lehtonen, and Mohamed M. F. Darwish. “Optimal Harmonic Mitigation in Distribution Systems with Inverter Based Distributed Generation.” Applied Sciences 11, no. 2 (January 15, 2021): 774. https://doi.org/10.3390/app11020774.

[4] Hsiao, Ying-Tung, Shu-Min Lin, Shyi-Ming Chen, and Chih-Ju Chou. “Multi-Objective Optimization of Planning Single-Tuned Harmonic Filter Utilizing Interactive Forest Algorithm.” Information Sciences 661 (March 2024): 120141. https://doi.org/10.1016/j.ins.2024.120141

[5] AbuZahew, Ahlam, Sanaa Salama, Ibrahim Mahariq, and Muhammet Server Firat. “Hybrid Genetic Algorithm-Particle Swarm Optimization Based Approach for Optimal Sizing and Placement of a Passive Filter,” June 12, 2024. https://doi.org/10.21203/rs.3.rs-4447556/v1.

[6] Abubakar, Ali, Reindorf Nartey Borkor, and Peter Amoako-Yirenkyi. “Stochastic Optimal Harmonic Suppression with Permissible Photovoltaic Penetration Level for Grid-Linked Systems Using Monte Carlo-Based Hybrid NSGA2-MOPSO.” Energy Systems, October 23, 2024. https://doi.org/10.1007/s12667-024-00700-9.

[7] Daramukkala, Pavankumar, and Kanungo Barada Mohanty. “Development of Robust Multi-Objective Mayfly Optimization Method for Passive Power Filter Design to Improve Power Quality.” 2024 IEEE 3rd International Conference on Electrical Power and Energy Systems (ICEPES), June 21, 2024, 1–6. https://doi.org/10.1109/icepes60647.2024.10653508.

[8] Karooyee, Mohammed, and S. Ali Khajehoddin. “Search Space Reduction for Fast Multi-Objective Optimization of Inverter LCL Filters.” 2024 IEEE Workshop on Control and Modeling for Power Electronics (COMPEL), June 24, 2024, 1–6. https://doi.org/10.1109/compel57542.2024.10614067.

[9] Faria, João, João Fermeiro, José Pombo, Maria Calado, and Sílvio Mariano. “Proportional Resonant Current Control and Output-Filter Design Optimization for Grid-Tied Inverters Using Grey Wolf Optimizer.” Energies 13, no. 8 (April 14, 2020): 1923. https://doi.org/10.3390/en13081923.

[10] Rahman, Syed, Irfan Khan, Muhammad Ashar Ayaz, and Atif Iqbal. “A Distributed Filtering-Based Comprehensive Framework to Optimize Grid Performance of Harmonic-Polluted Distribution Systems.” IEEE Transactions on Industry Applications 60, no. 6 (November 2024): 8159–67. https://doi.org/10.1109/tia.2024.3432848.

[11] Wang, Lei, Wen Zhou, Can Su, Jiawen Fan, Weikuo Kong, and Pan Li. “A Distributed Harmonic Mitigation Strategy Based on Dynamic Points Incentive of Blockchain Communities.” Energies 17, no. 11 (May 31, 2024): 2683. https://doi.org/10.3390/en17112683

[12] Yin, Xiaoqing, Hao Yi, Zebin Yang, Yang Liu, Wenbin Liu, and Yao Zhang. “An Two-Stage Harmonic Optimization Method For Optimal Sitting and Sizing of APFs in Distribution Networks to Minimize THDs and Harmonic Losses.” 2024 Energy Conversion Congress &Amp; Expo Europe (ECCE Europe), September 2, 2024, 1–6. https://doi.org/10.1109/ecceeurope62508.2024.10751801.

[13] Hamza, Shrouk A., Mariem Y. Yousef, A. A. Ali, Said M. El Masry, and Magdi A. Mosa. “Multi-Stage Framework for Optimal Incorporating of Inverter Based Distributed Generator into Distribution Networks.” Scientific Reports 14, no. 1 (June 10, 2024). https://doi.org/10.1038/s41598-024-62708-6.

[14] Lin, Shu-Min, Chih-Ju Chou, Chia-Te Lu, and Ying-Tung Hsiao. “Optimal Design of Single-Tuned Filters for Industrial Power Systems to Reduce Harmonic Distortion Using Forest Algorithm.” Journal of the Chinese Institute of Engineers 47, no. 4 (April 2, 2024): 414–22. https://doi.org/10.1080/02533839.2024.2334205.

[15] Alhaider, Mohammed M., Shady H. E. Abdel Aleem, Ziad M. Ali, and Ahmed M. Zobaa. “Harmonics Management and Hosting Capacity Enhancement: Optimal Double-Resistor Damped Double-Tuned Power Filter with Artificial Hummingbird Optimization.” Edited by Mohit Bajaj. PLOS ONE 19, no. 5 (May 10, 2024): e0303207. https://doi.org/10.1371/journal.pone.0303207.

[16] Li, Xiaotong, He Wang, and Jinling Li. “Power Quality Control and Multi-Objective Optimization Compensation Strategy for Grid-Connected Multi-Functional Inverters.” 2024 China International Conference on Electricity Distribution (CICED), September 12, 2024, 969–74. https://doi.org/10.1109/ciced63421.2024.10754248.

[17] Adnan, Myasar Mundher, Ashish Sharma, T. Saravanan, Reena R, R.K. Bhagat, Bhuvaneshwari Jolad, and K. Kumaraswamy Naidu. “Advanced Control Strategies for Multilevel Inverters in Renewable Energy Systems: Enhancing Power Quality and Efficiency.” Edited by K.P. Navven, J. Guru Jawahar, K. Ramya, and B. Pragathi. E3S Web of Conferences 591 (2024): 01013. https://doi.org/10.1051/e3sconf/202459101013.

[18] Pavón, Wilson, Michael Chamorro, Ismael Michala, Diego Guffanti, and Mohan Kolhe. “Optimizing Electrical Systems Stability: A Novel Lyapunov Framework for Harmonic Reduction and Power Quality Enhancement.” 2024 3rd International Conference on Power Systems and Electrical Technology (PSET), August 5, 2024, 473–81. https://doi.org/10.1109/pset62496.2024.10808376.

[19] Metry, Morcos, Wesam Rohouma, and Robert S. Balog. “Online Weight Factor Tuning for MPC Regulated Power Quality Compensators.” 2024 4th International Conference on Smart Grid and Renewable Energy (SGRE), January 8, 2024, 1–6. https://doi.org/10.1109/sgre59715.2024.10428815.

[20] Venkata Anjani Kumar Gaddam, and Manubolu Damodar Reddy. "TLBO Trained an ANN-Based DG Integrated Shunt Active Power Filter to Improve Power Quality." Journal of Advanced Research in Applied Sciences and Engineering Technology 43, no. 2 (April 17, 2024): 93–110. https://doi.org/10.37934/araset.43.2.93110.